



**Virtualization 2.0 is all about Manageability.**

**What You Should Look for in a**



**Monitoring Solution**

## **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 eG Innovations, Inc. All rights reserved. All trademarks, marked and not marked, are the property of their respective owners. Specifications subject to change without notice.

## Introduction

The first phase of virtualization, Virtualization 1.0, involved the use of virtual infrastructures primarily in staging and development environments. The emphasis during this phase was on making sure that virtualization provided many of the promised benefits including space consolidation, power savings, easy configuration and deployment, etc. From a performance standpoint, the emphasis was on ensuring that the virtualized infrastructure delivered performance in line with that obtainable from a purely physical infrastructure. Often, this was achieved by over-provisioning the virtualized servers. During this phase, the overall focus was on functionality more than performance, and in enabling newer service delivery opportunities such as the use of virtualization to support remote desktop applications.

The second phase of virtualization – Virtualization 2.0 – is already upon us. Administrators now have a choice of virtualization technologies and the much-researched hypervisor is now almost a commodity. Faced with shrinking budgets, administrators are looking for ways to achieve the maximum with limited hardware and software resources, through optimal resource allocation techniques and to plan proactively for future demands. Over provisioning of virtual infrastructures is therefore a thing of the past!



***“The market share war in the virtualization space will be fought at the management layer rather than at the hypervisor layer.”***

*Virtualization Management – The Battlefield,  
Virtual Strategy Magazine, July 2008*

Virtual infrastructures are becoming increasingly prevalent in production environments and they are being used to support critical business services. The emphasis is shifting from functionality to the core monitoring and management challenges that physical IT infrastructures have tackled for decades. Virtual infrastructures have promised high availability and reliability, and they now have to ensure minimal downtime and peak performance for critical business services that they support. The challenge here is that there are various layers of software – the applications, the protocol layers, the operating systems in the virtual machines (VMs) and the virtualization platform – that have to work together to ensure the proper functioning of the business service. Many of these software layers are outside the scope of the virtual infrastructure itself and knowing when a problem happens, whether it is being caused in the virtual infrastructure, or in the applications, or in the network will be crucial. The faster the problem can be diagnosed, the shorter the service downtime and better the overall service performance.

This combination of factors means that proactive monitoring and effective root-cause diagnosis across the entire infrastructure will gain in prominence in Virtualization 2.0. The acute shortage of

IT professionals skilled in working with virtualized environments means that a management solution must offer superior automation and root-cause diagnosis to enable administrators with limited expertise to be effective in spotting problems and taking the proper corrective action quickly. In this document, we define the key characteristics that a *Virtualization 2.0 Ready* monitoring solution should have, and we highlight how the eG Enterprise Suite from eG Innovations addresses all of these key characteristics.

## Virtualization 2.0 Ready Monitoring

The emphasis of monitoring and management in Virtualization 2.0 is shifting from virtual machine (VM) management to business service management, i.e., knowing how a business service is performing and which domains (network, server, VM, applications) are working and which are not. Hence it is no longer sufficient to just monitor the uptime or resource usage levels of virtual machines and physical servers\* and believe that the entire IT infrastructure is working well.

A *Virtualization 2.0 Ready* monitoring and management solution should be able to:

- Provide a single view of virtual and physical infrastructures
- Support multiple virtualization technologies
- Track physical resource availability, configuration and usage by VMs
- Provide an inside view of virtual machines with clear problem identification
- Automatically establish performance baselines and norms
- Perform automatic correlation for true root-cause diagnosis
- Scale as the infrastructure monitored grows
- Support for virtualized desktop environments
- Offer personalized views for the various stakeholders in an organization to enable collaborative management

The rationale for why these are key requirements for any *Virtualization 2.0 Ready* monitoring solution is provided below:

- **Provide a single view of virtual and physical infrastructures** – Even though virtual infrastructures are being used for many mission-critical applications, most enterprises are moving to virtualization only in a phased manner. For example, I/O intensive applications are still being hosted on non-virtual servers. Therefore, a business service may involve some applications that reside on physical machines and others that run on virtual machines. To provide an integrated view of the target infrastructure, the monitoring and management system needs to be able to manage infrastructures with a set of virtual and physical machines equally well, providing a single integrated interface across these differing technologies.
- **Support multiple virtualization technologies** – Administrators now have a choice of virtualization technologies based on their business needs and preferences. VMware® ESX, Citrix Xen, Microsoft Hyper-V, as well as different Unix options (Solaris Containers, AIX LPARs) all offer robust solutions for virtualization. Most large infrastructures will include a mix of these virtualization technologies, and it is important to have a single unified dashboard from where these different virtualization technologies can be monitored.

\* We use the term *virtual machine* to refer to a virtualized operating system. The term *physical machine* refers to an operating system running on a physical computer. The term *physical server* is used in this whitepaper to refer to a server with a hypervisor that is capable of hosting multiple virtual machines.

- **Track physical resource availability, configuration and usage by VMs** -- As deployment of virtual infrastructures proliferates, it is essential that administrators have a comprehensive view of the virtual infrastructure. While monitors designed for conventional physical machines can be installed and used on individual VMs, they have no specialized capabilities for virtualized environments. Knowing such things as how the hypervisor is performing, which VMs are powered on and what resources they are using, if the physical server has sufficient resources to handle its workload, whether the VMs are configured with sufficient resources, etc., are critical requirements that only a monitoring solution that is specialized for virtual infrastructures can deliver.


Many virtualization platforms support high availability and live migration configurations to provide reliability and failover for mission-critical applications. Administrators need to know whether these capabilities are working properly or whether any configurations need to be tuned (e.g., Are migrations happening too often? Why did a migration suddenly take place?).

- **Provide an inside view of virtual machines with clear problem identification** -- While most virtualization administrators understand the importance of tracking the resource usage levels (CPU, memory, disk, network) of each of the VMs on a physical server, very few can monitor what is going on within each virtual machine. This is because most Virtualization 1.0 monitoring solutions focused on capacity planning and provisioning. For capacity planning and provisioning, it is important to track the portion of a physical server's resources that each VM is taking up. This view, which provides insight into how a physical server's resources are used across all its VMs, is the "outside view" of a VM. While the outside view helps identify a resource-hungry VM, it falls short of providing additional information that is critical for problem diagnosis and further optimization. For instance, why is a specific VM taking up excessive resources? Is it because of a heavy workload? Or is it due to a malfunctioning application (e.g., a run-away job or a memory leak in one of the applications running in the VM)? To provide this information, an "inside" view of each VM is necessary. This view tracks such dynamics as end user activity, resource allocation for each application and the application mix running inside the VM guest operating system.

