



The eG Suite

***Enabling Real-Time Monitoring and
Proactive Infrastructure Triage™***

White Paper



Restricted Rights Legend

The information contained in this document is confidential and subject to change without notice. No part of this document may be reproduced or disclosed to others without the prior permission of eG Innovations Inc. eG Innovations, Inc. makes no warranty of any kind with regard to the software and documentation, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose.

Copyright

© Copyright 2003 eG Innovations. All rights reserved. eGurkha and eG ASPlite are trademarks of eG Innovations. All other trademarks, marked and not marked, are the property of their respective manufacturers. Specifications subject to change without notice

Introduction

In the recent past, the complexity of Internet/Intranet services has grown dramatically. Many new business models, new customer-focused services, and efficient on-line collaboration services have emerged that improve the overall operational efficiency of businesses. To support these new services, IT infrastructures has grown in complexity. Rather than supporting simple client-server applications, IT infrastructures are now designed to comprise of multiple inter-operating tiers. The front-end includes firewalls to safeguard against malicious attacks, web servers to handle user traffic, and load balancers to distribute traffic amongst all the web servers. The back-end has grown to be even more complex. While the web servers mainly act as HTML gateways that forward user requests, it is the middleware application servers hosting the business logic components that communicate with database servers, payment gateways, order processing systems, etc., to accomplish the business functions.

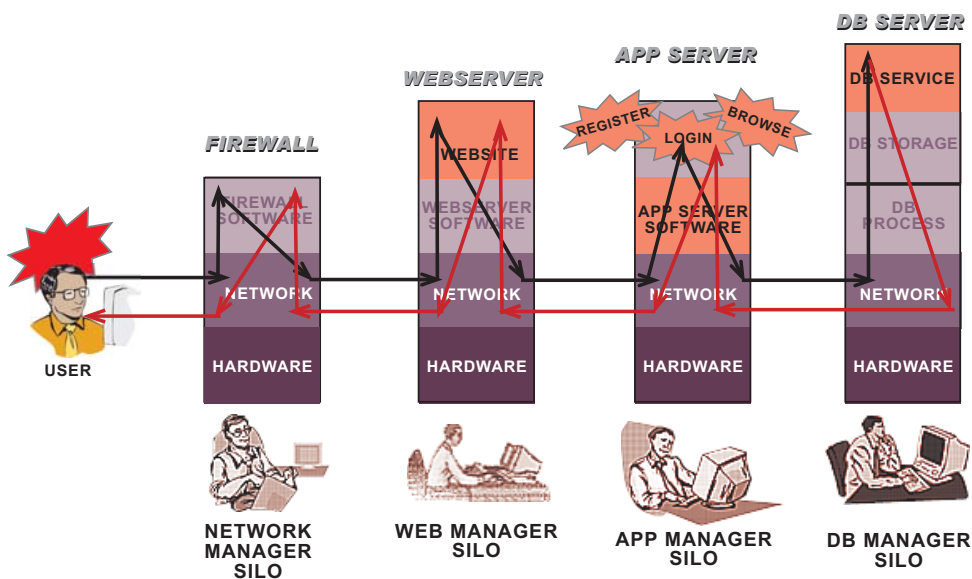


Figure 1: Silo based monitoring is no longer sufficient for managing multi-tier IT infrastructures

While the use of multi-tier architectures helps with respect to infrastructure scalability, it also poses interesting challenges for monitoring and management. For example, consider a user logging into a multi-tier web site (Figure 1). The user request is received by a web server, forwarded to the login application running on the middleware application server, which in turn accesses a backend database. If there is a problem with the database service (say, the database access is 50% slower than normal), it is likely that the login application will be affected, and that the web server will also be affected. In this example, a single problem has rippled and affected multiple infrastructure tiers resulting in a number of alarms - e.g., from the database server, application server, web server, etc. Since the end-to-end service involves multiple dependent application and network elements, a failure in one of the tiers (e.g., database) can affect the other tiers as well (e.g., web server, applications, etc). Consequently, problem identification and diagnosis in multi-tier infrastructures is a huge challenge.

Owing to the inter-dependencies between the different tiers of a multi-tier infrastructure, monitoring solutions that look at the target environment as a collection of distinct, independent silos are often not capable of assisting operators to quickly determine when and where problems originate. For example, in the example in Figure 1, a database problem could be impacting the performance of the web server and application server tiers. A monitoring solution based on the silo approach will indicate that problems exist in all the infrastructure tiers, but will not be capable of differentiating what and where the source of the problem is.

