



How eG Enterprise enables business continuity and high performance for key IT services with AIOps

Leveraging the power of machine learning and automation to make IT management proactive, fast, and simple



IT service quality is now critical to support business growth, customer satisfaction and employee efficiency

As organizations pivot towards digital transformation, they are becoming increasingly dependent on IT applications and infrastructures and the teams that support them. There are four key challenges to address:



Availability

More than ever before, application and business service availability and performance are now critical to core business operations. IT operations has the task of meeting demanding SLAs required by the business.



End user experience

Outages, slowdowns and poor user experience are unacceptable to both customers and internal employees and impact on customer experience, staff morale and ultimately business revenues.



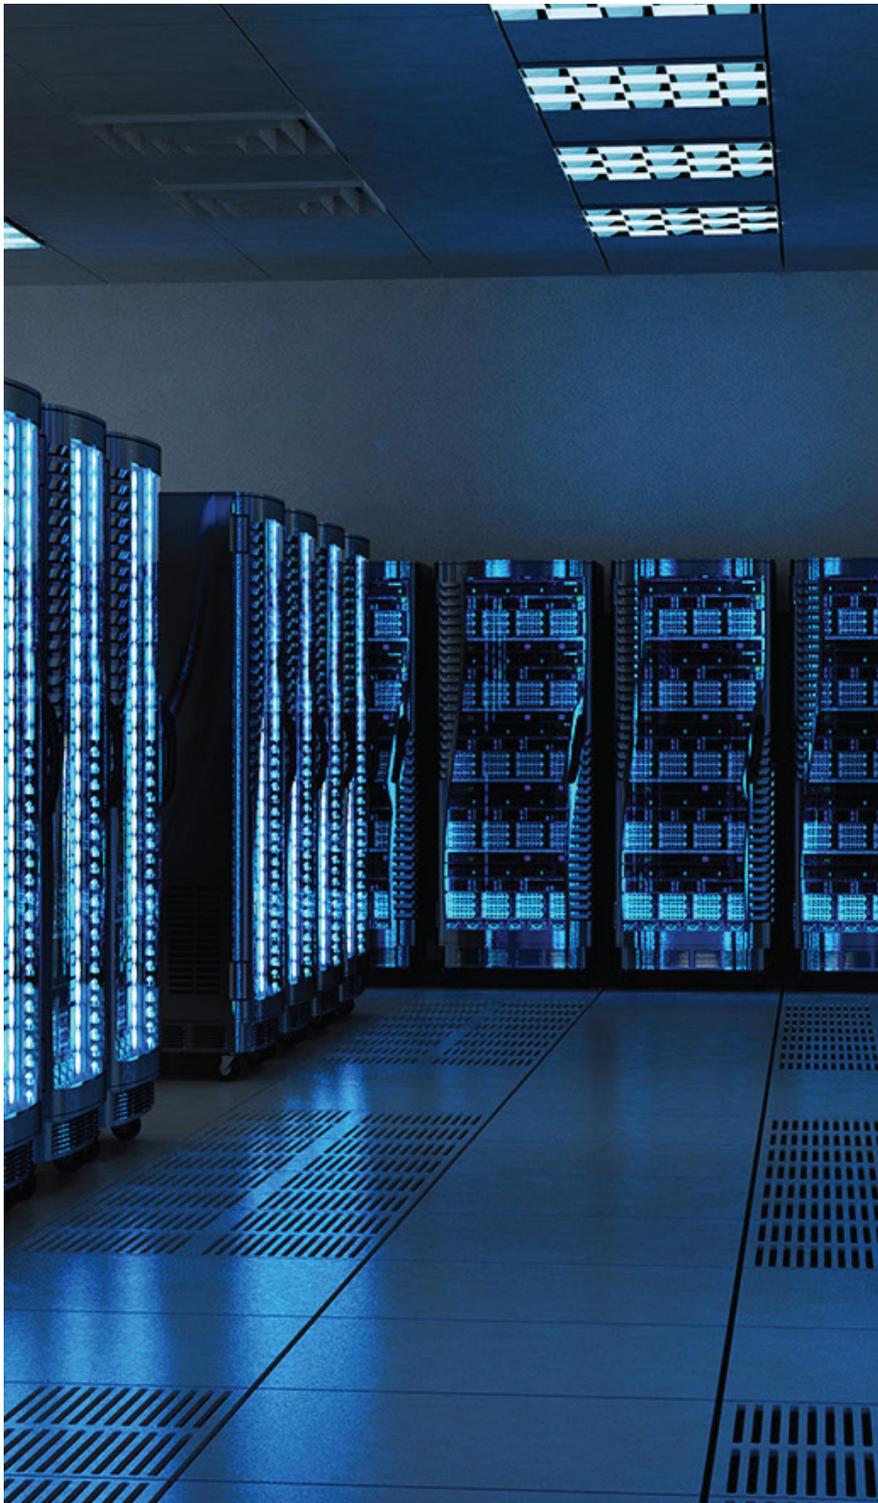
Complexity

Applications and infrastructures are becoming increasingly complex, dynamic and distributed including many specialized technologies and multiple deployment models including hybrid cloud, SaaS and on-premises models. IT operations teams must not only have expertise in these technologies, they also need to be able to analyze the large volume of metrics from the IT technologies.



Legacy team structures

Organizational silos still exist in practice with specialist teams owning domains of expertise and tools. This often leads to finger-pointing, dependence on experts, slow analysis and troubleshooting and increased MTTR.



Traditional monitoring is insufficient at scale

Monitoring tools trigger alerts when pre-determined thresholds are crossed. IT managers often have several common complaints about the alerting process:

Numerous false alerts

A false alert refers to situations in which the monitoring tool indicates there is a problem, but the IT manager determines that there is no real issue or impact. Thousands of alerts can distract and overwhelm IT administrators and prevent them from focusing on or identifying real issues that can impact IT service quality.

Tedious threshold setting

To avoid false alerts, IT managers must define threshold values for the different metrics collected. In a large enterprise, monitoring tools can collect millions of metrics. Having to set thresholds manually for every metric is a time-consuming and monotonous process. This can be time consuming and expensive.

Don't handle time-varying norms

