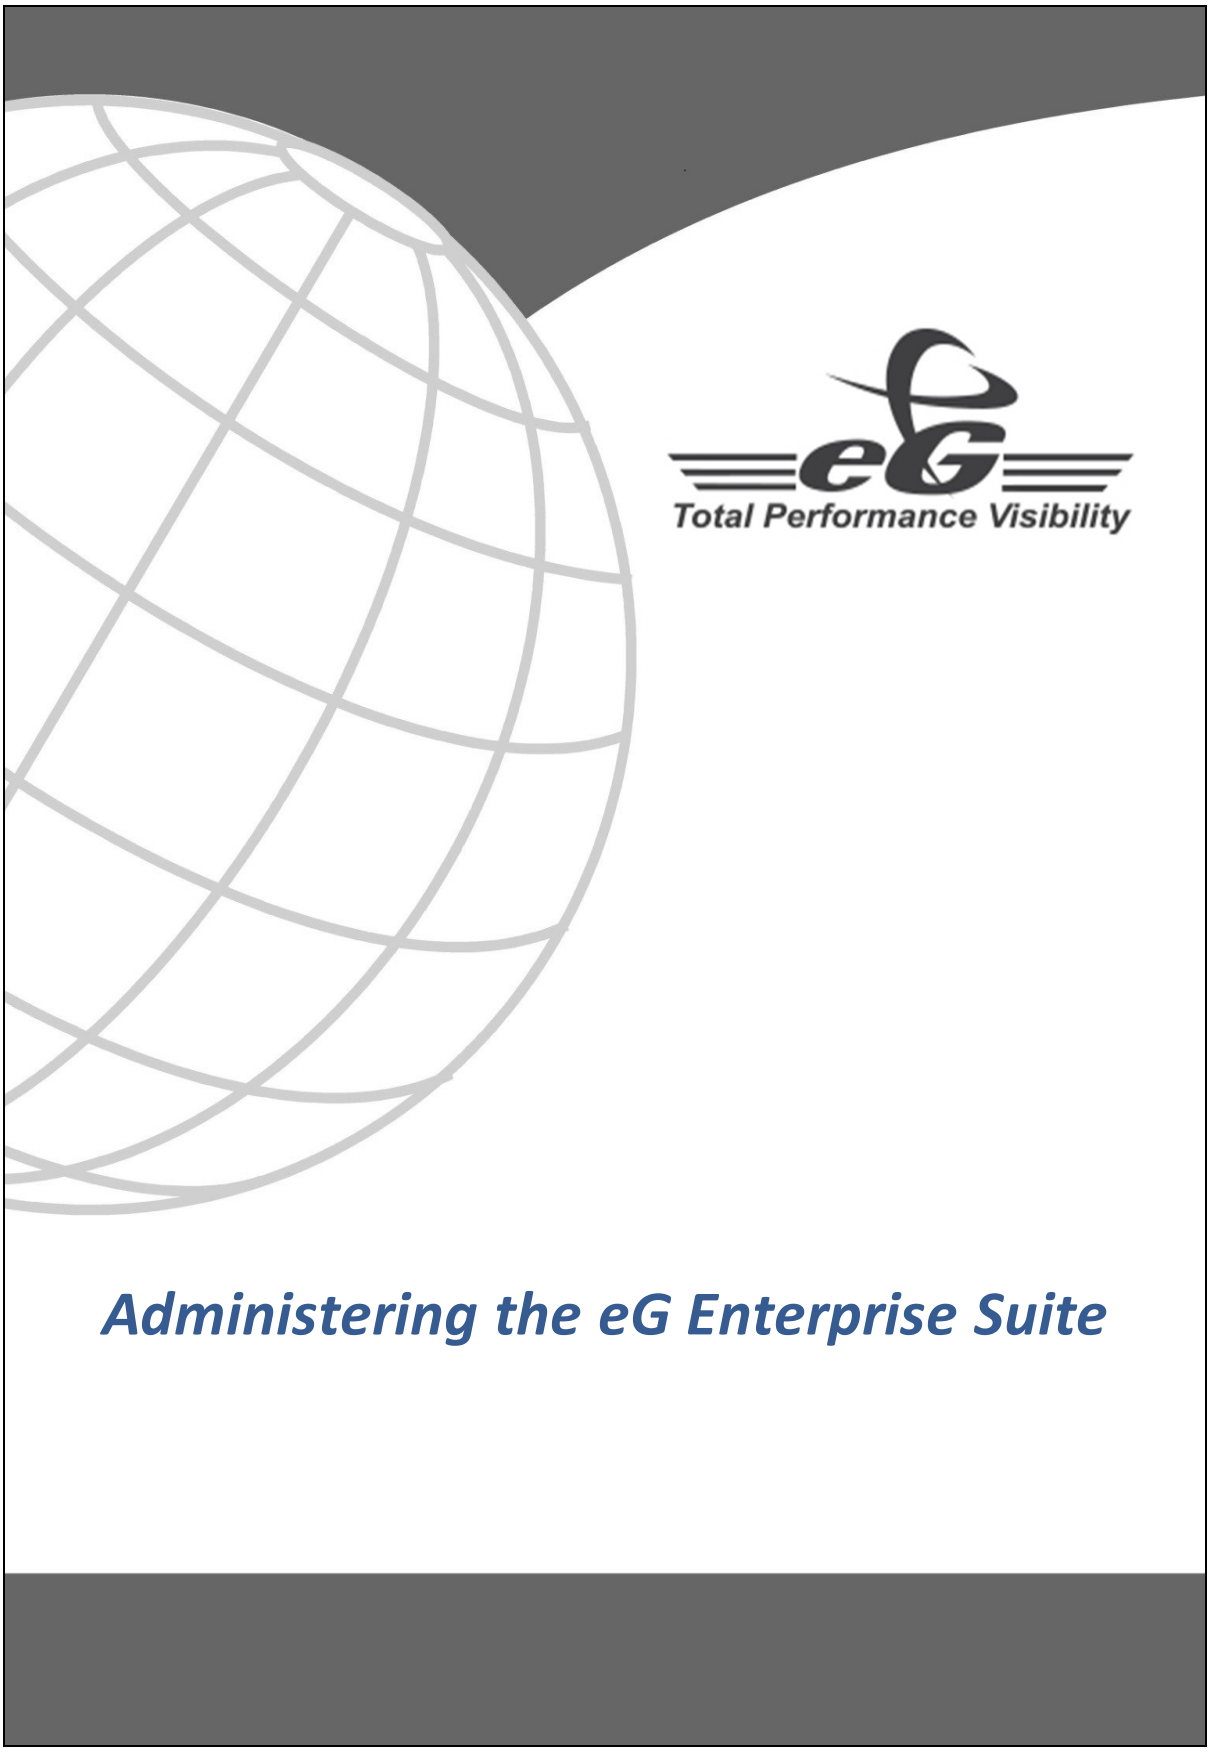

Administering the eG Enterprise Suite

Table of contents

Table of Figures

ABOUT THE GUIDE:

This document describes the architecture, installation, configuration, administration, and use of eG Enterprise.

DOCUMENT CONVENTIONS:

Conventions	Descriptions
Bold	Keywords, headings and field names
Italics	Document name, Menu sequence
Note:	Notes contain helpful suggestions.
Reference:	References to materials not covered in the manuals.

DOCUMENT FEEDBACK:

We recognize that the success of any product depends on its ability to address real customer needs, and are eager to hear from you regarding requests for enhancements to the products, suggestions for modifications to the product, and feedback regarding what works and what does not. Please provide all your inputs as well as any bug reports via email to support@eginnovations.com.

ABOUT EG ENTERPRISE:

The eG Enterprise Suite is a 100% web-based monitoring solution that offers integrated monitoring of all the infrastructure components (network, system, and application) that are involved in delivering a business service. The eG Enterprise Suite follows the traditional manager-agent architecture. While the eG manager controls what infrastructure elements are to be monitored and how frequently they are to be monitored, the eG agents are software components that perform the monitoring functions and report metrics in real-time to the management console. The eG manager takes care of storing the data reported by the agents, correlating between data from different agents to perform root-cause analysis, and providing data analysis and troubleshooting capabilities for users.

1

Overview

The dramatic increase in the complexity of IT infrastructures poses interesting monitoring and their management challenges. While IT infrastructures have grown in scale, their increased complexity is predominantly attributable to the several layers of software components that are used in these environments. A simple site itself may comprise of at least three tiers - a web server tier that front-ends user requests, an application server tier that hosts business logic components, and a database tier that hosts the business data. Owing to the inter-dependencies between the software component tiers, a problem in one tier (e.g., the database server) may ripple through and impact the other tiers (e.g., application server and web server). One of the challenges in monitoring and managing IT infrastructures effectively is to identify exactly which of the tiers is the cause of problem(s) in the IT infrastructure.

The last few years have also witnessed a radical shift in the way in which Internet servers are operated and managed. Large and small corporations and enterprises alike have begun to outsource the hosting of their servers with specialized Internet Data Centers (IDCs) and Application Service Providers (ASPs). While the hosting provider (IDC or ASP) is responsible for the hardware, and the network and software infrastructure, the actual service operating on the hosted servers is the responsibility of the customer. The presence of multiple, independent domains of control and responsibility poses interesting challenges in operating and maintaining outsourced Internet services:

- Since the performance of a service depends on the network, system, and application components that it uses, there is very often a lot of finger pointing when a problem occurs. Faced with severe competition, the hosting providers have had to expend a lot of resources in troubleshooting customer problems. Consequently, their support costs are high.
- A second complication in hosted environments is that different customer web sites and IT infrastructures can be hosted in the same network. Sometimes, different sites may even be supported on the same hardware (such a configuration is often referred to as shared hosting). Usage, performance, and availability measurements pertaining to a customer's IT infrastructures is perceived as being sensitive information that cannot be revealed or shared with other customers.
- In other cases, the different customer systems may be in different domains, probably using different IP address ranges. To protect the integrity of the customer environments, these systems may even have private, internal IP addresses that are not accessible from the open Internet. Consequently, a monitoring system for hosted environments must be capable of monitoring environments with multiple demilitarized zones, each with a set of IP private addresses.

Traditionally, monitoring systems have been viewed as a cost-center, being mostly used to improve the efficiency and internal operations of enterprises, corporate IT departments, and also IDCs. Since most monitoring systems are internal focused, hosting providers have used these systems primarily for their internal operations. Typically, customers of the hosting providers do not have a real-time view of the status and performance of their services and servers. Instead, they have to be content with weekly and monthly reports mainly focused on server and network usage.

Many existing monitoring solutions do not handle the challenges posed by the multi-domain nature of hosted environments. Moreover, they lack the ability to clearly demarcate whether a problem is caused in the customer domain or in the hosting provider's domain. Furthermore, faced with severe competition, many hosting providers are also looking to offer new, value-added monitoring and optimization services to their customers. Such a solution must be capable of:

- Offering real-time views of the status of a customer's hosted environment. The view displayed to a customer must be customizable for the specific customer - i.e., the customer must only be capable of viewing the status of the servers / applications that they have paid for. More importantly, real-time access to performance information can enable a customer to understand changes that are happening in their infrastructure in real-time and to react to these changes instantly so as to be able to offer optimal performance to their customers.
- Handling security issues across customers - i.e., one customer should not be able to view the status of another customer's environment.
- Clearly demarcating where a problem may originate in the hosted environment - i.e., whether in the customer domain or in the service provider domain. Such a capability can significantly decrease support costs for the service provider.
- Operating across customer environments in different IP address ranges, with multiple levels of firewalls between these environments.

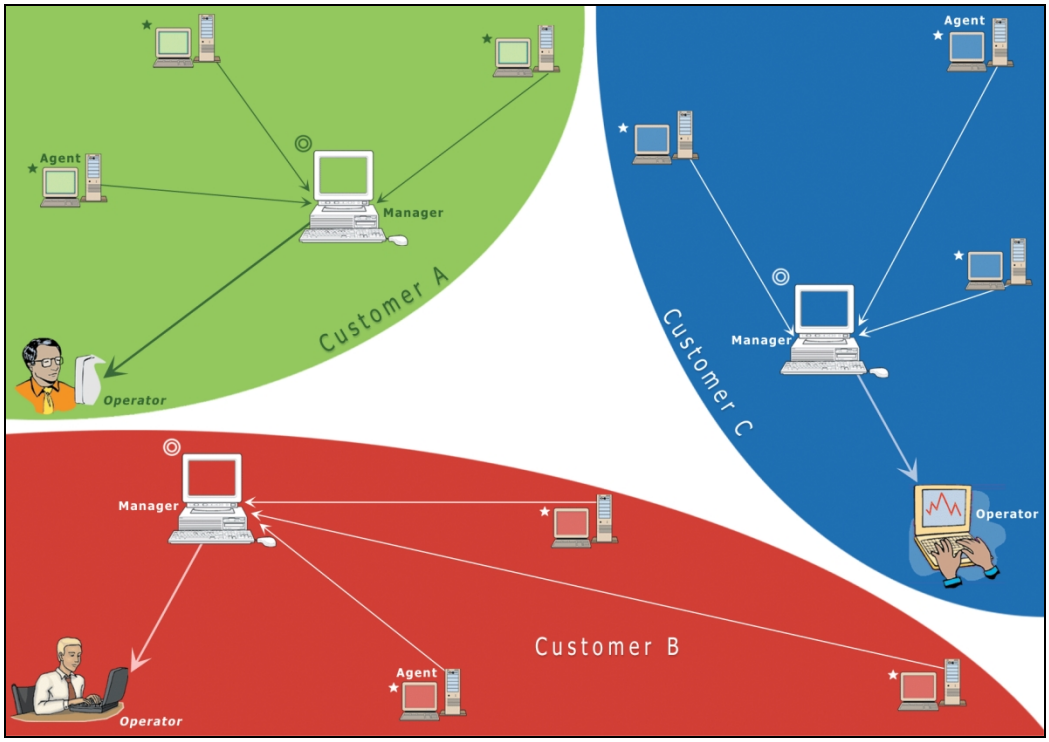

Figure 1.1: Retrofitting existing monitoring solutions for hosted environments

Many service providers are retrofitting existing monitoring solutions to meet these needs. As can be seen in Figure 1.1, in this architecture, service providers have to use separate managers for each customer supported by the hosted environment.

The drawbacks of this approach are :

- The need to own and operate multiple managers. First, separate hardware is required to host each manager. Second, the software costs - for the manager, the manager's database, etc., have to be borne individually by each customer. This need for multiple independent managers makes the overall solution proposed in Figure 1.1 very expensive.
- In the architecture of Figure 1.1, even if the same hardware is used to host different customers (i.e., shared hosting), the agents required per customer may have to be distinct, so as to preserve the security of each customer's data.

1.1 eG Enterprise

eG Enterprise is a virtual, private monitoring solution for hosted environments. eG Enterprise is virtual, because it does not involve a dedicated manager per customer. Instead, the cost of the manager component is amortized among all the customers of the hosted environment. Moreover, eG Enterprise is private because although the manager component is shared, this component is designed so as to preserve the privacy that is provided to customers in a dedicated solution such as the one in Figure 1.1. A web-based interface supported by the manager enables different customers to login to the central manager and obtain customized views of the monitored environment.

Many large enterprises too can benefit from eG Enterprise. Large enterprises often consist of multiple independent organizations that share a common network and server infrastructure. There are security constraints between these organizations. eG Enterprise's virtual monitoring architecture can be used to provide independent, personalized views for administrators of each of these organizations.

1.2 Distinguishing Features of eG Enterprise

The key features of eG Enterprise that make it a preferred monitoring solution for enterprises, IDCs, and ASPs include:

1. **Integrated monitoring** of the entire hosted infrastructure - ranging from networks to systems to applications.
2. **Automated single-click root-cause diagnosis:** Using sophisticated correlation techniques, the eG manager performs correlation in two phases to deduce the root-cause of problems. While the first phase involves correlation across the network, system, and application layers of an IT infrastructure component, the second phase involves correlation across the entire IT infrastructure, taking into account the interdependencies among components. The root-cause of problems can be reported to users over the web, via email, pager, etc. Unlike many competing solutions that require a lot of customization to perform correlation, eG Enterprise includes this correlation capability out-of-the-box, and requires no site-specific set-up and customization to enable the correlation.
3. **The virtual manager architecture** that allows hosting providers to offer revenue-producing monitoring services to customers. The main feature enabled by the virtual manager architecture is personalization. Every customer of a service provider (ASP/IDC/MSP) has exclusive access to specific infrastructure and application components being supported by the service provider. To provide personalized views of the hosted environment for every customer, the eG manager is designed as a number of virtual managers, one for each customer. A virtual manager corresponds to a specific customer, and provides customized, real-time views of the hosted infrastructure for each and every customer enabling him/her to remotely track

their on-line presence. To enable service providers to offer monitoring as a service, the eG manager allows subscription-based access for customers. At the same time, to make it simpler and less time consuming for service providers to support the monitoring service, the eG manager provides automated subscription tracking and alerting.

4. **Real-time monitoring of real web transactions:** The experience that an eBusiness offers to its customers is governed predominantly by how well the application components that support the eBusiness perform. While some products use emulated requests and log analyzers to monitor web transactions, eG Enterprise uses a novel web adapter technology to monitor real user transactions.
5. **Novel layered presentation model:** eG Enterprise's web presentation model is specifically tailored for hosted environments wherein the hosting provider is responsible for the hardware and network infrastructure, and the customer is responsible for the software applications. By depicting each IT infrastructure component as a collection of layers, and monitoring each of the layers independently, eG Enterprise is able to pinpoint which of the layers is the root-cause of problems. The isolation of problems that this layered presentation model enables is especially useful for clearly demarcating between problems in the service provider domain and the customer domain, and can significantly reduce support costs for the service provider.
6. **Centralized administration and update via a centralized web console:** A centralized user management module simplifies the creation and administration of custom views. The distributed operation of the eG agents can be controlled from a web-based administrative interface. Auto-discovery of components, configuration of the component topology, turning on and off individual tests, changing the frequency of a test and the test's parameters, updating the thresholds for every individual measurement, changing alarm policies, can all be made from an administrative interface.
7. **Scalable architecture:** eG Enterprise can scale the same way as the websites can scale to handle increased load. This is mainly due to the fact that eG Enterprise relies totally on web based mechanisms for both communicating and reporting.
8. Figure 1.2 shows the platforms that eG Enterprise supports.

Platform	Version
Solaris	7, 8, 9, or 10
Red Hat Enterprise Linux	3 (or above)
Windows	2008, Vista, 7, 8, 10, 2012
AIX	4.3.3, 5.x, 6.1, 7
HP-UX	10 and above
Free BSD	5.4
Tru64	5.1
openSUSE	11 (or above)
CentOS	5.2 (or above)
Oracle Linux	6.x (or higher)

Figure 1.2: Platforms supported by eG Enterprise

Also, agentless support is provided for the following platforms:

- Novell Netware
- AS/400
- OpenVMS
- Mac OS

9. Figure 1.2 summarizes the monitoring capabilities of eG Enterprise.

Component Type	Component Brand
Web servers	Apache, Microsoft IIS, IBM HTTP Server, Oracle HTTP Server, Sun Java Web Server, NGINX web server
Web application servers	WebLogic, ColdFusion, Sun Java Application server, Microsoft transaction server, WebSphere, SilverStream, Jrun, Orion, Tomcat, Oracle 9i OC4J, Oracle Forms Servers, Borland Enterprise Servers (BES), JBoss, Domino application server, GlassFish Enterprise Server
Database servers	Oracle, Oracle RAC, Microsoft SQL server, DB2 UDB, DB2 DPF, Sybase, MySQL, SQL clusters, Backup SQL, Intersystems Cache, PostGre SQL, Oracle RAC, DB2 DPF, SAP HANA, UniVerse Database Server
Network devices	Cisco routers, Cisco Catalyst switches, Baystack hub, Cisco VPN, Network nodes, Local director, Juniper SA and DX Device, 3COM CoreBuilder switch, Big- IP/F5 Load Balancer, Brocade SAN switches, Alcatel switches, Generic Fibre Channel switches, Cisco SAN switches, Cisco CSS, Cisco ASA, F5 Big-IP Local Traffic Manager (LTM), Coyote Point Equalizer, Coyote Point Load Balancer, Juniper EX Switch, Open VPN Access, InfoBlox
Microsoft Applications	Active Directory, BizTalk server, Windows Internet Name Service (WINS), DHCP server, MS Print server, MS Proxy Server, MS File server, ISA Proxy server, Microsoft System Management Server, Microsoft Dynamics AX, Windows clusters, MS SharePoint, FAST Search for SharePoint 2010/2013, Terminal Services Licensing server, Microsoft Dynamics CRM, Microsoft Project 2010, Microsoft DFS, Microsoft Lync, Microsoft Dynamics NAV server
Firewalls	Check Point Firewall-1, Cisco Pix, Netscreen Firewall, FortiGate Firewall, Microsoft Forefront TMG, Sonic Firewall, WatchGuard Firewall
Terminal servers	Citrix XenApp server, Microsoft RDS server
Other Citrix Products	Citrix Secure Gateway, Citrix Secure Ticketing Authority, Citrix Web Interface (NFuse), Citrix Access Gateway, Citrix Netscaler LB, Netscaler ADC, Citrix StoreFront, Citrix Branch Repeater, Citrix XenDesktop Director, Citrix XenDesktop Site, Citrix CloudBridge

Component Type	Component Brand
Citrix XenMobile	Citrix XenMobile MDM, Citrix ShareFile, Citrix AppController, Citrix Storage Zones
Email servers	Microsoft Exchange 2003/2007/2010/2013, Instant Messenger on the Exchange 2000 server, Lotus Domino R5, Sun Java Messaging, Qmail server, AsyncOS Mail, Postfix mail
Backup servers	Symantec Backup, Veeam Backup
Messaging servers	MSMQ, IBM MQ, FioranoMQ, Novell Groupwise, Tibco EMS
SAP	SAP ABAP server, SAP Internet Transaction server (ITS), SAP Web Application server, MaxDB, SAP BOBI, SAP Web Dispatcher
Virtual Infrastructures	VMware® ESX Servers 3/3.5/ESXi, Solaris Containers, Microsoft Virtual Server, Solaris LDoms, Citrix XenServer, VMware vCenter, Microsoft Hyper-V, AIX LPARs on IBM pSeries servers, IBM HMC server, Citrix Provisioning Server, Oracle VirtualBox, RHEV Server, RHEV manager, VDI-in-a-Box, KVM, Quality virtual desktop, Oracle VM server, Oracle VM Manager, Dockers
Connection Brokers	Virtual Desktop Manager, Leostream connection broker, XenDesktop Broker, VMware Horizon View, Oracle VDI Broker
SAN Storage Devices	Hitachi AMS, Hitachi USP, HP EVA StorageWorks Array, IBM DS RAID Storage, EMC CLARiiON, Dell EqualLogic, NetApp USD, EMC VNX Unified Storage, HP P2000 SAN, IBM Storwize V7000, Atlantis ILIO, QNAP NAS, Data Domain, Dell PowerEdge VRTX, IBM DS 8000, Dell Compellent, Nimble Storage, EMC XtremIO, HP 3PAR Storage
Other Operating Systems	Generic SNMP, Generic Netware, AS400 server, OpenVMS server, Mac OS
Siebel Enterprise	Siebel Web Server, Siebel Application Server, Siebel Gateway
Cloud Providers	Amazon EC2, vCloud Director, Microsoft Azure
Others, LDAP, DNS	LDAP and SunONE LDAP server, DNS and Windows DNS server, FTP, MTS, Event Logs, Tuxedo domain servers, Printers, Windows Domain Controller, NetApp filers and NetCache, SiteMinder Policy server, Radius server, COM+ server, Tcp server, ASP .NET server, Network File System on Solaris server and client, Network File System on Linux server and client, MS RAS server, MS Radius server, eDirectory server, Sun Ray server, HP Blade server, XUps, Endeca Search, Bluecoat AV, 2X Client Gateway, 2x Publishing Agent, 2X Terminal server, SunONE Directory Server, Cisco UCS manager, Egenera PAN Manager, Teratext Arbortext, Teratext Content Server, DoubleTake Availability, Marathon everRun, Microsoft RDS Licensing Server, Delta UPS, NTP server, IBM

Component Type	Component Brand
	Integration Bus, App-V Client and Server, Tuxedo PIA , CheckPoint Smart Appliance

Figure 1.3: IT infrastructure components monitored by eG Enterprise

2

System Architecture

Before getting into the details of how eG Enterprise is installed and configured, it is imperative for a user to understand the architecture of eG Enterprise. A thorough understanding of eG Enterprise architecture can enable the user to deploy and use eG Enterprise product effectively. This chapter delves into the details of eG Enterprise's architecture.

eG Enterprise follows the manager-agent architecture that has been widely used in the past for designing monitoring systems. While the manager is a software component that controls what elements are monitored and how frequently they are monitored, the agents are software components that perform the monitoring functions. Figure 2.1 depicts the main components of eG Enterprise and the following sections describe these components in detail.

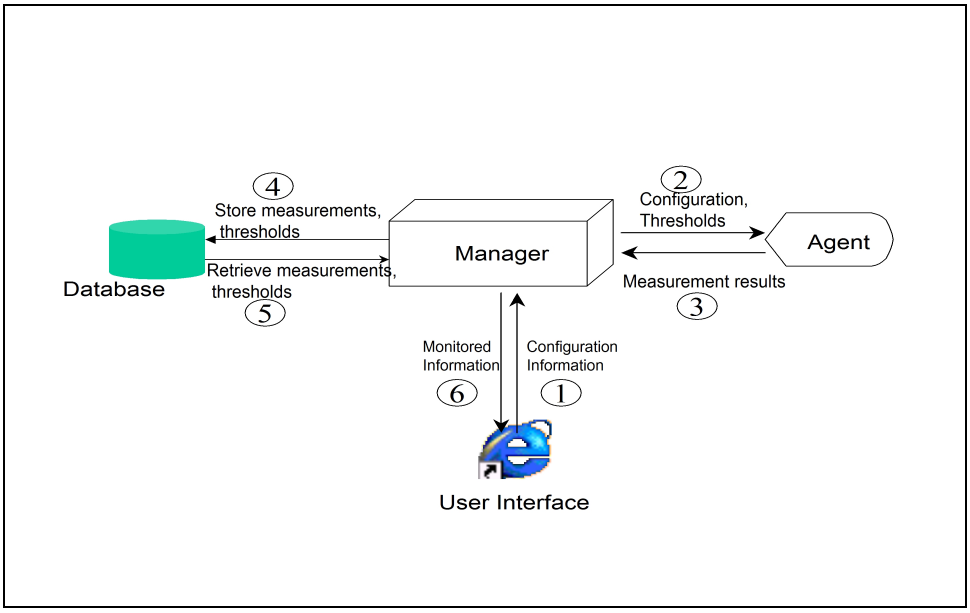

Figure 2.1: Main components of eG Enterprise

2.1 Manager

The eG manager is responsible for coordinating the functioning of the agents, analyzing the reports from the agents to determine whether any problems exist, and for handling user requests to eG Enterprise. The main functions of the manager are:

1. **Environment discovery:** The manager discovers the components, applications, and network elements that exist in the environment.
2. **Agent specification:** Based on the output of the discovery process, the manager specifies where agents must be deployed in the target environment and what tests each agent should run.

3. **Database storage** of measurement results returned by agents.
4. **Threshold computation** to determine the normal limits for each measurement being reported by the agents.
5. **Alarm correlation:** This involves diagnosis and reporting of the root cause of problems detected in the target environment.
6. **User interactions:** The manager handles all requests from users to eG Enterprise.
7. The eG manager comprises of two major components:
   - **Virtual managers:** The manager is designed as a number of virtual managers, one for each customer. A virtual manager corresponds to a specific customer, and is responsible for providing customized displays of the hosted environments for the customer. The virtual manager also handles license tracking for a customer and generation of alerts in real-time for the customer. The virtual manager uses a core set of functions supported by a second manager component called the main manager.
   - **Main manager:** This component implements the core set of functions of the manager such as the receipt and storage of the measurement results, threshold computation for the collected results, analysis of the stored data for trending and service-level audits, alarm correlation for root-cause diagnosis, user login, configuration of the user's virtual monitored environment, etc.

In eG Enterprise implementation, the virtual managers are optimally implemented within the context of the main manager process itself. Figure 2.2 depicts the virtual private manager architecture.

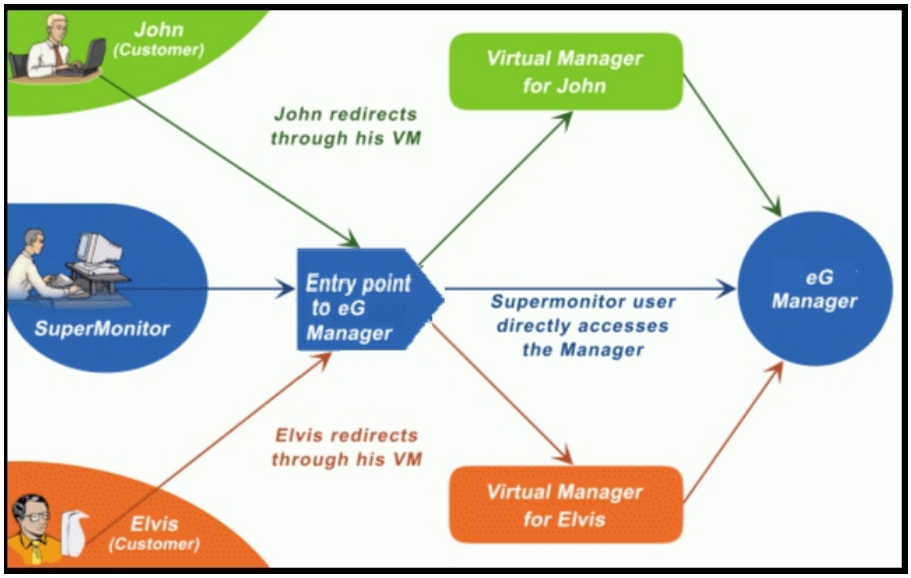

Figure 2.2: Architecture of the virtual, private manager

2.2 Agents

The agents monitor the environment by running periodic **tests**. The outputs of the tests are called **measurements**. A measurement determines the state of a network / system / application / service element of the target environment. For example, a **Process test** reports the following measurements:

1. Number of processes of a specific type executing on a system.
2. The CPU utilization for these processes
3. The memory utilization for these processes

Agents use different approaches for testing the target environment. The tests can be executed from locations external to the components that are responsible for the operation of the IT infrastructure. Agents that make such tests are called **external agents**. These agents take an external view of the IT infrastructure and indicate if the different services supported by the IT infrastructure are functioning properly or not.

Often external agents alone may not be sufficient to completely gauge the health of an IT infrastructure and to diagnose problems when they occur. For example, it may not be possible to measure the CPU utilization levels of a web server from an external location. To accommodate such situations, eG Enterprise uses **internal agents**. An internal agent runs on a component that supports the IT infrastructure and monitors various aspects pertaining to the component (e.g., CPU, memory, and disk utilization, the processes executing on it, and the applications).

For making measurements, eG agents support various mechanisms. The Simple Network Management Protocol (SNMP) continues to be the standard for monitoring network elements (routers, load balancers, WAP gateways, etc.). Besides monitoring network elements, eG Enterprise also manages systems and applications. SNMP is rarely supported at the application layer. Hence, for monitoring applications, eG Enterprise supports various other mechanisms:

1. **Emulated transactions:** By emulating typical transactions from clients to different applications, eG agents monitor various aspects of the component. For example, to measure the health of a web server, eG agents use an HttpTest that emulates user accesses to the web server. Depending on whether and when a response is received or not, as well as based on the status code returned by the web server in the Hyper Text Transport Protocol (HTTP) response returned by the web server, the eG agent assesses the availability of the web server and the response time for the request.
2. **SNMP data collection:** To monitor the various network elements and any other application components that support SNMP, eG agents support SNMP-based monitoring.
3. **OS-specific instrumentation:** Operating systems already collect a host of statistics regarding the health of the component and processes executing on it. For example, CPU, memory, and disk space utilizations, network traffic statistics, process-related measures can all be collected using operating system specific hooks. eG agents use these hooks to collect and report a variety of statistics of interest.
4. **Application specific adapters:** For monitoring specific applications, eG agents use custom adapters. One example of a custom adapter is the `web adapter`. eG's web adapter is designed to enable web sites to collect statistics regarding user accesses in real-time, without the need for explicit logging of requests by the web server. The web adapter is a layer that fits between the TCP/IP stack and the web server itself. It can be thought of as a passive probe that watches the requests received by the web server and the responses produced by the web server. By applying a fast, pattern-matching algorithm on the packets that flow by, the web adapter collects a variety of statistics regarding web sites and the transactions executed by users at these sites. Details of the statistics collected by the web adapter are provided in the eG Measurements Manual.

5. eG agents have been pre-programmed to execute specific tests for web servers, SSL servers, LDAP servers, DNS server, Database servers, and web application servers. Please see the eG Measurements Manual for details on the tests included in eG Enterprise.

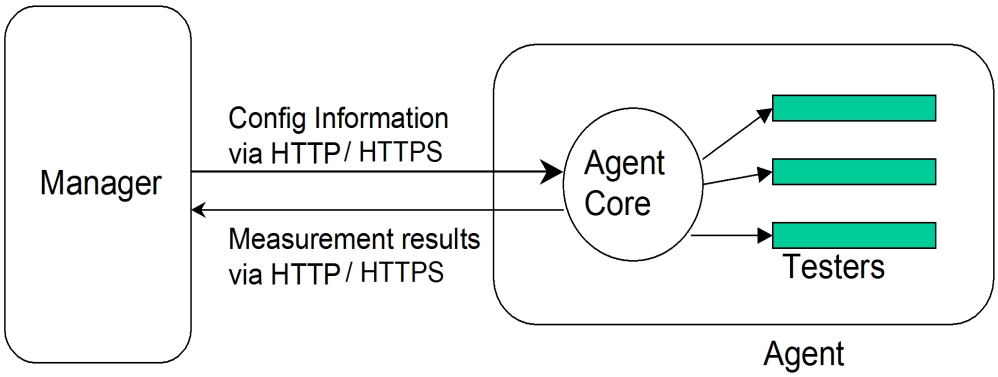

Figure 2.3: The manager-agent communication in the eG architecture

Figure 2.3 depicts the manager-agent interactions. All manager-agent communication happens over the HTTP or HTTPS protocol. The agent uses **tester** threads, each of which is responsible for a specific test. The main functions of the agent core are:

- To read configuration information from the manager and determine what tests are to be executed on a host.
- To periodically refresh the configuration information from the manager and determine if any of the testers needs to be stopped or restarted, or whether the configuration information for any of the tests needs to be changed.
- To read the threshold information from the manager and use it to determine whether the state of each measurement is normal or not
- To provide alarms to the manager in the event that the state of any measurement changes
- To upload measurement results back to the manager for permanent storage.

Figure 2.4 depicts the typical deployment architecture of eG Enterprise. The eG manager is installed on a component called the eG server. By default, an external agent is also hosted on this system. Internal agents are installed on all the other components being monitored in this environment. The configuration of external agents can be modified to suit the target environment. For example, in Figure 2.4, an external agent is located within each customer's network.

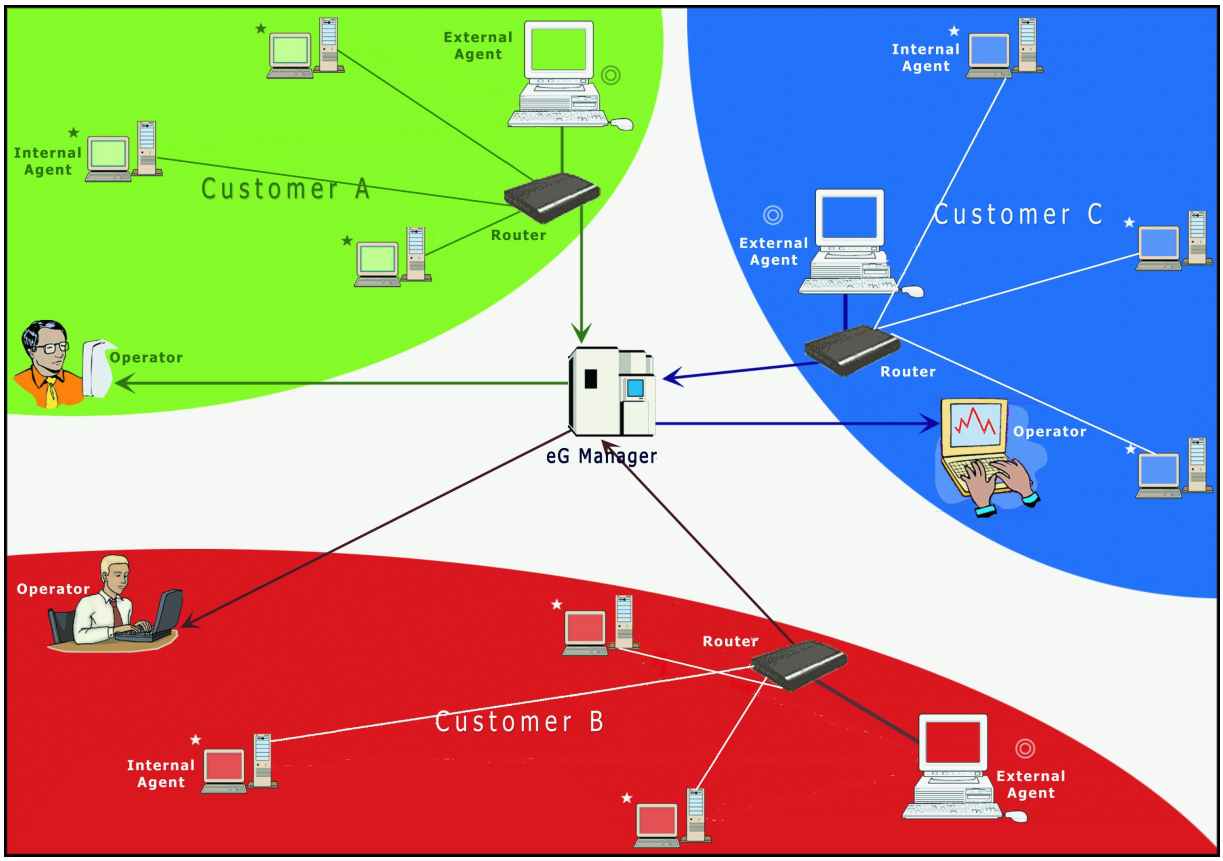

Figure 2.4: Typical deployment architecture of eG Enterprise

2.2.1 Customizing Applications for Monitoring by eG's Web Adapter

The HTTP protocol specification provides various status codes that are used by web applications to indicate error conditions. For example, a response code of 404 indicates that a specific page was not found on the web server. Likewise, a response code of 500 indicates a server-side processing error. As the Web has evolved to support a variety of complex applications that involve dynamic rather than static content, application developers have resorted to newer methods of informing users of application-level problems. For example, rather than returning a 404 response code to indicate missing content (which results in the browser throwing an error message), an application developer may choose to report the error by formatting it within a HTML page (i.e., by providing a 200 response code) and providing additional user-friendly error messages such as the email address of the web site's administrator. The side effect of this approach is that the large number of existing monitoring tools that primarily use the HTTP response code to detect application failures will not be able to effectively detect and report problem conditions.

To enable eG agents to detect and report on such application-specific problems, eG's web adapter allows applications to use a specific HTTP header variable called *Eg-Status* to report error conditions to it. To report a specific problem, application developers should assign corresponding error codes (following the HTTP protocol specification's response code convention). For example, to report an application-specific error, an application that uses Java technologies (JSP, Servlets, etc.) can incorporate the following code to set the Eg-Status header value while generating a HTTP response:

```
response.setHeader("Eg-Status", "500");
```

The eG web adapter searches the HTTP header of all responses generated by a web server. If the Eg-Status header exists, the value of this variable is used to override the HTTP response code value. This method allows application developers to indicate potential error conditions to the eG agents, without affecting the output being provided by their applications to users.

2.2.2 Agentless Monitoring

To support environments where administrators may not be interested in deploying agents on every server, eG Enterprise offers an agentless monitoring option. A remote data collector - called the Remote Agent - can be deployed on a central server and using agentless interfaces (such as SNMP, NetBios/Perfmon, Secure Shell, etc.), the remote agent can collect statistics on a number of servers/applications without needing agents to be installed on each and every server. Details of the agentless monitoring approach and the tradeoffs in using agentless monitoring as compared to agent-based monitoring are explained in Section **7.4**.

Note:

Typically, the remote agent should be installed and configured on the same operating system and locale as that of the servers that are monitored by that agent. In multi-lingual environments therefore, you would require a remote agent for every locale that is in use - for instance, in environments with servers that support both French and Japanese locales, you would require an exclusive remote agent for the French servers and another for the Japanese servers.

2.3 Database

The eG database is responsible for persistent storage of the measurement results. Separate tables are maintained for each of the tests being executed by eG agents. Besides the measurement tables, the database hosts threshold tables for each test. A threshold table indicates the upper and lower ranges of the threshold values for each measurement.

The database design provides a way to periodically purge old data from the database. The periodicity with which the data will be purged by the database is configurable by the user.

2.4 User Interface

A web-based user interface enables a user to interact with eG Enterprise. The recommended browser for the eG user interface is Internet Explorer 9.0 and above, Mozilla Firefox Version 16 and above (OR) Chrome. Broadly, the eG user interface allows a user to first customize the configuration of the eG agents (i.e., what components and services to monitor, how frequently to monitor, what specific tests to run, etc.) and subsequently to monitor the measurements made by the agents.

To avoid overwhelming users with the variety and amount of results being generated based on measurements made by the eG agents, the user interface presents the results of the measurements in a logical and coherent manner. The eG manager's interpretation of the state of each element of the IT infrastructure is first displayed before the results of the individual measurements are made available - e.g., by displaying graphs indicating the change in value of the measurement with time of day. An alarm window immediately highlights the

pending alarms in the target environment, prioritized based on the eG manager's assessment of the severity of the associated problems.

2.5 Scalability Options for the eG Manager

The eG manager runs as a Java process. The maximum heap memory that can be allocated to a 32-bit eG manager process is limited to 1.5 GB. The maximum heap memory allocation to a 64-bit eG manager process on the other hand, is limited to 3 GB.

Where a large number of components are to be monitored, you may want to allocate more memory heap to the eG manager process. In such a case, follow the steps discussed below on an eG manager on Windows:

1. Login to the eG manager host.
2. Edit the **mgrdebugon.bat** or **mgrdebugoff.bat** file in the **<EG_INSTALL_DIR>\lib** directory.
3. Search for the entry *JvmMx* in the file. You will then find an entry that reads as follows:

   -JvmMx 1024 -JvmMs 1024
4. The *JvmMx* and *JvmMs* specifications govern the heap memory allocations to the eG manager. Both these specifications will be set to *1024* (MB) by default. If you want to increase it to say, 2 GB (i.e., 2048 MB), change these specifications as indicated below:

   -JvmMx 2048 -JvmMs 2048
5. Finally, save the file, and run the **mgrdebugoff.bat** or **mgrdebugon.bat** file (as the case may be).
6. On a Unix manager, follow the steps below to modify the heap memory allocation:
7. Login to the eG manager host.
8. Edit the **catalina.sh** file in the **/opt/egurkha/manager/tomcat/bin** directory.
9. Search for the entry *Xms* in the file. You will then find an entry that reads as follows:

   -Xms256m -Xmx256m
10. The *Xms* and *Xmx* specifications govern the heap memory allocations to the eG manager. Both these specifications will be set to *256m* (i.e., MB) by default. If you want to increase it to say, 512 MB, change these specifications as indicated below:

    -Xms512m -Xmx512m
11. Finally, save the file.

    While overrding the default heap memory allocations to the eG manager process, ensure that the allocated heap memory should not be greater than the total memory capacity of the eG manager host.
12. Moreover, even if the physical server on which the eG manager is installed has more memory, since it is a single Java process, the eG manager cannot exploit the additional memory available on the server. To overcome this limitation, in eG Enterprise, the critical eG manager functions such as email alert management, threshold computation, trending, and database cleanup activities can all be run as separate Java processes (i.e., in addition to the core eG manager process).

13. For configuring this, follow the steps below:

    - Edit the **eg_services.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
    - In the **[EXEC]** section of the file, you will find the **ExecuteMailAsProcess**, **ExectureTrendAsProcess**, and **ExecuteThresholdAsProcess** flags set to **No** by default. This implies that by default the eG manager functions of email alert management, trend computations, and threshold computations will be performed by a single eG manager process only. Similarly, the **ExecuteCleanupAsProcess** flag in the **[MISC_ARGS]** section of the file is also set to **No** by default, indicating that database cleanup activities are also performed by default by the single eG manager Java process only. To spawn separate Java processes for each of the above-mentioned functions, set each of the flags mentioned above to **Yes**.
    - Also, these additional Java processes can be spawned with separate Java heap settings using the **Mail_java_options**, **Thresh_java_options**, and **Trend_java_options** parameters in the **[EXEC]** section. By default, the mail alerting, trend computation, and threshold calculation processes are configured with the following minimum and maximum heap settings:

      Mail_java_options=-Xrs -Xmx256M

      Thresh_java_options=-Xrs -Xmx384M -Xms256M

      Trend_java_options=-Xrs -Xmx364M -Xms256M

      While the value specified next to the **-Xmx** entry is the maximum memory that can be used by the corresponding process, the value specified next to the **-Xms** parameter represents the minimum memory setting of that process. For instance, the **Thresh_java_options** parameter is set to **Xrs -Xmx384M -Xms256M** by default. This implies that the process that computes thresholds is by default configured with a minimum memory of 256 MB (**-Xms256M**) and a maximum heap size of 384 MB (-**Xmx384M**). If required you can change the minimum and maximum memory values to suit the needs of the process in your environment.
    - Finally, save the file.

Removing these key functions from the core eG manager process makes additional memory available for the core eG manager functions including data reception and analysis, alarm correlation, and web-based access and reporting. This reconfiguration of the eG manager into separate Java processes allows the eG manager to make better utilization of available server hardware resources and thereby offers enhanced scalability. In turn, this allows customers to get more leverage from their existing investment in the hardware that hosts the eG manager.

You can also closely track the status of the threshold computation and trend computation processes by enabling logging for each of these processes using the **MANAGER SETTINGS** page (Configure -> Settings menu sequence) in the eG administrative interface. This will result in the creation of the **thresh_log** and **trend_log** files (in the **<EG_INSTALL_DIR>\manager\logs** directory), to which the details of the threshold and trend-related activities (respectively) will be logged as and when they occur. For tracking the email alerting process on the other hand, you will have to enable logging for this process using the **MAIL LOG DETAILS** section of the **MAIL SERVER - ADVANCED OPTIONS** page (Configure -> Mail Settings -> Advanced) in the eG administrative interface. Once logging is enabled, log files for this activity will be created in the **<EG_INSTALL_DIR>\manager\logs\egmailmanager** directory.

While these log files can report on the status of the threshold, trend, and email alerting related operations, they cannot provide pointers to why any of these operations failed. To easily troubleshoot the failure of these critical processes, the errors/exceptions raised by these individual processes should be captured. To enable this, do the following:

- Edit the **eg_services.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
- In the **[EXEC]** section of the file, you will find three flags that are set to **No** by default. These are as follows:

```
ExecMailLog=No
```

```
ExecThreshLog=No
```

```
ExecTrendLog-No
```

  These entries indicate that, by default, the errors/exceptions encountered by the email alerting, trending, and threshold computation activities are not captured by the eG Enterprise system. To ensure that exceptions related to each of these entries are logged, then set all the three flags above to **Yes**.

- This will result in the creation of the **mail_out**, **thresh_out**, and **trend_out** files (in the **<EG_INSTALL_DIR>\manager\logs** directory) to which the errors/exceptions related to the email alerting, threshold computation, and trend activities (respectively) are logged.

2.6 Self-Monitoring and Recovery

To ensure enterprise-class monitoring, the eG manager includes the capability to monitor its various components and to recover from failure of these components. When the eG manager is started, a separate eG recovery process is started. On Unix environments, this process is called *eGmon*. On Windows environments, this process executes as a service named *eGmon*.

This process periodically attempts to connect to the eG manager, access the various components of the manager, including the eG database. If it detects any problems during such access, the recovery process attempts to perform further diagnosis. The specific actions performed by the recovery process are as follows:

- If the eG manager is not accessible, the recovery process attempts to restart the eG manager. If it fails to restart the eG manager thrice in succession, the recovery process generates an alert message to the eG administrator (using the **MAIL SENDER ID** specified in the **Mail Configuration** settings of the administration interface).
- If the eG manager is accessible, the recovery process tests the connections from the eG manager to the database server that it uses. In the event it detects problems, it alerts the administrator of potential problems with the database server access. By connecting directly to the database server (i.e., without using any other eG manager components), the recovery process further determines whether the database access problem is being caused either because of a database failure or because the eG manager's pool of database connections is not sufficient to handle the current load on the manager.

When the eG manager is stopped manually, the eG recovery process is also shutdown.

To further improve its resilience to failures, eG Enterprise is architected in such a way that when an eG agent is not able to report measurements to the manager, it stores a local copy of the measurement results. When its connection to the manager is re-established at a later time, the agent uploads the saved measurement results to the manager, thereby ensuring that measurement results are not lost even if the manager/agent connection fails temporarily.

Note:

Handling of Old Data from the eG Agents:

Typically, if the network link between an eG agent and the manager goes down, the agent stores the metrics it collects locally and later once the link comes back up, the agent uploads the metrics to the eG management console. This design ensures that loss of monitoring data during network outages is minimized. A configuration setting on the eG manager governs how the eG manager handles old data being sent by an agent to it. This setting is the *OldDataIgnorePeriod* entry in the **eg_db.ini** configuration file in the **<EG_INSTALL_DIR>\manager\config** directory of the eG manager. If this entry is unavailable or if its value is -1, the eG manager chooses to process the old data being sent by the agent as if the data has been generated in real-time. Thus, all measurement results from the agent are analyzed and alerts generated by the eG manager if any abnormality is detected.

Some administrators may prefer not to have the eG manager process old data. For instance, suppose the network link has been down for 3 hours, and during this period, a process went down for a while and came back up. The eG agent's measurements would indicate the change in state of the process, and if the eG manager processes the old data, it would first generate an alert indicating that the process went down, followed by an almost immediate event indicating that the process has restarted. Administrators who do not wish to receive alerts for older data from the agents can define the period of time beyond which the eG manager determines that data being received from the agents is old data. For example, if the **OldDataIgnorePeriod** is set to 10, the eG manager will consider all data that has a timestamp earlier than 10 mins prior to its current time as old data, and state computations and alarm correlation are not performed using such data.

The eG agents too include self-monitoring capabilities. When the eG agent is started, a separate recovery process is also started. On Unix environments, this process is driven by a script **eGAgentmon**. On Windows environments, this script executes as a service named **eGAgentmon**.

Every 5 minutes, this script spawns a process named **java EgCheckAgent**, which checks if the agent is alive or not. If the agent is found to have stopped abnormally, then **java EgCheckAgent** process restarts the agent. The recovery process records all of the recovery actions it attempts and records the outcome (i.e., whether success or failure) of these actions in the agent log file, which is located in the **<EG_HOME_DIR>/agent/logs/error_log**). **EG_HOME_DIR** refers to the installation directory of the eG manager and agents (eg. **/opt/egurkha** on Unix, **C:\Program Files\eGurkha** on Windows).

However, note that if the eG agent is stopped manually, the agent recovery process is also shutdown.

So far we have highlighted the key components of eG Enterprise. The four stages in deploying eG Enterprise in the target environment are :

1. **Installation** of the eG manager and the agents. This stage mainly involves deployment of the software on the appropriate components, creating user accounts, and setting up the directory structures.

2. **Configuration** of the eG manager and the agents. In this stage, the environment is set up for the proper operation of eG Enterprise and the manager and agent processes are started.
3. Please refer to "*The eG Installation Guide*" for a detailed description on the above two stages.
4. **Administration** of the eG Enterprise system. At this stage, the user interacts with the eG manager through the eG user interface to determine where agents must be deployed, what tests these agents must run, how often the tests should run, etc.
5. **Monitoring** using the eG Enterprise system. At this stage, using the user interface, users can monitor various aspects of their IT infrastructure.
6. Figure 2.5 depicts the various stages involved in deploying eG Enterprise in a target environment.

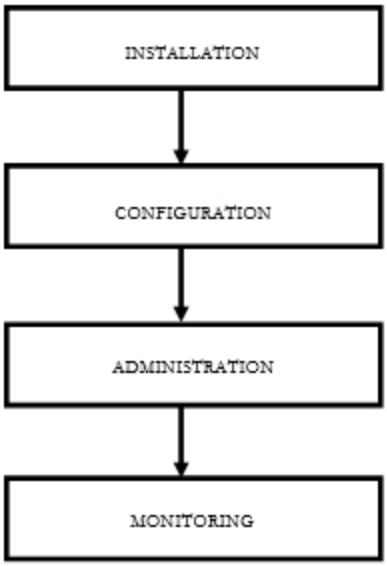

Figure 2.5: Stages involved in deploying eG Enterprise

For details of the first two steps, please refer to the *eG Installation Guide*. The rest of this manual focuses on the third and fourth steps.

3

Licensing

eG Enterprise follows the manager-agent model and a centralized licensing. The eG licensing policy is based on the following thumb rules:

a. eG Enterprise has pre-built rules that determine where the different agents should run. The eG manager is automatically associated with one external agent, by default. Additional external agents can be added as required. The eG license controls the maximum number of external agents that are allowed.

b. The number of internal agents necessary is equal to the number of unique IP addresses or nick names being managed. When the eG agent is installed on a host, the eG Enterprise system automatically determines what tests need to be run on the host. Hence, the eG agents do not require different licenses for managing web servers, application servers, database servers, LDAP and DNS servers, etc.

c. The eG license controls the total number of **Monitors** that can be used by an eG installation. The **Total Monitors** listing in the eG license indicates the total number of **basic** and **premium** monitors that the current installation of eG Enterprise is allowed to use, the total number of such monitors that are currently utilized, and the overall usage percentage.

3.1 Types of eG Monitoring Licenses

eG monitoring licenses are of two types:

- Server-based
- User-based

Customers can pick one of the license types listed above, depending upon the nature of their environment and their monitoring needs. A mix of server-based and user-based licenses can also be used – e.g., all the key server applications like Active Directory, SQL server, Java applications can be monitored with agent/agentless monitor licenses and thin-client servers can be monitored with named user licenses.

The sections that follow will discuss each of these license types in detail.

3.1.1 Server-based eG Monitoring License

By default, the eG monitoring license is server-based. For each server, a unique nick name is assigned and an eG monitor license is required. The number of eG monitors that can be deployed in the target infrastructure is controlled by the eG monitoring license.

eG Enterprise supports agent-based and agentless monitoring. The eG monitors can be inter-changeably used and the IT manager has complete flexibility in deciding which servers to monitor with agents and which ones to monitor without agents. Any combination of agent and agentless monitoring can be used and eG Enterprise still provides a consistent view of metrics across these different monitor types.

Since the license is deployed on the eG manager, there is no need to deploy any licenses on the agents. Furthermore, the eG license only controls the number of monitors that are deployed. It does not control which

servers are monitored. Hence, the IT manager can decide to monitor one set of servers this week and over the next week, can choose to deploy the monitoring on a completely different set of servers. As long as the total number of monitors being deployed is within the limit indicated in the license, your eG Enterprise system will function correctly.

Server-based licensing is ideal for environments where a large number of users access a few servers – eg., in server-based environments (e.g., Citrix XenApp) hosted on physical servers having 100-200 users per server. In this case, the cost of the monitoring license is amortized across the users, thus making server-based licensing a cost-effective choice.

Different types of monitors are supported by the eG server-based license. Each of these types have been elaborately discussed in the sub-sections below.

3.1.1.1 The OS Monitor

To monitor a server operating system, you need an **OS monitor**, which is referred to in the eG license as a **basic monitor**. Using a basic monitor license, you can monitor Windows, AIX, Solaris, Linux, HPUX, Netware, AS400, OpenVMS or Mac OS. This license can also be used for Microsoft File and Print servers. The OS monitor tracks system uptime, utilization of key OS resources such as CPU, memory, and disk, network traffic to and from the server, the performance of the TCP/IP stack, etc. It can also monitor application log files for exceptions, Windows and Unix system logs, and it can monitor the status and resource usage of key processes/services running on the system. Typically, the OS monitor is used to any staging/development systems that do not require in-depth monitoring of applications.

The eG licensing is not based on the hardware capabilities of the server being managed or on the specific operating system being monitored. This provides unparalleled flexibility to an IT manager. Therefore, she/he can be managing a Unix server on one day and a Windows servers on another, with the same eG monitoring license.

3.1.1.2 The Virtualization Monitor

To monitor a virtualization platform such as VMware vSphere, Citrix XenServer, Microsoft Hyper-V, Solaris LDoms AIX LPARs, etc., you need a **virtualization monitor**. In the eG license, this is referred to as a **premium monitor**. The number of licenses required is equal to the number of virtualized servers to be monitored. This implies that the licensing is not based on the number of CPU cores or sockets on the servers, its memory configuration, or the number of virtual machines (VMs) that are hosted on the server. With a single virtualization monitor, an administrator can monitor the hypervisor as well as the VMs. For the VMs, eG Enterprise provides its patent-pending In-N-Out monitoring view that allows the administrator to see the portion of physical resources that a VM is using as well as the portion of virtual resources that each application running inside the VM is consuming.

3.1.1.3 The Application Monitor

To monitor applications such as Oracle databases, Microsoft SQL server, web servers like Apache and IIS, Java application servers like Tomcat, JBoss, WebLogic and WebSphere, Citrix XenApp and Terminal servers, or any of the other 150+ applications that eG Enterprise supports, you will need an **application monitor** for each server to be monitored. In the eG license, this is referred to as a **premium monitor**. One premium

monitor is required for each server operating system being monitored (assuming one IP address per operating system). This means that eG's licensing is not for individual applications. If multiple applications run on the same system (e.g., IIS, SQL server and Tomcat all run on the same system), a single premium monitor license will suffice. Likewise, eG's licensing model allows IT managers great flexibility in deploying the monitoring. The eG Enterprise suite does not use the concept of knowledge modules or smart plugins for each application to be monitored. This means an IT manager can use an eG license to monitor an Oracle database on one day and reuse the same license to monitor a Citrix XenApp server on another day. Note that an application monitor includes the capabilities of an OS monitor, so a separate OS monitor license is not required if a server already has an application monitor license.

3.1.1.4 The eG External Monitor

To monitor network devices using SNMP, to track network connectivity to different servers and network devices, and to monitor applications from an external perspective, an eG **external monitor** is required. In the eG license, this maps to a **premium monitor** license. One external monitor license is required for every 50 targets being monitored from an external perspective.

3.1.2 User-based eG Monitoring License

The server-based licensing model is appropriate for applications that are licensed per server – e.g., web servers, databases, J2EE, infrastructure servers, etc. Applications that handle user accesses like Citrix XenApp and Microsoft Remote Desktop Services are often licensed per user/users, rather than per server. To align the monitoring solution's licensing with the application's licensing, eG Enterprise supports the following user-based licensing modes:

- Named user licensing
- Concurrent user licensing

Each of these user-based licensing modes are discussed in the sub-sections that follow.

3.1.2.1 Named User Licensing

By default, the eG monitoring license is server-based. For each server (with a unique nick name), an eG monitor license is required. This licensing model is appropriate for applications that are licensed per server – e.g., web servers, databases, J2EE, infrastructure servers, etc. Applications that handle user accesses like Citrix XenApp and Microsoft Remote Desktop Services are often licensed per user, rather than per server. To align the monitoring solution's licensing with the application's licensing, eG Enterprise supports a named user licensing option. This licensing option is applicable to **Thin Client** and **VDI** environments only. If this option is enabled, you can monitor any number of Citrix XenApp servers, Microsoft Terminal servers, 2X Terminal servers, and/or VDI servers in your environment **without any monitor licenses**, provided the total number of **unique users** who accessed all these servers over the last 90 days is within a licensed limit. If the **Named User** license is violated, the eG agents will stop monitoring all managed **Thin Client** or **VDI** components, and will no longer allow administrators to add/manage more components of such types, until additional **Named User** licenses are purchased. A mix of server-based and user-based licenses can also be used – e.g., all the key server applications like Active Directory, SQL server, Java applications can be monitored with agent/agentless monitor licenses and thin-client servers can be monitored with named user licenses.

Note:

- **Named User** licensing governs only the following component-types: Citrix XenApp, Citrix MF XP, Microsoft Terminal server, 2x Terminal server, VMware vSphere VDI, Citrix XenServer – VDI, Microsoft Hyper-V – VDI, RHEV Hypervisor– VDI, and Oracle VirtualBox. For monitoring all other component types, **Basic** and/or **Premium Monitor** licenses are mandatory! For instance, a typical Citrix environment may comprise of many Citrix XenApp servers, an Active Directory for user authentication, Profile servers, Licensing servers, Web Interface servers, and more. In this case, you can use **Named User licensing** for monitoring the Citrix XenApp servers alone – i.e., you can add any number of XenApp servers for monitoring **without any premium monitor licenses!** For monitoring the Active Directory, Profile, Licensing, and Web Interface servers, you have to have a license per server.
- Once the **Named User** licensing patch is applied, you will not be able to make any changes to the component types mentioned above, using eG's Integration Console plugin. In other words, you can no longer use the Integration Console interface to build new monitoring capabilities or modify the existing models/monitoring capabilities for the following component types - Citrix XenApp, Citrix MF XP, Microsoft Terminal server, 2x Terminal server, VMware vSphere VDI, Citrix XenServer – VDI, Microsoft Hyper-V – VDI, RHEV Hypervisor– VDI, and Oracle VirtualBox.
- If the **Named User license** capability is enabled, then the report by user flag that is available for some of the tests associated with the *VMware vSphere VDI*, *Citrix XenServer – VDI*, *Microsoft Hyper-v – VDI*, *RHEV Hypervisor – VDI*, and *Oracle VirtualBox* component-types, will be disabled. In other words, you cannot override the default status of this flag (which is **Yes**) during test configuration. In this case therefore, all such tests will always report the name of the user who is currently logged into each virtual desktop as the descriptor. Likewise, the aggregate user sessions flag that is available for a few other tests mapped to the above-mentioned component-types will also be disabled if the **Named User** licensing mode is activated. Since this flag is set to **No** by default, these tests will always report a set of metrics for every *username on guestname*; the default status of this flag too cannot be altered using the test configuration web page in the eG administrative interface. In the same manner, the reportbyCLIENTNAME flag that is available for the **Citrix Users** test and the **Citrix Applications** test mapped to the Citrix XenApp server will also be disabled if **Named User licensing** is applied. This means that the detailed diagnosis of these tests will **not include** the clientname column, which displays the host name of the client machine from which users accessed applications on the XenApp server.

3.1.2.1.1 Tracking the Usage of Named User Licenses

To enable users to track the usage of their **Named User** licenses, the **LICENSE USAGE** section of the **LICENSE INFORMATION** page (which appears when the **Click here to get the license details** link in the **LICENSE USAGE SUMMARY** section of the **Admin Home** page is clicked) displays the following against the head, **Named Users** (see Figure 3.1):

- **Allowed:** The total number of **Named Users** who are permitted by the eG license to access the managed **Thin Client** and **VDI** components in the environment;
- **Used**: The total number of unique users who actually accessed the managed **Thin Client** and **VDI** components during the last 90 days. This section additionally indicates the number of user licenses currently being utilized per user type ; this way, you can figure out who is using the maximum number of licenses – is it the VDI users? Citrix users? or Terminal users?

Note that the user types displayed here depend upon the types of components that are managed in your environment – for instance, in Figure 3.1 **you do not see the number of 'Citrix Users' because no Citrix components have been managed in your environment**.

- **Available**: The number of **Named User** licenses that are currently unused (i.e., **Available)**; this is the difference between the number of named users who are **Allowed** to access the managed components and the number of **Used** named user licenses; a value close to 0 indicates that your eG installation is about to violate the **Named User** license, and that you may have to obtain additional **Named User** licenses to avoid such an eventuality.
- **Usage (%)**: The percentage of **Allowed** named user licenses that are currently **Used**; this is a good indicator of how effectively the **Named User** licenses are being utilized by your environment and whether more of these licenses need to be obtained in the immediate future. Against the **VDI Users**, **Terminal Users**, and **Citrix Users** sub-sections, you can view the percentage of total **Allowed** licenses that are currently been utilized by Citrix, Terminal, and VDI users; in the event of repeated violations of the **Named User** license, you can use this break-up to figure out where the user density is high – in Citrix environments? VDI environments? or Microsoft Terminal environments? **Here again, the user types displayed depend upon the types of components that are managed in your environment**.

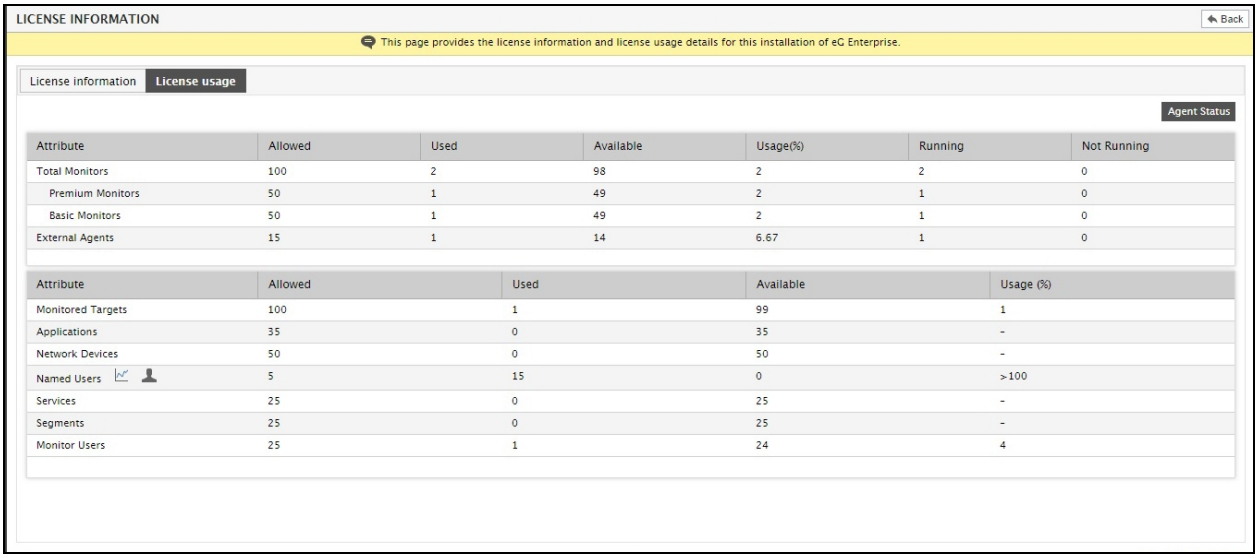

Figure 3.1: The LICENSE USAGE section of the License Information page displaying the usage details of the Named User licenses

To track how the named user licenses are utilized per day, click on the **GRAPH** icon against **Named Users** in 3.1.2. Figure 3.2 will then appear, which graphically depicts the number of users who accessed the managed thin client and VDI components every day during the last 2 weeks (by default). If the eG manager reports a license violation, you can use this graph to quickly figure out how many times in the last 2 weeks the violation has occurred and on which days.

Figure 3.2: The 3D graph depicting daily usage of the Named User licenses

You can change the timeline of this graph by picking a different **Duration**, and you can change the dimension of this graph by selecting a different option from the **Graph** drop-down.

To view the **NAMED USERS REPORT**, click on the icon near **Named Users** in 3.1.2. Figure 3.3 then appears listing the names of the unique users who accessed the managed Citrix XenApp, Microsoft RDS, 2X Terminal, and VDI components.

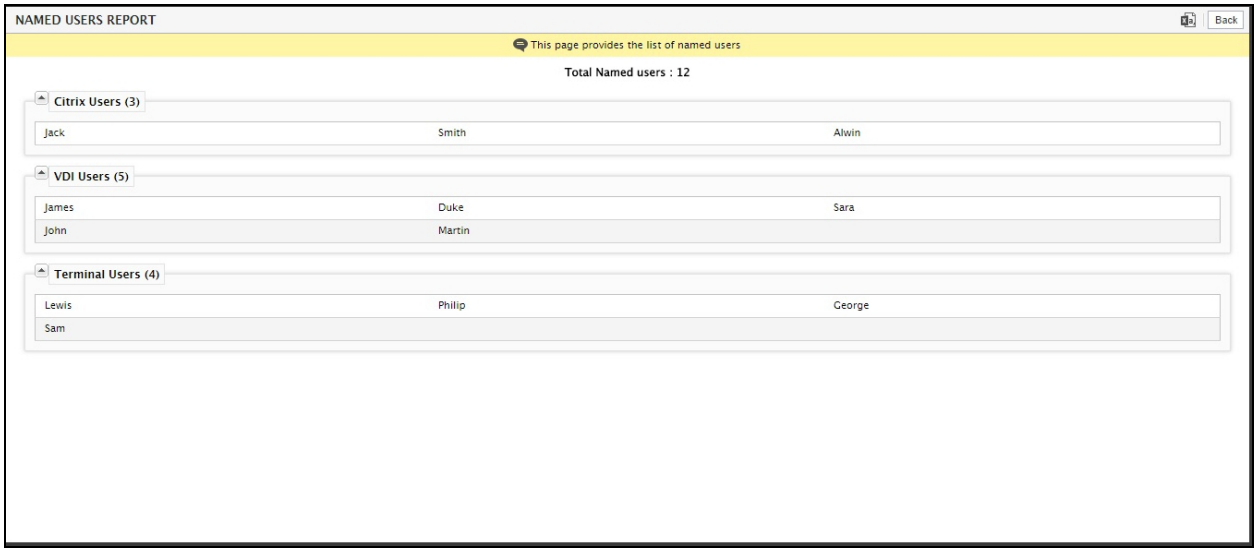

Figure 3.3: The Named Users Report

3.1.2.1.2 How does eG Enterprise Detect and Deal with a Named User License Violation?

If the **Named User** license is enabled, then at the end of every day, the eG Enterprise system automatically computes the total number of 'unique' users who accessed all the managed **Thin Client** and **VDI** components (of the types mentioned previously) in the environment during the last 90 days, and stores this user count in the eG backend. The solution then checks the user count records of the last 14 days for violations. When performing this check, if the solution finds that the total number of unique users on any day during the last 14 days exceeds the licensed number of **Named Users**, then a license violation is registered. If the solution finds that such a violation has occurred only once during the last 14 days, then the next time you log into the eG management console, the following message will appear warning you of the consequences of a continued violation:

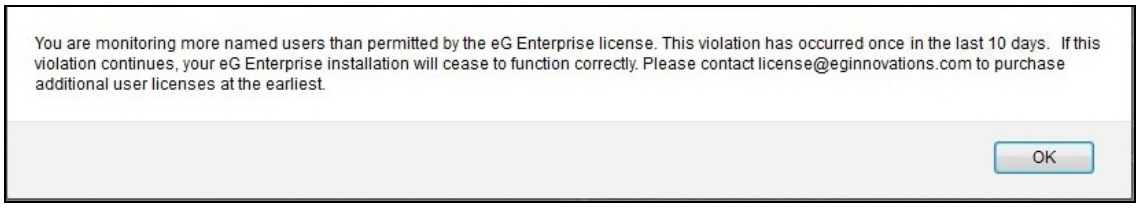

Figure 3.4: The warning message that appears if the Named User license is violated once in the last 14 days

A similar warning message (as depicted by Figure 3.4) will appear for every subsequent violation that is detected in the last 14 days, till the sixth violation. However, if the **Named User** license is violated for the seventh time around in 14 days, any subsequent attempt to login to the eG management console will result in the following message:

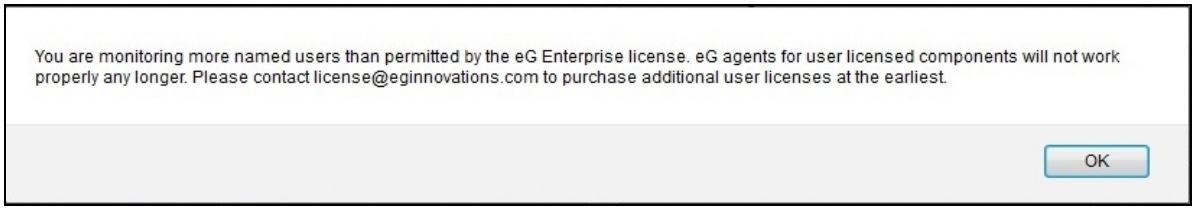

Figure 3.5: The message that appears when the Named User license is violated thrice in 7 days

Within 5 minutes of detection of the seventh violation (in 14 days), the eG agents will stop running the following tests:

- The **application-level** tests mapped to the managed **Citrix XenApp, Microsoft Terminal,** and **2X Terminal** servers;
- All the **inside-view** tests of the VDI components
- A few **outside-view** tests of the VDI components

This implies that the **host-level** tests of these components will continue to run and report metrics.

Moreover, after the seventh violation (in 14 days), you will not be able to add/manage any additional thin client or VDI components (of types mentioned previously) until you obtain additional **Named User** licenses.

The agents will start executing the above-mentioned tests within 5-10 minutes of the application of a new license allowing additional **Named Users**.

3.1.2.1.3 Named User Licensing FAQ

This section answers some of the frequently asked questions on named user licensing

1. **What is named user licensing? Why would I need it?**

   By default, the eG monitoring license is server-based. For each server (with a unique nick name), an eG monitor license is required. This licensing model is appropriate for applications that are licensed per server – e.g., web servers, databases, J2EE, infrastructure servers, etc. Applications that handle user accesses like Citrix XenApp and Microsoft Remote Desktop Services are often licensed per user, rather than per server. To align the monitoring solution's licensing with the application's licensing, eG Enterprise supports a named user licensing option. This licensing option is applicable to **Thin Client**and **VDI** environments only. If this option is enabled, you can monitor any number of Citrix XenApp servers, Microsoft Terminal servers, 2X Terminal servers, and/or VDI servers in your environment **without any monitor licenses**, provided the total number of **unique users** who accessed all these servers over the last 90 days is within a licensed limit.

2. **When is named user licensing ideal and when should I go for server-based licensing?**

   Named user licensing is ideal for environments that have a low user density per server. For example, when client-centric applications like Citrix XenApp were hosted on physical servers, having 50 to 100 users per server was common. In such scenarios, a server-based licensing model was more cost effective because the cost of the monitoring license could be amortized across the users using a server. As these applications have become virtualized, the number of users per server instance (i.e., per virtual machine) has decreased. In such situations, the **Named User** licensing option is more cost-effective.

   Based on the standard pricing of the eG Enterprise monitoring licenses, it is cost-effective to use named user licensing when supporting infrastructures with **less than 30 users** per server instance. For environments with more than 30 users per server instance, the server-based licensing option for the monitoring solution is more cost-effective.

3. **What application types does named user licensing cover?**

   Named User licensing can be used for the following application/virtualization types in eG Enterprise: Citrix XenApp, Citrix MF XP, Microsoft Terminal server, 2x Terminal server, VMware vSphere VDI, Citrix XenServer – VDI, Microsoft Hyper-V – VDI, RHEV Hypervisor– VDI, and Oracle VirtualBox.

4. **What about other application types?**

   All application types, other than the ones mentioned in the response to question 3 above, will use the server-based licensing model only.

5. **Can I use a mix of named user and server-based licensing in my environment?**

   Yes, you can. You can have the Citrix XenApp/Terminal/VDI components in your environment managed using named user licensing and all other applications managed using server-based licensing.

6. **Can I have some XenApp servers governed by named user licensing and some others using the**

server-based model?

No, this is not possible. If named user licensing is enabled, then all Citrix XenApp/Terminal/VDI components that are managed in your environment will use this licensing model only.

7. **Can I monitor the XenApp servers using user-based licensing and the VDI servers using the server-based model?**

   No, this is not possible. Once your eG installation enables the **Named User Licensing** model, it automatically applies to all the managed Citrix XenApp/Terminal/VDI components in your environment.

8. **I have managed a few Exchange servers in my environment. Later, I virtualized the Exchange servers, and enabled the user-based licensing capability of my eG installation. Will this licensing model apply to the Exchange servers as well?**

   No, it will not. The Exchange servers will continue to consume **premium monitor licenses**.

9. **Can a single eG manager handle both named user licenses and server-based licenses?**

   Yes. A single eG manager is capable of handling a mix of user-based and server-based licenses.

10. **How do I know whether/not named user licensing is enabled for my eG installation?**

    To know whether named user licensing is enabled or not, do the following:

    - Login to the eG administrative interface as admin with password admin.
    - Click on the **Click here to know more** link in the **License Usage Summary** section of the **Admin Home** page.
    - When the **LICENSE INFORMATION** page appears, click the **License Usage** tab page. Figure 3.6 appears. If you find a **Named Users** entry in Figure 3.6 with a positive, non-zero value in the **Allowed** column, it is a clear indicator that the named user licensing is enabled for your eG installation.
    - If you do not find such an entry in the **License Usage** tab page of Figure 3.6, you can be rest assured that the named user licensing capability is not enabled for your eG installation.

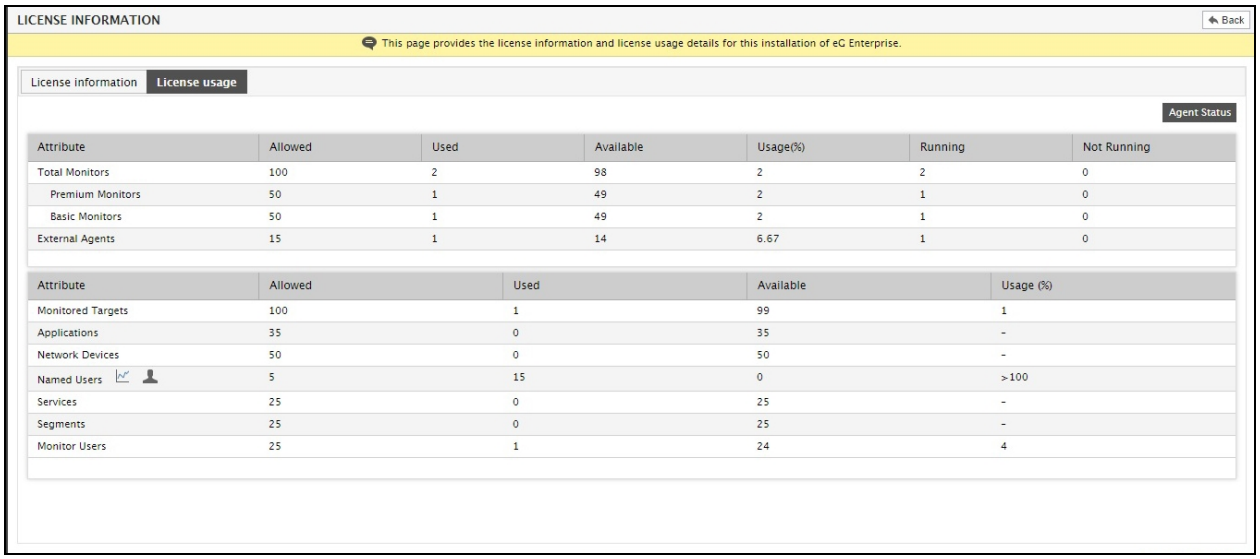

Attribute	Allowed	Used	Available	Usage(%)	Running	Not Running
Total Monitors	100	2	98	2	2	0
Premium Monitors	50	1	49	2	1	0
Basic Monitors	50	1	49	2	1	0
External Agents	15	1	14	6.67	1	0

Attribute	Allowed	Used	Available	Usage (%)
Monitored Targets	100	1	99	1
Applications	35	0	35	-
Network Devices	50	0	50	-
Named Users	5	15	0	>100
Services	25	0	25	-
Segments	25	0	25	-
Monitor Users	25	1	24	4

Figure 3.6: The Named Users entry in the License Usage tab page

11. **How do I enable named user licensing?**

    If your current eG installation does not manage any Citrix XenApp/Terminal/VDI components, then follow the procedure detailed below to enable named user licensing.

    - Contact sales@eginnovations.com with your request.
    - A new license with the **Named User** licensing capability will be generated and sent to you.
    - Upon receipt of the eG license, stop the eG manager, and copy the license to the **<EG_INSTALL_DIR>\bin** directory (on Windows; on Unix, this will be the **/opt/egurkha/bin** directory on Unix).
    - Restart the eG manager.

    If you have already managed a few Citrix XenApp/Terminal/VDI components in your eG installation, then contact sales@eginnovations.com for the way forward.

12. **What happens once the named user licensing capability is enabled?**

    If the **Named User** license is enabled, then at the end of every day, the eG Enterprise system automatically computes the total number of 'unique' users who accessed all the managed **Thin Client** and **VDI** components (of the types mentioned in the response to question 3) in the environment during the last 90 days, and stores this user count in the eG backend.

13. **If a single user logs into the three different servers, what would be the unique user count that eG reports?**

    Regardless of the number of servers accessed by a single user, eG Enterprise will report the unique user count as `1`.

14. **Can I detect a potential named user license violation, before it occurs?**

Yes, you can. For this, you need to access the **License Usage** tab page of the **LICENSE INFORMATION** page that appears when the **Click here to know more** link in the in the **License Usage Summary** section of the **Admin Home** page (see Figure 3.7).

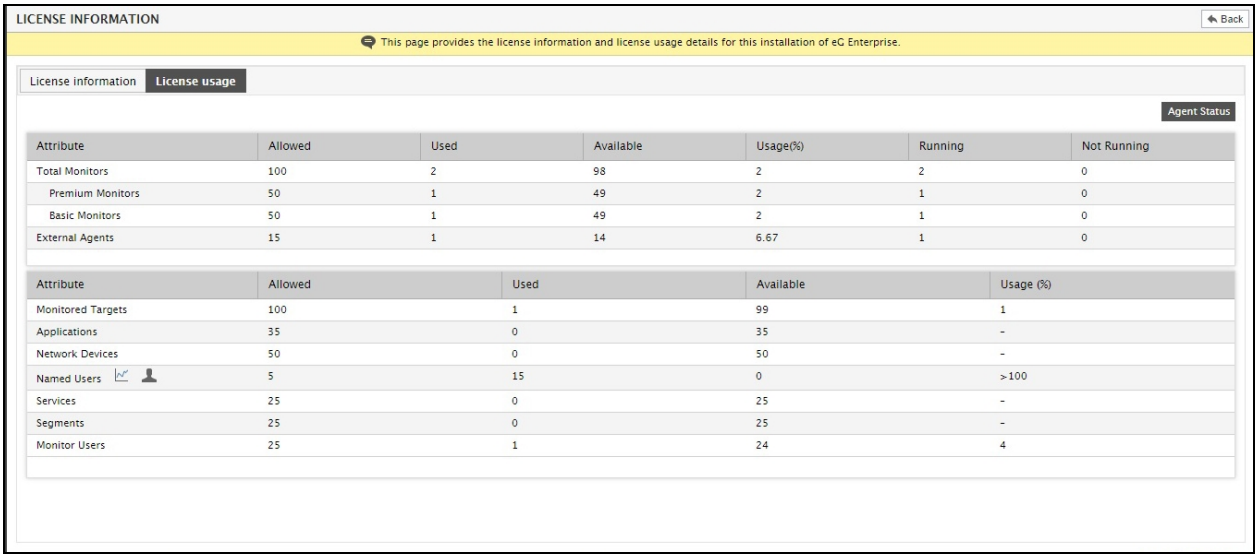

Attribute	Allowed	Used	Available	Usage(%)	Running	Not Running
Total Monitors	100	2	98	2	2	0
Premium Monitors	50	1	49	2	1	0
Basic Monitors	50	1	49	2	1	0
External Agents	15	1	14	6.67	1	0

Attribute	Allowed	Used	Available	Usage (%)
Monitored Targets	100	1	99	1
Applications	35	0	35	-
Network Devices	50	0	50	-
Named Users	5	15	0	>100
Services	25	0	25	-
Segments	25	0	25	-
Monitor Users	25	1	24	4

Figure 3.7: Tracking named user license usage

The **License Usage** tab page helps you continuously track named user license usage and proactively detect a potential violation. This section reveals the maximum number of **Named Users** that the eG installation **Allows**, and also reports the number and percentage of named user licenses that are currently utilized. By comparing the **Allowed** limit with the count of **Used** named user licenses, you can quickly figure out if license usage is optimal or is close to exhaustion. For instance, if the **License Usage** tab page reveals that 8 out of the 10 named user licenses granted to an eG installation have been used up, it implies that nearly 80% of the licenses are in use. This is a clear warning sign of an impending license violation!

You can also track how the named user licenses are used per day by clicking the **GRAPH** icon against **Named Users** in Figure 3.7. Figure 3.8 will then appear.

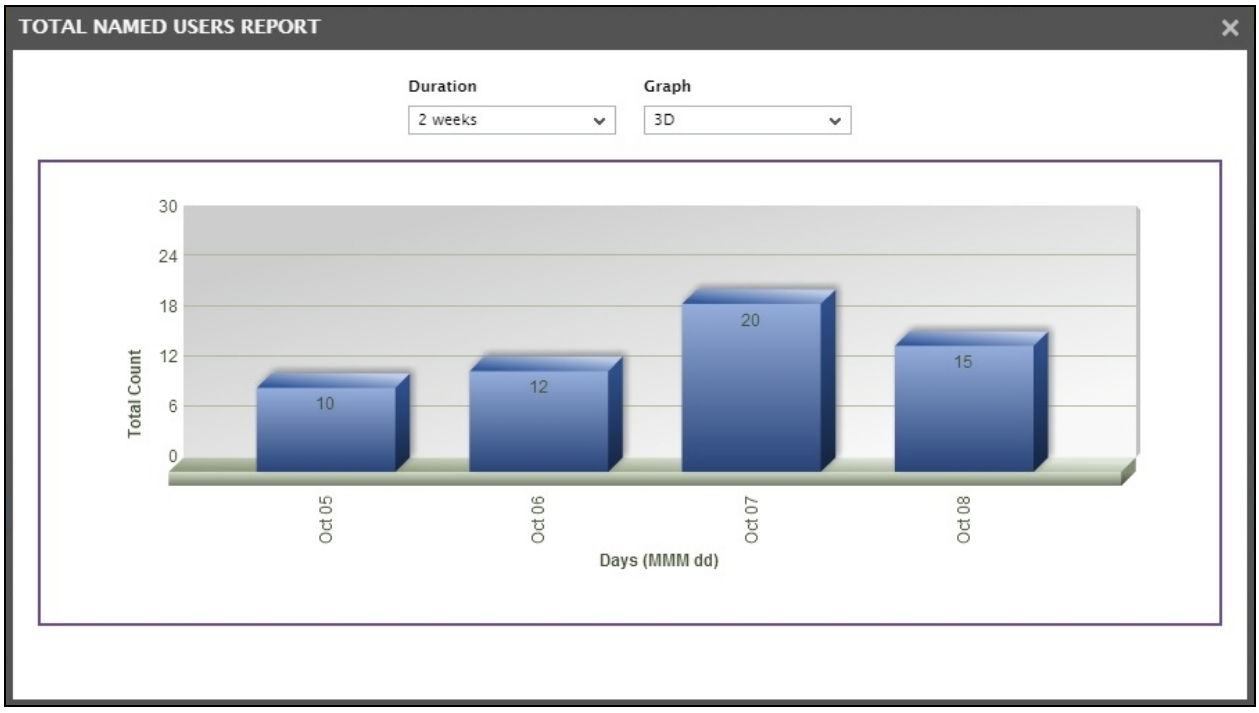

Figure 3.8: The 3D graph depicting daily usage of the Named User licenses

15. **What happens when the named user license is violated?**

If the **Named User** license is enabled, then at the end of every day, the eG Enterprise system automatically computes the total number of 'unique' users who accessed all the managed **Thin Client** and **VDI** components (of the types mentioned previously) in the environment during the last 90 days, and stores this user count in the eG backend. The solution then checks the user count records of the last 14 days for violations. When performing this check, if the solution finds that the total number of unique users on any day during the last 14 days exceeds the licensed number of **Named Users**, then a license violation is registered. If the solution finds that such a violation has occurred only once during the last 14 days, then the next time you log into the eG management console, the following message will appear warning you of the consequences of a continued violation:

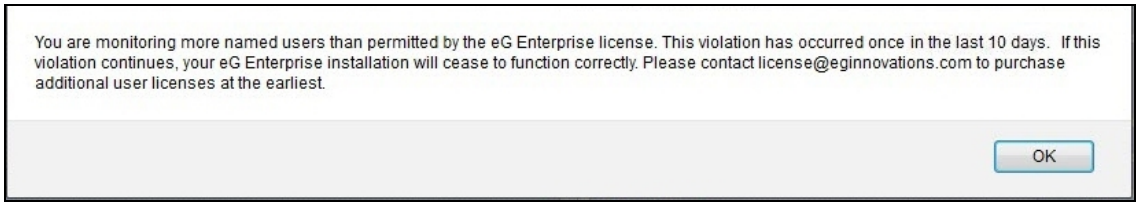

Figure 3.9: The warning message that appears if the Named User license is violated once in the last 14 days

A similar warning message (as depicted by Figure 3.9) will appear for every subsequent violation that is detected in the last 14 days, till the sixth violation. However, if the **Named User** license is violated for the seventh time around in 14 days, any subsequent attempt to login to the eG management console will result in the following message:

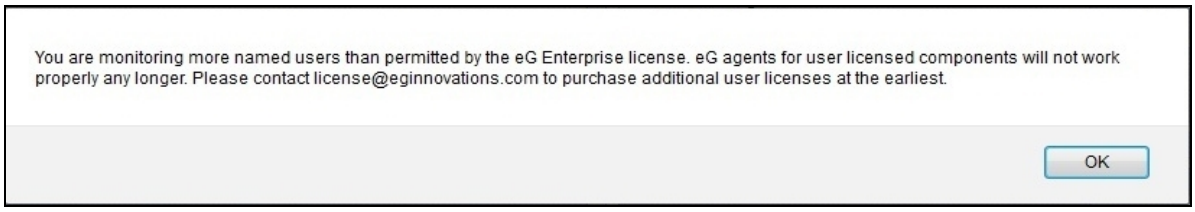

Figure 3.10: The message that appears when the Named User license is violated thrice in 7 days

Within 5 minutes of detection of the seventh violation (in 14 days), the eG agents will stop running the following tests:

- The **application-level** tests mapped to the managed **Citrix XenApp, Microsoft Terminal,** and **2X Terminal** servers;
- All the **inside-view** tests of the VDI components
- A few **outside-view** tests of the VDI components

This implies that the **host-level** tests of these components will continue to run and report metrics.

Moreover, after the seventh violation (in 14 days), you will not be able to add/manage any additional thin client or VDI components (of types mentioned previously) until you obtain additional **Named User** licenses.

If any of these messages pop-up when you login to the eG management console, then, you can navigate to the **License Usage** tab page to view more details of the violation.

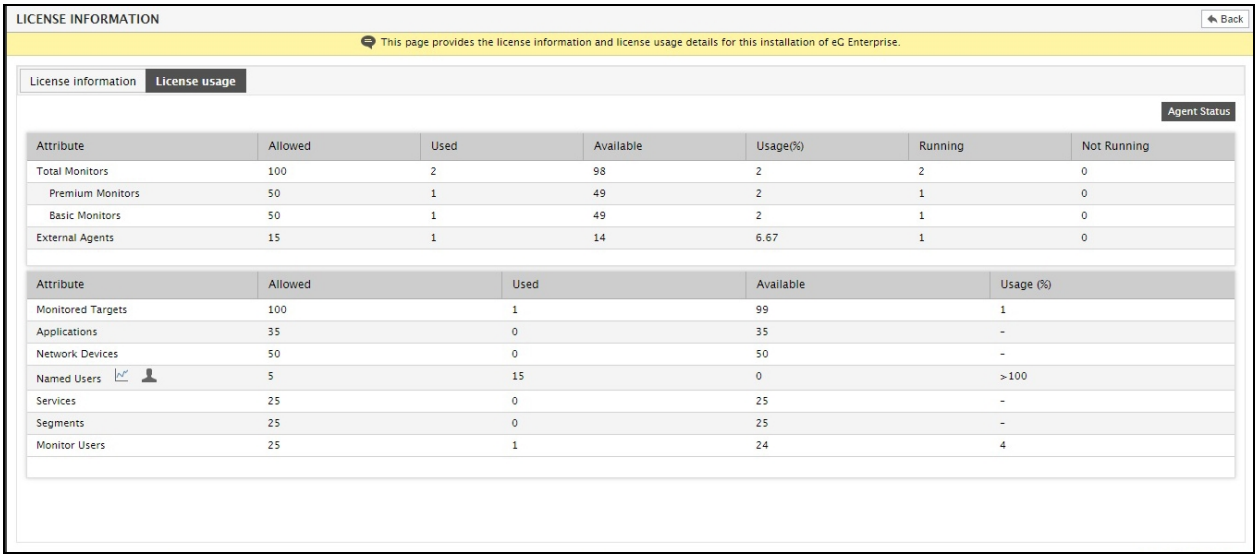

Figure 3.11: The License Usage tab page revealing a named user license violation

In the example of Figure 3.11, you can see that while only **5** named users were **Allowed** by the eG license, over **15** unique users have actually accessed the managed thin-client/VDI components during the last 90 days. Since the environment is supporting more than the allowed number of distinct users, the **Usage (%)** has been automatically set to **<100** in the case of this example (see Figure 3.11).

16. **If the user license is violated for the seventh time in 14 days, then, do all tests mapped to the**

managed thin client and VDI components stop running automatically?

No. The eG agents will stop executing the following tests if 7 out of 14 user count records report violations.

- The **application-level** tests mapped to the managed **Citrix XenApp, Microsoft Terminal,** and **2X Terminal** servers;
- All the **inside-view** tests of the VDI components
- A few **outside-view** tests of the VDI components

This implies that the **host-level** tests of these components will continue to run and report metrics.

17. **How soon will components start reporting metrics once additional named user licenses are applied?**

    The agents will start executing the tests mentioned in response to Question 16 above, within 5-10 minutes of the application of a new license allowing additional **Named Users**.

3.1.3 Concurrent User Licensing

The concurrent user licensing model is ideal for environments where a large number of users access the servers every day, but only a small subset of this user population accesses the servers concurrently – i.e., at the same time. For example, in a university, you could have thousands of students coming in every day; but, every time a class is in session, you will have a few students accessing their desktops simultaneously to attend the class. If you opt for the named user licensing model here, you will have to obtain licenses for all the users who log into their desktops each day; this may prove to be expensive. The concurrent user licensing model on the other hand, will be more cost-effective in such environments, as you will have to obtain a license for only those users who access their desktops concurrently. Like named user licensing, **Concurrent user** licensing too is applicable to **Thin Client** and **VDI** environments only. If this option is enabled, you can monitor any number of Citrix XenApp servers, Microsoft RDS servers, 2X Terminal servers, and/or VDI servers in your environment **without any monitor licenses**, provided the maximum number of users who access these servers every day during the last 14 days is within a stipulated limit. At configured intervals (default: 30 minutes), the eG manager automatically computes the total number of users who accessed all the managed **Thin Client** and **VDI** components in the environment during the last 30 minutes (by default), and stores this user count in the eG backend. The user count records so collected during the course of the day are compared at the end of the day to identify the maximum number of concurrent users for that day. This maximum number is compared with the licensed number of concurrent users to capture violations (if any). If this license is violated in any 7 out of the last 14 days, the eG agents will stop running the following tests:

- The **application-level** tests mapped to the managed **Citrix XenApp, Microsoft Terminal,** and **2X Terminal** servers;
- All the **inside-view** tests of the VDI components
- A few **outside-view** tests of the VDI components

This implies that the **host-level** tests of these components will continue to run and report metrics.

Moreover, after the seventh violation (in 14 days), you will not be able to add/manage any additional thin client or VDI components (of types mentioned previously) until you obtain additional **concurrent user** licenses.

The agents will start executing the above-mentioned tests within 5-10 minutes of the application of a new license allowing additional **Concurrent Users**.

3.2 Viewing License Information

Figure 3.12 shows the license information of the eG manager for a particular installation. To view the license information of the eG manager for a particular installation, click on the **Click here to know more** link in the **License Usage Summary** section of the **Admin Home** page.

The **LICENSE INFORMATION** opens with the **License Information** tab page selected by default (see Figure 3.12).

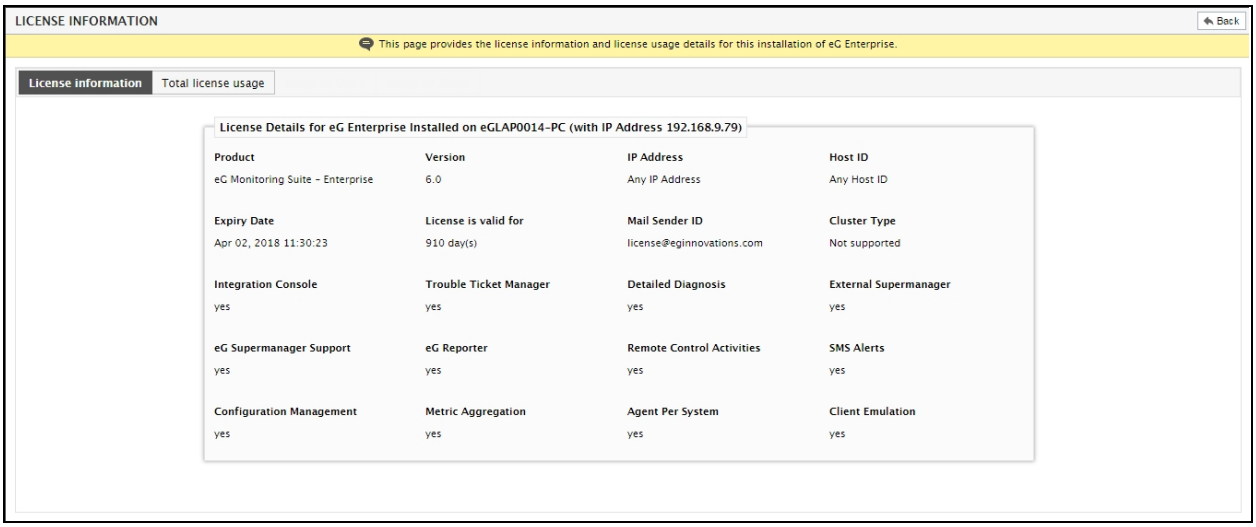

Figure 3.12: License information

The IP address, if specified, restricts the eG manager to a specific host. The Host ID, if specified, restricts the eG manager to a host that has a specific host ID. On Unix systems, run the command **hostid** to determine the host id of the system. The output of the command must match the host ID specified in the license for the eG manager to start. On Windows systems, look for the physical address specification in the output of the **ipconfig /all** command. The host ID specified in the license must match one of the physical addresses of the host (ignore any dashes (-) in the physical address).

Besides monitoring purposes, customers have the option of just using the eG Enterprise system as a reporting engine to obtain reports on historical trends and events. The **Product** name in case of a 'reporting-only' installation will read **eG Monitoring Suite - Reporter Only**. Offered as a lower cost option, eG Enterprise with the reporting-only capability still supports agent-based and agentless monitoring of over 85 different applications and seven different operating systems. The table below lists the key differences between a normal eG manager installation and a "reporting-only" installation.

	Normal manager	Reporter-only manager
1.	Users who are assigned the Admin role can	Users who are assigned the Admin role can login

	Normal manager	Reporter-only manager
	login to the eG administrative interface and make configuration changes.	to the eG administrative interface and make configuration changes.
2.	By default, users who are assigned Monitor, Admin, ServerAdmin, or SuperMonitor roles have access to the eG monitor interface. Similarly, if you create a new role with monitoring privileges, then all users who are assigned that role will be able to access the eG monitoring console.	No user can access the eG monitoring console; therefore, all the information that is typically available in the eG monitoring console, such as, the current problem list, alert history, component/service/segment/zone state, quick insight, live graphs, knowledge management, service and virtual topology previews, layer model representation, measurement data etc., will not be available to any user.
3.	When a user who is assigned the Monitor/Supermonitor role logs into the eG manager, the eG monitoring console automatically opens.	When a user who is assigned the Monitor/Supermonitor role logs into the eG manager, the eG Reporter console automatically opens.
4.	The eG manager can send out email/SMS alerts of issues to configured users.	The eG manager cannot send out email/SMS alerts.
5.	Normal alerts, unknown alerts, heartbeat mails, escalated alerts, and shift-based alerts, if enabled, will be sent.	Normal alerts, unknown alerts, heartbeat mails, escalated alerts, and shift-based alerts, even if configured, will not be sent by the eG manager.
6.	Since email/SMS alerts can be sent, configurations specific to email/SMS alerting such as alarm subject customization, alarm content definition, when the alerts are to be sent and how, etc., can be configured and enforced using the **MAIL SETTINGS** page in the eG admin interface.	Since email/SMS alerts cannot be sent, changes that you make to the look and feel of these alerts will not be effected.
7.	Detailed diagnosis, if enabled, can be accessed from the monitoring and reporter consoles.	Detailed diagnosis, even if enabled, will not be available for Snapshot reports.
8.	The settings defined in the **MONITOR SETTINGS** page of the eG administrative interface apply to the eG monitor and the eG Reporter interfaces.	Of the settings available for configuration in the **MONITOR SETTINGS** page, only the following settings are relevant to the Reporter-only manager: • the default language setting • the default date format setting • the timescale for the reporter graphs • the dimension of the reporter graphs (whether 3D or 2D) • the depth of lines in a 3D graph

	Normal manager	Reporter-only manager
		• Whether to show the daywise distribution reports (in the reporter home page) in minutes or percentage.
9.	A user with access to the eG monitor, reporter, and admin interfaces, can change the skin color of any/all the interfaces.	Since no user has access to the eG monitor interface, he/she can change the skin color of the Admin and/or Reporter interfaces only.
10.	A user with access to the eG monitoring console can set a default home page for the console.	Since no user has access to the eG monitor interface, none of the users registered with the Reporter-only manager will have the option of setting a default home page for the monitoring console.
11.	Logos can be configured for admin, monitor, and reporter interfaces.	Logos can be configured for all interfaces, but the monitor interface-specific configuration will not be considered, as no user can access the eG monitoring console.
12.	Administrators can generate audit log reports revealing the user activities on all eG interfaces.	Administrators cannot generate audit log reports for accesses to the eG monitor interface, as no user can access the monitor interface.

Besides the above, the license also governs the following additional eG features:

- Integration of the eG manager with external trouble ticketing systems
- Support for a supermanager that is empowered to handle multiple managers in an environment
- Support for an optional detailed diagnosis capability
- **The agent per system capability**: By default, if a host has multiple IP addresses, the eG Enterprise system requires one agent license for each IP address that is managed internally. Likewise, if multiple nicknames are used for the same IP address, a separate internal agent license is used for each unique nickname that has been specified. In many large environments, a single server has many IP addresses, each with different nicknames. The agent per system capability is intended to optimize the internal agent license usage in such large infrastructures. If this capability is enabled by the eG license, the administrator has the option of overriding the default eG agent licensing policy. For example, suppose a host A has two IP addresses 192.168.10.7 and 10.10.10.1, and that the first IP address 192.168.10.7 has already been managed in the eG Enterprise system. When adding the second IP address, 10.10.10.1, the administrator has the option of overriding eG's default internal agent licensing policy - in this example, the administrator can indicate that the internal agent for the IP address 10.10.10.1 is actually the one that is already associated with the IP address 192.168.10.7. By doing so, the administrator can ensure that a single agent license is sufficient to manage all the IP addresses and applications executing on a host.
- The eG Integration Console (for extending eG Enterprise to monitor custom applications)
- The eG manager's ability to control the agents from remote locations (Remote Control Activities)
- The eG Reporter (for providing a variety of useful reports)

- The license also controls the capability of the eG agents to run client emulation tests. eG Enterprise supports integration with two client emulation tools, namely, CitraTest and QA Wizard. These tools allow administrators to record user transactions to an application, and configure the eG agent to replay the recorded transaction to extract the required measures. For more details about the client emulation tools, refer to *The eG Client Emulation Guide*. eG agents will be able to run the emulated tests only if the **Client Emulation** flag in the eG license is set to **Yes**.
- The SMS alerting capability of the eG manager
- To ensure high availability of the eG monitoring solution, eG Enterprise offers a redundant manager option wherein a secondary manager can act as an active standby for the primary manager. This option is governed by the **Cluster Type** condition in the eG license. If the **Cluster Type** condition contains the value *Not Supported*, it indicates that the current installation of eG Enterprise supports a single manager only. If **Cluster Type** is set to *Active-Active*, then it indicates that manager redundancy has been enabled for that eG installation. A cluster can have only a single **primary manager** and a single **secondary manager.** An *Active-Active* cluster is one where both the primary and secondary managers can both have agents reporting measures to them during normal operation.

This page also depicts the number of days left for the license to expire (see Figure 3.12).

If you now click on the **Total license Usage** tab page in the **LICENSE INFORMATION** page, you will be able to evaluate whether the eG license usage is optimal or is excessive, and accordingly decide on future license requirements (see Figure 3.13).

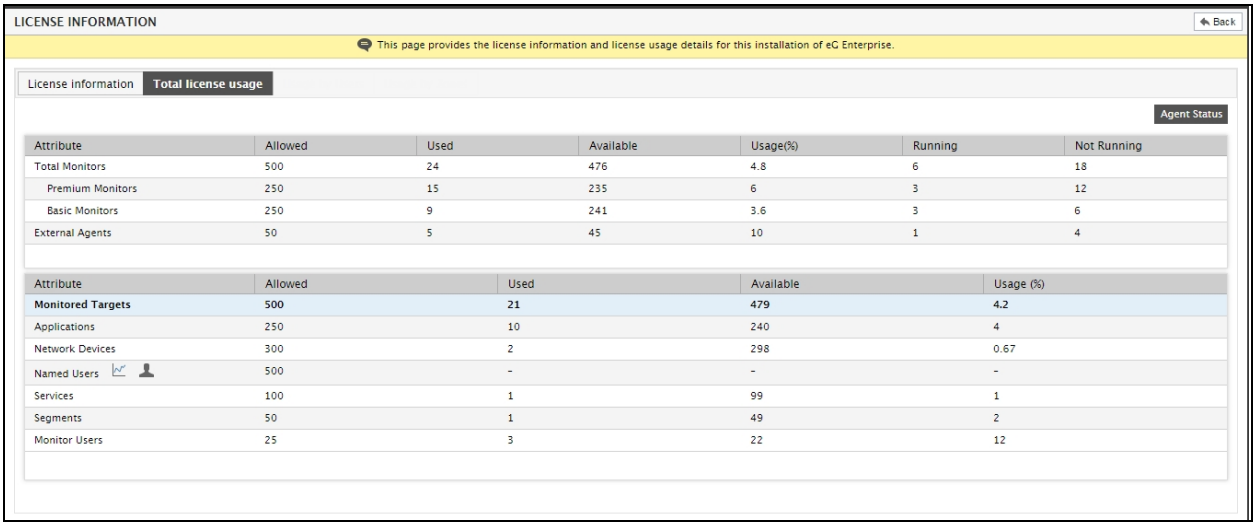

Attribute	Allowed	Used	Available	Usage(%)	Running	Not Running
Total Monitors	500	24	476	4.8	6	18
Premium Monitors	250	15	235	6	3	12
Basic Monitors	250	9	241	3.6	3	6
External Agents	50	5	45	10	1	4

Attribute	Allowed	Used	Available	Usage (%)
Monitored Targets	**500**	**21**	**479**	**4.2**
Applications	250	10	240	4
Network Devices	300	2	298	0.67
Named Users	500	-	-	-
Services	100	1	99	1
Segments	50	1	49	2
Monitor Users	25	3	22	12

Figure 3.13: License usage

- The **Total Monitors** listing indicates the total number of **basic** and **premium** monitors that the current installation of eG Enterprise is allowed to use, the total number of such monitors that are currently utilized, and the overall usage percentage. The number of **Monitors** running and not running is also indicated.
- The total number of **Premium Monitors** that an eG installation is allowed to use is automatically computed as the difference between the **Total Monitors** and **Basic Monitors**. For instance, if the eG license allows 20 monitors totally and 7 basic monitors, then 13 will be automatically set as the maximum number of

premium monitors that can be configured in the environment. If required, the customer can even have 19 basic monitors, reserving 1 premium monitor for the 1 external agent that is a must for every eG installation.

- The count and percent usage of **Named User** or **Concurrent User licenses** (if any).
- The license specifies the maximum number of **Services** that the manager can support, and the number of services currently managed;
- In addition, the license imposes a ceiling on the number of **Applications** and the number of unique IPs (refers to the number of **Monitored Targets**) that can be monitored. Therefore, if a user monitors 4 applications executing on a single host, the eG Enterprise system will count the number of applications as 4 and the number of monitored targets as 1.
- Restrictions on the number of **Users** permitted, the number of **Network Devices** that can be monitored, and the maximum number of topology **Segments** that can be configured are also specified in the license. The number of network devices and topology segments currently monitored are also displayed as part of the **Total license usage**.

The **Total license usage** tab page not only indicates the number of licenses utilized, but also leads you to the exact components/network devices/segments/services/users (as the case may be) that are using these licenses. To get to this usage break-up, click on the **Used** licenses count of any **Attribute** listed in the **Total license usage** tab page of Figure 3.13. For instance, clicking on the **Used** count of the **Premium Monitors** license attribute in Figure 3.13 will lead you to Figure 3.14, where you can view the complete list of components that are utilizing the premium monitor licenses.

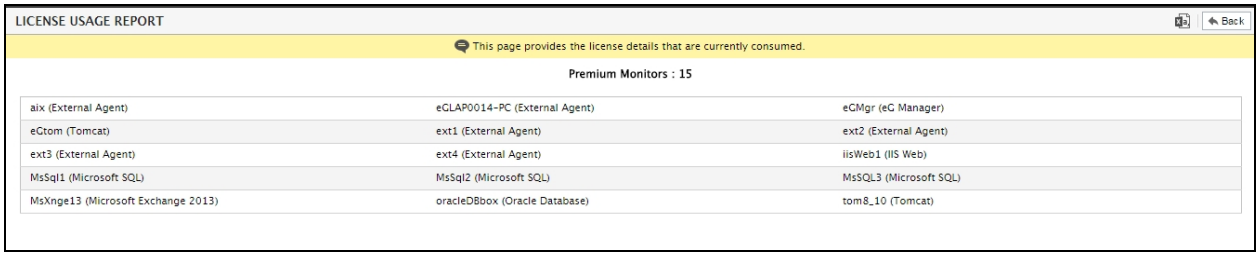

Figure 3.14: Drilling down to the components using the Premium Monitor licenses

Clicking on the **Agent Status** button in Figure 3.13 will lead you to the page that provides you with the status of the agents of a selected type.

Managed service providers (MSPs) deploy monitoring in a multi-tenant configuration, so multiple customers are hosted using a common eG management server. In such environments, MSPs need to be able to track monitoring license usage for each customer, so that usage billing can be done accordingly. For this purpose, the **LICENSE INFORMATION** page can be configured to report license usage at a granular level – i.e., report the number of licenses used by each user and zone configured in the management console. Besides being useful for billing, this information is also useful for planning future licensing requirements.

To enable user-wise license usage tracking, the MSP can do the following:

1. Edit the **eg_services.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory).
2. Set the **UserWiseCertReport** flag in the **[ZoneAndUserCertUsage]** section of the file to **Yes**.
3. Save the file.

Doing so will append a **Usage by Users** tab page to the **LICENSE INFORMATION** page (see Figure 3.15). This tab page will list all the users currently registered with the eG Enterprise system along with the count and percentage of licenses used by each user.

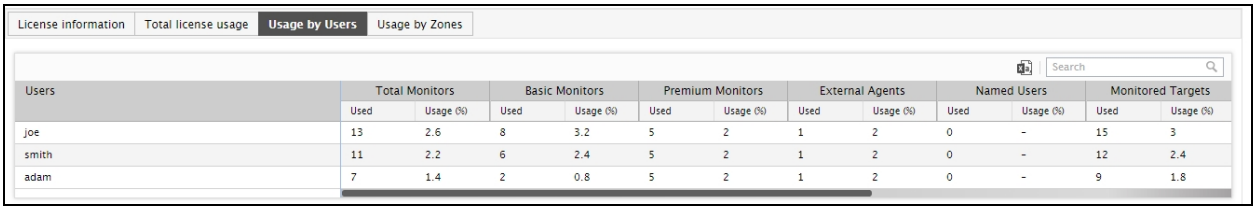

Users	Total Monitors		Basic Monitors		Premium Monitors		External Agents		Named Users		Monitored Targets	
	Used	Usage (%)	Used	Usage (%)	Used	Usage (%)	Used	Usage (%)	Used	Usage (%)	Used	Usage (%)
joe	13	2.6	8	3.2	5	2	1	2	0	-	15	3
smith	11	2.2	6	2.4	5	2	1	2	0	-	12	2.4
adam	7	1.4	2	0.8	5	2	1	2	0	-	9	1.8

Figure 3.15: Viewing license usage by users

To enable the **LICENSE INFORMATION** page to provide a zone-wise breakup of license usage, do the following:

1. Edit the **eg_services.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory).
2. Set the **ZoneWiseCertReport** flag in the **[ZoneAndUserCertUsage]** section of the file to **Yes**.
3. Save the file.

Doing so will append a **Usage by Zones** tab page to the **LICENSE INFORMATION** page (see Figure 3.16). This tab page will list all the zones that have been currently configured and the count and percentage of licenses used by the components in each zone.

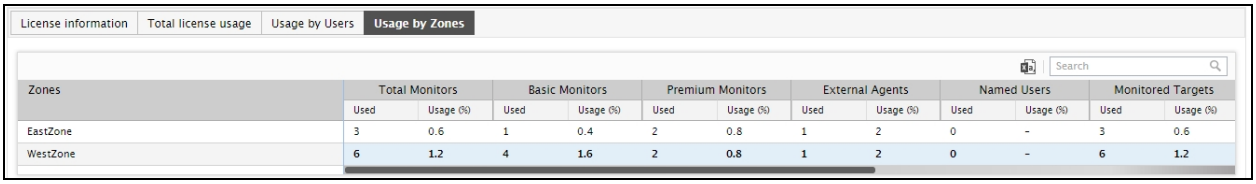

Zones	Total Monitors		Basic Monitors		Premium Monitors		External Agents		Named Users		Monitored Targets	
	Used	Usage (%)	Used	Usage (%)	Used	Usage (%)	Used	Usage (%)	Used	Usage (%)	Used	Usage (%)
EastZone	3	0.6	1	0.4	2	0.8	1	2	0	-	3	0.6
WestZone	6	1.2	4	1.6	2	0.8	1	2	0	-	6	1.2

Figure 3.16: Viewing license usage by zones

4

Getting Familiar with the Admin Interface

An important step in installing and configuring the eG manager and agent is to administer the eG Enterprise system. During this process, an administrator configures which components are monitored by eG Enterprise, how the components are interconnected, what tests are executed for each component, and how the measurements reported by the tests are to be interpreted. The administrator is also responsible for determining which users are allowed access to the eG Enterprise system. This chapter describes the various functions that an eG administrator can perform.

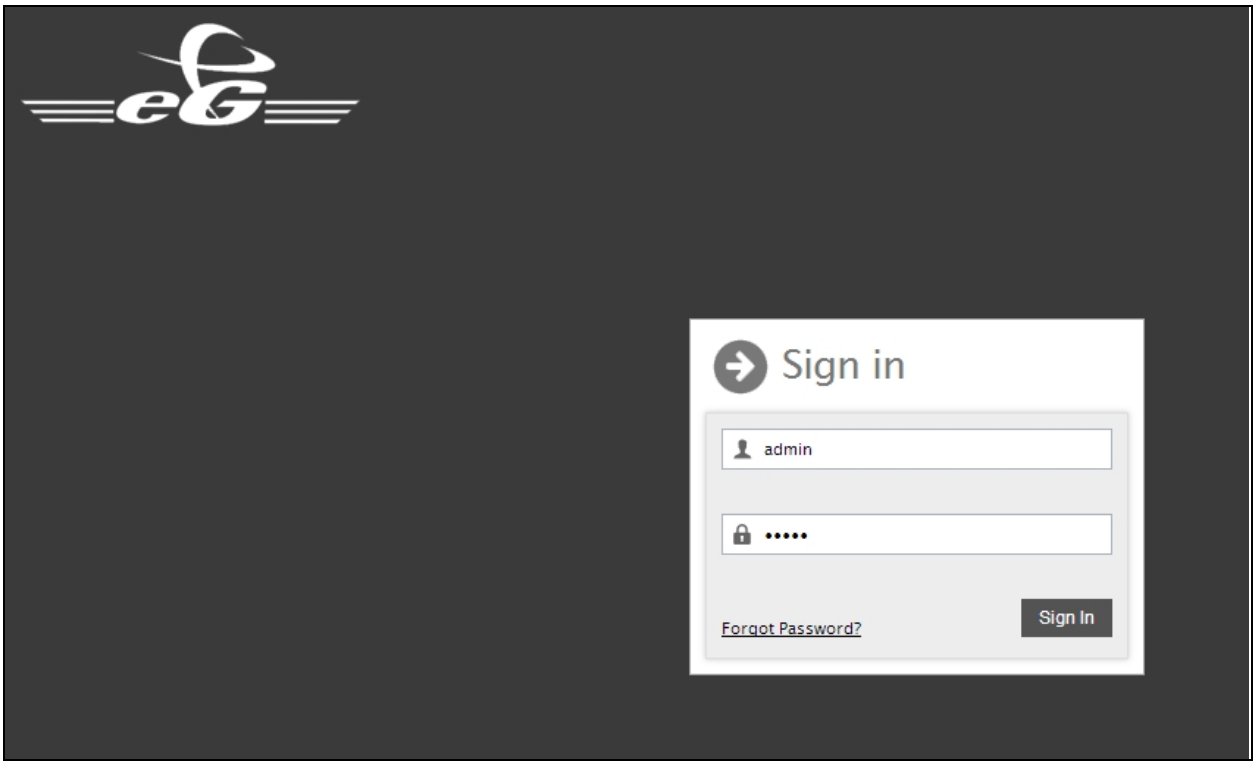

Figure 4.1: The eG login screen

To access eG Enterprise using a web browser, connect to the URL, if the manager is configured without SSL. If the manager is configured with SSL, connect to the URL. Figure 4.1 shows the eG login window. The user has to login from this window in order to access eG Enterprise. The eG Enterprise system is predefined with a default administrator account with a login of admin and password admin. An administrator can also login using any other **Username**, provided the role assigned to that user name allows him/her access to the eG administrative interface. To know more about user roles and user profile creating in the eG Enterprise system, refer to Section **6.1** and Section **6.3** of this manual.

Note:

While specifying the URL, please take care of the following aspects:

- If the host name was provided when installing the manager, use this name (and not the IP address) for accessing the user interface via the web browser.
- If the host name is provided, make sure that forward and reverse lookups for this name are enabled via the DNS service in the target environment.

If an administrator forgets the login **Password**, he/she can click on the **Forgot Password** link in Figure 4.1. Doing so invokes Figure 4.2 wherein the administrator would have to provide the **Username** for which the password details are required, and then click the **Get password** button to retrieve the password.

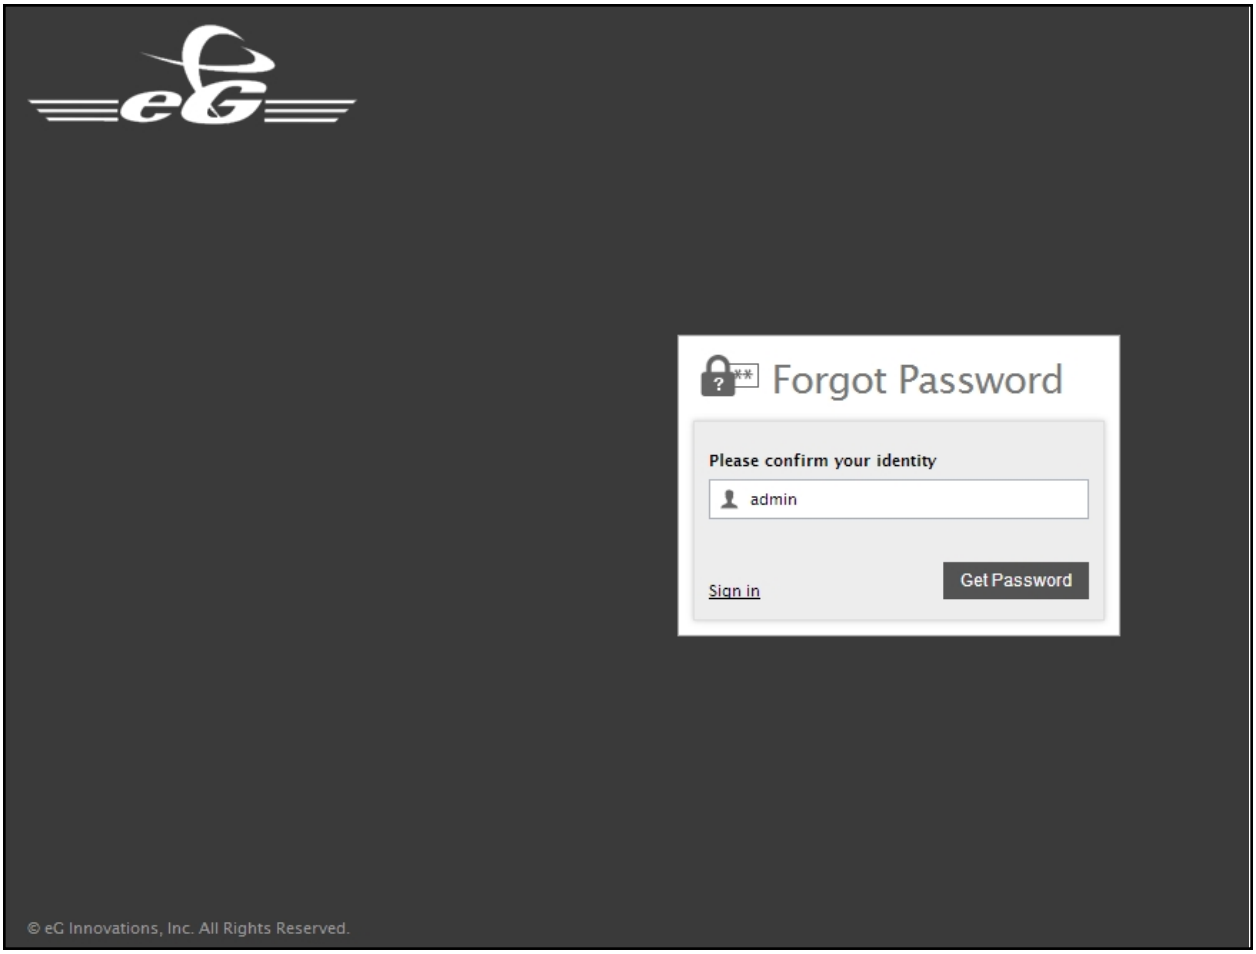

Figure 4.2: Retrieving password

If the **Username** specified is valid, then the password will be emailed to the user with the given **Username**.

Note:

The eG manager will be able to send password details by mail to a user, only if:

- The **Username** specified in Figure 4.2 has been configured to receive email alerts, and a valid email ID has been assigned to the user; please refer to Section **6.3** of this manual to know how to configure a user profile to receive email alerts of issues.

- The mail server has been properly configured to handle eG alerts. In the **MAIL SERVER SETTINGS** page of the eG administrative interface a valid mail host and **eG Administrator maild ID** should have been configured. Please refer to Section **5.2.1** of this manual for more details with regard to the same.

Once the administrator has logged in, the eG manager may indicate an expiry of the license as in Figure 4.2.

Note:

Sometimes, users may not want to login via the login interface provided by eG Enterprise. For instance, if a user is already logged into a web portal, he/she may not want to login again to gain access to the eG user interface; instead, they may want to directly connect to the eG management console from the portal. To access the eG web-based interface without logging in, you can use the following URL:

http://<eGmanagerIP>:<eGmanagerport>/final/servlet/com.egurkha.EgLoginServlet?uname=<username>&upass=<password>&accessKey=eGm0n1t0r

If the eG manager you want to connect to is SSL-enabled, then, use the following URL:

https://<eGmanagerIP>:<eGmanagerport>/final/servlet/com.egurkha.EgLoginServlet?uname=<username>&upass=<password>&accessKey=eGm0n1t0r

Make sure that you configure the URL with the correct <eGmanagerIP> and <eGmanagerport>. Also, ensure that the name of a user with rights to access the eG management console is provided against uname. You can, if you so need, provide the password of the given user against upass, or can leave the password blank. If the URL is not configured with a password, the eG Enterprise system will automatically pick the password that corresponds to the specified uname from the database. However, note that the 'accessKey' provided in the URL should not be changed.

A sample URL (with a blank password) is provided below:

https://192.168.10.21:7077/final/servlet/com.egurkha.EgLoginServlet?uname=admin&upass=&accessKey=eGm0n1t0r

4.1 The Admin Home Page

The first page that will be displayed on logging into the eG administrative interface is the **Admin Home** page (see Figure 4.3).

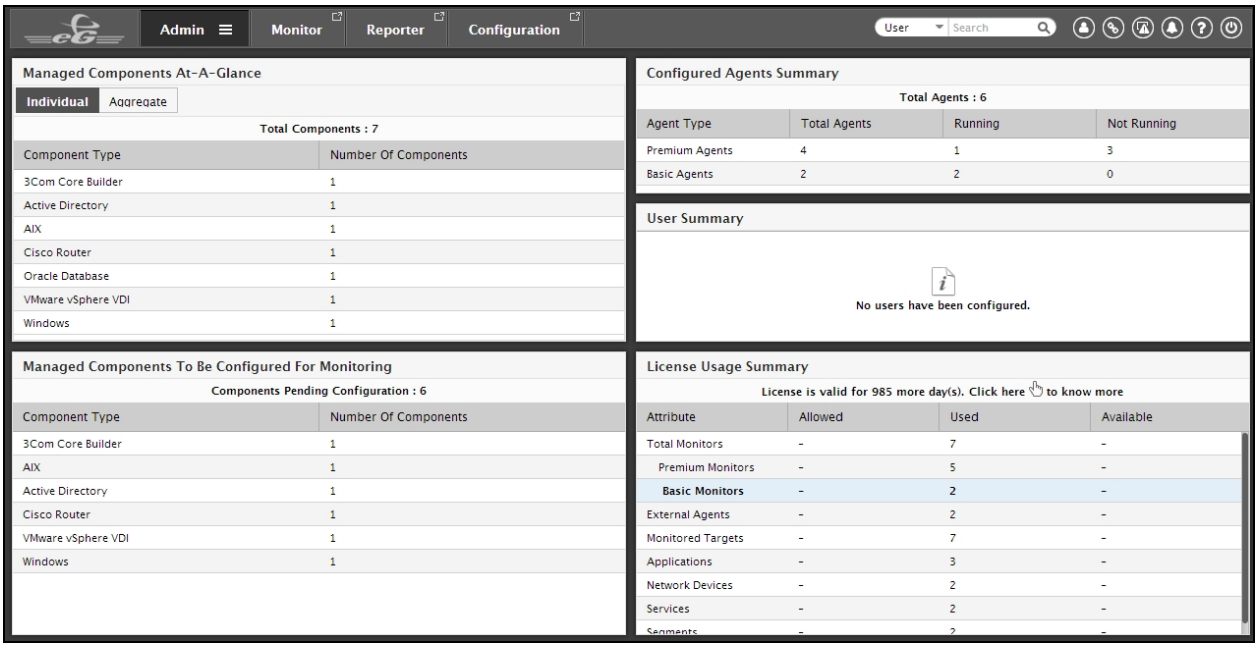

Figure 4.3: The admin home page

This page enables the administrator to understand, at a glance, the status of the eG monitoring system. The page reveals the following information:

- The **Managed Components At-A-Glance** section of the page provides a quick look at the managed infrastructure, by listing the types of components that are currently monitored and the number of components of each type that are being monitored. To know which components of a type are currently managed, click on the **Component Type** or the **Number of Components** corresponding to that type in Figure 4.3. This will lead you to the **COMPONENTS** page, where you can not only view the list of components managed, but can also add more components of that type.
- The **Configured Agents Summary** section in Figure 4.3 displays the total number of agents that have been configured for the environment. The number of agents of each type (Basic and Premium) that have been configured, and the number of agents that are currently running/not running are also indicated. Clicking on the agent type (basic/premium) or the number of agents of a type that are **Running** or **Not Running** will lead you to the **AGENTS – STATUS** page, where you can view the complete list of agents of that type and the current status of each.
- Below the **Configured Agents Summary** is the **User Summary** (see Figure 4.3). This section indicates the number of `Expired` user accounts, user accounts which are **Nearing Expiry** (i.e. the accounts that will expire within 7 days), and the number of local/domain users and domain groups that are currently active. Clicking on a row here will lead you to the **USER INFORMATION** page, from which you can identify the users whose subscription is currently active/has expired/is nearing expiry.
- For each component (i.e. network device or application) being monitored, eG includes a specialized model that dictates what tests must be run by the eG agent to monitor the component. Many of eG's tests are pre-configured i.e., do not require manual configuration. A few tests require explicit configuration. The third section in the **Admin Home** page provides the **Managed Components To Be**

Configured For Monitoring (see Figure 4.3). This section provides a component type-wise breakup of the number of components for which tests are yet to be configured. To know which components of which type are pending configuration, click on a specific **Component Type** or on the **Number Of Components** of a type. Doing so will lead you to the **LIST OF UNCONFIGURED TESTS** page, which lists the tests that are pending configuration for every component of that type. The **LIST OF UNCONFIGURED TESTS** page will also appear whenever administrators/users click the  icon to signout of the eG administrative interface . If administrators/users find this page distracting whenever it pops up during signout, they can hide this page from the eG administrative interface. To hide this page from appearing during signout, administrators/users can set the **ShowUnconfigTestsOnSignout** flag in the **<eG_INSTALL_DIR>\manager\config\eg_services.ini** file to **No**. By default, this flag will be set to **Yes**.

- Adjacent to the **Managed Components To Be Configured For Monitoring** section is the **License Usage Summary** section (see Figure 4.3). The eG license governs a wide variety of factors such as the number and type of agents that the installation can support, the number of applications that can be monitored, etc. The information provided by the **LICENSE USAGE SUMMARY** helps the administrator assess whether the eG licenses are being effectively utilized. Clicking on the **Click here** link will take you to the **LICENSE INFORMATION** page that provides the license details as well as its usage details.

4.2 The Admin Menu and Toolbar

The **Admin** menu is available as tiles and can be invoked by clicking the  icon adjacent to the tab labelled **Admin**. The tiles that appear and the options they offer are as follows:

- **Infrastructure:** This menu enables the administrator to discover, manage, and add/modify components to the eG Enterprise system, configure zones, services, segments, component groups, and component topologies.
- **Agents :** The options provided by this menu allow you to:
  - Configure tests to be executed on target components, so as to extract performance data from them
  - Configure additional external and remote agents for the environment and assign hosts to them
  - View the status of agents
  - Define auto-upgrade settings
  - Define general agent settings such as the following:
    - Define the manager-agent communication settings, and configure remote control commands; **remote control is a license-controlled capability and will be available only if your eG license enables it**.
    - Define how frequently detailed metrics are to be collected by the eG agents; **detailed diagnosis is a license-controlled capability and will be available only if your eG license enables it.**
    - Configure vCenter tasks and events; **tasks and events are of significance only if one/more VMware vCenter servers are being monitored**
- **Alerts:** Using the options provided by this menu, you can control the alerting capabilities of the eG Enterprise Suite by performing the following tasks:

- Define the alarm policies, based on which the eG Enterprise system will generate alarms
- Configure thresholds, which govern state changes
- Add/modify/delete maintenance policies, which when defined, can suppress the alarms pertaining to specific hosts/components for configured time periods
- Configure the mail server for sending email alerts
- Some environments may already be using network-monitoring systems such as HP OpenView, Tivoli NetView, etc., for monitoring their networks and systems. Administrators of such environments may desire that eG Enterprise's alarms be reported to their existing alarm consoles. By configuring the eG manager to send eG alarms as SNMP traps to one/more SNMP management consoles in an environment, you can enable eG Enterprise to support the integrated display and tracking of alarms from a single monitoring console. Using the options in the **Alerts** tile, you can configure the SNMP managers/trap receivers to which the eG manager needs to send SNMP traps and define the SNMP trap settings. In addition, you can upload SNMP MIBs of SNMP-enabled network devices that are not monitored out-of-the-box by the eG Enterprise system. If the **Integration Console** plugin is enabled for use by the eG license, then administrators can browse the new MIB tree online and create new SNMP-based tests for the new network devices.
- Assign agents to primary and secondary managers in Active-Active redundant manager cluster; **this option will be available only if your eG license enables the 'Redundant Manager' capability**.

➢ **Integration Console:** eG Enterprise includes extensive built-in monitoring capabilities for a majority of off-the-shelf applications. However, in any realistic environment, one may encounter applications that are not supported by the eG products. Moreover, administrators may prefer to extend eG's built-in application models to suit their needs and preferences (e.g., to add specific tests from the model). To support these capabilities, eG Enterprise includes an optional component called the **Integration Console**. This license-controlled module allows users to add new components for monitoring, include new layers for diagnosing specific components, and enhance eG's measurement capability to expose additional information relating to the managed components. The **Integration Console** tile will be available only if the eG license enables the **Integration Console** capability. If available, administrators can use the menu options in the tile to do the following:

- Create new tests and measures;
- Add new layers and associate the new tests with the new/existing layers;
- Add a new component-type, define its layer model, and associate new/existing tests with the component type;
- Take a backup of the eG Enterprise Suite

Reference:

To know how to extend eG's monitoring capabilities using the **Integration Console** plugin, refer to the Extending the Monitoring Capabilities of eG Enterprise.

➢ **User Management**: You can create/modify/delete new users and user roles using this menu option, and can view reports on user logins and user activity.

➢ **Settings**: You can configure the key settings that govern manager operations and how state and

performance data is displayed in the eG monitoring console. These include:

- Configuring the eG database cleanup periods
- Configuring general manager settings such as the threshold lookback period, user account lockout policy, password policy, what should be displayed in the manager message board (i.e., the **MANAGER NOTIFICATION** window), etc.
- Configuring general display settings for alarms and graphs in the eG monitoring console
- Configuring the measures to be included in the **Measures At-A-Glance** section of the **Monitor Dashboard**
- Effecting cosmetic changes to the eG admin and monitor interfaces by including custom logos and defining custom alert messages
- Choosing the measures for which Operation reports need to be generated in the eG Reporter; **the eG Reporter component will be available to you only if your eG license enables it**.

➢ **Miscelleneous**: You can view and analyze the manager log files, oversee license usage, check database properties, import/export configuration across managers, and many other administrative activities. You can also add/delete **eG Supermanagers**. Large enterprises often have thousands of devices, servers, and applications that have to be managed, and a single eG management console may not have the capacity to handle the entire enterprise. To support such enterprises, multiple eG managers may be needed. However, if each of these managers operates independently, they may not provide a common view of the entire enterprise. Hence, it could be very cumbersome to have the IT staff of the enterprise login to different eG management consoles to get a complete view of the status of the target infrastructure. A SuperManager is a manager of managers that provides a consolidated view of the status of the IT infrastructure that is being handled by different eG managers. The **eG Supermanager** is a 100% web-based component of the eG suite that provides a consolidated view across disparate eG managers.

Reference:

For a detailed installation and configuration procedure for the **eG Supermanager** and an elaborate discussion on how it works, refer to the document titled *The eG Supermanager*.

➢ **Audits :** eG Enterprise can be optionally configured to log every user action performed on the eG user interface. Using the **Audits** menu, a variety of reports can then be generated based on the details logged, so as to enable the administrator to audit the following:

- User logins to the eG Enterprise system
- Failed login attempts to the eG Enterprise system
- Configuration changes effected by users to the eG administrative interface
- User activities with respect to the eG monitoring console
- User accesses to the eG Reporter module

By default, the 'audit logging' capability of the eG manager is disabled. Therefore, the **Audits** menu will not be available by default.

At the right, top corner of the **Admin** interface, you will find a tool bar. The table below briefly describes the tools provided by this tool bar:

Tool	Tool Name	Purpose
	Current user	Move your mouse pointer over this tool to know who is currently logged in. Click on the tool to edit the profile of the user logged in.
	Quick links	Click here to view 'single-click' links to important or frequently accessed pages of the eG administrative interface. To know how to configure quick links, refer to Section of this document.
	Manager notification	Click here to pull down a message board, where the eG manager displays useful messages for the administrator. Such messages may intimate administrator of agents that may not be running, or components recently managed and awaiting test configuration.
	Quick alerts	You can click on this icon from anywhere in the eG management console to take a quick look at the current alarms.
	Help	Click here for a context-sensitive help page providing useful information pertaining to the page that is currently open.
	Signout	Click here to sign out of the eG administrative interface.

If your user profile has monitoring rights, then a **Monitor** tab will appear, clicking on which will enable you to login to the eG monitoring console, without having to log out of the admin interface. Similarly, if you have rights to generate reports, then a **Reporter** tab will appear, clicking on which will allow you to instantly switch to the eG Reporter interface (if your eG license enables it). By default, clicking on either of these tabs will open the monitor and reporter interfaces in the currently open window itself. If you want these consoles to open in a separate window instead, click on the  symbol next to the tab page name **Monitor** or **Reporter** as the case may be.

4.3 The Manager Notification Window

As soon as a user logs into the eG administrative interface, a **MANAGER NOTIFICATION** window (see Figure 4.4) automatically pops up. This window serves as a message board where critical messages of significance to an eG administrator will be published.

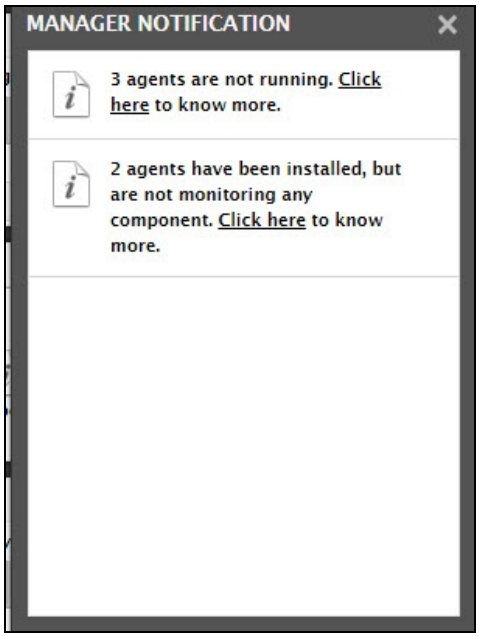

Figure 4.4: The Manager Notification window

An administrator can configure the type of messages that he/she wants displayed in the board. By default, the notification window alerts administrators to the following:

- eG license expiry
- Exhaustion of eG agent licenses
- Agents not running
- Components newly discovered
- Agents awaiting test configuration

The administrator can override this default setting by turning on/off specific messages. For this, the administrator has to select the **Manager** option from the **Settings** tile of the eG administrative interface, select the **Manager Notification** option from the **MANAGER SETTINGS** panel, and then toggle the status of the notification message types displayed in the right panel.

Reference:

For more information on how to configure the contents of the notification window, refer to 5.5.2 of this document.

At any given point in time, administrators can click on the **Click here** hyperlink accompanying a message in the **MANAGER NOTIFICATION** window (see Figure 4.4) for more details on a particular alert. For instance, Figure 4.4 indicates that 3 agents are not running. Clicking on the **Click here** hyperlink alongside this message in Figure 4.4 will lead the administrator to Figure 4.5, which displays the list of agents that are not running currently.

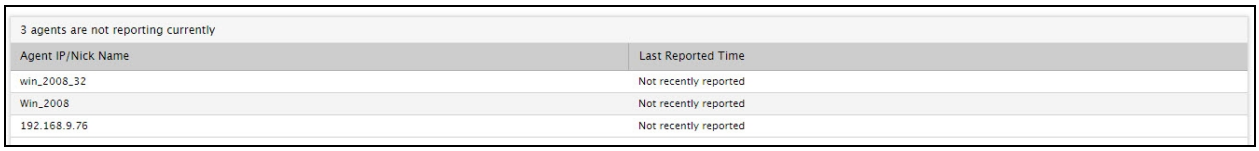

3 agents are not reporting currently

Agent IP/Nick Name	Last Reported Time
win_2008_32	Not recently reported
Win_2008	Not recently reported
192.168.9.76	Not recently reported

Figure 4.5: A page displaying the list of agents that are not running

To close the **MANAGER NOTIFICATION** window, you can either move your mouse pointer over the window or click the **X** button at its right, top corner. Once closed, you can invoke the **MANAGER NOTIFICATION** window yet again by clicking the button in the **Admin** toolbar (at the right, top corner of the eG admin interface).

4.4 Quick Alerts

eG Enterprise brings problems to the attention of administrators via multiple modes such as the eG monitoring console, emails, SMS, and SNMP traps. The latest addition to this list is the optional, **Quick Alerts**. The **Quick Alerts** mechanism introduced by eG Enterprise, saves administrators who might be making configuration changes using the eG admin interface or say, scheduling the delivery of reports using the eG Reporter, the time and trouble involved in switching to the eG monitoring console, everytime he/she needs a quick update on the problems affecting the infrastructure. This feature helps administrators track problems continuously by displaying the number and details of current alarms when the icon at the right, top corner of the eG management console is clicked.

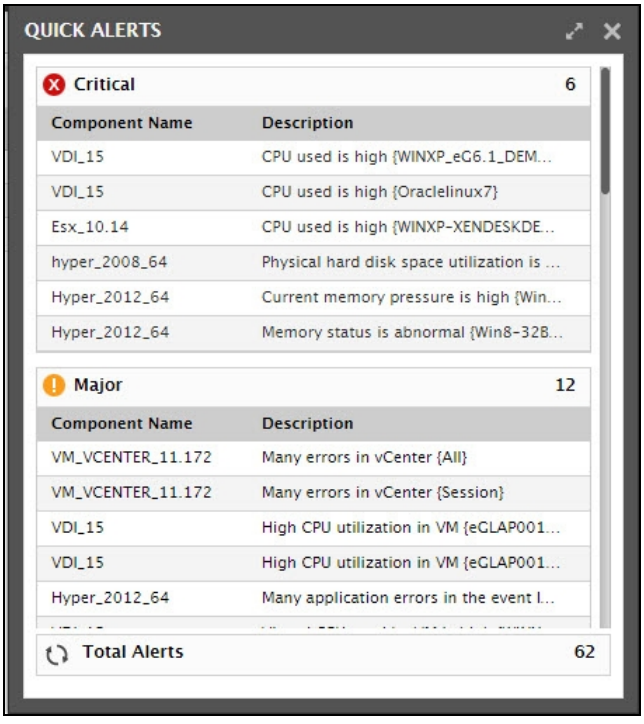

Figure 4.6: Quick Alerts

The **QUICK ALERTS** window that appears (see Figure 4.6) groups alarms by priority and displays the count and details of alarms of each priority. Clicking on an alarm here will lead the user to the layer model of the problem component, which will reveal the exact layer that is affected by the problem, the test that reported the problem, and the problematic measure.

This way, a user can receive instant updates on performance issues and can even drill down to ascertain the exact nature of the issue, regardless of which eG module he is logged into currently.

Note:

The **QUICK ALERTS** window will display the alarms pertaining to only those components that have been assigned to the user (who is currently logged into eG) for monitoring.

5

Basic Settings

Before managing components using the eG administrative interface, you need to define certain basic settings to enable the eG manager to function properly. This chapter discusses these settings in detail.

5.1 Configuring Manager-Agent Communication

By default, the eG agent-manager communication is neither authenticated, nor encrypted. This default setting can however be overridden by the administrator. For this, the procedure discussed below should be followed:

1. From the **Agents** tile, pick the **Settings** option.
2. When Figure 5.1 appears, pick the **Communication** node from the **AGENT SETTINGS** panel.

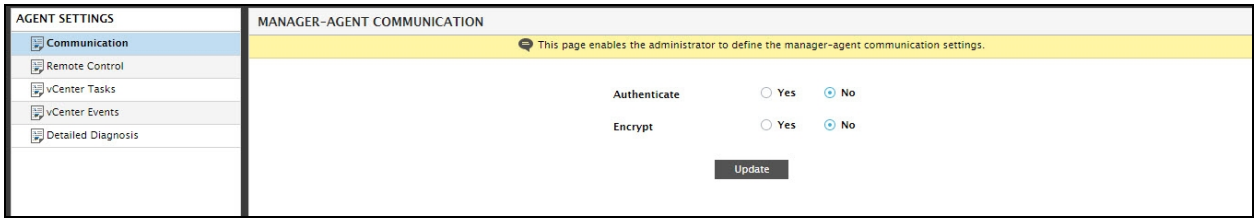

Figure 5.1: Settings for the eG manager - agent communication

3. A **MANAGER-AGENT COMMUNICATION** panel then appears in the right panel (see Figure 5.1). Using the **Authenticate** option, an administrator can turn on or off authentication of agents by the eG manager. The default option is `No`, indicating that authentication is disabled. An administrator may choose to turn on authentication by selecting the `Yes` option. Enabling authentication ensures that only a specified agent(s) is reporting metrics to the eG manager, and not any other malicious program.
4. The **Encrypt** option determines whether data communication from the agent to the manager is encrypted or not. The default option is `No`, indicating that encryption is off. You can switch on encryption by choosing the `Yes` option. Typically, you might want to enable encryption if measure data is to be transmitted over a HTTP connection. If using HTTPS already, then this option is NOT required. For a HTTP based eG manager/agent setup, this option provides a way for secure communication without the elaborate setup needed for SSL.

5.2 Configuring the Mail Settings

5.2.1 Configuring the Mail Server

You need to configure the mail server in your environment to allow the automatic generation and transmission of email alerts to specified recipients. Figure 5.2 depicts the configuration of mail settings for the eG manager. This page can be accessed by selecting the **Server Settings** option from the **Mail Settings** menu of the **Alerts** tile.

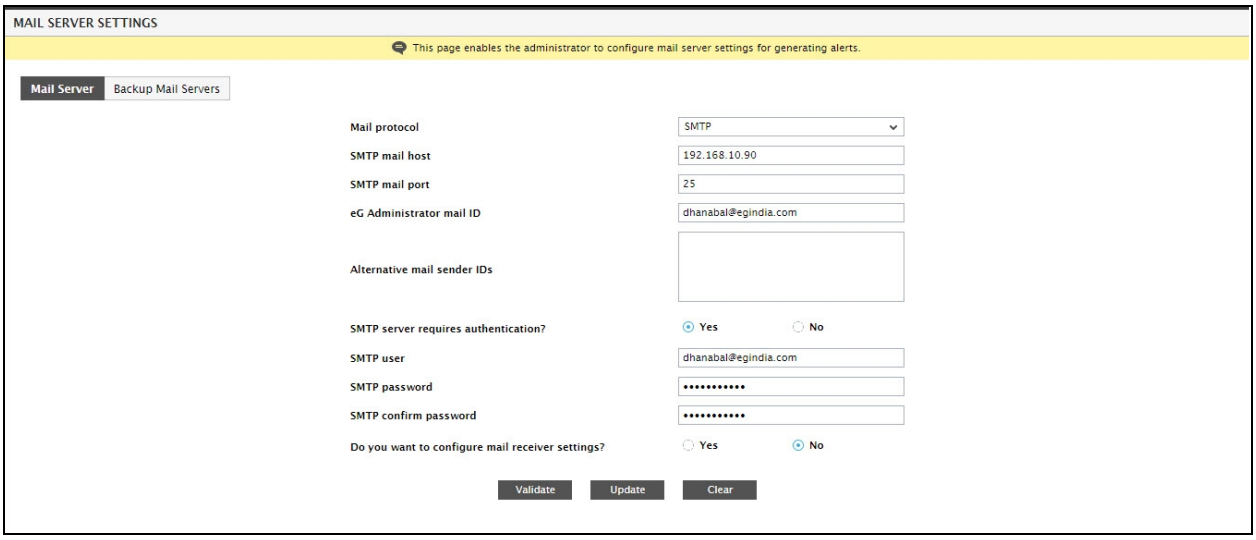

Figure 5.2: Mail Server Settings page

The protocol through which you wish to transmit or send the outgoing mail messages across the Internet Protocol (IP) networks has to be selected from the **Mail protocol** list box. The **SMTP** option would be selected by default in this list box. If the mail server through which you wish to send the mail messages is **SSL-enabled**, then select, **SMTP-SSL** from the **Mail protocol** list box. If your mail server offers enhanced security and provides certificate based authentication, select the **SMTP-TLS** option from the **Mail protocol** list.

The identity (IP address or host name) of the mail server to be used by the eG manager for generating alarms has to be entered in the **SMTP mail host** text box. The port at which the mail host listens has to be provided in the **SMTP mail port** text box. The entry in the **eG Administrator mail ID** text box will be the mail ID from which the alarms are generated to eG users.

In MSP environments typically, different support groups are created to address performance issues relating to different customers. These support groups might prefer to receive problem intimation from customer-specific mail IDs instead of the global admin mail ID, so that they can instantly identify the customer environment that is experiencing problems currently. Moreover, this way, every support group will be enabled to send status updates on reported issues directly to the concerned customer, instead of overloading the admin mailbox. To facilitate this, the **MAIL SERVER SETTINGS** page allows the administrator to configure multiple **Alternative Mail sender IDs** - normally, one each for every customer in case of an MSP environment. While configuring multiple sender IDs in the space provided, ensure that you press the **Enter** key on your keyboard after every mail ID. This way, every ID will occupy one row of the text area. Later, while creating a new user, the administrator can select one of these configured sender IDs from the **Mail sender** list in the **ADD USER** page, and assign it to the new user. This ensures that all email alerts received by the user are generated by the chosen ID only.

If the mail server requires users to login before sending mails, then select the **Yes** option against the **SMTP server requires authentication?** field. By default, authentication is set to **No**. Upon selecting **Yes**, you will be required to provide a valid **SMTP user** name and **SMTP password** for logging into the mail server. Confirm the password by retyping it in the **SMTP confirm password** text box.

To safeguard from spam, some mail servers are configured so that they will allow mails to be sent from a system only if that system is also used to receive mails. To allow the eG manager to use such mail servers to

send email alerts, additional configuration is needed. In such a case, select the **Yes** option against the **Do you want to configure mail receiver settings?** field. By default this field is set to **No**. When you enable this authentication to **Yes**, you need to specify the following details in the corresponding text boxes:

- **Mail receiver ID:** Specify the login name to be used for receiving mails.
- **Mail receiver password:** The password of the mail receiver needs to be specified here.
- **Port used for receiving mails:** The port number on the mail server to which the mail manager connects needs to be provided here.
- **Protocol for receiving mails:** Mention the protocol used for receiving mails. The protocol can be either POP3 or IMAP.
- **Server for receiving mails**: Specify the server to which the mail manager will connect to receive mails.

Sometimes, alarm mails may not be received by the configured recipients. When such an anomaly occurs, administrators typically spend hours to determine the reason for the non-delivery of emails. One of the most common causes for non-delivery of email alerts is the improper configuration of the mail server for the eG manager. For instance, an incorrect IP address specified against **SMTP host** in Figure 5.2 or invalid credentials provided against **SMTP user** and **SMTP password** can halt the generation and transmission of email alerts. To enable administrators to spot and fix such configuration issues before the eG manager even attempts to send out email alerts, a **Validate** button is provided in the **MAIL SERVER SETTINGS** page of Figure 5.2. Clicking on this button instantly verifies the correctness of the values configured in the **MAIL SERVER SETTINGS** page and promptly indicates discrepancies to the administrator. This way, administrators need not have to wait for delivery failures to occur to isolate configuration issues.

The location of this **Validate** button and the information it validates will vary depending upon the status of the following flags in Figure 5.2:

- **Does SMTP server require authentication?**
- **Do you want to configure mail receiver settings?**

The table below explains how the status of the aforesaid flags influences the location and the functionality of the **Validate** button:

Status of the Authentication flag	**Status of the Mail receiver settings** flag	**Location of Validate** button	Function of **Validate** button
No	No	Adjacent to **eG Administrator Mail ID** text box	Validates the **SMTP mail host** and the **SMTP mail port** fields
Yes	No	Next to the **SMTP Confirm Password** text box	Validates the **SMTP mail host**, **SMTP mail port**, and the user credentials that are provided for the SMTP server authentication
No	Yes	Next to the **Server for receiving mails** field	Validates the **SMTP mail host**, **SMTP mail port**, and the complete **Mail receiver**

Status of the Authentication flag	**Status of the Mail receiver settings** flag	**Location of Validate** button	Function of **Validate** button
			configuration
Yes	Yes	Next to the **Server for receiving mails** field	Validates the **SMTP mail host**, **SMTP mail port**, user credentials for authenticating SMTP server access, and the complete **Mail receiver** configuration

Note:

The mail ids provided in the **eG Administrator mail ID** and the **Alternative Mail sender IDs** fields will **not** be validated using the **Validate** option. If an incorrect mail id is provided in these fields, delivery failures are bound to occur.

If, upon clicking the **Validate** button, the corresponding information is validated, the message confirming the success of the validation will appear as shown in Figure 5.3.

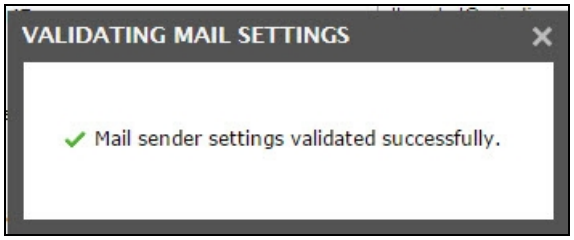

Figure 5.3: Mail Settings Validation pop up window

If the validation is unsuccessful, then a message to that effect would appear.

Finally click the **Update** button in Figure 5.2 to register the changes.

5.2.2 Configuring a Backup Mail Server

The eG manager must to be configured to use a mail server for routing email alerts to users. If this mail server fails for any reason, then important problem notifications may not reach administrators. In turn, this causes performance issues to remain undetected (and hence, unresolved!).

eG Enterprise v6 allows administrators to configure more than one mail servers for routing email alerts to users. When an alert is generated, the eG manager will first attempt to send out an email alert using the primary mail server. If it is unable to do so, then the eG manager will automatically try and send the email alerts using each of the configured backup mail servers in sequence, until it succeeds. This ensures that no problem goes unnoticed by administrators, even if one mail server is unavailable. Moreover, the next time an email alert needs to be sent out, the eG manager intelligently picks the mail server that successfully sent out alerts during the last attempt and uses that server first to process the alert.

To add a new backup mail server, do the following:

1. First, click the **Backup mail servers** tab page in Figure 5.2.
2. When Figure 5.4 appears, click the **Add** button therein to add a new backup mail server.

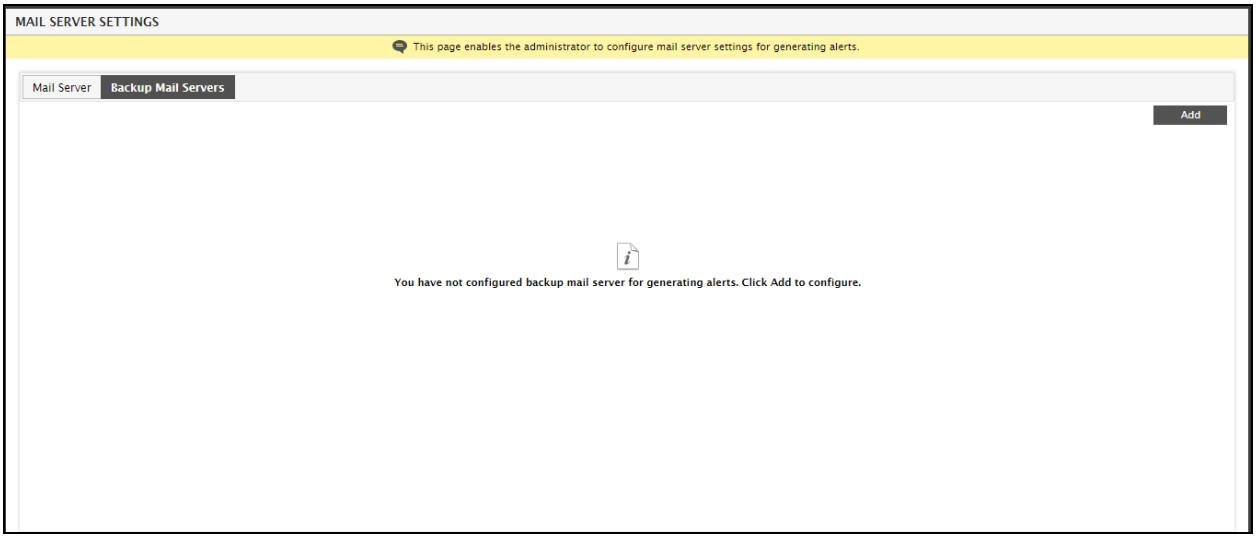

Figure 5.4: Adding a new Backup mail server

3. Figure 5.5 will then appear, where you can provide the details of the backup mail server.

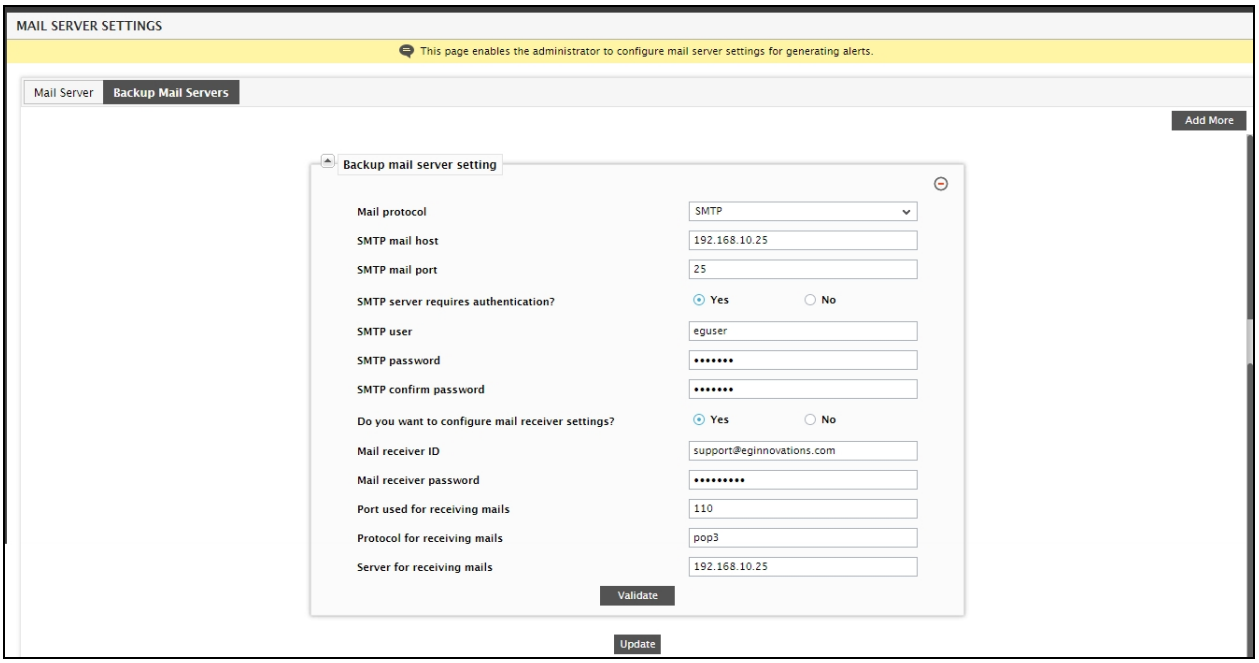

Figure 5.5: Configuring the details of the new backup mail server

4. To know how to configure the parameters in Figure 5.5, refer to 5.2.1 above.
5. Click the **Update** button to save the changes.
6. To add another backup mail server, click the **Add More** button in Figure 5.5.

7. To remove a backup mail server that has already been added, just click the encircles '-' button indicated by Figure 5.5 above.

5.2.3 Configuring the Mail Alert Settings

To configure the mail alert settings, select the **Alert Settings** from the **Mail Settings** menu of the **Alerts** tile. Figure 5.6 will then appear.

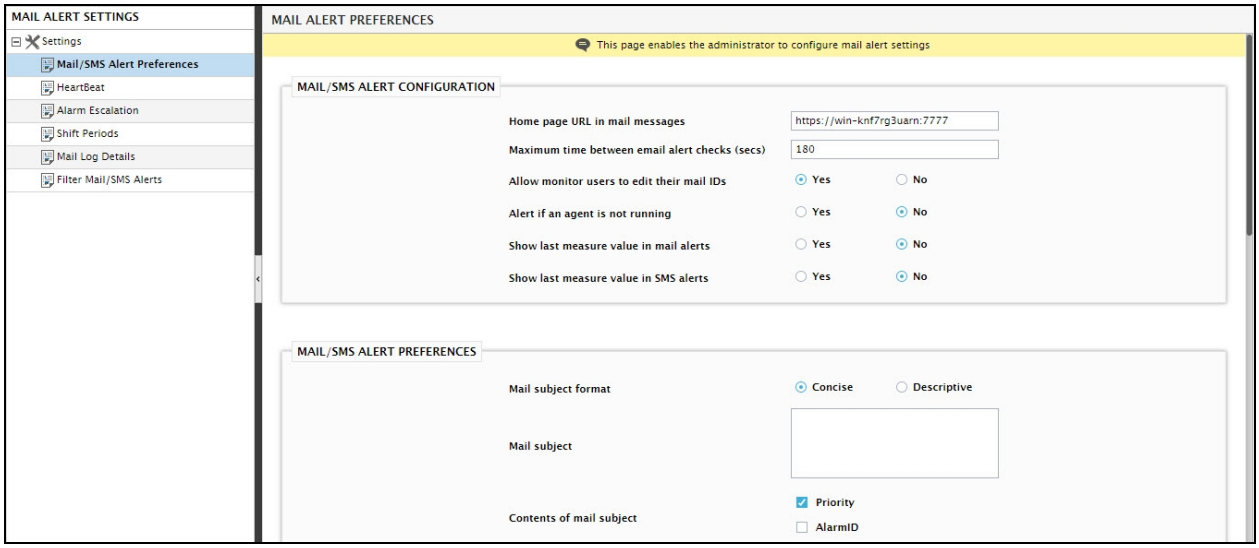

Figure 5.6: The MAIL ALERT SETTINGS page

As can be inferred from Figure 5.6, the **MAIL ALERT SETTINGS** panel to the left provides a **Settings** tree-structure – each node of the tree represents a collection of parameters that an administrator can set, in order to control the email alerting function of the eG manager. The contents of the right panel will change in context to the node chosen from the **Settings** tree.

The sub-sections below discuss each of these nodes elaborately.

5.2.3.1 Mail/SMS Alert Preferences

Click on the **Mail/SMS Alert Preferences** node to define when and how email alerts should be sent. The right panel will change to display the options shown in Figure 5.7.

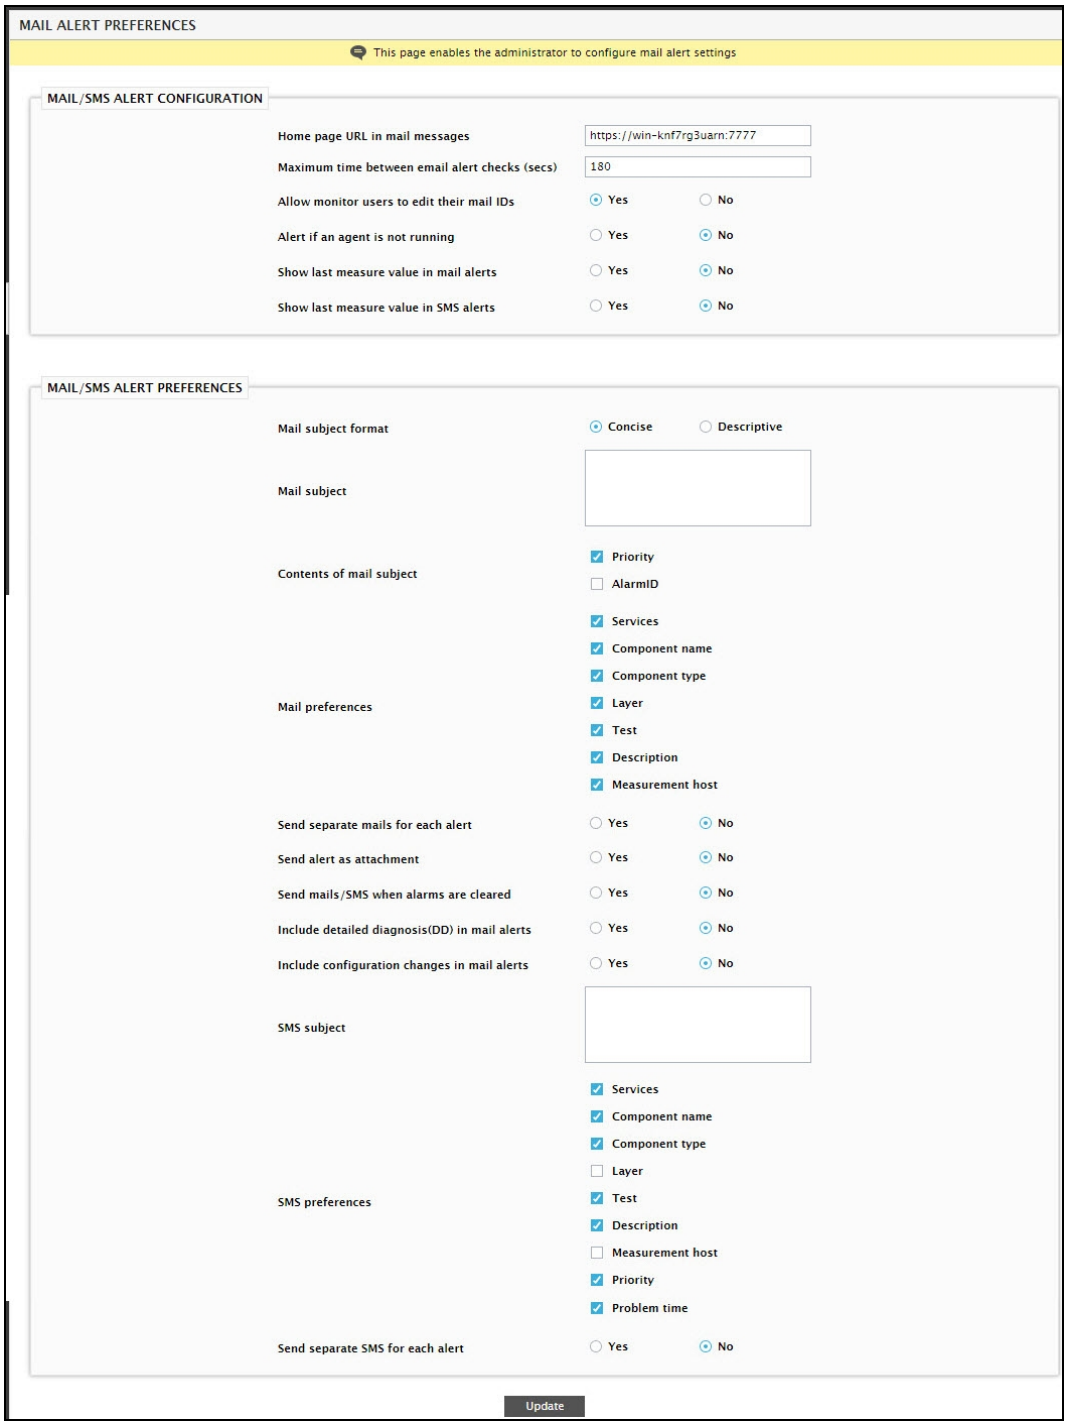

Figure 5.7: The Mail/SMS Alert Configuration section of the Mail Alert Preferences page

In the **MAIL/SMS ALERT CONFIGURATION** section of Figure 5.7 above, you can configure the following:

- **Home page URL in mail messages:** If a user has been configured to receive email alerts in the **HTML MESSAGE MODE**, then alarms sent by mail to that user will carry a hyperlink named **Home** at the right top corner. The destination of the hyperlink can be configured in this text box. By default, clicking on the **Home**

link will connect you to the eG manager and open the login screen. By providing a specific URL here, you can ensure that monitor users are lead to the specified URL upon clicking the **Home** hyperlink. **Note that the URL specified here will appear on the title bar of the eG management console as soon as you login.**

- **Maximum time between email alert checks:** By default, the eG Enterprise system sends email alerts every 3 minutes (i.e., 180 seconds). In some environments, one/more tests could be run very frequently, say every 1 minute or 30 seconds. If these tests report abnormalities, then the user will receive an email alert of the issue only at the 3rd minute. What's worse, if the issue detected by the test is resolved by the time the email alert is sent out by the eG manager, then the user might not even know that a problem had occurred - critical problems could thus go unnoticed! To avoid this, you can reduce the frequency at which the eG manager sends out email alerts. For this, you need to override the default value (of 180 seconds) displayed against the **Maximum time between email alert checks** text box.
- **Allow monitor users to edit their mail IDs:** By default, the password and e-mail/mobile no. settings in the **USER PROFILE** page are editable. However, for security purposes, if the password and e-mail/mobile no. settings need to be rendered non-editable, then, set this flag to **No**.
- **Alert if an agent is not running:** You can configure the eG Enterprise system to send out email alerts when an agent stops running. To do so, set this flag to **Yes**.

Figure 5.8 depicts the email alert received by the eG administrator, when some of the agents are not reporting measures to the manager.

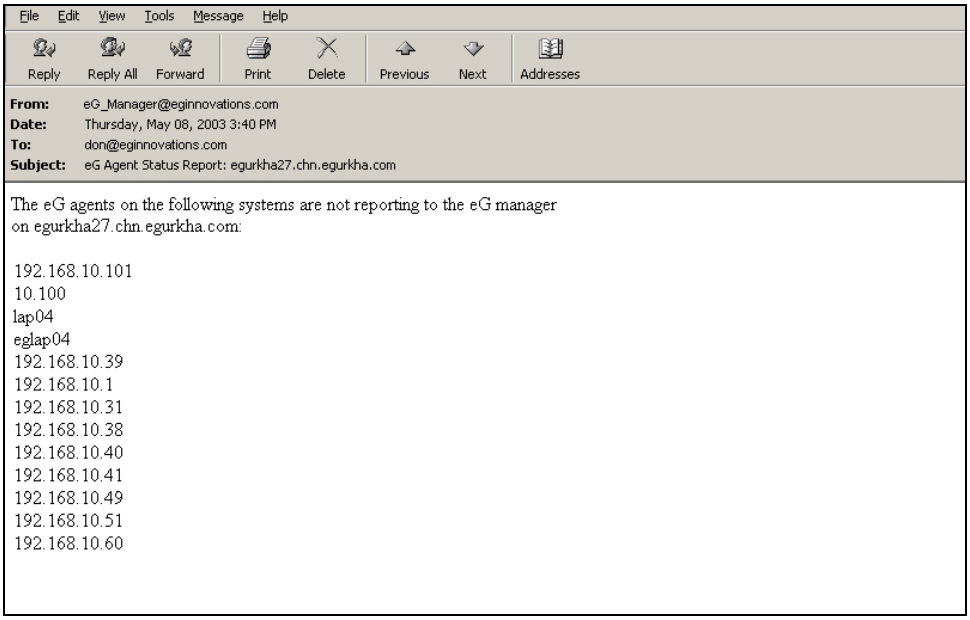

Figure 5.8: Email alert received by the administrator when agents do not report measures to the manager

Note:

An eG agent can be configured to run specific tests once a day or once every few hours. You can configure the eG manager to exclude tests that are infrequently run when it determines whether an agent is running or not. To do this, modify the value of **NotReportingCutoffFactor** in the **[MISC_ARGS]** section of the **eg_services.ini** file. By default, tests running with measure period of greater than 20 minutes are not considered by the eG manager for determining if an agent is running or not.

- **Show last measure value in mail alerts:** If you want the email alert to display the value that was last reported for the problem measure, then set this flag to **Yes**. If this is done, then the test name, measure description, and last measure value will by default be included in every email alert. Accordingly, in the **Mail preferences** sub-section of the **MAIL/SMS ALERT PREFERENCES** section (see Figure 5.7), the **Test**, **Measure**, and **Last measure value** check boxes will appear selected, but will be rendered non-editable. If you do not want the last measure value in your email alert, then set the **Show last measure value in alerts** flag to **No**. In this case therefore, the **Last measure value** check box will not even be available for selection in the **Mail preferences** sub-section. As for the **Test** and **Description** check boxes, they will appear selected by default, but administrators can uncheck the check boxes if they do not want to include this information in the email alerts.
- **Show last measure value in SMS alerts:** If you want the SMS alerts to display the value that was last reported for the problem measure, then set this flag to **Yes**. If this is done, then the test name, measure description, and last measure value will by default be included in every SMS alert. Accordingly, in the **SMS preferences** sub-section of the **MAIL/SMS ALERT PREFERENCES** section (see Figure 5.7), the **Test**, **Measure**, and **Last measure value** check boxes will appear selected, but will be rendered non-editable. If you do not want the last measure value in your SMS alert, then set the **Show last measure value in SMS alerts** flag to **No**. In this case therefore, the **Last measure value** check box will not even be available for selection in the **SMS preferences** sub-section. As for the **Test** and **Description** check boxes, they will appear selected by default, but administrators can uncheck the check boxes if they do not want to include this information in the email alerts.

5.2.3.2 Mail/SMS Alert Preferences

To configure the subject and format for email/SMS alerts, use the **MAIL/SMS ALERT PREFERENCES** section (see Figure 5.9).

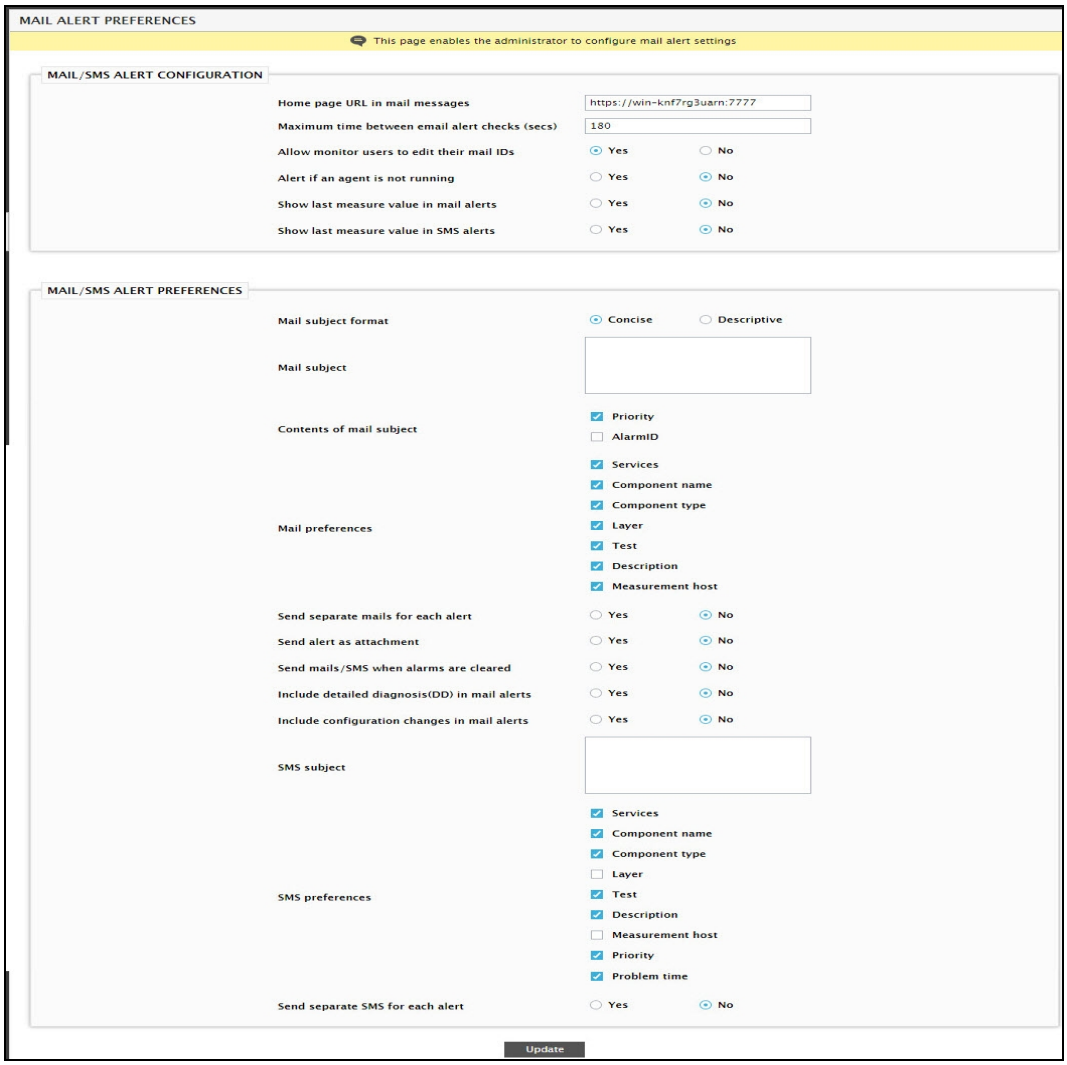

Figure 5.9: Mail/SMS Alert preferences section of the Mail Alert Preferences page

The administrator can customize the subject of the alarm mails by specifying an appropriate subject. Towards this end, the administrator will first have to indicate whether he/she needs to provide a simple, brief subject, or a more descriptive subject. To provide a short and crisp subject, select the **Concise** option against **Mail subject format**. In which case, the administrator has to specify the mail subject in the **Mail subject** text box (see Figure 5.10).

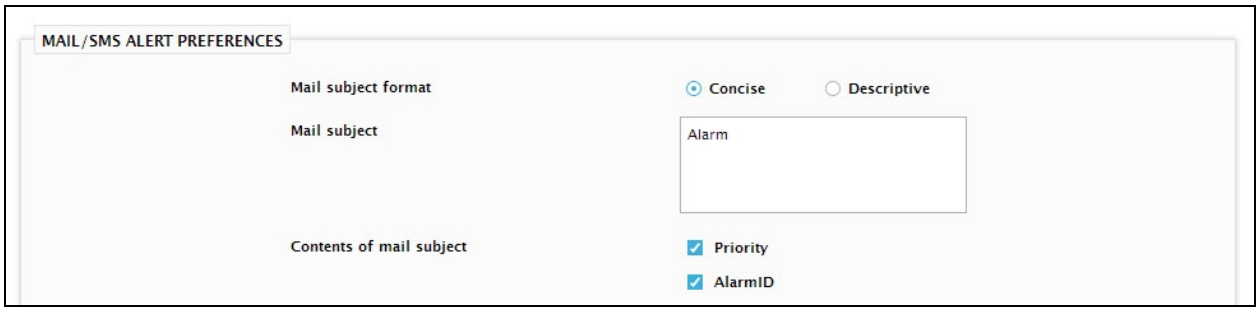

Figure 5.10: Building a 'Concise' mail subject

By default, the **Concise** mail subject will include the alarm priority - accordingly, the **Priority** check box will be selected by default in the **Contents of mail subject** section. If you want to exclude the priority from the mail subject, simply deselect the **Priority** check box. If required, you can also add the **AlarmID** to the mail subject. Everytime an alarm is generated, the eG Enterprise system automatically assigns a unique **AlarmID** to it. To include this ID in the **Concise** mail subject, select the **AlarmID** check box.