As virtualization goes mainstream, it will no longer be sufficient to just plan and provision virtual infrastructures correctly. Production environments are dynamic, and when problems occur it is important to determine what is causing the problem. Is the physical server running out of capacity? Is it a VM not having sufficient resources because it was not correctly provisioned? Is it a malfunctioning application inside the VM? The answer to these questions will determine who is responsible for fixing a problem – is it the VM administrator, or is it the application administrator/expert? Only a monitoring solution capable of presenting both the "outside" and "inside" views can provide this richness of information.

- **Automatically establish performance baselines and norms** -- Often, the emphasis of monitoring is just problem diagnosis. When there is a problem, administrators want to know what is wrong. While problem detection is easy (if your monitoring system does not alert you, your users will!), isolating the problem and determining what the true root-cause of a problem is can be a challenge. Establishing performance baselines and norms was important in the non-virtual world. This is even more important in a virtualized world, since the number of "moving parts" is much higher (hypervisor, VMs, applications, migration, etc).

Understanding what has changed and when is critical to quickly zooming in on the root-cause of a problem. The ability to establish these norms automatically is important in many ways. Administrators do not always know what is "normal" in their environment. The norm also varies from one server to another based on its sizing. Experts who understand what norms need to be adopted are few and not readily available. Even for such experts, setting norms for each and every server can be an arduous task. Hence, it is important to have the right automation built into the monitoring system to automatically determine what the performance baselines are for the infrastructure.



***"Infrastructure virtualization management is moving to the forefront. The market is shifting from a focus on consolidated servers to distributed virtualization management. The value proposition is also changing, from saving hardware money to increasing agility."***

*Gartner (Thomas Bittman & Philip Dawson)*

*"Server Virtualization Trends, Best Practices, and the Future", Nov 2007*

This capability is also key to being able to monitor your infrastructure proactively. The monitoring solution should be able to compare current performance with respect to the baseline and be able to generate alerts well before a failure happens. This provides administrators with precious advanced notice that can help them avert potentially serious failures in business service performance.

- **Perform automatic correlation for true root-cause diagnosis** -- While auto-baselining can provide proactive alerts, analyzing these alerts and determining where the root-cause of a problem lies is a huge challenge. Effective root-cause diagnosis is critical to reducing the downtime of business services and enhancing operational efficiency (so expert staff spend less time fire fighting).

Root-cause diagnosis in a physical infrastructure is a huge challenge. The addition of virtualization just makes the problem even harder. To understand why, consider a business service supported by a typical multi-tier configuration of application

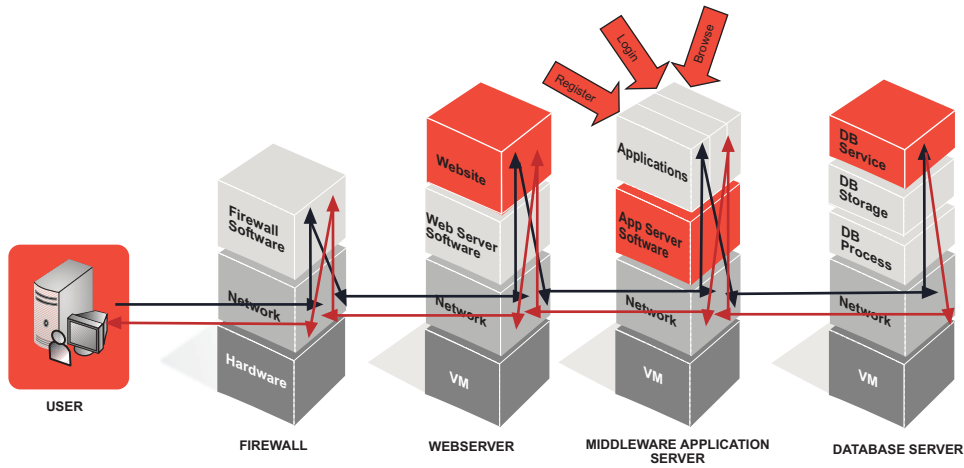


Figure 1: Problem diagnosis in a multi-tier infrastructure. A single business service involves multiple tiers of inter-dependent applications. Hence, a problem in one tier can impact all the other tiers. Root-cause diagnosis must account for these inter-dependencies.



**“As enterprises continue to embrace virtualization, IT staffers see a need for better management tools for this key new technology. 75% of IT managers call virtualization management important to their operations, according to a survey of 100 IT managers. Among them, 39% call virtualization management ‘very important’.”**

*Laurianne McLaughlin, CIO Magazine, Nov 2007*

tiers (Figure 1). In this example, the user accesses the service through a firewall. User requests are forwarded by the web server to a middleware application server. The application server performs the business logic, accessing a backend database to get the data for analysis. If the database server were to slow down suddenly by 50%, since the application server depends on the database for its functioning, the application server will become slower than normal. This in turn will result in the web server appearing to be slow and the end-user response will be poor. In this case, a problem in one application tier impacted all the other tiers that depend on it. Diagnosing a problem in a multi-tier architecture requires an understanding of the inter-dependencies that exist among applications in the underlying infrastructure, and then using these inter-dependencies to determine where the root-cause of a problem lies and what the effects are.

In this example, if the applications supporting the business service were running on physical machines, we would have concluded based on the scenario above that the database server was the root cause of the performance problem. However, the applications are in fact running in a virtual

infrastructure – i.e., inside VMs. In this example, suppose the Oracle database server is running on the same physical server as a Citrix application and a media server (Figure 2). A sudden increase in accesses to the media server can cause disk accesses on the physical server to increase to the point that disk access becomes a bottleneck.

At this stage, queries handled by the database server start to take longer and longer. Thus the database slowdown in Figure 1 may actually be caused by a sudden increase in workload to the media server in Figure 2. In this case, the root-cause of the problem is a disk bottleneck on the physical server caused by an increase in workload for the media server application.