The right threshold for a metric to trigger an alert can vary wildly depending on the context of what is normal behavior and the usage level for each and every system. Certain metrics also vary based on time of day or day of month. For instance, network traffic could be high during office hours but would be highly unusual at 3am on a Sunday.

Reactive, not proactive

IT managers must be able to detect problems in advance. Often IT monitoring is reactive – i.e., alerts are triggered after a service SLA has been breached. Monitoring tools must provide IT managers early warning indicators, so preemptive action can be taken.

AIOPs

AIOPs (Artificial Intelligence for IT Operations) refers to the application of machine learning technology that enhance IT operations analytics.

How eG Enterprise's machine learning, intelligent correlation and advanced analytics help organizations:

- Actionable alerts** from millions of metrics, logs and signals collected in real-time
- Proactively learn** about issues that may affect service delivery down the line
- Remediate application and infrastructure issues** quickly and accurately
- Drive continuous improvement** in service quality and operational efficiency

AIOPs is at the core of eG Enterprise's monitoring capabilities

From its inception, eG Innovations has incorporated machine learning technologies to build a powerful AIOPs platform that overcomes the limitations of traditional monitoring approaches. Automatic root-cause diagnosis, the holy grail of IT operations management, has been a key focus of our eG Enterprise suite.

Based on years of experience with legacy tools that implemented if-then-else rules, our architecture is designed to support rules-free correlation, thereby minimizing the effort required to configure and operate the monitoring solution.

Our patented root cause diagnosis engine was the 1st technology built to be cloud and virtualization-aware. The layered stack model paradigm that eG Enterprise uses is at the core of its ability to assimilate millions of metrics to yield meaningful alerts.

With these varied capabilities, eG Enterprise pin-points and prioritizes underlying issues, avoids alert storms and distractions from secondary symptoms. IT teams can monitor volumes of metrics beyond what a human can tackle, support business-critical applications with minimal training and without needing specialist help, and handle large-scale and dynamic environments.

The following capabilities form the foundation of eG Enterprise's AIOPs capabilities:

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

Embedded domain expertise is key for effective monitoring

eG Enterprise monitors your entire IT landscape – from hardware to application, from network to storage, from virtualization to cloud and containers.