To provide an elaborate subject, which should include the names and/or types of components to which the alarms pertain, select the **Descriptive** option. In this case, administrators will be allowed to "build" the entire mail subject, by first specifying how the subject should begin in the **Start of mail subject** text box (see Figure 5.11). Then, he/she should choose the **Contents of mail subject** by selecting the desired check boxes. A typical mail subject can include one/more of the following:

- **Name**: The name of the problem component
- **Component Type**: The name of the problem component-type
- **Layer**: The name of the problem layer
- **Test**: The test that reported the problem measure(s)
- **Description**: The problem descriptor (if any), and a brief description of the problem
- **Priority**: The problem severity (Critical/Major/Minor)
- **Last measure value**: The last value reported by the problem measure; this check box will appear only if the **Show last measure value in alerts** option in the **MAIL SERVER - ADVANCED OPTONS** page is set to **Yes**.
- **AlarmID**: You can choose to include this ID in your email subject by selecting this check box.

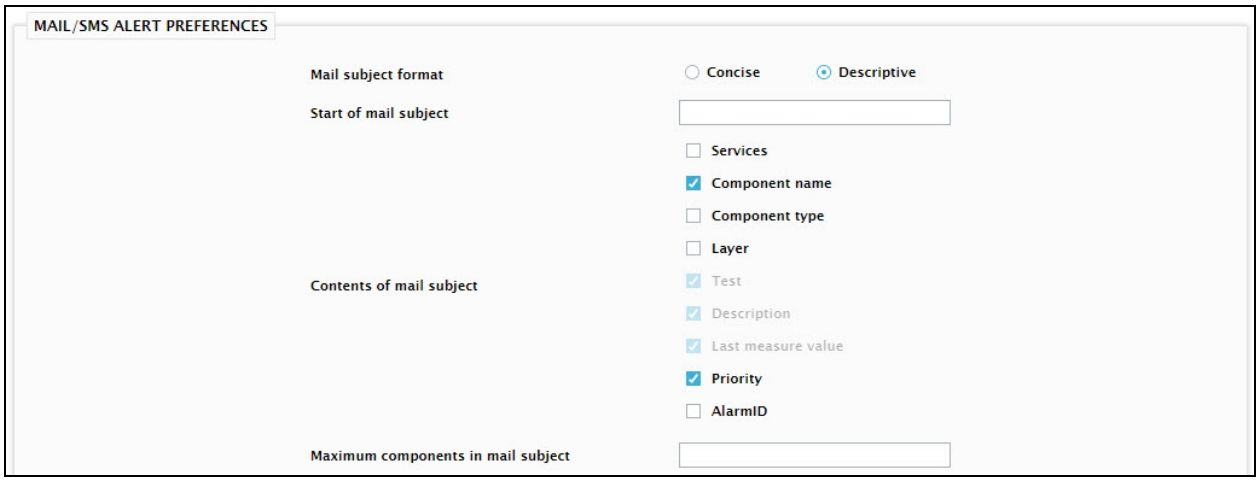

Figure 5.11: Building a 'Descriptive' mail subject

For instance, assume that, except the **AlarmID** check box, all the other check boxes in the **Contents of mail subject** section are selected. Now, say that the **DiskSpace** test mapped to the **Operating System** layer of a component named **Event50** has detected a **Critical** space crunch in the **C** drive of that component. The corresponding email alert will therefore carry the following subject:

Critical,Event50,Host system,Operating System,DiskSpace,Capacity is low{C},98

From the above example, it is clear that the above email alert is of the following format:

<Priority>,<Component Name>,<Component Type>,<Layer>,<Test>, Alarm description {Descriptor},Last measure value

The maximum number of problem components to feature in the mail subject should also be indicated in the **Maximum components in mail subject** text box.

Note:

The **Descriptive** mail subject discussed above will take effect only if the following conditions are fulfilled:

- The user profile must be configured to receive email alerts for **New** alarms.
- The **Send separate mails for each alert** flag should be set to **Yes**.

In this case, note that the **Maximum components in mail subject** setting becomes irrelevant. On the other hand, if the mail subject is set to **Descriptive**, the user profile is configured to receive alerts for the **Complete** list of alarms, and the **Send separate mails for each alert** flag is set to **No**, then the subject of the resulting email alert will indicate only the *<ComponentName> <ComponentType>*. In this case however, the **Maximum components in mail subject** setting gains significance.

Similarly, if the mail subject is set to **Descriptive**, the user profile is configured to receive alerts for the **New** list of alarms, and the **Send separate mails for each alert** flag is set to **No**, the resulting email alert will not display the variables chosen from the **Contents of mail subject** section; instead, a mail subject of the format, *<ComponentName> <ComponentType>,* will accompany each mail alert generated. Here again, the **Maximum components in mail subject** setting gains significance.

Besides the mail subject, the contents of an email alert can also be customized by selecting the relevant check boxes from the **MAIL/SMS ALERT PREFERENCES** page of Figure 5.9. The **Last measure value** check box will appear in this section only if the **Show last measure value in alerts** flag is set to **Yes**. You will find this flag in the **MAIL ALERT CONFIGURATION** page (see Figure 5.7) is clicked.

Note:

If you select the **Description** check box in the **Mail preferences** section, the **Test** and **Measure** (if available) check boxes will get selected automatically. Similarly, deselecting any one of the **Description**, **Test,** or **Measure** check boxes will automatically disable the other two options. This indicates that if an alarm description is contained in an email alert, the details of the problem test, and the last measure value (if the **Measure** check box appears in the **MAIL/SMS ALERT PREFERENCES** page) will appear.

If multiple alarms are generated simultaneously, then the eG Enterprise system, by default, sends a single email alert comprising of all the alarm information. Accordingly, the **Send separate mails for each alert** flag is set to **No** by default. To ensure that a separate email is sent for every alarm, select the **Yes** option.

Note:

The "separate email/SMS alert" flag setting will take effect only if a user is configured to receive email/SMS alerts for the **New** alarms. For users who are configured to receive the **Complete List** of alarms, details of multiple alarms will continue to be clubbed in a single email/SMS regardless of the flag setting.

Figure 5.12 depicts a sample email alert sent out by the eG manager.

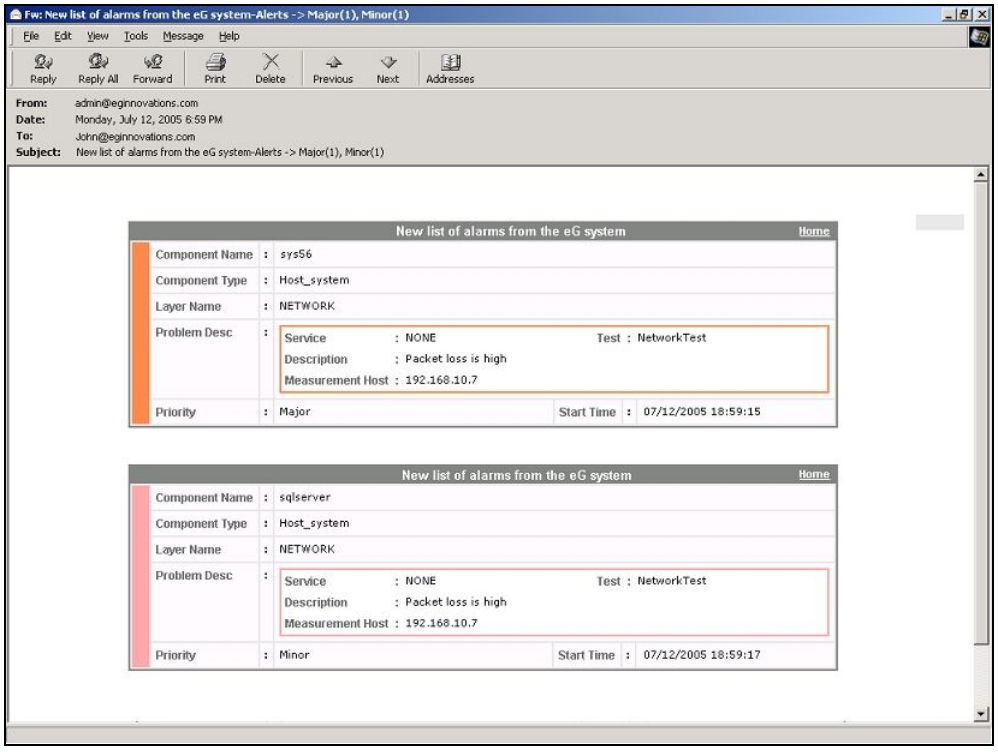

Figure 5.12: An email alert with the default mail subject

Note:

- If the eG agent detects an issue in a test **Descriptor**, then the email alert will also include the descriptor name. If another descriptor of the same test suffers a performance setback later, then users who are configured to receive only **New**alarms, will receive a single email containing the information pertaining to both the alarms. Moreover, the start **Date** and time for both the problems, as indicated by the alert, will be that of the first alarm.
- Similarly, when the alarm priority changes - say, from *Critical* to *Major* - the email alert for the *Major* alarm will carry the **Date** and time of the original *Critical* alarm.

In some environments, the mail servers used for processing email alerts might not be adequately configured to support certain data types/fonts used in the email content. In such cases, alarm information included in the body of the mail could appear distorted / misaligned. To avoid this, the eG manager can be configured to send alerts as attachments, so that users can open the alarm information and view it using any program that they choose.

To achieve this, set the **Send alert as attachment** flag in Figure 5.9 to **Yes**. By default, this flag is set to **No**. Note that both HTML and text alerts can be sent as attachments. Figure 5.13 depicts a sample HTML alert attached to an email.

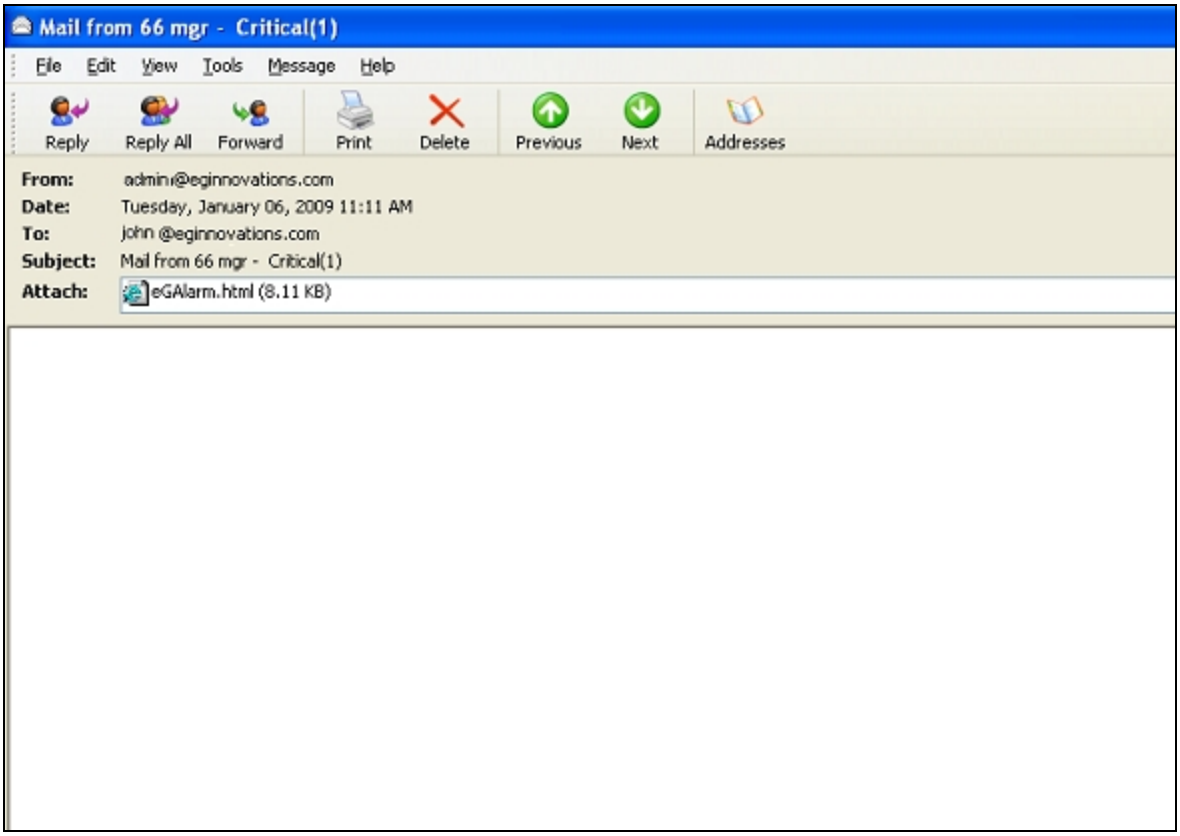

Figure 5.13: Alert sent as an attachment

The eG Enterprise system can also be configured to send email alerts/SMS when a problem gets fixed. This can be done by setting the **Send mails/SMS when alarms are cleared** flag in Figure 5.9 to **Yes**. A sample normal alert is shown below:

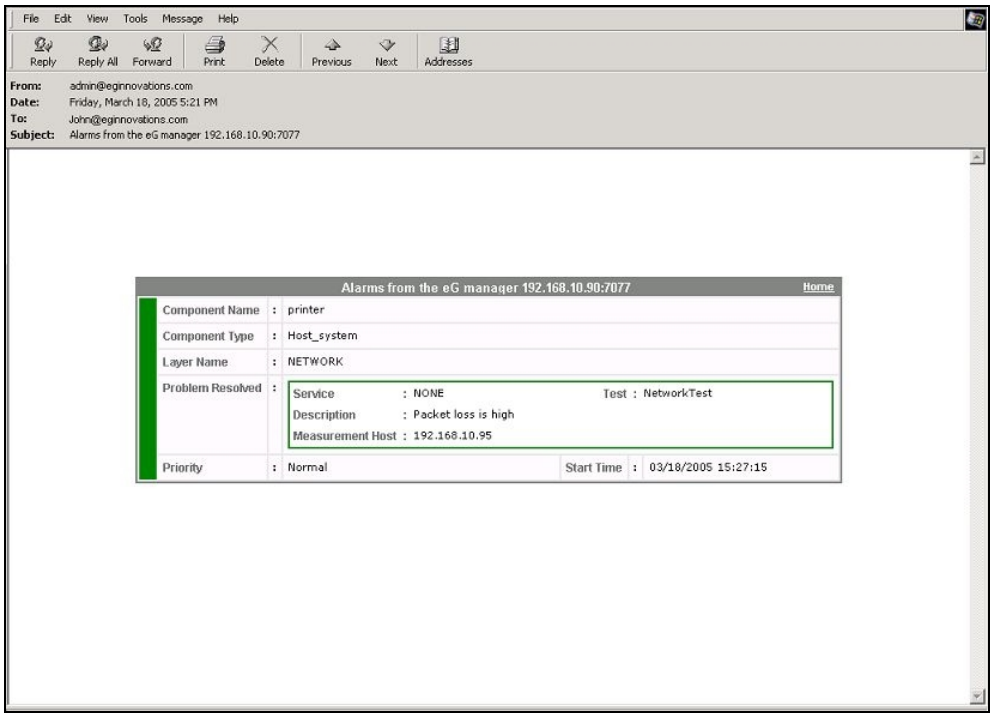

Figure 5.14: A normal alert sent by the eG Enterprise system

Besides the above-mentioned information, the email alerts sent by the eG Enterprise system can also be configured to include the detailed diagnosis of the problem measures. For instance, an email alert indicating excessive CPU utilization on a host can also be configured to list the top 10 CPU-consuming processes on that host. This way, administrators can easily perform further diagnosis without having to login to the monitor interface; the required information will be available in the email itself.

To ensure that the detailed diagnosis information accompanies the alarm details, set the **Include detailed diagnosis (DD) in mail alerts** flag to **Yes**.

Note:

Once this flag is switched on, then users who are configured to receive the **New** list of alarms will begin receiving email alerts with DD. However, in the case of users who have been configured to receive the **Complete** list of alarms, the eG manager will only send email alerts without DD.

Figure 5.15 depicts a sample HTML email alert with DD information.

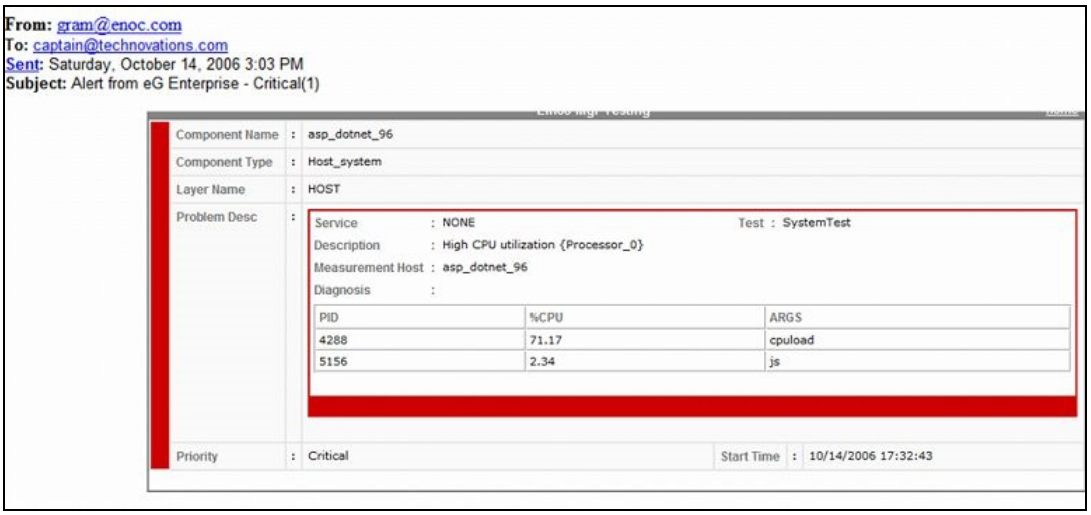

Figure 5.15: A sample email alert with Detailed Diagnosis information

You can also make sure that configuration changes are also intimated to administrators via email alerts. For this, set the **Include configuration changes in mail alerts** flag to **Yes**.

Similarly, for alarm information transmitted via SMS, an administrator can provide a subject of his/her choice in the **SMS subject** text box. By default, the contents of the SMS will include the *alarm priority*. The user can customize the other details to be displayed in the mobile phone's console by selecting the relevant check boxes in the **SMS preferences** section. By default, the **Component Name**, **Component Type**, **Description**, **Services**, **Test**, **Priority**, and **Problem time** are selected. If the **Measurement host** should also be displayed, then the administrator must select the **Measurement host** check box.

If multiple alarms are generated simultaneously, then the eG Enterprise system, by default, sends separate SMS alerts for every alarm. Accordingly, the **Send separate SMS for each alert** flag in Figure 5.9 is set to **Yes** by default. To ensure that a single SMS alert comprises of information pertaining to all the simultaneous alarms, set this flag to **No**.

Note:

The "separate email/SMS alert" flag setting will take effect only if a user is configured to receive email/SMS alerts for the **New** alarms. For users who are configured to receive the **Complete List** of alarms, details of multiple alarms will continue to be clubbed in a single email/SMS regardless of the flag setting.

Note:

- The **SMS subject** and **SMS preferences** fields will appear only if the eG license enables the SMS alerting capability. If the **LICENSE INFORMATION** screen displays a **No** against **SMS Alerts**, then the SMS subject and settings cannot be configured.
- To stop receiving email alerts, the administrator can remove the **eG Administrator mail ID** from the **MAIL/SMS SETTING** page (see Figure 5.2). For that, he/she would have to remove the mail host first, and then the **eG Administrator mail ID**, and finally, click the **Update** button.

5.2.3.3 Heartbeat Configuration

If you click on the **HeartBeat** node in the **Settings** tree in the left panel of Figure 5.6, a **HEART BEAT** section (see Figure 5.16) is provided for configuration in the **MAIL/SMS ALERT PREFERENCES** page.

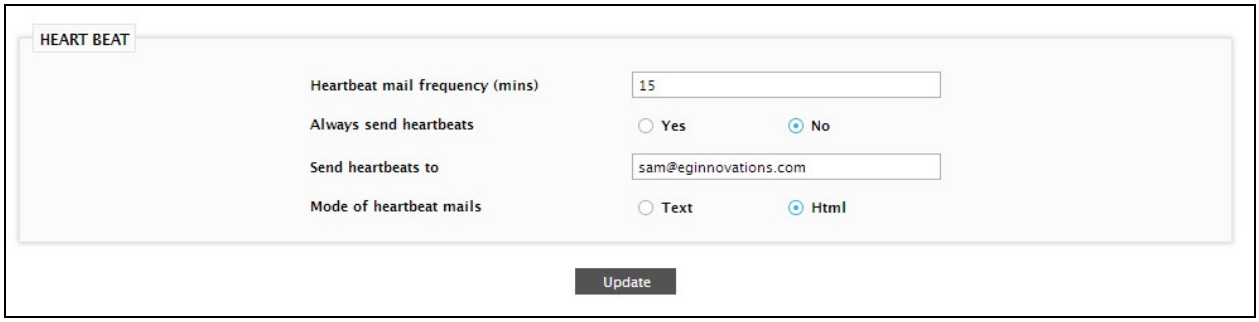

Figure 5.16: Heartbeat section of the Mail Alert Preferences page

Some administrators rarely login to the eG Enterprise system, preferring instead to rely on email alerts to be notified of problems. In such situations, if the email system that is used by the eG Enterprise system fails, administrators will not be notified of problems. A "heartbeat" function is supported by eG Enterprise to let administrators know that the email alerting functionality is operational. To configure heartbeat mails, specify the following in the **HEART BEAT** section of Figure 5.16:

- **Heartbeat mail frequency:** Specify (in minutes) how frequently the eG Enterprise system needs to send out heartbeat mails, against this parameter. If this parameter is left unconfigured, then the eG Enterprise system will not send any heartbeat mails.

  Note:

  The **Heartbeat mail frequency** should be greater than or equal to the **Maximum time between email alert checks**.

- **Always send heartbeats:** This parameter and the rest of the parameters in this section will appear, only if a valid integer value is provided against the **Heartbeat mail frequency** text box. If the heartbeat mails are to be sent out regardless of whether or not email alerts have been sent during the specified frequency, then, set the **Always send heartbeats** flag to **Yes**. By setting this parameter to **No**, you can make sure that the heartbeat mails are sent only if email alerts are not received by the users during the set **Heartbeat mail frequency**.
- **Send heartbeats to**: To ensure that the heartbeat mails are sent to specific individuals, provide their mail ids as a comma-separated list in this text box.
- **Mode of heartbeat mails**: You can also choose between sending the mails as **Html** or **Text**, by selecting the corresponding option from this section.

Figure 5.17 depicts a sample heartbeat mail.

Figure 5.17: A sample heartbeat mail

5.2.3.4 Alarm Escalation

To ensure the continuous availability of mission-critical IT services, it is essential that problems be detected at the earliest and remedial action be initiated immediately. Naturally, the performance of an IT operations team is assessed by its ability to proactively isolate problems and by the speed with which the identified issues are fixed. As most IT operations teams are required to support strict service level guarantees, problems that remain unnoticed or unresolved for long periods of time could result in service level violations, warrant severe penalties, and ultimately even impact the reputation of the service provider.

The eG Enterprise suite, with its patented correlation technology and its multi- modal (email/SMS/pager/console) problem alerting capability accurately identifies potential issues in the monitored environment, and intimates the concerned IT operators before any irredeemable damage is done. To enable IT managers to proactively track the performance of their operations teams, eG Enterprise also includes a time-based alarm escalation capability. With this capability, when a problem remains unresolved for a long time period, the eG Enterprise manager automatically escalates the alarm to one or more levels of IT managers. The alarm escalation is based on a pre-defined escalation period, which is configured by the administrator of eG Enterprise.

The escalations are personalized for each user - i.e., each user in the eG Enterprise system is associated with multiple levels of managers. When an alert that has been sent to a user is not resolved within the escalation period, the alert is forwarded to the first level of management. If the problem remains unresolved for another escalation period, the second level of management is informed, and so on. By hierarchically escalating problems to IT managers, eG Enterprise ensures that the management staff stays informed of the state of the mission-critical IT services they control, and that they can intervene in a timely manner to ensure quick and effective resolution to key problems.

The settings for escalation can be defined in the **ALARM ESCALATION** section (see Figure Figure 5.18) of the **MAIL ALERT PREFERENCES** page. To access this section, click on the **Alarm Escalation** node in the **Settings** tree in the left panel of Figure 5.6.

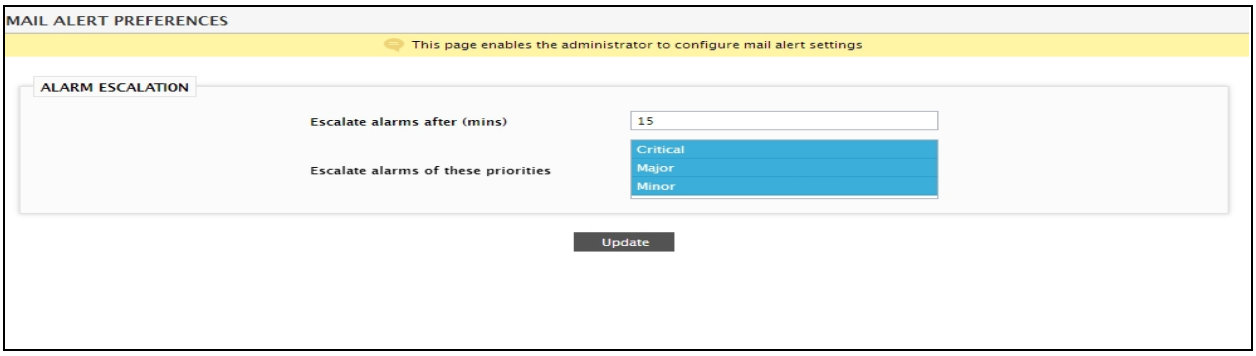

Figure 5.18: Alarm Escalation section of the Mail Alert Preferences page

- **Escalate alarms after:** Here, specify the duration beyond which the eG Enterprise system needs to escalate a problem to the next level.

  Note:

  The duration specified against the **Escalate alarms after** text box, should be greater than or equal to the **Maximum time between email alert checks**.

- **Escalate alarms of these priorities**: This parameter will appear only if a valid integer value is provided in the **Escalate alarms after** text box. You can indicate the type of problems that need to be escalated, by selecting one/more of the alarm priorities listed here.

  Note:

  For alarm escalation to work effectively, the **Escalate alarms of these priorities** specification should be the same as or a subset of the **Alarms by mail / SMS**setting (seeFigure 6.61) for a user. For example, if a user is configured to receive email/SMS alerts for **Critical**issues only and the **Escalate alarms of these priorities** parameter is set to Major and Minor, then the **Critical**alerts will not be escalated at all. However, if the **AlarmEscalationType** is set to Critical, Major, and *Minor*, or simply Critical, then the **Critical**alerts will be escalated.

Figure 5.19 depicts a sample escalation mail:

Figure 5.19: An escalated email alert

Note:

- An alarm acknowledgement is an assurance from the user – typically, a help desk executive - that theproblem in question is being looked into and will be resolved in time. If an acknowledged alarm isescalated, you may sometimes be unnecessarily inviting the help desk manager's intervention to resolve a problem that is already been investigated by a help desk executive, and which could verywell be resolved by that executive if he/she is given the time. To avoid such meaningless escalations,some administrators may prefer not to escalate acknowledged alarms. In order to stop escalating the acknowledged alarms, set **the StopEscalationOnAlarmAck** flag in the **[MISC_ARGS]** section of the **eg_services.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory) to **Yes**. Thisway, the escalation mails will not be sent for the acknowledged alarms until the help desk executivechooses otherwise.
- If administrators choose to escalate acknowledged alarms, they can additionally configure the eG manager to send the acknowledgement description as part of the email/escalation alert. This way, help desk managers can understand that the problem is being looked into, know who is looking into it, and demand status updates from that specific executive. If you wish to include the acknowledgment description in the email and escalation alarm mails, then set the **AckDetailsWithEmailAlerts** flag in the **[MISC_ARGS]** section of the **eg_services.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory) to **Yes**. By default, turning this flag on will ensure that the descriptions related to the **Top 3** acknowledgments based on the acknowledged time are included in the email/escalation mails. This default number can however be overridden. For this, set the value of your choice against the **NoOfAcknowledgementsInMail** flag in the **[MISC_ARGS]** section of the **eg_services.ini** file. If both the **StopEscalationOnAlarmAck** and **AckDetailsWithEmailAlerts** flags are set to **Yes**, then the **StopEscalationOnAlarmAck** flag takes precedence and the alarm escalation stops for an acknowledged alarm.

5.2.3.5 Shift Period Configuration

Some environments - especially the ones that span geographies - could have operators working in shifts; for instance, an MSP environment could comprise of one/more user groups, which might work only in the nights, in order to provide help-desk services to the customers in a particular geographic region. These users naturally, would want to receive email alerts of issues only during their working hours; during the rest of day, they may prefer to be alerted via SMS. To facilitate this, eG Enterprise allows you to configure shift periods for individual users. For this, click on the **Shift Periods** node of the **Settings** tree in the left panel of Figure 5.6.

Separate shift periods can be configured for receiving email alerts, SMS alerts, and escalation mails. By default, this capability is bundled with the eG Enterprise suite. Accordingly, the **Allow shift period configuration** flag is set to **Yes** by default in the **SHIFT PERIODS CONFIGURATION** section (see Figure 5.20) of the **MAIL ALERT PREFERENCES** page. You can then proceed to configure the **Maximum number of day-shift combinations** as well. By default, a maximum of 5 day-shift pairs can be configured per user. You can override this default setting by choosing a different number from the **Maximum number of day-shift combinations** list.

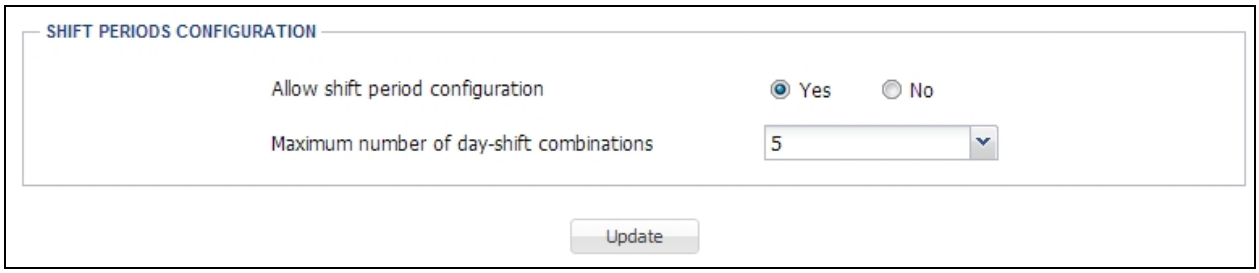

Figure 5.20: The Shift Period Configuration section of the Mail Alert Preferences page

In environments where shifts are not relevant, shift period configuration would be meaningless. In such environments therefore, you can disable this capability by setting the **Allow shift period configuration** flag to **No**. When this is done, the **Maximum number of day-shift combinations** list will automatically disappear.

5.2.3.6 Mail Log Details

The eG manager can be configured to send out email alerts of performance issues to specified recipients. However, in environments where the mail server is not properly configured, critical problem information might not reach users on time, thus causing problems to persist, aggravate, or be unnecessarily escalated to higher authorities. It is hence necessary to continuously track the status of the email activity generated by the eG manager and also provide administrators with alternative access to problem information (in the event of failure of the mail system). To ensure this, eG Enterprise allows mail logging. To enable mail logging, first, click the **Mail Log Details** node of the **Settings** tree in the left panel of Figure 5.6. The right panel will then change to display the **MAIL LOG DETAILS** section as depicted by Figure 5.21.

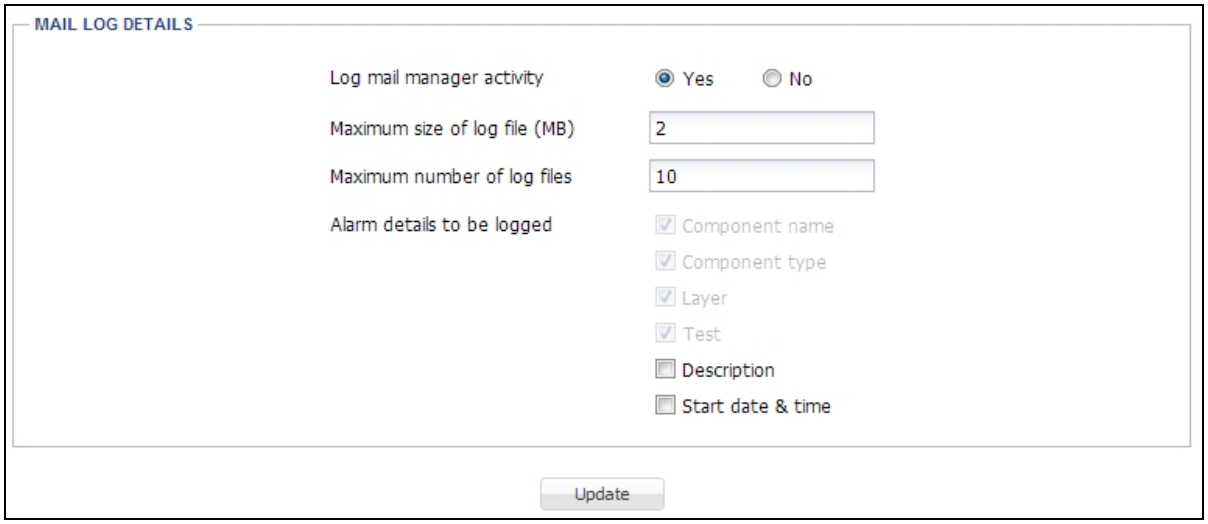

Figure 5.21: The Mail Log Details section of the Mail Alert Preferences page

By default, the mail logging capability is disabled. To enable the capability, all you have to do is to enable a flag called **Log mail manager activity** in the **MAIL LOG DETAILS** section as shown in Figure 5.21. By setting this flag to **Yes**, you can ensure that a log file is created in the **opt/egurkha/manager/log/egmailmanager** directory, where the following alarm details are logged by default:

- Whether the email alert was successfully sent or not
- The severity (critical/major/minor) of the problem for which the alert was raised
- The problem component
- The problem component-type
- The problem layer
- The test that reported the problem measurement

If required, you can modify the default settings, by selecting/deselecting the relevant check boxes in this section. For instance, by selecting the **Description** and **Start date & time** check boxes, you can make sure that the mail log also records the details of problem descriptors, a brief description of the problem, and the problem date and time.

Assume that the Processes test has raised an alert indicating that a critical process is not running on the event log server. The default log entry that corresponds to this event would be: Mail send failed [Critical,Event92,Event Log,Application Processes,Processes.

If you have enabled the **Description** and the **Start date & time** check boxes, the log entry would be:

Mail send failed [Jan.21.2008 14:21:28,Critical,Event92,EventLog,ApplicationProcesses,Processes {+firefox -> Process not running}.

Besides, to ensure that a single log file is not overloaded with problem details or does not grow enormously in size, you can trigger the creation of additional log files as soon as the size of a log file exceeds a pre-configured limit. This ceiling can be set using the **Maximum size of the logfile(MB)** text box. Once this limit is reached, the eG manager creates a new log file, copies the old data to it, and starts logging the latest information to the older log file. In environments where there is excessive mail activity, this can result in a large number of log files, which might in turn consume too much space on the eG manager host. In such a case, you can conserve space on the eG manager host using the **Maximum number of log files** configuration. If such a limit is configured, then the eG manager will continue creating new log files only till such time that the said file limit is reached. Beyond this point, no additional log files will be created; instead, the eG manager will overwrite the currently open log file with the newer problem information.

5.2.3.7 Enabling/Disabling Email/SMS Filtering

By default, a user receives email/SMS alerts for all issues pertaining to all components assigned to him/her. In some circumstances, the user may not want to receive all of these alarms. For instance, in a large, multi-tier infrastructure, a user may be monitoring all the applications and network devices involved in supporting a business service. However, the user may have primary responsibility only for some of the components supporting the business service (e.g., a network administrator's primary responsibility is to monitor the network devices). In such cases, while the user may want to view the status of all the components of the business service, he/she may want to receive email or SMS alerts pertaining to specific components of the infrastructure alone (e.g., network devices).

To enable such selective alerting, eG Enterprise provides administrators with the option to configure the eG manager to **not send out email/SMS alerts related to specific layers/components/component-types/tests for specific users**.

By default, the ability to **filter mail/SMS alerts** is disabled. To enable it, first, click the **Filter Mail/SMS Alerts** node of the **Settings** tree in the left panel of Figure 5.6. This will open the **FILTER MAIL/SMS ALERTS** section in the right panel (see Figure 5.22). Here, set the **Allow mail filter configuration** flag to **Yes**. This will allow **Admin** users to configure email/SMS filters for user profiles. To allow non-admin users to alter the email/SMS settings defined by the **Admin** user, set the **Allow non-admins to update** flag to **Yes**.

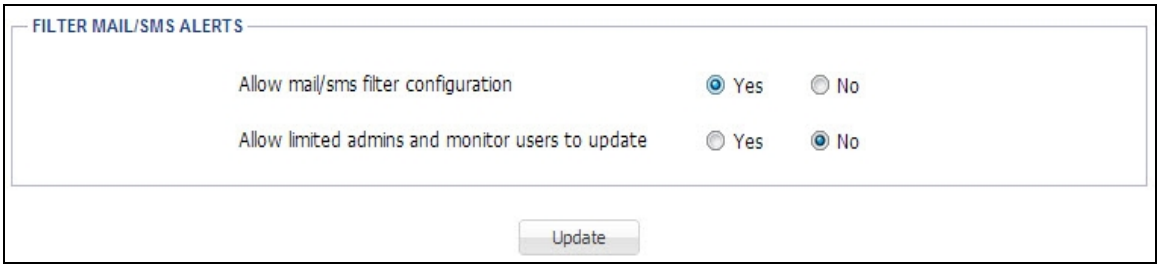

Figure 5.22: Filter Mail Alerts section of the Mail Alerts Preferences page

To update the changes, click the **Update** button in Figure 5.22.

Note:

Like email activity, the eG manager can also be optionally configured to maintain logs of the alarms generated by it. In some environments, this could serve as a tool for effective 'postmortem' analysis of problem situations - i.e., administrators can use the logs to be informed of problems that might have occurred during a period of their non-availability. Auditing of alarms reveal when an alarm was generated, the nature of the problem reported by the alarm, what were the problems clubbed together, and which problem is the source of a set of related problems. In some other environments, alarm logging could be a 'must have'; such environments may host applications that read the alarm information stored in the logs for converting them into email/SMS alerts.

To enable alarm logging, do the following:

- Edit the **eg_services.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
- In the **[MISC_ARGS]** section of the file, set the **NEED_ALARMLOG** parameter to **true**, to enable alarm logging.
- To indicate what type of alarm information needs to be stored in the logs, use the **ALARMLOG_FORMAT** parameter. By default, the alarm priority, the problem layer, the alarm ID, the problem description, the problem component name, the component type, and the last measurement time are logged in the log file. Accordingly, the default **ALARM_LOG** format is: *PRIORITY:LAYER:PROBLEM_DESC:ALARM_ID:COMP_NAME:COMP_TYPE:MSMT_TIME*
- Indicate the separator that separates alarm details against the **ALARMLOG_SEPRT** parameter.
- Also, specify how frequently (in milliseconds) the alarm log configurations need to be refreshed against the **ALARM_LOG_REFRESH_INTERVAL** parameter. If this parameter is set to 0 or is left blank, then, by default, the alarm log configurations will be refreshed every 30 minutes - i.e., 1800000 milliseconds.
- Besides, to ensure that a single log file is not overloaded with problem details or does not grow enormously in size, you can trigger the creation of additional log files as soon as the size of a log file exceeds a pre-configured limit. This ceiling can be set using the **ALARM_LOG_MAX_SIZE** parameter.

Once this limit is reached, the eG manager creates a new log file, copies the old data to it, and starts logging the latest information to the older log file. In environments where tto many alarms are generated, this can result in a large number of log files, which might in turn consume too much space on the eG manager host. In such a case, you can conserve space on the eG manager host by specifiying the maximum number of log files that can be created using the **ALARM_LOG_MAX_FILES** configuration. If such a limit is configured, then the eG manager will continue creating new log files only till such time that the said file limit is reached. Beyond this point, no additional log files will be created; instead, the eG manager will overwrite the currently open log file with the newer problem information.

- Finally, save the **eg_services.ini** file.

Besides intimating users of problems with components and their subsequent return to normalcy, eG Enterprise can also be configured to send out emails / SMS when the state of a component becomes **UNKNOWN**. To configure unknown state mails/SMS, do the following:

1. Edit the **eg_services.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory to set the values of the parameters in the **[UNKNOWN_STATUS_REPORTING]** section, as shown below:

```
UnknownStateMail=No
UnknownStateSMS=No
UnknownStateInfoMail=
UnknownStateInfoSMS=
UnknownStateMailList=
UnknownStateSMSList=
DefaultUnknownStatePeriod=
```

   In order to receive unknown state mails, set the **UnknownStateMail** parameter to **Yes**. The default is **No**. Similarly, specify **Yes** against **UnknownStateSMS** parameter to be able to receive SMS alerts when the state of a component changes to Unknown. The eG Enterprise system can also be configured to send out unknown state alerts even if the state of a test descriptor changes to Unknown, by setting the **UnknownStateInfoMail** parameter to **Yes**. To receive SMS alerts to that effect, set **UnknownStateInfoSMS** to **Yes**. You can even specify the users who should receive the email/SMS alerts by providing a comma-separated list of mail ids and mobile numbers (as the case may be) against the **UnknownStateMailList** and **UnknownStateSMSList** parameters, respectively. The eG Enterprise system will email/SMS the configured users only when a component remains in the Unknown state for the duration specified in the **DefaultUnknownStatePeriod** parameter. This duration can also be set specific to a test, by inserting the test name as a parameter in the **[UNKNOWN_STATUS_REPORTING]** section and providing a value against it, as shown below:

```
OraTableSpacesTest=60
```

   The **DefaultUnknownStatePeriod** will automatically apply to all those tests which do not have a specific unknown state period or which have been misspelt in the **[UNKNOWN_STATUS_REPORTING]** section.

2. Finally, save the **eg_services.ini** file.
3. Figure 5.23 depicts a typical unknown state mail alert.

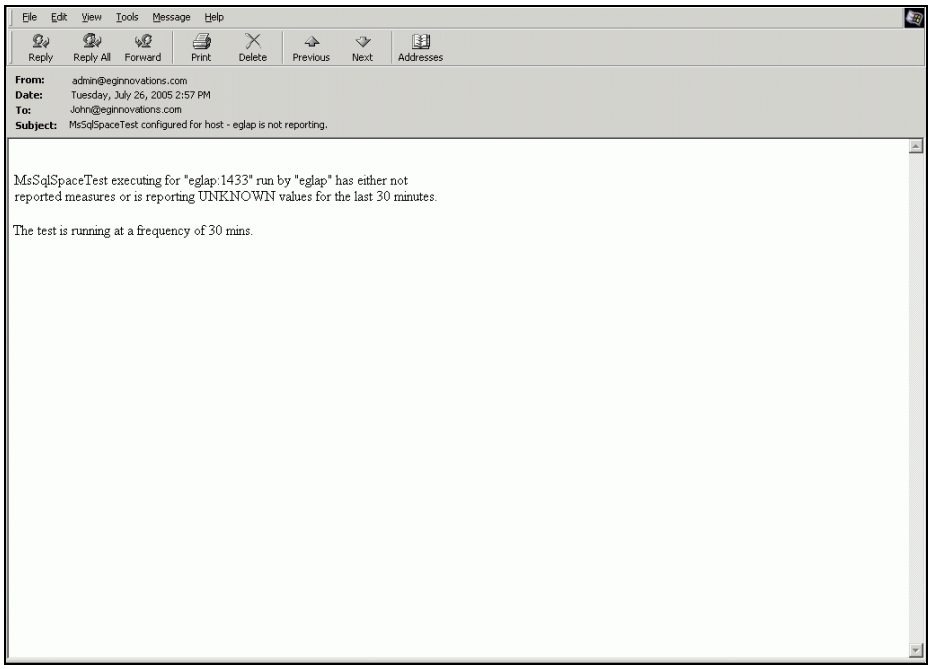

Figure 5.23: A sample unknown state mail alert

The eG manager also has the ability to monitor itself and to generate alarms to the administrator when any of the eG manager processes fail, or when the database server used by the eG manager is inaccessible. In these cases, the eG manager alerts the administrator by generating email alerts to the **eG Administrator mail ID**.

5.3 Configuring the Database Settings

Figure 5.24 depicts how the eG database settings can be modified. To access this page, select the **Database** option from the **Settings** tile.

The eG database stores:

- The instantaneous measurements reported by the agents,
- Detailed diagnosis measures (if applicable),
- Hourly, daily and monthly trends that summarize the measurements made in the past, and
- Audit logs
- The measure data used for generating reports
- Configuration data and details of configuration changes (if the eG license enables Configuration Management)

To ensure that database size does not grow continuously, the eG manager automatically deletes old entries from the database. The **Database Purge Periods** section shown in Figure 5.24 govern when the manager deletes old entries from each of the eG database tables. You can also schedule day-end activities such as trend computation, scheduled mail generation, and cleanup, by specifying the time in the **Day-end activities start at this time** field.

Note:

- The **DATABASE PURGE PERIODS** configured for the hourly, daily, and monthly trend data automatically applies to hourly, daily, and monthly capacity computations as well.
- Typically, capacity computations are performed for a test only if the **Capacity Planning** flag is explicitly enabled for that test using the **CUSTOM SETTINGS - CAPACITY PLANNING** page (see Figure 2.26 in *The eG Reporter* document) of the eG administrative interface. By default, whenever the eG database cleanup process runs, it checks the status of the **Capacity Planning** flag of tests and cleans up the old capacity computations of only those tests for which **Capacity Planning** is enabled **at the time of database cleanup**. Sometimes, this flag could have been turned on for a few tests for a while, and later turned off. If say, the flag had been turned off before the database cleanup process runs, then the capacity values calculated during this period will remain in the database until the **Capacity Planning** capability is switched on again for those tests; this is because, cleanup will ignore the old capacity computations for those tests for which the **Capacity Planning** flag is set to **Disabled** at the time of cleanup. This stale data therefore may end up occupying critical database space for an infinite period of time! To avoid this, the eG Enterprise system provides you with the option to 'force the cleanup of old capacity data'. By default, this capability is disabled. This is why, by default, the **CanForceCapacityCleanup** flag in the **[MISC_ARGS]** section of the **eg_services.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory) is set to **false**. To enable this capability, set this flag to **true** and save the **eg_services.ini** file. Once this is done, then, the next time the database cleanup process executes, it will automatically delete the old capacity computations of all tests, regardless of the current status of the **Capacity Planning** flag - i.e., regardless of whether/not **Capacity Planning** is **Enabled** for those tests.

To optimize accesses to the database, the eG manager uses connection pooling. By using a pre-established set of connections and multiplexing requests over these connections, the eG manager ensures that individual connections are not established and closed for each request. The **Database Connection Pool Settings** section in Figure 5.24 governs the initial and the maximum number of connections in the connection pool. Consider increasing the number of connections in the connection pool as the number of components monitored or the number of users accessing the system increases. By default, the **Maximum connections** parameter (see Figure 5.24) is set to 100.

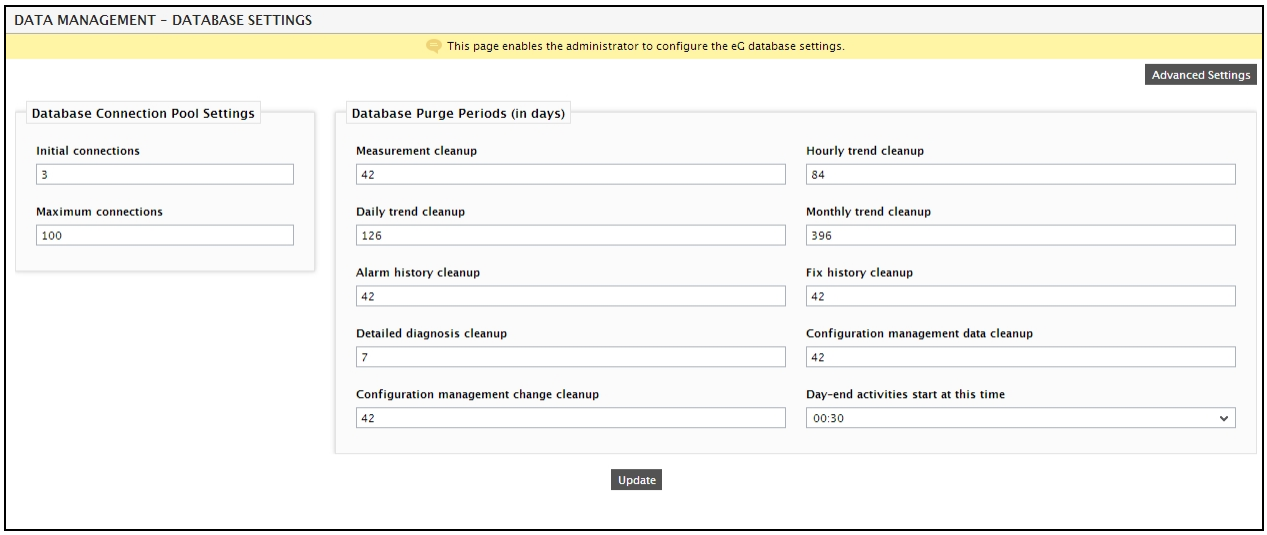

Figure 5.24: Configuring the database settings

If you want to configure cleanup frequencies for the detailed diagnosis data reported by specific tests, click the **Advanced Settings** button in Figure 5.24. Figure 5.25 then appears displaying the default detailed diagnosis cleanup frequencies for tests executing on **All** managed component-types (by default). In the default scenario therefore, if a large number of components have been managed in your environment, then, a long list of tests will appear in Figure 5.25, making it difficult for you to locate the tests of interest to you. In such situations, you can narrow your search for tests by picking a particular **Component type** from Figure 5.25; this way, only those tests that execute on the chosen type and their default diagnosis frequencies will be listed. You can override the default frequencies by changing them. To save the changes, click the **Update** button in Figure 5.24.

Note:

Some tests may be common to more than one component type - for e.g., Network test, operation system-level tests, etc. Such tests will be listed under each of these **Component type**s. However, once the diagnosis cleanup frequency of such a test is changed for one component type, it will automatically change the diagnosis cleanup frequency for this test across all other associated component types as well.

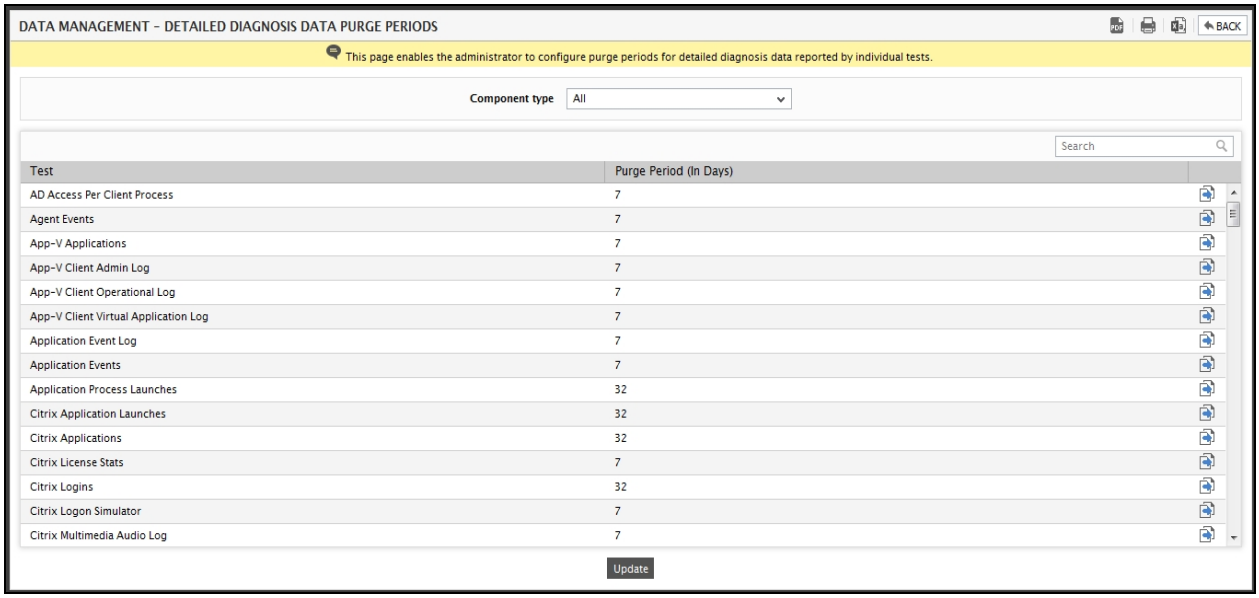

Figure 5.25: Configuring the cleanup frequency for detailed diagnosis data pertaining to specific tests

Also, if you want the frequency that you have configured for one test to be applied to one/more other tests, regardless of the component-type to which they are mapped, you need not manually change the frequency of each test to this effect. All you need to do instead is follow a simple sequence of mouse clicks to effortlessly update multiple tests across component types with the diagnosis cleanup frequency of a single test. For this, click on the button alongside the *source* test - i.e., the test with a diagnosis frequency that needs to be applied to other tests. This will invoke Figure 5.26.

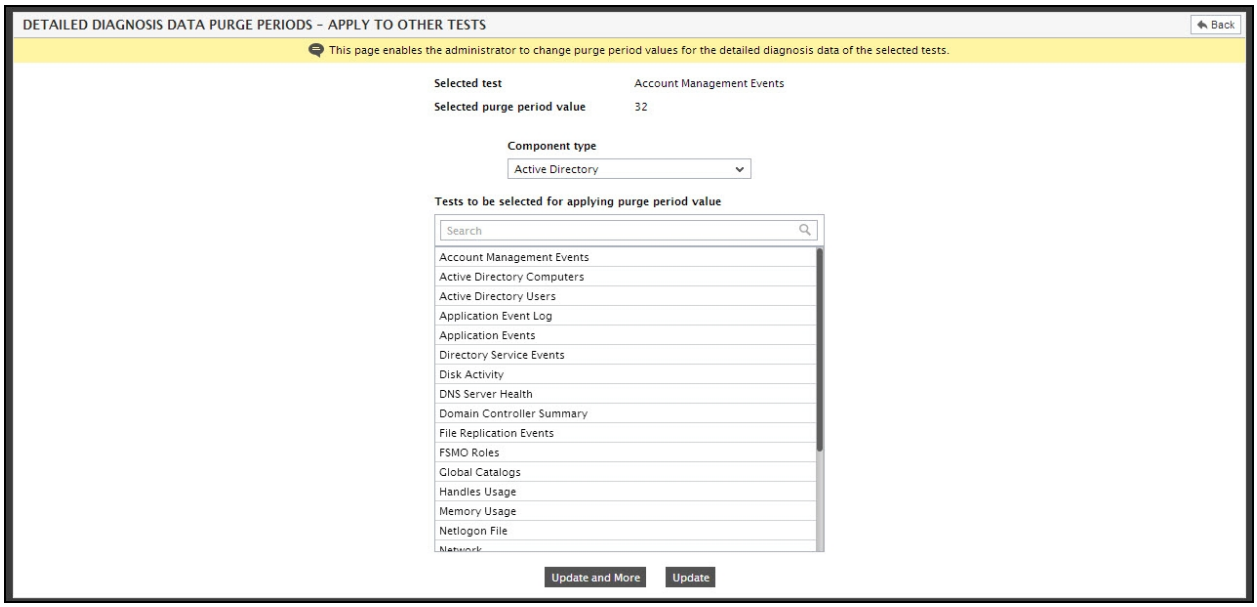

Figure 5.26: Applying the diagnosis cleanup frequency configured for a test to other tests

By default, the **All** option will be chosen from the **Component type** list in Figure 5.26. Accordingly, the list of tests displayed below will pertain to all managed component types, by default. You can filter this tests list by selecting a particular component type from the **Component type** list; this ensures that the tests associated with the chosen type alone are displayed in Figure 5.26.

From the tests list in Figure 5.26, you can pick the tests to which the cleanup frequency of the *source* test needs to be applied. If you do not want to pick another source test, then simply click the **Update** button to save the changes and exit the page. If you want to pick a different source test and apply its cleanup frequency to a few other tests, then click the **Update and More** button.

If you wish to print the purge period of the tests, then, you can do so by clicking the **Print** icon in Figure 5.25. Likewise, you can export the purge period of the tests as an **Excel** file or a **PDF** file by clicking the appropriate icons provided in Figure 5.25.

Note:

If a test is picked as a *source*, then, such a test will not be available for selection in the tests list of Figure 5.27, regardless of the component-type you choose.

Note:

- The **eg_services.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory comprises of an **IndexRebuild** flag (in its **[MISC_ARGS]** section). If this flag is set to **YES**, then the trend manager, upon execution, will automatically initiate an index recreation process. If this process is to be performed manually, then set the **IndexRebuild** flag to **NO**.
- A file named **eg_indextables.ini** also exists in the **<EG_INSTALL_DIR>\manager\config** directory, which consists of the following entries:

- **MaxIndexTime**, which indicates the duration (in minutes) for which auto index creation will run. If this parameter is set to 60, it means that the auto index creation will run for 60 minutes, i.e., 1 hour.
- **ReBuildFrequency**, which governs how frequently the auto index creation process should run. If this parameter is set to 2, it means that the auto index creation process will run every two days.

Note:

- The eG manager keeps track of all the tables that it has created in its database in the **eg_db.ini** and **eg_dbase.ini** configuration files. If these files are truncated during the operation of the eG manager (e.g., because the eG manager system ran out of disk space and the configuration files could not be written to disk), the eG manager recreates the measurement tables in the database. The *RecreateTables* flag in the **eg_db.ini** configuration file in the **<EG_INSTALL_DIR>\manager\config** directory controls this operation. If this flag is yes, the database tables are recreated. If this flag is no, the tables are not recreated. The *RestartOnCorruption* flag controls how the eG manager functions when it detects a table in its database for which entries are missing in the **eg_db.ini** configuration file. If this flag is set to yes, the eG manager automatically restarts itself when it detects a table in the database but entries for the table are missing in the **eg_db.ini** file.
- If the eG manager has not run the data purging/cleanup operation for a long time, or if the data retention period has been reduced and a significant amount of data needs to be cleaned up from the eG database, cleaning up a lot of data at one time can impose a lot of stress on the eG database, slowing down user accesses and the performance of the eG manager. Administrators can control the number of days of data that the eG manager can purge during one run of its day-end data purging process. The *MaxDaysToCleanOnce* setting in the **[MANAGER_SETTINGS]** section of the **eg_db.ini** file controls the number of days of measurement data that is purged during one run of the day-end data purging process. A value of 10 indicates that at most 10 days of data is purged during one run. A value of -1 indicates that there is no limit on the amount of data that can be purged during one run.

Note:

- By default, the trend manager uses the connections in the connection pool for execution. To ensure that the trend manager runs on a new connection every time, but not from the connection pool, set the **TrendWithDedicatedConnection** flag in the **[MISC_ARG]** section of the **eg_services.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory to **yes**. By default, this flag is set to **no**.
- By default, the database cleanup process is executed as a thread within the eG Tomcat process. Accordingly, the **ExecuteCleanupAsProcess** flag in the **[MISC_ARGS]** section of the **eg_services.ini** file is set to **no**, by default. If you set it to **yes**, cleanup will run as a separate java process.
- To ensure that the day-end data purging process is throttled according to the availability of connections in the connection pool, a **PoolThrottleLimit** can be set in the **[MANAGER_SETTINGS]** section of the **eg_db.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory). By default, this is set to 40 (percent). This means that if connection pool usage exceeds 40%, then the cleanup process will be spaced out in such a way, so as to not impact connections to the database. You can however, change this throttle limit based on the normal connection pool usage in your environment.

5.4 Configuring Detailed Diagnosis

This **Detailed Diagnosis** option will be available in the left panel of the **AGENT SETTINGS** page (that can be accessed by selecting **Settings** option from the **Agents** tile) only if the eG license supports the detailed diagnosis capability. If not, a message stating that the license does not allow this feature will appear when you access this page.

If the detailed diagnosis capability is enabled, eG agents will perform more detailed measurements periodically. For example, an agent measures the CPU usage of a host. With the detailed diagnosis capability, the agent can determine and report the top 10 CPU consuming processes on the host. The functioning of the detailed diagnosis capability is determined by two settings, which can be configured using the **DIAGNOSIS SETTINGS** screen (see Figure 5.27):

- **Diagnosis period during normal operation**: In this text box, specify the periodicity with which the eG agents need to provide a detailed diagnosis of a measure, regardless of its state (i.e. Normal, Critical, Major, or Minor). By default, this text box will contain the value **1**.

  Example:

  Assume that the diagnosis period during normal operation is specified as 4. Also, assume that a test that provides a detailed diagnosis of its measures, runs every 60 seconds. This means that the eG agents will generate detailed measures every fourth time that the test runs, i.e. at the end of 60*4 = 240 seconds (or 4 minutes).

- **Diagnosis period during abnormal operation**: In this text box, specify the frequency with which the eG agents should report detailed diagnosis measures if they detect a problem (i.e. state is Critical, Major, or Minor). By default, this text box will contain the value **1**.

  Example:

  Assume that the diagnosis period during abnormal operation is set to 2. In such a case, if an agent detects a problem with a measurement, it starts reporting detailed diagnosis measures every second time that it performs the test.

If you specify 0 as the normal diagnosis period, but set the abnormal diagnosis period to a valid value, then the eG agents will disregard the normal frequency, and will generate detailed measures only according to the abnormal diagnosis period. The vice-versa also holds good. If both the frequencies are set to 0, then no detailed measures will be generated. In such a case, the **Application**, **Session**, and **User Thin Client Reports** cannot be generated.

Clicking the **Update** button in Figure 5.27 will register the changes made to the system.

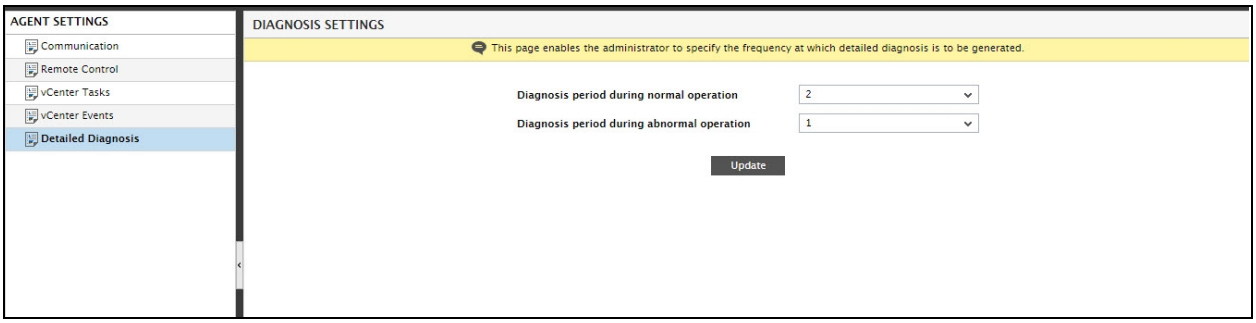

Figure 5.27: Configuring the frequency for detailed diagnosis

5.5 Defining Manager and Monitor Settings

Using the **Manager** and **Monitor** options of the **Settings** tile, you can do the following:

- Define the default settings for a few critical operations of the eG manager
- Enable logging for eG manager functions such as trending, thresholding, cleanup, etc.
- Configure the default settings for the monitor and reporter interfaces
- Defining custom logo and messages for the login, **Admin**, **Monitor**, **Reporter**, and **Configuration** management interfaces

5.5.1 Configuring Monitor Settings

To configure the **Monitor** settings, select the **Monitor** option from the **Settings** tile. When Figure 5.28 appears, a **MONITOR SETTINGS** tree will appear in its left panel with the following nodes:

- **General** - This node and its sub-nodes provide you with a number of options that will enable you to effortlessly control the look, feel, and basic operations of the eG monitoring console. Each of these sub-nodes have been dealt with in great detail later in this section.
- **Measures At A Glance** - The Monitor dashboard can be optionally configured with a **Measures At-A-Glance** section. This section can be designed to provide you with a heads up on how a few critical performance parameters are currently faring, so that the hot spots in the environment can be identified almost as soon as you login to the monitoring console. This node and its sub-nodes enable you to configure the metrics that need to feature in the **Measures At-A-Glance** section. Each of these sub-nodes have been dealt with elaborately later in this section.

5.5.1.1 Configuring General Monitor Settings

This section describes each sub-node of the **General** node and the settings that can be defined using that sub-node.

1. Clicking on the **Alarms** node invokes the **ALARMS** page in the right panel (see Figure 5.28).

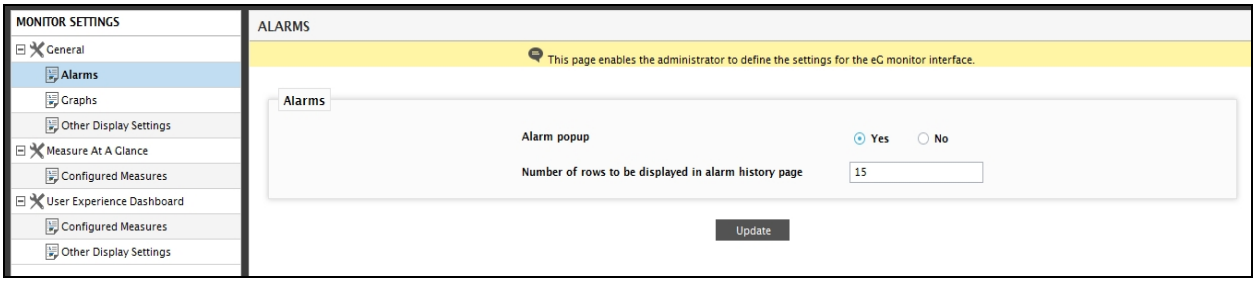

Figure 5.28: Alarms page

2. In Figure 5.28, specify the following:

   - **Alarm popup**: By default, as soon as a user logs into the eG monitoring console, a **CURRENT ALARMS** window will pop-up, listing all the current performance issues in the target environment. This default behavior is governed by the **Alarm popup** flag, which is set to **Yes** by default. If you do not want the **CURRENT ALARMS** window to pop up upon logging in to the eG monitoring console, then set the **Alarm popup** flag to **No**.
   - **Number of rows to be displayed in the alarm history**: The **Event History** page in the monitor interface provides a detailed history of events that have occurred in the target environment. Instead of allowing too many rows of events to crowd a single page of this display, you can spread the information across multiple pages. To ensure this, you can specify the maximum number of rows to be listed in one page of the event display, in the **Number of rows to be displayed in the alarm history page** text box (see Figure 5.28).

3. Click on **Update** button in Figure 5.28 to implement the changes.

4. If you click on the **Graphs** subnode, the **GRAPHS** page will appear (see Figure 5.29).

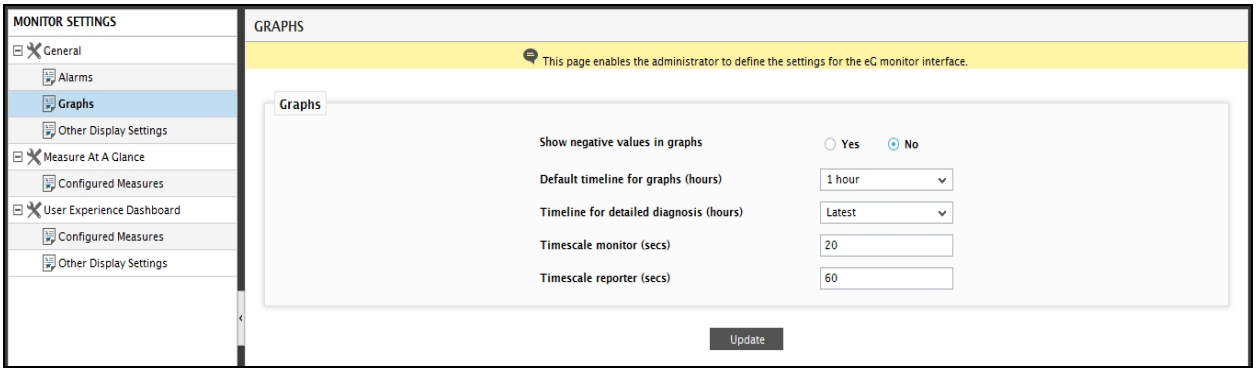

Figure 5.29: Graphs page

5. In this page, you can customize the graph displays in the monitor and reporter interfaces (see Figure 5.29), as discussed below:

   - **Show negative values in graphs** : eG Enterprise reports negative values when the eG agent is unable to collect data. By plotting these negative values in a graph, you can accurately determine the times when the eG agent was unable to extract performance data from target components. By default, the eG manager ignores these negative values while plotting graphs. However, if you want to view the

periods when the eG agent succeeded in collecting data vis-a-vis the periods when the agent failed to collect data, then you can turn on the plotting of negative values in graphs, by setting the **Show negative values in graphs** flag to **Yes**.

- **Use lines in 3D graph to show depth** : Lines in 3D graphs enhance the depth of your graph display. To achieve this, set the **Use lines in 3D graph to show depth** flag to **Yes** (see Figure 5.29).
- **Default timeline for graphs:** The default duration for measure graphs in the eG monitor interface is 1 hour. You can override this default setting by specifying a higher or lower time period (in hours) in the **Default timeline for graphs** text box.
- **Timeline for detailed diagnosis:** The default duration for detailed diagnosis in the eG monitor interface is 1 hour. You can override this default setting by specifying a higher or lower time period (in hours) in the **Time line for detailed diagnosis** text box.
- **Timescale monitor**: By default, the graphs in the monitor interface plot values averaged over every 20 seconds of the specified timeline. For instance, to plot the values of a measure gathered over an hour, by default, 180 data points will be plotted in the graph, one for every 20 seconds of data. If the default time scale remains as 20 seconds, then, longer timelines will result in a large number of data points been plotted on the graph; this, in turn, provides administrators with deeper insights into measure behavior. However, sometimes, administrators might require less granular information on the graph, so that they are able to read and analyze the graphs better. To facilitate this, eG Enterprise permits administrators to specify a custom time scale for graphs in the **Timescale Monitor** text box.
- **Timescale reporter**: Similarly, the default timescale for graphs in eG Reporter is 60 seconds. To change this setting, you need to alter the **Timescale reporter** specification in this page.

6. Click on the **Update** button to implement the changes as depicted by Figure 5.29.
7. Using the **Other Display Settings** sub-node, you can customize the other displays in the monitor interface (see Figure 5.30).

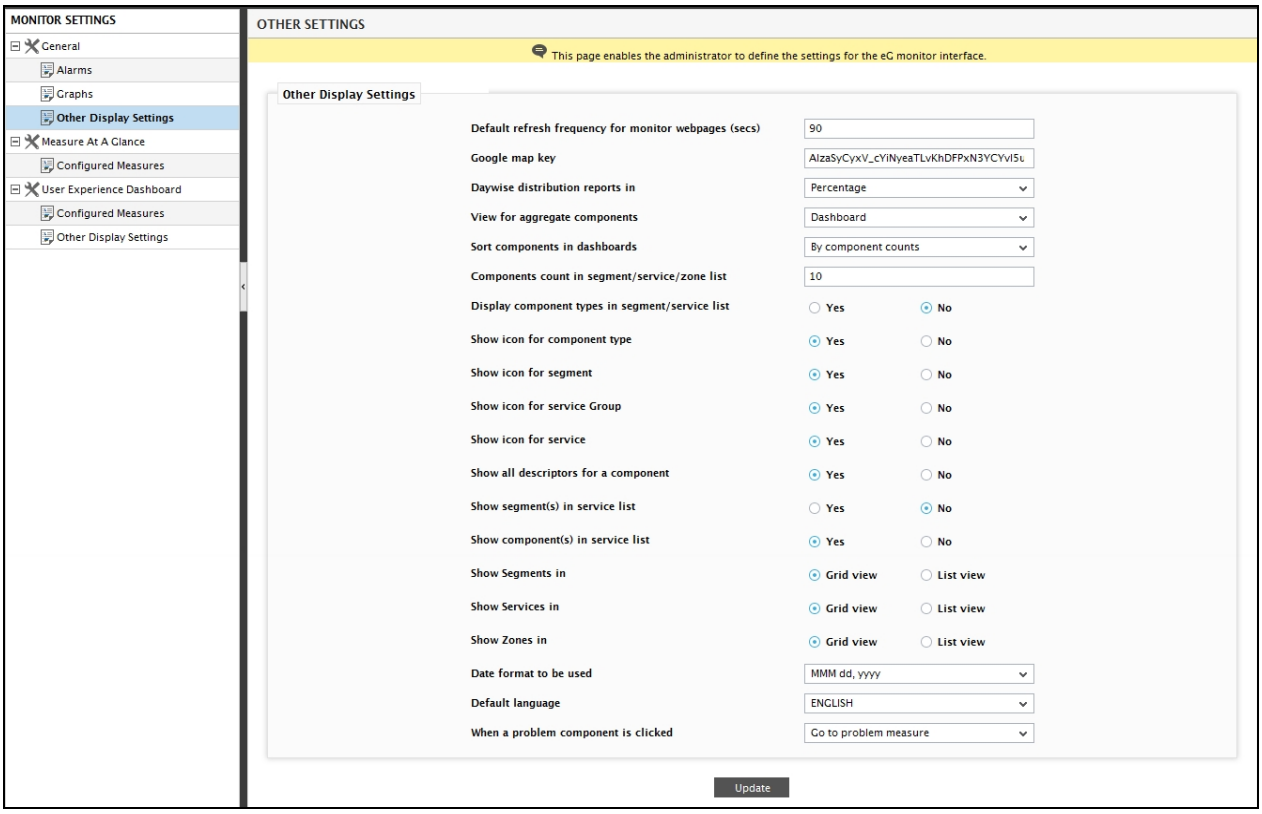

Figure 5.30: Other display settings page

8. The settings that can be configured using the **Other Display Settings** page are as follows:

   - **Default refresh frequency for monitor webpages (secs):** Indicates how often the web pages of the eG monitor module need to be refreshed. By default, this is set to 90 seconds.
   - **Google map key:** Administrators of large infrastructures spanning geographies might prefer to monitor their infrastructure by viewing it as smaller, more manageable business units. In eG parlance, these business units are termed ZONES. A zone can comprise of individual components, segments, services, and/or other zones that require monitoring. Typically, zones can be used to represent the status of the IT infrastructure in a specific geographic location. eG Enterprise allows you to drill down on a geographic map to visually figure out the exact geographic area where a zone operates, and instantly evaluate the performance of the different zones spread across the different locations worldwide. The geographic map display for zones is achieved through an integration of the eG Enterprise management console with the Google maps service. To enable the integration, a Google Maps API key is required. This key is bundled into the eG manager, and is automatically displayed here.
   - **Daywise distribution reports in** : By default, the **Event Duration Analysis** section in the **Reporter Dashboard** and **Executive Reports** depicts the duration (in minutes) for which the target environment/entity (as the case may be) had experienced performance issues every day during the given **Timeline**. Instead of the problem 'minutes' per day, if you want this day-wise event distribution to indicate the 'percentage of time' every day the target entity suffered performance degradations, then set the **Daywise distribution reports in** flag to **Percentage**.

- **View for aggregate components:** This option will be available only if the eG license enables **Metric Aggregation**. If so, then, you can pick an option from this list to indicate the default view for aggregate components in the eG monitoring console. If the **Dashboard** option is chosen, then, clicking on an aggregate component in the monitoring console will invoke the aggregate dashboard. Selecting the **LayerModel** option on the other hand, will lead you to the layer model of the aggregate component that you click on.
- **Sort components in dashboards**: By default, in the **Components At-A-Glance** section of the Monitor and Zone dashboards, the component-types are sorted in the descending order of the values in the `Count` column of the section. If you want to change this default sort order - i.e., if you want the contents of this section to be arranged in the alphabetical order of the component-types by default - set the **Sort components in dashboards** flag to **By component types**.
- **Components count in segment/service/zone list**: The **SEGMENT LIST**, **SERVICE LIST**, and **ZONE LIST** pages in the eG monitoring console reveal the current state of all the segment, services, and zones (respectively) that have been configured in the target environment. To make sure that the lists are not cluttered, by default, against each segment/service/zone displayed in these pages, only the top-10 components associated with that segment/service/zone will be displayed. This is because, the **Component count in segment/service/zone list** is set to **10** by default. The top-10 component list is arrived at by sorting all the components associated with a segment/service/zone on the basis of their current state and then arranging them in the alphabetical order of the component names. If you want more components to be displayed in these pages, then, you can specify a value of your choice in the **Component count in segment/service/zone list** text box. To display all components, specify *-1* here.
- **Display component types in segment/service list** : Some administrators might prefer to view the abnormal component-types associated with a problem segment/service, rather than component names, in the **SERVICE LIST** or **SEGMENT LIST** pages. To achieve this, such administrators can set the **Display component types in segment/service list** flag to **Yes**.
- **Show icon for component type** : By default, an icon prefixes every component listing in the **COMPONENT LIST** page of the eG monitor interface, representing the component-type. In target environments where a large number of components are managed, these icons could cramp the component display. To avoid this, you can set the **Show icon for component type** flag to `No`. Doing so removes the icons from the **COMPONENT LIST** page, and replaces them with the names of the corresponding component-types.
- **Show icon for segment**: By default, an icon prefixes every segment listing in the **SEGMENT LIST** page of the eG monitor interface, representing the segment. In target environments where a large number of segments are managed, these icons could cramp the display. To avoid this, you can set the **Show icon for segment** flag to `No`. Doing so removes the icons from the **SEGMENT LIST** page.
- **Show icon for service group**: By default, an icon prefixes every service group listing in the **SERVICE GROUPS LIST** page of the eG monitor interface, representing the service group. In target environments where a large number of service groups are configured, these icons could cramp the display. To avoid this, you can set the **Show icon for service group** flag to `No`. Doing so removes the icons from the **SERVICE GROUP LIST** page.
- **Show icon for service**: By default, an icon prefixes every service listing in the **SERVICE LIST** page of the eG monitor interface, representing the service. In target environments where a large number of

services are configured, these icons could cramp the display. To avoid this, you can set the **Show icon for service** flag to No. Doing so removes the icons from the **SERVICE LIST** page.

- **Show all descriptors for a component:** By default, when you click on a particular layer in the monitoring model of a component, all the tests mapped to that layer, and all descriptors that have been enabled for the descriptor-based tests (if any), will be displayed in the **Tests** panel of the layer model page. Some tests support a large number of descriptors. For instance, the UserProfileTest reports the profile size of each and every user to a Citrix or Terminal server. When such descriptors are displayed, they will naturally crowd the **Tests** panel. To enhance the 'look and feel' of the layer model page, you can switch off descriptor display by default. To ensure this, set the **Show all descriptors for a component** flag to **No**.
- **Show segment(s) in service list**: By default, against each service displayed in the **SERVICE LIST** page in the eG monitoring console, only the components (top-10, by default) associated with that service will be displayed. This is why, the **Show segment(s) in service list** flag is set to **No** by default. If you want the segments associated with a service to also be displayed alongside the service name, then set this flag to **Yes**.
- **Show component(s) in service list:** By default, against each service displayed in the **SERVICE LIST** page in the eG monitoring console, the components (top-10, by default) associated with that service will be displayed. Accordingly, the **Show component(s) in service list** page is set to **Yes** by default. To hide the component list, set this flag to **No**.

  Note:

  If required, you can set both the **Show segment(s) in service list** and **Show component(s) in service list** flags to **No**. If this is done, then neither the segments nor the components associated with a service will be displayed against the service name in the **SERVICE LIST** page. Instead, you will only find a list of fully-configured services and their state in this page.
- **Show Segments in**: Indicate how segments are to be displayed by default in the eG monitoring console. Select the **List view** option if the segments are to be displayed in a vertical list. By default, the **Grid view** option is chosen indicating that the segments will be listed in a grid format.
- **Show Services in**: Indicate how services are to be displayed by default in the eG monitoring console. Select the **List view** option if the services are to be displayed in a vertical list. By default, the **Grid view** option is chosen indicating that the services will be listed in a grid format.
- **Show Services in**: Indicate how services are to be displayed by default in the eG monitoring console. Select the **List view** option if the services are to be displayed in a vertical list. By default, the **Grid view** option is chosen indicating that the services will be listed in a grid format.
- **Show Zones in:** Indicate how zones are to be displayed by default in the eG monitoring console. Select the **List view** option if the zones are to be displayed in a vertical list. By default, the **Grid view** option is chosen indicating that the zones will be listed in a grid format.
- **Date format to be used**: The default date format for the eG user interface is MMM dd, yyyy. You can change this default format depending upon the country you are in, by selecting a different format from the **Date format to be used** list

- **Default language**: The eG manager provides multi-language support, but the default language is 'ENGLISH'. To configure the eG manager to support a different language by default, select the language of your choice from the **Default language** list.
- **When a problem component is clicked:** Select an option from this list to indicate what should happen when a problem component is clicked in the eG monitoring console. By default, clicking on a problem component leads you the layer model page, where you can directly view the problem layer, problem test, and problem measure. Accordingly, the **Go to problem measure** option is chosen from this list by default. On the other hand, if you want only the layer model of the problem component to appear when you click on it, then, pick the **Go to layer model** option from this list. In this case, you will have to manually drill down from the layer model to the problem measure.

9. Finally, click the **Update** button to register the changes as depicted by Figure 5.30.

5.5.1.2 Configuring Measures At-A-Glance

By default, the Monitor Dashboard provides a **Measures At-A-Glance** panel, which allows users to view the min/max values of critical measurements updated in real-time. Users can thus receive instant status updates on sensitive performance parameters, and can also accurately determine, at a single glance, the pain points of an infrastructure. To configure the measures which need to be listed in the **Measures At-A-Glance** section, you can use the **Measures At A Glance** node in the **MONITOR SETTINGS** tree-structure and its sub-nodes.

1. First, expand the **Measures At A Glance** node in the **MONITOR SETTINGS** tree in the left panel. Then, click the **Configured Measures** sub-node (see Figure 5.31).
2. The **CONFIGURED MEASURES** page will then appear in the right panel (see Figure 5.31).

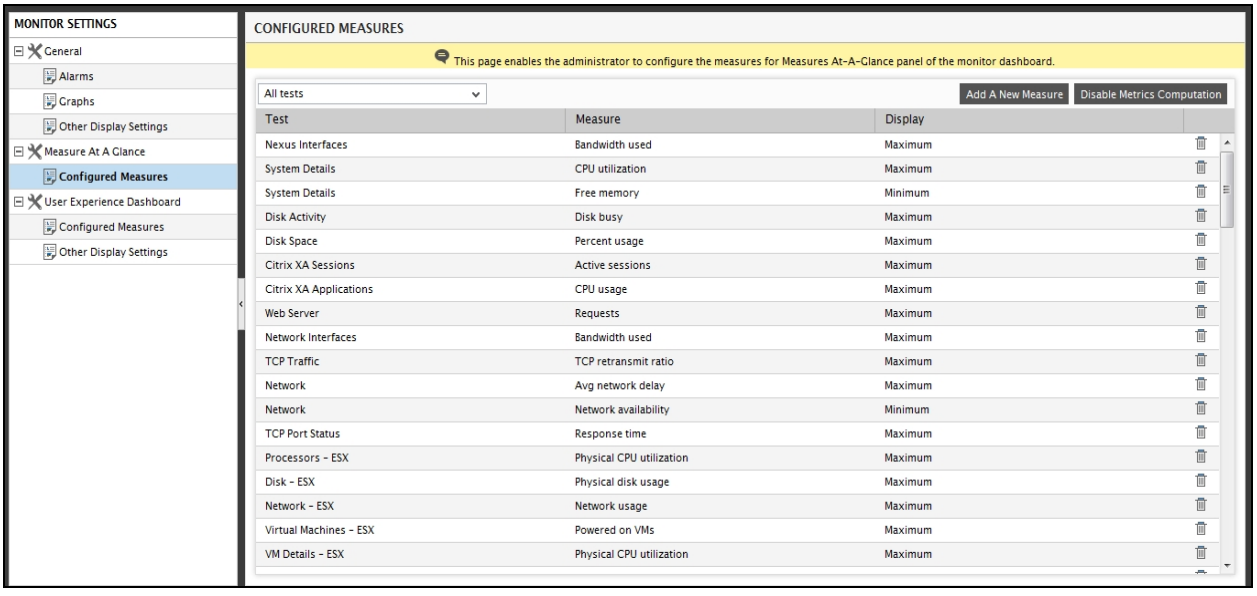

Figure 5.31: The MEASURES AT-A-GLANCE CONFIGURATION page

3. Since the **Measures At-A-Glance** section appears by default in the Monitor Dashboard, the **CONFIGURED MEASURES** page of Figure 5.31 displays the complete list of metrics pre-configured for that section. If you

want to hide this section from the Monitor Dashboard, just click the **Disable Metrics Computation** button in Figure 5.31.

4. On the other hand, if you to add a new measure to the section, click on the **Add New Measure** button in Figure 5.31. An **ADD A NEW MEASURE** pop up window will then appear (see Figure 5.32).

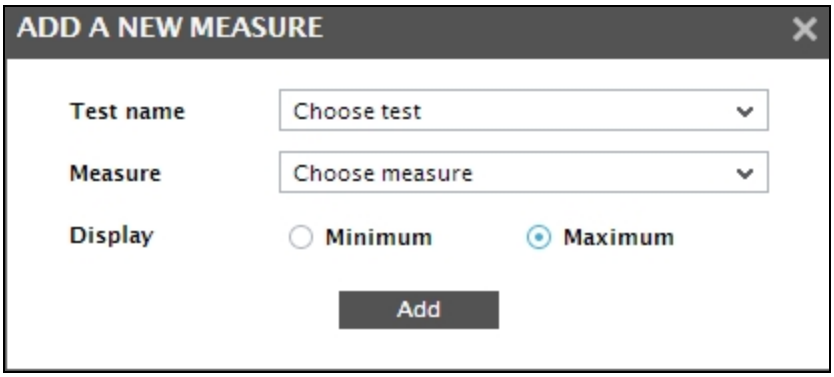

Figure 5.32: Adding a new measure

5. To add more measures to the **Measures At-A-Glance** section, first select a test from the **Test name** list in Figure 5.32. All the measures that are reported by the chosen test will then populate the **Measures** list. From this list, select the measure that should appear in the **Measures At-A-Glance** section of the Monitor Dashboard. Then, indicate whether the maximum/minimum value of the measure should be displayed in the **Measures At-A-Glance** section, by selecting either the **Maximum** or the **Minimum** flag against the **Display** section. To add the measure, click on the **Add** button therein. In the same way, you can add multiple measures for display in the **Measures At-A-Glance** section.

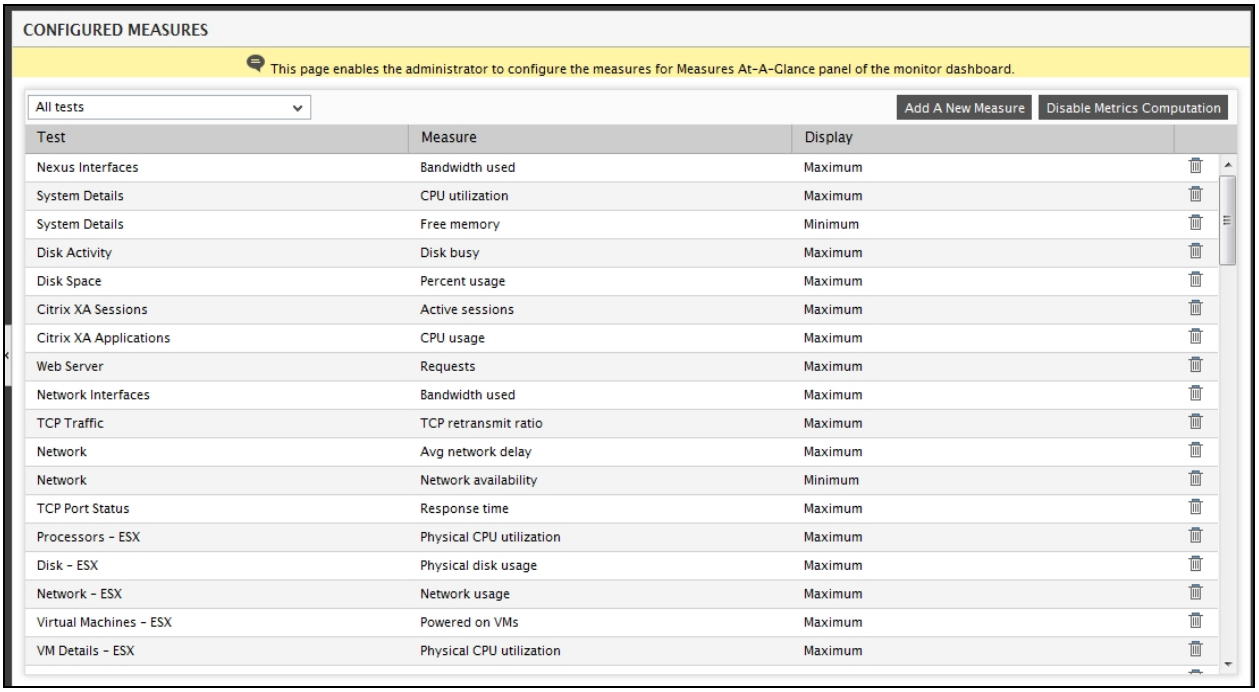

Figure 5.33: Viewing the existing measures

6. You can view a pre-configured list of measures and their corresponding min/max settings in the right panel of Figure 5.33. To delete any of the measures that pre-exist, click on the **Delete** button corresponding to that measure.

5.5.1.3 Configuring the User Experience Dashboard

The User Experience Dashboard provided by the eG Enterprise Suite helps end-users themselves to view the performance metrics related to their access to the Citrix/VDI/Terminal server infrastructure. Using this dashboard, end users can easily determine when they see a slowdown, is the problem being caused by connectivity to the Citrix/VDI infrastructure, by any application(s) that they are using within a Citrix session, or by the Citrix infrastructure itself. If a performance problem is in the interconnecting network or in one of the applications the user has launched, the user can initiate corrective action (e.g., kill the offending process, contact the local network team, etc.) to alleviate the issue. To configure the measures which need to be highlighted in the User Experience Dashboard, you can use the User Experience Dashboard node in the **MONITOR SETTINGS** tree-structure and its sub-nodes.

1. First, expand the **User Experience Dashboard** nodein the **MONITOR SETTINGS** tree in the left panel. Then, click the **Configured Measures** sub-node (see Figure 5.34).
2. The **CONFIGURED MEASURES** page will then appear in the right panel (see Figure 5.34).

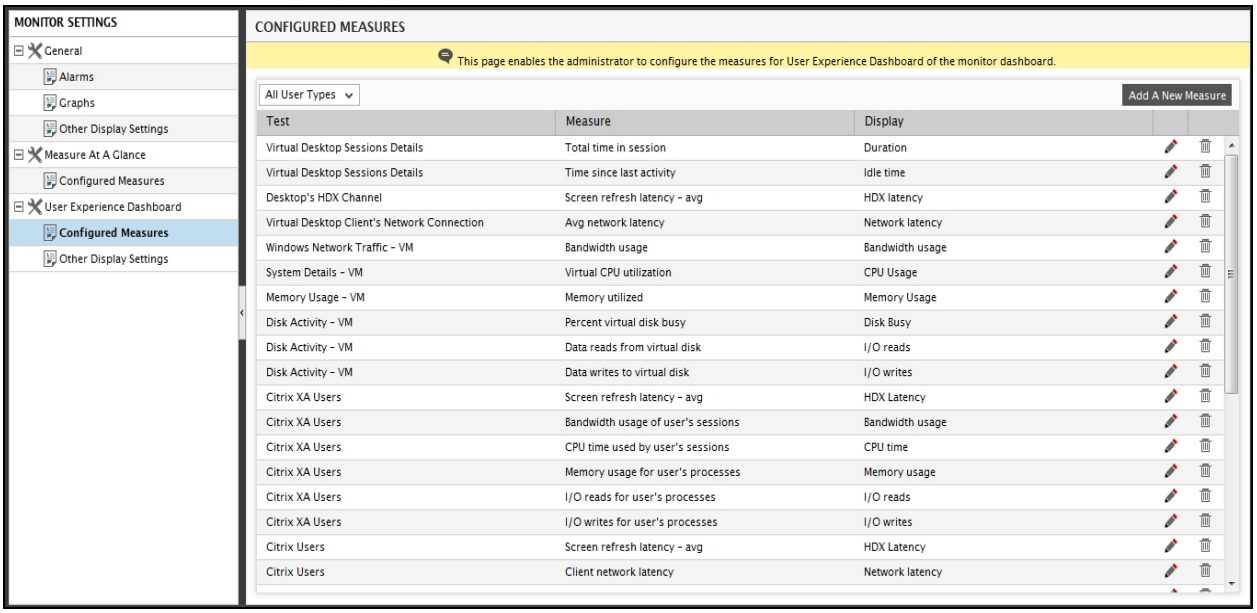

Figure 5.34: The measures configured for the User Experience Dashboard

3. You can configure the User Experience Dashboard with a different set of measures for different types of users (i.e., VDI/XenApp/Terminal) in your environment. By default, the **All User Types** option (See Figure 5.34) is chosen from the list box indicating that the **CONFIGURED MEASURES** page displays the complete list of metrics pre-configured for display in the User Experience Dashboard, irrespective of the types of users in your environment. If you wish to view the measures that are relevant only to your *VDI/XenAPP/XenApp7/Terminal* servers, then you can select the appropriate option from the list box.

4. On the other hand, if you to wish to view the performance of a measure that is not pre-defined in this section, click on the **Add New Measure** button in Figure 5.34. An **ADD A NEW MEASURE** pop up window will then appear (see Figure 5.35).

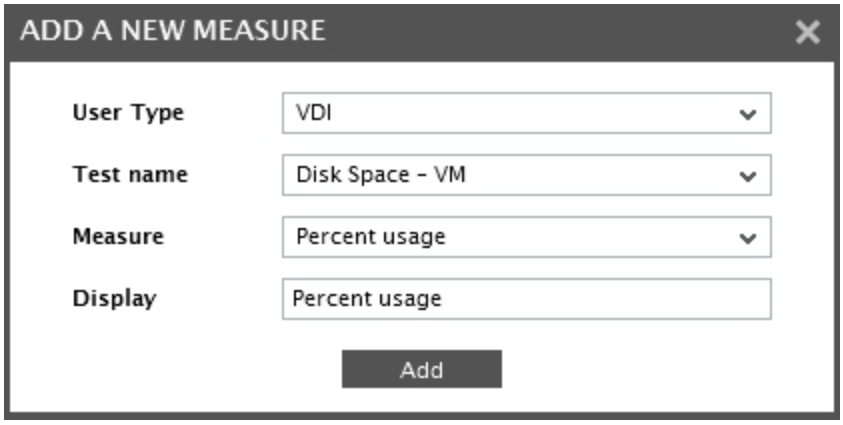

Figure 5.35: Adding a new measure for display in the User Experience dashboard

5. To add more measures to the **User Experience Dashboard** section, do the following:
   - First select a user type from the **User Type** list of Figure 5.35.
   - Then,select a test from the **Test name** list in Figure 5.35.
   - All the measures that are reported by the chosen test will then be populated in the **Measures** list.
   - From this list, select the measure that should appear in the **User Experience Dashboard.**
   - Then, specify the display name of the measure in the **Display** text box i.e., you can override the default name of the measure and provide the name of your choice for the chosen Measure.
   - To add the measure, click the **Add** button in Figure 5.35. Likewise, you can add multiple measures

of your choice for display in the User Experience Dashboard.

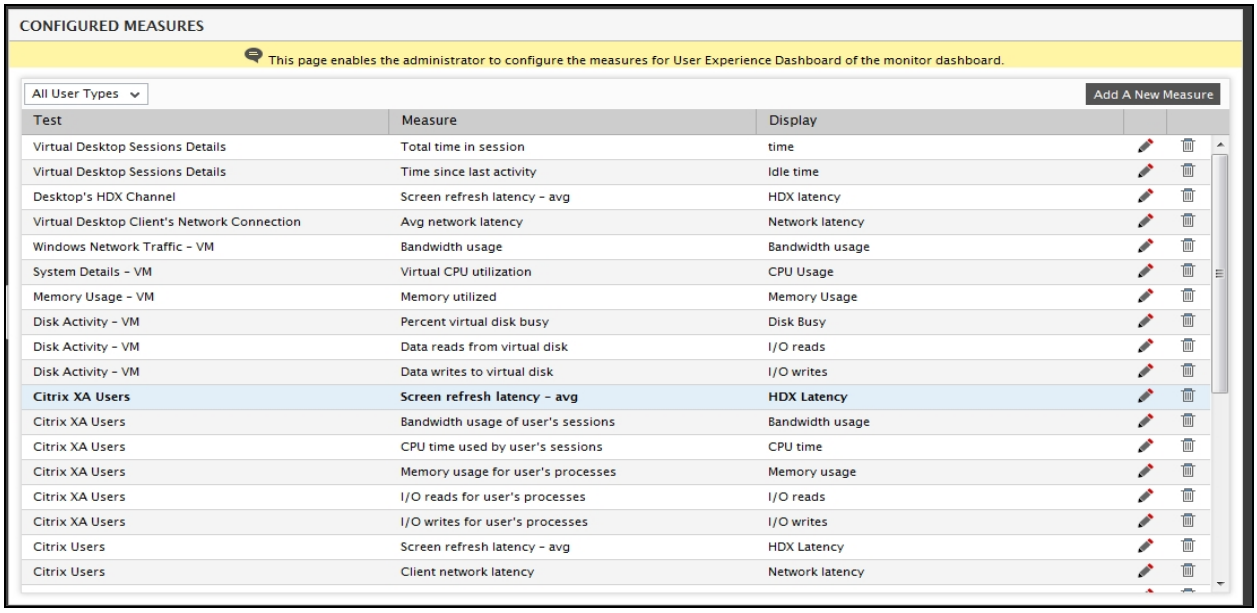

Figure 5.36: Viewing the existing measures configured for the User Experience dashboard

6. You can view a pre-configured list of measures and their corresponding tests in the right panel of Figure 5.36. To delete any of the measures that pre-exist, click on the **Delete** icon corresponding to that measure. To modify the display name of a pre-defined measure, you can click the **Modify** icon against that particular **Test/Measure** combination. Figure 5.37 will then appear.

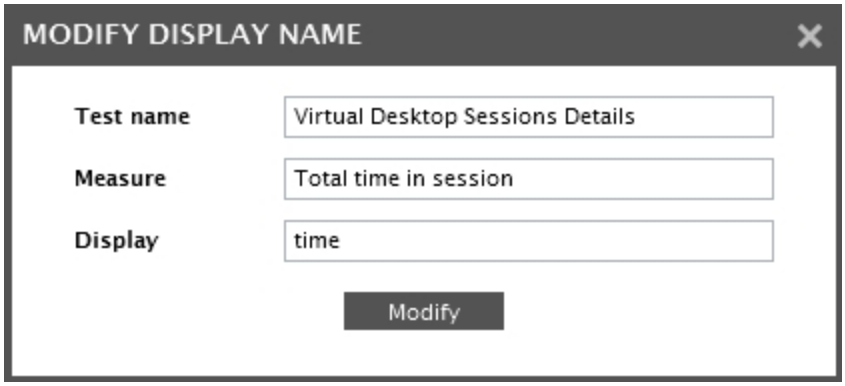

Figure 5.37: Modifying the display name of a measure configured for the User Experience dashboard

7. The **Test name** and the **Measure** text box will be auto populated when you click the **Modify** icon in Figure 10. By default, the default display name of the measure will be displayed in the **Display** text box. You can provide the display name of your choice in the **Display** text box and click the **Modify** button to save the changes.
8. Using the **Other Display Settings** sub-node under the **User Experience Dashboard** (see Figure 5.38), you can customize the display of the User Experience Dashboard in the eG monitor interface of your environment.

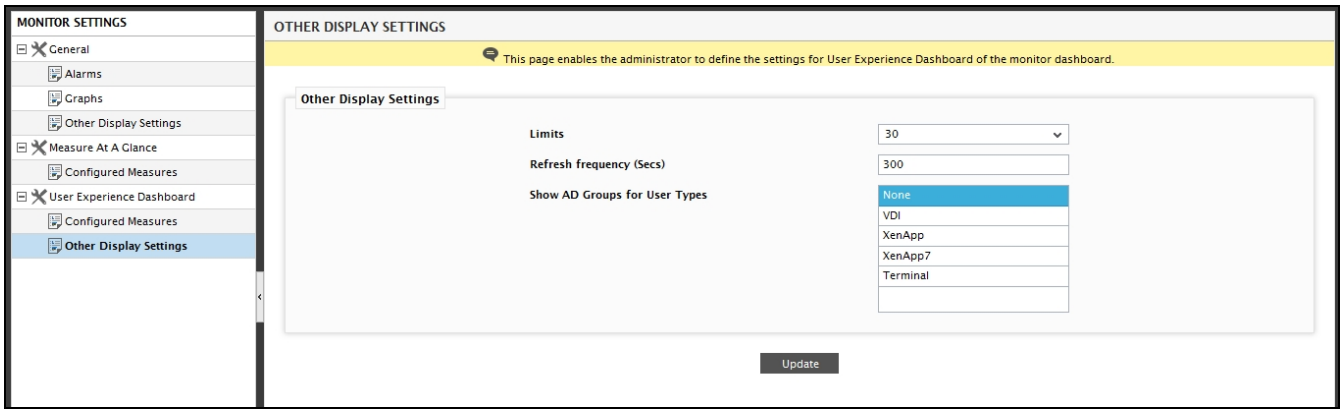

Figure 5.38: Configuring other display settings of the User Experience dashboard

9. By default, the **User Experience Overview** dashboard will list a maximum of *30* desktops/XenApp/Terminal servers in the monitored environment. If you wish to alter the maximum number of desktops/XenApp/Terminal servers that should be displayed in the overview dashboard, then you can set your own value using the **Limits** list.
10. Specify the frequency (in seconds) at which the User Experience Dashboard should auto-refresh in the **Refresh Frequency** text box. By default, this is *300* seconds.
11. By default, the User Experience dashboard can be generated for users of VDI/XenApp/Terminal server environments. If the users of the VDI/XenAPP/Terminal servers are part of an Active Director group, then eG Enterprise Suite can group such users/desktops in terms of the Active Directory Group in the User Experience Dashboard. To achieve this, a **Show AD Groups for User Types** list is provided (see Figure 5.38). If you select XenApp from this list, then the users logging into the XenApp server will be grouped based on the Active Directory groups and the User Experience dashboard will be displayed based on each Active Directory Group. By default, **None** will be selected from this list.
12. Clicking the **Update** button in Figure 5.38 will register your changes.

5.5.2 Configuring Manager Settings

To configure manager settings, select the **Manager** option from the **Settings** tile. Figure 5.39 will then appear with a **MANAGER SETTINGS** panel to its left. Clicking on any option in the **MANAGER SETTINGS** panel will invoke a list of parameters in the right panel that you can configure.

1. Clicking on the **General** option will open a **GENERAL** page in the right panel (see Figure 5.39).

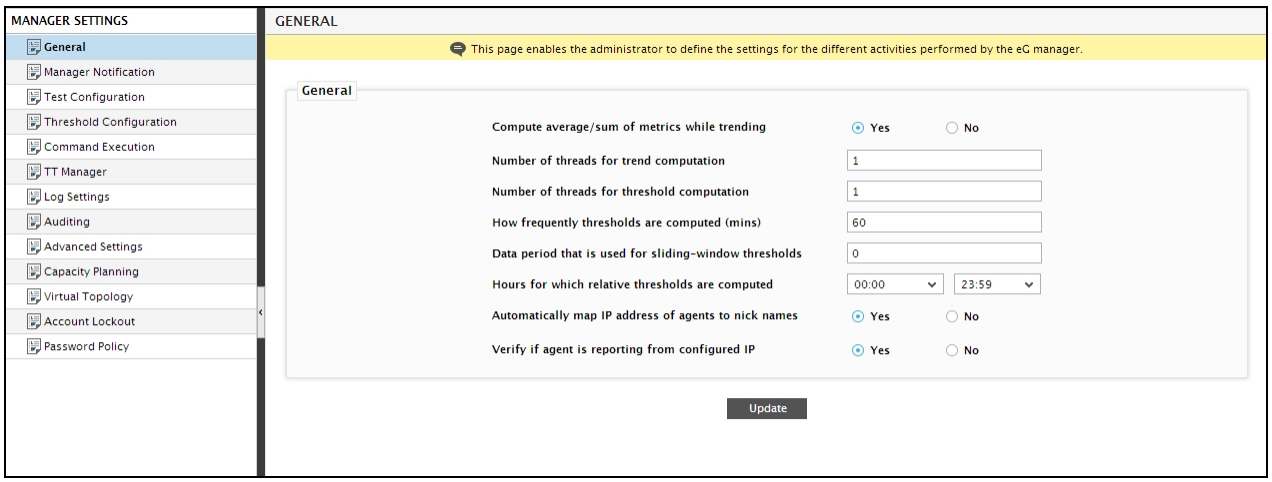

Figure 5.39: Configuring the General manager settings

2. Using the **GENERAL** page, you can configure the following settings:

   - The trend graph option in the eG monitoring console allows users to view and analyze the computed trend values over time. By default, during trending, the eG manager computes the upper and lower bounds of every metric using statistical quality control techniques. In addition, the manager also computes the average and sum of every metric while trending, and stores these values too in the database. Accordingly, the **Compute average/sum of metrics while trending** flag is set to **Yes** by default. This ensures that users have the option to plot min-max, average, or sum values of a chosen metric in a trend graph, and are able to better assess and plan the current and future performance of the target components. However, if, for some reason, you do not want the eG manager to perform average and sum computations while trending, you can indicate that by setting the **Compute average/sum of metrics while trending** flag to **No**.
   - The default value **1** in the **Number of threads for trend computation** text box indicates that by default, the eG manager uses a single thread for computing trend values. To speed up the day-end trending activity, you can configure the eG manager to spawn more threads, by increasing the **Number of threads for trend computation**.
   - The default value **1** in the **Number of threads for threshold computation** text box indicates that by default, the eG manager uses a single thread for computing thresholds. To speed up the threshold computation process, you can configure the eG manager to spawn more threads, by increasing the **Number of threads for threshold computation**.
   - Typically, the eG manager computes thresholds on an hourly basis. Accordingly, the default value in the **How frequently thresholds are computed** text box is 60 (minutes). However, in case of highly dynamic environments (eg., trading environments), where significant data changes occur at regular intervals, you might want to compute thresholds more frequently, say every half an hour, so as to capture the smallest of performance variations. On the contrary, in fairly static environments, where performance data is not likely to change frequently or dramatically, just a few threshold computation points would suffice - in other words, thresholds can be computed less frequently. To alter the

frequency of threshold computation to suit your environment, simply modify the **How frequently thresholds are computed** setting.

- Since the default threshold frequency is 1 hour (i.e., 60 minutes), by default, thresholds are computed for distinct one-hour data periods - eg., between 7 and 8 PM, between 1 and 2 PM, etc. However, in some environments, data transition might not always occur at such definite time windows. For instance, dramatic data changes could occur in an environment between 6.45 and 7.45 PM, and not during the precise 7 - 8 PM slot. To address the unique thresholding needs of such environments, you can off-set the data period for threshold computation by a fixed duration (in minutes) specified in the **Data period that is used for sliding-window thresholds** text box. For example, to ensure that the data collected during 6.45 PM and 7.45 PM is considered for threshold computation, specify 15 (minutes) in the **Data period that is used for sliding-window thresholds** text box.
- By default, relative thresholds are computed for the whole day. This is why the `From` and `To` specifications in the **Hours for which relative thresholds are computed** configuration indicate 24 hours by default. Sometimes, administrators might decide to compute thresholds for the working hours only - say between 9 AM and 6 PM - so as to avoid computing thresholds for periods of inactivity in the environment, and thus reducing the unnecessary stress on the eG backend. In such a case, you can modify the **From** time and the **To** time accordingly.
- When an agent attempts to download from the manager the details of managed components and tests to be executed, by default, the manager first determines the IP of the agent that is requesting for the information. Then, the manager verifies the IP to nickname mapping, identifies all the nick names that map to that IP address, and allows the agent to download those sections of the **eg_agents.ini** file that have either the agent's IP address or the nicknames that correspond to it. Accordingly, the **Automatically map IP address of agents to nick names** flag is set to **Yes** by default. In a managed service provider environment however, the same IP address could exist in two different customer networks. Similarly, if you have DHCP enabled environments, the IP address could change frequently. In such cases therefore, you can set this flag to **No**.
- By default, the manager checks the nickname to IP mapping of every agent that talks to it. Accordingly, the **Verify if agent is reporting from configured IP** parameter is set to **Yes** by default. However, in case of DHCP enabled environments, where the IP constantly changes, this has to be set to **No**, as the agent will otherwise fail.
- Finally, click the **Update** button,

3. Clicking on the **Manager Notification** option in the **MANAGER SETTINGS** panel will invoke a **MANAGER NOTIFICATION** page (see Figure 5.40) in the right panel.

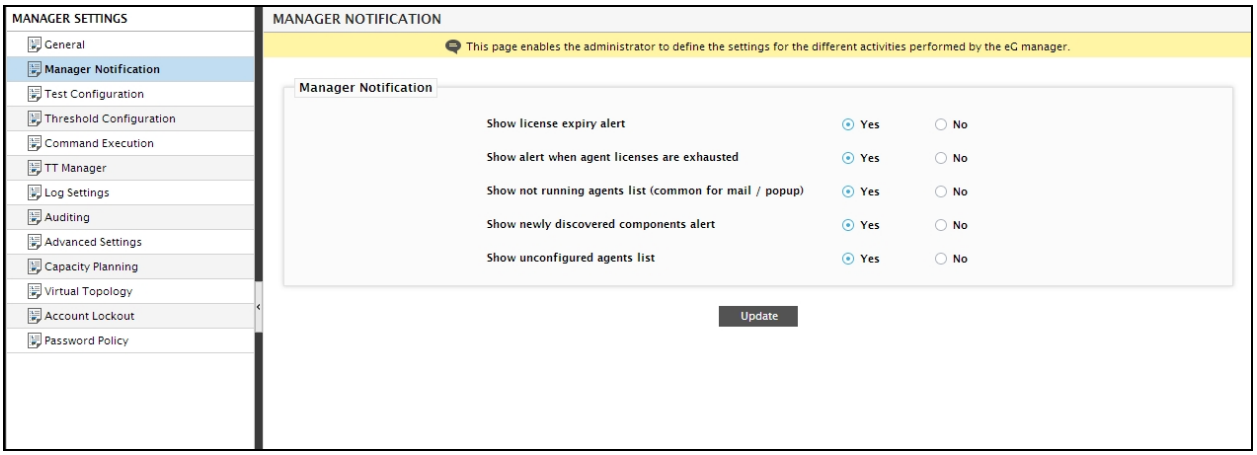

Figure 5.40: Configuring Manager Notification

4. Using the **MANAGER NOTIFICATION** page, you can configure the type of alerts you want displayed in the **MANAGER NOTIFICATION** window that pops out as soon as you login to the eG administrative interface or when you click the **Manager Notification** icon ( ) in the Admin tool bar available at the right, top corner of the Admin interface.

   - If you want the **MANAGER NOTIFICATION** window to alert you to the impending expiry of the eG license, then set the **Show license expiry alert** flag to **Yes**. If this is done, then, starting from 7 days before the license expiry, the **MANAGER NOTIFICATION** window will display a license expiry alert every day.
   - If you want the **MANAGER NOTIFICATION** window to alert you once your eG installation runs out of agent licenses, then set the **Show alert when agent licenses are exhausted** flag to **Yes**.
   - To ensure that the **MANAGER NOTIFICATION** window alerts you to agents that are not running, set the **Show not running agents list** flag to **Yes**. **Note that if this flag is turned on, you will also be alerted by email to agents that are not running.**
   - Whenever the eG manager discovers/re-discovers new components/devices, you can optionally configure the manager to display the list of the recently discovered components/devices in **MANAGER NOTIFICATION** window. To enable this capability, set the **Show newly discovered components alert** flag to **Yes**.
   - Many times you may install an eG agent on a host, but may not use that agent to monitor any component/device in your environment. To make sure that the **MANAGER NOTIFICATION** window informs you about the existence of such 'unused' agents, set the **Show unconfigured agents list** flag to **Yes**. Once this is done, then, the **MANAGER NOTIFICATION** window will display the count of the unconfigured (i.e., unused) agents. Clicking on the count will lead you to a page that lists the agents that are yet to be used for monitoring.
   - Finally, click the **Update** button.

5. If you click on the **Test Configuration** option in the **MANAGER SETTINGS** panel, a **TEST CONFIGURATION** page will appear in the right panel (see Figure 5.41).

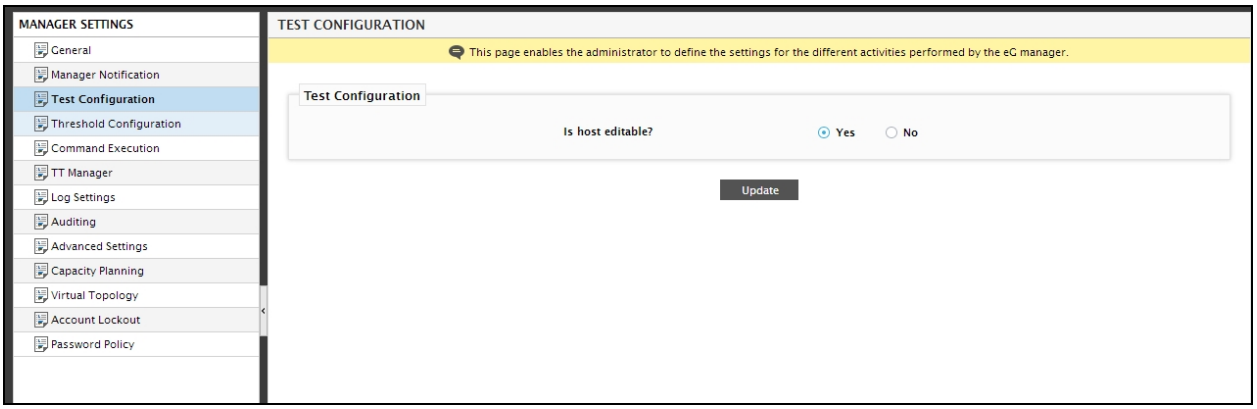

Figure 5.41: Configuring the test configuration settings

6. Using the **TEST CONFIGURATION** page, you can define the following:

   - By default, the **HOST** parameter of a test cannot be modified during test configuration. Accordingly, the **Is host editable?** flag is set to **No** by default. If you want to change the value of the **HOST** parameterat the time of test configuration, then set this flag to **Yes**.
   - Finally, click the **Update** button.

   Note:

   If the **Is host editable?** flag is set to **Yes**, then the **HOST** parameter for all tests – both internal and external – will become editable. However, this capability, when enabled, is most useful when configuring external tests, as it allows administrators the flexibility to run the test on a remote host and collect metrics.

7. Clicking on the **Threshold Configuration** option in 5.5.2 will invoke a **THRESHOLD CONFIGURATION** page in the right panel (see Figure 5.41).

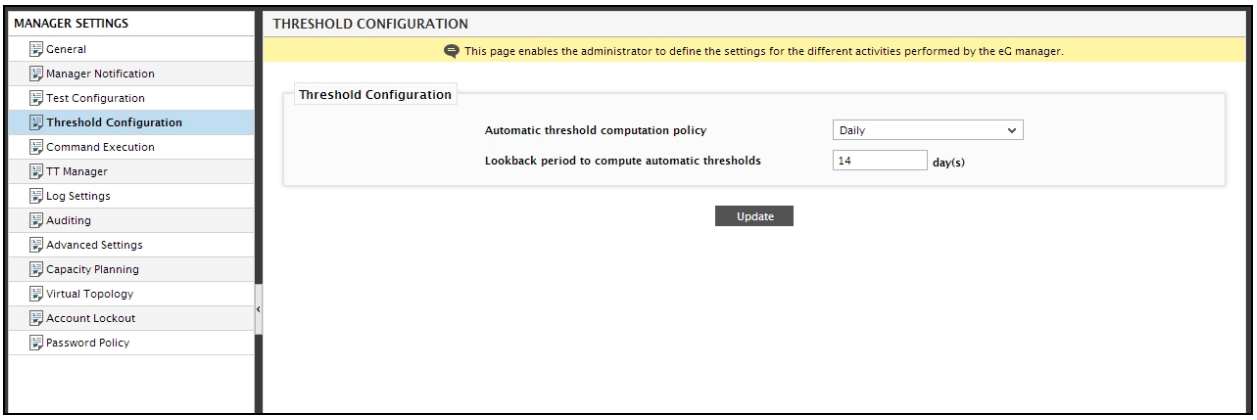

Figure 5.42: Configuring the Threshold configuration settings

8. eG Enterprise is capable of automatically computing the thresholds of performance using historical data. This auto-thresholding capability eliminates the need for determining the norms of performance manually in dynamic environments where measure values vary with time. eG Enterprise uses tried and tested statistical quality control techniques to analyze past values of the metrics, and automatically sets the

upper and lower bounds for each of the metrics based on this analysis. With the help of the **THRESHOLD CONFIGURATION** page, you can configure the 'past values' to be considered for automatic threshold computations.

- By default, to perform the threshold computation, the values reported by a measure during each day of the last 14 days is used. This is why, the **Automatic threshold computation policy** is set to **Daily**, and the **Lookback period to compute automatic thresholds** is set to **14**, by default. According to these default settings, the threshold for a measure for 8 AM to 9 AM on August 18, will be automatically computed using the actual values reported by that measure during 8 AM to 9 AM on each of the 14 days prior to August 18.

  In some environments, mandatory operations - for example, data backup operations or virus scanning operations - may occur on one/more servers once a week. This in turn may increase the activity levels on those servers during that time window. In some other environments, such routine operations may be scheduled to take place at a specific time every month. If threshold computations in such environments are based on **Daily** data collections, then the resulting thresholds may not reflect the weekly/monthly deviations in usage, thus causing false alerts. In such a case, you can set the **Automatic threshold computation policy** to **Weekly** or **Monthly** (as the case may be), and specify the number of past weeks/months to be considered for computing thresholds in the **Lookback period to compute automatic thresholds**. For instance, say that the **Automatic threshold computation policy** is set to **Weekly** and the **Lookback period** is configured as **2**. In such a case, to compute thresholds for 8 AM to 9 AM on August 18, the actual measure values reported on the same day and time in the last **2** weeks will be considered - this will be 8 AM - 9 AM on August 11 and 8 AM - 9 AM on August 4. Similarly, assume that the computation policy is set to **Monthly** and the **Lookback period** is set to **2**. In such a case, to compute thresholds for 8 AM to 9 AM on August 18, the actual measure values reported on the same date and time but in the last **2** months will be considered - this will be 8 AM - 9 AM on July 18 and 8 AM - 9 AM on June 18.

  Note:

  If the **Monthly** computation policy is chosen and the date for which thresholds are to be computed is not available in the previous month - say, thresholds are to be computed for July 31, but the previous month of June has only 30 days - then, the data collected during the last day of the previous month will be used for threshold computation. In the case of our example therefore, the data collected on June 30 will be used.

- Finally, click the **Update** button.

9. Click on the **Command Execution** option in the **MANAGER SETTINGS** panel to view the **COMMAND EXECUTION** section in the right panel (see Figure 5.43).

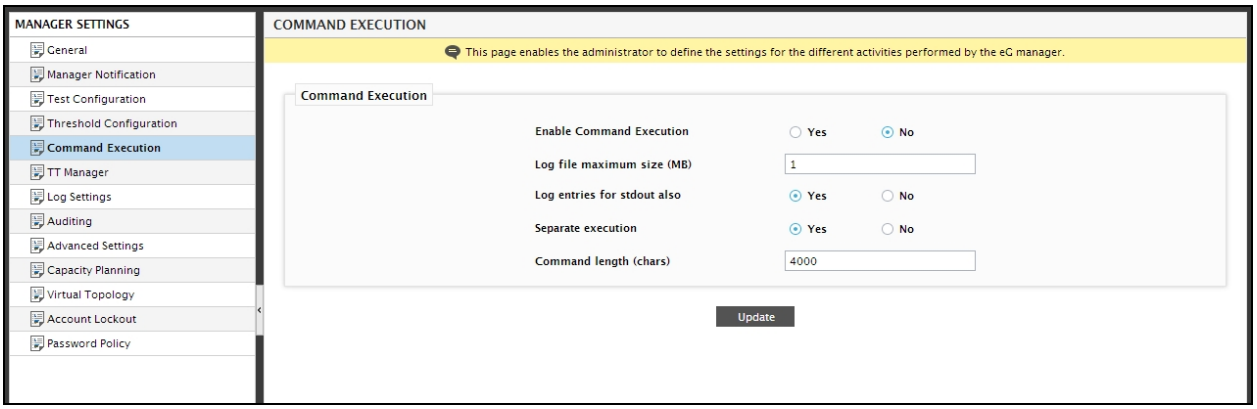

Figure 5.43: Configuring command execution on alert generation

10. Like email IDs / mobile numbers, you can associate one/more custom scripts with users to the eG Enterprise system. Whenever alarms are raised/modified/closed for a specific user, the associated custom script will automatically execute, so that the details of the alarms are routed to third-party customer relationship management systems or TT systems, and trouble tickets automatically created (or closed, as the case may be) for the corresponding user. The custom scripts thus provide a mechanism by means of which eG alerts are integrated into CRM/TT systems. These custom scripts can be configured in addition to or instead of email / SMS alerts. The eG manager will invoke this custom script with a fixed set of parameters, namely - *ComponentType*, *ComponentName*, *LayerName*, *Desc*, *StartTime*, *Priority*

    Note:

    This capability is also supported in a redundant eG manager configuration. In case of the redundant manager configuration, if the primary manager is up and running, it will perform script execution. If the primary manager is down for any reason, the secondary manager will perform script execution.

11. Using the **COMMAND EXECUTION** page, you can configure the details of these scripts as follows:

    - First, to enable the command execution capability, set the **Enable Command Execution** flag to **Yes**. By default, this is set to **No**.
    - To track the status of the script execution and to troubleshoot issues with the same, use the **mailexec.log** file that is automatically created in the **<EG_INSTALL_DIR>\manager\logs** directory when the custom scripts execute. You can specify the maximum size up to which this log file can grow, in the **Log file maximum size** text box. When the file reaches the specified size limit, the details originally logged in the **mailexec.log** file will be moved to another log file named **mailexec.log.1**, and the newer information will be logged in the **mailexec.log** file instead. This log rotation mechanism helps ensure that the log file does not grow beyond control. The default maximum size is 1 MB.
    - You can even indicate the type of information you want logged in the **mailexec.log** file. By default, the log files capture both the errors and the standard output of the specified **Command**; accordingly, the **Log entries for stdout also** flag is set to **Yes** by default. If you want to capture only the errors, set the **Log entries for stdout also** flag to **No**.
    - To execute the script for every alert that is generated, set the **Separate execution** flag to **Yes**. This means that when the script executes, the details of only a single alert will be included in the script

output. Given below is an extract from the **mailexec.log** file, when the **Separate execution** flag is set to **Yes**.

```
31/08/2011 11:06:51 USER admin COMMAND echo PARAMS -ComponentType "Host system" -ComponentName "esx150" -LayerName "Operating System" -Desc "-|Processors - Esx|Usage|Physical CPU usage of ESX server's processor is high|Processor 3|192.168.8.67" -StartTime "Aug 31, 2011 11:05:18" -Priority "Critical"

31/08/2011 11:06:51 INFO -ComponentType "Host system" -ComponentName "esx150" -LayerName "Operating System" -Desc "-|Processors - Esx|Usage|Physical CPU usage of ESX server's processor is high|Processor 3|192.168.8.67" -StartTime "Aug 31, 2011 11:05:18" -Priority "Critical"
```

As you can see, for a single alert, two lines have been logged in the **mailexec.log** file.

The first line is the script invocation. It displays the user who will be receiving the alarm intimation, the syntax of the command that is being executed by the script, and the params - i.e., the input parameters/arguments - that the command takes while executing. In the sample provided, *echo* is the command executed by the script. Therefore, the params tag in our extract is followed by the input parameters required by the *echo* command. As you can see, every parameter of the *echo* command consists of two components: the parameter name and its value at runtime. While the param name begins with a - (hyphen), its runtime value is enclosed within "double quotes". Take for instance the parameter, *-ComponentType "Host system"*. Here, *-ComponentType* is the parameter, and the *"Host system"* is its value.

The second line logs the output of the command - in the case of our example, this will be the output of the *echo* command. The command output typically begins with the tag info, and will be followed by the details of the alert being sent. Like its input, the output of the *echo* command too is a combination of the parameter name and its value at runtime. While the parameter name indicates what type of information is being sent, the actual information itself is contained within "double quotes" and forms the parameter value. The parameters included in the output of the *echo* command have been discussed in the table below:

Parameter	Description
ComponentType	The problem component type
ComponentName	The name of the problem component
LayerName	The name of the problem layer
Desc	A brief description of the problem; typically, for the *echo* command, a problem description includes the following: • the site affected; if the alarm does not pertain to any site, only a '-' will appear in the output, as is the case in our sample output;

Parameter	Description
	• the problem test - this is *Processors - Esx* in our example; • the problem measure - this is *Usage* in our example; • the alarm description - in the case of our sample, this is: *Physical CPU usage of ESX server's proce*ssor *is high*; • the problem descriptor (if any); in the case of our example, this is *Processor 3*. For non-descriptor-based tests, a '-' will appear here. • the measurement host - this is *192.168.8.67* in the case of our example • the last measurement value - this is not available in the case of our sample. This value will appear in the description only if the ***Show last measure value in alerts*** flag in the **MAIL ALERT PREFERENCES** page (Alerts -> Mail Settings -> Alerts) is set to ***Yes***.
StartTime	The problem date/time
Priority	The problem severity

If the **Separate execution** flag is set to **No** on the other hand, the script will be executed only once for all alerts raised simultaneously. Given below is an extract from the **mailexec.log** file, when the **Separate execution** flag is set to **No**.

```
30/08/2011 12:13:27 USER admin COMMAND echo PARAMS -ComponentType "Microsoft SQL"
- ComponentName "sql100:1433" - LayerName "MS SQL Service" - Desc "-
|MsSqlNet|Availability|SQL Server unavailable|master|192.168.8.77" -StartTime "Aug
30, 2011 12:10:26" -Priority "Critical" # -ComponentType "Windows" -ComponentName
"win77" - LayerName "Windows Service" - Desc "-
|WindowsServices|Availability|Service not up|eGAgentMon|win77" -StartTime "Aug 30,
2011 12:10:41" -Priority "Critical" # -ComponentType "Host system" -ComponentName
"esx150" -LayerName "Operating System" -Desc "-|Processors - Esx|Usage|Physical
CPU usage of ESX server's processor is high|Processor 3|192.168.8.67" -StartTime
"Aug 30, 2011 12:11:18" -Priority "Critical" # -ComponentType "Host system" -
ComponentName "vdi136" - LayerName "Network" - Desc "- |Network -
Esx|Availability|Network interface of the ESX server is down|vmnic1|192.168.8.67"
- StartTime "Aug 30, 2011 12:11:23" - Priority "Critical"
30/08/2011 12:13:27 INFO - ComponentType "Microsoft SQL" - ComponentName
"sql100:1433" - LayerName "MS SQL Service" - Desc "- |MsSqlNet|Availability|SQL
Server unavailable|master|192.168.8.77" - StartTime "Aug 30, 2011 12:10:26" -
Priority "Critical" # -ComponentType "Windows" -ComponentName "win77" -LayerName
"Windows Service" - Desc "- |WindowsServices|Availability|Service not
up|eGAgentMon|win77" -StartTime "Aug 30, 2011 12:10:41" -Priority "Critical" # -
ComponentType "Host system" -ComponentName "esx150" -LayerName "Operating System"
-Desc "-|Processors - Esx|Usage|Physical CPU usage of ESX server's processor is
high|Processor 3|192.168.8.67" - StartTime "Aug 30, 2011 12:11:18" - Priority
"Critical" # - ComponentType "Host system" - ComponentName "vdi136" - LayerName
"Network" -Desc "-|Network - Esx|Availability|Network interface of the ESX server
is down|vmnic1|192.168.8.67" - StartTime "Aug 30, 2011 12:11:23" - Priority
"Critical"
```

In this case again, only two lines will be logged in the **mailexec.log**. The first line will display the user name, command syntax, and the command params (with their corresponding values). Since multiple alerts will be clubbed in the command output, the first line will include the params for all alerts. '#' (hash) will be used to separate the parameter-value pairs of one alert from the other.

The second line, which begins with the tag info, will display the combined output of all email alerts generated simultaneously - each alert in the output will also be separated by the hash (#) symbol.

Note:

The **Separate execution** flag setting will take effect only if a user is configured to receive **New** alarm intimations alone - i.e., if the **Type of notification** is set to **New** for a user in the **ADD USER** page. For users who are configured to receive the **Complete List** of alarms, details of multiple alarms will always be clubbed in a single script execution, regardless of this flag setting.

- Specify the maximum permissible length of the command in the **Command length** text box. The command line here includes the command syntax, its input arguments (if any), and the value of each argument. By default, the command line can have a maximum of 4000 characters. You can alter this default setting by specifying a length of your choice in the **Command length** text box. If the **Separate**

execution flag is set to **Yes**, then, if the **Command length** is violated, the command will be truncated at the end of the parameter value that is closest to the configured **Command length**. If the **Separate execution** flag is set to **No**, then, upon a **Command length** violation, the command will be truncated at the end of the complete alert specification that is closest to the configured length.

- Finally, click the **Update** button.

12. The eG Manager supports a command line interface, called the **TT MANAGER CLI**, that can be configured to automatically execute TT (Trouble Ticketing) system-specific commands as and when alarms are added, modified, or deleted in eG Enterprise. This interface offers a way of communication between the eG Manager and a TT system. Use the options provided by the **TT MANAGER** page to control the behavior of this command line interface. To access this page, click on the **TT Manager CLI** option in the **MANAGER SETTINGS** panel (see Figure 5.44)

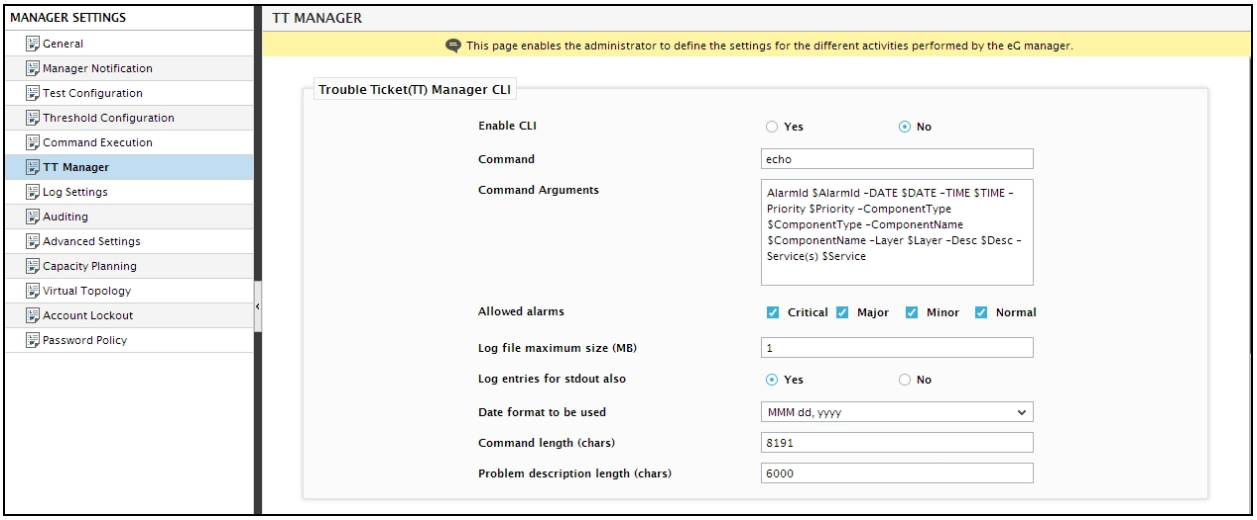

Figure 5.44: Configuring the TT Manager CLI

Note:

This section will appear only if the **Trouble Ticket Integration** capability is enabled by the eG license.

13. Using the **TT MANAGER** page, you can configure the following:

- Set the **Enable CLI** parameter to **Yes** to enable this capability.
- In the **Command** text box, **echo** is displayed by default, indicating that the eG manager will execute an *echo* command by default to communicate with the TT system.
- The **Command Arguments** text box displays the default input parameters that the **echo** command takes during execution. These default parameters are as follows:

  AlarmId $AlarmId -DATE $DATE -TIME $TIME -Priority $Priority -ComponentType $ComponentType -ComponentName $ComponentName -Layer $Layer -Desc $Desc -Service $Service

  As you can see, each parameter is represented by a qualifier and a variable name. While the qualifier is typically prefixed by a hyphen (-), the variable name is prefixed by a $ symbol. These variables will be substituted by actual values during runtime. Using the qualifiers, you will be able to tell what value

follows. For instance, at runtime, the parameter –Priority $Priority could appear as *–Priority Critical*. This implies that the **Priority** of the problem is **Critical**.

- By selecting the required check boxes against **Allowed alarms**, you can indicate the alarm priorities for which the eG manager needs to execute the specified command.
- To track the status of the command execution and to troubleshoot issues with the same, use the **ttexec.log** file that is automatically created in the **<EG_INSTALL_DIR>\manager\logs** directory. You can specify the maximum size up to which this log file can grow, in the **Log file maximum size** text box. When the file reaches the specified size limit, the details originally logged in the **ttexec.log** file will be moved to another log file named **ttexec.log.1**, and the newer information will be logged in the **ttexec.log** file instead. This log rotation mechanism helps ensure that the log file does not grow beyond control.
- You can even indicate the type of information you want logged in the **ttexec.log** file. By default, the log files capture both the errors and the standard output of the specified **Command**; accordingly, the **Log entries for stdout also** flag is set to **Yes** by default. If you want to capture only the errors, set the **Log entries for stdout also** flag to **No**.
- From the **Date format to be used** list box, select the format in which the date/time of the problem should be reported in the command output.
- Specify the maximum permissible length of the command in the **Command length** text box. By default, the command line can have a maximum of 8191 characters. You can alter this default setting by specifying a length of your choice in the **Command length** text box. If the actual command length exceeds the specified limit, then the output will not return the list of affected services and the detailed diagnosis information; instead, an empty string will appear next to the *–Services* qualifier and the *-DD* qualifier. If the command length continues to exceed the specified limit even after truncating the services list and the DD, the command execution will return an error.
- Specify the length of the problem description in the **Problem description length** text box. If the actual problem description exceeds the specified length, the characters that fall beyond the specified limit will be truncated.
- Finally, click the **Update** button.

14. Clicking on the **Log Settings** option in the left panel of Figure 5.44 will open the **LOG SETTINGS** page in the right panel (see Figure 5.45).

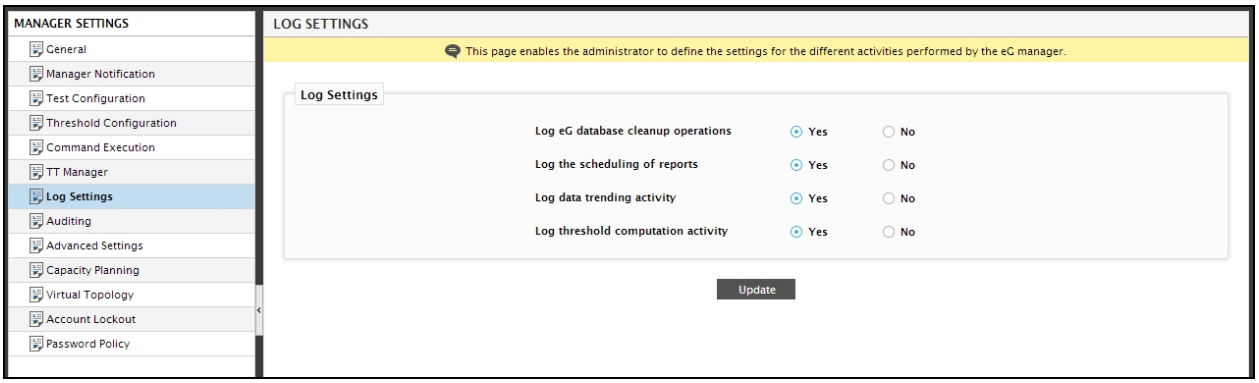

Figure 5.45: Configuring the Log settings

15. In the **LOG SETTINGS** section of this page, the following settings can be defined:

- The eG manager runs database cleanup operations at pre-configured frequencies. By default, the eG manager logs the details of a cleanup in the **cleanup_log** file in the **<EG_INSTALL_DIR>\manager\logs** directory. The details such as when the cleanup started, when it ended, the duration of the cleanup, what was cleaned up, errors during cleanup, etc., are logged, so that administrators are enabled to efficiently troubleshoot issues (if any) during cleanup. If, for some reason, you want to disable this logging activity, set the **Log eG database cleanup operations** to **No**.
- Once key reports are scheduled to be emailed to specific recipients, the eG manager, by default, creates a **schedule_ log** file in the **<EG_ INSTALL_ DIR> \manager\logs** directory, to which the success/failure of the report scheduling activity is logged. If you do not want to maintain this log file, then you can disable logging of report scheduling activities, by setting the **Log the scheduling of reports** flag to **No**.
- The eG manager, by default, logs the details of the day-end trend activity, so that administrators know when, how, and for how long trending occurred; this information enables administrators to troubleshoot issues with trending. A **trend_log** file is created in the **<EG_INSTALL_DIR>\manager\logs** directory for this purpose. However, if need be, trend logging can be disabled, by setting the **Log data trending activity** flag to **No**.
- By default, the eG Enterprise system performs threshold logging - i.e., the ability to record threshold activities in a log file. A **thresh_log** file is created in the **<EG_INSTALL_DIR>\manager\logs** directory, to which all threshold-related processes are logged. However, if need be, threshold logging can be disabled, by setting the **Log threshold computation activity** flag to **No**.
- Finally, click the **Update** button.

16. To enable the **Audit log** capability of eG Enterprise, click on the **Auditing** option in the **MANAGER SETTINGS** panel. An audit log can be best described as a simple log of changes, typically used for tracking temporal information. The eG manager can now be configured to create and maintain audit logs in the eG database, so that all key configuration changes to the eG Enterprise system, which have been effected via the eG user interface, are tracked. The eG audit logs reveal critical change details such as what has changed, who did the change, and when the change occurred, so that administrators are able to quickly and accurately identify unauthorized accesses/modifications to the eG Enterprise system.
17. Clicking on the **Auditing** option reveals an **AUDITING** page in the right panel (see Figure 5.46).

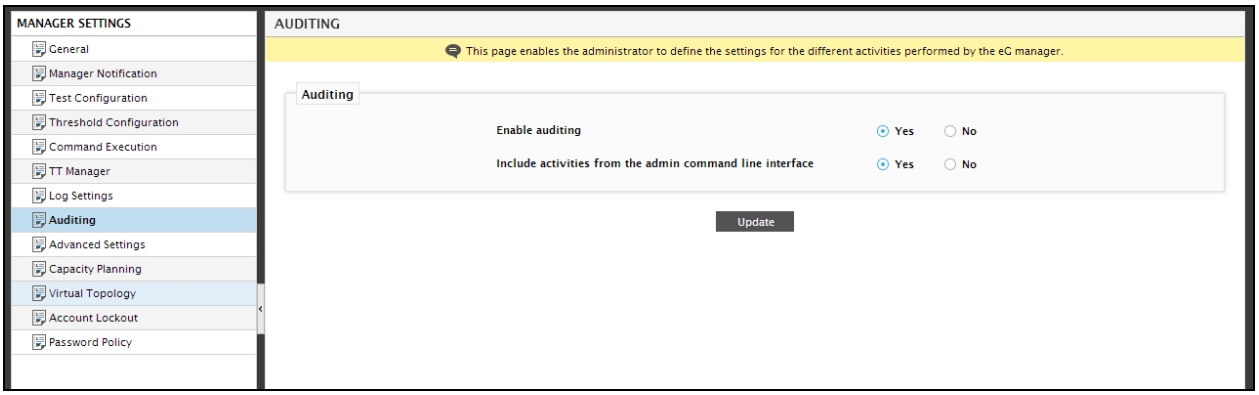

Figure 5.46: Enabling the auditlogging capability

18. Define the following in the **AUDITING** page:

- By default, the eG manager does not have audit logging capabilities. To enable the eG manager to perform audit logging, set the **Enable auditlog** flag to **Yes**. By default, this flag is set to **No**.
- Setting the **Enable auditlog** flag to **Yes** will allow you to include/exclude admin CLI (command lineinterface) activities in the eG manager's audit logging purview. The eG Enterprise Suite provides a command-line interface (CLI) which allows any automation tool that pre-exists in the target environment or a script to communicate with the eG manager and execute simple commands on the manager to perform critical configuration tasks. This integration minimizes user intervention in the configuration of the monitoring system. By default, the eG manager also performs audit logging for configuration activities executed via the command line. Accordingly, the **Include activities from the admin command line interface** flag is set to **Yes** by default. To ensure that audit logging is not performed for admin CLI activitie, set this flag to **No**.
- Finally, click the **Update** button.

19. For configuring advanced manager settings, click on the **Advanced Settings** option in **MANAGER SETTINGS** panel. This will result in the display of the **ADVANCED SETTINGS** page (see Figure 5.47).

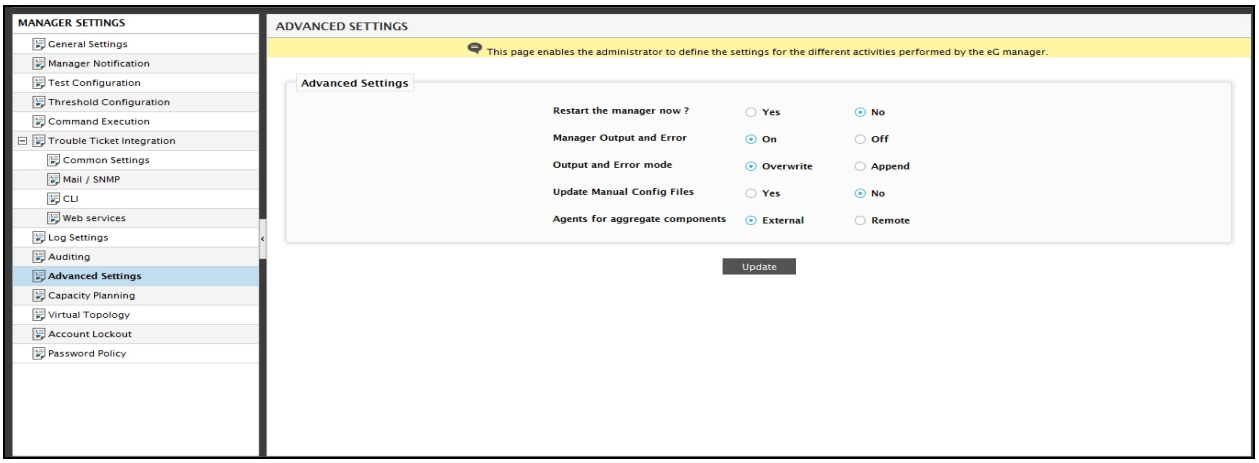

Figure 5.47: Configuring advanced manager settings

20. You can configure the following in the **ADVANCED SETTINGS** page:

- You can now restart the eG manager from the eG administrative console itself. To restart the eG manager immediately, set the **Restart the manager now?** flag to **Yes**. By default, it is set to **No**.
- To enable logging of the errors and output related to the eG manager's activities, set the **Manager Output and Error** flag to **Yes**. This automatically creates the **managererr.log** and the **managerout.log** files in the **<EG_INSTALL_DIR>\manager\logs** directory. While the errors are logged in the **managererr.log** , the output of manager operations is logged in the **managerout.log**. Using these log files, you can quickly and easily track the status of the eG manager's activities, capture the errors that occur, and analyze the reasons for these errors/failures.
- Sometimes, after logging manager-related errors and output for a while, you may switch off logging for a brief period, and then switch it back on. Likewise, for some reason, you may decide to restart an eG manager, which is actively logging manager errors and output. In both these cases, new error and output messages will emerge, waiting to be written to the log files that pre-exist. By setting the **Output and Error mode** flag to **Overwrite**, you can ensure that the new messages completely replace the

existing contents of the log files. Setting this flag to **Append** on the other hand, ensures that the new messages are only appended to the old contents.

- Typically, whenever a user manually changes a configuration file (.ini file) of the eG manager, he/she would have to restart the eG manager for the changes to take effect. To make sure that changes to configuration files are effected **even without a manager restart**, then, soon after you make modifications to one/more configuration files, set the **Update Manual Config Files** flag in Figure 5.47 to **Yes**, and click the **Update** button. This will automatically update the eG manager with the configuration changes. Once the updation is complete, the status of the **Update Manual Config Files** will automatically switch to **No**.

  Note:

  The **Update Manual Config Files** flag cannot be used to update the eG manager with changes to language files. Gor changes to language files to take effect, you will always have to restart the eG manager.

- The Metric Aggregation capability of the eG Enterprise Suite allows administrators to group one or more components of a particular type and monitor that group as a single logical entity termed as an aggregate component. This aggregate component provides a collective view of the performance of all the components that are grouped together. Aggregates can be used to represent all the components of a specific type, or all components of a type in a geography or in a business unit.

  Different organizations may want to view performance in different ways - some by type, some by geographies, and some by business unit. To provide administrators with maximum flexibility, and help them to configure any number of aggregate components without being constrained by the licensing of the eG monitors, eG Enterprise performs metric aggregation using external agents. By default, the **Agents for aggregate components flag** is set to *External* indicating that the default external agent bundled with the eG Enterprise suite is capable of monitoring an aggregate component. This helps the administrators to add any number of aggregate components depending on the views of their infrastructure they desire without needing additional licenses. In case, if you wish to monitor the aggregate components using a remote agent, then set this flag to *Remote*.

  Note:

  If you are already monitoring aggregate components using an external agent and now, if you wish to monitor the same aggregate components using a remote agent, then, you are required to purchase the premium licenses accordingly. If enough licenses are not available, then a message as shown in Figure 5.48 will be displayed and the external agent will continue to monitor the aggregate components.

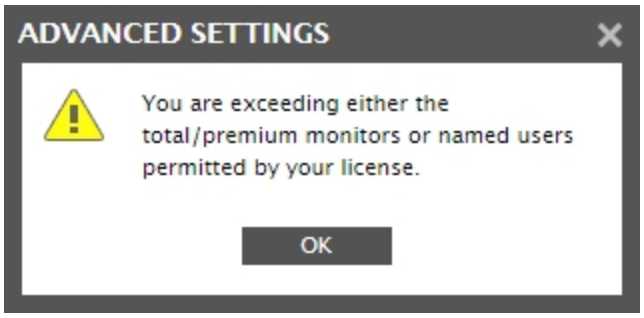

Figure 5.48: A message displaying that there are not enough premium licenses available

- Finally, click the **Update** button.

21. To enable administrators to easily and effectively study the historical trends in performance and accurately assess future capacity requirements, the eG Reporter offers a dedicated **Capacity Planning Reports** category. By default, eG Reporter does not allow users to generate **Capacity Planning** reports. Clicking on the **Capacity Planning** option in the **MANAGER SETTINGS** panel will invoke the **CAPACITY PLANNING** page in the right panel (see Figure 5.49).

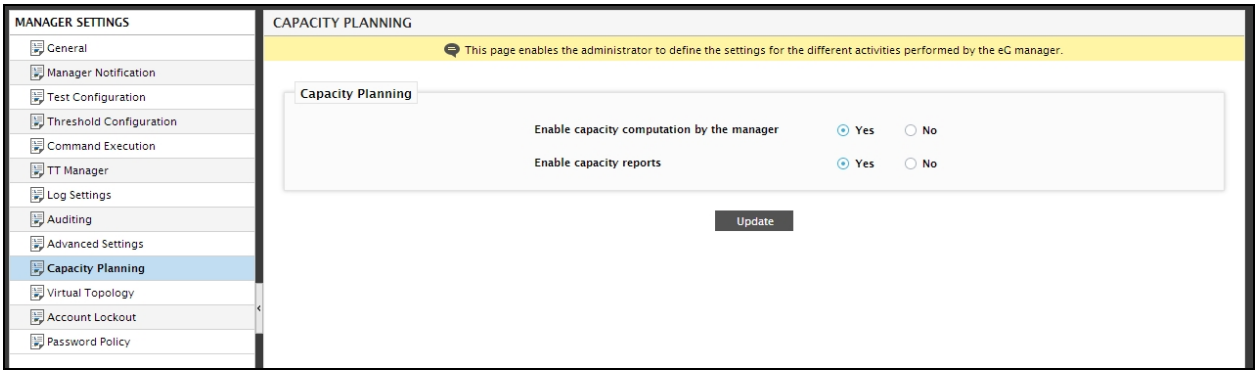

Figure 5.49: Configuring capacity planning

22. To enable the generation of this category of reports, do the following in the **CAPACITY PLANNING** page:

- Reports generated using trend data can accurately reveal the past load and performance trends, using which future trends in load and usage can be determined, potential sizing inadequacies predicted, and the future capacity prudently planned. This trend data is typically computed by applying the *Avg, Sum, Max, Min*, and *Percentile* functions on the measure data. To enable the eG manager to compute the trend data by applying the aforesaid functions, set the **Enable capacity computation by the eG manager** flag to **Yes**. By default, this is set to **No**.

- To enable capacity report generation, set the **Enable capacity reports** flag to **Yes**. If you do not want to generate any capacity planning reports, then, set this flag to **No**. If this flag is set to **Yes** and the **Enable capacity computation by the eG manager** flag to **No**, then the eG manager will not compute the *Avg*, *Sum, Max, Min*, and *Percentile* values of the measures, but you can still attempt to generate capacity planning reports. In this case, the **Custom** and **System** capacity planning reports, which rely solely on trend data, will not display any results. Also, the **Cumulation**, **Correlation**, and **Prediction** reports can be generated using raw measure data only and not the computed trend data. On the other hand, if both the flags are

set to **Yes**, then all reports can be generated.

- Finally, click the **Update** button.

23. To configure the automatic generation of a virtual topology, click the **Virtual Topology** option in the **MANAGER SETTINGS** panel of Figure 5.49. The right panel will then change to reveal the **VIRTUAL TOPOLOGY** page (see Figure 5.50).

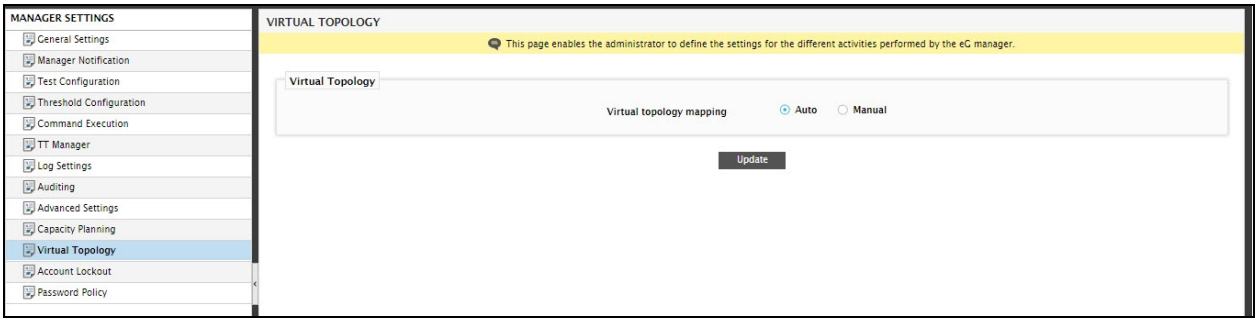

Figure 5.50: Configuring Virtual Topology settings

24. eG Enterprise is capable of automatically determining the mapping of applications to virtual hosts. This is why, the **Virtual topology mapping** flag in Figure 5.50 is set to **Auto** by default. This default setting ensures that eG Enterprise intelligently discovers which managed applications are actually executing on the VMs of a virtual host and automatically maps these applications to that virtual host in the **Virtual Topology** preview in the eG monitoring console. This mapping of applications to virtual hosts is important for root-cause diagnosis - for example, a problem with the virtual host (e.g., excessive disk slowdowns) can impact the performance of all the applications running on that host's virtual machines.

    If you set the **Virtual topology mapping** flag to **Manual**, then, when adding applications for monitoring in an environment where one/more hypervisors are already monitored, you will have to explicitly indicate whether/not the application being added is executing on a virtual host. If so, you will have to additionally pick the managed virtual host with which that application is associated.

25. In order to protect the eG Enterprise system from misuse by malicious users, the eG manager automatically locks out a user if he/she consistently fails to login to the eG management console. You can set when the lockout should occur, how long a user should remain locked out, and can even disable the locking capability, using the **Account Lockout** option in the **MANAGER SETTINGS** panel of Figure 5.50. Selecting this option invokes the **ACCOUNT LOCKOUT** page in the right panel, as depicted by Figure 5.51.

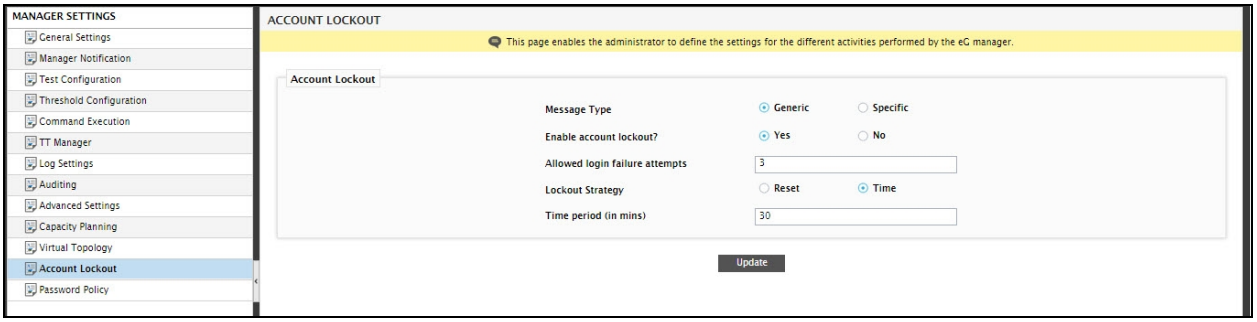

Figure 5.51: Defining Account Lockout Policies

26. Here, specify the following:

- First, choose the type of error message that you want the eG manager to display if a user login fails. If you want the message to clearly indicate the reason for the failure, then set the **Message type** flag to **Specific**. This is useful if you want to troubleshoot the login failure (see Figure 5.52).

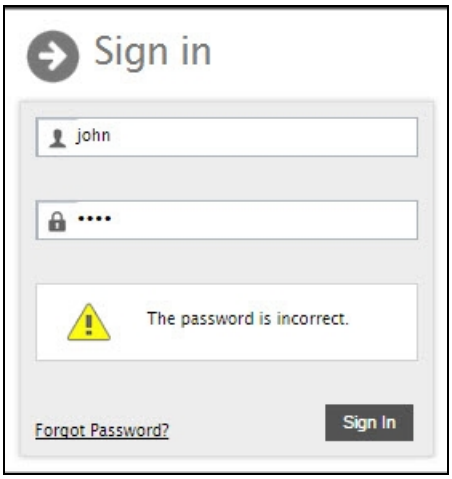

Figure 5.52: A sample login screen with a specific error message

High-security environments on the other hand, may want to be discreet about why a user login was unsuccessful, so as to discourage attempts by unscrupulous users to gain access through devious means. In such a case, its best that the **Message type** flag is set to **Generic**. In this case, when a user login fails, the eG manager will provide only a general failure message, with no specific pointers to why it failed (see Figure 5.53).

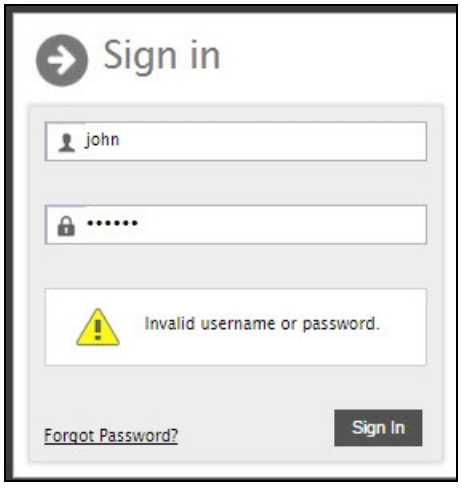

Figure 5.53: A sample login screen with a generic error message

- By default, account locking is enabled for the eG Enterprise system. This is why, the **Enable account lockout?** flag is set to **Yes** by default. If you want to disable this capability, set this flag to **No**. If this is done, then a user can try to login to the eG management console any number of times, without consequence.
- If the account lockout feature is enabled for an eG manager – i.e., if the **Enable account lockout?** flag is

set to **Yes** - then the following settings become applicable:

- Specify the number of unsuccessful login attempts beyond which a user account (registered with eG) will be locked. Specify this number against the **Allowed login failure attempts** text box. By default, the value 3 will be displayed here, indicating that a user account will be locked as soon as that user's third consecutive login attempts fails.
- Mention what the **Lockout strategy** is. A locked user account can be unlocked/released in one of the following ways:
- You can select **Time** as the lockout strategy and set a time duration (in minutes) beyond which a locked user account will be automatically released in the **Time period** text box.
- You can set **Reset** as the lockout strategy if you want a locked account to be released only by an **Admin** user. In this case, the **Admin** user will have to login to the eG administrative interface, access the **LOCKED ACCOUNTS** page in the eG administrative interface (by selecting the **Locked Accounts** menu option from the **User Management** tile), and unlock chosen user accounts.

  Note:

  - If you select **Time** as the **Lockout strategy**, then a user whose account is locked can either wait for the time specified in Figure 5.51 for an automatic release or request for an **Admin** user's intervention to unlock his/her account. However, if **Reset** is the **Lockout strategy**, then a user can have his/her account released only by contacting the **Admin** user.
  - The **Lockout strategy** set does not apply to users with **Admin** privileges to the eG Enterprise system. If an **Admin** user's account gets locked, it will automatically unlock in 1 minute, thus enabling that user to try and login again.
- Finally, click the **Update** button.

27. You can also define a policy for the password you set for local (not domain) users to the eG management console. For this, select the **Password Policy** option from the **MANAGER SETTINGS** panel in Figure 5.51. This will invoke the **PASSWORD POLICY** page in the right panel (see Figure 5.54).

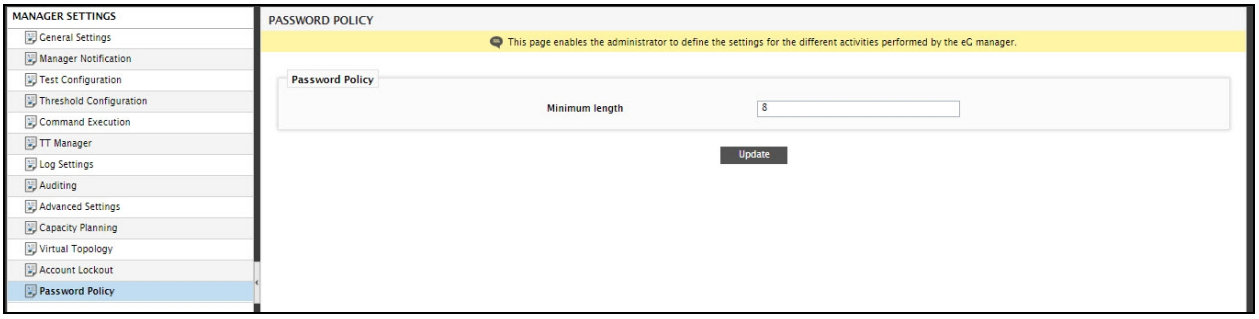

Figure 5.54: Configuring the password policy

28. In Figure 5.54, enter the **Minimum length** for user passwords. When creating a new local user to the eG Enterprise system, the password you specify for the new user should be at least 8 characters long by default. If you want this minimum length changed, use the **Minimum length** parameter in Figure 5.54. Finally, click the **Update** button.

5.5.3 Configuring Logo/Messages

The **Logo / Messages** menu option in the **Settings** tile, provides administrators with options to configure a logo on the right hand side of the eG login, **Monitor**, **Reporter**, and **Configuration Management** interfaces, and to configure audible alerts and custom messages.

Note:

You can include a logo in the **REPORTER** and **CONFIGURATION MANAGEMENT** interfaces, only if your eG license enables the **eG Reporter** and **Configuration Management** capabilities (respectively).

1. To configure a custom logo to be displayed in the eG login screen, click on the **Login Screen** option in the **LOGO & MESSAGES** panel to the left of Figure 5.55. The right panel will then change to display a **LOGIN SCREEN** page.

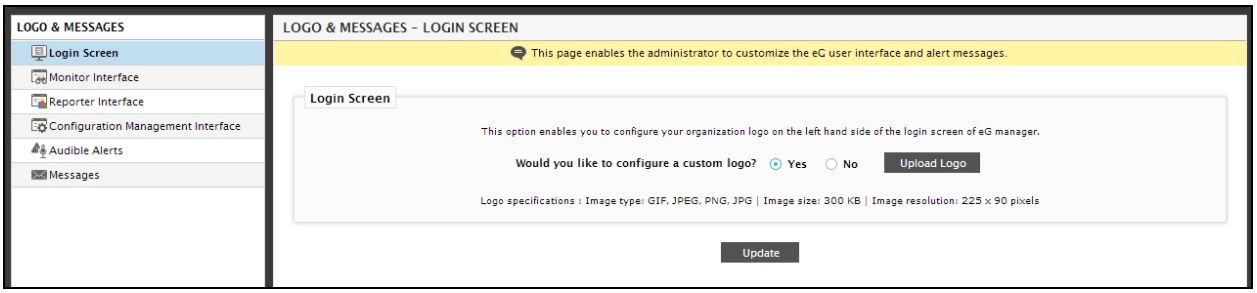

Figure 5.55: Configuring custom logo for the Login Screen

2. Then, set the **Would you like to configure a custom logo for the eG login interface?** to **Yes**. Next, proceed to upload the logo by clicking the **Upload Logo** button. Note that the logo you upload should fulfill the **Logo specifications** displayed in Figure 5.55.
3. Clicking on the **Upload Logo** button will invoke Figure 5.55. Specify the full path to the image file to be uploaded using the **Browse** button therein and click the **Upload** button. Once you return to Figure 5.55, click the **Update** button.

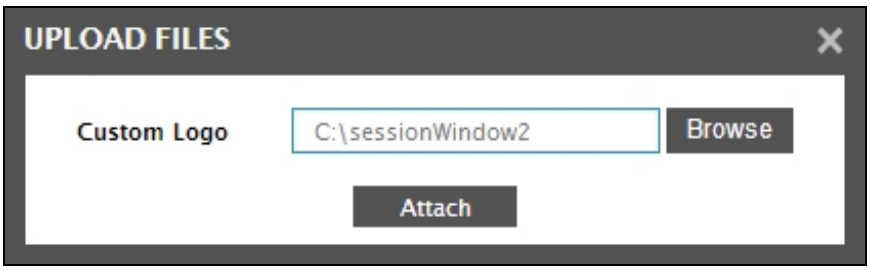

Figure 5.56: Specifying the path for the Logo to be uploaded in the monitor interface

4. To configure a logo for the **Monitor** interface, click on the **Monitor Interface** option in the **LOGO & MESSAGES** panel to the left of Figure 5.55. Then, from the right panel, **Choose the logo type preferences**. If you want to proceed with the **Default** image, set this flag to **Default**. If you do not wish to configure a custom logo for the monitor interface, set this flag to **None**. To define a custom logo, set this flag to **Custom** and click the **Upload Logo** button. In the window that pops up (see Figure 5.56), specify the full path to the image file and

click the **Upload** button. Once you return to Figure 5.57, click the **Update** button.

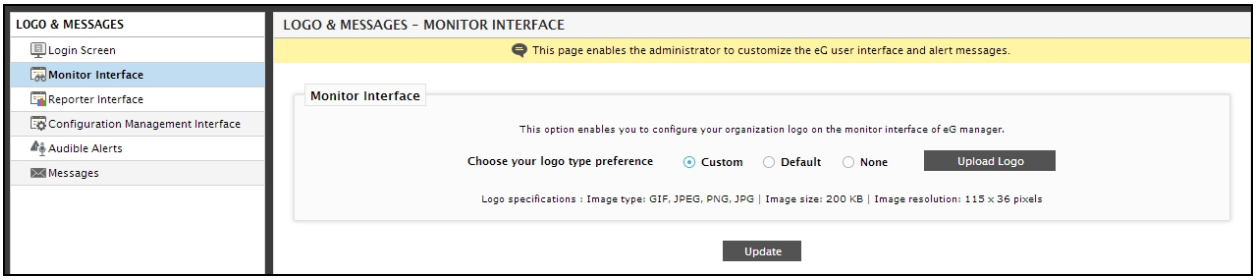

Figure 5.57: Configuring Logo for the Monitor interface

5. For configuring a custom logo for the reporter interface, click on the **Reporter Interface** option in the **LOGO & MESSAGES** panel to the left of Figure 5.55. Then, from the right panel, **Choose your logo type preferences**. If you want to proceed with the **Default** image, set this flag to **Default**. If you do not wish to configure a custom logo for the reporter interface, set this flag to **None**. To define a custom logo, set this flag to **Custom** and click the **Upload Logo** button. In the window that pops up (see Figure 5.58), specify the full path to the image file and click the **Upload** button. Once you return to Figure 5.58, click the **Update** button.

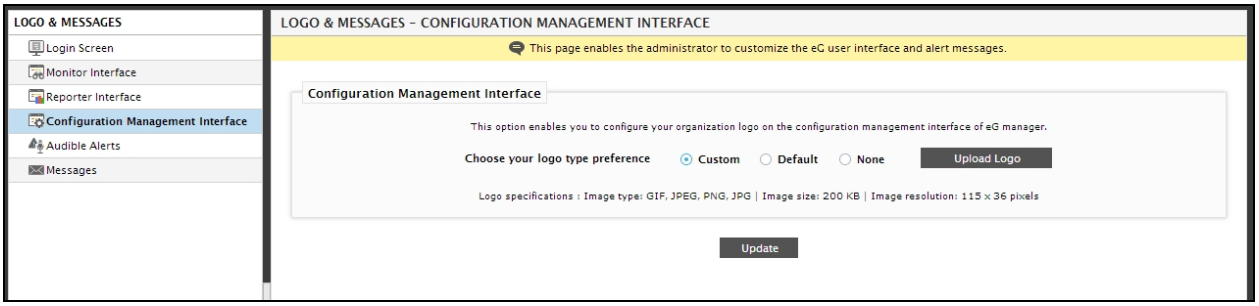

Figure 5.58: Configuring a custom logo for the Configuration Management interface

6. For configuring a custom logo for the configuration management interface, click on the **Configuration Management Interface** option in the **LOGO & MESSAGES** panel to the left of Figure 5.55. Then, from the right panel, **Choose your logo type preferences**. If you want to proceed with the **Default** image, set this flag to **Default**. If you do not wish to configure a custom logo for the configuration management interface, set this flag to **None**. To define a custom logo, set this flag to **Custom** and click the **Upload Logo** button. In the window that pops up (see Figure 5.56), specify the full path to the image file and click the **Upload** button. Once you return to Figure 5.59, click the **Update** button.

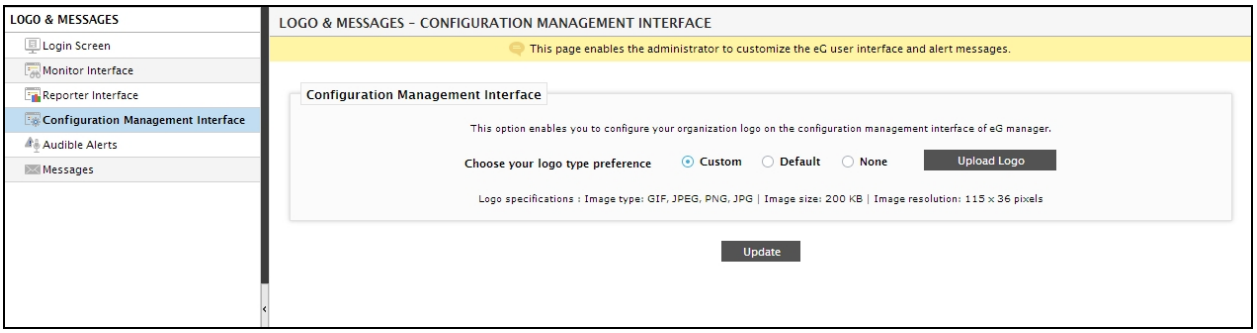

Figure 5.59: Configuring a custom logo for the Configuration Management interface

7. You can also associate audio files of your choice with alarms. To achieve this, first, click on the **Audible Alerts** option in the **LOGO & MESSAGES** panel to the left of Figure 5.55.

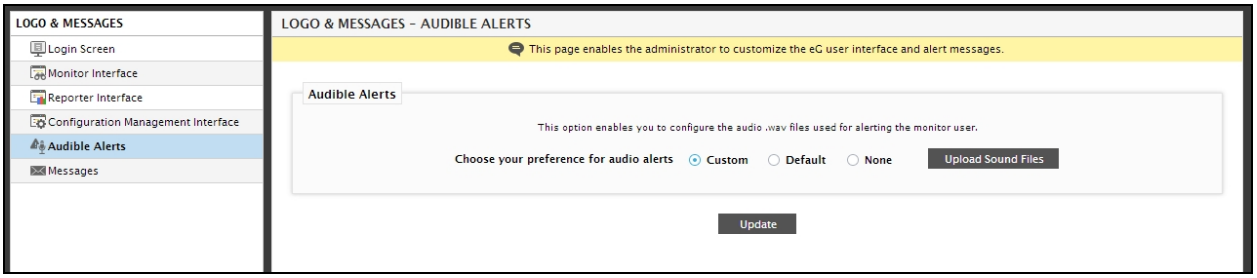

Figure 5.60: Associating audio files with alarms

8. From the **AUDIBLE ALERTS** page (see Figure 5.60) that appears in the right panel, **Choose your preference for audit alerts**. If you want to proceed with the **Default** audio, set this flag to **Default**. If you do not wish to configure any audio files, set this flag to **None**. To associate an audio file with the alarms, set this flag to **Custom** and click the **Upload Sound Files** button. In the window that pops up (see Figure 5.61), use the **Browse** button to specify the full path to the audio files to be associated with **Critical**, **Major**, and **Minor** alarms. Then, click the **Upload** button.

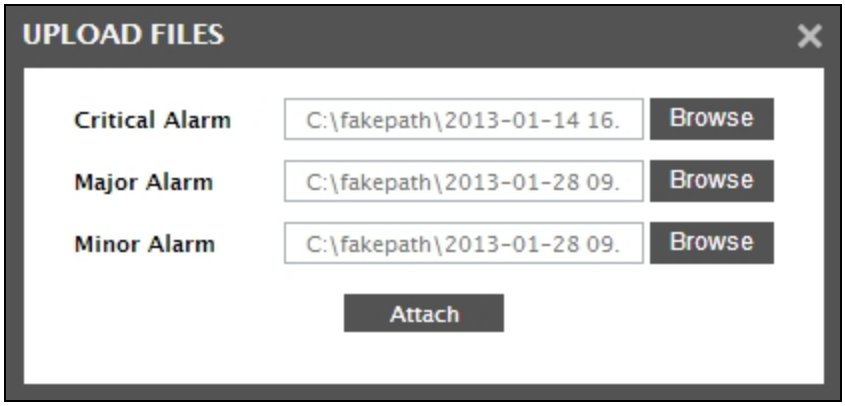

Figure 5.61: Associating audio files with alarms

9. Once you return to Figure 5.61, click the **Update** button.
10. The **MESSAGES** page (see Figure 5.62) that appears when the **Messages** option in the **LOGO & MESSAGES** panel is clicked enables the administrator to display customized messages.

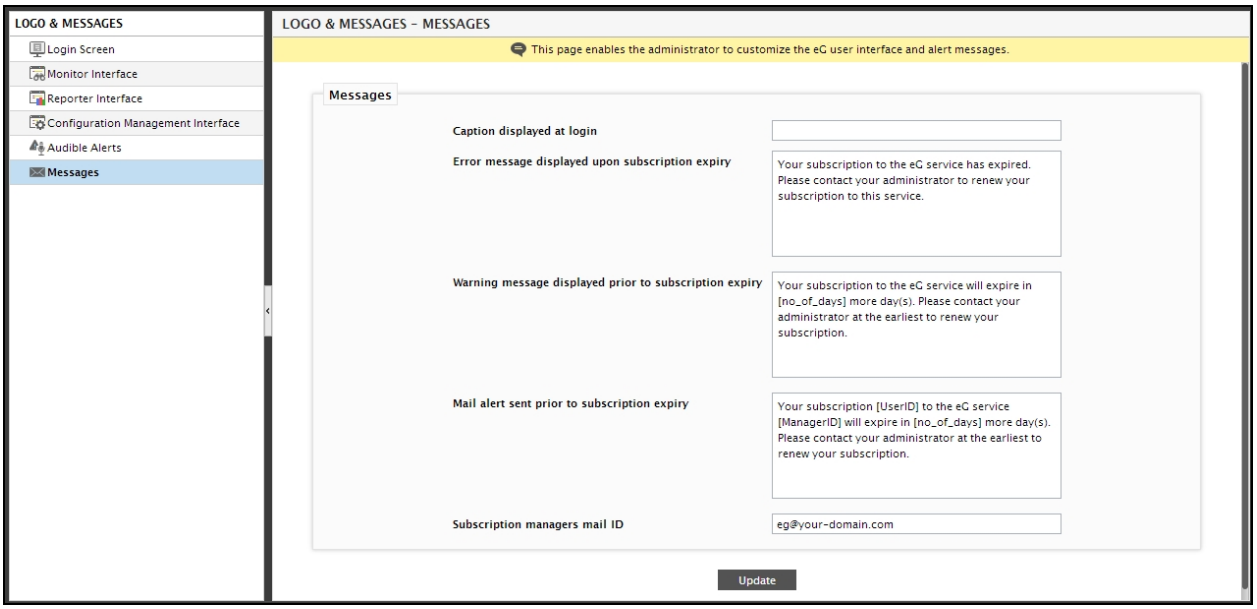

Figure 5.62: Configuring custom messages

11. This page enables administrators to specify a custom caption to be displayed in the login screen. This caption can be specified in the **Caption displayed at login** text box. T he Error message displayed upon subscription expiry text box contains the default message that would appear upon expiry of a subscription. The contents of this text box can be changed if required. The Warning message displayed prior to subscription expiry text box contains the warning message that would appear prior to the subscription expiry. As you can see, this message contains a variable [no_of_days]. This variable indicates the number of days left for the subscription to expire and hence should remain unchanged. In order to alert the user about the impending subscription expiry without him/her having to login to the user interface, eG sends an email to the user ID. The body of this email can be customized using the **Mail alert sent prior to subscription expiry** text box as depicted in the Figure 5.62. In the default contents of this text box, the variable [UserID] refers to the ID of the user whose subscription is about to expire, [ManagerID] refers to the ID of the manager that is being used by the user, and [no_of_days] refers to the number of days left for the subscription to expire. The Subscription manager's mail ID field indicates the mail ID from which the warning messages prior to the subscription expiry are sent to different users.

12. Finally, click the **Update** button.

6

User Management

In large enterprises, the management may not want to grant unrestricted administrative/monitoring access to all users in the environment – hence, such environments may not allow just about any user to login using the default 'admin' and 'supermonitor' logins; they may want a user's access rights to be aligned with his/her organizational responsibilities; and some may want each user to be able to view the status of only those components that have been specifically provisioned for them;

The User managemet helps the administrator to add, modify, and delete the user roles/domains/users.

The **User Management** tile of the **Admin** tile menu enables an **Admin** user to perform the following:

- Add, delete and modify new user roles
- Adding/Modifying/Deleting a Domain
- Add a new user
- Delete an existing user
- Change the profile for any other user
- Add/modify/delete a domain
- Release locked accounts
- Change the password of the administrator
- View different reports corresponding to the users

6.1 Adding New User Roles

In large enterprises, the IT staff have clearly demarcated roles and responsibilities. The help desk staffs are responsible for handling user complaints and their main concern when a user calls about a problem is to determine whether the user call pertains to a problem that the other operations staff is already working on. The domain experts and service managers are responsible for the early detection, diagnosis and fixing of problems with the networks, servers, applications, and services they control. While the domain experts are interested in the detailed performance metrics relating to the IT infrastructure, the executive managers are interested in high-level service level reports that detail if the IT infrastructure is meeting the service expectation of their users. To support these varying requirements of the IT operations staff, eG Enterprise supports different user roles. The user roles define the rights and responsibilities that any user of the eG Enterprise system has. Each user in the eG Enterprise system is assigned to a user role.

By default, eG Enterprise embeds two users namely, *admin* and *supermonitor*. The *admin* user reserves the administrative rights to the monitored environment, and also receives an unrestricted view of the monitored environment. Only users with the privileges of the *admin* user can add new users or new roles to the eG Enterprise system. The *supermonitor* user cannot perform administrative tasks, but is authorized to monitor the performance of the entire environment.