From this example, it should be clear that root-cause diagnosis technologies for virtual environments need to go beyond how




Figure 2: Why root-cause diagnosis in a virtual infrastructure is even harder than in a physical infrastructure: Oracle, Citrix, and Media Server applications are hosted on VMs residing on the same physical server. A sudden surge in requests to the media server causes excessive disk reads on the physical server, thereby slowing down the performance of the Oracle database server!

they operate in a physical world. For true root-cause diagnosis, VMs running on each physical server must be auto-discovered. Applications running inside each of the VMs need to be detected and the monitoring system should automatically determine which applications coexist on the same physical server. This information is then used to determine where the root cause of a problem lies.

- **Scale as the infrastructure monitored grows --** As virtualization penetrates the enterprise, a large deployment will have hundreds of physical servers and thousands of VMs that require monitoring. In fact, as virtualization for desktops becomes popular, the ratio of VMs to physical servers could be as high as 30:1. The monitoring solution must be able to scale to handle such large infrastructures.

Virtualized Application Server Environments	Virtual Desktop Environments
Few VMs (<10) per physical server	30-40 VMs per physical server
VMs are mostly powered on all the time	VMs are powered on/off dynamically
Monitoring is mostly from the VM perspective - which VMs are on, what resources are they using	Monitoring is needed from the user perspective (who is logged in, what resources are they using)
In-depth application monitoring is required (Citrix, Oracle, etc.)	In-depth monitoring of the applications on the desktop is not required

Figure 3: Differences in monitoring requirements between virtualized application server environments and virtual desktop environments. A Virtualization 2.0 Ready monitoring solution should be able to handle these differing environments.

The extent of the automation determines the cost savings that the monitoring solution offers. Reduced downtime directly contributes to a business' top-line. Further, by pinpointing the root-cause of a problem, a monitoring solution can save endless hours of finger pointing that goes on in most IT organizations. This results in cost savings from enhanced operational efficiency and reducing the man hours spent in routine fire-fighting.

- **Support for virtualized desktop environments --** Virtual Desktop Infrastructure (VDI) is being viewed as a viable alternative to Citrix- and terminal server-based remote access technologies. For situations where each user requires his/her own desktop as opposed to shared access to an operating system (e.g., for software development or to run a legacy application), VDI is being viewed as the technology of choice for remote access.

	Virtualization 1.0	Virtualization 2.0
Monitoring of physical servers: the hypervisor, service console	✓	✓
Auto discovery of VMs and tracking of their up/down status	✓	✓
Outside view of the VMs: what physical resources is each VM taking up?	✓	✓
Detecting VM bottlenecks - CPU ready time, throttled time, balloon memory, disk latencies, SAN, etc.	✓	✓
Support for multiple, heterogeneous virtualization technologies from a single management console		✓
Inside view of VMs to understand how applications are consuming the resources taken up by the VM		✓
Scalable architecture to handle hundreds of physical servers, thousands of VMs		✓
Monitoring Live Migration and understanding when/why it happens		✓
Automatic baselining of performance and understanding norms		✓
Correlation between VM and physical server performance to understand where the bottleneck lies		✓
Automated root-cause diagnosis by correlating business service performance, network status, application performance, VM and physical server performance all from a single console		✓
Monitoring of the virtual desktop ecosystem - VMs, physical servers, connection brokers, datastores, terminal servers, etc.		✓
Personalized views for different stakeholders (domain experts) in an organization		✓

Figure 4: The changing monitoring and management requirements in a Virtualization 2.0 deployment.

Virtual desktop environments have different characteristics than environments where VMs are used to host server applications such as databases and web servers (*Figure 3*). VDI environments also have an ecosystem of new application technologies, such as the connection brokers, terminal access controllers, etc. A *Virtualization 2.0 Ready* monitoring solution should be capable of handling the diverse monitoring requirements of virtual server and virtual desktop environments.

- Offer personalized views for the various stakeholders in an organization to enable collaborative management**
  - Different stakeholders responsible for supporting a business service may need different views of the monitored infrastructure. Virtualization administrators, application experts, database admins, infrastructure architects, helpdesk personnel, and capacity planners may require different views of the infrastructure in keeping with their roles and responsibilities. The monitoring system must be flexible, providing each stakeholder with views that are aligned with their roles in the organization.

*Figure 4* summarizes the new considerations that IT organizations deploying virtualization in production need to take into account when determining their monitoring and management needs.

## Organizational Process Challenges in Virtualization 2.0

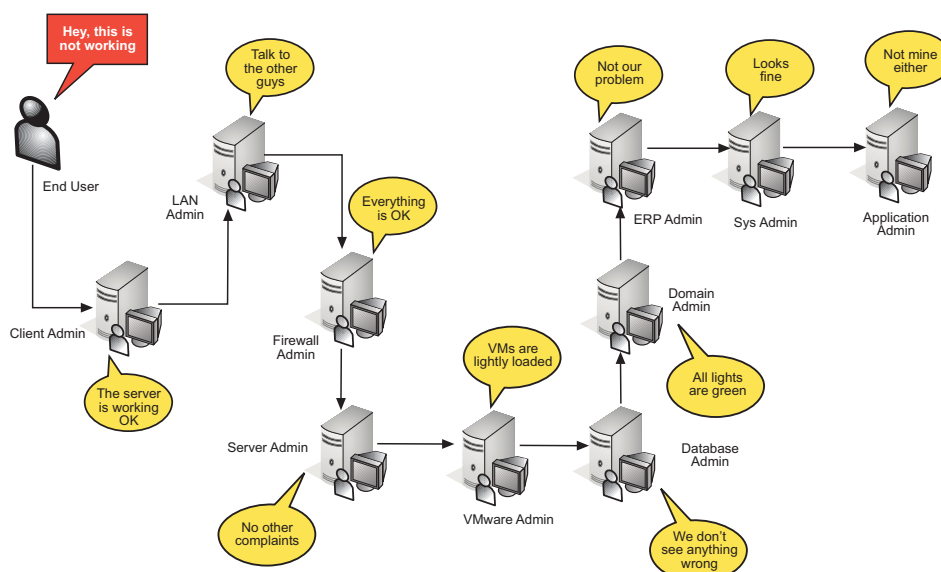
While the previous discussion focused on the monitoring requirements for Virtualization 2.0, it is equally important to understand that Virtualization 2.0 also impacts the core of most organization's operational processes.

Most organizations handle VM provisioning in much the same manner as physical server procurement has typically been done. Business units and application owners specify the sizing of the virtual machines they need, and the appropriate VMs are provisioned by the virtualization group that handles the physical servers on which the VMs are setup. However, the virtualization

group usually does not have any information or visibility into what applications are being hosted inside the VMs. When the physical servers are over provisioned and fewer VMs are executed in parallel, this siloed approach, wherein the virtualization group and the application teams do not interoperate, is sufficient.