There is no single protocol or mechanism to monitor all these heterogeneous tiers. When architecting eG Enterprise, we placed a lot of expertise on “what to monitor” rather than “how to monitor”. Which metrics and signals to collect is determined based on a deep understanding of how each technology tier works and what are its failure modes. Vendor-recommended best practices and industry expert recommendations also go into determining the metrics that eG Enterprise collects.

Many monitoring tools are limited to collecting metrics exposed via technology APIs. eG enterprise is designed to understand and ingest millions of data points from:

- Application APIs
- OS commands and APIs
- Log files, syslog
- SNMP
- Web services functions
- Traces collected via byte-code instrumentation

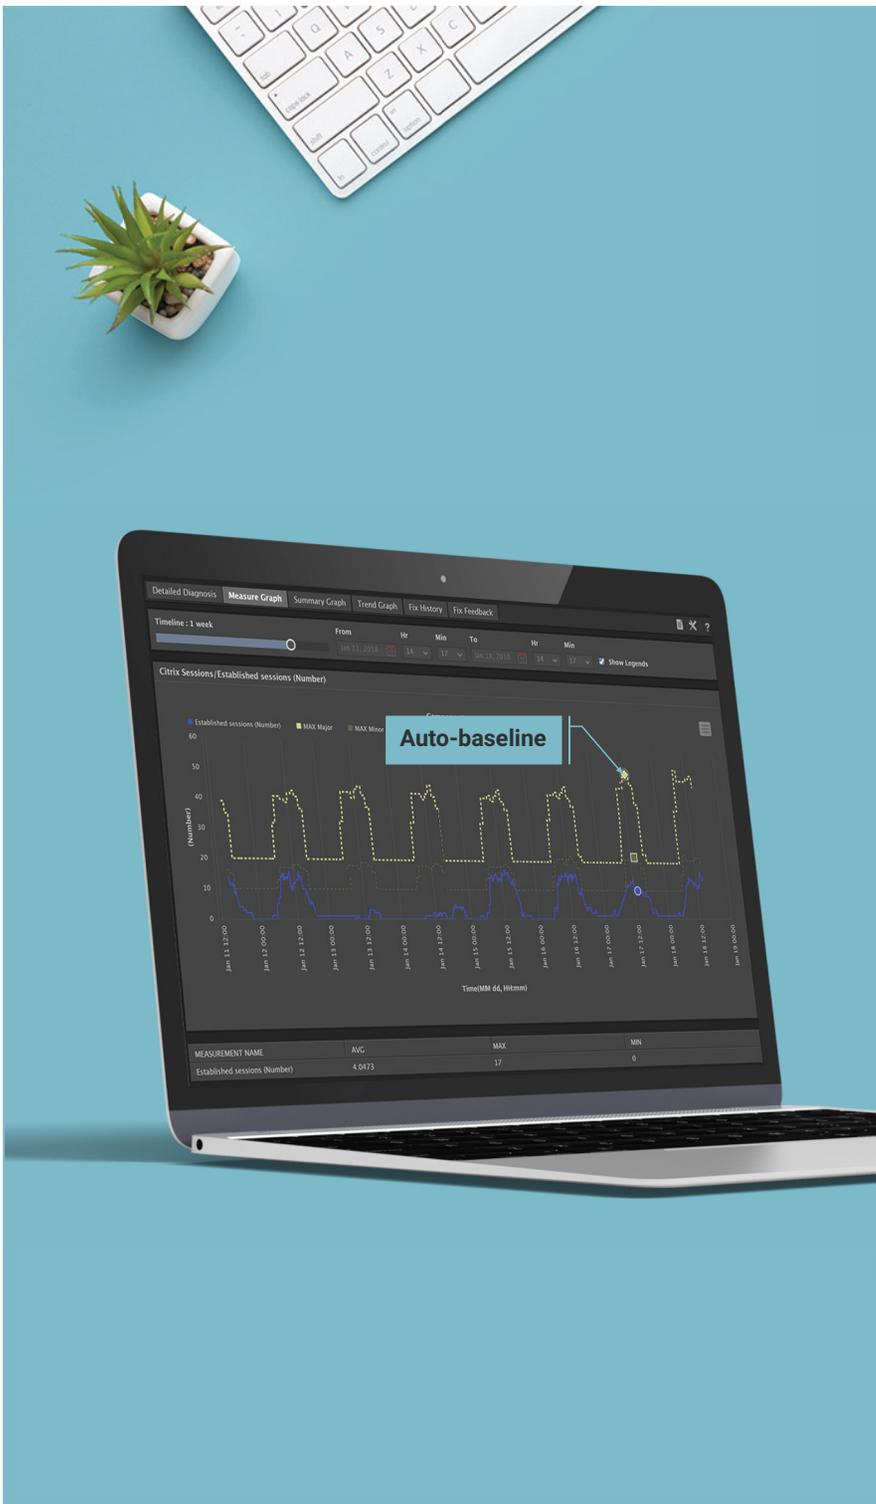
Domain-aware monitoring ensures that eG Enterprise has answers to the toughest IT issue of today, which is users asking “*why is my application slow?*”. Whether it is blocked threads in a Java application, database query not using indexes, or an abnormally high queue on an individual LUN in the storage layer, eG Enterprise has ready answers.