To view the default roles and to create, modify, or delete new roles, do the following:

1. From the **User Management** tile, select the **Roles** menu option. Figure 6.1 will then appear.

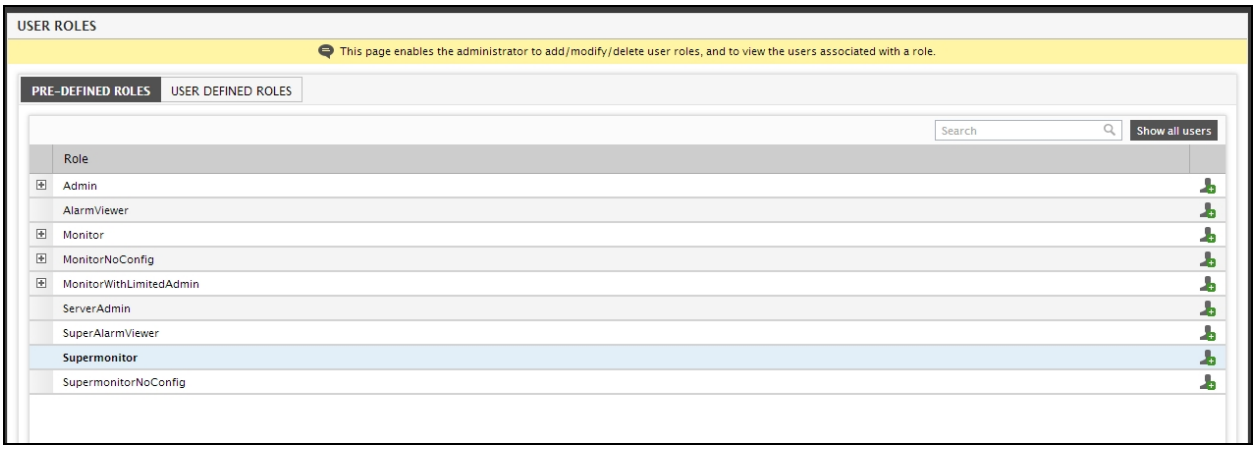

Figure 6.1: The list of default roles

2. As can be inferred from Figure 6.1, the **PRE-DEFINED ROLES** tab page of the **USER ROLES** page opens by default. This tab page displays the default roles pre-defined by the eG Enterprise system. These are as follows:

   - **Admin :** Users who are assigned administrative rights become the super-users of the system. Such users can choose what hardware and application servers are to be monitored by the system, where the agents should be executed to monitor the hosted environment, what tests these agents should run, how often these tests should be executed, and can view the status of the entire monitored infrastructure. The administrative user also has the rights to add, delete, and modify user roles and individual user profiles. The default *admin* user is assigned the **Admin** role only.
   - **Monitor :** If the eG license enables the **eG Reporter** and **Configuration Management** capabilities, then **Monitor** users will have access to the monitoring, reporting, and configuration management consoles of eG Enterprise. In these consoles, the monitor user can view the details pertaining to only those components/segments/services/zones/service groups that have been explicitly assigned to him/her. Each monitor user is associated with an email address to which alarms pertaining to the assigned elements will be forwarded. The user's profile also includes information regarding his/her alarm preferences - whether alarms have to be forwarded in text or HTML mode, whether a complete list of alarms has to be generated each time a new alarm is added, or whether the new alarm alone should be sent via email, etc. Each monitor user is associated with a subscription period. eG Enterprise allows the monitor users to access the system until this period only.
   - A **Supermonitor** user has an unrestricted view of the monitored infrastructure. He/she can receive alarms pertaining to the whole infrastructure that has been configured by the administrative user. The default *supermonitor* user is assigned the **Supermonitor** role only. A **Supermonitor** user is allowed access to the reporting and configuration management modules as well, provided the eG license enables the **eG Reporter** and **Configuration Management** capabilities.
   - **AlarmViewer**: This role is ideal for help desk personnel. The users vested with **AlarmViewer**

permissions can login to the monitor interface, and perform the following functions:

- View the details of alarms associated with the specific components and services assigned to them
- Provide feedback on fixes for the alarms
- View feedback history
- Change their profile

Like **Monitor** users, users with this role can only monitor the components assigned to them.

- **SuperAlarmViewer**: Users with the **SuperAlarmViewer** role have all the privileges of the **AlarmViewer** role. In addition, users with the **SuperAlarmViewer** role have access to all the components being monitored.
- **ServerAdmin**: The users who have been assigned the **ServerAdmin** role have all the administrative rights of an **Admin** user, except **the right to user management**. Similarly, like a **Supermonitor** user, a **ServerAdmin** user can monitor the complete environment, and even change his/her profile.
- **MonitorNoConfig**: The users who have been assigned the **MonitorNoConfig** role will have access to the eG monitoring and reporting interfaces only, and not the eG **Configuration Management** interface.
- **SupermonitorNoConfig**: Users with **SupermonitorNoConfig** privileges will have unrestricted access to the monitoring and reporting consoles only - such users will not be able to access the configuration management console, even if the eG license enables this capability.
- **MonitorwithLimitedAdmin**: Administrators can create additional users with administrative privileges to configure the monitoring for the components that are assigned to them. These users can configure tests, thresholds, alarm policies and maintenance policies for the components in their purview. The MonitorwithLimitedAdmin role included in eG Enterprise can be used to create such users. This capability allows delegated administration, which is a key requirement for many enterprises and service providers.

3. If too many roles are listed in this page, you can quickly search for a particular role using the **Search** text box in this page. Specify the whole/part of the role name to search for in the `Search` text box. All role names that embed the specified string will then appear in this page (see Figure 6.2).

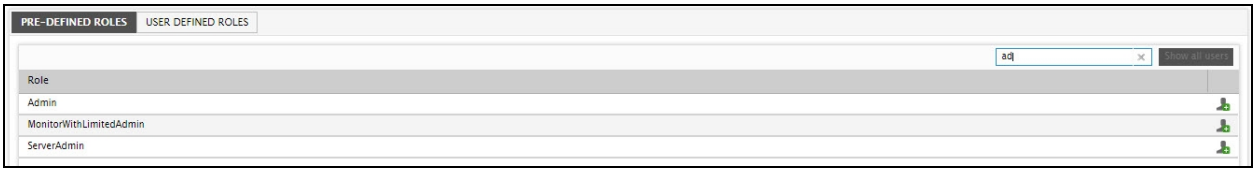

Figure 6.2: Searching for a role

4. Roles that have already been assigned to specific users are highlighted by a '+' symbol preceding the role names. If you want to view the users who have been assigned a role, click on the '+' button that pre-fixes the role. This will expand the role to reveal the users (see Figure 6.3).

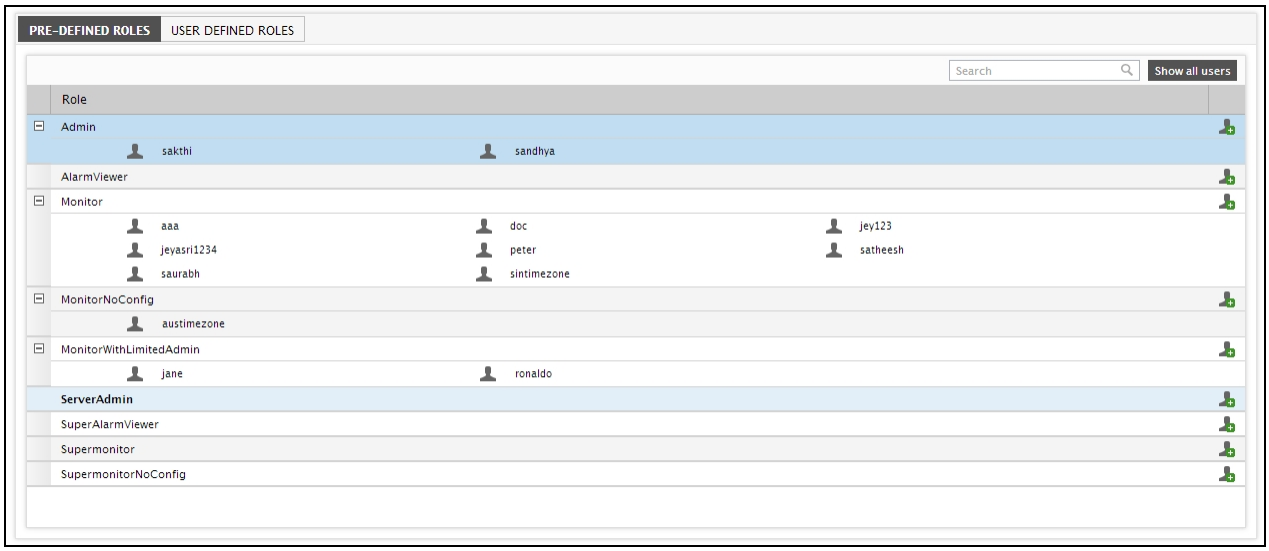

Figure 6.3: Users who are assigned a particular role

5. Clicking on a user name in Figure 6.3 will lead you to the **MODIFY USER** page, using which you can modify the profile of that user.

6. If you want to view at one shot, which users have been assigned which roles, just click the **Show all users** button next to the **Search** text box in Figure 6.3. Figure 6.4 will then appear. To hide the users list that accompanies all roles, click on the **Hide all users** button next to the **Search** text box in Figure 6.4.

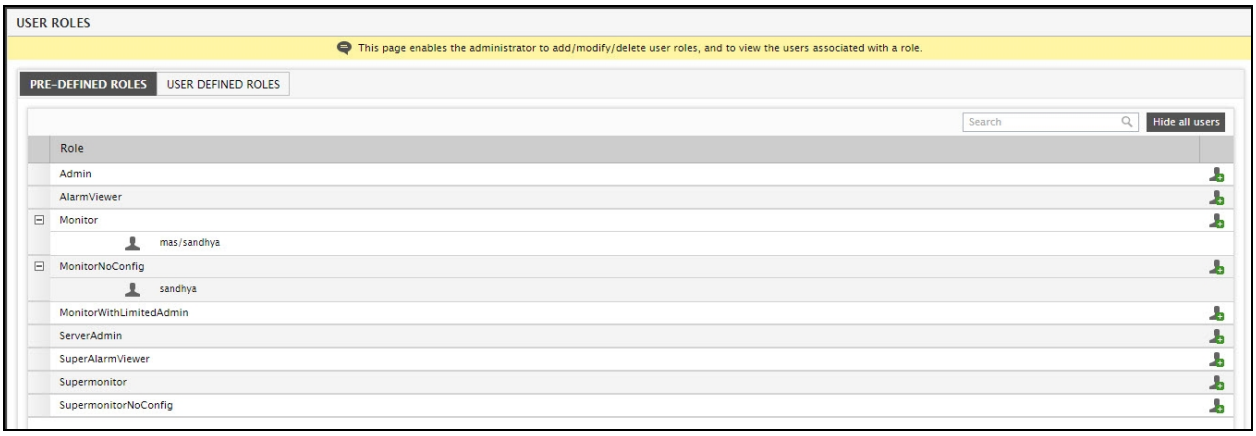

Figure 6.4: Showing which user has been assigned which role

7. To add a new user for a role, just click the **Add User** icon corresponding to that role in Figure 6.4. This will lead you to the **ADD USER** page, where you will find the chosen role automatically displayed against the **User role** list. You can then proceed to create a new user who is assigned that role.

8. To add a new role on the other hand, first, switch to the **USER DEFINED ROLES** tab page by clicking on it. If any custom roles pre-exist, they will be listed in that appears. If no custom roles exist, then a message to that effect will be displayed here. To create a new role, click the **Add New Role** button in Figure 6.5.

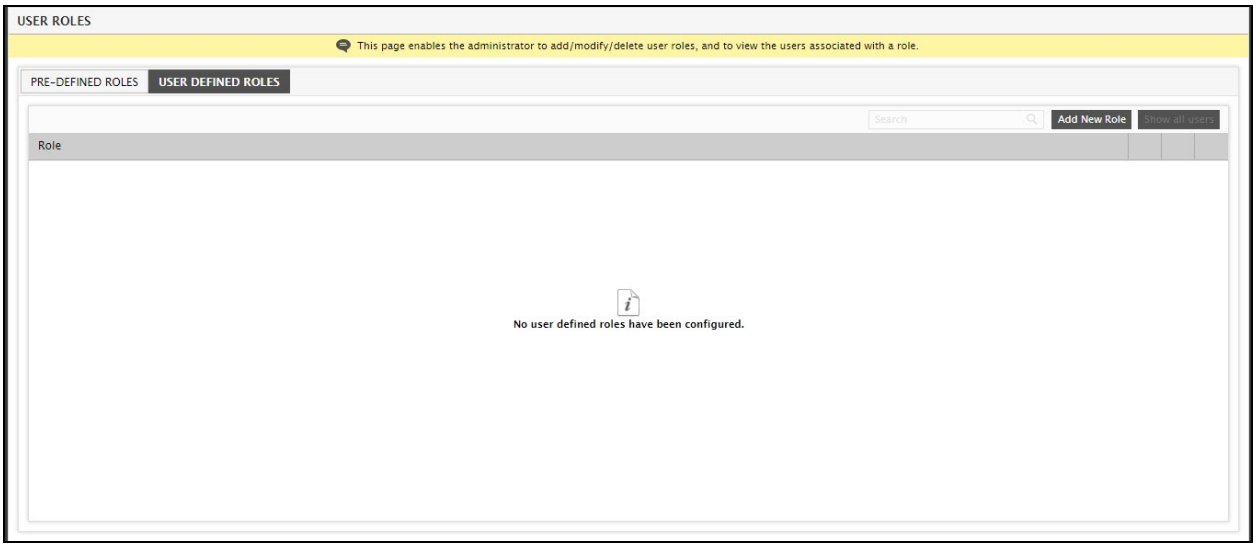

Figure 6.5: The User Defined Roles tab page indicating that no custom roles pre-exist

9. Figure 6.6 will then appear.

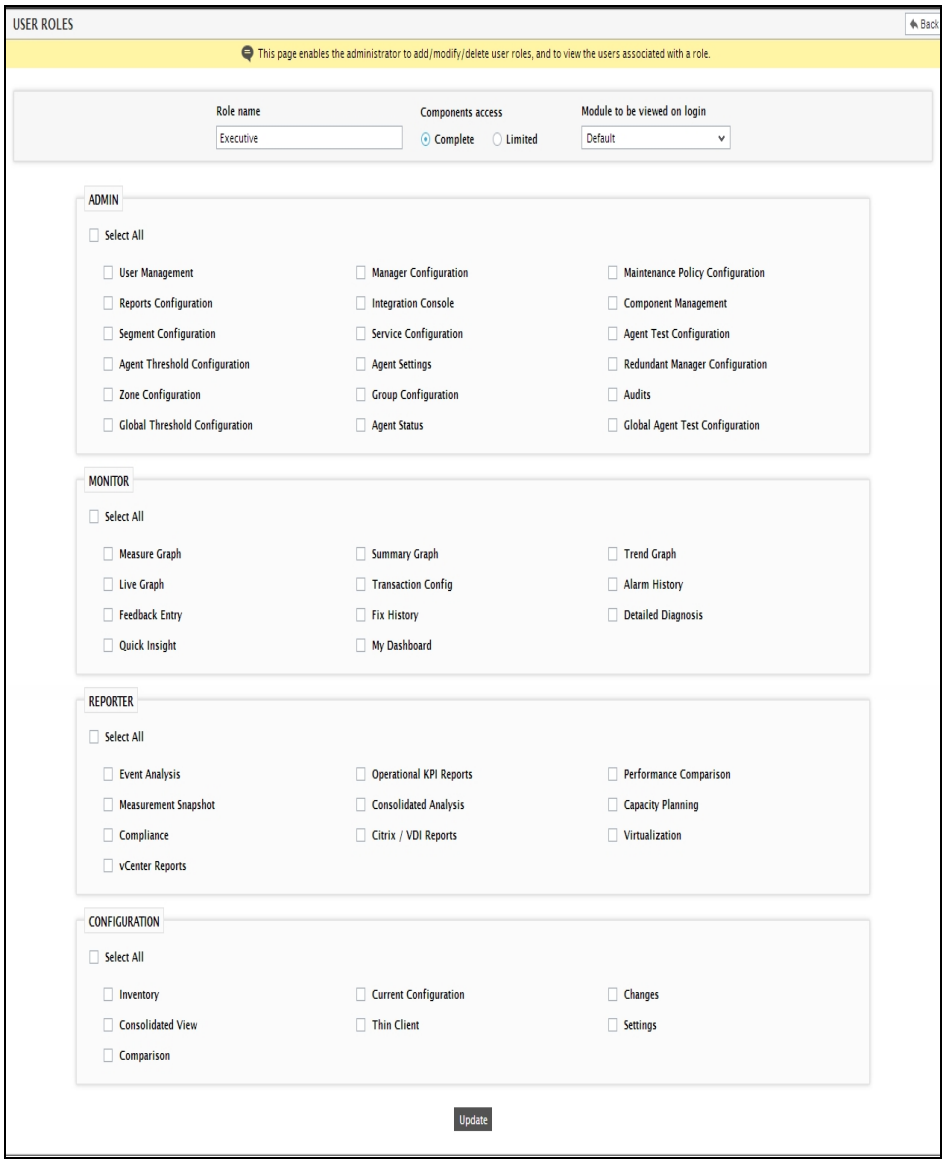

Figure 6.6: Creating a new role

10. In Figure 6.6, provide a name for the role against **Role name**.
11. Next, indicate whether **Limited** or **Complete Components access** is to be granted to the new role. Selecting the **Limited** option restricts the new role's access to specific components/segments/services/zones that have been configured in the infrastructure. This is ideal for MSP environments, which cater to the hosting requirements of multiple customers. By assigning a role that allows only **Limited** component access to each of its customers, the MSP can ensure that every customer has access to only those infrastructure elements that are specific to his/her hosted environment.
12. Typically, if a user has access to the **Admin** and **Monitor** modules, by default, when the user logs in he/she would have access to the **Admin** module. Likewise, if a user has access to the **Monitor** and **Reporter** modules, the **Monitor** module would be the default module when the user logs in. This default behavior can be altered by selecting an option from the **Module to be viewed on login** list. For instance, in Figure 6.6

above, the user role Executive has been granted both monitoring and reporting rights - i.e., a user who is assigned the Executive role will be able to access both the **Monitor** and the **Reporter** modules. By default, the **Module to be viewed on login** for this role is set to **Default**; this implies that the **Monitor** module will be the default module for the Executive user upon login. However, you might have granted extensive report-generation rights to the Executive role and limited monitoring rights, and hence, might prefer to set **Reporter** as the default module. In such a case, to grant the Executive role primary access to the **Reporter** module and not the **Monitor** module, select the **Reporter** option from the **Module to be viewed on login** list.

13. In any monitored environment typically, administrators alone have the right to make configuration changes using the eG administrative interface. Monitor users on the other hand have no access to the administration console. In large enterprises, multiple distinct administration teams may use the same eG Enterprise manager for their monitoring. These teams would require the ability to configure the monitoring for the servers they operate. To address such environments, eG Enterprise includes the capability to configure users with limited administration rights. For instance, a separate role can be created to allow monitor users with just the permissions to configure tests that should be executed on their servers, or to change the thresholds that can be applied for monitoring their servers. This is why, as soon as the **Limited** option is chosen, all the check boxes except the **Agent Test Config**, **Agent Threshold Config**, and **Maintenance Policy Config** check boxes, are grayed out in the **Admin** section of Figure 6.6. This implies that user roles with **Limited** component access can only perform one/more of the following administrative functions:

    - Configuring tests pertaining to the components assigned to them
    - Configuring the thresholds related to the components under their monitoring purview
    - Suppressing the alerts related to the components assigned to them by configuring maintenance policies

    On the other hand, if the **Complete** option is chosen, it implies that the user role has access to all the monitored elements in the infrastructure, and can be granted any administrative/monitoring privilege as the administrator deems fit.

14. If administrative privileges need to be assigned to the new role, then select the privileges from the **ADMIN** section. To assign all the admin privileges to a role, select **Select All**. As stated earlier, if the **Limited Components access** option is chosen, then except the **Agent Test Config**, **Agent Threshold Config**, and **Maintenance Policy Config** check boxes in this section, all other check boxes will be disabled.

15. To provide the new role with access to all the features of the eG monitor interface, select the **Select All** check box in the **MONITOR** section. To grant specific monitoring rights to the role, select the individual monitor modules from the **MONITOR** section.

16. If the eG license enables the **eG Reporter**, then a **REPORTER** section will appear in Figure 6.6. If the new role has access rights to all the **REPORTER** modules, then click on the **Select All** checkbox in the **REPORTER** section. To restrict access to specific reporter modules, select the required modules from the **REPORTER** section.

17. Similarly, if the eG license enables **Configuration Management**, then a **CONFIGURATION** section will appear in Figure 6.7. If the new role has access rights to all the **Configuration Management** modules, then click on the **Select All** checkbox in the **CONFIGURATION** section. To restrict access to specific modules, select the required modules from the **CONFIGURATION** section.

18. Finally, click the **Update** button. Figure 6.8 will then appear, displaying the newly added role.

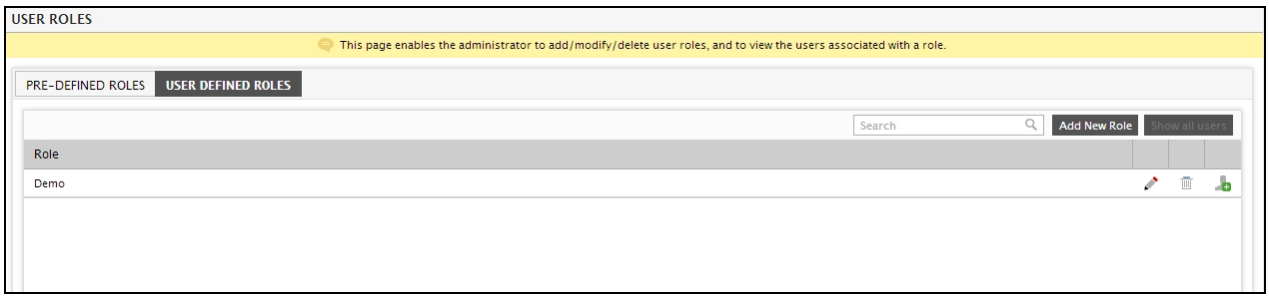

Figure 6.7: The newly created role being displayed in the list of roles

19. Note that while the **PRE-DEFINED ROLES** can neither be deleted nor modified, the user-defined role that was newly added can be modified by clicking on the **Modify** icon (i.e., the 'pencil' icon) corresponding to that role in Figure 6.8. To delete a particular role, use the **Delete** icon (i.e., the 'trash can' icon) against that role in Figure 6.7. However, note that if any of the user-configured role has been assigned to any new user registered with the eG Enterprise system, then such roles cannot be deleted; therefore the **Delete** icon corresponding to such roles will be disabled. You can even create a new user for a role instantly, by clicking on the **Add User** icon corresponding to that role. This will lead you to the **ADD USER** page, where you will find the chosen role automatically displayed against the **User role** list. You can then proceed to create a new user who is assigned that role.

6.2 Adding/Modifying/Deleting a Domain

The eG administrative interface provides administrators with a wide variety of options to manage user information. Be it user creation, modification, deletion, or simply viewing user information, any type of user-related activity can be performed quickly and easily using the eG administrative console. Typically, when an eG user logs into the eG Enterprise system, the login is validated by the eG database, which stores the user information. However, in large IT environments that span multiple domains, the Active Directory server functions as the central repository for information related to users spread across domains, and also authenticates domain user logins. To ensure that the AD server continues to be the central authority for validating domain user logins to the eG Enterprise system and not the eG database (as in the case of local users), administrators might want the eG manager to integrate with AD.

To enable this integration, the eG administrative interface allows the following:

1. Automatic discovery / manual creation of one/more domains and sub-domains (if any) configured in the target environment;
2. Addition of a domain user to the eG Enterprise system, and validating the user's logins by automatically connecting to the associated domain.

Since step 2 above has already been discussed in Section **6.3** of this document, let us proceed to discuss how step 1 can be performed - i.e., how domains can be created in the eG Enterprise system.

1. To achieve this, first, select the **Domains** option from the **User Management** tile in the **Admin** tile menu. Figure 6.8 then appears, using which, you can create multiple parent and child domains.

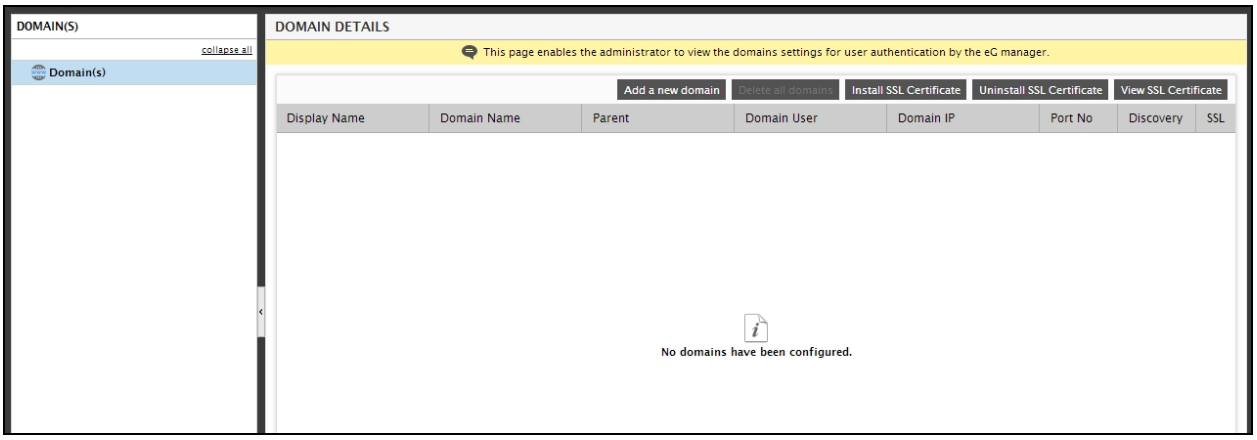

Figure 6.8: The DOMAIN DETAILS Page

As you can see, Figure 6.8 consists of two panels - a left panel that hosts a tree structure, and a context-sensitive right panel, the contents of which vary according to the node chosen from the tree. Typically, the parent domains that you configure will be the primary nodes of the tree-structure in the left panel, and the sub-domains will be the sub-nodes. By default, the global **Domain(s)** node will be chosen from the tree. Accordingly, the right panel will display the complete details of all the parent and child domains that pre-exist in the eG Enterprise system. If no domains pre-exist, then the message depicted by Figure 6.8 will be displayed in the right panel.

Let us now proceed to create domains. The first step towards creating multiple domains is to create a *parent domain*. Before attempting to create a parent domain, you would have to choose between the following:

- Automatically discovering the IP address and port number of the domain server
- Manually configuring the IP address and port number of the domain server

Automatic domain discovery is recommended if you are not certain about the IP and port number on which the AD server functions, or if the IP address of the AD server is configured to change frequently (for eg., in a DHCP environment). On the other hand, you might opt for manual domain configuration, if the IP/Port number of the AD server is static.

The sections that follow will discuss both these approaches in great detail.

6.2.1 Automatically Discovering Parent and Child Domains

The eG manager is capable of **automatically discovering only that domain in which it has been deployed**. To auto-discover the parent domain in which the eG manager operates, follow the steps given below:

1. Click on the **Add a new domain** button in the right panel of Figure 6.8.
2. Doing so displays the domain configuration parameters in the right panel (see Figure 6.9).

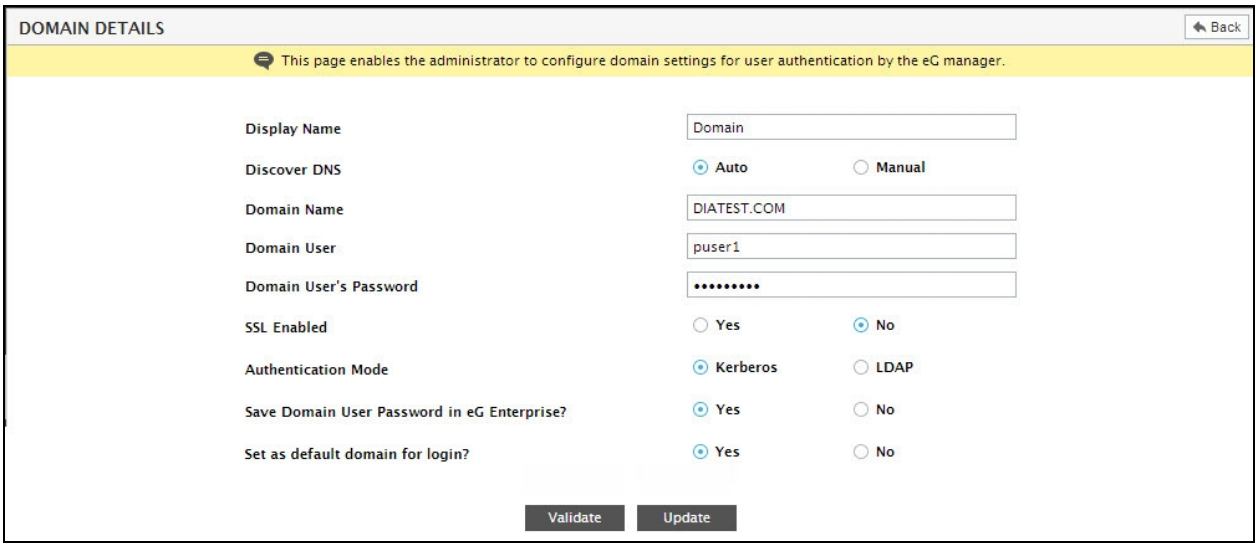

Figure 6.9: Automatically discovering the eG manager's domain

3. To auto-discover the eG manager's parent domain, specify the following in Figure 6.9:

   - First, provide a **Display Name** for the domain in the right panel of Figure 6.9.
   - Next, indicate whether or not the eG manager needs to auto-discover the IP/Port number of the AD server. To auto-discover the domain, set the **Discover DNS** settings flag to **Auto**.

     Note:

     Note that **only the domain in which the eG manager is deployed can be auto-discovered**.
   - Next, specify the fully-qualified **Domain Name**.
   - To connect to the AD server and access the domain user information stored within, the eG manager requires a domain user's privileges. To facilitate this connection, provide a valid domain user's name and password against **Domain User** and **Domain User's Password**.
   - Then, indicate whether the AD server is SSL-enabled or not, by setting the **SSL** flag to **Yes** or **No**, as the case may be. If the **SSL** flag is set to **Yes**, then you will have to follow the procedure discussed in the **Appendix** below to ensure that the eG manager is able to communicate with the AD server over SSL.
   - Next, indicate how accesses to the AD server are to be authenticated - using **Kerberos** or **LDAP**. **Kerberos** is a computer network authentication protocol which works on the basis of "tickets" to allow nodes communicating over a network to prove their identity to one another in a secure manner. **Kerberos** is ideal for AD environments with high security considerations. The **Lightweight Directory Access Protocol** on the other hand, is an application protocol for accessing and maintaining distributed directory information services over an Internet Protocol (IP) network. The **LDAP authentication mechanism** is best suited for environments with not very high security constraints.
   - Next, indicate whether the **Domain User's Password** that you have provided here for enabling the eG manager to connect to the AD server, should be saved in eG Enterprise or not. To save the password, set the **Save Domain User Password in eG Enterprise?** flag to **Yes**. If this is done, then, the specified **Domain User's Password** will be automatically encrypted and saved to the **eg_authenticate.ini** file, which will be available in the **<EG_MANAGER_INSTALL_DIR>\manager\config** directory. On the other hand, if the

Save Domain User Password in eG Enterprise? flag is set to **No** instead, the password will not be saved to the **eg_authenticate.ini** file. If the password is not saved, then every time the eG manager attempts to connect to the AD server - say, when validating/registering domain user profiles configured on the eG manager (using the **ADD USER** page) with the AD server - you will be prompted for the **Domain User's Password**.

- Also, indicate whether/not the domain being configured should be set as the default domain at the time of login. To set the new domain as the default domain, set the **Set as default domain for login?** flag to **Yes**. If this is done, then the next time a user attempts to log into the eG management console by typing his/her user name in the login page, the **Domain** selection will instantly change from Local to the domain that you have set as the default. This capability is most useful in environments where the eG manager integrates with only one domain. By setting this domain as the default, administrators can save users the trouble of selecting a **Domain** every time he/she tries to login.
- In virtual environments where **LDAP** is used to authenticate access to the AD server, administrators may want to keep track on specific user information for e.g., location, vendor etc of the users accessing their environment through the AD server. For example, in addition to viewing the user experience with their virtual environment, if administrators are able to view the location of the user, it would help them troubleshoot location specific issues at the earliest. This approach would definitely help administrators improve the overall performance of their environment. To view such user specific information in the eG monitoring console, administrators should do the following:
  - First, set the **Discover User Details from AD** flag to **Yes**. By default, this flag is set to **No**.

    Once this flag is set to **Yes**, the user specific information will automatically be populated in the **ADUserDetails.ini** file that is located in the *<eG_INSTALL_DIR>/manager/config* location.

    If the **Discover User Details from AD** flag is set to **Yes**, then an additional **Update User Details from AD** option will appear in the **What would you like to do?** list in the right panel as shown in Figure 10. Clicking the **Update** button will immediately integrate the user information from the domain to the **ADUserDetails.ini** file available in the **<EG_INSTALL_DIR>\manager\config** directory.

    Note:

    By default, the user information available in the domain will be integrated with the **ADUserDetails.ini** file once in 7 days. If you wish to override this default, setting, then you can do the following:

    - Edit the **eg_services.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory).
    - Set the **ThreadFrequency** parameter in the **[ADUSERDETAILS_THREAD_SETTINGS]** section of the file to a frequency of your choice.
    - By default, the information will be integrated every **Sunday**. If you wish to override this default day, then you can change the **DayToRun** parameter to the day of your choice.
    - Save the file.

    The user specific information so updated can be viewed in the eG monitoring console in the following features offered by the eG Enterprise Suite:

- User Experience Dashboard
- Current Alarms
- Layer model page of the tests where users are the descriptors of the tests

To view the user specific information in the User Experience Dashboard, you have to edit the **<eG_ INSTALL_ DIR>/manager/config/eG_ enduserdetails.ini** file with the procedure mentioned below:

First, if you want to view user specific information in the User Experience dashboard for VDI environments, then you have to set the **VDI:ShowUserLocations** flag under the **[GEO_ LOCATION_SETTINGS]** location to **true**.

Set the **XenApp7:ShowUserLocations** flag to **true** if you want to view the User Experience dashboard for Citrix XenApp 7 and above environments. Set the **XenApp:ShowUserLocations** flag to **true** if you want to view the User Experience Dashboard for Citrix XenApp servers.

Once you have set the **ShowUserLocations** flag to true accordingly, you have to specify the format of the user details that were discovered from the AD server and populated in the ADUserDetails.ini file. This can be achieved using the *<UserType>:Format=Vendor-City-CompanyName* where *UserType* can be *VDI* or *XenApp7* or *XenApp*. For example, if you want the User Experience Dashboard for VDI environments, then you can specify the format as:

VDI:Format=Vendor-City-CompanyName

Then, you have to specify the separator using which the user details can be separated into columns while being displayed in the User Experience dashboard against the Separator field. By default, the separator is hyphen (-).

Once you have specified the format and separator, you have to provide the display name for all user specific information that you have mentioned against the format field. By default, the display name for certain user specific information that you have mentioned in the *UserType:Format* section will be specified under the **[USER_DESKTOP_METRICS]** section.

If you want to include the address of the user in the dashboard, then you have to follow the following format:

;GeoDetails:Address~$~Address

In the above example, the format is *GeoDetails:<Internalname specified in the UserType:Format section>~$~D<Displayname>*.Once you have provided all the entries, the final step in this process is to specify the columns that should appear in the User Experience Dashboard. By default, the eG Enterprise provides out of the box support to display the Vendor, City and Company of the users in the User Experience Dashboard. If you want to include the address of the VDIusersinthe dashboard,then you have to append the *GeoDetails*:Address~$~Address format to the "VDI=" section under the **[USER_DESKTOP_ METRICS]** section as shown below:

VDI=EsxLoginTest:New_ logins:LoginTime:DATE:Logon Time, GeoDetails:Address~$~Address

To view the user specific information in the current alarms and the layer model, you have to append the user specific test to the **[Show_ User_ location]** section of the **eg_ dashboardConfig.ini** file which is situated in the **<eG_INSTALL_DIR>/manager/config** location. By default, the entry in this section should be in the following format:

<Testname>:iconDesktopUser

If you have to view the user specific information, then you have to remove the semicolon(;) in front of the <*Testname*>.

Once you have configured the necessary files, it is mandatory for you to restart the eG manager to effect your changes.

- Then, to verify the correctness of your specifications, click the **Validate** button. Figure 6.10 will then appear indicating whether/not the **Display Name, Domain Name**, **Domain User**, and **Domain User Password** values that you have provided are indeed valid.

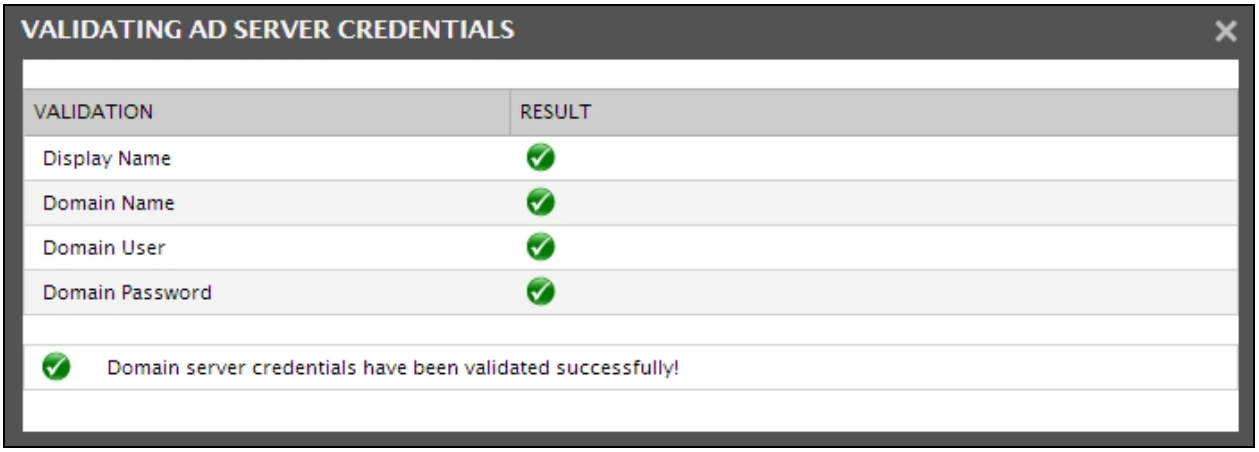

Figure 6.10: Validating domain server credentials

- Since the eG manager auto-discovers the IP/Port of the AD server, you will not be prompted to manually specify the same. Therefore, simply click the **Update** button to add the new domain.
- Once the parent domain is auto-discovered and added to the eG Enterprise system, a message to that effect will appear (see Figure 6.11). The tree in the left panel will also change to reflect the addition of the parent domain.

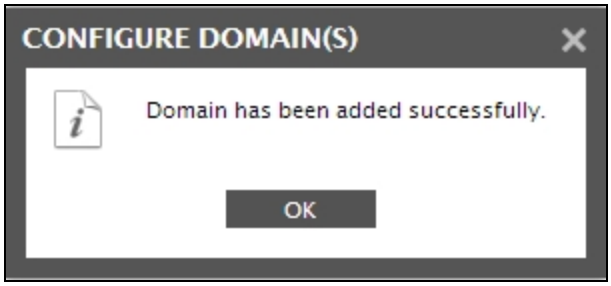

Figure 6.11: A message box informing the successful addition of a domain

4. By default, only the **parent domain of the eG manager can be auto-discovered**; the child domains (if any) under this parent domain will not be auto-discovered. This is because, the **AutoDiscoverChildDomains** flag in the **[MISC_ARGS]** section of the **eg_services.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory) is set to **false** by default. If need be, you can configure the eG manager to automatically discover the child domains along with the eG manager's parent domain; to achieve this, set the **AutoDiscoverChildDomains** flag to **true**. In this case therefore, the child domains (if any) will also be automatically discovered along with the eG manager's parent domain and will be displayed as sub-nodes of the parent domain's node, as depicted by Figure 6.12 below.

5. To view the details of all domains (both parent and child) that have been configured, click on the **Domains** node tree. The right panel will then change to display a tabular column, where you can view the configuration of all the domains that you have created (see Figure 6.12).

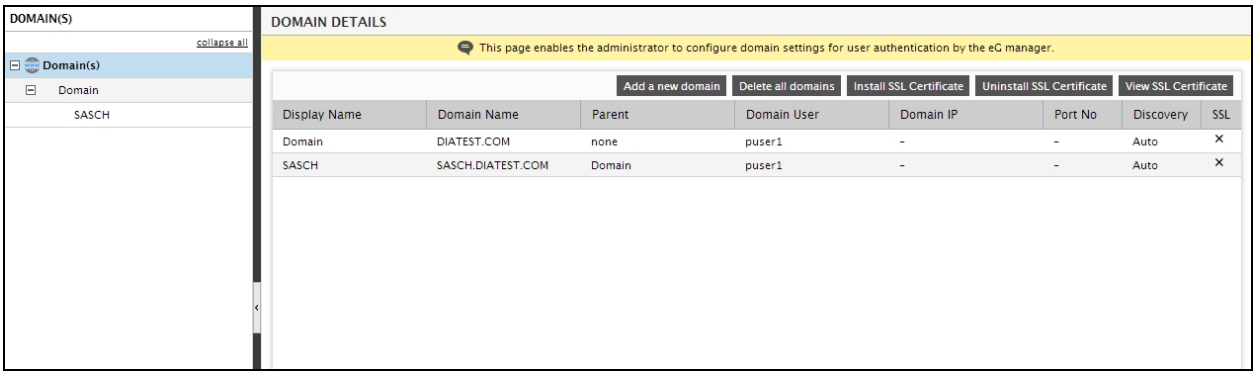

Figure 6.12: The details of both the parent and child domains

6. At any time, you can view the configuration of the auto-discovered parent domain by just clicking on the node representing that domain in the **DOMAIN(S)** tree. The right panel will then change to display the parent domain's current configuration (see Figure 6.13).

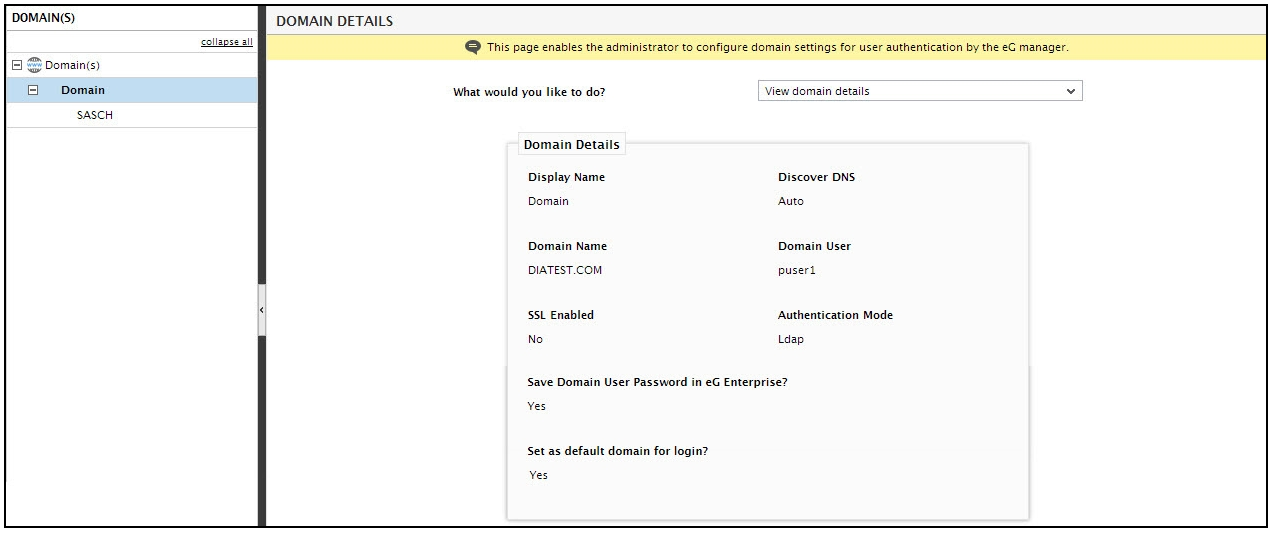

Figure 6.13: Viewing the parent domain's configuration

7. To modify the configuration of a parent domain, just select the **Modify domain details** option from the **What**

would you like to do? list in Figure 6.13.

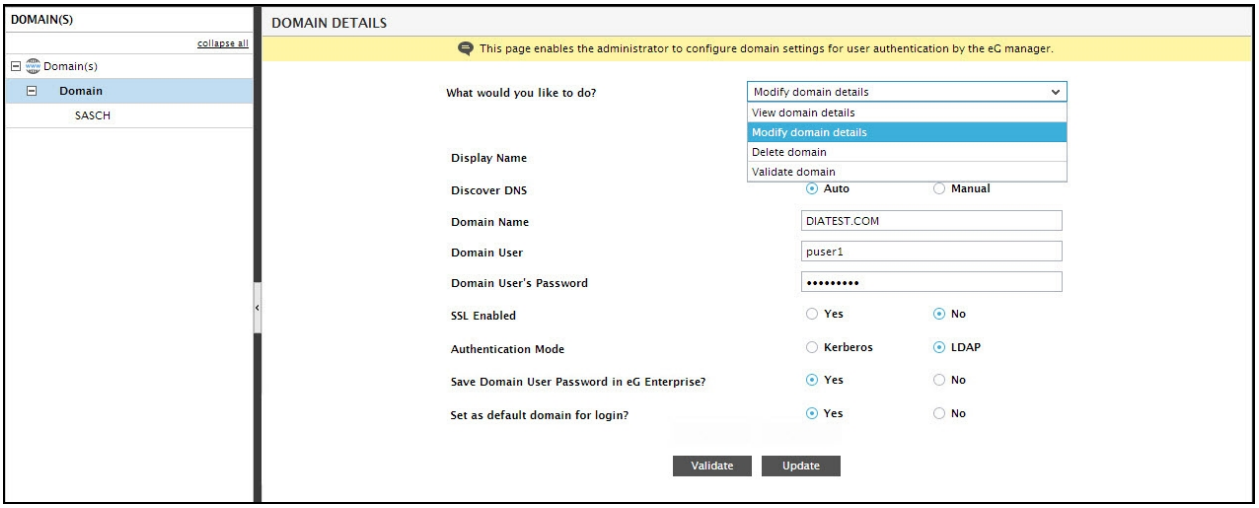

Figure 6.14: Figure 6. 14: Selecting the Modify option of an auto-discovered parent domain

8. The right panel will once again change to display the parent domain's current configuration, but in an editable mode (see Figure 6.15).

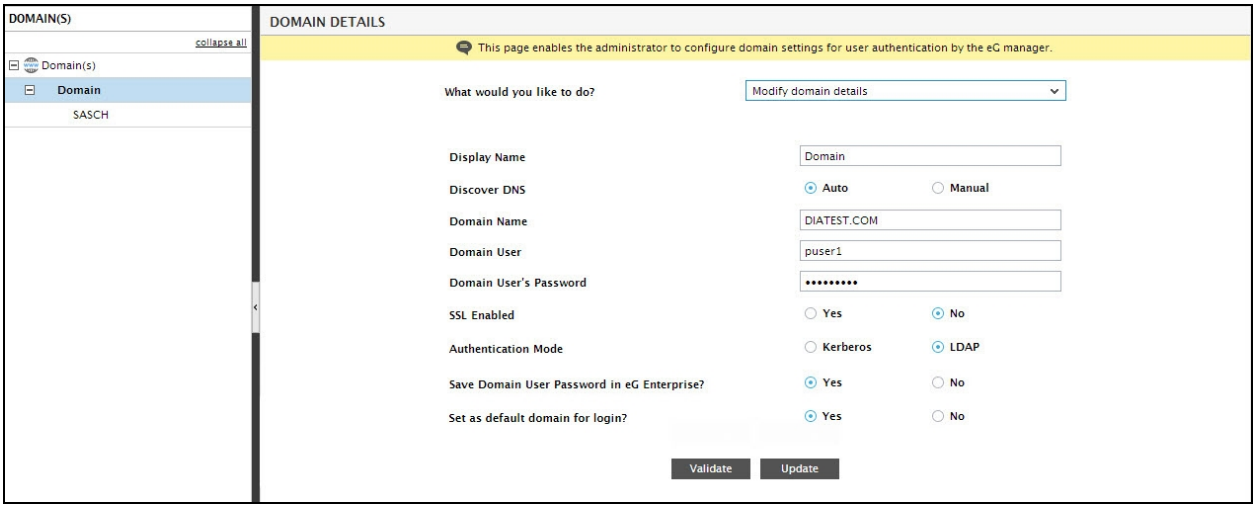

Figure 6.15: Modifying the parent domain's configuration

9. Except the **Display Name**, all other details of the parent domain can be modified. Once you are done with your changes, click the **Update** button in the right panel of to save the changes.

   Note:

   Whenever the configuration of a parent / child domain is modified, make sure that you restart the eG manager.

10. To view the details of an auto-discovered child domain, click on the node representing the child domain in the tree structure. The details of the chosen node will be displayed in the right panel (see Figure 6.16).

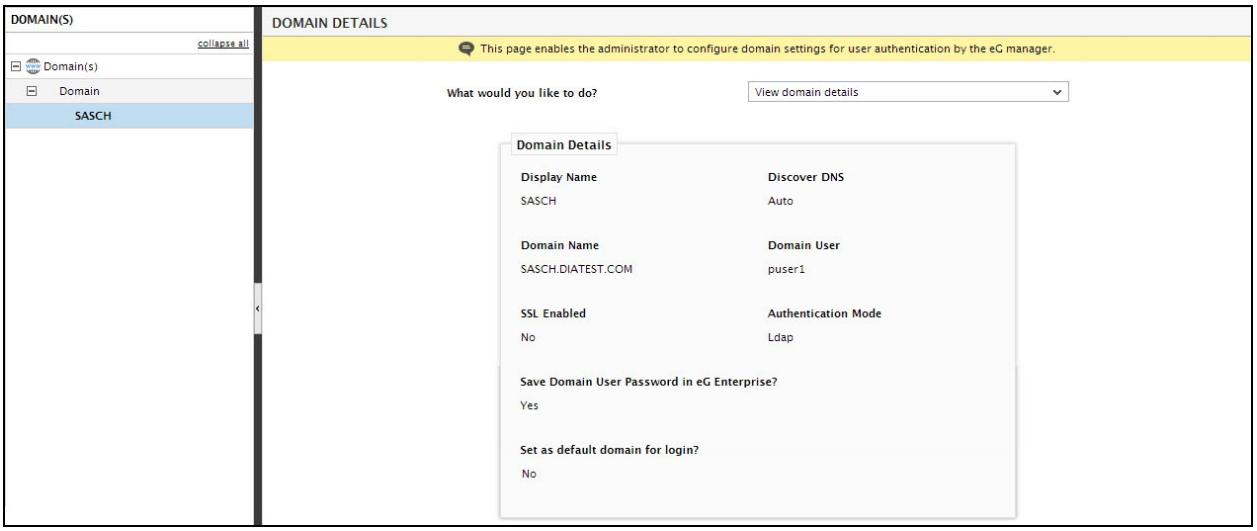

Figure 6.16: Selecting the option to view the details of the auto-discovered child domain

11. Unlike an auto-discovered parent domain, an auto-discovered child domain **cannot be modified**. However, you can delete an auto-discovered child domain. For this, just select the **Delete domain** option from the **What would you like to do?** list as depicted by 6.2.1.

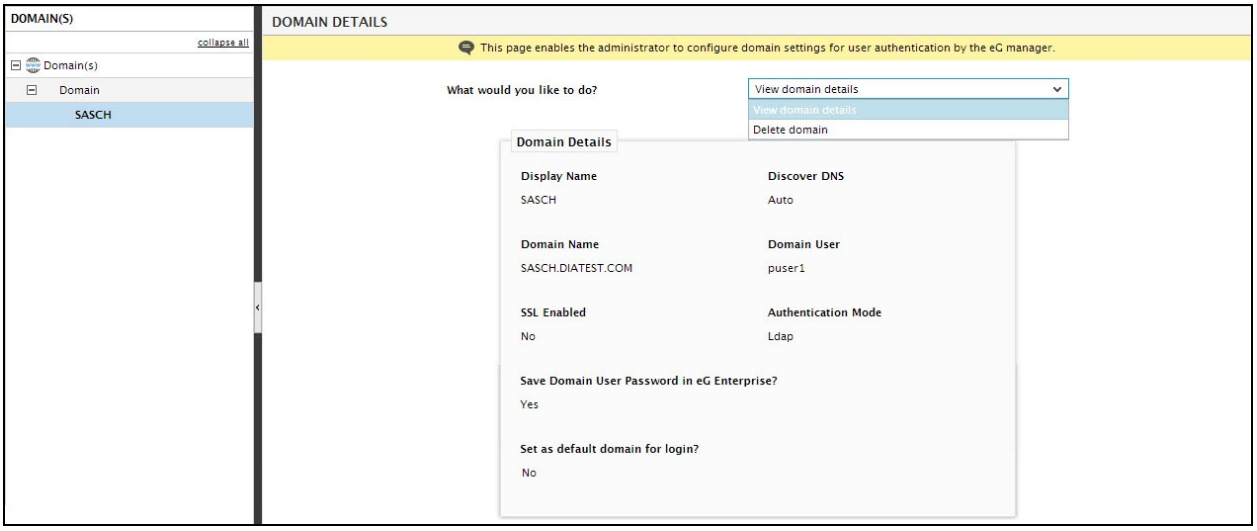

Figure 6.17: Selecting the Delete option of an auto-discovered child domain

12. Doing so will invoke the message box of Figure 6.18, which will request for your confirmation to delete the child domain. Click the **Delete** button in Figure 6.18 to confirm deletion.

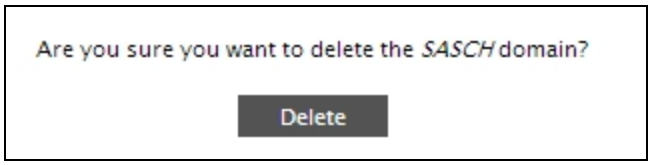

Figure 6.18: Deleting an auto-discovered child domain

13. While deleting a child domain will delete only that domain, deleting a parent domain will delete all its child

domains as well. Therefore, to delete a parent domain and all its child domains, first, click on the parent domain node in the **DOMAIN(S)** tree and select the **Delete** option from the **What would you like to do?** list. A message box requesting your confirmation to delete the parent domain will appear. Click the **Delete** button in the message box to confirm deletion.

Note:

- Ensure that the eG manager is restarted after deleting a domain.
- Deleting an auto-discovered parent domain automatically deletes all its discovered sub-domains as well.
- A domain can be deleted only if no user registered with eG belongs to that domain.

14. If you have chosen **Yes** against the **Discover User Details from AD**, then, an additional **Update AD User Details Now** option will appear in the **What would you like to do?** list.

6.2.2 Manually Configuring Parent and Child Domains

If you want to manage users spread across multiple domains, then, all domains, except the eG manager's domain (which can be auto-discovered), will have to be manually configured using the eG administrative interface.

Follow the steps below to manually add parent and child domains:

1. Click on the global **Domain(s)** node in the tree structure in the left panel, and then click the **Add a new domain** button in the right panel.
2. Figure 6.19 then appears displaying the parameters to be configured for creating a new domain.

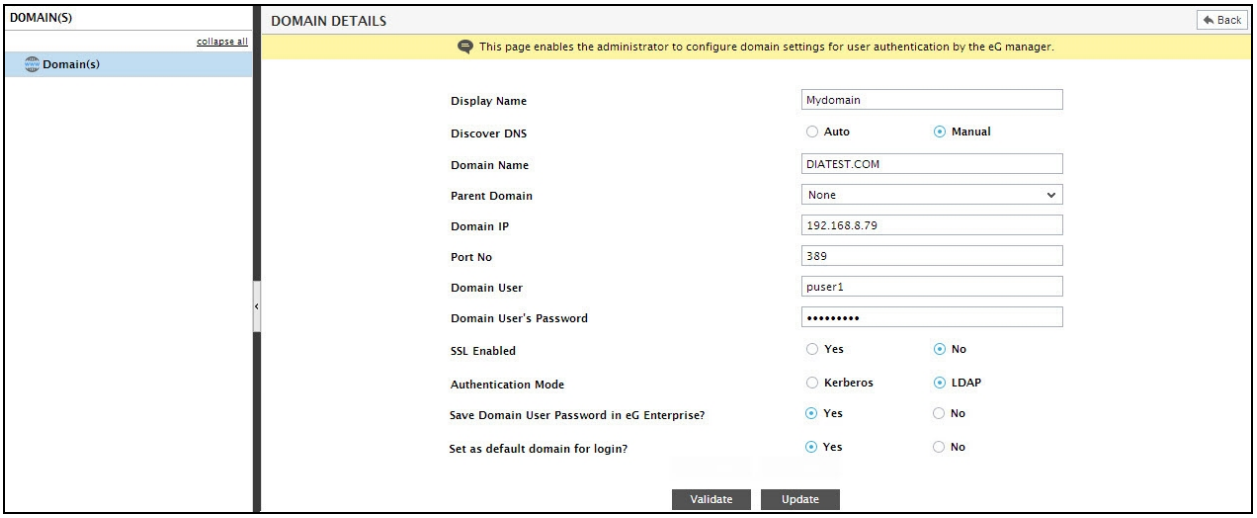

Figure 6.19: Manually configuring a domain

3. In the right panel of Figure 6.19, specify the following to create a parent domain:
   - Provide a **Display Name** for the new domain.
   - To manually configure the IP address and port number of the domain server, set the **Discover DNS** flag to **Manual**.

- Next, specify the fully-qualified **Domain Name**.

  Note:

  eG Enterprise disallows **Domain Name** duplication- i.e., you cannot assign the domain name of an existing parent/child domain to a new domain.

- If the domain has an alias name in the target environment, you can set the **Does domain have an alias?** flag t **Yes**. By default, this flag is set to **No**. If there exists an alias name, then, you can specify the other name of the domain in the **Domain Alias Name** text box.
- To add a parent domain, set the **Parent Domain** parameter to *None*.
- Since auto-discovery of DNS is disabled, you need to manually specify the **Domain IP** and **Port No** of the AD server.
- To connect to the AD server and access the domain user information stored within, the eG manager requires a domain user's privileges. To facilitate this connection, provide a valid domain user's name and password against **Domain User** and **Domain User's Password**.
- Then, indicate whether the AD server is SSL-enabled or not, by setting the **SSL** flag to **Yes** or **No**, as the case may be. If the **SSL** flag is set to **Yes**, then you will have to follow the procedure discussed in the **Appendix** below to ensure that the eG manager is able to communicate with the AD server over SSL.
- Next, indicate how accesses to the AD server are to be authenticated - using **Kerberos** or **LDAP**. **Kerberos** is a computer network authentication protocol which works on the basis of "tickets" to allow nodes communicating over a network to prove their identity to one another in a secure manner. **Kerberos** is ideal for AD environments with high security considerations. The **Lightweight Directory Access Protocol** on the other hand, is an application protocol for accessing and maintaining distributed directory information services over an Internet Protocol (IP) network. The **LDAP authentication mechanism** is best suited for environments with not very high security constraints.
- Next, indicate whether the **Domain User's Password** that you have provided here for enabling the eG manager to connect to the AD server, should be saved in eG Enterprise or not. To save the password, set the **Save Domain User Password in eG Enterprise?** flag to **Yes**. If this is done, then, the specified **Domain User's Password** will be automatically encrypted and saved to the **eg_authenticate.ini** file, which will be available in the **<EG_MANAGER_INSTALL_DIR>\manager\config** directory. On the other hand, if the **Save Domain User Password** flag is set to **No** instead, the password will not be saved to the **eg_authenticate.ini** file. If the password is not saved, then every time the eG manager attempts to connect to the AD server - say, when validating/registering domain user profiles configured on the eG manager (using the **ADD USER** page) with the AD server - you will be prompted for the **Domain User's Password**.
- Also, indicate whether/not the domain being configured should be set as the default domain at the time of login. To set the new domain as the default domain, set the **Set as default domain for login?** flag to **Yes**. If this is done, then the next time a user attempts to log into the eG management console by typing his/her user name in the login page, the **Domain** selection will instantly change from Local to the domain that you have set as the default. This capability is most useful in environments where the eG manager integrates with only one domain. By setting this domain as the default, administrators can save users the trouble of selecting a **Domain** every time he/she tries to login.

- Next, indicate whether you wish to discover and display specific details of a user (for e.g., location, address etc) who is part of the domain in the eG monitoring console. To discover and display the user information in the eG monitoring console, set the **Discover User Details from AD** flag to **Yes**. By default, this flag is set to **No**.
- Then, to verify the correctness of your specifications, click the **Validate** button. Figure 6.20 will then appear indicating whether/not the **Display Name, Domain Name**, **Domain IP**, **Port No**, **Domain User**, and **Domain User Password** values that you have provided are indeed valid.

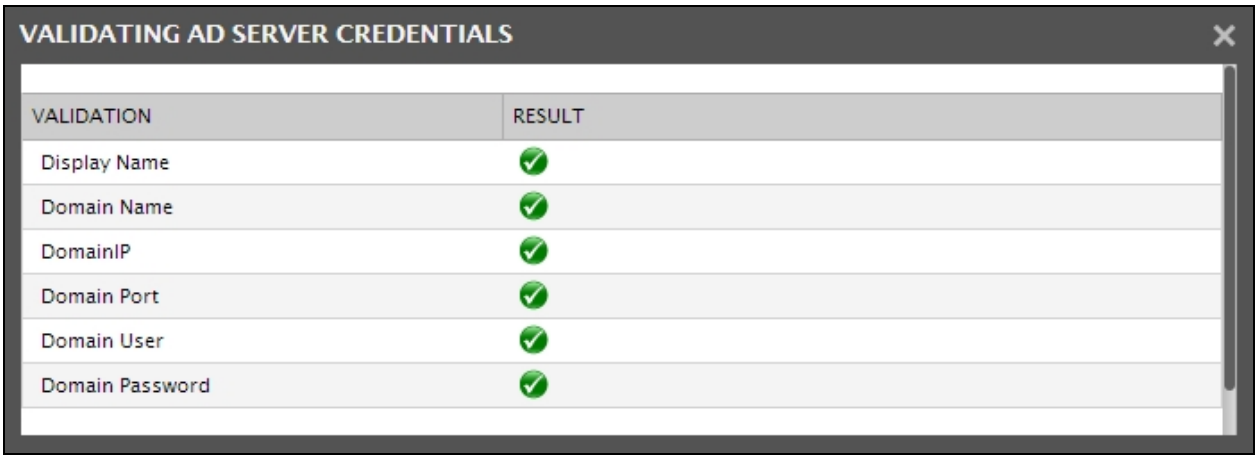

Figure 6.20: Validating the specifications of a domain that has been manually configured

- Click the **Update** button to add the new domain.
- Once the parent domain is added to the eG Enterprise system, a message to that effect will appear (see Figure 6.21). The tree in the left panel will also change to reflect the addition of the parent domain.

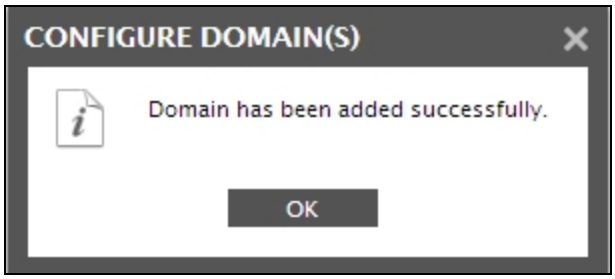

Figure 6.21: The tree structure indicating that another parent domain has been added

4. A parent domain that is created manually can be viewed or modified the same way as an auto-discovered parent domain. Therefore, follow the procedure described in steps 5 – 9 in Section **6.2.1** above to know how to view/edit the details of a parent domain. To know how to delete a manually configured parent domain, follow the same procedure described in step 13 of Section **6.2.1** above.
5. For all parent domains that are created manually, sub-domains also need to be manually created. To do so, follow the steps given below:
   - Right-click on the node representing the manually configured parent domain in the **DOMAIN(S)** tree, and pick the **Add Sub-domain** option from the **What would you like to do?** list in the right panel (see

Figure 6.22) .

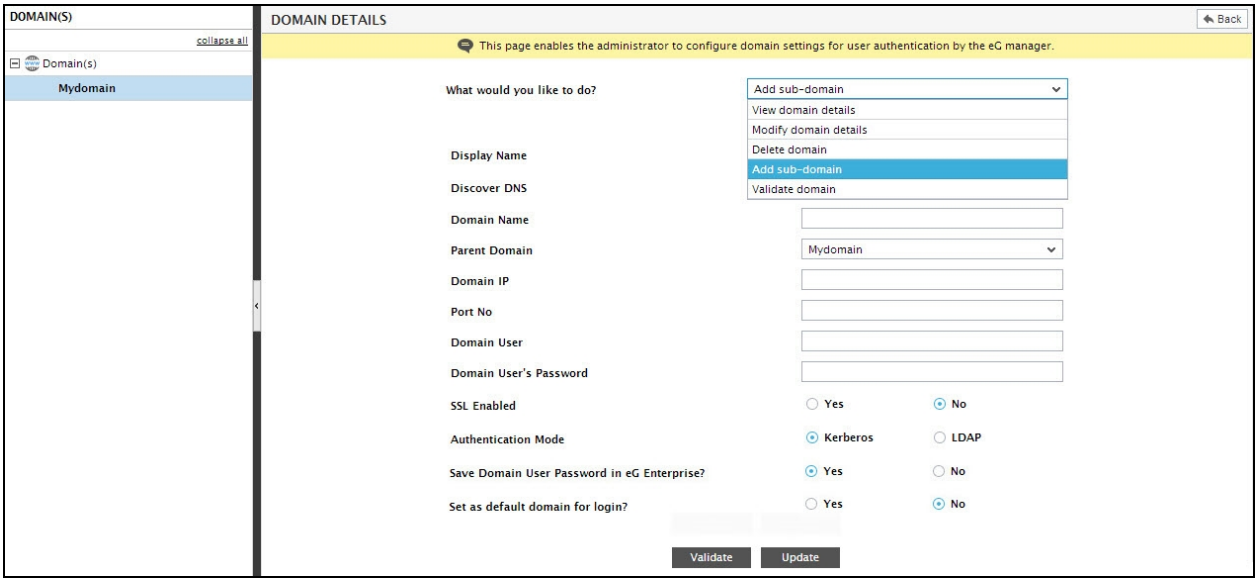

Figure 6.22: Selecting the 'Add Sub-domain' option

- When Figure 6.23 appears, first provide a **Display Name** for the sub-domain.

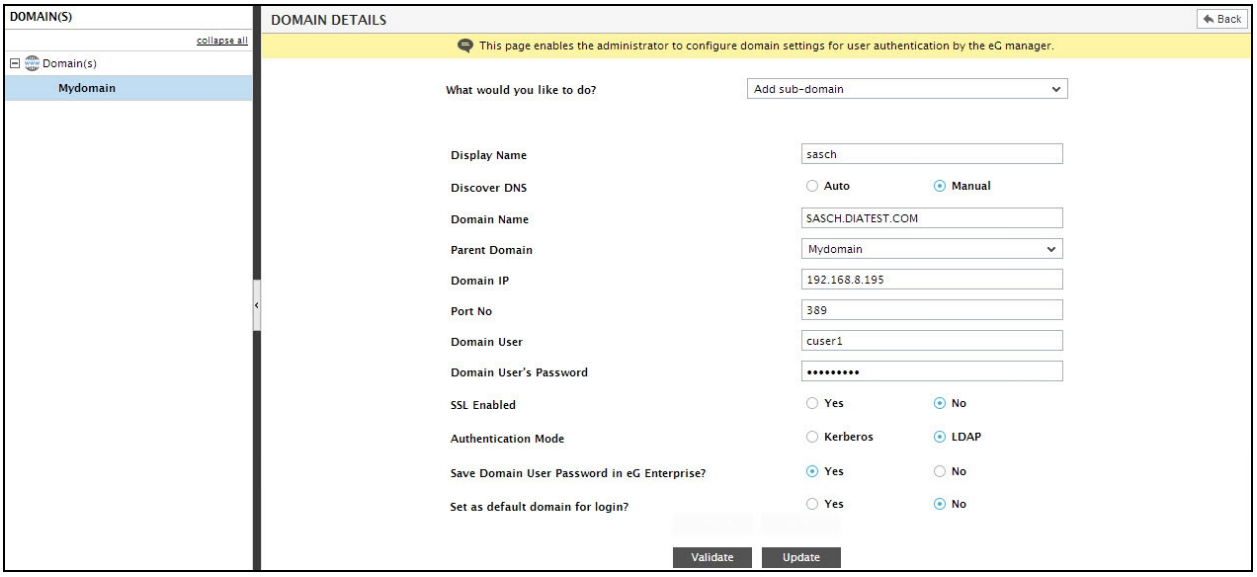

Figure 6.23: Manually adding a child domain

- Set the **Discover DNS** flag to **Manual**.

  Note:

  Note that if a parent domain is configured manually, then its sub-domains cannot be auto-discovered - i.e., you should not set the **Discover DNS** flag to **Auto** while configuring such a sub-domain.

- Provide the fully-qualified **Domain Name**.

- Next, from the **Parent Domain** list, select the parent domain under which this sub-domain is to be created.
- Since auto-discovery of DNS is disabled, you need to manually specify the **Domain IP** and **Port No** of the AD server.
- To connect to the AD server and access the domain user information stored within, the eG manager requires a domain user's privileges. To faciliate this connection, provide a valid domain user's name and password against **Domain User** and **Domain User Password**.
- Then, indicate whether the AD server is SSL-enabled or not, by setting the **SSL** flag to **Yes** or **No**, as the case may be. If the **SSL** flag is set to **Yes**, then you will have to follow the procedure discussed in the **Appendix** below to ensure that the eG manager is able to communicate with the AD server over SSL.
- Next, indicate how accesses to the AD server are to be authenticated - using **Kerberos** or **LDAP**. **Kerberos** is a computer network authentication protocol which works on the basis of "tickets" to allow nodes communicating over a network to prove their identity to one another in a secure manner. **Kerberos** is ideal for AD environments with high security considerations. The **Lightweight Directory Access Protocol** on the other hand, is an application protocol for accessing and maintaining distributed directory information services over an Internet Protocol (IP) network. The **LDAP authentication mechanism** is best suited for environments with not very high security constraints.
- Next, indicate whether the **Domain User Password** that you have provided here for enabling the eG manager to connect to the AD server, should be saved or not. To save the password, set the **Save Domain User Password in eG Enterprise?** flag to **Yes**. If this is done, then, the specified **Domain User Password** will be automatically encrypted and saved to the **eg_authenticate.ini** file, which will be available in the **<EG_MANAGER_INSTALL_DIR>\manager\config** directory. On the other hand, if the **Save Domain User Password in eG Enterprise?** flag is set to **No** instead, the password will not be saved to the **eg_authenticate.ini** file. If the password is not saved, then every time the eG manager attempts to connect to the AD server - say, when validating the domain configuration using the eG manager or when validating/registering domain user profiles configured on the eG manager (using the **ADD USER** page) with the AD server - you will be prompted for the **Domain User Password**.
- Also, indicate whether/not the domain being configured should be set as the default domain at the time of login. To set the new domain as the default domain, set the **Set as default domain for login?** flag to **Yes**. If this is done, then the next time a user attempts to log into the eG management console by typing his/her user name in the login page, the **Domain** selection will instantly change from Local to the domain that you have set as the default. This capability is most useful in environments where the eG manager integrates with only one domain. By setting this domain as the default, administrators can save users the trouble of selecting a **Domain** every time he/she tries to login.
- In virtual environments where **LDAP** is used to authenticate access to the AD server, administrators may want to keep track on specific user information for e.g., location, vendor etc of the users accessing their environment through the AD server. For example, in addition to viewing the user experience with their virtual environment, if administrators are able to view the location of the user, it would help them troubleshoot location specific issues at the earliest. This approach would definitely help administrators improve the overall performance of their environment. To view such user specific information in the eG monitoring console, administrators should do the following:

- First, set the **Discover User Details from AD** flag to **Yes**. By default, this flag is set to **No**.

  Once this flag is set to **Yes**, the user specific information will automatically be populated in the **ADUserDetails.ini** file that is located in the *<eG_INSTALL_DIR>/manager/config* location.

  If the **Discover User Details from AD** flag is set to **Yes**, then an additional **Update User Details from AD** option will appear in the **What would you like to do?** list in the right panel as shown in Figure 10. Clicking the **Update** button will immediately integrate the user information from the domain to the **ADUserDetails.ini** file available in the **<EG_INSTALL_DIR>\manager\config** directory.

  Note:

  By default, the user information available in the domain will be integrated with the **ADUserDetails.ini** file once in 7 days. If you wish to override this default, setting, then you can do the following:

  - Edit the **eg_services.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory).
  - Set the **ThreadFrequency** parameter in the **[ADUSERDETAILS_THREAD_SETTINGS]** section of the file to a frequency of your choice.
  - By default, the information will be integrated every **Sunday**. If you wish to override this default day, then you can change the **DayToRun** parameter to the day of your choice.
  - Save the file.

  The user specific information so updated can be viewed in the eG monitoring console in the following features offered by the eG Enterprise Suite:

  - User Experience Dashboard
  - Current Alarms
  - Layer model page of the tests where users are the descriptors of the tests

- Then, to verify the correctness of your specifications, click the **Validate** button.
- Once the specifications are validated, click the **Update** button to add the new domain.
- Once a sub-domain is manually added, a message to that effect will appear. The tree will also change to reflect the addition of the sub-domain.

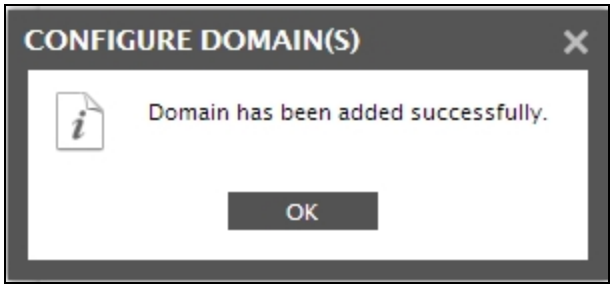

Figure 6.24: A message indicating that the sub-domain has been successfully created

6. Similarly, you can add multiple child domains to a parent domain. In fact, you can even add sub-domains to a child domain.
7. Also, unlike an auto-discovered sub-domain where changes cannot be made to domain details, you can modify a manually-configured sub-domain. For this, select the node representing the sub-domain from the **DOMAIN(S)** tree, and choose the **Modify domain details** option from the **What would you like to do?** list in the right panel.

   Note:

   Whenever the configuration of a parent / child domain is modified or deleted, make sure that you restart the eG manager.
8. Also, you have the option of deleting a sub-domain. For this, select the sub-domain node from the **DOMAIN (S)** tree, and pick the **Delete domain** option from the **What would you like to do?** list (see Figure 6.25).

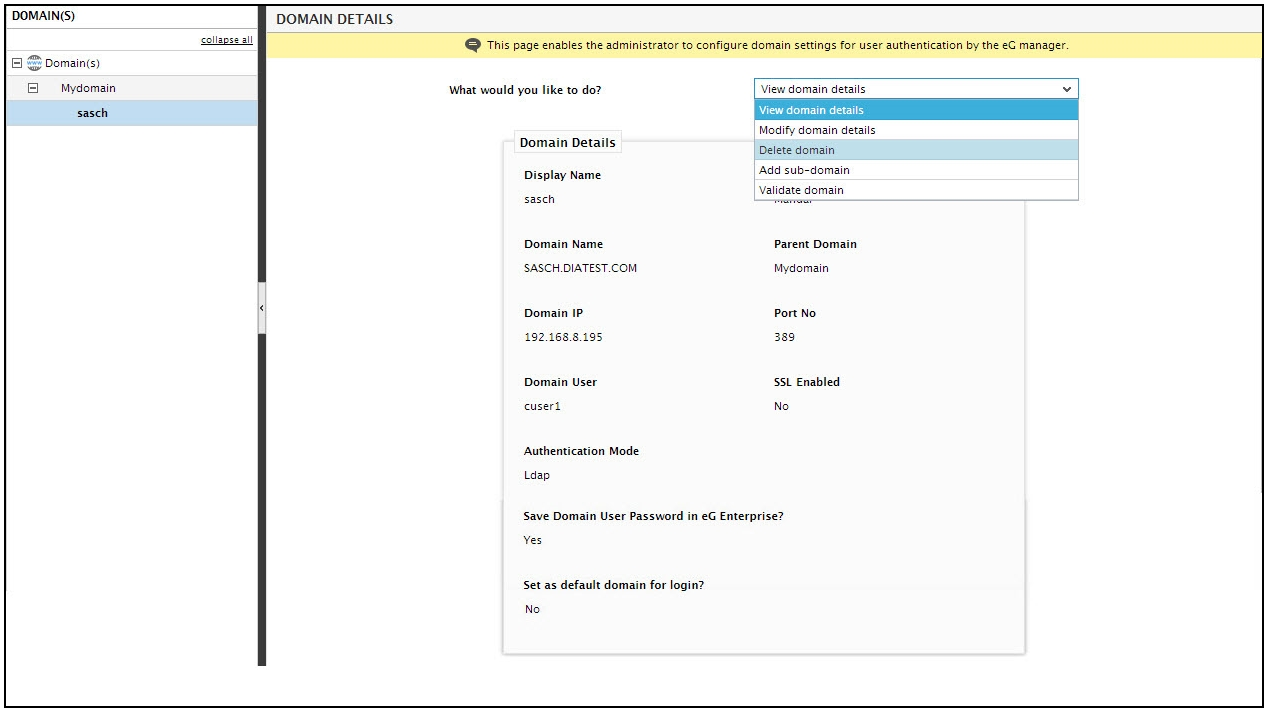

Figure 6.25: The menu list displaying the Delete option

9. A message box depicted by Figure 6.26 will appear requesting your confirmation to delete the chosen sub-domain. Click the **OK** button to proceed with the deletion.

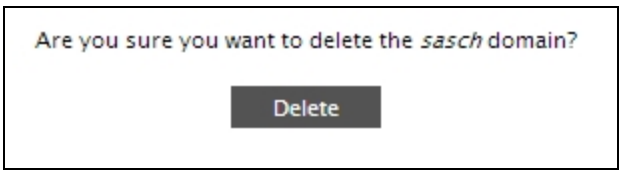

Figure 6.26: Confirming the deletion of a manually created sub-domain

Note:

Ensure that the eG manager is restarted after deleting a parent/child domain.

10. To simply view the names and current configuration of the parent and child domains that have been created using the eG administrative interface, just click on the global **Domain(s)** node in the tree-structure. Figure 6.27 will appear.

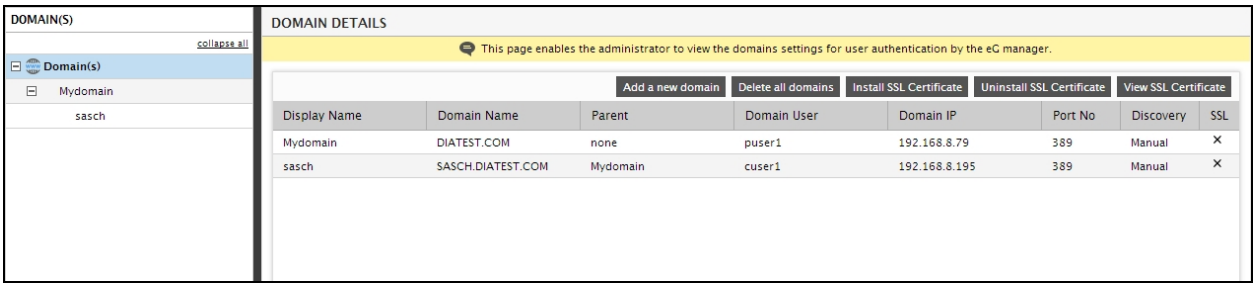

Figure 6.27: Viewing the names and current configuration of all domains

11. You can delete all displayed domains at one shot by simply clicking on the **Delete all domains** button in the right panel of Figure 6.27.

    Note:

    Discovery of AD and KDCs is an on-going process - a configurable time period is used to determine for how long discovered AD/KDC information is cached by the eG manager. The default period is 15 minutes. To override this default setting, do the following:

    - Edit the **eg_services.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory)
    - Set the **ADRediscovery** parameter in the **[MISC_ARGS]** section of the file to a duration (in minutes) of your choice.
    - Save the file

12. If the **Discover User Details from AD** flag is set to **Yes**, then an additional **Update User Details from AD** option will appear in the **What would you like to do?** list in the right panel as shown in Figure 10. Clicking the **Update** button will immediately integrate the user information from the domain to the **ADUserDetails.ini** file available in the **<EG_INSTALL_DIR>\manager\config** directory.

    By default, the user information available in the domain will be integrated with the **ADUserDetails.ini** file once in 7 days. If you wish to override this default, setting, then you can do the following:

    - Edit the **eg_services.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory)
    - Set the **ThreadFrequency** parameter in the **[ADUSERDETAILS_THREAD_SETTINGS]** section of the file to a frequency of your choice.
    - By default, the information will be integrated on every **Sunday**. If you wish to override this default, day, then you can change the **DayToRun** parameter to the day of your choice.
    - Save the file.

6.2.3 Validating Parent/Child Domain Configuration

As demonstrated already, the eG Enterprise system provides administrators with a **Validate** option that helps them check the correctness of the domain configuration when creating that domain, and enables them to make changes to the configuration on-the-fly. This way, the solution prevents the creation of domains with incorrect/invalid details!

Sometimes however, as part of a routine maintenance exercise or owing to a policy requirement, administrators may make some significant changes in the AD environment post the eG manager-AD integration - for example, the domain name can be changed, the domain can be migrated to another server with a different IP address, the login names of domain users can be modified, and so on. Some changes may also occur inadvertently - for instance, a user account may expire or may be locked out, network connection between the eG manager and the AD server could become flaky, etc. Such changes are bound to affect the AD-eG manager integration, causing issues in manager accesses to the AD server, domain user registration with the eG Enterprise system, and even user logins. Therefore, when users to the eG Enterprise system complaint of issues related to the integration, administrators need to rapidly initiate investigations in order to diagnose the reason for this occurrence.

To facilitate this preliminary prognosis, the eG administrative interface provides the **Validate a domain** option. Using this option, administrators can quickly check a registered domain's accessibility and the correctness of the connection details provided at the time of domain configuration, from the eG manager itself. In addition, with the help of this page, administrators can quickly view the user groups that are available in a domain and even check the validity of domain user accounts that they intend to add to the eG Enterprise system, without having to physically login to the AD server.

To use this page, do the following:

1. Select the domain that needs to be investigated from the **DOMAIN(S)** tree, and pick the **Validate domain** option from the **What would you like to do?** list in the right panel (see Figure 6.28).

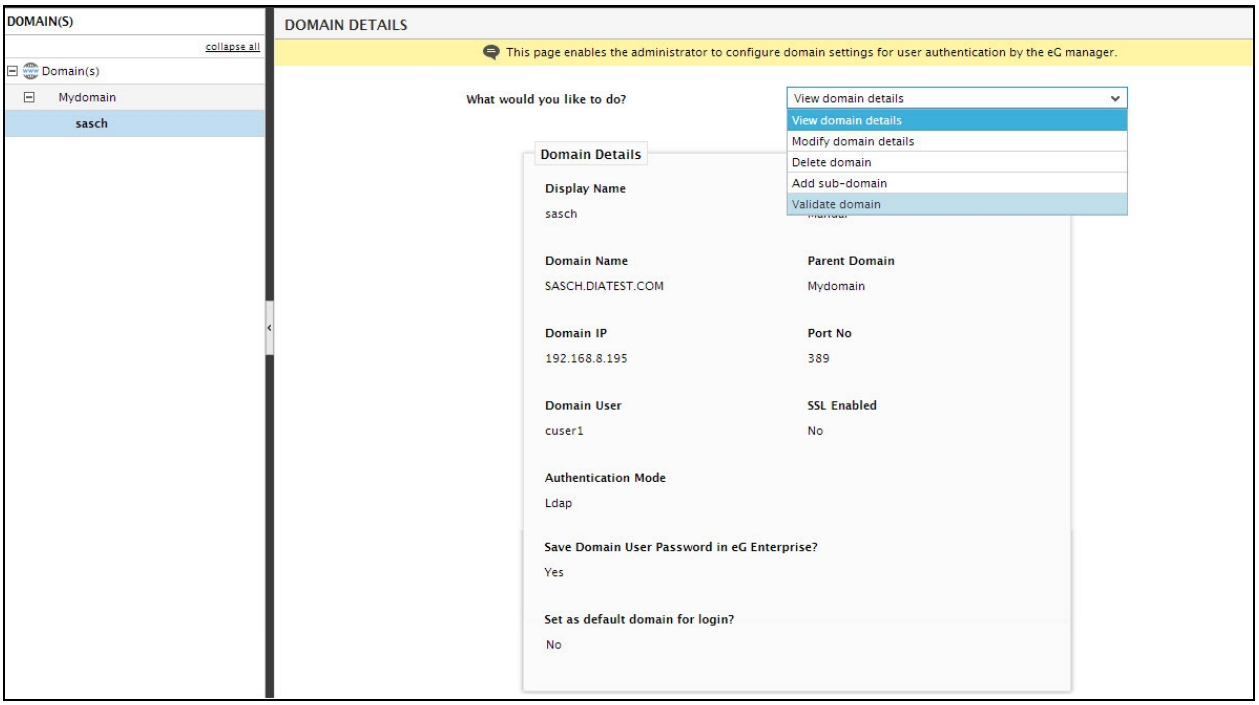

Figure 6.28: Selecting the Validate option from a domain's right-click menu

2. Figure 6.29 will then appear. First, from the **What would you like to validate?** drop-down, select the option that indicates what is that you wish to validate. By default, the **Is this domain reachable?** option is chosen from this list. When users complaint that they are unable to connect to a domain, then, you can select this option to verify whether the domain name that you had provided at the time of domain configuration is still valid or not. If an auto-discovered domain is chosen for validation from the tree-structure, then, selecting the **Is this domain reachable?** option displays the **Display Name** and the fully-configured **Domain Name** of the selected domain. On the other hand, if a manually added domain is chosen for validation from the tree-structure, then, selecting the **Is this domain reachable?** option displays the **Display Name,** the fully-configured **Domain Name**, the **Domain IP**, and the **Port No** of that domain. Click the **Connect** button to check whether the displayed domain is reachable or not. If the domain is reachable, then a message to that effect will appear. If not, then the reasons for the failure will also be indicated.

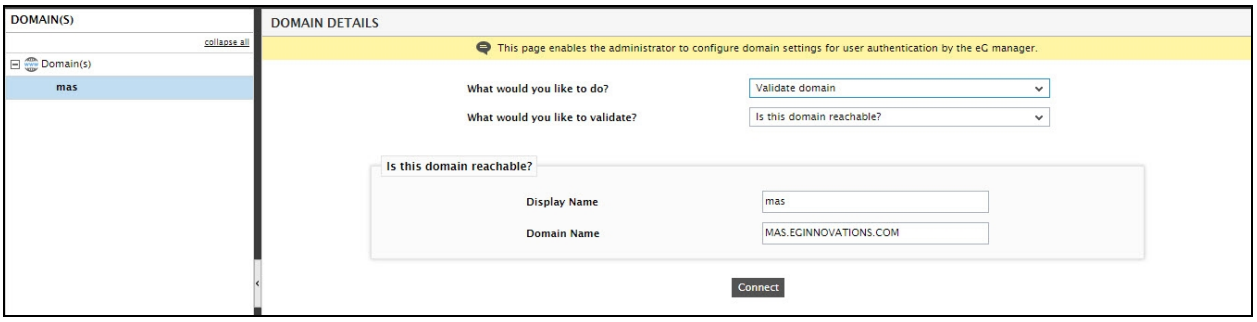

Figure 6.29: Checking whether or not an auto-discovered domain is reachable

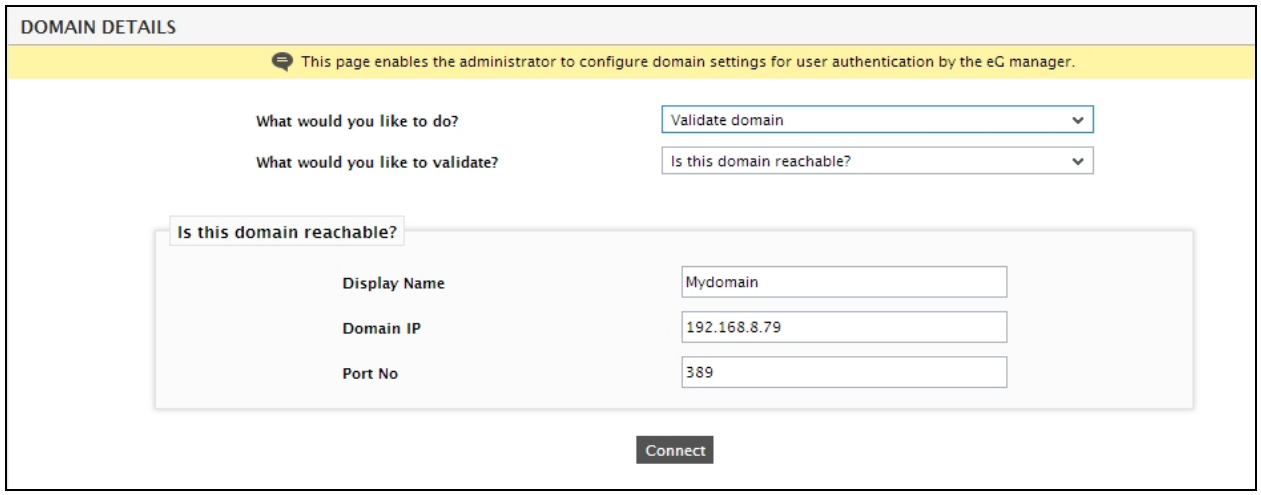

Figure 6.30: Checking whether or not a manually added domain is reachable

3. Sometimes, a domain may be reachable over the network – i.e., the reachability check performed at step 2 above may return positive results – but, one/more users in that domain may still not be able to login to the eG management console. One of the probable reasons for this could be a problem in communication between the eG manager and certain domain controllers configured in specific sites in that domain. To troubleshoot such issues in communication, an administrator can choose the **Troubleshoot communication with domain controller** option from the **What would you like to validate?** list. Once this is done, the **Display Name** and fully-qualified **Domain Name** of the chosen domain will be displayed (see Figure 6.31).

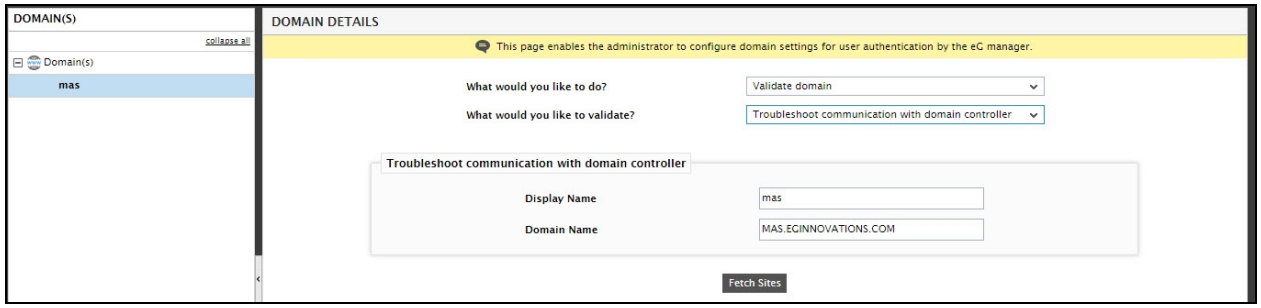

Figure 6.31: Troubleshooting communication with domain controller

4. Click on the **Fetch Sites** button in Figure 6.31 to know which sites are operating in that domain. This will bring up a **Select an Active Directory Site** drop-down, which will be automatically populated with the sites configured in the domain (see Figure 6.32).

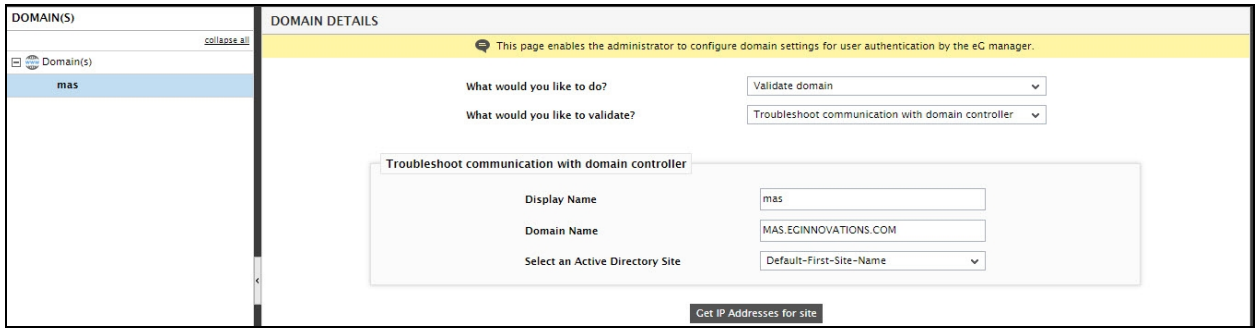

Figure 6.32: Fetching AD sites configured in a domain

5. To know which domain controllers are configured in a site and to verify communication with each domain controller, select a site from the **Selet an Active Directory site** list and click the **Get IP Addresses for site** button (see Figure 6.32). The **IP Address, Host Name**, and **Port** of each domain controller operating within the selected site will then be displayed in a table, as depicted by Figure 6.33.

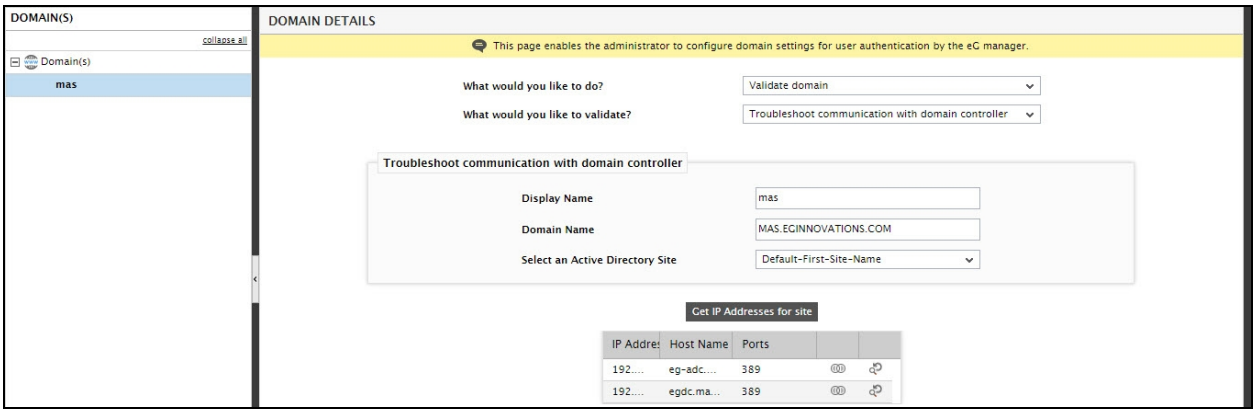

Figure 6.33: Getting IP address of controllers in site

6. To know if the eG manager is able to establish a socket connection with the IP address of a domain controller in the table, click the **Bind** icon corresponding to that domain controller. If the manager is able to communicate successfully, then a message box shown by Figure 6.34 will appear confirming the same. If the bind is unsuccessful, it is indicative of an issue in communication between the eG manager and the domain controller.

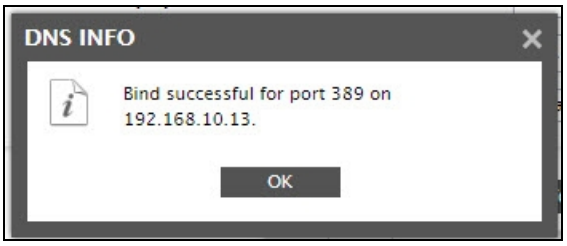

Figure 6.34: A message stating that the Binding was successful

7. Likewise, you can click on the **Reverse Lookup** icon corresponding to a domain controller in the table to check whether/not the DNS server is able to correctly resolve the host name of the controller to its IP address. If this lookup is successful, then a message box shown by 6.2.3 will appear confirming the same. If the reverse lookup is unsuccessful, it could mean that an improper DNS configuration could be the reason behind the communication issue between the eG manager and the domain controller.

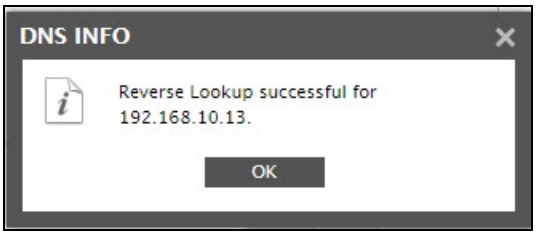

Figure 6.35: A message stating the Reverse lookup was successful

8. If the password of the **Domain User** is changed post domain configuration, then the eG manager will no longer be able to connect to the AD server for creating/validating domain user logins. If users complaint, then administrators can select the **Is this domain connection credential valid?** option from the **What would you like to validate?** list to verify the validity of the **Domain User Password**. Soon after selecting this option, the **Domain User** and **Domain User's Password** will be displayed. Click the **Validate** button in Figure 6.36 to check validity. The resulting message will indicate whether the displayed connection credentials are valid or not.

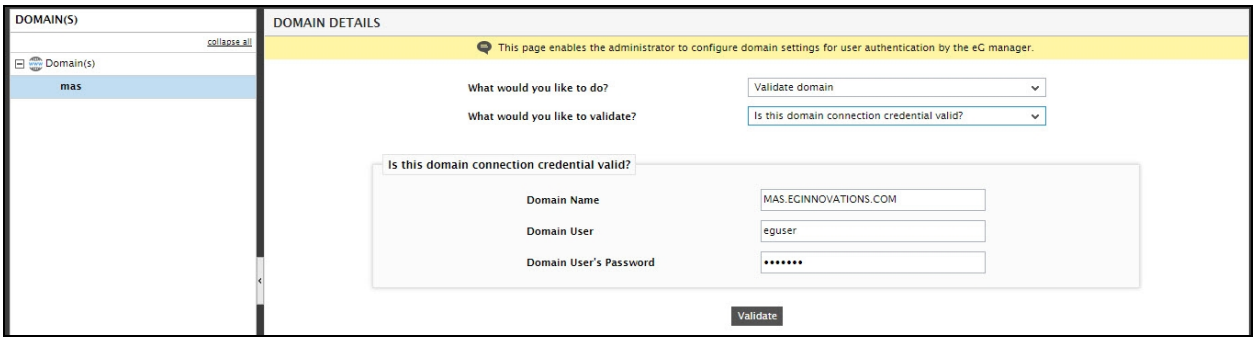

Figure 6.36: Checking the validity of the domain connection credentials

9. Before attempting to register a domain user with the eG Enterprise system, you may want to check whether the user really exists in that domain. For this, select the **Does the user exist in this domain?** option from the **What would you like to validate?** list. Upon selecting this option, the chosen **Domain Name** will be displayed. Enter the name of the user who needs to be checked in the **User Name** text box. Finally, click the **Validate** button. The resulting message will indicate whether the user exists in the domain or not, and if not, suggests a solution for the same (see Figure 6.37).

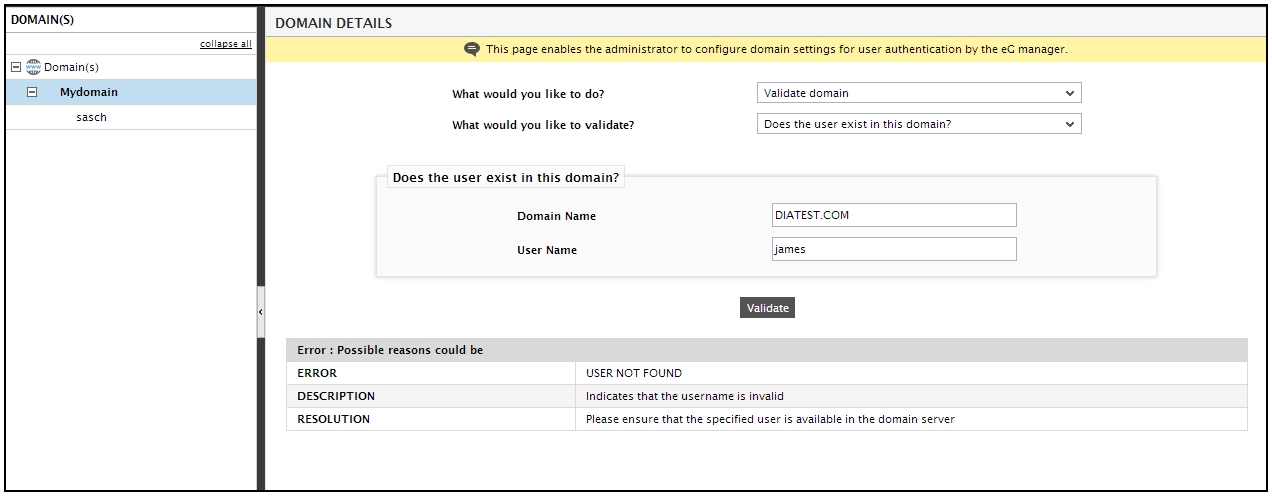

Figure 6.37: Checking whether the user exists in the domain or not

10. Domain user logins to the eG Enterprise system may also fail if one of the following is/has become invalid:
    - Domain name
    - User name

- User password

To know which one of the above parameters is invalid, select the **Is the user able to login to domain?** option from the **What would you like to validate?** list. Once the chosen **Domain Name** is displayed, enter the login credentials of the user to be verified, and click the **Login** button. The resulting message indicates whether the login was successful or not.

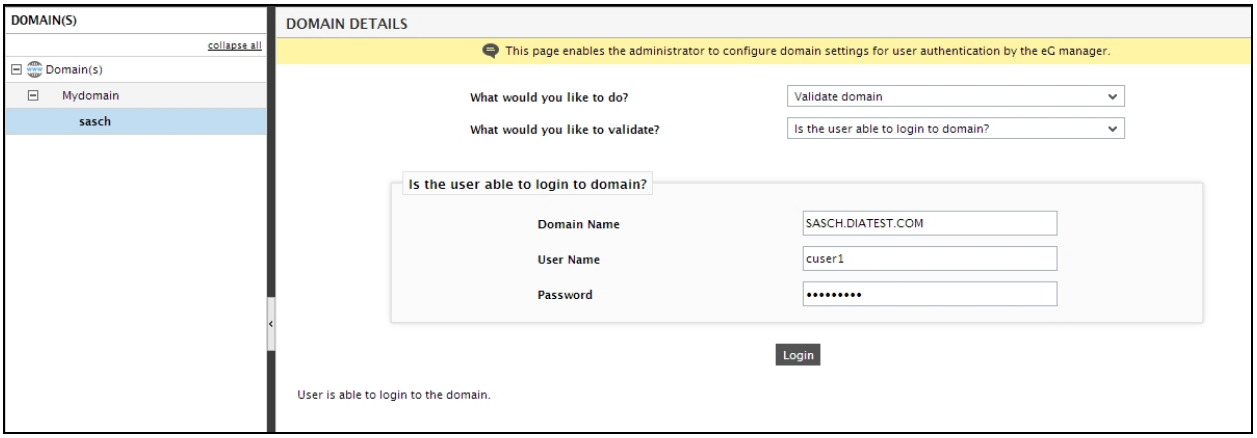

Figure 6.38: Checking whether the user is able to login to the domain

11. The first step to registering a domain user group with the eG Enterprise system is finding which user groups exist in the domain. For this, select the **Enumerate domain user groups** option from the **What would you like to validate?** drop-down list and click the **Enumerate** button. All user groups available in the chosen domain will then be listed (see Figure 6.39).

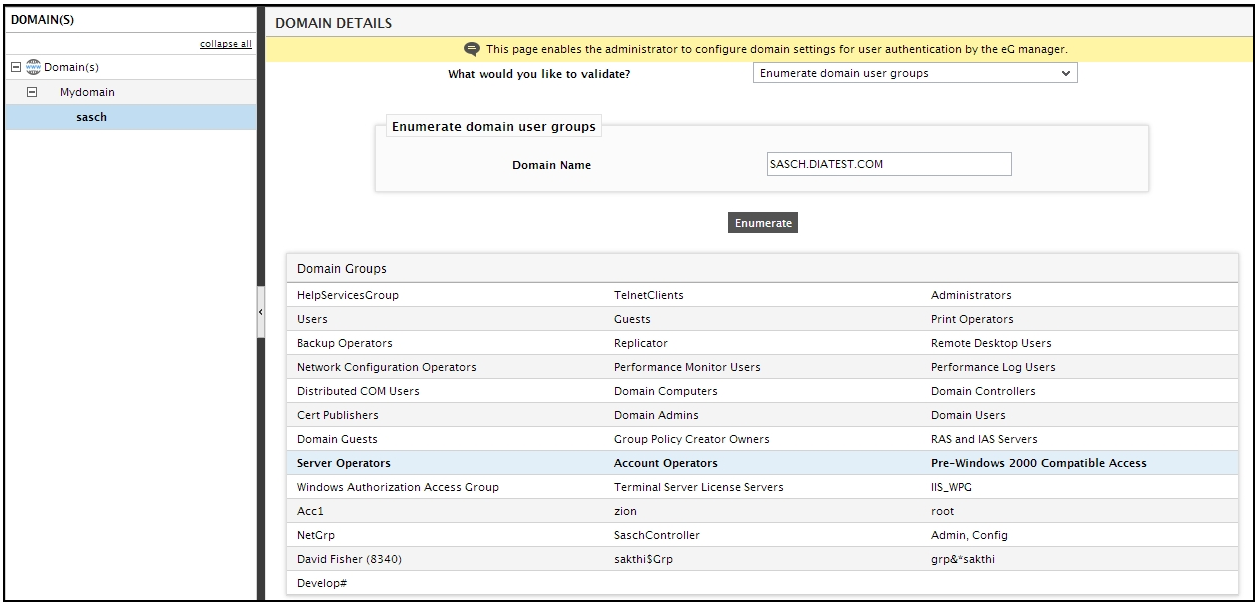

Figure 6.39: Enumerating domain user groups

6.2.4 SSL-Enabling the eG Manager and AD Communication

If the AD server with which the eG manager integrates is SSL-enabled, then before attempting the integration, you will have to SSL-enable the eG manager and AD communication. The broad steps in this process are as follows:

- Copy the SSL certificate of the AD server to the eG manager host,
- Import the certificate to the eG manager.

The sub-sections that follow will discuss each of the steps above elaborately.

6.2.4.1 Copying the SSL Certificate of the AD Server to the eG Manager Host

To achieve this, follow the instructions furnished below:

1. Login to any Windows host in the domain.
2. Follow the menu sequence, *Start -> Run*, and enter **mmc** in the **Run** text box (see Figure 6.40).

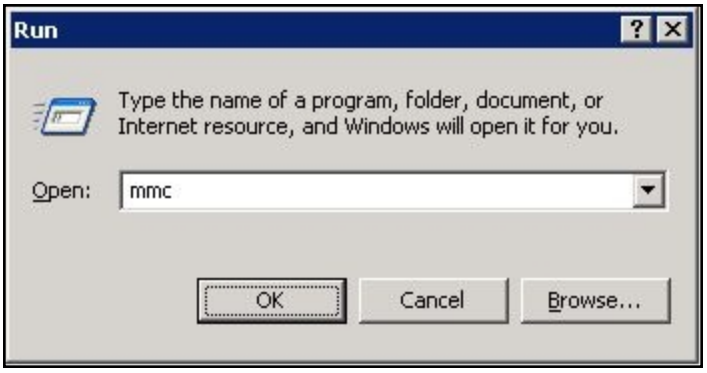

Figure 6.40: Executing mmc

3. A snap-in **Console** will then appear (see Figure 6.41).

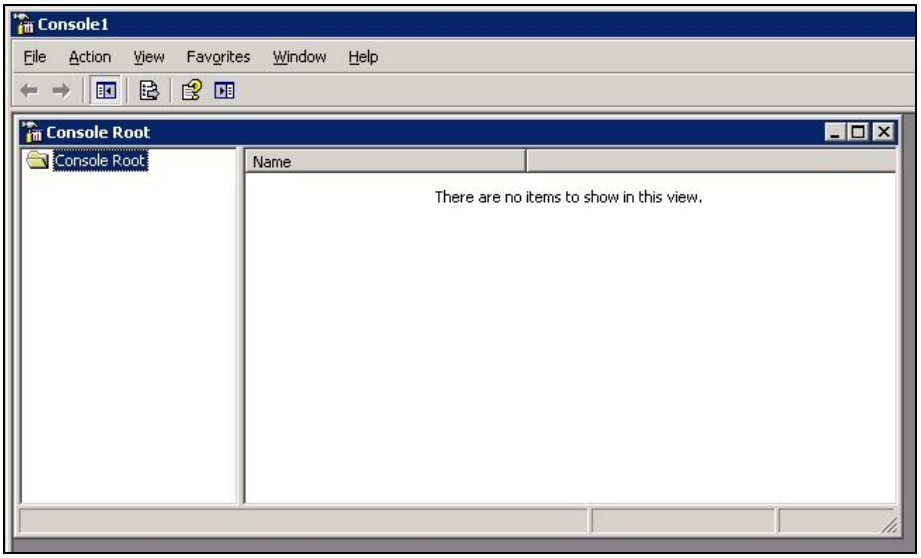

Figure 6.41: The Snap-in Console

4. Follow the *File -> Add/Remove Snap-in* menu sequence as depicted by Figure 6.42.

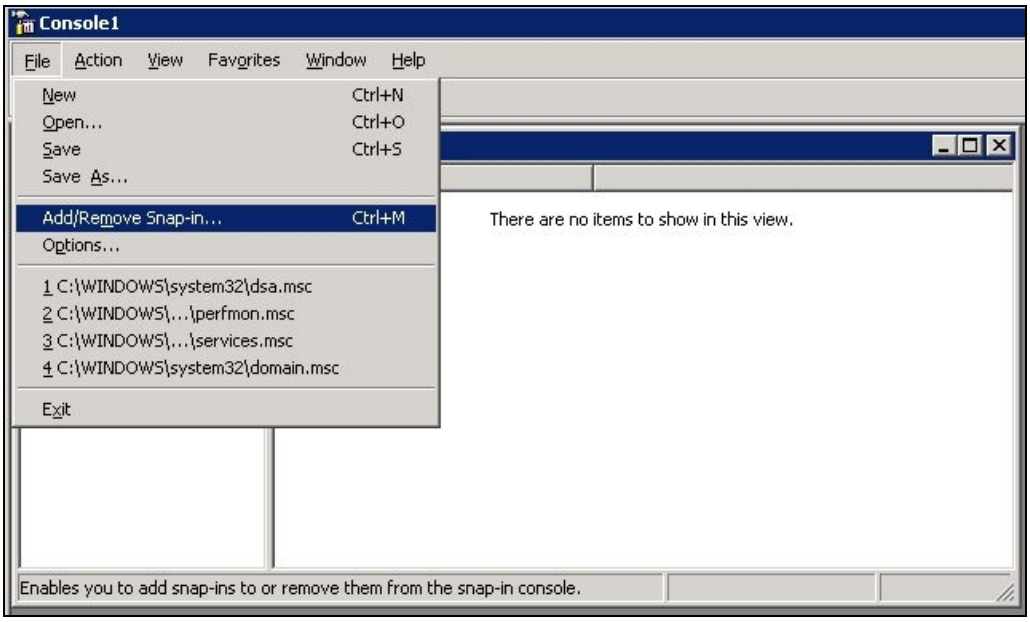

Figure 6.42: Selecting the Add/Remove Snap-in option

5. Figure 6.43 will then appear. Click the **Add** button in Figure 6.43 to add a snap-in.

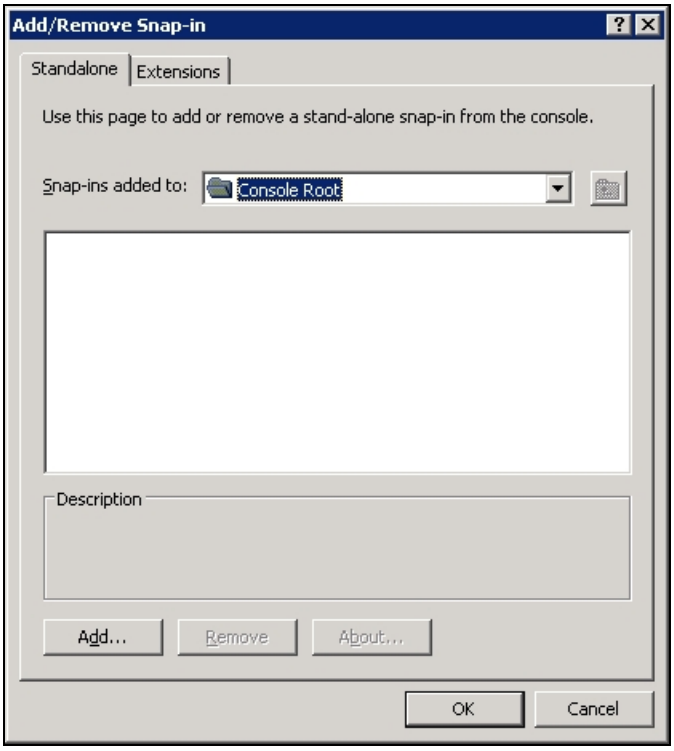

Figure 6.43: Clicking on the Add button

6. Figure 6.44 will then appear displaying the list of standalone snap-ins. Select the **Certificates** option from the **Available standalone snap-ins** list, and click the **Add** button in Figure 6.44.

Figure 6.44: Selecting the Certificates option

7. This will invoke Figure 6.45 from which you need to select the **Computer account** option. Then, click the **Next** button to move on.

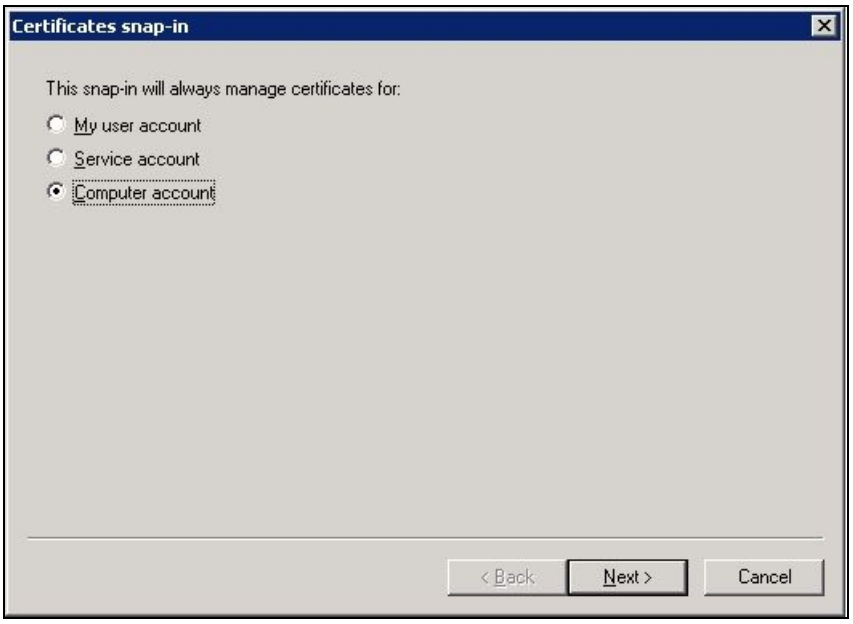

Figure 6.45: Selecting the Computer account option

8. When Figure 6.46 appears, indicate whether the AD server is located on the local host or on a remote computer. If the AD server is available on the local host itself, then, select the **Local computer** option

followed by the **Finish** button. On the other hand, if the domain server exists on a remote computer, then indicate the name of the remote host in the **Another computer** text box and then click the **Finish** button.

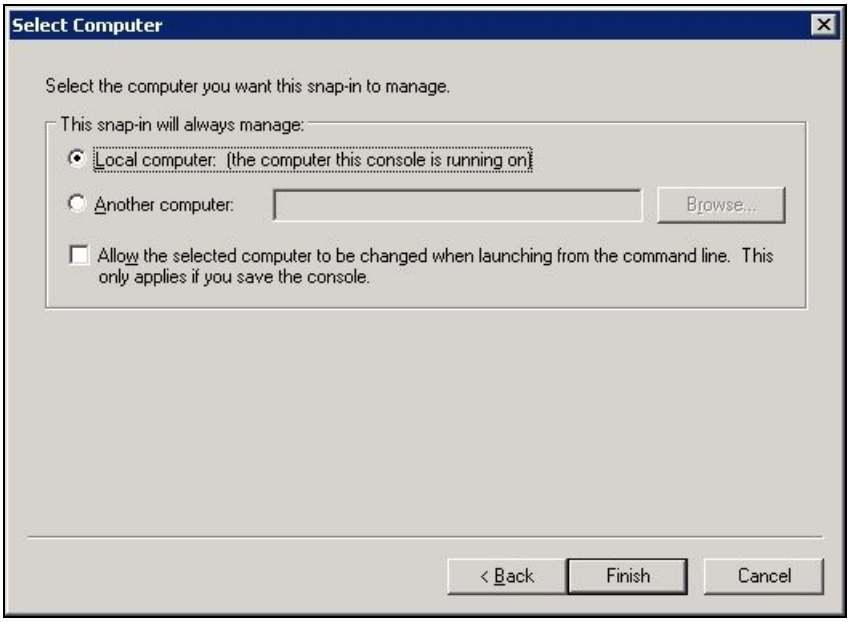

Figure 6.46: Indicating the location of the AD server

9. Once the **Finish** button is clicked, Figure 6.47 will appear displaying the **Certificates** snap-in that was added.

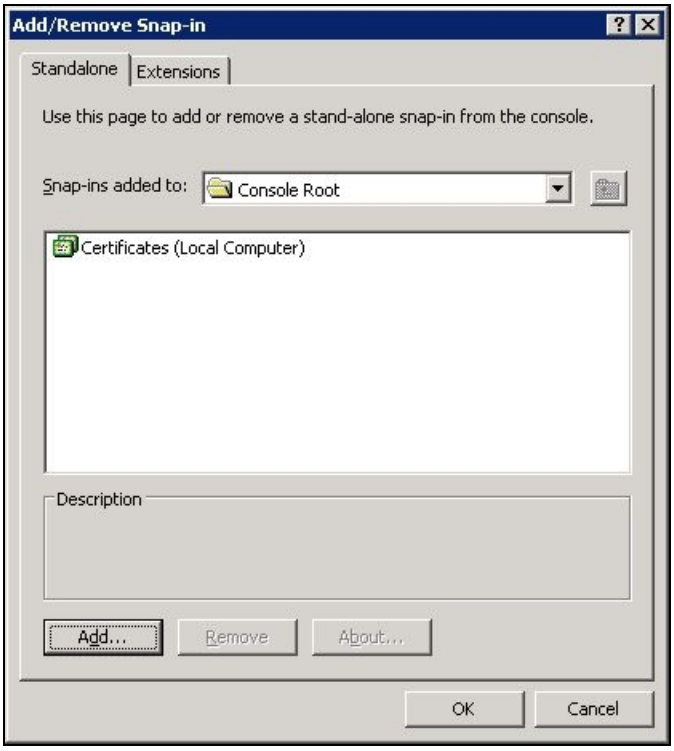

Figure 6.47: The Certificates snap-in that was added

10. Click on the **OK** button in Figure 6.47. This will lead you back to the **Snap-in Console**, which now displays

the **Certificates** snap-in that was added.

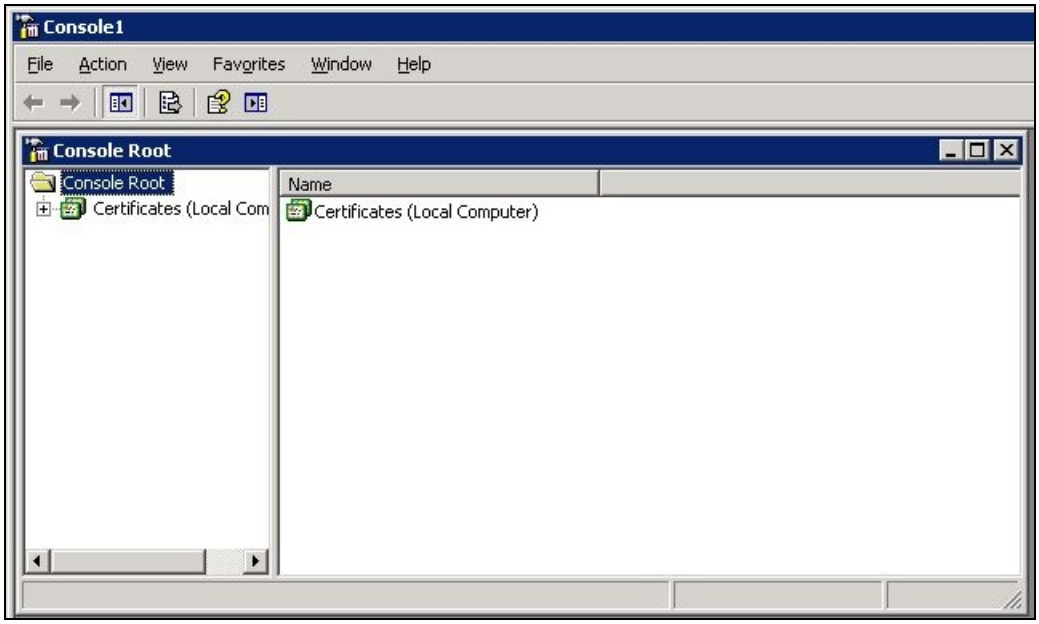

Figure 6.48: The Snap-in Console displaying the Certificates snap-in that was added

11. Next, expand the **Console Root** node in the tree-structure in the left panel of Figure 6.49, and then, expand the **Certificates (Local Computer)** sub-node. A **Personal** sub-node will then appear, which when expanded, will reveal the **Certificates** sub-node. Click on the **Certificates** sub-node to view the complete list of ertificates on the domain server (see Figure 6.49)

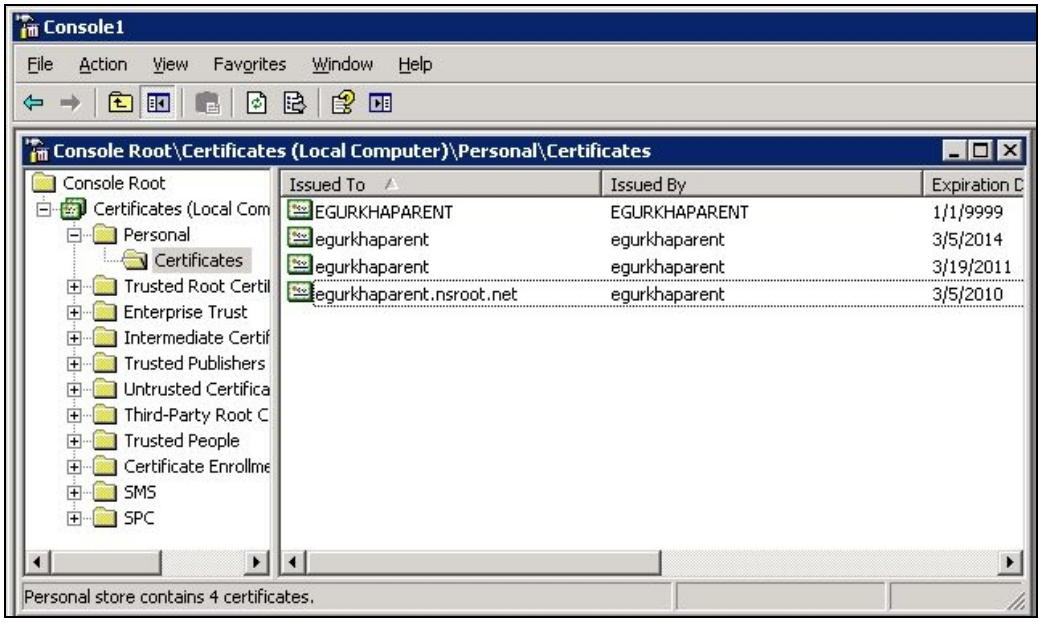

Figure 6.49: Viewing the certificates on the domain server

12. Browse the list to identify the SSL certificate of the AD server. Once identified, attempt to export the certificate to the local host (i.e., the local host). For this purpose, select the certificate from the right-panel

of Figure 6.49, right-click on the selection, choose the **All Tasks** menu, and pick the **Export** option (see Figure 6.50).

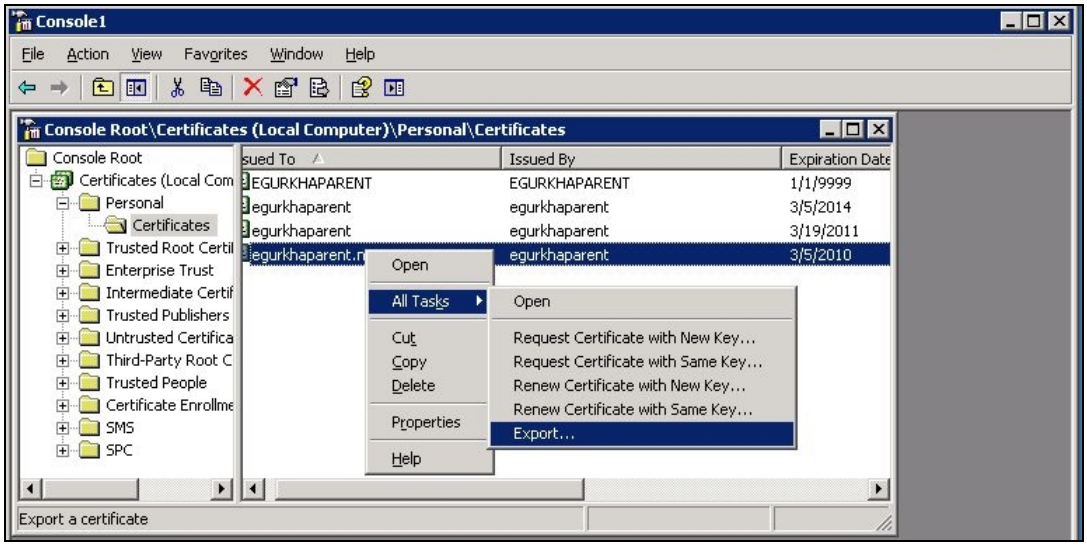

Figure 6.50: Exporting the SSL certificate of the AD server

13. Figure 6.51 will then appear welcoming you to the Certificate Export Wizard. Click the **Next** button in Figure 6.51 to continue exporting.

Figure 6.51: The Certificate Export Wizard's Welcome screen

14. In Figure 6.52 that appears, click the **Next** button to continue.

Figure 6.52: Clicking the Next button to continue

15. Select the **DER encoded binary X.509 (.CER)** option from Figure 6.53 as the export file format, and click the **Next** button to continue.

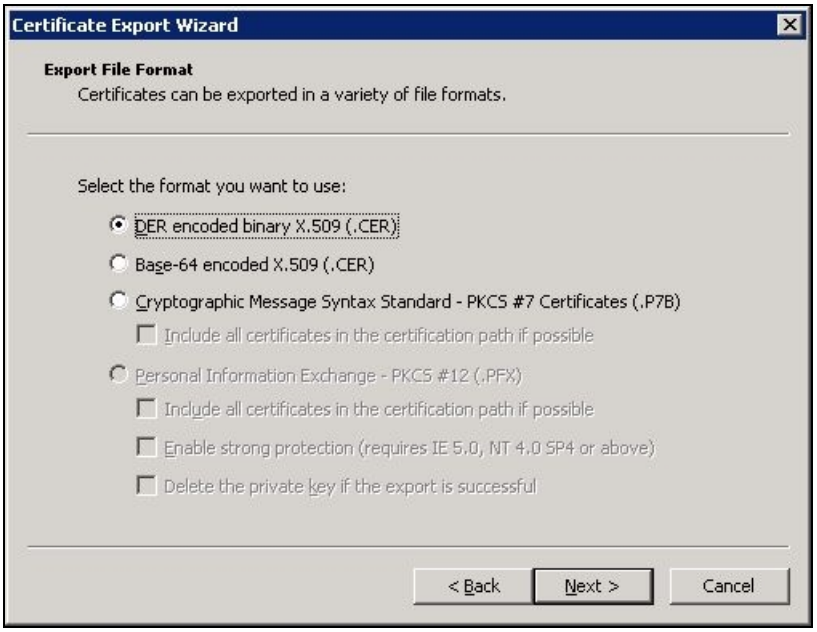

Figure 6.53: Selecting the export file format

16. Next, specify the name of the file you want to export and also indicate the directory to which the file is to be exported. You can use the **Browse** button in Figure 6.54 to specify the destination directory of the exported file. Then, click the **Next** button in Figure 6.54 to continue.

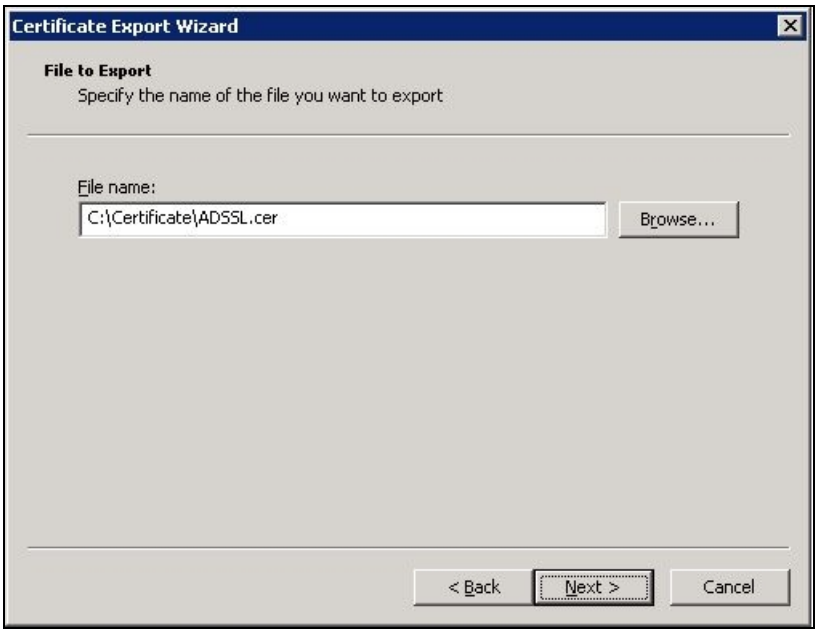

Figure 6.54: Specifying the name and destination of the exported file

17. When Figure 6.55 appears, click the **Finish** button to complete the export procedure. Once the file is exported successfully, a message box displaying a message to this effect will appear.

Figure 6.55: Finishing the export

18. Finally, copy the exported file from the local Windows host to any folder on the eG manager host.

6.2.4.2 Importing the SSL Certificate to the eG Manager

The steps in this regard are as follows:

1. Click the global **Domain(s)** node in the **DOMAIN(S)** tree of Figure 6.56. Then, click on the **Install SSL Certificate** button in the right panel.

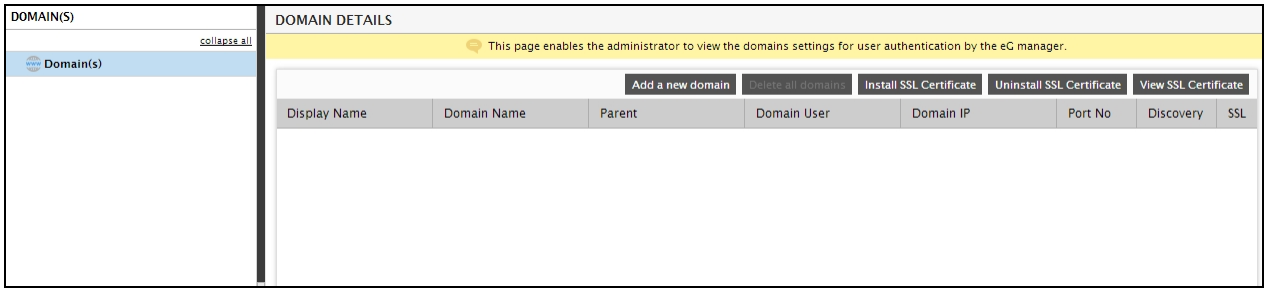

Figure 6.56: Clicking the Ínstall SSL Certificate' button

2. A **SSL CERTIFICATE INSTALLATION** page then appears (see Figure 6.57).

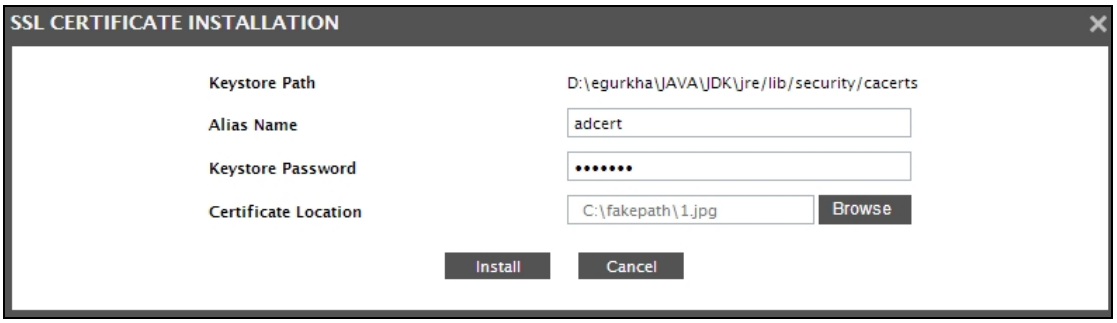

Figure 6.57: The SSL Certificate Installation popup

3. Here, specify the following:
   - **Keystore Path:** Specify the full path to the certificate file that the JDK used by the eG manager checks for trusted certificates
   - **Alias name:** Provide an alias name for the certificate being imported.
   - **Keystore password:** The default keystore password provided by Java is changeit. Provide this password against **Keystore password**.
   - **Certificate Location:** Specify the full path to the SSL certificate that was copied to the eG manager using the procedure discussed in Section **6.2.4.1**. You can use the **Browse** button in Figure 6.57 to specify the path.
4. Finally, click the **Install** button in Figure 6.57 to install the SSL certificate on the eG manager.
5. In the same way, you can install many SSL certificates on the eG manager and enable its SSL communication with many domain servers in the target environment.
6. You can view all the SSL certificates so installed by clicking the **View SSL Certificate** button in the right panel of Figure 6.56.

6.2.4.3 Uninstalling the SSL Certificate

At any given point in time, you can disable SSL communication between the eG manager and AD, by uninstalling the SSL certificate. The steps to be followed are:

1. Click the global **Domain(s)** node in the **DOMAIN(S)** tree of Figure 6.56. Then, click on the **Uninstall SSL Certificate** button in the right panel.
2. The **UNINSTALL SSL CERTIFICATE** page then appears.

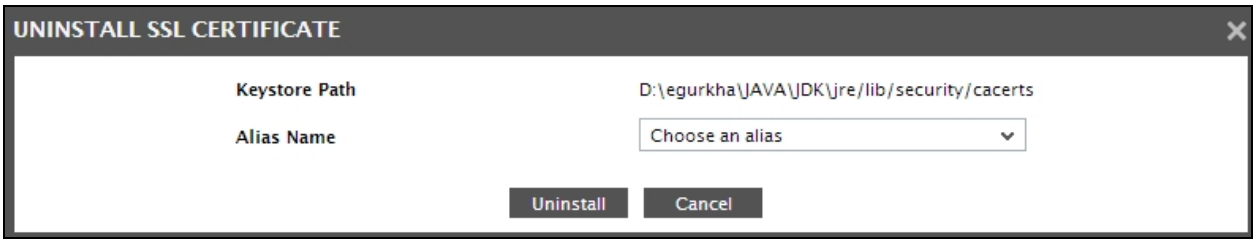

Figure 6.58: Uninstalling the SSL Certificate from the eG manager

3. Choose the **Alias Name** of the certificate to be uninstalled, and then click on the **Uninstall** button in Figure 6.58.

Troubleshooting eG Integration with Active Directory

If you have difficulty in validating domain users or are unable to login to the eG manager as a domain user, do the following:

1. Make sure that the eG manager is using **JDK 1.5**.
2. Next, go to the command prompt on the eG manager host and do the following:
   - First, set the classpath of the eG manager using the following command:

     set classpath=<EG_INSTALL_DIR>\lib\eg_manager.jar;<EG_INSTALL_DIR>\lib\jaas.jar;%classpath%
   - Next, execute the following command:

     java com.eg.KerberosAuthentication <EG_INSTALL_DIR>\manager\config\egAD_<domain>.ini <domainIP> <domainUser> <domainPass> <ValidUser> <UserBase>

     For example:

     java com.eg.KerberosAuthentication c:\egurkha\manager\config\egAD_chn.egurkha.com.ini 192.168.10.5 egtest egurkha2007 Raja DC=CHN,DC=EGURKHA,DC=COM
   - This command, upon execution, will report an exception if there is a problem connecting to the domain. If no connection errors have occurred, then an output similar to the sample output displayed below will appear.

```
The target Domain IP Address = 192.168.10.5
The connect username is = egtest
The connect password is = xxxxxxxxxx
The search username is = Raja2
The userBase is = DC=CHN,DC=EGURKHA,DC=COM
The logged in user is egtest@CHN.EGURKHA.COM
0
The logged in user is egtest@CHN.EGURKHA.COM
false
```

The penultimate line of the resulting output will display the logged in user name. The last line of the output will indicate whether the user name passed to the command above (i.e, **<ValidUser>**) is valid or not. If valid, you will find **true** in the last line, and if invalid, **false** will be displayed therein.

6.3 Adding New Users

Figure 6.59 illustrates how an administrator can add a new user to eG Enterprise. The first step toward this is to select the **Add User** menu option from the **User Management** tile to access the **ADD USER** page of Figure 6.59. To add a new user using this page, an administrator has to specify the following in Figure 6.59:

1. The **BASIC INFORMATION** tab page opens by default. From this page, select a **User role** to be assigned to the new user.
2. The eG administrative interface provides administrators with a wide variety of options to manage user information. Be it user creation, modification, deletion, or simply viewing user information, any type of user-related activity can be performed quickly and easily using the eG administrative console. Typically, when an eG user logs into the eG Enterprise system, the login is validated by the eG database, which stores the user information. However, in large IT environments that span multiple domains, the Active Directory server functions as the central repository for information related to users spread across domains, and also authenticates domain user logins. To avoid the confusion that might arise when using both the eG manager and the AD server for user authentication in such multi-domain environments, administrators might want the eG manager to integrate with AD; this ensures that the eG manager serves as the single, central, secure console for automatically authenticating logins by eG users, regardless of the size of the environment or the domain to which the user belongs. The first step towards implementing this integration is the creation of the domains and sub-domains. Use the Users -> Configure Domains menu sequence to configure the domains. For a detailed domain creation procedure, refer to Section of this document. Subsequent to domain creation, if you attempt to create a new user using this page, you will be prompted to indicate the **User authentication** mode that applies to the new user. If you are creating a domain user/group, whose login requests are to be authenticated by the Active Directory, then select the **Domain** option. If you are creating a user who is local to the eG Enterprise system, and whose login requests are to be authenticated by the eG database, select the **Local** option. Upon choosing the **Domain** option, you will have to select the domain to which the user belongs from the **Domain** drop-down in Figure 6.59.

   Reference:

   To know how to create domains, refer to Section **6.2** of this document.

   Then, indicate what you want to create - whether a domain **User** or a domain user **Group** - by picking the relevant option from the **Operation** section.

   To create a domain **User**, do the following:

   - Set the **Operation** flag in Figure 6.59 to **User**.
   - If, at the time of registering that domain with the eG Enterprise system, you had set the **Save Domain User Password in eG Enterprise?** flagto **No**, then, upon selecting the **Domain** here, you will be prompted to re-enter the **Domain User Password** (see Figure 6.59). Without this password, the eG manager will not be able to connect to the domain server and validate domain user accounts. Provide the password and click the **Submit** button in Figure 6.59 to proceed.

- Then, specify the ID of the new user in the **User ID** text box, and click the **Validate** button (see Figure 6.60). When this is done, the eG manager immediately connects to the Active Directory server and verifies whether the user is a valid domain user or not. If the user is not a valid user, then an error message to that effect appears. On the other hand, if the user is indeed a valid domain user, then the eG manager allows you to proceed with the user creation (see Figure 6.60). However, you cannot provide a password for the domain user. This is because, the credentials of the domain user are configured in and maintained by the Active Directory server; eG Enterprise therefore, will neither reveal nor allow you to modify the password of the domain user, thus ensuring data integrity. Moreover, subsequently, when you log into the eG management console as a domain user, you will have to make sure that you prefix the user name with the domain name in the format: <<domainName>>/<<Username>> (or <<domainName>>\<<Username>>). Every time a domain user logs into the eG Enterprise system, the login will be authenticated by the Active Directory server that manages the users in that domain.

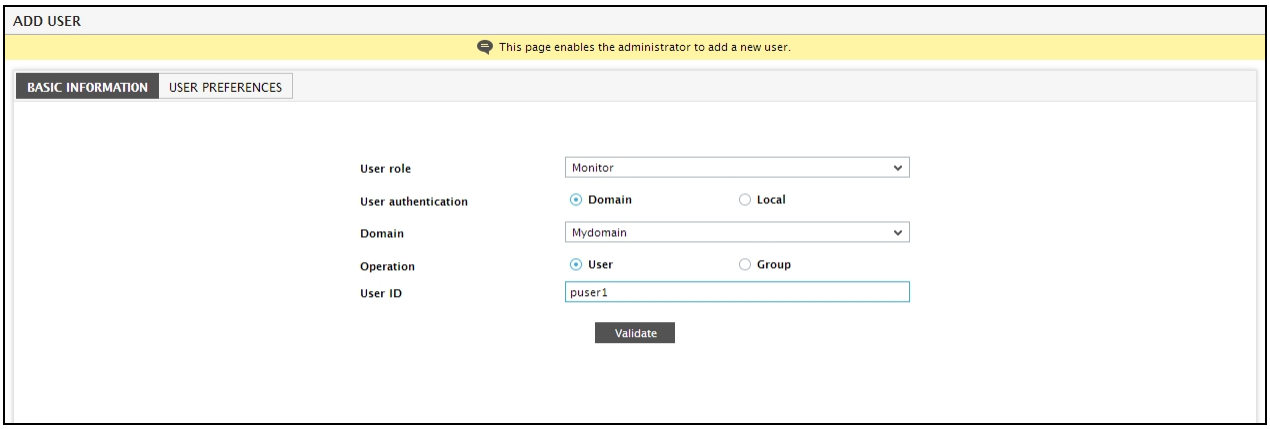

Figure 6.59: Validating a domain user name

3. Apart from individual domain users, you can also create domain user groups using the **ADD USER** page. Once a domain group is added to the eG Enterprise system, all domain users who belong to that group will be able to login to the eG Enterprise console, even if their domain credentials have not been explicitly registered in the eG system. Moreover, the access rights, privileges, and monitoring scope defined for the group will automatically apply to the users in the group, thereby saving the time and drudgery of configuring multiple user profiles - one each for every user in an Active Directory group. To create a domain user group, do the following:

   - Set the **Operation** flag in Figure 6.60 to **Group**.
   - Then, proceed to specify the **Group Name**. A domain in AD may consist of many organizational units (OUs). Each OU may be associated with a set of domain user groups. You can quickly browse the OUs in the chosen **Domain** to locate the user group of interest to you, by clicking the button to the right of the **Group Name** box. Figure 6.60 will then appear.

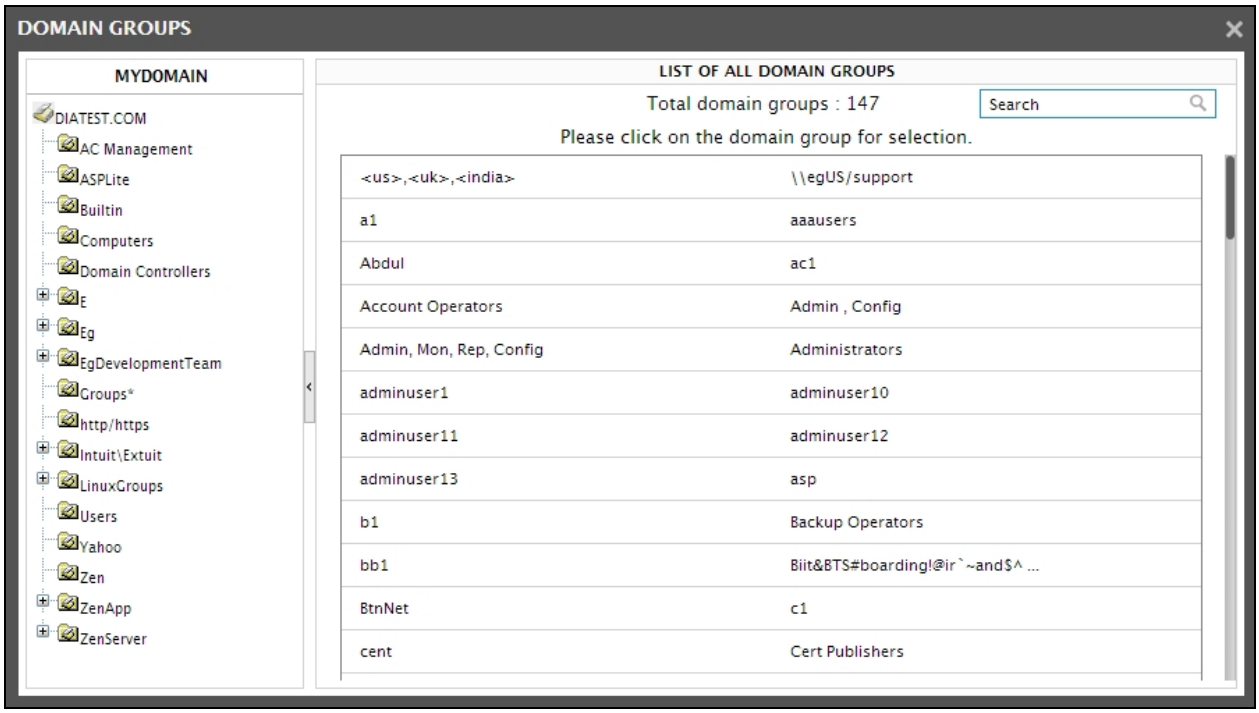

Figure 6.60: Choosing a domain user group

- The left panel of Figure 6.60displays a tree structure – while the OUs configured in the chosen **Domain** appear as the nodes of the tree, the sub-units within an OU appear as the sub-nodes. You can expand an OU node to view the sub-units within. Clicking on an OU/sub-unit in the tree will list all the domain user groups associated with that OU/sub-unit in the right panel. Click on a domain group in the right panel to register that group with the eG Enterprise system. The selected domain user group then appears against **Group Name** , as depicted by Figure 6.60 below. All users who are part of this AD group will now be allowed access to the eG Enterprise system. The rights and privileges (eg., role, expiry date, email/SMS alert settings, alarm acknowledgement/deletion rights, etc.) defined for the chosen group will govern all users who belong to that group. This saves administrators the trouble of defining separate profiles for each domain user in a group.

  Note that the group is not associated with any 'password'. This implies that while a *group* itself cannot login to the eG management console, a user who belongs to the group can login using the credentials defined for him/her in the AD server. At the time of login, the group user should provide his/her name in the format: *<DomainName>\<UserName>*. Everytime a group user logs into the eG management console, the solution automatically connects to the AD server to validate the login.

  Note:

  - eG Enterprise can be integrated with Active Directory only if the eG manager is installed using JDK 1.5 or higher. If not, you will not find any of the above-mentioned options in the eG administrative interface.
  - If a domain user group is registered with the eG Enterprise system, and a profile is later created in eG for a particular domain user in that group, then, when that user logs into the eG management console, the user-level settings will override the group-level settings.

- If a domain user belongs to more than one AD group that is created in the eG Enterprise system, then, when that user logs in, the solution provides him/her with a list of domain groups to choose from. Selecting a group from the list enables the user to automatically inherit the access rights and monitoring scope defined for that group.

4. Upon choosing the **Local** option, on the other hand, you will be prompted to specify the following:
   - Specify a unique **User ID**.
   - Provide a **Password** for the new user, and then, confirm the password by retyping it in the **Retype password** text box. This is because, in case of users who are local to the eG Enterprise system, it is the eG database which maintains the user information, and not the Active Directory. Therefore, whenever a local user is created using this page, a password has to be explicitly provided, so that both the user name and password of the local user credentials are stored in the eG database. Moreover, when a local user logs into the eG management console, his/her Username need not be pre-fixed by the domain name. The **User ID** and **Password** that the local user provides while logging in will be validated by the eG database that manages the local users.

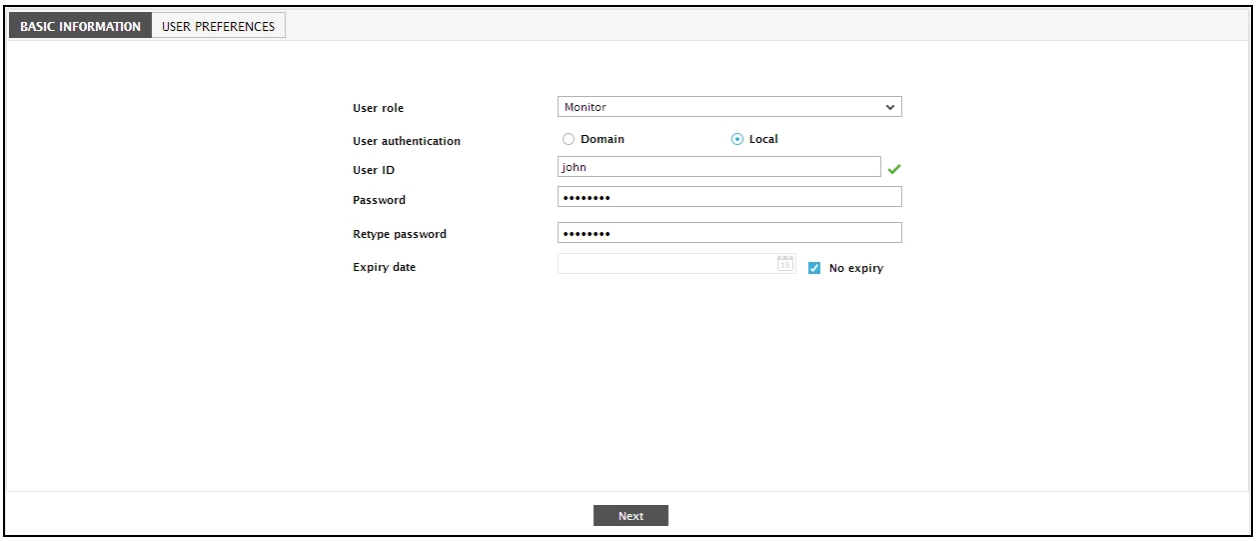

Figure 6.61: Adding a new local user

The rest of the user creation steps are common to both the authentication modes - domain and local - and to both domain users and domain user groups. The next step in user creation is to provide an **Expiry date** for the new user. One more added feature of the eG Enterprise suite is that it checks the validity of the user. A user is granted permission to monitor the services associated with him/her only for a stipulated period of time. Clicking on the Calendar button next to the **Expiry date** label will result in the display of a calendar from which the administrator can choose the validity date for a new user. Beyond this date, the user is regarded as an invalid user. Optionally, you can click on the **No expiry** check box, if a new user has to remain valid for an indefinite period of time.

5. Then, click the **Next** button in Figure 6.61. Doing so will automatically open the **USER PREFERENCES** tab page depicted by Figure 6.62.

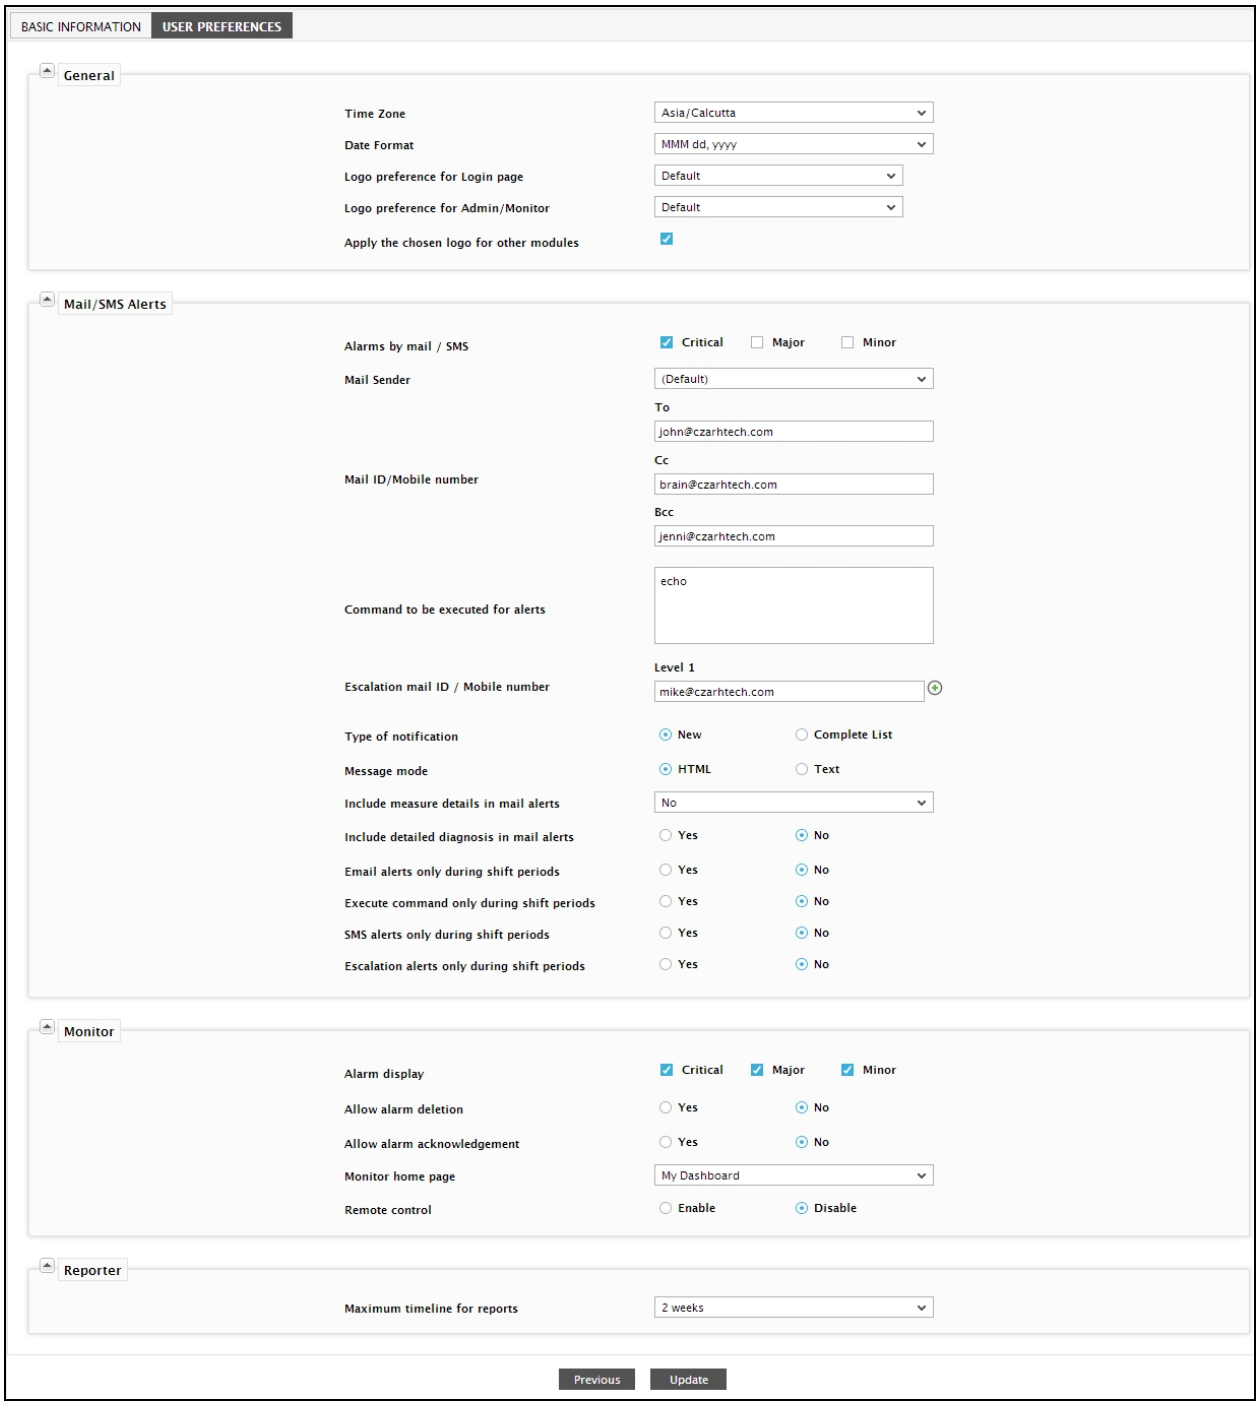

Figure 6.62: Defining user preferences

6. In the **General** section of Figure 6.62, specify the following:

- **Time Zone**: eG Enterprise is often deployed to manage servers in different geographies and time zones. For example, a large enterprise may have a central eG Enterprise management console to which agents from different locations can be reporting. In a managed service provider environment, multiple customer infrastructures can be monitored from the same eG manager. In such situations, users (administrators in different geographies, customers of an MSP in different regions) prefer to see the performance metrics reported in their respective time zones. eG Enterprise allows time zones to be associated to each user's profile. By default, all users are associated with the local time zone of the location where the eG manager is hosted. However, an administrator can change the time zone preferences of a user to suit that user's requirements. For this, when creating a user profile, the administrator can pick a **Time zone** for that user. When that user logs into the eG Enterprise console, all the metrics, alerts, and reports that the user accesses will be displayed in the respective local time zone. This new capability ensures that eG Enterprise users receive a completely 'local' experience, regardless of which part of the world the eG manager is located in.

  To understand time zone association better, take the case of an example. Assume that the user elvis, who has received the email depicted by 6.3, has been configured to use the time zone, Asia/Calcutta. We can thus conclude that the Start Time of the problem indicated by 6.3 above has been determined based on the time settings of the Asia/Calcutta zone.

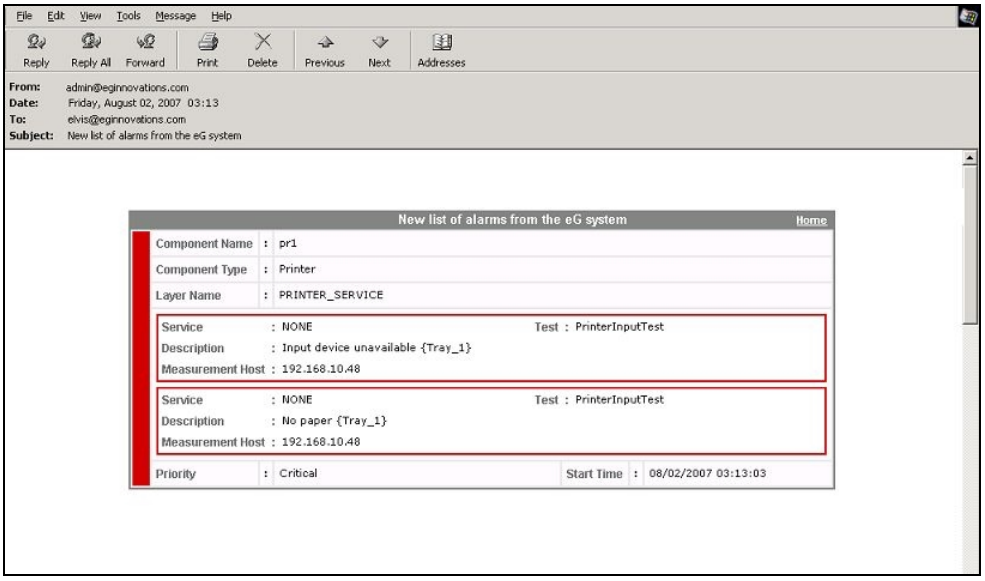

Figure 6.63: Email alert of user elvis

Now, assume that a user named john has been configured to receive email notification of the same problem. However, john has been assigned the Asia/Singapore time zone. Figure 6.64 depicts the email alert of user *john*.

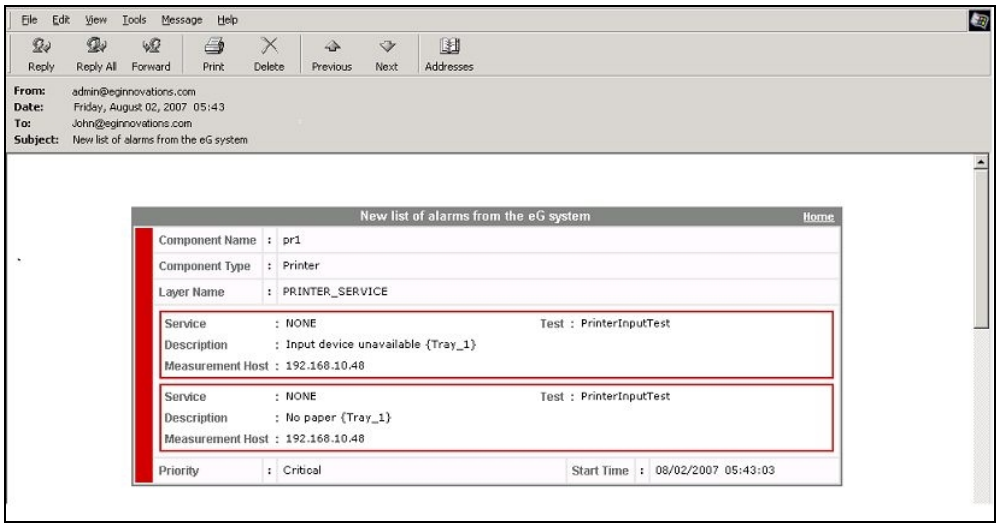

Figure 6.64: Email alert for user john

In Figure 6.64, you can see that while the other problem details remain the same as in Figure 6.64, the **Start Time** has changed to reflect the time settings of the *Asia/Singapore* zone.

Note:

While configuring a time zone, remember the following:

- When normal mails are generated by the eG manager, the **Start Time** displayed in such mails will also be based on the **TIME ZONE** setting for the corresponding user.
- If a user is configured to use multiple email IDs (i.e., a comma-separated list of mail IDs has been provided in the **TO**, **CC**, and/or **BCC** columns), then the **TIME ZONE** specification for that user applies to all the configured email IDs. In other words, every user can have a separate **TIME ZONE**, but every mail ID configured for a user cannot have a separate **TIME ZONE**.
- When mail alerts are being escalated, the time zone settings will be derived from the user account that the alarms pertain to. In other words, each escalation level will **NOT** have a separate **TIME ZONE** - the time zone setting for the user account will apply to escalated mails as well.
- Any alerts generated by the eG manager to report an unusual situation with the eG manager itself (e.g., database not working, agent not running, etc.) will not be affected by this **TIME ZONE** settings. All such alerts will be generated in the eG manager's local time zone setting.

- **Date format:** The default date format for the eG user interface is MMM dd, yyyy. This date format can be changed depending upon the country in which the user being created lives, by selecting a different format from the **Date format** list. Whenever this user logs in, the eG user interface will display dates in the chosen format only. This is particularly useful in MSP environments, where customers of the MSP could be separated by geographies and may require performance and problem reports of their hosted enivonments to be delivered in the date format that applies to their geography.
- **Logo preference for Login page, Logo preferences for Admin/Monitor module,** and **Apply the chosen logo**

for other modules: Typically, using the **Logo/Messages** menu option in the **Settings** tile, an administrator can configure a custom logo for the login screen and for every module that is enabled in the eG license – i.e., the Admin, Monitor, eG Reporter, and/or the eG Configuration Management module. MSP environments typically create a user profile in eG for each customer environment they host. While administrative rights to the eG Enterprise system will generally lie with the MSP, the customers may be granted access to one/more of the other consoles. Each of these customers may want to have their company's logo appear when they login into a console, as opposed to the MSP's logo. The **General** section of the **USER PREFERENCES** tab page enables the configuration of user-specific logos – one for every user registered with the eG Enterprise system – so that MSP customers see their company logo in the modules they have access to.

Figure 6.65: Configuring a custom logo for a user

By default, the **Default** logo will be assigned to the login page and the Admin/Monitor modules. The **Default** logo is the logo configured for the login page and for the Admin/Monitor using the**Logo/Messages** option in the **SETTINGS** panel of the **MONITOR SETTINGS** page. To define a custom logo for a user, select the **Custom** option against the module for which you want to set the custom logo (see Figure 6.65). Then, click the button adjacent to the list box to upload the custom logo to the eG manager. 6.3 will then appear.

Figure 6.66: Uploading the logo to the eG manager

The type, size and resolution of the logo image should match with the specifications mentioned in 6.3. Click the **Browse** button in 6.3 to browse for the custom logo image, and then click the **Upload** button to upload the image. When the user logs into a specific module of the manager, the logo customized for that user for that module will get displayed.

If the logo chosen for the Admin/Monitor module has to be applied to all the other modules of the eG management console (i.e., the eG Reporter and Configuration Management modules), then select the **Apply the chosen logo for other modules** check box in Figure 6.65.

7. In the **Mail/SMS Alerts** section of Figure 6.62, specify the following:

- **Alarms by mail / SMS:** The eG manager is capable of alerting users as and when problems occur. The alarms are classified into critical, major, and minor. By choosing one or more of the check boxes corresponding to the **Alarms by mail /SMS** field, a user can indicate his/her preference in terms of the priority of problems for which he/she wishes critical priority alarms alone and not the other types. If no alarm priority is chosen, then the user will not receive alerts by email / SMS.
- **Mail Sender: This option appears if at least one alarm priority is chosen from the 'Alarms by mail / SMS' section.** By default, eG Enterprise sends email alerts from the **eG Administrator Mail ID** configured in the **MAIL SERVER SETTINGS** page in the eG administrative interface. In MSP environments typically, different support groups are created to address performance issues relating to different customers. These support groups might prefer to receive problem intimation from customer-specific mail IDs instead of the global admin mail ID, so that they can instantly identify the customer environment that is experiencing problems currently. Moreover, this way, every support group will be enabled to send status updates on reported issues directly to the concerned customer, instead of overloading the admin mailbox. To facilitate this, the **MAIL SERVER SETTINGS** page allows the administrator to configure multiple **Alternative mail sender IDs** - normally, one each for every customer in case of an MSP environment. Moreover, while creating a new user, the administrator can select one of these configured sender IDs from the **Mail Sender** list and assign it to the new user, so that all email alerts received by the user are generated by the chosen ID only. Moreover, the **EG ADMINISTRATOR MAIL ID**specified in the **MAIL/SMS SETTINGS** page will also be added to the **Mail Sender** list in this page, and will be the default selection.
- **Mail ID/Mobile number: This option will appear only if at least one check box is selected from the 'Alarms by mail / SMS' section.** The **Mail ID/Mobile number**option allows a mail account(s) / mobile number(s) to be associated with a user. When multiple mail IDs are specified, an administrator can specify which mail address(es) need to be in the **To**: field of the mail alarm and which ones should be in the **Cc***:* and **Bcc:** fields. In the same way, you can even provide mobile numbers in the **To:, Cc:,**and **Bcc:**fields.

  If a mobile number(s) is specified then a compact alarm report that is ideal for a mobile phone console is generated. The first line of this report comprises of the following information, separated by slash (/).

  - IP address and port of the problem component
  - The component type

  The second line of this report would consist of the following information (separated by slash):

  - The name of the test that generated the problem measure(s)
  - The name of the problematic measure
  - The name of the service; this would be NONE if the problem component does not host any service

  Given below is a sample report transmitted via SMS:

  192.168.10.8:7077/Web_server

  ProcessTest/Num_procs_running/HTTPD/NONE

In the above example:

192.168.10.8:7077, represents the IP address and port of the component which has encountered a problem

Web_server is the type of component

ProcessTest is the name of the test that generated the problem measures

Num_procs_running is the name of the problematic measure

HTTPD is the name of the descriptor

NONE denotes that the web server does not host any service

Note:

eG alarms will be forwarded to a mobile phone only if the **NowSMS Lite** SMS manager or the **eG SMS Manager** has been installed in the network, and the eG manager has been configured to work with the SMS manager.

- **Command to be executed for alerts:** Like email IDs / mobile numbers, you can associate one/more custom scripts with users to the eG Enterprise system. Whenever alarms are raised/modified/closed for a specific user, the associated custom script will automatically execute, so that the details of the alarms are routed to third-party customer relationship management systems or TT systems, and trouble tickets automatically created (or closed, as the case may be) for the corresponding user. The custom scripts thus provide a mechanism by means of which eG alerts are integrated into CRM/TT systems. These custom scripts can be configured in addition to or instead of email / SMS alerts. To associate the command that executes the custom script with a specific user's profile, specify the command in the Command text box of Figure 6.61.

  Note:

  The **Command to be executed for alerts** text box will appear only if the **Enable Command Execution** flag in the **COMMAND EXECUTION** section of the **MANAGER SETTINGS** page (that appears when the **Manager** option is chosen from the **Settings** tile) is set to **Yes**.

- **Escalation mail ID / mobile number:** To ensure the continuous availability of mission-critical IT services, it is essential that problems be detected at the earliest and remedial action be initiated immediately. Naturally, the performance of an IT operations team is assessed by its ability to proactively isolate problems and by the speed with which the identified issues are fixed. As most IT operations teams are required to support strict service level guarantees, problems that remain unnoticed or unresolved for long periods of time could result in service level violations, warrant severe penalties, and ultimately even impact the reputation of the service provider.

  The eG Enterprise suite, with its patented correlation technology and its multi-modal (email/SMS/pager/console) problem alerting capability accurately identifies potential issues in the monitored environment, and intimates the concerned IT operators before any irredeemable damage is done. To enable IT managers to proactively track the performance of their operations teams, eG Enterprise also includes a time-based alarm escalation capability. With this capability, when a problem remains unresolved for a long time period, the eG Enterprise manager automatically

escalates the alarm to one or more levels of IT managers. The alarm escalation is based on a pre-defined escalation period, which is configured by the administrator of eG Enterprise.

The escalations are personalized for each user - i.e., each user in the eG Enterprise system is associated with multiple levels of managers. When an alert that has been sent to a user is not resolved within the escalation period, the alert is forwarded to the first level of management. If the problem remains unresolved for another escalation period, the second level of management is informed, and so on. By hierarchically escalating problems to IT managers, eG Enterprise ensures that the management staff stays informed of the state of the mission-critical IT services they control, and that they can intervene in a timely manner to ensure quick and effective resolution to key problems.

The **Escalation mail ID** section ofFigure 6.62 is where the different levels of escalation need to be specified. A comma-separated list of mail IDs/mobile numbers can be specified for the **Level 1**field to indicate the first level of escalation. You can, if you so desire, define additional support levels by clicking on the '+' sign that appears at the end of the **Level 1**text box. This way, issues that remain unresolved even at **Level 1**will be escalated to **Level 2**and so on. You can create upto a maximum of 5 escalation levels. To delete a newly added level, click on the '-' sign at the end of the corresponding **Level**text box.

Note:

Alarm escalation will work only if you configure the following:

- The duration beyond which the eG Enterprise system needs to escalate a problem to the next level
- The alarm priorities to be escalated

Both these parameters can be configured using the **ALARM ESCALATION** section ins the **MAIL ALERT PREFERENCES** page that appears when the **Alert Settings** option is selected from the **Mail Settings** menu of the **Alerts** tile.

Note:

By default, where multiple levels of escalation are configured, the eG manager does not consider the changes that may occur in the priority of an alarm between two levels of escalation. For instance, if the priority of an alarm changes from Major to Critical after the first level escalation alert is sent, the second level escalation alert will be sent for the **Critical** alarm only. Recipients of the second level escalation alerts will hence have no knowledge of the original state of that alarm. Similarly, recipients of the first level escalation alerts will not know that the alarm priority has changed. In the absence of complete problem information, the recipients of escalation alerts may not be able to perform effective problem diagnosis and provide accurate solutions. To avoid this, you can now configure the eG manager to reset escalation levels if a state transition occurs. This way, if the alarm priority changes after the first level escalation, the escalation cycle will begin all over again – i.e., escalation alerts related to the modified alarm will first be sent to the first level recipients and then the next level and so on.

- **Type of notification**: By choosing the **New** option, an administrator can indicate to eG Enterprise that when alerting a user via email/SMS, the system should send the details of newly added alarms

only. On the other hand, if the **Complete** option is chosen, the user will receive a complete list of current alarms every time a mail/SMS message is generated.

- **Message mode**: This option governs the format in which an alarm is reported in an email message. If the HTML option is chosen, the alarm details are formatted as HTML text whereas the **Text** option formats the alarm details as plain text.

  Note:

  If **HTML** is chosen as the **Message mode**, then alarms sent by mail will carry a hyperlink named **HOME** at the right top corner. The destination of the hyperlink can be configured using the **eg_services.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory. The **[MISC_ARGS]** section of the **eg_services.ini** file contains a **MailHomeURL** parameter that is left blank by default. In this case, clicking on the **HOME** link will connect you to the eG manager and open the login screen. By providing a specific URL against **MailHomeURL**, you can ensure that monitor users are lead to the specified URL upon clicking the **HOME** hyperlink.

- **Include measure details in mail alerts**: By default, the **No** option is chosen from the **Include measure details in mail alerts** list, indicating that the email alerts to a user will not include any measure details. However, if you want the email alerts to a user to include a time-of-day graph of the problem measure plotted for the last 1 hour (by default), then, pick the **Graph** option from this list. If you want the email alerts to a user to include the data plotted in a 1-hour measure graph, then, pick the **Data** option from this list.

- **Include detailed diagnosis in mail alerts**: By default, this flag is set to **No**. This implies that, by default, the detailed diagnosis (if available) of the problem measure will not be sent along with the email alerts to users. If you want the email alerts to a specific user to include detailed diagnosis information as well, then, set the **Include detailed diagnosis in mail alerts**flag to **Yes**. This information will enable users to move closer to the root-cause of the problem condition.

- **Email alerts only during shift periods:** Some environments - especially the ones that span geographies - could have operators working in shifts; for instance, an MSP environment could comprise of one/more user groups, which might work only in the nights, in order to provide help-desk services to the customers in a particular geographic region. These users naturally, would want to receive email alerts of issues only during their working hours; during the rest of day, they may prefer to be alerted via SMS. To facilitate this, eG Enterprise allows you to configure shift periods for individual users. Separate shift periods can be configured for receiving email alerts, SMS alerts, and escalation mails.

  For instance, if you want to indicate on which days and at what times a user needs to receive **email alerts of issues**, then he/she should first enable the **Email alerts only during shift periods** flag, by setting it to **Yes**.

  - **Email alerts only during shift periods:** Some environments - especially the ones that span geographies - could have operators working in shifts; for instance, an MSP environment could comprise of one/more user groups, which might work only in the nights, in order to provide help-desk services to the customers in a particular geographic region. These users naturally, would want to receive email alerts of issues only during their working hours; during the rest of day, they may prefer to be alerted via SMS. To facilitate this, eG Enterprise allows you to configure

shift periods for individual users. Separate shift periods can be configured for receiving email alerts, SMS alerts, and escalation mails.

For instance, if you want to indicate on which days and at what times a user needs to receive **email alerts of issues**, then he/she should first enable the **Email alerts only during shift periods** flag, by setting it to **Yes**.

Note:

In environments where shifts are not relevant, the **Email alerts only during shift periods** flag may be meaningless. You can therefore ensure that this flag does not appear in the **USER PREFERENCES** tabpage by following the steps given below:

- Select the **Alert Settings** option from the **Mail Alerts** menu of the **Alerts** tile.
- Select the **Shift Periods** node from the **Settings** tree in the left panel of the page that appears.
- When the **SHIFT PERIODS CONFIGURATION** section appears in the right panel, set the **Allow shift period configuration** flag to **No**. By default, this flag is set to **Yes**.
- Finally, register the changes by clicking the **Update** button in that page.

Upon setting the flag to**Yes,** you will be required to specify the **Days** on which the user should receive email alerts; also, in the **Shifts** field alongside, you need to mention the specific time periods on the chosen **Days** at which the user should receive email alerts (seeFigure 6.67):

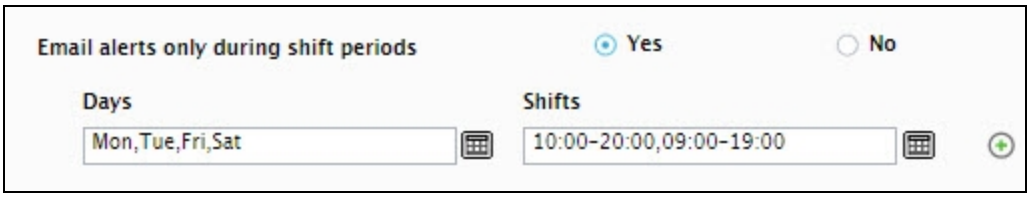

Figure 6.67: Enabling the 'Email alerts only during shift periods' flag

To select one/more **Days**, do the following:

- First, click on the **Calendar** control ( ) next to the **Days** field.
- From the **DAYS** list that pops out (see Figure 6.68), which lists the days of the week, select the days on which email alerts need to be sent to the user.

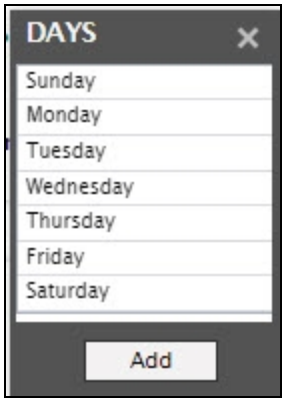

Figure 6.68: The DAYS list

- To choose more than one day from the list, select a day by clicking on the left mouse button, and then, with the **Ctrl** button on your keyboard pressed, click on another day to select it. Similarly, multiple days can be selected. To add your selection to the **Days** field, click the **Add** button in the **DAYS** list. You will thus return to the **USER PREFERENCES** tabpage where the selected days will be listed against the **Days** field (seeFigure 6.67).

- Next, using the **Shifts** field, provide the specific time periods at which email alerts should be sent out to the user on the chosen days. For that, do the following:
  - First, click on the **Calendar** control (🖩) next to the **Shifts** field. Doing so invokes the **SHIFTS** window, wherein you can specify a **From** time and **To** time for your shift (seeFigure 6.69). **Ensure that the shift timings correspond to the Time zone chosen for the user**.

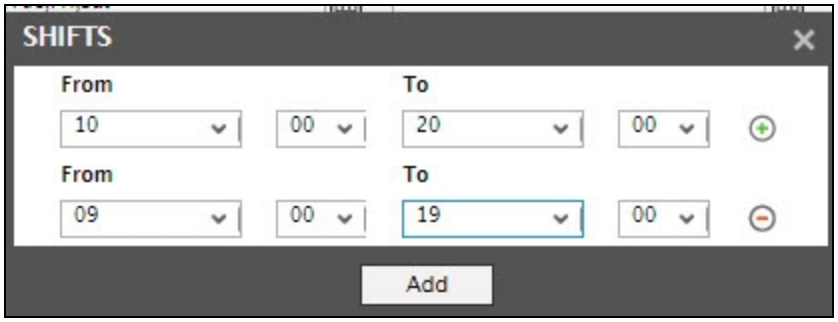

Figure 6.69: Specifying the shift time periods

  - To provide an additional time slot, click on the circled '+' button at the end of the first row. Another row then comes up wherein you can provide one more time period (seeFigure 6.69). In this way, you can associate a maximum of 5 shift periods with the chosen **Days**.
  - To remove a shift period from the **SHIFTS** window, simply click on the circled '-' button against the corresponding specification. Finally, to add these time periods to the **Shifts** field, click on the **Add** button in the **SHIFTS** window. You will thus return to the **USER PREFERENCES** tabpage, where you can find the time period(s) that you specified appear against the **Shifts** field (seeFigure 6.67).

With that, one **Day-Shift** specification is complete.