Service Monitoring - The Need

To further compound matters, as IT infrastructures have grown in complexity, it has also become impossible for a single operator/administrator to be responsible for the entire infrastructure. Typically, the maintenance team for a large infrastructure comprises of application developers - those who develop the applications, and domain experts like the WebLogic administrator, network administrator, database administrator, etc. While the application developers and domain experts are responsible for putting together the infrastructure, yet another group, the service managers are responsible for the 24*7 operation and performance of the end-to-end service.

Since a service involves multiple, heterogeneous applications and network devices, it is not reasonable to expect that a service manager has expertise in all the domains involved in the service (web, network, database, application, etc.). While there are a variety of tools available for managing specific applications and network devices in-depth, these tools are mainly appropriate for the domain experts and application developers, who are interested in optimizing the performance of the infrastructure components under their control, and in advanced troubleshooting and diagnosis - e.g., which SQL query is consuming too many resources, which java component is leaking memory, etc. On the other hand, the service managers are primarily interested in keeping the service running with good quality of service. Their interest is mainly in determining when a problem happens, which domain is the cause of the problem - is it the network? is it the server? is it the database?. By knowing this, a service manager can quickly determine when a problem happens which domain expert/application developer to hold responsible for solving a problem.

The term "triage" refers to the process by which a service manager can rank the current status of an IT infrastructure in importance and priority, and sort them based on their need for immediate action. In today's environment, the infrastructure triage process is very cumbersome and time-consuming. When a problem is reported, the service manager has to bring together all the domain experts and application developers to review the problem report and analyze which domain(s) could be causing the problem. The fact that the domain expert/application developer could each be using disparate tool sets, with widely different user interfaces makes the triage process extremely complicated and time-consuming (see Figure 2). A rule of thumb is that it takes eight hours on an average to find out the cause of a problem, and that over 80% of the time to repair is actually spent in problem isolation.

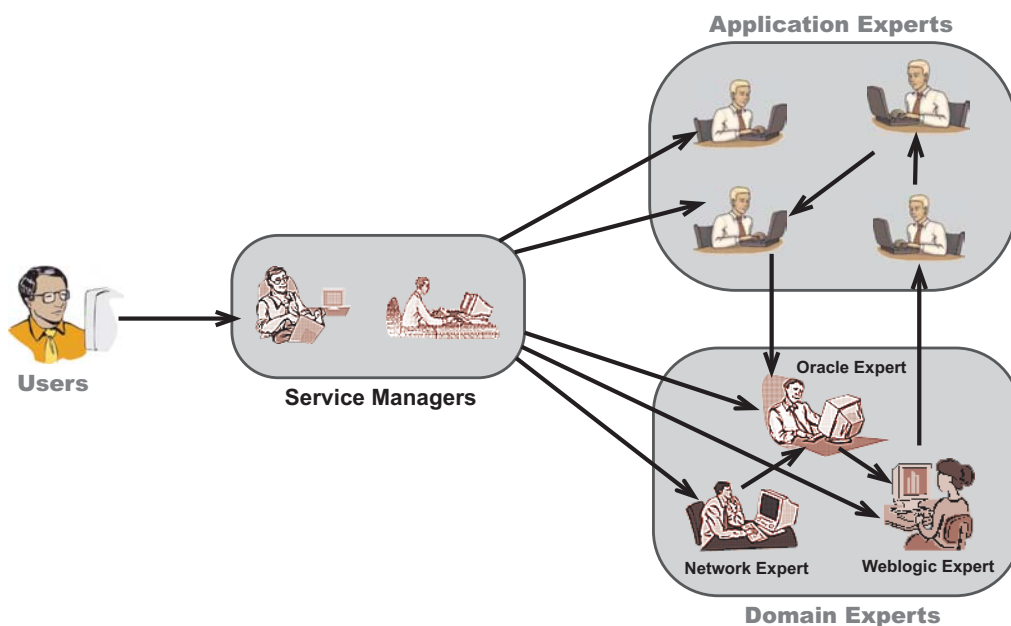


Figure 2: Infrastructure triage is often cumbersome and time-consuming



In order to effectively maintain and manage IT infrastructures, service managers require monitoring and management solutions that can enable them to determine the following in real-time:

- **How is the service performing?**

The service managers should be able to obtain service quality reports that can quantify the performance being delivered to users of the infrastructure services;

- **If there is a problem, which domain is the cause of the problem - is it the network? server? database? application?**

It is critical for the service manager to be able to quickly pin-point when a problem happens, which domain could be the cause of the problem, and what the potential problem could be. Depending on the nature of the problem, the service managers themselves should be able to correct simple/often recurring problem situations. For more complex issues, the problem reports provided to the service managers should allow them to quickly hand over the problems to the appropriate domain expert or application developer for immediate troubleshooting and resolution.