70%

of users stated that eG Enterprise helped *identify the cause of complex application performance* problems that involve multiple domains.

Source: TechValidate survey



Auto-baselining for proactive problem detection and anomaly detection

Thousands to millions of metrics may be collected by the monitoring solution in any medium to large scale IT infrastructure. Manually sifting through metrics and configuring thresholds is possible but time consuming and requires expertise. This manual approach isn't sufficient for metrics that may be time-varying – e.g., the number of users connected to a server. Thresholds may also need to be configured differently on different servers, accounting for differences in their capacities and configuration.

eG Enterprise's **auto-baselining technology** uses machine learning to determine normal bounds of metrics. This technology tracks time of day, day of week behaviour of each metric and uses past history to estimate dynamic thresholds for each metric. Administrators can choose the granularity with which they apply this derived intelligence to allow thresholds to be automatically and dynamically fine-tuned. Auto-baselines can be computed for minimum and/or maximum values - the configuration depends on the specific metric and administrators can override the setting.

Minimal expertise is required to configure the solution. If a metric generates too many alerts, administrators have an option to configure "leniency levels" that allow for some degree of variation from the normal baselines. Different leniency levels can be set for minor, major and critical alert limits.

Administrators can also specify static limits that control when auto-baselines are applied. This is particularly useful in non-production environments where for example, you don't want a few users logging into a lightly loaded server to trigger an alert.

Auto-baselining is a key to making the monitoring solution proactive and easy to use:

- **Early warning alerts:** Administrators get alerts to anomalies and abnormal usage of resources, unusual traffic patterns before end users are impacted.
- **Less manual work:** Administrators no longer need to configure thresholds one by one, for each metric, and for each system.

 **84%**

of users reported that eG Enterprise helped them avoid **application performance outages** by identifying problems before end users were affected.

Source: TechValidate survey



Automatic problem investigation and diagnosis

eG Enterprise checks key parameters regarding your applications and infrastructure in a manner similar to how medical tests check key parameters of humans. These tests are designed to impose minimal overhead on the target environment.

If anomalies or problems are detected, additional diagnosis is automatically performed to collect more details about a problem.

Deep domain expertise is required to determine what additional diagnosis is required. The diagnostic checks vary from one application to another and from system to system. For example, when a large number of IOPS are observed on a database server, the diagnosis collects information check about the top queries that are causing IOPS. The execution plans of these queries provides administrators the details they need to troubleshoot the problem effectively.



When a problem happens, you often don't have time to investigate. You may need to reboot a system and the problem might go away. But then the same problem could occur again ... Hence, collecting detailed diagnosis automatically when a problem occurs is vitally important.

With eG Enterprise's detailed diagnosis:

- When a large number of IOPS are observed on a database server, the diagnosis collects information about the top queries that are causing IOPS.
- When a user's logon to a virtual desktop is slow, the diagnosis gathers additional insights into what could have caused the slowness: is it profile loading, GPO processing, Active Directory authentication, logon scripts, etc.?

This is a key for enhanced user experience: accurate diagnosis, faster resolution and enhanced service performance.