To add another **Day-Shift** specification, just click on the circled '+' button at the end of the first row Figure 6.67. Another row will then appear, where you can specify a few more **Days** and **Shifts**. This way, a number of **Day-Shift** specifications can be associated with a user (Figure 6.70).

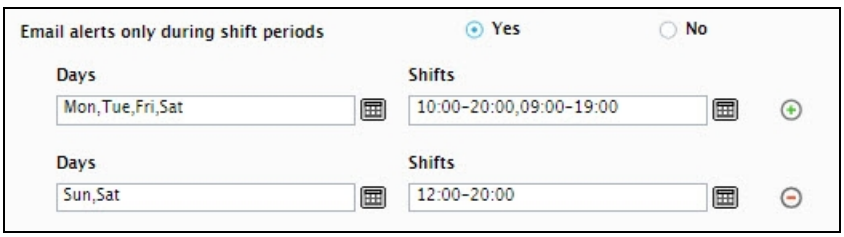

Figure 6.70: Configuring multiple Day-Shift combinations

This number is configurable, and can be any number between or equal to 1 and 10. To configure this number, go to the **SHIFT PERIODS CONFIGURATION** section of the **MAIL ALERT PREFERENCES**page that appears when you click on the **Alert Settings** option of the **Mail Settings** menu in the **Alerts** tile. In the **SHIFT PERIOD CONFIGURATION** section, select the **Maximum number of day-shift combinations**.

To delete a particular **Day-Shift** specification from Figure 6.70, simply click on the circled '-' button in Figure 6.70.

- **SMS alerts only during shift periods:** Similar to email alerts, SMS alerts can also be configured to be sent out only during specified time periods on specific days of the week. The first step towards this is to enable the **SMS alerts only during shift periods** flag by selecting the **Yes** option in the **USER PREFERENCES** tab page. Using the **Days** and **Shifts** fields that appear subsequently, you can configure one/more **Day-Shift** combinations in the same manner as discussed for email alerts (see Figure 6.71).

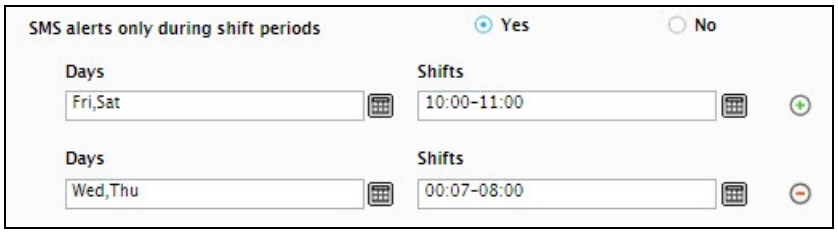

Figure 6.71: Configuring shift periods for SMS alerts

Note:

In environments where shifts are not relevant, the **SMS alerts only during shift periods** flag is meaningless. You can therefore ensure that this flag does not appear in the **USER PREFERENCES** page by following the steps given below:

- Select the **Alert Settings** option from the **Mail Alerts** menu of the **Alerts** tile.
- Select the **Shift Periods** node from the **Settings** tree in the left panel of the page that appears.
- When the **SHIFT PERIODS CONFIGURATION** section appears in the right panel, set the **Allow shift period configuration** flag to **No**. By default, this flag is set to **Yes**.
- Finally, register the changes by clicking the **Update** button in that page.

- **Escalation alerts only during shift periods:** Like email and SMS alerts, the eG manager can be configured to send escalation mails/SMS' also at pre-defined days and time slots. To enable this capability, first, turn on the **Escalation alerts only during shift periods** flag by selecting the **Yes** option in the **USER PREFERENCES**page. As before, this will bring up the **Days** and **Shifts** fields (see Figure 6.72), using which you can configure the days on which and the times at which alerts are to be escalated to the specified individuals. The procedure for configuring the **Day-Shift** combinations is the same as that for email and SMS alerts (see Figure 6.72).

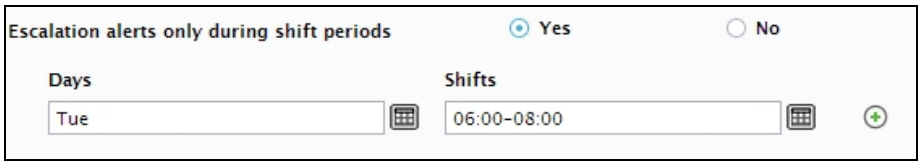

Figure 6.72: Configuring shift periods for escalation mails/SMS

Note:

In environments where shifts are not relevant, the **Escalation alerts only during shift periods** flag is meaningless. You can therefore ensure that this flag does not appear in the **USER PREFERENCES** page by following the steps given below:

- Select the **Alert Settings** option from the **Mail Alerts** menu of the **Alerts** tile.
- Select the **Shift Periods** node from the **Settings** tree in the left panel of the page that appears.
- When the **SHIFT PERIODS CONFIGURATION** section appears in the right panel, set the **Allow shift period configuration** flag to **No**. By default, this flag is set to **Yes**.
- Finally, register the changes by clicking the **Update** button in that page.

- **Execute commands only during shift periods:** As mentioned already, the **Command to be executed for alerts** text box in Figure 6.62 can be configured with the command to run a custom script. This script will be automatically executed when an alarm is newly created/modified/deleted for a specific user, and will route the details of the alarms to third-party customer relationship management systems or TT systems. If need be, you can have this custom script run only during specific time slots on specific days. To enable shift-based execution of the specified **Command**, set the **Execute commands only during shift periods** flag to **Yes.** As before, this will bring up the **Days** and **Shifts** fields, using which you can configure the days on which and the times at which the specified **Command** needs to run. The procedure for configuring the **Day-Shift** combinations is the same as that for email and SMS alerts.

8. Next, in the **Monitor** section:

   - **Alarm display**: By selecting one or more options provided against this field, you can associate specific alarm priorities with the user being created. When this user later logs into the eG monitor interface, alarms of the chosen priorities alone will be displayed in the **CURRENT ALARMS**window of the monitor interface.
   - **Remote control**: Using this option, you can indicate whether the new user is to be provided access to the managed components. This capability, when enabled, allows monitor users to remotely manage and control components from a web browser itself. By default, this capability will be **Disabled** for a new user. To enable this capability for a particular user, select the **Enable** option. Doing so invokes a **Remote command execution** list (see Figure 6.73), using which you need to indicate whether the new user is authorized to execute any command remotely, or is only allowed to choose from a pre-configured list of commands.

Figure 6.73: Selecting a command for remote execution

- **Allow alarm deletion:** To allow the new user to delete alarms from the **CURRENT ALARMS** window in the eG monitor interface, select the **Yes** option from the **Allow alarm deletion** section.
- **Allow alarm acknowledgement:** Optionally, specific users can be configured to acknowledge an alarm displayed in the eG monitor interface. By acknowledging an alarm, a user can indicate to other users that the issue raised by an alarm is being attended to. In fact, if need be, the user can even propose a course of action using this interface. In such a case, a user with Admin or Supermonitor privileges (roles) can edit the acknowledgement by providing their own comments/suggestions on the proposed action. The acknowledgement thus works in three ways:
  - Ensures that multiple members of the administrative staff do not unnecessarily invest their time and effort in resolving a single issue;
  - Serves as a healthy forum for discussing and identifying permanent cures for persistent performance ills;
  - Indicates to other users the status of an alarm

  To enable the alarm acknowledgement capability for the new user, select the **Yes** option from the **Allow alarm acknowledgement** section in this page.
- **Monitor Home Page:** By default, the **Monitor Dashboard** appears as the home page of the eG monitoring console - i.e., as soon as a user logs into the monitoring console, the **Monitor Dashboard** appears as the first page by default. eG Enterprise however, allows administrators to set any page they deem fit as the **Monitor Home Page** for individual users to the eG monitoring console. This way, every user, upon logging into the eG monitor interface, is enabled to view straight up the information that interests him/her the most, thereby saving time and minimizing the mouse clicks that may be required to navigate to that information!

  The home page preference is typically driven by the monitoring needs of specific users and the roles assigned to them. For instance, a service manager, who is responsible for minimizing/eliminating service outages, would want to know on login how all the critical services in the environment are performing currently, and which services are in an abnormal state. For this purpose, administrators may want to set the **Service List** as the home page of such users.

9. In the **Reporter** section, specify the following:
   - **Maximum timeline for reports**: Typically, eG Enterprise permits multiple users to simultaneously access the eG Reporter console and generate a wide variety of reports spanning any timeline of their choice. While this imparted tremendous flexibility in report generation, when concurrent users generate reports for broad time periods, report generation could slow down a little. In order to avoid this, administrators can set the maximum timeline for which each user can generate reports, by selecting an option from the **Maximum timeline for reports**list in the **ADD NEW USER** page. By default, **1 month** will be selected here. This implies that the user being created can generate reports for a

maximum timeline of **1 month** only, by default. The other options in this list are as follows: 1 day, 2 days, 3 days, 4 days, 5 days, 6 days, 1 week, 2 weeks, 3 weeks, and 4 weeks. Besides ensuring that unauthorized users are denied access to more historical information than necessary, this timeline restriction also greatly reduces the strain on the eG database.

Note:

The default users - *admin* and *supermonitor* - are not governed by this maximum timeline setting; these two users therefore can generate reports for any timeline.

10. In the **Associate segments/services/service groups/zones/components** section, indicate the following:

- **Auto- associate to other users:** eG Enterprise allows administrators to assign specific segments/services/components/zones to a new user for monitoring. If one/more other existing users share the same assignment, then you can automatically associate all the infrastructure elements chosen for one user with other users to the eG Enterprise system. To achieve this, first select the **Auto-associate to other users** check box. Doing so invokes an **Available user(s)** list from which you canselect the users to whom the same set of segments/services/components/zones need to be assigned.

Note:

The **Associate segments/services/service groups/zones/components**section will appear only when the **User role** chosen allows access to **Limited** components in the monitored environment. If the role chosen allows **Complete** components access, this section will not appear.

Note:

At any given point in time, you can close any section that is available for configuration in the **USER PREFERENCES** tab page by clicking the down arrow button that precedes the section name. For instance, the **Reporter** section can be closed by clicking the 'down-arrow' button adjacent to the section name, **Reporter**.

Upon clicking the **Update** button in the **USER PREFERENCES** tab page of Figure 6.62, Figure 6.74page will appear, using which you can associate one/more infrastructure elements such as zones/services/segments/components with the new user.

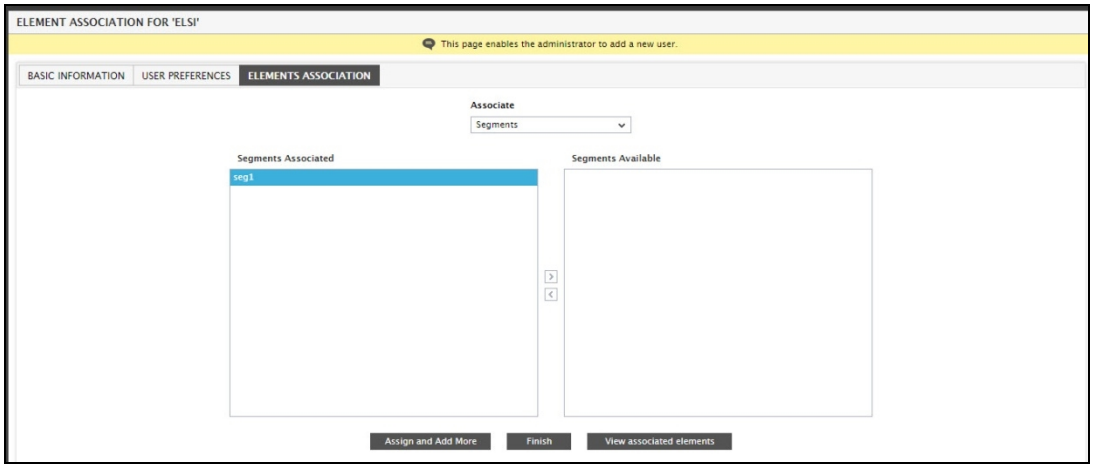

Figure 6.74: Associating a segment with a user

In order to manage large infrastructures in a more effective and efficient manner, administrators might prefer to break down the infrastructure into smaller, more manageable units known as **Zone**s, and monitor these individual units. A zone can typically comprise of individual components, segments, services, and/or other zones that require monitoring. Using this page, administrators can associate a new user with specific zones that have been configured in the environment. This association ensures that the user is authorized to monitor all the components, segments, services, and/or other zones that form part of the selected zone.

Another characteristic feature of the eG Enterprise product is that users can be associated with specific services and service groups. A service can be a group of applications or network devices that work together to deliver certain services to the end-user. It is characterized by a group of servers that belong to a segment(s) or a group of independent servers. In this case, a user associated with one service cannot view the details pertaining to the other.

Large organizations may have multiple services grouped under different business units. There may hence be a need to represent groups of services as an entity. To address this requirement, eG Enterprise allows the configuration of service groups in the eG admin interface. The service groups so configured can also be assigned to specific users.

Besides the above infrastructure elements, virtual machines can also be assigned to specific users. This feature is particularly useful for cloud service providers, who often need to provision a VM on-demand for any customer who requests for it. These cloud consumers (i.e., customers) are only concerned with the availability and internal health of those VMs that the service provider has provisioned for them, as typically, they will have no knowledge of the virtual servers on which the VMs operate. This means that these customers may not require monitoring access to the whole virtual server as such. By configuring user-VM mappings using the eG administrative interface, the cloud service providers will not only be able to track who is using which VM, but will also be able to provide a customer with the ability to view in real-time the status, overall performance, and problems related to only the VMs (and not the virtual servers) that were specifically launched on the cloud for him/her.

To assign one/more of these infrastructure elements to a user for monitoring, follow the steps given below:

1. First, choose the type of element to be associated with the new user from the **Associate** list in Figure 6.75. This list displays all the infrastructure element types that have been configured in the environment. To associate a segment for instance, select the **Segment** option from this list.
2. Choosing the **Segment** option lists all fully-configured segments in the environment in the **Segments Available** list. From this list, select the segment to be associated with the new user, and then click the **<** button. This will transfer the selection to the **Segments Associated** list in this page. Similarly, you can disassociate segment from a user by selecting the segments from the **Segments Associated** list and clicking the **>** button. If you want to update the current association and continue adding more elements to the user view, click on the **Assign and Add More** button.
3. Similarly, you can associate zones and/or services to a new user (see Figure 6.75).

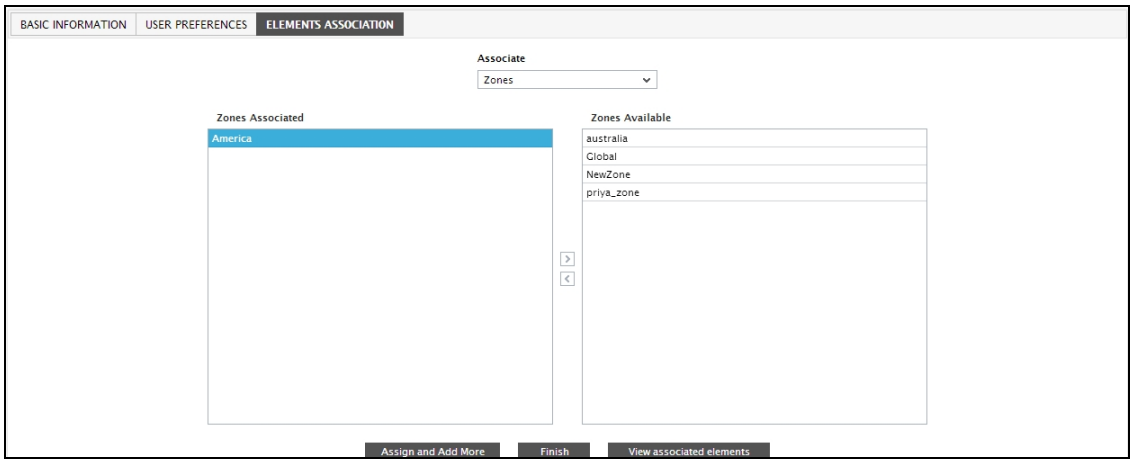

Figure 6.75: Associating services/zones to a new user

4. To associate independent components with a user, first select the **Components** option from the **Associate** list, and then select a **Component type** (see Figure 6.76).

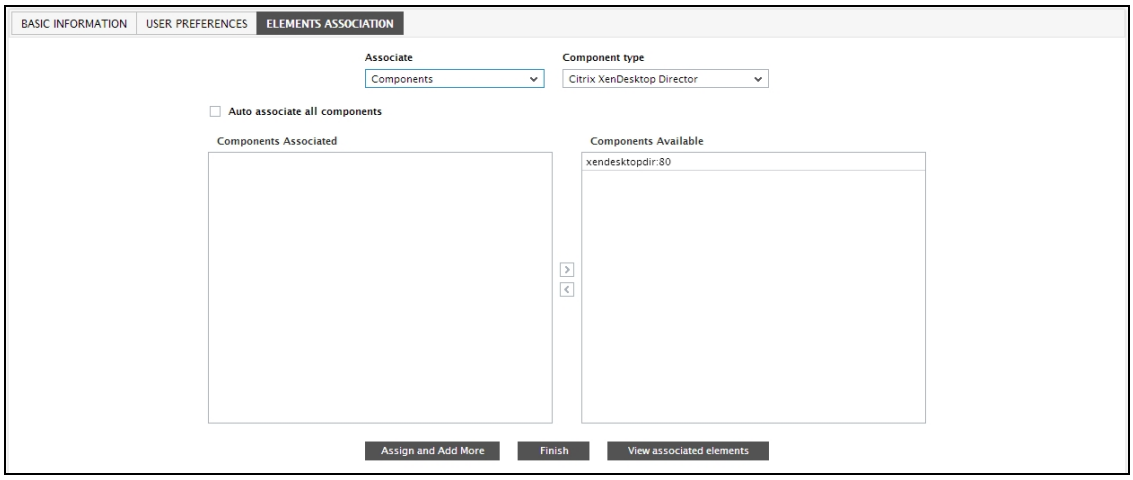

Figure 6.76: Viewing the components of a type

5. All independent components of the chosen type will be displayed in the **Components Available**list in Figure 6.76. To associate specific components, select them from the **Components Available** list, and then click the <button. Instead, if you want to associate all components of the chosen type, simply click the **Auto associate all components**check box as indicated byFigure 6.77. Doing so automatically transfers all the components displayed in the **Components Available** list to the **Components Associated**list (see Figure 6.77).

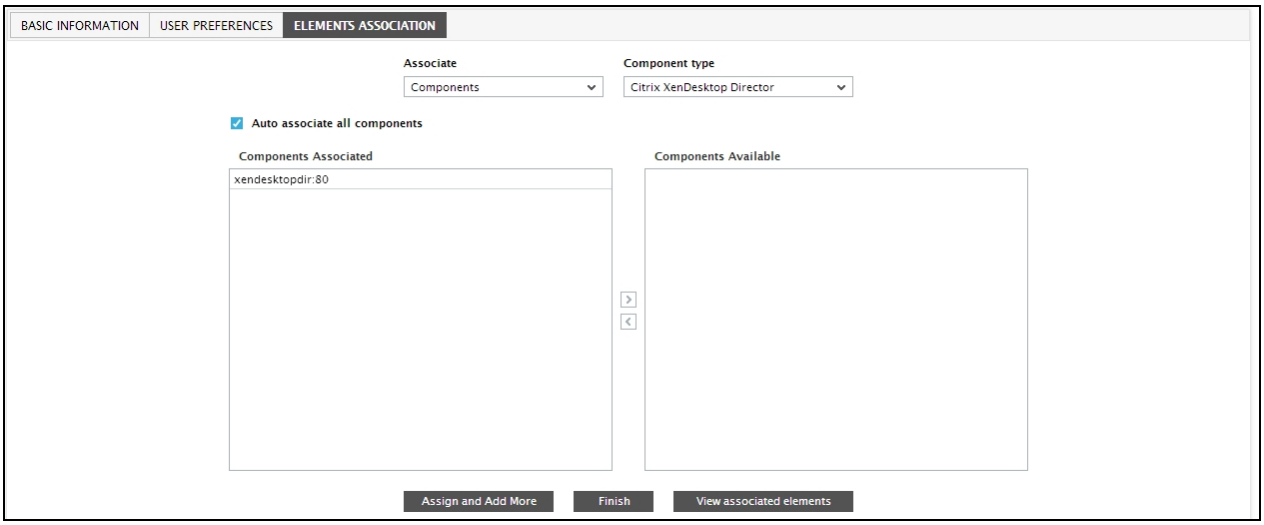

Figure 6.77: Associating all components of a chosen type with a user

6. If you want to assign one/more VMs to users, select **Components** from the **Associate** list, and then pick **Virtual Machine** as the **Component type**. This will automatically populate the **Hypervisor Types**list (see Figure 6.78) with all the managed hypervisors in the environment. Select a hypervisor from this list; **note that at any given point in time only a single hypervisor type can be chosen from this list**.

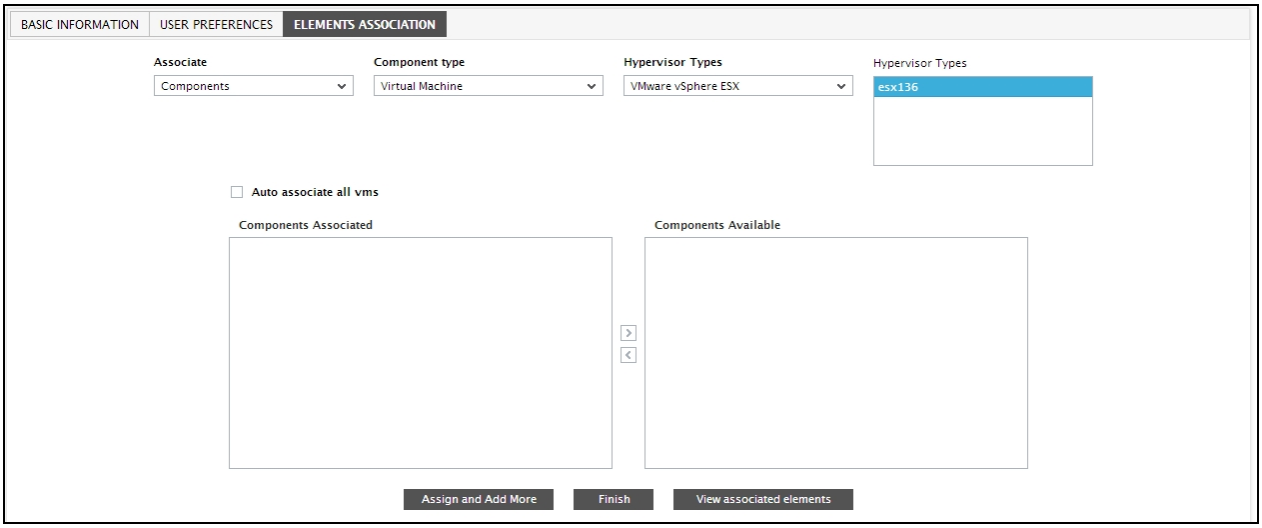

Figure 6.78: Mapping VMs to a user

7. Doing so will instantly populate the **Hypervisor Types**list box with all the managed virtual hosts of the chosen **Hypervisor Types**. From the **Hypervisor Types**list box, select the hosts that have been configured with the VMs to be assigned to the user. The VMs that the eG agent auto-discovers from the chosen virtual hosts will then be displayed in the **Components Available**list. From this list, select the VMs that are to be assigned to the user for monitoring and click the **<** button. The chosen VMs will then be moved to the **Components Associated**list. To associate all the VMs displayed in the **Components Available**list with the user at one go, select the **Auto associate all vms** check box in Figure 6.78.
8. To disassociate one/more VMs that were previously mapped to a user, select the VMs from the **Components Associated** list and click the **>** button.

9. If you are done with associating elements to a user, then you can save all your previous associations and exit this page by clicking the **Finish** button. To add more elements, click the **Assign and Add More** button.
10. At any given point in time during the element association, you can click the **View associated elements** button in Figure 6.78 to view the elements that have been associated with the user in question.

Note:

- Independent components that belong to a zone that is associated with a user, will be automatically removed from the **Components Available**list ofFigure 6.77.
- Newly added/managed components belonging to the selected component type do not get associated with the new user immediately. Since this association is mapped as part of the discovery process, there might be a latency equal to the rediscovery period before an association between users and components is updated. If the rediscovery period has not been specified, there will be a latency equal to one day.

Note:

- The **ELEMENTS ASSOCIATION** tab page (see Figure 6.78) will appear only if the role assigned to the new user allows access to **Limited** components in the monitored environment. If the new user is assigned a role that allows **Complete** access, then this page will not appear.
- Say, a user was assigned a role that allowed **Limited** component access. Assume that a segment named *seg-a* and a service named *online_shop* were assigned to this user. If the user role is now modified to allow **Complete** component access, the access rights of the user will change accordingly - i.e., the user will now have access to all the managed elements in the infrastructure. Say, the user role is now modified once again to allow **Limited** component access. When this is done, the corresponding user profile will also change, and the segment (*seg-a)* and service (*online_shop*) that were originally associated with this user will be automatically reassigned. This indicates that when the access rights of a user role goes from **Limited** to **Complete** and then back to **Limited**, eG Enterprise retains the original assignments of the corresponding user and applies the same eventually.

Note:

Typically, the user activity in high-security environments is periodically audited to ensure compliance with set standards and to enable the swift detection of unauthorized accesses. One of the requirements of such audits is a report that provides a consolidated list of users to the target environment as on the current date, the application(s) they have access to, and the details of the access privileges granted to each user with respect to that application. Such a report enables both the administrators and the auditors to determine if any user has been allowed access to more areas than necessary, thus enabling them to fine-tune their firewall rules.

As part of this exercise, if administrators want to generate a report for tracking users to the eG Enterprise application alone, then they can enable the **user logging** capability of the eG manager. To enable this capability, do the following:

- Edit the **eg_services.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
- Set the **UserAudit** flag in the **[MISC_ARGS]** section to **Yes** (default is **No**).

- Save the **eg_services.ini** file.

Enabling user logging results in the creation of a **user_ log** file in the **<EG_ INSTALL_ DIR>\manager\config\logs** directory. By default, the access permissions of all 'active' user accounts registered with the eG Enterprise system as on the current date, are logged in this file every day. If you want the file to log the permissions of "expired" user accounts too, then, set the **LogExpiredUsersPrivilege** flag in the **[MISC_ARGS]** section of the **eg_services.ini** file to **Yes**.

The format of the entries in the **user_log** file is as follows:

ApplicationNumber, User Id, User Role, Access Permissions
In the format, *ApplicationNumber* is a unique identification number that you need to manually assign to the eG Enterprise application. To define an *ApplicationNumber* for eG Enterprise, edit the **eg_specs.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory, and provide any string/number against the **ApplicationNo** parameter (in the **[MISC_ARGS]** section). **Note that if an Application Number is not defined for eG users in the eg_specs.ini file, then user logging will not occur!**

UserId refers to the name of the user registered with the eG Enterprise system.

User Role represents the role assigned to the user.

Access Permissions is a brief description of the specific permissions that have been granted to the user.

A sample log entry is provided below:

100, john, testConfigure, testConfigure users have the following abilities: [Admin] configure tests; configure thresholds.

Here, *100* is the *ApplicationNumber*, *john* is the *UserId*, *testConfigure* is the *User Role* assigned to *john*, and finally, user *john* is permitted to *configure tests* and *configure thresholds* using the eG administrative interface.

6.4 Cloning an Existing User's Profile

In large environments characterized by numerous users, you may have many users who play the same role or shoulder the same responsibilities in the organization; such users may also require that they be assigned more or less the same monitoring rights and privileges to the eG Enterprise system. In such circumstances, instead of repeatedly performing the redundant exercise of configuring a separate profile for each user, administrators can quickly 'clone' or 'copy' the monitoring settings of a particular user to create another profile. This saves administrators the time and effort involved in manually configuring and registering multiple user profiles with similar monitoring preferences.

To 'clone' a specific user's profile, do the following:

1. Select the **Clone User** option from the **User Management** tile.
2. In Figure 6.76 that then appears, **Choose an existing user to clone**.

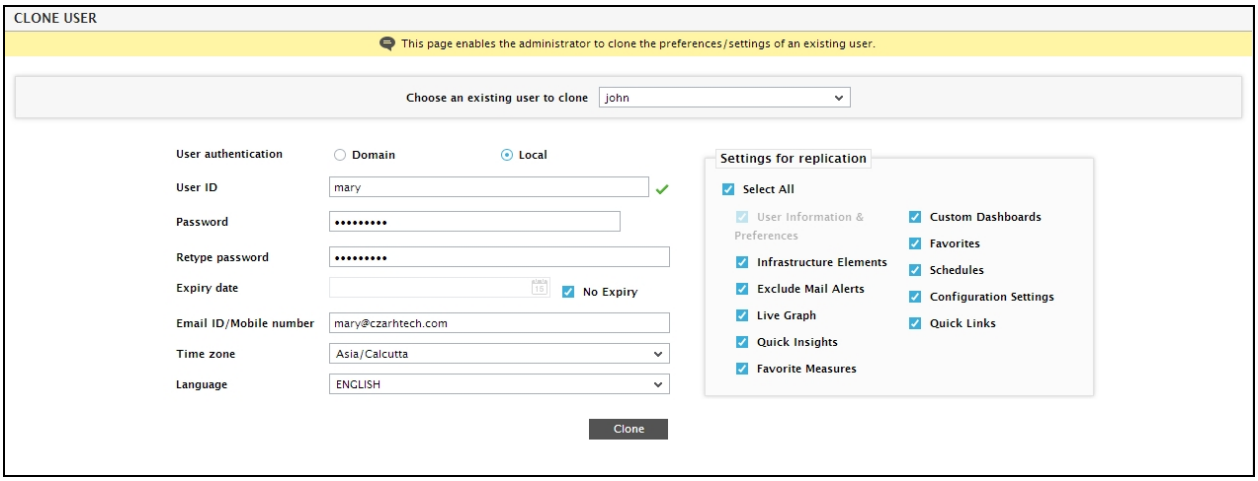

Figure 6.79: Cloning an existing user

3. Then, proceed to provide the settings that are specific to the new user. The first in that list is the **User authentication** mode. Here, indicate whether the user being created is a **Domain** user or a **Local** user. By default, the **Local** option is chosen. In this case, proceed as discussed below:

   - Specify the **User ID**.
   - Enter the **Password** of the user.
   - **Retype** the **Password**.
   - Against **Expiry date**, indicate the date on which the user account will expire. Alternatively, you can select the **No expiry** check box if the user account being created should never expire.
   - To configure the new user to receive email/SMS alerts of problem conditions in the environment, provide a comma-separated list of mail IDs and/or mobile numbers against the **Email ID / Mobile number** text box. If no **Email ID / Mobile number** is provided here, then alarm information will not be sent as emails/SMS to this user.

     Note:

     Even if the original user has not been configured to receive email/SMS alerts of issues, the 'clone' can be configured to receive the same. In this case therefore, the 'clone' will by default receive email/SMS notifications of issues, regardless of their priority - i.e., the 'clone' will be alerted to Critical, Major, and Minor problems via email/SMS.

     On the other hand, if the original user is configured to receive email/SMS alerts of issues, then, all the settings related to these alerts (with the exception of the email IDs and mobile numbers to which the alarms are to be sent) will apply to the clone as well, by default. These settings include the following:

     - The email IDs to which the email alerts are to be copied (**Cc** and **Bcc**)
     - The alarm priorities for which email/SMS alerts are to be sent
     - The **Mail Sender ID**
     - The **Escalation Mail ID/Mobile number** (if any)
     - The **Type of notification** (**New** or **Complete List**)

- The **Message Mode** (**HTML** or **Text**)
- Whether or not to **Include measure details in mail alerts**
- Whether or not to **Include detailed diagnosis in mail alerts**
- **Time Zone** configuration
- Email/SMS alerts and escalation alerts during shift periods (if specified)

- By default, the original user's **Time Zone** will be chosen here. You can, if you need, set a different **Time Zone** for the clone.
- By default, the original user's choice of **Language** will be chosen here. If needed, you can pick a different language from this list for the clone.

4. If the **User authentication** mode is set to **Domain** on the other hand, you will have to proceed as described below:

- Pick the **Domain** to which the user belongs. If, at the time of registering that domain with the eG Enterprise system, you had set the **Save Domain User Password to the eG Enterprise System**to **No**, then, upon selecting the **Domain** here, you will be prompted to re-enter the **Domain User Password**. Without this password, the eG manager will not be able to connect to the domain server and validate domain user accounts. Provide the password and click the **Submit** button in to proceed.

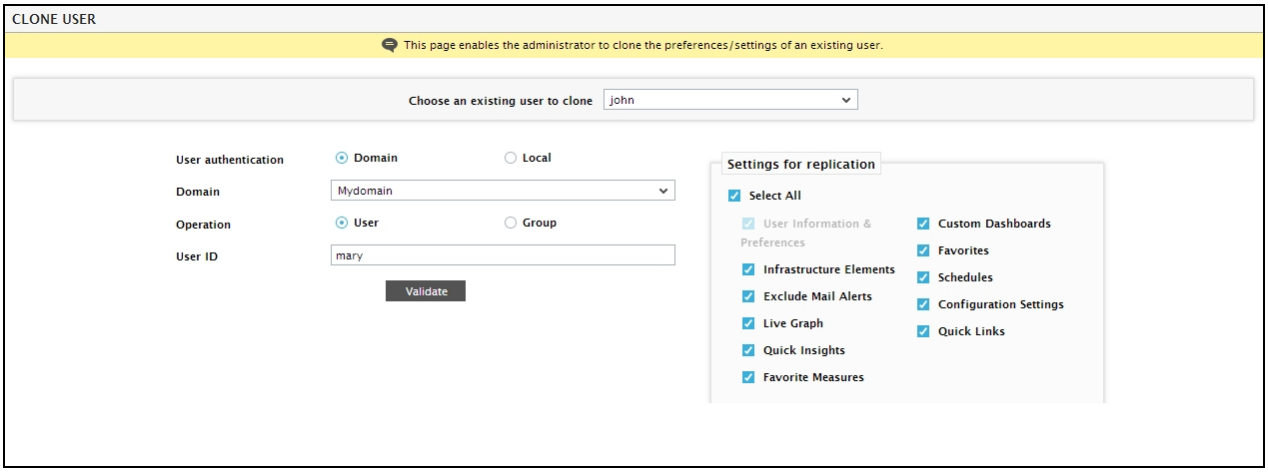

Figure 6.80: Providing the domain admin password

- Then, select an **Operation** to indicate what you want to create - a domain **User** or a domain **Group**.
- Specify the **User ID** of the new user/group, and click the **Validate** button to validate that user/group with the AD server. If the **User ID** is successfully validated, additional options will appear in the **CLONE USER** page as indicated by Figure 6.81.

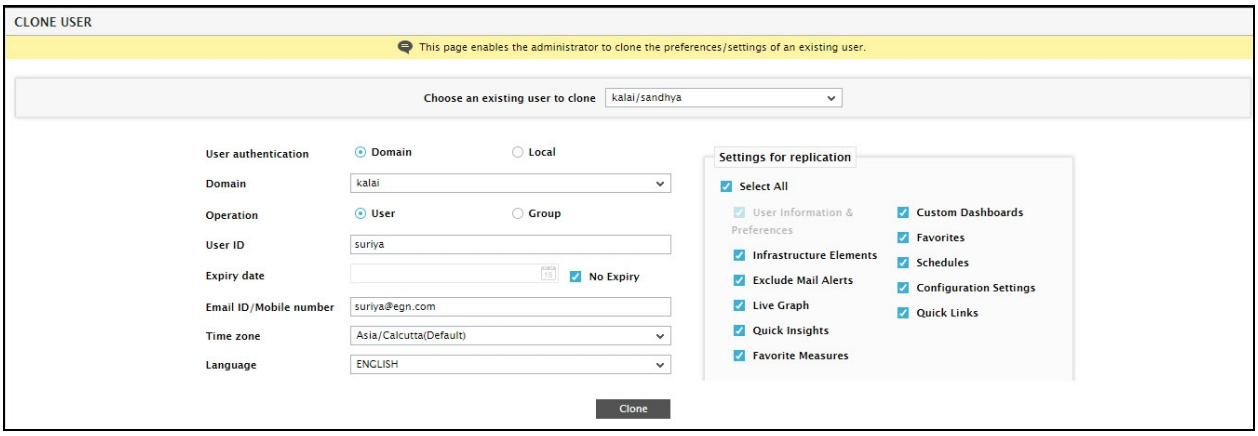

Figure 6.81: Options that appear after validating a cloned domain user

- Against **Expiry date**, indicate the date on which the user account will expire. Alternatively, you can select the **No expiry** check box if the user account being created should never expire.
- To configure the new user to receive email/SMS alerts of problem conditions in the environment, provide a comma-separated list of mail IDs and/or mobile numbers against the **Email ID / Mobile number** text box. If no **Email ID / Mobile number** is provided here, then alarm information will not be sent as emails/SMS to this user.

  Note:

  Even if the original user has not been configured to receive email/SMS alerts of issues, the 'clone' can be configured to receive the same. In this case therefore, the 'clone' will by default receive email/SMS notifications of issues, regardless of their priority - i.e., the 'clone' will be alerted to Critical, Major, and Minor problems via email/SMS.
  On the other hand, if the original user is configured to receive email/SMS alerts of issues, then, all the settings related to these alerts (with the exception of the email IDs and mobile numbers to which the alarms are to be sent) will apply to the clone as well, by default. These settings include the following:

  - The email IDs to which the email alerts are to be copied (**Cc** and **Bcc**)
  - The alarm priorities for which email/SMS alerts are to be sent
  - The **Mail Sender ID**
  - The **Escalation Mail ID/Mobile number** (if any)
  - The **Type of notification** (**New** or **Complete List**)
  - The **Message Mode** (**HTML** or **Text**)
  - Whether or not to **Include measure details in mail alerts**
  - Whether or not to **Include detailed diagnosis in mail alerts**
  - **Time Zone** configuration
  - Email/SMS alerts and escalation alerts during shift periods (if specified)

- By default, the original user's **Time Zone** will be chosen here. You can, if you need, set a different **Time Zone** for the clone.
- By default, the original user's choice of **Language** will be chosen here. If needed, you can pick a different language from this list **for** the clone.

5. Then, to indicate which settings of the original user need to be replicated to the new user, select the relevant check boxes from the **Settings for replication** section. To replicate all settings, select the **Select All** check box.

   Note:

   - If you choose to replicate the **Live Graphs**, **Quick Insight** views, and **Favorite** reports to the new user, then this implies that those live graphs, quick insight views, and favorites that have been created by the original user and those that are shared (by other users) with the original user, will now be available to the clone as well.
   - By default, the **User Information & Preferences** check box is selected. **You cannot deselect it**. This means that the following settings of the original user are automatically replicated to the clone:
   - The **User role**
   - The **Maximum Timeline for Reports**
   - The **Skins** chosen for the **Admin**, **Reporter**, **Monitor**, and **Configuration Management** modules
   - The **Monitor Home Page** setting
   - The **Refresh Frequency**
   - Whether to **Allow alarm deletion** or not
   - Whether to **Allow alarm acknowledgement** or not
   - The **Alarm display**
   - Enabling/disabling **Remote control**; if enabled, the **Remote command execution** setting

6. Finally, click the **Clone** button to clone the chosen user.

6.5 Filtering Email/SMS Alerts

By default, a user receives email/SMS alerts for all issues pertaining to all components assigned to him/her. In some circumstances, the user may not want to receive all of these alarms. For instance, in a large, multi-tier infrastructure, a user may be monitoring all the applications and network devices involved in supporting a business service. However, the user may have primary responsibility only for some of the components supporting the business service (e.g., a network administrator's primary responsibility is to monitor the network devices). In such cases, while the user may want to view the status of all the components of the business service, he/she may want to receive email or SMS alerts pertaining to specific components of the infrastructure alone (e.g., network devices).

To enable such selective alerting, eG Enterprise provides administrators with the option to configure the eG manager to **not send out email/SMS alerts related to specific layers/components/component-types/tests for specific users**.

By default, the ability to **filter mail/SMS alerts** is disabled. To enable it, do the following:

- Select the **Alert Settings** option from the **Mail Settings** menu of the **Alerts** tile.
- From the **Settings** tree in the left panel of the page that appears, select the **Filter Mail/SMS Alerts** node.
- When the **FILTER MAIL/SMS ALERTS** section opens in the right panel, set the **Allow mail/sms filter configuration** flag to **Yes**.
- Click the **Update** button to save the changes.

Once this is done, when adding/modifying a user, an additional **MAIL/SMS ALERTS FILTERING** tab page will appear.

You can filter the email/SMS alerts to be sent out to a user by clicking on this tab page, provided the following are in place:

- The user in question is configured with a mail ID and/or mobile number;
- The user is mapped to at least one infrastructure element (i.e., component/zone/segment/service);

If the user profile fulfills the above-stated requirements, then you can proceed to filter the email/SMS alerts of the user, by following the steps given below:

1. Clicking on the **MAIL/SMS ALERTS FILTERING** tab page will open Figure 6.82.

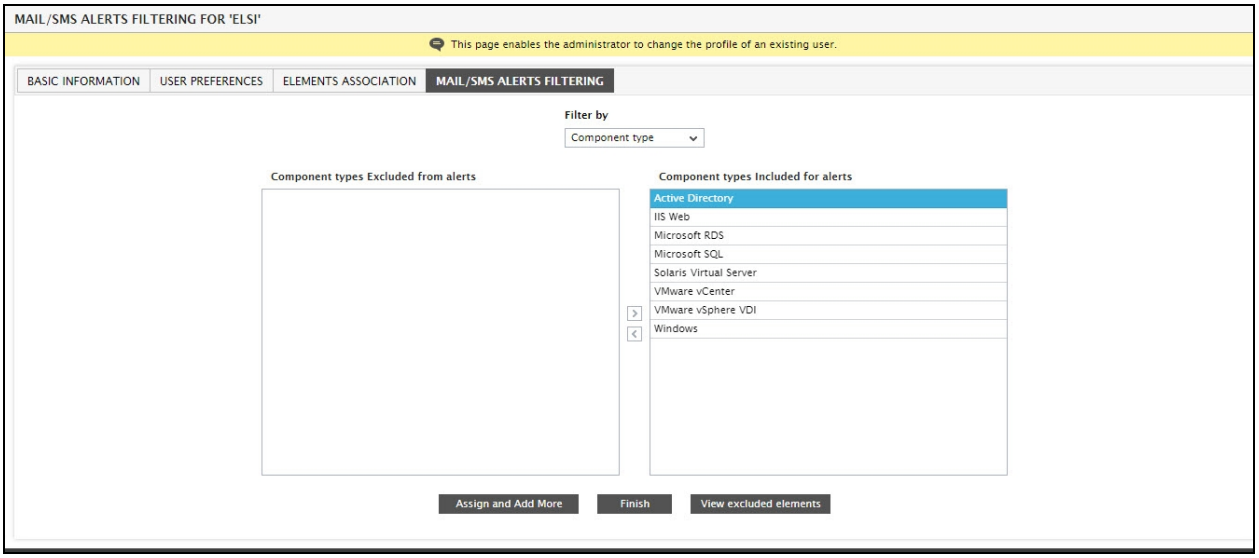

Figure 6.82: Selecting component type in the Filter By list box

2. Choose a filter type from the **Filter By** list box, based on which the email/SMS alerts to the chosen user are to be filtered. The available filter options are as follows:
    - Component type
    - Component
    - Layers
    - Tests
    - Descriptors

3. If you choose **Component type** from the **Filter By** list box, all the component types that are associated with the selected user will appear in the **Component types Included for alerts** list as depicted by Figure 6.82. Now, proceed as follows:

   - From the **Component types Included for alerts**list of Figure 6.83, select the component types for which you wish to suppress the mail alerts, and then click on the **<** button.
   - Clicking on **<** button will list all the chosen components in the **Component types Excluded from alerts** list box as depicted by Figure 6.83. You can even double-click on a component-type in the **Component types Included for alerts** list to immediately transfer it to the **Component types Excluded from alerts** list.

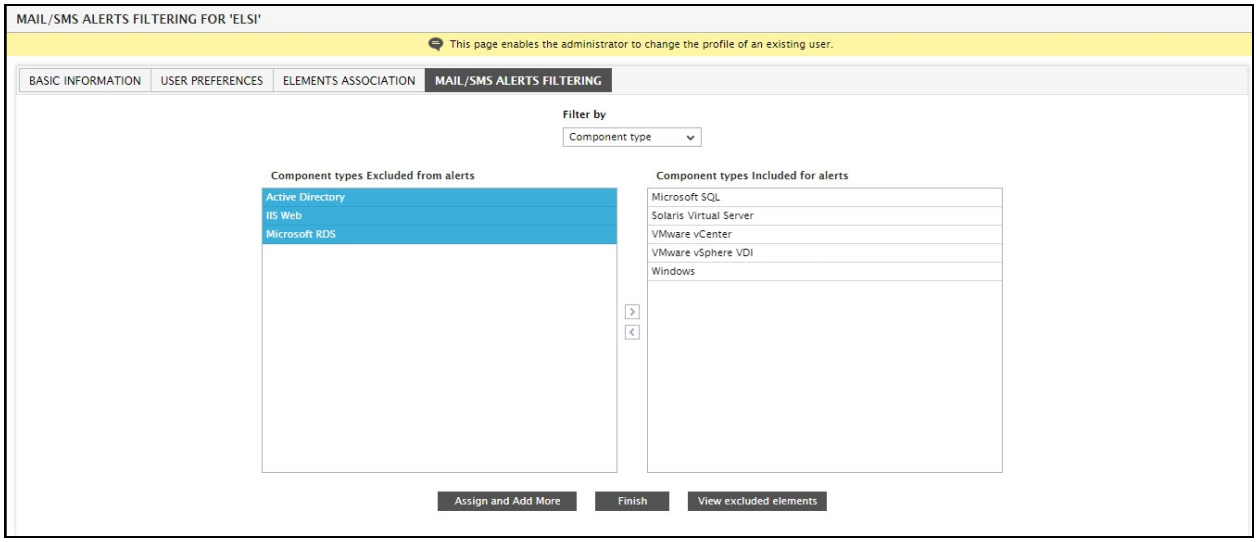

Figure 6.83: Excluding the specific component type

   - To include the excluded components, select the components from **Component types Excluded from alerts** list box and click on the **>** button. Alternatively, you can double-click on a component-type in the **Component types Excluded from alerts** list to immediately transfer it to the **Component types Included for alerts** list.
   - Then, click on the **Finish** button if you are done with filtering. On the other hand, if you want to proceed with excluding more elements, click on the **Assign and Exclude More** button.

4. If you want to make sure that the user does not receive email/SMS alerts for one/more components of a specific type, then, choose **Component** from the **Filter By** list box.

   Note:

   If email/SMS alerts for a particular component-type have already been excluded for the chosen user by selecting **Component type** from the **Filter By** list box, then this excluded component-type will not be available for selection in the **Component type** list that appears upon selecting **Component** from the **Filter By** list.

5. Then, proceed as follows:

   - If hundreds of components have been managed in a target environment, then selecting the specific components for which mail/SMS alerts should not be sent could prove to be a tedious task. To

minimize the time involved in this exercise, eG Enterprise system allows you to additionally filter the components list on the basis of the option chosen from the **View By** list. By default, the **Component** option is chosen from the **View By** list. In this case therefore, all managed component types will be available for selection in the **Component type** list (seeFigure 6.84). On the other hand, if you choose the **Segment, Service, or Zone** option from the **View By** list, you will have to pick a particular **Segment, Service, or Zone** (as the case may be) to which the components for which mail/SMS alert filters are to be defined. In this case, the **Component type** list will be populated with all the component types that are part of the chosen **Segment, Service, or Zone**.

Note:

The **View By** list will appear only when the chosen user is associated with a **Service**, **Segment** or **Zone.**

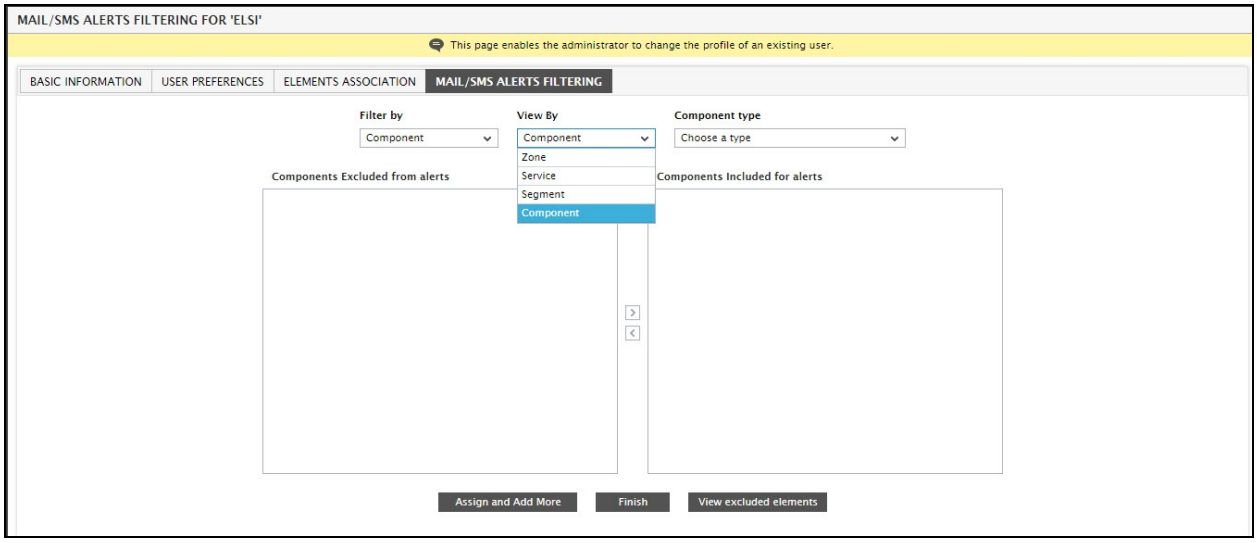

Figure 6.84: Selecting an option from the 'View By' list

- From the **Component type** list, choose a component type. Doing so will list all the components of that type in the **Components Included for alerts** list box. Now, select the components to be excluded from the this list box, and click on the **<** button. You can even double-click on a component in the **Components Included for alerts** list to immediately transfer it to the **Components Excluded from alerts** list. Finally, click on the **Finish** button. If you want to add more elements for which alerts are to be filtered, click on the **Assign and Exclude More** button instead (see Figure 6.84).

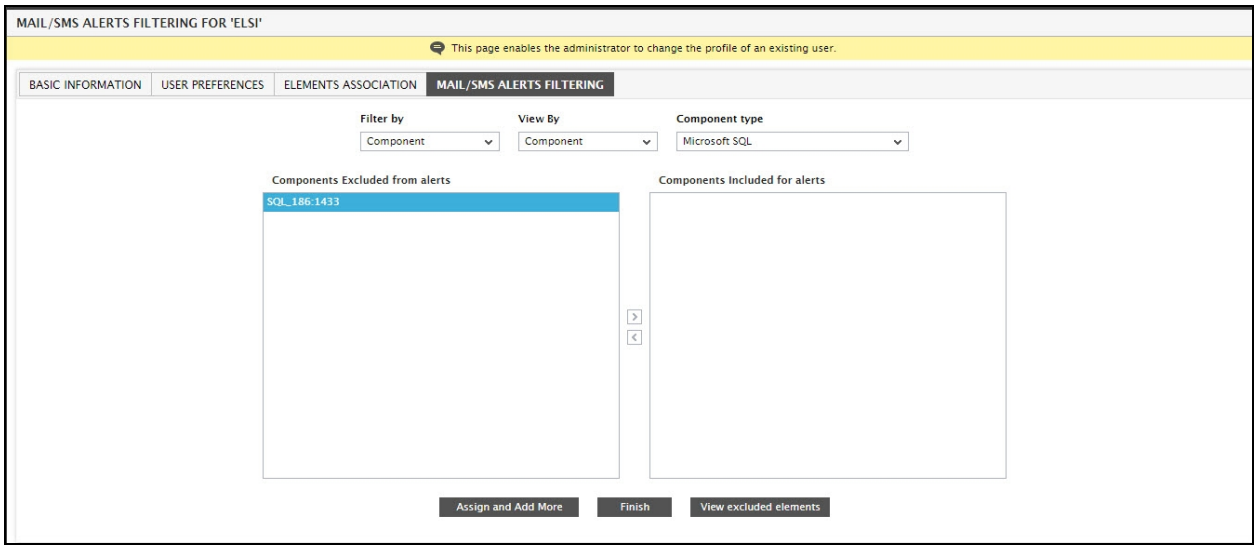

Figure 6.85: Excluding mail/SMS alerts for specific components of a chosen type

6. Likewise, you can suppress email/SMS alerts for specific layers for a user. For this, choose **Layers** from the **Filter By** list box. To further narrow-down the **Component type** list, pick the **Component, Segment, Service, or Zone** option from the **View By** list. The **Component type** list will be populated according to the option chosen from the **View By** list. Select any component-type from the **Component type** list (see Figure 6.86).

   Note:

   If email/SMS alerts for a particular component-type have already been excluded for the chosen user by selecting **Component type** from the **Filter By** list box, then this excluded component-type will no longer be available for selection in the **Component type** list that appears upon selecting **Layers** from the **Filter By** list.

7. This will list all the layers of the chosen component type in the **Layers Included for alerts** list box. Now, proceed as follows:

   - If you select **All** from the **Component type** list box, then the layers pertaining to all assigned components will be available for selection in the **Layers Included for alerts** list.

     Note:

     If email/SMS alerts for a particular component-type have already been excluded for the chosen user by selecting **Component type** from the **Filter By** list box, then all the layers pertaining to this excluded component-type will no longer be available for selection in the **Layers Included for alerts** list, if the **Component type** is set to **All**.

   - Next, select the layer to be excluded from the **Layers Included for alerts** list and click on the **<** button. Alternatively, you can double-click on a layer in the **Layers Included for alerts** list to immediately transfer it to the **Layers Excluded from alerts** list. Click on the **Finish** button to exit the mail/SMS alert filtering process, or click the **Assign and Exclude More** button to exclude more elements (see Figure 6.86).

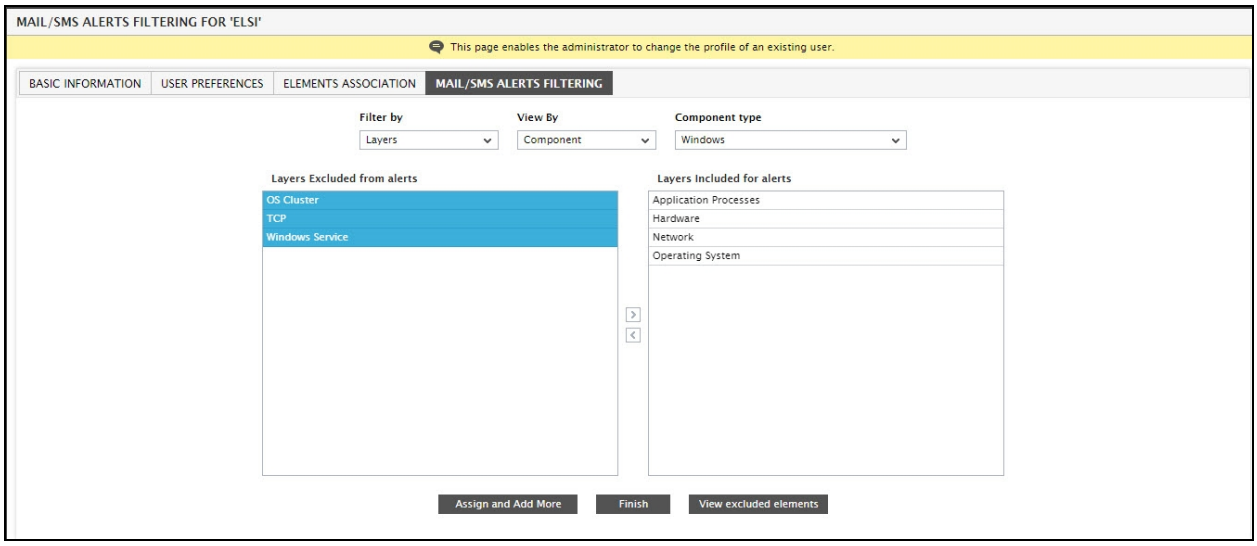

Figure 6.86: Excluding specific layers of a component type

8. Similarly, you can suppress email/SMS alerts for specific tests for a chosen user. For this, choose **Tests** from the **Filter By** list box, pick any option from the **View By** list to further filter the **Component type** list, and then select any component-type from the **Component type** list (see Figure 6.87).

   Note:

   If email/SMS alerts for a particular component-type have already been excluded for the chosen user by selecting **Component type** from the **Filter By** list box, then this excluded component-type will no longer be available for selection in the **Component type** list that appears upon selecting **Tests** from the **Filter By** list.

9. This will list all the tests of the chosen component type in the **Tests Included for alerts** list box. Now, proceed as follows:

   - If you select **All** from the **Component type** list box, then the tests pertaining to all assigned components will be available for selection in the **Tests Included for alerts** list.

     Note:

     If email/SMS alerts for a particular component-type have already been excluded for the chosen user by selecting **Component type** from the **Filter By** list box, then all the tests pertaining to this excluded component-type will no longer be available for selection in the **Tests Included for alerts** list, if the **Component type** is set to **All**.

   - Next, select the test to be excluded from the **Tests Included for alerts** list and click on the **<** button. Alternatively, you can double-click on a layer in the **Tests Included for alerts** list to immediately transfer it to the **Tests Excluded from alerts** list. Click on the **Finish** button to exit the mail/SMS alert filtering process, or click the **Assign and Exclude More** button to exclude more elements (see Figure 6.87).

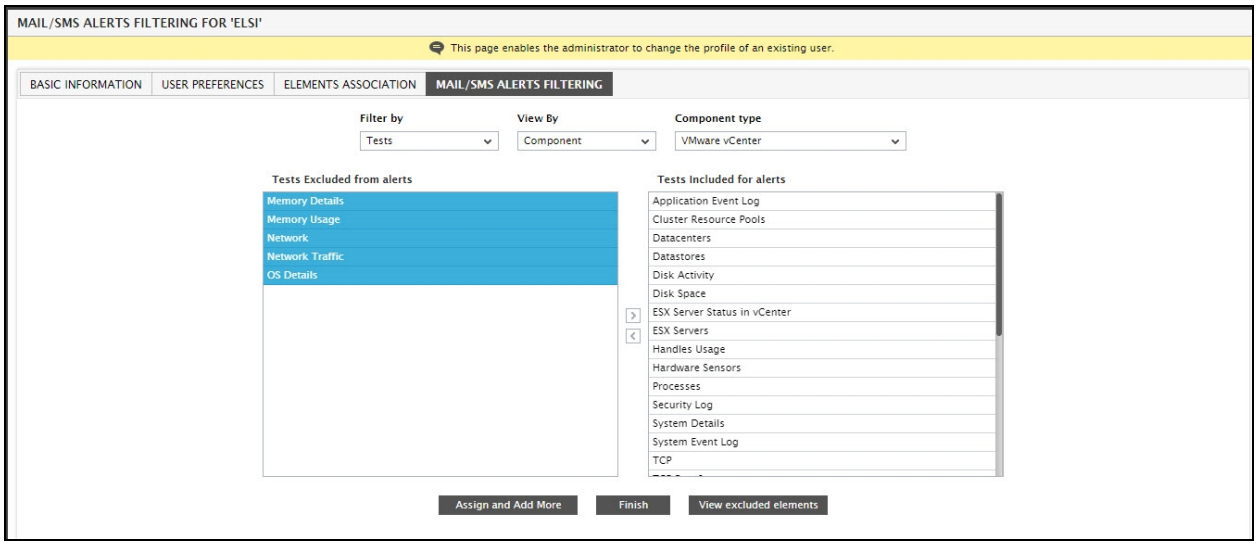

Figure 6.87: Excluding specific tests of a component type

10. In the same, you can make sure that email/SMS alerts are not sent out for specific descriptors of a test. For this, choose **Descriptors** from the **Filter By** list box. To further narrow-down the **Component type** list, pick the **Component, Segment, Service, or Zone** option from the **View By** list. The **Component type** list will be populated according to the option chosen from the **View By** list. Select any component-type from the **Component type** list (see Figure 6.86).

    Note:

    If email/SMS alerts for a particular component-type have already been excluded for the chosen user by selecting **Component type** from the **Filter By** list box, then this excluded component-type will no longer be available for selection in the **Component type** list that appears upon selecting **Descriptors** from the **Filter By** list.

11. This will list all the descriptor-based tests associated with the chosen component type in the **Test** list and all managed components of the chosen type in the **Component** list (see Figure 6.88). Pick a **Test** and a **Component** from the respective lists. Then, proceed as follows:

    - All the descriptors that are currently active for the chosen **Component** and **Test** will be available for selection in the **Descriptors Included for alerts** list.

      Note:

      If email/SMS alerts for a particular component-type have already been excluded for the chosen user by selecting **Component type** from the **Filter By** list box, then all the descriptors pertaining to this excluded component-type will no longer be available for selection in the **Descriptors Included for alerts** list, if the **Component type** is set to **All**.

    - Next, select the descriptor to be excluded from the **Descriptors Included for alerts** list and click on the **<** button. Alternatively, you can double-click on a layer in the **Descriptors Included for alerts** list to immediately transfer it to the **Descriptors Excluded from alerts** list. Click on the **Finish** button to exit the mail/SMS alert filtering process, or click the **Assign and Exclude More** button to exclude more elements (see Figure 6.86).

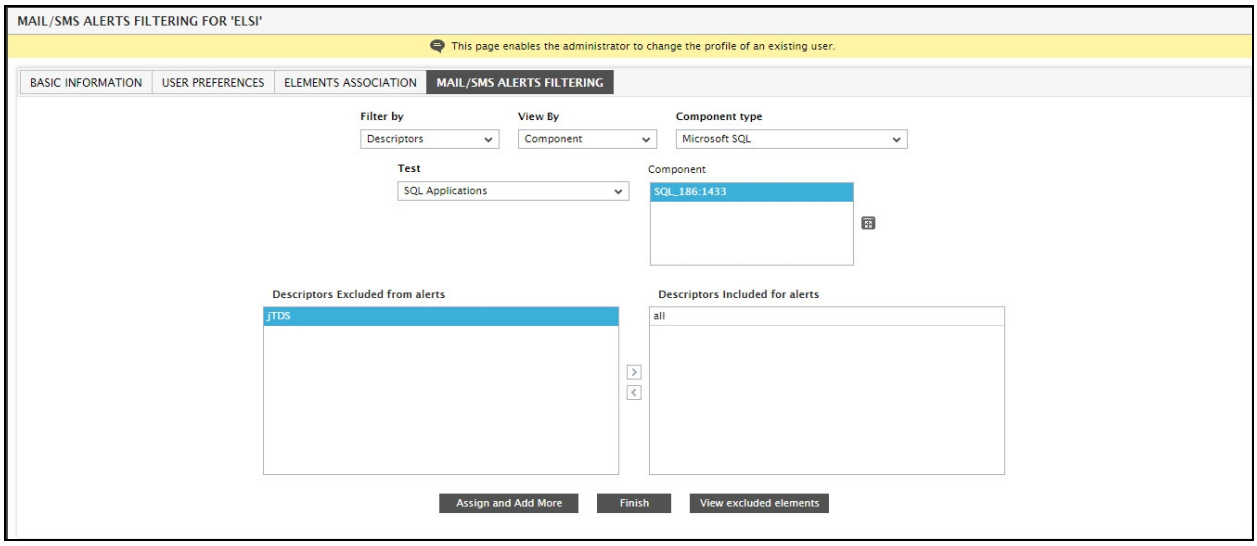

Figure 6.88: Excluding specific layers of a component type

6.6 Deleting Users

Figure 6.89 depicts how an existing user of eG Enterprise can be removed. This page can be obtained by selecting the **Delete Users** option from the **User Management** tile. In this page, the list of users currently existing in eG Enterprise (except the default *admin* and *supermonitor* users) is displayed. Each user id displayed is accompanied by the corresponding user role. By selecting one or more user accounts and clicking the **Delete** button, the corresponding user accounts can be deleted.

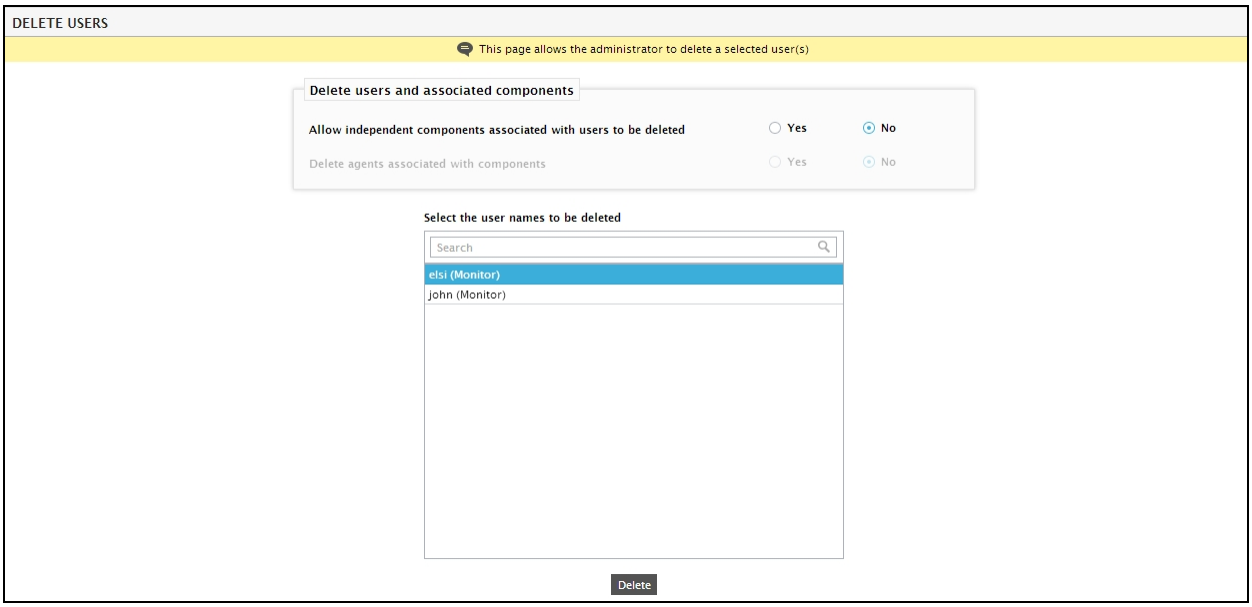

Figure 6.89: Deleting an existing user

If the user list is very long, then deleting a particular user would pose quite a challenge. You can then use the **Search** text box, wherein you can specify the whole/part of the user name to be deleted. Then, click on the

'magnifying glass' icon next to the text box. Doing so automatically selects all the user IDs that embed the specified search string. You can delete all the selected user IDs, or deselect a few of them, based on your requirements.

By default, when a user is deleted from this page, the user account is deleted. However, the components mapped to that user or the agents monitoring these components are not deleted. This is because the **Allow components associated with users also to be deleted** flag in Figure 6.89 is set to **No** by default. Likewise, the **Delete agents associated with components** flag in Figure 6.89 is also disabled **by** default.

Sometimes, administrators may prefer to have the components and agents associated with a user to be automatically deleted when the corresponding user account is deleted. For example, when the eG manager is being deployed by a managed service provider (MSP), each user account corresponds to a customer of the MSP and typically, the components assigned to a particular user are not assigned to other users. In such cases, when a user account is deleted, the administrator has to manually locate the components assigned to the user and delete or unmanage these components one by one. This can be a laborious and monotonous task. Likewise, external or remote agents may be configured specific to each user and these agents also have to be decommissioned when the user account is being deleted. eG Enterprise simplifies this task by providing the administrator with the option to automatically delete the components and agents associated with a specific user account when the user account is being deleted. To achieve this, do the following:

1. First, set the **Allow components associated with users also to be deleted** flag in Figure 6.90 to **Yes**. This will enable the **Delete agents associated with components** flag. To automatically delete the agents associated with the components monitored by the user being deleted, set this flag also to **Yes**.

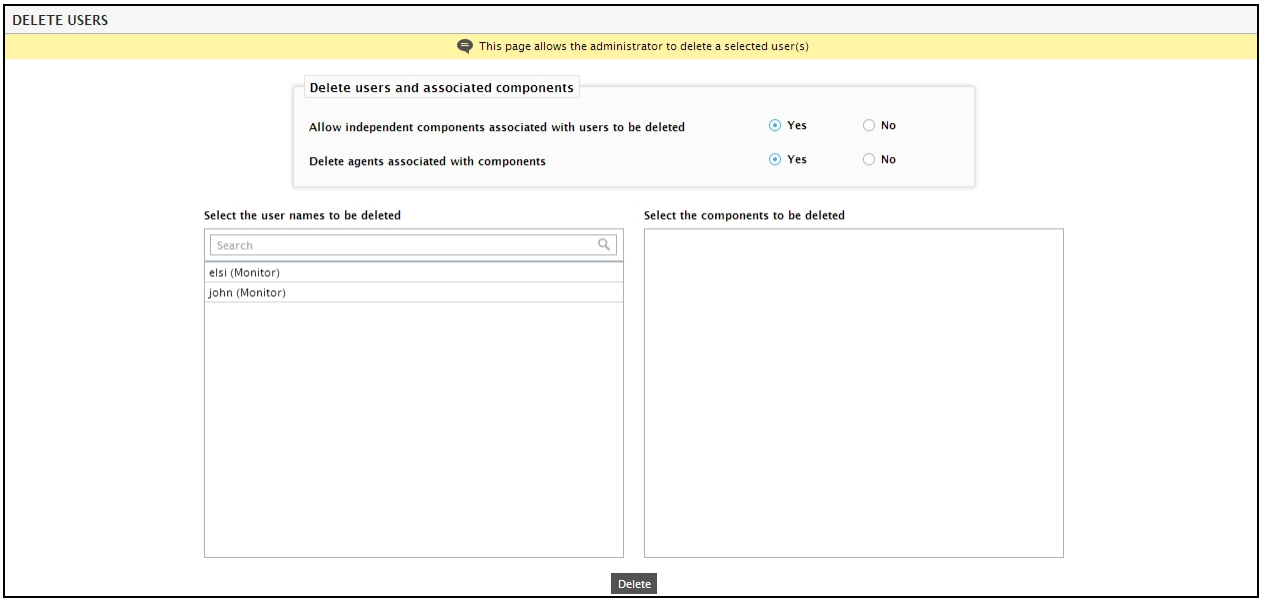

Figure 6.90: Deleting the one/more components associated with a user, when deleting the user

2. All the user accounts registered with the eG Enterprise system will then be listed in the **Select the user names to be deleted** list in Figure 6.90. Selecting a user name from this list box, will automatically display the *independent components* (i.e., the components that are not included in a segment/service/zone) associated with that user in the **Select the components to be deleted** list. From here, you can select

one/more components that need to be deleted along with the chosen user profile. Similarly, you can choose multiple users and multiple components for deletion.

3. If no independent components are associated with the user chosen for deletion, then, the following message will appear, as soon as the user is chosen.

Figure 6.91: A message box indicating that no independent components are associated with a chosen user

4. To delete the selected users and components, click on the **Delete** button in Figure 6.90.
5. If the **Delete agents associated with components** flag is set to **Yes**, then clicking the **Delete** button in Figure 6.90 will also delete the external and remote agents monitoring the components chosen from the **Select the components to be deleted** list.

Note:

- Only *independent components* assigned to a user can be deleted along with the user profile.
- If a single component is assigned to more than one user, then eG Enterprise will not even make that component available for deletion until all associated users are chosen for deletion. For instance, if the component, *apache:80*, is assigned to both user *john* and user *elvis*, then this component will not be available for deletion in the **Select the components to be deleted list** list, if only user *john* is selected from the **Select the user names to be deleted** list. Both *john* and *elvis* will have to be chosen for deletion, for *apache:80* to be listed in the **Select the components to be deleted list** list.
- The eG Enterprise system deletes only those remote and external agents that have been assigned 'exclusively' to the deleted component. In other words, if the same remote/external agent is mapped to more than one component, and such a component(s) is not assigned to the user being deleted, then such an agent will not be deleted by the eG Enterprise system, even if the **Delete agents associated with components** flag is set to **Yes**.

6.7 Changing the User Profile

Figure 6.92 displays the page that enables the administrator to change the profile of a user. To access this page, select the **Modify User** option from the **User Management** tile. In the **MODIFY USER** page, select the **User ID** of the user whose profile you want to modify. The details of the chosen user will then be displayed. You can now make changes to the following:

- The **BASIC INFORMATION** pertaining to the chosen user, which includes the role assigned to the user, the user password (if it's a **Local** user), and expiry date of the user subscription.
- All **USER PREFERENCES** – eg., time zone, email ID/mobile number, maximum timeline of reports, etc.

- Associate new elements/disassociate associated elements using the **ELEMENTS ASSOCIATION** tab page;
- The **MAIL/SMS ALERT FILTERS**, if email/SMS filtering is enabled.

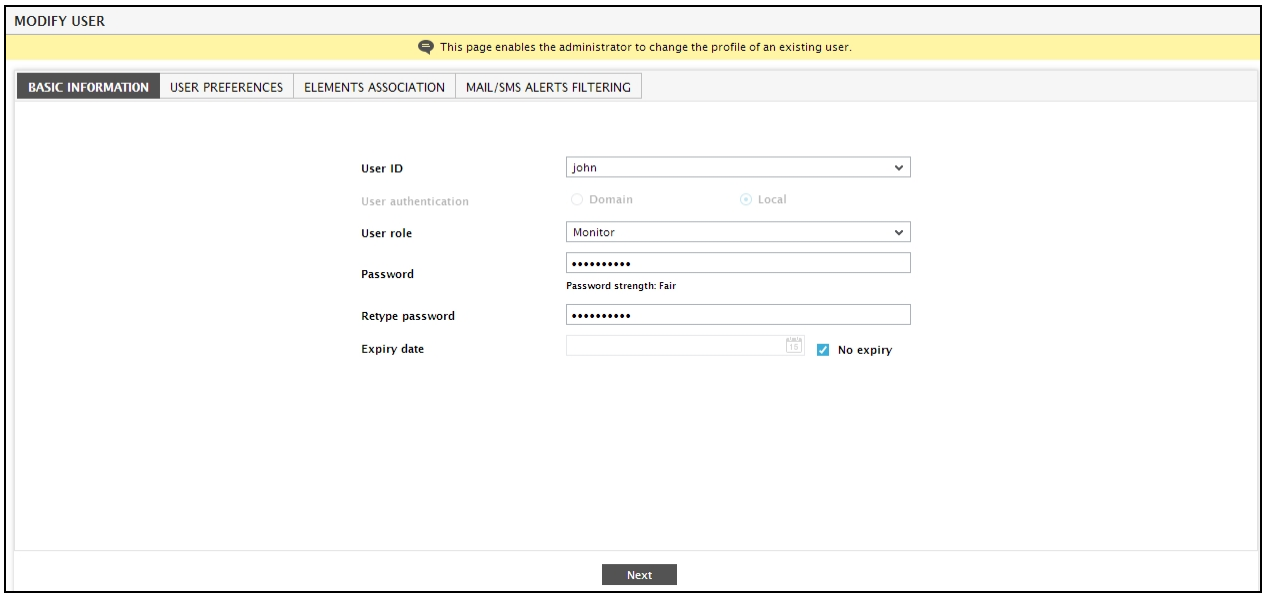

Figure 6.92: Changing the user's profile

Note:

Newly added/managed components belonging to the selected component type do not get associated with the new user immediately. Since this association is mapped as part of the discovery process, there might be a latency equal to the rediscovery period before an association between users and components is updated. If the rediscovery period has not been specified, there will be a latency equal to one day.

6.8 Locked Out User Accounts

The Account Locking feature enables the eG manager to protect the eG Enterprise system from malicious users. If tis lockout feature is enabled for the eG manager, then, if a user's attempt to login to the eG management console fails a configured number of times, the eG Enterprise system will automatically 'lock' that user account. In this case, the user will not be able to login until:

a. The expiry of a configured period of time, or;
b. The administrator manually unlocks the account using the eG administrative interface

Reference:

To know how to enable/disable the account lockout capability and configure its settings, refer to Section **5.5.2** of this document.

The **LOCKED ACCOUNTS** page in the eG administrative interface enables option (b) above. Using this page, an administrators can unlock locked user accounts. To access this page, select the **Locked Accounts** option from the **User Management** tile. Figure 6.93 then appears listing all user accounts that are currently locked.

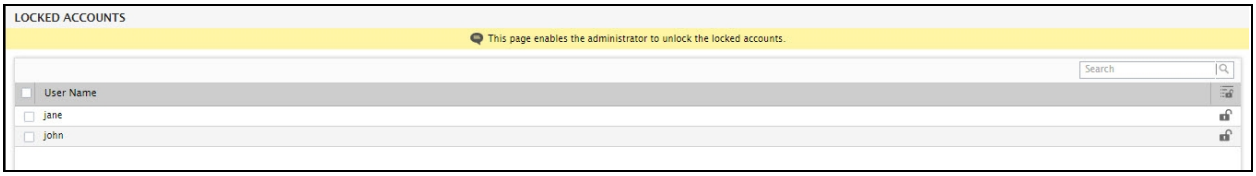

Figure 6.93: The LOCKED ACCOUNTS page listing the locked user accounts

To unlock a specific user account, just select the check box corresponding to that user account in Figure 6.93 and click the **Unlock** icon against that account, as depicted by Figure 6.94

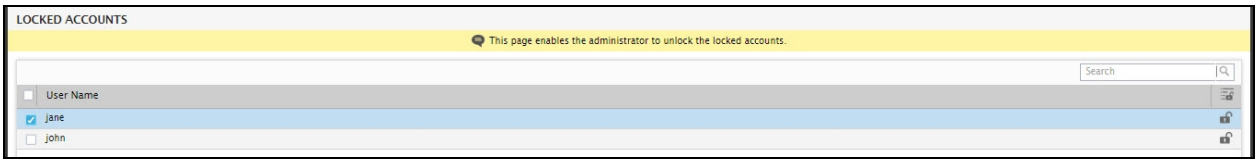

Figure 6.94: Unlocking a user account

Doing so will invoke the message box shown by Figure 6.95, which requests your confirmation to unlock the selected account. Click **Yes** in Figure 6.95 to proceed with the unlocking.

Figure 6.95: A message box requesting your confirmation to unlock the chosen account

To unlock all locked accounts at one go, simply select the check box in the header row of the table listing locked accounts in Figure 6.96. Then, click the **Unlock Selected** icon in the header row.

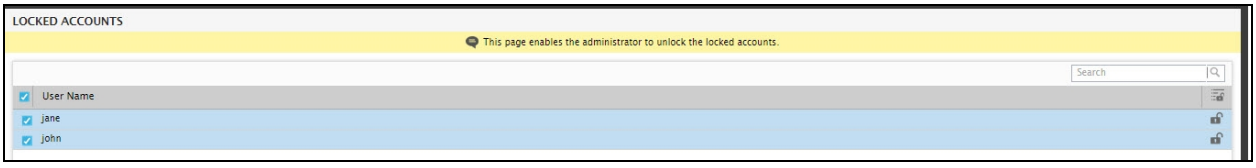

Figure 6.96: Unlocking all locked accounts simultaneously

This will invoke Figure 6.97, which requests your confirmation to unlock all the chosen accounts. Click **Yes** in Figure 6.97 to proceed with the unlocking.

Figure 6.97: A message box requesting your confirmation to unlock all chosen accounts

6.9 Viewing Details of Logged In Users

In environments where the eG management console is accessed by multiple concurrent users, administrators might want to know how many users are logged in currently, and the total session load on the eG manager. This information would enable administrators to audit user accesses to the eG manager and instantly identify unauthorized accesses (if any) that are currently active on the manager. The **USER SESSION INFORMATION** page (see Figure 6.98) that appears upon selecting the **Logged In Users** option from the **Reports** menu of the **User Management** tile, provides the information that will help administrators analyze currently active user sessions on the eG manager. This page lists the total number of active sessions and the number of 'distinct' users who are currently using the eG management console. Besides, the page also provides details of every active user session; these details include the **HOST IP** from which a user is logging in, the name of the user, the role assigned to the user, the time of login, and the last accessed time. To refresh this list so that the latest session details are available to the administrators, click on the **Refresh** button.

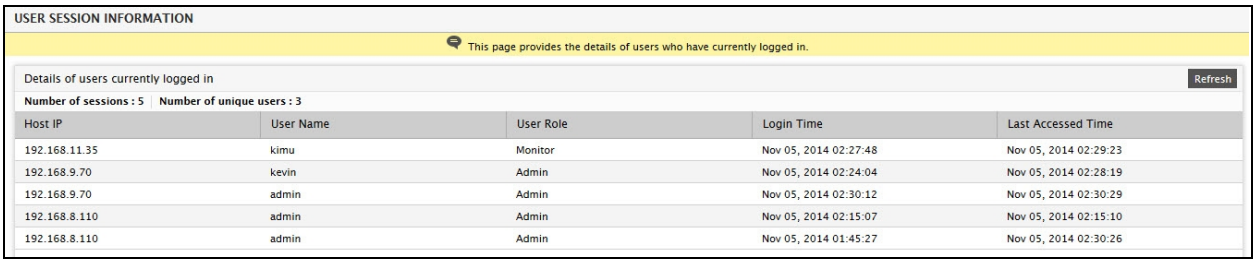

Figure 6.98: Viewing session information

6.10 Viewing User Details

The administrators can view the details of users registered with the eG manager using the **User Detail** option in the **Reports** menu of the **User Management** tile (see Figure 6.99). By default, the **User category** is set to All, indicating that the details of all configured users will be displayed here. The other options provided by the **User category** list are as follows:

- The **Nearing Expiry Users** option: Use this option to view the details of only those users who are about to expire in the next 7 days.
- The **No Expiry Users** option: Use this option to view the details of only those users whose subscription will never expire.

- The **Expired Users** option: Use this option to view the details of only those users whose subscription has already expired.
- The **Active Local Users** option: Use this option to view the details of 'local' users whose subscription is still active on the eG management console.
- The **Active Domain Users** option: Use this option to view the details of 'domain' users whose subscription is still active on the eG management console.
- The **Active Domain Groups** option: Use this option to view the details of users belonging to domain groups whose subscription is still active on the eG management console.

Once a **User category** is chosen, details of all users who belong to that category will be displayed. These details include the user role, the date at which the user subscription will expire, the time zone, and date format assigned to the user. Clicking on the '+' symbol alongside a user name will reveal the zones, services, segments and the components included in the user's view. Using the **Modify** icon (i.e., the 'pencil' symbol) against a user name in Figure 6.99, you can modify the profile of a user.

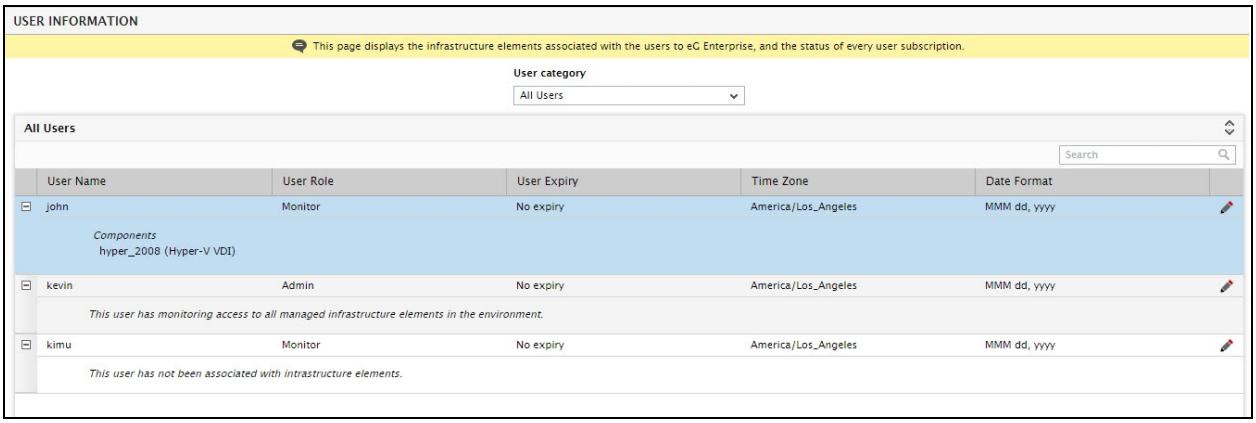

Figure 6.99: The eG user interface depicting the details of the existing users

6.11 User Registration Report

Figure 6.100 depicts how the administrator can view the reports pertaining to the users who have registered to the eG Enterprise system. To access this page, select the **Registration History** option from the **Reports** menu of the **User Management** tile.

The **Last week** option in Figure 6.100 enables the administrator to view the list of users who have registered to the eG Enterprise system in the previous week. The Last month option enables the administrator to view the list of users who have registered to the eG Enterprise system in the previous month. The **Any period** option enables the administrator to view the list of users who have registered to the eG Enterprise system in the specified period. Clicking on the **Go** button displays the report corresponding to the chosen option. The registration date as well as the validity date corresponding to the user is displayed along with the user ID (see Figure 6.100). You can also print these details by clicking on the **Print** button.

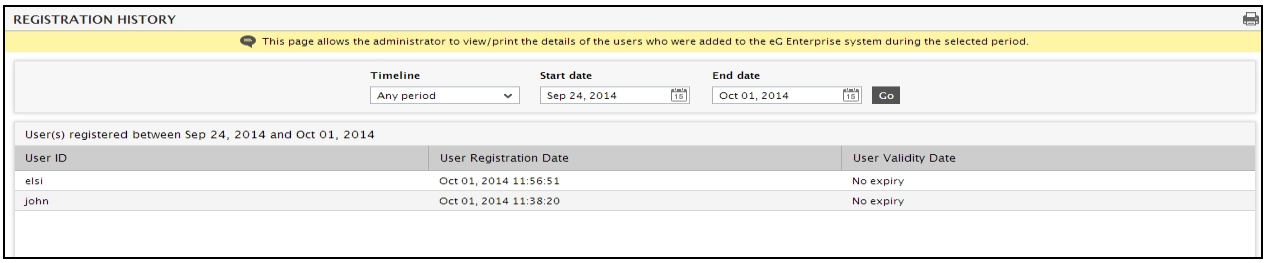

Figure 6.100: A sample user registration report

6.12 Account Expiry Report

Figure 6.101 depicts how the administrator can view the list of users of the eG Enterprise system whose subscription period exhausts within a week. The details pertaining to the user such as the user ID, validity date, and the number of days left for the subscription to expire are presented in the **EXPIRY INFORMATION** page that appears when the **Account Expiry** option is chosen from the **Reports** menu of the **User Management** tile.

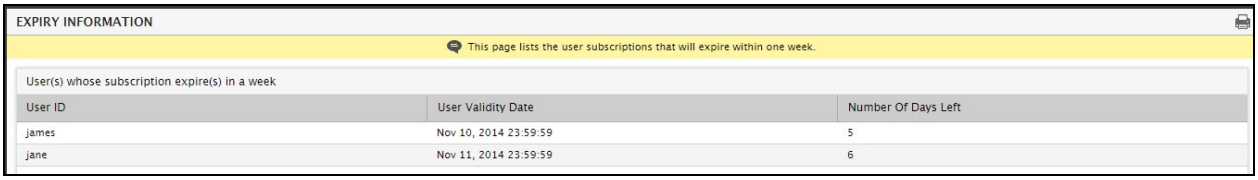

Figure 6.101: A sample subscription status report

7

Managing Components Using the eG Admin Interface

The flowchart below explains the basic steps for managing components using the eG administrative interface.

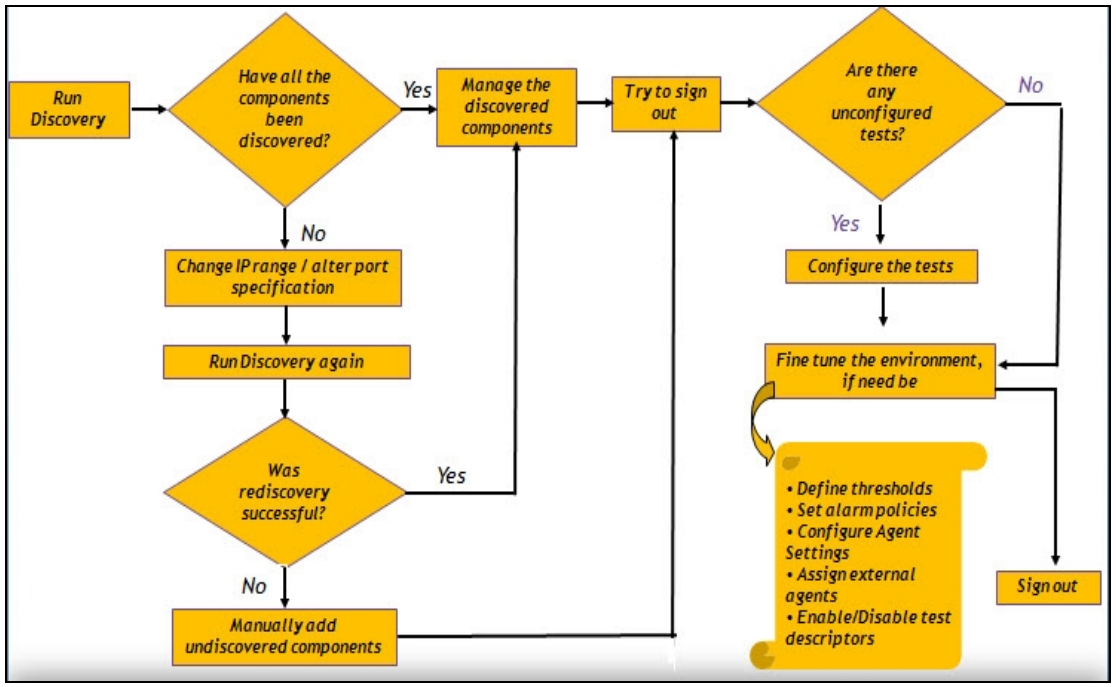

Figure 7.1: The Manage/Unmanage flow chart

Step 1: Start Discovery.

Step 2: Check if all the components are automatically discovered. If the components are not automatically discovered, then change the IP range/ Alter port, and run discovery again.

Step 3: Now manage the discovered components. The Discovered components have to be explicitly managed for the eG agents to monitor them.

Step 4: If the components are still not discovered, then manually add the components to the eG Enterprise system.

Step 5: Try to sign out of the eG administrative interface.

Step 6: The unconfigured test window appears. Now, configure the tests manually.

Step 7: Every test takes certain input parameters for execution; these parameters vary according to the purpose of the test and may typically indicate:

- How often the test is to be executed;
- On which host should the test run;
- What should be monitored – for eg., a test that checks the availability of TCP ports on a host will take a list of TCP ports to check for as a parameter;
- Credentials required (if any) for accessing the target host or for running certain commands built-into the test, and MORE.

Step 8: Finally, signout of the eG administrative interface.

Now let us discuss each step in detail in this chapter and the ones that follow.

7.1 Discovering Components

eG Enterprise is capable of automatically discovering the components in the target environment. To perform this discovery, administrators can use either the central eG manager or the eG agents installed on the target hosts. Both these discovery methodologies are briefly discussed below:

- **Discovery using the eG manager**: If the eG manager is used to discover the targets for monitoring, then such a discovery will typically be based on the port number(s) on which the components are listening. In case of network devices or components that do not listen on any port, the eG manager uses SNMP for discovery. When discovery is triggered, the eG manager uses a unique port scanning technique/SNMP (as the case may be) to discover all the components/network devices that are in configured IP ranges, regardless of whether agents are installed on them or not. Typically, an agent is installed on a host only if the performance of one/more applications executing on that host interests the administrator.

  Since the eG manager discovers components unmindful of whether agents are monitoring them or not, this discovery process might end up discovering a wide variety of applications that the administrator might not even be interested in monitoring! Secondly, manager discovery is based on the assumption that the eG manager has access to all the components in the target environment. In the real world however, this might not be the case. The manager could be behind a firewall, and might hence be denied access to many/all components in the target environment; in this case therefore, with manager-based discovery, a number of components could go undiscovered.

- **Discovery using the eG agents**: As stated earlier, typically, an agent is installed on a host only when an administrator is interested in monitoring one/more applications executing on that host. In large IT environments with thousands of components, it is often difficult for administrators to manually track where agents are installed, and manage only the corresponding hosts and applications. In such environments therefore, administrators might prefer to run a discovery procedure that automatically discovers only those components that they are “interested in monitoring” - i.e., an auto-discovery procedure that can discover only those applications which are executing on the hosts where agents are installed; this ensures that the eG administrative interface is not crowded with a wide variety of applications that an administrator might not even be interested in monitoring. This can be achieved only **if discovery is performed using the eG agents**.

  The agents use many intelligent techniques to discover the applications that may require monitoring. Typically, the agent employs a simple port scanning technique to discover the applications executing

on its host. In addition, the agent is also capable of logging into (using configured login credentials) an application so discovered to gather information related to the remote applications it interacts with frequently, so that such remote applications are discovered as well. This way, the agent automatically discovers applications running on even those hosts that do not have an agent installed on them! Command-based discovery of applications is also supported, which is useful while discovering non-port-based applications. However, note that **agents cannot automatically discover network devices operating in an environment**.

Since every agent performs the discovery and communicates the results to the eG manager, the location of the eG manager will not in any way impact the discovery process - i.e., even if the eG manager is behind a firewall and is not able to access the components in the target environment, the eG agent will be able to promptly detect the additions/removals in the environment everytime it rediscovers, and will be able to update the eG manager with this knowledge.

Besides individual components, the agent can also be used to automatically discover the component topology - i.e., the physical/logical inter-relationships that exist between discovered components. This ability, when enabled, helps administrators draw an almost accurate segment topology with minimal user intervention and time! Also, since the auto-discovered topology depicts the 'real' dependencies that prevail in a 'real ' environment, it consequently enhances the eG Enterprise system's innate ability to auto-correlate performance and proactively detect performance issues.

The sections to come discuss the steps to be followed to trigger a discovery using the eG manager or the eG agents.

7.1.1 Discovering Components using the eG Manager

To configure manager-based component discovery, click on the icon available in the **Admin** tab. Then, select the **Discovery** option from the **Components** menu in the **Infrastructure** tile.

Figure 7.2 then appears. As you can see, Figure 7.2 displays two panels - a left panel comprising of a **DISCOVERY** tree structure, which consists of nodes and sub-nodes that enable you to quickly navigate the discovery-related options, and a context-sensitive right panel that changes according to the node chosen from the tree.

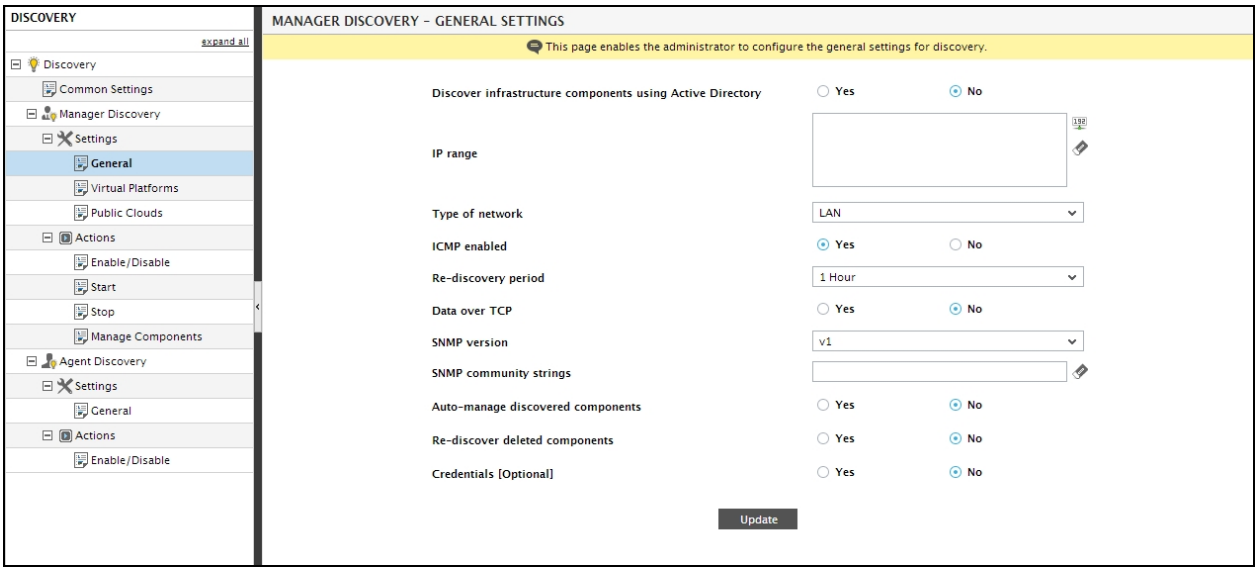

Figure 7.2: The DISCOVERY tree and the context-sensitive right panel

To hide the tree, click on the left-arrow button indicated by Figure 7.2. To unhide the tree, click on the right-arrow button that appears thereafter.

Before attempting to perform discovery using the eG manager, you will have to configure the parameters discussed in the sections that will follow.

7.1.1.1 Enabling/Disabling Manager Discovery

The first step to performing component discovery using the eG manager is to check whether/not the manager is capable of discovering components. For this, select the **Enable/Disable** sub-node of the **Actions** node in the **DISCOVERY** tree in the left panel of Figure 7.2. The right panel will then change as depicted by Figure 7.3.

Figure 7.3: Checking whether/not manager discovery is enabled

If the manager discovery capability is disabled, then the **Manager discovery** check box in Figure 7.3 will be deselected. The eG manager can be used for performing component discovery only if this capability is enabled. Therefore, to turn on this capability, select the **Enabled?** check box in Figure 7.3. A message box depicted by Figure 7.4 will then appear requesting your confirming to enable the capability. Click **Yes** in the message box to confirm enabling. Doing so will invoke a message box shown by Figure 7.5, which informs you of the success of the operation.

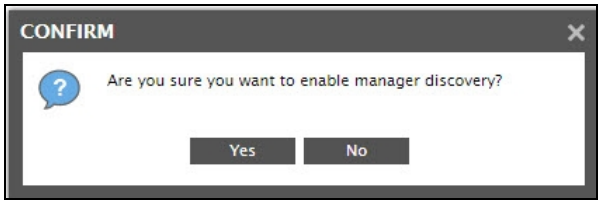

Figure 7.4: A message box requesting your confirmation to enable manager discovery

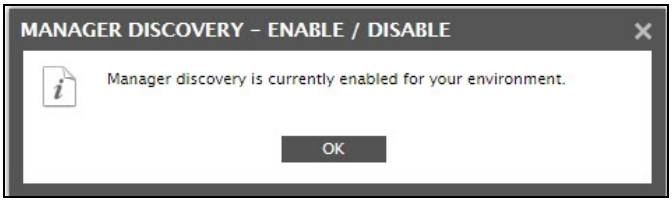

Figure 7.5: A message box that appears once manager discovery is successfully enabled

You can disable manager discovery at any point in time by simply deselecting the **Enabled?** check box in Figure 7.3.

7.1.1.2 Configuring Common Settings

Once manager discovery is enabled, proceed to define the discovery settings that are common to both the eG manager and eG agent-based discovery processes.

To configure these settings, do the following:

1. Select the **Common Settings** sub-node under the **Discovery** node of the **DISCOVERY** tree in Figure 7.6.

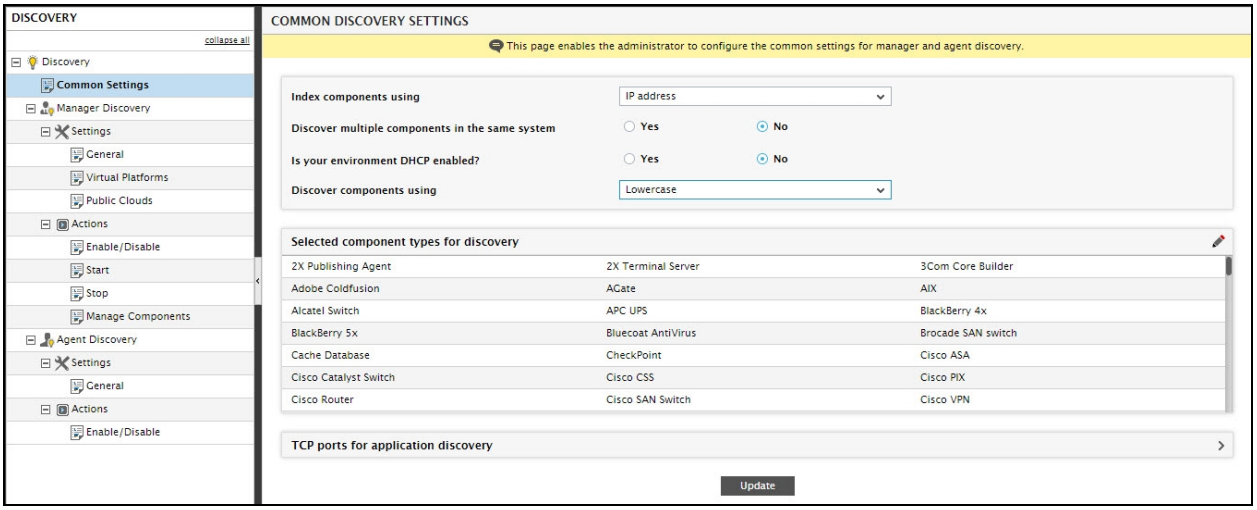

Figure 7.6: Configuring Common Discovery Settings

2. Then, in the right panel, pick an option from the **Index components using** list to indicate what you want to set as the nick name of the discovered components – the IP address of the components or the host name? In DHCP environments typically, you may want the host name to be set as the nick name of the components, as the IP address may keep changing.

3. In the eG Enterprise system, priorities are pre-assigned to different applications for the purpose of auto-discovery. By default, the eG manager/agent discovers applications based on the priorities/weightages that have been pre-configured for them. This is because, the **Discover multiple components in the same system** flag is set to **No** by default. In this case, the discovery process first looks for high priority applications on a host. As soon as high priority applications are discovered on a host, discovery will stop; this means that all low priority applications on that host will be ignored by the discovery process. However, if no high priority applications are discovered on a host, then the discovery process will look for low priority applications on that host. Weights are pre-assigned to low priority applications based on how important it is to monitor them. If only low priority applications are available on a host, then the eG manager will auto-discover only that application on the host that has been assigned the highest weightage; all other low priority applications on that host will hence be ignored. This capability ensures that eG Enterprise's discovery process only discovers the key components that administrators are likely to be interested in monitoring.

   However, if the **Discover multiple components in the same system** flag is set to **Yes**, then the eG manager will auto-discover all applications on a host, regardless of the priorities/weights assigned to them.

   Note:

   The high priority and low priority applications are pre-configured in the **eg_services.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory. While the the high priority applications are pre-defined in the **[HIGH_PRIORITY_COMPONENTS]** section of the file, the low priority applications are pre-defined in the **[LOW_PRIORITY_APPLICATIONS]** section. If you want to add a component to the list of high priority applications, then append an entry of the following format in the **[HIGH_PRIORITY_COMPONENTS]**section of the file.

   InternalName_of_ComponentType=HIGH

For instance, if you want the Citrix XenDesktop Director component to be set as a high priority component during discovery, then append the following entry to the **[HIGH_PRIORITY_COMPONENTS]**section:

Citrix_XcXenDesktop_Director=HIGH

Likewise, to set a component as a low priority component, append an entry of the following format to the **[LOW_PRIORITY_APPLICATIONS]** section of the **eg_services.ini** file:

InternalName_of_ComponentType=Weightage_Assigned

The Weightage_Assigned should be a value between 1 and 10, where 1 is the lowest and 10 is the highest. For instance, if you want the Citrix XenDesktop Director component to be set as a low priority component with the weightage 10, then the entry in the **[LOW_PRIORITY_APPLICATIONS]** section should be as follows:

Citrix_XcXenDesktop_Director=10

Finally, save the **eg_services.ini** file.

4. Next, indicate whether/not the target environment is DHCP-enabled. By default, the **Is your environment DHCP-enabled?** flag is set to **No**. If the environment is DHCP-enabled, then set this flag to **Yes**. Doing so ensures that discovery is performed using the host names of applications and not their IP addresses.

   Note:

   - If the **Is your environment DHCP-enabled?** flag is set to **Yes**, then the **Index components using** flag will automatically change to **Host name**, if it had been set to **IP address** earlier.
   - The selection from the **Index components using** drop-down and the **Is your environment DHCP-enabled** flag setting do not just apply to general component discovery, but also to the following types of discovery that can be performed using the eG manager:
     - the discovery of vSphere/ESX servers via vCenter;
     - the discovery of IBM pSeries servers via HMC server;
     - the discovery of Public clouds;
     - the discovery of RHEV servers via the RHEV manager

5. The **Discover components using** flag becomes relevant only if the **Hostname** option is chosen from the **Index components using** list. In this case, from the **Discover components using** drop-down, select the case to be applied to host/nick names during discovery. If you want host/nick names of discovered components to be in lower case only, then select the **Lowercase** option. If this is done, even if a host/nick name is in upper case during discovery, it will be automatically changed to the lower case. Select the **Uppercase** option if you want the host/nick names of discovered components to be in upper case only. By selecting this option, you can make sure that even host/nick names that are in lower case during discovery are automatically changed to the upper case. Select the **Anycase** option, if you want discovery to dynamically decide at runtime which case to apply. If this option is selected, then the very first time discovery runs post the selection, it will automatically determine whether the host/nick names discovered are in upper case/lower case. The case so determined will be automatically applied to all host/nick names that are discovered subsequently. For instance, if, soon after selecting the **Anycase** option, discovery runs and discovers a host/nick name in the upper case, the case of all host/nick names discovered subsequently will automatically change to uppercase.

6. Then, proceed to select the types of components that you want the discovery process to auto-discover. The component types selected by default will be listed in the **Selected component types for discovery** section. To override this default selection, click the **Configure components for discovery** icon (i.e., the 'pencil' icon) at the far end of the section. Figure 7.7 will then appear. Select a **Category** from Figure 7.7. All component types under the chosen **Category** that have been selected for discovery will be listed in the **Components** list of Figure 7.7. To exclude one/more components from the discovery, just deselect those components from the **Components** list as depicted by Figure 7.8. Then, click the **Apply** button.

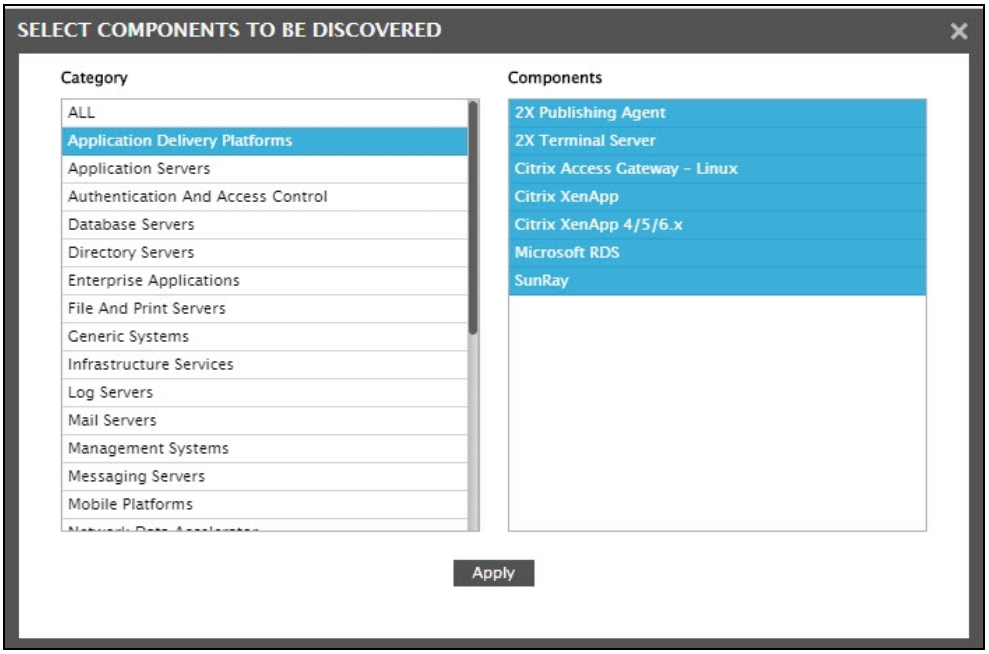

Figure 7.7: Component types selected by default

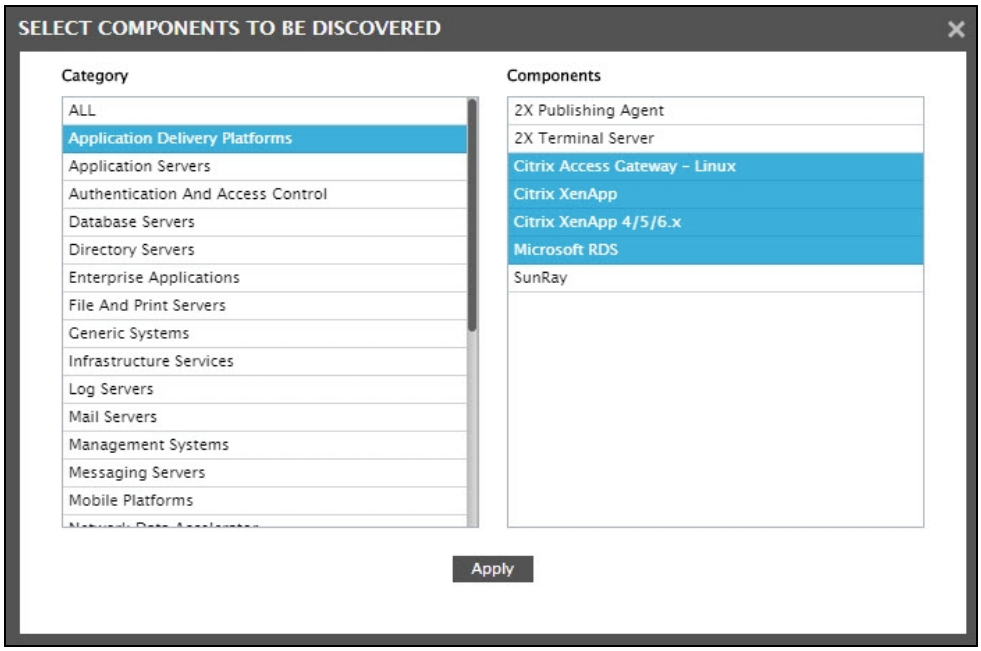

Figure 7.8: Deselecting components to be excluded from discovery

7. In the same way, you can configure components from more categories for discovery.
8. Next, proceed to configure the ports for discovery. Discovery of components is typically based on the port number(s) on which the components are listening. Each component type supported by eG Enterprise is mapped to a set of port numbers. You can, if need be, change the default port preferences mapped to each component type. For instance, the default port mapped to an Apache Tomcat server is 8080. If your environment comprises of a few Tomcat servers that listen on port 8080 and a few others that listen on port 8088, the discovery process will by default discover only those Tomcat servers that run on port 8080. To ensure that the discovery process automatically discovers the Tomcat servers that listen on 8080 and the ones that listen on 8088, do the following:
   - Click the > button at the far end of the **TCP ports for application discovery** to expand the section as depicted by Figure 7.6.
   - Keep scrolling down the **TCP ports for application discovery** section until you see the **Tomcat** server entry. Click on the port specification against **Tomcat** to edit it.
   - The ports will become editable as indicated by Figure 7.9. To make sure that the Tomcat servers listening on 8080 and 8088 are discovered, provide both port numbers as a comma-separated list (see Figure 7.9)

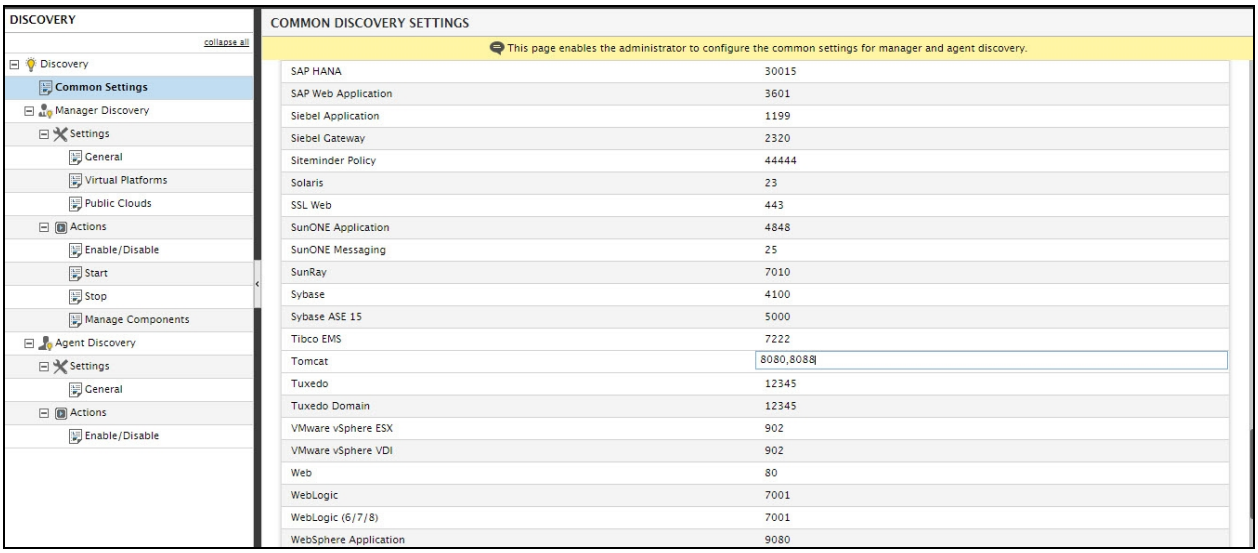

Figure 7.9: Modifying the port specification of a Tomcat server

9. Finally, click the **Update** button in the **COMMON DISCOVERY SETTINGS** panel in the right to save the changes.

7.1.1.3 Configuring General Settings

After indicating what types of components listening on which ports need to be discovered, proceed to define the **General** discovery settings.

For this, click the **General** sub-node under the **Settings** node in the **DISCOVERY** tree in the left panel of Figure 7.2. Accordingly, the right panel will display the **MANAGER DISCOVERY – GENERAL SETTINGS** page (see Figure 7.2).

To start the discovery process, you must specify the IP address range(s) that characterize the target environment in the **GENERAL SETTINGS** page. You can provide the IP range for discovery in one of the following ways:

- You can manually specify the IP range for discovery;
- You can instruct the eG manager to automatically discover the IP range from the target environment;
- If your eG manager integrates with an Active Directory server, then, you can discover the range of IP addresses available in a particular subnet of a chosen domain;

To manually specify the IP range, do the following:

- Set the **Discover infrastructure components using Active Directory** flag to **No**.
- Provide an IP address manually in the **IP range** text area as depicted by Figure 7.10. While many environments have a single IP address range, some environments may involve components at different locations, with each location using a different IP address range. To support these different types of environments, the eG administrative interface permits the administrator to enter one or more IP address ranges that characterize the target environment, one below the other.

  Note:

Both IPv4 and IPv6 ranges are supported. However, a single IP range can include either IPv4 addresses or IPv6 addresses only, and not a combination of both.

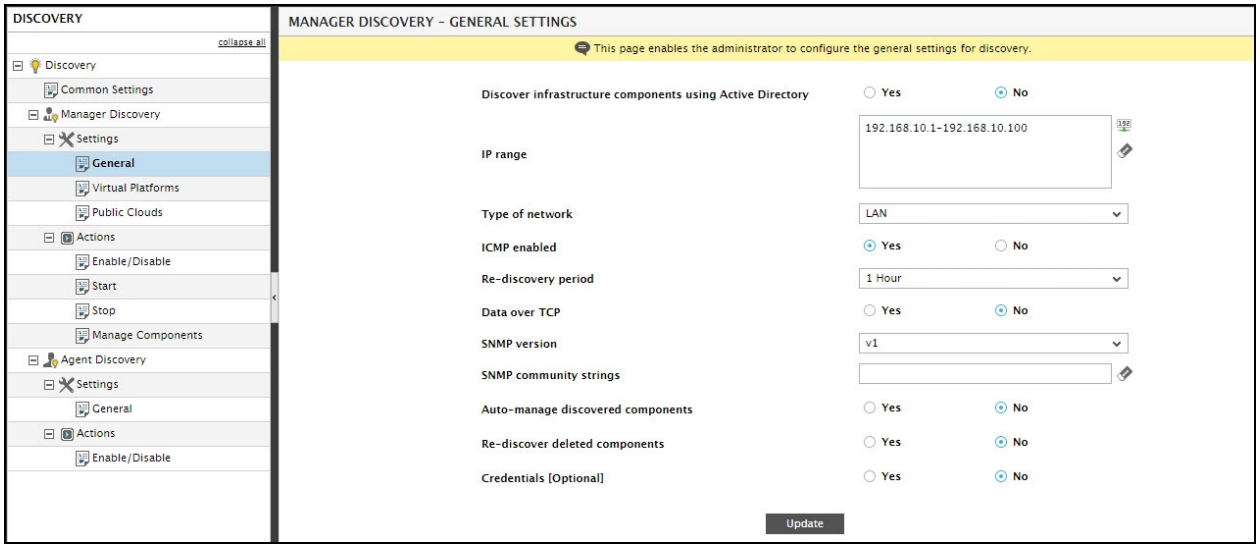

Figure 7.10: Manually specifying the IP range for discovery

To make sure that the eG manager auto-discovers the IP range available in the target environment, do the following:

- Set the **Discover infrastructure components using Active Directory** flag to **No**.
- Click the **Fetch IP range for this environment** icon next to the **IP range** text area in Figure 7.11.
- This will automatically populate the **IP range** text area with the range of IP addresses available in the target environment.

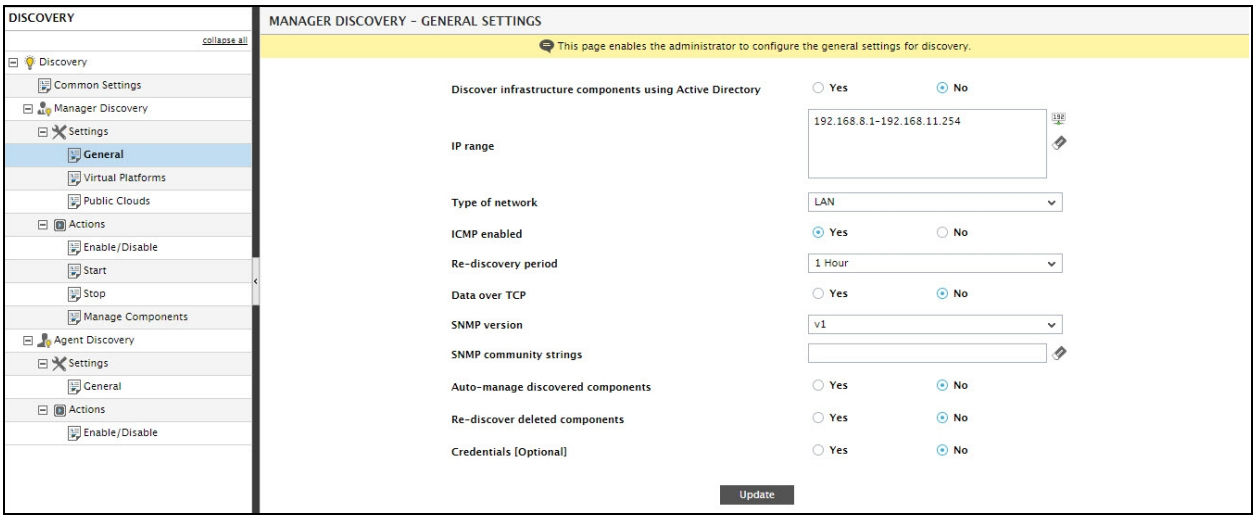

Figure 7.11: Automatically discovering the IP range of the target environment

If you want to fetch the IP range from the AD server with which the eG manager integrates, do the following:

- Set the **Discover infrastructure components using Active Directory** flag to **Yes** (see Figure 7.12).
- Pick a domain from the **Active Directory Domain** list (see Figure 7.12).
- The sites in the chosen domain will then be available for selection. **Select an Active Directory Site** from the list (see Figure 7.12).
- Next, pick one/more subnets from the **Select SubNet(s) From Site** list (see Figure 7.12).
- This will automatically populate the **IP range** text area with the range of IP addresses in the chosen subnets (see Figure 7.12).

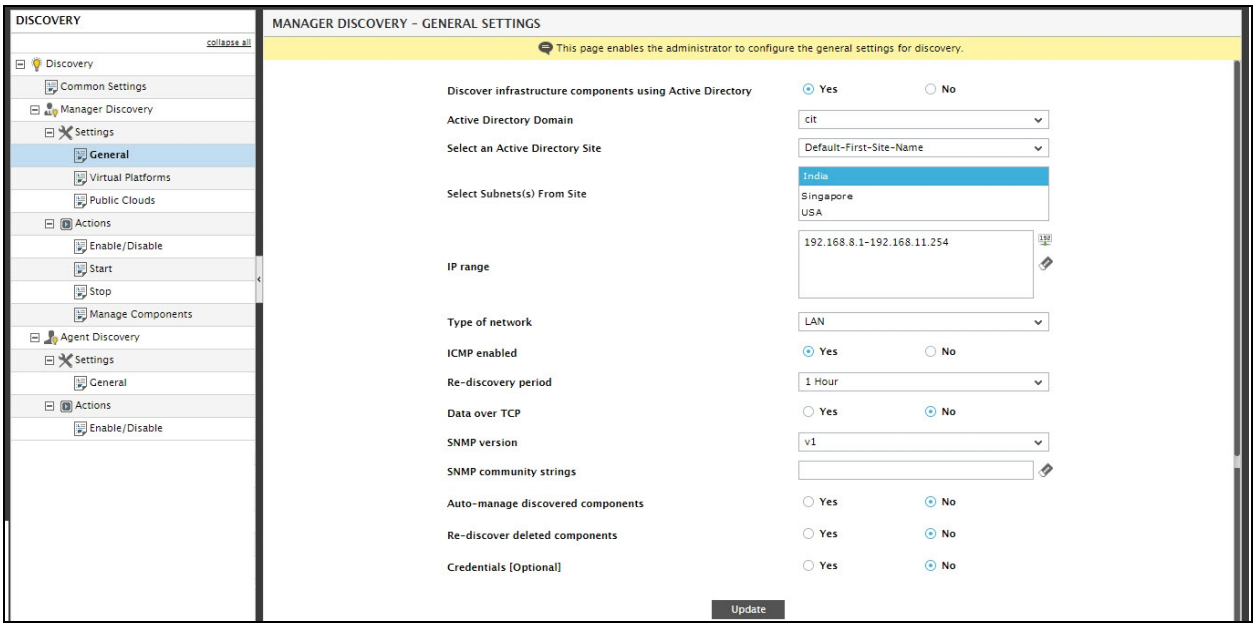

Figure 7.12: Fetching the IP range from the Active Directory server

Next, from the **Type of Network** list box in Figure 7.12, choose the type of network being discovered. The default choice is a Local Area Network (LAN), wherein network delays are of the order of a few tens of milliseconds. To support networks that span multiple geographic locations, an administrator can change the network type to a Wide Area Network (WAN) in which case, the network delays are of the order of seconds. By default, the eG manager attempts to use the Internet Control Message Protocol (ICMP) to figure out which hosts are available in the target environment. Once the hosts that are available are detected, the manager usx`xes a simple port scanning technique to discover the components running on these hosts. For security reasons, some environments may not allow ICMP traffic through their networks. To inform the eG manager of such a restriction, use the **ICMP enabled** option. If this value is set to "No", the eG manager directly uses the port scanning technique for each of the hosts in the specified IP address ranges.

The **Re-discovery period** in Figure 7.12 determines the frequency with which the discovery process executes. This frequency governs how quickly the eG manager is able to discover components that may have been newly added to the target environment. The default is 1 hour.

By default, in an IT environment, all data transmission occurs over UDP. Some environments however, may be specifically configured to offload a fraction of the data traffic – for instance, certain types of data traffic or traffic pertaining to specific components – to other protocols like TCP, so as to prevent UDP overloads. In

such environments, you can instruct the eG manager to conduct the discovery-related data traffic over TCP (and not UDP). For this, set the **Data over TCP** flag in Figure 7.12 to **Yes**. By default, this flag is set to **No**.

Then, proceed to indicate how network devices are to be discovered by the eG manager. Typically, SNMP is used to discover network devices. The version of SNMP used by network devices in the target environment should first be selected from the **SNMP version** list in Figure 7.12. The **SNMP community strings** using which the network device discovery is to be performed should then be provided as a comma-separated list. This specification can be cleared at any point in time by clicking the **Clear the community string** icon (i.e., the 'eraser' icon) next to the **SNMP community strings** text box.

By default, components discovered by the eG manager are not automatically monitored by eG Enterprise; you have to explicitly manage the discovered components for eG to begin monitoring them. This is why, the **Auto-manage discovered components** flag in Figure 7.12 is set to **No** by default. If you want eG to automatically manage the discovered components, then set this flag to **Yes**. Selecting **Yes** automatically invokes the **Include/Exclude Host IP/Name** field, where you can indicate which components you want auto-managed upon discovery, and which ones need not be. To auto-manage only specific components, select **Include** from the drop-down list, and specify a comma-separated list of IP addresses/host names that you want auto-managed – e.g., *egora,sqlserver,xenapp*. IP/host name patterns can also be provided using wildcard characters – for e.g., *192.168.10.*,192.168.*.121,*.168.9.**. In these cases, only the configured components or those that match the configured patterns will be auto-managed upon discovery; all other discovered components will have to be manually managed.

Likewise, to auto-manage all discovered components, except a few specific components, select the **Exclude** option from the drop-down, and provide a comma-separated list of IP/host names that need not be auto-managed. Here again, IP/host name patterns can also be provided.

By default, components that are deleted will not be discovered again. Accordingly, the **Re-discover deleted components** flag in Figure 7.12 is set to **No** by default. If you want the eG Enterprise system to re-discover deleted components, then set this flag to **Yes**.

Typically, in an IT environment, administrators may want to monitor the performance of their critical Windows servers. However, they may not be as interested in monitoring Windows desktops. This is why, administrators may want to configure the eG manager to perform 'selective' discovery – i.e., discover the Windows servers alone and ignore the Windows desktops.

To determine the type of Windows operating system (Server Edition or Desktop Edition) that runs on the hosts in the given IP range, the eG manager should connect to each host using the privileges of a valid user to the domain in which those hosts exist. To configure the credentials of a domain user, set the **Credentials** option to **Yes** in Figure 7.12. This will invoke the following parameters:

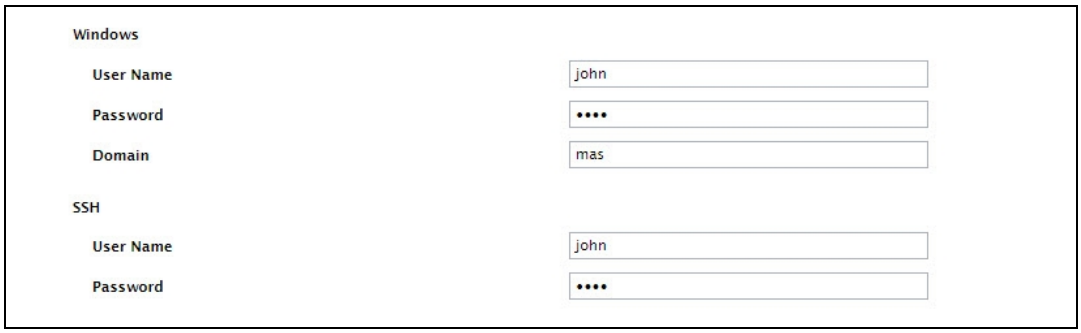

Figure 7.13: Configuring credentials for discovery

Here, enter the **Domain** name, and the **User Name** and **Password** of a valid domain user in the **Windows** section of Figure 7.13.

Finally, click the **Update** button in Figure 7.12. This will start the discovery process.

7.1.1.4 Starting/Stopping the Discovery

Once the **Update** button in Figure 7.12 is clicked, control will automatically switch to the **Start** sub-node under the **Actions** node in the **DISCOVERY** tree in the left panel of Figure 7.12. This indicates that discovery has begun. You can track the progress of the discovery and can also identify the types of components that have been discovered using the page that appears in the right panel when the **Start** node is selected (see Figure 7.14).

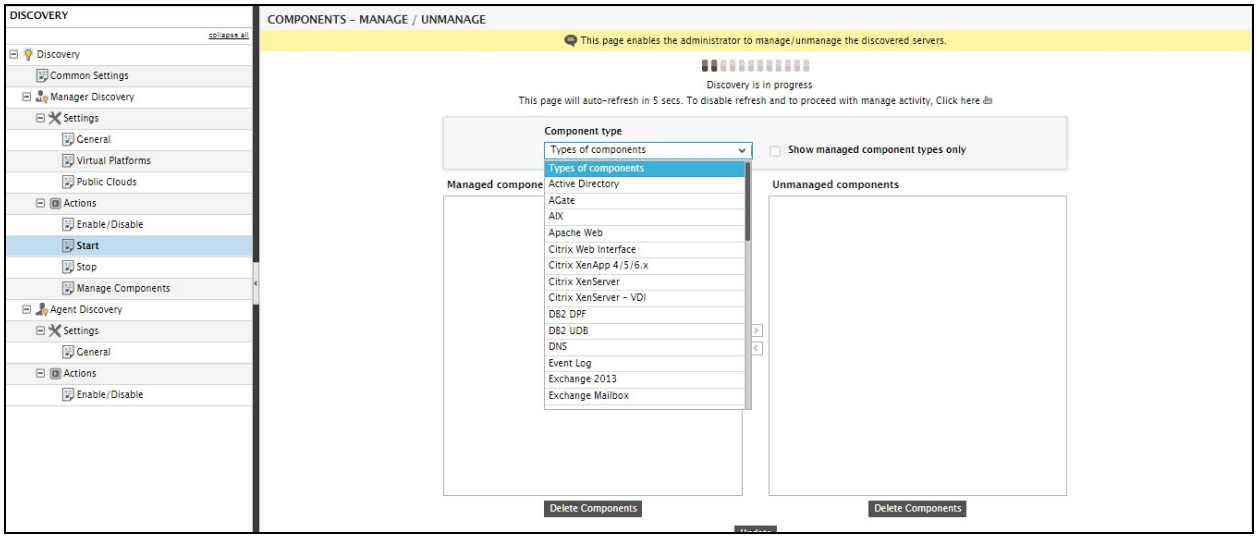

Figure 7.14: Discovery in progress

The **Component type** list in Figure 7.14 lists the types of applications/devices that the discovery process has discovered.

You can stop the discovery process any time after it is started by clicking the **Stop** sub-node under the **Actions** node in the **DISCOVERY** tree (see Figure 7.14).

7.1.1.5 Discovering Virtual Platforms

By default, the eG manager attempts to discover virtual platforms using the port-scanning technique (that was described earlier in this chapter). In the case of some virtual platforms, this technique may not be effective. In case of a few others, security considerations may discourage administrators from opting for this technique. To address these concerns, eG Enterprise allows administrators the flexibility to discover virtual platforms without running the port scans. How this is done differs from one virtual platform to another. The sub-sections below elaborate on how each virtual platform can be discovered using the eG manager.

Discovering vSphere/ESX Servers

eG Enterprise is capable of automatically discovering the ESX servers in the environment using the eG manager or the eG agent that is monitoring the ESX server. Since both these approaches employ a port-scanning technique to discover ESX servers, they might not be suitable for high-security environments where the firewall has been configured to block traffic to and from specific ports. Likewise, in environments that are spread across multiple sub-nets, discovery performed using the eG manager or the agents could increase the bandwidth consumption. To address these concerns, eG Enterprise provides you with the option to directly connect to one/more VMware vCenter installations in your environment to perform ESX discovery. The additional benefit that accrues in this process is that, when one/more ESX servers discovered using vCenter are managed, then eG Enterprise automatically uses the same vCenter server to collect performance metrics related to the ESX servers. In other words, eG Enterprise auto-configures the tests pertaining to the managed ESX servers with the details of the vCenter server used for their discovery; this way, the solution minimizes the time and effort involved in manual test configuration, and quickly starts collecting metrics from vCenter. Similarly, if the vCenter details need to be modified for any reason, then, you will not be required to manually reconfigure each test for this purpose; changing the configuration of the vCenter server in the **DISCOVERY** page will automatically update all the tests that have been configured to collect metrics from that vCenter server.

To discover the vSphere/ESX hosts, select the **Virtual Platforms** node from the **DISCOVERY** tree in the left panel of Figure 7.15. Then, choose the **vSphere/ESX Hosts** option from the **Choose a virtual platform to discover** drop-down in the right panel (see Figure 7.15). If one/more vCenter servers have already been added to enable the discovery of vSphere/ESX servers, the details of the same will be displayed in the right panel as shown by Figure 7.15.

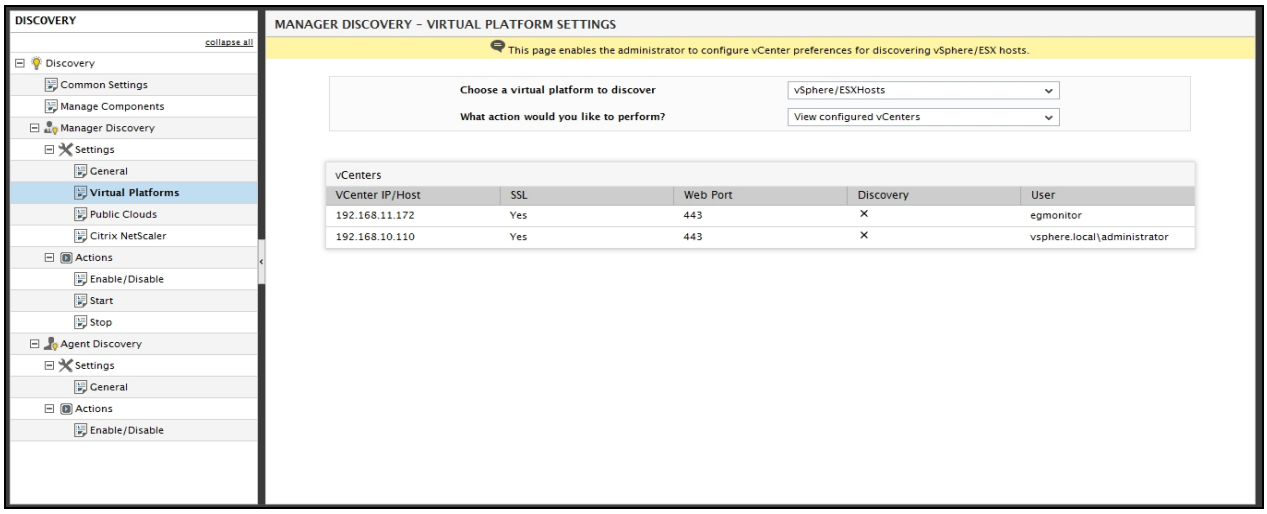

Figure 7.15: Discovering vSphere/ESXHosts

The first step towards using vCenter for ESX discovery and performance monitoring is to configure the eG manager with the details of the vCenter server(s), and mark the server(s) that will be used for discovery.

Prior to vCenter creation, you might have to increase the memory settings of the eG manager. This is because, VI APIs are memory intensive, and therefore, the default memory setting of 128 MB for the eG manager might not be sufficient for performing ESX discovery via vCenter. To override this default setting, select the Configure memory settings for discovery option from the What would you like to perform? drop-down list in the right panel of Figure 7.15. Figure 7.16 will then appear.

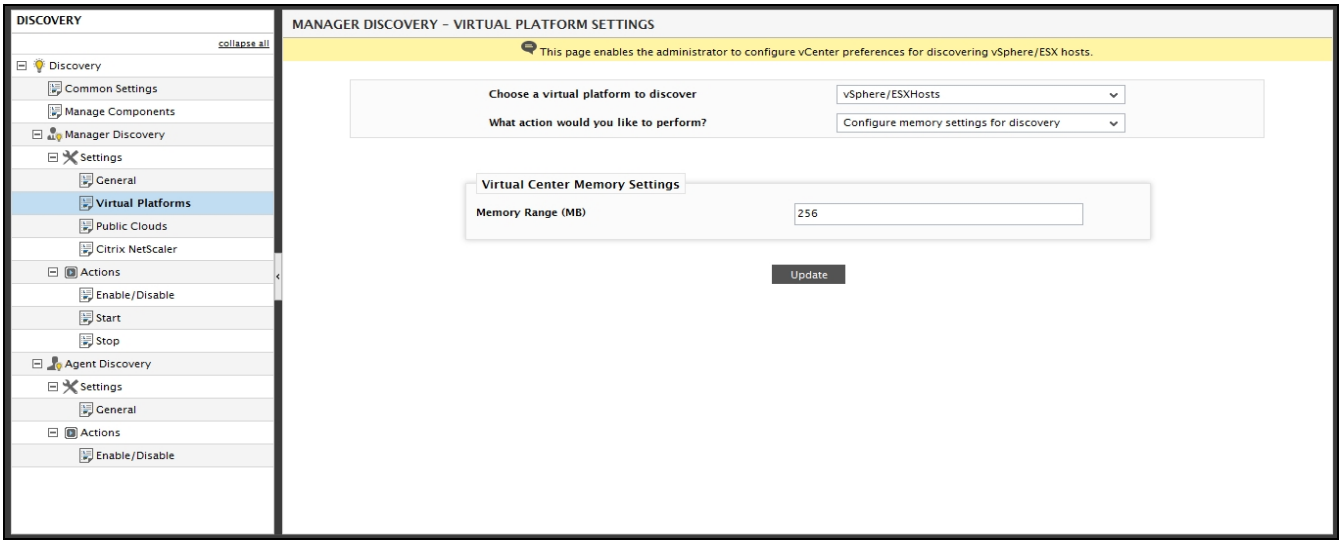

Figure 7.16: Overriding the default memory settings of the eG manager

You can specify any value between 256 MB and 1024 MB in the **Memory Range (MB)** text box in Figure 7.16, and click the **Update** button therein.

Next, you can proceed to add a new vCenter. To add a new vCenter server, do the following:

1. Select the **Add new vCenter** option from the **What action would you like to perform?** drop-down in the right panel (see Figure 7.17).

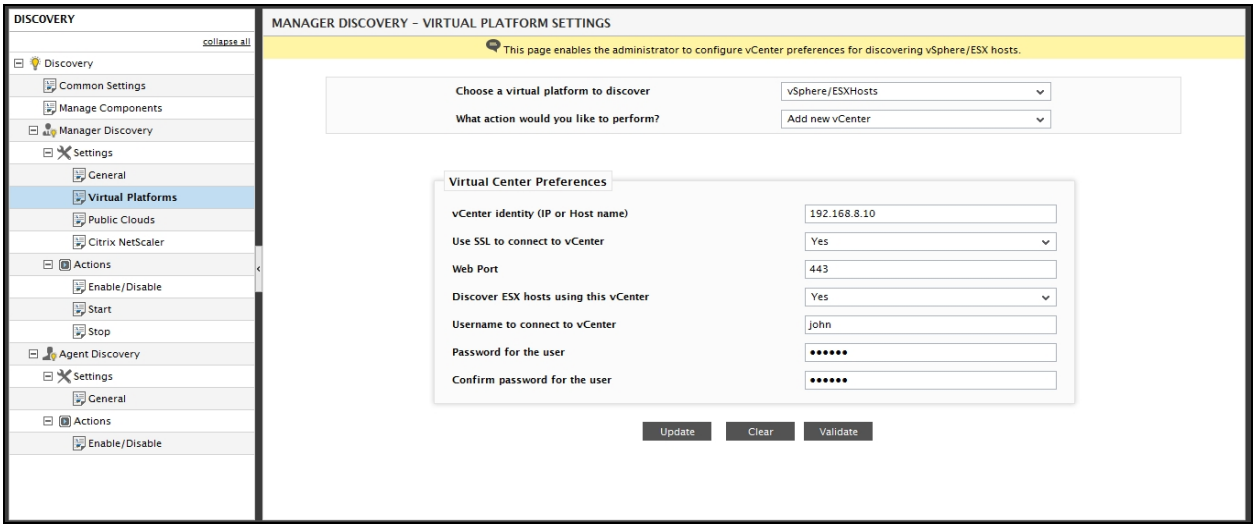

Figure 7.17: Adding a vCenter configuration

2. Then, specify the following in right panel of Figure 7.17:

   - Specify the IP or host name of the vCenter in the **vCenter identity** text box.

     Note:

     When providing the IP address of the vCenter server against **vCenter identity**, you can specify either the IPv4 or the IPv6 address of the vCenter server.

   - Then, indicate whether the eG manager is to connect to vCenter using SSL or not by selecting the **Yes** or **No** option from the **Use SSL to connect to vCenter** list. By default, this flag is set to **Yes**.

   - By default, in most virtualized environments, vCenter listens on port 80 (if not SSL-enabeld) or on port 443 (if SSL-enabled) only. This implies that while discovering ESX servers using vCenter, eG Enterprise communicates with vCenter via the default ports 80 or 443, depending upon the SSL-enabled status of vCenter. Accordingly, the **Web Port** parameter in Figure 7.17 is set to 443 by default, if the **SSL** flag is set to **Yes**, and displays the default value 80 if the **SSL** flag is set to **No**. In some environments however, the default ports 80 or 443 might not apply. In such a case, in the **Web Port** text box, specify the exact port at which vCenter in your environment listens.

   - Typical virtualized environments may consist of multiple vCenter installations, each managing a different set of ESX servers. To enable the eG manager to automatically discover those ESX servers that are managed by the vCenter being added, then set the **Discover ESX hosts using this vCenter** flag to **Yes**. If not, then set it to **No**.

   - In order to be able to discover ESX servers using a vCenter server, the eG manager needs to connect to vCenter using the credentials of a valid user to vCenter. Provide the user name and password of such a user in the **Username to connect to vCenter** and **Password for the user** text boxes. This user typically requires **Administrator** or **Virtual Machine Administrator** privileges. However, if you cannot expose the credentials of such a user owing to security constraints, then, you can use the credentials of a user with 'Read-only' privileges to vCenter. If such a user pre-exists, then, provide

the name and password of that user in the text boxes mentioned above. Otherwise, assign the 'Read-only' role to a local/domain user to vCenter, and provide the name and password of this user in the **Username to connect to vCenter** and **Password for the user** text boxes. The steps for achieving this are detailed in the *Monitoring VMware Infrastructures* document.

- Confirm the password of the user by retyping it in the **Confirm password for the user** text box.
- To clear all the configured details, click on the **Clear** button in Figure 7.17. To start discovery instead, click on the **Update** button.
- Clicking on **Update** brings up Figure 7.18, which requests you to confirm whether you want to start discovery based on the specifications provided. Click the **OK** button to begin discovery.

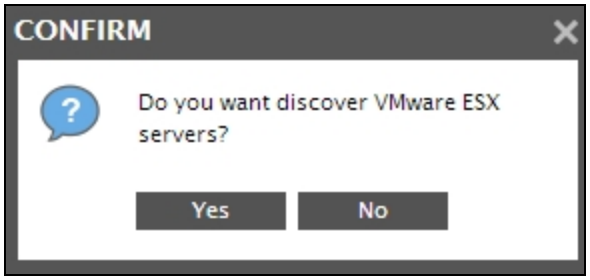

Figure 7.18: Starting ESX discovery

- Regardless of the discovery approach used (discovery using an IP range or using vCenter), triggering ESX discovery will lead you straight to the **COMPONENTS - MANAGE / UNMANAGE**. To know how to manage/unmanage components using this page, refer to Section **7.2** of this document.

To view the vCenter servers so added, select the **View configured vCenters** option from the **What action would you like to perform?** drop-down. Figure 7.19 will then appear listing the vCenter servers that have been configured for discovering ESX servers.

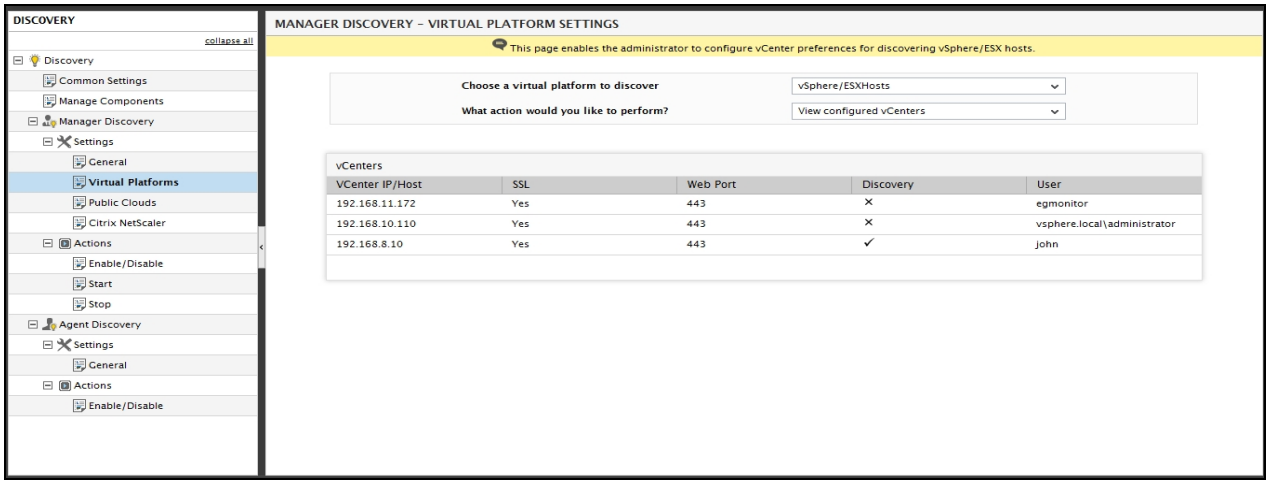

Figure 7.19: Viewing the configured vCenter servers

To modify a vCenter configuration, do the following:

1. Select the **Modify Configured vCenters** option from the **What action would you like to perform?** drop-down. Figure 7.20 will then appear, using which you can modify the vCenter configuration. For that, first, select the IP/host name of the vCenter to be modified from the **vCenter identity (IP or Host name)** list. The details of the chosen vCenter will then be displayed against the appropriate fields. You can modify any of the displayed details and update the changes by clicking the **Update** button.

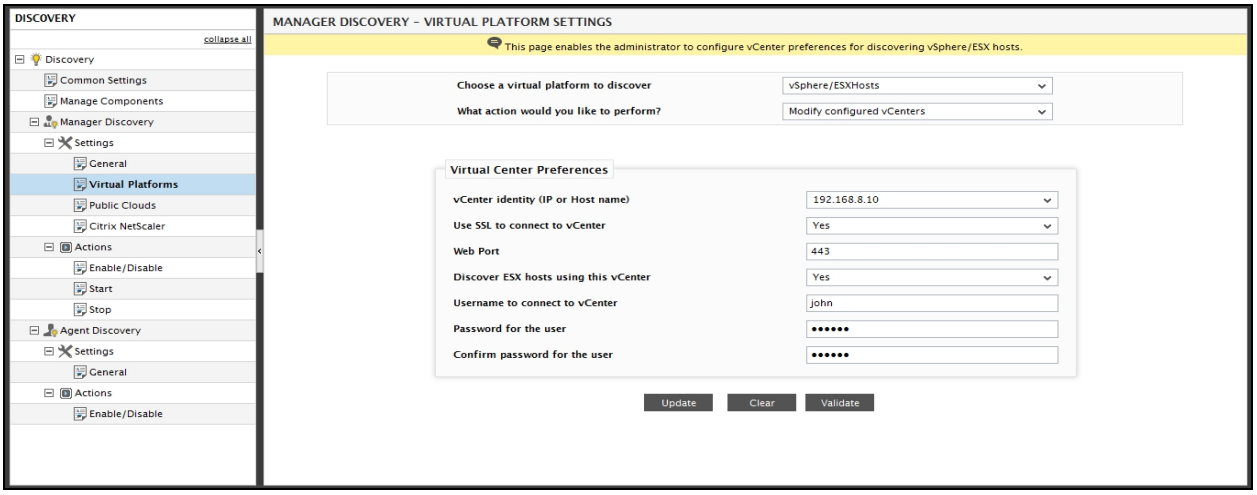

Figure 7.20: Modifying the vCenter configuration

To delete a particular vCenter, select the **Delete Configured vCenter** option from the **What action would you like to perform?** drop-down. Figure 7.21 then appears listing the available vCenter servers.

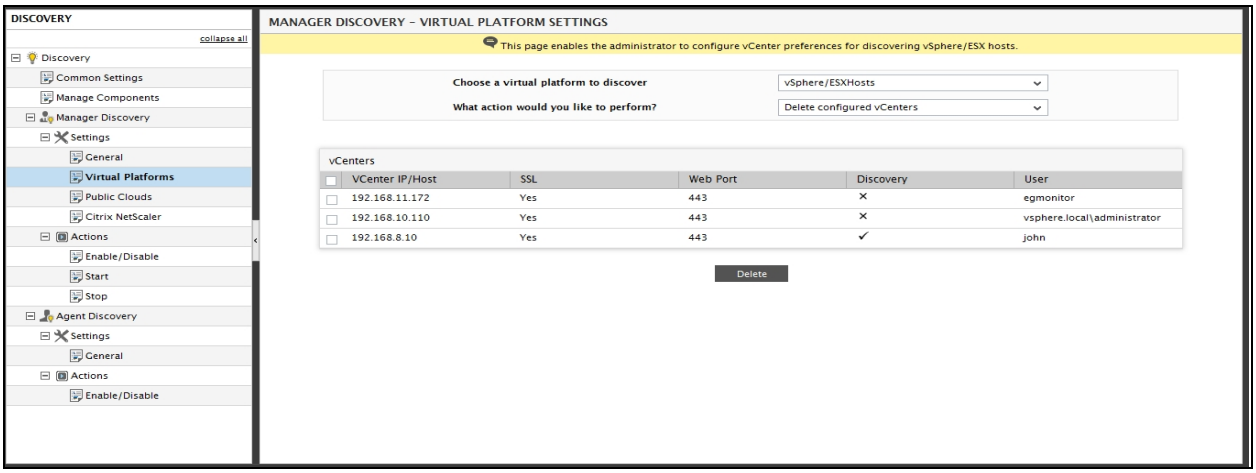

Figure 7.21: Deleting a vCenter

Select the vCenter servers to be deleted by selecting the check boxes corresponding to the vCenter configuration in Figure 7.21. To mark all the listed vCenter servers for deletion simultaneously, simply select the top-most check box in the column of check boxes. To delete the marked vCenter servers, click on the **Delete** button in Figure 7.21.

Often administrators may want to view the remote agents that are currently monitoring the discovered vSphere/ESX hosts with ease!To identify the remote agents that are assigned to the vCenters while discovering the vSphere/ESX hosts, do the following:

- Select the **View remote agent** mapping option from the **What would you like to perform?** drop-down list. Figure 7.22 will then appear.
- As shown in Figure 7.22, all the vCenters will be listed along with the count of the remote agents and the name of the remote agents.

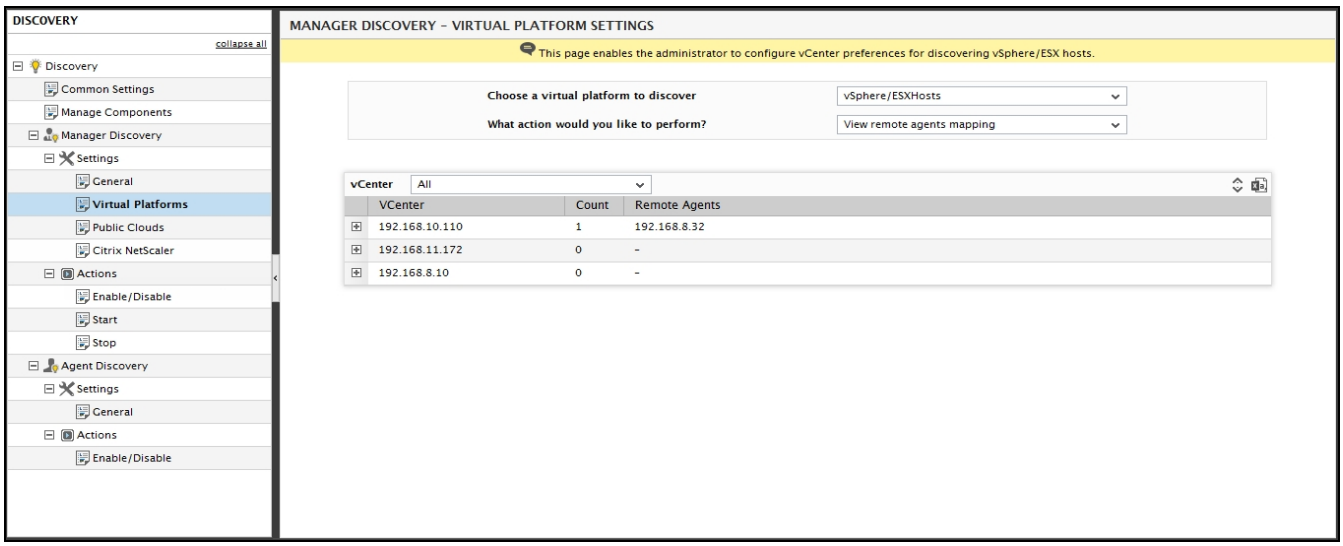

Figure 7.22: The remote agents associated with each vCenter

- Further drilling down each vCenter, the components associated with each remote agent will be listed (see Figure 7.23).

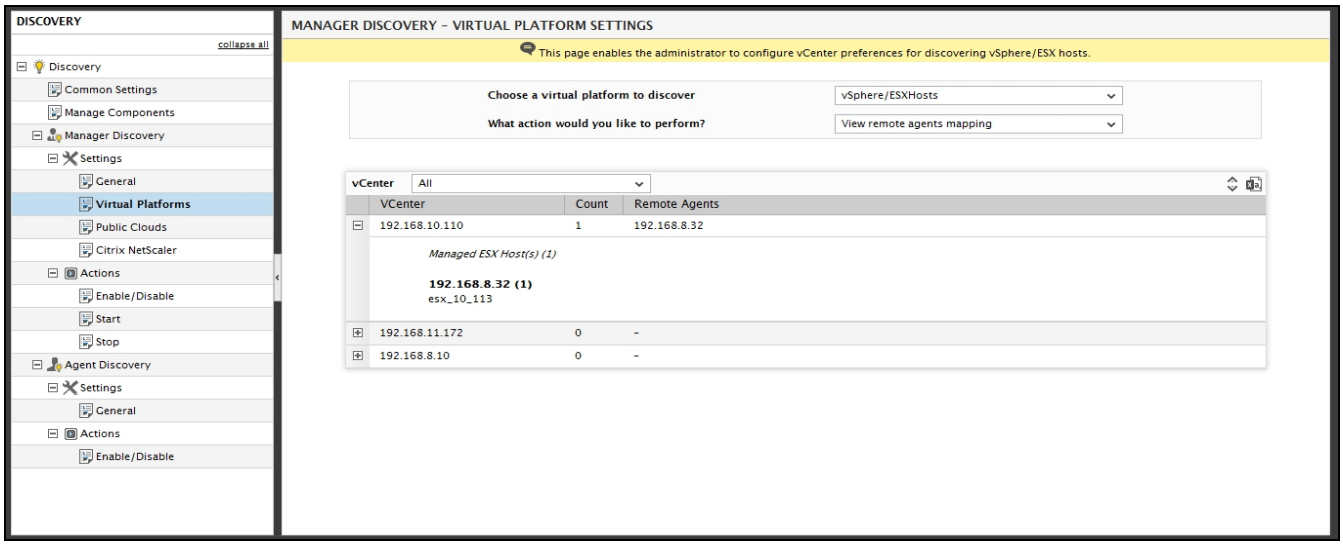

Figure 7.23: Listing the vSphere/ESX hosts within the vCenter associated with each remote agent

Discovering IBM pSeries Servers

eG Enterprise is capable of monitoring the AIX LPARs that are operating on IBM pSeries servers using a patented 'In-N-Out' monitoring approach. The first step towards implementing this approach is to discover the IBM pSeries servers in the environment. To save administrators the trouble of manually adding each IBM pSeries server to be monitored, the eG management console facilitates the configuration of an HMC (Hardware Management Console) server, using which the IBM pSeries servers can be automatically discovered. The **HMC** node in the **DISCOVERY** tree enables this. Using the sub-nodes of the **HMC** node, you can configure a new HMC server for discovery purposes, modify the HMC server configuration if required, view the HMC configuration, and delete the same.

To discover the pSeries Servers, select the **Virtual Platforms** sub-node under the **General Settings** node of the **DISCOVERY** tree. Then, choose the **pSeries Servers** option from the **Choose a virtual platform to discover** drop-down in the right panel, as shown by Figure 7.24.

To add a new HMC server, do the following:

1. Select the **Add new HMCs** option from the **What action would you like to perform?** drop-down.
2. In the right panel, specify the **IP address of the Hardware management console** to be used for the discovery.

   Note:

   Only an IPv4 address can be specified against **IP address of the Hardware management console**.

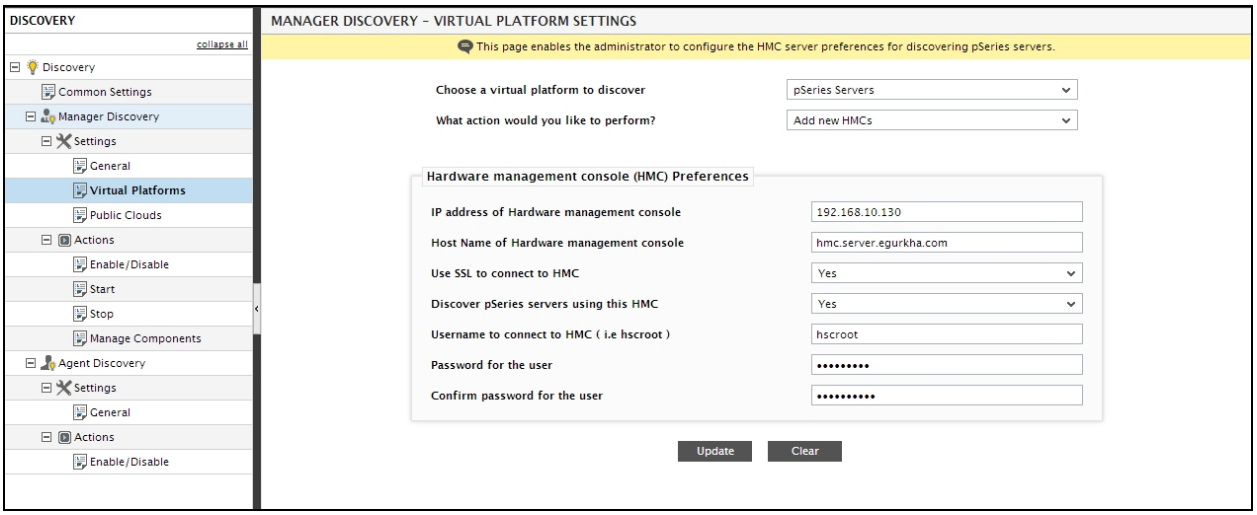

Figure 7.24: Adding an HMC server

3. Then, mention the **Host Name of the Hardware management console**.
4. Next, indicate whether/not the eG manager needs to use SSL to connect to HMC. To use SSL, set the **Use SSL to connect to HMC** flag to **Yes**. Otherwise, set it to **No**.
5. If more than one HMC server is deployed in your environment, you may want to discover the IBM pSeries servers using one of the HMC servers alone; you may still want to feed the eG manager with the details of the other HMC server, so that it can be used for discovery some time later. This is why, while configuring an HMC server using this page, you need to indicate whether/not the HMC server being configured is to be used for discovery or not. To use the HMC server for discovery, set the **Discover pSeries servers using this HMC** flag to **Yes**. Set the flag to **No** if you do not want to discover pSeries servers using this HMC server.

6. Provide a valid user name for connecting to the HMC server in the **Username to connect to HMC** text box.
7. Specify the **Password** of the HMC user.
8. Confirm the password by retyping it in the **Confirm password for the user** text box.
9. Finally, click the **Update** button to save the changes.
10. A message box will then appear requesting your confirmation to proceed with the discovery. Click the **OK** button to discover the pSeries servers, or click the **Cancel** button to cancel discovery.

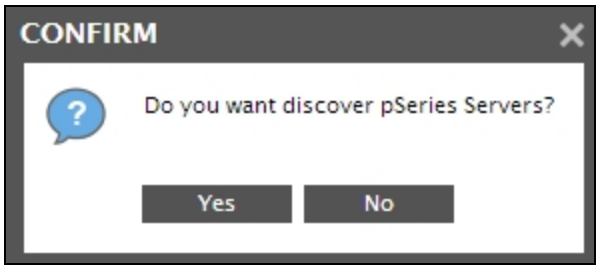

Figure 7.25: Discovering the pSeries servers

To view the HMC servers that have been configured, pick the **View configured HMCs** option from the **What action would you like to perform?** drop-down in Figure 7.26. The details of the HMC servers that pre-exist will then appear.

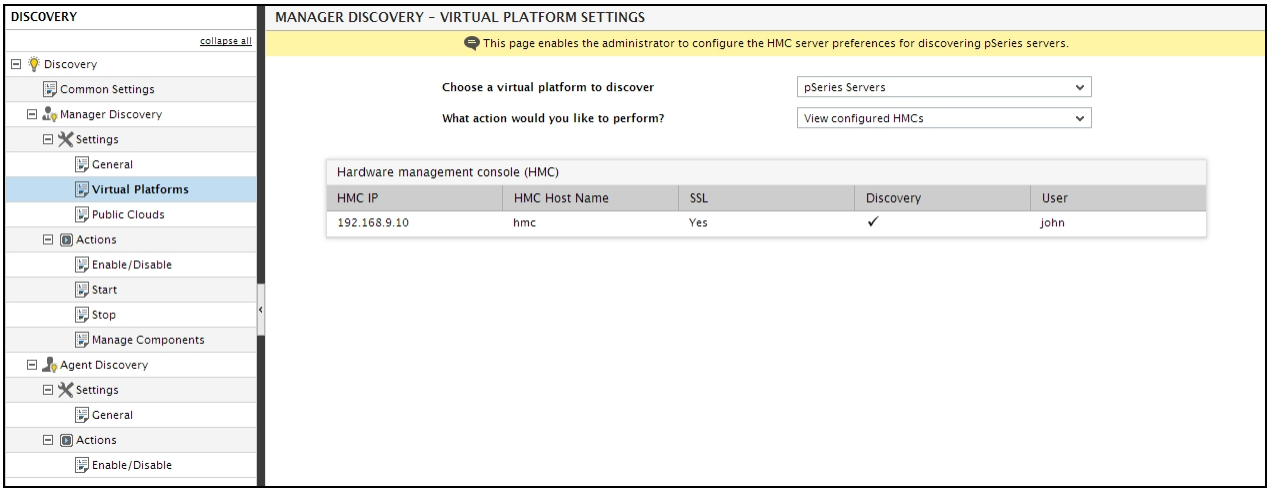

Figure 7.26: Viewing the HMC server configuration

To modify the configuration of an HMC server, do the following:

1. Click on the **Modify configured HMCs** option from the **What action would you like to perform?** drop-down, as depicted by Figure 7.27.
2. Pick the **IP address of the Hardware management console** to be modified from the right panel. The details of the chosen HMC will then be displayed. Make the required changes, and click the **Update** button to save them.

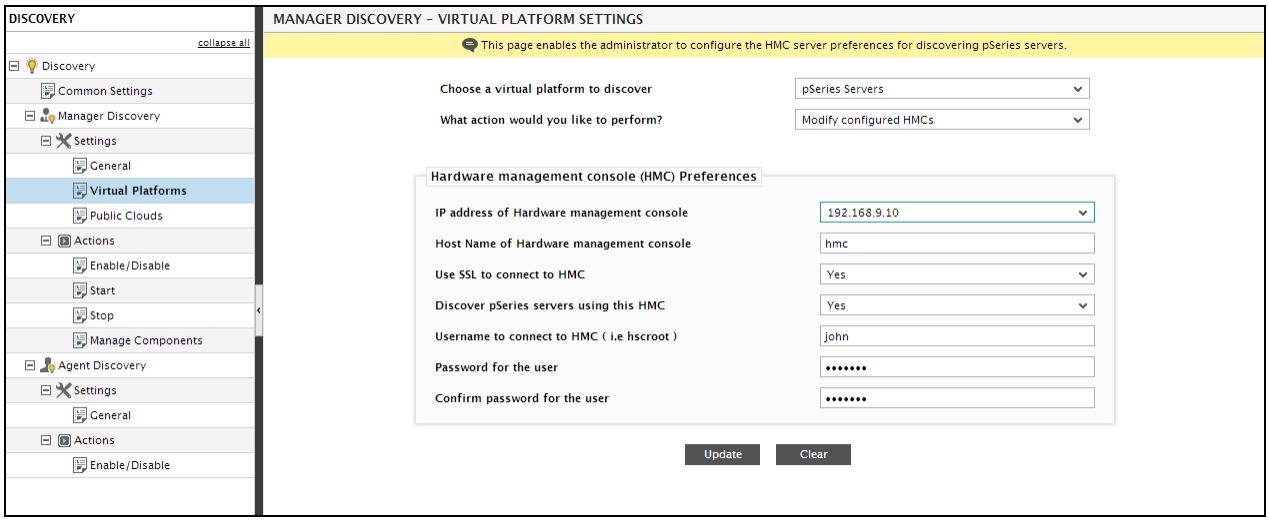

Figure 7.27: Modifying the HMC server configuration

To delete a HMC configuration, do the following:

1. Click on the **Delete configured HMCs** option from the **What action would you like to perform?** drop-down.
2. The details of all existing HMC servers will then be displayed in the right panel. To mark a HMC server for deletion, simply select the check box alongside the details of that HMC. To mark all listed HMC servers for deletion, simply select the check box that precedes the column lable **HMC IP**.

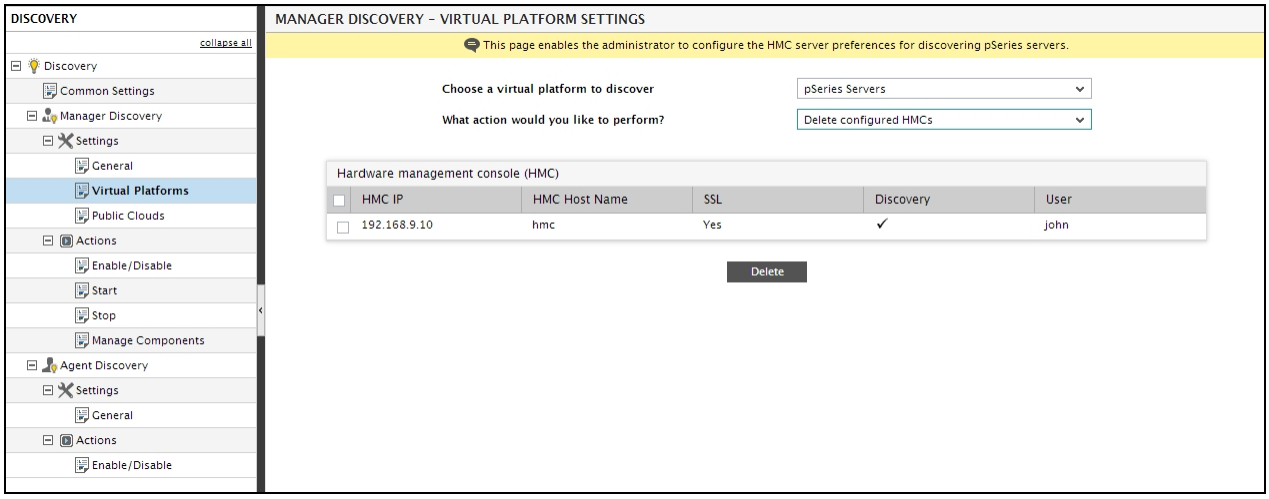

Figure 7.28: Deleting an HMC server

3. Finally, click the **Delete** button to delete the selection.

Discovering RHEV Servers via the RHEV Manager

eG Enterprise embeds the ability to monitor the RHEV servers in an environment and the VMs operating on those servers. The first step to monitoring the RHEV servers is to identify those servers that require monitoring. The eG Enterprise system can auto-discover the RHEV servers in a virtualized environment using the RHEV managers that manage those servers. To enable this auto-discovery, you first need to enable the

eG manager to connect to the RHEV manager in the environment. The **Virtual Platforms** sub-node in the **DISCOVERY** tree can be used for this purpose. This is how:

To discover the RHEV Servers, select the **Virtual Platforms** sub-node of the **General Settings** node in the **DISCOVERY** tree in the left panel of Figure 7.29. Then, choose the **RHEV Hypervisors** option from the **Choose a virtual platform to discover** drop-down in the right panel of Figure 7.29.

1. First, you need to configure the RHEV manager with which the eG manager should connect for automatically discovering the RHEV servers that require monitoring. For this, select the **Add new RHEV Managers** option from the **What action would you like to perform?** drop-down.

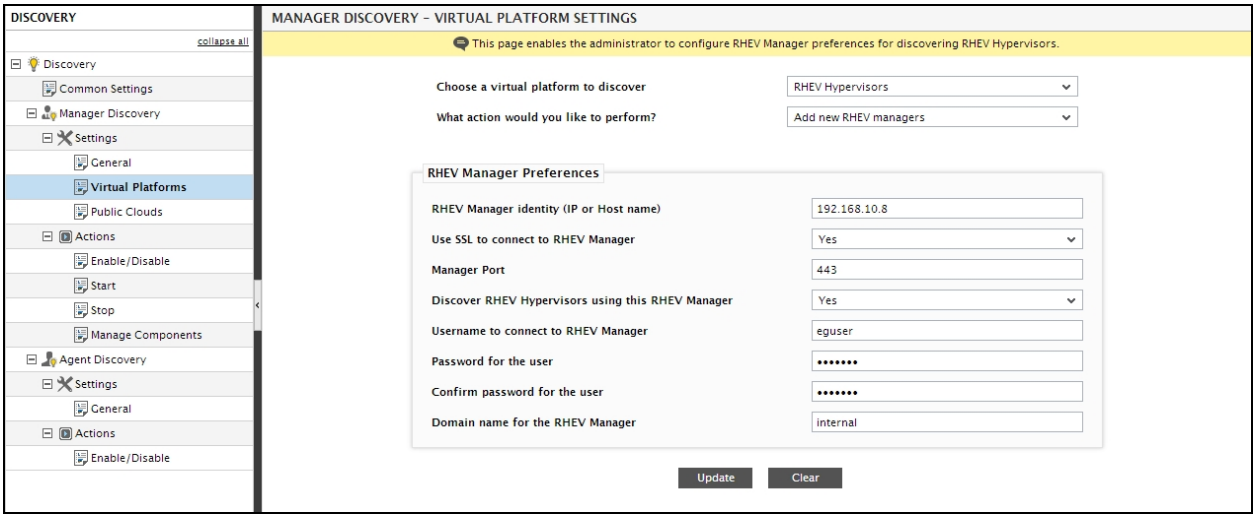

Figure 7.29: Adding the RHEV Manager to be used for discovering the RHEV servers

2. In the right panel of Figure 7.29, specify the following:
   - **RHEV Manager Identity**: Specify the IP address/host name of the RHEV manager in your environment.

     Note:

     If you choose to specify the IP address of the RHEV manager against **RHEV Manager Identity**, then make sure that you specify only an IPv4 address.
   - **Use SSL to connect to the RHEV Manager**: Set this flag to **Yes** if the RHEV manager in your environment is SSL-enabled. Otherwise, set this flag to **No**.
   - **Manager Port**: If the RHEV manager is SSL-enabled, then 8443 will be displayed here by default. On the other hand, if the manager is not SSL-enabled, the default **Manager Port** will be 8080. If the RHEV manager in your environment listens on a different SSL or non-SSL port, then make corresponding changes to the default setting.
   - **Discover RHEV Hypervisors using this RHEV Manager**: If you want to discover RHEV servers in your environment using this RHEV manager, set this flag to **Yes**. If you only want to use this RHEV manager to obtain the *outside view* of VMs, set this flag to **No**.
   - **Username to connect to RHEV Manager** and **Password for user**: Specify the credentials (i.e., user name and password) of a user who has been assigned the **RHEVMUser** role. If no such user pre-exists, then create

a special user for this purpose, assign the **RHEVMUser** role to this user, and provide his/her login credentials here.

- **Confirm password for user**: Confirm the password of the **RHEVMUser** by retyping it here.
- **Domain name for the RHEV manager**: Specify the name of the domain to which the RHEV manager belongs.

3. Then, click the **Update** button in Figure 7.29. A message box requesting your confirmation to proceed with the discovery will then appear (see Figure 7.30. Click the **OK** button in the message to trigger the discovery. Clicking the **Cancel** button will save the details of the RHEV manager, but will not start the discovery. In this case, you can use the RHEV manager configuration so saved to perform discovery at a later point in time.

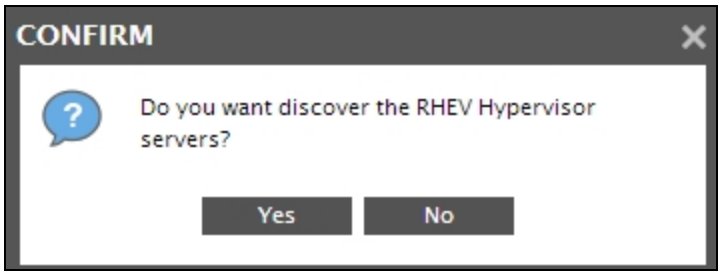

Figure 7.30: A message box requesting confirmation to discover the RHEV servers

4. This way, you can add the details of multiple RHEV managers to the eG Enterprise system.
5. To view the details of all the RHEV managers that are added, choose the **View configured RHEV managers** option from the **What action would you like to perform?** drop-down in the right panel of Figure 7.31. The right panel will then change to display the names and other details of the RHEV managers that have been configured.

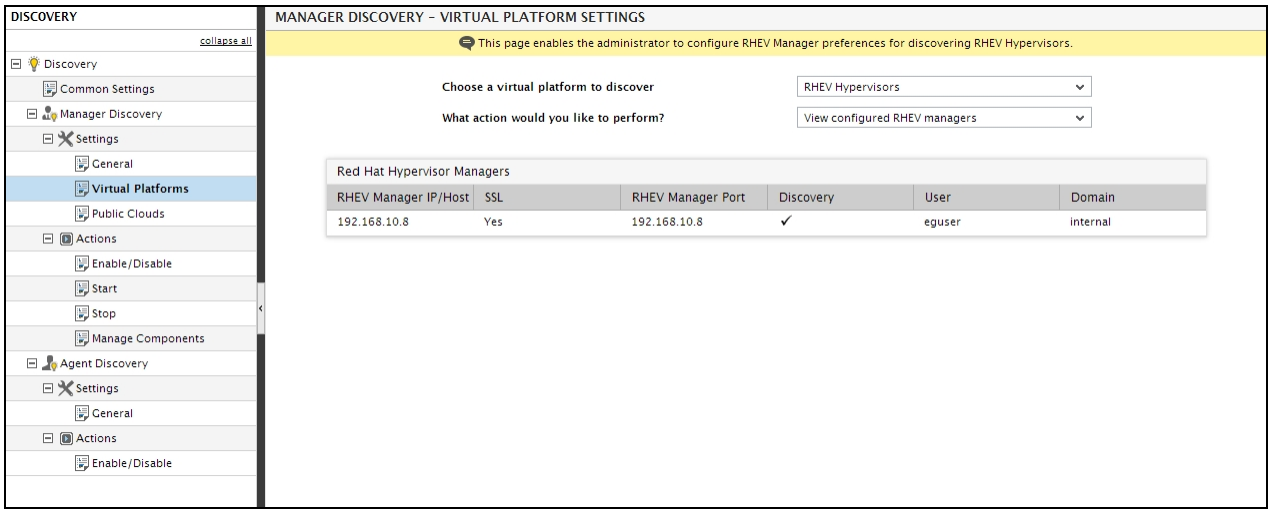

Figure 7.31: Viewing the details of existing RHEV manager

To modify the details of an RHEV manager, select the **Modify configured RHEV Managers** option from the **What action would you like to perform?** drop-down. The right panel will then change as depicted by Figure 7.32. From the **RHEV Manager Identity** list, select the IP address of the RHEV manager to be modified. Upon selection of

the RHEV manager's IP address, the other parameters in the right panel will be populated with the corresponding details. You can change any of the displayed details to suit your needs. Finally, click the **Update** button to save the changes.

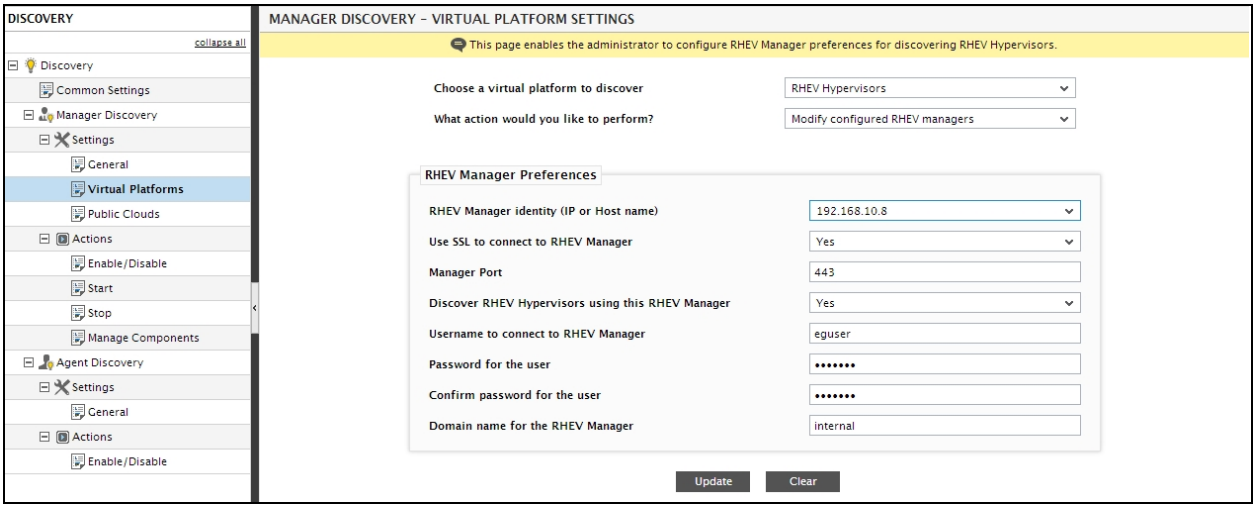

Figure 7.32: Modifying the RHEV Manager's configuration

To delete an RHEV manager, select the **Delete Configured RHEV Managers** option from the **What action would you like to perform?** drop-down. Select the check box corresponding to the RHEV manager to be deleted and click the **Delete** button to delete it. To delete all the RHEV managers at one shot, select the check box just before the column heading, **RHEV Manager IP / Host**, and click the **Delete** button.