Keeping in mind that the service manager may not have the expertise or the time to sift through tons of data, the monitoring solution should be simple to use, and effective – i.e., enable the service manager to perform his/her tasks without needing to spend a lot of time and effort.

- **Where are the potential bottlenecks in service delivery and how can the service performance be optimized?**

Even when the service is performing as expected, there may be periodic trends in service usage that could point to potential future problem situations. The ideal monitoring solution for a service manager will provide proactive indicators of system bottlenecks that if corrected in advance could avoid future performance bottlenecks.

Proactive Infrastructure Triage™ using the eG suite

The eG suite is a comprehensive real-time monitoring and proactive infrastructure triage solution that addresses the key requirements of IT infrastructure service managers. Figure 3 illustrates how the eG solution operates. While network monitoring solutions focus on the network elements alone, and silo-based application monitors focus on individual applications, the eG suite takes a holistic view of the entire IT infrastructure. Taking the end-user perspective, the eG suite tracks the service performance in terms of availability, response times, and usage. To complement the service level monitoring (which can reveal potential problems with the service) and to further triage a service problem, the eG suite tracks the health of the individual IT infrastructure components including network devices, servers, applications, etc. Specialized monitors for over fifty popular application platforms, support for most common Microsoft Windows and Unix server operating environments, and coverage of basic network monitoring requirements, ensures that the eG suite provides comprehensive insights into an IT infrastructure's performance.

The eG suite is targeted primarily at the service managers. To understand the function of the eG service manager which is the central component of the eG architecture, let us draw an analogy to the function of a general physician handling a medical complaint from a patient. Most of the time, the general physician is the first point of contact for the patient. In many cases, the general physician him/herself is able to deduce where the problem is and prescribe a remedy. In more complex cases, the general physician directs the patient to an expert (e.g., dentist, eye specialist, neurologist, etc.) who will be able to correct the problem. The eG service manager performs a similar function for IT infrastructures. A service manager can use the eG manager to determine how the infrastructure services are performing. When a problem is detected, the eG manager provides the next level of detailed diagnosis. In a majority of cases, using this information, the service manager can proceed to fix the problem. In cases where they do not have the necessary access or where additional troubleshooting/expertise is necessary, the service manager can forward the problem on to the appropriate domain expert. Through its infrastructure triage capability, the eG manager helps a service manager determine which domain(s) is the cause of the problem, using which the service manager can forward the problem on to the appropriate domain expert (Figure 3).

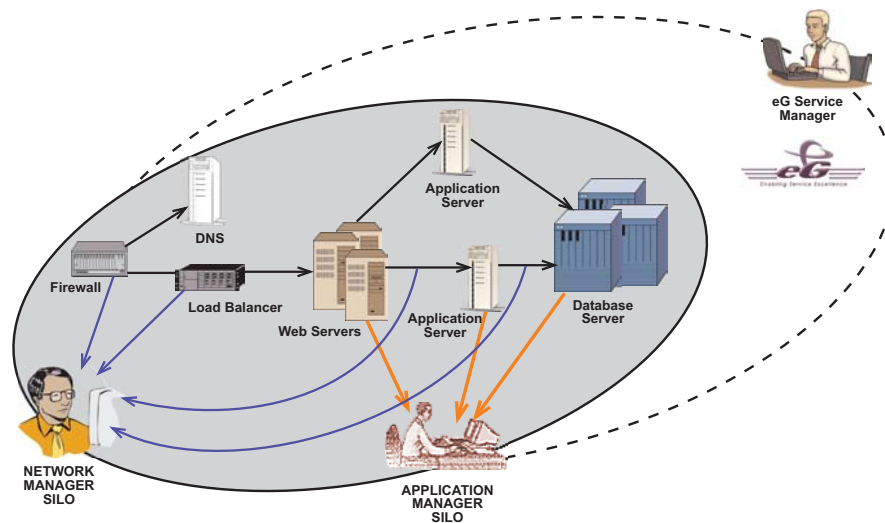


Figure 3: Automatic infrastructure triage with the eG suite eliminates finger – pointing

Figure 4 summarizes the benefits that accrue from using the eG suite for IT service management:

- Service managers can greatly benefit from the infrastructure triage capabilities – they can quickly figure out which expert to contact in the event of a problem.
- By reducing the finger-pointing between domain experts and application developers, the eG suite ensures that problems are resolved faster – well before users can notice them, thus enabling higher service availability and improved user satisfaction.
- With the eG suite in place, service managers too have to spend less time in problem resolution, and hence, they can focus their time and energies on more productive activities.
- By ensuring that the domain experts and application developers are involved in troubleshooting problems that are only relevant to their areas of expertise and responsibility, the eG suite ensures that these experts are efficiently used.
- Since it provides a 100% web-based interface, the eG suite facilitates collaborative management in multi-domain environments – for example, in a hosted environment, the service manager is responsible for the network and server infrastructure, but the application layer is the responsibility of the user. In such situations, both the service manager and users can access the eG manager and obtain a consistent view of the status of their infrastructure. Since it clearly quantifies the performance across the different tiers and layers of the infrastructure, the eG suite enables service managers and users to quickly figure out whether a problem relates to the user domain or the service provider domain. This powerful capability can ensure that users can monitor their own applications and they need not even call the service manager in the event that a problem is being caused by a fault in their application(s). The consequent reduction in support calls to the service manager can result in a significant cost saving in such multi-domain environments.

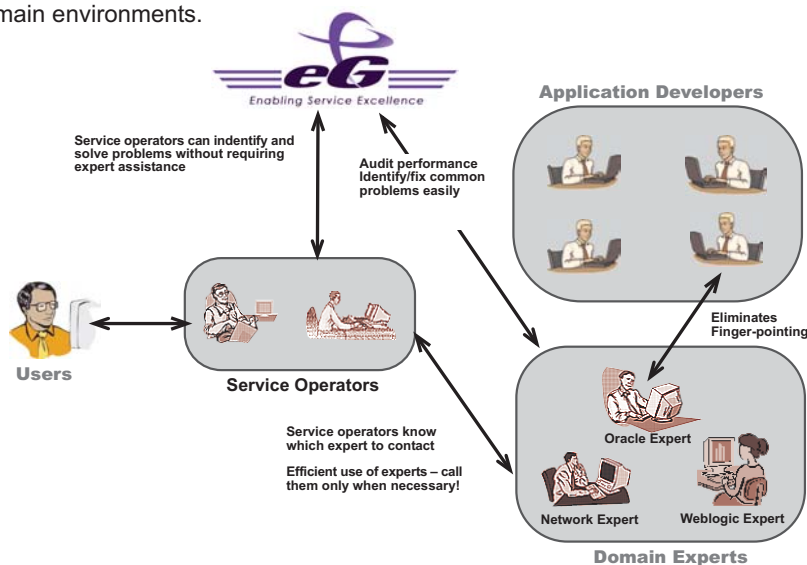


Figure 4: Benefits of using the eG suite for IT service management

A service monitoring solution like the eG suite is intended to augment current monitoring and maintenance practices, not radically change them. For instance, domain experts will still need to use silo-based expert tools like network sniffers, database tuning tools, source code optimization solutions, etc., for fine-tuning the performance of the infrastructure components they control. By providing a high degree of visibility for service managers into the functioning of the different domains of an IT infrastructure, the service monitoring solution enables more effective streamlining and efficient operation of the IT infrastructure.

The eG Difference

Having highlighted the need for a service monitoring and infrastructure triage solution in IT infrastructures, in this section, we will focus on what makes the eG suite the preferred solution for most service managers. The key characteristics of the eG suite and how they benefit customers are discussed below:

- **Scalable, 100% WEB-BASED architecture:**

Although it uses the conventional manager/agent architecture that is widely used by most management systems, the eG suite is unique in its use of web technologies. The eG architecture itself is built along the lines of multi-tier web architectures and hence supports small and large IT infrastructures equally well. All communications between the manager and agents use HTTP/HTTPS. The key advantage of this approach is that it permits the manager and agents to be in different physical locations, possibly separated by multiple demilitarized zones. In fact, the agents can even reside within private Intranets and still be managed by an eG manager in a central location. This architecture is ideally suited for large enterprises and managed service provider environments. Many IT infrastructures have virtual private networks deployed between the managed environment and the network operations center simply to allow secure access to the monitored servers. By innovatively using the web protocols (HTTP/HTTPS) and agent polling technology, eG's 100% web based architecture offers an easy to deploy solution at a much lower-cost for monitoring and managing your IT infrastructures across geographically disparate networks.

- **Single agent technology:**

As alluded to earlier, the eG suite includes extensive monitoring capabilities for networks, servers, and applications. The table below (Figure 5) summarizes the variety of IT infrastructure components monitored by the eG agents.