72%

of users reported that eG Enterprise helped reduce the time to **identify the root cause of application performance.**

Source: TechValidate survey



End-to-end auto-correlation for root-cause analysis

When a problem occurs, it is essential to be able to determine what caused a problem. This is where **Root Cause Analysis (RCA)** comes in. Accurate and effective root-cause diagnosis saves organizations hours of time they spend in troubleshooting calls and in finger-pointing. Good RCA capability ensures that IT operations teams get a prioritized list of alerts – the highest priority points to the root-cause of problems, while the lowest priority often denotes the effects.

There are two reasons why root-cause analysis in modern IT environments is hard:

- Most IT organizations operate in silos, with separate teams responsible for various parts of the infrastructure and applications. Often these teams use different set of tools to monitor the health of different domains: cloud, virtualization, networking, storage and applications.
- Modern applications use multi-tier architectures where one tier depends on another tier to function. Inter-dependencies between tiers means if one tier has a problem, it affects other dependent tiers as well and it becomes difficult to determine where a problem started.

eG Enterprise implements an industry first virtualization and cloud-aware auto-correlation engine. This patented technologies uses various types of dependencies to learn about cause effect relationships:

- There are application-to-application interdependencies (e.g., a web server depending on a database server).
- Application components depend on network elements to communicate and hence, there are application-to-network dependencies.
- Application components often run on virtual machines (VMs) and VMs are hosted on physical machines.

Evaluating alerts from different tiers and considering the different inter-dependencies, eG Enterprise automatically pinpoints where the root cause of a problem lies. This allows IT operations teams to differentiate the cause of a problem from its effects and resolve performance issues in minutes rather than in days.



70%

of users stated that eG Enterprise helped *identify the cause of complex application performance* problems that involve multiple domains.

Source: TechValidate survey

Top-to-bottom auto-correlation for root-cause analysis

End-to-end correlation helps pinpoint which of many tiers in an infrastructure is responsible for a problem. The next challenge is to determine which layer is causing the issue: is it the operating system, or network connectivity, or the application infrastructure (e.g., JVM), or the application code.

There may be different administrators responsible for each tier – e.g., the system team manages the operating system, while the application operations team is responsible for the application infrastructure. Effective root-cause diagnosis is a key to determining which team has to work to resolve a problem.

eG Enterprise uses a stack/layer model representation of each tier to perform root-cause diagnosis.

- Metrics are mapped to layers. Layers are stacked hierarchically in such a way that the layers above depend on the layers below for their functioning.
- The state of a layer is based on the state of metrics that map to it.
- Alerts for a tier are correlated across layers, considering the hierarchical model for that tier. When two layers have problems of the same severity, the layer above is assigned a lower priority as it is likely to be an effect.

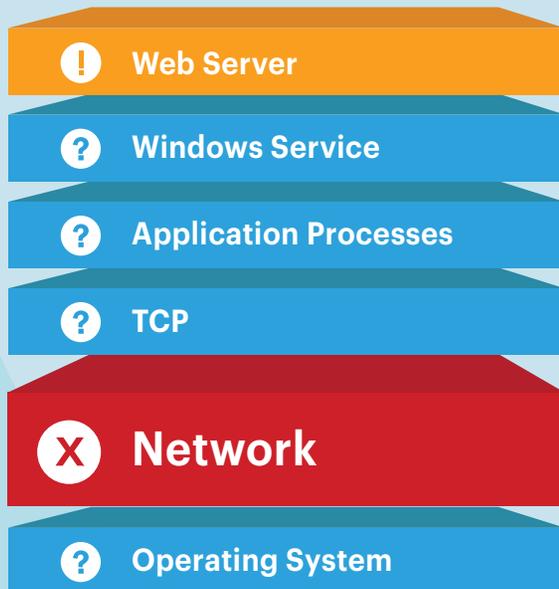
Performance alerts are correlated temporally, i.e., the monitoring tool learns and evolves its diagnosis as it gets more information about every layer and every tier.

The use of layer models allows clear demarcation and assignment of problems. For example, if a web server has an issue in the web server and network layers, the network layer which is lower in the hierarchy has the higher priority problem and hence, the problem is assigned to the network team for resolution.

Effective and automated root-cause analysis improves user experience in many ways. Service uptime is enhanced, SLAs are met. Problems will rarely occur after being resolved, so repeated outages are avoided. IT operations efficiency is also improved as IT admins can spend time on productive activities rather than in firefighting.