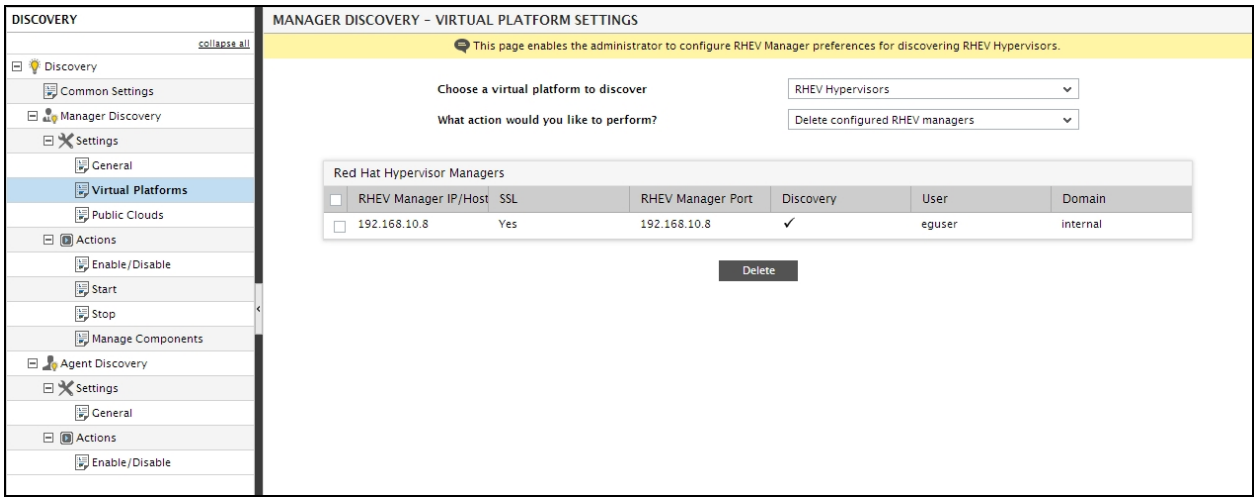

Figure 7.33: Deleting an RHEV manager

7.1.1.6 Discovering Public Clouds

eG Enterprise is a 'Cloud-Ready' monitoring system that provides monitoring `From` the cloud, `of` the cloud, and `For` the cloud! As part of its ongoing efforts to deliver performance management **FOR** the cloud, the solution currently lends out-of-the-box monitoring support for the **AWS EC2 Public Cloud** and its regions. To this

end, the solution offers two specialized monitoring models - the *AWS EC2 Cloud* model and the *AWS EC2 Region* model. The *AWS EC2 Cloud* monitoring model provides you with proactive updates on the overall health and status of the cloud and points you to unavailable regions/availability zones and resource-hungry instances in the cloud. To zoom into the health of specific regions and the instances operating within those regions, use the *AWS EC2 Region* model.

The first step to monitoring an *AWS EC2 Region* is to automatically discover the regions available for monitoring in an AWS EC2 Cloud infrastructure. The **Public Clouds** node in the **DISCOVERY** tree structure facilitates the auto-discovery of these regions. To auto-discover one/more regions, do the following:

1. First, you need to configure the AWS EC2 Cloud account with which the eG manager should connect for automatically discovering the regions that require monitoring. For this, select the **Add new AWS EC2 cloud** option from the **What action would you like to perform?** drop-down in the right panel (see Figure 7.34).

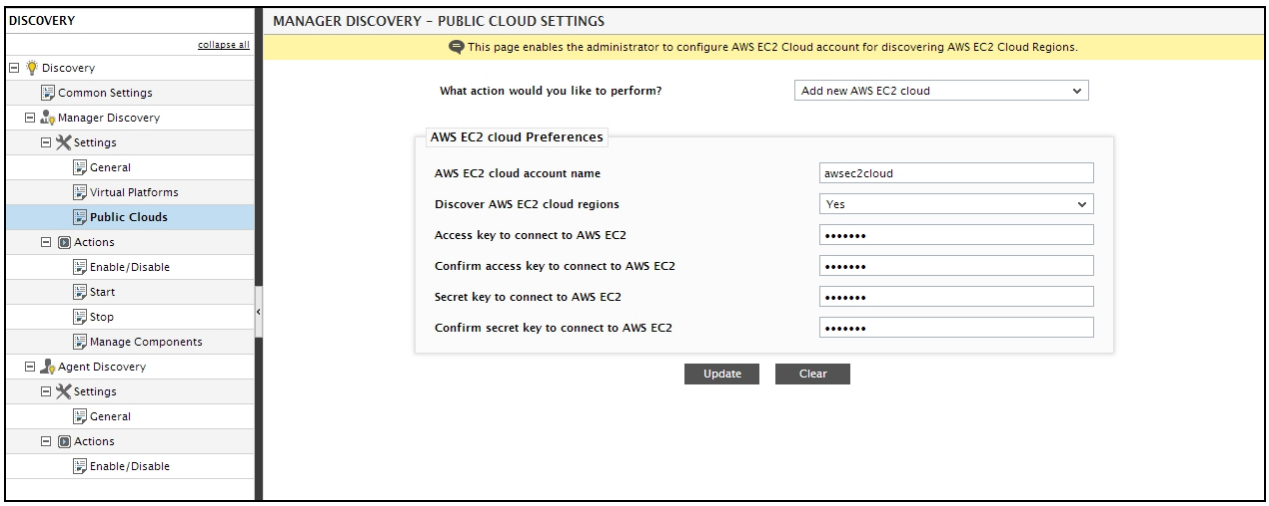

Figure 7.34: Configuring the AWS EC2 Cloud account for discovery

2. In the right panel of Figure 7.34, specify the following:

   - **AWS EC2 cloud account name**: Specify the name of a valid AWS EC2 cloud account with which the solution should connect for discovering the regions. To obtain a cloud account, do the following:
     - Sign up for a new AWS account from the **http://aws.amazon.com/ec2/** page.
     - Provide the details of the user for whom you wish to create the AWS account.
     - Based on the AWS EC2 Regions, you will be requested to choose the pricing for an instance to be deployed in the AWS EC2.
     - Once the payment is made, the user will be automatically signed in to the AWS account.

     Specify the name of this account in the **AWS EC2 cloud account name** text box.
   - **Discover AWS EC2 cloud regions**: If you want to discover the regions configured in the specified cloud account, set this flag to **Yes**.
   - **Access key to connect to AWS EC2** and **Confirm access key to connect to AWS EC2**: To monitor an AWS EC2 region, the eG agent has to be configured with the "access key" of a user with access to the AWS account specified in the **AWS EC2 cloud account name** text box. To obtain an access key, login to

the account specified in the **AWS EC2 cloud account name** text box, and request for an 'access key'. You will then be provided with an 'access key' and 'secret key'. Specify the "access key" in the **Access key to connect to AWS EC2** text box and confirm it by retyping it in the **Confirm access key to connect to AWS EC2** text box.

- **Secret key to connect to AWS EC2** and **Confirm secret key to connect to AWS EC2**: When you request for an "access key" by logging into the AWS EC2 cloud specified, you will be provided with an access key and its corresponding secret key. Enter the secret key in the **Secret key to connect to AWS EC2** text box and confirm it by retyping it in the **Confirm secret key to connect to AWS EC2** text box.

3. Finally, click the **Update** button in Figure 7.34.
4. A message box requesting your confirmation to proceed with the discovery will then appear (see Figure 7.35). Click the **OK** button in the message to trigger the discovery. Clicking the **Cancel** button will save the details of the AWS EC2 cloud account, but will not start the discovery. In this case, you can use the cloud account so saved to perform discovery at a later point in time.

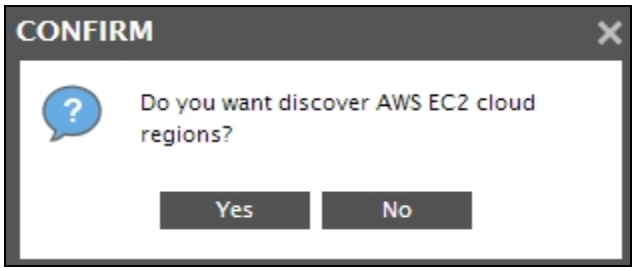

Figure 7.35: A message box requesting confirmation to discover the cloud regions

5. This way, you can add the details of multiple AWS EC2 cloud accounts to the eG Enterprise system.
6. To view the details of all the cloud accounts that are added, select the **View configured AWS EC2 clouds** option from the **What action would you like to perform?** drop-down in the right panel of Figure 7.36. The right panel will then change to display the names and other details of the cloud accounts that have been configured.

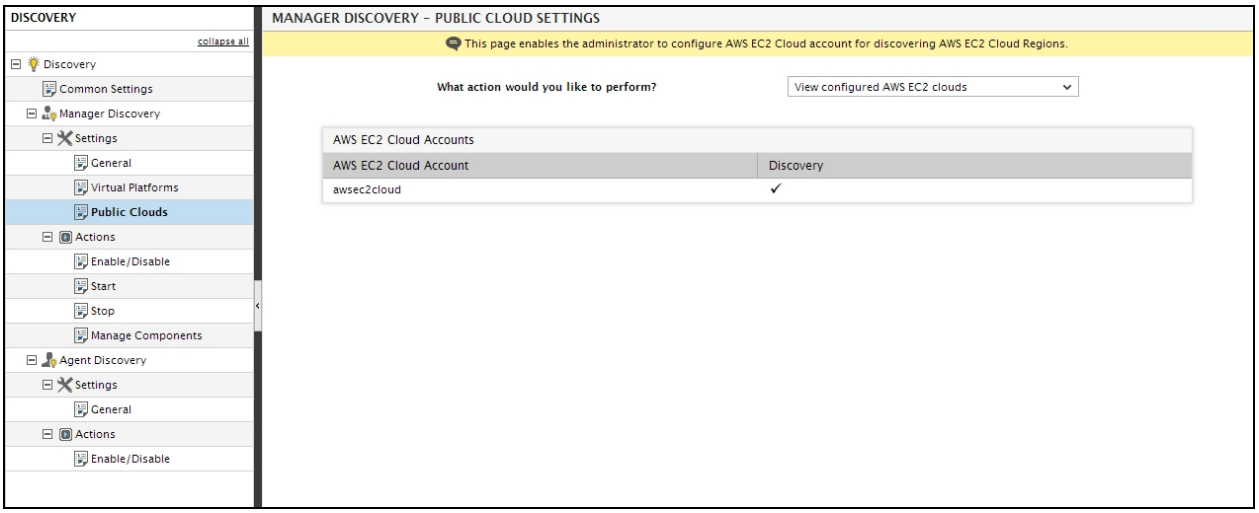

Figure 7.36: Viewing the details of existing AWS EC2 cloud accounts

To modify the details of a cloud account, pick the **Modify configured AWS EC2 clouds** option from the **What action would you like to perform?** drop-down (see Figure 7.37). From the **AWS EC2 cloud account name** list in the right panel, select the cloud account to be modified. Upon selection of the cloud account, the other parameters in the right panel will be populated with the corresponding details. You can change any of the displayed details to suit your needs. Finally, click the **Update** button to save the changes.

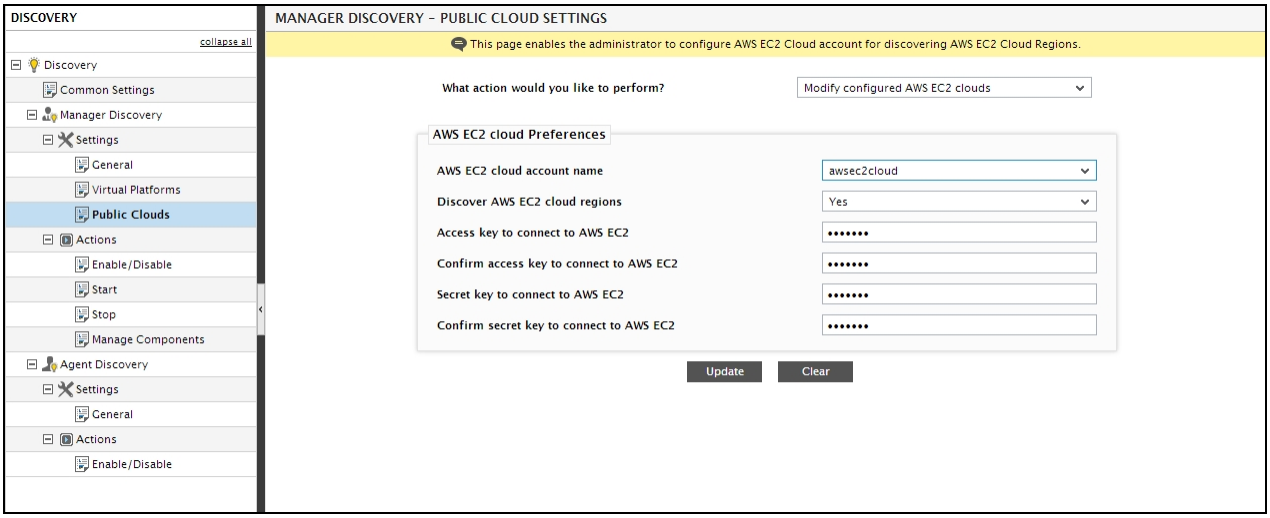

Figure 7.37: Modifying the cloud account's configuration

To delete a cloud account, choose the **Delete configured AWS EC2 clouds** option from the **What action would you like to perform?** drop-down. Select the check box corresponding to the cloud account to be deleted and click the **Delete** button to delete it. To delete all the displayed cloud accounts at one shot, select the check box just before the column heading, **AWS EC2 cloud account**, and click the **Delete** button.

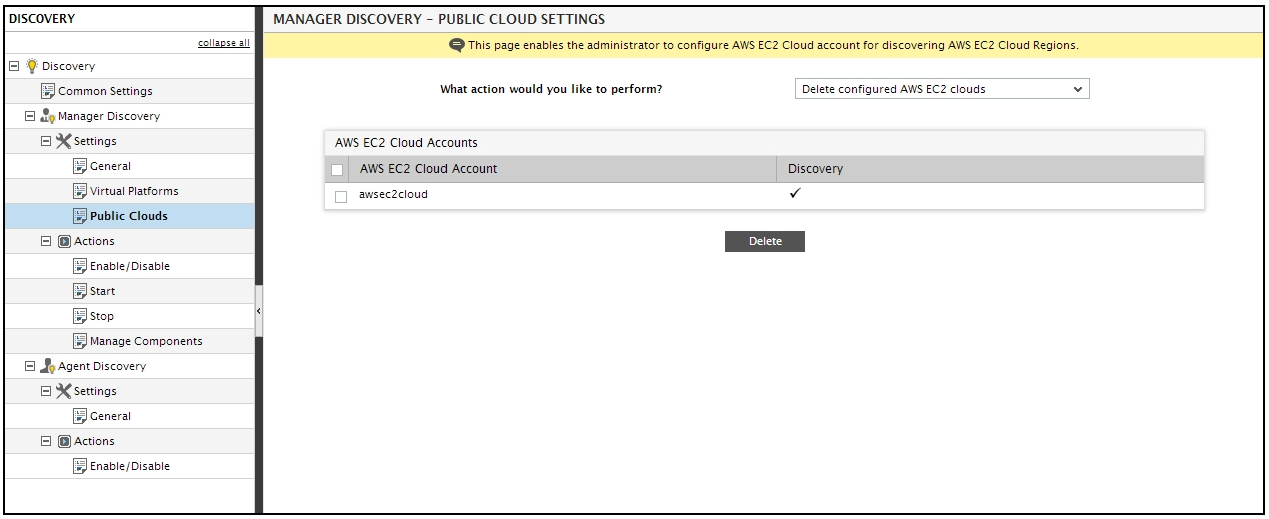

Figure 7.38: Deleting an AWS EC2 cloud accout

Discovering Citrix StoreFront Servers

eG Enterprise is capable of automatically discovering the Citrix StoreFront servers in the environment using the eG manager or the eG agent that is monitoring the Citrix NetScaler. eG Enterprise provides you with the option to directly connect to one/more Citrix NetScaler installations in your environment to perform Citrix StoreFront discovery. The additional benefit that accrues in this process is that, when one/more Citrix StoreFront servers discovered using NetScaler are managed, then eG Enterprise automatically uses the same Citrix NetScaler to collect performance metrics related to the Citrix StoreFront servers. In other words, eG Enterprise auto-configures the tests pertaining to the managed Citrix StoreFront servers with the details of the Citrix NetScaler used for their discovery; this way, the solution minimizes the time and effort involved in manual test configuration, and quickly starts collecting metrics from Citrix NetScaler. Similarly, if the Citrix NetScaler details need to be modified for any reason, then, you will not be required to manually reconfigure each test for this purpose; changing the configuration of the Citrix NetScaler in the **DISCOVERY** page will automatically update all the tests that have been configured to collect metrics from that Citrix NetScaler.

To discover the Citrix StoreFront servers via the Citrix NetScaler server, do the following:

1. Select the **Citrix NetScaler** option from the **Settings** node of the **Discovery** tree (see Figure 7.39).
2. Then, select the **Add new NetScaler** option from the **What action would you like to perform?** dropdown list.

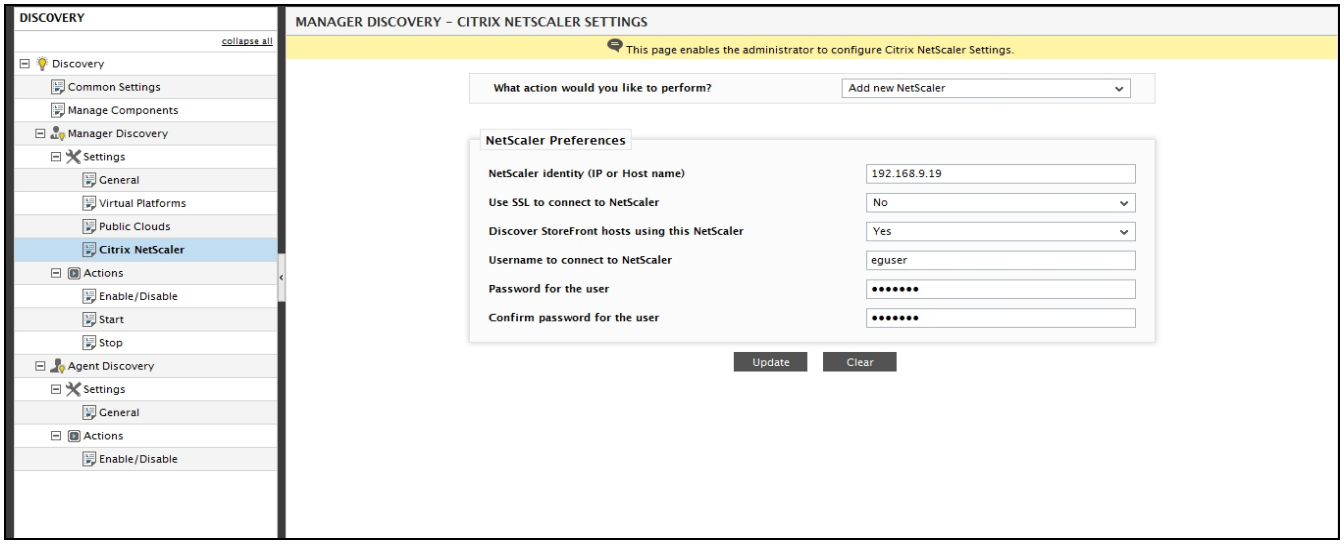

Figure 7.39: Configuring the Citrix NetScaler for discovering the Citrix StoreFront servers

3. Then, specify the following in the **NetScaler Preferences** section that appears (see Figure 7.39).
   - Specify the IP or host name of the NetScaler in the **NetScaler Identity (IP or Host name)** text box.
   - Then, indicate whether the eG manager is to connect to the Citrix NetScaler using SSL or not by selecting the **Yes** or **No** option from the **Use SSL to connect to NetScaler** list. By default, the **Yes** option is chosen from this list.
   - An IT environment may consist of multiple Citrix NetScaler installations, each managing a different set of Citrix StoreFront servers. To enable the eG manager to automatically discover those Citrix StoreFront servers via the Citrix NetScaler server being added for monitoring, the **Discover StoreFront hosts using this NetScaler** list is set to Yes, by default. If you do not want to discover

the Citrix StoreFront servers, select No from this list.

- In order to discover Citrix StoreFront servers using the Citrix NetScaler, the eG manager should connect to the Citrix NetScaler using valid user credentials. Provide the user name and password of such a user in the **Username to connect to NetScaler** and **Password for the user** text boxes. You can also use the credentials of a user with *'Read-only'* privileges to the Citrix NetScaler. If such a user pre-exists, then, provide the name and password of that user in the text boxes mentioned above. Otherwise, assign the *'Read-only'* role to a local/domain user to the Citrix NetScaler, and provide the name and password of this user in the **Username to connect to NetScaler** and **Password for the user** text boxes.
- Confirm the password of the user by retyping it in the **Confirm password for the user** text box.
- To clear all the configured details, click the **Clear** button. To start discovery instead, click the **Update** button.
- Upon clicking the **Update** button, Figure 7.40 will appear requesting you to confirm whether you want to start discovery based on the specifications provided. Click **Yes** button to begin discovery.

Figure 7.40: Confirm the discovery of the Citrix StoreFront servers

- Regardless of the discovery approach used (discovery using an IP range or using the Citrix NetScaler), triggering the Citrix StoreFront discovery will lead you straight to the **COMPONENTS - MANAGE / UNMANAGE** page.

4. To view the Citrix NetScaler server configured for discovering the Citrix StoreFront servers, select the **View Configured NetScalers** option from the **What action would you like to perform?** dropdown list (see Figure 7.39). Then the list of Citrix NetScaler servers that have been configured for discovering the Citrix StoreFront servers will appear as shown in Figure 7.41.

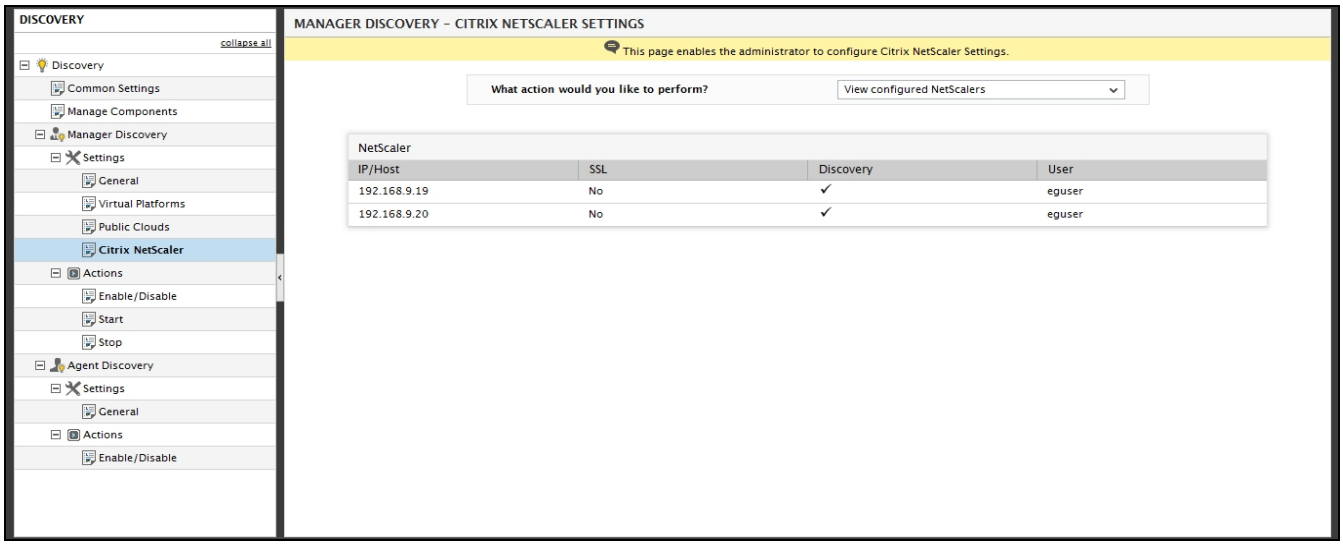

Figure 7.41: Viewing the configured Citrix NetScaler servers

5. To modify a Citrix NetScaler configuration, select the **Modify Configured NetScaler** option from the **What action would you like to perform?** dropdown list. Figure 7.42 will then appear displaying the existing configuration of the Citrix NetScaler. You can modify any of the displayed details and update the changes by clicking the **Update** button.

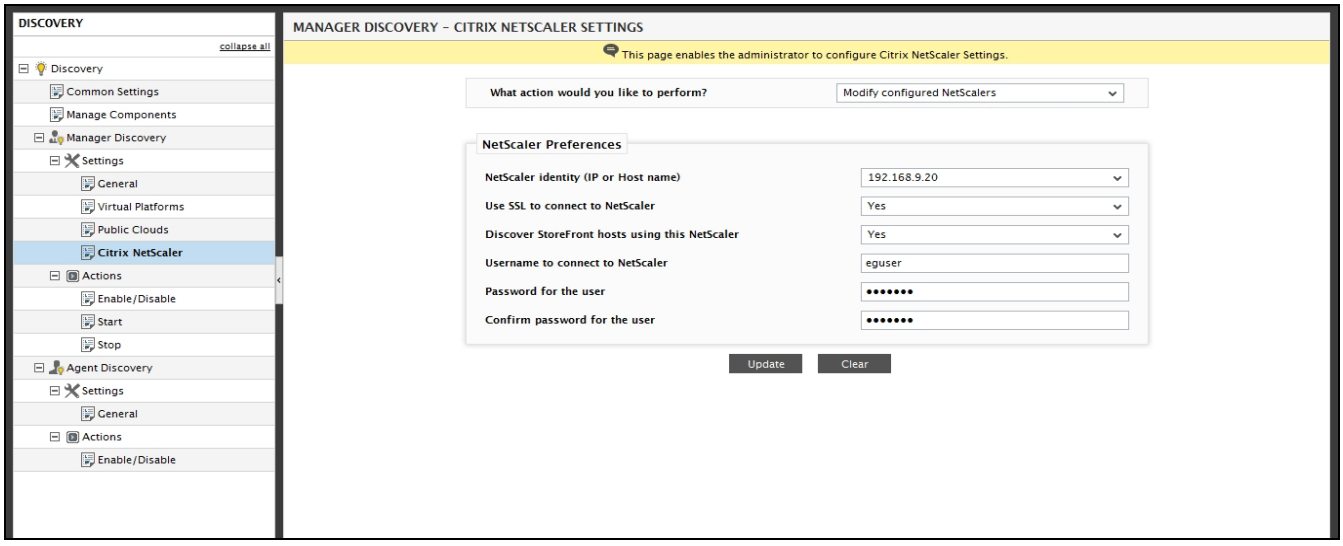

Figure 7.42: Modifying the credentials of the configured Citrix NetScaler server

6. To delete a particular Citrix NetScaler server, select the **Delete Configured NetScaler** option from the **What action would you like to perform?** dropdown list from Figure 7.39. Figure 7.43 then appears listing all the configured Citrix NetScaler servers. To delete a configured Citrix NetScaler server, select the check box corresponding to it and click the **Delete** button. To mark all the listed Citrix NetScaler servers for deletion simultaneously, simply select the top-most check box in the column of check boxes and click on the **Delete** button.

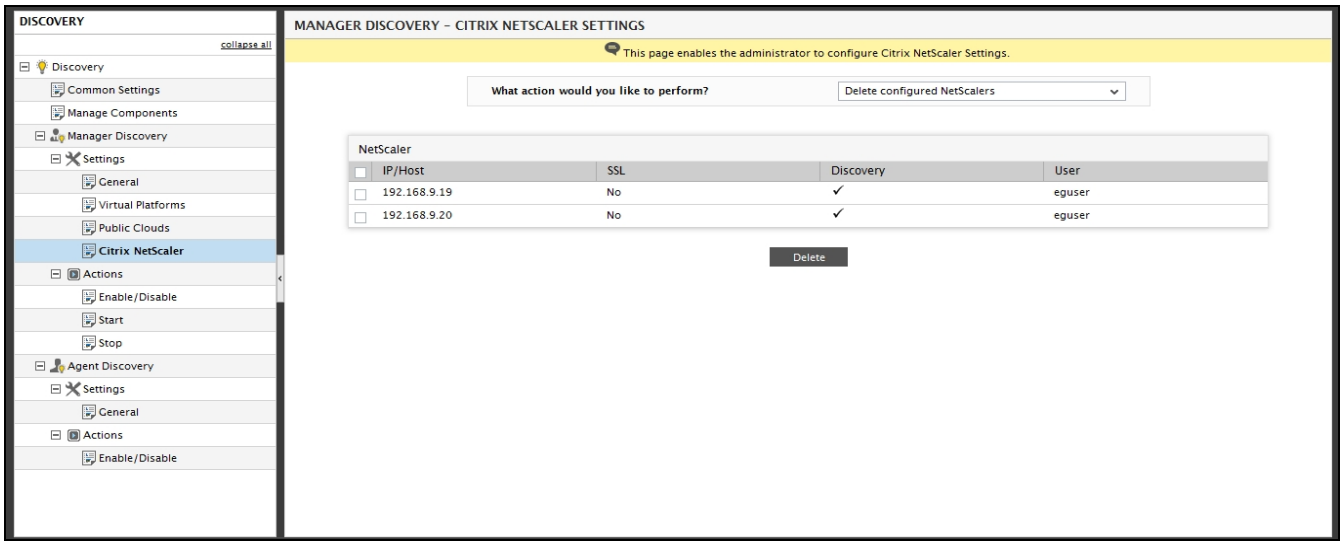

Figure 7.43: Deleting the Citrix NetScaler servers configured for discovery

7.1.2 Discovering Components Using the eG Agents

7.1.2.1 Enabling/Disabling Component Discovery by eG Agents

By default, it is the eG manager which discovers components in the target environment. To enable the eG agents to perform component discovery instead, do the following:

1. Select the **Enable/Disable** sub-node under the **Actions** node of the **Agent Discovery** node of the **DISCOVERY** tree of Figure 7.9.
2. The right panel will then change as shown by Figure 7.44.

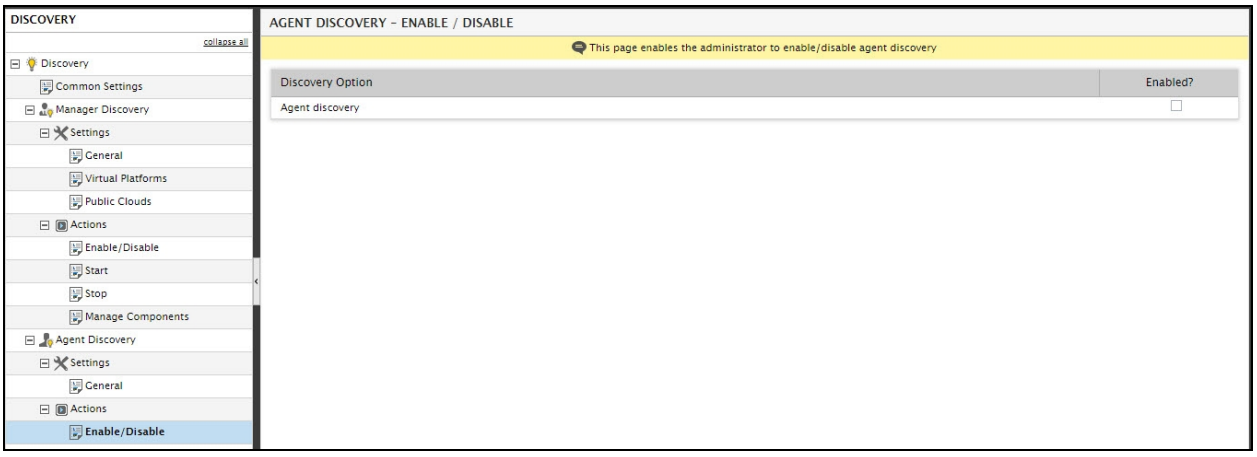

Figure 7.44: Enabling agent discovery

3. As is evident from Figure 7.44, **Agent discovery** is disabled by default. To enable it, just click the check box against **Agent discovery**.
4. Doing so brings up the message box depicted by Figure 7.45, which prompts you to confirm whether/not

you want to enable component discovery by eG agents. Click **Yes** to confirm enabling.

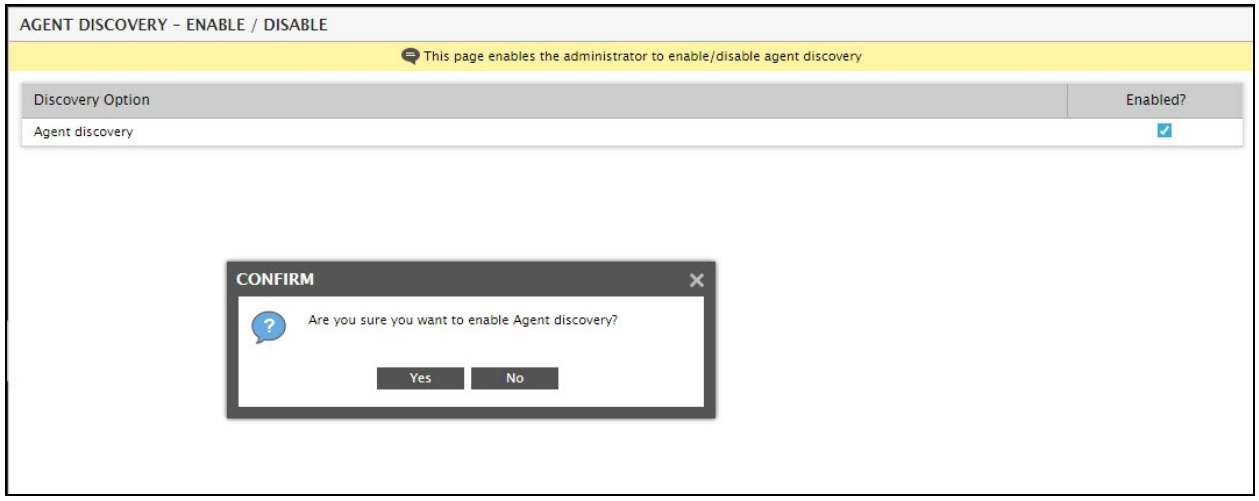

Figure 7.45: A message box requesting your confirmation to enable component discovery using eG agents

5. Clicking **Yes** in the message box of Figure 7.45 also enables the eG agents' capability to auto-discover inter-application dependencies – in other words, the capability to auto-discover the 'topology'. Accordingly, as soon as **Yes** is clicked in the message box above, a **Topology discovery** flag will appear (as shown by Figure 7.46), which will be selected by default.

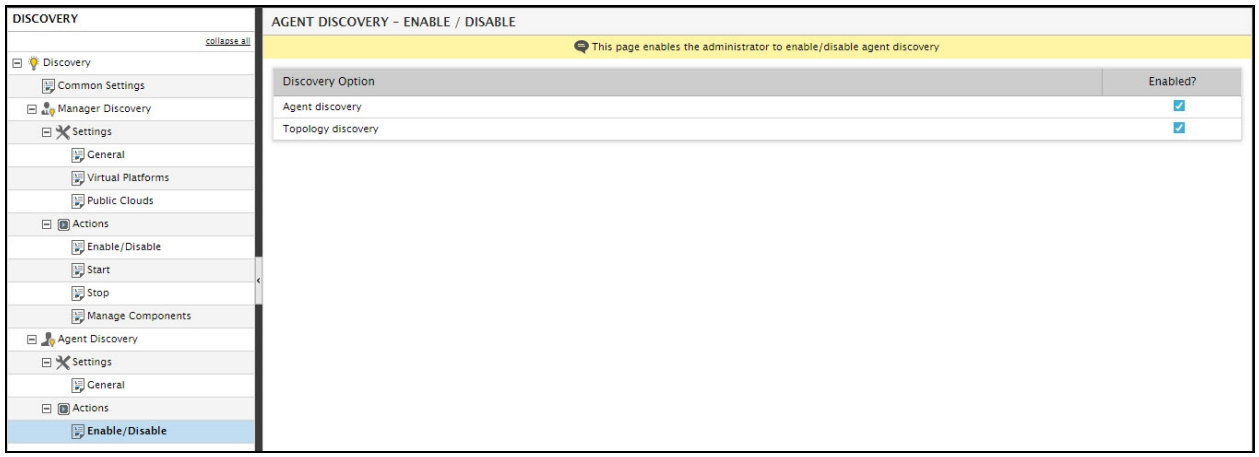

Figure 7.46: Enabling automatic topology discovery

6. If you want to disable topology discovery, just deselect the **Topology discovery** check box in Figure 7.46.
7. If you want to disable component discovery by eG agents, then just deselect the **Agent discovery** check box in Figure 7.46. **Note that once 'Agent discovery' is disabled, 'Topology discovery' will automatically be switched off.**

7.1.2.2 Configuring General Settings for Agent Discovery

To configure the settings that govern component discovery by eG agents, do the following:

1. Click the **General** sub-node under the **Settings** node of the **Agent Discovery** node in the **DISCOVERY** tree of Figure 7.46.
2. An **Agent Discovery Settings** page will then appear in the right panel (see Figure 7.47).

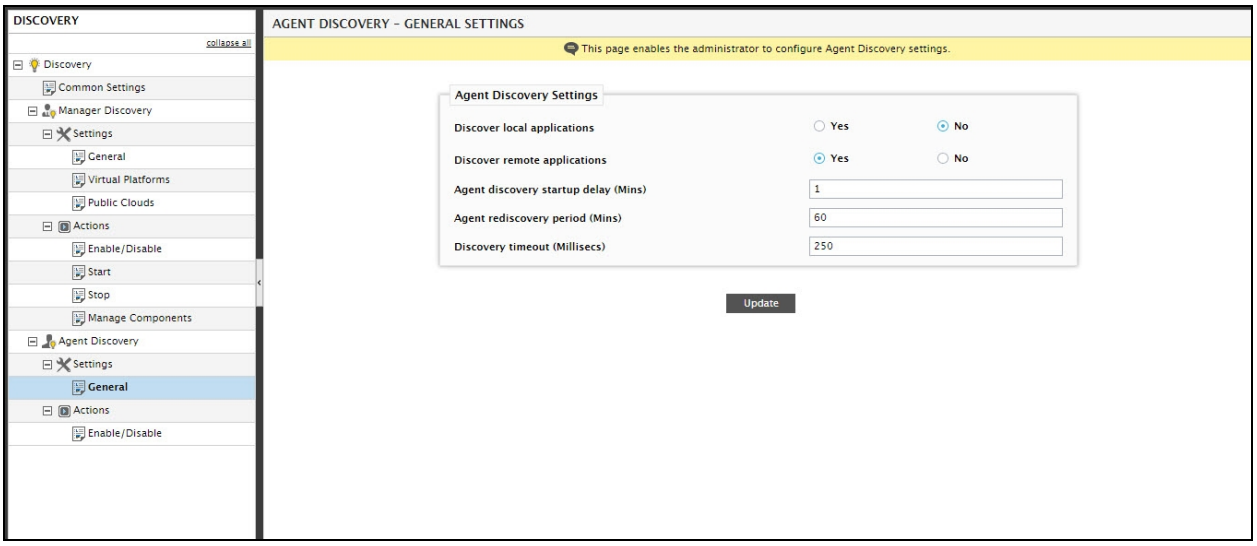

Figure 7.47: Defining settings for agent discovery

3. By default, the eG agent will discover only those applications that are running in the local system. Accordingly, the **Discover local applications** flag is set to **Yes** by default. If you do not want the eG agent to discover local applications, set this flag to **No.** Likewise, by default, the eG agent will not discover related applications running on remote hosts. This is why, the **Discover remote applications** flag is set to **No** by default. If you want the eG agent to discover remote applications, set this flag to **Yes**. If you set both flags to **No**, then agent discovery will not take place, even if the discovery capability is explicitly enabled.
4. In the **Agent discovery startup delay (Mins)** text box in the right panel, indicate the duration (in minutes) for which the agent needs to wait before beginning discovery.
5. To indicate how frequently the agent should rediscover components, set a time period (in minutes) against the **Agent rediscovery period (Mins)** field.
6. You can also instruct the eG agent to timeout if no components are discovered from an agent host beyond a specified duration. This duration needs to be mentioned in the **Discovery timeout (Millisecs)** text box.
7. Finally, click the **Update** button.

Based on these **General** settings and the **Common** settings discussed in Section Section **7.1.1.2** of this document, the eG agents then proceed to perform discovery.

Note:

Typically, if environment discovery is performed using the eG agent, then the agent is first installed and started so as to initiate the discovery procedure. The discovered components are then pushed to the eG manager by the agent. When this happens, the agent also checks the manager for applications that require monitoring. However, since the administrator would not have managed any component at this stage, the agent will have no information to download. Therefore, the agent remains inactive for 1 minute, after which it polls the manager yet again for downloadable information. The eG agent determines this poll frequency using

an exponential backoff algorithm - this is typically 1.5 times the previous poll period - for instance, if the agent polls the manager on the 3rd minute but finds nothing to download, then it will sleep and poll the manager again on the 4.5th minute (3*1.5). At this rate, if no component is managed by the administrator for a long time, it would be over 1 hour before the agent starts reporting metrics to the manager

To avoid this, by default, during the first 2 days after an eG agent is started, the agent polls the manager more frequently, so that administrators can start viewing performance results soon after configuration of the agent. During this default period (2 days), the agent checks every 10 minutes (approximately) for instructions from the manager. After this period, the sleep time begins to grow in the same manner mentioned previously. This default behavior is governed by the **IncreaseAgentSleepTimeAfter** flag in the **[AGENT_SETTINGS]** section of the **eg_tests.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory). By default, this is set to **2** (days) from the agent start time. To ensure that the agent poll frequency remains at 10 minutes for a more number of days since agent startup, change the value of the **IncreaseAgentSleepTimeAfter** flag.

7.1.2.3 Automatically Discovering the Topology Using the eG Agent

Whether an infrastructure is virtual or physical, inter-dependencies exist between applications. For example, a web server uses a middleware application server, and an application server relies on a database server. eG Enterprises uses this inter-dependency information for root-cause diagnosis – so administrators can determine where exactly the problem lies and where the effects are.

eG Enterprise includes the capability to auto-discover inter-application dependencies. This auto-discovery reduces the time and effort involved in setting up the performance monitoring solution and also reduces the human errors that can be involved in manual specification of inter-application dependencies.

In the eG Enterprise system, segment/service **topology definitions** embed the **inter-application dependencies**. Discovery of this topology information is initiated by the agents.

By default, the ability of the eG agent to **automatically discover topology** is enabled. However, this default setting will take effect only if the **eG agent has the ablity to automatically discover components**. This is because, topology discovery cannot be performed without component discovery. This is why, as soon as the **Agent discovery** flag is turned on (see Figure 7.46), the **Topology discovery** flag is also enabled.

If **Topology discovery** is enabled, then, in the right panel of Figure 7.47, a **Topology Discovery Settings** section will appear below the **Agent discovery settings** section. The settings that can be defined in the **Topology Discovery Settings** section are depicted by Figure 7.48.

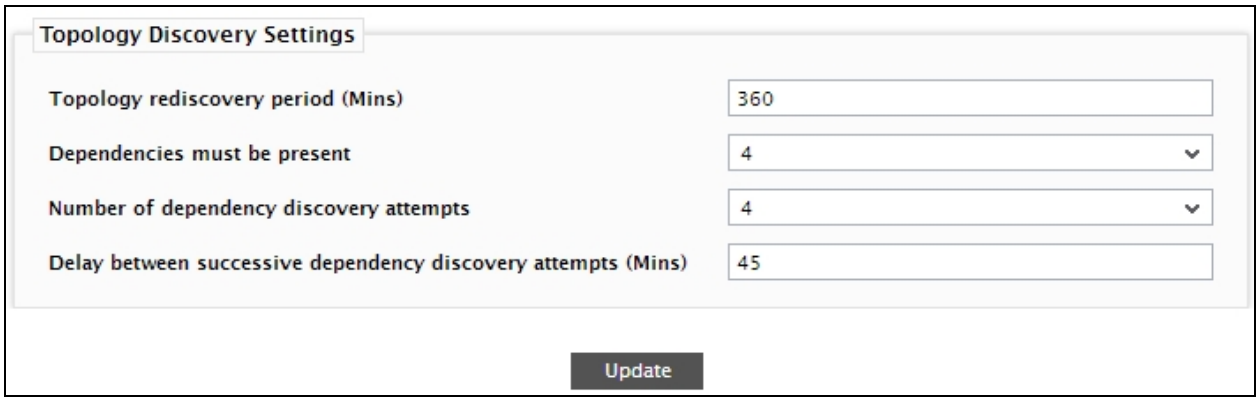

Figure 7.48: Defining topology discovery settings

Once automatic topology discovery is enabled, the eG agent will run the **netstat** command on the target host every 45 minutes (by default) to determine which applications are operating on the host and which port they listen to. This default frequency can be overridden using the **Delay between successive dependency discovery attempts (Mins)** parameter in **Topology Discovery Settings** section of Figure 7.48. By default, this parameter is set to **45** by default, indicating that the default discovery frequency is 45 minutes. To override this default frequency, specify a different duration (in minutes) against the **Delay between successive dependency discovery attempts (Mins)** parameter.

Whenever the eG agent on a host runs **netstat**, it retrieves a list of ports that are operating on that host. While some of these TCP ports may be standard listening ports - i.e., TCP ports at which the applications executing on that host listen for requests from remote hosts/applications - a few other TCP ports may be local ports created dynamically on the host for a temporary purpose. To clearly differentiate between listening ports and local ports, the eG agent does the following:

- By default, the eG agent compares the ouput of four consecutive executions of the **netstat** command on a host. If a port number is repeated in all the four **netstat** outputs by default, then that port number is counted as a 'server listening port'. This default behavior is governed by the **Dependencies must be present** and the **Number of dependency discovery attempts** parameters in the **Topology Discovery Settings** section of Figure 7.48. As can be inferred from Figure 7.48, both these parameters are set to **4** by default. The **4** against the **Number of dependency discovery attempts** parameter indicates the maximum number of consecutive **netstat** outputs to be considered for identifying the server listening ports. The **4** against the **Dependencies must exist** parameter indicates the minimum number of **netstat** outputs in which a port number should appear for it to be considered as a 'server listening port'. You can change this if you need. For instance, if you set the **Dependencies must exist** parameter to 3 and let the **Number of dependency discovery attempts** parameter to remain at **4**, then the eG agent will count the port numbers that appear in at least **3** out of **4** consecutive **netstat** outputs as active listening ports.
- Once the listening ports are identified, the agent then closely observes traffic to and from a 'server listening port', identifies the remote applications that frequently connect via this port, and thus automatically discovers the inter-relationships that exist between applications in an IT infrastructure. The interdependencies that are so discovered are then sent to the eG manager. On the other hand, all those port numbers that do not conform to the specification governed by the **Dependencies must be**

present and the **Number of dependency discovery attempts** parameters are counted as local ports. All traffic to local ports are hence disregarded for the purpose of topology auto-discovery.

By default, the whole cycle of operations - beginning with isolating the listening ports to discovering inter-dependencies to reporting the discovery to the eG manager - takes 180 minutes (as indicated by the default value **4** against the **Dependencies must be present** setting) to complete. After which, the eG agent will wait for another 60 minutes (i.e., 1 hour, by default) to rediscover the topology. In other words, all the aforesaid activities will be performed again by the eG agent every hour after the first cycle is complete. Like other default settings, you can also override the frequency with which topology is rediscovered by the eG agent. For this, use the **Topology rediscovery period (Mins)** parameter, which is set to **360** by default.

7.2 Managing and Unmanaging Components

Soon after component discovery, proceed to manage/unmanage the discovered components. Component management is a necessary step in monitoring because, administrators of IT infrastructures may not want to monitor all components that are discovered. For instance, an IT infrastructure may have a number of staging and testing components, in addition to their production components. An administrator may want to have only the production components monitored by eG Enterprise. To enable administrators monitor discovered servers selectively, eG Enterprise provides the **COMPONENTS – MANAGE / UNMANAGE** page. For managing discovered systems, administrators can use the **SYSTEMS – MANAGE/UNMANAGE** page. The sub-sections that follow discuss how these pages can be used for managing servers/systems.

7.2.1 Managing/Unmanaging Servers

The **COMPONENTS - MANAGE/UNMANAGE** page allows you to manage/unmanage servers. To access this page, do one of the following:

a. Select the **Manage/Unmanage/Delete** option from the **Components** menu in the **Infrastructure** tile, (OR)

b. Select the **Manage Components** sub-node under the **Actions** sub-node of the **Manager Discovery** node in the **DISCOVERY** tree of Figure 7.47. To access the **DISCOVERY** tree, select the **Discovery** option from the **Components** menu of the **Infrastructure** tile.

Figure 7.49 will then appear, using which you can manage certain server applications and unmanage the others. For each component type that an administrator chooses from the **Component type** list box, he/she can view the components being managed and those that are unmanaged. The component-types listed in the **Component type** list box are the ones that have been discovered by the discovery process. To filter the component-types listed in the **Component type** drop-down, you can use the **Show managed component types only** check box. By default, this check box will be unchecked, indicating that the **Component type** drop-down will list all discovered component types by default. Select the **Show managed component types only** check box if you want the **Component type** drop-down to list only those component types with at least one managed component. This way, you can filter out all those component types without any managed components.

Also, in the **Component type** list, any host that responds to ICMP ECHO messages but which does not have an eG-managed application executing on it will be listed as a **Generic server**. For a generic server, an eG agent monitors the server's CPU, memory, disk, and network statistics. In addition, if desired, specific application processes executing on the server can be monitored. If a user wants to monitor a system for its CPU and

memory utilization, but not individually monitor the applications executing on it, then he/she can manage that system as a **Generic server**.

Note:

eG's discovery process automatically discovers network routers alone. Other network devices, such as switches, hubs, etc. have to be manually added for monitoring, via the administrative interface.

Once a **Component type** is chosen, all newly discovered components of that type will populate the **Unmanaged components** list in Figure 7.49. When eG Enterprise discovers a component, it will mark the component as being unmanaged and prefix it with an asterix (*) to indicate that this component has been newly discovered.

To manage one/more of the newly discovered components of a chosen **Component type**, select the components of interest from the **Unmanaged components** list, and click the **<** button in Figure 7.49. This will transfer the selection to the **Managed components** list. To unmanage a managed component, select the component from the **Managed components** list andclick the **>** button in Figure 7.49. Finally, click the **Update** button.

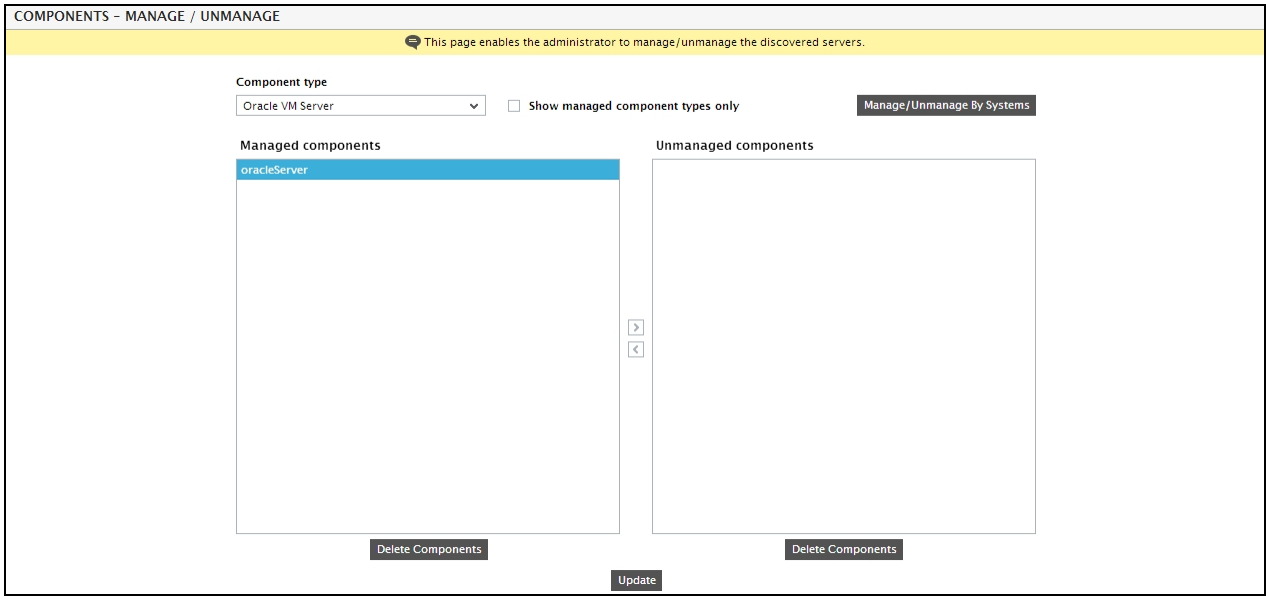

Figure 7.49: Configuration of components that are to be managed by eG Enterprise

You can also permanently delete a managed/unmanaged component using Figure 7.49. For this purpose, select that component from the **Managed components** or **Unmanaged components** list (as the case may be), click the **Delete Components** button below the corresponding list, and click the **Update** button.

Note:

- By default, deleted components cannot be automatically discovered by the eG manager's rediscovery process. You can however override this default behavior when configuring manager discovery. To achieve this, follow the steps below:
- Select the **Discovery** option from the **Components** menu of the **Infrastructure** tile.

- In the **DISCOVERY** tree of the page that then appears, the **General Settings sub-**node under the **Manager Discovery** node will be chosen by default. Accordingly, the right panel will display the **MANAGER DISCOVERY - GENERAL SETTINGS** page (see Figure 7.10).
- In the right panel, the **Re-discover deleted components** flag will be set to **No** by default. To make sure that the discovery process rediscovers deleted components, set this flag to **Yes**.
- Finally, click the **Update** button to save the changes.
- eG Enterprise disallows the deletion of components that are part of segments/services.

Note:

If your eG license enables **Metric Aggregation**, then, as soon as you manage the very first component of a type, a corresponding *aggregate component type* will be automatically created by the eG Enterprise system. For instance, as soon as you manage the first *IIS web* server in your environment using the **MANAGE / UNMANAGE** page, a corresponding *IIS Web Aggregate* component-type will be dynamically generated by the eG Enterprise system. You can add one/more *aggregate components* of the *aggregate types* so generated by following the procedure discussed in theMetrics Aggregation chapter of this document.

On the other hand, if you do not create any *aggregate component* for an *aggregate type*, then, as soon as you delete/unmanage all 'non-aggregate' components of the non-aggregate type that corresponds to it, the *aggregate type* will cease to exist. In other words, if you unmanage/delete all the managed *IIS web* server components in your environment, then, the corresponding *IIS Web Aggregate* component-type will be automatically removed from the eG Enterprise system, if no *aggregate component* of that type pre-exists.

7.2.2 Managing an Oracle Database Server

While most components are identified uniquely by a host and port number, Oracle database servers are identified uniquely by a combination of host, port number, and an instance identifier (referred to as SID). While eG's discovery process can auto-discover an Oracle database server, it cannot auto-discover the SID(s) of that server. This is why, any newly discovered Oracle database server will not have an SID suffix. If such an Oracle database server is later managed using the **COMPONENTS – MANAGE/UNMANAGE** page, you will have to manually change the profile of that server in the eG Enterprise system to configure its SID. If this is not done, eG Enterprise cannot monitor the Oracle database server.

This is why, if an Oracle database server is managed without its SID, the **MANAGER NOTIFICATIONS** window keeps alerting the user to the absence of the SID and prompts the user to configure the SID for the server (see Figure 7.50).

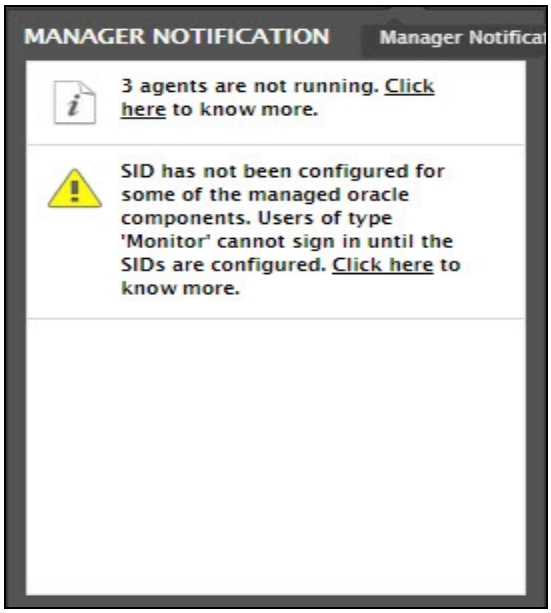

Figure 7.50: The Manager Notification window alerting users to the absence of the SID of a managed Oracle database server

To configure the SID of the managed Oracle database server, do the following:

- Click the **Click here** hyperlink in the SID-related alert message displayed in the **MANAGER NOTIFICATION** window of Figure 7.50.
- Figure 7.51 will then appear listing all Oracle database servers that have been managed without an SID.

Figure 7.51: List of managed Oracle database servers without an SID

- To configure the SID, click on any of the Oracle database servers listed in Figure 7.51. The profile of that Oracle database server will then appear as depicted by Figure 7.52.

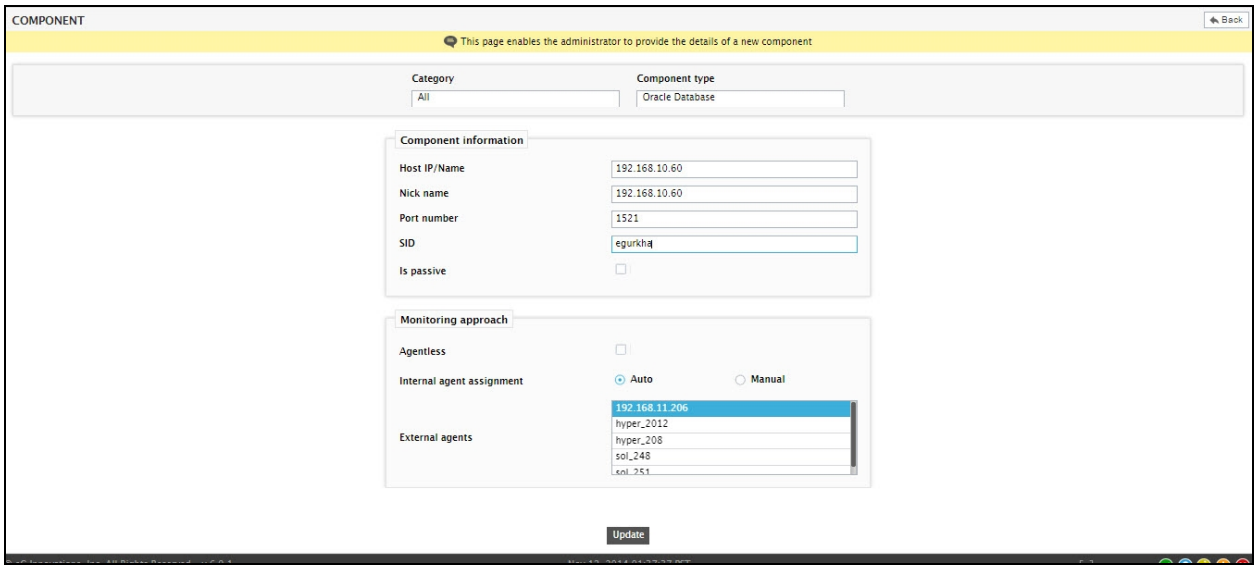

Figure 7.52: Configuring the SID of the Oracle database server

- In Figure 7.52, enter the **SID** of the Oracle database server. You can also configure multiple SIDs as a comma-separated list.
- Finally, click the **Update** button.
- If, for any reason, you do not want to configure the SID of a managed Oracle database server, you can unmange that server instead. For this, click on the **Click here** link in Figure 7.51. Figure 7.53 will then appear. From the **Managed components** list of Figure 7.53, pick the Oracle database server to be unmanaged and click the **>** button to unmanage it. Finally, click the **Update** button in 7.2.2.

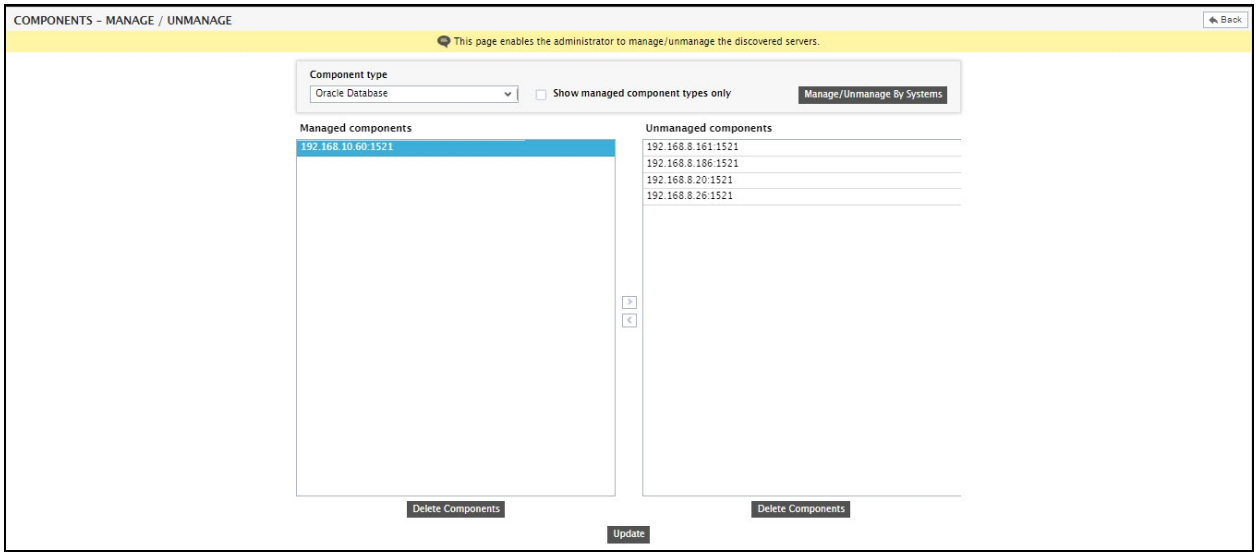

Figure 7.53: Selecting the Oracle database server that is to be unmanaged

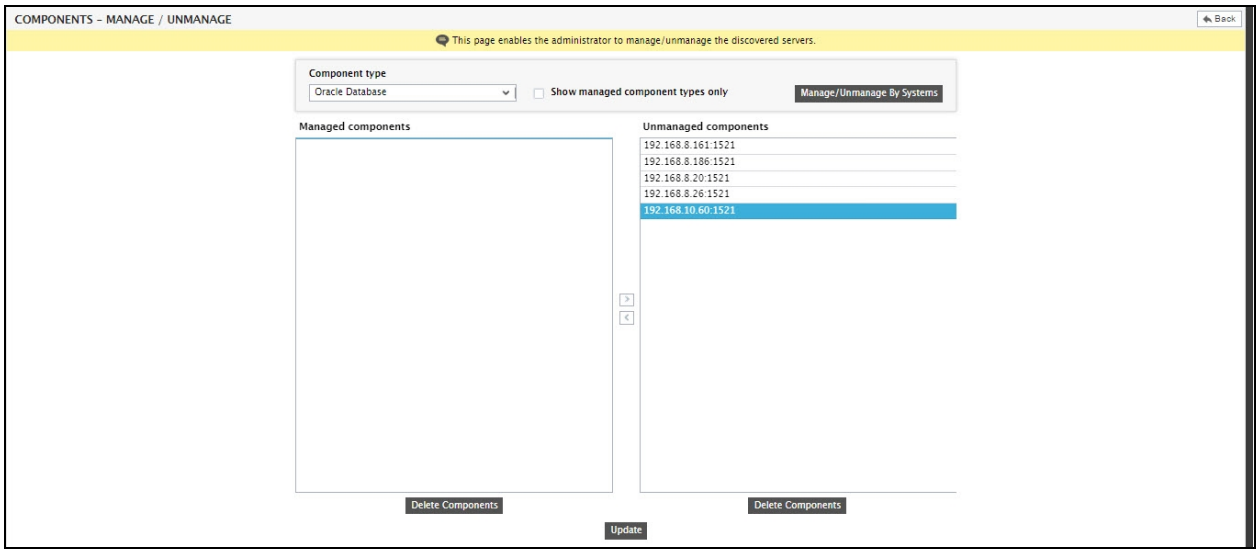

Figure 7.54: Unmanaging the Oracle database server

7.2.3 Managing/Unmanaging Systems

Using the **COMPONENTS - MANAGE/UNMANAGE** page, administrators can manage/unmanage only those components that belong to a chosen type. In environments where a large number of components of varied types are to be managed/unmanaged, using the **COMPONENTS - MANAGE/UNMANAGE** page could prove to be time-consuming. eG Enterprise therefore offers the **SYSTEMS - MANAGE/UNMANAGE** page (see Figure 7.55).

To access this page, first go to the **COMPONENTS – MANAGE/UNMANAGE** page by selecting the **Manange/Unmanage/Delete** option from the **Components** menu of the **Infrastructure** tile. Figure 7.53 will then appear. Click the **Manage/Unmanage By Systems** button in Figure 7.53 to open the **SYSTEMS – MANAGE/UNMANAGE** page (see Figure 7.55).

To manage components using the **SYSTEMS - MANAGE/UNMANAGE** page, follow the steps given below:

1. First, select the **Host / Nick name** to be managed from Figure 7.55. All components that are currently managed using the chosen host/nick name will be listed in the **Managed components** list. Components newly discovered with the chosen host/nick name will appear in the **Unmanaged components** list. To manage a newly discovered element, select it from the **Unmanaged components** list and click the **<** button (see Figure 7.55). This will transfer the selection to the **Managed components** list (see Figure 7.56). Likewise, you can unmanage a managed component by selecting it from the **Managed components** list and clicking the **>** button. Finally, click the **Update** button.

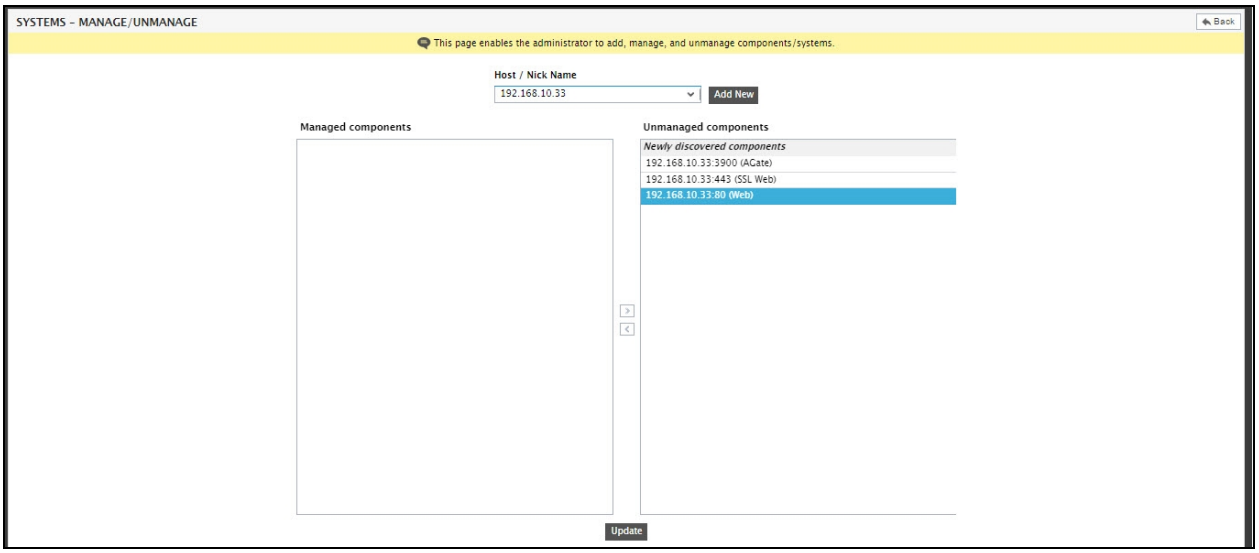

Figure 7.55: Selecting the component to manage

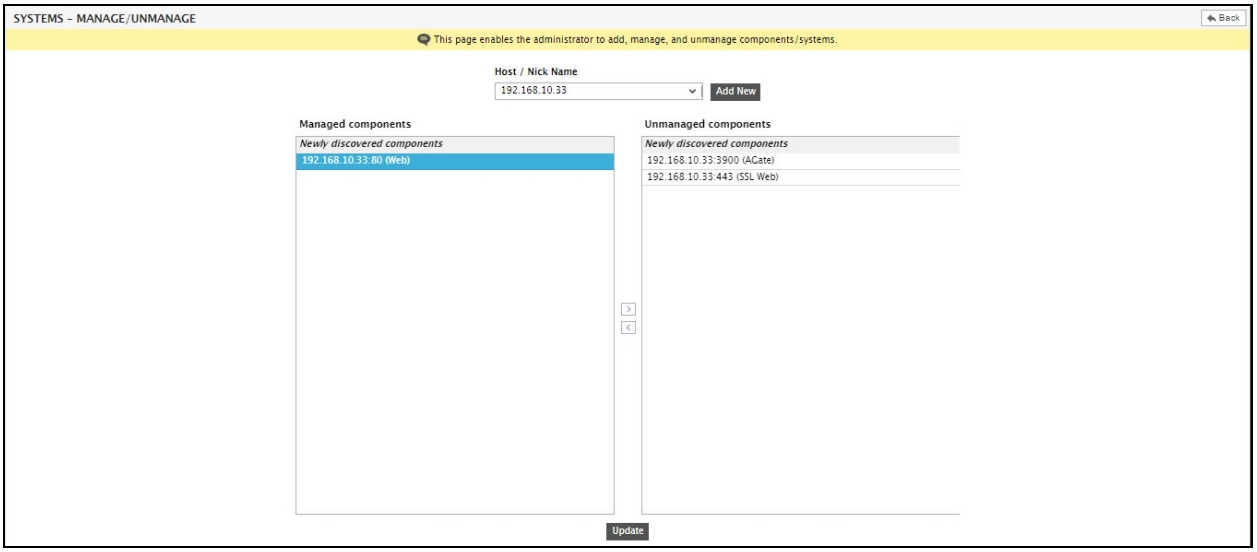

Figure 7.56: Managing the chosen component

2. If the component that you want to manage is not auto-discovered by the eG Enterprise system, it will not appear in the **Managed components** list. In this case, you have the option to manually add that component to the eG Enterprise system. For this, once you select the **Host/Nick Name** of the component that you want to manage from Figure 7.55, click the **Add New** button therein. Figure 7.57 will then appear.

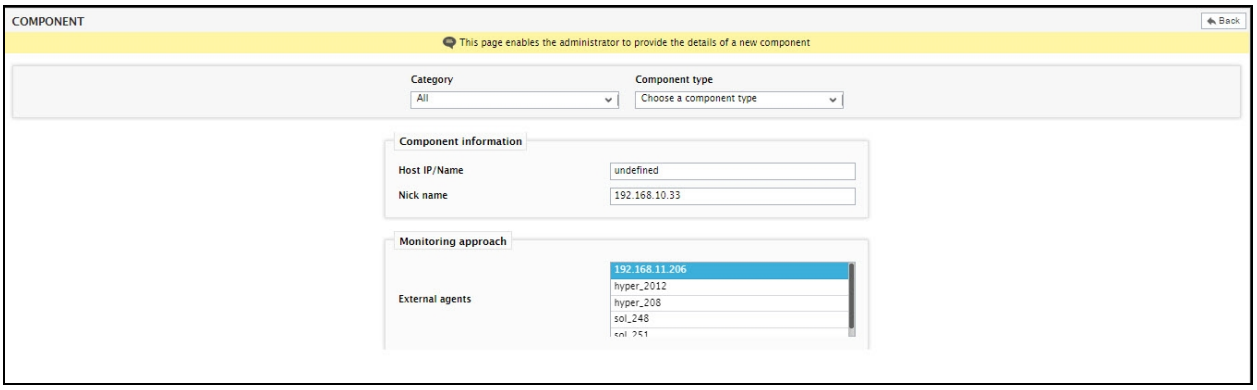

Figure 7.57: The COMPONENT page that appears when the Add New button in the SYSTEMS – MANAGE/UNMANAGE page is clicked

3. To add a new component with the **Host/Nick name** chosen from Figure 7.55, first select the **Component type** to which the new component belongs from Figure 7.57. Then, provide a **Host IP/Name** for the new component. This way, provide all other specifications that are relevant to the component being added using Figure 7.58, and click the **Add** button therein to add the new component.

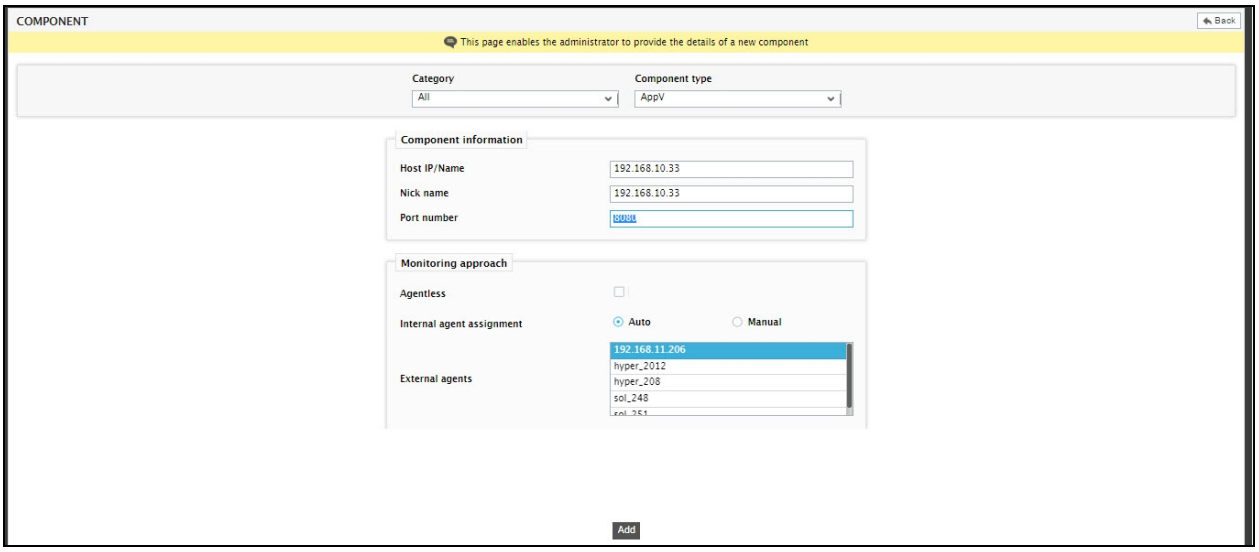

Figure 7.58: Adding a new component with the host/nick name chosen from the SYSTEMS – MANAGE/UNMANAGE page

Reference:

For complete details on how to manually add a new component using the eG administrative interface, refer to Section **7.3** of this document.

4. A message box depicted by Figure 7.59 will then appear informing you of the successful addition of the component. In addition, the message box will also prompt you to confirm whether/not you want to add more components using the same host/nick name. Clicking **Yes** will take you back to Figure 7.57, using which you can add another component. Clicking **No** in Figure 7.51 will lead you to Figure 7.60, where you will find the newly added component appear in the list of **Managed components** for the selected **Host/Nick name**.

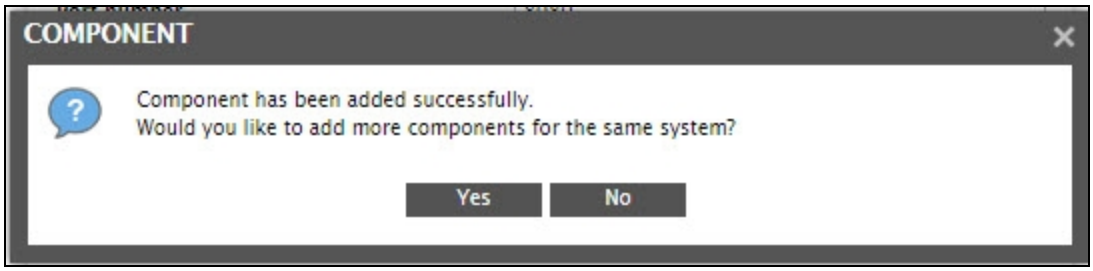

Figure 7.59: A message box requesting you to confirm whether/not you want to add more components for the chosen host/nick name

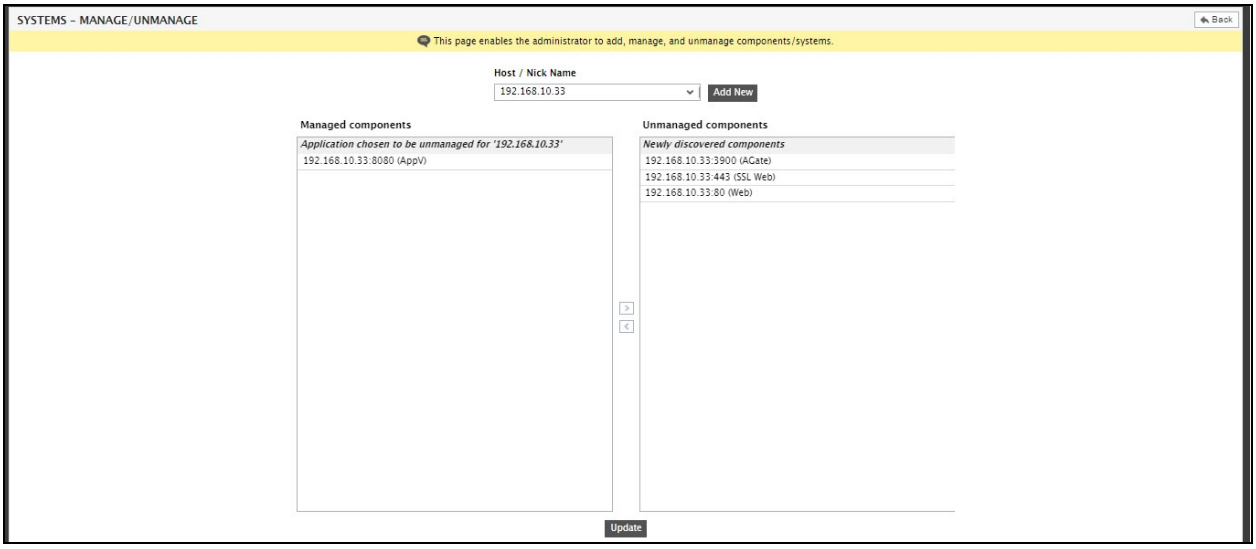

Figure 7.60: The newly added component appearing in the list of Managed components

7.3 Adding/Modifying/Deleting Components

7.3.1 Adding Components

In the cases in which eG Enterprise's discovery process is not able to discover a specific component or set of components (like network node, load balancers etc.), the **Add/Modify** option in the **Components** menu of the **Infrastructure** tile permits the user to explicitly add the component for monitoring by eG Enterprise. After choosing the **Add/Modify** option, first select the **Category** to which the component to be added belongs. In eG Enterprise, components are grouped into categories based on their functionality – for instance, Oracle and Microsoft SQL servers are grouped under the category Database servers and VMware vSphere and Citrix XenServers are grouped under the category Virtualization Platforms. If a specific **Category** is chosen, then only those component types that belong to the selected category will be available for selection in the **Component type** list of Figure 7.61. By default, All is chosen as the **Category**, which is why all component-types monitored out-of-the-box by eG Enterprise are by default available for selection in the **Component type** list.

Next, select the type of component to be managed from the **Component type** list (see Figure 7.61). By default, the **Component type** list includes all those component types that are supported out-of-the-box by the eG Enterprise Suite. To filter the **Component type** list, you can use the **Show managed component types only** check

box. By default, this check box will be unchecked, indicating that the **Component type** drop-down will by default list all component types that are supported out-of-the-box by the eG Enterprise Suite. Select the **Show managed component types only** check box if you want the **Component type** drop-down to list only those component types with at least one managed component. This way, you can filter out all those component types for which not even one component is managed.

For each type of component chosen from the **Component type** list, the administrator can add a new component for monitoring by choosing the **Add New Component** option Figure 7.61.

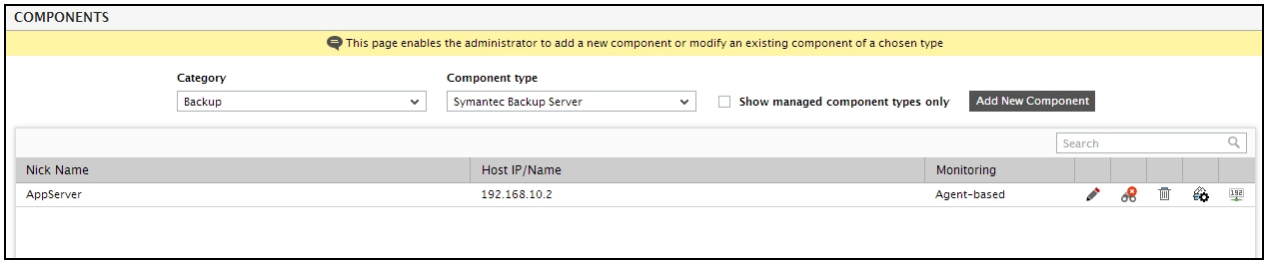

Figure 7.61: Manually adding a new component for monitoring by eG Enterprise

Figure 7.62 then appears, using which you can build a profile for the new component by specifying a variety of parameters.

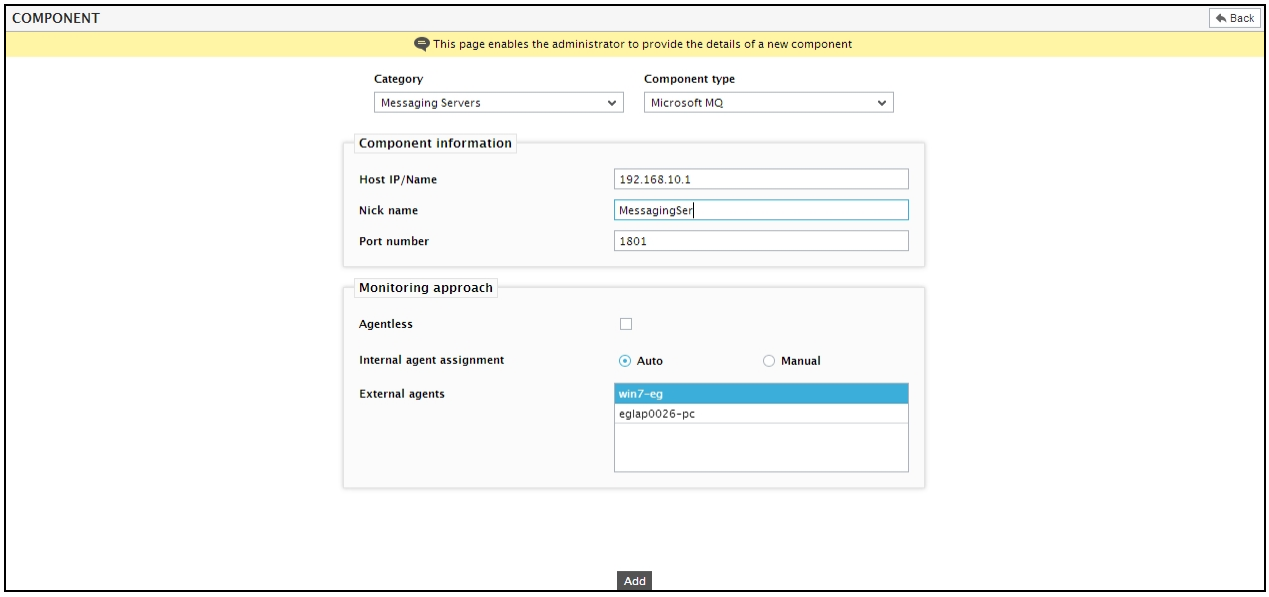

Figure 7.62: Configurations to be specified when adding a new component

These parameters may differ from one type of component to another. Figure 7.62 depicts the standard set of parameters that need to be defined when adding a new component. This includes the following:

- **Host IP/Name:** Specify the IP address or the 'fully qualified host name' of the component being added.

  Note:

If you choose to specify the IP address of a component against **Host IP/Name**, then remember that this IP address can be an IPv4 or an IPv6 address.

- **Nick name:** A nick name is a logical name that is associated with the host being monitored. The nick name specified here should match the ones specified while installing the eG agent.

  Note:

  - While specifying a nick name for a component, ensure that you do not use the same name as that of the *internal* name of the component-type. For instance, you cannot add a component of type *Generic_server* with the nick name *Generic_server*. To know the internal component type names, follow the procedure discussed in Page 285.
  - If required, both the **Host IP/Name** and **Nick name** parameters of a component can be configured with the IP address of that component. However, if the **Host IP/Name** parameter has been configured with the target component's IPv6 address, then you cannot configure the **Nick name** parameter also with that IPv6 address; in this case, the **Nick name** should be some other logical name by which you want to identify the component.

- **Port number:** The default port on which the component listens will be displayed. You can change this, if required. If a component type has multiple default ports, then the first port in the comma-separated list displayed/configured against the component type in 7.1.1 will be displayed here. For components that are not associated with ports, the port number has to be specified as NULL.

- **Agentless:** By selecting/deselecting the **Agentless** checkbox (see Figure 7.62), indicate whether the component being added is to be monitored in an agentless or agent-based manner. By default, this will be set to **No**.

  Reference:

  For more details on eG's agentless monitoring, refer to Section **7.6** below.

  Note:

  If two components have the same IP-nickname combination but different monitoring modes, then eG will not be able to monitor those components properly - i.e., if one component is managed as agent-based and the other, agentless, then eG will not be able to monitor both the components.

- **Internal agent assignment:** Next, indicate whether the eG Enterprise system's **Internal agent assignment** is **Auto** or **Manual**. This option will be available to you only if **Agentless** support is set to **No**. By default, the **Internal agent assignment** flag will be set to **Auto**.

  Reference:

  For more details on the **Internal Agent Assignment** parameter, refer to Section **7.7** below.

- **External agents:** From the **External agents** box that lists all the external agents that have been configured in the environment, select the external agent(s) that will monitor the component being added from an external perspective. By default, the eG agent on the eG manager's host serves as the external agent for the entire monitored environment. If required, you can configure additional external agents. To know how, refer to Section **7.4** of this document.

In the case of some components, a few additional parameters may have to be configured. For instance, when adding an Oracle database server, the name of the Oracle instance has to be specified in the **SID** text box. Multiple SIDs can also be provided as a comma-separated list.

Note:

If multiple **SID**s are specified for an Oracle database server, then eG Enterprise system will represent each such SID as a separate Oracle database server. For instance, if an Oracle server, 192.168.10.100:1521,has been added with the **SID**s – eg and egdemo – then, when this specification is updated, two new Oracle database servers will be added to the eG Enterprise system – namely, the server 192.168.10.100:1521:eg and the server 192.168.10.100:1521:egdemo.

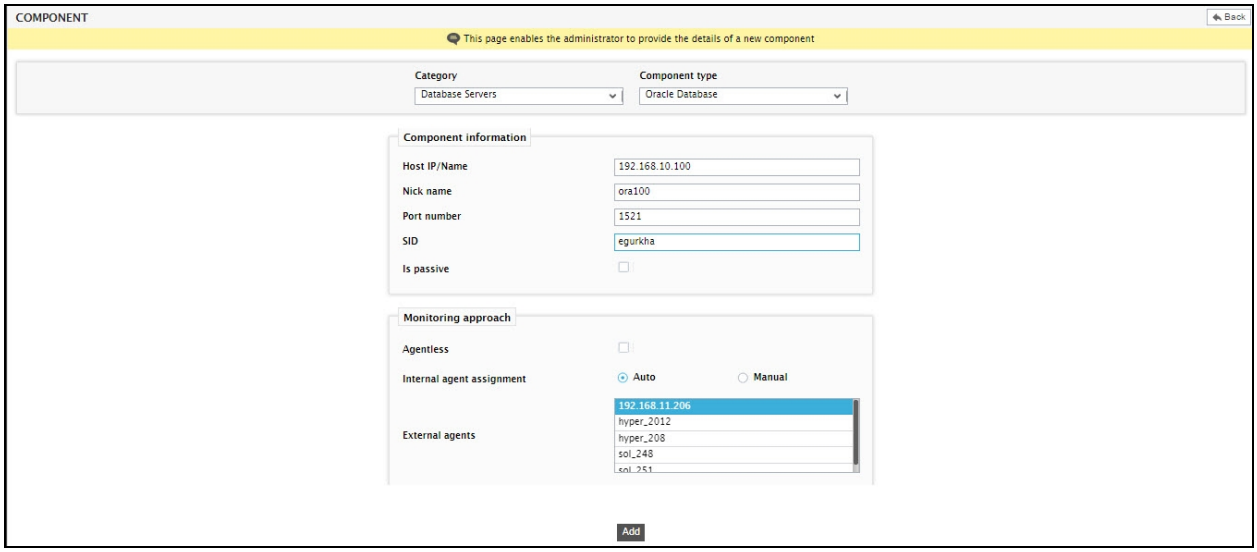

Figure 7.63: Parameters that need to be configured when adding an Oracle database server

Moreover, where cluster monitoring is supported, eG Enterprise enables users to indicate at the time of component addition, whether the component being added is an active or passive server in the cluster. Currently, eG Enterprise is capable of monitoring the following cluster services: Oracle, MS SQL, Active Directory, WebSphere MQ, and Exchange 2000/2003. This is why, while adding an *Oracle database*, *Microsoft SQL*, *Active Directory*, *WebSphere MQ,* or *Exchange 2000/2003* component, the **COMPONENT** page includes an additional **Is passive** flag (see Figure 7.63). This flag is unchecked by default, which implies that the component being added is an active server in a cluster, by default. If this flag is checked, it indicates that the server being added is a passive server in a cluster. In this case, no alerts will be generated if the server is not running. Measures will be reported as "Not applicable' by the agent if the server is not up.

In the case of the Microsoft web servers (e.g., IIS web or IIS SSL servers), an additional field called **MTS enabled** appears. eG's discovery process does not automatically discover Microsoft Transaction Servers (MTSs). Using the **MTS enabled** option, an administrator can specify whether an MTS server is executing on the same host as a Microsoft web server (IIS or IIS SSL). If required, the administrator can associate an MTS server with the host that executes the Microsoft IIS web server or SSL server.

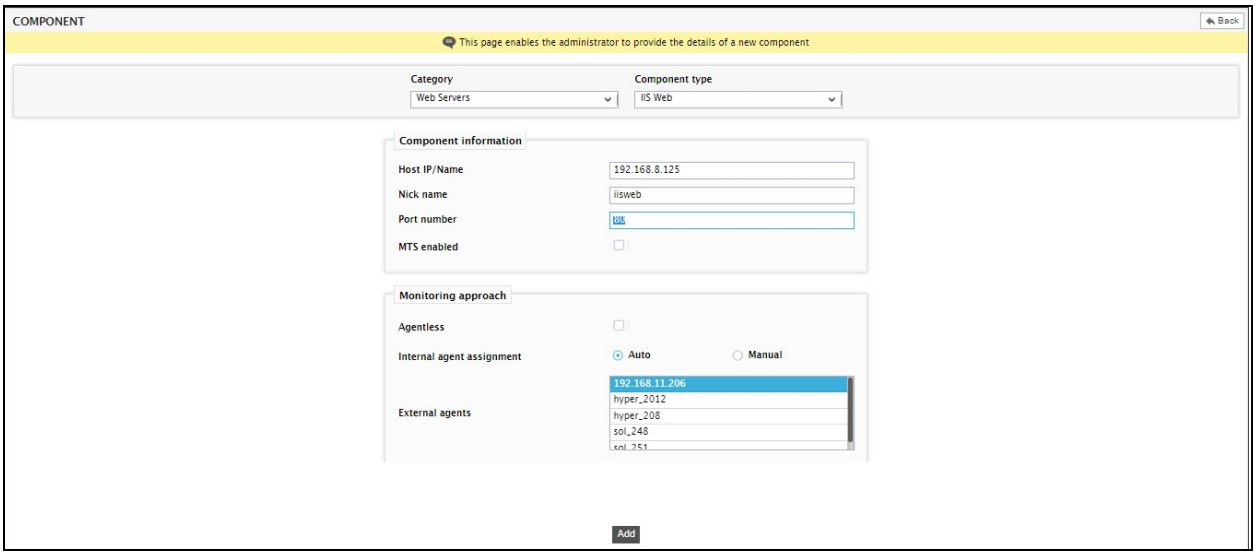

Figure 7.64: Parameters to be configured when adding an IIS web server

Once values are provided for all parameters, click the **Add** button in Figure 7.64 to add the new component. As soon as the **Add** button is clicked, the eG manager checks the existing IP address to host name mapping of the component. If the specified host name or nick name does not exist in the existing mapping, the user is alerted to ensure that additional care is taken while specifying the host name (see Figure 7.65). Note that Figure 7.65 allows you to choose between assigning the specified host name to the new component alone, and assigning it to all components with the specified IP. Choose the required option and click on the **Update** button. To change the hostname, click on the **Back** button to return to the **COMPONENT** page.

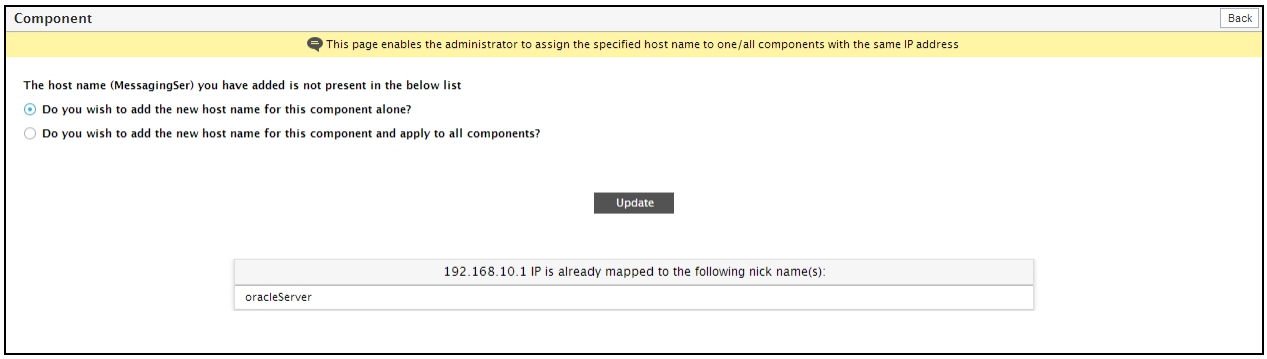

Figure 7.65: Hostname verification

7.3.2 Modifying Component Details

At any given point in time, you can modify the profile of a component using the following steps:

1. Select the **Add/Modify** option from the **Components** menu of the **Infrastructure** tile.
2. When Figure 7.66 appears, select a component **Category** and then the **Component type** to which the component to be modified belongs.

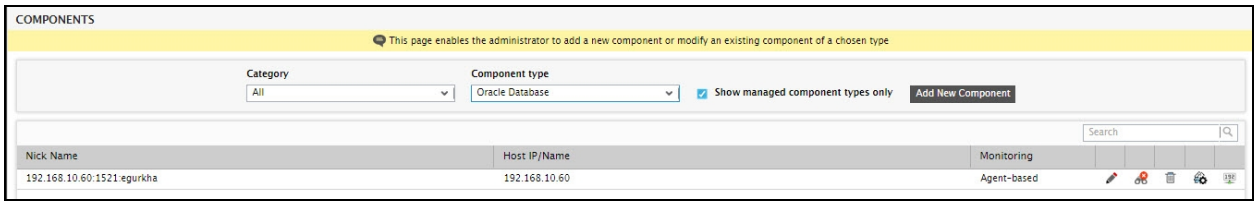

Figure 7.66: Selecting the component type to be modified

3. If components of the chosen type pre-exist in the eG Enterprise system, the same will be displayed as depicted by Figure 7.66.
4. To modify the details of a component, just click the icon corresponding to that component.
5. Figure 7.67 will then appear.

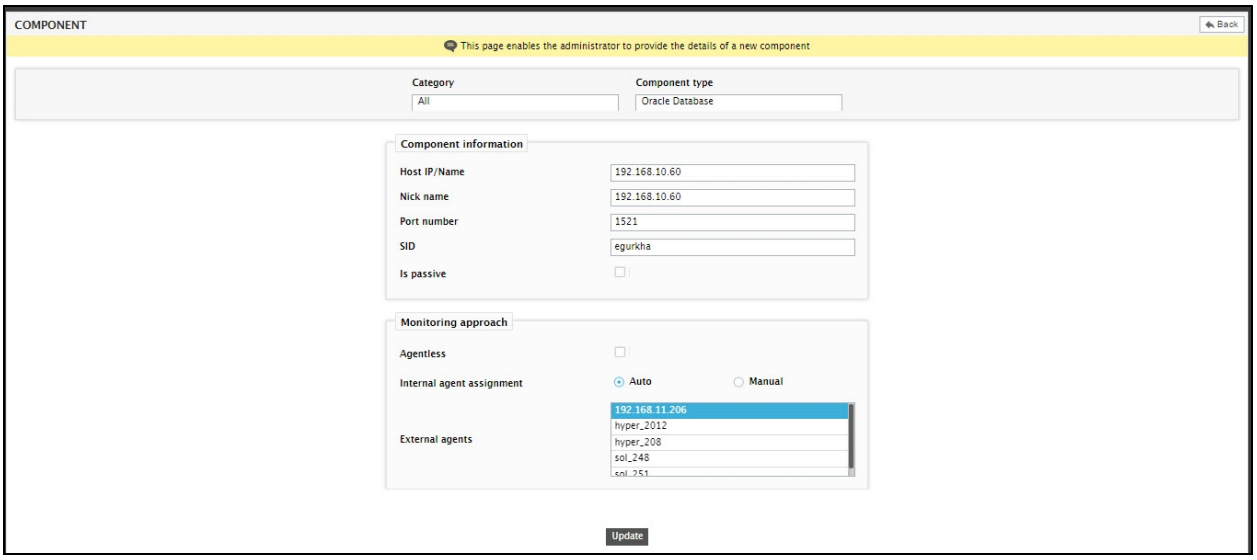

Figure 7.67: Modifying the details of a component

6. Using Figure 7.67, any of the displayed component details can be changed, except the **Host IP/Name**.

   Note:

   - While specifying a nick name for a component, ensure that you do not use the same name as that of the *internal* name of the component-type. For instance, you cannot add a component of type *Generic_server* with the nick name *Generic_server*. To know the internal component type names, follow the procedure discussed in Page 285.
   - If required, both the **Host IP/Name** and **Nick name** parameters of a component can be configured with the IP address of that component. However, if the **Host IP/Name** parameter has been configured with the target component's IPv6 address, then you cannot configure the **Nick name** parameter also with that IPv6 address; in this case, the **Nick name** should be some other logical name by which you want to identify the component.

7. Finally, click the **Update** button in Figure 7.67 to save the changes.

7.3.3 Changing the Host IP/Name of Component

To change the **Host IP/Name** of a component, do the following:

1. Select the **Add/Modify** option from the **Components** menu of the **Infrastructure** tile.
2. When Figure 7.68 appears, select a component **Category** and then the **Component type** to which the component to be modified belongs.

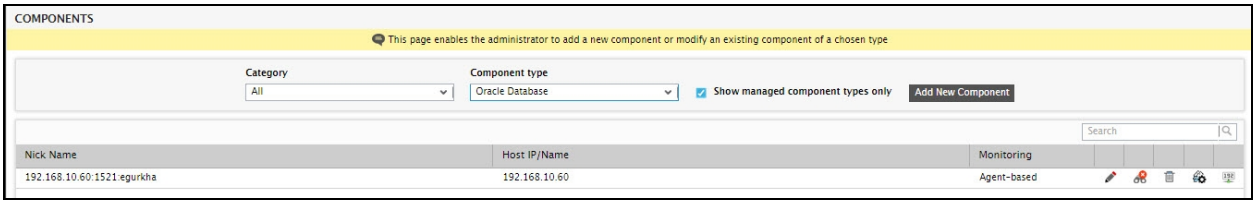

Figure 7.68: Selecting the component type to be modified

3. If components of the chosen type pre-exist in the eG Enterprise system, the same will be displayed as depicted by Figure 7.68.
4. To change the IP address or host name of a component, just click the icon corresponding to that component.
5. Figure 7.69 will then appear.

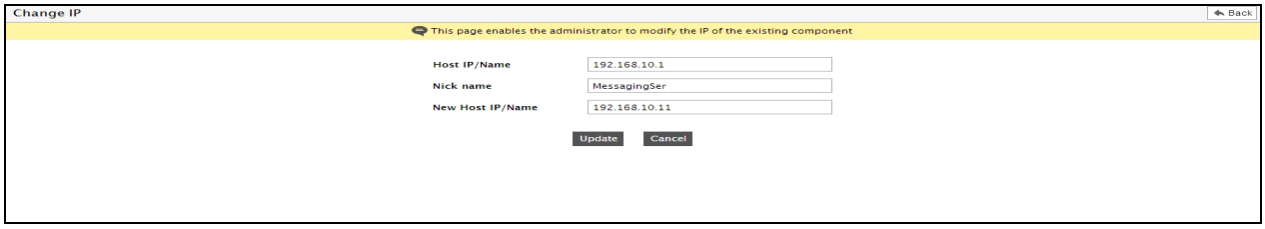

Figure 7.69: Changing the Host IP/Name of a component

6. In the **New Host IP/Name** text box of Figure 7.69, specify the new IP address/fully qualified host name of the component.

   Note:

   If you choose to specify the IP address of a component against **New Host IP/Name**, then remember that this IP address can be an IPv4 or an IPv6 address.

7. Finally, click the **Update** button in Figure 7.69 to save the changes.

7.3.4 Unmanaging a Component Added Manually

To unmanage a component that was added manually using the **Add/Modify** menu option, do the following:

1. Select the **Add/Modify** option from the **Components** menu of the **Infrastructure** tile.
2. When Figure 7.70 appears, select a component **Category** and then the **Component type** to which the component to be unmanaged belongs.

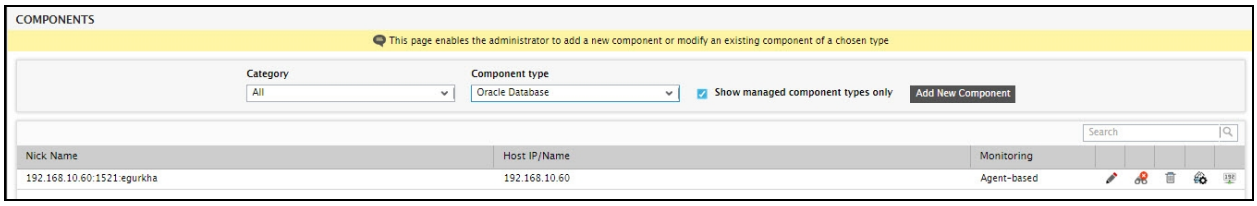

Figure 7.70: Selecting the component type to be modified

3. If components of the chosen type pre-exist in the eG Enterprise system, the same will be displayed as depicted by Figure 7.70.
4. To unmanage a component, click the icon corresponding to that component.
5. Figure 7.71 will then appear. Select the component to be unmanaged from the **Managed components** list of Figure 7.71 and click the **>** button to unmanage it. This will transfer the selection to the **Unmanaged components** list as depicted by Figure 7.72.

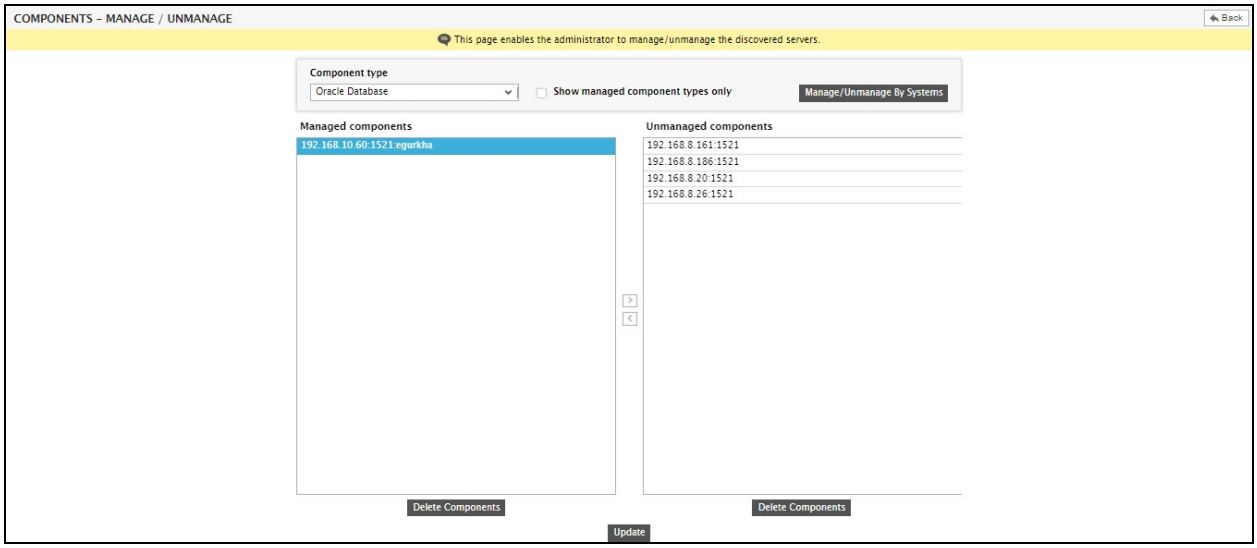

Figure 7.71: Selecting the component to be unmanaged

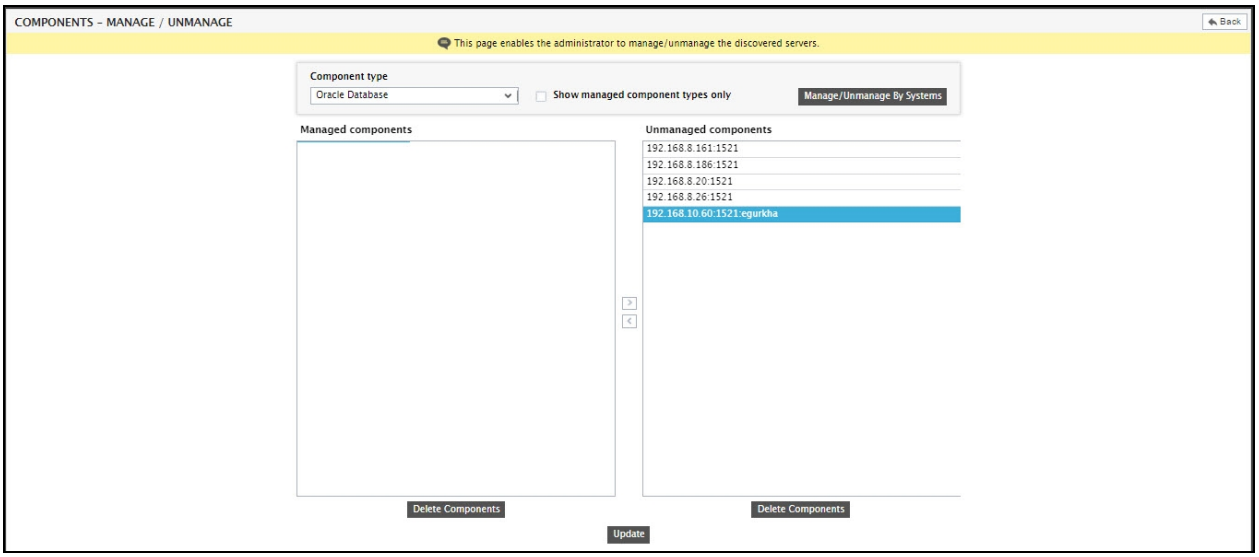

Figure 7.72: Unmanaging a component

6. Finally, click the **Update** button in Figure 7.72 to save the changes.

7.3.5 Deleting a Component Added Manually

To delete a component that has been manually added to the eG Enterprise system using the **Add/Modify** option, do the following:

1. Select the **Add/Modify** option from the **Components** menu of the **Infrastructure** tile.
2. When Figure 7.73 appears, select a component **Category** and then the **Component type** to which the component to be deleted belongs.

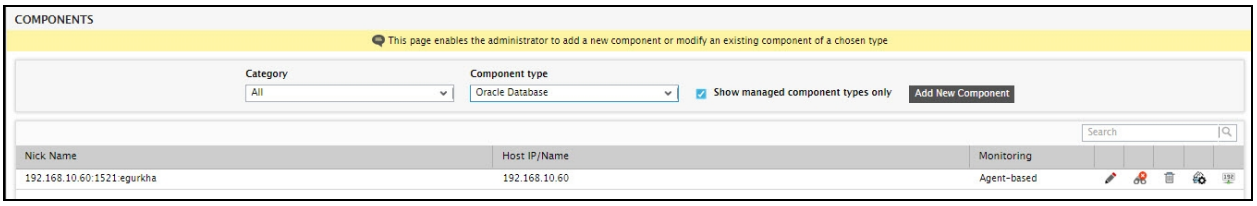

Figure 7.73: Selecting the component type to be modified

3. If components of the chosen type pre-exist in the eG Enterprise system, the same will be displayed as depicted by Figure 7.73.
4. To delete a component, click the icon corresponding to that component.
5. A message box will then appear requesting your confirmation to delete the component (see Figure 7.74).

Figure 7.74: A message box requesting your confirmation to delete the component

6. Click the **Yes** button to delete the component or the **No** button to cancel deletion.

Note:

- By default, deleted components cannot be automatically discovered by the eG manager's rediscovery process. You can however override this default behavior when configuring manager discovery. To achieve this, follow the steps below:
- Select the **Discovery** option from the **Components** menu of the **Infrastructure** tile.
- In the **DISCOVERY** tree of the page that then appears, the **General Settings sub-**node under the **Manager Discovery** node will be chosen by default. Accordingly, the right panel will display the **MANAGER DISCOVERY - GENERAL SETTINGS** page (see Figure 7.10).
- In the right panel, the **Re-discover deleted components** flag will be set to **No** by default. To make sure that the discovery process rediscovers deleted components, set this flag to **Yes**.
- Finally, click the **Update** button to save the changes.
- eG Enterprise disallows the deletion of components that are part of segments/services.

7.4 Asset Management

eG Enterprise now allows administrators to record asset information for every application, device, or server being managed. When an application, server, or device experiences performance degradation, through the Alarms window in the eG monitoring console, a help-desk person has single-click access to the asset information of each problematic application, server or device. Administrators can also configure the eG manager so that asset information can be included in email alerts sent out to users. By making useful asset details easily available to help desk staff and administrators, eG Enterprise helps minimize troubleshooting time and improves the efficiency of IT operations.

To record asset information, do the following:

1. Select the **Add/Modify** option from the **Components** menu of the **Infrastructure** tile.
2. Select a **Category** and then a **Component type** from the **COMPONENTS** page (see Figure 7.73) that then appears.
3. All managed components of the chosen type will then be listed as depicted by Figure 7.73.
4. To configure the asset details of a managed component, just click the icon corresponding to that component in Figure 7.73.

5. This will lead you to Figure 7.75

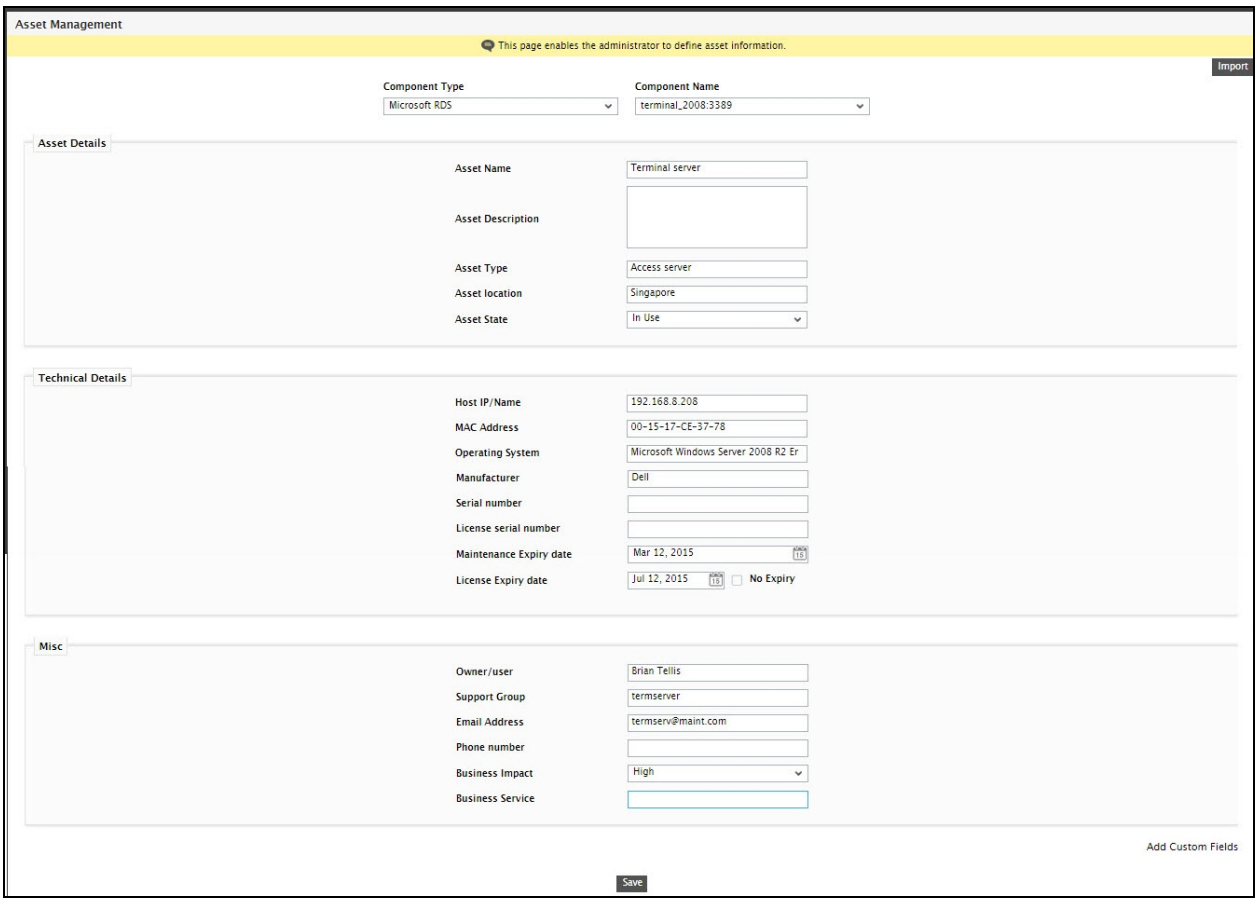

Figure 7.75: Recording asset information of a component

6. Using Figure 7.75, you can record details such as the name of the asset, description, type, location and state. Additional details on manufacturer, serial number of the asset, maintenance information and license information can also be recorded. Ownership details including the person to be contacted in the event of an issue can also be recorded.

7. If required, you can customize the asset details to capture more asset-related parameters that what is provided by default by the eG Enterprise system. For this purpose, custom fields can be added to the **Asset Management** page of Figure 7.75. To add custom fields, click the **Add Custom Fields** option at the bottom, right corner of Figure 7.75.

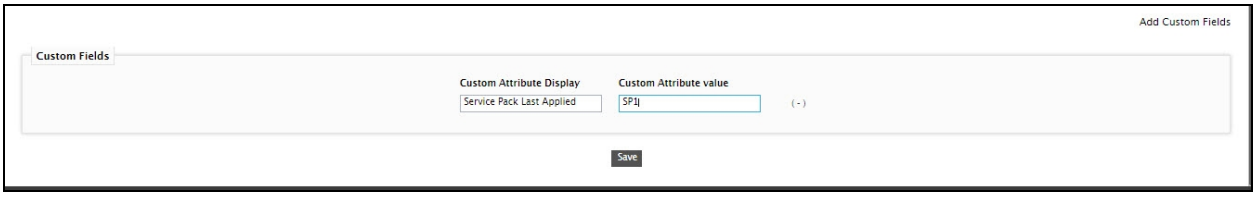

Figure 7.76: Adding custom fields

8. Type the name of the new attribute that you want to capture against the **Custom Attribute Display** text box of Figure 7.76. Enter the value of new attribute in the **Custom Attribute value** text box.

9. In the same way, you can add more asset parameters and configure their values.
10. At any given point in time, you can delete an attribute-value pair by clicking the **(-)** icon corresponding to that pair in Figure 7.76
11. Finally, click the **Save** button to register the changes.
12. Asset information can also be mass imported into eG Enterprise from CSV files. For this, click the **Import Asset Details** button at the right, top corner of Figure 7.75. Figure 7.77 will then appear using which you can indicate the full path to the CSV/XLS file containing asset details.

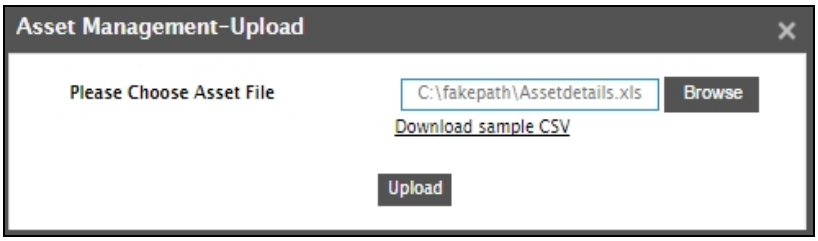

Figure 7.77: Importing an XLS file with asset details

13. The asset details should be provided in a specific format in the CSV/XLS file you upload. To know the format, download a sample CSV file by clicking the **Download sample CSV** link in Figure 7.69. An **AssetDetails.csv** file will then be downloaded.
14. Once the file is downloaded, open the file, save it using any other name you want to (if required) and as any file type (CSV/XLS) of your choice, and then open the file. In the file, you will find a series of columns, each of which represents a standard asset-related parameter – e.g., OS/DEVICE/APPLICATION, HOSTNAME, PORT, ASSETNAME, DESCRIPTION, etc. Replace the sample entries under each of the column heads with the details of the assets you want to manage. Each row should contain details of one asset.
15. For an asset, besides the default/standard asset information, you can also maintain custom properties. To enable you to specify these custom properties, columns labeled ATTR2 and ATTR3 are provided in the sample CSV file. Replace these column labels with the names of the properties that you want to additionally capture for an asset. **However, before that, make sure that you add custom fields to your ASSET MANAGEMENT page using steps 7 and 8 above. Make sure that the count of custom fields you add is equal to the count of custom attributes you specify in the sample file. Also, when labeling the columns for custom properties in the sample file, take care to assign the same names as the custom fields to these columns.**
16. Click the **Upload** button in Figure 7.77 to upload the file to the eG manager and import its contents into the asset details repository.

7.5 Configuring External Agents

7.5.1 Adding/Modifying/Deleting External Agents

The external availability and responsiveness of every component managed by the eG Enterprise system is measured using an external agent. A single component can be associated with one or more external agents. By default, the eG agent operating on the eG manager's host serves as an external agent for all monitored

applications/devices in an environment. If needed, an administrator can configure additional external agents and assign them to components.

To add more external agents, first select the **External Agents** option from the **Agents** tile.

1. Figure 7.78 then appears. The list of external agents (including the default external agent that is executing on the eG manager host) that have been configured in the environment is available here.

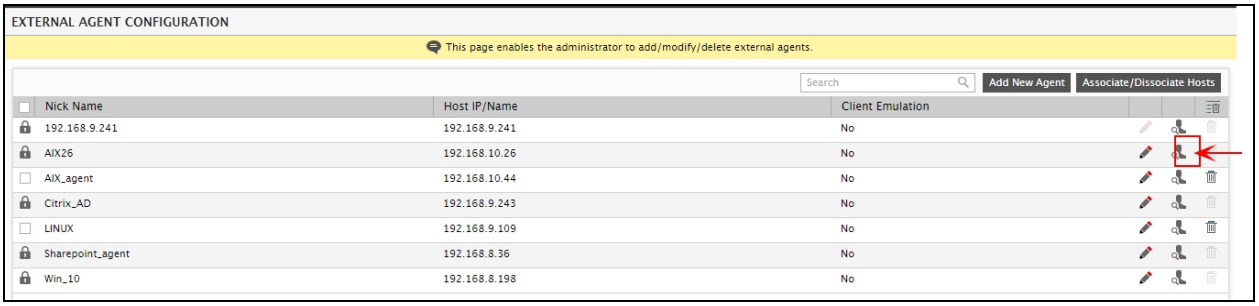

Figure 7.78: Configuring external agents

New agents can be configured via the **Add New Agent** option in Figure 7.78. The IP address/fully qualified host name of the system where the agent is to be configured has to be specified against **Host IP/Name** as shown by Figure 7.79.

Note:

If you choose to configure an external agent with its IP address (instead of its host name), then remember that this IP address can be an IPv4 or an IPv6 address.

2. Next, provide a **Nick name** for the external agent.

Note:

If required, you can configure both the **Host IP/Name** and **Nick name** of an external agent with its IP address. However, if the **Host IP/Name** parameter has been configured with an IPv6 address, then you cannot configure the **Nick name** parameter also with the IPv6 address; in this case, the **Nick name** should some other logical name by which you want to identify the external agent.

eG Enterprise supports client emulation tests. These tests which emulate user accesses to different services/applications, use third party client emulation tools like CitraTest and Itexis AppsMon to playback pre-recorded user actions to a service/application. Based on the output of these client emulation tools, the eG tests report the availability and performance of the services being monitored. The client emulation tests are executed by the eG external agents. When configuring an external agent, the administrator must decide whether the agent should be configured to run client emulation tests or not. The selection **Client emulation** in Figure 7.79 will only appear if the eG license has Client Emulation capability enabled. Note that a dedicated external agent is required for executing the client emulation tests - i.e., an external agent that is configured to execute the client emulation tests cannot execute any other test.

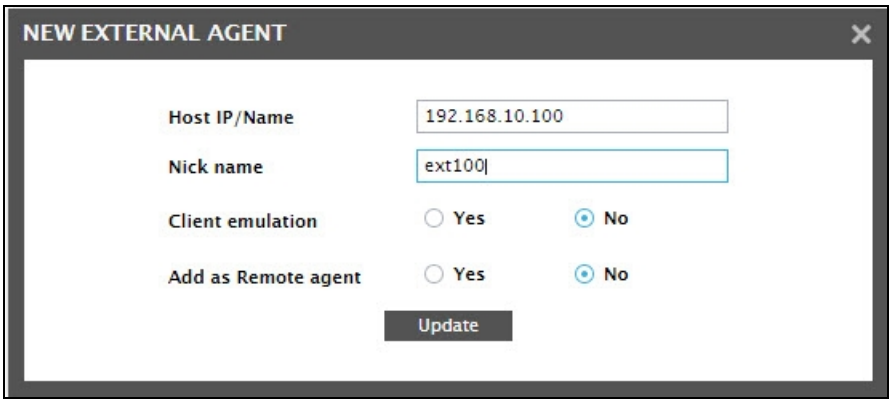

Figure 7.79: Adding new external agents

3. If the external agent being added should also serve as remote agent, then set the **Add as Remote agent** flag to **Yes**. This will automatically add a remote agent with the same **Host IP/Name** and **Nick name** as the external agent being configured. This way, you can configure an external and a remote agent at one shot, without having to follow separate procedures to achieve the same.

   Note:

   An external agent for which **Client emulation** has been enabled cannot operate as a remote agent. In other words, as soon as the **Client emulation** flag is set to **Yes** for an external agent, the **Add as Remote agent** flag will automatically disappear from the **NEW EXTERNAL AGENT** window.

4. Finally, click the **Update** button to add the new external agent.

To modify an external agent's configuration, do the following:

1. Select the **External Agents** option from the **Agents** tile.
2. 7.5.2 that then appears lists the external agents that have been configured for the environment.
3. To modify the details of any of these external agents, click the icon corresponding to it. Figure 7.70 will then appear.

   Note:

   You cannot modify the details of the default external agent on the eG manager host.

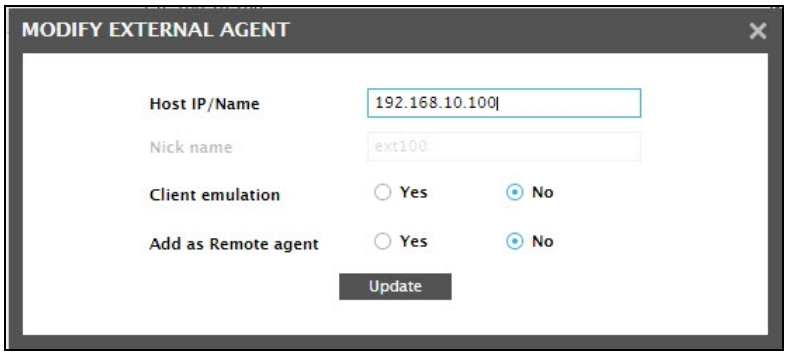

Figure 7.80: Modifying an external agent's configuration

4. Using Figure 7.80, you can change the **Host IP/Name** of an external agent, but cannot modify its **Nick name**. You can also change the status of the **Client emulation** and **Add as Remote agent** flags.

   Note:

   - An external agent for which **Client emulation** has been enabled cannot operate as a remote agent. In other words, as soon as the **Client emulation** flag is set to **Yes** for an external agent, the **Add as Remote agent** flag will disappear from the **MODIFY EXTERNAL AGENT** window.
   - Once you set the **Add as Remote agent** flag to **Yes** for an external agent, you cannot turn this flag off (i.e., set it to **No**) for that agent in the **Mo**

5. Finally, click the **Update** button to register the changes.

To delete an external agent, do the following:

1. Select the **External Agents** option from the **Agents** tile.
2. Figure 7.80 that then appears lists the external agents that have been configured for the environment.

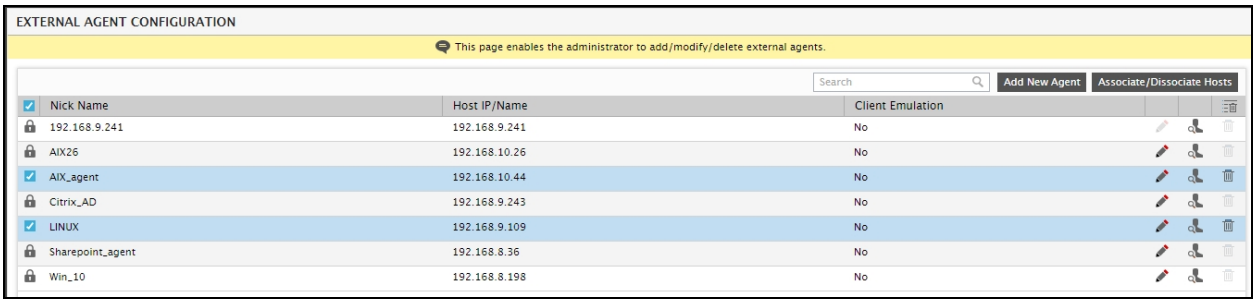

Figure 7.81: Deleting external agents

3. To delete a particular external agent, just click the (i.e., the 'trash can' icon) corresponding to that agent in Figure 7.81.
4. This will invoke the message box depicted by Figure 7.82. Click the **Yes** button in the Figure 7.82 to confirm the deletion of the chosen external agent.

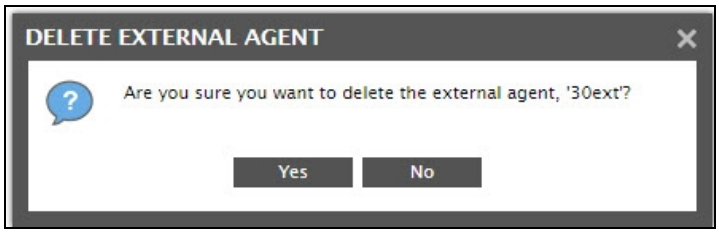

Figure 7.82: A message box requesting your confirmation to proceed with the deletion of the external agent

5. To delete multiple external agents at one shot, select the check box corresponding to such agents as depicted by Figure 7.81. Then, click the icon adjacent to the column head **Client Emulation** in Figure 7.81 to delete the selected external agents.

   Note:

   - You cannot delete the default external agent on the eG manager host.
   - You cannot delete external agents that have already been assigned to one/more hosts or components.

7.5.2 Assigning External Agents to Hosts

You can assign an external agent to multiple hosts simultaneously, so that all applications managed on that host are automatically associated with that external agent. For this purpose, do the following:

1. Select the **External Agents** option from the **Agents** tile.
2. To assign hosts to any external agent, click the **Associate/Disassociate Hosts** button in Figure 7.78 that then appears. This will open Figure 7.83.

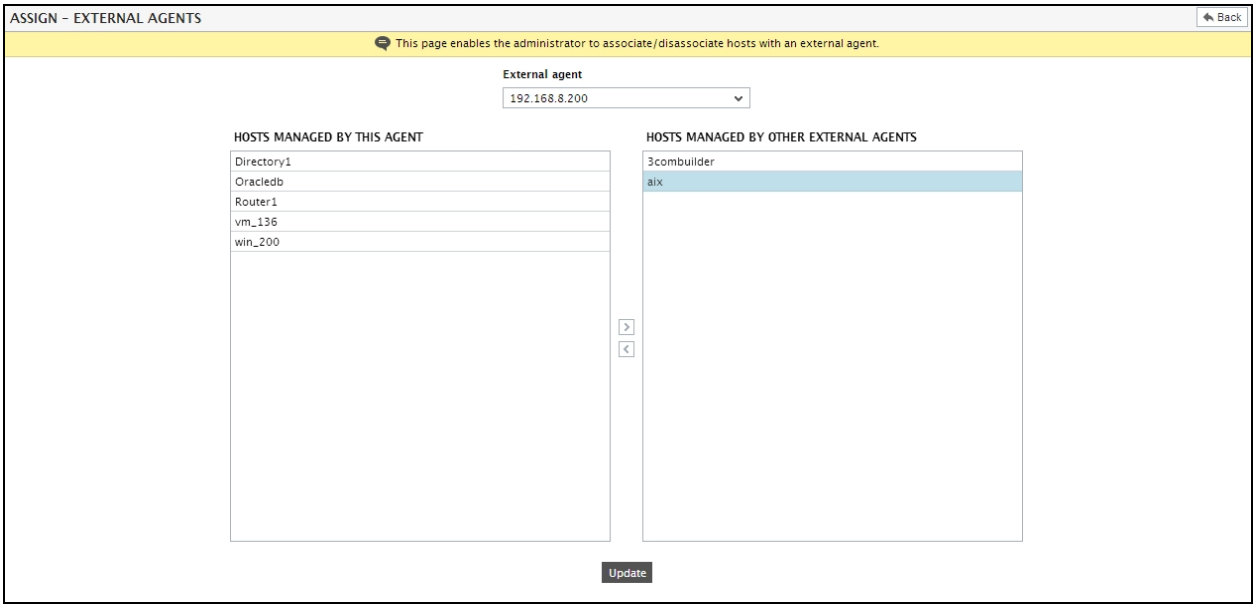

Figure 7.83: Associating hosts in the environment with an external agent

3. From the **External agent** list in Figure 7.83, select the external agent for which you want to assign hosts.

4. The hosts that are currently managed by the chosen agent will appear in the **HOSTS MANAGED BY THIS AGENT** list box (see Figure 7.83). The other hosts, which may be associated with other external agents, appear in the **HOSTS MANAGED BY OTHER EXTERNAL AGENTS** list box.
5. To assign the selected external agent to multiple hosts, select the hosts from the **HOSTS MANAGED BY OTHER EXTERNAL AGENTS** list box and click the **<** button.
6. To disassociate the external agent from hosts with which they are currently associated, select the hosts from the **HOSTS MANAGED BY THIS AGENT** list box and click the **>** button.
7. The **ASSOCIATE>>** button enables the administrator to add a host(s) to the existing list of hosts that are being targeted by the external agent under consideration. Similarly, the **<< Disassociate** button allows the user to move the selected host to the list of hosts that are not monitored by this agent.
8. Finally, click the **Update** button.
9. On the other hand, if you want to assign hosts to a particular external agent, click the button that corresponds to that agent, as indicated by Figure 7.84.

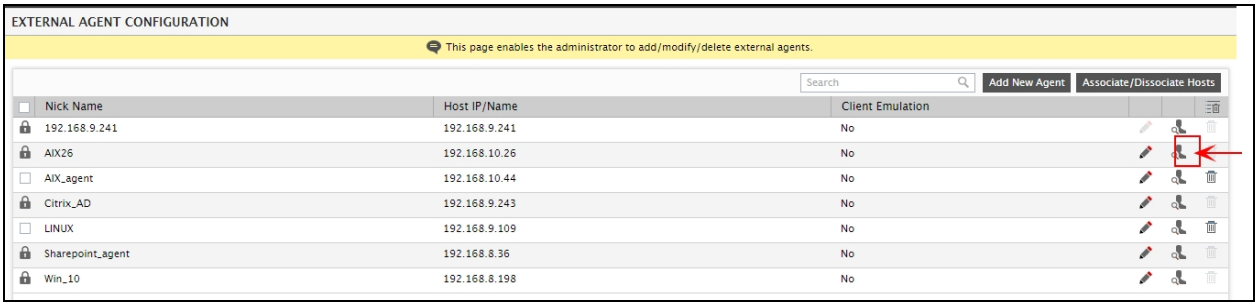

Figure 7.84: Clicking on the Associate Hosts button of a particular external agent

7.6 Agentless Monitoring

The agent-based approach to monitoring requires one internal agent per system that is to be monitored. As operating systems and applications have evolved, they have incorporated newer instrumentation mechanisms that allow monitoring of these environments from remote locations - i.e., external to the system or application that is to be monitored. The main advantage of this remote monitoring approach is that it does not require an agent to be installed on every system that is to be monitored - hence, the name "*agentless monitoring*" for this approach. The table below summarizes the tradeoffs between agentless and agent-based monitoring.

Agentless Monitoring	Agent-based Monitoring
Easier to implement, maintain, and upgrade as there is no need to deploy agents on all the servers	More difficult to implement as it requires agents on critical servers
The nature of statistics that can be reported depends on the instrumentation levels and on the access rights provided for monitoring from a remote location	Typically, used to collect a wealth of availability, performance, usage metrics; can provide tight integration with the operating system and applications monitored

Agentless Monitoring	Agent-based Monitoring
Mostly used for problem detection and alerting; provides limited information for diagnosis	Can provide critical information for problem diagnosis and remote control/problem correction
Significant network overheads; all the monitoring is done over the network	This configuration optimizes network usage; most of the data collection and analysis is done on the servers and only processed metrics are transmitted; hence, this approach is very efficient in its network usage
Higher security risk - monitoring from a remote location must be allowed; some ports (e.g., Netbios/SSH) need to be opened for remote access	No security risk - the agent does not listen on a port, no firewall rules need to be changed, no new ports need to be opened for the monitoring

Unlike many existing solutions, eG Enterprise does not require that IT administrators choose between one of these contrasting approaches at the time of installing/deploying eG Enterprise. Instead, when adding any new application/system for monitoring, administrators have a choice of whether to use agent-based or agentless monitoring for the application or system under consideration. For instance, an administrator can choose to monitor the most critical servers in an agent-based manner, and to monitor the less critical servers in the staging/development environment in an agentless manner.

7.6.1 How does Agentless Monitoring Work in eG Enterprise?

Agentless monitoring is implemented in eG Enterprise using **Remote Agents.** A remote agent is capable of monitoring a number of systems and applications remotely, i.e., without requiring an agent to be locally installed on the system that is to be monitored.

Figure 7.85 depicts how a single remote agent interacts with and extracts measures from different servers across platforms.

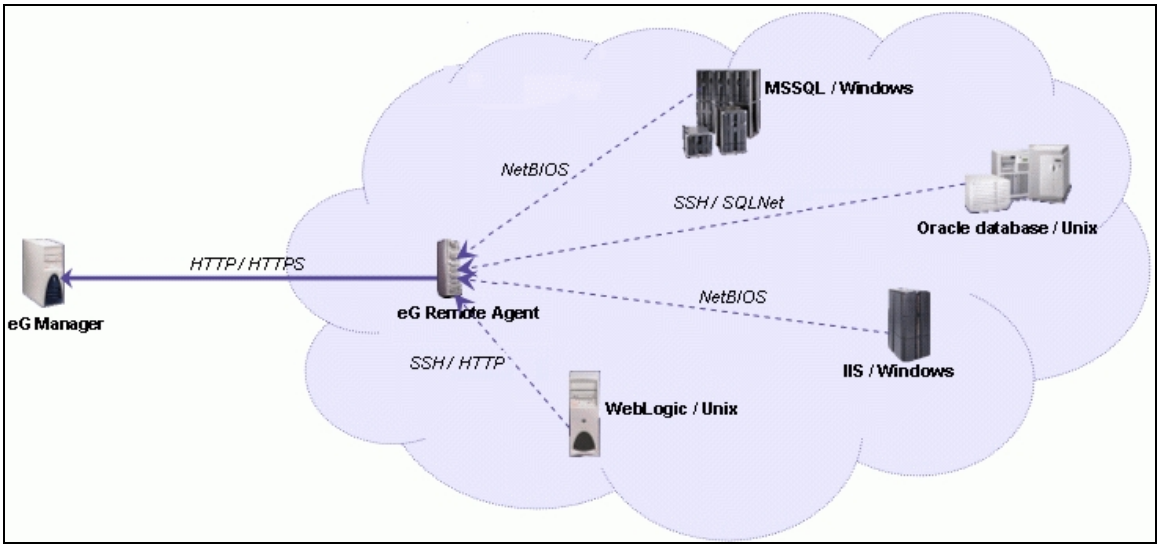

Figure 7.85: How do remote agents work?

For monitoring Microsoft Windows systems and applications, a remote agent uses Netbios/perfmon to communicate with the operating system/applications. For monitoring Unix systems, secure shell (SSH) is used. In addition, for specific applications, the remote agent uses application-specific protocols to communicate with the application (e.g., SQLNet for Oracle databases, HTTP for WebLogic and WebSphere application servers, JDBC for Sybase, etc.).

As is the case with internal and external agents, the configuration on the eG manager dictates which agent functions as an internal agent, or a remote agent, or an external agent. Since eG Enterprise uses a single agent model, the same agent installation can function as any of the agent types. In fact, the same agent could function both as an external agent and as a remote agent. With the external agent functionality, the agent performs black-box testing (typically, exercising the application being monitored or generating synthetic transactions). When functioning as a remote agent, the agent polls the target server for critical internal metrics ranging from system CPU, memory, disk space to application process execution status.

Currently, remote agents are supported on the following servers:

- Windows 2008,
- Windows 2012
- Windows 7
- Windows 8
- Windows 10
- HPUX
- AIX
- Solaris

Note:

Remote agents that monitor VMware vSphere, Citrix XenServer, and Microsoft Hyper-V virtualization platforms can operate on Windows, Linux, or Solaris hosts.

The number of remote agents allowed for a target environment and the number of systems that can be monitored in an agentless manner are controlled by the eG license.

Note:

The following capabilities of eG Enterprise are not available for servers or applications that are managed in an agentless manner:

- Detailed diagnosis
- Automatic corrective script execution
- Remote control actions
- The eG web adapter for in-depth web server monitoring on Unix and Microsoft environments - i.e., web transaction monitoring for web servers is not available with agentless monitoring.

7.6.2 Pre-requisites for Monitoring Components in an Agentless Manner

When deploying a remote agent, it is essential to ensure that the remote agent can communicate using multiple agentless mechanisms (see Figure 7.85) with the target environment (appropriate firewall rules must be configured for this purpose).

The sub-sections below describe the pre-requisites that need to be fulfilled to enable agentless monitoring of different types of applications.

7.6.2.1 General Pre-requisites

1. At least one remote agent should be available for performing agentless monitoring. By default, the eG agent on the eG manager host serves as both a remote agent and external agent for the environment. If required, you can configure more remote agents. For this purpose, you can follow the procedure detailed in Section 7.8.3 of this document.
2. If you intend to use SNMP for monitoring a component in an agentless manner, then make sure that the SNMP port 161 is opened on the target component.

7.6.2.2 Pre-requisites for Monitoring a Windows System/Application in an Agentless Manner

For monitoring a Microsoft Windows environment in an agentless manner, the following requirements should be fulfilled:

1. The **UDP** ports 137 and 138 and the **TCP** ports 139 and 445 should be opened on the target Windows sytem to enable the remote agent to communicate with that system.
2. The remote agent must be installed and running with the permissions of the domain administrator. To ensure this, do the following before starting a remote agent:
   - Open the **Component Services** window (using the menu sequence, Settings -> Control Panel -> Administrative Tools -> Component Services) of the remote agent host. Right-click on the **eGurkhaAgent** service listed therein, and select **Properties** from the shortcut menu.

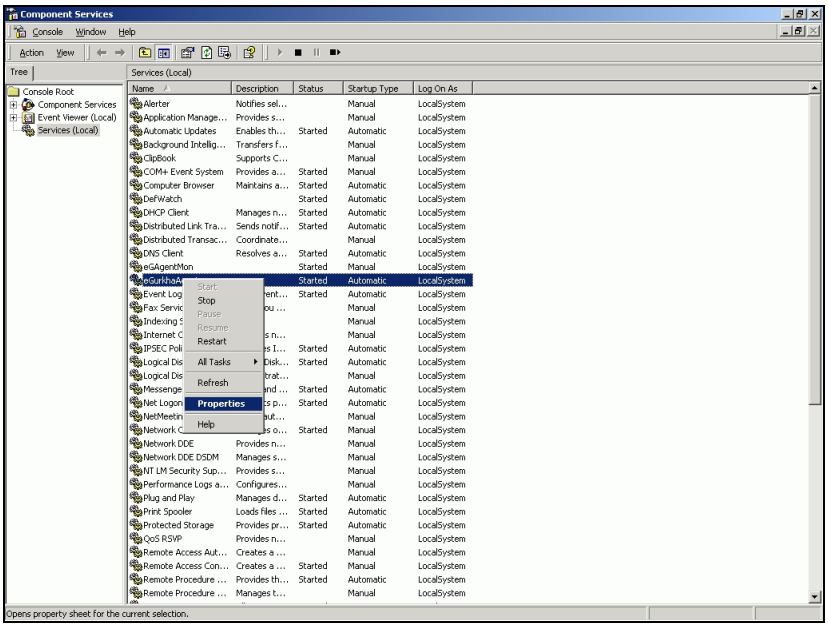

Figure 7.86: Selecting properties from the eGurkhaAgent service's shortcut menu

- Select the **LogOn** tab from the **Properties** dialog box that appears. Then, choose the **This account** option from the **Log on as** section of, and provide the *Domainname**Username* of the domain administrator in the adjacent text box. Provide the **Password** of the domain administrator, and confirm the password by retyping it in the **Confirm Password** text box.

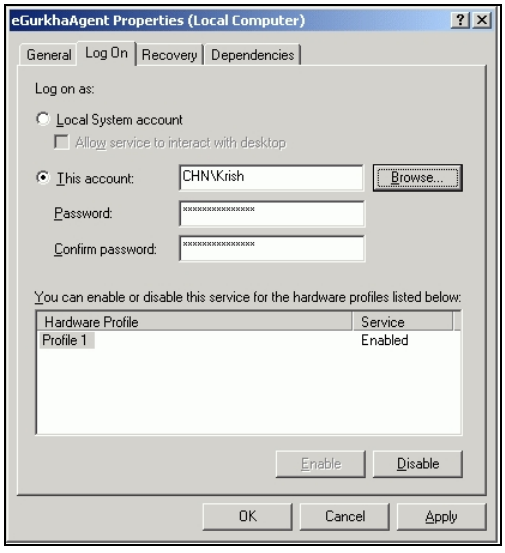

Figure 7.87: Changing the Log On account

- Finally, click the **Apply** button in and then the **OK** button to register the changes.
- Hence, typically, at least one remote agent is necessary for each independent domain being monitored.

Note:

The above-mentioned steps are applicable only if the remote agent and the target Windows host are in the same domain. However, if the remote agent is within a domain, but the target host operates out of the domain or within a workgroup, then perform the following broad steps to ensure that the remote agent monitors the host:

- Create a new user on the target host and assign *local administrator* privileges to that user.
- Create a new user on the remote agent's host with the same credentials as the user created on the monitored host; assign *local administrator* privileges to this user also.
- Ensure that the remote agent runs with the permissions of this new user.

The detailed steps in this regard are available in Chapter 7 of this document - i.e., the chapter on *Troubleshooting*.

3. Ensure that the target Windows system consists of the default share named **ADMIN$**. If this share does not exist on a target, then the remote agent will not be able to connect to that system or collect metrics from it. To check whether **ADMIN$** pre-exists, do the following:
   - Login to the target system and go to the command prompt.
   - Type the command **net share**; this command will list all the default and user-configured shares on the system
   - If **ADMIN$** is not listed, it is a definite indicator that the system does not consist of the **ADMIN$** share.

   The way forward is to manually configure the **ADMIN$** share on the target system. To do so, issue the command **net share ADMIN$** from the command prompt of the system. This will create the **ADMIN$** share. For more information on **net share**, check out the URL:

   http://www.microsoft.com/resources/documentation/windows/xp/all/proddocs/en-us/net_share.mspx?mfr=true

4. Also, the remote agent connects to the Windows host to be monitored using **NetBIOS** - a protocol that allows applications on different computers to communicate within a local area network. If **NetBIOS** is not enabled on a Windows host, then eG Enterprise cannot monitor that host in an agentless manner.

   To enable **NetBIOS** on a Windows 2000 host, follow the steps given below:
   - Follow the menu sequence: Start -> Programs -> Accessories -> Communications -> Network and Dial-up Connections
   - Figure 7.88 will then appear. Right-click on the **Local Area Connection** and then select **Properties** from the shortcut menu.

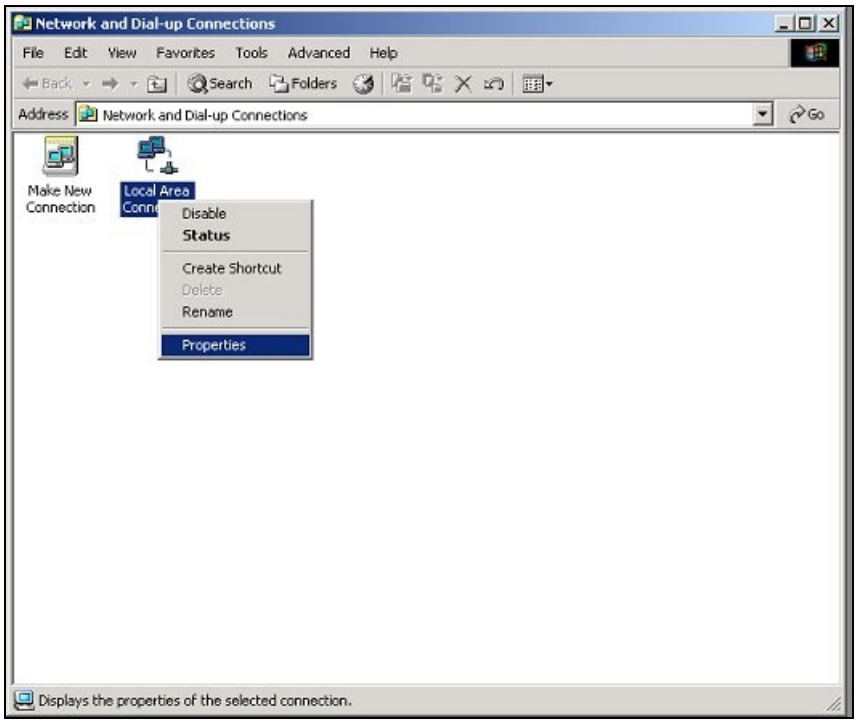

Figure 7.88: Selecting the Properties option of the Local Area Connection option

- Figure 7.89 that appears next displays the properties of the **Local Area Connection**. Double-click on the **Internet Protocol (TCP/IP)** option indicated by Figure 7.89 to view its properties.

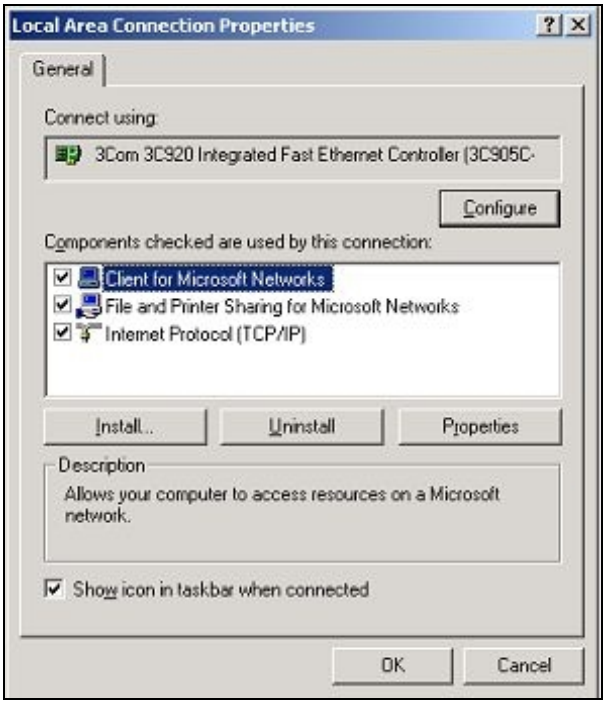

Figure 7.89: The General tab of the Properties dialog box

- Click on the **Advanced** button within the **General** tab of the **Properties** dialog that appears next.

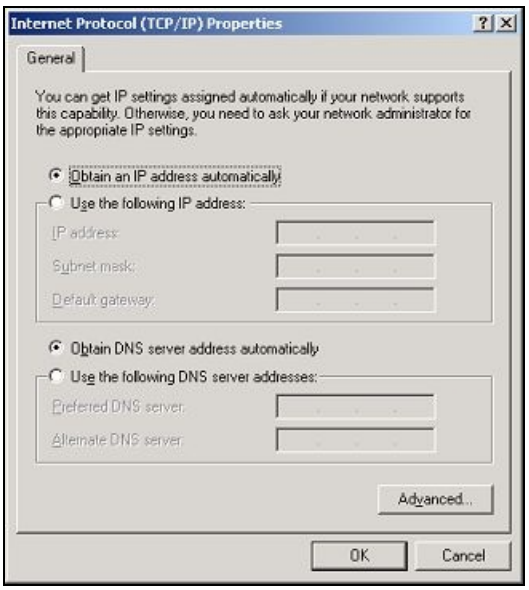

Figure 7.90: Properties of the Internet Protocol (TCP/IP)

- Click on the **WINS** tab of the **Advanced TCP/IP Settings** dialog box, and select the **Enable NetBIOS over TCP/IP** option within. Finally, click the **OK** button therein to register the changes.

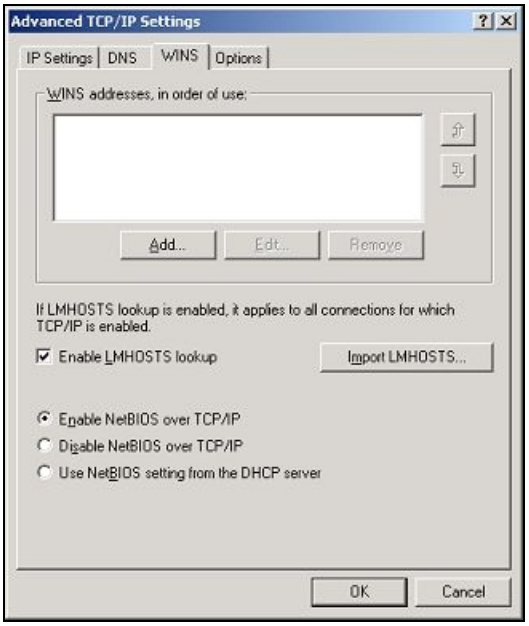

Figure 7.91: Enabling NetBIOS

To enable **NetBIOS** on a Windows 2003 host, follow the steps given below:

- Follow the menu sequence: Start -> Programs -> Accessories -> Communications -> Network Connections
- Right-click on the **Local Area Connection** option and select **Properties**.

- Then, double-click on the **Internet Protocol (TCP/IP)** under the **"This connection uses the following items:"** list box.
- Click on the **Advanced** button within the **General** tab of the **Properties** dialog box that appears next.
- Click on the **WINS** tab of the **Advanced TCP/IP Settings** dialog box, and select the **Enable NetBIOS over TCP/IP** option within. Finally, click the **OK** button therein to register the changes.

To enable **NetBIOS** on a Windows XP host, follow the steps given below:

- Follow the menu sequence: Start -> All Programs -> Accessories -> Communications -> Network Connections
- Right-click on the **Local Area Connection** option and select **Properties**.
- Then, double-click on the **Internet Protocol (TCP/IP)** under the **"This connection uses the following items:"** list box.
- Click on the **Advanced** button within the **General** tab of the **Properties** dialog box that appears next.
- Click on the **WINS** tab of the **Advanced TCP/IP Settings** dialog box, and select the **Enable NetBIOS over TCP/IP** option within. Finally, click the **OK** button therein to register the changes.

To enable **NetBIOS** on a Windows 7 system, follow the steps given below:

- Click **Start** and then click **Network**.
- Click on the **Network and Sharing Center**.
- Click **Manage Network Connections**.
- Right-click on the **Local Area Connection** and select **Properties**.
- Select **Internet Protocol version 4 (TCP/IPv4)**.
- Click the **Advanced** button under the **General** tab page.
- Click the **WINS** tab page.
- Click **Enable NetBIOS over TCP/IP**.
- Click **OK** and **Exit** the settings.

5. Additionally, to remotely monitor a Windows 7 host in particular, you need to start the **Remote registry service**.
6. By default, remote Windows systems/applications are monitored using **Perfmon**. To ensure that the remote agent is able to collect metrics from the target Windows systems using **Perfmon**, the following pre-erquisites should be additionally fulfilled:
   - PerfMon should have at least READ access to the **Perflib\LanguageID** subkey on the remote computer (which allows external access to PerfMon). The Perflib\LanguageID subkey is located in the following Registry path: **HKEY_ LOCAL_ MACHINE\Software\Microsoft\WindowsNT\CurrentVersion\Perflib\LanguageID** . The LanguageID is a numeric code for the spoken language of the installed operating system. A computer with a LanguageID of 009 (the English LanguageID) has the following **Perflib\Language** subkey: **HKEY_ LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\Perflib\009**

- The Disk Performance Statistics Driver (diskperf) should exist on the target computer; allow READ access explicitly to the user account for the following registry key and all subkeys: **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Diskperf**
- The monitored computer should be able to connect to IPC$. The following registry entry enables connecting to IPC$:
    - Hive: **HKEY_LOCAL_MACHINE\SYSTEM**
    - Key: **CurrentControlSet\Services\LanmanServer\Parameters**
    - Name: **AutoShareWks**
    - Type: **REG_DWORD**
    - Value: **1**
- At least READ access should be granted to the following registry subkey (allowing it to remotely connect to the Windows registry): **HKLM\SYSTEM\CurrentControlSet\Control\SecurePipeServers\winreg**. This permission determines who can remotely connect to a registry. If this subkey does not exist, all users can remotely connect to the registry. To remotely connect to a registry, a user must have at least READ access to the winreg subkey on the target computer.
- At least READ access should be granted to the following registry keys on the remote computer:
    - HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\SecurePipeServers\winreg
    - HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\WindowsNT\CurrentVersion\Perflib
- To monitor Windows 2000 and Windows XP, the user name must have access granted by the following group policies:
    - Profile single process
    - Profile system performance

    Both group policies are security settings that you can set from the **Local Policies => User Rights** option in the **Administrative Tools** of the **Control Panel**.
- To monitor Windows XP, if the system root is on an NTFS partition, the user name must have at least READ access to the following two files:
    - %SystemRoot%\System32\Perfc009.dat
    - %SystemRoot%\System32\Perfh009.dat

7.6.2.3 Pre-requisites for Monitoring a Unix System/Application in an Agentless Manner

For monitoring Unix systems/applications in an agentless manner, ensure that the pre-requisites outlined below are fulfilled:

1. By default, the remote agent communicates with a target Unix system via **Rexec**. In this case, make sure that port 512 is opened on the target to enable this communication. If **SSH** is the mode of communication between the remote agent and target system, then make sure that port 22 is opened on the target.

2. If you want the remote agent to connect to a target host via **Rexec**, then make sure that the **rexecd** daemon on the target host is started. This daemon provides the server function for the **rexec** command. To check the status of this daemon, open the **/etc/inetd.conf** file on the target host and look for the **rexecd** line in it. If this line is uncommented, it indicates that the daemon has been started. If the line is commented, then it indicates that the daemon has not been started. In this case, uncomment the line, save the **inetd.conf** file, and then run the **refresh -s inetd** command or the **kill -1 InetdPID** command to inform the **inetd** daemon of the changes to its configuration.
3. If you want the remote agent to communicate with a target host via SSH, then you should also pick an encryption type/mode for the SSH connection. The options here are: **Password authentication** and **Key-based authentication**.

   If you want the SSH connection to use the **Password Authentication** mode, then, first make sure that **Password Authentication** is enabled on the target host. For this, follow the steps below:

   - Login to the Unix host to be monitored.
   - Edit the **sshd_config** file in the **/etc/ssh** directory.
   - Check whether the **Password Authentication** flag in the **sshd_config** file is set to **no**. If so, set it to **yes**.
   - Then, save the file and restart/signal the SSH daemon (eg., using **kill -1 <SSHD PID>**)

   If you want the SSH connection to use the **Key-based Authentication** mode instead, then first make sure that **Key- based authentication** is enabled on the target host. For this purpose, you will require a public/private key pair. A public/private key pair is available in the **/opt/egurkha/agent/sshkeys/** directory of the eG agent. While the private key is available in the file named **id_rsa**, the public key is contained within the file **authorized_keys**. The pass phrase associated with the default private key is **eginnovations**. You now have the option to proceed with the default keys or generate a different key pair.

   If you decide to go with the keys bundled with the eG agent, do the following:

   - To enable key-based authentication, the private key should remain in the **/opt/egurkha/agent/sshkeys/** directory of the eG agent, and the public key should be copied to the remote Unix host to be monitored. To achieve this, first login to the Unix host to be monitored as the eG user.
   - Create a directory named **.ssh** in the **<USER_HOME_DIR>** on the host using the command: **mkdir ~/.ssh**.
   - Next, copy the **authorized_keys** file from the **/opt/egurkha/agent/sshkeys/** directory on the eG agent host to the **<USER_HOME_DIR>/.ssh** directory on the Unix host. Make sure that the permission of the **.ssh** directory and the **authorized_keys** file is **700**.

   On the other hand, if you want to generate a new key pair, then do the following:

   - Login to the Unix host to be monitored as an eG user.
   - From the **<USER_HOME_DIR>,** execute the command: **ssh-keygen -t rsa**. Upon executing the command, you will be requested to specify the full path to the file to which the key is to be saved. By default, a directory named **.ssh** will be created in the **<USER_HOME_DIR>**, to which the key pair will be saved. To go with the default location, simply press **Enter**.

     ```
     Generating public/private rsa key pair.
     Enter file in which to save the key (/home/egurkha/.ssh/id_rsa):
     ```

   - Next, you will be prompted to provide a pass phrase. Provide any pass phrase of your choice.

```
Enter passphrase (empty for no passphrase): eginnovations
Enter same passphrase again: eginnovations
```

- If the key pair is created successfully, then the following messages will appear:

```
Your identification has been saved in /home/egurkha/.ssh/id_rsa.
Your public key has been saved in /home/egurkha/.ssh/id_rsa.pub.
The key fingerprint is:
09:f4:02:3f:7d:00:4a:b4:6d:b9:2f:c1:cb:cf:0e:e1 dclements@sde4.freshwater.com
```

- The messages indicate that the private key has been saved to a file named **id_rsa** in the **<USER_HOME_DIR>/.ssh**, and the public key has been saved to a file named **id_rsa.pub** in the same directory. Now, rename the **id_rsa.pub** file as **authorized_keys**.
- Then, to enable key-based authentication, login to the eG agent host.
- Copy the **id_rsa** file from the **<USER_HOME_DIR>/.ssh** directory of the target Unix host to the **<EG_INSTALL_DIR>\agent\sshkeys** directory on the eG agent host.
- Finally, log into target Unix host once again (as the eG user) and lock the file permissions down to prevent other users from being able to read the key pair data, using the following commands:

chmod go-w ~/
chmod 700 ~/.ssh
chmod go-rwx ~/.ssh/*

7.6.3 Configuring Additional Remote Agents

By default, the eG agent on the eG manager host will function as both the external and remote agent for the monitored environment. If required, you can add more remote agents. The procedure for this is discussed below:

1. Select the **Remote Agents** option from the **Agents** tile.
2. The page that appears next will list the remote agents that pre-exist (if any). By default, if no additional remote agents are yet to be configured in the environment, then Figure 7.92 will display the IP address of the default remote agent on the eG manager host alone. However, **note that this default remote agent cannot be modified or deleted**.

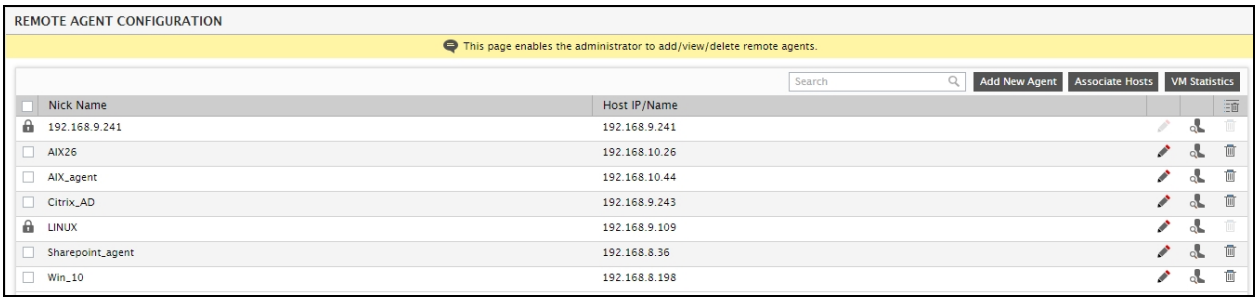

Figure 7.92: The existing remote agents

3. Click on the **Add New Agent** button in Figure 7.92 to add a new remote agent. Doing so will invoke Figure

7.93, wherein the **Host IP/Name** and the **Nick name** of the remote agent need to be specified. If you want to add an external agent with the same IP/host name and nick name as the remote agent that is being configured, set the **Add as External agent** flag in Figure 7.93 to **Yes**.

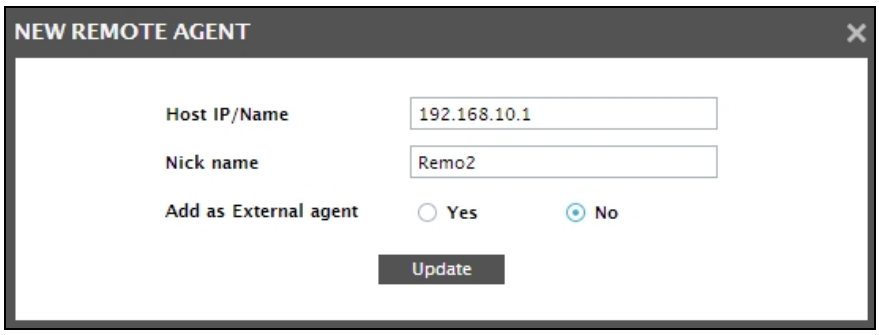

Figure 7.93: Adding a new remote agent

Note:

- You can configure a remote agent with its IP address or its 'fully qualified host name'.
- If you choose to configure a remote agent with its IP address (instead of its host name), then remember that this IP address can be an IPv4 or an IPv6 address.
- If required, you can configure both the **Host IP/Name** and **Nick name** parameters with the IP address of the remote agent. However, if the **Host IP/Name** has been configured with an IPv6 address, then you cannot configure the **Nick name** with that IPv6 address; in this case, the **Nick name** should be some other logical name using which you want to identify the remote agent.

4. Finally, click the **Update** button in Figure 7.93 to register the changes.
5. Figure 7.94 then appears displaying the newly created remote agent. To modify the configuration of a remote agent, click the **Modify** icon (i.e., the 'pencil' icon) corresponding to that agent in Figure 7.94.

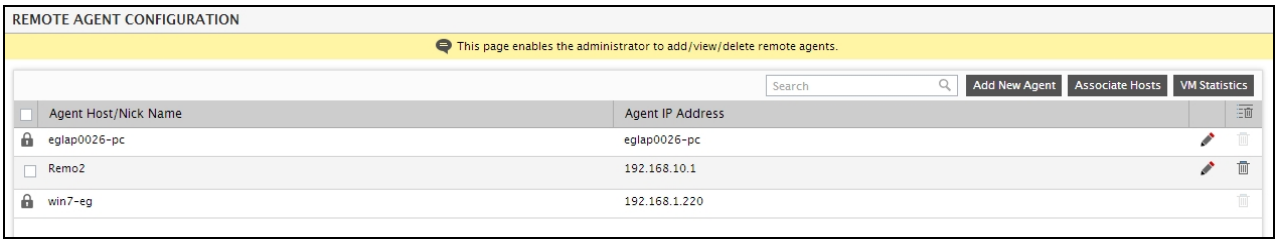

Figure 7.94: The newly configured remote agent appearing in the remote agents list

6. Doing so brings up Figure 7.95, using which you can modify the **Host IP/Name** and **Nick name** of the remote agent. If you had not added an external agent earlier with the IP and nick name of the remote agent, you can do so in the **Modify** mode by setting the **Add as External Agent** flag to **Yes**. However, if you have already added such an external agent, then you cannot modify that flag setting once again in the **Modify** mode. After making the necessary changes, click on the **Update** button in Figure 7.95 to register the changes.

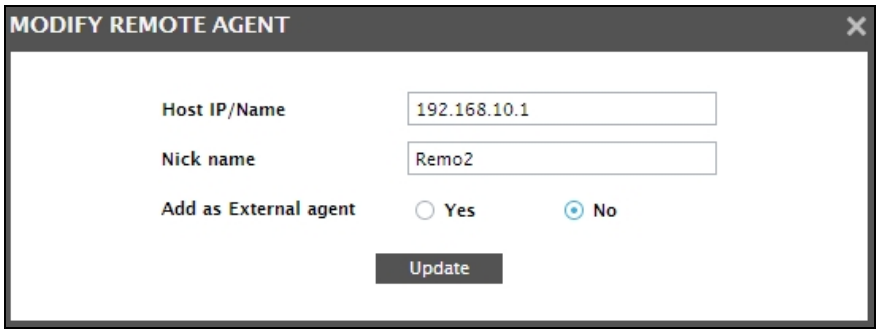

Figure 7.95: Modifying the IP address of the remote agent

After configuring a remote agent, assign the hosts to be monitored by the remote agent. You can either assign a remote agent to multiple hosts simultaneously, or can do it for one host component at a time.

To assign a remote agent to multiple hosts at one shot, do the following:

1. First, click on the **Associate Hosts** button in Figure 7.94.
2. When Figure 7.96 appears, select the remote agent in question from the **Remote agent** drop-down. All the hosts that are currently not managed by the chosen remote agent will appear in the **OTHER HOSTS** list. From this list, select the hosts that are to be managed by the chosen remote agent, and click the **<** button to move the selection to the **HOSTS MANAGED BY THIS AGENT** list.

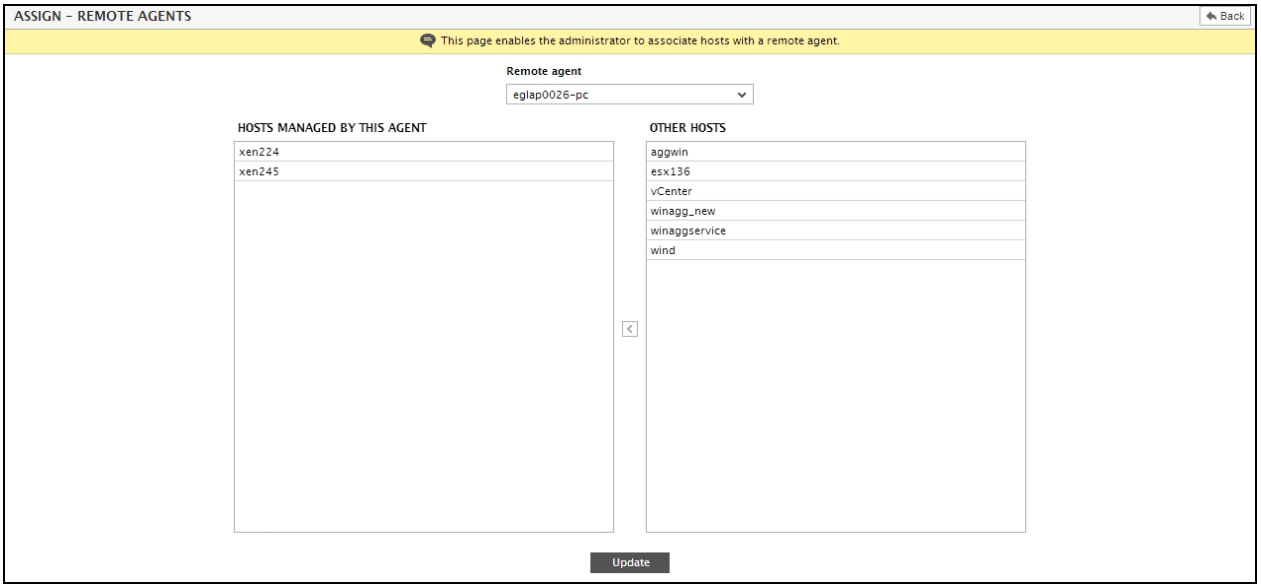

Figure 7.96: Assigning hosts to a remote agent

3. Finally, click the **Update** button to save the changes.
4. Alternatively, you can click the button corresponding to a remote agent in Figure 7.86 and then assign hosts to that agent using Figure 7.96.

Note:

It is recommended that a remote agent be used to support up to 10 targets.

Note:

- Typically, the remote agent should be installed and configured on the same operating system and locale as that of the servers that are monitored by that agent. In multi-lingual environments therefore, you would require a remote agent for every locale that is in use - for instance, in environments with servers that support both French and Japanese locales, you would require an exclusive remote agent for the French servers and another for the Japanese servers.
- A remote agent can monitor only those hosts that have the same architecture as the remote agent. For instance, a remote agent executing on a 32-bit host can monitor only another 32-bit host. Similarly, a remote agent deployed on a 64-bit host can only monitor other 64-bit hosts.

In some environments, administrators may have configured one/more remote agents, but might not have assigned any hosts to them - i.e., these remote agents might not be monitoring any host in the environment. Administrators might want to remove such "unused" agents from the eG Enterprise system. To enable administrators to isolate such unused remote agents and delete them, the **REMOTE AGENT CONFIGURATION** page marks every used remote agent with a 'lock' icon (see Figure 7.97) and the unused remote agents with a check box. To delete the unused agents, simply select the check box corresponding to such agents, and click the 'trash can' icon against them. **Remote agents that are in use (i.e., agents with the lock symbol) can be modified, but not deleted**.

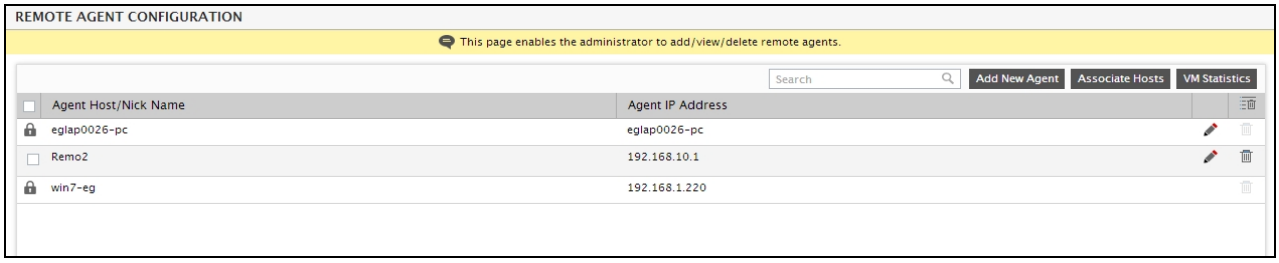

Figure 7.97: Deleting remote agents not monitoring any host

Reference:

To know how to assign a remote agent for each component at a time, refer to Section **7.6.4**.

7.6.4 How to Manage Components in an Agentless Manner?

While adding a new component using the eG administrative interface, the eG Enterprise system requests the administrator to indicate whether the component being added supports agentless monitoring or not (see Figure 7.62). This option appears only if at least one remote agent has been configured for the target environment. Otherwise, agent-based monitoring is used by default.

To manage a new component in an agentless manner, do the following:

1. Click the **Agentless** check box in Figure 7.98.

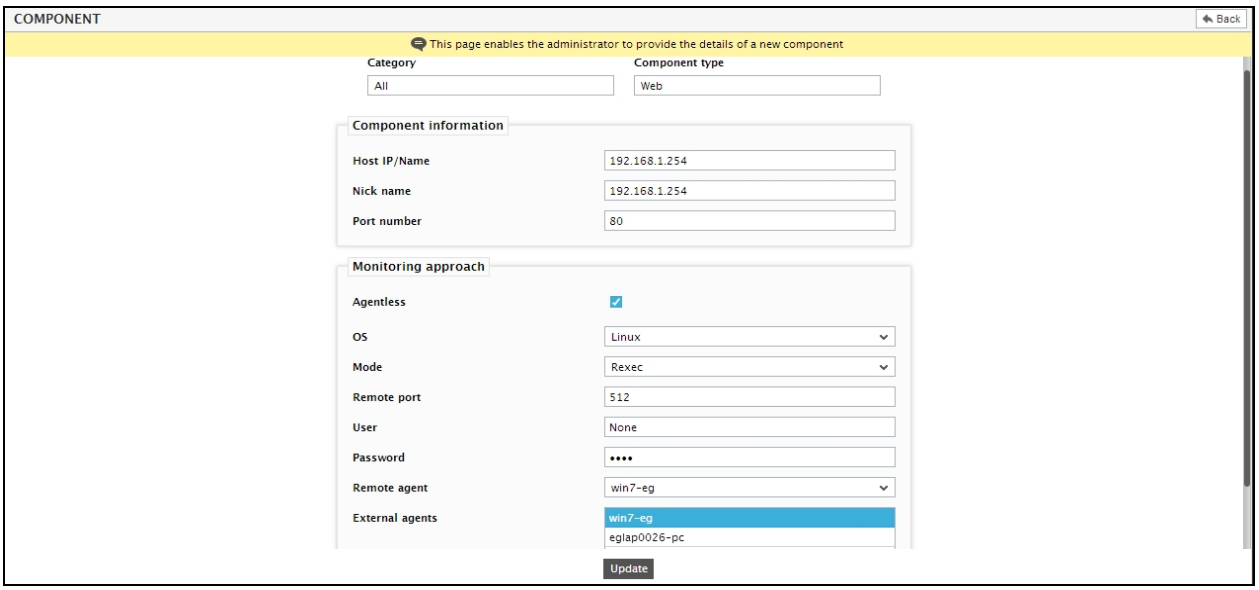

Figure 7.98: Enabling agentless monitoring for a Linux component with Rexec as the Mode

2. Next, select the operating system of the component to be monitored from the **OS** list box (see Figure 7.98). Remote agents can extract measures from all the eG-supported platforms (Windows 2000, Windows 2003/XP, Windows 2008/Windows 7/Windows Vista, Windows 8/2012, Linux, Solaris, HPUX, AIX, Tru64). If a Unix operating system is selected (i.e. Linux/Solaris/HPUX/AIX/Tru64) as the **OS**, then the following sequence of steps will apply:

   - As soon as any of the Unix platforms is selected from the **OS** list box, **Rexec**is displayed against **Mode**. **Rexec** executes commands remotely on a Unix host. Before proceeding with **Rexec**, make sure that the Rexec-related pre-requisites detailed in Section **7.6.2.3** are fulfilled.
   - The default port number of the **Rexec** command is 512, which will be displayed against the **Remote port** field.
   - The **rexec** command requires a valid user name and password at the remote host to connect to it. Therefore, provide the necessary login credentials against the **User** and **Password** text boxes.

     Note:

     While connecting to a remote host, **Rexec** sends passwords in clear text, and hence can be used only in secure environments. **SSH** is a more secure alternative
   - Alternatively, you can choose **SSH** as the **Mode** (see Figure 7.99). Before proceeding with the **SSH** mode, make sure that the SSH-related pre-requisites detailed in Section **7.6.2.3** are fulfilled.

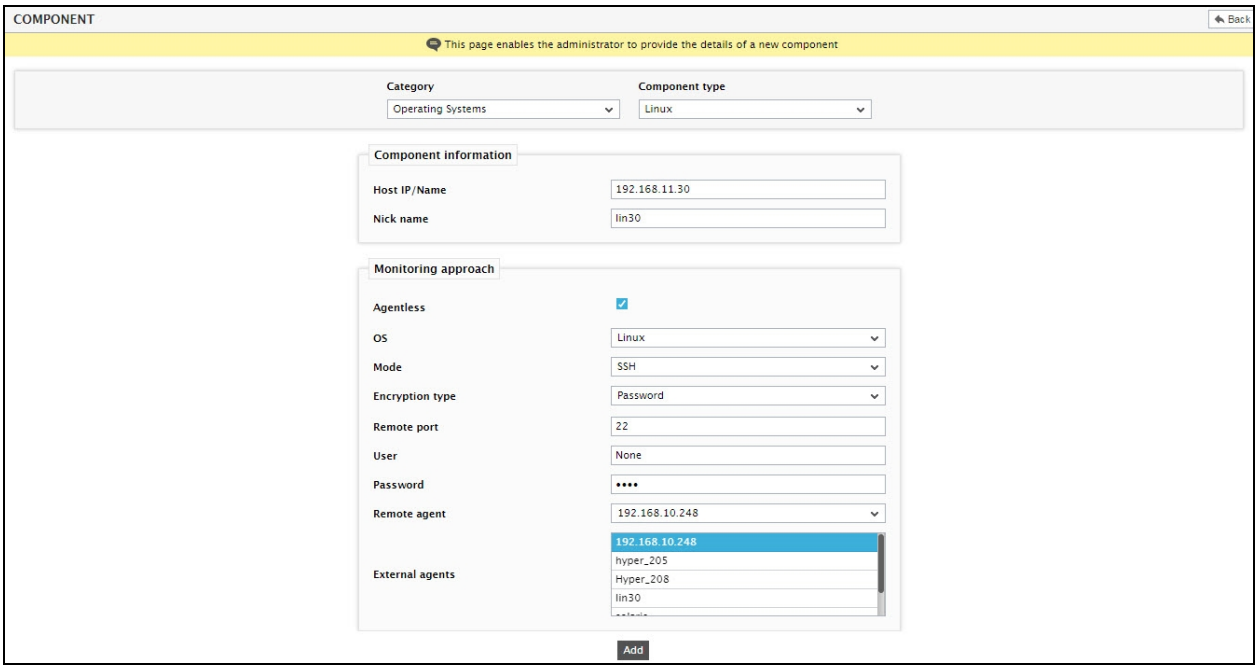

Figure 7.99: Enabling agentless monitoring of a Linux component with SSH as the Mode

- As stated in the **Pre-requisites** section (Section **7.6.2.3**), if **SSH** is chosen as the **Mode**, then you will also have to indicate the **Encryption type** for your SSH connection. By default, **Password** is set as the encryption type. If need be, you can select **Keybased** as the **Encryption type**.
- If **Encryption type** is set to **Keybased**, you will have to specify a **Key file name**. Here, specify the full path to the private key file **id_rsa**, which will be available in the location, **<EG_AGENT_INSTALL_DIR>\agent\sshkeys**. For instance, if the eG agent has been installed in the D:\eGurkha folder, your **Key file name** will be: D:\eGurkha\agent\sshkeys\id_rsa
- The port number for secure shell access has to be specified against **Remote port** field. By default, 22 is the port at which the SSH server listens.
- Regardless of the **Encryption type** chosen, provide the valid **User** name and **Password** using which the SSH connection to the host is to be established. Against **User**, provide the name of the user as who you logged into the target Unix host for copying the **authorization_keys** file (see point 2 of Section **7.6.2.3**). Against **Password**, specify the pass phrase that is associated with the private key file, **id_rsa**. If you are using the default public/private keys bundled with the eG agent, **eginnovations** will be the pass phrase.