Component Type	Component Brand
Operating systems	Windows NT, 2000, 2003 server, AIX, HPUX, Red Hat Linux, Solaris (SNMP-based support for Novell Netware and other operating systems)
Web servers	Apache, iPlanet/SunONE, Microsoft IIS, IBM HTTP Server, Oracle HTTP Server
Web application servers	WebLogic, ColdFusion, ATG, iPlanet/SunONE, Microsoft transaction server, WebSphere, SilverStream, JRun, Tomcat, Oracle 9i OC4J, Oracle Forms Servers, Borland Enterprise Server
Database servers	Oracle, Microsoft SQL server, DB2 UDB, Sybase, MySQL
Network devices	Cisco routers, Cisco Catalyst switches, Baystack hub, Local director, any MIB-II compliant device
Microsoft applications	Active Directory, BizTalk server, Windows Internet Name Service (WINS), Domain controller, FTP server, DNS server, DHCP server, Print server, Proxy Server, File server, Event logs
Firewalls	Check Point Firewall-1, Cisco PIX
Thin-client servers	Citrix MetaFrame, Microsoft Terminal server
Email servers	Microsoft Exchange, Lotus Domino R5, SunONE/iPlanet messaging server
Messaging servers	MSMQ, WebSphere MQ, FioranoMQ
Others	Tuxedo domain servers, Network printers, NetApp filers and NetCache, Novell Groupwise

Figure 5: IT infrastructure components monitored by the eG suite

Most silo-based monitoring solutions require one agent module per application that is monitored. With such a model, separate agent licenses need to be purchased depending on various parameters like the deployment platform, the types and number of applications monitored, the number of CPUs, etc. In contrast, the eG suite offers a powerful single agent licensing policy for its agents. As per this policy, a single eG agent can monitor all the applications executing on a server Assuming one IP address per server. . Moreover, agent licenses are not tied to operating systems or node-locked, thereby allowing operators to pick and choose where they want to deploy the agents, and to even dynamically change the location of the agents. Furthermore, the agent licensing is also not tied to the hardware capabilities of the server being monitored by the agent. Its simple and cost-effective agent licensing model makes the eG suite an attractive solution for IT infrastructure monitoring.

- **Real-time, PROACTIVE MONITORING of the TRUE end-user experience:**

The experience that an IT infrastructure offers to its users is governed predominantly by how well its application components perform. Many monitoring tools use emulated requests to monitor web transactions. The drawbacks of such emulation-only techniques are:

- ▶ This approach cannot be used to monitor critical transactions such as payment, registration, etc.
- ▶ Moreover, since they merely sample the functioning of the target environment, these emulation techniques typically detect and report problems only when they are severe enough to impact the end user performance, i.e., they are useful mainly for reactive monitoring.

In order to avoid the drawbacks of the emulation-only approach, eG agents deploy a proprietary web-adaptor technology that enhances vanilla web servers with the capability to track and report various metrics relating to individual web sites and even web transactions in real-time. The monitoring is done in an implementation-independent manner, as a result of which eG agents are able to monitor Java (Servlets, EJB, JSPs) and other non-Java implementations (ASP, PHP, CGI, etc.) with equal felicity. Since it is able to monitor real-user transactions to web servers in real-time, eG’s web adaptor technology enables the agents to proactively monitor and quantify all anomalies that may occur in an IT infrastructure. Figure 6 below shows real-time monitoring of web transactions for a web site. In this example, the user registration transaction is experiencing a problem, and the percent errors is 100%, implying that users are not able to register via the web site.

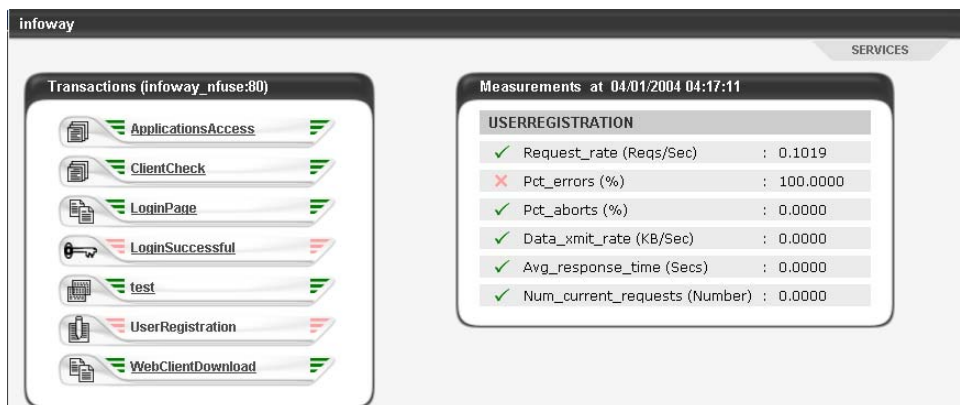


Figure 6: Tracking real user transactions to a web site

- **Automatic thresholding**

All the metrics collected by an agent are subjected to thresholding – i.e., comparing their values with pre-defined upper or lower bounds to determine if there is any abnormality. Many monitoring solutions require administrators to specify the thresholds for every measurement. Explicitly configuring thresholds for each and every metric being collected can be a laborious process – spanning days or even months for large IT infrastructures. To simplify the configuration process for in an IT infrastructure, the eG suite includes a unique automatic threshold computation capability. In this approach, the eG manager computes the thresholds to be used dynamically, using tried and tested statistical quality control techniques to analyze past values of the metrics and to automatically set the upper and lower bounds for each of the metrics, using the historical data. Since the values of the metrics vary from time to time, the historical thresholds are also time-varying. This ensures fast and easy setup of the eG system as administrators do not have to configure thresholds for each and every metric (Figure 7).