 **60%** of users found that *layered modelling of applications* to be one of the most useful capabilities of eG Enterprise.

Source: TechValidate survey



Root-Cause

Auto-discovery and dependency mapping

eG Enterprise uses embedded domain expertise to automatically discover application landscapes. Application specific expertise (e.g., Citrix, SAP, web application domain knowledge), application APIs, configuration information and application-specific commands are used by eG Enterprise to discover application topologies and inter-dependencies. Dynamic dependencies like VM to physical machine dependencies are discovered using virtualization APIs.

This makes the deployment and configuration of the monitoring system very easy. Manual overrides to the discovered topology are possible, should the auto-discovered information need to be modified. A unified overview of the entire topology enables any user to understand the dependencies beyond their own domain expertise enabling different teams within the business to work in synchronous harmony.

This auto-topology discovery capability addresses a number of challenges that traditional monitoring tools have had:

- Manually creating topology graphs is challenging. Most administrators may not even be aware of their service topologies. Furthermore, manually discovering these dependencies can be time consuming.
- At the same time, some dependencies are dynamic – e.g., an application and its VM can move between physical nodes in a virtualization cluster. Hence, to be effective monitoring tools must be smart enough to auto-discover the applications and infrastructure they are to monitor.
- Modern applications are deployed on virtualized and containerized environments where inter-application communications can happen within a physical server itself. Tools that rely on packet sniffing for topology discovery cannot function well in such environments.

Unlike many other solutions, the auto-topology discovery module is an integral part of eG Enterprise and does not require any additional license.



70%

of users stated that eG Enterprise helped *identify the cause of complex application performance* problems that involve multiple domains.

Source: TechValidate survey

Metric aggregation and service quality indicators

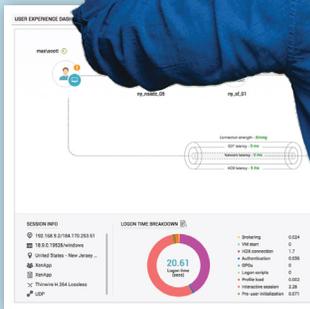
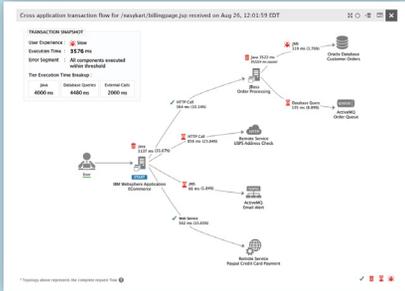
Analytics is a key tenet of eG enterprise. While monitoring of individual KPIs is important, there are several situations where it is more important to look at combinations of KPIs or to look at KPIs across multiple systems being monitored. eG Enterprise implements metric analysis and interpretation in several ways:

- **Performance ratings:** A performance rating is a single quantifiable rating/index that is representative of multiple metrics. Performance ratings can be defined for any component monitored. A weightage can be assigned for each metric assigned to a rating – so different priorities can be assigned to different metrics. A rating takes a value between 0 and 100, with 100 indicating a good rating and 0 indicating a poor rating.

The advantage of defining ratings is that administrators can see a simple and clear rating of overall performance, instead of having to dig through multiple different metrics. Organizations can set SLAs based on the overall user experience rating they are striving to achieve.

- **Aggregated metrics:** Production systems often have groups of systems working together to handle the workload – e.g., a web server cluster. In such a case, it is more important to look at the web server cluster as an aggregate, rather than looking at individual servers. eG Enterprise allows administrators to define aggregated metrics that represent the overall health of a cluster or a group of systems. Aggregated metrics provide enterprise-wide views of IT demand, quality and utilization.
- **Conditionally aggregated metrics:** Unlike simple aggregation where the metrics from multiple systems are combined together using min, max, average or sum functions, conditional aggregation applies conditions against metrics from each system in the group. For instance, an aggregated metric can reflect the percentage of servers in a group that have CPU utilization above 80%. Different alert levels can also be set – e.g., generate a minor alert when 20% of the servers have utilization above 80%, a major alert when 50% have utilization above 80% and a critical alert when 80% of servers have utilization above 80%. Using this capability, administrators can set complex conditions based on which they are alerted.