3. On the other hand, if a Windows operating system (Windows 2000/2003/2008/NT/XP/7/8/2012) is selected from the **OS** list box, then the **Mode** will change to **Perfmon.** Before selecting **Perfmon** as the **Mode**, make sure that the pre-requisites detailed in Section **7.6.2.2**

Note:

Besides **SSH** and remote **Perfmon**, agentless monitoring can be done using SNMP and externally accessible APIs/protocols (e.g., SQL for Oracle, MQ API for IBM MQ, HTTP for WebLogic, Web Service, CLI for storage devices etc.).

4. From the list of remote agents configured in the, select the remote agent that will monitor the component being added (see Figure 7.100).

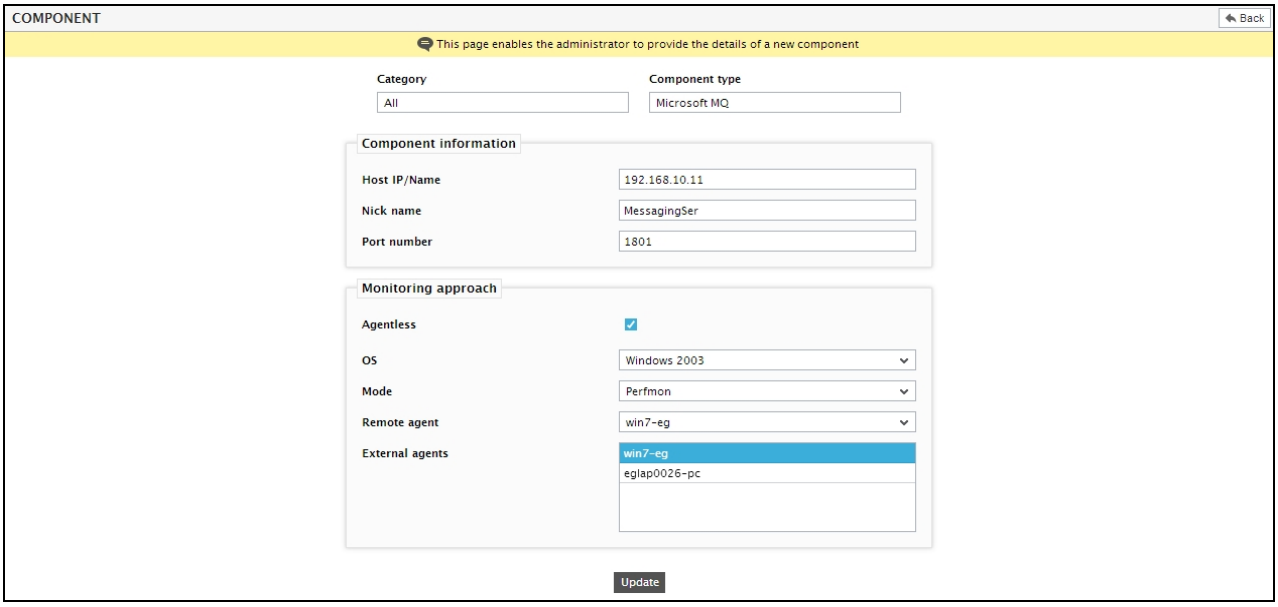

Figure 7.100: Enabling Agentless support for a Windows component

5. Finally, associate one or more **External agents** with the component, and click the **Update** button to complete the configuration.

Note:

- Like other agent types, remote agents too communicate with the eG manager through HTTP/HTTPS.
- A single agent can function as a remote agent as well as an external agent. However, an external agent for which client emulation has been enabled cannot function as a remote agent.

Note:

It is recommended that a remote agent be used to support up to 10 targets.

7.6.5 Applications Supported for Agentless Monitoring

The table below summarizes the applications that are supported for agentless monitoring.

Application/ Component	Solaris	Red Hat Linux	Win 2008/2012	AIX	HPUX	Remarks
Active Directory			✓			
AIX server				✓		
AIX LPARs on IBM pSeries Servers	✓	✓	✓			
Apache Web server	✓	✓		✓	✓	Web adapter not sup-

Application/ Component	Solaris	Red Hat Linux	Win 2008/2012	AIX	HPUX	Remarks
						ported
ASP .Net server			✓			
Atlantis ILIO		✓	✓			
AWS EC2 Cloud			✓			
BizTalk			✓			
Backup SQL			✓			
BlackBerry server			✓			
Borland Enterprise Server	✓	✓	✓	✓	✓	
Coldfusion	✓	✓				
Checkpoint Firewall	✓	✓	✓	✓	✓	
Checkpoint Smart Appliance	✓	✓	✓	✓	✓	
Cisco UCS Manager	✓	✓	✓	✓	✓	
Citrix Access Gateway			✓			
Citrix Branch Repeater	✓	✓	✓	✓	✓	
Citrix CloudBridge	✓	✓	✓	✓	✓	
Citrix Farm server			✓			
Citrix NetScaler ADC	✓	✓	✓	✓	✓	
Citrix Secure Gateway			✓			
Citrix STA			✓			
Citrix Web Interface (Nfuse)			✓			
Citrix License server			✓			
Citrix Netscaler VPX	✓	✓	✓	✓	✓	
Citrix NetScaler HDX Insight	✓	✓	✓	✓	✓	
Citrix NetScaler Web Insight	✓	✓	✓	✓	✓	
Citrix XenServer		✓	✓			
Citrix XenMobile MDM	✓	✓	✓	✓	✓	
Citrix XenMobile 10	✓	✓	✓	✓	✓	
Citrix ShareFile	✓	✓	✓	✓	✓	
Citrix AppController	✓	✓	✓	✓	✓	
Citrix StoreFront			✓			

Application/ Component	Solaris	Red Hat Linux	Win 2008/2012	AIX	HPUX	Remarks
Citrix Storage Zones	✓	✓	✓	✓	✓	
COM+ Applications			✓			
Data Domain	✓	✓	✓	✓	✓	
DB2 DPF	✓	✓	✓	✓	✓	
DB2 Universal Database	✓	✓	✓	✓	✓	
Dell Compellent	✓	✓	✓	✓	✓	
Dell EqualLogic	✓	✓	✓	✓	✓	
Dell PowerEdge VRTX	✓	✓	✓	✓	✓	
Delta UPS	✓	✓	✓	✓	✓	
DHCP			✓			
DNS	✓	✓	✓	✓	✓	
Dockers	✓	✓	✓			
Domino Application Server	✓	✓	✓	✓	✓	
Domino Mail Server	✓	✓	✓	✓	✓	
Double Take Availability	✓	✓	✓	✓	✓	
eDirectory	✓	✓	✓	✓	✓	
Egenera PAN Manager	✓	✓	✓	✓	✓	
EMC CLARiiON	✓	✓	✓	✓	✓	
EMC VNX Unified Storage	✓	✓	✓	✓	✓	
EMC XtremIO	✓	✓	✓	✓	✓	
FAST Search for SharePoint 2010			✓			
FioranoMQ Server	✓	✓	✓			
FTP	✓	✓	✓	✓	✓	
Glassfish Enterprise Server	✓	✓	✓	✓	✓	
Generic server	✓	✓	✓			
Generic Mail server	✓	✓	✓	✓	✓	
Generic Netware server	✓	✓	✓	✓	✓	
Generic Storage RAID	✓	✓	✓	✓	✓	

Application/ Component	Solaris	Red Hat Linux	Win 2008/2012	AIX	HPUX	Remarks
Generic SNMP server	✓	✓	✓	✓	✓	
Groupwise Internet Agent (IA)	✓	✓	✓	✓	✓	
Groupwise Message Transfer Agent (MTA)	✓	✓	✓	✓	✓	
Groupwise PostOffice Agent (POA)	✓	✓	✓	✓	✓	
HPUX server					✓	
Hitachi SAN AMS			✓			
Hitachi SAN USP			✓			
HP Blade Server	✓	✓	✓	✓	✓	
HP EVA StorageWorks	✓	✓	✓	✓	✓	
HP P2000	✓	✓	✓	✓	✓	
HP 3PAR	✓	✓	✓	✓	✓	
IBM MQ server	✓	✓	✓	✓	✓	
IBM HTTP server	✓	✓		✓	✓	Web adapter not supported
IBM DS Raid Storage	✓	✓	✓	✓	✓	
IBM DS 8000	✓	✓	✓	✓	✓	
IBM Integration Bus	✓	✓	✓	✓	✓	
IIS SSL server			✓			Web adapter not supported
IIS Web server			✓			Web adapter not supported
IBM Storwize	✓	✓	✓	✓	✓	
Informix		✓	✓	✓	✓	
InfoBlox	✓	✓	✓	✓	✓	
Intersystems Cache	✓		✓			
ISA Proxy server			✓			
Java applications	✓	✓	✓	✓	✓	
JBoss Server	✓	✓	✓	✓	✓	
JRun Server	✓	✓	✓			
LDAP	✓	✓	✓	✓	✓	
Leostream connection broker		✓				
Linux	✓					

Application/ Component	Solaris	Red Hat Linux	Win 2008/2012	AIX	HPUX	Remarks
MaxDB server	✓	✓	✓	✓	✓	
Microsoft Azure	✓	✓	✓	✓	✓	
Microsoft Lync			✓			
Microsoft Project 2010			✓			
Microsoft SharePoint			✓			
Microsoft Dynamics AX			✓			
Microsoft Dynamics NAV			✓			
MS File server			✓			
MS FTP server			✓			
MS Print server			✓			
MS Proxy server			✓			
MS Radius server			✓			
MS RAS server			✓			
MS SQL 7/ 2000 / 2005 / 2008 / 2012 / 2014			✓			
MSMQ server			✓			
MS Transaction server			✓			
MS Exchange 5.5 server			✓			
MS Exchange 2000/2003 server			✓			
MS Exchange Instant Messenger			✓			
MS Systems Management server			✓			
MySQL	✓	✓	✓			
Netscape Application server	✓		✓	✓	✓	
	✓	✓	✓	✓		
NetApp USD	✓	✓	✓	✓	✓	
NexentaStor	✓	✓	✓	✓	✓	
NGINX web server	✓	✓	✓	✓	✓	
Nimble Storage	✓	✓	✓	✓	✓	
NTP Server			✓			
OpenVMS	✓	✓	✓	✓	✓	

Application/ Component	Solaris	Red Hat Linux	Win 2008/2012	AIX	HPUX	Remarks
Oracle 9i application server	✓	✓	✓	✓	✓	
Oracle 10G application serv er	✓	✓	✓	✓	✓	
Oracle HTTP server	✓	✓	✓	✓	✓	
Oracle Database	✓	✓	✓	✓	✓	
Oracle Forms						
Oracle RAC	✓	✓	✓	✓	✓	
Oracle VDI Broker	✓	✓	✓			
Oracle VirtualBox	✓	✓	✓			
Oracle VM Server	✓	✓	✓			
Oracle VM Manager	✓		✓			
Orion	✓	✓	✓	✓	✓	
Qmail	✓	✓		✓	✓	
QNAP NAS	✓	✓	✓	✓	✓	
Postgre SQL	✓	✓	✓	✓	✓	✓
Radius	✓	✓		✓	✓	
RHEV Hypervisor		✓	✓			
RHEV Manager		✓	✓			
SAP ABAP Instance	✓	✓	✓	✓	✓	
SAP BOBI	✓	✓	✓	✓	✓	
SAP Internet Transaction server (ITS)	✓	✓	✓	✓	✓	
SAP Netweaver Web application server	✓	✓	✓	✓	✓	
Siebel application server	✓	✓	✓	✓	✓	
Siebel Gateway server	✓	✓	✓	✓	✓	
Siebel Web server	✓	✓	✓	✓	✓	
Silverstream application server	✓	✓			✓	
SiteMinder Policy server	✓	✓			✓	
Solaris	✓					
Sonic Firewall	✓	✓	✓	✓	✓	
SSL Web server	Partial	Partial		Partial	Partial	Web adapter not sup-ported
Sun Java application		✓	✓		✓	

Application/ Component	Solaris	Red Hat Linux	Win 2008/2012	AIX	HPUX	Remarks
server						
Sun Java Messaging server		✓	✓			
Sun Java Directory Server		✓	✓			
Sybase Database	✓	✓	✓	✓	✓	
Symantec Backup			✓			
Tomcat server	✓	✓	✓	✓	✓	
Tibco EMS			✓			
Terminal Services Licensing server			✓			
VNX Unified Storage	✓	✓	✓	✓	✓	
Voyager Transaction Processor			✓			
Voyager Load Balancer			✓			
Voyager Front End			✓			
vCloud Director	✓		✓			
VMware vSphere ESX		✓	✓			
VMware vCenter			✓			
Watchguard Firewall	✓	✓	✓	✓	✓	
Microsoft RDS			✓			
WebLogic	✓	✓	✓	✓	✓	
WebSphere	✓	✓	✓	✓	✓	
WINS			✓			
Windows DNS			✓			
Windows Domain Controller			✓			
Windows Generic server			✓			
Windows server			✓			
2X Terminal Server			✓			
2X Client Gateway			✓			
2X Publishing Agent			✓			

7.7 Internal Agent Assignment

By default, if a host has multiple IP addresses, the eG Enterprise system requires one agent license for each IP address that is managed internally. Likewise, if multiple nicknames are used for the same IP address, a

separate internal agent license is used for each unique nickname that has been specified. In many large environments, a single host has many IP addresses, each with different nicknames. The agent per system capability is intended to optimize the internal agent license usage in such large infrastructures. If this capability is enabled by the eG license, the administrator has the option of overriding the default eG agent licensing policy. For example, suppose a host A has two IP addresses 192.168.10.7 and 10.10.10.1, and that the first IP address 192.168.10.7 has already been managed in the eG Enterprise system. When adding the second IP address, 10.10.10.1, the administrator has the option of overriding eG's default internal agent licensing policy - in this example, the administrator can indicate that the internal agent for the IP address 10.10.10.1 is actually the one that is already associated with the IP address 192.168.10.7. By doing so, the administrator can ensure that a single agent license is sufficient to manage all the IP addresses and applications executing on a host.

The **COMPONENT** page of the eG administrative interface facilitates such an internal agent assignment (see Figure 7.101).

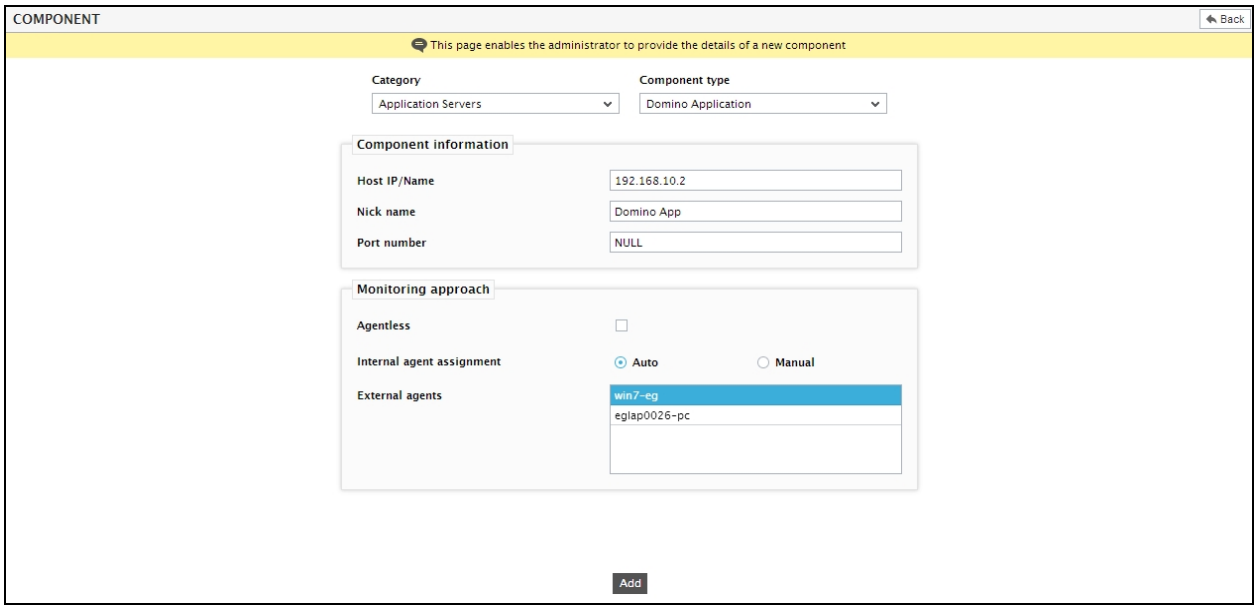

Figure 7.101: Assigning an internal agent to a new component

As you can see, Figure 7.101 displays an **Internal agent assignment** field.

Note:

The **Internal agent assignment** field will appear only if the following conditions are fulfilled:

- The eG license enables the **Agent per system**flag
- The **Agentless**flag in Figure 7.101 is deselected. In other words, users will not have the option of mapping an IP to an internal agent, if agentless monitoring is enabled.

To map the IP being added to an existing internal agent using Figure 7.101, do the following:

1. By default, the **Auto** option against the **Internal agent assignment** field will be selected. This indicates that by default, eG maps every configured IP/nick name with a separate internal agent. To manually define the IP-internal agent association, select the **Manual** option.

2. Upon choosing **Manual**, an additional **Internal agent** list box will appear (not shown in Figure 7.101). From this list box, select the internal agent that needs to be associated with the IP of the component being created.
3. Finally, associate one or more **External agents** with the component, and click the **Add** button to complete the configuration.

7.8 Configuring Tests

For each component type, eG Enterprise has a set of predefined tests. These tests can be configured via the **Tests** menu in the **Agents** menu. While many of eG Enterprise tests do not require any manual configuration, some tests require explicit, manual configuration. You can either configure such tests for all components of a particular type at one shot, or provide component-specific configuration. The sections that will follow will discuss both these test configuration options.

7.8.1 Default Test Configuration

To configure tests for all components of a type, pick the **Default Configuration** option from the **Tests** menu of the **Admin** tile. Figure 7.102 will then appear. Pick the **Component type** to which the default test configuration applies, indicate whether the test to be configured is a **Performance** test or a **Configuration** test by picking an option from the **Test type** list, and then, select the test to be configured by default from the **Test name** list.

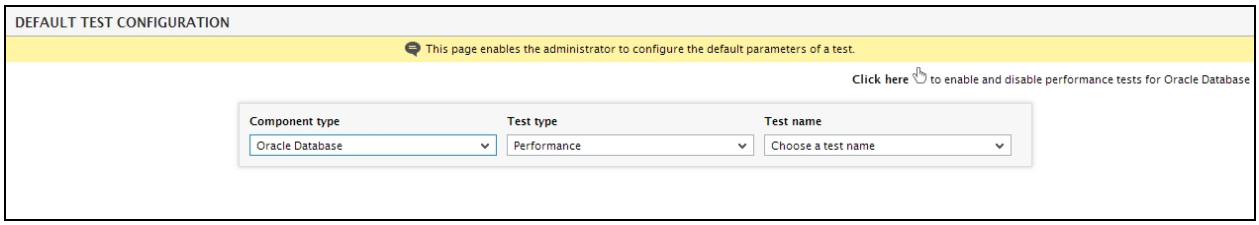

Figure 7.102: Selecting the test that needs to be configured by default for a chosen component-type

The default parameters of the chosen test will then appear as indicated by Figure 7.103.

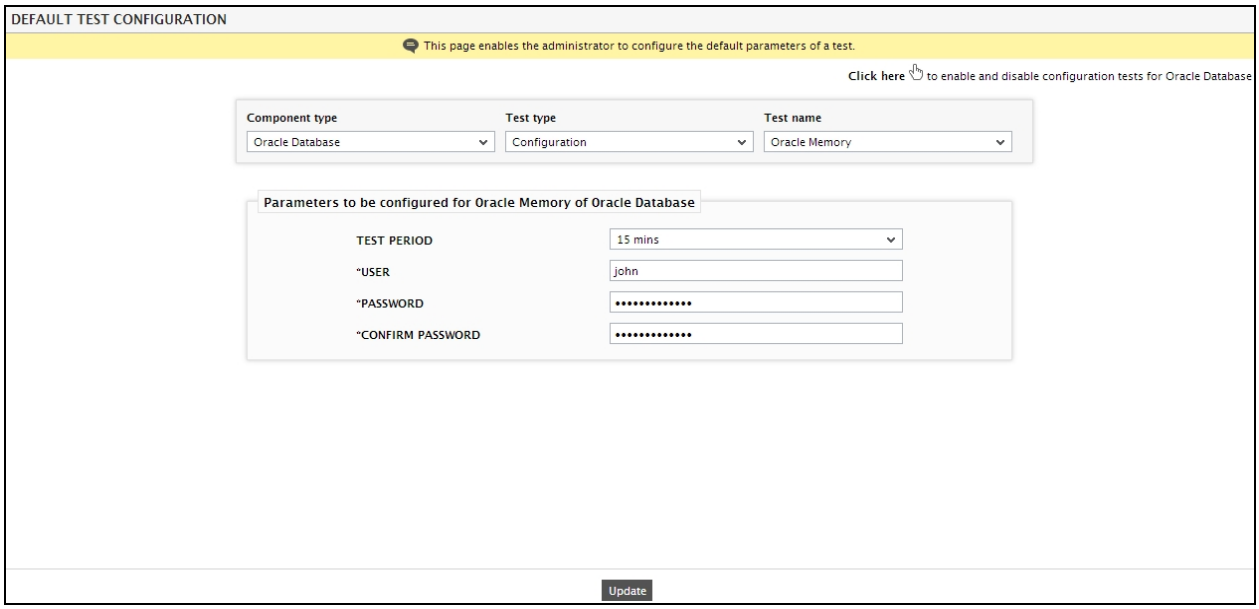

Figure 7.103: Configuring the default parameters of a test for the Oracle database component-type

The parameters prefixed by a * (asterisk) denote those parameters that require manual configuration. Once the default parameters are configured, click the **Update** button in Figure 7.103 to save the changes to the test configuration.

Some tests take certain special default parameters. For instance, the process parameter of the **Processes** test. To ensure that the eG agent monitors only a specific set of processes for all components of a particular type, you can configure a comma-separated list of default process patterns for the Processes test of that component type in the **DEFAULT TEST CONFIGURATION** page. However, in case of components where `PROCESS` definitions vary with the operating system on which the component is deployed (eg., Terminal server, Microsoft Print server, Domino mail server, etc.), the `PROCESS` parameter is set to `{EG_RUNTIME}` by default.

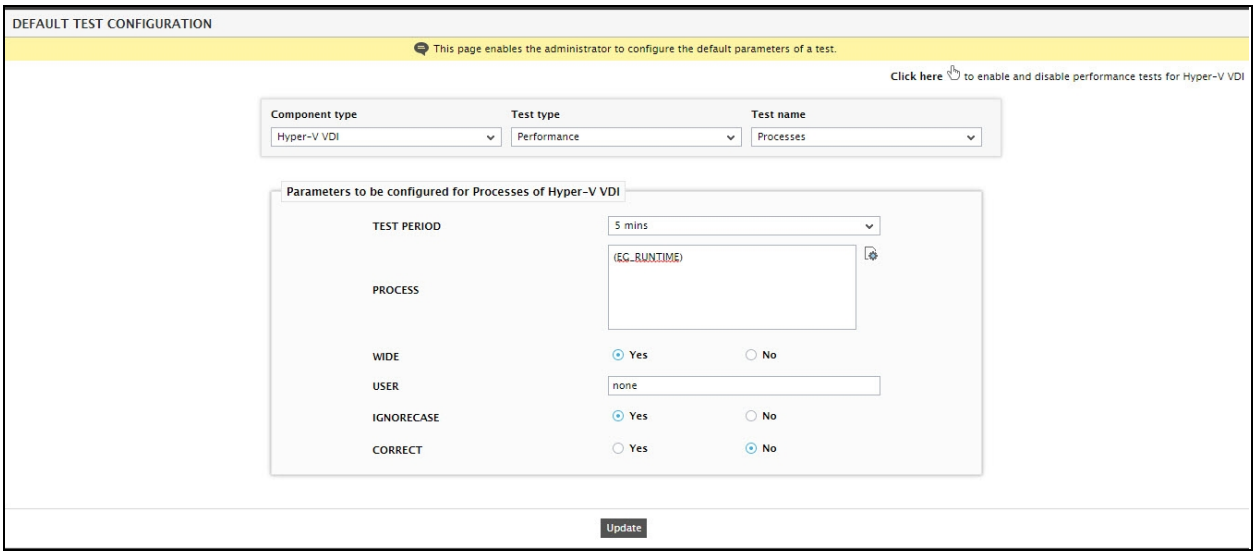

Figure 7.104: The default test parameters of Processes test

The variable **{EG_RUNTIME}** implies that the Processes test automatically determines the process pattern to be monitored at run-time, depending upon the operating system on which the test executes. To view the OS-specific process patterns that are associated with a component-type by default, and to make changes to them if required, lick on the **Configure OS Patterns** ( ) button next to **{EG_RUNTIME}** (see Figure 7.104). An **OS SPECIFIC PATTERN CONFIGURATION** window (see Figure 7.105) then pops out, which lists the **Operating systems** on which the component executes, and the default **PROCESS pattern** configurations that correspond to every operating system.

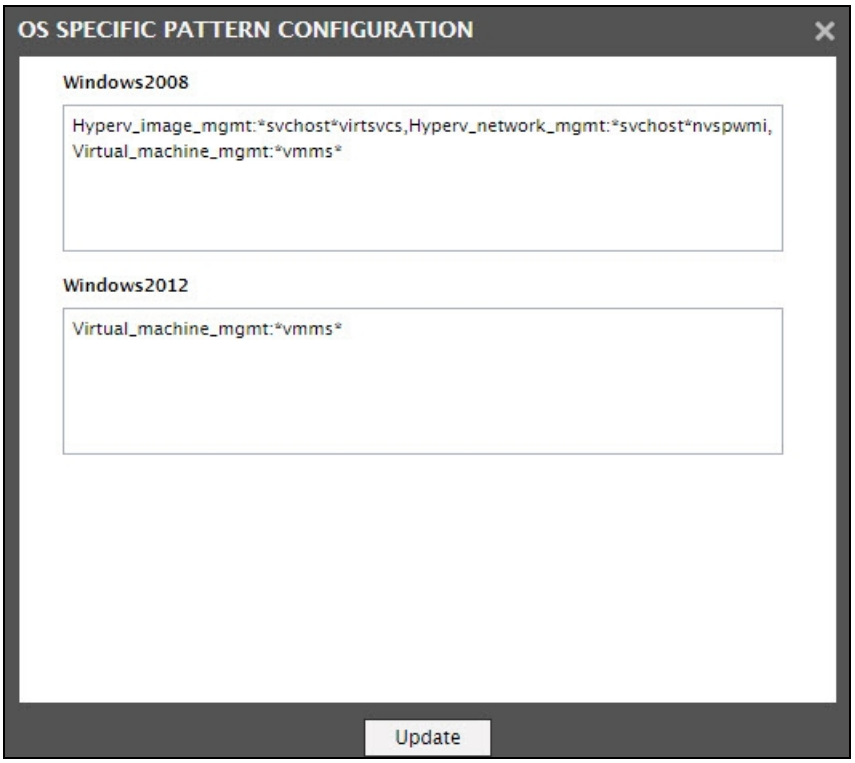

Figure 7.105: The OS-specific process configurations

If need be, additional process patterns can be configured per operating system, in the format: *processName:processPattern*. *processName* is the unique name that you need to assign to a pattern, and *processPattern* is the pattern of processes that you want to monitor. Make sure that every additional *processName:processPattern* pair is separated by a comma (,). The default configurations too can be modified or removed, as required. However, if you want to add a new operating system and/or a new component type and set its default process patterns, follow the steps detailed below.

1. Edit the **eg_tests.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
2. In the **OS_PROCESSPATTERN** section of the file, provide an entry for the new operating system, in the following format:

   {Operating_system}:{Internal_component_type}:{Internal_testname}={processName:processPattern}

   For instance, if the default process patterns for the Windows XP version of a Terminal server are to be set, then the specification could be:

   *WindowsXP:Terminal_server:ProcessTest=TermSvr:*svchosts**

   Here, *WindowsXP* refers to the *Operating_system*, Terminal_server is the eG internal name for the Terminal server, and ProcessTest is the eG internal name for the Processes test. Note that multiple processName:processPattern pairs can be provided as a #- separated list - i.e., *TermSvr:*svchosts*#Term:*svc**.

To know the internal component-type names and test names, refer to the **eg_lang*.ini** file in the **<EG_INSTALL_DIR>\manager\config**, where * is the language code that represents the language preference that you have set using the **USER PROFILE** page. In this file, the component types, measure names, test names, layer names, measure descriptions, and a wide range of other display information are expressed in a particular language, and are mapped to their eG equivalents. Search the file for the component-type and test of interest to you. For **INSTANCE**, to know the internal name for the Processes test, search the **[TEST_NAME_MAPPING]** section of the file for Processes. This will reveal the internal test name that maps to the Processes test. Similarly, search the **[TYPE_NAME_MAPPING]** section to figure out the internal component type name.

3. Finally, save the **eg_tests.ini** file.

Besides the process parameter of Processes test, the **{EG_RUNTIME}** variable appears against the process parameter of the Windows Processes test, and the servicename parameter of the Windows Services test. While the default OS-specific configurations of the Windows Processes test can be overridden in the same manner as discussed for the Processes test, for the Windows Services test however, it needs to be done a little differently. In case of the Windows Services test, clicking on the **Configure OS Patterns** button next to **{EG_RUNTIME}** (see Figure 7.106) reveals a list of **Operating Systems** and the **SERVICENAME**s that will be monitored by default in every operating system on which the test executes (see Figure 7.107).

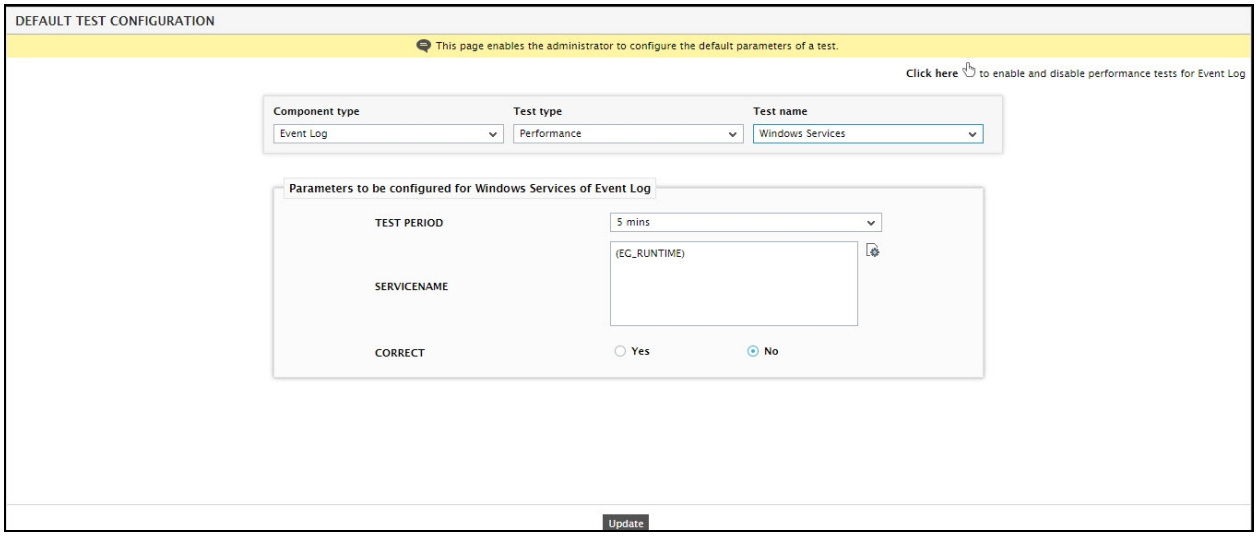

Figure 7.106: The default parameters of the Windows Services test

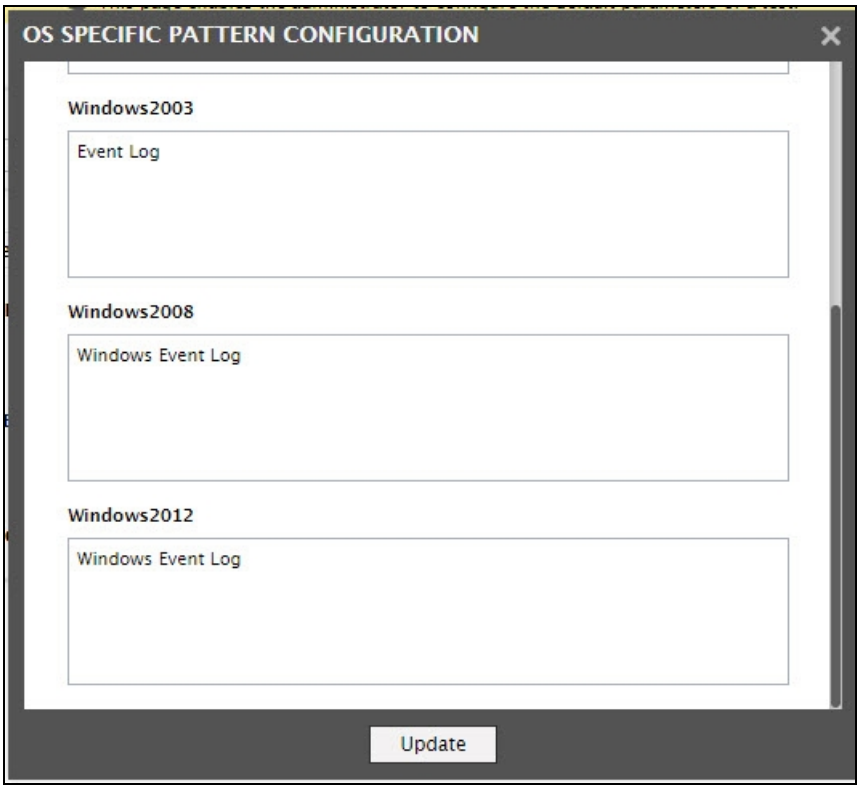

Figure 7.107: The OS-specific service configurations

If need be, additional services can be configured for monitoring per operating system. Make sure that every additional specification is separated by a comma (,). The default configurations too can be modified or removed, as required. However, if you want to add a new operating system and/or a new component type and set its default service configurations, follow the steps detailed beloow:

1. Edit the **eg_tests.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
2. In the **OS_PROCESSPATTERN** section of the file, provide an entry for the new operating system, in the following format:

   {Operating_system}:{Internal_component_type}:{Internal_testname}={serviceName}

   For instance, if the default service names that correspond to a Terminal server on Windows 2000 are to be set, then the specification would be:

   Windows2000:Terminal_server:WinServiceTest=Terminal Services

   Here, *Windows2000* refers to the *Operating_system*, Terminal_server is the eG internal name for the Terminal server, and WinServiceTest is the eG internal name for the WindowsServices test. When configuring multiple service names for an operating system, use **#** as the separator.

   To know the internal component-type names and test names, refer to the **eg_lang*.ini** file in the **<EG_INSTALL_DIR>\manager\config**, where * is the language code that represents the language preference that you have set using the **USER PROFILE** page. In this file, the component types, measure names, test

names, layer names, measure descriptions, and a wide range of other display information are expressed in a particular language, and are mapped to their eG equivalents. Search the file for the component-type and test of interest to you. For **INSTANCE**, to know the internal name for the WindowsServices test, search the **[TEST_NAME_MAPPING]** section of the file for WindowsServices. This will reveal the internal test name that maps to WindowsServices. Similarly, search the **[TYPE_NAME_MAPPING]** section to figure out the internal component type name.

3. Finally, save the **eg_tests.ini** file.

Note:

- Typically, the OS-specific process/service configurations of Windows 2003 automatically apply to Windows XP as well.
- The **{EG_RUNTIME}** variable and the adjoining **Configure OS Patterns** button will be available only while modifying the default parameters of the Processes, Windows Processes, or Windows Services test. In other words, you can view or modify the OS-specific configurations for any of the above-mentioned tests, only if you choose the **Default Configuration** option of the tests. On the other hand, while configuring one of these tests for a particular component of a type (i.e., if you choose the **Specific Configuration** option), the eG agent automatically identifies the operating system on which the said component is executing, and directly displays the default process/service configuration (as the case may be) that corresponds to that operating system in the test configuration page. In this case therefore, the **{EG_RUNTIME}** variable will not be visible.

7.8.2 Component-specific Test Configuration

The user interface also allows for a test to be configured differently for each component.

To configure a test for a specific component of a type, pick the **Specific Configuration** option from the **Tests** menu in the **Agents** tile. Figure 7.108 will then appear.

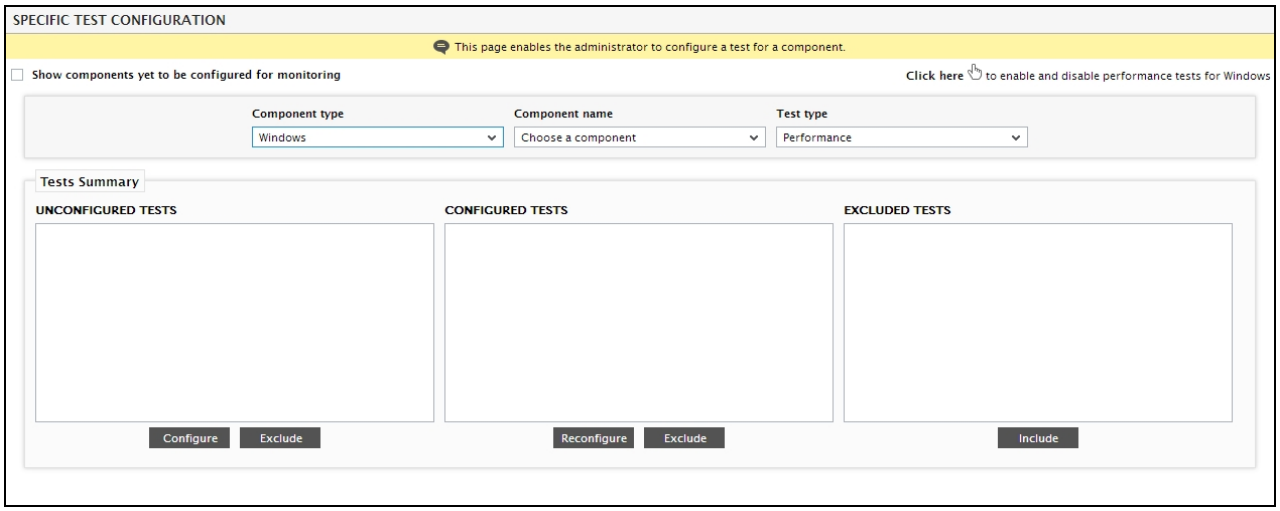

Figure 7.108: Viewing the configuration states for all the tests pertaining to specific component

Then, follow the steps given below:

1. All components of the chosen **Component type** will then populate the **Component name** list. Select the component for which a test is to be configured.
2. First, from the **Component type** list, choose the type of component. By default, all managed component types will populate the **Component type** list. You can filter this list (if required) to list only those types of components for which no test has been configured yet. For this, click the **Show components yet to be configured for monitoring** flag in Figure 7.108. This will condense the **Component type** list, thus helping you quickly pick the component for which a test is to be configured.

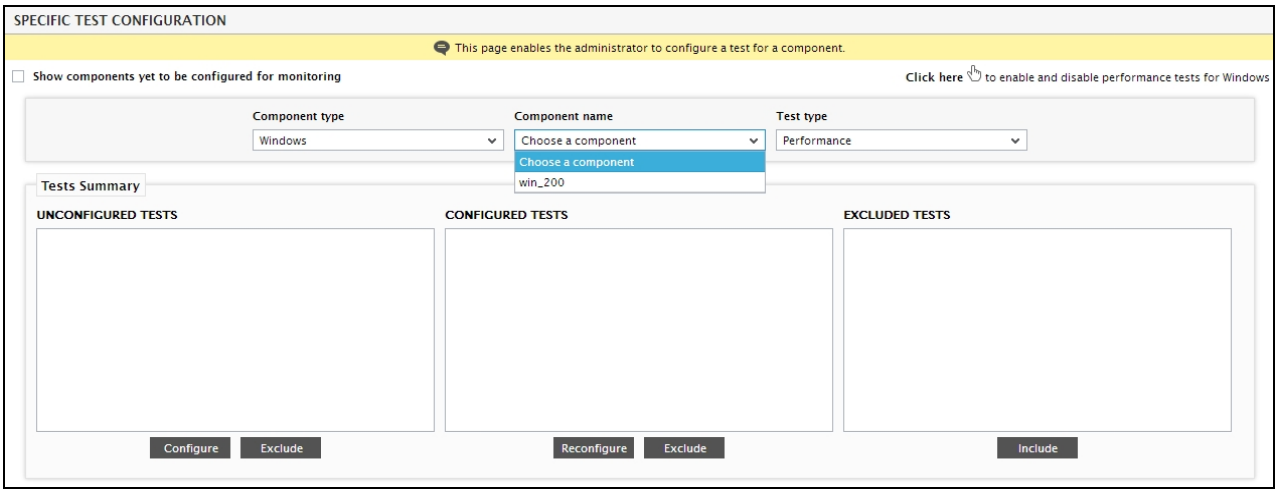

Figure 7.109: Selecting a Component type and then a Component

3. Then, select the type of test to be configured - i.e., whether a **Performance** test or a **Configuration** test. The **Configuration** option will be available only if the eG license enables the **Configuration Management** capability.
4. Doing so will invoke an **AGENT SUMMARY** section, using which you can instantly determine which internal agent and external agent is monitoring the chosen **Component**.
5. Below the **AGENT SUMMARY**, you will find a list of **UNCONFIGURED TESTS**, **CONFIGURED TESTS**, and **EXCLUDED TESTS** for the chosen **Component**. The **UNCONFIGURED TESTS** list displays those tests that have been enabled for the chosen component-type, but are yet to be configured for this component. To configure any of the listed tests, select the test of interest from the **UNCONFIGURED TESTS** list, and click the **Configure** button. The parameters of the chosen test will then appear as depicted by Figure 7.110 will appear.

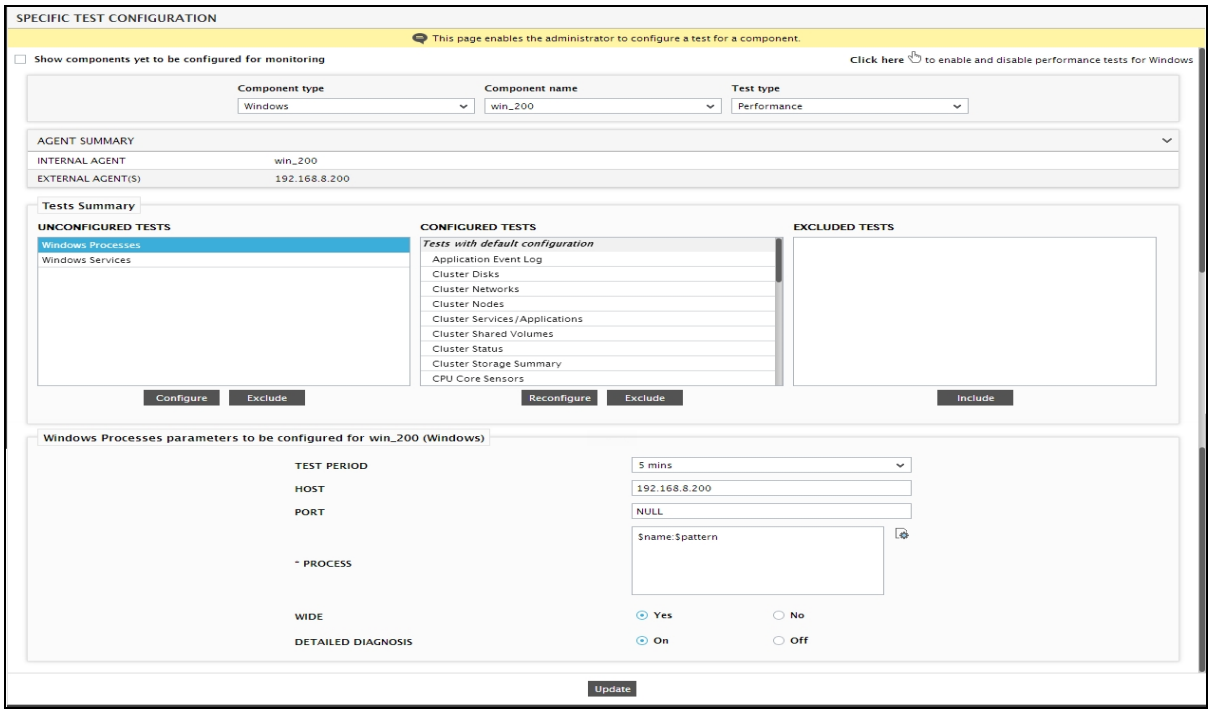

Figure 7.110: Configuring an unconfigured test for a specific component

6. Here again, the parameters prefixed by a * (asterisk) denote those parameters that require manual configuration. Once the parameters are configured, click the **Update** button in Figure 7.110 to save the changes to the test configuration.

   Note:

   When changing default configurations of tests, the values with "$" indicate variables that will be replaced by eG Enterprise according to the specific component being managed (e.g., if the URL is "http://$hostName:$port", for a specific web server running on the host www.abc.com on port 80, the URL parameter will be set to "http://www.abc.com:80").

7. If multiple components of the same component type are awaiting configuration, then an **Apply to other components** button will appear in the test configuration page (see Figure 7.111). Click on this button if you want the test configuration to be applied to one/more components of the chosen type.

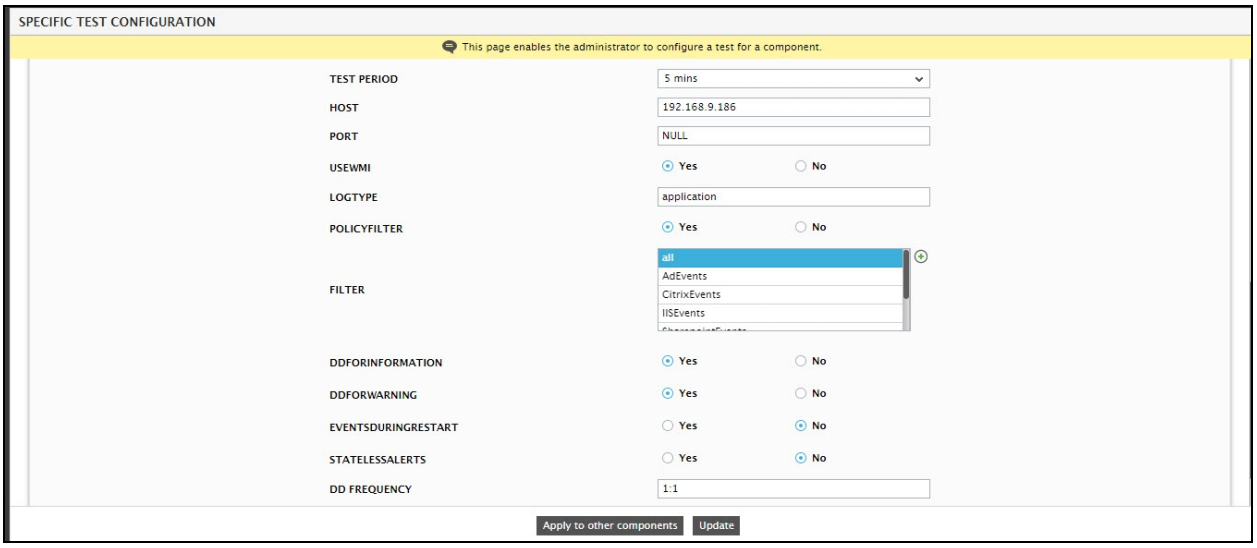

Figure 7.111: The parameters associated with the Terminal Authentication test for a specific Microsoft Terminal server

8. Figure 7.112 will then appear. Using Figure 7.112, you can choose the **Parameters** that are to be applied. To apply all the **Parameters**, select the check box that precedes the column name **Param Name**. Then, select the components to which the chosen parameters are to be applied from the **Existing Components** list, and click < button to transfer the selection to the **Components to be configured** list. Finally, click the **Apply** button. This ensures that the values passed to the selected parameters are automatically applied to the components in the **Components to be configured** list. This saves the time and labor involved in repeating the same set of configurations for every component of a particular component type.

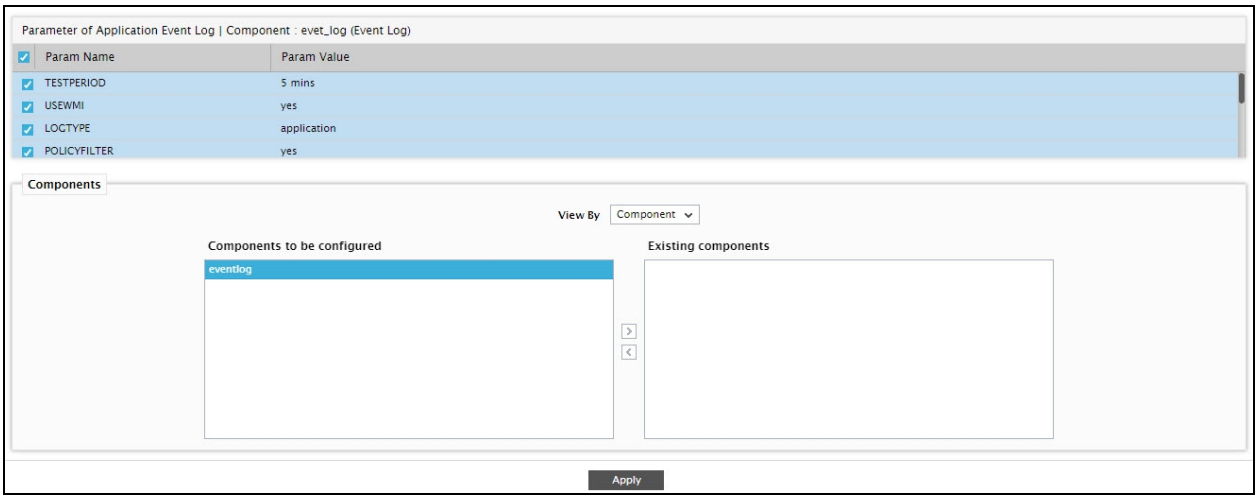

Figure 7.112: Choosing the parameters and components to which the configuration is to be applied

9. Let us now return to the **SPECIFIC TEST CONFIGURATION** page of Figure 7.110 and explore the other options it provides. The **CONFIGURED TESTS** list of Figure 7.110, displays those tests that have already been configured - such tests are typically of two types - tests that take the default configuration, and tests that have been specifically configured for the chosen **Component**. Accordingly, the **CONFIGURED TESTS** list box will list tests under two categories - **Tests with specific configuration** and **Tests with default configuration**. If

required, you can reconfigure any of these tests, by picking the test of interest from the **CONFIGURED TESTS** list and clicking the **Reconfigure** button. This will once again display the parameters of the chosen test, allowing you to edit the values of any of the parameters.

10. The **EXCLUDED TESTS** list displays tests that have been excluded for a specific component - i.e., tests that will not execute for a selected **Component**. To exclude a test that has not yet been configured, pick the test from the **UNCONFIGURED TESTS** list, and click on the **Exclude** button. To exclude a test that has already been configured, pick the test from the **CONFIGURED TESTS** list, and click the **Exclude** button. If more than one component of the chosen type has been managed, then, clicking on the **Exclude** button will lead you to 7.8.2, using which you can exclude the selected test for the other components of that type as well.

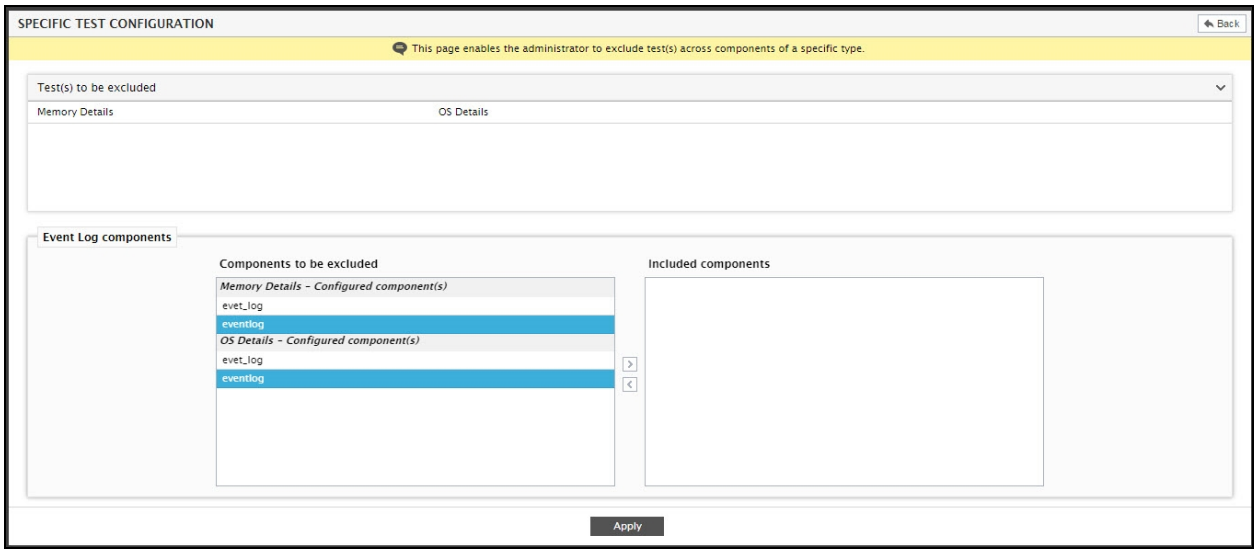

Figure 7.113: Excluding a test for more than one component

11. The **INCLUDED COMPONENTS** list of 7.8.2 will list all the components (of the chosen component type) for which the chosen test is currently available. The components in this list will be grouped according to the state in which the test (chosen for exclusion) exists with respect to a component. For instance, if the test chosen for exclusion is in the **UNCONFIGURED STATE** for a few components, then such components will be listed under the group **<TestName> - Unconfigured component(s)** in the **INCLUDED COMPONENTS** list. The **COMPONENTS TO BE EXCLUDED** list on the other hand, will by default display that component for which the test is currently being chosen for exclusion. To exclude the test for the other components, pick them from the **INCLUDED COMPONENTS** list regardless of their state and click the **<<** button; this will transfer the selection to the **COMPONENTS TO BE EXCLUDED** list. If you want to include the test for one/more of the components chosen for exclusion, then, pick the components from the **COMPONENTS TO BE EXCLUDED** list and click the **>>** button. Finally, click the **Apply** button to save the changes. You will then be lead to Figure 7.114, where you can view the test that you just excluded in the **EXCLUDED TESTS** list.

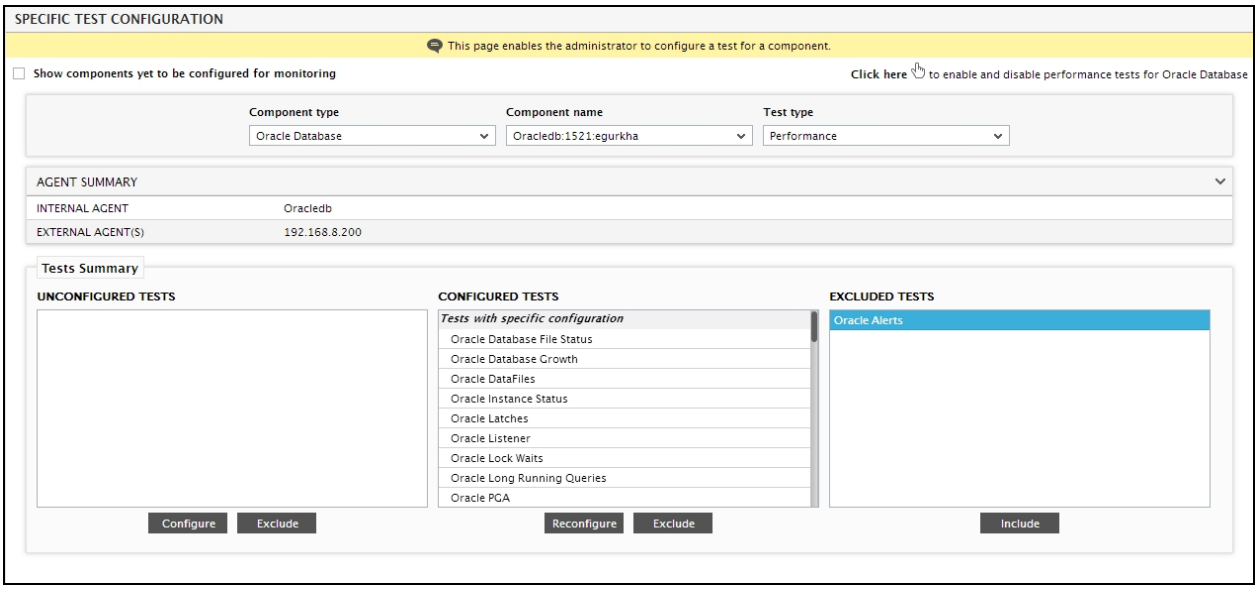

Figure 7.114: Selecting the OracleExtents to be included for its execution

12. At any given point in time, you can make sure that an excluded test starts executing on a component, by including that test. For that, select the test from the **EXCLUDED TESTS** list, and click the **Include** button. If the chosen test was previously excluded for more than one component of the chosen type, then, clicking the **Include** button will lead you to Figure 7.115, using which you can include the chosen test for all such components at one shot. While the **COMPONENTS TO BE INCLUDED** list will display that component for which the test has just now been chosen for inclusion, the **EXCLUDED COMPONENTS** list will list all those components (of the chosen type) for which the test was previously excluded. The components in the **EXCLUDED COMPONENTS** list will be grouped according to the state in which the test (chosen for exclusion) exists with respect to a component. For instance, if the test chosen for exclusion is in the **UNCONFIGURED STATE** for a few components, then such components will be listed under the group **<TestName> - Unconfigured component(s)** in the **EXCLUDED COMPONENTS** list. To include the test for one/more excluded components, select the components of interest from the **EXCLUDED COMPONENTS** list and click the **<<** button. To make sure that the test stays excluded for any component, pick the component from the **COMPONENTS TO BE INCLUDED** list and click the **>>** button. Finally, click the **Apply** button to save the changes.

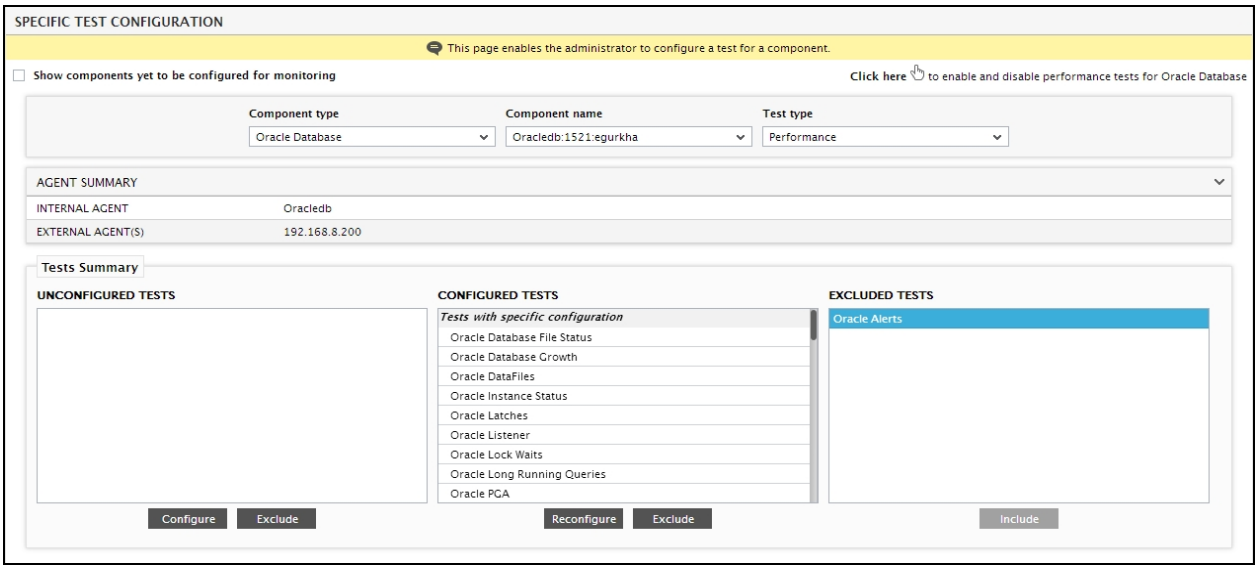

Figure 7.115: Selecting the components for which the chosen test is to be excluded

Note:

- If a specific test configuration has been applied to a component, the default parameter values will no longer apply. The user has to change the parameters manually. By excluding and then including the component, the user can re-ensure that the default parameters apply for a component.
- Normally, a few/all the tests associated with a component share the same set of parameters. To ensure that administrators do not perform the redundant task of configuring each such test separately, you can configure the eG Enterprise system in such a way that if one of these tests is configured, the other related tests get configured automatically. For instance, take the case of the FnSessionTest and the FnHaStatusTest that the eG external agent executes on a Fortigate Firewall. These two tests support the same parameters as the FnSystemTest. To ensure that these two tests are automatically configured, as soon as the FnSystemTest is configured, do the following:

  In order to do this:

  - Navigate to the directory **<EG_INSTALL_DIR>\manager\config** directory and edit the **eg_testparam.ini** file.
  - Go to the **[TEST_RELATIONS]** section, where entries for tests that are related to one another pre-exist.
  - To set a relationship between the Fortigate Firewall tests in our example, append the following entry to the **[TEST_RELATIONS]** section:

    FnSystemTest=FnSessionTest, FnHaStatusTest

    This ensures that, once you configure the FnSystemTest in the administrative console, values provided for the FnSystemTest parameters, are automatically applied to the FnSessionTest and FnHaStatusTest, configuring those two tests as well. However, such a test relationship will not affect the **TESTPERIOD** parameter. If for some reason, different tests executing on a component are set to execute at a different **TESTPERIOD**, then, even if an entry exists for such tests in the **[TEST_RELATIONS]** section, the **TESTPERIOD** for the tests will continue to be different. On the contrary, if

you also want the **TESTPERIOD** of one test to be automatically applied to the other tests executing on the same component, you will have to follow the steps given below:

- Edit the **eg_testparam.ini** file.
- In the **[COMMON_TEST_RELATIONS]** section of the file, provide an entry for all the tests that will share the same set of parameters,including the **TESTPERIOD**.This specification should be in the following format.

  <Comma-separated list of related tests>=None

  For instance, for the Fortigate Firewall example above, the entry would be:

  FnSessionTest,FnSystemTest,FnHaStatusTest=None

  This means that if any one of the three tests specified is configured, then all the parameters (including **TESTPERIOD**) of the configured test will be automatically applied to the other two tests, thereby configuring those two tests also.

7.8.3 Enabling / Disabling Tests

To enable or disable a test, select the **Enable/Disable** option from the **Tests** menu of the **Agents** tile**.** From Figure 7.116 that appears, an administrator has to first choose the type of component (from the **Component type** list box) and the type of test (from the **Test type** list box) for which the tests have to be either enabled or disabled. Remember that the **Component type** list box lists only those component-types that have been managed by eG Enterprise. Ext, pick the **Test type**. By default, **Performance** will be displayed as the **Test type**. Where the **Configuration Management** capability has been enabled, **Configuration** will also be displayed as the **Test type**.

While all the tests that have been enabled - i.e., tests that will execute - for the chosen **Component type** will be listed in the **ENABLED TESTS** list of Figure 7.116, those that are disabled - i.e., those that will not execute for the chosen **Component type** - will be listed in the **DISABLED TESTS** list.

To enable one/more disabled tests, select the tests from the **DISABLED TESTS** list box and click on the **<** button. This will transfer the selection to the **ENABLED TESTS** list, as shown by Figure 7.117.

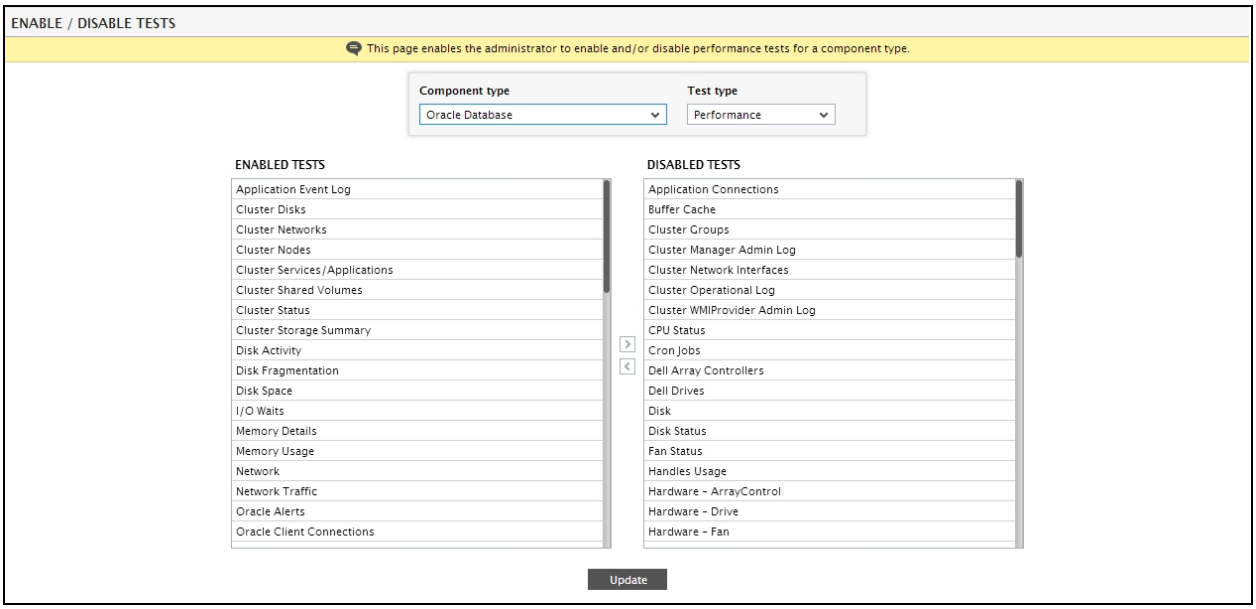

Figure 7.116: Selecting the tests to be enabled for Oracle Database component type

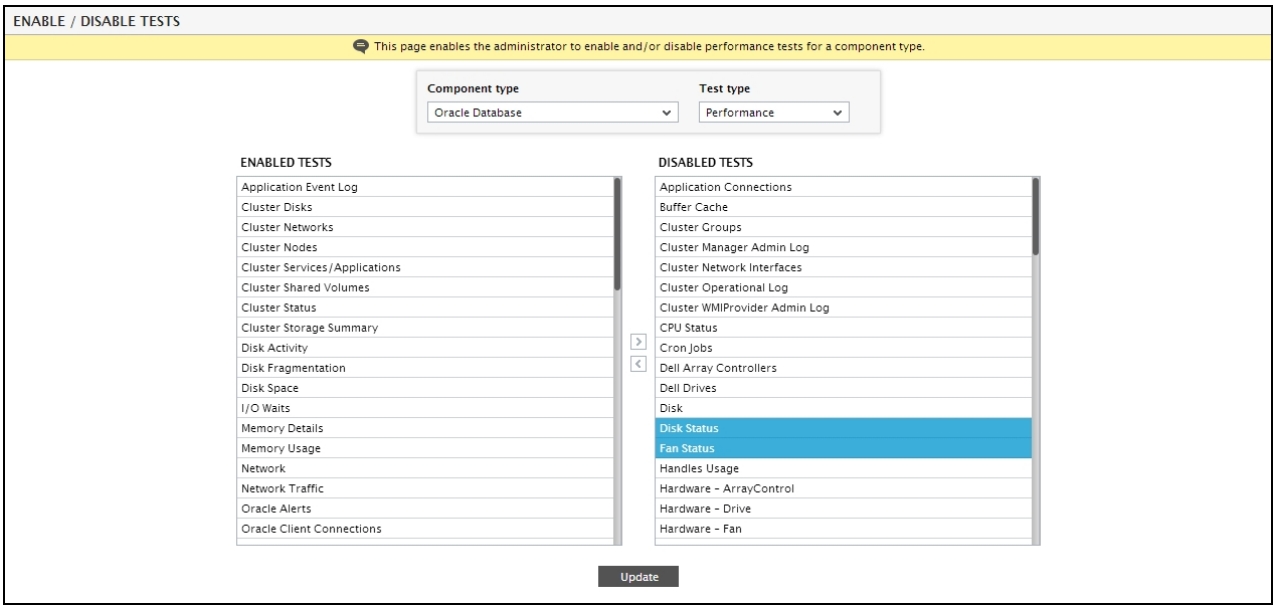

Figure 7.117: Enabling the tests for the Oracle database component type

Finally, click the **Update** button to save the changes.

To disable one/more tests, pick the tests from the **ENABLED TESTS** list, click the **>** button in Figure 7.117, and click the **Update** button.

Note:

Tests that are disabled for a component-type will not be available for configuration in the **DEFAULT CONFIGURATION** and the **SPECIFIC CONFIGURATION** pages.

7.8.4 Enable / Disable Detailed Diagnosis

The detailed diagnosis capability, if enabled, allows eG agents to generate more detailed measurements periodically. Before enabling this capability for a specific component, you must ensure that the corresponding component-type is configured with this capability; **note that if this capability is disabled for a component-type, then the 'SPECIFIC TEST CONFIGURATION' page of any component of that type will not support the DETAILED DIAGNOSIS parameter.**

To enable/disable the detailed diagnosis capability for a component-type, do the following:

1. Select the **Enable/Disable DD** option from the **Tests** menu of the **Agents** tile.
2. Figure 7.118 will then appear. The **DD ENABLED TESTS** list of Figure 7.118 will list all those tests for which detailed diagnosis has already been enabled. The **DD DISABLED TESTS** list on the other hand, will display tests for which detailed diagnosis is disabled.
3. To enable DD for a test, pick the test from the **DD DISABLED TESTS** list, click the **<** button, and then click the **Update** button. This will transfer the selection to the **DD ENABLED TESTS** list.

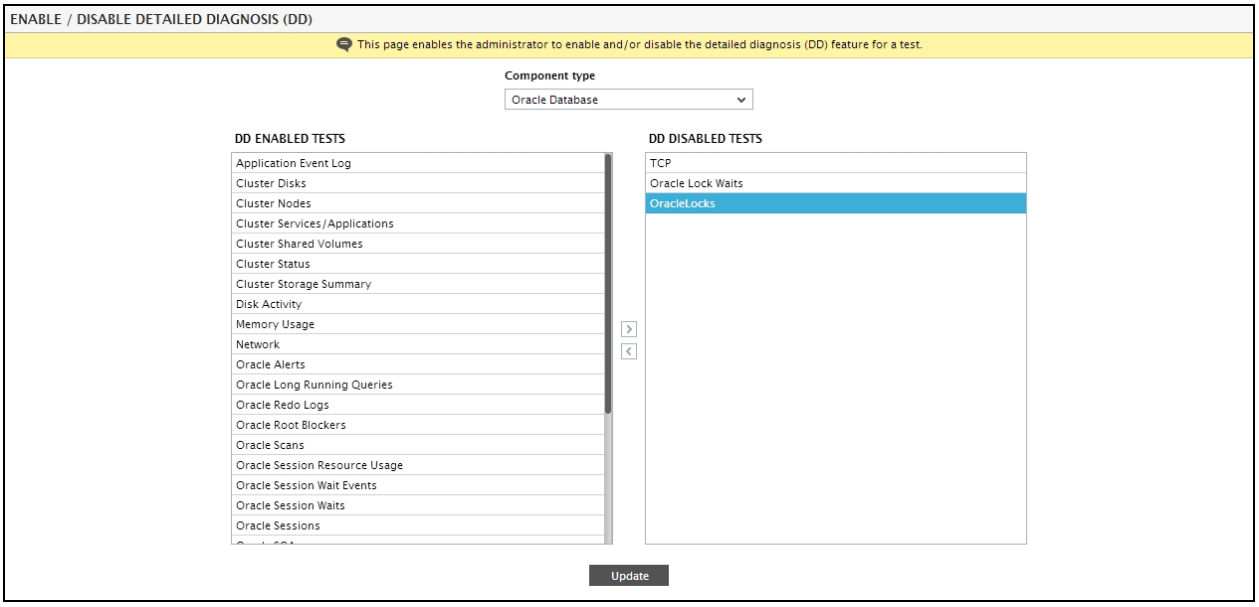

Figure 7.118: Selecting the test for which DD is to be enabled

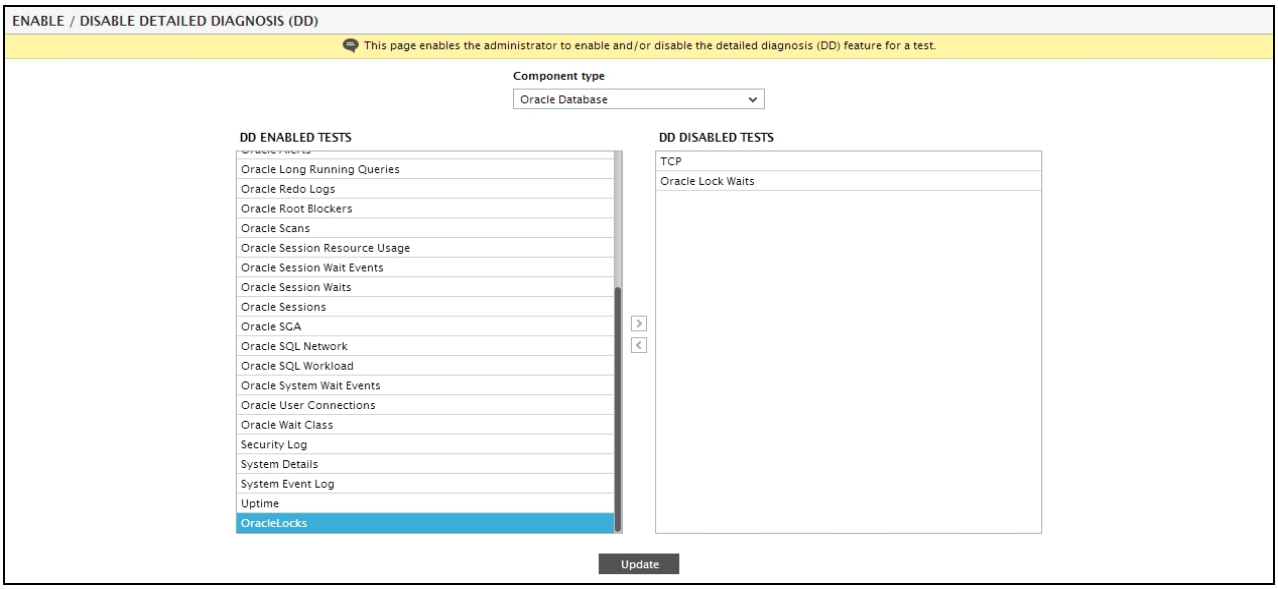

Figure 7.119: Enabling DD for a test

4. Similarly, to disable DD for a test, pick the test from the **DD ENABLED TESTS** list, click the <<button to transfer the selection to the **DD DISABLED TESTS** list, and click the **Update** button.

   Note:

   The option to selectively enable/disable the detailed diagnosis capability will be available, only if both the normal and abnormalfrequencies of the detailed diagnosis have not been configured as 0. If so, then the detailed diagnosis capability will be automatically switched off. In such a case, the **Enable / Disable DD** option will not be available in the **Tests** menu of the **Agents** tile.

7.8.5 Including / Excluding a Test for Multiple Components

The **Enable / Disable Tests** option in the **Tests** sub-menu of the **Agents** menu enables administrators to enable / disable specific tests for a component-type as a whole - tests so disabled will be unavailable for all components of the chosen type. Sometimes however, administrators may want a test to execute for only a few components of a type, and disable the same test for a few other components of that type. The **Include / Exclude Components** option in the **Tests** menu of the **Agents** tile facilitates such selective test execution. When this option is chosen, Figure 7.120 will appear. To exclude a test for specific components of a type, do the following using this page:

- Select a **Component type** from Figure 7.120.
- Pick a **Test type** (whether **Performance** or **Configuration**).
- The **Test name** list will then be populated with all tests (of the chosen **Test type**) that are associated with the chosen **Component type**. From this list, choose the tests to be excluded / included.

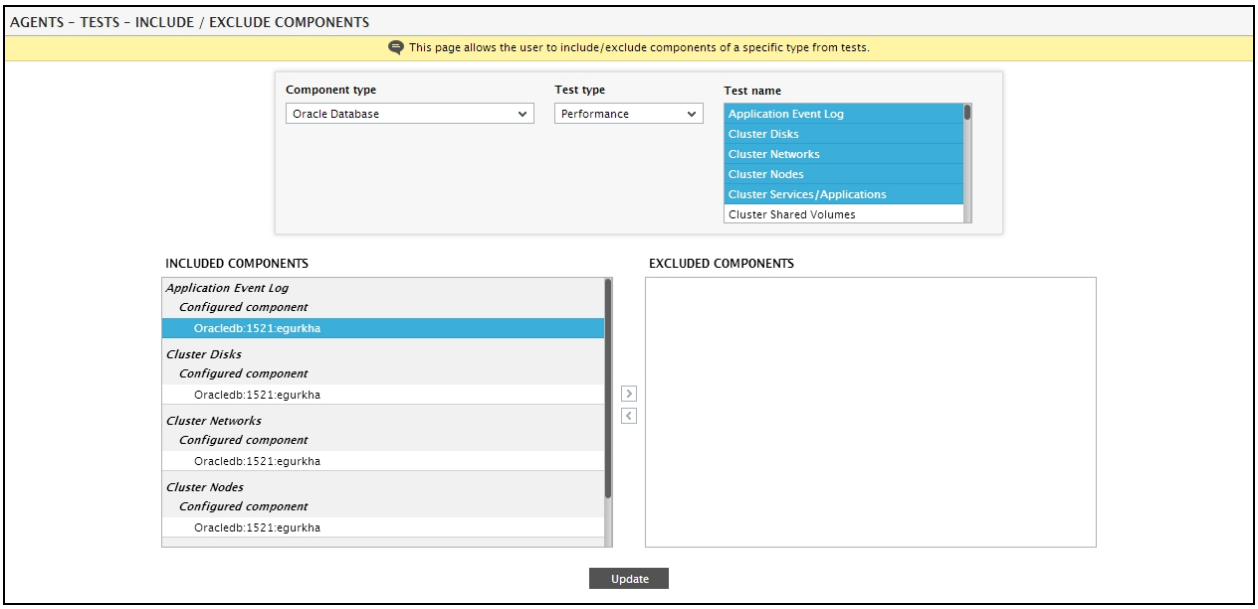

Figure 7.120: Selecting the components for which tests are to be excluded

- All managed components of the chosen type for which the selected tests are currently enabled will then appear in the **INCLUDED COMPONENTS** list. Besides the test name, the **INCLUDED COMPONENTS** list will also indicate the state of the test with respect to each component - for instance, if one of the chosen tests has already been configured for one/more components of the selected type, then such components will appear under the **Configured components** section in the **INCLUDED COMPONENTS** list. If the test is to be disabled / excluded for any of the components in the **INCLUDED COMPONENTS** list, then, select the components to be excluded as depicted by Figure 7.120. Then, click the **>** button.
- This will transfer the selection to the **EXCLUDED COMPONENTS** list (see Figure 7.121). Finally, click the **Update** button. This will ensure that the chosen test no longer executes for the components in the **EXCLUDED COMPONENTS** list.

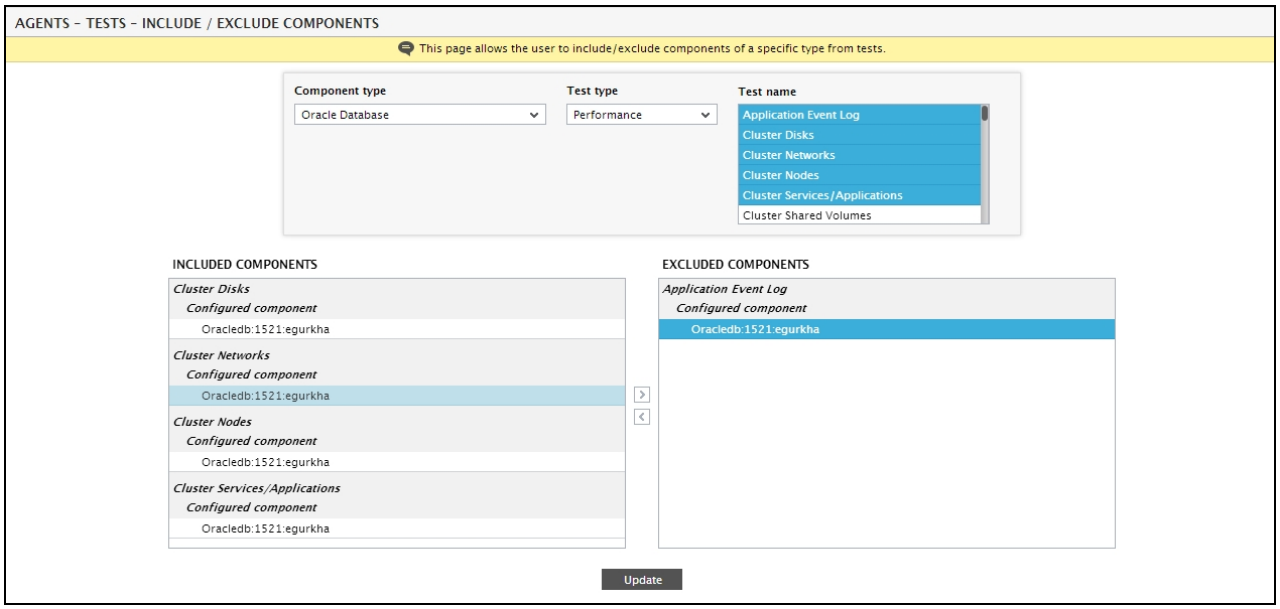

Figure 7.121: Excluding a test for selected components

- On the other hand, if you want to enable the execution of a test for one/more components of a chosen type, then, pick the components to be included from the **EXCLUDED COMPONENTS** list, and click the **<** button. Finally, click the **Update** button.

7.8.6 Configuring Tests for an Oracle Database Server

As already mentioned, eG agents can collect a variety of statistics regarding Oracle database servers. To monitor an Oracle database server, eG Enterprise requires that a special database user account be created in each Oracle database instance to be monitored. Figure 7.122 depicts the parameters to be configured for the Oracle Database File Status test. While configuring the Oracle tests, a **Configure a database user** button ( ) is provided alongside the **USER** parameter.

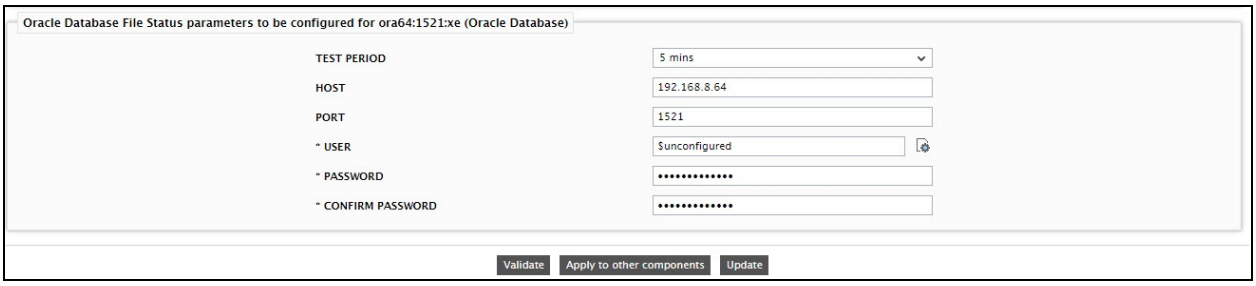

Figure 7.122: Parameters to be configured for Oracle Database File Status test

Clicking on the button will take you to the page shown by Figure 7.123. In Figure 7.123, the host name, port number and the side of the Oracle server to be configured appear in the corresponding location and cannot be changed. So, proceed to select a **DB VERSION**. **`If the user chooses to monitor Oracle 7.x, it is mandatory to provide the password of the SYS user in the`** SYS PASSWORD **`text box.`** To add a new user, specify the database administrator's login and password. While creating a new user, the account has to be associated with a default tablespace where the user's data is hosted and a temporary tablespace which is used for buffering, sorting etc. The identities of the default and temporary tablespaces have to be input in Figure 7.123.

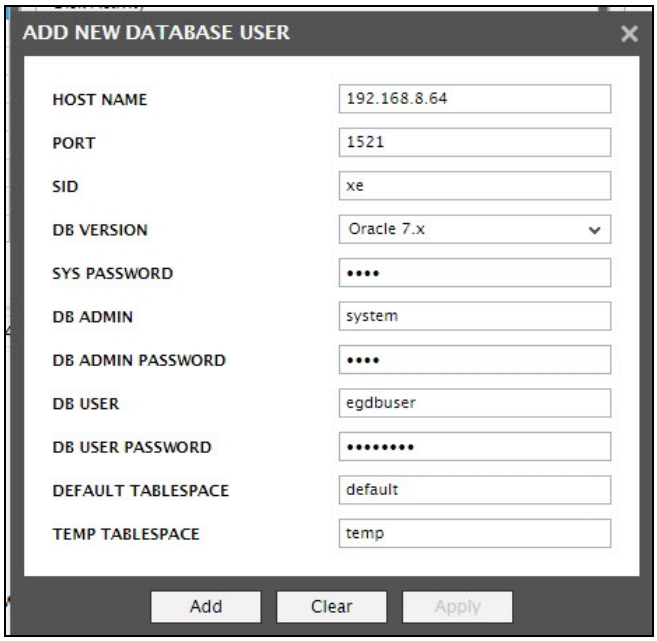

Figure 7.123: Configuring a database user for an Oracle database server v7.x

If you choose **Oracle 8.x - 11g** as the **DB VERSION** (see Figure 7.124, it implies that the Oracle database server being monitored is of a version between 8.x and 11g. In this case, you do not need a **SYS PASSWORD** for creating a new database user. Just provide the **DB ADMIN** and **DB PASSWORD** and then proceed to specify the credentials of the new database user against the **DB USER** and **DB USER PASSWORD** text boxes. Then, as before, specify the identities of the default and temporary tablespaces and click **Add** to add the new database user.

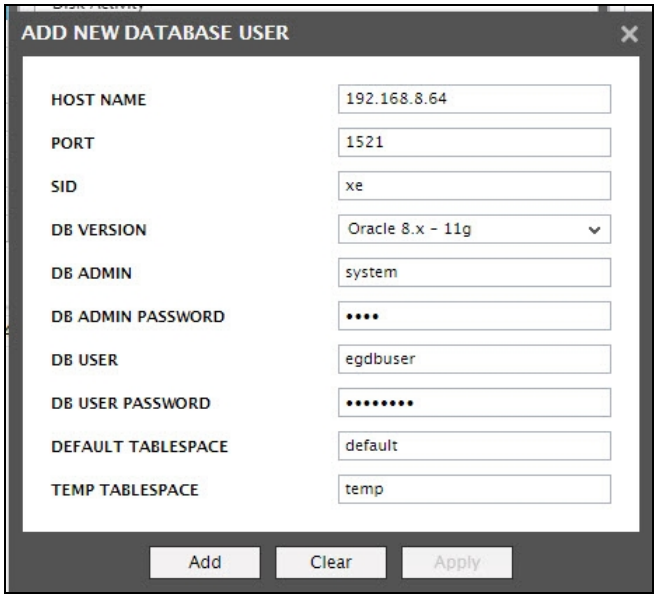

Figure 7.124: Configuring a database user for an Oracle database server v8.x to 11g

Choosing **Oracle 12c** as the **DB VERSION** requires that you pick the type of **DB CONNECTION** (see 7.8.6). The **DB CONNECTION** can either be **PDB** or **CDB**, depending upon the archicture of the Oracle 12c server. On the surface, a CDB or a Container database, may seem very similar to a conventional Oracle database, as it contains most of the working parts you will be already familiar with (controlfiles, datafiles, undo, tempfiles, redo logs etc.). It also houses the data dictionary for those objects that are owned by the root container and those that are visible to all PDBs.

Since the CDB contains most of the working parts for the database, the PDB or the Pluggable Database only needs to contain information specific to itself. It does not need to worry about controlfiles, redo logs and undo etc. Instead it is just made up of datafiles and tempfiles to handle it's own objects. This includes it's own data dictionary, containing information about only those objects that are specific to the PDB.

Once the **DB CONNECTION** is set, just provide the **DB ADMIN** and **DB PASSWORD** and then proceed to specify the credentials of the new database user against the **DB USER** and **DB USER PASSWORD** text boxes. Then, as before, specify the identities of the default and temporary tablespaces and click **Add** to add the new database user.

Figure 7.125: Configuring a database user for Oracle database server v12c

7.8.7 Associating/Disassociating Tests

eG Enterprise provides a specialized monitoring model for every component type it monitors out-of-the-box. The tests to be run on a component, the measures to be reported by each test, and the thresholds governing the state of every measure, are pre-defined in the monitoring model of that component type. If this built-in monitoring template needs to be modified to include new tests or exclude existing tests, then previously, administrators had to obtain an **Integration Console** license, and use the **Integration Console** module it enables to effect the change. But not any more! eG Enterprise now provides a specialized interface using which the same can be achieved outside of the **Integration Console** itself! This easy-to-use, select-and-click interface helps administrators associate any pre-defined test with any monitoring model supported out-of-the-box by eG Enterprise. In the same way, administrators can delink any test from any built-in monitoring model. Because of this capability, administrators no longer have to obtain the **Integration Console** license, just for associating/disassociating tests from components.

To associate/disassociate tests from a component, do the following:

1. Select the **Associate/Disassociate** option from the **Tests** menu of the **Agents** tile.
2. Figure 7.126 will then appear. From the **Component type** drop-down in Figure 7.126, select the component type for which tests are to be associated/disassociated.

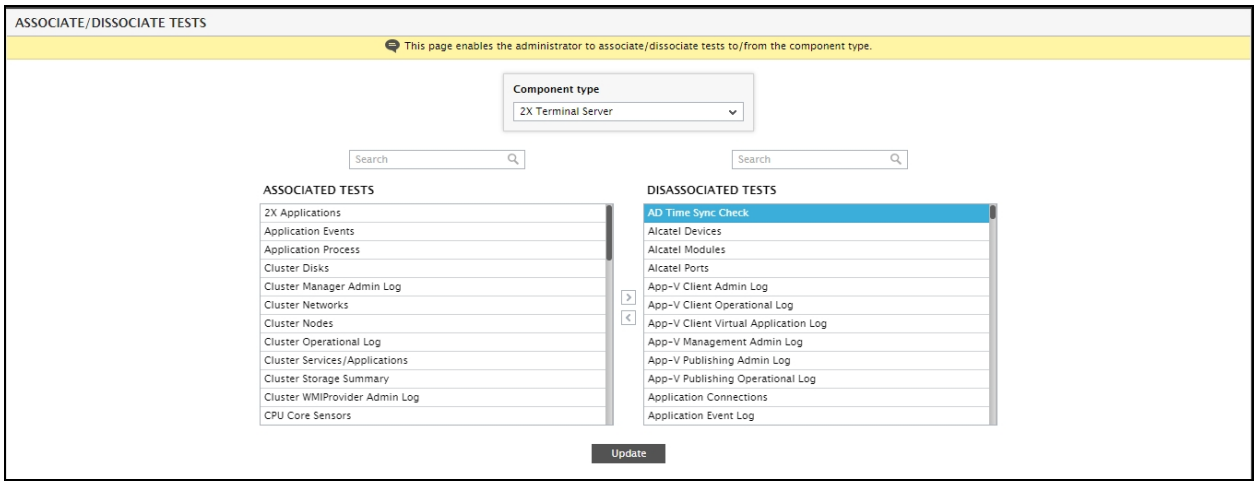

Figure 7.126: Selecting the test to be disassociated from a component type

3. To include additional tests for the chosen component type, first select the test(s) to be associated with the component type from the **DISASSOCIATED TESTS** list. Then, click the **<** button in Figure 7.126.
4. This will transfer the selection to the **ASSOCIATED TESTS** list, as shown by Figure 7.127.

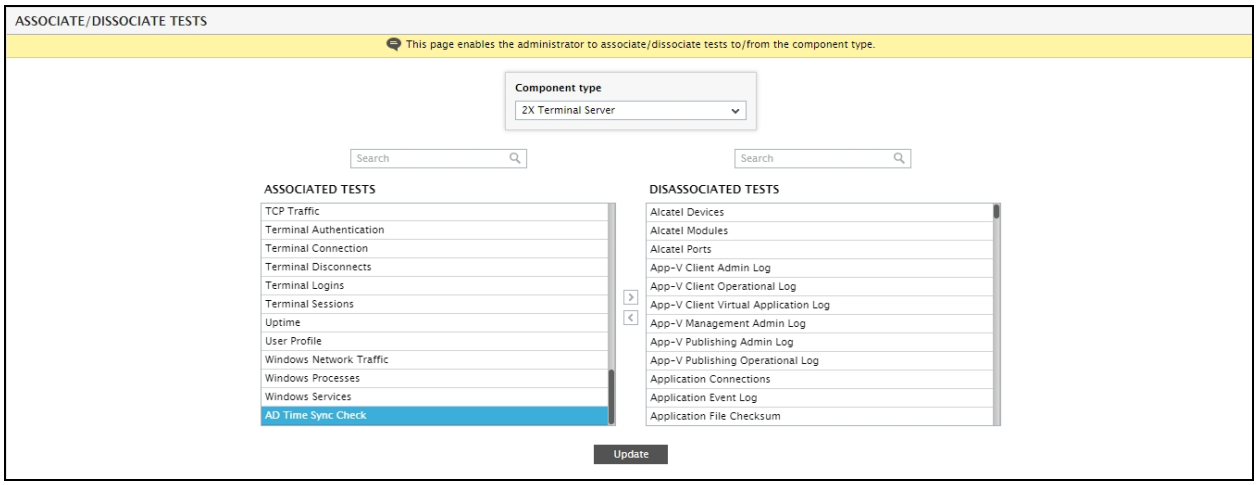

Figure 7.127: Associating the chosen test with the selected component type

5. Finally, click the **Update** button in Figure 7.127 to save the changes.
6. Likewise, to disassociate a test from the chosen component type, select the test from the **ASSOCIATED TESTS** list of Figure 7.127, click the **>** button, and then click **Update**.

7.8.8 Viewing Tests

In environments where a large number of components are being monitored, naturally, a large number of tests would have been configured. It is therefore very difficult for administrators of such environments to verify all test configurations pertaining to a single component/component-type. In order to ease their pain, eG Enterprise provides a single interface for reviewing the test configurations pertaining to a chosen component-type or component. To access this interface, select the **View** option from the **Tests** menu of the **Agents** tile. Doing so will invoke Figure 7.128.

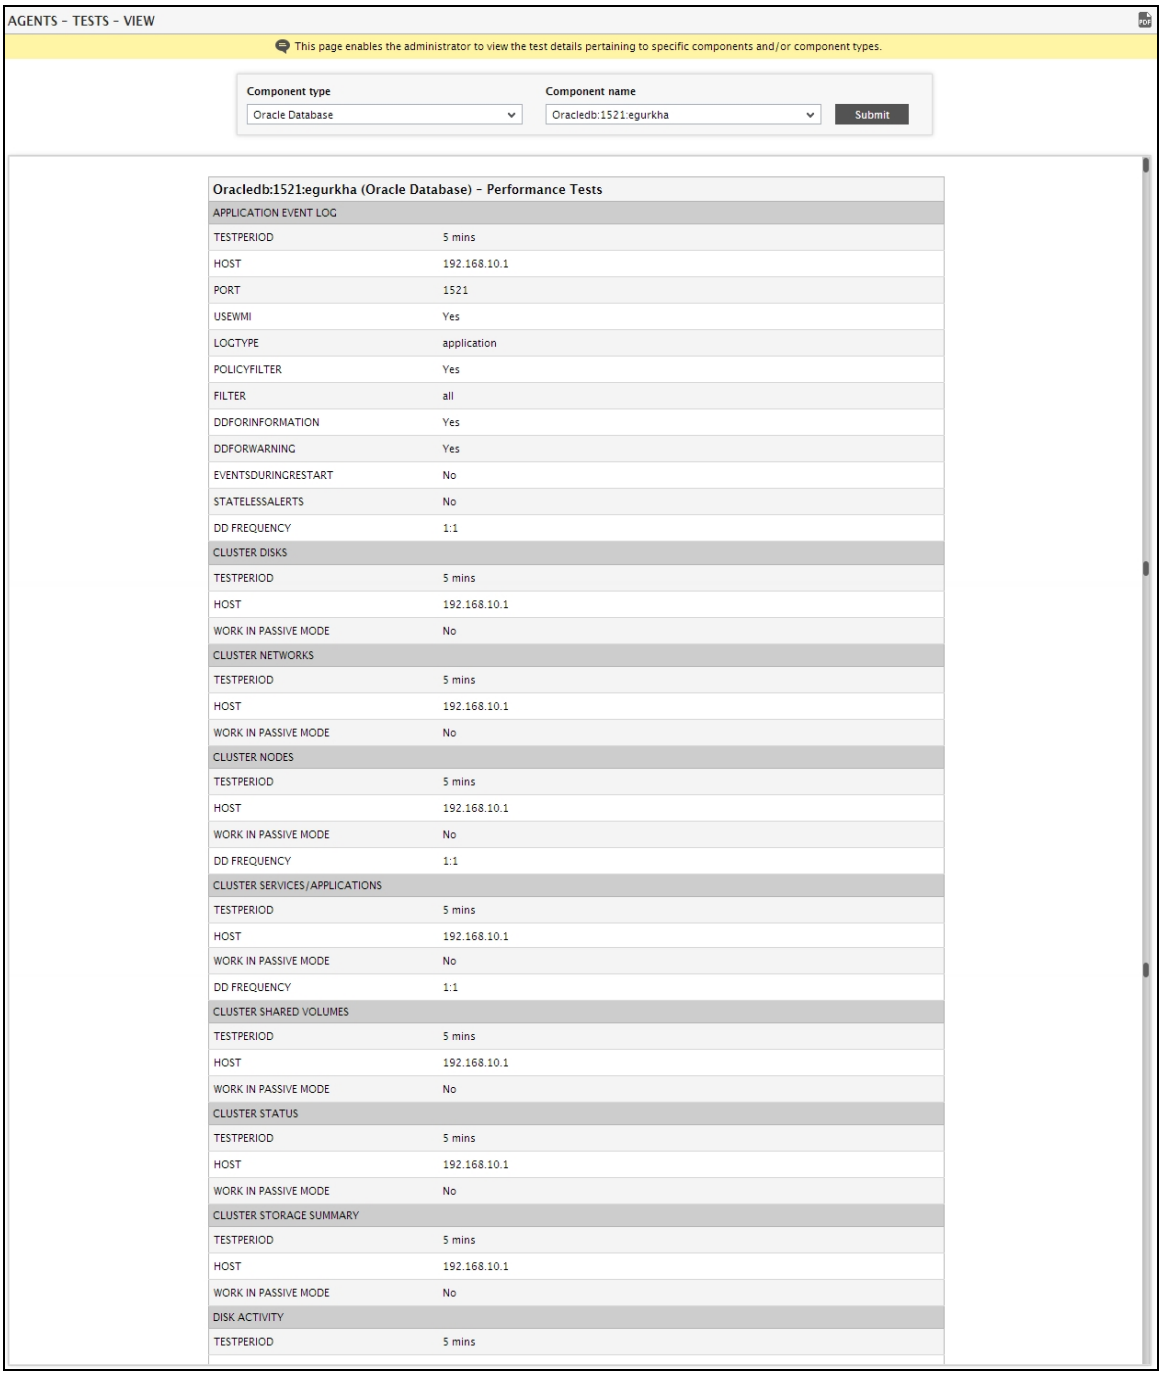

Figure 7.128: Viewing the test configurations pertaining to all managed components

By default, the **All** option will be selected in the **Component type** and **Component** lists, thus displaying the test configurations for every monitored component. To view the test configurations of all components of a particular type, select a **Component type** from the list and click the **Submit** button. Similarly, to view the test configurations of a specific component, select the corresponding **Component** from this page and click the **Submit** button. A **Print** icon is also made available in this page, to enable you to print and file the test configurations for future reference.

7.8.9 Configuring Test Descriptors

As already indicated, in the eG terminology, a descriptor represents a unique attribute that identifies a set of measurements returned by a test. For example, the different disk partitions on a system are different descriptors of the DiskSpaceTest. Likewise, the different tablespaces used by an Oracle database instance are descriptors of the OraTablespacesTest.

The eG manager provides administrators the capability to selectively turn on or off specific test descriptors, thereby controlling the operation of individual tests. When a test descriptor is disabled, the eG manager instructs the agent to stop reporting any further measurements for this test descriptor. All the prior measurements stored in the eG database for this test descriptor are also removed at this time. Using this page, at any later time, the administrator can choose to enable data collection for a descriptor that has been disabled earlier.

To enable or disable specific test descriptors, first choose the **Enable/Disable Descriptors** option from the **Tests** menu of the **Agents** tile. Figure 7.129 then appears. From the **Component type** list in Figure 7.129, select the type of component for which descriptors are to be enabled/disabled. Only those component-types that have been managed by eG Enterprise will be available in this list box. Once the component type is chosen, the **Component name** list box is populated with the list of components belonging to the chosen component type. Only the currently monitored components are shown in this list box. The administrator has to next choose the component for which a specific descriptor has to be enabled or disabled. The list of tests that are running against the selected component then populates the **Test** list. Select the test for which descriptors are to be enabled/disabled.

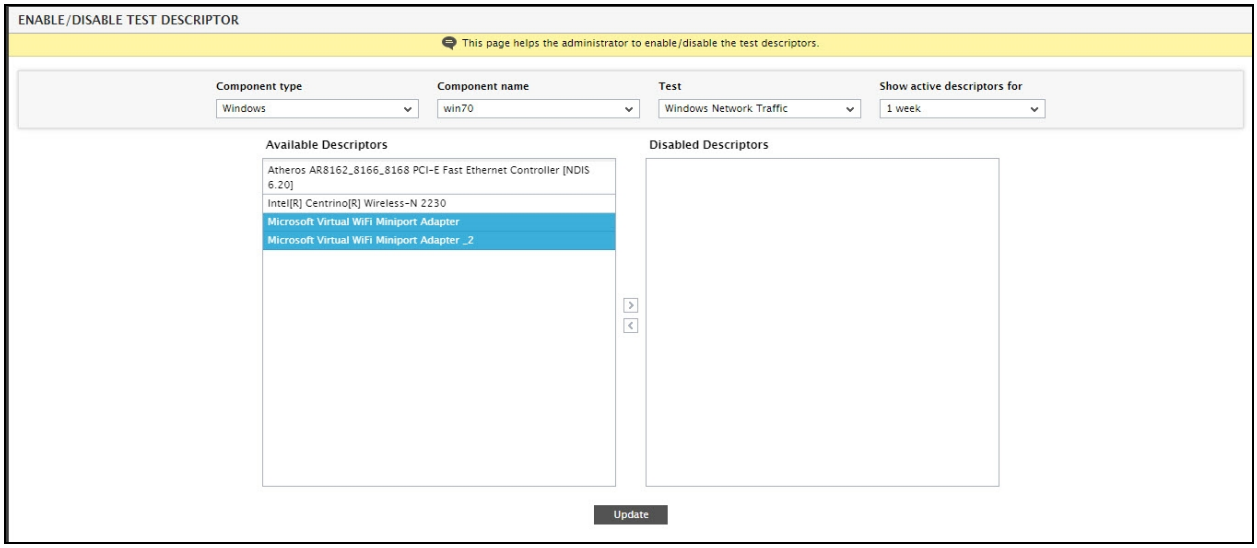

Figure 7.129: Selecting the descriptors to be disabled

You can further filter your descriptor-list by choosing to view only those descriptors that were active during a selected time period. For this, select a time period from the **Show active descriptors for** drop-down. By default, the **1 week** option will be chosen from this drop-down list. If you select the **Any** option, then all the descriptors supported by the chosen **Test**, regardless of how long they were active, will be made available for enabling/disabling.

The **Available Descriptors** list will then list all the active descriptors that fulfill the filter conditions described above. From the **Available Descriptors** list, select the descriptors to be enabled and click the **>** button in Figure 7.129. This will transfer the selection to the **Enabled Descriptors** list (see Figure 7.130).

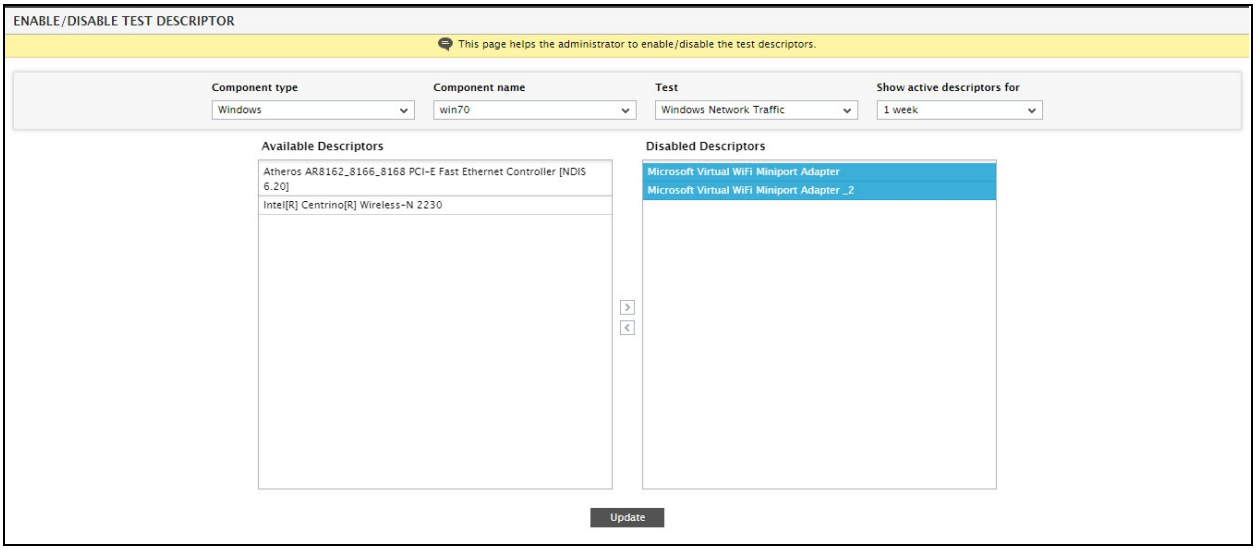

Figure 7.130: Disabling chosen descriptors

To disable one/more enabled descriptors, select them from the **Enabled Descriptors** list and click the **<** button in Figure 7.130. Finally, click the **Update** button.

8

Performance Rating Tests

Typically, in order to know the overall performance of a server, IT administrators should keep track of all metrics reported by eG Enterprise for that server. An executive on the other hand, may want to quickly determine server status, without having to look through hundreds of metrics. This is where performance ratings help.

Administrators can define a performance rating as an aggregate metric that is based on a number of other metrics collected and reported by an eG Enterprise agent. Consider this similar to APDEX rating or Customer Satisfaction Index (CSI). For example, the user experience of a Citrix user can be defined based on the logon time of the user, the screen refresh latency of the user, the profile size of that user and the network latency seen by that user. By looking at a single metric that takes a percentage value between 0 and 100, an executive can easily determine if that user is happy or dissatisfied. Likewise, a stress rating for a server can be based on its CPU utilization, memory utilization, disk space available and disk activity level.

With a Performance Rating, administrators can:

- quantify the overall performance of an entity using a sub-set of metrics that eG Enterprise reports for that entity
- ascertain, at a glance, whether performance of that entity is within desired levels or not
- receive a single alert when the performance rating dips, instead of a flood of alarms, thereby enabling you to rapidly detect performance degradations
- instantly pinpoint which parameter/measure caused the overall performance of the entity to slide
- highlight the key performance indicator of a server, user, service on dashboards

There are 3 out-of-the-box performance rating tests available:

- Citrix User Experience Rating
- Citrix XenApp User Experience Rating
- VDI User Experience Rating

These out-of-the-box tests do not require any additional license. You may just want to enable the test from the **ENABLE/DISABLE TESTS** lists provided you have managed atleast *one* Citrix/VDI server in your infrastructure. For example, if you have managed the *Citrix XenApp 7.x* server, then you can enable the *Citrix User Experience Rating* test from the **DISABLED TESTS** lists (see Figure 8.1)

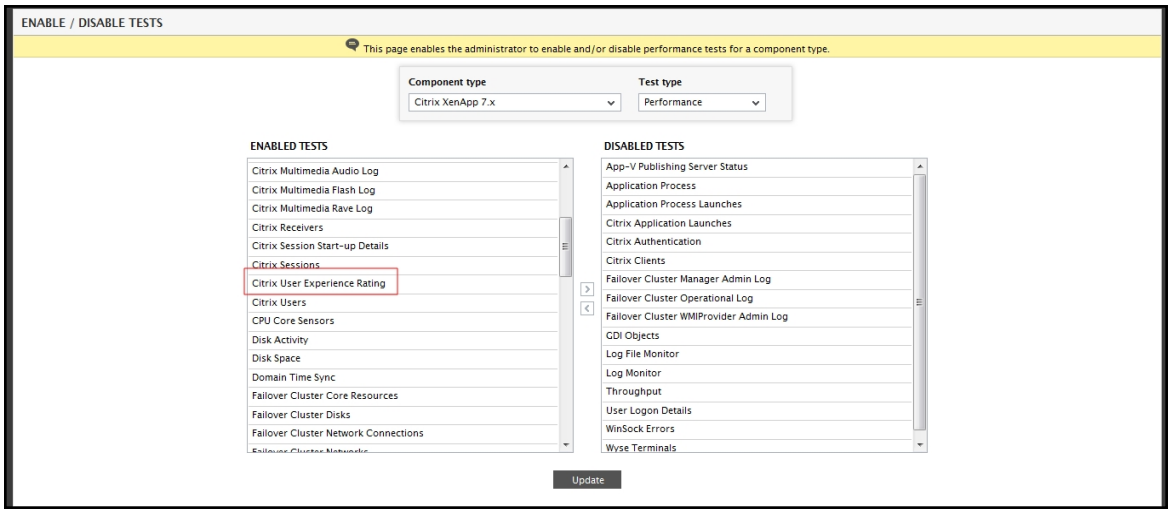

Figure 8.1: Enabling a pre-configured Performance Rating test

Alternately, you can add a new **Performance Rating** test to suit your requirements. In order to avail such flexibility, the *Metric Aggregation* capability should be enabled in your eG manager license (see Figure 8.2).

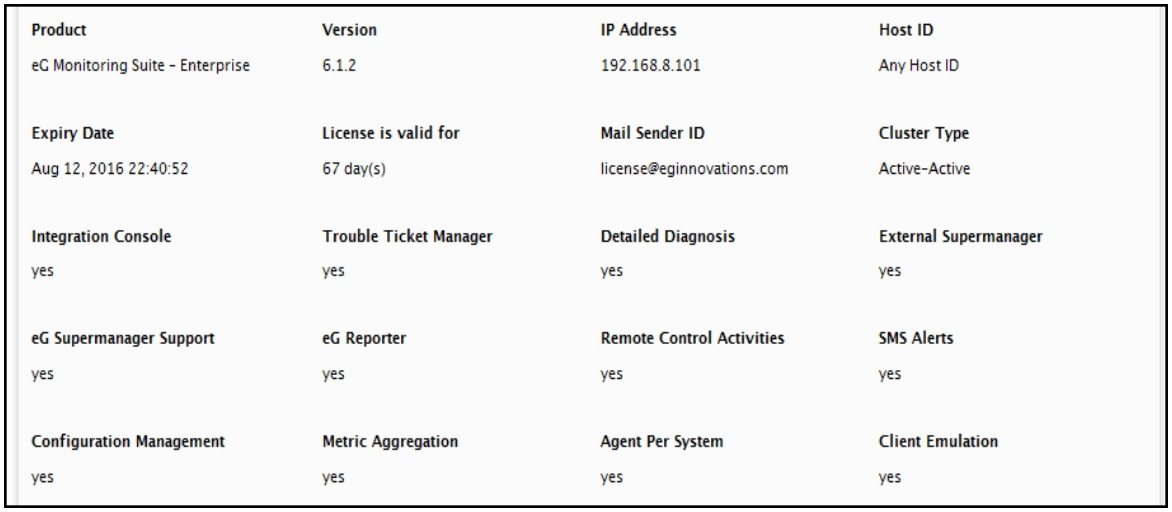

Figure 8.2: Enabling the Metric Aggregation capability

Once the *Metric Aggregation* capability is enabled, the *Performance Rating* capability will be available in the eG Admin user interface as shown in Figure 8.3.

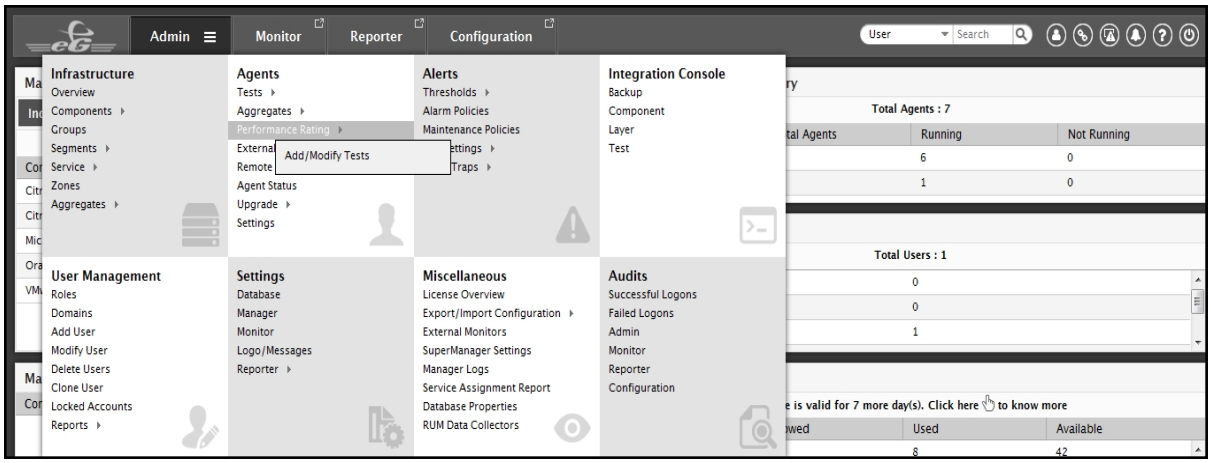

Figure 8.3: The Performance Rating capability in the eG administrative interface

Note:

A separate eG agent is *not* required for collecting metrics of the **Performance Rating** tests. The same eG agent (internal or remote) that monitors the component to which the **Performance Rating** test is associated, can be used.

8.1 Adding a new Performance Rating Test

To ascertain the overall performance of a disk on a Microsoft Windows server, a help desk executive has to traverse through the metrics of the Disk Activity test and the Disk Space test. The help desk executive finds it difficult to view each and every measure of both these tests separately. Let us now help him/her in adding a new Performance Rating test so that he/she can figure out the overall performance of the disk at a single glance. To add a new Performance Rating test, follow the steps mentioned below:

1. Follow the menu sequence: *Admin -> Performance Rating -> Add/Modify Tests.* Figure 8.4 will then appear stating that there are no user-defined Performance Rating test in your environment.

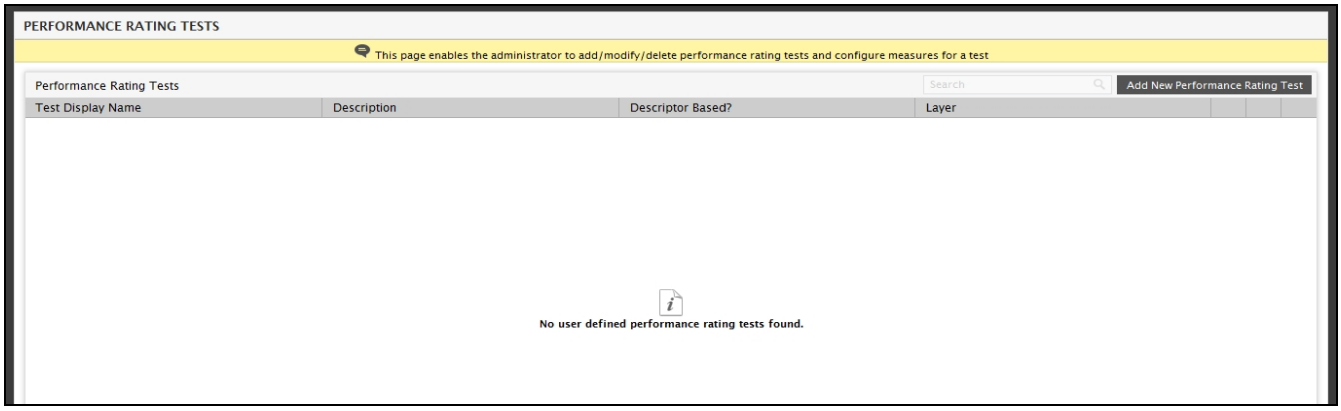

Figure 8.4: A message stating no user defined Performance Rating tests are found in your environment

2. By default, if you have managed Citrix/VDI servers in your environment and enabled the user-related tests, then, the Performance Rating tests that are pre-configured and bundled with the eG Enterprise Suite

will automatically appear in the **PERFORMANCE RATING TESTS** page as shown in Figure 8.5.

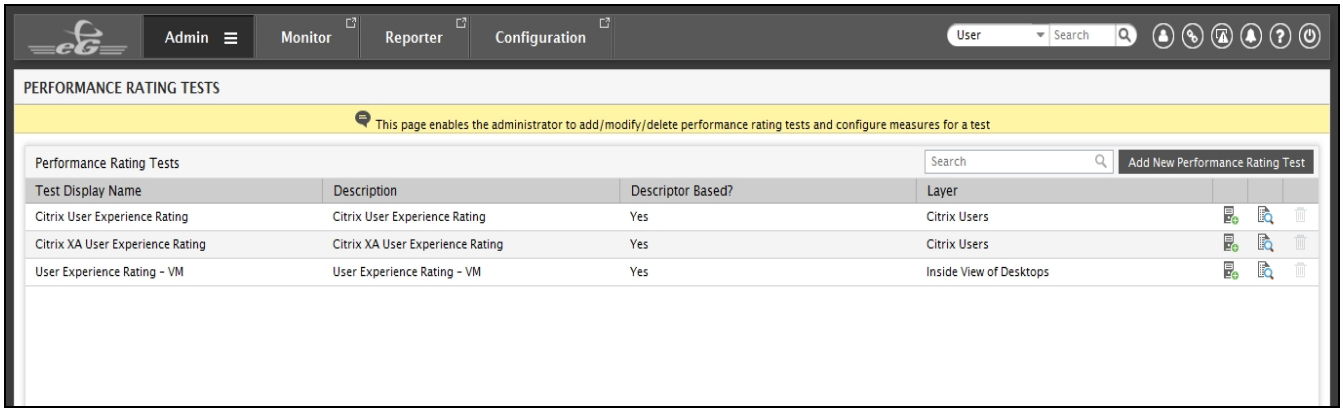

Figure 8.5: The list of pre-configured Performance Rating tests

3. If you wish to view the component types to which a pre-configured test is associated with, then click the icon against a chosen Performance Rating test (see Figure 8.5). Figure 8.6 will then appear listing all the component types to which the test is associated with.

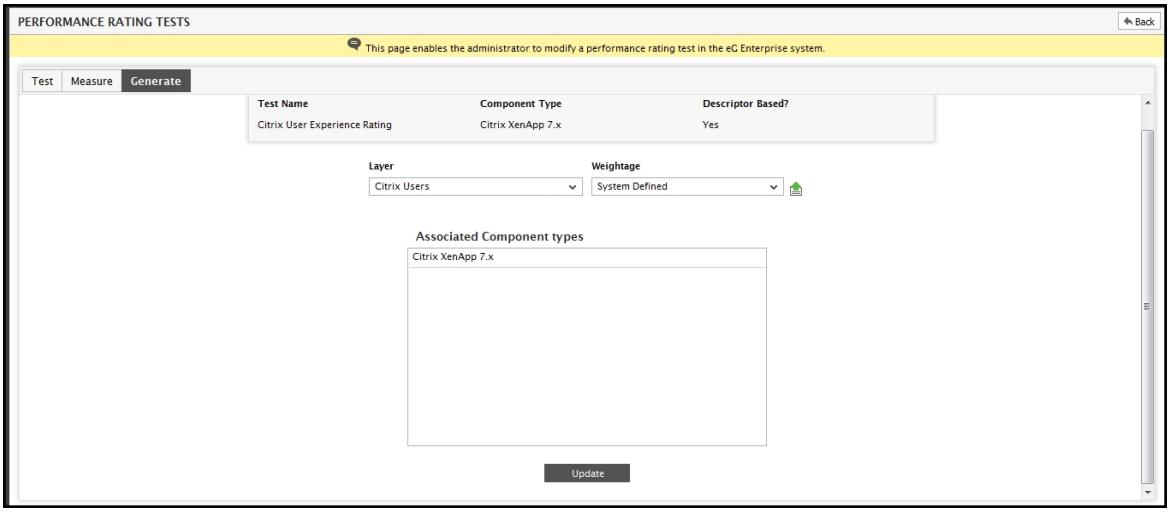

Figure 8.6: The component types to which a pre-configured Performance Rating test is associated with

Note:

- You cannot modify a pre-configured Performance Rating test. Instead, you will be allowed to view the configuration of the per-configured test by clicking the icon (see Figure 8.5) against each pre-configured test.
- You cannot delete a pre-configured Performance Rating test.

4. To add a new Performance Rating test, click the **Add New Performance Rating Test** button in Figure 8.4.
5. Figure 8.7 will then appear.

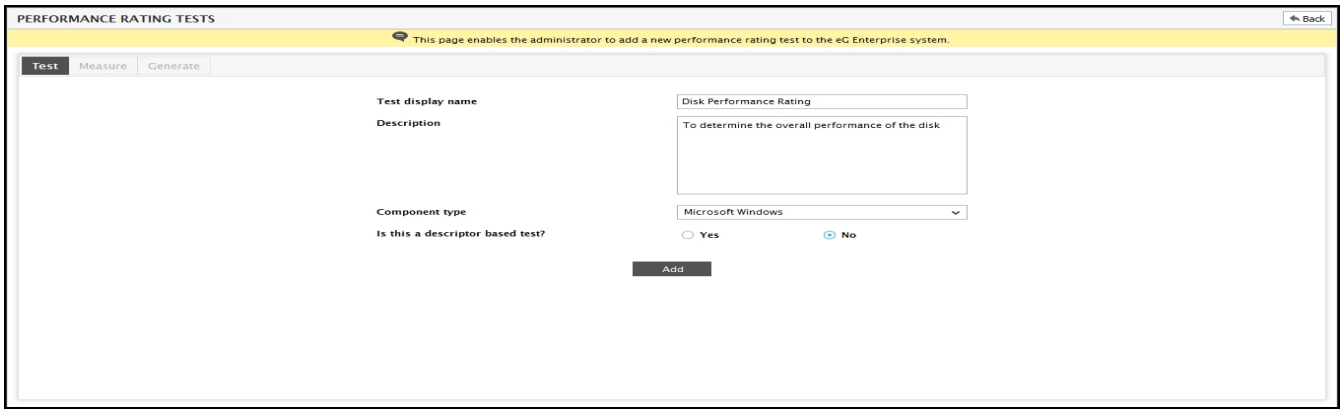

Figure 8.7: Adding a new Performance Rating test

6. Here, specify the name of the Performance Rating test that you wish to add in the **Test display name** text box.
7. Once the test name is specified, you need to provide a description of the test in the **Description** text box.
8. Select the component type to which you wish to associate your test from the **Component type** list.

   Note:

   The **Component Type** list will be populated only with the component types that are managed in your infrastructure.
9. Next, specify whether the Performance Rating test of your choice is descriptor based or not by selecting **Yes** or **No** against the **Is this test a descriptor based test?** flag. By default, this flag is set to **No**.
10. If the **Is this a descriptor based test?** flag is set to **No**, then, clicking the **Add** button will reveal Figure 8.9.
11. Alternately, if you want your Performance Rating test to derive its descriptors from a test associated with the chosen **Component type**, then, set the **Is this a descriptor based test?** flag to **Yes**. The **Derive descriptors from the test** list as shown in Figure 8.8 will then appear listing all the tests applicable to the chosen **Component type.**

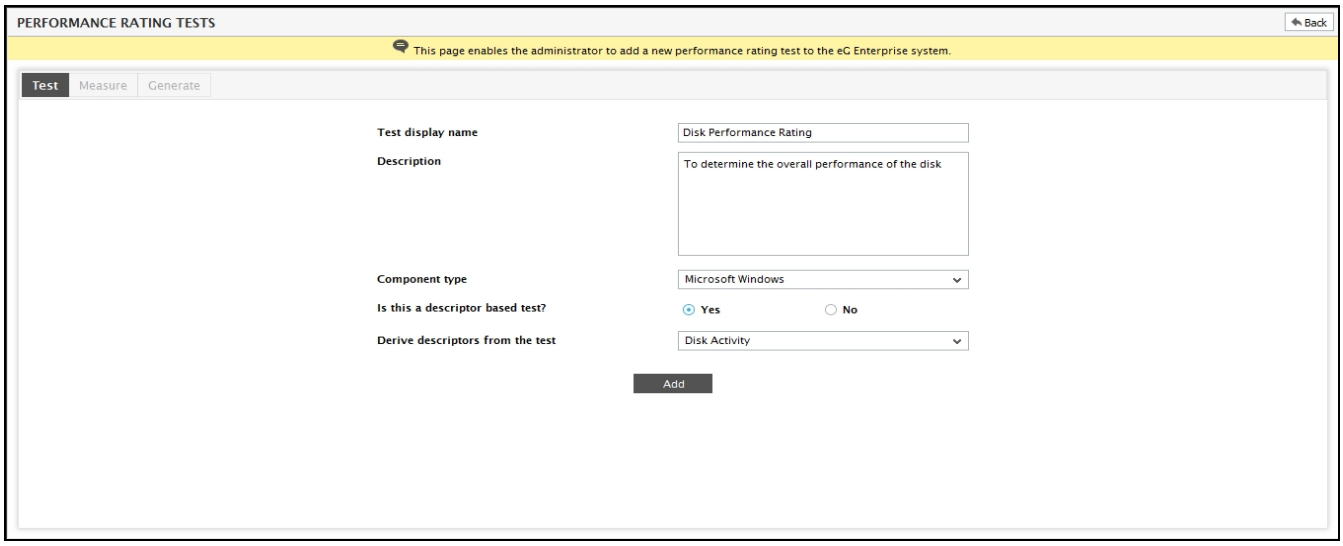

Figure 8.8: Choosing a test whose descriptors will be applicable for the newly created Performance Rating test

12. Select a test from the **Derive descriptors from the test** list and click the **Add** button. Figure 8.9 will then appear.

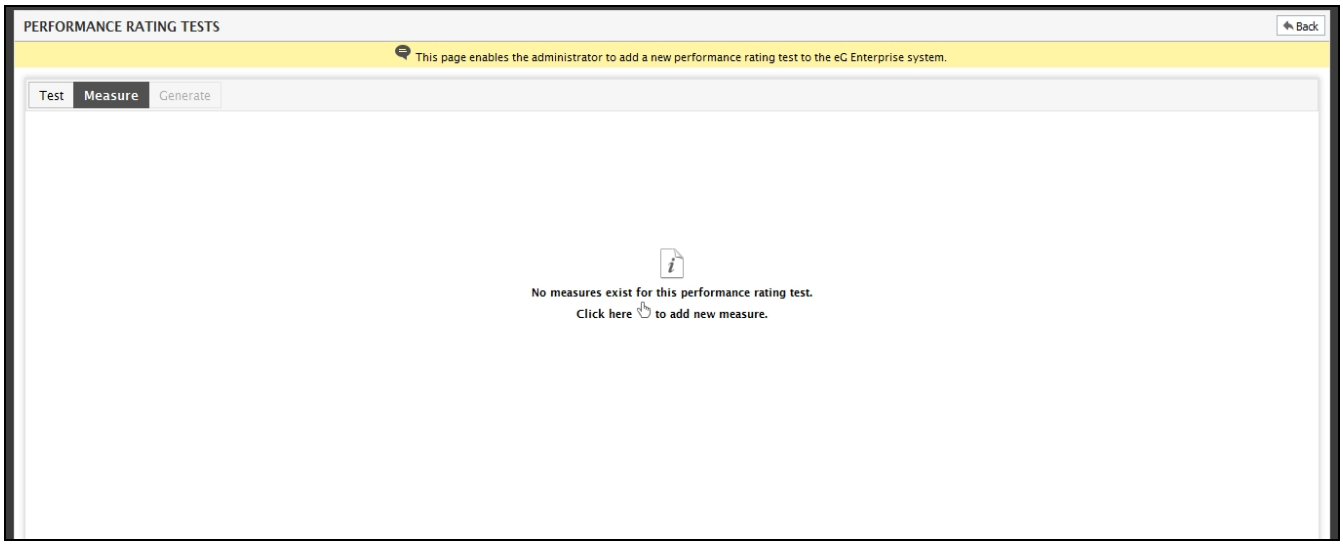

Figure 8.9: The Measure tab where no measures are associated with the Performance Rating test

13. By default, no measure will be associated with the test. To add measures for the test, click the **Click here** link in Figure 8.9.

14. Figure 8.10 will then appear using which you can configure the measures for your test.

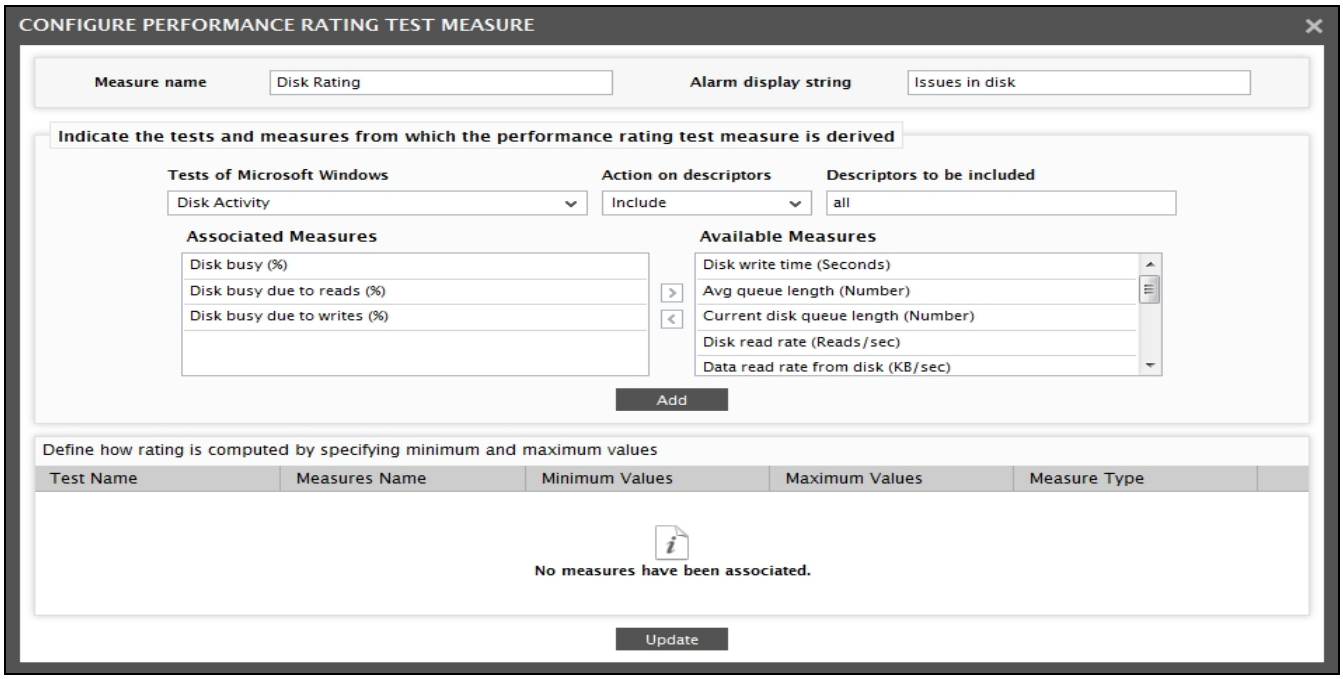

Figure 8.10: Configuring measures for the Performance Rating test

15. Specify the following in Figure 8.10.
16. First, specify the name of the measure that you wish to add in the **Measure name** text box.
17. In the **Alarm display string**, specify the alert message that is to be displayed when a threshold violation is detected for the measure that you are configuring.

    Note:

    The **Alarm Display String** text box is limited to *32* characters only. Therefore, specify a short alert message that does not exceed 32 characters.

18. The tests associated with the component type that you have chosen from the **Component Type** list box of Figure 8.10 will now be populated in the **Tests of** list box. To select the measures for your test, it is first necessary to select the test to which the measures are associated with. Select the test from the **Tests of** list.
19. Now all the measures pertaining to the chosen test will be listed in the **Available measures** list. To include the measures of your choice, select the measures from the **Available measures** list and click the **<** button. Your selection will now be moved to the **Associated measures** list.
20. If the test that you have chosen from the **Tests for** list is descriptor based, then you can either include all the descriptors of the test or include only a few of the descriptors. This can be achieved using the **Action on Descriptors** list. By default, the *Include* option will be chosen from the **Action on Descriptors** list and *all* option will be displayed in the **Descriptors to be included** text box. If you wish to include only a few of the descriptors of the chosen test, then you can do so by specifying a comma-separated list of descriptors in the **Descriptors to be included** text box.
21. Likewise, if you wish to exclude certain descriptors, then you can do so by selecting the *Exclude* option

from the **Action on Descriptors** list and then specify a comma-separated list of descriptors in the **Descriptors to be excluded** text box. By default, *none* will be displayed in the **Descriptors to be included/Descriptors to be excluded** text box.

Note:

You can also include/exclude descriptors based on specific patterns. By default, the *_*on*_* pattern is supported.

22. Once you have chosen the test, measures and the descriptors to be included/excluded, click the **Add** button in Figure 8.10.
23. Your selection will now be listed in the section below as shown in Figure 8.11. Now, specify the maximum/minimum values against each measure based on which the Performance Rating will be computed. The Minimum Values specified should be in ascending order and the Maximum Values specified should be in descending order. By default, the Minimum Values/Maximum Values should be in the format: *Critical/Major/Minor*.

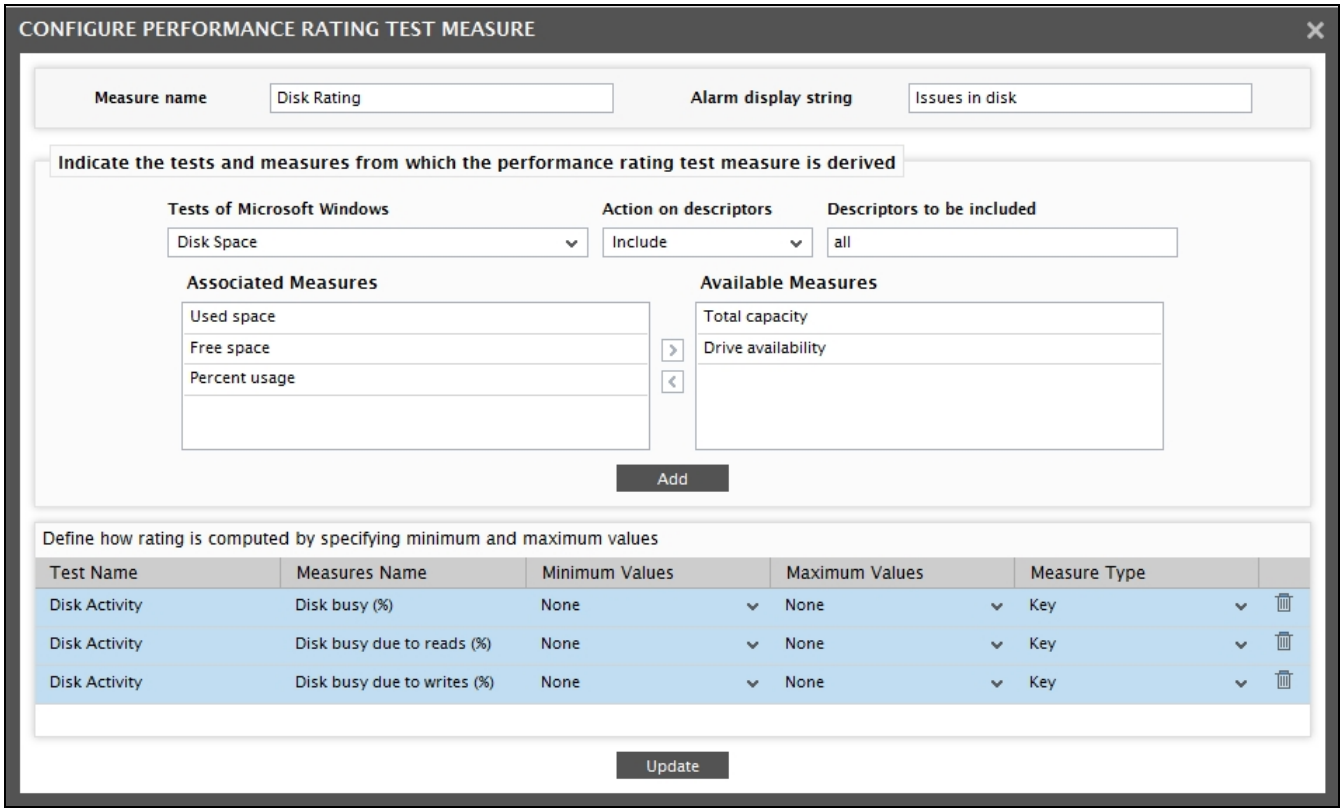

Figure 8.11: Listing the chosen measures

24. Next, specify whether the measure should carry a higher weightage in the calculation of the Performance Rating or a normal weightage by selecting either *Key* or *Non-key* from the **Measure Type** column.

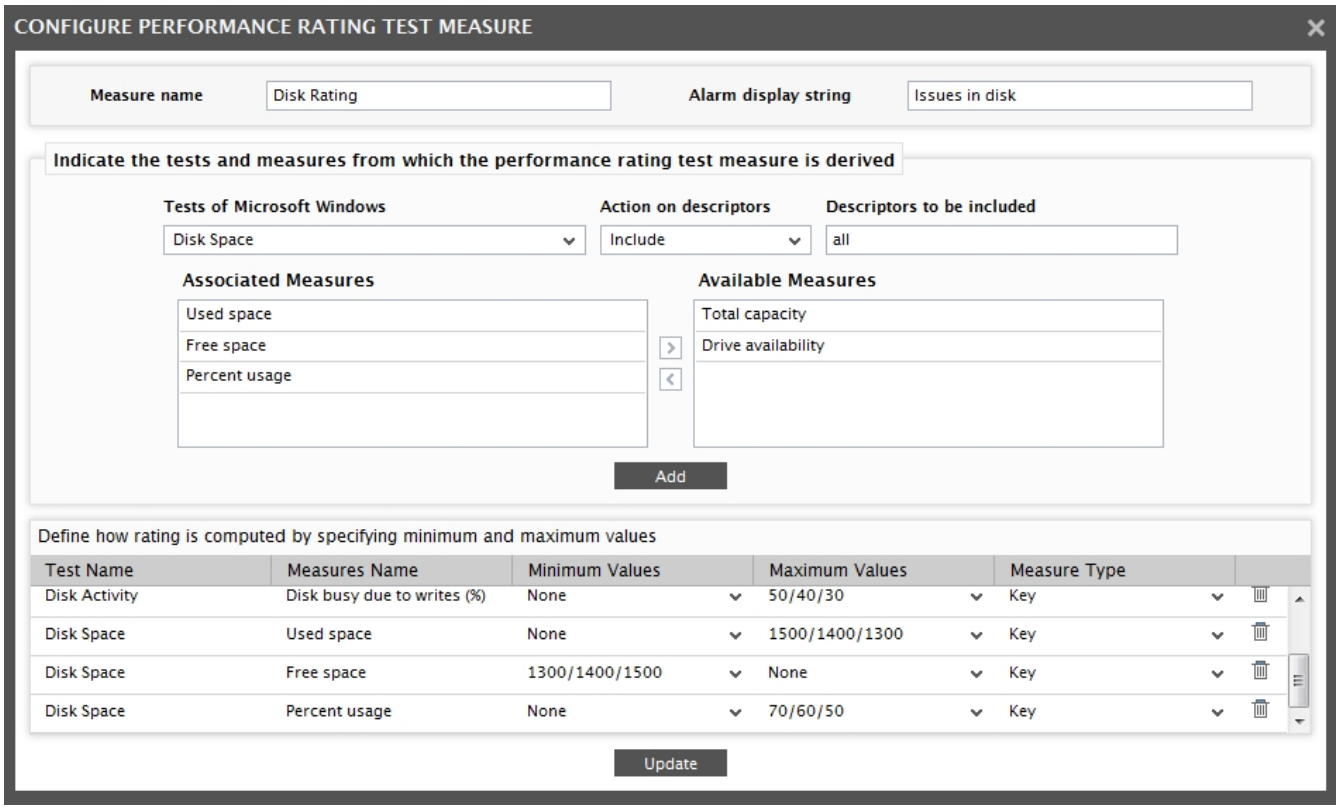

Figure 8.12: Setting the Minimum and Maximum values for the chosen measures

25. If you wish to delete your selection, then you can do so by clicking the button (see Figure 8.12).

    Note:

    By default, the Maximum Values and Minimum Values column will be *None*. While adding the measures for calculating the Performance Rating, only either the *Maximum Values* or the *Minimum Values* column can be *None*. If both the columns are *None*, Figure 8.13 will appear.

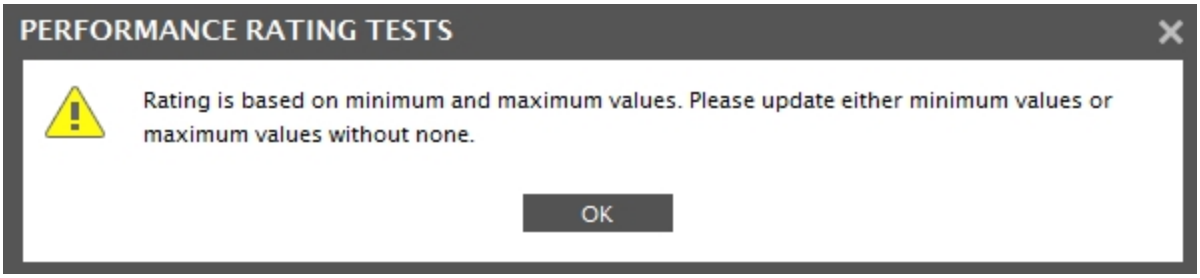

Figure 8.13: A message stating both the minimum values and maximum values cannot be none

26. Clicking the **Update** button in Figure 8.12 will lead you to the **Generate** tab as shown in Figure 8.14.

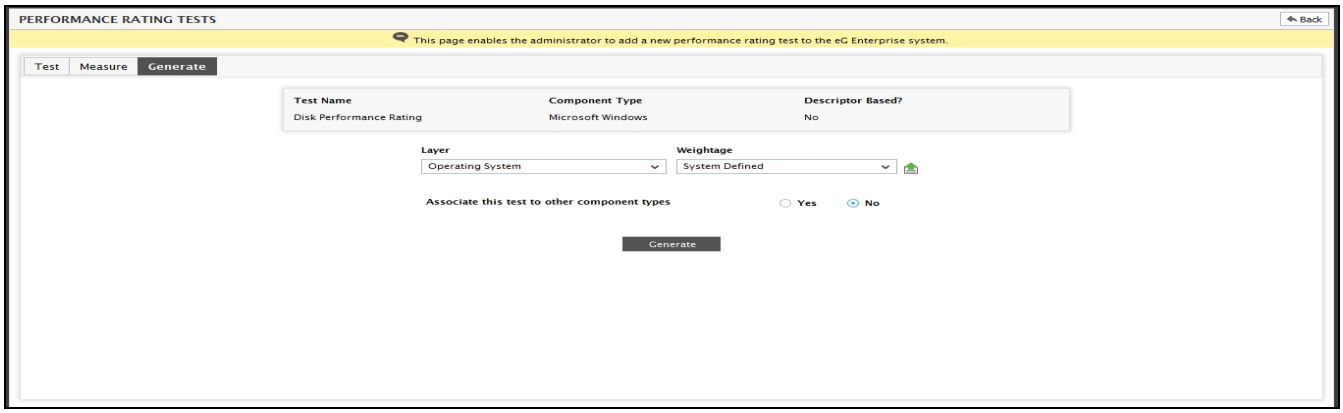

Figure 8.14: Associating the Performance Rating Test to a layer

27. In Figure 8.14, select the layer to which your test should be associated with.

28. Next, choose the weightage that needs to be applied to the measures for calculating the overall performance of your measure. The weightage to be applied can either be **System Defined** or **Custom Defined**.

29. If you have chosen **System Defined** from the **Weightage** list, then upon clicking the icon adjacent to the list, Figure 8.15 will appear listing the default system defined weightage values for the measures that are to be configured.

SYSTEM DEFINED WEIGHTAGE FOR MEASURES

MEASURE VALUE	WEIGHTAGE
Normal	1
Minor	None (Weightage of normal measure will be applied)
Major	None (Weightage of normal measure will be applied)
Critical	0

Figure 8.15: The default weightage

30. If you have chosen **Custom Defined** option from the **Weightage** list, then you would be required to provide your own wieghtage values for key measures in the **Weightage for Key measures** section and non-key measures in the **Weightage for non-key measures** section as shown in Figure 8.16.

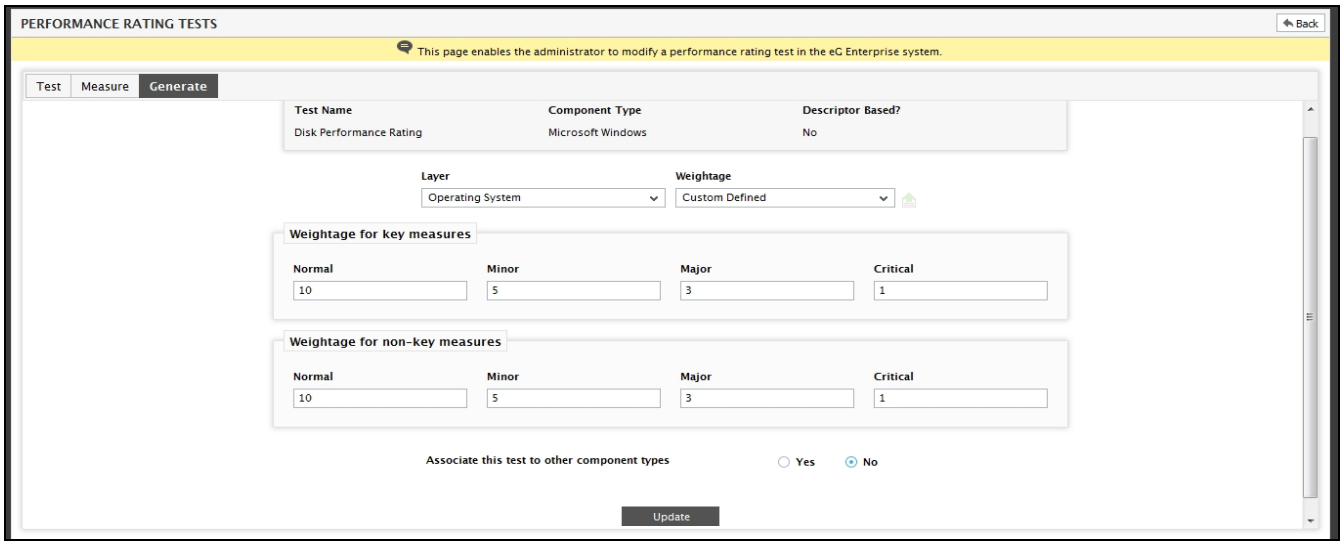

Figure 8.16: Applying your own weightage for Key and Non-key measures

Note:

If you have chosen the **Weightage** as **System Defined**, then the weightage applied for both *Key* measures and *Non-key* measures will be the same.

31. If you wish to associate your Performance Rating test to other component types, set the **Associate this test to other component types** flag to **Yes**. By default, this flag is set to **No**. If the flag is set to **Yes**, then the **Available Component Types** list will list all the component types that are associated with the chosen test and layer. Select the component types to which you wish to associate the test and click the **<** button. Your selection will now be transferred to the **Associated Component Types** list.

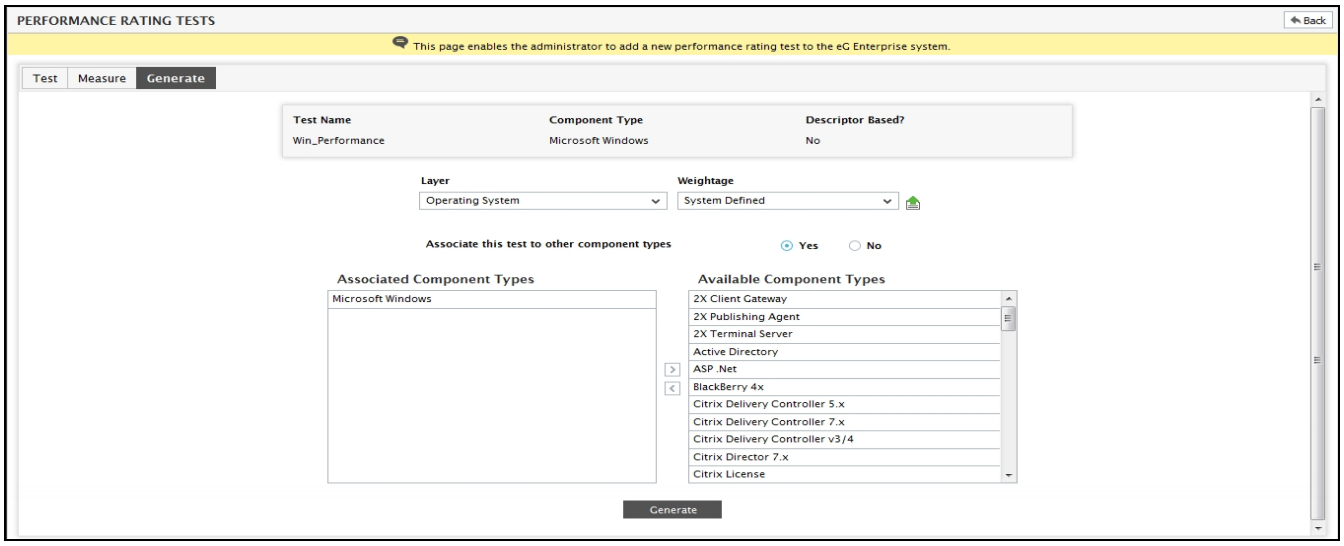

Figure 8.17: Associating the Performance Rating test to other component types

32. Clicking the **Generate** button in Figure 8.17 will configure the Performance Rating test.

Note:

The **Available Component Types** list will list all the component types that are associated with the test based on which the Performance Rating test is derived and the layer to which the Performance Rating test is associated with. The user should therefore intelligently associate the test to other component types. If the user chooses a component type that does not support the chosen layer, then the Performance Rating test will not report metrics.

33. Next, logout of the eG administrative interface and proceed to login back to the eG Enterprise Suite as a monitor user. You can now notice the newly configured test in the layer and component type that you have associated with.

9

Viewing the Managed Infrastructure

While monitoring large environments using eG Enterprise, administrators may want to instantly figure out the total number of managed components in the environment across categories / component-types to assess the load on the eG manager and the database. Some other times, they may require a quick summary of the number of managed components within a particular zone/service/segment. To enable such administrators and top-level executives to receive an overview of the managed infrastructure based on chosen criteria, the eG administrative interface provides the **MANAGED INFRASTRUCTURE** page.

This page appears when the **View** option is chosen from the **Components** menu of the **Infrastructure** tile. The page provides administrators with a quick look at the managed infrastructure as a whole, based on a chosen service / zone / segment.

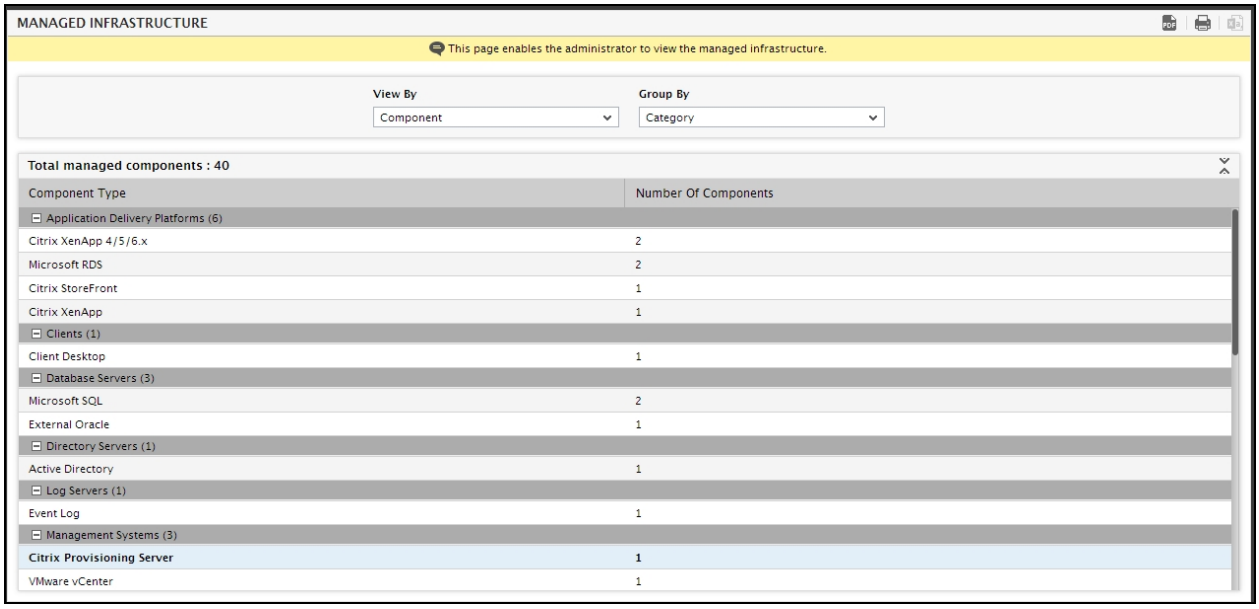

Figure 9.1: Viewing the Managed Infrastructure

Once you are in this page, do the following to view the information you require:

1. By default, this page groups the managed components on the basis of component categories and then displays the number of managed components under each component category. Accordingly, by default, the **View By** list is set to **Component** and **Category** is chosen from the **Group By** list. To view the managed components within a particular service/zone/segment instead, select the corresponding option from the **View By** list. The impact of this selection on the contents of the page is been discussed below:

- **Zone**: If you select this option from the **View by** list, then, you will have to select a particular zone from the **Zone** list; in this case, this page will display the number and category/type of components that have been included in the given zone.
- **Service**: If you select this option from the **View by** list, then, you will have to select a particular service from the **Service** list; in this case therefore, the page will display the number and category/type of components that are engaged in the delivery of the chosen service.
- **Segment**: If you select this option from the **View by** list, then, you will have to select a particular segment from the **Segment** list; in this case therefore, the page will display the number and category/type of components that are part of the chosen segment.

2. If the component list is to be grouped according to a pre-configured component category - eg., Database Servers, Web Servers, Web Application Servers - select **Category** from the **Group By** list. To group the resulting component list on the basis of a component-type, select **Type** from the **Group By** list. If you want the infrastructure to be grouped according to the component names, then select **Name** from the **Group By** list.
3. If you click on a managed component-type displayed under a particular category in Figure 9.1, then Figure 9.2 will appear displaying the complete details of every component of that type.

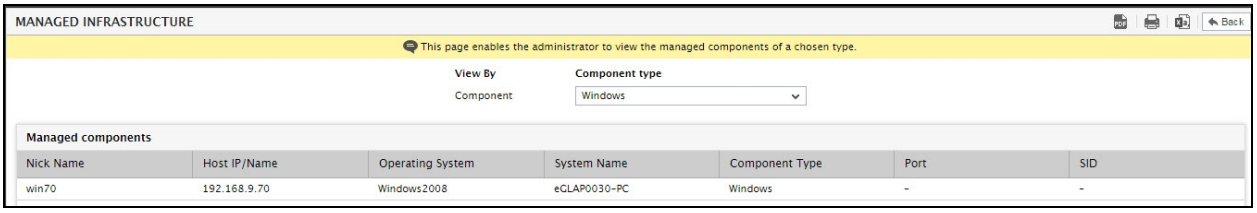

Figure 9.2: Viewing the details of the components of a particular type

4. These details include the following:
   - The **Nick Name** of the component;
   - The **Host IP/Name** of the component;
   - The **Operating System** on which the component excutes;
   - The **System Name** - i.e., the name of the system that hosts the component;
   - The **Component Type**;
   - The **Port** at which the component listens, if applicable;
   - The **SID** of the component, if applicable
5. Using the icons provided at the right, top corner of the component list, you can print the list, save it as a PDF document, or save it as a CSV file.
6. To know which internal/remote agent and external agent have been assigned to a particular component, click on the **Nick Name** of that component in Figure 9.2. This will lead you to Figure 9.3, where the monitoring details of that component will be displayed.

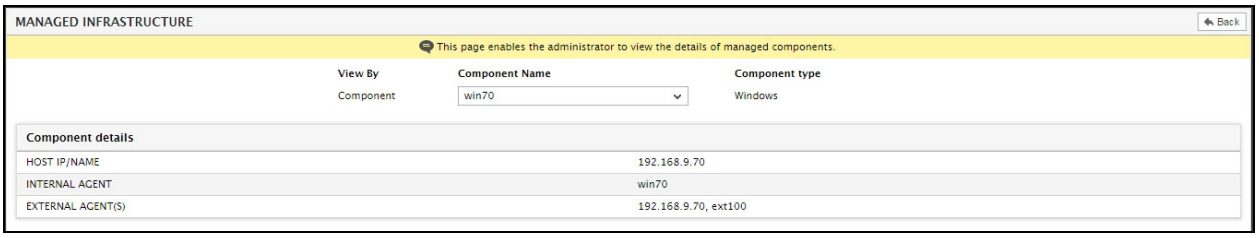

Figure 9.3: The details of a particular component chosen

10

Thresholds, Alarm Policies and Maintenance

Proactive alarm generation and their automatic suppression during maintenance periods are key capabilities of the eG Enterprise system.

Figure 10.1 briefly describes when and how alarms are triggered by eG Enterprise.

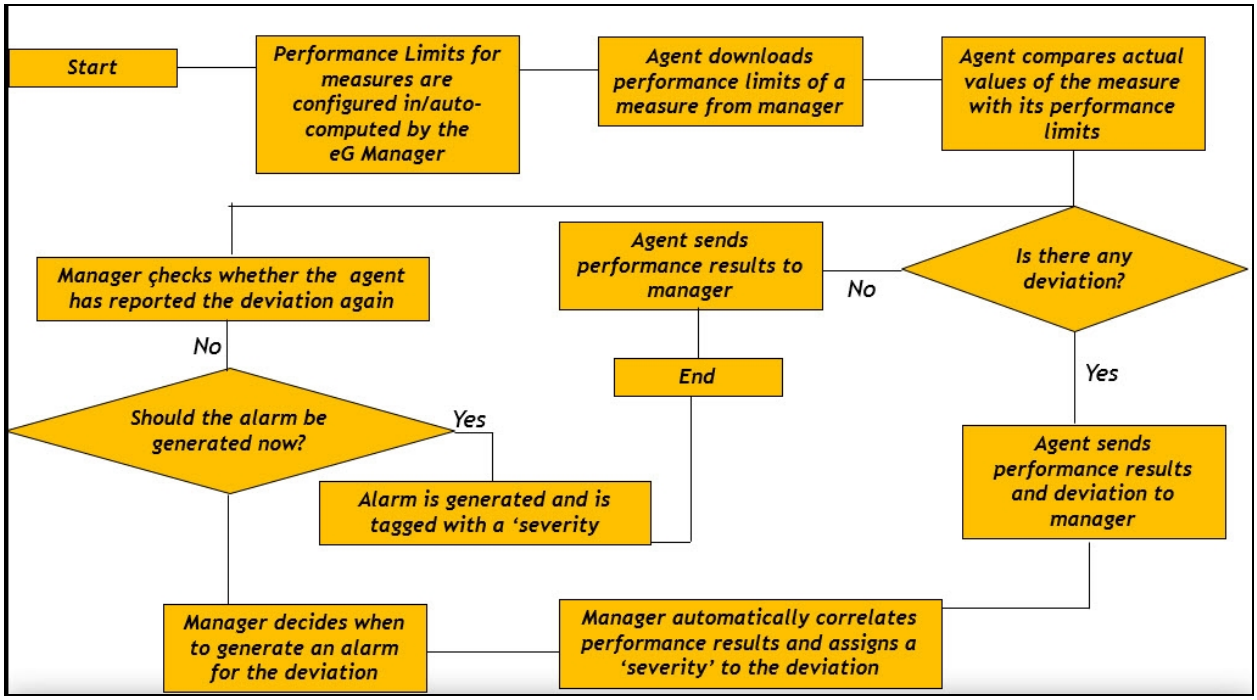

Figure 10.1: How does eG generate alarms?

As is evident from Figure 10.1, the eG agents periodically download pre-configured/auto-computed standards of performance (per measure) from the eG manager. These performance standards are called 'thresholds'. The eG agents then compare actual performance results with the standards. Deviations, if any, are noted and reported to the eG manager. The eG manager then, using a patented, virtualization-aware correlation algorithm, automatically correlates the performance results, intelligently triages the deviation reports, and accordingly assigns a severity/priority to each deviation. Based on the 'alarm policy' assigned to each measure, the eG manager then determines whether/not the deviation is severe enough for an alarm to be generated. If so, the manager triggers an alarm.

Alarms so generated can also be automatically suppressed when components are shut down for maintenance.

In this chapter, we will be discussing the concepts and procedures related to alarm generation and suppression in eG Enterprise. Broadly, this chapter will discuss:

- Alarm policy settings
- Threshold configuration
- Maintenance policy definitions

10.1 Alarm Policies

eG Enterprise includes four predefined alarm policies indicated in Figure 10.2. These alarm policies are used by the system to determine when to generate alarms. For instance there might be an instantaneous surge in the CPU utilization of a system. While an instantaneous surge may not be a problem, a set of periodic surges or a persistent increase in the measurement may indicate a problem. To differentiate between these different scenarios, eG Enterprise uses two parameters *Window size* and *Number of crossings*. The **Window size** represents the number of measurement values that are considered in determining the current state of a measurement. A crossing indicates a measurement being in violation of its threshold (i.e. a measurement value being lower than its lower threshold value, or a measurement value being higher than its upper threshold value). The **Number of crossings** denotes the number of times a measurement has crossed its threshold. Each of eG's alarm policies defines a window size and number of crossings. For example, an **immediate** policy has a **Window size** of 1 and **Number of crossings**. This means only one measurement value is considered in determining the state of a measurement. If the current value exceeds the upper threshold limit, the measurement is said to be in an abnormal state (since number of crossings is 1).

As its name indicates, this policy is ideal for cases where the administrator needs to be immediately alerted when an anomaly occurs. Metrics such as network / application availability can be monitored using this policy. For some other metrics, an administrator may not wish to be bothered about a sporadic threshold violation and may prefer to be alerted if a problem remains for a period of time. The **standard** alarm policy can be ideal for this, as it has a window size of 6, with number of crossings as 4.

You can modify any of the pre-configured policies or define custom policies using the eG admin interface. For this, select the **Alarm Policies** option from the **Alerts** tile. This will open Figure 10.2, which will be split into two distinct panels. The **Pre-defined alarm policy** panel will display all the default alarm policies. The **User defined alarm policy** anel will display all custom alarm policies.

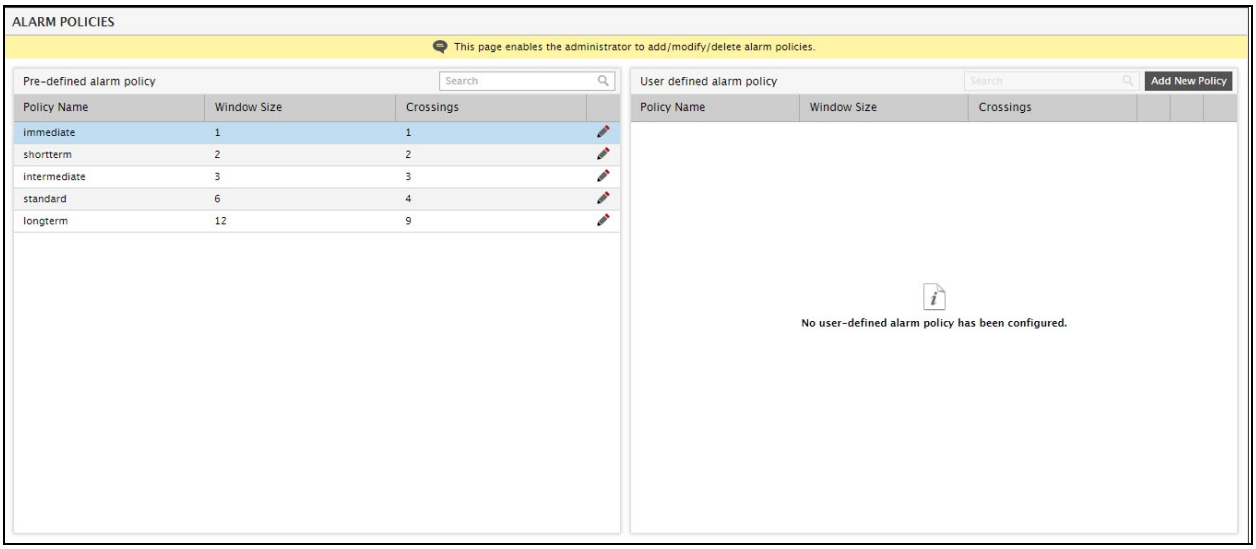

Figure 10.2: Predefined alarm policies of eG Enterprise

By default, the pre-defined alarm policies will be arranged in the **Ascending** order of their window sizes. You can change the sort order to **Descending**, by simply clicking on the column head, **Window Size**, in the **Pre-defined alarm policy** panel. You can even sort the contents of the **Pre- defined alarm policy** panel in the ascending/descending order of any of the other columns displayed therein, by just clicking on the corresponding column head.

You can modify any of the default policies by clicking on the button corresponding to that policy, but you cannot delete the default policies.

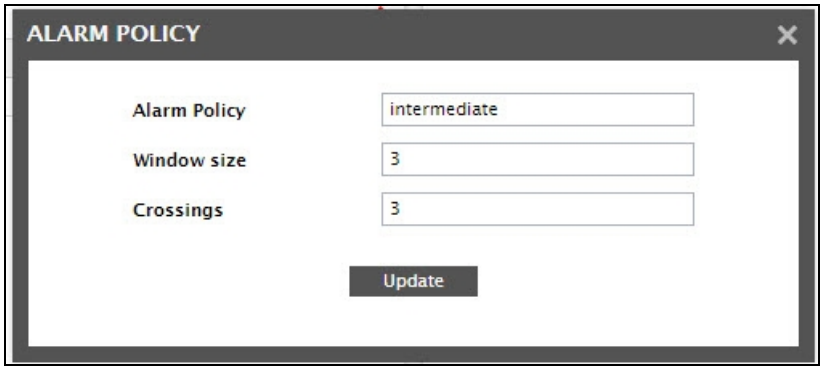

Figure 10.3: Modifying an existing alarm policy of eG Enterprise

To help you understand the **Window size** and **Number of crossings** concepts better, Figure 10.4 explains them with the help of an illustration.

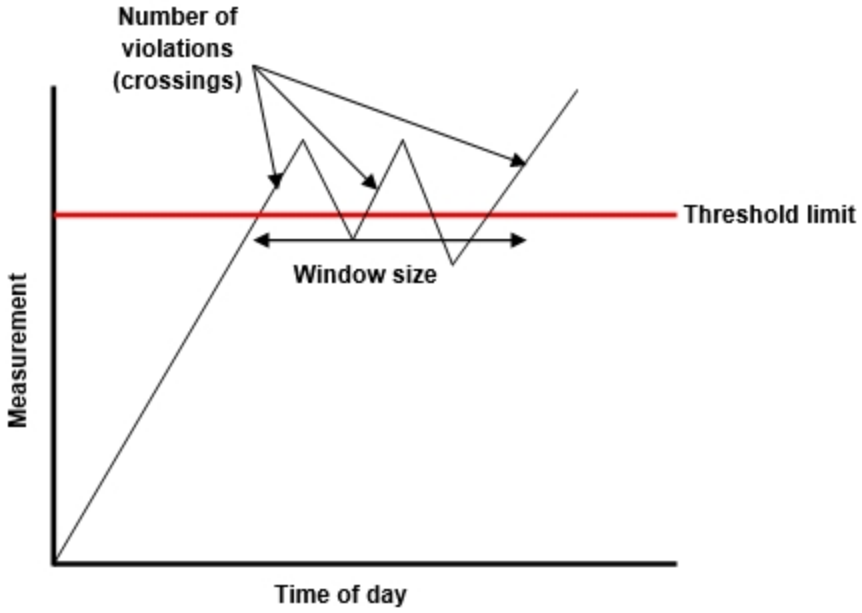

Figure 10.4: Graphical explanation of the concepts of Window size and Number of crossings

New policies can be added using the **Add New Policy** button in the **User defined alarm policy** panel of Figure 10.2. Figure 10.5 shows the details to be specified for the new alarm policy. Once you define your own policies, then those will also be listed in the **AGENTS - ALARM POLICIES** page under the **USER DEFINED ALARM POLICY** section depicted by Figure 10.5.

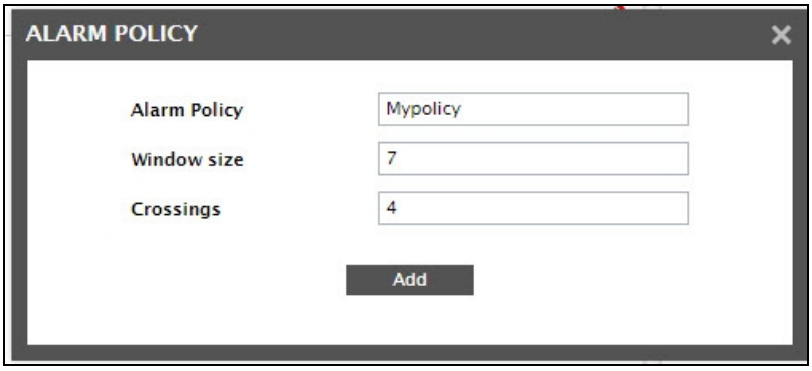

Figure 10.5: Adding a new alarm policy through the user interface

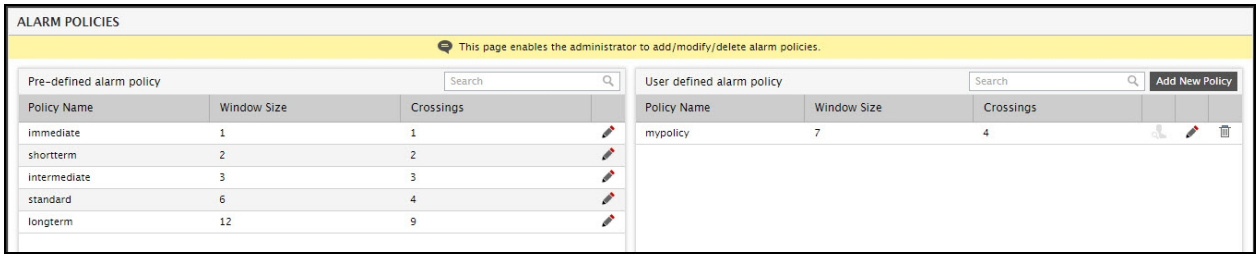

Figure 10.6: List of default and user-defined policies

Unlike default policies, a user-defined policy can be modified and/or deleted. Deleting an existing policy is possible using the button corresponding to a **User defined alarm policy** in Figure 10.6.

Until a newly added policy is associated with a measure, the icon (the **Show Associates** icon) corresponding to that policy (in the **User defined alarm policy** panel) will be disabled. As soon as the policy is associated with a measure, this icon will become enabled. You can click on this icon to view the measures that are currently associated with the corresponding policy (see Figure 10.7).

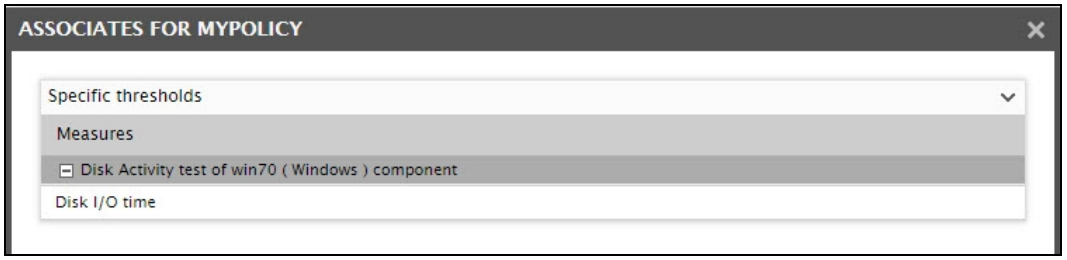

Figure 10.7: Viewing the measures associated with an alarm policy

This information is typically useful when attempting to determine which measures will be impacted when modifying an alarm policy.

Note that alarm policies that are assigned to one/more measures cannot be deleted.

To quickly locate a default policy for modifying, provide the whole/part of the alarm policy name to search for in the **Search** text box in the **Pre-defined alarm policy** panel, and click the 'magnifying glass' icon next to it. To quickly locate a user-defined policy for modifying/deleting, provide the whole/part of the alarm policy name to search for in the **Search** text box in the **User defined alarm policy** panel, and click the 'magnifying glass' icon next to it.

10.2 Thresholds

Thresholds govern the state of a measure. A threshold is characterized by an upper and/or lower limit of performance for the chosen measure. Whenever the threshold is violated, the state of the corresponding measure becomes 'abnormal'.

10.2.1 Types of Thresholds

eG Enterprise supports the following thresholding capabilities:

- Static
- Automatic
- Auto-static
- None

10.2.1.1 Static thresholding

For many metrics, thresholds can be set statically. For instance, based on the service level expectations and agreements, IT managers can set thresholds for metrics such as network availability, CPU usage, and latency. Application availability and response time can also be handled in the same manner. For example, availability should be 100% whenever the metric is measured. If not, a violation should be detected. Likewise, a network latency of several seconds is usually an indicator of a problem, no matter what time of day the measurement is made at.

To enable administrators to set static baselines for time-invariant measures such as the ones discussed above, the eG Enterprise system includes the **static thresholding** capability.

To illustrate how static thresholding works, consider the example of the CPU utilization of a host. The CPU utilization measure should never exceed a prescribed limit. Therefore, absolute threshold limits have to be explicitly defined for the CPU utilization measure in the **Absolute Maximum** and **Absolute Minimum** columns of Figure 10.8. The graph in Figure 10.8 depicts the absolute threshold values of the CPU utilization measure and its actuals.

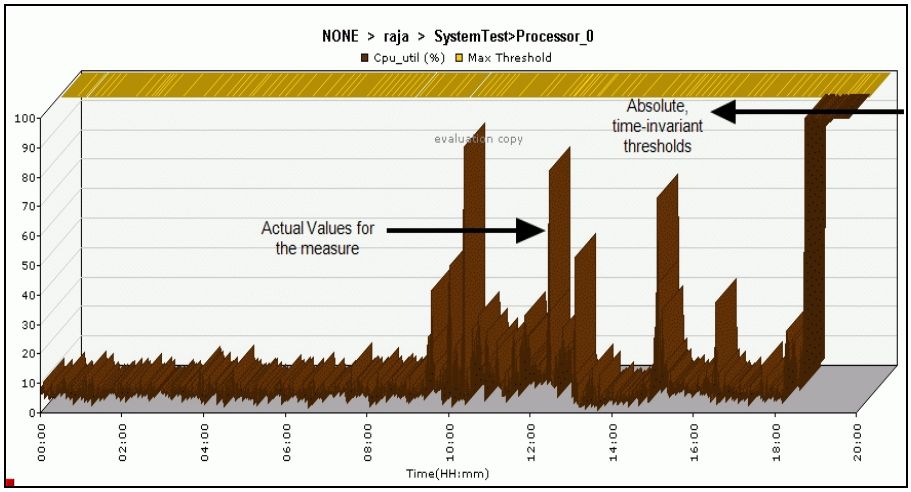

Figure 10.8: Measure graph of the CPU utilization measure indicating the set absolute thresholds and the actual values

Often, there is a need to set different threshold levels to map to different levels of severity of problems. The eG Enterprise system offers three levels of thresholds that correspond to the three alarm priorities - **Critical**, **Major**, and **Minor**. The user has to specify three maximum and/or three minimum threshold values in the format: *Critical/Major/Minor*. While the maximum thresholds are to be provided in the descending order, minimum thresholds have to be specified in the ascending order. For example, take the case of the Percent usage measure of Figure 10.8. This measure reports the percentage of disk space that has been utilized. The user can set a single Maximum threshold of say, 98, and expect to be alerted when the disk utilization crosses 98%. Alternatively, the user can also set multiple maximum thresholds, thereby instructing the eG Enterprise system to send different types of alerts at various levels of disk usage - in other words, the user can instruct the eG Enterprise system to trigger a **Minor** alert if the disk utilization crosses 50%, a **Major** alert if the disk utilization crosses 75%, and a **Critical** alert if it falls beyond 90%.