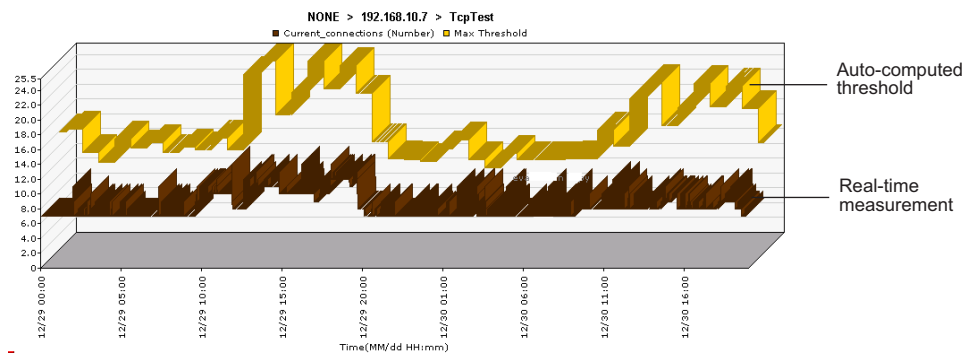


Figure 7: Automatic Thresholding capability of the eG suite

- **Automatic infrastructure triage:**

To ensure that IT infrastructures operate with minimum downtime, it is critical to perform problem detection and diagnosis instantly and accurately. Correlation of various problems reported at the network, system, and application layers is critical for speedy and accurate problem diagnosis. Most application monitoring solutions do not include any specialized correlation capability – manual analysis of the collected data is essential to determine the root-cause of problems. In contrast, the eG suite uses a novel, patented correlation and automatic infrastructure triage technology. To implement this capability, the eG manager incorporates a series of heuristics that take into account the configured site topologies and pre-built models of different network and application components. By automatically correlating across the network, system, and application layers, the eG suite is able to accurately identify and report the root-cause of problems. For the example in Figure 6, where different transactions of a web site are failing, Figure 7a depicts the service topology – i.e., the data flow/dependency between the different applications and network components involved in providing this service. The color coding in the figure denotes the current status of all the applications/network components involved in delivering the service. It is obvious from the color coding that although there are many components that are experiencing problems, the root-cause of the problem is the Oracle database server.

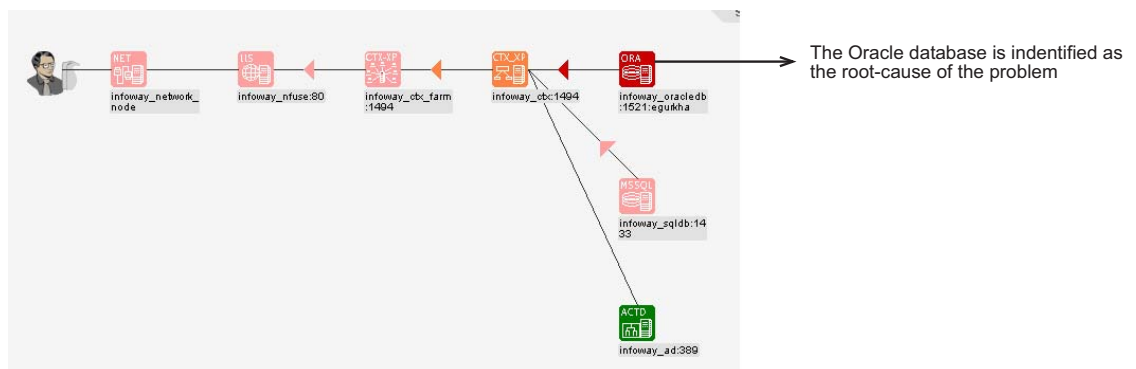


Figure 8a: Topology representation of a problem

Further drilldown into the Oracle database server reveals that the real problem is due to one of the tablespaces having run out of the allocated space (see Figure 8b). Due to its ability to represent the service inter-dependencies as a service topology graph, and its ability to model the different applications as a set of hierarchical layer, the eG manager is even able to automatically analyze the current state of the infrastructure components and provide automatic analysis that pin-points to the root-cause of problems. In Figure 8c, the eG alarm window clearly pin-points that the root-cause of the service failure in Figure 6 is the Oracle database tablespace issue. This prioritized information is made available via SMS, email, over the web, or via SNMP traps, thereby ensuring that service managers can quickly triage their IT infrastructure from any where, at any time.