But with Virtualization 2.0, organizations seek better return on investment for virtualization technologies, and deploy more complex applications inside virtual environments. Now it is no longer sufficient for the virtualization group to remain oblivious to the resource requirements of the application groups and their VMs. For instance, two memory-intensive applications hosted on the same physical server may contend for the same resources, thereby impacting each other's performance. Of course, by strictly partitioning the resource usage of each of the VMs, the virtualization group can offer performance guarantees. This has two key disadvantages. First, strict partitioning reduces the possibility of resource sharing across VMs, thereby limiting the consolidation benefits that virtualization offers. Second, due to limitations in the virtualization technologies, not all resources can be completely isolated across virtual machines; e.g., disk I/O. Hence, Virtualization 2.0 requires that virtualization groups of organizations play a more active role in how VMs are provisioned, including understanding what applications are to be hosted in each VM, what assumptions have been made regarding their workloads and resource requirements, and how the workload of different applications varies over time and with load. All of these details are essential for effectively balancing load and optimizing resource usage a virtual infrastructure. For example, by hosting a memory-intensive application and a CPU-intensive application on the same physical server, the virtualization group can make best use of the available resources (rather than by hosting all CPU intensive applications on the same physical server).

Yet another problem that virtualization administrators have to contend with under Virtualization 2.0 is finger-pointing and problem diagnosis (*Figure 5*). A single business service often spans multiple application and network tiers, so when a problem occurs it is unclear what caused the problem; i.e., is it the network? The application? The database? The server? In a virtualized



*Figure 5: An illustration of why monitoring an IT infrastructure as silos does not suffice. Finger-pointing across silo administrators takes endless hours, resulting in high downtime for the business service.*

infrastructure, there are additional possibilities for where the problem could lie: in a VM? In the physical server? In the hardware? In the virtual network interface? In the SAN?

Since most administrators already have silo tools for monitoring and management, there is no common dashboard from where the entire infrastructure can be monitored and diagnosed. Virtualization administrators will need to get accustomed to working in a multi-silo organization where finger-pointing is common. Monitoring and management solutions that provide deep visibility into every layer of every tier of the infrastructure and serve as a common dashboard for all the different administrators in an organization can go a long way in ensuring that Virtualization 2.0 environments operate properly.

## The eG Enterprise Solution for Virtualization 2.0

Many of the management challenges with virtualized environments are similar to those found in existing server-based computing infrastructures, such as Citrix and Microsoft Terminal Services. While Citrix and terminal services environments involved multiple users accessing a single operating system on a physical machine, virtual environments involve multiple operating systems sharing the resources of a physical server. In either case, a single malfunctioning application could impact the performance of all other applications sharing the common resources of the server.

The eG Enterprise Suite™ from eG Innovations has been deployed by enterprises worldwide for monitoring mission-critical IT infrastructures. The eG Monitor for VMware® Infrastructures™ (the eG VM Monitor™), part of the eG Enterprise Suite, provides capabilities essential to fulfilling the monitoring requirements of Virtualization 2.0 discussed above. These include:

- Integrated dashboard for virtual and physical infrastructure monitoring
- Tracking the performance of virtual and physical machines in a virtual infrastructure
- In-N-Out Monitoring™ of virtual infrastructures for deep diagnostics
- Monitoring every layer of every infrastructure tier
- Automatic baselining of the target infrastructure
- Problem demarcation and automatic root-cause diagnosis
- Scalability of the monitoring solution
- Virtual desktop monitoring and reporting
- Personalized role-based views for different stakeholders

Each of these capabilities of the eG Enterprise Suite is discussed in more detail below.

- **Integrated dashboard for virtual and physical infrastructure monitoring** -- With monitoring technology that is capable of handling virtual and non-virtual servers, eG Enterprise offers a single integrated dashboard to allow the entire infrastructure to be managed end-to-end. Agent-based monitoring is supported using a single agent technology that is licensed per operating system monitored, irrespective of the applications executing on the operating system. As the name implies, the licensing is independent of the operating system (i.e., Windows, Solaris, AIX, Linux, HPUX, etc.) and is also independent of

the hardware capabilities of the server being monitored (i.e., the same license works for a one processor or a four processor system). The single agent license allows unparalleled flexibility in the way the monitoring system is set up and used – for example, an administrator could be using the agent to monitor Oracle on Solaris at one time, and could later use the same agent license to monitor Microsoft SQL on Windows.

Agentless monitoring (using Windows Management Instrumentation – WMI, SNMP, or Secure Shell) is also supported and administrators have the flexibility to decide which servers and applications they would like to monitor with agents and which ones to monitor in an agentless manner.

Metrics about the virtual and physical infrastructure, reported by the agents are analyzed by a management console and real-time alerts and historical reports made available to administrators.

- **Tracking the performance of virtual and physical machines in a virtual infrastructure** -- As indicated earlier, agents designed for physical infrastructures cannot be directly used for virtual infrastructures. To obtain visibility into the performance of the physical machines of a virtual infrastructure and to track the behavior of VMs supported on these physical



*“The tools needed to manage virtual environments are going to have to quickly evolve to deal with what amounts to a reinvention of enterprise computing. So even if the established systems management vendors can add support for virtual machines today there is some doubt about how quickly they will be able to keep up with a rapidly changing enterprise computing landscape.”*

*eWeek System Management Blog,*

*“Virtualization Management War Begins in Earnest”, June 2008*

machines, the monitoring solution must communicate with and extract key metrics from the hypervisor. The eG VM Monitor enhances eG Enterprise with the ability to monitor virtual infrastructures. Using APIs and command line interfaces supported by hypervisors for VMware ESX, Citrix XenServer, Solaris Logical Domains (LDOMS), Solaris Containers, etc., eG agents and agentless monitors extract metrics that indicate the usage levels of the physical server and the physical resources that each of the VMs consume. Thus, eG Enterprise offers a solution capable of handling the heterogeneous hypervisor options that have become available in Virtualization 2.0.