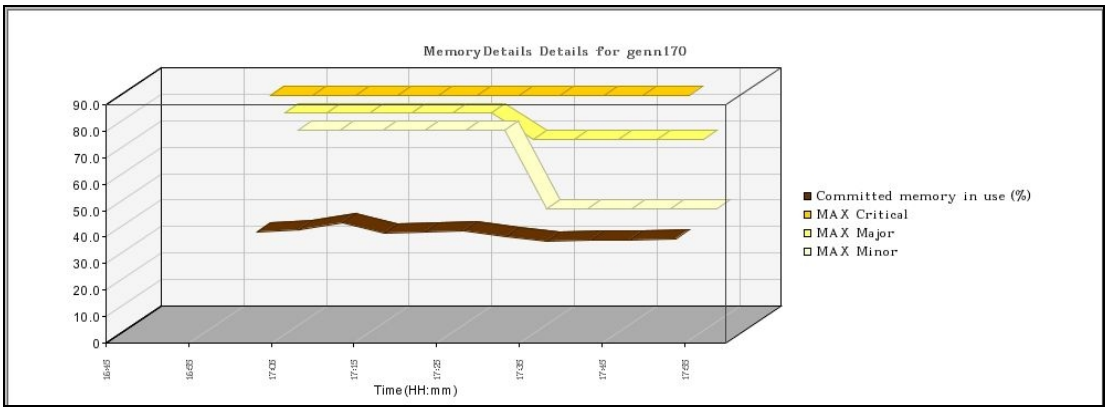

Figure 10.9: Multiple thresholds set for the 'Committed memory in use' measure of a Windows server

Multiple levels of threshold settings allows proactive alarms to be generated when a metric is slightly out of conformance, and a severe alarm to be generated when the problem worsens. This provides ample opportunity to the user to identify and attack a problem early in its life cycle.

In the case of the Percent usage measure in our example, the Maximum thresholds can be defined as "99/97/95". Since an absolute minimum threshold is not required for the Percent usage measure, it can remain as "none".

According to this specification, if the Maximum threshold of 99 is violated, then a **Critical** priority alarm will be generated. This is indicative of a critical issue with the host. Similarly, if the value of this measure crosses the Maximum threshold of 97, then a **Major** priority alarm will be generated. This is indicative of the existence of a major issue with the host. Likewise, a value beyond the Maximum threshold of 95 will result in a **Minor** priority alarm.

Similarly, if you take the case of the Free memory measure, where the minimum threshold is of significance, a sample minimum threshold setting for the measure could be, 200/300/400 (MB). In this case, the lower the threshold that is violated, higher will be the alarm priority.

Note:

If the threshold for a measure is set to *-/-/-*, then, it implies that such a threshold need not be computed for that measure. For instance, if you set the **Static Maximum** threshold of a measure to *-/-/-*, it means that for that measure static maximum thresholds need not be computed. This is why, when you revisit the **AGENT - THRESHOLDS** page to simply view or modify the threshold specifications of the same measure, you will find that the **Static Maximum** threshold specification displays *none* instead.

Thresholds can also be set based on industry standard best practices. For example, a rule of thumb when tuning an Oracle database server is that the database dictionary cache hit ratio should be 90% or more. If the hit ratio falls below this value, it indicates a need to tune the database server. This is another example where a threshold is set statically, without considering the time of day when the measurement is being made. eG Enterprise includes pre-specified threshold values for many metrics based on industry standard best practices.

10.2.1.2 Automatic Thresholds

In infrastructures where a metric varies with time, a static threshold value cannot serve as a reliable basis for judging performance. For example, consider a web server hosting a web site. The number of TCP connections to the web site (i.e. the current connections measure of the TCP test in the figure below) could be rather high on a particular day and low on another. Similarly, it could be high during the working hours and low during the nights. In such situations where measurement values change with the time of the day, it is very difficult to set accurate maximum and minimum limits manually. In such cases, the threshold value for this metric also has to be time variant.

Even when a metric is not time variant, its value may change from one server to another. For example, a high-end datacenter server may be able to handle hundreds of users, whereas a low-end standard server may be able to handle only a few tens of servers. In such cases too, it is extremely laborious and time consuming to determine what the normal values are for each and every server.

To handle such situations, eG Enterprise includes an **automatic thresholding** capability. Using past history of the values of the metric, eG Enterprise uses 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. In this approach, for example, the threshold values for a metric between 9am-10am tomorrow are based on the value of the metric for the same time period over the past days (the number of days to be looked at in the past is configurable).

Note:

You can configure how far back the manager should check for past history when computing automatic thresholds for a measurement. The default look back period is 14 days (i.e., 2 weeks). You can change this value, if required. For this, do the following:

- Select the **Manager** option from the **Settings** tile.
- From the **MANAGER SETTINGS** panel to the left of the page that appears, select the **Threshold Configuration** option.
- Specify a **Lookback period to compute automatic thresholds** in the right panel, and click the **Update** button to register the changes.

With eG's auto-thresholding capability, like the metric value, the threshold also is time varying. Whenever a deviation from this auto baseline (threshold) is detected, an alert is triggered. Since the baseline is set automatically, using this technique ensures that administrators are informed of problems well before they become critical enough to impact the end user experience.

Automatic thresholding is ideal for time varying metrics such as number of requests to a web server, the workload on a database server, queue lengths of requests waiting for processing, etc.

Even when thresholds are set automatically, an IT manager may want to choose a leniency factor for the thresholds. For example, an IT manager may want to allow for a 10% deviation from the norm. To accommodate such requests, eG Enterprise allows administrators to set a "sensitivity slider" for automatic thresholds. To configure the leniency factor, you need to specify the slider as a multiple of the auto-computed threshold value computed. For example, consider the case of the "Free memory" measure, which is an indicator of the amount of free memory available on a server. Assume that on one of the managed servers, the

free memory is known to decrease consistently and then grow back up (e.g., the operating system frees memory periodically). In such a scenario, the free memory threshold will be violated often (since the value decreases consistently), and this will result in a number of false alerts. In such a situation, the eG administrator can set the threshold to be a multiple of the auto-computed baseline - for example, if the minimum threshold is set to 30% of auto, it implies that the administrator has introduced a 30% leniency. That is, alerts are generated only if the free memory is 30% lower than what is the normal value. This capability allows administrators to fine-tune eG's automatic thresholding capability to suit their specific requirements.

Like static thresholds, multiple automatic threshold values should only be set - one each for every alarm priority. Let us take the example of the Free memory measure. Say, that administrators wanted to be alerted to the erosion of Free memory on a target server, at various stages. While they wanted proactive minor alerts to be generated if the free memory was 30 % lower than normal, a major alert was required for a 50% reduction in free memory, and a critical alert for an alarming 70% depletion of the memory resources. To ensure this, your **Minimum Automatic Threshold** setting should be: 30% of auto,50% of auto, 70% of auto.

The measure graphs provided by eG Enterprise's monitor interface can bring out the differences between static and automatic thresholding, more clearly. The graph in Figure 10.10 depicts the threshold limits that were automatically assigned to the Current connections measure in the example discussed above. Notice that the statistical data is very periodic and the threshold that is automatically computed by eG Enterprise follows the same pattern as the measurement values.

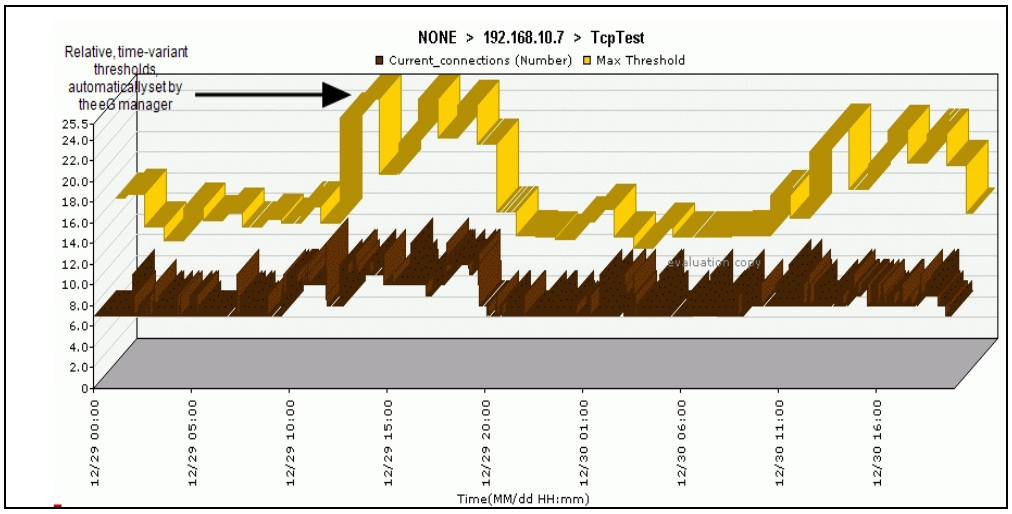

Figure 10.10: Measure graph of the Current_connections measure indicating the relative thresholds and the actual values

10.2.1.3 Auto-Static Thresholds

Automatic thresholds are ideal for metrics that are time variant. Often, the same metric may vary significantly from one server to another and from time to time. Consider a staging environment with a web server. Typically, there is no load on the web server and the automatic threshold is set accordingly. When someone logs in, the threshold will be breached and an alert may be raised by the system. This is a false alert because one user logging in does not signify a situation of interest to an IT manager. This scenario shows that while

automatic thresholding reduces the effort involved in configuring the monitoring tool (because IT managers do not have to configure thresholds for every metric and server), it does not eliminate false alerts.

Therefore, eG Enterprise allows IT managers to use a combination of static and automatic thresholds. A static threshold applied along with an automatic threshold provides a realistic boundary that has to be crossed before an alert is to be triggered. An IT manager can now configure an absolute maximum and an automatic maximum threshold for a metric. eG Enterprise compares the actual measurement value with the higher of the two maximum thresholds, and generates an alert only when the higher threshold is violated. In the example of the staging web server, the IT manager can set a static maximum of 100 requests in a measurement period (or a similar number representing a reasonable load). Once this is done, only if the actual load exceeds 100 requests in a measurement period, will an alert be generated, even if the auto-computed threshold is less than 100. If the auto-computed threshold is greater than 100, this value is used as the actual threshold.

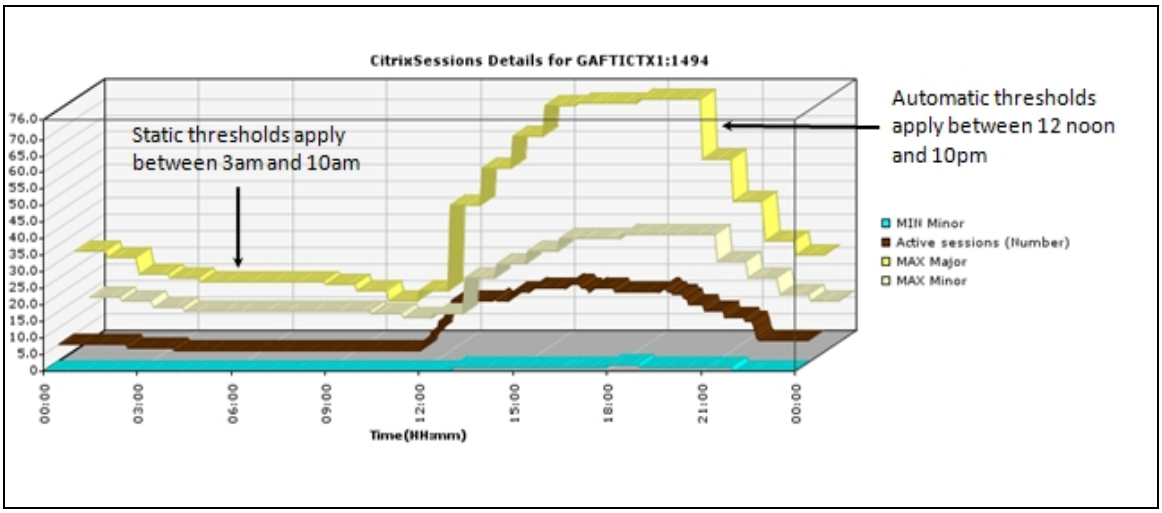

Figure 10.11: An auto-static combination threshold applied to the 'Active Sessions' measure of the CitrixSessions test

As in the case with the maximum thresholds, if a static minimum and an automatic minimum threshold are specified, then eG Enterprise will generate alarms only when the current value falls below the lower of the two threshold settings.

10.2.1.4 None

If the threshold policy for a measurement is none, an eG agent will stop tracking the state of this measurement (i.e. The agent will continue to collect values for this measurement but will not generate any alarms relating to this measurement). Even in the case of static and automatic thresholding, eG Enterprise allows the minimum and maximum threshold values to be “none”.

10.2.2 How to Configure Thresholds?

The broad steps for configuring thresholds are as follows:

- Figure out whether you want to override the default thresholds or configure component-specific thresholds;
- Define thresholds for every measure of the test, depending upon the type of thresholds suitable for each measure;

Each of these broad steps are dealt with elaborately in the sections that follow.

10.2.2.1 Choosing Between Default and Specific Thresholds

Default thresholds typically apply to all components of a chosen type. So, choose the **Test** option from the **Default** sub-menu of the **Thresholds** menu in the **Alerts** tile, only if you want your threshold settings to be shared by all components of a type. In this case, Figure 10.12 will appear.

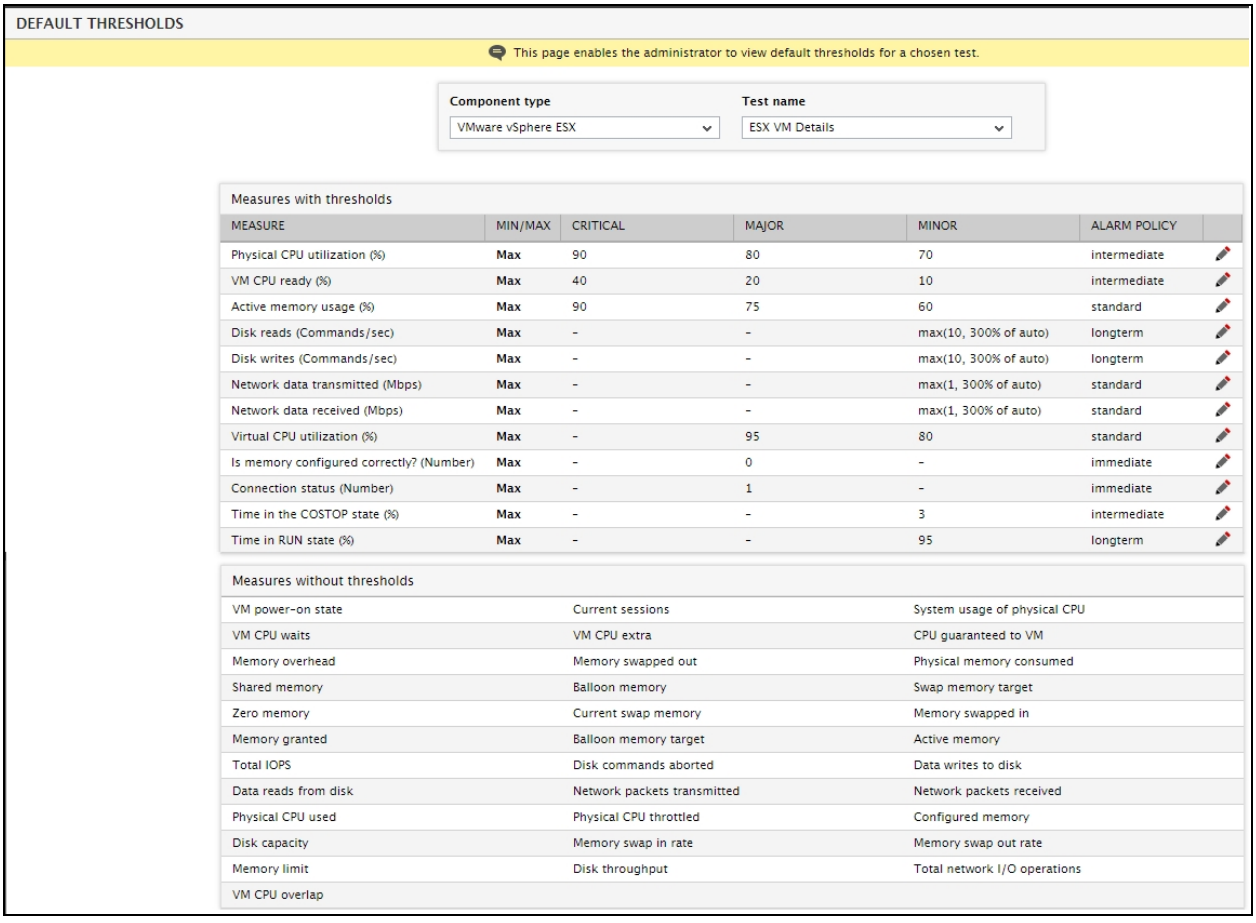

Figure 10.12: The DEFAULT THRESHOLDS page

From Figure 10.12, select the **Component type** to which the default thresholds apply and then pick the **Test name** for which the default thresholds need to be configured. All measures related to the chosen test will then

be listed. While some measures of a test may be pre-configured with thresholds, some may not be associated with any thresholds at all. Accordingly, the measure list in Figure 10.12 will be split into two sections – **Measures with thresholds** and **Measures without thresholds**. Each measure in the **Measures with thresholds** section will be accompanied by its current threshold configuration and alarm policy setting. You can change these configurations by clicking the icon corresponding to a measure in Figure 10.12. If you want to define thresholds for a measure in the **Measures without thresholds** section, simply click on that measure name in that section.

Note:

Changes made to the default thresholds of a test will automatically apply to all component types with which that test is associated. **Therefore, exercise caution while making changes to these threshold settings.**

The default thresholds set for a test will automatically apply to all descriptors supported by the test. Sometimes however, administrators may want to define the same threshold values for a group of closely-related descriptors. For instance, consider a virtualized environment where the first three letters of the VM name represent the operating system of the VM – lin for Linux VMs, win for Windows VMs and so on. In such an environment, the administrator may want to configure the same upper and lower bounds of performance for the ESX VM Details test of all the Windows VMs alone – i.e., for all VMs with names that begin with win. To achieve this, the administrator can quickly define a descriptor pattern that groups all VMs that begin with win and can configure thresholds for that pattern, so that the threshold settings govern all those descriptors that match the configured pattern – in this case, the Windows VMs.

For defining a descriptor pattern and configuring its default thresholds, select the **Descriptor** option from the **Default** sub-menu of the **Thresholds** menu in the **Alerts** tile. This will open Figure 10.13. If descriptor patterns have already been defined and their thresholds configured, Figure 10.13 will list these patterns. You can modify the threshold configuration of any of the existing patterns by clicking the button that corresponds to that pattern, or click the button to remove a pattern and its threshold specifications.

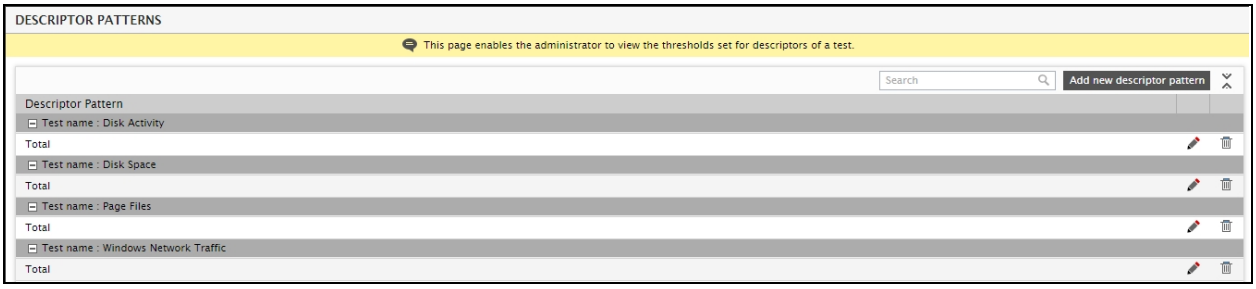

Figure 10.13: List of descriptor patterns that pre-exist

To add a new pattern, click the **Add new descriptor pattern** button in Figure 10.13. This will open Figure 10.14. In Figure 10.14, enter the **Descriptor pattern**. The pattern can be an expression of the form *expr* or expr or *expr or expr*. A leading '*' signifies any number of leading characters, while a trailing '*' signifies any number of trailing characters. In the case of our example, the descriptor pattern should look for VM names that begin with the string, win; so, enter win* in the **Descriptor pattern** text box. Then, click the icon corresponding to a measure in **Measures with thresholds** section of Figure 10.14 to change the threshold settings of that measure for a descriptor pattern. If you want to define thresholds for a measure in the **Measures without thresholds** section, simply click on that measure name in that section. Finally, click the **Update** button in Figure 10.14.

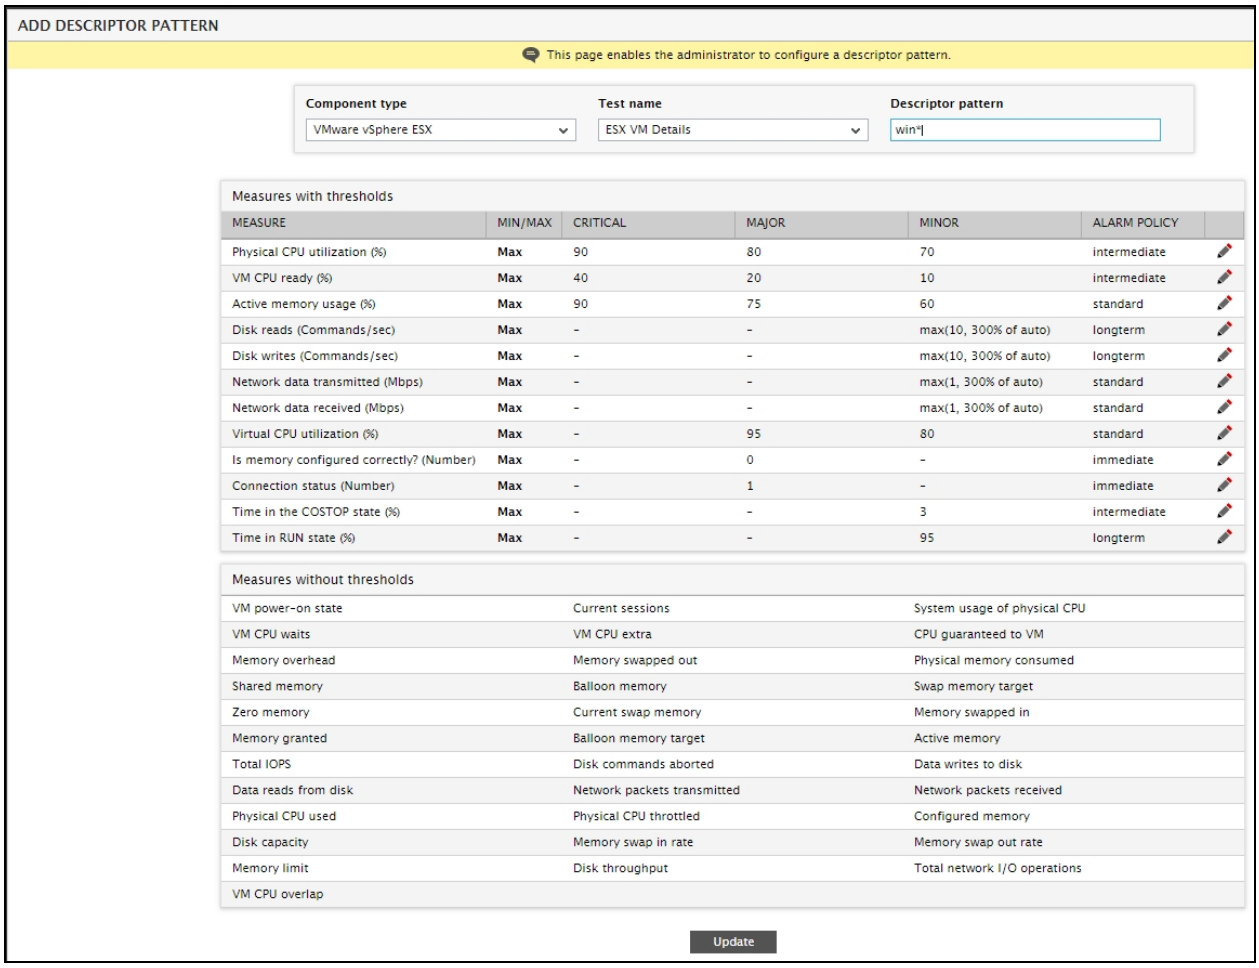

Figure 10.14: Configuring default thresholds for a descriptor pattern

Note:

Changes made to the default thresholds of a descriptor pattern will automatically apply to all component types with which the chosen test is associated and reports measures for a descriptor that matches the specified pattern. **Therefore, exercise caution while making changes to these threshold settings.**

Specific thresholds on the other hand apply to a chosen component only. To configure those thresholds that will govern the state of a particular component alone, select the **Specific** option from the **Thresholds** menu of the **Alerts** tile. Figure 10.15 will appear in this case.

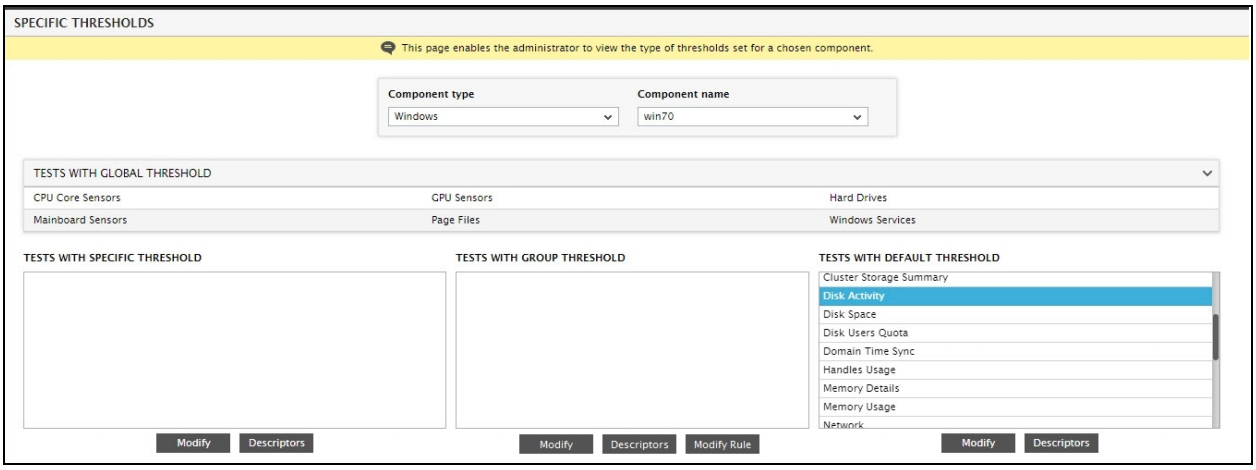

Figure 10.15: The SPECIFIC THRESHOLDS page

From Figure 10.15, select the **Component type** and then the specific **Component name** for which thresholds are to be configured. If global thresholds apply to any tests mapped to the chosen component, such tests will be listed in the **TESTS WITH GLOBAL THRESHOLD** section. To alter the threshold configuration of such tests, click on the test name in the **TESTS WITH GLOBAL THRESHOLD** section. This will allow you to modify the default thresholds of that test. **To know more about global thresholds, refer to** Section **10.4 of this document.**

Default thresholds will apply to all tests of a component, until specific thresholds are explicitly configured for one/more tests of that component. This is why, if a component is chosen for threshold configuration for the very first time from the **Component name** list of Figure 10.15, all tests mapped to that component will by default appear in the **TESTS WITH DEFAULT THRESHOLD** list box of Figure 9.4. To configure component-specific thresholds for a test in this list box, select the test and click the **Modify** button. Figure 10.16 will appear, which will list the measures reported by that test. While some measures of a test may be pre-configured with thresholds, some may not be associated with any thresholds at all. Accordingly, the measure list in Figure 10.16 will be split into two sections – **Measures with thresholds** and **Measures without thresholds**.

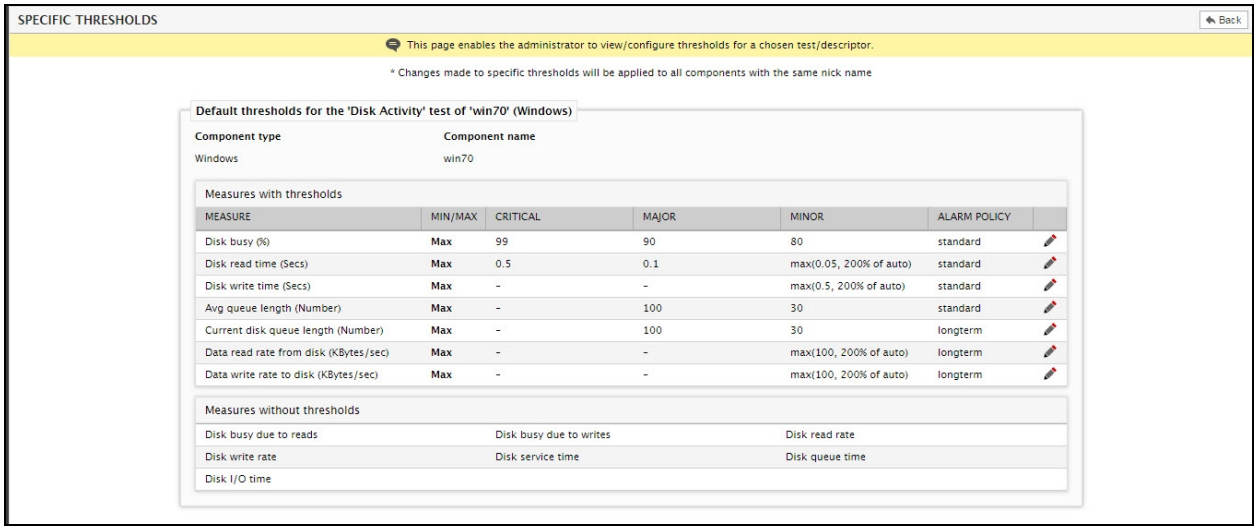

Figure 10.16: Configuring specific thresholds for a test mapped to a component

Each measure in the **Measures with thresholds** section will be accompanied by its current threshold configuration and alarm policy setting. You can change these configurations by clicking the icon corresponding to a measure in Figure 10.16. If you want to define thresholds for a measure in the **Measures without thresholds** section, simply click on that measure name in that section.

Note:

Changes made to the threshold settings of a test for a specific component will automatically apply to all those components with the same nick name as that component and to which the chosen test is mapped. **Therefore, exercise caution while making changes to these threshold settings.**

If you click on the **Descriptors** button in Figure 10.16 after selecting a test from the **TESTS WITH DEFAULT THRESHOLD** list box, a separate **Descriptors with component threshold settings** will automatically pop up as depicted by Figure 10.17. This section will display all descriptors that are currently active for the chosen component-test combination. To configure component-specific thresholds for a descriptor, click on that descriptor in Figure 10.17.

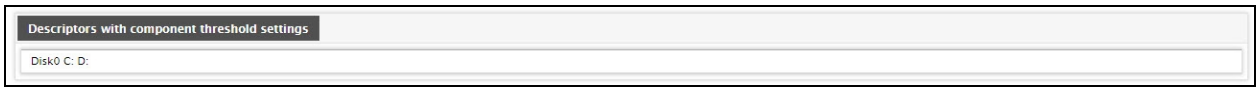

Figure 10.17: Viewing descriptors of a test

Figure 10.18 will appear revealing the default threshold configuration of each measure reported for that descriptor. Here again, you can click the icon corresponding to a measure in the **Measures with thresholds** section to change that measure's threshold configuration. Similarly, you can click on a measure name in the **Measures without thresholds** section of Figure 10.18 to configure descriptor-specific thresholds for it.

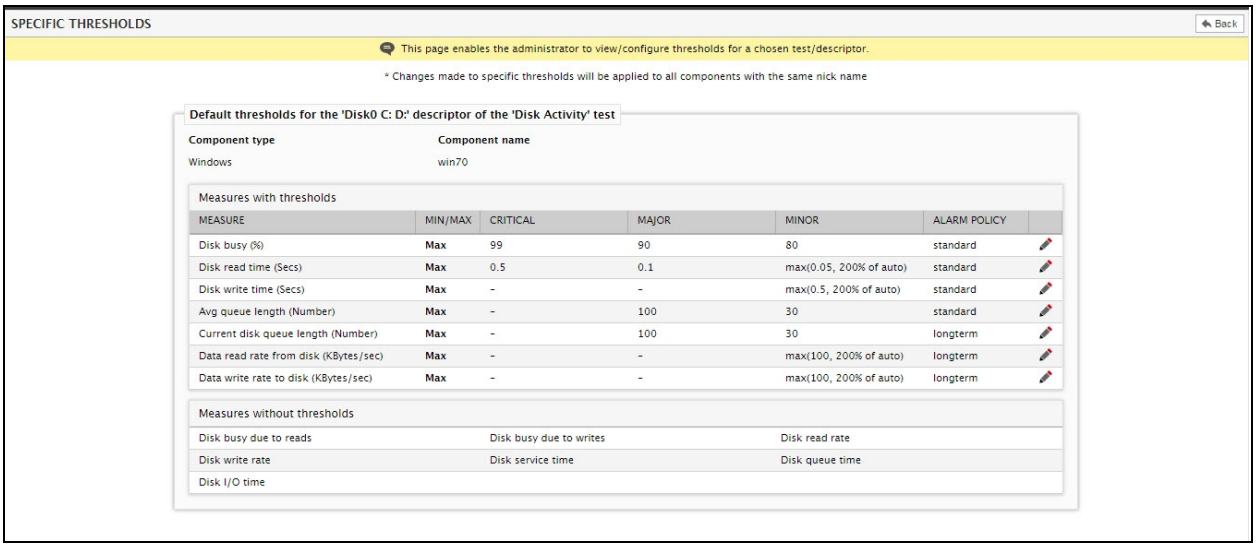

Figure 10.18: Viewing the default threshold configuration of the measures of a descriptor related to a specific component

Note:

You can also choose not to compute thresholds for specific descriptors of a test - for instance, you might not want to compute thresholds for the "all" descriptor of the Event Log Test, and say, the "word" and "excel" descriptors of Processes Test. In order to achieve this, do the following:

- Edit the **eg_tests.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
- In the **[NOTHRESHOLD_DESCRIPTORS]** section of the file, provide the descriptors to be disabled during threshold computation, in the following format:

```
[NOTHRESHOLD_DESCRIPTORS]
ProcessTest=word,excel
EventLogTest=all
```

- Finally, save the **eg_tests.ini** file.
- Once this is done, the next time you attempt to modify the threshold definition for Processes test for a specific component, the **Descriptors with component threshold settings** section of Figure 10.17 will not display the "word" or the "excel" descriptors. Similarly, for the Event Log test, the "all" descriptor will not be listed in the **Descriptors with component threshold settings** section. You are thus prevented from setting thresholds for these descriptors. However, if default thresholds or component-level thresholds (as the case may be) are configured for the test, such thresholds will continue to be applied to the excluded descriptors, and they will continue to appear in the monitor interface.

If specific thresholds have already been configured for one/more tests mapped to the selected component, then such tests will be listed in the **TESTS WITH SPECIFIC THRESHOLDS** list box of Figure 10.15. To make further changes to the specific threshold settings of any of these tests, select the test from the **TESTS WITH SPECIFIC THRESHOLD** list box and click the **Modify** button. If you want to make changes to the thresholds of the descriptors of any of these tests, just select the test from the **TESTS WITH SPECIFIC THRESHOLD** list box and click the **Descriptors** list.

Thresholds can also be configured and applied for a group of components. If such group thresholds apply to the chosen component, then all tests governed by these group thresholds will be listed in the **TESTS WITH GROUP THRESHOLD** list box of Figure 10.15. You can modify these group threshold settings for a specific test alone by selecting a test from this list box and clicking the **Modify** button below it. Once you modify the thresholds set for a test from this list, then the test will automatically move to the **TESTS WITH SPECIFIC THRESHOLDS** list i.e., specific thresholds will take precedence over the group threshold set for the test. If you wish to modify the Threshold rule governing the chosen test alone, then you can click the **Modify Rule** button. If you want to make changes to the thresholds of the descriptors of any of the tests listed in the **TESTS WITH GROUP THRESHOLD** list box, just select the test from the list and click the **Descriptors** list.

10.2.2.2 Defining the Thresholds of a Measure

Once you decide for what you want to define thresholds – whether it is for a component-type (default thresholds), a specific component (specific thresholds), a descriptor (descriptor thresholds), or a descriptor pattern – and select the test for which thresholds are to be defined, proceed to configure the upper and lower limits of performance for every measure of that test. The procedure for the same is as follows:

1. To define thresholds for any measure listed in the **Measures with thresholds** section of the **DEFAULT THRESHOLDS** (see Figure 10.12), **SPECIFIC THRESHOLDS** (see Figure 10.15), or the **DESCRIPTOR THRESHOLDS** page (see Figure 10.14 or Figure 10.18), first click on the icon corresponding to that measure. To define thresholds for a measure listed in the **Measures without thresholds** section of these web pages, simply click on the measure name. For instance, let us take the case of the Free memory measure reported by the **System Details** test. By default, no thresholds are set for this measure. Clicking on the measure name will hence lead you to Figure 10.19 below.

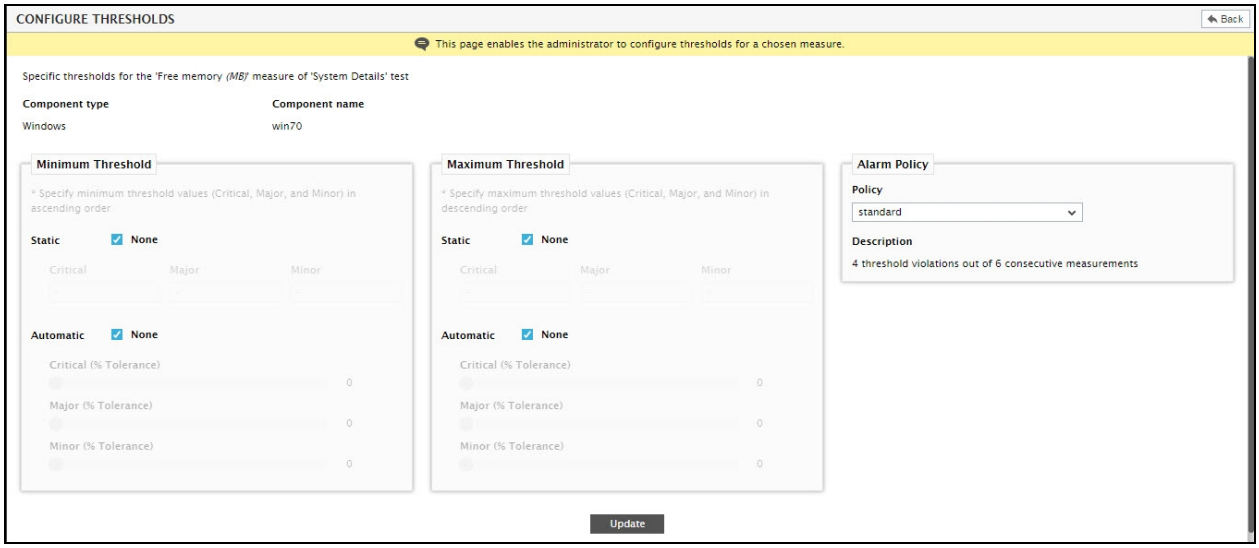

Figure 10.19: Configuring the Minimum thresholdfor a measure

2. As can be inferred from Figure 10.19, the Free memory measure has no **Minimum** Threshold; nor a **Maximum Threshold**. Accordingly, the **None** check box is selected against **Static** and **Automatic** in both the **Minimum Threshold** and **Maximum Threshold** sections.

   In the real world, any Windows/Unix host should have adequate free memory at all times. Excessive memory consumption by a host can often lead to the erosion of free memory, thus causing the performance of the host to degrade with time. This is why, if the amount of free memory on a host starts dwindling, it is a problem condition that should be brought to the attention of the administrator. To enable the eG Enterprise system to register a drop in the amount of free memory on a host as an 'abnormality', a **Minimum Threshold** will have to be set for the Free memory measure. A **Minimum Threshold** is violated when the actual value of a measure falls below the configured threshold value. Hence, it is only appropriate that we set a **Minimum Threshold** for the Free memory measure.

3. The first step to setting a **Minimum Threshold** is to indicate the type of threshold that should be set – static? automatic? or auto-static?. A **Static Minimum Threshold** should be configured with absolute/fixed values, and is hence, ideal for environments where load does not vary with time. To configure a **Static Minimum Threshold**, first deselect the **None** check box against **Static** in the **Minimum Threshold** section (see Figure 10.20).

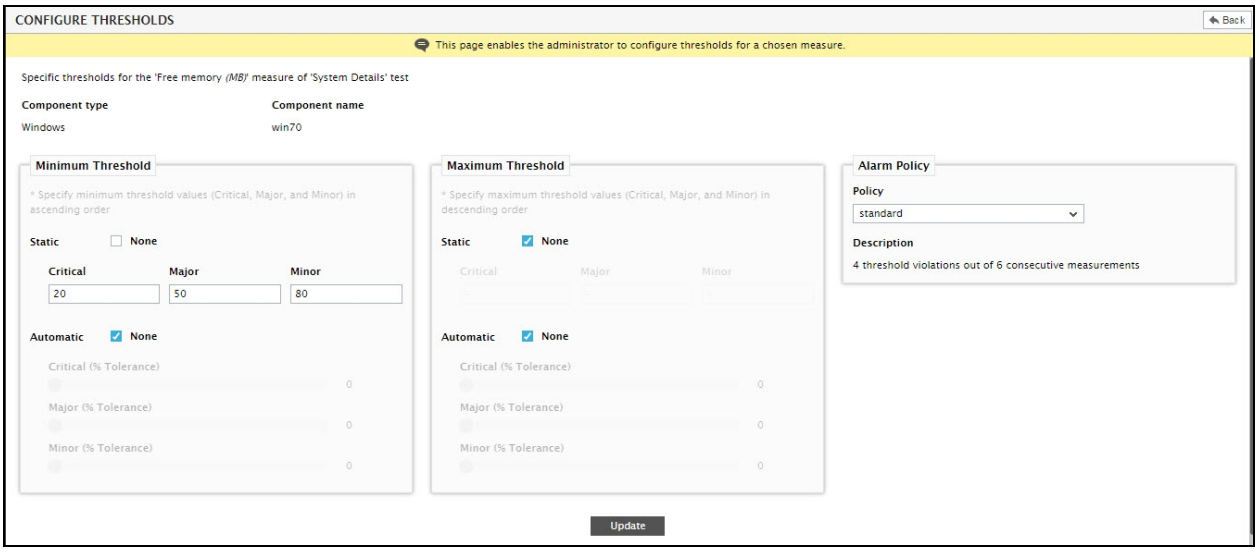

Figure 10.20: Configuring a Static Minimum Threshold for the Free memory measure

4. Multiple levels of **Static** thresholds need to be set – one each for every alarm priority (Critical, Major, and Minor). Therefore, specify a value (in MB) in the **Critical**, **Major**, and **Minor** text boxes in Figure 10.20. Also, note that the **Minimum Threshold** specifications should be in the ascending order. In other words, the **Critical** threshold should be configured with the lowest value and the **Minor** threshold with the highest value. This ensures that the eG Enterprise system sends out a **Minor** alert when the Free memory dips slightly, a **Major** alert when it falls a little more, and a **Critical** alert when Free memory falls steeply. Also, if you so need, you can even skip threshold levels – i.e., you need not specify threshold values for all the three (Critical/Major/Minor) levels; if required, you can omit threshold specifications for one/two levels.

5. Now that the **Static Minimum Threshold** has been set for the Free memory measure, click the **Update** button to register the changes.

6. Instead of manually specifying the thresholds, you can also configure the eG Enterprise system to automatically compute the thresholds for Free memory. Automatic threshold computation is ideal for environments in which load is dynamic, owing to which, administrators often struggle to figure out the norms of performance of a measure. To set an **Automatic Minimum Threshold** for the Free memory measure, instead of the **Static Minimum Threshold**, first select the **None** check box against **Static** to disable **Static** thresholding, and then deselect the **None** check box against **Automatic** to enable **Automatic** thresholding (see Figure 10.21).

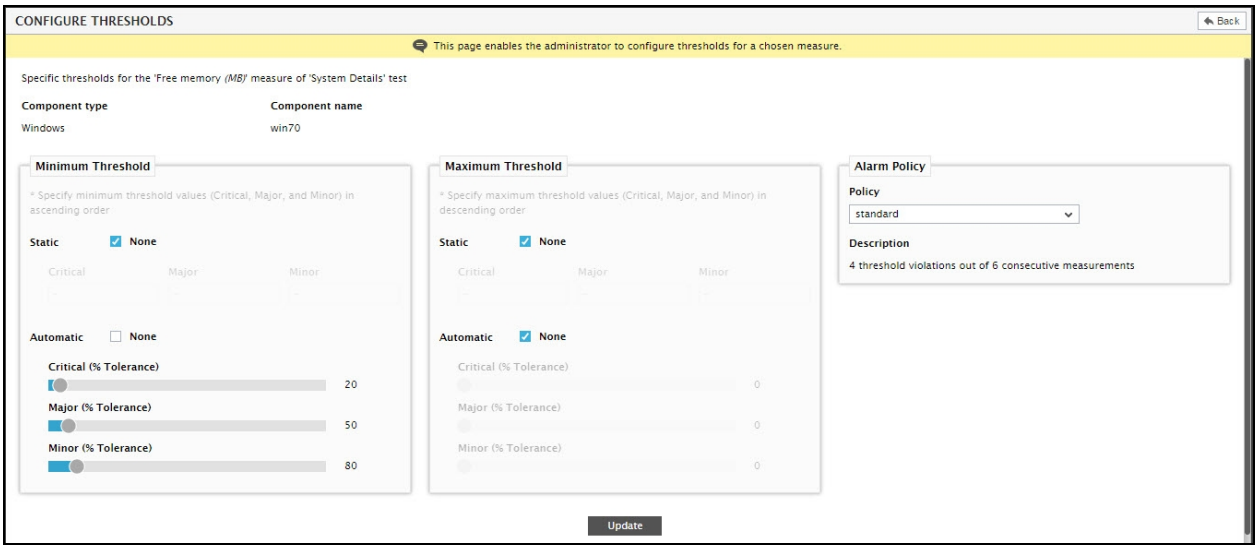

Figure 10.21: Configuring Automatic Minimum Thresholds for the Free memory measure

7. Using the 'slider controls' provided by the **CONFIGURE THRESHOLDS** page, you can even weave a leniency factor into your auto-computed performance limits. For instance, you may want the eG Enterprise system to ignore a condition where the Free memory is 20% of its auto-computed threshold. However, you would want the monitoring system to alert you if the Free memory value is lower than 20% of the auto-computed baseline. In this case, you can use the slider controls to factor in the 'tolerance level' of 20% into the auto-computed baselines. Like the **Static Thresholds**, multiple levels of **Automatic Thresholds** can also be set. For the Free memory measure in our example, let us make sure that a **Critical** alert is generated if Free memory is less than 20% of the auto-computed thresholds, a **Major** alert is generated if the amount of Free memory is lower than 50% of the auto-computed baseline, and a **Minor** alert is generated if the actual Free memory value is less than 80% of the auto-computed limit. For this purpose, first click on the knob on the slider control below **Critical** and drag the knob until the reading touches 20. Similarly, toggle the slider controls under **Major** and **Minor** in Figure 9.14 and make sure that the tolerance level is set to 50 and 80 respectively. **Note that your tolerance levels should also be in the ascending order.** Here again, you can skip one/two tolerance levels, if you do not want to provide any specification for them.
8. With that, we have learned how to configure **Automatic Minimum Thresholds** for a measure.
9. Alternatively, you can even configure your **Minimum Thresholds** to include a combination of **Static** values and **Automatic** baselines (see Figure 10.22). To achieve this, first deselect the **None** checkbox against **Static** to enable static thresholding. Then, manually specify absolute values in the **Critical**, **Major**, and **Minor** text boxes. Next, enable **Automatic** thresholding by deselecting the **None** check box against **Automatic**, and then use the slider controls to provide **Critical**, **Major**, and **Minor** tolerance levels.

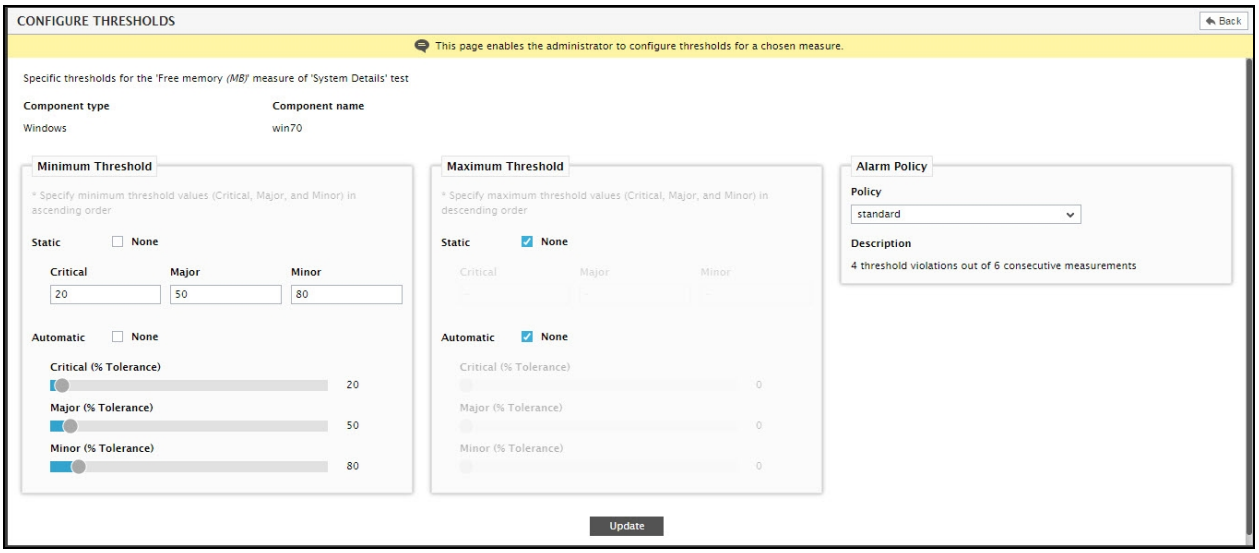

Figure 10.22: Configuring Auto-static Minimum Thresholds for the Free memory measure

10. When **Auto-static Minimum Thresholds** are configured, then, at any given point in time, eG Enterprise compares the actual performance results with the lower of the two threshold specifications, and thus isolates violations.

11. Once you have configured thresholds for a measure, assign an **Alarm Policy** to that measure by picking a **Policy** from the **Alarm Policy** section of Figure 10.22. Alarm policies indicate when alarms are to be generated for a measure.

Reference:

To know more about **Alarm Policies**, refer to Section **10.1** of this chapter.

Note:

As can be inferred fromFigure 10.22, the thresholds set for every measure is accompanied by an **Alarm Policy** specification. This specification indicates when alarms should be generated by the eG manager. The priority that will be assigned to such an alarm depends upon the threshold configuration and its corresponding alarm policy specification. By default, the following rules are applied when determining the alarm priority, if the number of violations in a time window matches the alarm policy specification (e.g., 4 out of 6):

- If all violations are critical, then alarm priority would be critical
- If all violations are major, then the alarm priority would be major
- If all the violations are minor, then the alarm priority is minor
- If the number of critical violations is greater than the number of major, and the number of critical violations is greater than the number of minor violations, then the alarm priority is critical

- If the number of major violations is greater than or equal to the number of critical violations, and the number of major violations is greater than the number of minor violations, then the alarm priority is major
- In all other cases, the alarm priority is minor

12. In the same way, if required, you can also set a **Maximum Threshold** for the Free memory measure. However, it is pointless to set an upper bound for the Free memory measure, because the higher the Free memory, the better the performance of the host. If you want to save the threshold settings for the Free memory measure, click the **Update** button in Figure 9.15.
13. Now, let us take a different example to understand how a **Maximum Threshold** is to be set. Let us take the case of the Disk I/O time measure reported by the **Disk Activity** test. The Disk I/O time measure reports the average time taken by a disk to perform read/write operations. A sudden spike/steady increase in the value of this measure therefore is a sign of a problem condition. To ensure that administrators are alerted to surges in the value of this measure, you need to configure a **Maximum Threshold** for this measure. A **Maximum Threshold** is violated if the actual value of a measure exceeds/grows beyond the maximum threshold level.
14. Like in the case of **Minimum Thresholds**, you can configure **Static Maximum Thresholds**, **Automatic Maximum Thresholds**, or **Auto-static Maximum Thresholds** for a measure. Figure 10.23 depicts how **Static Maximum Thresholds** can be configured for the Disk I/O time measure.

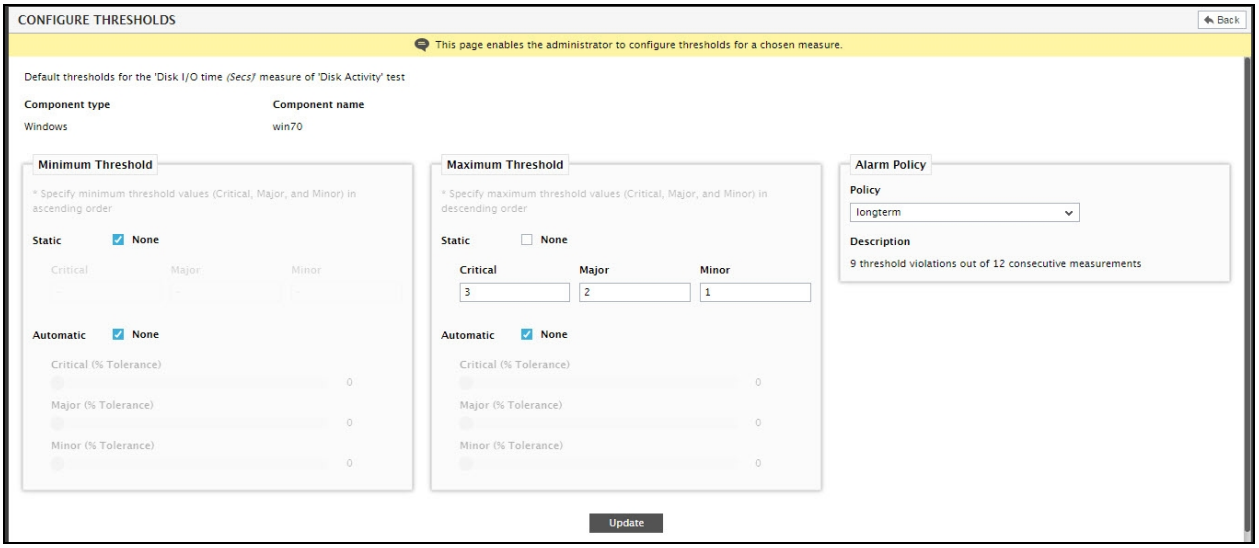

Figure 10.23: Configuring Static Maximum Thresholds for the Disk I/O time measure

15. The steps to be followed for configuring **Static Maximum Thresholds** for the Disk I/O time measure are as follows:
    - First, deselect the **None** check box against **Static** in the **Maximum Threshold** section, as depicted by Figure 10.23. This will enable **Static thresholding**.
    - Then, proceed to specify **Critical**, **Major**, and **Minor** threshold values in the respective text boxes. For the Disk I/O time measure, these threshold values should be in seconds. When specifying the **Maximum Threshold**, note that the **Critical**, **Major**, and **Minor** threshold values should be provided in the descending

order – i.e., the **Critical** threshold should be configured with the highest value and the **Minor** threshold should be configured with the lowest value. In the case of our example, let us configure the eG Enterprise system to change the state of the Disk I/O time measure to **Critical**, if the value of this measure exceeds 3 seconds, change the measure state to **Major**, if the value of this measure exceeds 2 seconds, and change the measure state to **Minor**, if disk read/write operations take over 1 second to complete. Accordingly, set **Critical** to 3, **Major** to 2, and **Minor** to 1, as shown by Figure 10.23.

- If you so desire, you can skip one/two threshold levels. For example, you can configure **Critical Static** and **Major Static** thresholds, and omit the **Minor Static**.

16. If disk activity is very dynamic in your environment, then it may not be advisable to go with static thresholds, which do not change with time. In such cases, it is best practice to disable static thresholding and enable **Automatic thresholding** instead. To auto-compute the baselines for the Disk I/O time measure, first select the **None** check box against **Static** in the **Maximum Threshold** section, to disable static thresholding. Then, deselect the **None** check box against **Automatic** to enable auto-thresholding (see Figure 10.24).

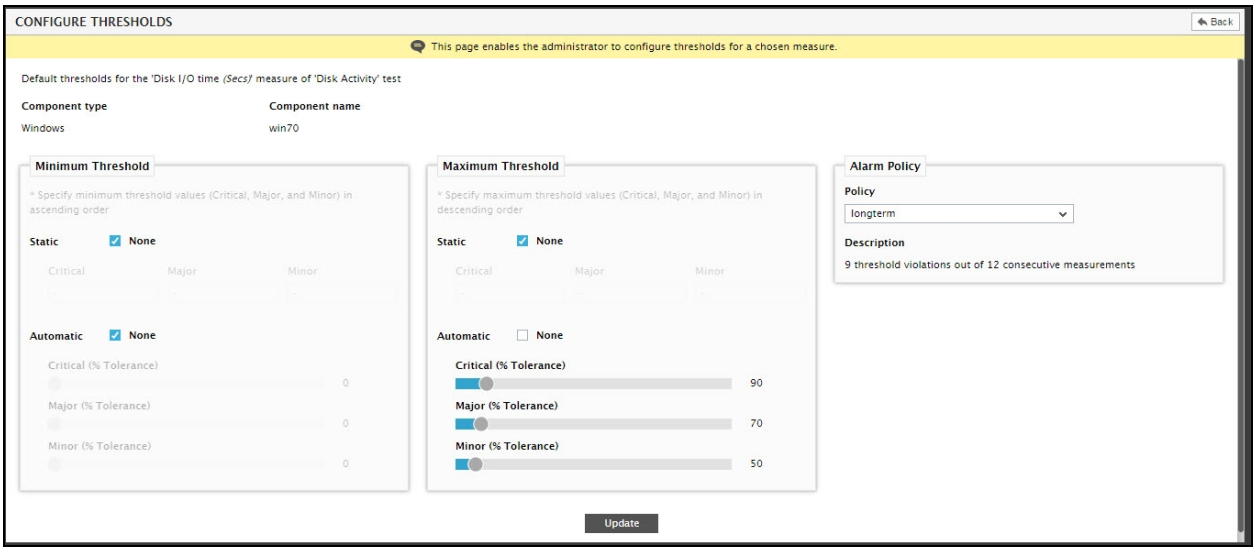

Figure 10.24: Configuring Automatic Maximum Thresholds for the Disk I/O time measure

17. Then, follow the steps below:

    - Use the slider controls that correspond to **Critical**, **Major**, and **Minor** to set tolerance levels for your auto-computed thresholds (see Figure 10.24). This is necessary if you want to induce the eG Enterprise system with some degree of tolerance towards problem conditions. For instance, you may want the eG Enterprise system to tolerate deviations from the auto-computed norms until the Disk I/O time measure reports a value that is over 50% of the auto-computed norm. If the '50 %' limit is crossed, then, you may want the system to generate a **Minor** alert for the Disk I/O time measure. This ensures that the eG Enterprise system disregards – i.e., tolerates - all violations in the range of 1% - 50%. Similarly, you may want eG to generate a **Major** alert if the Disk I/O time exceeds 70% of its auto-computed threshold, and a **Critical** alert if the same measure exceeds 90% of its auto-computed baseline. To configure these tolerance levels, use the **Critical**, **Major**, and **Minor** slider controls in Figure 10.24. Click on the knob in the slider control and keep dragging it to the right or left until the value you

desire is displayed alongside. Dragging the knob to the right will increase the tolerance level and dragging it to the left will reduce it. Like the **Static Thresholds**, the **Automatic Thresholds** should also be in the descending order.

- If you so desire, you can skip one/two tolerance levels. For example, you can set a tolerance level for **Critical** and none for **Major** and **Minor**.

18. If your environment is characterized by long spells of disk inactivity followed by a sudden increase in disk activity, then a combination of **Static** and **Automatic** thresholding – i.e., **Auto-static Thresholding** – is most suitable for the Disk I/O time measure. To configure auto-static maximum thresholds for the Disk I/O time measure, first, deselect the **None** check box against **Static** and **Automatic**, as depicted by Figure 10.25. This will enable both capabilities.

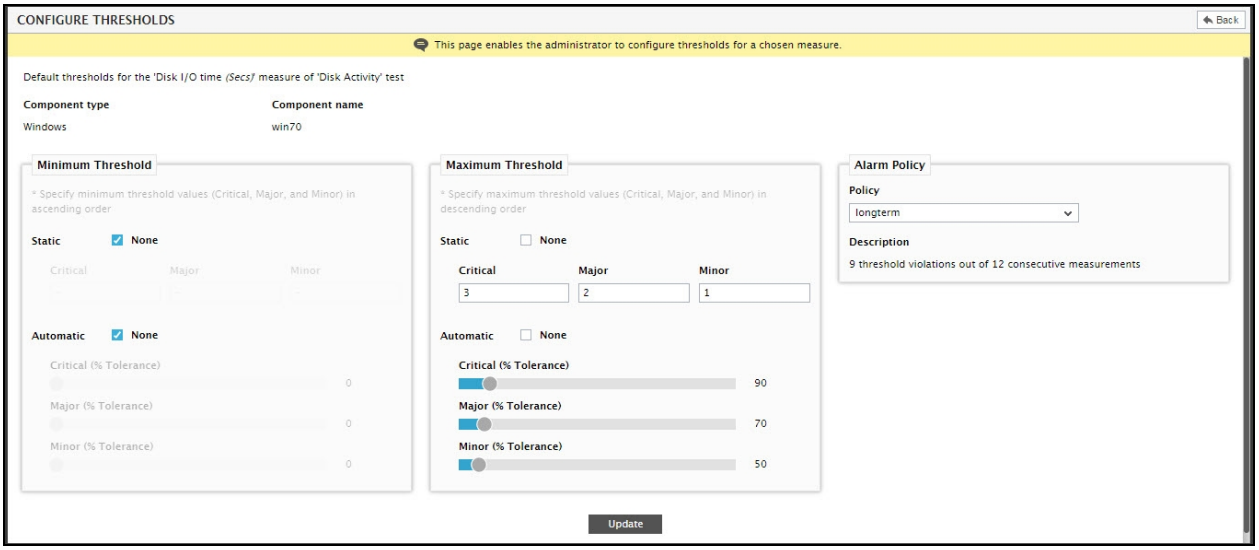

Figure 10.25: Configuring Auto-static Maximum Thresholds for the Disk I/O time measure

19. Then, provide absolute values, in descending order, against the **Critical**, **Major**, and **Minor** text boxes. Likewise, provide a tolerance %, in descending order, using the **Critical**, **Major**, and **Minor** slider controls.

20. When, both automatic and static maximum thresholds are available for a measure, then, at any given point in time, eG will compare the actual value of the measure with the higher of the two (static and automatic) threshold values, and thus detect deviations.

21. Once you are done with configuring thresholds for a measure, assign an **Alarm Policy** to the measure by selecting a **Policy** from the **Alarm Policy** section of Figure 10.25.

Reference:

To know more about alarm policies, refer to Section **10.1** of this document.

Note:

As can be inferred fromFigure 10.25, the thresholds set for every measure is accompanied by an **Alarm Policy** specification. This specification indicates when alarms should be generated by the eG manager. The priority that will be assigned to such an alarm depends upon the threshold configuration and its corresponding alarm policy specification. By default, the following rules are applied when determining the

alarm priority, if the number of violations in a time window matches the alarm policy specification (e.g., 4 out of 6):

- If all violations are critical, then alarm priority would be critical
- If all violations are major, then the alarm priority would be major
- If all the violations are minor, then the alarm priority is minor
- If the number of critical violations is greater than the number of major, and the number of critical violations is greater than the number of minor violations, then the alarm priority is critical
- If the number of major violations is greater than or equal to the number of critical violations, and the number of major violations is greater than the number of minor violations, then the alarm priority is major
- In all other cases, the alarm priority is minor

22. Finally, click the **Update** button to register the changes.
23. Once threshold specifications are updated, Figure 10.26 will appear allowing you to review your changes.

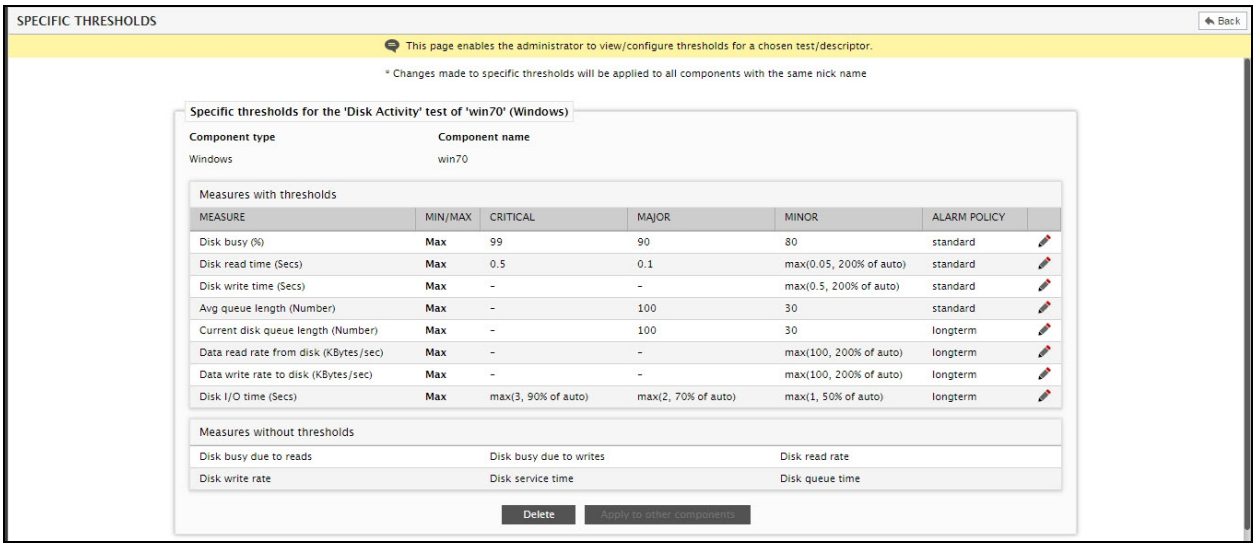

Figure 10.26: Reviewing changes to threshold settings

24. You can proceed to configure thresholds for more measures listed in Figure 10.26 using the same procedure described above. Alternatively, you can also remove all the threshold configurations you see in Figure 10.26 by simply clicking the **Delete** button therein.
25. When configuring component-specific thresholds, Figure 10.26 will additionally allow you to apply the threshold specifications of one component to one/more other components of the same type. To do so, click the **Apply to other components** button in Figure 10.26.

    Note:

    The **Apply to other components** button will be enabled only if more than one component of a type has been managed in the eG Enterprise system.

26. This will invoke Figure 10.27. To quickly get to the components to which the threshold configuration needs to be replicated, use the **View By** drop-down in Figure 10.27. By default, the **Component** option will be

chosen from the **View By** list. Accordingly, all other managed components that are of the same type as that component whose threshold configuration is to be copied, will be displayed in the **Existing components** list box in Figure 10.27. Where too many managed components of the same type exist, you can further narrow your search, by selecting one of the following options from the **View By** list:

- **Zone:** Select this option from the **View By** list and then select the specific **Zone** to which the components of interest to you belong. This will populate the **Existing components** list box with those zone components that are of the same type as the component whose threshold settings are to be copied.
- **Segment**: Select this option from the **View By** list and then select the specific **Segment** to which the components of interest to you belong. This will populate the **Existing components** list box with those segment components that are of the same type as the component whose threshold settings are to be copied.
- **Service**: Select this option from the **View By** list and then select the specific **Service** to which the components of interest to you belong. This will populate the **Existing components** list box with those service components that are of the same type as the component whose threshold settings are to be copied.

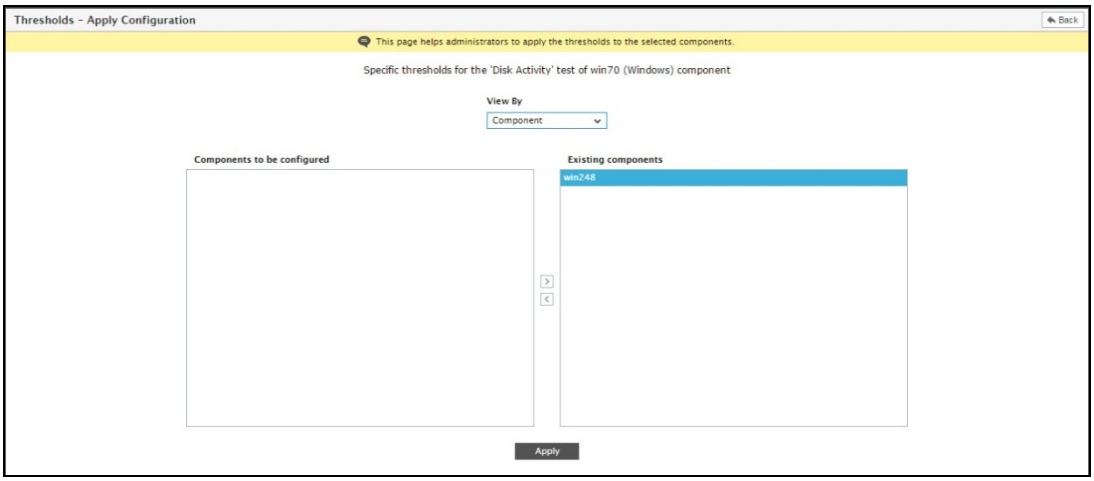

Figure 10.27: Selecting the components to which the threshold configuration is to be replicated

27. Next, from the **Existing components** list, select the components to which the threshold configuration is to be copied. Then, click the **<** button in Figure 10.27.
28. This will transfer the selection to the **Components to be configured** list (see Figure 10.28).

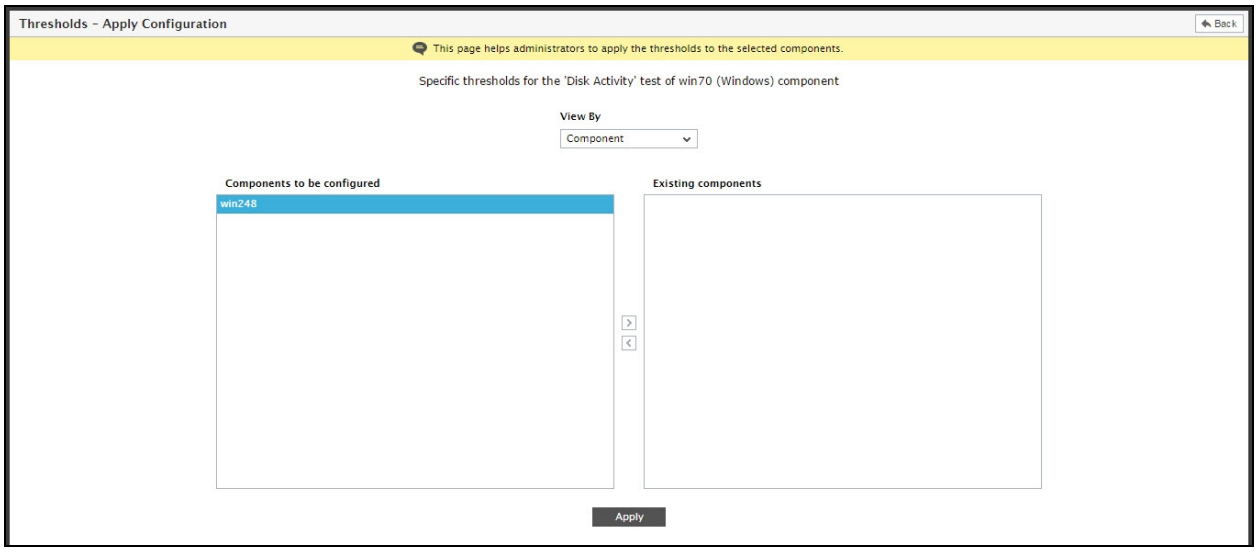

Figure 10.28: Applying the threshold configuration of one component to other components of the same type

29. Finally, click the **Apply** button in Figure 10.28.

Note:

By default, the eG agent initiates the process of test execution and threshold downloads, simultaneously. Though these key processes begin together for a test, at least a minimal difference in their completion times is inevitable. In some cases, this discrepancy could cause real-time deviations to go undetected. Take the case of the UptimeTest for instance. Typically, this test is used to monitor the up time of critical Windows/Unix servers in a target environment, and to alert administrators when a reboot occurs. Now, assume that the **Rebooted** measure of this test has a maximum threshold of 0; the value 0 for this measure indicates that the server has not been rebooted during the last measurement period. If the server were to reboot during a measurement period, then this value would become 1, compelling the eG agent to generate an alert. Say, when the UptimeTest ran for the first time, the test completed execution and reported measures before the thresholds were downloaded. The agent therefore, was unable to detect any abnormalities and hence generated no alerts. To ensure that critical state-changes are not missed, administrators can configure the eG Enterprise system in such a way that tests reporting ultra-critical measures are started after the corresponding thresholds are downloaded. The **eg_tests.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory) consists of an **[IMMEDIATE_THRESHOLDS]** section, which facilitates this configuration. Here, by default, the *UptimeTest* is set to *yes*. This indicates that the UptimeTest executes only after the related thresholds are downloaded by the eG agent. You can append more tests to this list (as indicated below), so that such tests are executed after the eG agent downloads their threshold values.

```
[IMMEDIATE_THRESHOLDS]
UptimeTest=yes
ProcessTest=yes
```

Alternatively, you can configure the thresholds of all tests to be downloaded prior to test execution, by setting the *AllImmediateThresholds* flag in the **[AGENT_SETTINGS]** section of the **eg_tests.ini** file to *yes*. By default, this flag is set to *no*.

10.3 Threshold Groupings

As large infrastructures may have a huge number of components, configuring thresholds for each of the components in the infrastructure can be a laborious process. In order to reduce the workload of the administrators, eG Enterprise allows the creation of threshold component groups and threshold rules. While threshold rules enable the reusability of threshold settings, the threshold component groups, which typically contain a set of components, facilitate the easy and instant application of thresholds to multiple components in an infrastructure. This saves the time and labor involved in individual threshold assignment.

This section explains how to configure threshold component groups and threshold rules, and how to associate a threshold rule with a component group, using an example. In this example, a threshold rule for **Disk Activity** test will have to be applied to a group of Windows servers in an infrastructure. The first step towards this is to create the Windows server group.

10.3.1 Creating a Threshold Component Group

A threshold component group is nothing but a group of components that share the same threshold settings for a particular test / tests. By grouping component, you can ensure that a threshold rule that is assigned to the server group automatically applies to all the components within that group.

To create a group of Windows servers, do the following:

1. Select the **Group** option from the **Thresholds** menu of the **Agents** tile. Figure 10.29 will then appear.

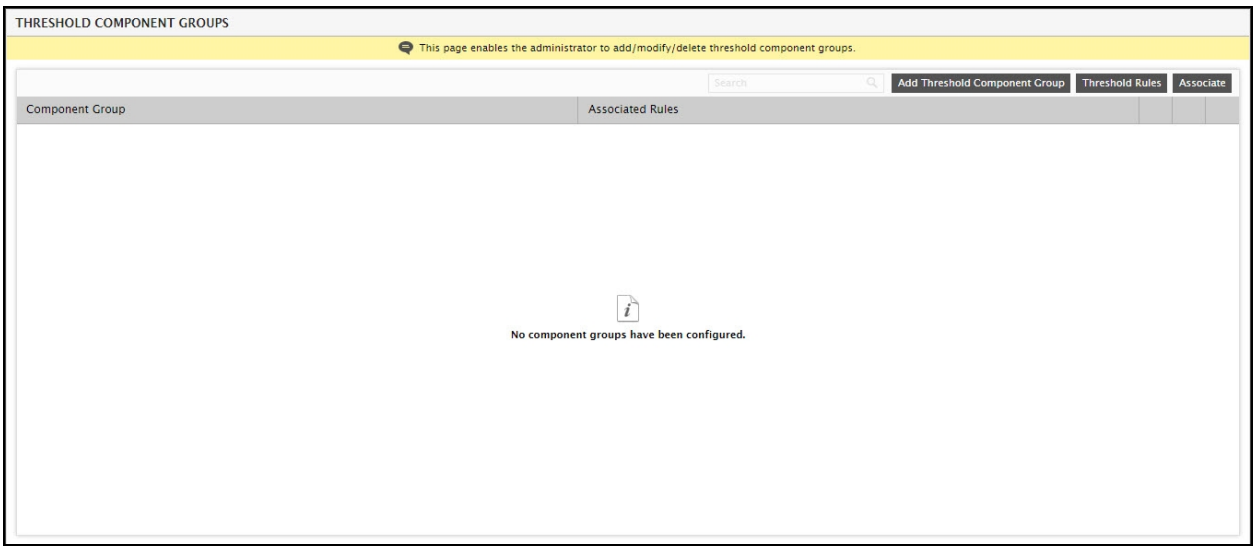

Figure 10.29: Clicking on the Add Threshold Component Group button

2. To create a Windows server group, first click the **Add Threshold Component Group** button in Figure 10.29.
3. When Figure 9.29 appears, first pick the **Component type** to which the components to be added to the group belong. For the purpose of our example, select Windows as the **Component type**. Now, specify the name of the new group in the **Component group name** text box. The **Disassociated components** list box in Figure 10.30 will list all the Windows servers that are not part of any existing group. From this list, select the ones that need to be included in group that is being currently created, and click the **<** button in Figure 10.30.

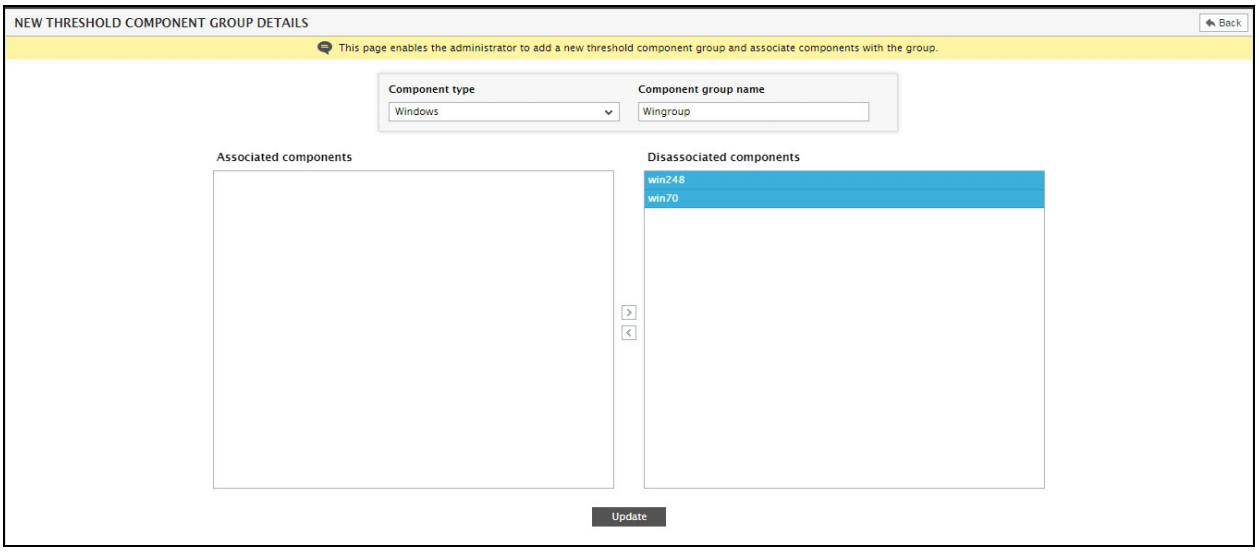

Figure 10.30: Creating a web server group

4. Doing so, will transfer the selection to the **Associated components** list (see Figure 10.31). Similarly, to disassociate servers from the group, select them from the **Associated components** list (see Figure 10.31) and click on the **>** button.

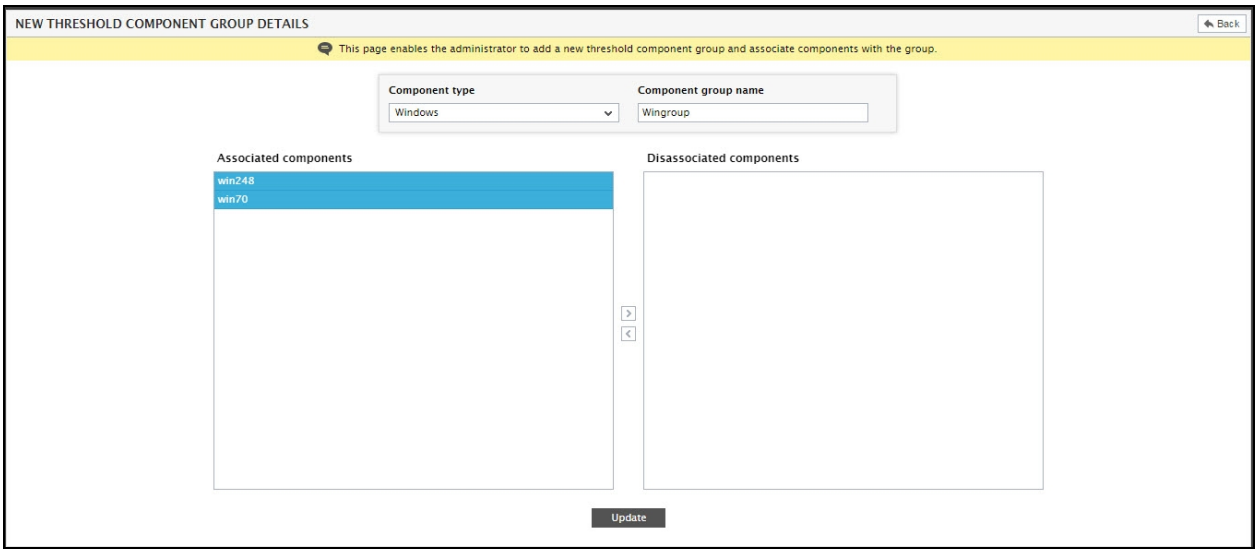

Figure 10.31: Transferring the selection

5. Finally, click the **Update** button. Figure 10.32 will then appear displaying the newly created group.

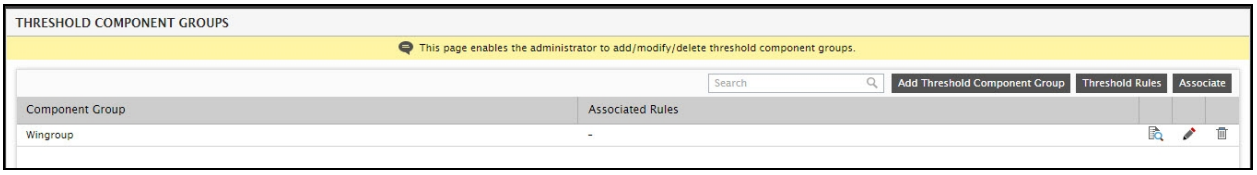

Figure 10.32: New group successfully added

6. To modify the displayed group, click on the icon corresponding to the group in Figure 10.32. To delete the group, click on the icon corresponding to the group. You can also view the composition of a group by clicking the button against the group in Figure 10.32.

10.3.2 Creating a Threshold Rule

The next step is to create a threshold rule. A threshold rule constitutes the typical threshold settings such as the maximum and minimum threshold values, the threshold policy, the alarm policy, etc., for every measure of a test. To create a threshold rule for the **Disk Activity** test in our example, do the following:

1. Click on the **Threshold Rules** button in Figure 10.32. Figure 10.33 will then appear.

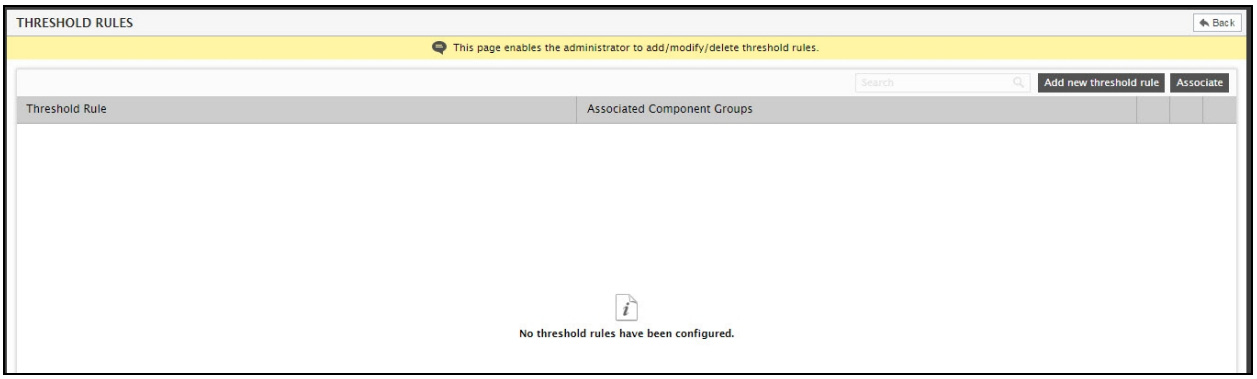

Figure 10.33: A message indicating that no threshold rules pre-exist

2. To add a new rule, click on the **Add Threshold Rule** button in Figure 10.33. Figure 10.34 will then appear. To configure a threshold rule for the **Disk Activity** test in our example, first select Windows as the **Component** type, pick **Disk Activity** from the **Test name** list, and specify a unique name for the new rule in the **Rule name** text box of Figure 9.33. Then, click the icon corresponding to any measure displayed in Figure 9.33 to configure maximum and/or minimum thresholds for that measure (see Figure 10.35).

   Reference:

   To know how to configure thresholds for a measure, refer to Section Section **10.2.2.2** of this chapter.

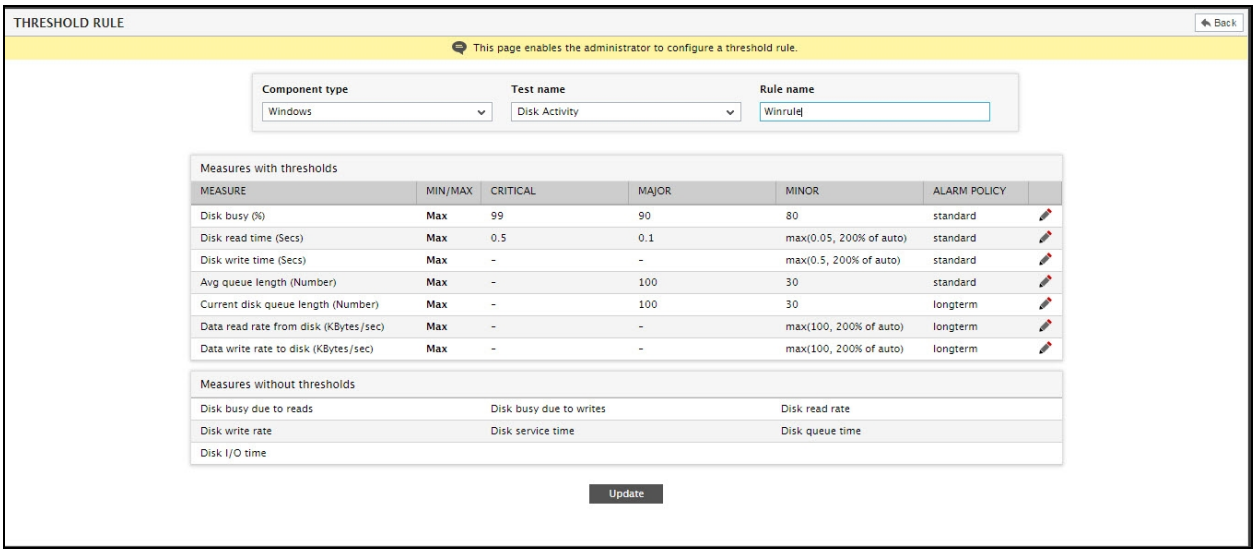

Figure 10.34: Creating a threshold rule

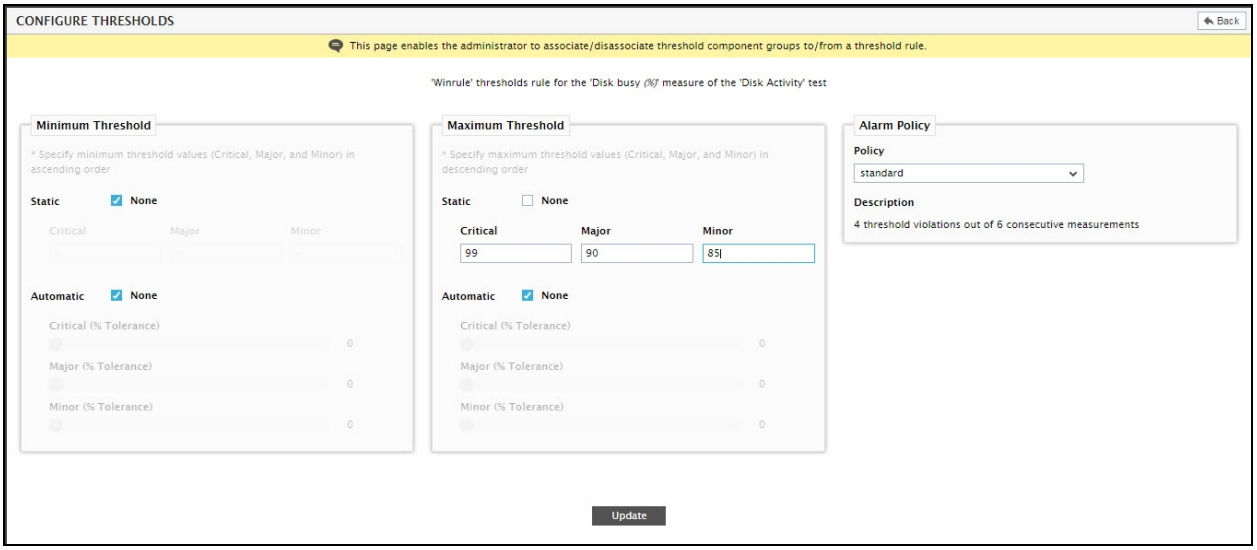

Figure 10.35: Configuring thresholds for a measure of the Disk Activity test in our example

3. Finally, click the **Update** button in Figure 10.35.
4. Upon successfully updating, Figure 10.36 will appear listing the newly created threshold rule.

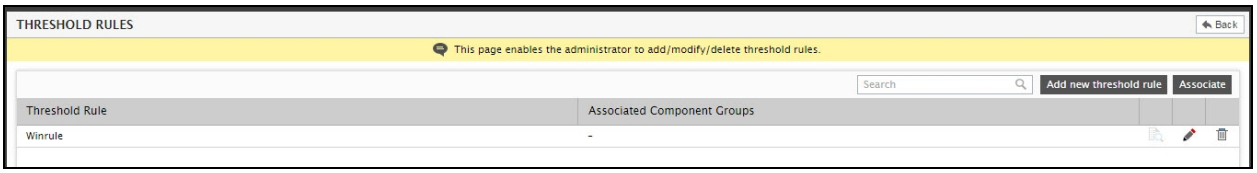

Figure 10.36: The newly created threshold rule displayed in the THRESHOLD RULES page

10.3.2.1 Associating a Threshold Rule with Threshold Component Groups

Next, proceed to associate the configured threshold rule with the Windows server group that was earlier created, so that the thresholds set within the rule automatically apply to each of the Windows servers in the group. To achieve this, do the following:

1. Click on the **Associate** button in Figure 10.36. Figure 10.37 will then appear.

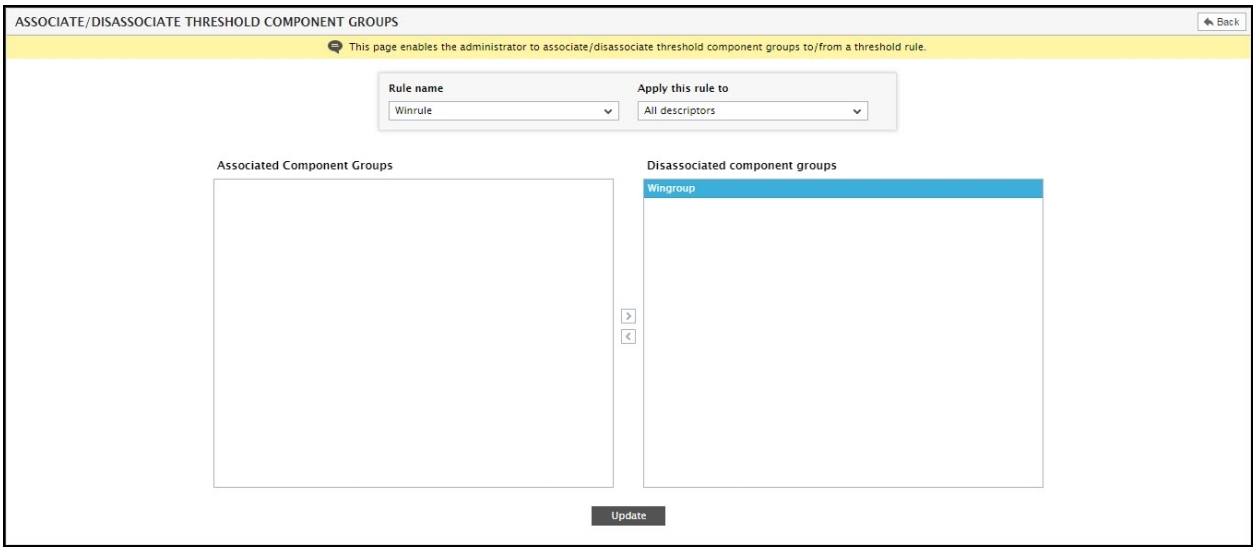

Figure 10.37: Selecting the web server group to be associated with the threshold rule

2. From Figure 10.37, select the **Rule name** that you want to associate with the Windows server group you created earlier. By default, the **Apply this rule to** list box will have the *All Descriptors* option chosen, indicating that by default, the chosen rule will be applied to all the descriptors that have been enabled for the **Disk Activity** test in our example. To ensure that the threshold settings defined in the rule govern the state of specific descriptors only, select the *Specific Descriptors* option from the **Apply this rule to** list box of Figure 10.37. This way, administrators are saved the trouble of manually repeating the threshold configurations for every descriptor of a test.
3. Next, select the web server group from the **Disassociated component groups** list, and click the **<** button to assign the chosen threshold rule to the selected threshold component group.
4. When this is done, the selected threshold component group will move to the **Associated Component Groups** list as shown by Figure 10.38.

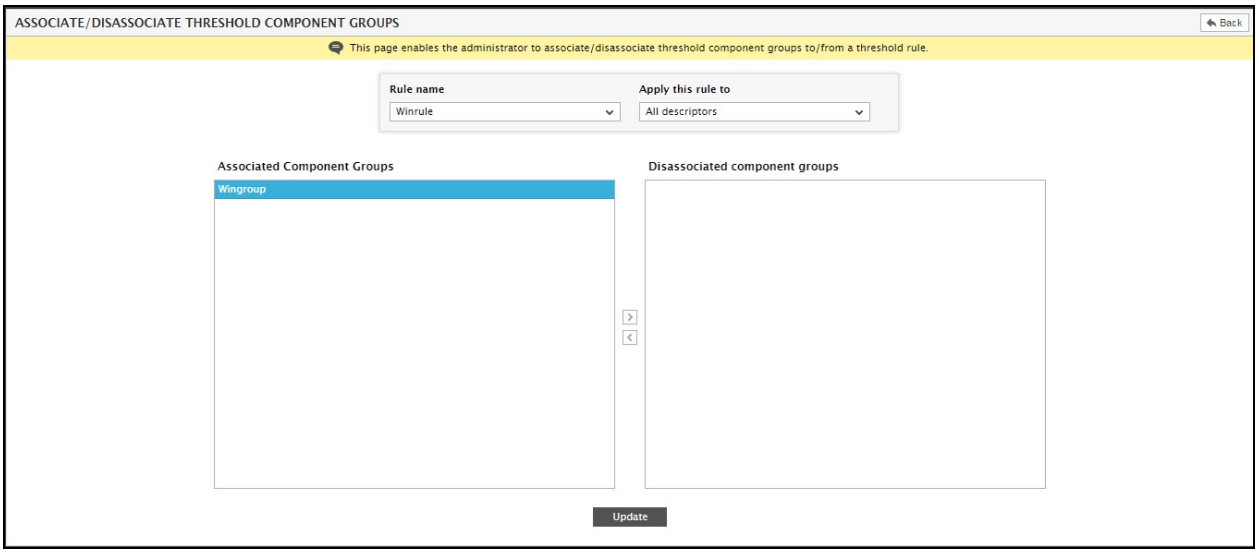

Figure 10.38: Associating the web server group with the threshold rule

5. To disassociate the web server group from the rule, select it from the **Associated Component Groups** list and click the **>** button in Figure 10.38.
6. Finally, click the **Update** button in Figure 10.38.
7. If the *Specific Descriptors* option had been selected from Figure 10.38, then clicking on the **Update** button will invoke 10.3.2 using which specific descriptors can be associated with the threshold rule. In the case of the **Disk Activity** test in our example, the descriptors are the disk partition on the Windows servers in the threshold component group. To associate descriptors with a rule, choose the required descriptors from the **Disassociated Descriptors** list in Figure 10.39, and click on the **<** button to associate them with the threshold rule displayed against **Rule name**. To disassociate descriptors, select them from the **Associated Descriptors** list and click on the **>** button.

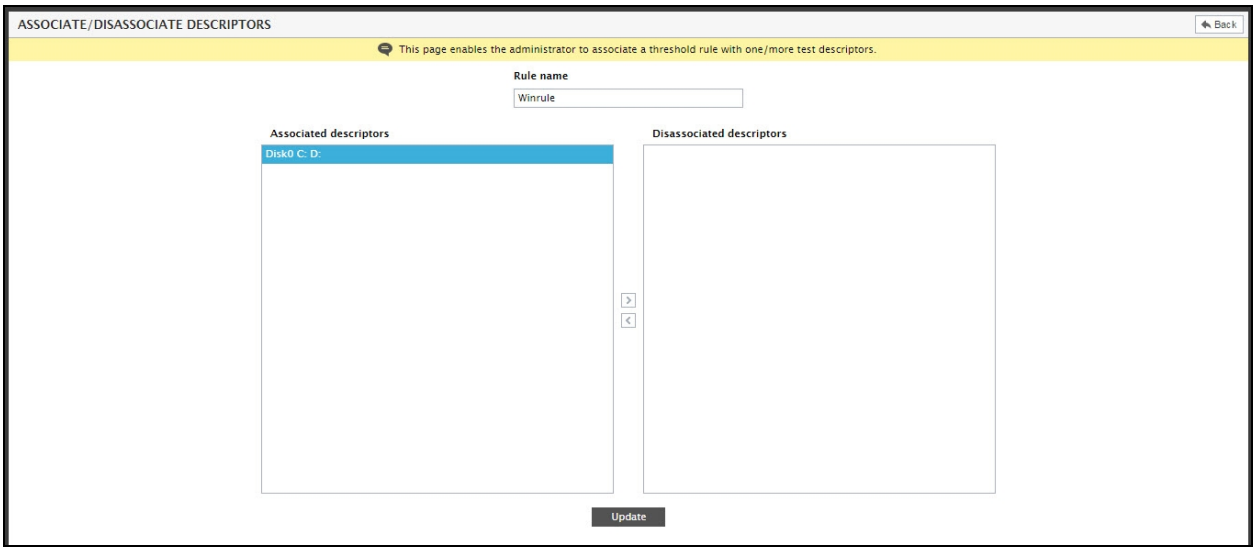

Figure 10.39: Associating specific descriptor with the threshold rule

8. Finally, click the **Update** button in Figure 10.39.

Note:

By default, the **updateDefaultThreshToDb** parameter in the **[MISC_ARGS]** section of the **eg_services.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory) is set to **False**. This ensures that the database is not updated with the changes made to the default thresholds. To make sure that all default threshold changes are reflected in the database, set the **updateDefaultThreshToDb** parameter to **True**. However, the default setting for the **updateSpecificThreshToDb** parameter in the `[MISC_ARGS]` section is **True**. This makes sure that the database is, by default, updated with the changes made to the specific thresholds. To make sure that all specific threshold changes are not reflected in the database, set the **updateSpecificThreshToDb** parameter to **False**.

Note:

- If a descriptor threshold is specified, this overrides all other threshold settings for the specific descriptor. Likewise, if a specific threshold is set for a component, it takes priority over the group and default threshold settings. If neither specific nor descriptor thresholds are set, the group threshold settings (if present) will apply. Otherwise, the default threshold settings will apply.
- By default, a server can belong to only one threshold component group at a time. Accordingly, the **AllowComponentsInMultipleGroups** flag in the **[MISC_ARGS]** section of the **eg_services.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory) is set to **false** by default. If you want to add a server to multiple threshold component groups, then set the **AllowComponentsInMultipleGroups** flag to **true**. With this setting, it will be possible to define different component groups - e.g., one for each test - and to map this group to different threshold rules. A caveat to note is that when creating threshold component groups, the administrator must explicitly take care to ensure that the same component is not associated with multiple threshold rules for the same test.

10.4 Global Thresholds

While monitoring large environments, some tests executed by the eG agent report statistics on hundreds of descriptors. For example, the User Profile test reports the profile size of each and every Citrix or Terminal server user of a server. Likewise, the Windows Services Status Test reports on the availability of each and every service of a Windows system. For such tests, storage of the threshold values for each hour for each descriptor can result in significant disk space usage in the eG database. In order to enable administrators to optimize database usage for tests that do not use the automatic threshold computation (i.e., relative thresholding) capability, eG Enterprise offers the **Enable/Disable Global Thresholds** page.

To access this page, select the **Global** option from the **Thresholds** menu of the **Alerts** tile. Figure 10.40 will then appear.

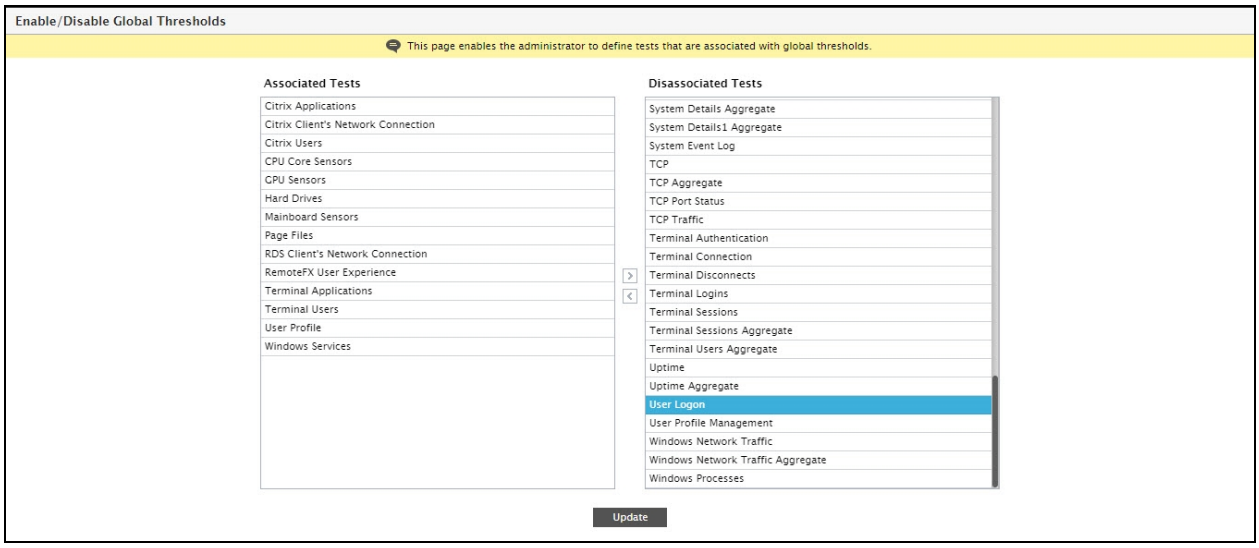

Figure 10.40: Selecting tests for which GLOBAL THRESHOLDS apply

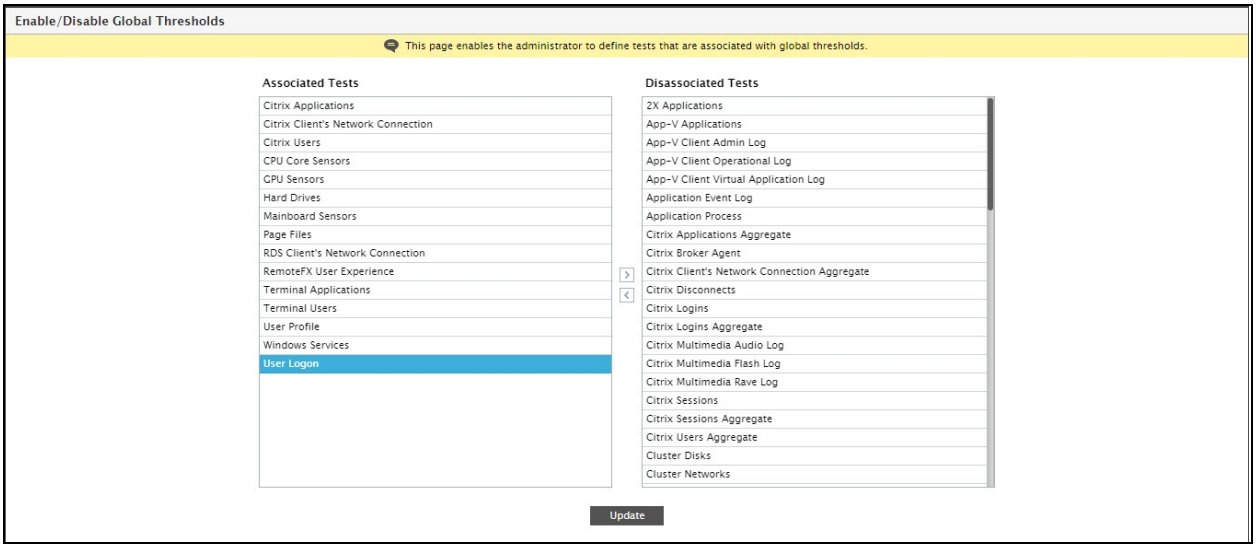

Figure 10.41: Transferring the test for which Global thresholds should be applied

The **Disassociated Tests** list in Figure 10.40 displays the tests mapped to all the managed components in the target environment. From this list, select those tests for which thresholds need not be stored in the database. For example, if the **User Logon** test has to be configured to not store thresholds in the database, select **User Logon** test from the **Disassociated Tests** list as shown by Figure 10.40, click on the **<** button, and move the selected test to the **Associated Tests** list. Once one/more tests are so associated (see Figure 10.41), the eG manager retrieves thresholds for such tests from their configuration files, and does not store any thresholds for these tests in the eG database. As a tradeoff, the thresholds for these tests apply to all the servers being managed, and cannot be set specifically for every server. Hence, **Specific Thresholds** cannot be computed for these tests. Moreover, when configuring the default threshold for these tests, the threshold policy has to be either "absolute" or "none". The threshold settings can be set differently for different descriptors of a test for which Global threshold is set. This setting reduces the number of false alerts that the monitoring system can

generate. For example, the inside view of a VM reports the CPU usage per processor and also an overall summary value. By setting descriptor -wise thresholds, the thresholds can be disabled for all processors of a VM and only the summary value can be set to generate alerts – this way an administrator is alerted only if the VM as a whole is experiencing high CPU utilization.

10.5 Viewing Thresholds

At any point in time, administrators can view the configured thresholds by selecting the **View** option from the **Thresholds** menu of the **Agents** tile. Doing so will invoke Figure 10.42.

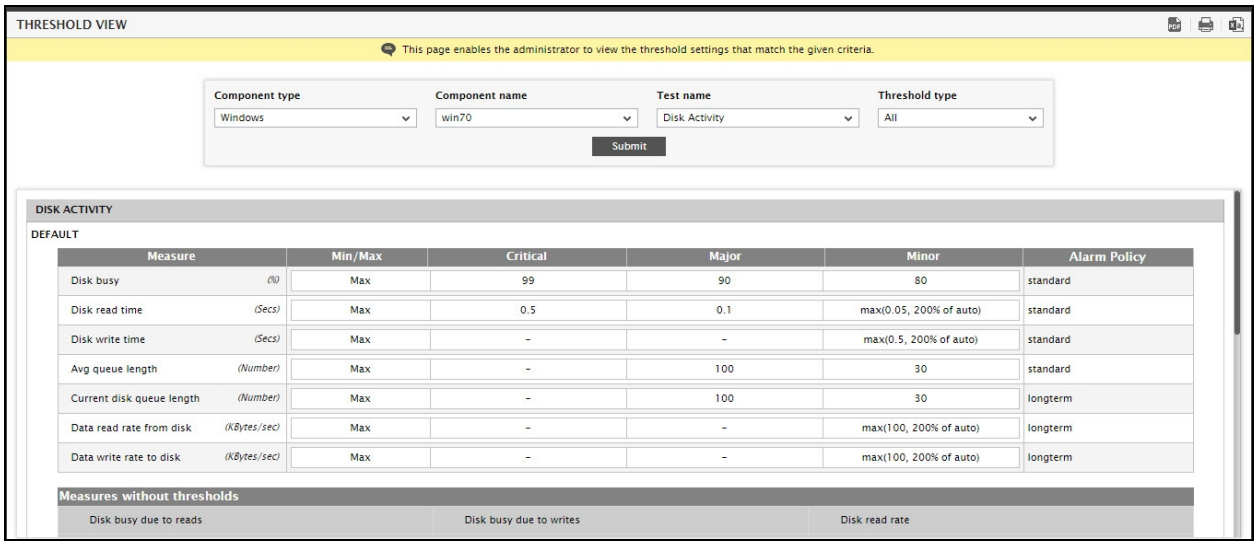

Figure 10.42: Viewing configured thresholds

To review threshold settings, do the following:

1. Select a **Component type**.
2. All components of the chosen type will then populate the **Component name** list. Select the component for which threshold specifications are to be reviewed. To view the threshold settings for all managed components of the chosen type, select *All* from the **Component name** list.
3. All the tests that run on the chosen component will then be listed in the **Test name** drop-down list. Pick the test for which threshold specifications are to be viewed from this list. To view the threshold settings of all tests mapped to the chosen component, select *All* from the **Test name** list.
4. Then, from the **Threshold type** drop-down. select the type of thresholds you want to view – the options are *Default*, *Specific*, *Descriptor*, *Group*, *Global*, and *All*. To view the threshold settings of all threshold types, pick All from this list.
5. Then, click the **Submit** button to view the threshold settings that fulfill the specified criteria.
6. You can save the threshold details displayed as a PDF file, by clicking the PDF icon at the right, top corner of Figure 10.42.

7. To print the threshold information, click the icon at the right, top corner of Figure 10.42.
8. To save the threshold details as a CSV file, click the icon at the right, top corner of Figure 10.42.

10.6 Maintenance

During times when an IT infrastructure is under maintenance, it is but natural that few/all of the monitored components are rendered unavailable. This in turn, could cause the monitoring tool to generate a plethora of alarms indicating a "non-existent" problem situation.

Similarly, some administrators might want network-related alerts to be suppressed during the non-working hours of a company, as such alerts are inevitable during that time of the day. Likewise, if administrators are deliberately bringing the web server process down for a brief period, they might consider insignificant the alerts pertaining to the *Processes* test on the web server during that period.

In order to prevent the meaningless generation of alarms during routine maintenance activities, eG Enterprise allows administrators to:

- Define maintenance policies based on the periodicity of the maintenance procedures performed on the environment
- Group the defined policies
- Associate the groups with components/hosts/tests in the target environment

These features enable administrators to switch off eG alarms for specific components/tests during maintenance periods.

The sections to come discuss how to create, modify, and delete policies.

10.6.1 Defining a Quick Maintenance Policy and Associating it with Hosts / Tests / Components

eG Enterprise provides a single, easy-to-use interface for creating a maintenance policy and associating the policy to specific infrastructure elements.

To create a maintenance policy, do the following:

1. Select **Maintenance Policies** from the **Alerts** menu. This will open the **MAINTENANCE POLICIES** page. If maintenance policies pre-exist, then this page will reveal the list of maintenance policies, a summary of the policy specifications, and the components/hosts/descriptors/tests to which each policy applies (see Figure 10.43).

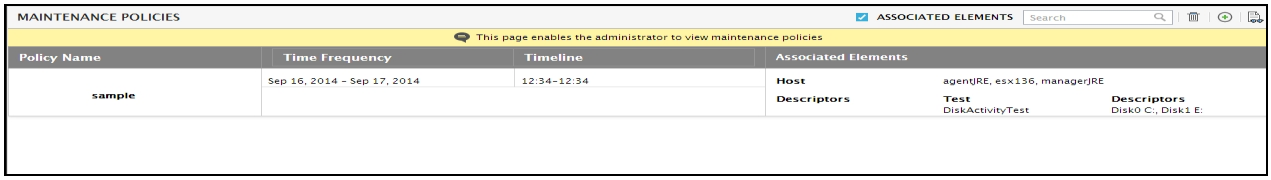

Figure 10.43: The MAINTENANCE POLICIES page

2. If you do not want to view the list of elements to which a policy applies, then simply uncheck the **ASSOCIATED ELEMENTS** check box at the top of the **MAINTENANCE POLICIES** page.
3. To add a new policy, click on the encircled '+' button in the policy tool bar at the right, top corner of Figure 10.43.
4. In Figure 10.44 that appears, provide a name for the new policy against the **Policy name** text box.
5. From the **Frequency** list box of Figure 10.44, select any of the following options:

   - **First day of month:** If the maintenance procedure is performed on the first day of every month, then, select this option.
   - **Last day of the month:** If the maintenance procedure is performed on the last day of every month, then, select this option.
   - **Daily:** If the maintenance procedure is performed every day, then, select this option.
   - **Day of week:** If the maintenance procedure is performed on a particular day in a week, then, select this option. When this option chosen, a list box appears alongside that allows you to specify the exact day of the week on which maintenance is to be performed.
   - **Date:** If the maintenance procedure is performed on a specific date, then, select this option. Then, using the **Calendar** option that appears alongside, specify the exact date on which maintenance is to be performed.

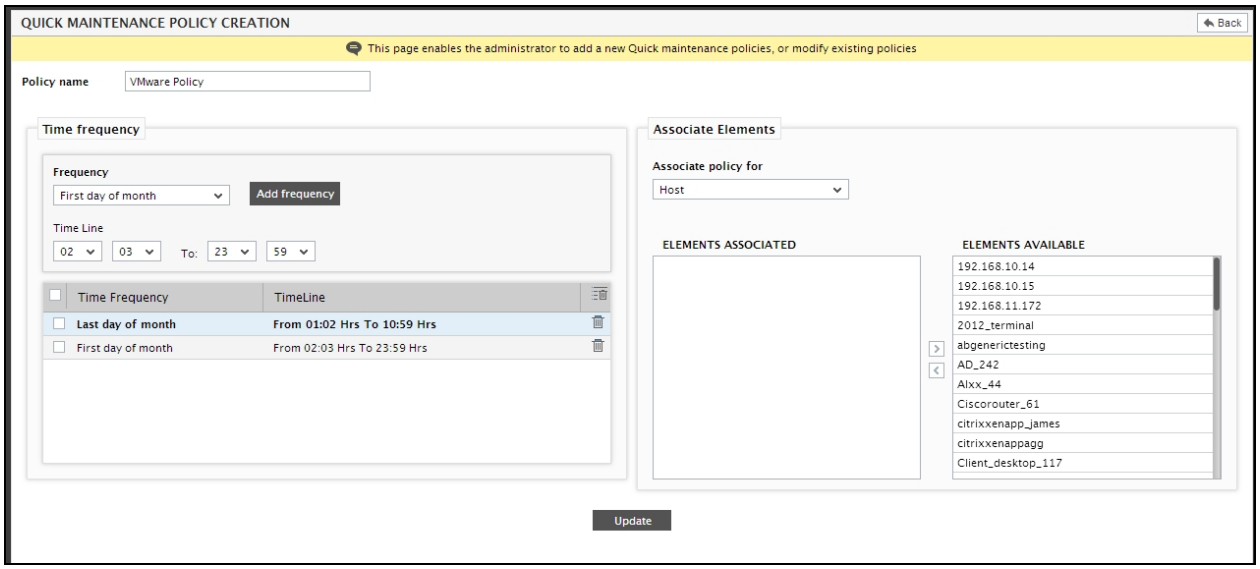

Figure 10.44: Creating a Quick maintenance Policy

6. Next, specify the duration of maintenance using the `From` and `To` fields against **Time Line** (see Figure 10.44).
7. Then, add these specifications using the **Add Frequency** button (see Figure 10.44). Similarly, multiple specifications can be added as part of a single policy.

8. To delete a particular entry, select it from the list by clicking on the check box corresponding to it in Figure 10.44, and click the **Delete** icon corresponding to it. To mark all the entries for deletion, simply select the check box at the top of the specification table (i.e., the check box alongside the column heading **Time Frequency**), and then click the **Delete** icon.

9. Now, using the **Associate Elements** section, you can associate the created policy to a `Host`, **Component**, **Test**, **Test for Host**, **Descriptor** , **Descriptor for Component**, or, **Test for Component.** To achieve this, select the relevant option from the **Associate policy for** list in Figure 10.45.

10. If the **Host** is chosen, then the list of hosts being monitored will appear in the **ELEMENTS AVAILABLE** list (see Figure 10.45).

11. Then, click on the **<** button to add the policies selected from the **ELEMENTS AVAILABLE** list to the **ELEMENTS ASSOCIATED** list (see Figure 10.45). You can also add individual policies to the **ELEMENTS ASSOCIATED** list by just double-clicking on the policy name in the **ELEMENTS AVAILABLE** list . If a maintenance policy is associated with a host, then such a policy will suppress the alerts generated by all the layers of all the applications executing on that host (i.e., applications managed by eG Enterprise that share the same IP address-nick name combination).

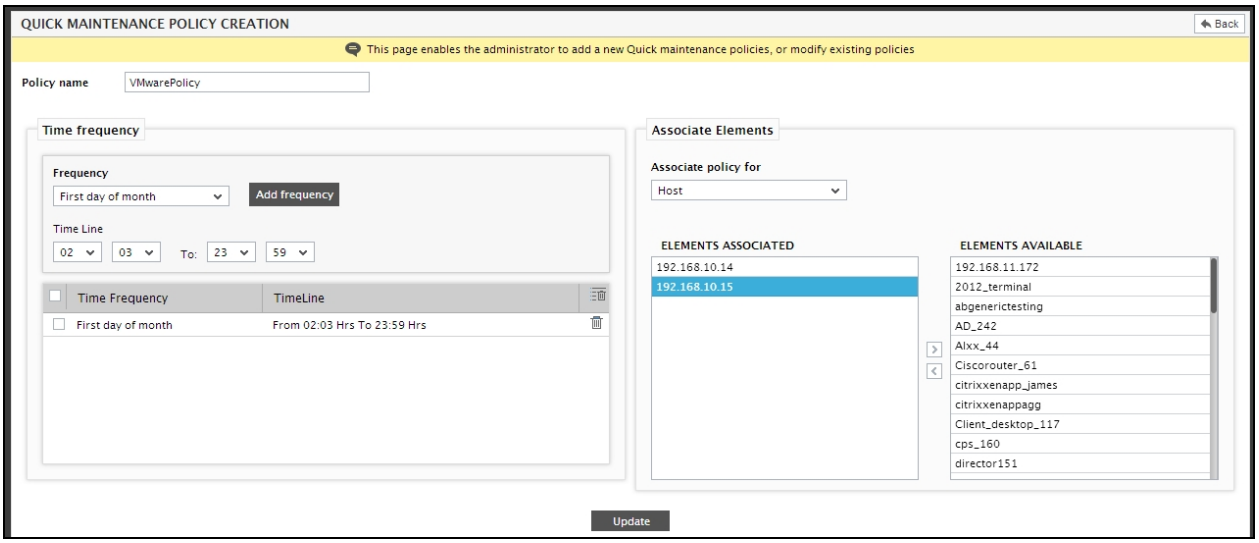

Figure 10.45: Associating a policy with a host

12. To remove an entry from the **ELEMENTS ASSOCIATED** list, select it from the list, click on the **>** button, or just double-click on the entry you need to remove.

13. Besides hosts, a maintenance policy can also be associated with one/more components. By associating a maintenance policy with a component, administrators can automatically suppress those alerts that are raised on the application-specific layers of that component - in other words, such a policy will not govern the alerts related to the host layers of that component.

    To map a maintenance policy to a component, first, select the **Component** option from the **Associate policy for** list. To save you the time and trouble involved in selecting the required components from a broad component window, eG Enterprise provides you with multiple filter options. You can narrow your search

for components using any one of the filter criteria available in the **Selectcomponents by** list. These criteria are as follows:

- Zone
- Segment
- Service
- Component type

14. If you want to associate one/more components included in a particular zone with a policy, then select **Zone** from the **Select components by** list, and then choose the desired zone from the `Zone` list as depicted in Figure 10.46. This will list all the components in the chosen zone in the **ELEMENTS AVAILABLE** list. Select the component(s) from this list and click the **<** button, or double-click on the desired component to add it to the **ELEMENTS ASSOCIATED** list. Finally, click the **Update** button.

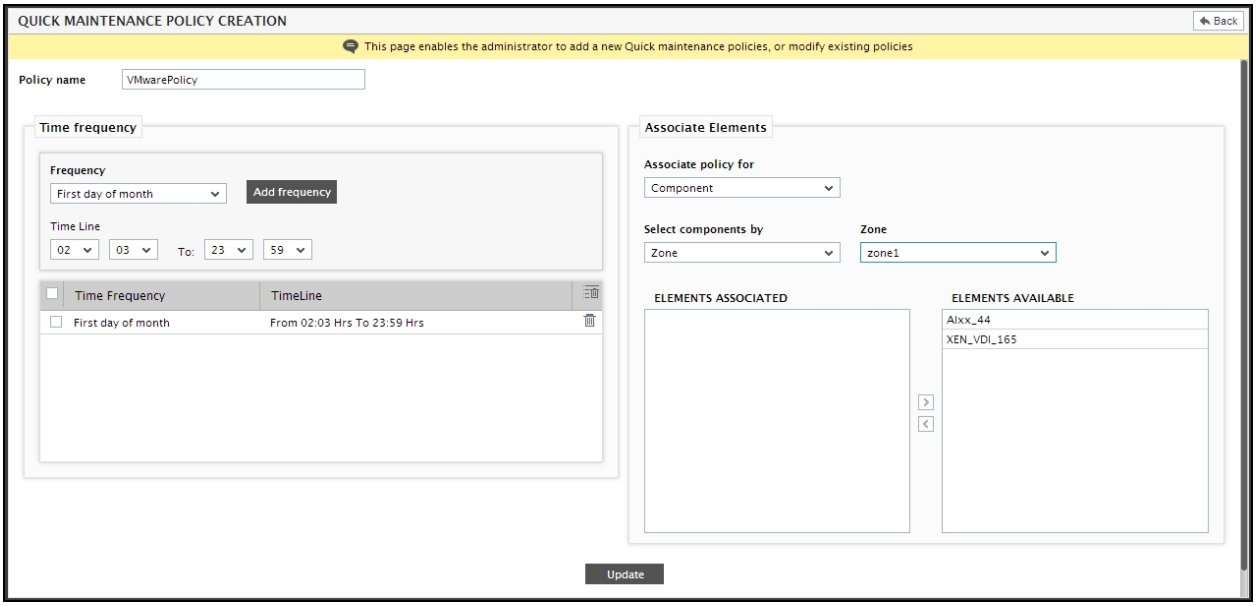

Figure 10.46: Associating a policy to the components in a particular zone

15. Similarly, if your maintenance policy has to be applied to one/more components that are part of a segment, select **Segment** from the **Select components by** list, and choose the desired segment from the **Segment** list (Figure 10.46). Doing so ensures that all the components of the chosen segment are listed in the **ELEMENTS AVAILABLE** list. Select the component from this list and click the **<** button, or double-click on a component to add it to the **ELEMENTS ASSOCIATED** list. Finally, click the **Update** button.

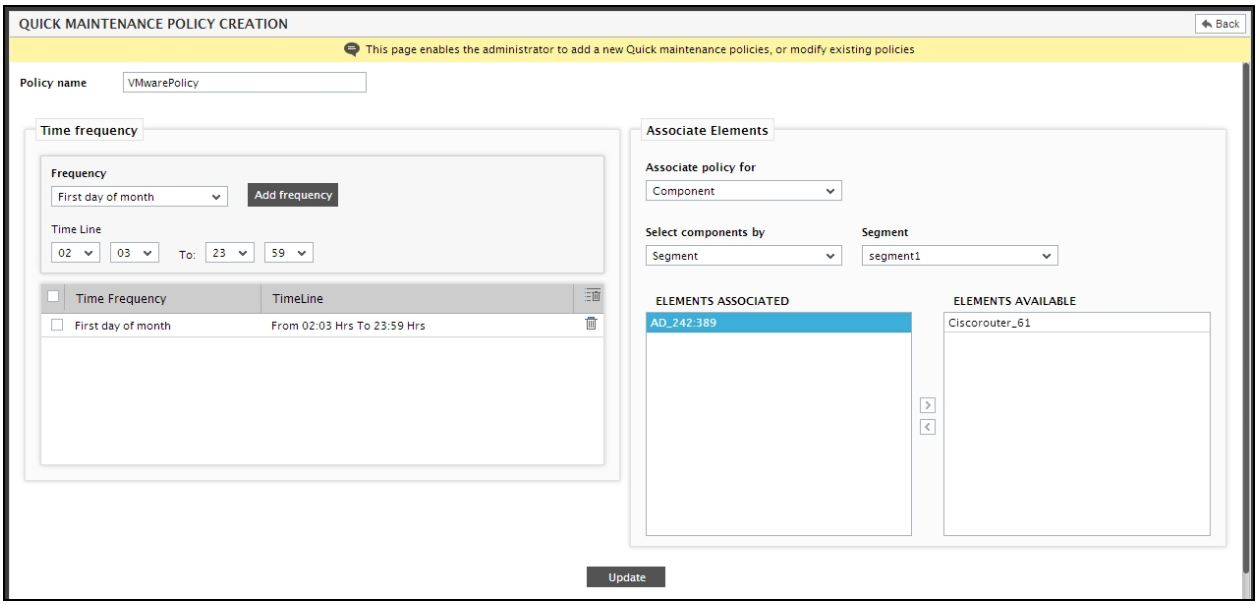

Figure 10.47: Associating a policy with the components in a particular segment

16. If **Service** is selected from the **Select components by,** you will be allowed to pick a **Service**, and then associate a maintenance policy with one/more components engaged in the delivery of that service as depicted in Figure 10.47. Once a service is chosen, the service components will be listed under **ELEMENTS AVAILABLE.** Select the component(s) of interest to you from this list and click the **<** button to transfer your selection to the **ELEMENTS ASSOCIATED** list. Alternatively, you can double-click on individual components to move them to the **ELEMENTS ASSOCIATED** list. Finally, click the **Update** button.

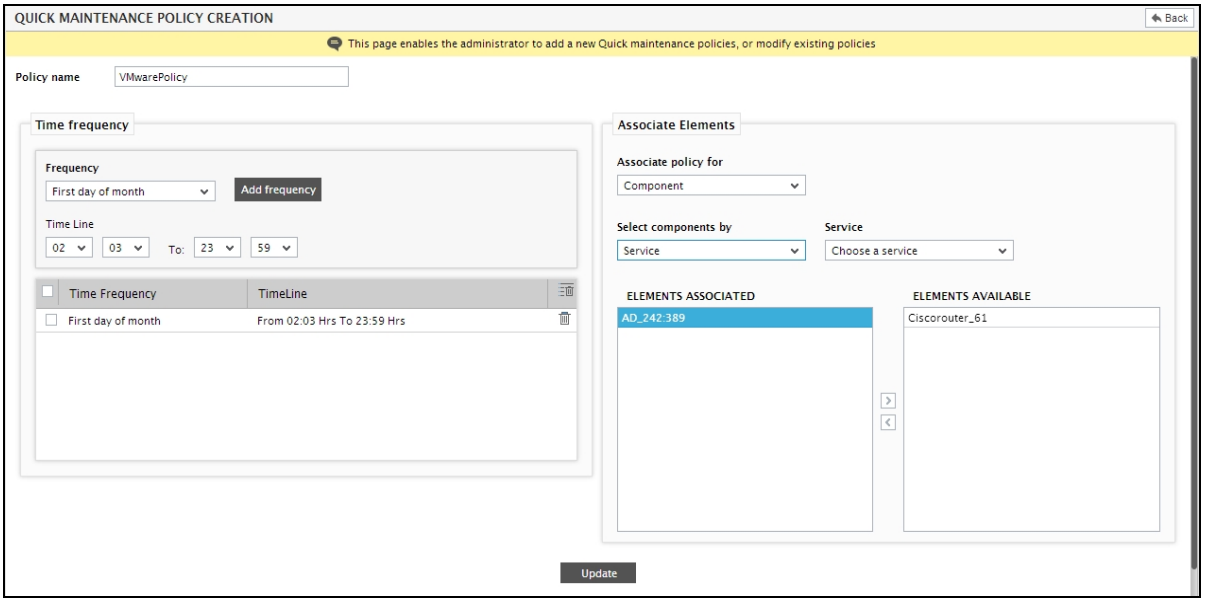

Figure 10.48: Associating a policy with the components engaged in the delivery of a service

17. Likewise, to associate one/more components of a specific type with a policy, first select **Component Type** from the **Select components by.** Next, select the desired component type from the **Component Type** field

(see Figure 10.48). This in turn results in the display of all managed components of that particular type in the **ELEMENTS AVAILABLE** list. Select the component from this list and click the **<** button or double-click on individual component to associate the selected component to the **ELEMENTS ASSOCIATED** list. Finally, click the **Update** button.

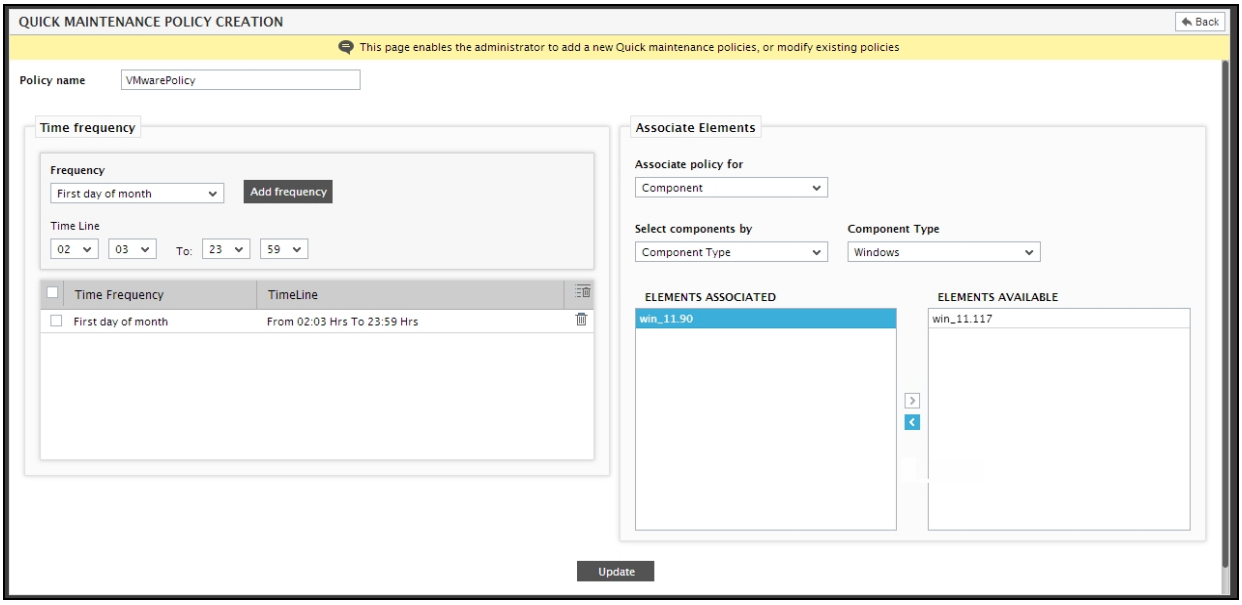

Figure 10.49: Associating a policy with the components of a particular type

18. Now, if **Test** is chosen from the **Associate policy for** list in Figure 10.49, then the **ELEMENTS AVAILABLE** list will, by default, be populated with the complete list of tests that are enabled and are running for all the managed components in the environment (see Figure 10.49). In environments where a multitude of components are being monitored, a long list of tests may appear in the **ELEMENTS AVAILABLE** list. Administrators may have to scroll down the list relentlessly to pick the tests of interest to them. To condense the **ELEMENTS AVAILABLE** list so that administrators have a less number of tests to choose from, use the **Filter By** drop-down. This drop-down is populated with the complete list of managed component types in the environment. Select a particular component type from the **Filter By** list so that the **ELEMENTS AVAILABLE** list displays only those tests that execute on the managed components of the chosen type.

    To suppress alerts generated by specific tests, first, select the required tests from this list and click the **>** button. Alternatively, you can double-click on individual tests to add them to the **ELEMENTS ASSOCIATED** list. Finally, click the **Update** button. A maintenance policy, once associated with a test, will automatically apply to all components on which the chosen test executes.

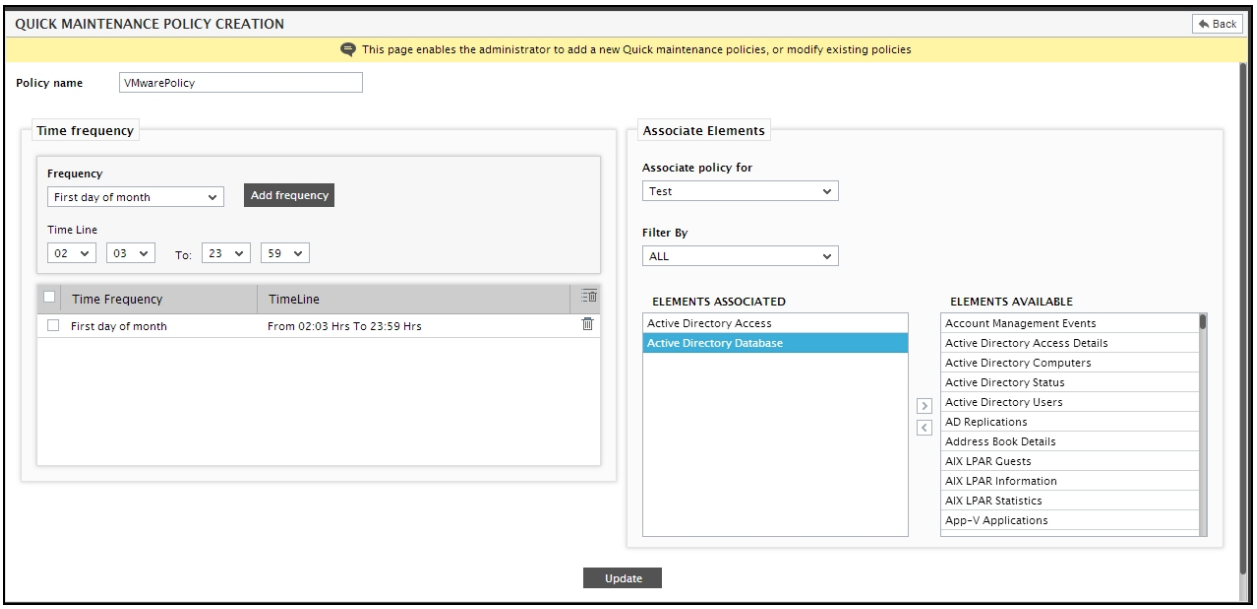

Figure 10.50: Associating a policy with a test

19. Maintenance policies can also be set for individual descriptors. By selecting the **Descriptor** option from the **Associate Policy for** list, you can assign maintenance policies to specific descriptors, and thus suppress the alerts pertaining to such descriptors across all the components to which they apply. For instance, all the managed MS SQL servers in an environment could have been shutdown as part of routine maintenance. To avoid been repeatedly alerted to the non-availability of the *SqlServer* process on multiple servers, administrators may want to suppress these process-related alerts across all the monitored SQL servers in the environment. The **Descriptor** option enables the administrator to achieve this. If the **Descriptor** option is chosen, all the descriptor-based tests will be available for selection in the **Test** list that appears. Selecting a test from this list will populate the **ELEMENTS AVAILABLE** list with the descriptors that currently exist for the chosen test. Select the descriptor of interest to you from the list and click the **<** button, or double-click on the descriptor to add it to the **ELEMENTS ASSOCIATED** list. Once this is done, the maintenance policy will be automatically assigned to all components that report measures for the chosen descriptorFinally, click the **Update** button.

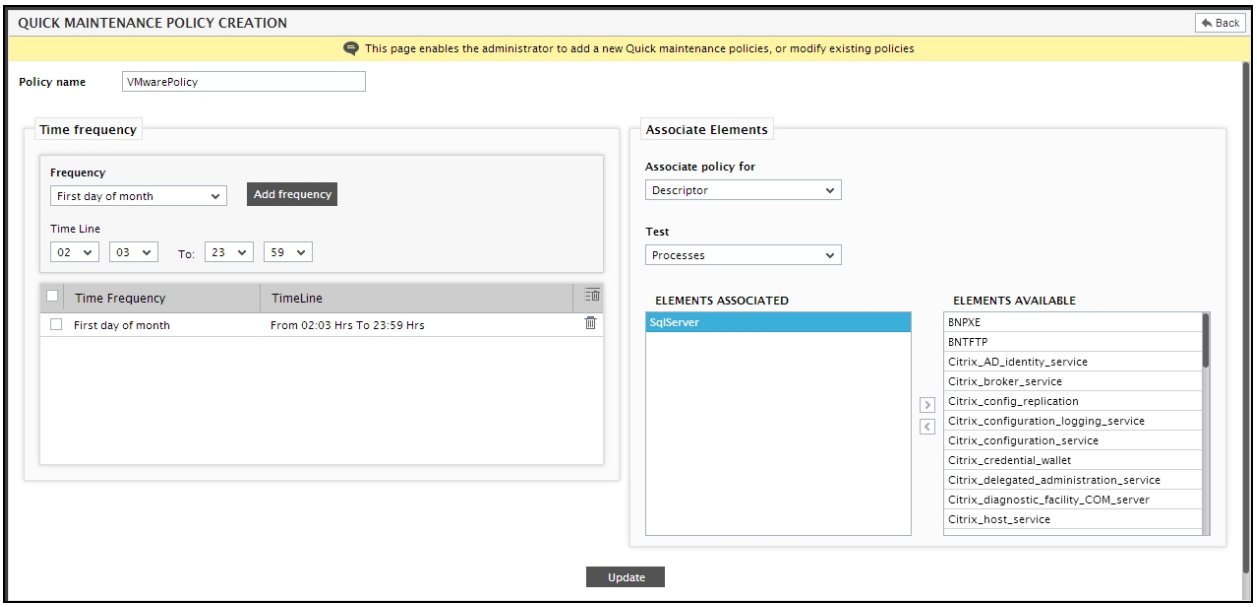

Figure 10.51: Associating a policy with one/more descriptors

20. You can also suppress alerts pertaining to one/more descriptors of a specific component alone. For instance, if one/more VMs on a particular virtual host are stopped for routine maintenance, then, you may not want the eG Enterprise system to unnecessarily alert you to the non-availability of such VMs. To achieve this, you need to first select the **Descriptor For Component** option from the **Associate policy for** list. Next, select the **Component** of interest to you (in the case of our example, this will be the virtualized component that hosts the VMs). Once the descriptor-based tests mapped to the chosen component populate the **Test** list, pick the test for which alarm suppression should occur. All the descriptors supported by that test will then appear in the **ELEMENTS AVAILABLE** list. From this list, choose the descriptors for which alerts are to be suppressed and click the **<** button to transfer the selection to the **ELEMENTS ASSOCIATED** list. Alternatively, you can double-click on your selection in the **ELEMENTS AVAILABLE** list to perform the transfer. Finally, click the **Update** button to associate the maintenance policy being configured with the descriptors in the **ELEMENTS ASSOCIATED** list. If you want to automatically apply the maintenance policy to all the managed components that support the chosen descriptors, click the **Apply To All Servers** button.

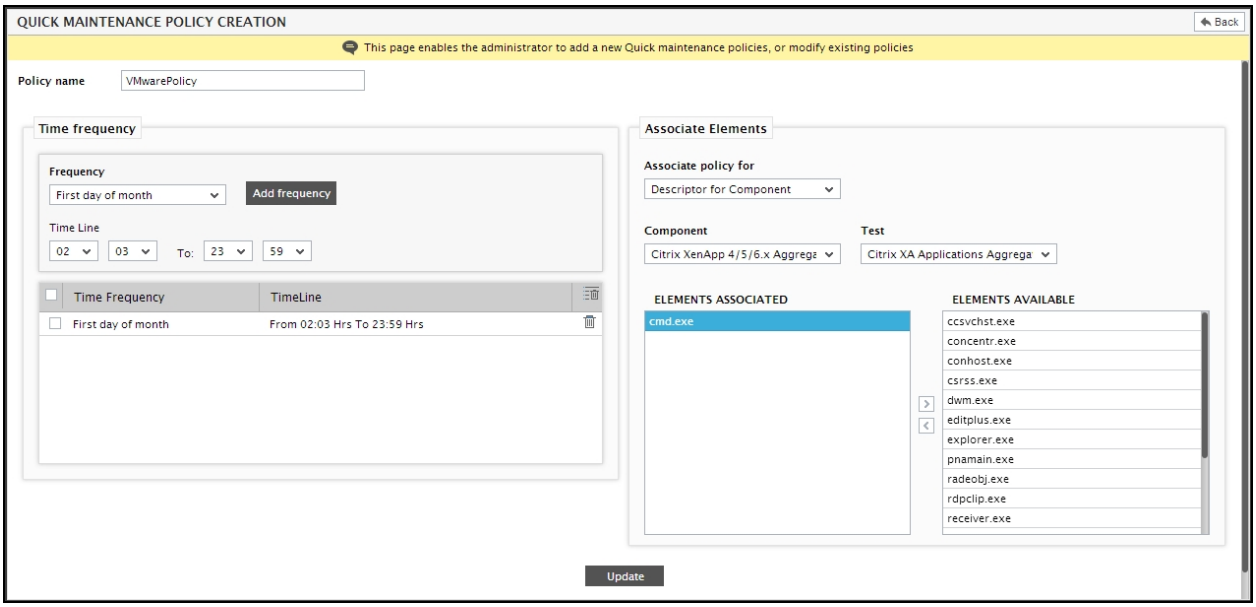

Figure 10.52: Assigning a maintenance policy to the descriptors of a specific component

21. If the **Test for Host** option is chosen from the **Associate policy for** list, then all the host-level tests that are actively reporting measures will be available for selection in the **Test** list that appears (see Figure 10.52). When you select a **Test** from this list, all the managed systems for which the chosen test reports measures will be made available for selection in the **ELEMENTS AVAILABLE** list. In large environments where tens of components are being monitored, you may have a very long list of systems to choose from. In such a case, you can filter your **ELEMENTS AVAILABLE** list further by picking a component type from the **Filter by type** list. By default, the **ALL** option is chosen from the **Filter by type** list. Selecting a particular component type from this drop-down will populate the **ELEMENTS AVAILABLE** list with only those hosts on which the components of the chosen type are executing. Select a host from this list and click the < button or double-click on individual test to associate the selected test to the **ELEMENTS ASSOCIATED** list. Finally, click the **Update** button.

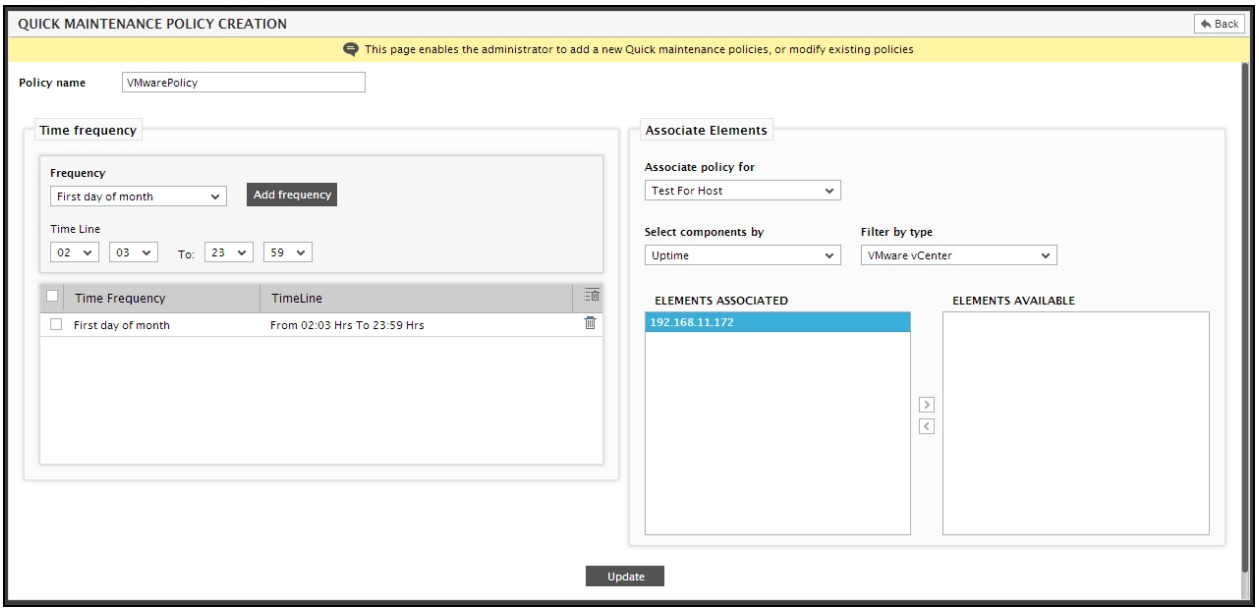

Figure 10.53: Associating a policy with the host for a particular test

22. Likewise, if the **Test for Component** option is chosen from the **Associate policy for** list, then all the application-level tests that are actively reporting measures will be available for selection in the **Test** list that appears (see Figure 10.54). Typically, when you select a test from the **Test** list, all the managed components to which that test is applicable will be displayed in the **ELEMENTS AVAILABLE** list. In environments where a large number of components are monitored, the **ELEMENTS AVAILABLE** list may be too long, making selection difficult. In such cases, you can run a quick search on the **ELEMENTS AVAILABLE** list to easily locate the specific component(s) of interest to you. For this purpose, use the **Search** text box. Specify the whole/part of the component name to search for in the **Search** text box and click the 'magnifying glass' button alongside. Doing so will instantly populate the **ELEMENTS AVAILABLE** list with all those managed components with names that contain the specified **Search** string. With a smaller list of components to select from, you can comfortably pick and choose the components that you want to associate with the maintenance policy from the **ELEMENTS AVAILABLE** list, and click the **<** button (or double-click on individual test) to associate the selection to the **ELEMENTS ASSOCIATED** list. Finally, click the **Update** button. If a maintenance policy is associated with a test that executes on a particular component, then the alerts generated by that test when executing on that component alone will be suppressed.

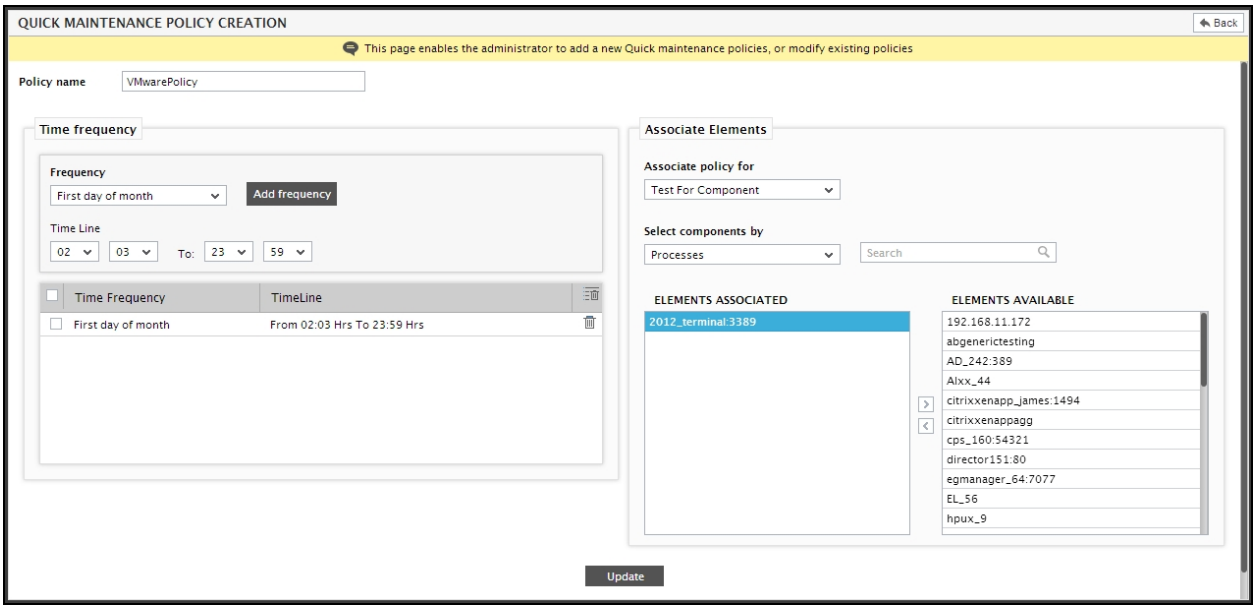

Figure 10.54: Associating a policy with a test that applies to a chosen component

23. Multiple maintenance policies can be created in the same manner discussed above.

Note:

- If a user with **Limited** administrative access - i.e., a user who is only authorized to configure tests, thresholds, and/or maintenance policies for components associated with them - creates a quick maintenance policy, then such a policy can be associated with a **Component**, a **Test for Component**, or a **Descriptor for Component** alone.
- When users with **Complete** component access logs into the eG management console, the quick maintenance policies created by users with **Limited** administrative access will be displayed throughout the user interface in the following format:

  PolicyName (Name_of_user_who_created_the_policy)

- A user with **Limited** administrative access can view in the tree structure (in the **MAINTENANCE POLICY** page) only those quick maintenance policies that he/she has created.

10.6.2 Modifying an Existing Quick Maintenance Policy

To modify an existing policy, do the following:

1. Move your mouse pointer over the policy to be edited in the **MAINTENANCE POLICIES** page. This will reveal a **Modify Policy** icon and a **Delete** icon. To modify that policy, click on the **Modify Policy** icon (see Figure 10.55).

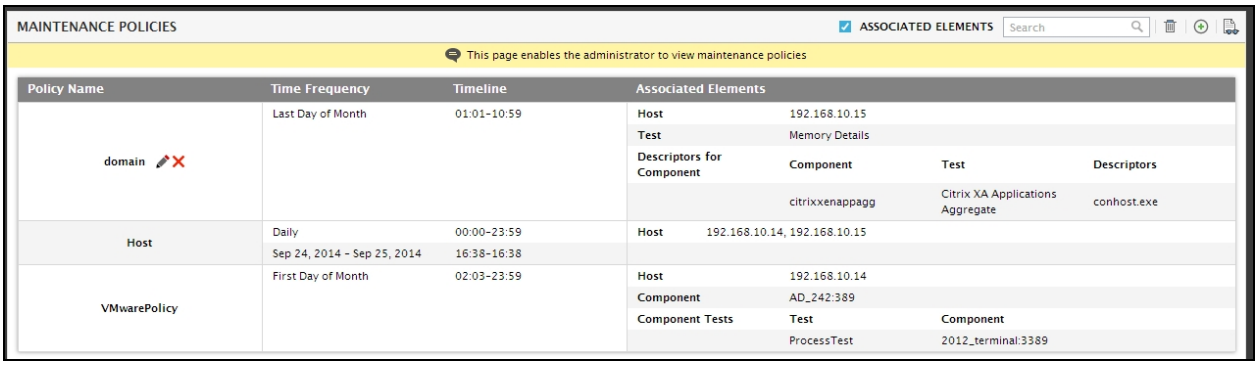

Figure 10.55: Modifying a Quick Maintenance policy

2. A **MODIFY QUICK MAINTENANCE POLICY** page then appears where you can add new time frequencies, remove existing time frequencies, associate new elements, and/or disassociate existing elements from the policy.
3. Finally, register the changes made by clicking on the **Update** button.

10.6.3 Deleting a Quick Maintenance Policy

To delete a particular policy, do the following:

1. Move your mouse pointer over the policy to be removed in the **MAINTENANCE POLICIES** page. This will reveal a **Modify Policy** icon and a **Delete** icon. To delete a policy, click on the **Delete** icon (see Figure 10.55).
2. eG Enterprise will then request your confirmation to proceed with the deletion of the chosen policy (see Figure 10.56). Click the **OK** button in the message box to confirm deletion. Otherwise, click on the **Cancel** button.

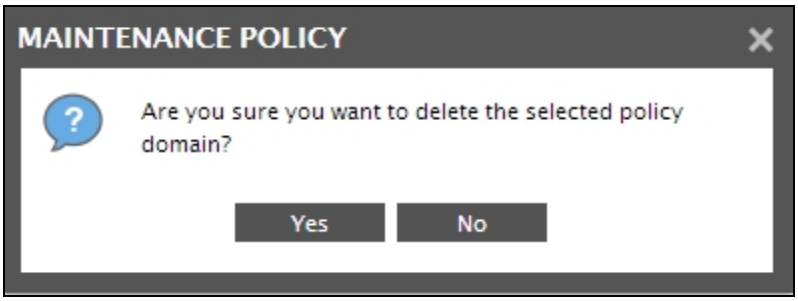

Figure 10.56: Confirmation to delete an Existing Policy

3. To delete multiple policies at one shot, click the **Delete Policies** icon (that resembles a trash can) in the maintenance tool bar at the right top corner **MAINTENANCE POLICIES** page. Figure 10.57 will then appear. Select the policies you want to delete and click the **Delete** button therein.

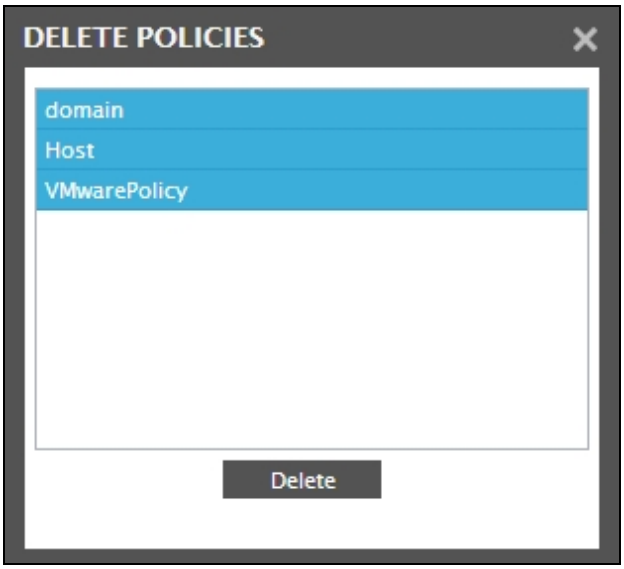

Figure 10.57: Deleting multiple Quick Maintenance Policies

4. eG Enterprise will then request your confirmation to proceed with the deletion of the chosen policies (see Figure 10.58). Click the **OK** button in the message box to confirm deletion. Otherwise, click on the **Cancel** button.

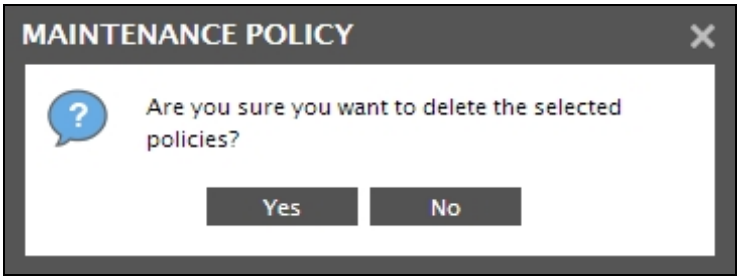

Figure 10.58: Confirmation to delete the policy

10.6.4 Maintenance Analysis

Maintenance is a mandatory task in many large IT infrastructures. Maintenance procedures would be regularly performed on the components in such infrastructures to ensure peak performance. Also, to minimize false alarms that might be raised during maintenance, maintenance policies might have been configured for a wide majority of these components using the eG administrative interface. The key need of the administrators of such large environments therefore is the knowledge of which components are under maintenance, at what level (i.e., host/component/test-level), and when the next maintenance cycle is scheduled for these components, so that they can promptly alert service managers and domain experts about the non-availability of their critical applications/devices during the maintenance periods. Moreover, since new components will continue to be added to such infrastructures, it is also the responsibility of the administrator to analyze and identify those components which need to brought under the purview of maintenance.

To provide administrators with quick and easy access to maintenance-related information and to enable them to efficiently review and analyze maintenance schedules, eG Enterprise provides a Maintenance policy

analysis page. This page serves as a single, central interface using which administrators can view and analyze the maintenance policies configured and applied across the environment.

To analyze maintenance details using this page, do the following.

1. First, pick a criterion for analysis from the **Analysis By** list. The options are: **All**, **Host**, **Component**, **Test**, **Test for Host**, **Test for Component**, **Descriptor**, **Descriptor for Component**, **Policy**, and **Time Constraints**.
2. For instance, if the **All** option is chosen, then the complete details of all policies configured – i.e., the policy names, the policy specification, and the elements they are applied to - will be displayed (see Figure 10.59).

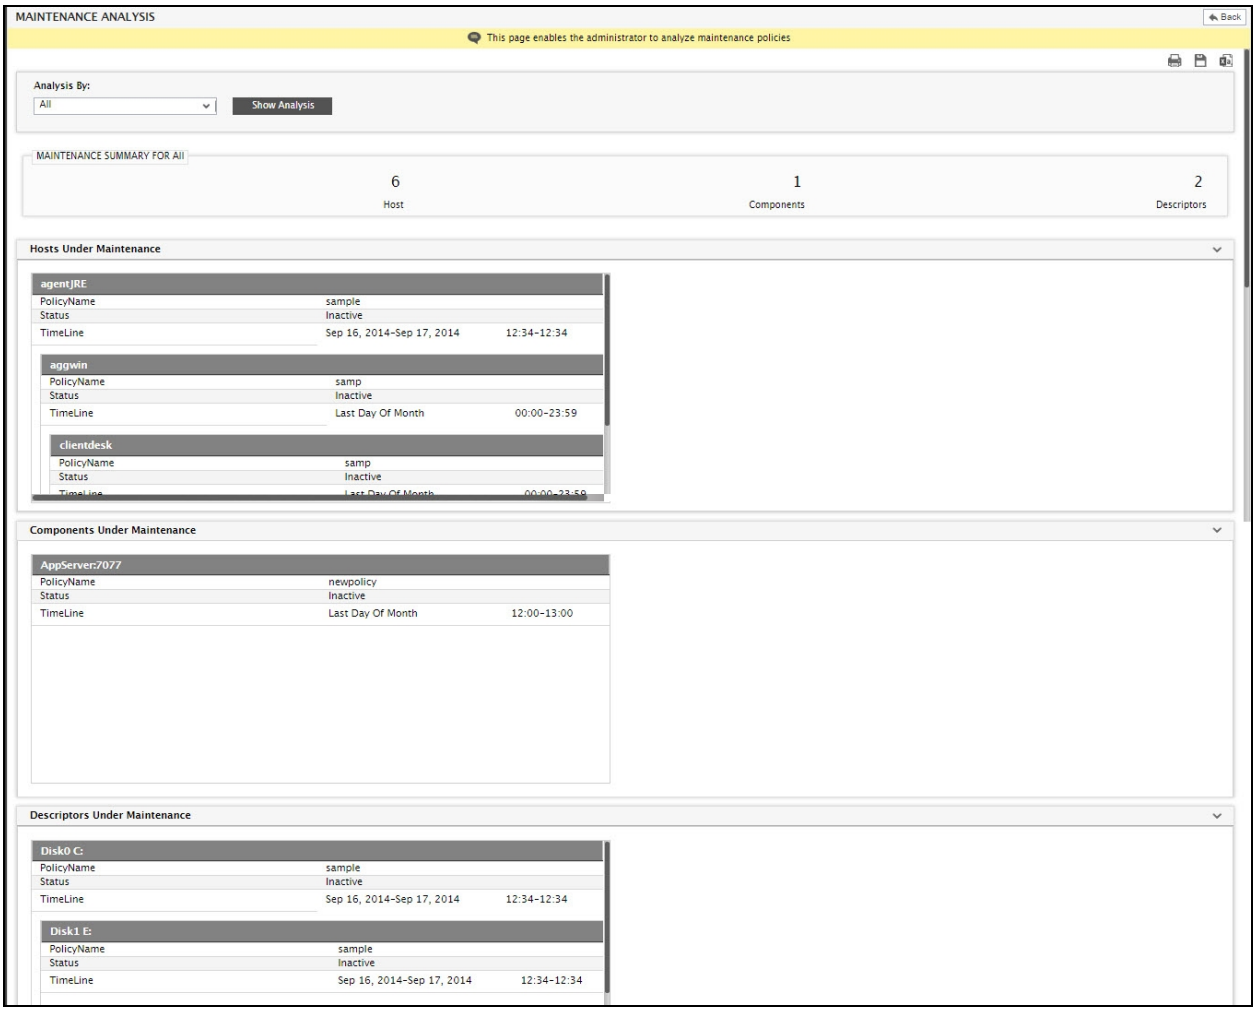

Figure 10.59: Maintenance analysis by All elements

3. Likewise, if the **Hosts** option is chosen from the **Analysis By** list, then you can either pick the **Host** for which you want to perform maintenance analysis or perform the analysis across **All** hosts. To pick a particular host, you can use the **Search** text box. Here, you can specify the whole/part of the host name to search for and click the 'magnifying glass' icon to search for hosts that match the specified search string. If such hosts are found, they will be displayed in the **Host** list, from which you can quickly select the host of interest to you. If you then click the **Show Analysis** button, the complete details of all policies that apply to the chosen host will appear (see Figure 10.60).

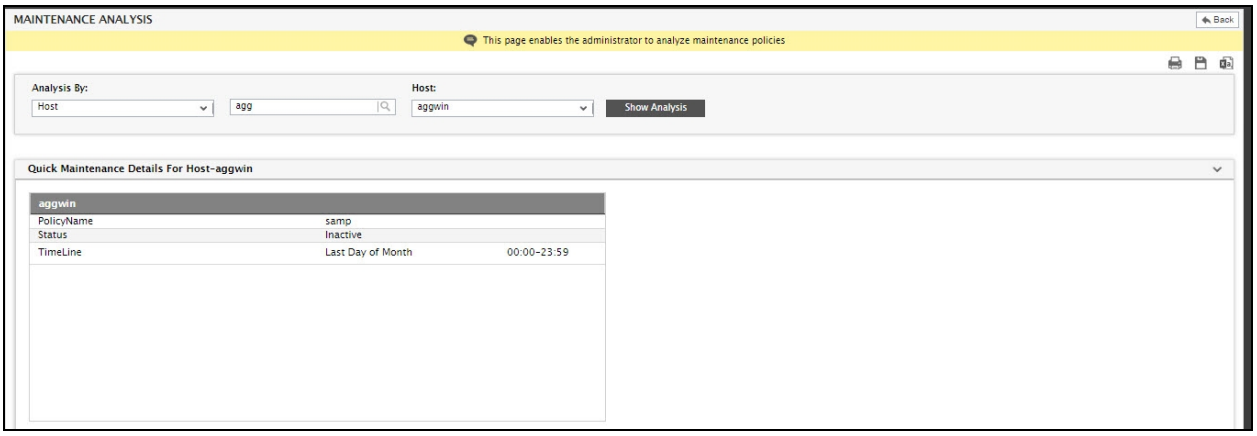

Figure 10.60: Maintenance analysis of a chosen Host

4. To analyze across hosts, pick the **All** option from the **Host** list. In this case, all hosts under maintenance will be displayed and the policies that are associated with each host will be revealed (see Figure 10.61).

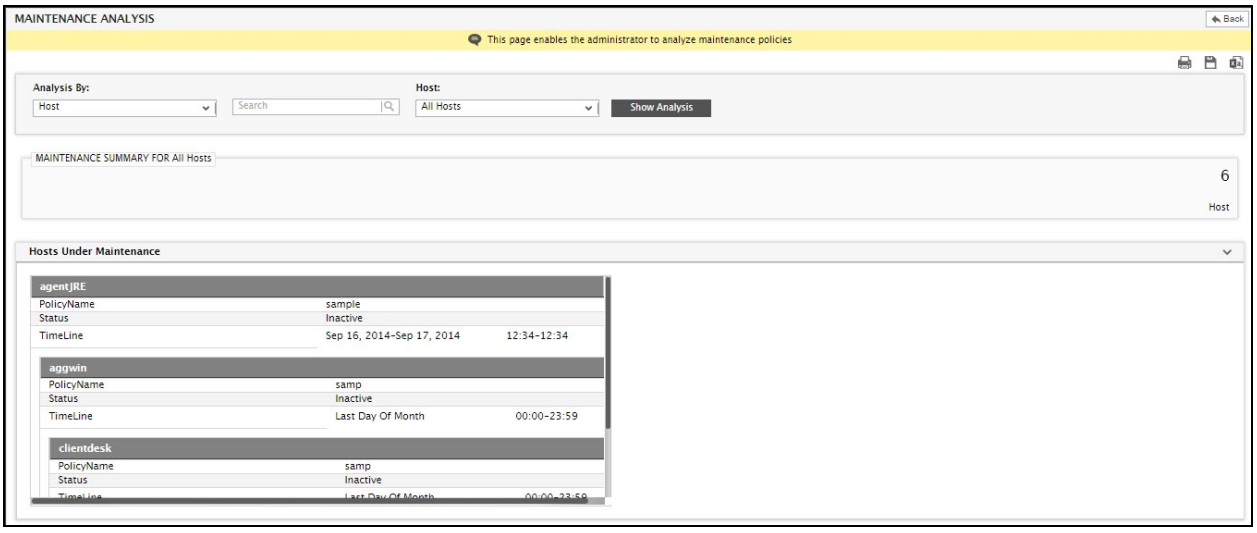

Figure 10.61: Maintenance analysis across Hosts

5. This way, you can view the policy details for one/all **Components**, **Test**, **Tests for Host**, **Tests for Component**, **Descriptors**, **Descriptors for Component**.

6. In addition, you can also view the details of a particular policy. For this, first pick **Policy** from the **Analysis By** list. Then, select a particular policy from the **Policy** list. Optionally, you can also use the **Search** text box to quickly locate the policy of interest to you. In this case, specify the whole/part of the policy name to search for and click the 'magnifying glass' icon to search for policies that match the specified search string. If such policies are found, they will be displayed in the **Policy** list, from which you can quickly select a policy. Then, click the **Show Analysis** button. Figure will then appear reviewing the number and names of elements to which the policy applies (see Figure 10.62).

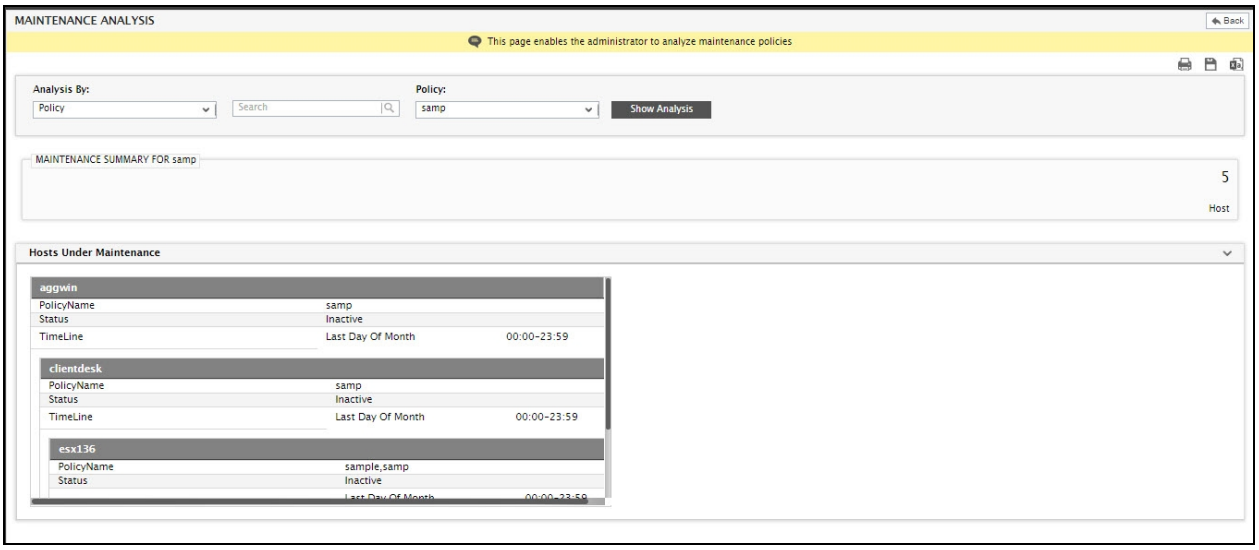

Figure 10.62: Performing maintenance analysis of a particular policy

7. You can also select the **Time Constraints** option from the **Analysis By** list to view the details of those policies that are applicable during a timeline selected from the **Time Constraints** list (see Figure 10.63).

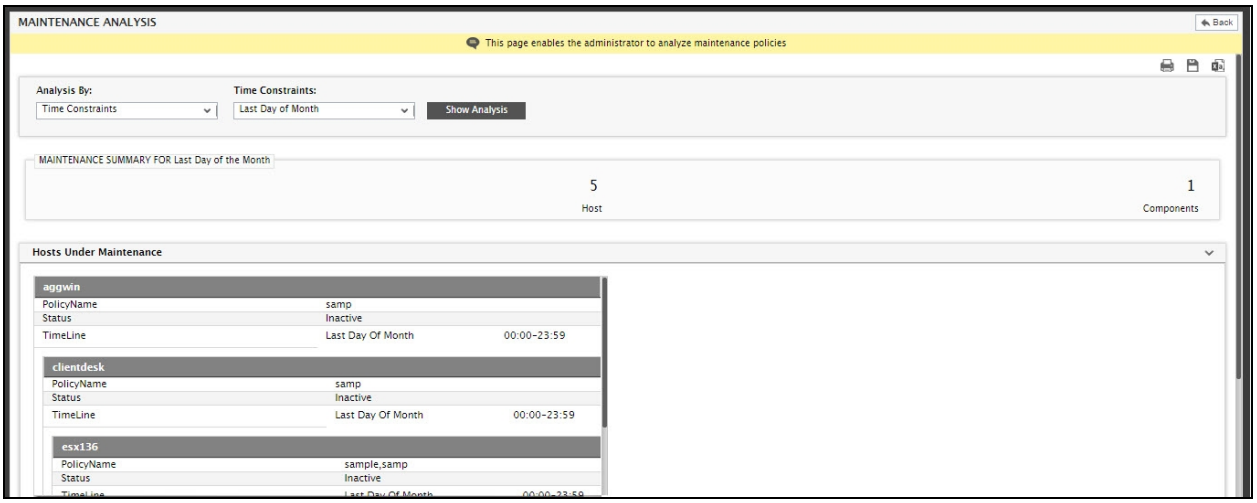

Figure 10.63: Maintenance analysis of policies to which a chosen Time Constraint applies

8. At any given point in time, you can **Save** the results of the analysis as a PDF file or a CSV file or print the same using the corresponding icons at the right, top corner of the **MAINTENANCE ANALYSIS** page.

11

Segment Topology

eG Enterprise's patented correlation technology is dependent on the specification of topology information that indicates how components are interconnected and which components rely on others for their functioning. The interconnections can represent either physical connections (e.g., a web server connected to a network router) or logical dependencies (e.g., a web server using a web application server). Each interconnection is associated with a direction. The direction signifies cause-effect relationships (if any) between the components being connected together.

Having configured the components that eG Enterprise should monitor in the target environment, the administrator has to next configure the component topology of the target environment. The component topology is comprised of one or more segments. A segment is a logical entity that represents a collection of components and the interdependencies between them. Segments can be chosen to represent user groups, organizational domains, or physical locations in the target environment. For example, the topology of a banking environment may include two segments - one representing all the components in the bank's NY branch and another representing the bank's CA branch.

A segment topology can either be manually configured or auto-discovered by the eG Enterprise system.

The sections that follow will discuss both these options in detail.

11.1 Manually Configuring a Segment Topology

To view the segments that are manually configured, follow the Infrastructure -> Segments -> Configured option. This results in Figure 11.1 where the segment list that pre-exists can be viewed. To modify a segment, click on the  icon corresponding to that segment in Figure 11.1. To delete a segment, click the  icon corresponding to it. To delete multiple segments simultaneously, select the check box against each segment and click the  icon. To delete all the segments configured in the infrastructure, select the check box against the **Segment Name** label and click the  icon.

If a large number of segments have been configured, locating the segment to be modified/deleted would become quiet a challenge. This page therefore allows you to quickly search for a segment, by first specifying the whole/part of the segment name in the **Search** text box, and then clicking the  icon next to the text box. All segments with names that embed the specified string will then appear.

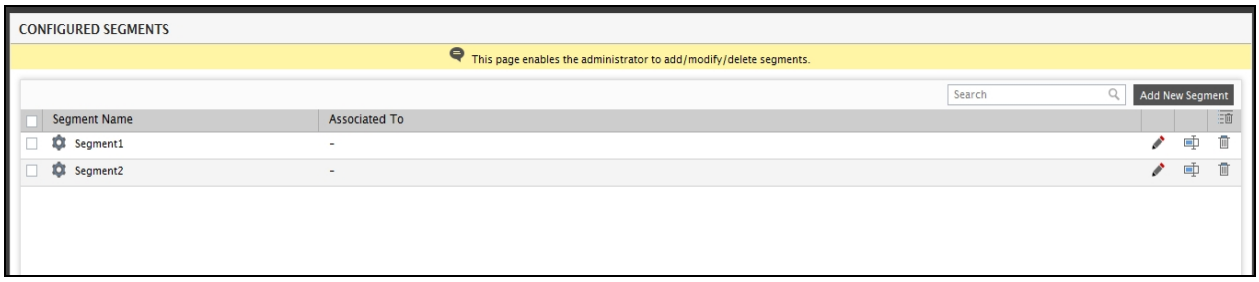

Figure 11.1: Currently configured segments in the environment

You can even change the name of a segment by clicking the icon corresponding to it. Provide the **New name for the selected segment** in Figure 11.2 that appears and click the **Rename** button therein to register the new name. This will automatically change the segment name across the eG user interface.

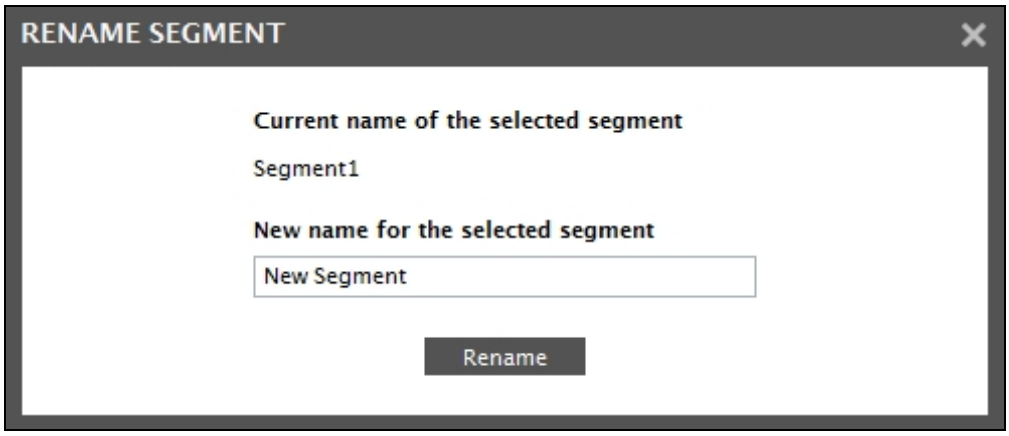

Figure 11.2: Renaming a segment

Note:

Since renaming changes the name of a segment across the user interface, you will no longer be able to access performance data of the segment for the period prior to the name- change using the monitoring/reporting consoles. Therefore, exercise caution when renaming a segment.

A new segment can be added by using the **Add New Segment** button as shown in Figure 1. The screen depicted by Figure 11.3 appears wherein the name of the segment has to be specified.

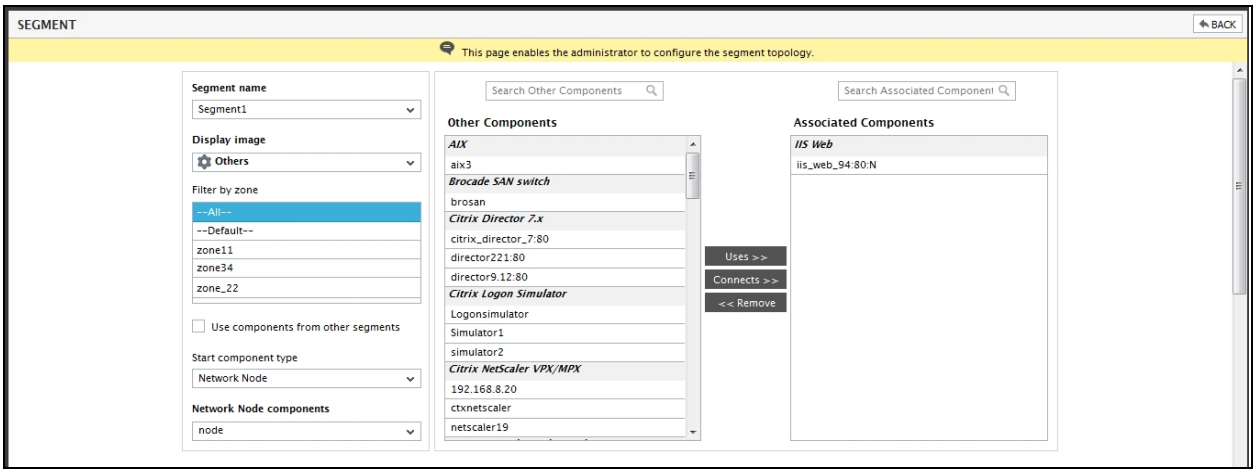

Figure 11.3: Adding a new segment

Next, select the **Display image** for the segment.

If you want your segment to contain a component that is already within a zone, then select that zone from the **Filter by zone** list. By default, the **--All--** option is displayed indicating that all components in the monitored infrastructure are available for creating a new segment.

By default, components that are already a part of other segments will not be available for inclusion in a segment. However, if a single component features in multiple segments, this default setting can be overridden - i.e., while configuring such segments, you can switch on the **Use components from other segments** flag in Figure 3. By default, this check box is deselected. By selecting it, you can ensure that components that are already a part of other segments are also available for inclusion in the new segment.