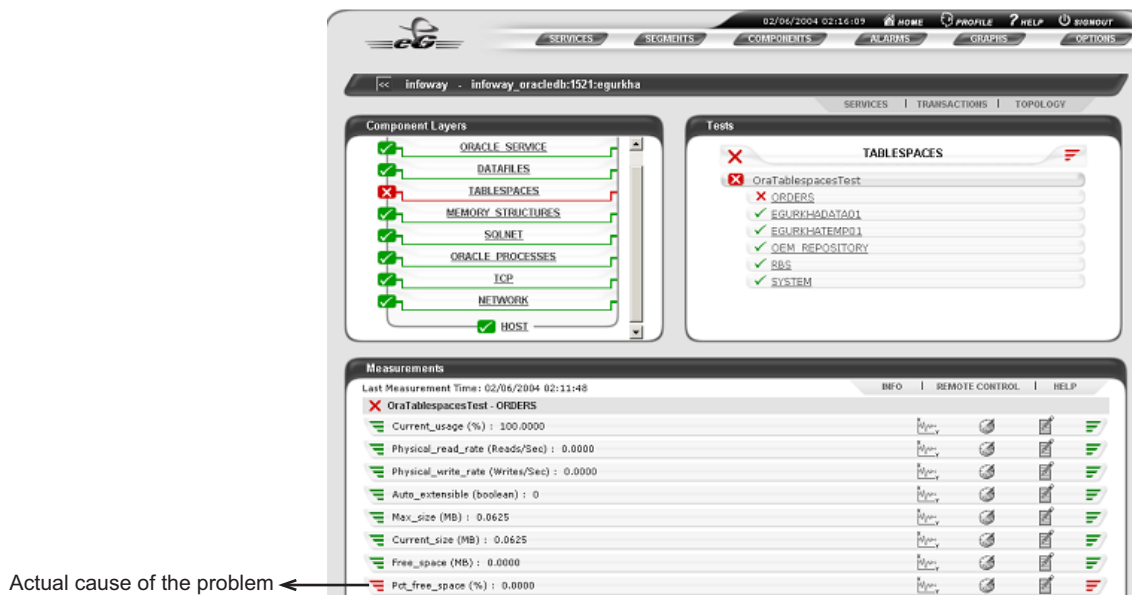


Figure 8b: The layer model showing that the Tablespaces layer in the Oracle database is the cause of the problem



Figure 8c: eG's alarm window showing automatic prioritization of alerts

● Simple and Fast Provisioning

Since only one agent needs to be installed per server, thresholds can be auto-determined, pre-defined models determine what metrics need to be collected by each agent (depending on what applications are monitored by the agent), the deployment of the eG suite is done very rapidly. A browser based interface ensures a near zero learning curve for users. Moreover, since eG's auto-triage technology does not involve setting up elaborate correlation rules and circuits, users can get the eG system up and running in a matter of hours, not weeks or months. Moreover, since the eG agents are auto-upgradable from the central manager, elaborate reconfiguration/reinstallation of the software is not necessary when new versions are released or support for new IT infrastructure components is introduced.

- **Real-Time and Post-Facto Analysis**

Besides laying extensive emphasis on real-time monitoring and troubleshooting, the eG suite also includes elaborate reporting and analysis capabilities for off-line analysis and proactive capacity planning. Different types of reports can be generated for the different levels of management in an organization. Operations reports provide in-depth insights across network, system, and applications thereby providing clear indicators of performance bottlenecks and trends that could be affecting infrastructure performance. Users have the flexibility to customize the operation reports to suit their individual needs and preferences (Figure 9).