- In-N-Out Monitoring™ of virtual infrastructures for deep diagnostics** -- Recognizing that Virtualization 2.0 requires detailed drill downs inside the VMs, eG Enterprise incorporates a patent-pending In-N-Out Monitoring technology (Figure 6). A virtualization-aware eG agent deployed on a physical server (e.g., VMware ESX) can be used to monitor the VM kernel, the service console (if appropriate) and the individual VMs. Agentless monitoring is also supported as an option for heavy-weight hypervisors like VMware ESX, while it is mandatory for light-weight hypervisors that do not have the service console, like VMware® ESXi servers.




Figure 6: In-N-Out Monitoring technology: A single eG agent can obtain an outside view of a VM as well as an inside view of each VM.

This “outside” view of a VM indicates relative usage of physical server resources (CPU, memory, disk, network) for each VM. As discussed earlier, while the outside view of a VM is useful in determining which VMs are resource hogs, this view does not provide in-depth insights needed for further diagnosis to determine which applications are consuming the resources. To complement the outside view, eG Enterprise provides an “inside” view of a VM, which highlights the relative resource consumption levels of the applications running inside the VM. While the outside view indicates the portion of physical resources a VM consumes, the inside view reveals the relative usage levels for the applications running inside the VM. This inside view of a VM is critical for effective root-cause diagnosis.

As shown in Figure 6, eG Enterprise obtains the inside view using the same agent that monitors the outside view. While Figure 6 shows an example of agent-based monitoring, Figure 7 illustrates how eG Enterprise also can support monitoring of virtualized environments in an agentless manner.

Administrators can even choose which servers they wish to monitor with agents and which ones to be monitored without agents. In either case, agents are not required to be installed on the VMs to obtain the inside view.

VMs are known to suffer from clock skews and hence, metrics collected from within a VM have to be analyzed with care. Rather than use metrics collected from within a VM in absolute terms, eG Enterprise uses the inside view to compare the resource usage levels across applications running in the VMs. Other key indicators of VM and application bottlenecks such as memory leaks, handle leaks, disk partitions filling up, VMs being rebooted unexpectedly, etc., can also be obtained using the inside view of a VM.

There are several advantages in using a single agent to obtain the inside and outside views of a VM. Installing additional agents inside the VMs to get the inside view can be a cumbersome process. This is especially true in a virtual desktop environment, where tens of VMs could be running on the same physical server. Installing agents inside the VMs could be wasteful in physical resource usage, and would incur substantial additional licensing costs. The deployment time is also reduced since only a single agent needs to be installed per physical server to obtain the inside and the outside view.

For monitoring virtualized environments, the eG Enterprise suite is licensed per physical server, irrespective of the number of sockets or CPUs on the server, and the number of VMs hosted on it.



**Compounding the infrastructure management problem is the fact that an application running on one VM may be depending on another application running on another VM (e.g., a web server may depend on a backend database server) to support a business service.**

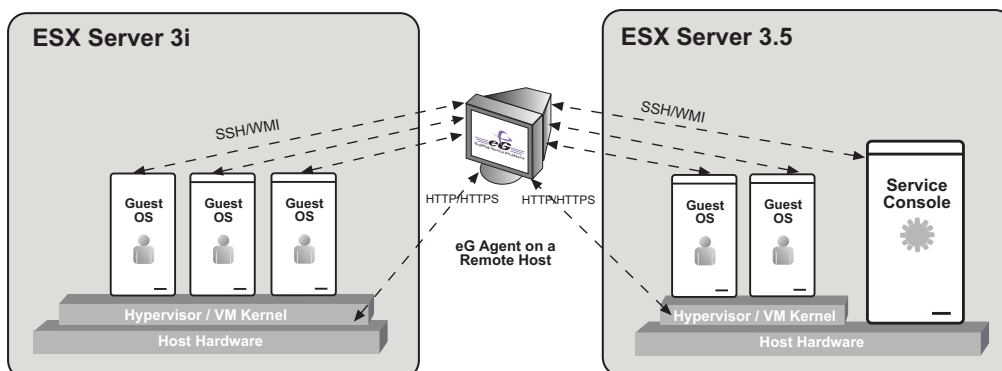


Figure 7: Agentless monitoring of virtual environments using eG Enterprise. The same agent can be used to monitor the virtual servers, the physical environment, as well as applications (Citrix, Oracle, SQL, Web servers, etc.).

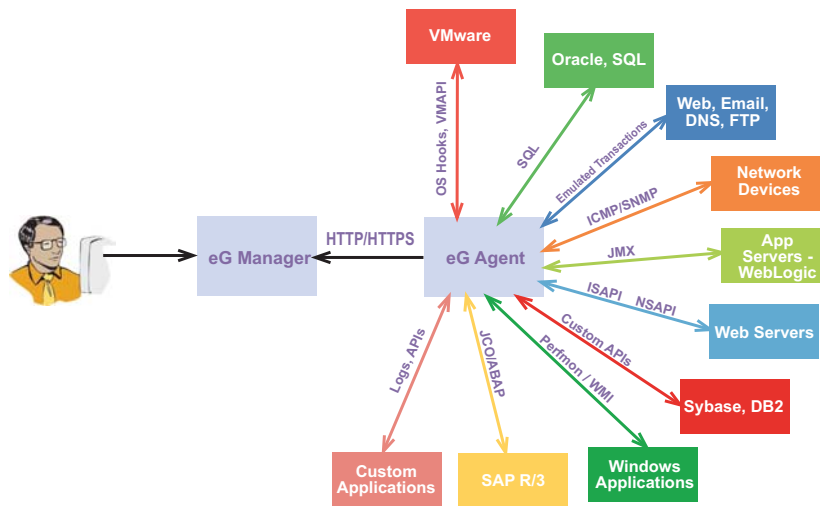


Figure 8: The eG single agent architecture

- A single agent license for Microsoft, Linux, Sun Solaris, HPUX, IBM, AIX, VMware, Tru64
- A single price, regardless of OS or server configuration - 2, 4, 8, 16 CPUs
- A single agent for monitoring any application
- A single price to manage multiple applications on the same server
- Auto-upgradeable
- Agentless monitoring option
- 100%web-based - HTTP/HTTPS

● **Monitoring every layer of every infrastructure tier** -- As its name suggests, the eG single agent is capable of monitoring a variety of networking, operating system and application technologies. Out of the box, the eG agent supports 10+ operating systems and virtualization platforms and over 85 common applications, including Citrix, terminal services, database servers (Oracle, SQL and Sybase), web servers (IIS and Apache), Active Directory, messaging servers, and Java application servers (WebLogic and WebSphere). *Figure 8* provides an overview of the eG single agent architecture. If in-depth monitoring of applications running inside the VMs is necessary, additional agents need to be installed on the VMs.