Thresholds, dashboards and alerts can be triggered based on aggregated metrics, just as they would on basic metrics.



Drive continuous improvement with bottleneck detection, forecasting and capacity planning

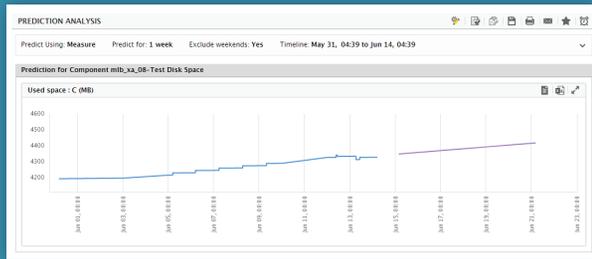
Monitoring tools are not just used for real-time monitoring and problem diagnosis. The wealth of data they collect can be used for intelligent analysis and to right-size and optimize the infrastructure and applications being monitored. eG enterprise embeds multiple capabilities in this area:

- KPI analysis and bottleneck detection:** Analysis of historical alerts is used to provide a clear indicator of which tiers and which layers are responsible for past issues in the infrastructure. Drilldowns allow for deeper analysis of bottlenecks.
- Forecasting:** You can use built-in prediction mechanisms and forecasting techniques implemented in eG Enterprise to automatically compute how a metric is likely to change in the future. Using this capability, you can determine when in the future the current resource capacity may get exhausted.
- Right-sizing:** For virtualized environments, eG Enterprise analyzes the performance of VMs over time and provides recommendations on which VMs are over-sized, which ones are under-sized, which VMs have not had significant activity in the past and are likely candidates for being powered off, etc.
- Capacity planning:** Trend analysis and outlier detection is used to analyze traffic patterns and resource utilization trends. Use these insights to determine how to plan effectively for future growth. For example, when deploying digital workspaces, you can determine how many more systems are needed to accommodate a thousand additional users.

KPI analysis



Forecasting



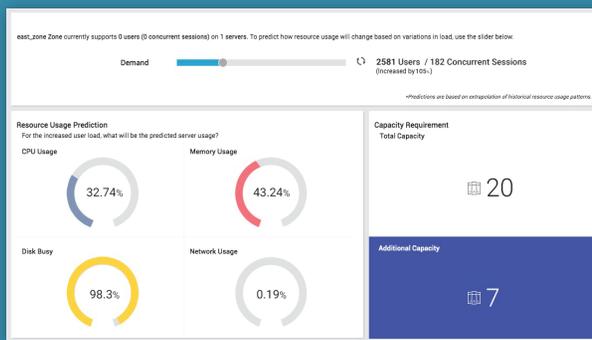
Right-sizing

VM Right Sizing Report - VMware vSphere ESX - Oversized VMs

HOST	VM	VCPUS USAGE (%)			AVG READY TIME(%)	NUMBER OF VCPUS	IS VM OVERSIZED?	REMARKS
		MIN	MAX	AVG				
VMwareESX15	BugTracker(9...	1.05	43.89	2.83	0.06	1	No	-
VMwareESX15	CBR-ESX-VPX(B...	0.33	0.42	0.36	0.06	2	Yes	vCPU Utilization
VMwareESX15	CentOs-7-0	0.63	1.09	14.96	0.08	1	No	-
VMwareESX15	DBZ10.5(11.11)	4.35	100	1.69	0.25	1	No	-
VMwareESX15	eG-Manager (1...	0.58	12.61	8.15	0.05	1	No	-
VMwareESX15	Excaltur2 (11...	1.36	31.11	20.73	0.09	2	Yes	vCPU Utilization
VMwareESX15	IBM IMT (6.8)	5.27	58.79	0.86	0.2	2	No	-
VMwareESX15	Informix (11.19)	0.48	1.9	5.3	0.4	1	No	-
VMwareESX15	Impressario (1...	0.33	26.63	3.1	0.11	1	No	-

Get insights for right sizing: Is the VM oversized?