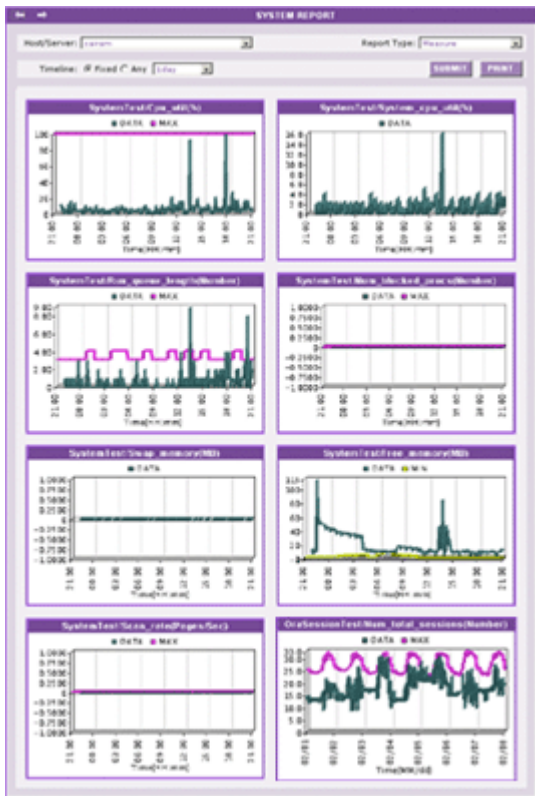


Figure 9: An operations report showing critical system metrics of all

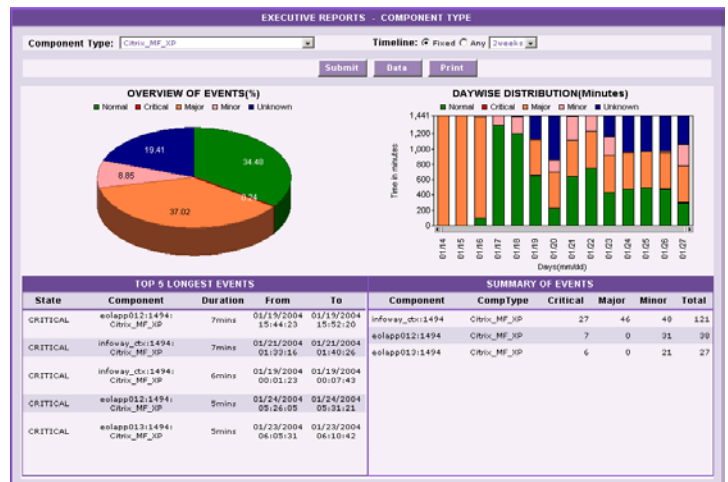


Figure 10: An executive Report summarizing the performance Citrix Metaframe servers in an IT infrastructure

Executive reports for management executives (Figure 10) offers comprehensive health reports that summarize the overall state of each of the infrastructure components. By reviewing a report of a server's health, an executive can determine what percentage of the time was the server's operation trouble-free. By comparing the performance reports of the different components, executives can quickly determine where the problem-prone areas of their infrastructure are. Comparison of performance across time periods can also provide indications of whether the infrastructure performance is improving over time.

Summary

This whitepaper has outlined how the eG suite makes IT infrastructure monitoring and triage easy, effective, and efficient. IT administrators and service managers can use the eG suite at all stages of the software lifecycle. In the development phase, the eG suite can be used in conjunction with load/stress testing tools and helps fine-tune application performance. In the deployment stage, the eG suite is used to ensure that the developed applications are meeting the performance expected of them. In the maintenance stage, the eG suite ensures that the IT infrastructure and its services are meeting the service levels of expected of them through its 24*7 monitoring and instantaneous troubleshooting capabilities.

eG in the software Lifecycle

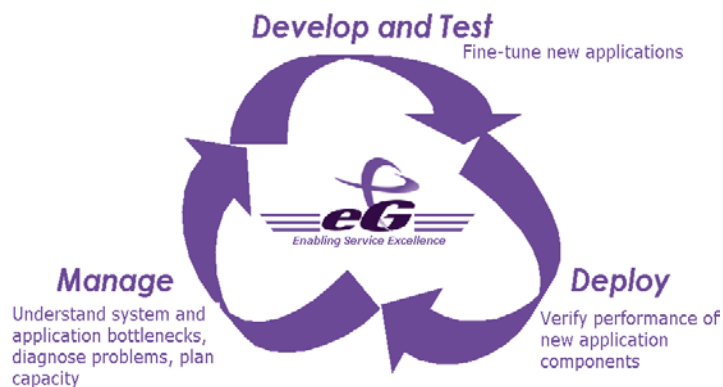


Figure 11: How the eG suite helps at different stages of the software lifecycle

About eG Innovations

eG Innovations is the leading provider of enterprise-class monitoring and management solutions for IT Infrastructure. The company's 100% web-based monitoring solutions are especially suited for mission-critical infrastructures where proactive monitoring, rapid diagnosis, and instant recovery are critical. Customers worldwide use the eG solutions to improve the quality of their services thus increasing their competitive positioning, lowering their operational costs, and optimizing the usage of their infrastructures.

For More Information

eG Innovations, Inc

33, Wood Ave, South, Suite 600, Iselin, New Jersey 08830. USA

Ph: (866) 526 6700

Email : info@eginnovations.com

Web : www.eginnovations.com