● **Automatic baselining of the target infrastructure** -- To determine the norms of the target infrastructure, eG Enterprise includes the ability for the eG manager to automatically baseline each and every metric that is collected. History is used as a guide, and trusted statistical quality control techniques are used to determine the norms of every metric collected. Any deviation of a metric from the norm is flagged as a proactive alert.

● **Problem demarcation and automatic root-cause diagnosis** -- Perhaps the most significant challenge that Virtualization 2.0 poses from a monitoring and management standpoint is effective root-cause diagnosis. When a problem happens, where is the real cause? Is it the network? Database?



**Since multiple virtual machines share the physical resources of the servers they are hosted on, a malfunctioning application inside one VM could affect the performance of other VMs on the same physical server.**

Application? Virtual machine? Physical server? eG Enterprise handles the root-cause diagnosis problem in a simple and elegant way. There are three main aspects to the solution:

- **Layer model representation** -- eG Enterprise represents each infrastructure component (network device, application, physical server) as a collection of functional protocol layers. The layers are arranged hierarchically, and the representation itself is similar to the OSI model of the protocol stack. While the OSI model was theoretical, eG Enterprise uses practical models to represent the current state of each infrastructure component.

*Figure 9* depicts the layer model of a VMware ESX server. Each layer is mapped to a number of metrics and the

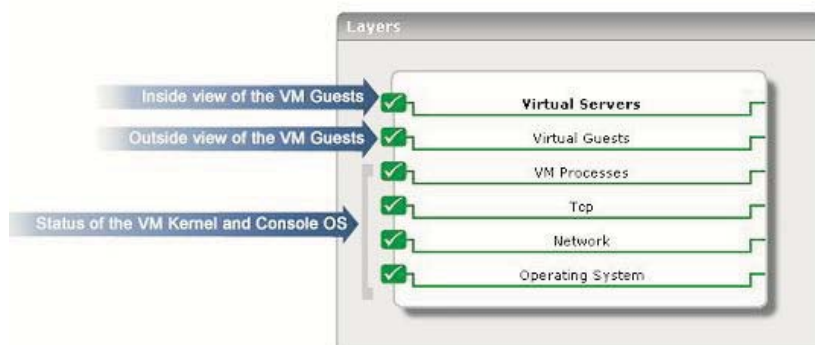


Figure 9: Layer model of a VMware ESX server. The layers clearly demarcate where the problem lies – i.e., in the physical server, the network, or the VMs.

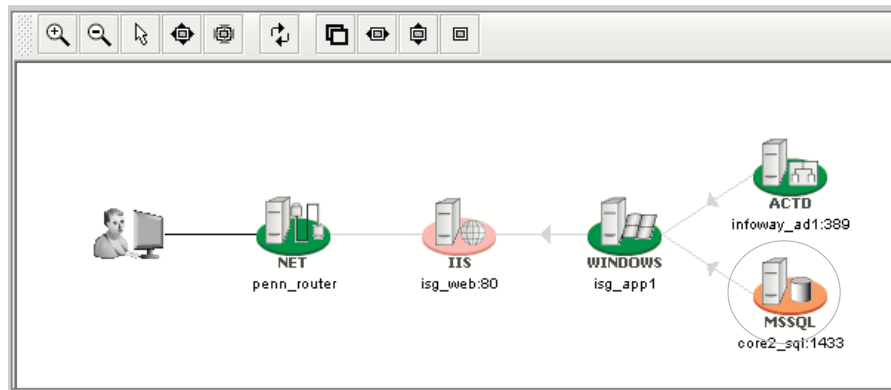


Figure 10: Topology of a business service showing components involved in supporting the service and inter-dependencies between them. eG Enterprise's automated root-cause diagnosis process has detected that a severe Microsoft SQL problem is impacting the web server.

state of a layer is determined based on the state of the metrics that are mapped to it. The state of a layer is correlated with that of other layers below it in the hierarchy. By representing the status of the physical server and the VMs as part of the same model, eG Enterprise automatically correlates their performance.

- **Using business services topologies for problem demarcation and root-cause detection** -- eG Enterprise represents service topologies to depict the components involved in supporting a business service and the inter-dependencies between them. By analyzing and correlating the status of each of the components in real-time using inter-dependency information that is available in the service topologies, eG Enterprise is able to help IT administrators with triage; i.e., to determine which domain is the potential cause of a problem.

Figure 10 shows the topology of a business service. In this example, the service is currently experiencing issues, and a look at the topology reveals that both the IIS web server and the SQL database are having problems. Based on the service topology and the inter-dependencies

contained in it, eG Enterprise has indicated that the SQL database has the more severe problem.

The business service topology representation is intuitive. Even a lesser-skilled help desk person can use the service topology view to spot that the problem in this case is the SQL database server.

- **Using VM auto-discovery for root-cause diagnosis** -- eG Enterprise auto-discovers the VMs that are running on each physical server and the applications executing on each of these VMs. This information is updated dynamically to keep track of any live migration that may be triggered in the virtual infrastructure. The application-to-VM and VM-to-physical server mappings thus discovered are used to further refine the root-cause diagnosis automatically. In Figure 10, the SQL server resides on a VMware ESX server that is also hosting a Windows application server and an IIS web server.

Figure 11 shows the virtual topology map. From the dependencies shown, eG Enterprise is able to automatically deduce that the problem in the SQL database server is being caused by an issue on the VMware ESX server.

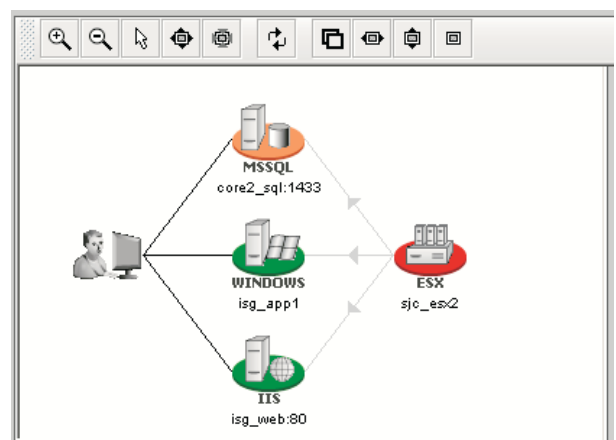


Figure 11: Virtual machine mappings are auto-discovered and a virtual topology map of VM to physical server dependencies is automatically created. This information is used by eG Enterprise to refine its root-cause diagnosis.