Capacity planning



“

With eG Innovations, we are reducing system maintenance and support costs, avoiding incremental IT spending and eliminating system downtime across the hospital. Performance and prediction reports help us optimize IT spending and save \$100,000 per year. ”



Peter Bakker
Manager, ICT GHZ



Automatic correction and remediation

An important way to enhance IT operational efficiency is automatic problem remediation. When the cause of a problem is well understood, then appropriate auto-correction can be implemented. For instance, if a service that is in automatic mode is not running, it can be restarted by the monitoring tool. Likewise, when a disk drive fills up, temporary files in well-known locations can be deleted. Automated actions like these are built into the eG Enterprise universal monitoring technology.

Automatic correction can also be done at the application level. For such actions, we are taking a domain-specific view. Future releases of eG Enterprise will embed corrective actions specific to each domain of interest (e.g., digital workspaces, enterprise applications, etc.). Corrective actions can be initiated well before service level violations occur, thereby preventing IT events from impacting business.

Details of the actions taken and the results will be made available in the web console, so administrators are in the know of such actions.

To support scenarios where IT administrators may prefer to execute their own corrective scripts, eG Enterprise provides a way for administrators to define and upload scripts to the system. These scripts are executed by the agents as and when a threshold condition is violated. The scripts embed the intelligence to analyze the thresholds, determine the faulty metrics and perform the appropriate action. With the right auto-correction mechanisms in place, you can increase service uptime and lower mean time to repair (MTTR).



“

The ROI of eG Innovations exceeded our expectations. Not only did we get those chargebacks reduced, we also benefitted IT as a whole because help desk tickets weren't created and time management was a lot. ”

Peter Dinh

Senior Virtualization Engineering Lead, eBay Inc.



Why is an AIOps platform mandatory?

Every business and organization appreciates that keeping mission critical systems running 24/7 and ensuring continuous availability and performance for key applications is now core to enabling their business to function.

Digital transformation and trends such as remote working and offshoring mean IT teams are being asked to deploy and manage increasing numbers of applications and services. Trends such as cloud, container and virtualization technologies mean that the core infrastructure is increasingly complex and distributed.

eG Enterprise helps IT teams to overcome their growing challenges by enabling them eliminate error-prone and tedious manual work, after-the-event troubleshooting, and move to automated proactive workflows, thereby addressing problems before users or services are impacted.

eG Enterprise's AIOps, machine learning and patented root cause analysis technologies enables IT teams to:

- ✓ See a clear, prioritized alert window of issues and warnings for their attention
- ✓ Be proactively notified of the true root cause of outages, service degradations and application failures
- ✓ Accurately pinpoint and remediate performance and availability issues rapidly
- ✓ Drive continuous improvement in both internal business process and service quality and operational efficiency
- ✓ Integrate the eG Enterprise single pane of glass with your IT service desk tools like ServiceNow, PagerDuty, AutoTask etc. and your DevOps tools such as Ansible, Github, JIRA and many others.



How eG Innovations will support you on your journey to AIOps adoption

Thousands of businesses and organizations, including many in heavily regulated sectors such as healthcare, finance and government, are already leveraging the capabilities of eG Innovations. We understand the responsibility of delivering critical applications 24/7 whether it be to support neo-natal intensive care units, banking ATMs or Microsoft 365 applications.

- **A single platform:** The eG Enterprise AIOps platform has been built from the ground up by our vastly experienced R&D team. It is a single platform that has grown and evolved organically to meet the changing requirements of organizations around the world. This way, you do not have to deal with multiple disparate acquired technologies and consoles from a vendor.

Our technology experts stay abreast of the latest technology developments, so our products can offer immediate support for new releases and products from IT vendors with whom we've integrated.

- **Unbeatable support:** eG Innovations has offices and technical support staff around the globe to provide a follow-the-sun support model for customers. You can reach your local helpdesk by email or phone, or submit a ticket to our global helpdesk. Our helpdesk teams have engineers certified on our products and with a deep understanding of modern application and infrastructure landscapes.
- **Established and dependable:** For over 20 years, eG Innovations has been supporting businesses with their IT performance monitoring requirements. Continuous feedback from our global clients has helped evolve the product into a platform that meets the most challenging and complex of setups and technologies.

When you choose eG Innovations to provide your AIOps platform, you will be working with a company that has a proven track record of serving its customers first.



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.

For more information

Visit: www.eginnovations.com | Contact: info@eginnovations.com



www.eginnovations.com/it-monitoring/free-trial