**“With the proliferation of virtualization comes significant management challenges. Virtual environments increase configuration, monitoring and deployment complexity for administrators. Traditional software for managing standalone machines comes up short in a virtualized environment. Combating “virtual sprawl” has turned into an overhead nightmare...As complexity increases, enterprises large and small will need better tools to manage virtual environments across storage, applications and the desktop. The big platform players have taken notice and the race is on.”**

Digital Goggles, July 18, 2008

TYPE	COMPONENT NAME	DESCRIPTION	LAYER	START TIME
Host system	sjc_esx2	CPU usage is high {console}	Operating System	04/25/08 04:19
Host system	core2_sql	Disk reads are slow {Disk0 C: E:}	Operating System	04/25/08 04:18
VMware VDI	nyc_vdi	High disk space usage {Kevin on ISGSYS8-XP:C}	Virtual Desktop	04/25/08 04:18
Host system	sjc_esx2	Many queued commands to storage device {vmba0}	Operating System	04/25/08 04:20
Microsoft SQL	core2_sql:1433	Slow response	MS SQL Service	04/25/08 04:19
IIS Web	isg_web:80	Response is slow {TransferBalances}	Web Transactions	

Figure 12: The alarm window of eG Enterprise clearly highlights where the root-cause of a problem lies.

Time	PID	%CPU	ARGS
04/25/08 09:58:16	10843	36	samba -backup 192.168.10.105 isg-web core-sql 192.168.10.102 -archive /opt/samba/today
	10875	36	samba -backup 192.168.10.105 isg-web core-sql 192.168.10.102 -archive /opt/samba/today
	17702	30.7	samba -backup 192.168.10.105 isg-web core-sql 192.168.10.102 -archive /opt/samba/today
	11209	0.9	/usr/lib/vmware/hostd/vmware-hostd /etc/vmware/hostd/config.xml -u -a
	19224	0.7	/opt/egurkha.demo/jre/bin/i386/native_threads/ja-Xrs -Dsun.net.inetaddr.ttl=900 EgMainAgent -manager 65.160.98.187 -port 80 -dir /opt/egurkha/ssl false -highSecurity false -proxyHost 192.168.10.105 -proxyPort 808
04/25/08 09:53:40	10843	36	samba -backup 192.168.10.105 isg-web core-sql 192.168.10.102 -archive /opt/samba/today
	10875	36	samba -backup 192.168.10.105 isg-web core-sql 192.168.10.102 -archive /opt/samba/today

Figure 13: Detailed diagnosis provided by the eG agent highlights the cause of the CPU spike on the service console. Several Samba backup jobs have been triggered in the middle of the day on the service console of the VMware server, and this is impacting the performance of business services that rely on the virtual infrastructure.

The exact problem in this scenario is depicted by Figures 12 and 13. Figure 12 shows the alarm window that maintains a list of outstanding alarms. The top-most alarm is the most critical: this signifies the root-cause of the problem. In this case, the alarm window clearly indicates that high CPU usage on the console operating system of the VMware ESX server is causing a response time slowdown for the business service.

Figure 13 provides an additional level of diagnosis. This figure indicates that run-away back up jobs on the ESX server are causing high CPU on the service console. In turn, this is causing the SQL server's disk accesses to be slow and thereby affects the business service's response time.

The above series of Figures 10 through 13 highlights how eG Enterprise can be used to troubleshoot a complex virtual infrastructure supporting critical business services. As virtual infrastructures continue to become commonplace, automation of the monitoring, analysis, and troubleshooting process offers several key benefits:

- ✓ Rapid diagnosis ensures lower downtime, directly translating into minimal business impact
- ✓ Proactive alerts ensure that problems are detected and corrected before users notice, thereby ensuring better customer satisfaction and productivity

- ✓ Quick problem identification ensures that only the right administrators are informed of a problem. This reduces finger-pointing among IT staff and ensures that less time and money is wasted on fire-fighting.
- ✓ By streamlining problem diagnosis, eG Enterprise empowers lesser- skilled helpdesk personnel to handle routine troubleshooting tasks. Infrastructure experts can thus spend time on more productive tasks, including capacity planning and optimization. This leads to significant operational efficiency.

**In a multi-tier infrastructure, the inter-dependencies among applications could mean that a problem in one application (e.g., a database server) could affect the performance of all the other applications (e.g., web servers, application servers) that are involved in the service delivery chain.**

- **Scalability of the monitoring solution** -- This is an important requirement for Virtualization 2.0. eG Enterprise is used in production environments to monitor hundreds of VMware ESX servers and thousands of VMs. For monitoring virtual environments, eG Enterprise offers various options. Agents can be installed on each of the servers and, since an agent is monitoring just the server on which it is installed, this approach scales easily. When agentless monitoring is used, eG Enterprise allows multiple remote data collectors to be set up. As each data collector starts to reach its monitoring limit, additional collectors can be added easily. The monitoring can be done by connecting to each physical server (e.g., VMware ESX) or by relying on the statistics collected by existing silo monitoring solutions like VMware's Virtual Center. Administrators can specify how the monitoring system is set up. The number of monitoring choices that eG Enterprise offers makes it one of the most flexible and scalable solutions in the industry.
- **Virtual desktop monitoring and reporting** -- The monitoring and reporting needs for VDI infrastructures are very similar to those for Citrix and Terminal Services infrastructures. The questions that administrators of these infrastructures need to answer include how many users are logged in, which user is logged into which VM, what resources each user is taking up, what applications are the users running, who the most frequent users are, and who the top resource consuming users are. Virtual infrastructure administration and monitoring tools (e.g.,

VMware VirtualCenter and Citrix XenCenter) do not take the user perspective and instead report resource usage for the individual VMs. On the other hand, using the inside view of the VMs, eG Enterprise is able to map VMs to users who are logged into the VMs, and all the metrics collected and reports generated by eG Enterprise report resource usage levels and login access patterns specific to the users who are using the VDI environment. Furthermore, the monitoring is done independent of the connection broker technology that is employed in the target infrastructure. The user oriented monitoring and reporting approach that eG Enterprise adopts for VDI infrastructures makes it the monitoring technology of choice for any enterprise looking for a VDI monitoring solution.

- **Personalized role-based views for different stakeholders** -- eG Enterprise includes pre-defined roles for monitoring staff, helpdesk personnel, executive staff, and capacity planners. Custom roles can also be created to suit the needs of the organization. Each user is assigned a specific role in keeping with the tasks that the user performs, and the user's view is created to provide the user with access to monitor, administer, and report on just the parts of the infrastructure that he/she is responsible for. Users may also share views of the common infrastructure components and business services that they help manage. By providing a common dashboard across heterogeneous components and services, eG Enterprise offers a platform that facilitates collaborative management among the IT operations teams.

A direct comparison of *Virtualization 2.0 Ready* requirements and the corresponding capabilities of eG Enterprise appears on the following page.

The table below summarizes why eG Enterprise is a *Virtualization 2.0 Ready* monitoring solution.

Virtualization 2.0 Ready Requirement		eG Enterprise Capability
1 .	Ability to handle a mix of physical and virtual infrastructures	<ul style="list-style-type: none"> <li>✓ Support for monitoring 10+ operating systems and 80+ applications on physical and virtual machines as well as a mix of virtualization platforms (VMware ESX, Citrix Xen, Solaris LDOMs and Containers, etc.)</li> </ul>
2 .	Support for heterogeneous virtual infrastructures	<ul style="list-style-type: none"> <li>✓ Provides out-of-the-box monitoring support for VMware Virtual Infrastructure, Citrix XenServer, Solaris Containers and Logical Domains (LDOMs), and Microsoft Virtual server environments</li> </ul>
3 .	Visibility into physical server and virtual machine configuration and performance	<ul style="list-style-type: none"> <li>✓ "Outside" view of performance of each VM</li> <li>✓ Monitoring of the virtualization platform – the hypervisor, VM kernel, service console</li> <li>✓ Monitoring of virtualization platform features such as Live Migration and High Availability</li> </ul>
4 .	Inside view of VMs with problem identification	<ul style="list-style-type: none"> <li>✓ Only monitoring solution that can provide an "outside" and an "inside" view of the virtualized environment using a single agent. The inside view is critical for root-cause diagnosis – to know which application inside the VM is faulty.</li> </ul>
5 .	Baseline metrics automatically	<ul style="list-style-type: none"> <li>✓ Uses past performance and statistical quality control techniques to automatically determine the norms of performance of every metric</li> <li>✓ Proactively alerts when these thresholds are violated</li> </ul>
6 .	Automatic correlation for pinpointing the root-cause of a problem	<ul style="list-style-type: none"> <li>✓ Correlation across virtual machines and between virtual machines and physical machines</li> <li>✓ Correlation across protocol layers to identify problematic layers</li> <li>✓ Correlate between applications responsible for business service delivery</li> <li>✓ Provide single-click diagnosis with root-cause information</li> </ul>
7 .	Scalability of the monitoring solution	<ul style="list-style-type: none"> <li>✓ Highly scalable, 100% web-based architecture</li> <li>✓ Agent-based and agentless monitoring flexibility</li> <li>✓ Integration with virtualization platform monitors like VirtualCenter; can also monitor physical servers directly to avoid any single point bottlenecks</li> </ul>
8 .	Support for virtualized desktop environments	<ul style="list-style-type: none"> <li>✓ Monitor of user activity, application mix, access patterns</li> <li>✓ Reports revealing the overall effectiveness of your virtual desktop environments – most frequent users, login/logout times for audit, applications accessed by users, top resource consumers</li> </ul>
9 .	Personalized role-based views for different stakeholders	<ul style="list-style-type: none"> <li>✓ Roles can be defined to restrict access to users based on their roles</li> <li>✓ Personalized views can be created for each user limiting their view to the portions of the infrastructure that they are responsible for.</li> </ul>

## Conclusion

This document outlines the key management challenges that must be overcome as the use of virtualization continues to increase in production enterprise environments. Virtualization 2.0 identifies fundamental changes that are needed in terms of how virtualized environments can be monitored most effectively and efficiently.

Over the last several years, eG Innovations has enhanced the eG Enterprise Suite to address the challenges that resource sharing in IT infrastructures poses. This document highlights

recent changes in the eG Enterprise Suite that are intended to make it the monitoring solution of choice as enterprises continue to look for solutions to handle the second phase of virtualization deployments -- Virtualization 2.0 -- and the phases that are sure to follow.

To obtain an evaluation of the eG Enterprise Suite or to view an online demonstration, visit <http://www.eginnovations.com/web/vmware.htm> or email [sales@eginnovations.com](mailto:sales@eginnovations.com).



***“A business user requires transparent virtualization of everything into global business applications where usage rather than the infrastructure connects the dots among them. Monitoring can’t be an after-thought. And after all, how can you manage what you can’t see?”***

***The more virtualized an organization becomes, the more complicated the monitoring becomes.”***

*Oded Noy, InternetEvolution, May 2008*



## About eG Innovations

eG Innovations, Inc. ([www.eginnovations.com](http://www.eginnovations.com)) is a global provider of IT performance monitoring and triage solutions for both virtual and physical infrastructures. The company's patented technologies provide proactive monitoring of every layer of every tier in the infrastructure, thereby enabling rapid diagnosis and recovery in enterprise and service provider networks.

By ensuring high availability and optimum performance of mission-critical business services, eG Innovations' solutions help enhance customers' competitive positioning, lower operational costs and optimize the performance of their infrastructures. The company has customers in 14 countries, including organizations of all sizes in government, banking/finance, telecom, healthcare, manufacturing and service industries.

---

### USA

#### **eG Innovations, Inc.**

33 Wood Ave. South, Suite 600  
Iselin, NJ 08830  
Ph: (866) 526 6700

### SINGAPORE

#### **eG Innovations Pte Ltd**

33A Tanjong Pagar Road  
Singapore 088456  
Ph : (65) 6423 0928

### UNITED KINGDOM

#### **eG Innovations UK Ltd.**

3 Grange Road, Camberley  
Surrey, GU15 2DH  
Ph: +44 (0) 1276 501590

### INDIA

#### **eG Innovations Pvt Ltd**

2, Murali Street, Mahalingapuram  
Chennai 600 034  
Ph : (91) 44 2817 2801

---

**Email : [sales@eginnovations.com](mailto:sales@eginnovations.com)**

**Web : [www.eginnovations.com](http://www.eginnovations.com)**