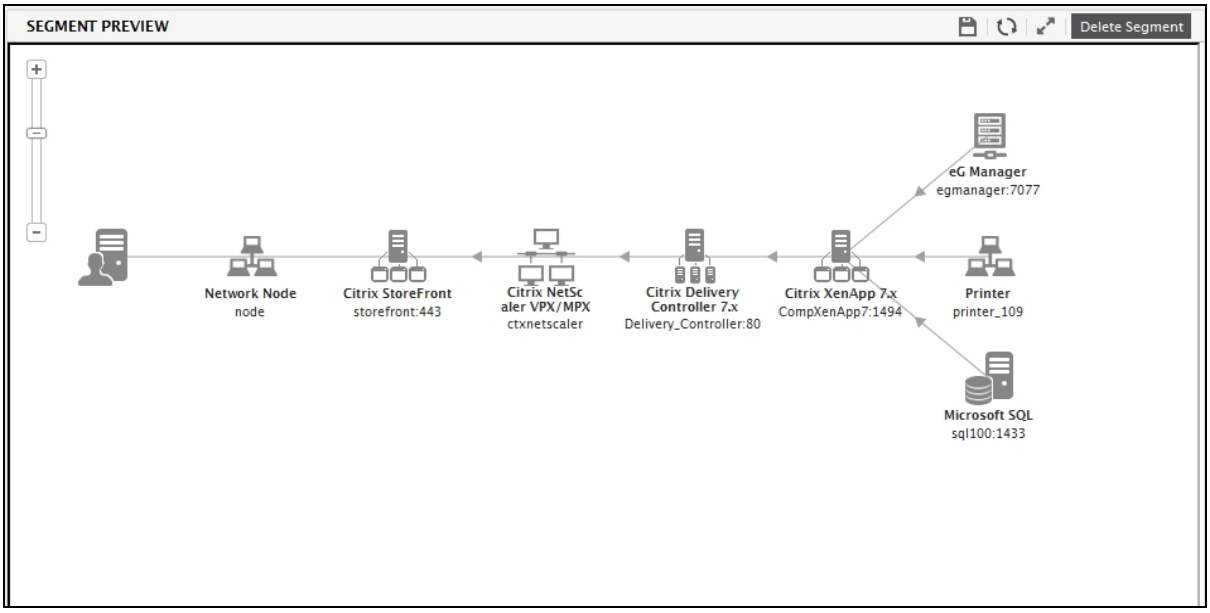

Figure 11.4: Preview of the configured segment topology

The interconnections between components of a segment must be configured from left-to-right, with the left most node(s) representing the entry points through which users access the web services offered in the target environment. As is evident from Figure 4, the segment is configured step-by-step, with all the dependencies of one of the components in the environment/zone being specified at each step. Using the **Start component type** and the **components** list that follows it, select the component type and the component (of that type) with which the dependency begins. Then, from the **Other Components** list, select the component with which the dependency ends. In other words, if a web server is dependent on an Oracle database server for its functioning, then select the web server component using the **Start component type** and the **components** list, and pick the Oracle database server from the **Other Components** list. Then, pick the type of dependency that binds the chosen components by clicking one of the two buttons – **Uses** or **Connects**. Figure 4 indicates that eG Enterprise supports two types of dependencies: the **Connects** dependency, which typically indicates flow of data (e.g., a physical connection between a web server and a network router), and the **Uses** dependency, which refers to a logical dependency (e.g., a web server and a web application server, a web application server and a database server, etc.). The key difference between the two forms of dependencies is that when one component **Uses** another, problems with the latter component can actually reflect in problems with the former component. With the **Connects** dependency, there are no such cause-effect relationships between the two components. This way, build the entire segment topology.

Figure 4 depicts the topology graph of a segment configured via the user interface. In this example, users access the target environment via a network node. The network node is connected to a Citrix Storefront. The storefront server forwards incoming client requests to a Citrix NetScaler, which in turn routes the requests to a Delivery Controller 7.x. The Delivery controller then routes the requests to XenApp 7.x server in the farm. A Microsoft SQL server, Printer and an eG Manager are used in the backend to provide support services.

Note:

- All the components that are not a part of any segment are considered as independent components.
- An independent component can be added to more than one segment.
- Only managed component-types are available in the Type of components list box.

The **Delete Segment** button in Figure 5 enables the administrators to delete any configured segment. At each step, as the dependencies are specified by the administrator, the segment preview as in Figure 5 displays the current segment configuration. The feedback provided by the segment preview can be used by the administrator to refine the segment specification.

The **Remove** option in Figure 4 can be used to delete any of the configured dependencies from the segment.

While the topology depicted by Figure 5 is fairly small and simple, large, mission-critical environments, on the other hand, are characterized by several load balanced groups of servers providing web, middleware, and database functionality. The servers in the group often perform the same functions and have the same set of dependencies on other infrastructure components. For a segment with tens of servers in a group, the segment topology representation can quickly get very cumbersome. To handle such environments, eG Enterprise allows the configuration of Groups. By grouping all components that share the same inter-relationships in a topology as one, administrators can ensure that large topology representations do not appear cluttered! To know how to configure Groups, refer to Section . To know how to use a Group in a topology configuration, let us once again consider the *Segment1* topology depicted by Figure 5 above. As you can see the Citrix XenApp server in the topology supports a **Uses** dependency with the following components:

- A Microsoft SQL server
- A Printer
- An eG Manager

Let's say that these 4 components have been grouped under a group named *new_group*. Now, to build this group into the topology, do as follows:

1. Since it is the Citrix XenApp server that is dependent on the group, first select *Citrix XenApp* from the **Start component type** list, as depicted by 11.1. Then, select the XenApp server from the **Citrix XenApp components** list.

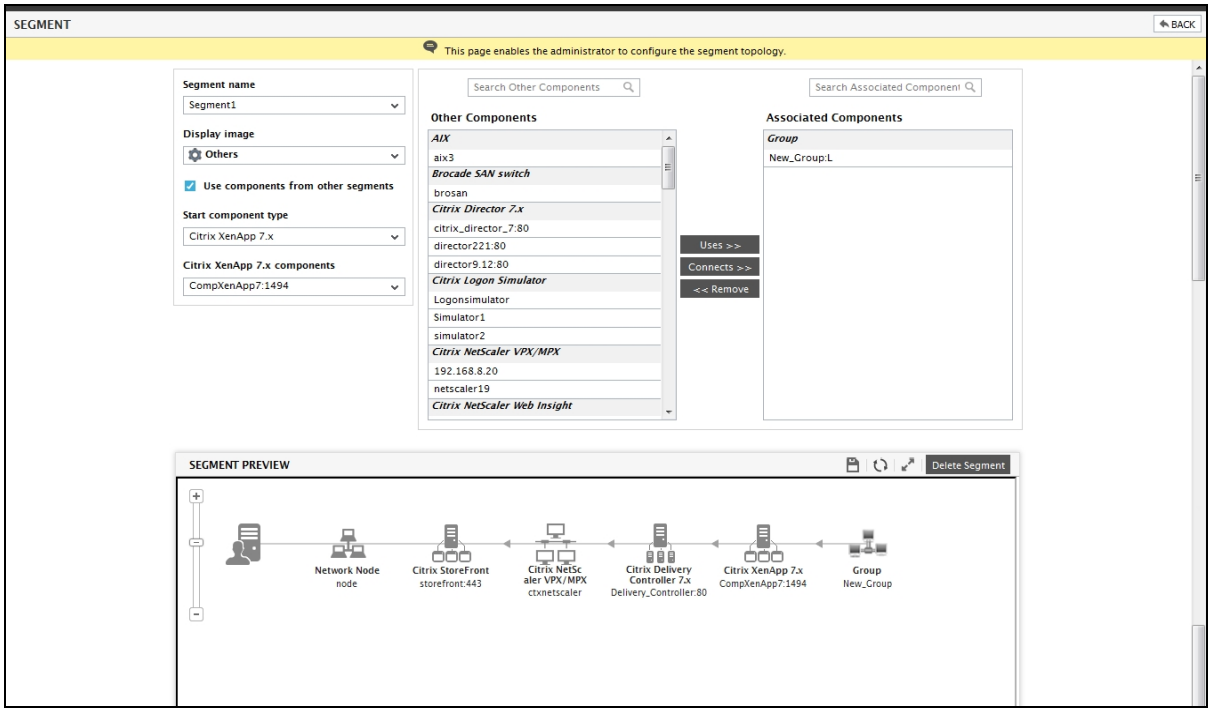

Figure 11.5: Selecting the group to be included in a segment topology

2. The group, *New_Group*, will now appear in the **OTHER COMPONENTS** list. To associate the group with the Citrix XenApp 7.x server, select it from the **OTHER COMPONENTS** list, and click the **Uses** button.
3. The **SEGMENT PREVIEW** depicted by 11.1 clearly indicates that the group, *New_Group*, has been associated with the Citrix XenApp 7.x server.

   Moreover, in order to help you view and analyze the configured topology better, the **SEGMENT PREVIEW** section (see 11.1) provides a special toolbar comprising of easy-to-use viewing options. These options have been discussed in the table below:

Tool	Purpose
	Helps to zoom in and zoom out the topology representation
	You can change the position of any server representation in your topology by cliking, dragging, and dropping that representation. Once done, you can click on this icon to save the change in position.
	As stated earlier, you can reposition any component in the **SEGMENT PREVIEW** by clicking on it and dragging it to the desired position. To restore the component(s) so shifted to their original positions, click on this icon.
	Click here to view the topology in a new window.

11.2 Auto- Discovering a Segment Topology that is Not Associated with a Zone

The eG agent has the ability to **auto-discover individual components and the component topology**. The ability to auto-discover the relationship between components, when enabled, helps administrators draw an almost accurate segment topology with minimal user intervention and time! By default, the ability of the eG agent to **automatically discover topology** is enabled. However, this default setting will take effect only if the **eG agent has the ablity to automatically discover components**. This is because, topology discovery cannot be performed without component discovery. To enable the eG agent to perform component discovery, follow the steps detailed in Section **7.1.2.1** of this document. If both component discovery and topology discovery capabilities of the eG agent are enabled, then, follow the Infrastructure -> Segments -> Discovered menu sequence to auto-discover a segment. If no auto-discovered segments pre-exist, Figure 11.6 will appear. In such a case, check the 'discovery status' displayed at the right, top corner of Figure 11.6 to determine whether/not topology discovery has been turned on. If not, then click the **Click here** link alongside to instantly switch to the **DISCOVERY** page to enable topology auto-discovery. If you then come back to Figure 11.6, auto-discovered topologies will be listed.

You can save any of the listed topologies as a **segment** using this page, or can modify/delete the discovered topologies.

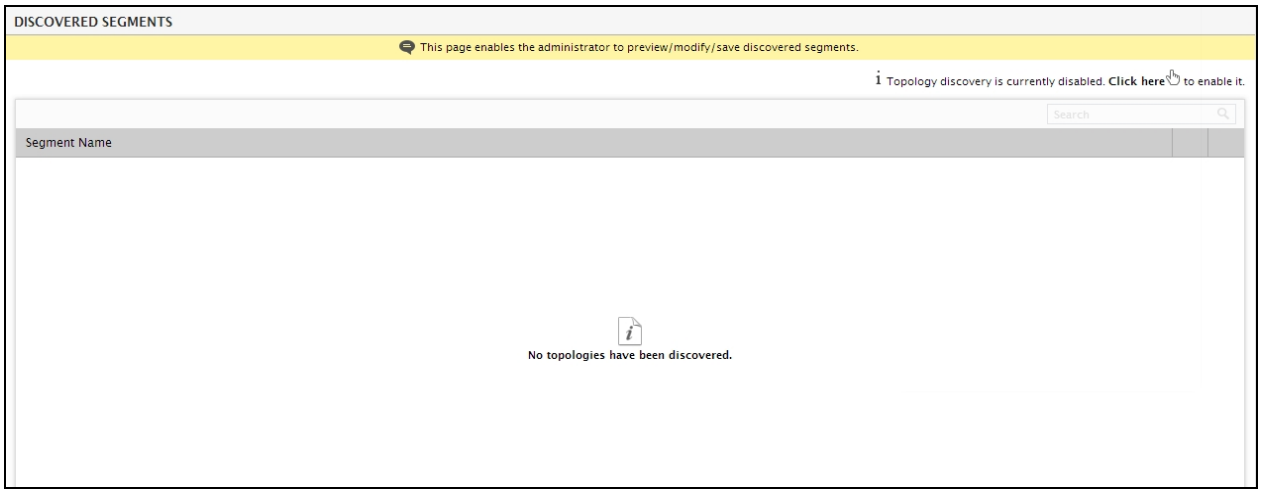

Figure 11.6: The AUTO TOPOLOGY page

To modify the topology of a segment, do the following:

1. Click the **Edit** icon corresponding to the discovered topology.
2. Figure 11.7 will then appear providing a preview of the discovered topology.

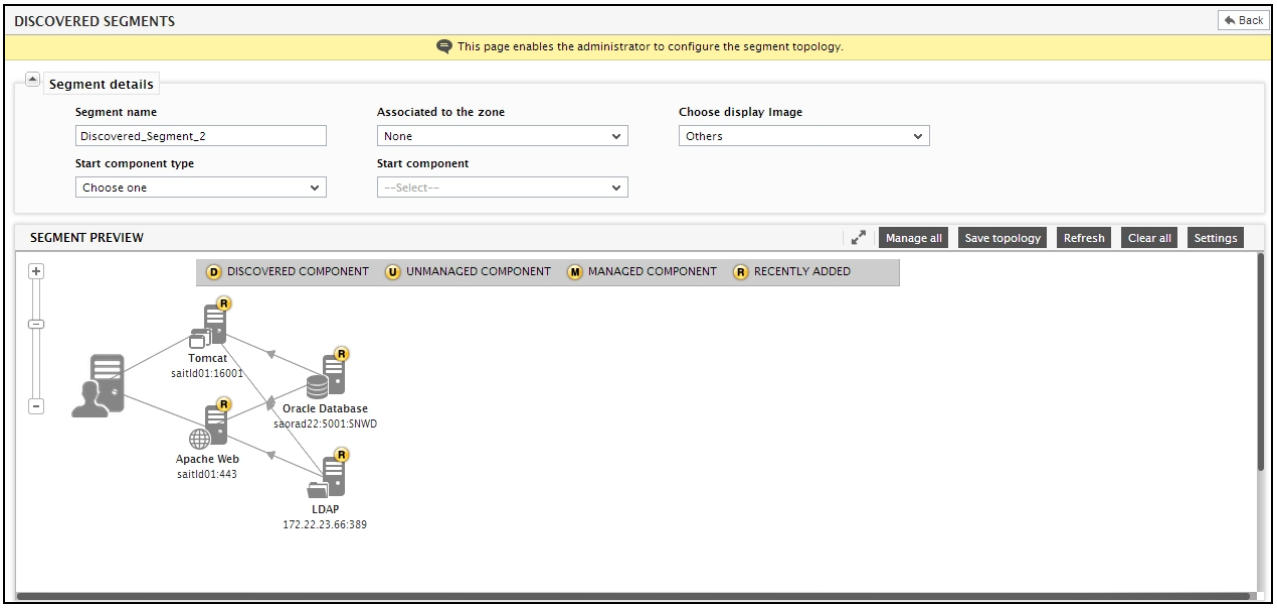

Figure 11.7: The auto-discovered segment topology

3. If for some reason, any segment component was not automatically discovered by the eG agent, you can alter the segment topology 'on-the-fly' to include that component(s). For instance, say, the LDAP server in Figure 11.7 talks to 2X Terminal server. Since the eG agent has not auto-discovered the inter-connection between the LDAP server and the 2X Terminal server, let us manually weave this relationship into the segment topology. To do this, first, right-click on the segment component with which the 2X Terminal server interacts - in the case of our example, it is the LDAP server. Doing so will bring up a shortcut menu as depicted by Figure 11.8.

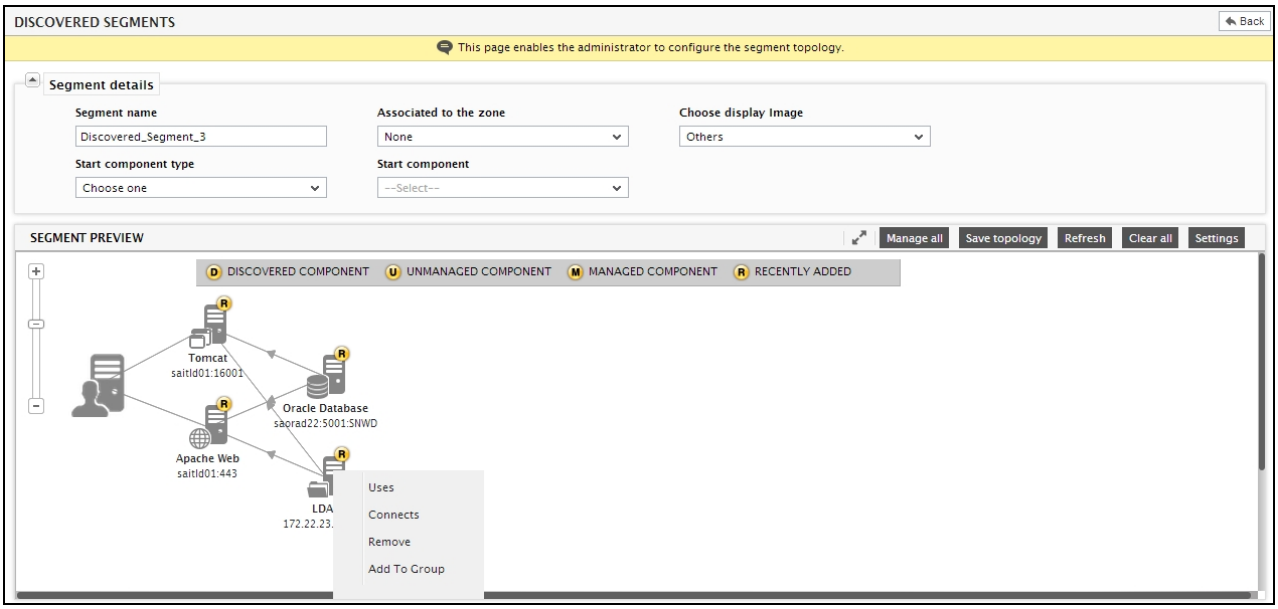

Figure 11.8: Modifying the auto-discovered topology to include a new component

4. To establish a **Uses** connection between the LDAP and 2X, pick the **Uses** option from the shortcut menu of Figure 3. This will invoke Figure 11.9.

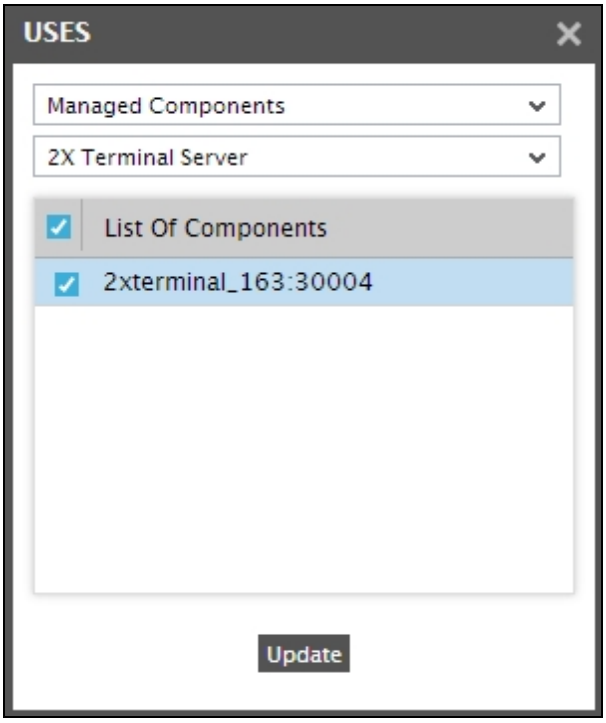

Figure 11.9: Selecting the component to be added to the auto-discovered topology

5. To an auto-discovered topology, you can add any managed component or a component that has just been discovered and is yet to be managed. To add a managed component to the topology, select the **Managed Components** option from Figure 11.9. For adding a component that has only been discovered and not

managed, select the **New** option from Figure 11.9. The 2X Terminal server in our example has already been managed. Therefore, pick the **Managed Components** option from Figure 11.9, and select *2X Terminal server* from the list of component types. All managed 2X Terminal servers will then appear in the **LIST OF COMPONENTS** section. Click on the check box alongside the LDAP server that you want to add to the topology. The topology will instantly change to reflect the addition (see Figure 11.10).

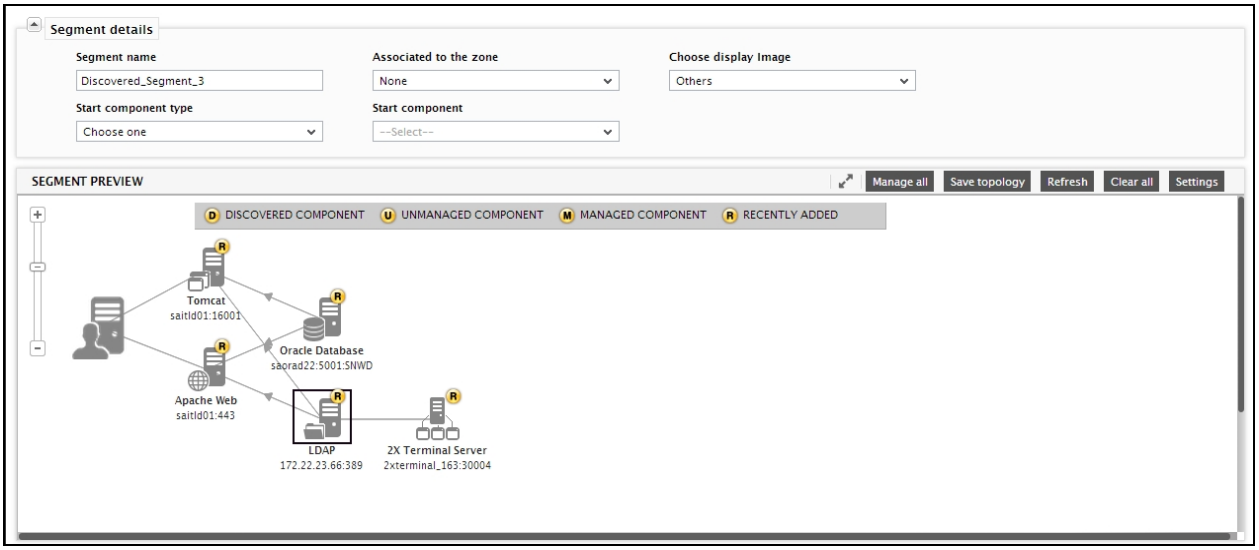

Figure 11.10: A new component added to an auto-discovered topology

6. Likewise, you can add multiple components to an auto-discovered topology.
7. In addition to components, component groups can also be added to the topology.
8. Besides adding components/component groups, you can also remove one/more components/groups from an auto-discovered topology. For instance, consider the auto-discovered topology of Figure 11.11.

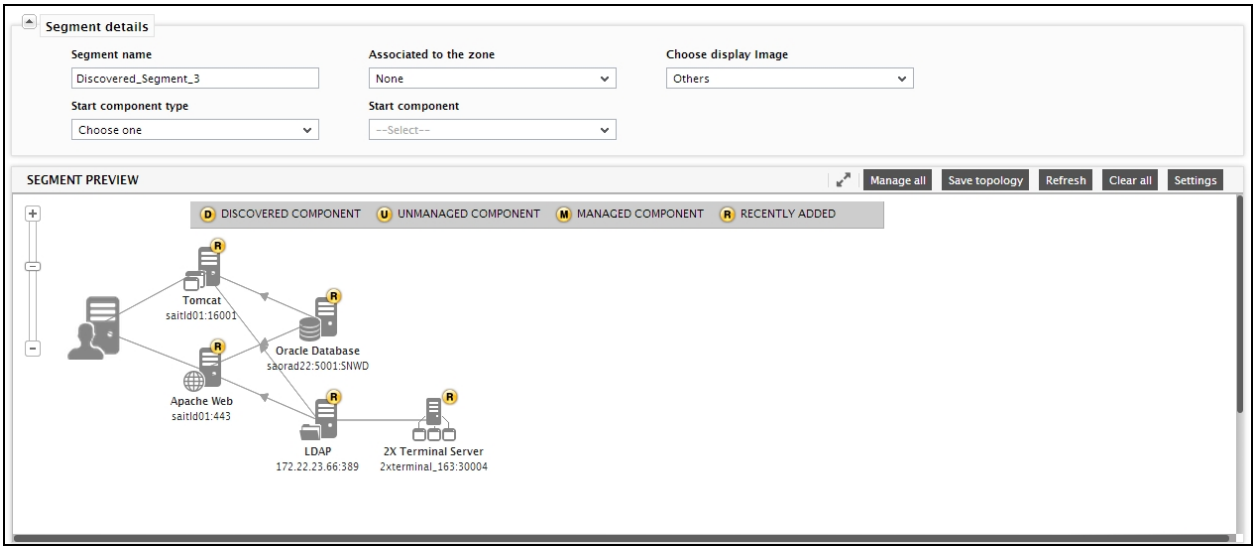

Figure 11.11: Modifying an auto-discovered topology by removing a component from it

9. Assume that you want to remove the 2X Terminal server component from the topology. For this, first, right-

click on the *2X Terminal server* server to be removed, and pick the **Remove** option from the shortcut menu (see Figure 11.12).

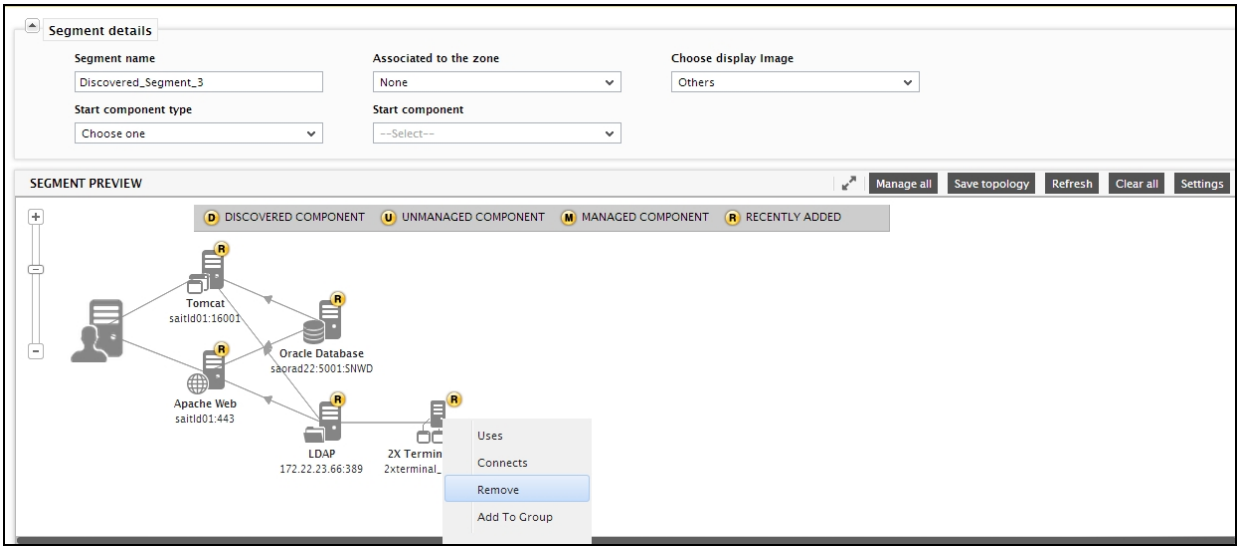

Figure 11.12: Removing a component from the auto-discovered topology

10. The change will be immediately reflected in the topology, as depicted by Figure 11.13 below.

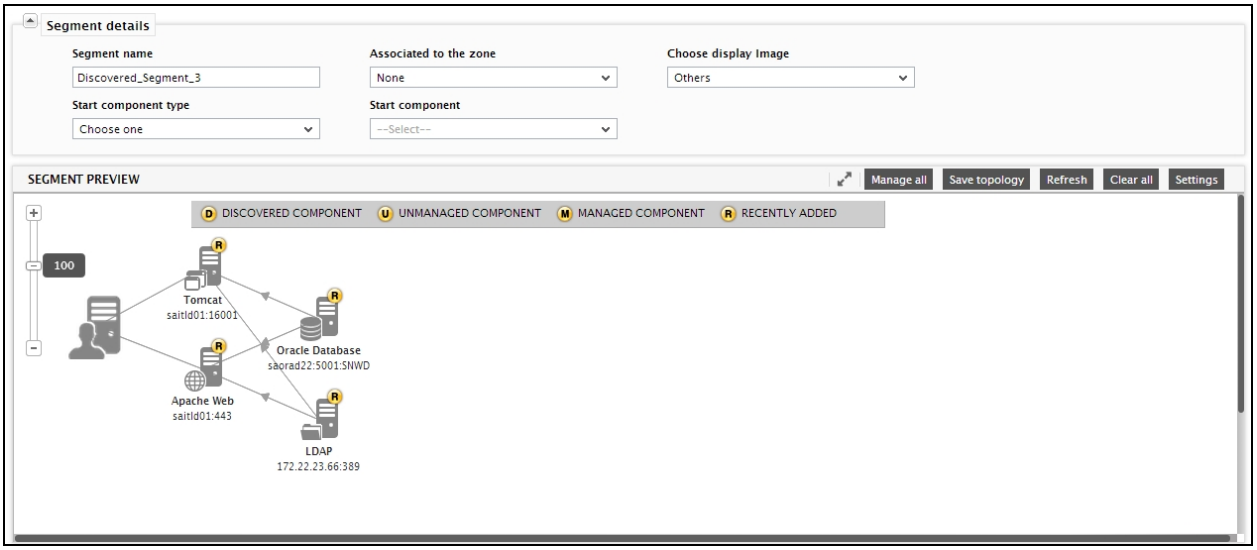

Figure 11.13: A changed auto-discovered topology

11. You can then proceed to save the topology. However, **note that a segment topology can contain only 'managed components'.** In the case of the example illustrated by Figure 11.14, the Tomcat server is yet to be managed. To be able to save the topology, you first need to manage this server. For this, right-click on the Tomcat server in the segment topology of Figure 11.14, and pick the **Manage** option (see Figure 11.14)

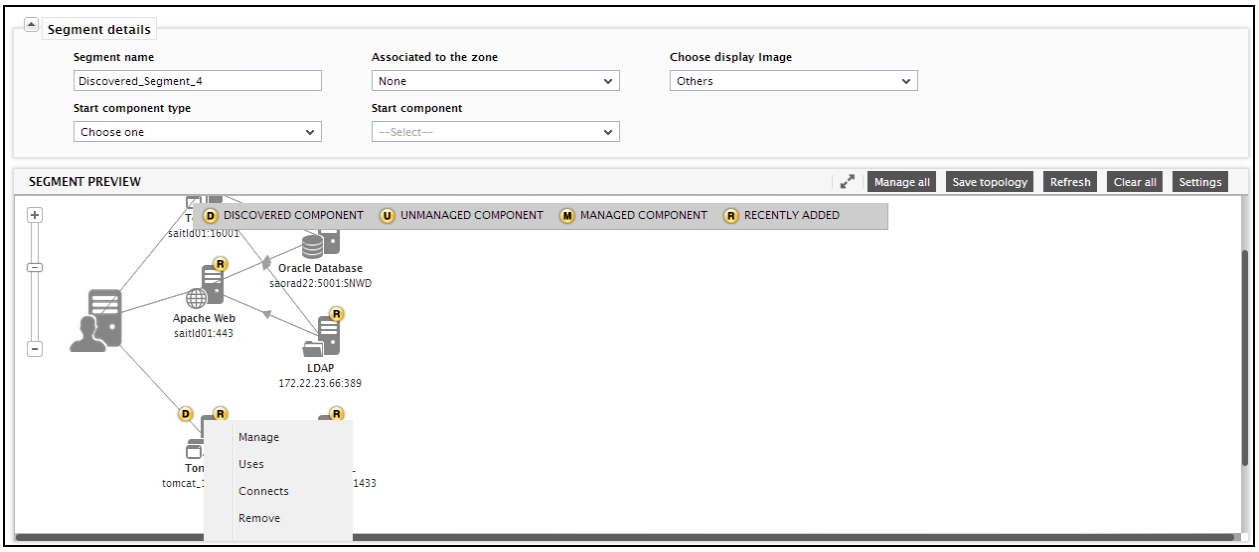

Figure 11.14: Managing a newly discovered server in an auto-discovered topology

12. Figure 11.15 will then appear, using which you can change the configuration of the server being managed. The **Nick name** of the server, the status of the **Agentless** flag and the **Internal agent assignment** flag (if available), and the **External agent** assigned to the server can be changed. Click the **Add** button in Figure 11.15 to add the server.

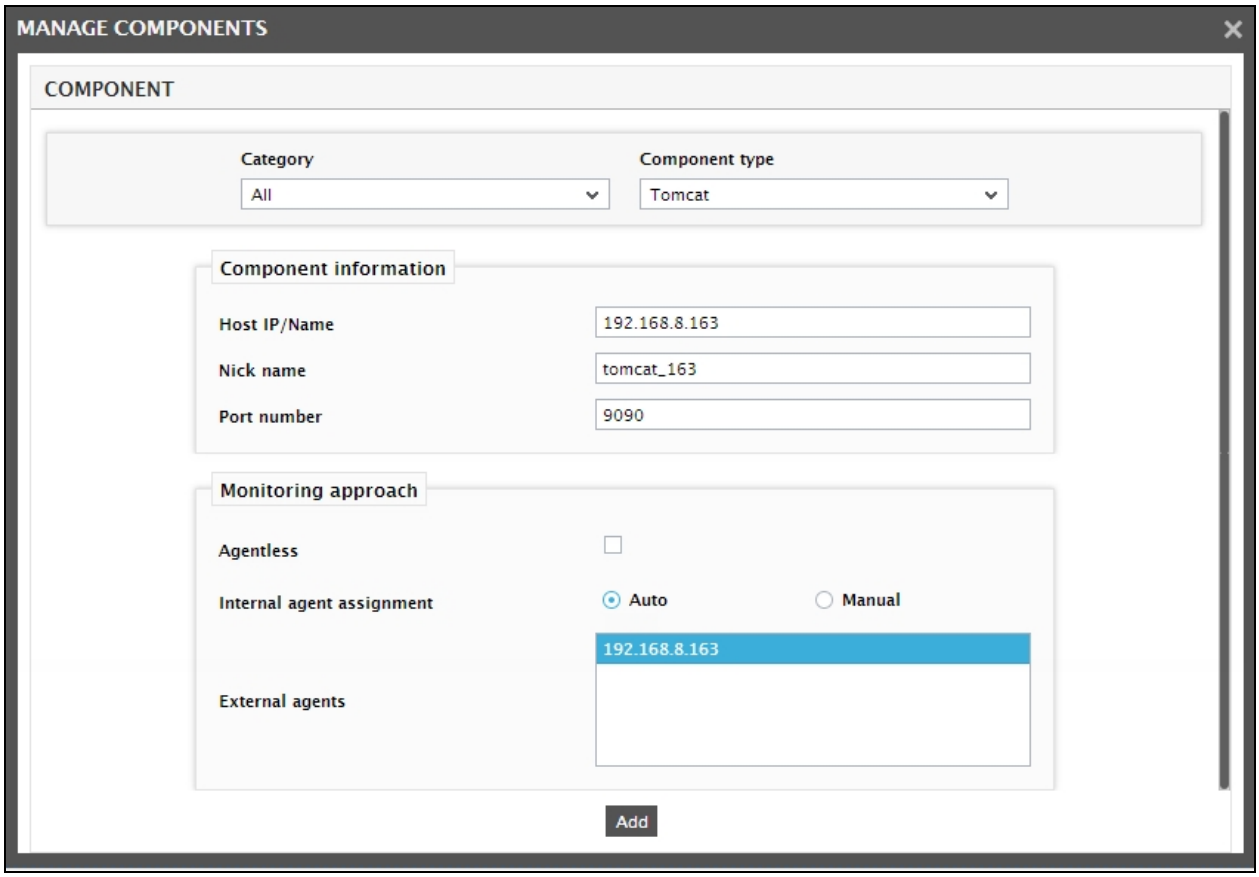

Figure 11.15: Managing the new server

13. With that, the unmanaged server has been managed. The topology display too will change to this effect (see Figure 11.16).

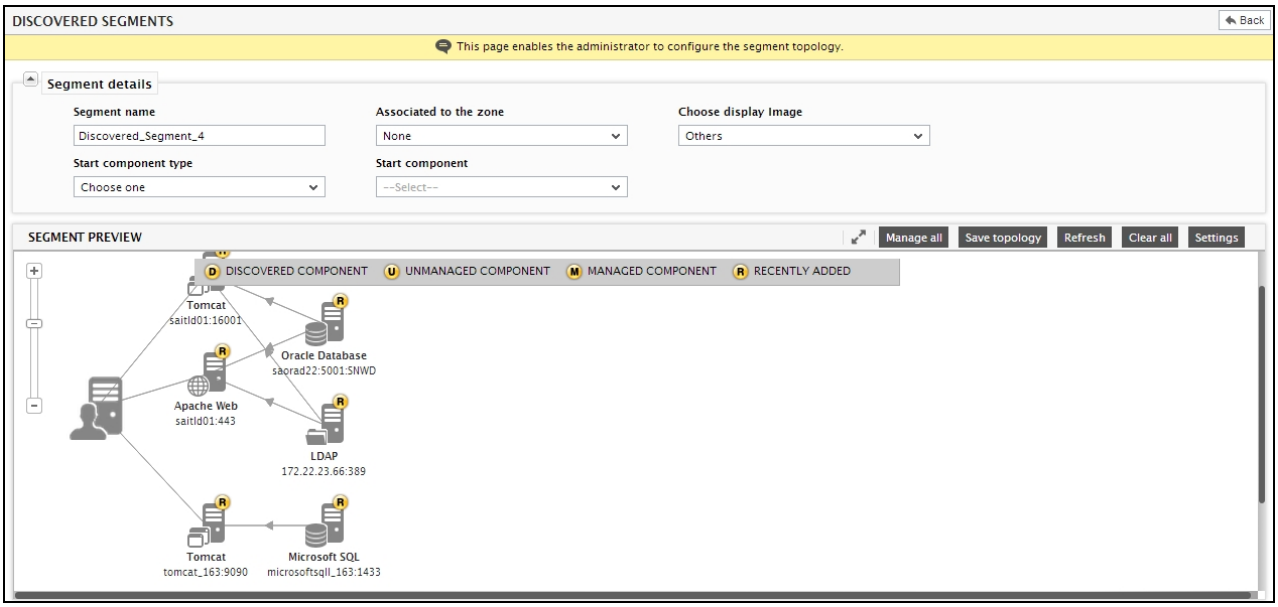

Figure 11.16: The auto-discovered topology after an unmanaged server is managed

14. You can now save the topology by clicking the **Save topology** button above the segment preview in Figure 11.17.

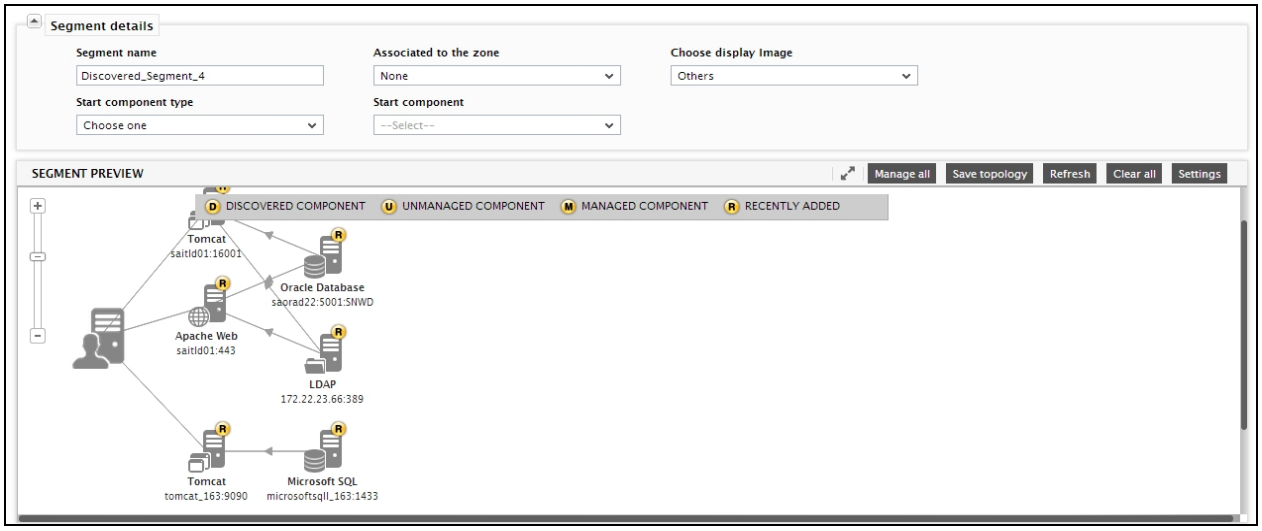

Figure 11.17: Saving the topology

15. You can click on the **Refresh** button in Figure 11.17 to refresh a topology, so that it displays recent additions/removals. To clear a topology, select the **Clear all** button in Figure 11.17.

16. In a topology with a limited number of newly discovered components, you can right-click on each unmanaged component to manage it, as discussed previously. However, when many components in a topology are unmanaged, managing each component one after another, will be a cumbersome task! To

save the time and trouble involved in this exercise, the **DISCOVERED SEGMENTS** page provides you with the option to manage all the newly discovered components in the topology at one shot! For that, you need to click the **Manage all** button in Figure 11.17, pick the components you want to manage from the window that then pops up, and click the **Add/Manage** button to manage your selection.

11.3 Configuring Groups

To create a group, do the following:

1. Follow the menu sequence: Infrastrucuture -> Groups menu.
2. The **COMPONENT GROUPS** page that appears next will display all the existing groups. If no groups pre-exist, then a message indicated by Figure 11.18 will appear.

Figure 11.18: A message indicating that no groups are currently available

3. To add a new group, click the **Add New Group** button in Figure 11.18.
4. Figure 11.19 will then appear. Here, specify the name of the group that is to be created in the **Group name** text box.

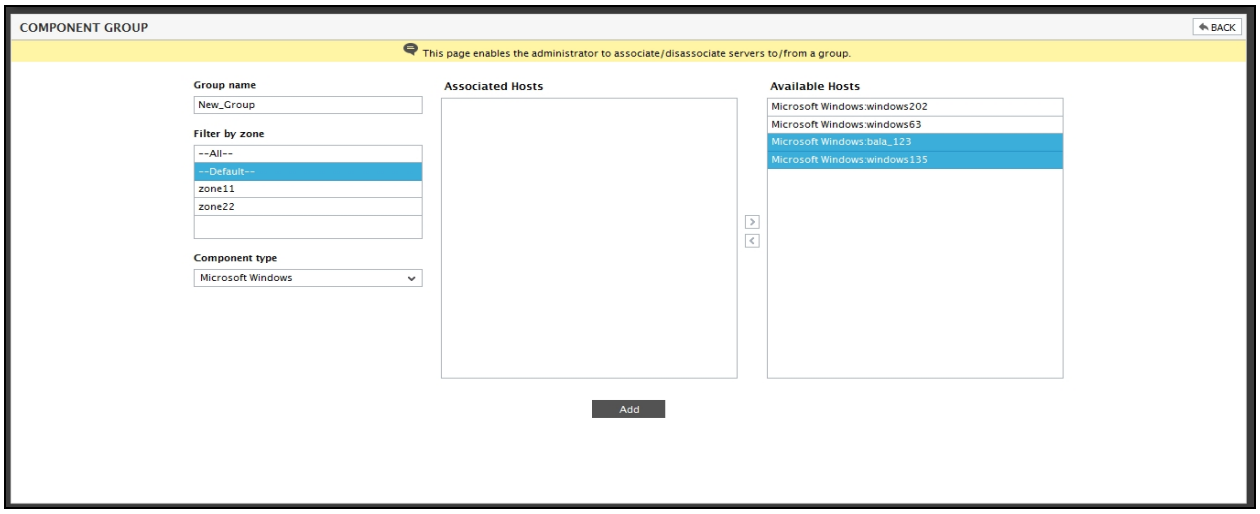

Figure 11.19: Providing the name of the group

5. If you want your group to contain a component that is already within a zone, then select that zone from the **Filter by zone** list. By default, the **--All--** option is displayed indicating that all components in the monitored infrastructure are available for creating a new group.
6. The **Component type** list will automatically be populated with all the component types that are associated with the chosen zone. Pick the component type of your choice from this list. The **Available Hosts** list will then be populated with the components of the chosen component type (see Figure 3). Select the hosts that you wish to add to the new group and click the **<** button. This will transfer your selection to the **Associated Hosts** list.

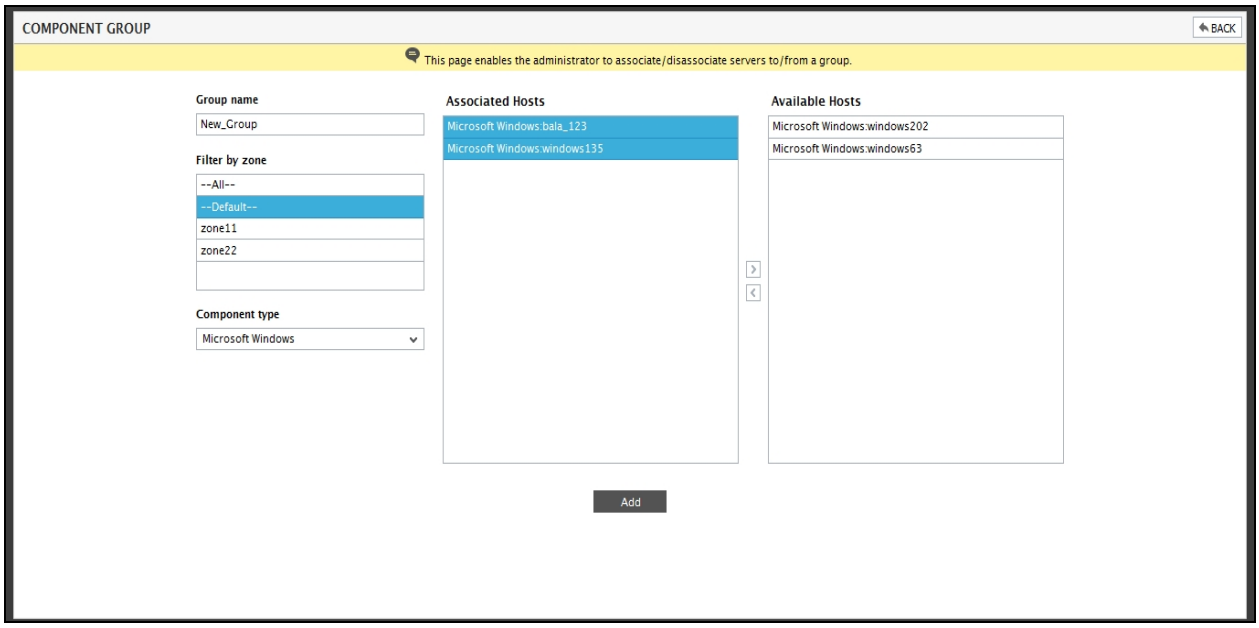

Figure 11.20: Selecting the components to be associated with the new group

7. To dissociate a host from the group, select the host from the **Associated Hosts** list and click the **>** button.
8. Finally, click the **Add** button in Figure 11.20.
9. You will then return to Figure 11.21, which will list the newly created group. To modify the group, click on the icon against the group name. To delete a group, click the icon against the group. To mark all the displayed groups for deletion, select the check box against the **Group Name** label and click the icon.

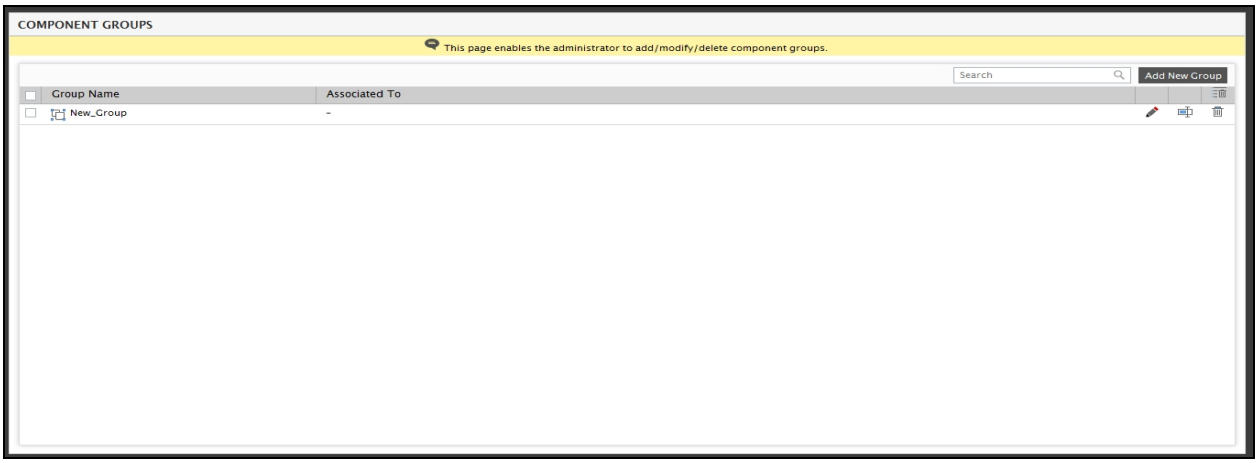

Figure 11.21: The newly created group been displayed in the LIST OF GROUPS page

10. If a large number of groups have been configured, locating the group to be modified/deleted would become quiet a challenge. This page therefore, allows you to quickly search for a group, by first specifying the whole/part of the group name in the **Search** text box, and then clicking the icon next to the text box. All groups with names that embed the specified string will then appear.

12

Configuring Services

So far we have not only seen how individual components are discovered and managed, but also how they can be interlinked into different segments. The different components working together deliver services for end users. The eG manager allows users to add one or more services for monitoring. A service can comprise just one or more independent components offering an end-user service (e.g., a web service offered by a web server).

Alternatively, a service can be a number of components working in conjunction. For example, a mobile payment service could involve a user accessing the service via an SMS gateway, which in turn hands of the request to a middleware application server. The application server could use a database for data storage / verification and rely on an external payment gateway. The collection of components and their interdependencies constitutes a service. Different business processes of an organization could be represented as a service. Note that a single segment topology can support multiple services and a single service may involve a subset of the components in a segment.

A service that can be associated only with a web server or a web application server is a web site. eG agents have the capability to monitor the performance of specific business functions supported by a web site. Such business functions are called *transactions* in the eG Enterprise system.

The topology configured in the previous section deals only with web servers and not with services. The next step in administering eG Enterprise is to configure services. To configure a service, an administrator has to first choose the **Services** sub-menu of the **Infrastructure** menu. Through the **Topologies** option of the **Service** sub-menu, the administrator can configure services.

Clicking on the **Topologies** option results in the currently configured services being displayed (see Figure 12.1). If the services listed include web sites, then such entries will be followed by the word 'Site' within brackets (see Figure 12.1). The existing details about the service can be modified using the  icon corresponding to a service. To delete a service, just click the  icon corresponding to it. If more than one service is to be deleted, mark the services for deletion by selecting the check boxes corresponding to them. Then, click on the  icon at the right, top corner of the page. To mark all the displayed services for deletion, click on the check box corresponding to the **Service Name** column label in Figure 12.1, and then click the  icon.

If a large number of services have been configured, locating the service to be modified/deleted would become quiet a challenge. This page therefore allows you to quickly search for a service, by first specifying the whole/part of the service name in the **Search** text box, and then clicking the  icon next to the text box. All services with names that embed the specified string will then appear.

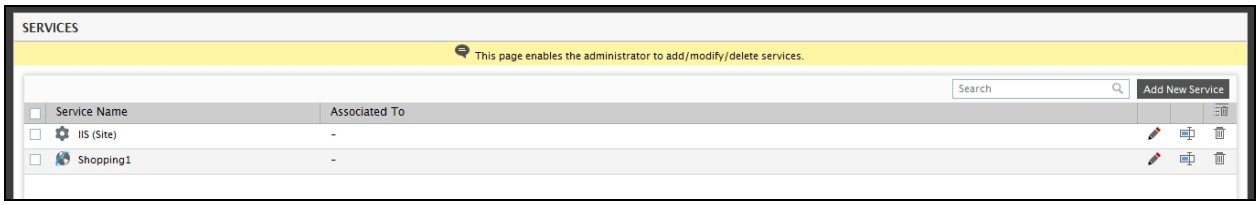

Figure 12.1: List of services configured

You can even change the name of a service by clicking on the **Rename** icon corresponding to it. Provide the **New name for the selected service** in Figure 12.2 that appears and click the **Rename** button therein to register the new name. This will automatically change the service name across the eG user interface.

Note:

Since renaming changes the name of a service across the user interface, you will no longer be able to access performance data of the service for the period prior to the name-change using the monitoring/reporting consoles. Therefore, exercise caution when renaming a service.

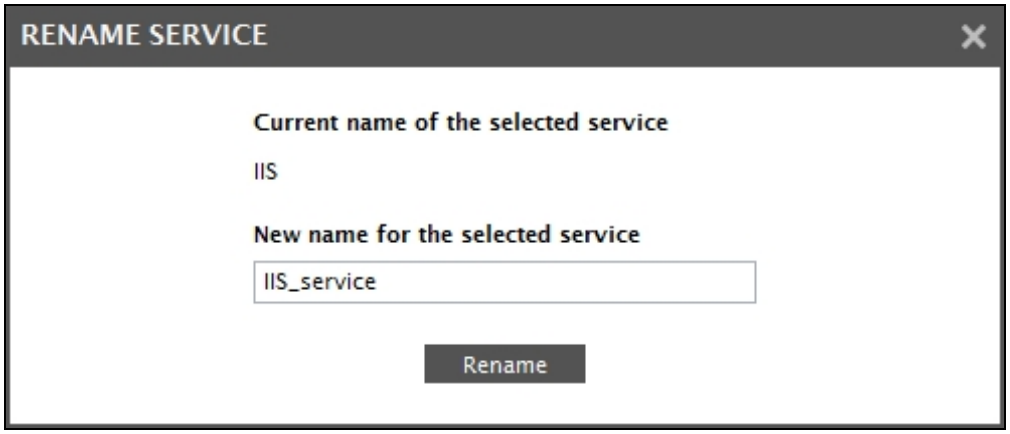

Figure 12.2: Renaming a service

New services can be added for monitoring by eG Enterprise using the **Add New Service** button in Figure 12.1.

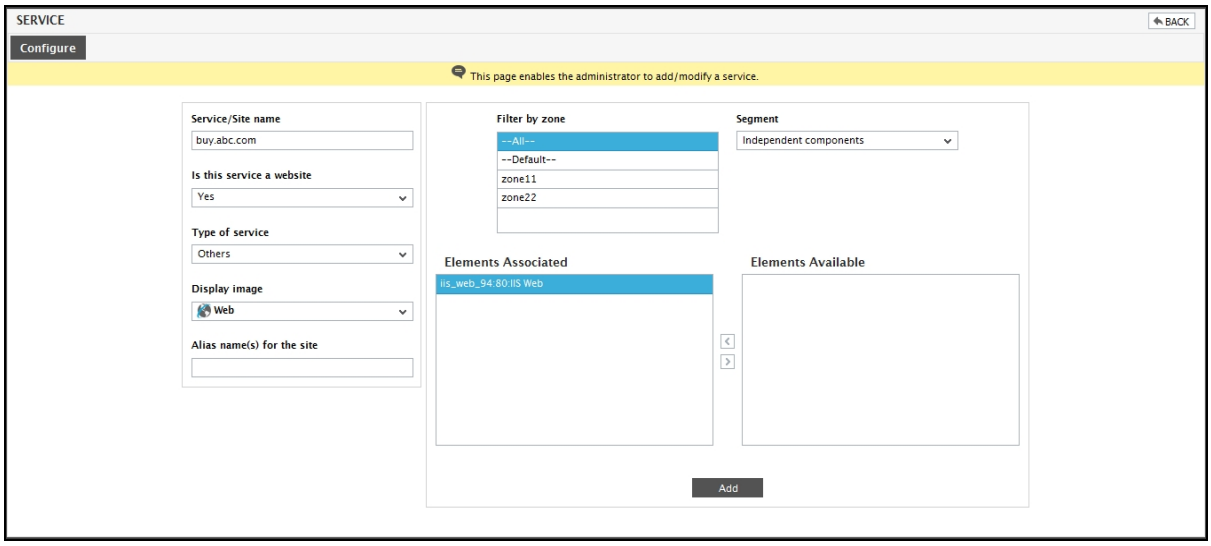

Figure 12.3: Adding a new web site

1. Using the **SERVICE** page (see Figure 12.3), an administrator can add a new web site / any other service. Figure 12.3 illustrates how the administrator can add a web site. The administrator has to specify the name of the new web site in the **Name of the Service/Site** text box. For e.g., when accessing a URL of the form http://www.abc.com/, the service name should be provided as www.abc.com.

   Note:

   You cannot add a service with the name *new*.

2. From the **Is this service a website?** list box, select **Yes** as shown in Figure 12.4 to configure a web site service. This indicates that the specified service is to be added as a web site. Then, **Choose a display image** for the service.

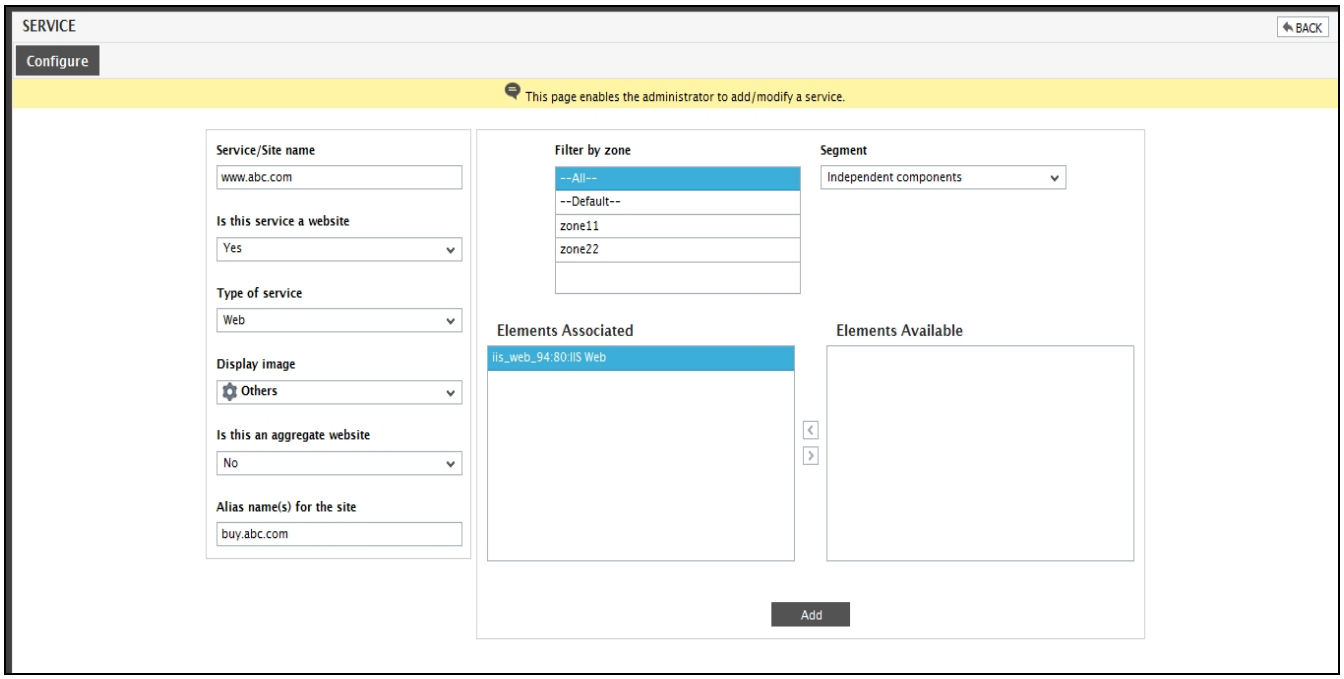

Figure 12.4: Configuring a web site

3. A single site can be addressed by various other names in the environment (e.g., www.abc.com may also be accessed as www.abc.com:80, abc.com, us.abc.com, 172.169.10.20 etc.). These names (or IP address:port combinations) can be specified in the **Alias name(s) for the site** text box. This input field is optional. To ensure that all requests to a website are captured, it is essential to ensure that all the alias names for a site are specified accurately. The administrators can specify a maximum of six alias names, each of which should be comma separated.

   While multiple alias names can be specified for a site, in the monitor interface, all the statistics pertaining to this web site will be reported using its site name that is specified in the **Name of the Service/Site** text box.

4. If you want your Service/Site to contain a component/segment that is already within a zone, then select that zone from the **Filter by zone** list. By default, the **--All--** option is displayed indicating that all independent components/segments in the monitored infrastructure are available for creating a new service/site. The **Segment** list of Figure 12.4 contains the list of fully configured segments in the target environment that contains at least a single web or application server. The site can be associated with any of these segments. If you choose Independent components from the **Segment** list, then all the independent components in the managed infrastructure will be listed in the **Elements Available** list. To associate the components to the Service/Site, select the component from the **Elements Available** list and click the **<** button. This will transfer your selection to the **Elements Associated** list. To remove an associated component, click the **>** button.

5. Click the **Add** button to add the service. The topology will then appear in a new **Topology** tab as shown in Figure 12.5.

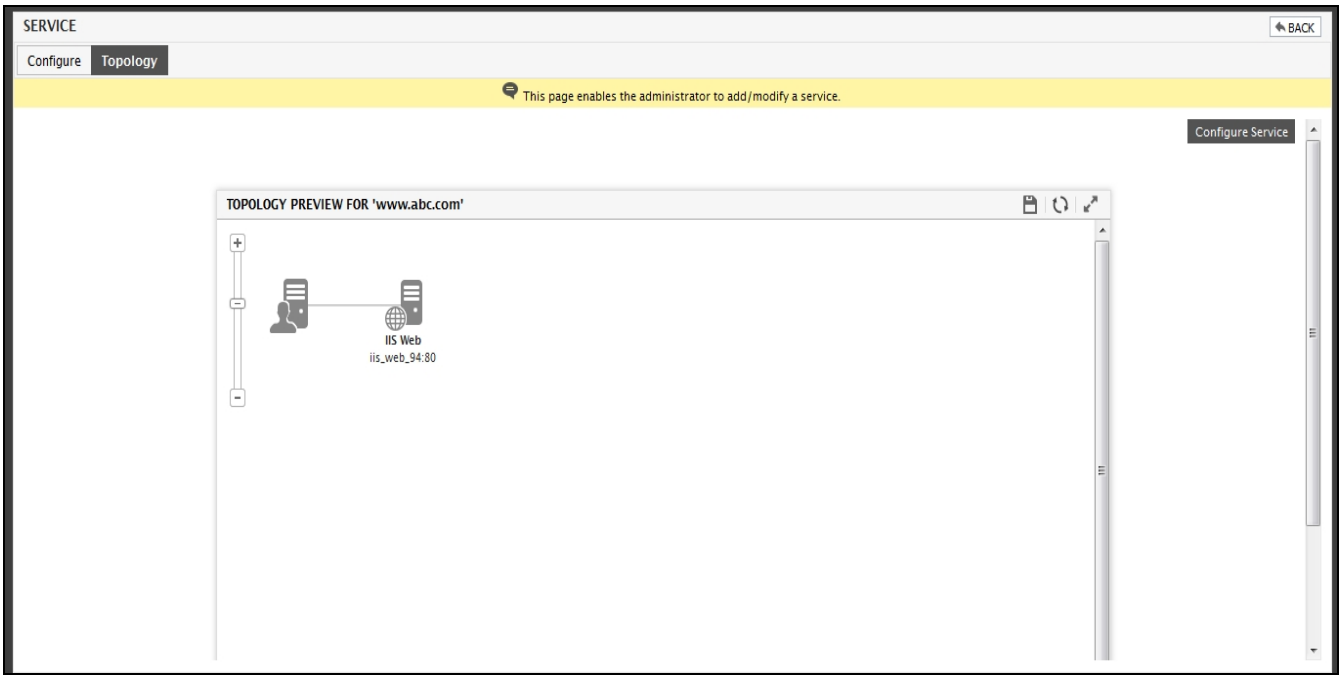

Figure 12.5: Viewing the Service topology

If you have chosen to add both segments as well as independent components from the **Segment** list of Figure 12.3, the Service topology as shown in Figure 12.6 will appear.

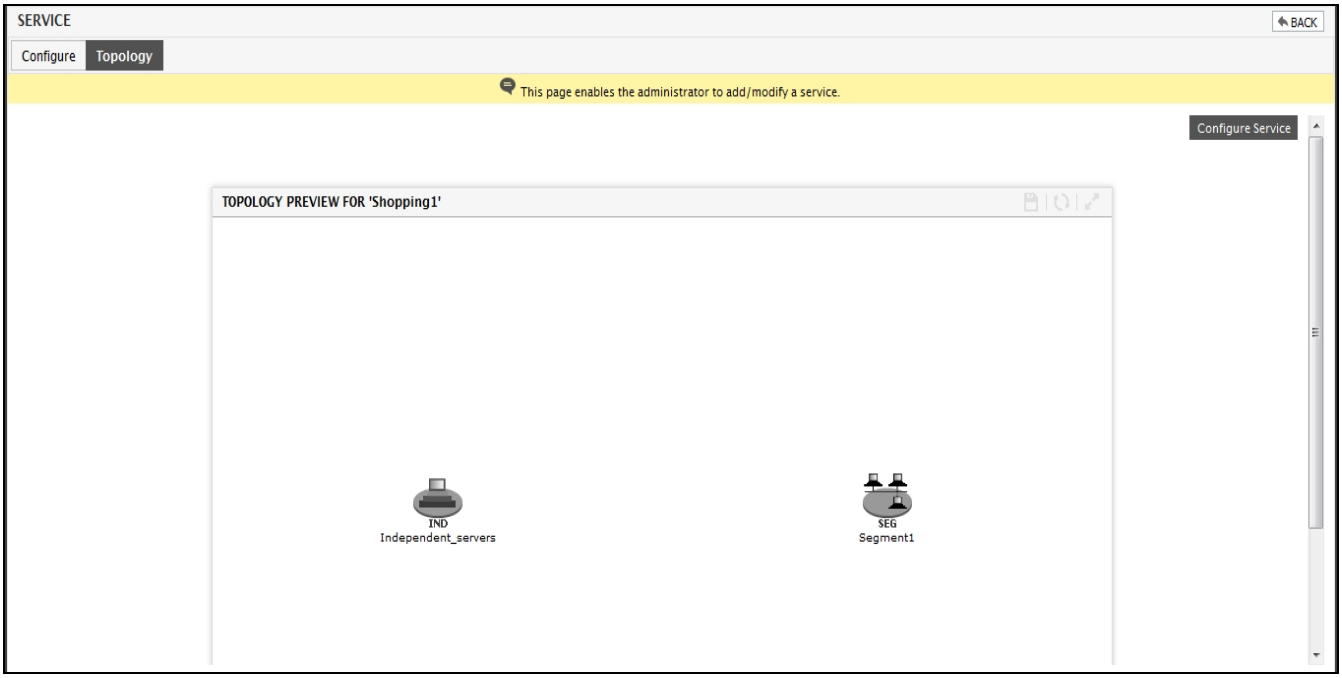

Figure 12.6: The Service topology that appears when a segment is associated with the service

Clicking on the segment will lead you to the topology of the segment as shown in Figure 12.7.

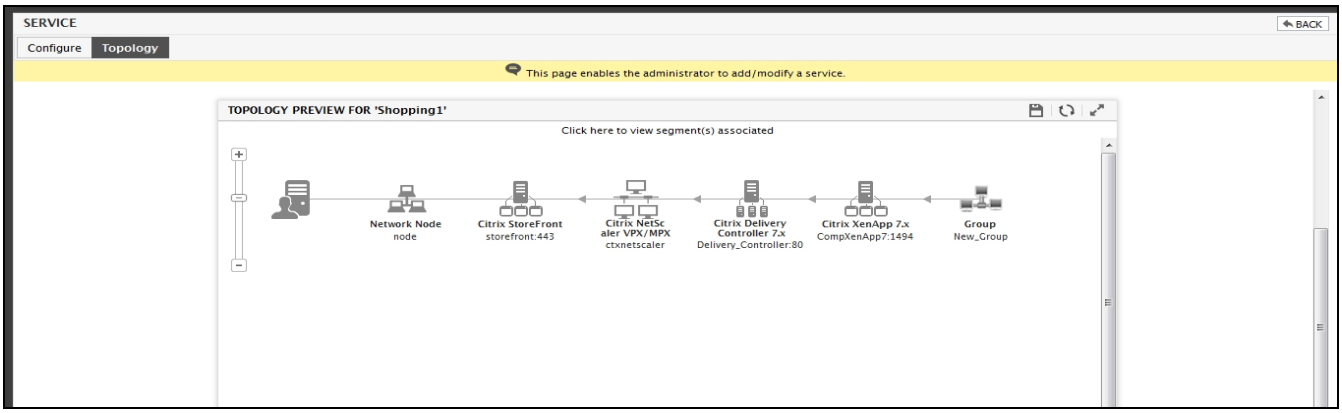

Figure 12.7: The topology of the segment associated with the service

Note:

By default, when a segment associated with a service is modified (by adding/removing components from the segment), the change will be automatically reflected in the service configuration. To ensure that such an automatic service updation does not occur, set the **Update_Service** parameter in the **[MISC]** section of the **eg_topology.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory) to **no** (by default, this will be set to **yes**).

12.1 Configuring Web Transactions

For each web site that has been configured, eG Enterprise has the ability to monitor individual transactions that happen via the web site. To configure a transaction for a web site, first, choose the **Transactions** option from the **Services** menu of the **Infrastructure** tile. Then, from the **Site name** list, pick the site for which you want to configure transactions. Doing so invokes all transactions that have already been configured for the chosen site (see Figure 12.8). You can delete any of the listed transactions by clicking on the **Delete** icon (i.e., the 'trash can') corresponding to that transaction in Figure 12.8. To delete two/more transactions at one shot, select the check boxes against those transactions in Figure 12.8 and click the **Delete Selected** icon that is available right below the **Add New Transaction** button in Figure 12.8.

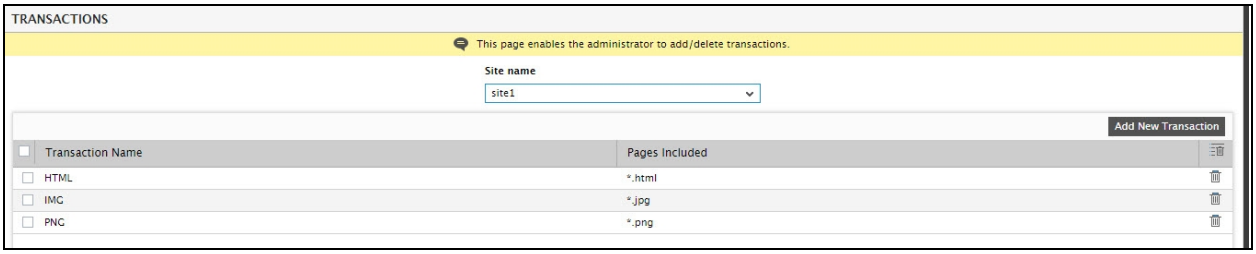

Figure 12.8: Configured transactions for a web site

Using the **Add New Transaction** button, you can add a new transaction for a web site. This will lead you Figure 12.9.

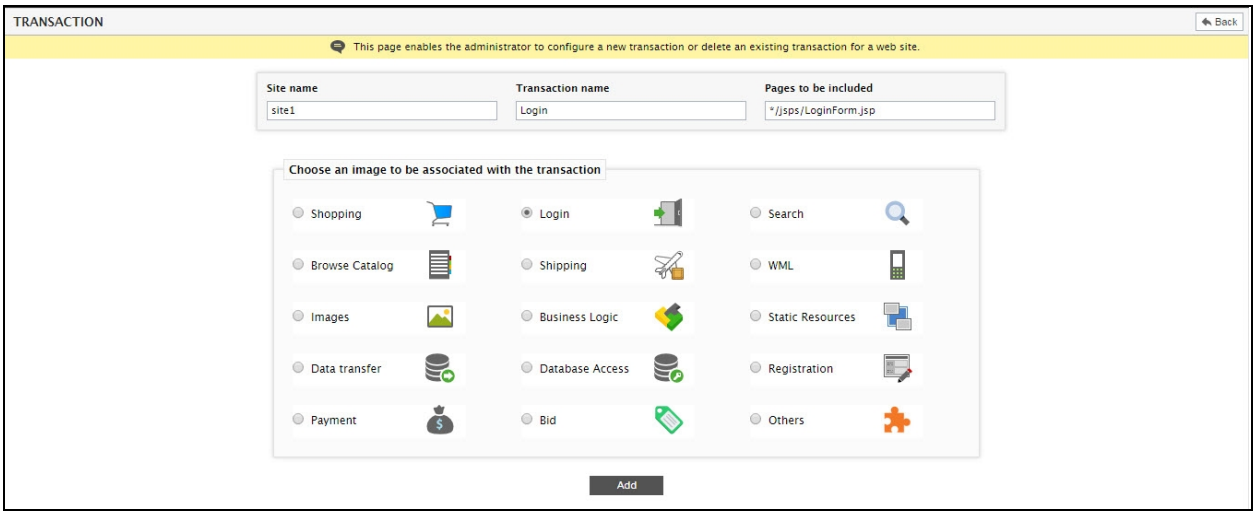

Figure 12.9: Details of transactions configured for a web site - buy.abc.com

Figure 12.9 depicts how a new transaction can be configured for a web site. First, a unique **Transaction name** has to be provided. Then, corresponding to each transaction, you have to specify the **Pages to be included** – i.e., one or more (comma-separated) regular expression patterns. Each pattern refers to a set of pages that constitute the transaction. There are two criteria that an administrator can use to define transactions that must be monitored:

Administrators can configure transactions to reflect the key operations performed by users of the web site. By monitoring individual transactions, web site operators can determine patterns of user accesses to individual transactions. Moreover, errors and response time issues with individual transactions can be monitored.

Note:

While mentioning the **Pages to be included**, ensure that the page names are prefixed by a slash (/) or an asterisk (*). If not, no measurements will be gathered from such pages.

Transactions can also be configured so as to differentiate between requests to the front-end web server and requests to the backend. For example, considering a web site that uses the iPlanet application server, all requests to the back-end application server can be represented by the pattern **/cgi-bin/gx.cgi** where * denotes zero or more characters. Using this approach, a web site operator can track requests sent to the back-end independent of requests targeted at the front-end web server and detect problems associated with the back-end easily.

12.2 Configuring Service Groups

Large organizations may have multiple services grouped under different business units. There may hence be a need to represent groups of services as an entity. To address this requirement, eG Enterprise allows the configuration of service groups in the eG admin interface, and represents the real-time state of the service groups in the eG monitoring interface.

To configure a service group, do the following:

1. Follow the menu sequence, Service -> Groups in the **Infrastructure** tile.
2. If no service groups pre-exist, then, a message to that effect will be appear in Figure 12.10. If groups pre-exist, then the groups will be listed as depicted by Figure 12.10.

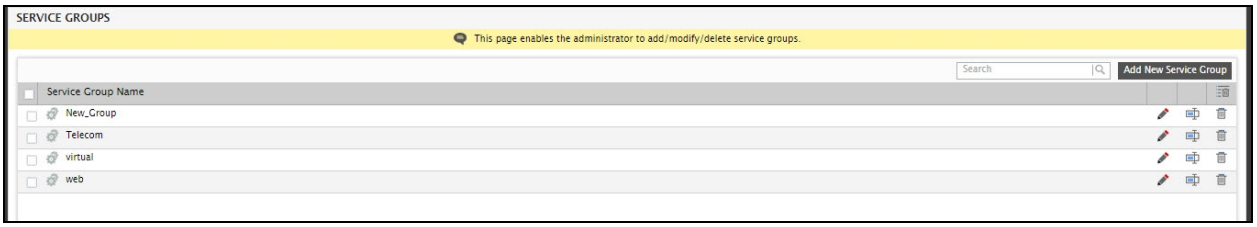

Figure 12.10: Creating a new Service Group

3. The configuration of an existing service group can be modified using the **Modify** icon (represented by the 'pencil') corresponding to it. To delete a service group, just click the **Delete** icon (i.e., the trash can) corresponding to it. If more than one service group is to be deleted, mark the groups for deletion by selecting the check boxes corresponding to them. Then, click on the **Delete Selected** icon at the right, top corner of the page. To mark all the displayed groups for deletion, click on the check box corresponding to the **Service Group Name** column label in Figure 12.10, and then click the **Delete Selected** icon. If a large number of groups have been configured, locating the group to be modified/deleted would become quiet a challenge. This page therefore allows you to quickly search for a service group, by first specifying the whole/part of the group name in the **Search** text box, and then clicking the 'magnifying glass' icon next to the text box. All service groups with names that embed the specified string will then appear.
4. Clicking on the **Add New Service Group** button in Figure 12.10 will invoke Figure 12.11. Here, specify the **Service group name**. Then, the list of all configured services across the infrastructure will be displayed in the **EXISTING SERVICES** list box in Figure 12.11. From this list box, select the services that are to be grouped under the new service group.

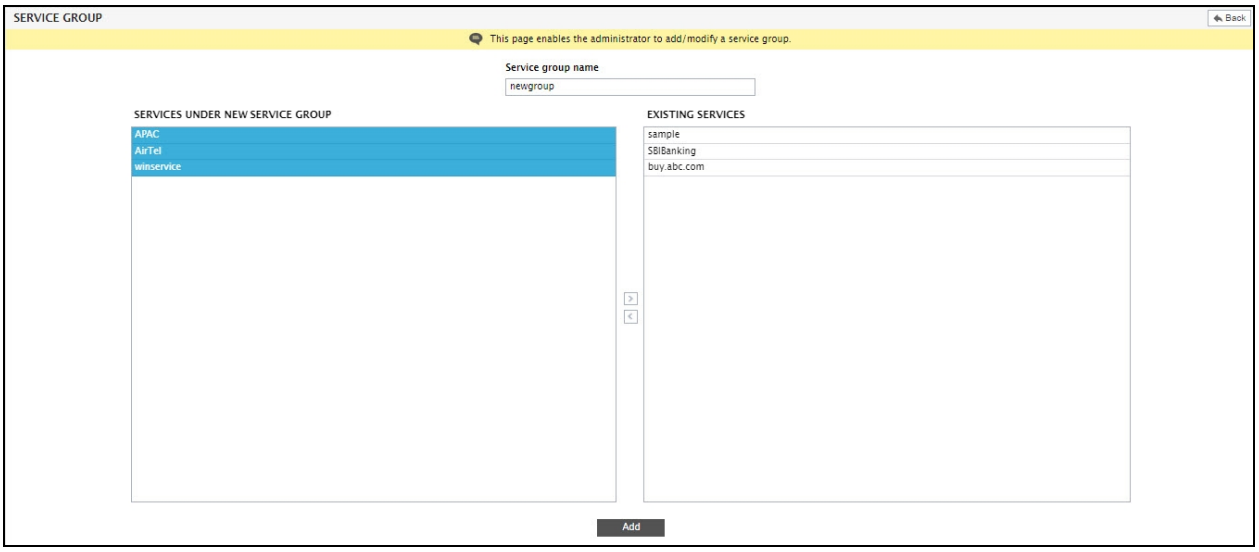

Figure 12.11: Configuring a service group

5. Next, click on the < button to add the chosen services to the new group. This will result in the display of the selected services in the **SERVICES UNDER NEW SERVICE GROUP** list box (see Figure 12.11).

6. To remove the service from the service group, select the services from the **SERVICES UNDER NEW SERVICE GROUP** list box and click on the **>** button as shown in Figure 12.11.
7. This in turn will shift the selected services back to the **EXISTING SERVICES** list box.
8. Finally, click on the **Add** button to add the service group.

   Note:

   A single service can be added to multiple service groups.

13

Configuring Zones

Large infrastructures spanning geographies can pose quite a monitoring challenge owing to the number of components involved and their wide distribution. Administrators of such infrastructures might therefore prefer to monitor the infrastructure by viewing it as smaller, more manageable business units. In eG parlance, these business units are termed **ZONES**. A zone can typically comprise of individual components, segments, services, and/or other zones that require monitoring. For example, in the case of an infrastructure that is spread across the UK, USA, and Singapore, a zone named *USA* can be created consisting of all the components, segments, and services that are operating in the US branch alone. The *USA* zone can further contain an East-coast zone and a West-coast zone to represent infrastructure and services being supported on the two coasts of the US.

While a service/segment contains a group of inter-related components with inter-dependencies between them, a zone contains a group of components, services, segments, or zones that may/may not have inter-dependencies.

To create a zone, do the following:

1. Follow the menu sequence: Infrastructure -> Zones.
2. A **ZONES** page will appear listing the existing zones (see Figure 13.1). If no zones pre-exist, then a message to that effect will appear here.

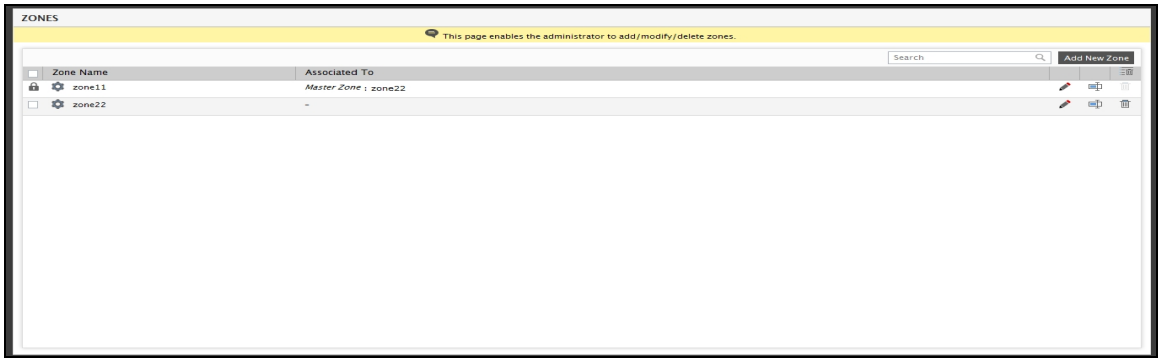

Figure 13.1: List of existing zones

To modify a zone, click on the icon corresponding to that zone in Figure 13.1. To delete a zone, click on the icon corresponding to it. To delete multiple zones at one shot, first mark the zones for deletion by selecting the check boxes corresponding to them. Then, click on the icon that is available just below the **Add New Zone** button in Figure 13.1. To mark all the displayed zones for deletion, click on the check box corresponding to the **Zone Name** column label in Figure 13.1, and then click the icon. If a large number of zones have been configured, locating the zone to be modified/deleted would become quiet a challenge. This page therefore allows you to quickly search for a zone, by first specifying the whole/part of

the zone name in the **Search** text box, and then clicking the icon in the text box. All zones with names that embed the specified string will then appear.

You can even change the name of a zone by clicking on the icon corresponding to it. Provide the **New name for the selected zone** in Figure 13.2 that appears and click the **Rename** button therein to register the new name. This will automatically change the zone name across the eG user interface.

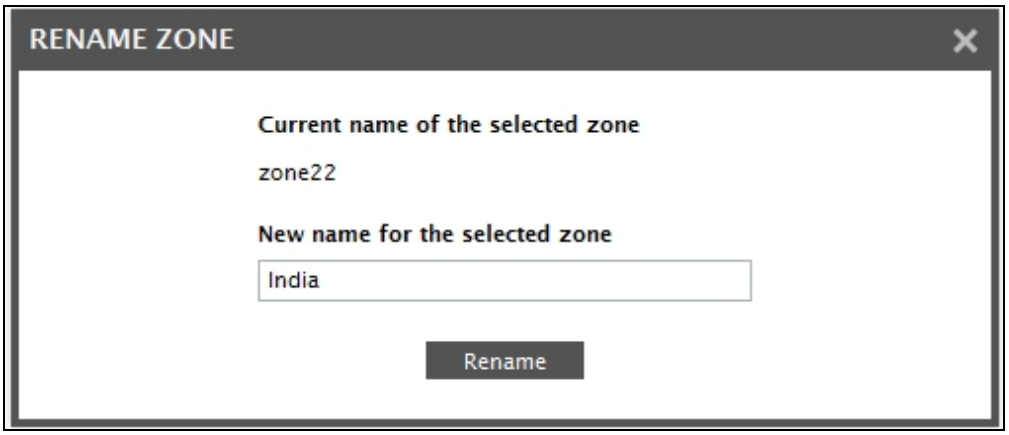

Figure 13.2: Renaming a zone

Note:

Since renaming changes the name of a zone across the user interface, you will no longer be able to access performance data of the zone for the period prior to the name-change using the monitoring/reporting consoles. Therefore, exercise caution when renaming a zone.

3. To create a new zone, click on the **Add New Zone** button in Figure 13.1.
4. Figure 13.3 will then appear. Specify the name of the zone in the **Zone name** text box.
5. Pick a display image for the zone from the **Display image** list.
6. Then, select the **Type of element** that you want to add to the zone – this can be a component/segment/service/service group/component group/zone. If you choose Component from the Type of element list, then the Component type list will appear.
7. Select the component type of your choice from the Component type list.

8. 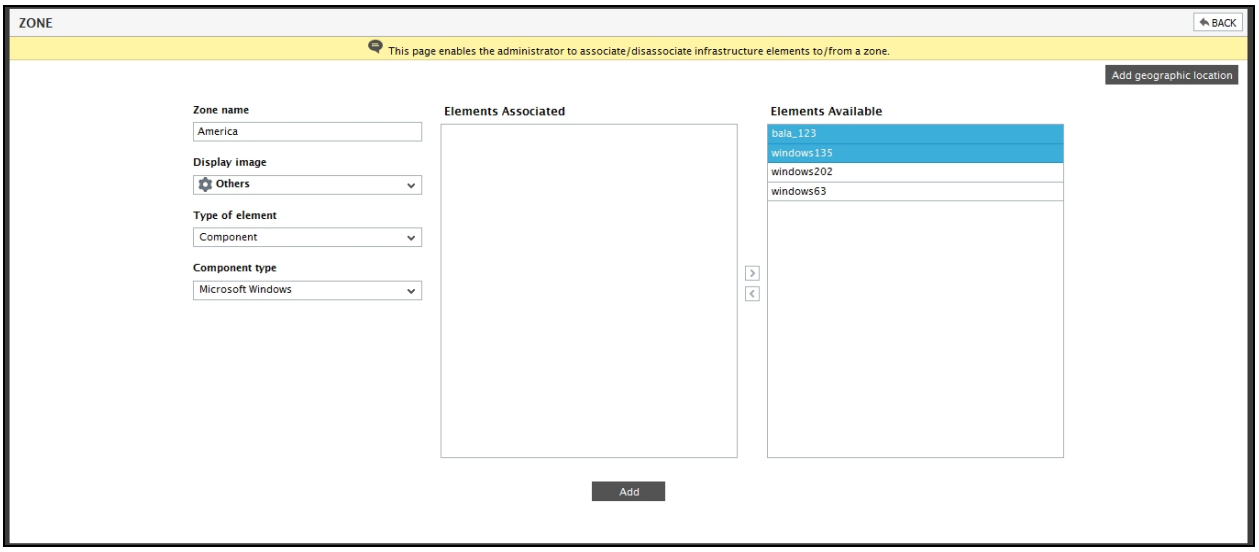

Figure 13.3: Selecting elements for association

9. All elements of the chosen **Type of element** will then be listed in the **Elements Available** list of Figure 13.3. An element can be a component of a particular component type/ segment/ service/group/ zone.

10. From the **Elements Available** list, select the elements to be added to the new zone and click the **<** button in Figure 13.3. This will transfer the selection to the **Elements Associated** list as depicted by Figure 13.4.

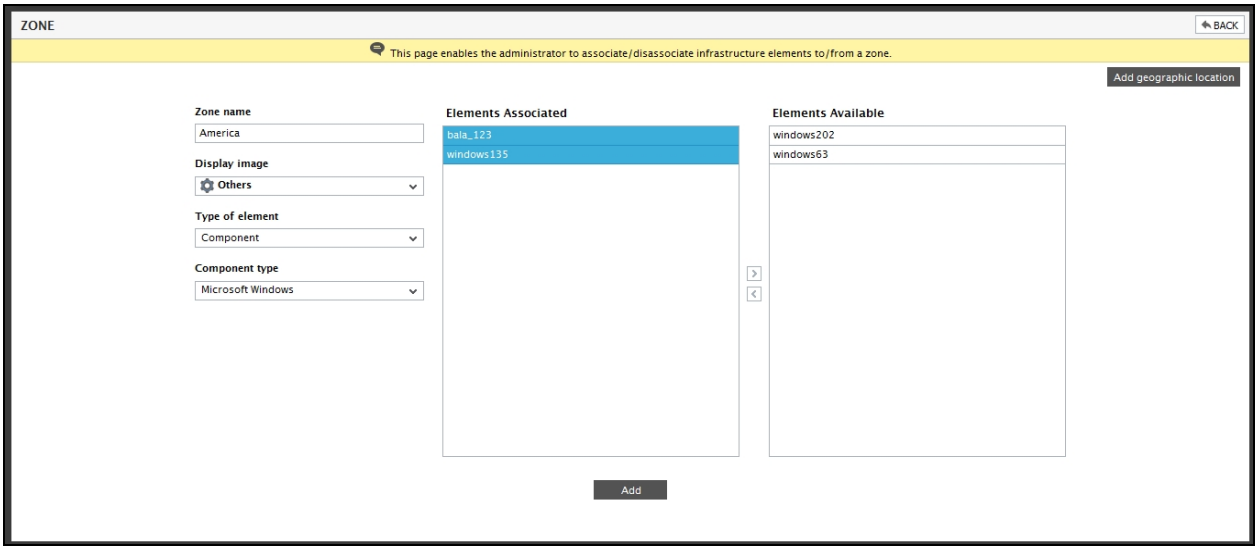

Figure 13.4: Associating elements with a zone

11. Similarly, multiple components/services/segments/zones can be associated with a new zone.

12. To disassociate an element from the zone, select the element from the **Elements Associated** list and click the **>** button in Figure 13.4.

13. Typically, zones can be used to represent the status of the IT infrastructure in a specific geographic location. eG Enterprise allows you to drill down on a geographic map to visually figure out the exact geographic area where a zone operates, and instantly evaluate the performance of the different zones

spread across the different locations worldwide. To enable such an analysis, you first need to indicate the geographic location of the configured zones using the built-in map interface of eG Enterprise To access this interface, click on the **Add geographic location** button in Figure 13.3. Figure 13.5 then appears revealing a world map. The geographic map display for zones is achieved through an integration of the eG Enterprise management console with the Google maps service.

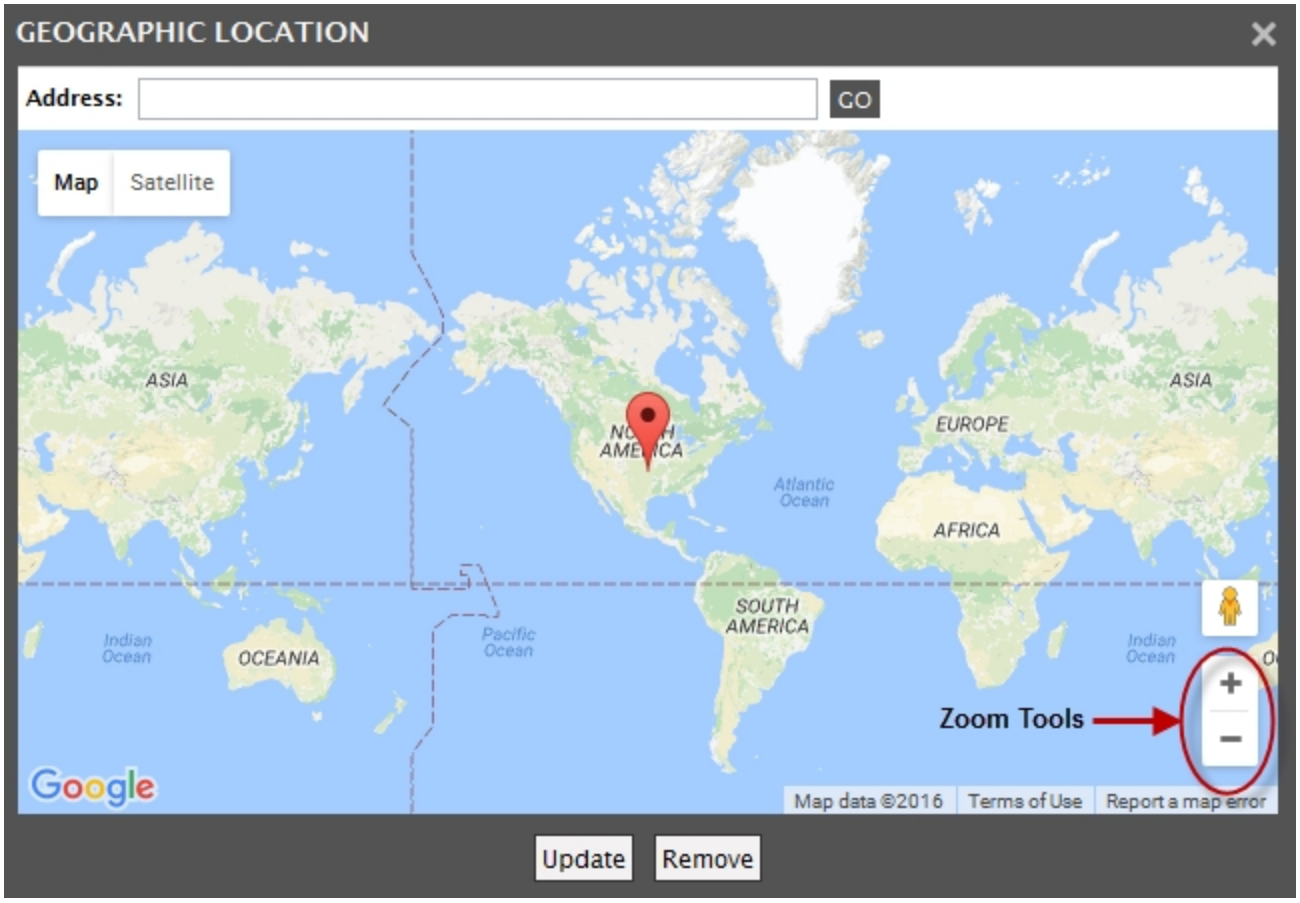

Figure 13.5: The eG Enterprise map interface

Note:

By default, eG Enterprise allows you to mark the location of a zone using its map interface. If, for some reason, you want the map option disabled, then, do the following:

- Edit the **eg_ui.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory
- By default, the **mapEnabled** parameter in the **[ZONE_MAP]** section of the **eg_ui.ini** file is set to **yes**. To disable maps, set this flag to **no**; doing so ensures that the **Add geographic location** link no longer appears in Figure 13.3.
- Finally, save the **eg_ui.ini** file.

Note:

If you are unable to view the zone map even after enabling it, then, follow the troubleshooting steps provided in the *Troubleshooting* chapter of this document to resolve the issue.

14. Use the arrow buttons indicated by Figure 13.5 to scroll right, left, up, or down, and the '+' and '-' buttons in the interface to zoom the map in and out, respectively. Now, to mark the geographic location of say, the *east_coast* zone, which represents the eastern coast of USA, first zoom into the **North American** continent in the world map. For that, move a little to the left of the map by clicking on the left arrow button in Figure 13.5. Then, to zoom into the country of USA in North America and view all its states as well, keep clicking on the '+' button in Figure 13.5 until USA and its states are visible in the map (see Figure 13.6).

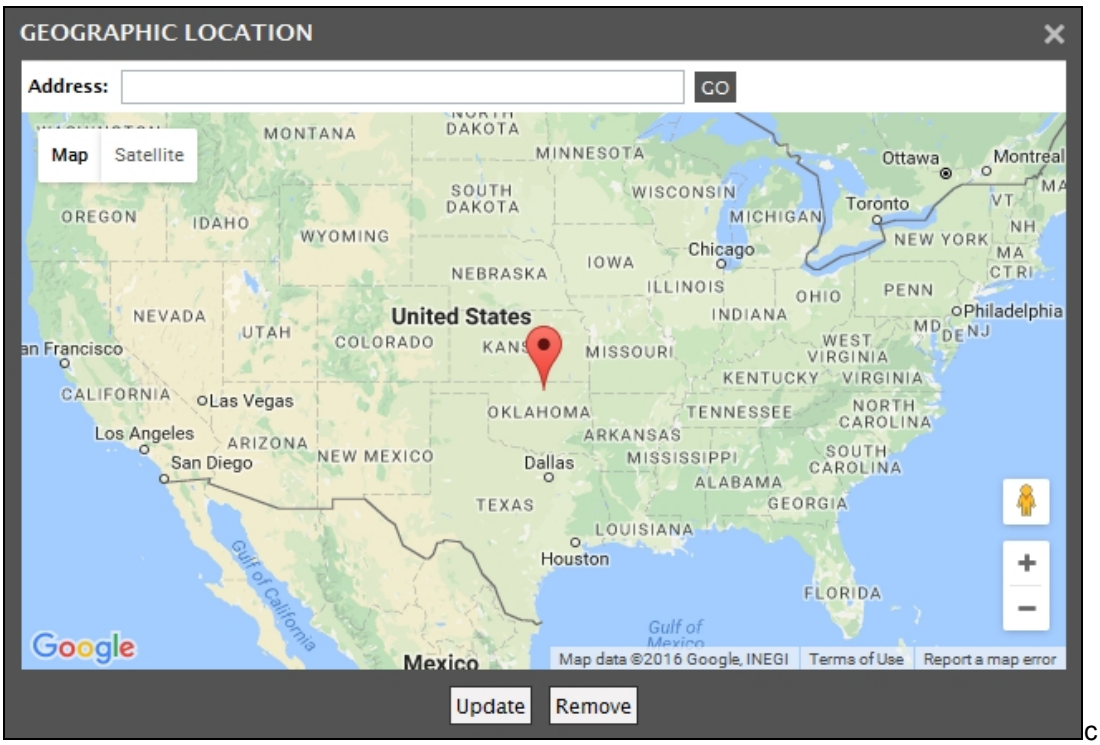

Figure 13.6: Zooming into the map to view USA and its states

15. Next, place your mouse pointer on any state in the east coast of USA, say New York, and then click on that state. To save the marking marked by the location pointer, click the **Update** button. To delete it, click the **Remove** button in Figure 13.7. To simply close the map interface, just click on the **X** button at the right, top corner of Figure 13.7.

16. Once you exit the map interface, you will return to the **ZONE** page. If you have marked the location of the zone in the map and have updated the marking, then the **Add Geographic location** link in the **ZONE** page will change to **Change / Remove geographic location**, as depicted by Figure 13.8. At any later point in time, you can make location changes in the map by clicking on this link.

Figure 13.7: Making location changes to the map

17. Once the zone configuration is complete, click on the **Back** button at the top of Figure 13.7 to open the **ZONES** page. The newly created zone will now appear here (see Figure 13.8).

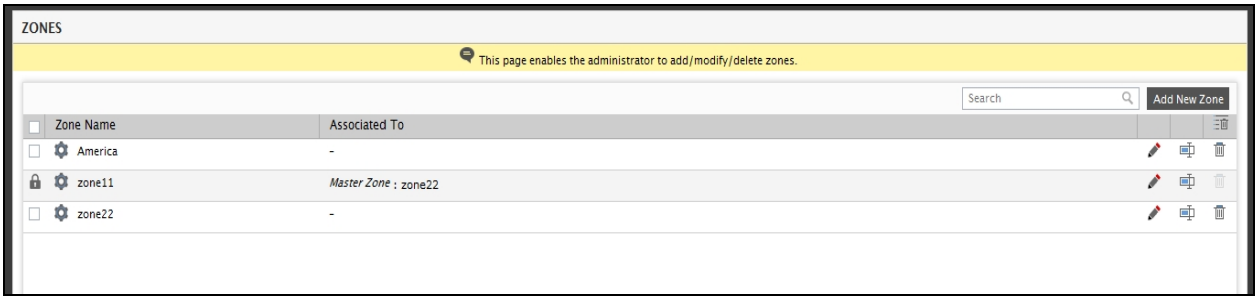

Figure 13.8: The newly created zone appearing in the LIST OF ZONES page

18. If any of the listed zones is a member of another zone, then such 'child' zones will be preceded by a 'lock' symbol.

Note:

A zone can be associated with multiple zones. For example, a zone named **newyork** can be added to the *east_coast* zone and also to the *USA* zone. However, a parent zone cannot be added to a child zone. For example, you can add **newyork** to *east_coast*, but subsequently, you cannot add *east_coast* to **newyork**.

14

Metric Aggregation

eG Enterprise typically monitors every component of a type, separately. However, sometimes, business owners may require aggregate metrics about their infrastructure. For instance, Citrix administrators might want to know the total number of users who are currently logged into all the Citrix servers in a farm, so that sudden spikes in the load on the farm (as a whole) can be accurately detected. Similarly, Windows administrators might want to figure out the average CPU usage across all the Windows servers in an environment, so that they can better plan the capacity of their Windows load-balancing clusters.

To provide such a consolidated view, eG Enterprise embeds a license-controlled **Metric Aggregation** capability. This feature, when enabled, allows administrators to group one or more components of a particular type and monitor the group as a single logical component, broadly termed as an *aggregate* component.

Figure 14.1 depicts how metrics aggregation works.

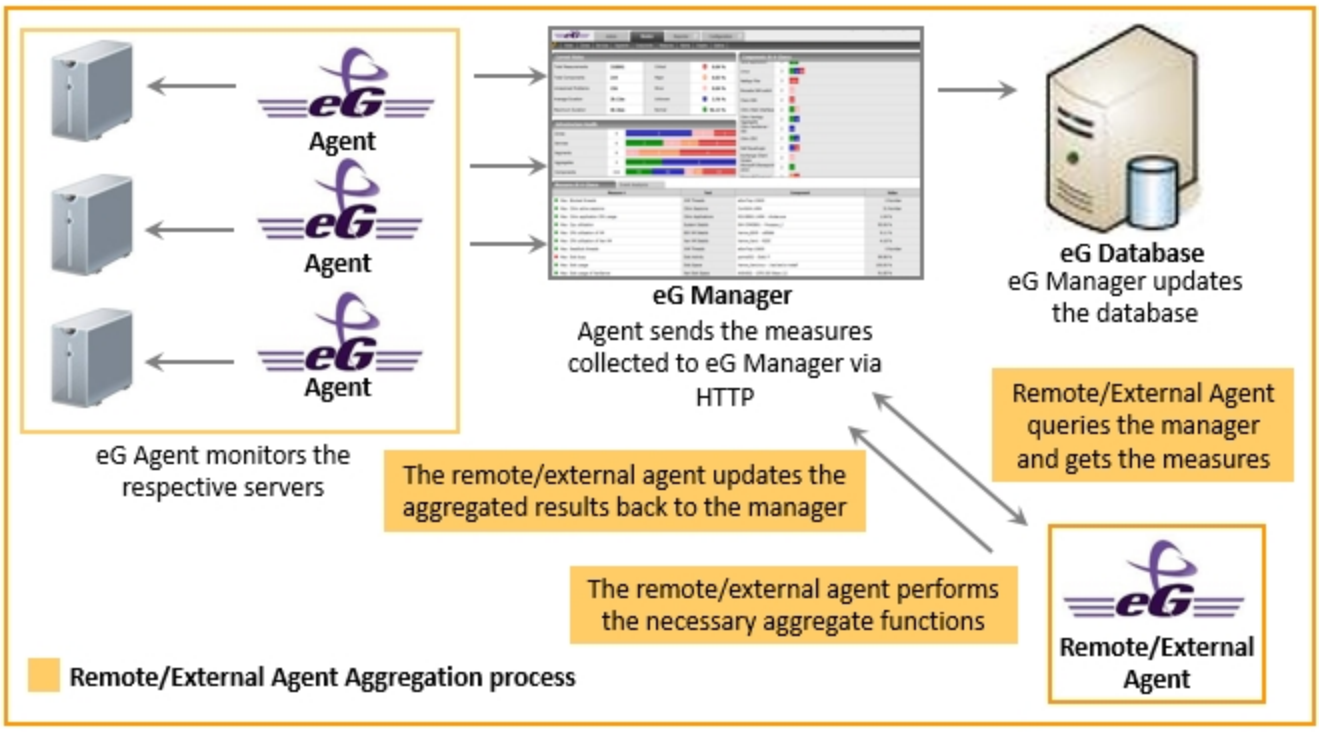

Figure 14.1: How metrics aggregation works?

Typically, every eG agent that is monitoring a component grouped under an *aggregate component*, reports metrics collected from the corresponding component to the eG manager. The eG manager consolidates the metrics so received and stores them in the eG backend. To aggregate these metrics, eG Enterprise **either requires that an external/remote agent be installed.**

Note:

- The remote agent that performs metrics aggregation can be installed on any remote host in your environment.
- No special privileges need be granted to a remote agent for performing metrics aggregation.
- It is recommended that you monitor a maximum of 10 *aggregate components* alone using a single remote agent.

Every time an aggregate test runs, this remote/external agent queries the eG manager for the metrics collected from the member components of the managed *aggregate component*. On receipt of the metrics, the remote/external agent applies pre-configured aggregate functions on the metrics, and updates the eG manager with the aggregated performance results.

Separate thresholds need to be set for the aggregated metrics to track deviations in the consolidated performance. The state of the *aggregate component* is governed by these exclusive thresholds, and not by the state of the components within the group.

Metrics Aggregation is a license-controlled capability. If remote agents are used to perform metric aggregation, one/more **premium monitor licenses** would be required for implementing this capability. However, if an external agent is used to perform metric aggregation, no license is required for implementing this capability.

Using this **Metric Aggregation** capability, administrators can perform the following:

- Effectively assess the collective performance of a group of components of a particular type
- Easily study load and usage trends of server farms (or groups) as a whole
- Accurately detect resource inadequacies or unusual load conditions in the component group or farm
- Compare and correlate the performance of the member components with that of the *aggregate component*, so that the reasons for performance issues with the *aggregate component* can be precisely determined;

Note:

Configuration tests are not applicable for *Aggregate Components*.

14.1 Adding Aggregate Components

The very first time the administrator manages/adds a component of a type, the eG Enterprise system dynamically creates a corresponding *aggregate component type*. For instance, if an *IIS web* server component is added to the eG Enterprise system for the first time, a component type named, *IIS Web Aggregate* is automatically created alongside.

To add a component of the dynamic *aggregate component type*, do the following:

1. In the eG administrative interface, select the **Add/Modify** option from the **Aggregates** menu of the **Infrastructure** tile.
2. Figure 14.2 then appears. Figure 14.2 is characterized by a tree-structure in the left panel, and a context-sensitive right panel, which changes based on the node chosen from the tree. Each node in the tree represents a category of infrastructure elements, namely - managed components, segments, services, (component) groups, and zones - which typically provide the components that can be aggregated.

Expanding the **Segments** node for instance will display all the fully-configured segments in the environment as sub-nodes. Likewise, you can expand the **Services**, **Groups**, or **Zones** node in the tree to view the services, groups, and zones (respectively) that have been registered with the eG Enterprise system. While clicking on the **Expand All** link below the tree will expand all nodes at one shot, the **Close All** link, when clicked, will close all expanded nodes simultaneously.

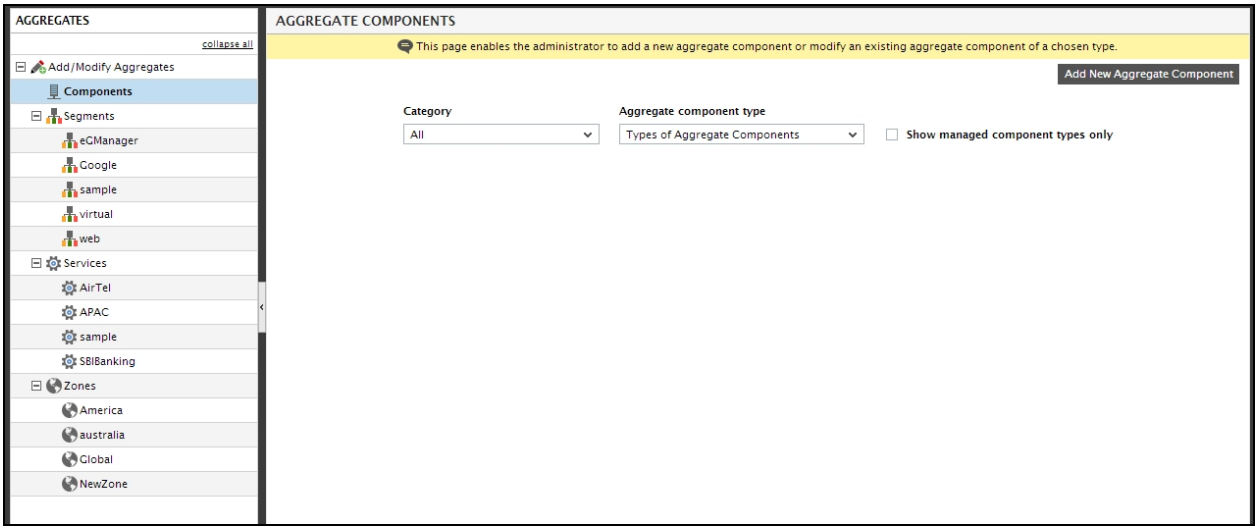

Figure 14.2: The ADD/MODIFY AGGREGATE COMPONENTS page

By clicking on the nodes of the tree, you can create an *aggregate component* using any managed component across the environment, or using the components that are included in a particular segment, service, zone, or component group.

The sections that follow will discuss each of these procedures in detail.

14.1.1 Creating an Aggregate Component Using Any Managed Component of a Type in the Target Environment

By default, the **Components** node is chosen from the tree in the left panel of Figure 14.2. With the **Components** node chosen, you can create *aggregate components* using any of the managed components from across the environment. This is why, by default, the **Aggregate component type** list in the right panel of Figure 14.2 will display all the *aggregate component types* that were dynamically created by the eG Enterprise system whenever components of the corresponding 'non-aggregate' types were added/managed using the eG administrative interface (see 14.1.1).

Note:

If you have created a new component type using the Integration Console plugin of eG, then, once you add a component of this new type, then eG Enterprise will dynamically create an aggregate component type that corresponds to the new type as well. Moreover, the **Aggregate component type** list in the right panel of Figure 14.2 will include new type. This means that you will be able to create an aggregate component of even custom component types using Figure 14.9.

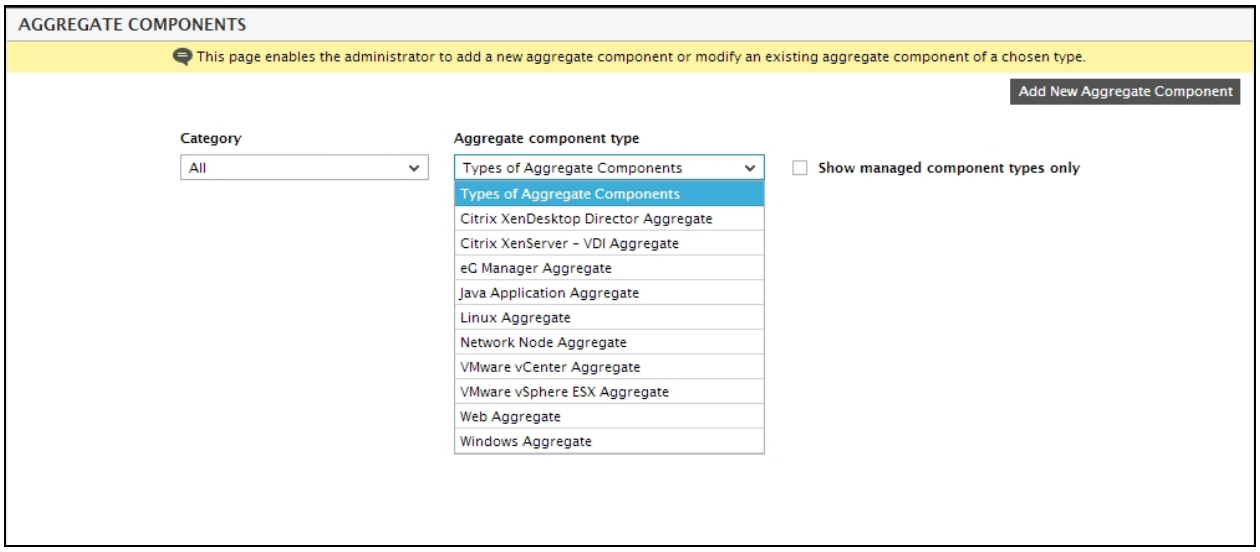

Figure 14.3: Viewing the component types that are available for selection in the Aggregate typelist

In other words, if an Active Directory server was added/managed using the eG administrative interface, it would have resulted in the automatic creation of an Active Directory Aggregate component-type; this aggregate component type will be available for selection in the **Aggregate component type** list in the right panel of 14.1.1.

To begin creating an *aggregate component*, do the following:

1. Pick a **Category** from the right panel of Figure 14.4.
2. Pick an **Aggregate component type**. If too many components of various types are managed in an environment, then the **Aggregate type** list in such environment will be very long; this is because, for every 'non-aggregate' component-type that is managed, a corresponding **Aggregate type** will be automatically created. Selecting an **Aggregate type** from this long list will obviously be difficult. To condense this list, you can choose to view in the list only those **Aggregate component types** for which at least a single *aggregate component* has already been created/managed. For this, select the **Show managed component types only** check box in Figure 14.4.
3. Once an **Aggregate component type** is selected, all the *aggregate components* of that type will be displayed therein. If no such components pre-exist, then a message to that effect will appear (see Figure 14.4).

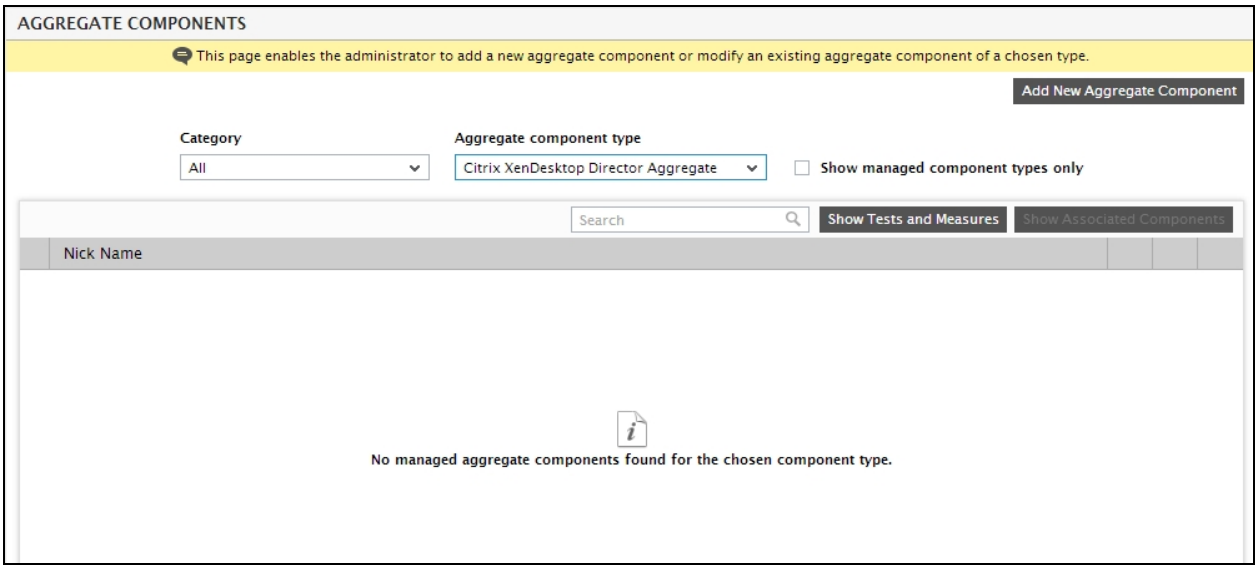

Figure 14.4: The message that appears when no components of a chosen aggregate type pre-exist

4. To add a new aggregate component of the chosen type, click on the **Add New Aggregate Component** button in Figure 14.4. This will invoke Figure 14.5, where you can provide the details of the new aggregate component.
5. Since an *aggregate component* is just a logical component and not a physical one, you need not assign an IP address to that component. Therefore, just assign a **Nick name** to the component (see Figure 14.5).
6. Next, since only a remote agent can perform metric aggregation, choose the **Remote agent** to be assigned to the aggregate component being added.
7. If web servers are aggregated, then an additional **Site Support** flag will appear in Figure 14.5. By default, this flag is set to **No**. You can set the flag to **Yes**, if you want to create an *aggregated web site* using the web aggregate component being created. To know more about creating *aggregated web sites*, refer to Section Section **14.3** of this document.
8. By default, the **AVAILABLE COMPONENTS** list in Figure 14.5 will display all the individual components that are candidates for inclusion as member components of the **Aggregate component type** chosen. For instance, when creating an *IIS Web Aggregate* component, all managed components of type *IIS Web* will be displayed by default in the **AVAILABLE COMPONENTS** list. If one/more of these individual components are included in a service, segment, zone, or component group, an additional **Filter by** list will be available (as indicated by Figure 14.5). By selecting an option from the **Filter by** drop-down, you can filter the **AVAILABLE COMPONENTS** list, so that it displays only those individual components that belong to a selected **Segment**, **Service**, **Zone**, or component **Group**. To associate one/more of these individual components with the *aggregate component* being added, select them from the **AVAILABLE COMPONENTS** list, and click the **<** button in Figure 14.5. This will transfer the selection to the **ASSOCIATED COMPONENTS** list (see Figure 14.6) On the other hand, to disassociate any of the member components of an *aggregate component*, select them from the **ASSOCIATED COMPONENTS** list and click the **>** button.

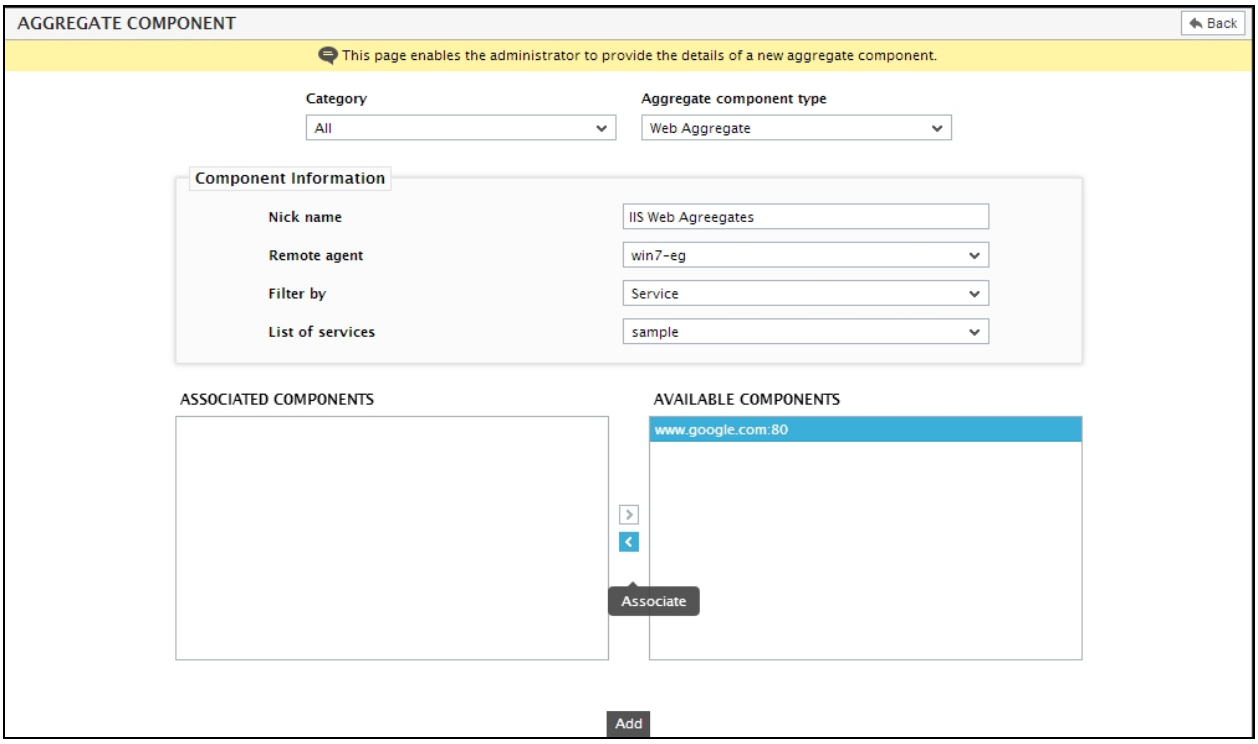

Figure 14.5: Selecting the components to be aggregated

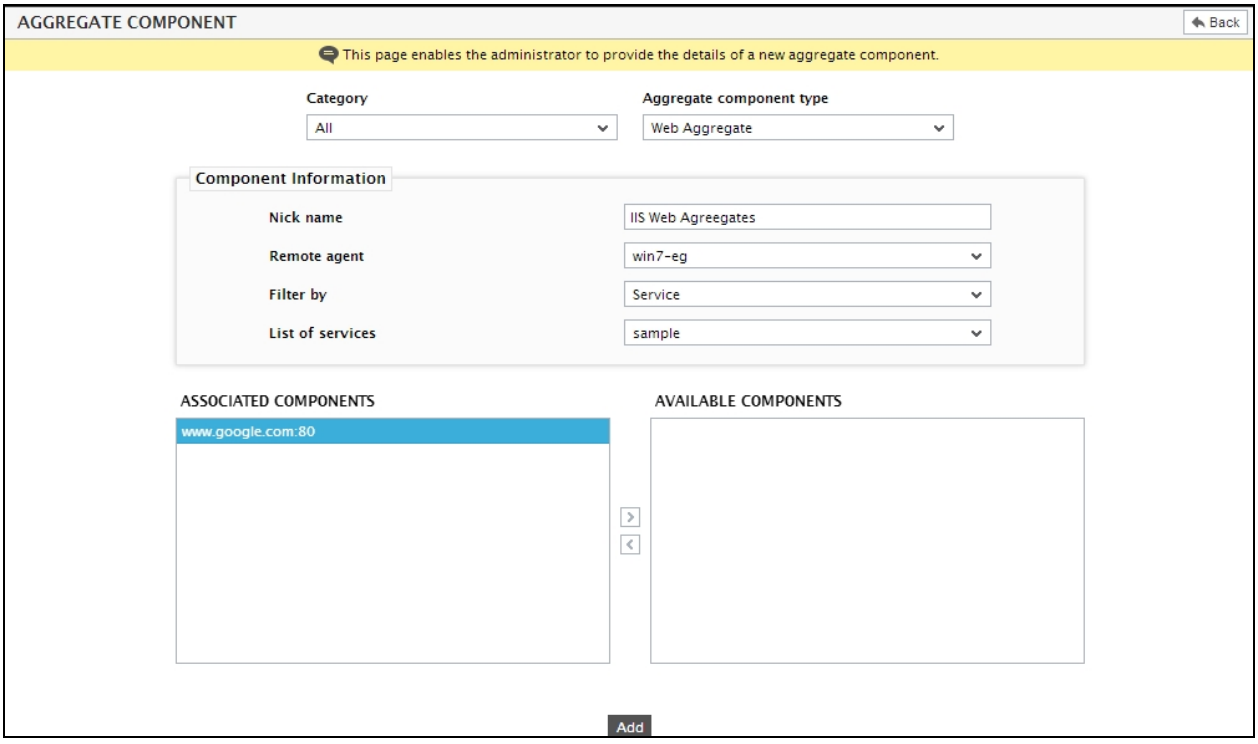

Figure 14.6: Associated components with an aggregate component type

9. To add the newly created aggregate component, click the **Add** button in Figure 14.6. Figure 14.7 will then

appear displaying the aggregate component that was just created.

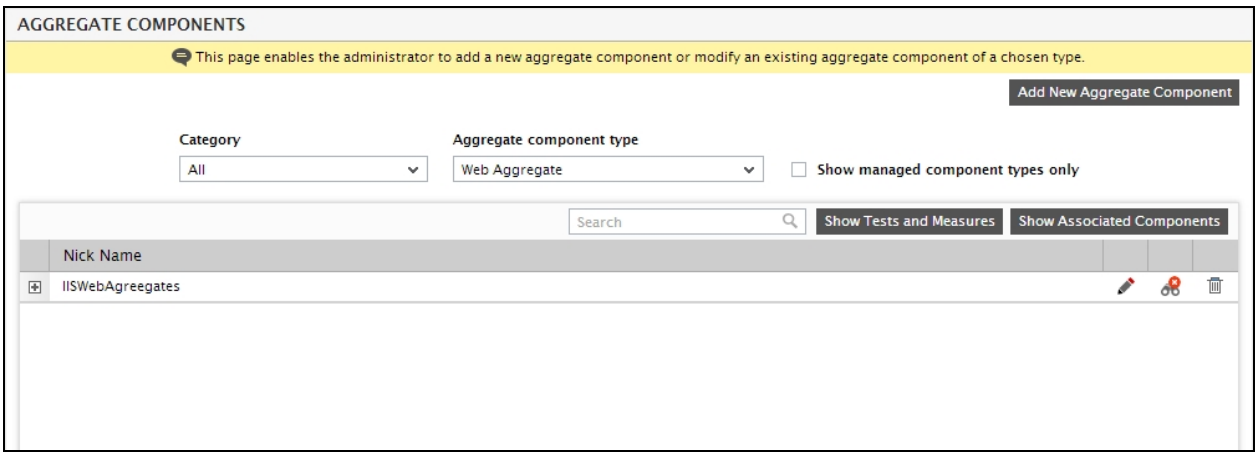

Figure 14.7: The aggregate component that was newly created being listed

10. In the same way, multiple *aggregate components* of a type can be added.
11. You can modify the details of the new component by clicking the **Modify** icon (i.e., the 'pencil' icon) against that component in Figure 14.7. Likewise, you can delete the newly added component by clicking the **Delete** icon (i.e., the trash can icon) against it. You can even unmanage an aggregate component by clicking the **Unmanage** icon corresponding to that component in Figure 14.7. This will lead you the **AGGREGATE COMPONENTS – MANAGE/UNMANAGE** page, using which you can unmanage that aggregate component.
12. Every aggregate component type so created is automatically associated with a default set of tests. To view the tests that are available by default for a particular aggregate component type and the measures they report, click on the **Show Tests and Measures** button in Figure 14.7. Figure 14.8 will then appear.

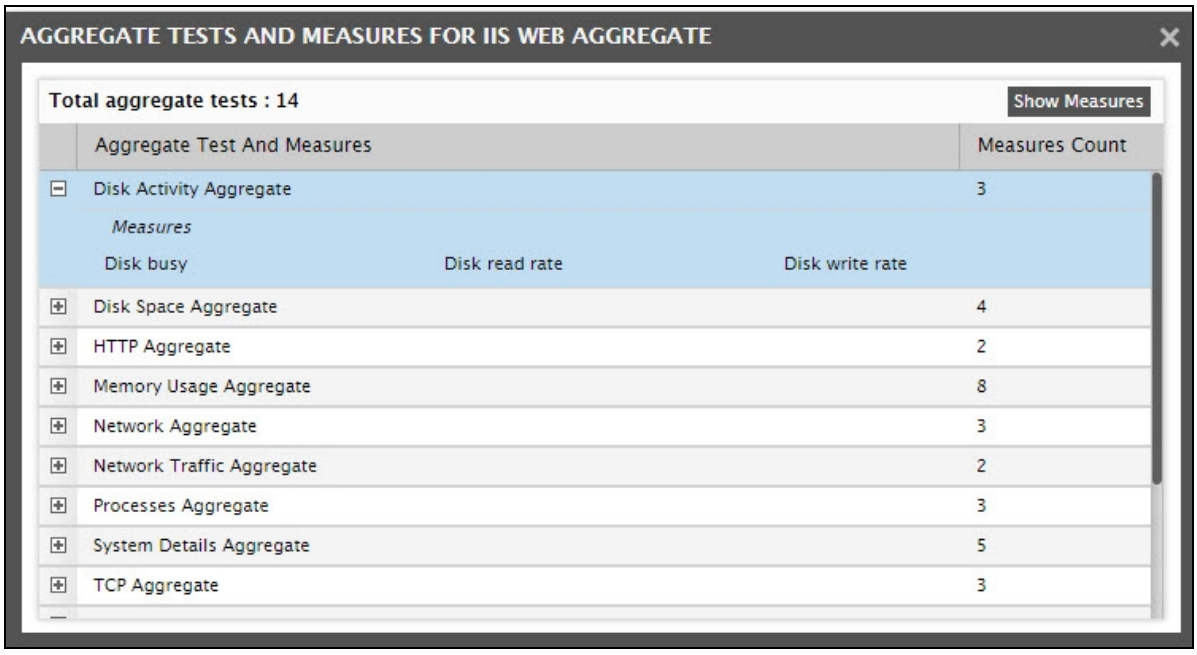

Figure 14.8: Viewing the tests that are available by default for a particular aggregate component type

13. To view the measures that will be reported by a particular test mapped to an aggregate component type, expand the test node as indicated by Figure 14.8. To view all the measures that will be reported by all tests mapped to an aggregate component type, click the **Show Measures** button in Figure 14.8. Figure 14.9 will then appear.

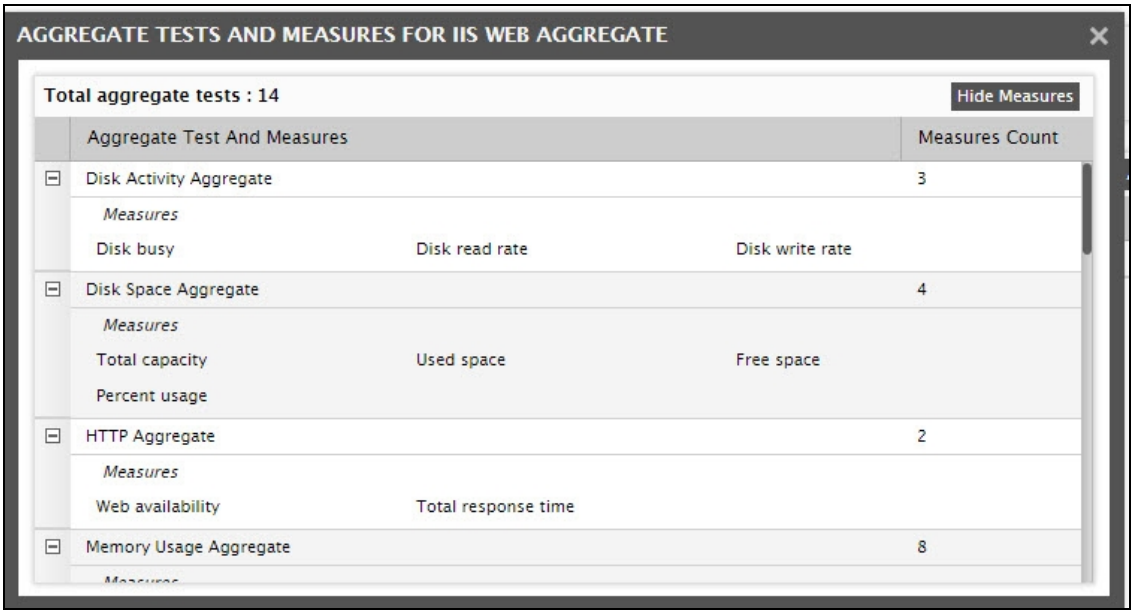

Figure 14.9: Measures of all aggregate tests mapped to an aggregate component type displayed

14.1.2 Creating an Aggregate Component Using the Components that are Part of a Segment/Service/Zone/Component Group

If you want to aggregate the components of a type that is included as part of a segment/service/zone/component group, then expand the corresponding node in the tree and click on the segment/service/zone/component group of interest to you.

For instance, to create an *aggregate component* using the individual components of a type in a service, do the following:

1. Expand the **Services** node in the tree structure and click on the service of interest to you (see Figure 14.10).
2. Depending upon the types of components that are part of the chosen service, the eG Enterprise system automatically identifies the *aggregate component types* that can be created using those components and displays the same in the right panel.

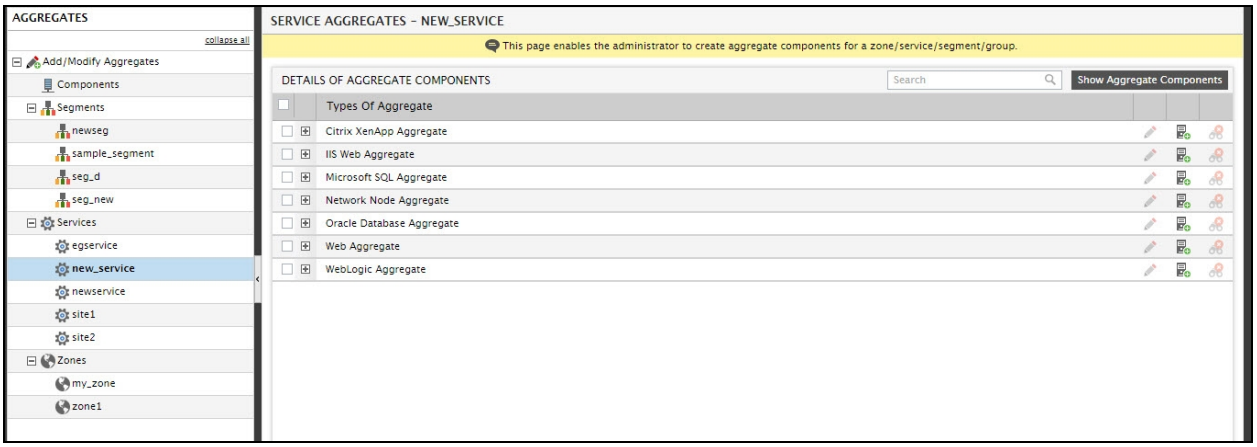

Figure 14.10: The aggregate component types that can be created using the components of a chosen segment

3. To create an *aggregate component* of one of the listed types, click on the **Add** icon corresponding to that type in Figure 14.10. Figure 14.11 will then appear. Provide a **Nick name** for the aggregate component and assign a **Remote Agent** to it using Figure 14.12.

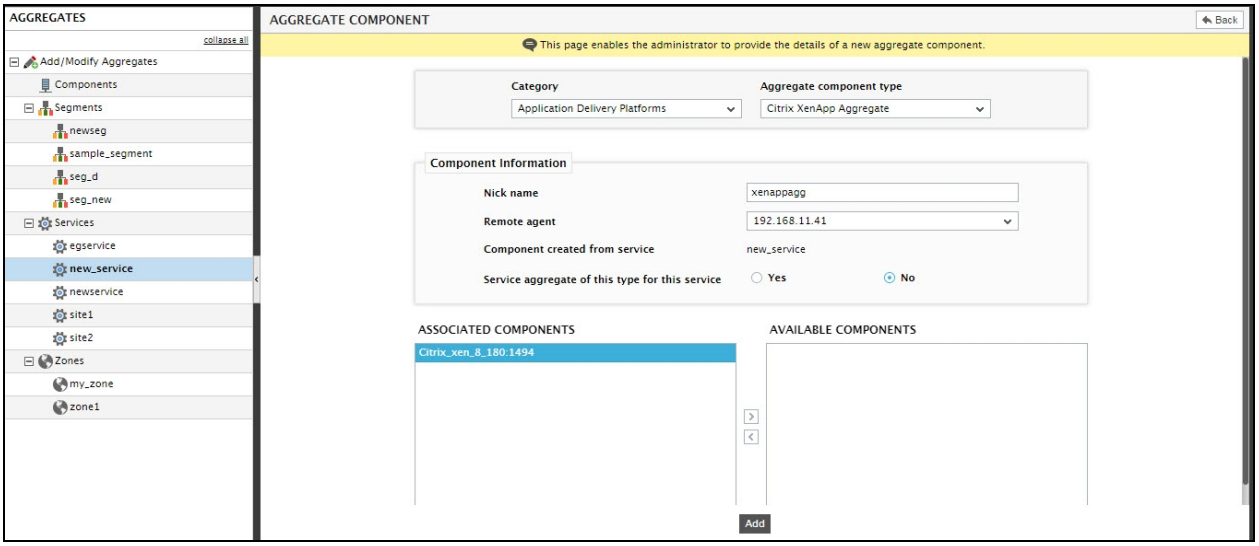

Figure 14.11: Adding an aggregate component using the components in a segment

4. Then, indicate whether/not the aggregate component being created is the **Service aggregate** for all components of that type in that service. In other words, indicate whether the aggregate component being created will represent the aggregated state of the components of that type in that service. This needs to be done because, where a service contains many components of a type, you can create any number of aggregate components of that type from the service. In this scenario, to accurately indicate service health, the eG Enterprise system must know which aggregate component of a type best represents the health of the service components of that type. This is why, ideally, any aggregate component that includes all components of a particular type within a service should be set as the **Service Aggregate of this type for this service**.

5. To set an aggregate component as the **Service aggregate**, set the **Service Aggregate of this type for this service** flag to **Yes**; otherwise, set it to **No**.

6. All the service components that can be included as members of the *aggregate component* of the chosen **Aggregate component type** will be automatically displayed in the **ASSOCIATED COMPONENTS** list - i.e., will be automatically associated with the **Aggregate component type**. For instance, when creating an *IIS Web Aggregate* component, all components in the chosen service that are of type *IIS Web* will be automatically associated with the *IIS Web Aggregate* component. You can disassociate one/more of these components by selecting them from the **ASSOCIATED COMPONENTS** list and clicking the **>** button in Figure 14.11. To add them back to the *aggregate component*, select them from the **DISASSOCIATED COMPONENTS** list and click the **Associate** button.
7. To add the newly created aggregate component, click the **Add** button in Figure 14.11.

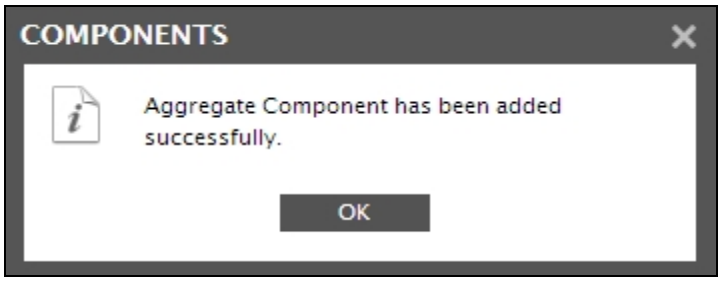

Figure 14.12: A result page indicating the successful addition of an aggregate component

8. When the message box of Figure 14.12 appears, click the **OK** button.
9. This will take you to Figure 14.13 where the newly created aggregate component will be listed.

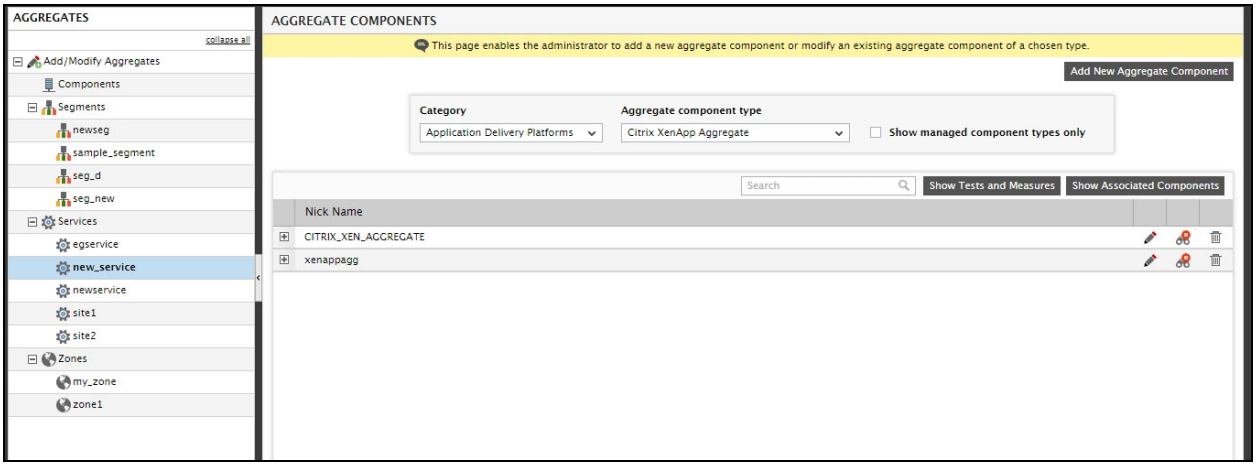

Figure 14.13: The aggregate component newly created from the service

10. Every aggregate component type so created is automatically associated with a default set of tests. To view the tests that are available by default for a particular aggregate component type and the measures they report, click on the **Show Tests and Measures** button at the top of Figure 14.13.
11. This way, aggregate components can be created from a segment, zone, and even a component group.

When creating an aggregate component from a zone, you can automatically associate the component being added to that zone. For this, while creating the aggregate component, you will have to select the **Auto-associate this component to the zone <zone_name>** flag to **Yes** (see Figure 14.14).

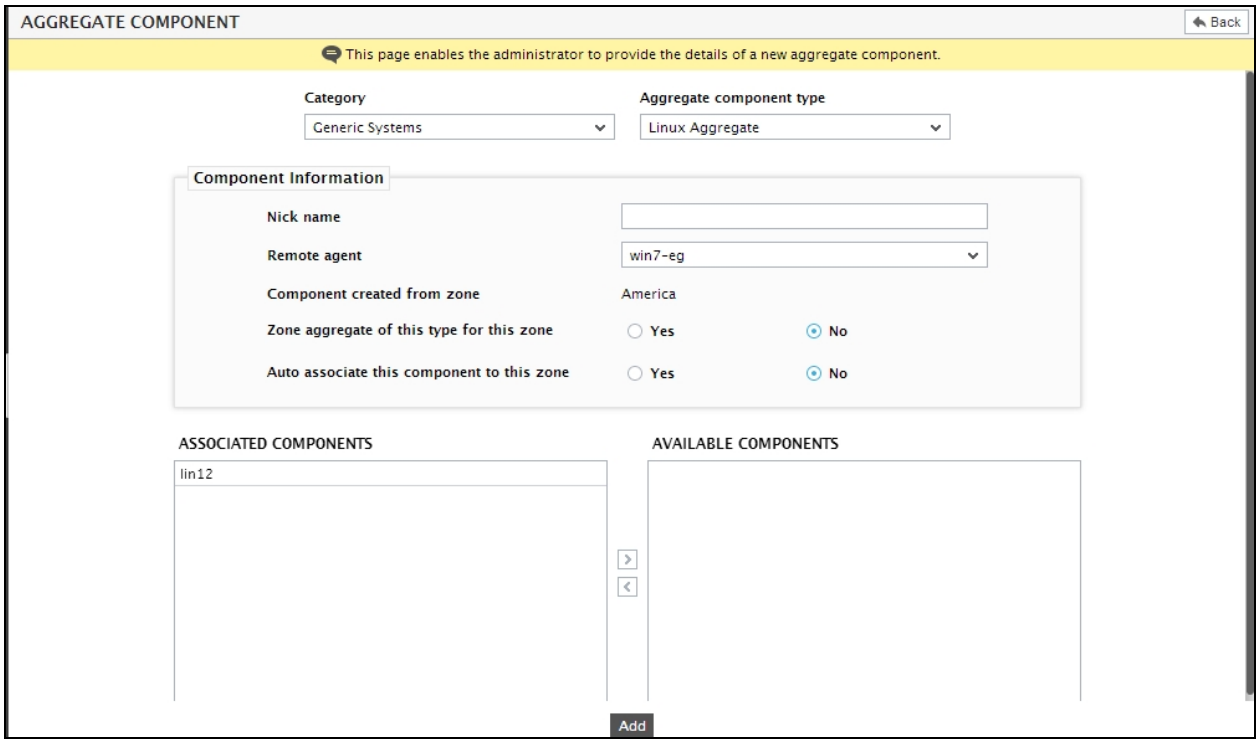

Figure 14.14: Creating an aggregate component from a zone

Also, when creating *aggregate components* from segments/services/zones/component groups, you can save yourselves the trouble of adding the *aggregate components* one at a time. Instead, you can simultaneously add multiple components - one of each *aggregate component type* - that can be supported by the segment/service/zone/component group that is chosen.

For instance, to create multiple aggregate components from the individual components of a segment at one shot, do the following:

1. Expand the **Segments** node in the tree structure and click on the segment of interest to you (see Figure 14.15).
2. Depending upon the types components that are part of the chosen segment, the eG Enterprise system automatically identifies the *aggregate component types* that can be created using those components and displays the same in the right panel (see Figure 14.15).

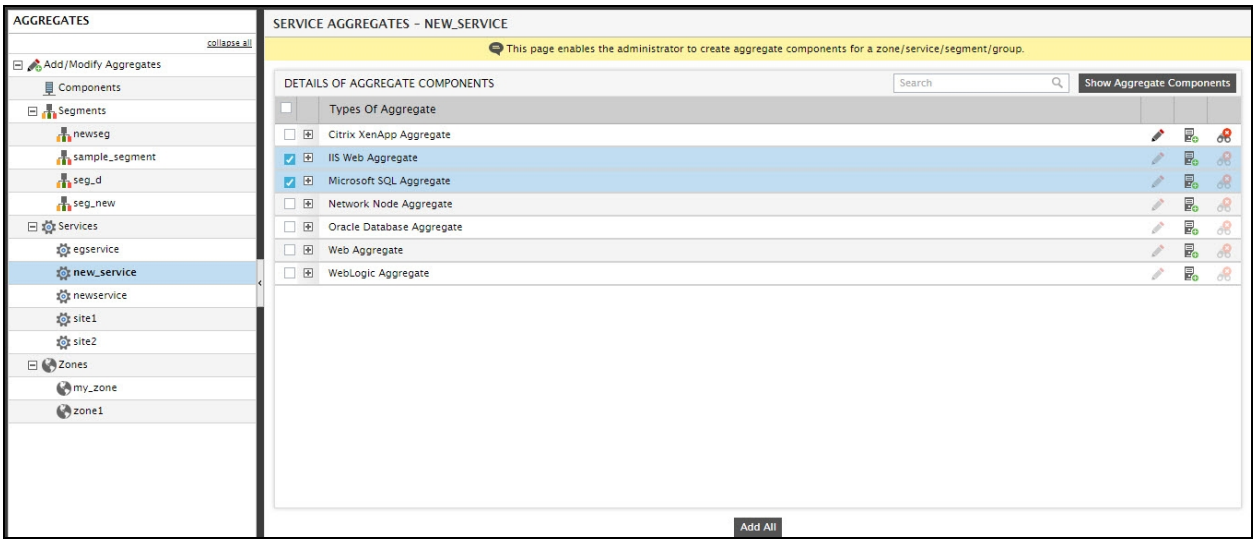

Figure 14.15: List of aggregate component types that a segment can support

3. From this component list, select the types of aggregate components you want to create by selecting the check boxes alongside the aggregate component types (see Figure 14.15).
4. Then, click the **Add All** button in Figure 14.15.
5. If this is done, then the eG Enterprise system will auto-assign a nick name and a remote agent to the aggregate components of each of the aggregate type chosen, and will then display this auto-configuration for your confirmation (see Figure 14.16). Click the **Add All** button in Figure 14.16 to automatically add the displayed components.

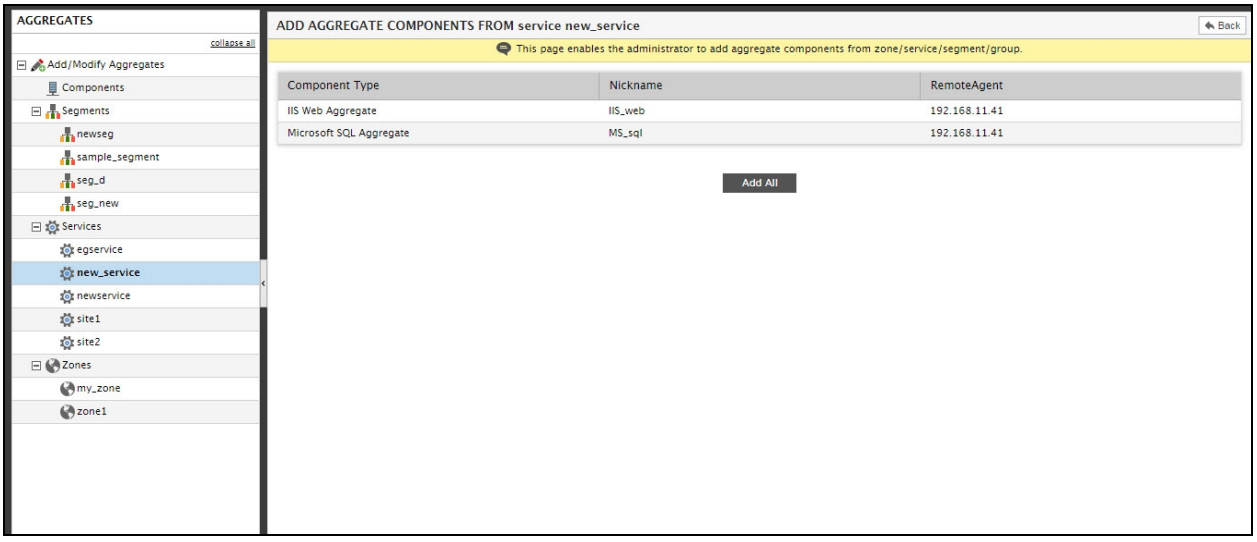

Figure 14.16: The aggregate components that can be created from the segment

14.2 Managing/Unmanaging Aggregate Components

All aggregate components that are manually added using the procedure discussed above are automatically managed by the eG Enterprise system. If you want to temporarily exclude any of these aggregate

components from your monitoring scope, or if you want to permanently delete an aggregate component from the eG Enterprise system, you can do so with the help of the **AGGREGATE COMPONENTS - MANAGE / UNMANAGE** page. To access this page, select the **Manage / Unmanage / Delete** option from the **Aggregates** menu of the **Infrastructure** tile. Doing so will invoke Figure 14.17.

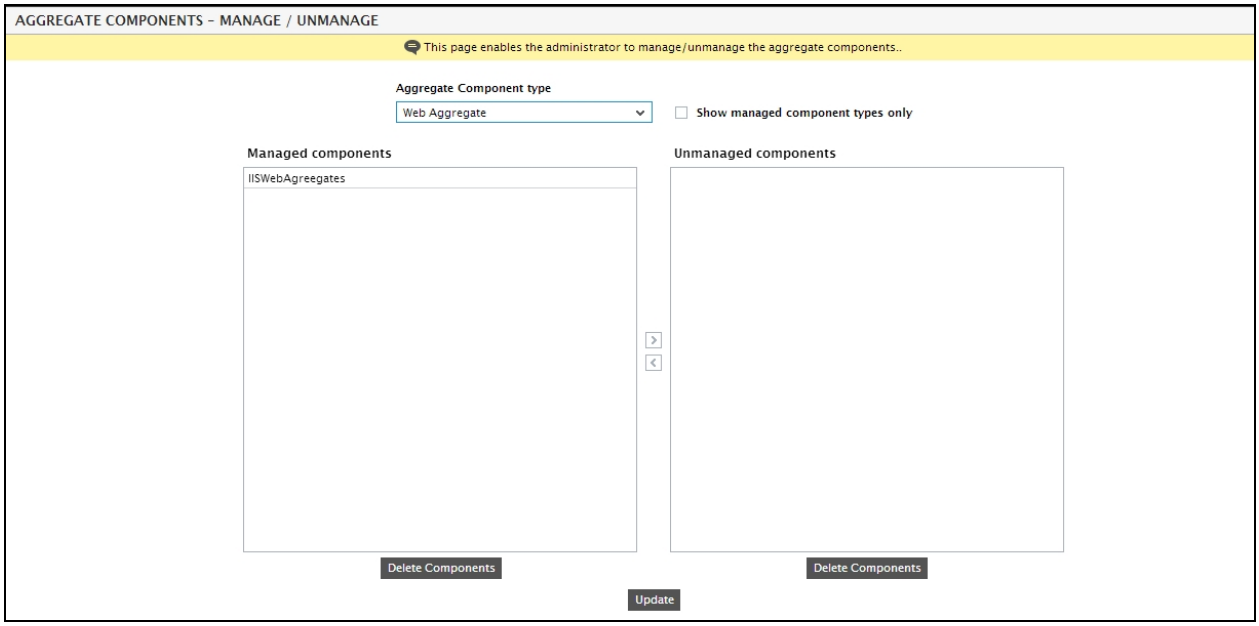

Figure 14.17: The AGGREGATE COMPONENTS - MANAGE/UNMANAGE page

To manage/unmanage aggregate components using Figure 14.17, do the following:

1. Select the **Aggregate Component type** for which component(s) need to be managed/unmanaged. By default, the **Aggregate Component type** list will contain all those *aggregate component types* with one/more managed and/or unmanaged aggregate components. In large environments, this list may be very long, making component-type selection painful! To ensure that this list is compact so that a component-type can be easily selected from it, click on the **Show managed aggregate component types only** check box. This check box, when selected, condenses the **Aggregate component type** list, so that it displays only those component-types with at least one managed *aggregate component*. This way, all types with only 'unmanaged aggregate components' will be automatically excluded from the list.
2. Upon selecting an **Aggregate Component type**, all aggregate components of that type that are currently monitored by the eG Enterprise system will be listed in the **Managed Components** list box. Likewise, those aggregate components of the chosen type that are currently excluded from monitoring will be listed in the **Unmanaged components** list box.
3. To unmanage one/more aggregate components, select them from the **Managed Components** list (as depicted by Figure 14.18) and click the **>** button. An attempt to unmanage an aggregate component will be followed by a message box that explains the implications of the same (see Figure 14.19).

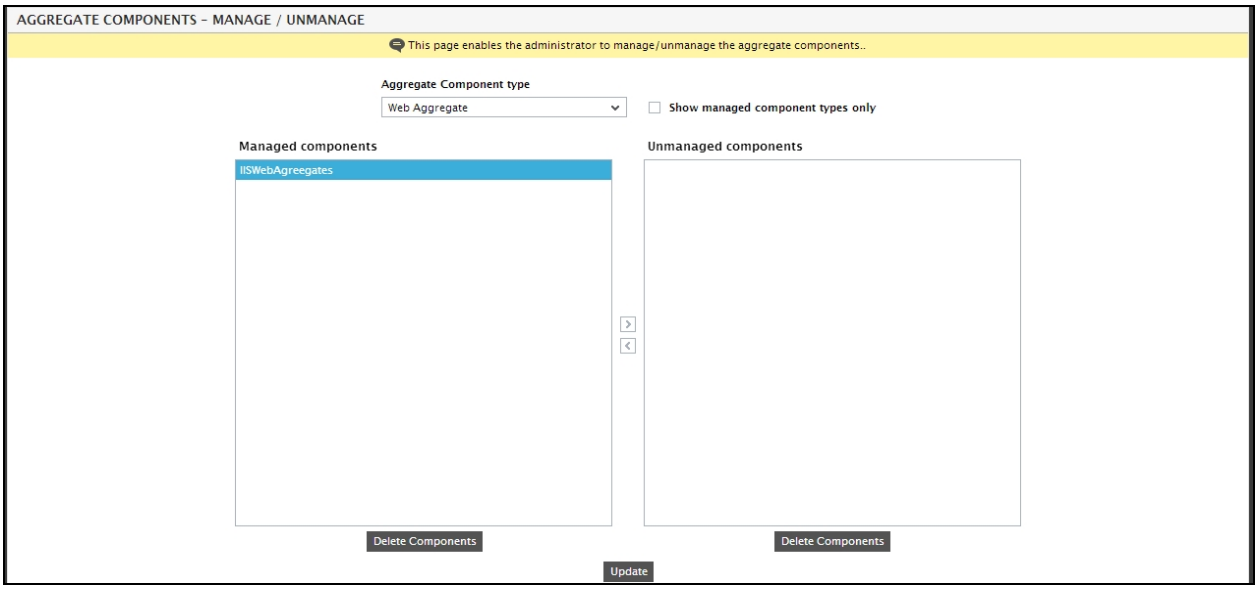

Figure 14.18: Selecting the managed aggregate components to be unmanaged

Figure 14.19: A message box informing the administrator that unmanaging a component will result in the loss of all the configuration information related to that component

4. Click the **OK** button in Figure 14.19 to proceed with unmanaging the aggregate component.
5. This will transfer the components chosen for unmanaging to the **Unmanaged Components** list (see Figure 14.20).

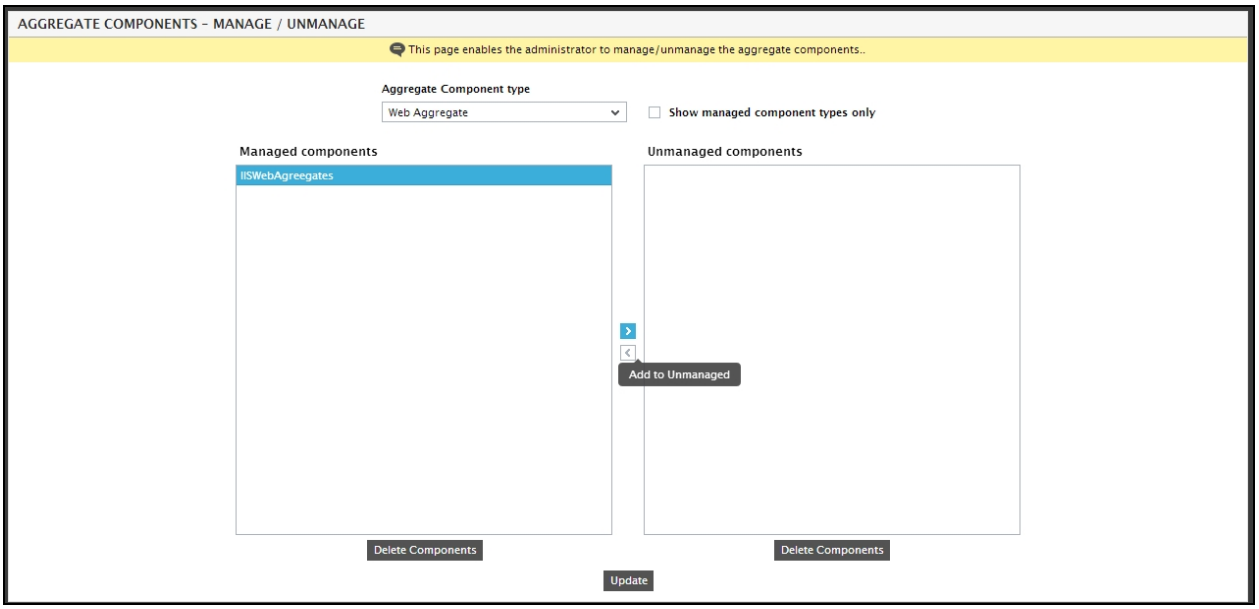

Figure 14.20: Unmanaging the chosen aggregate components

6. Finally, click the **Update** button.
7. Similarly, to manage one/more aggregate components that were previously unmanaged, pick them from the **Unmanaged Components** list, click the **<** button, and then click the **Update** button.
8. If for some reason, you want to delete a managed/unmanaged aggregate component, then, select the component from the **Managed Components** or **Unmanaged Components** list box (depending upon where the component is listed) and click the **Delete Components** button corresponding to that list box. Then, click the **Update** button. If a managed component is being deleted, then the message box depicted by Figure 14.21 will appear informing you of the consequences of deletion. To go ahead with the deletion, click the **OK** button in the message box. To quit deleting, click the **Cancel** button.

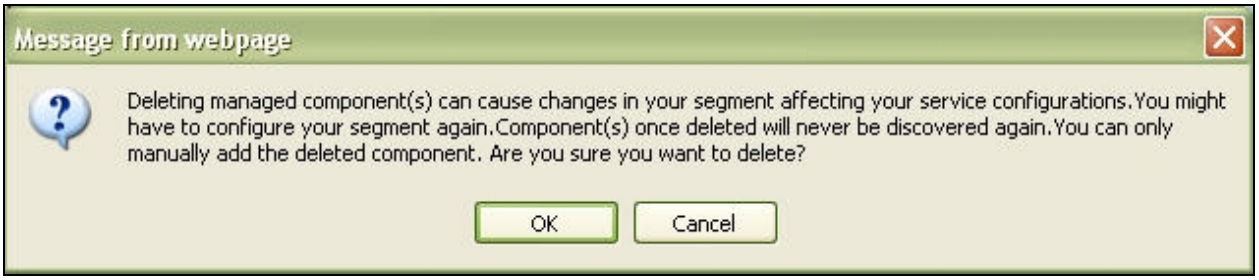

Figure 14.21: The message box requesting your confirmation to delete a managed component

14.3 Configuring an Aggregated Web Site

Typically, if more than one web server is involved in delivering a web site service, the eG Enterprise system will report site-related and transaction-related metrics separately for each web server mapped to that site - this way, the eG Enterprise system accurately reveals how each of the web servers mapped to a site handles the requests to it, and thus points you to the 'unhealthy' web servers (if any).

Sometimes however, service managers may require a consolidated view of performance of a web site across all the web servers that it overlays. This view enables administrators to determine the total request load on the site and the total number (and nature) of errors encountered by the site, regardless of which web server is processing the requests. With the help of this view, administrators can focus on the performance of the web site alone without being distracted by the web servers, periodically check web site usage and fine-tune the site to improve usage, accurately measure the user experience with the web site, and promptly detect its spoilers!

To provide this consolidated web site view, the eG administrative interface allows the creation and maintenance of *aggregated web sites*. An aggregated web site is typically delivered by an *aggregate web server*, which has **Site Support** enabled - in other words, when adding an *aggregate web server* component, you need to set the **Site Support** flag to **Yes**, if you intend to create an *aggregated web site* later using that component. Also, before attempting to create an *aggregated web* site, you need to make sure that at least one 'non-aggregated web site' pre-exists in the eG Enterprise system; this is because, web site aggregation aggregates the measures captured from live transactions to a 'stand-alone' (i.e., non-aggregate) web site only. In short, an *aggregate web site* cannot be created if the following conditions are not fulfilled:

- If you do not have any *aggregate web server* components with **Site Support** enabled;
- If not even a single 'non-aggregate' web site service pre-exists in the environment

Once the aforesaid pre-requisites are fulfilled, proceed to create an *aggregate web site* as discussed below:

1. Follow the Service -> Topologies menu sequence in the **Infrastructure** tile.
2. The **SERVICE** page then appears, listing all the services that pre-exist. To add a new service, click on the **Add New Service** button.
3. Figure 14.22 then appears. Provide the **Name of the service** and set the **Is this service a web site** flag to **Yes**. Select a service type and also choose a display image for the service. To add an aggregate web site, set the **Is this an aggregate website** flag to **Yes** and click the **Add** button.

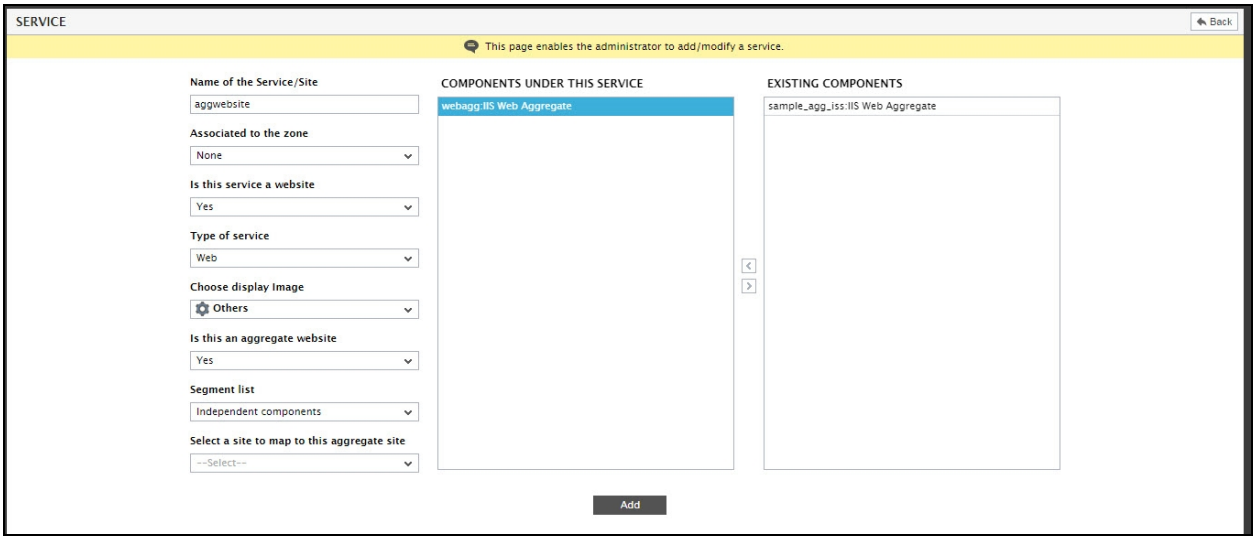

Figure 14.22: Figure 13. 22: Adding an aggregate web site

Note:

The **Is this an aggregate web site?** flag will appear only if the following conditions are fulfilled:

- You must have configured at least a single *aggregate web server* component with **Site Support** enabled;
- At least one 'non-aggregate' web site service should pre-exist in the environment

4. Using Figure 14.22, you can associate the aggregate web site being created with an *aggregate web server* component. For this, first indicate whether the *aggregate web server* component of interest to you is part of a segment or is an independent component. In case of the latter, select the **Independent Components** option from the **Segment list**. All independent *aggregate web server* components will then appear in the **EXISTING COMPONENTS** list. In case of the former, select the corresponding segment from the **Segment list**. All the *aggregate web server* components that are part of the chosen segment will then populate the **EXISTING COMPONENTS** list. From this list, select the *aggregate web server* component(s) you want to associate with the web site. Then, click the **<** button. This will transfer the selection to the **COMPONENTS UNDER NEW SERVICE** list. Similarly, you can disassociate components from the site by selecting them from the **COMPONENTS UNDER NEW SERVICE** list and clicking the **>** button.
5. Next, from the **Select a site to map to this aggregate site** list, select a 'non-aggregate' web site that is to be mapped to the aggregate site being configured.

   Note:

   For best results, it is recommended that the aggregate web server component you attach to an aggregate web site includes at least two of the web servers that have been associated with the 'non-aggregate' site mapped to that 'aggregate web site'.
6. Finally, click the **Update** button.
7. Once this is done, then all the web transactions that have been configured for the 'non-aggregate' site will be automatically available for the aggregate site as well. **Additional transactions cannot be configured for the aggregate site**. For the aggregate web site therefore, the eG Enterprise system will aggregate the metrics collected per transaction across all the web servers that fulfill both the conditions discussed below:
   - The web servers that are associated with the 'non-aggregate' web site mapped to the 'aggregate web site'
   - The web servers that are included as part of the *aggregate web server* component that is associated with the aggregate web site

14.4 Adding/Modifying/Deleting Aggregate Tests

By default, when an aggregate component is created, all tests mapped to the non-aggregate equivalent of that component will be automatically mapped to the aggregate component as aggregate tests. For instance, if a Windows Aggregate component is added, then all tests originally mapped to a Windows component will be auto-assigned to the Windows Aggregate component as aggregate tests – eg., Disk Activity Aggregate test, System Details Aggregate test etc. The aggregate functions applied on the measures reported by these default aggregate tests are also pre-defined.

If required, you can create a new aggregate test by adding a subset of the measures of the related non-aggregate test to it. You can change the display names of these measures and can even apply a mathematical function of your choice on the measures. For example, you can create, a Disk Activity 1

Aggregate test, which will report only the Disk busy measure of the Disk Activity test. By default, for an aggregate component, the aggregate disk busy value is computed by applying the Average function on the Disk busy measure. You can however, choose to apply the Maximum function on the measure instead.

To create a new aggregate test, do the following:

1. Select the **Add/Modify Tests** option from the **Aggregates** menu of the **Infrastructure** tile. Figure 14.23 then appears.

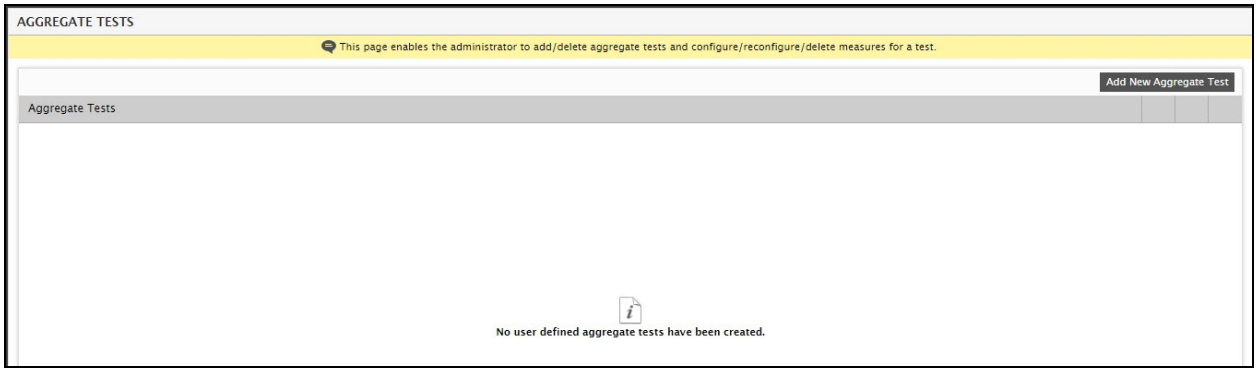

Figure 14.23: A message indicating that no user-defined aggregate tests pre-exist

2. If no user-defined aggregate tests pre-exist, then Figure 14.23 will display a message to that effect. To add a new test, click the **Add New Aggregate Test** button in Figure 14.23. This will open Figure 14.24.

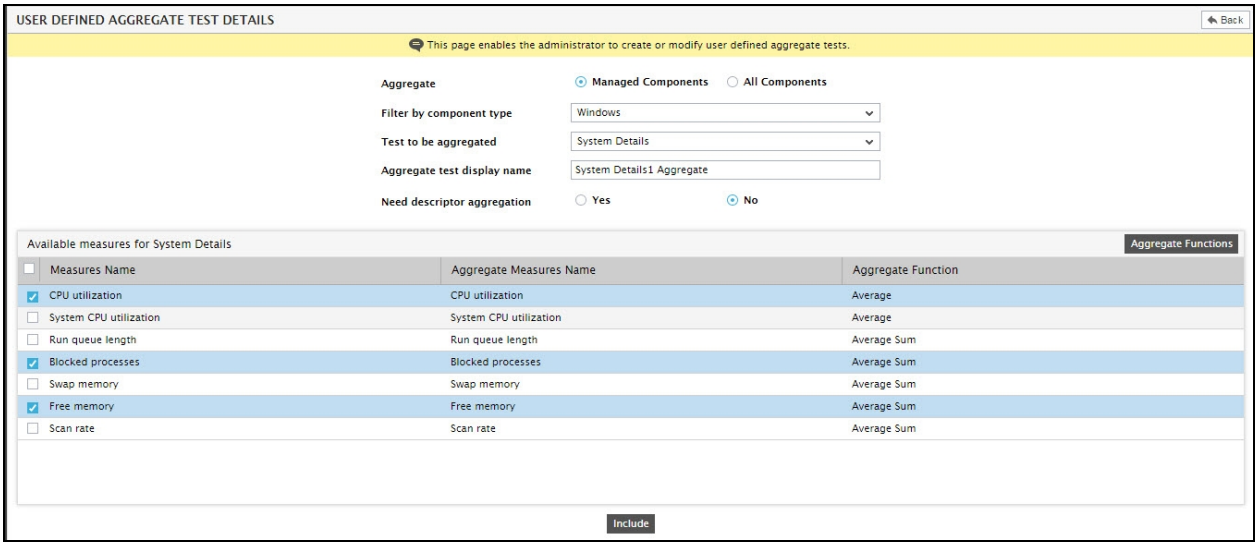

Figure 14.24: Adding a new aggregate test

3. If you want to create a new aggregate test for a managed component type, set the **Aggregate** flag in Figure 14.24 to **Managed Components**. In this case, the **Filter by component type** list will display only those types of which at least one component is currently managed. On the other hand, if you want to create a new aggregate test for any component type (managed or unmanaged), then set the **Aggregate** flag to **All Components**. In this case, the **Filter by component type** list will display all component types that are supported out-of-the-box by the eG Enterprise Suite.

4. Pick a component type from the **Filter by component type** list.
5. Select the **Test to be aggregated**.
6. Provide a display name for the new aggregate test in the **Aggregate test display name** text box.
7. **The 'Need descriptor aggregation' flag gains relevance only if the new aggregate test being added is a descriptor-based test**. If you want the new aggregate test to report measure values that are aggregated across descriptors, then set the **Need descriptor aggregation** flag to **Yes**. If you want the new aggregate test to report measures that are aggregated per descriptor, then set this flag to **No**.
8. The **Available measures** section will then display all measures that are reported by the **Test to be aggregated**. Against every measure name, the **Aggregate Measure Name** set by default for that measure and the **Aggregate Function** that is applied by default when aggregating that measure will be displayed.
9. You can override the default **Aggregate Measure Name** and **Aggregate Function** of any measure listed in the **Available measures** section. To change an **Aggregate Measure Name**, just click it. The name will now become editable (see Figure 14.25). Make the changes you deem fit.

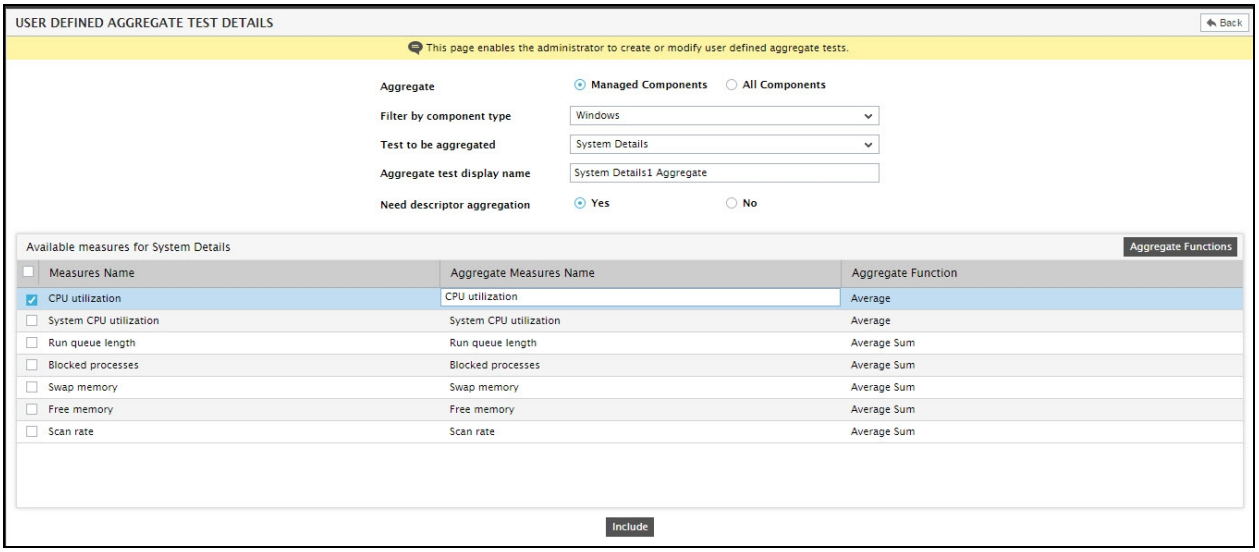

Figure 14.25: Changing the display name of an aggregate measure

Likewise, to change the **Aggregate Function**, just click on the aggregate function you want to change. An **Aggregate Function** drop-down list will now be available to you. Select any aggregate function you prefer from this list (see Figure 14.26). If your selection is a mathematical function such as Average, Minimum, Maximum, Sum, or Average Sum, then proceed as described from step 10 to step 20 below. On the other hand, if your selection is a non-mathematical function such as Condition or Multi-condition, then refer to the section on Conditional Aggregation and learn how and when these two options are to be used.

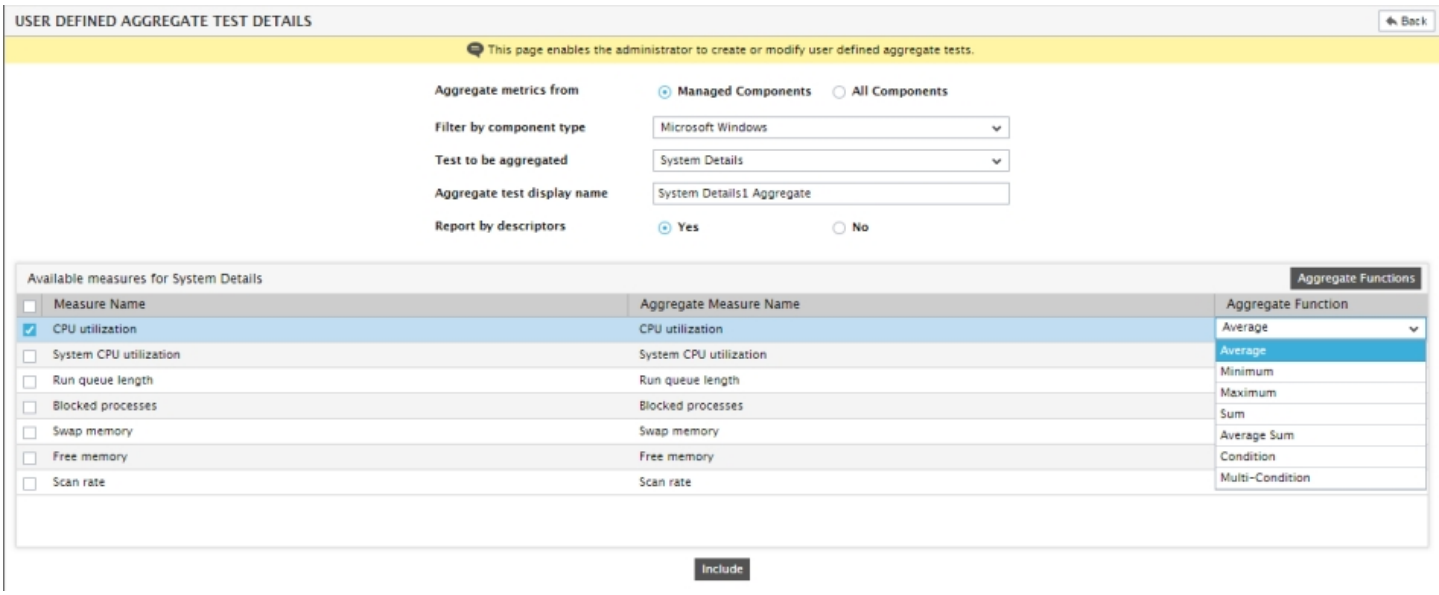

Figure 14.26: Changing the aggregate function to be applied on a measure

10. You can, if you so need, add only a few of the measures listed in the **Available measures** section to the aggregate test being created. For instance, let us try to create a System Details 1 Aggregate test that reports only the CPU utilization, Free memory, and Blocked processes measures. For this, first select the check box corresponding to each of these measures, as depicted by Figure 14.24. Then, click the **Include** button therein.
11. Figure 14.27 will then appear providing a summary of your selection.

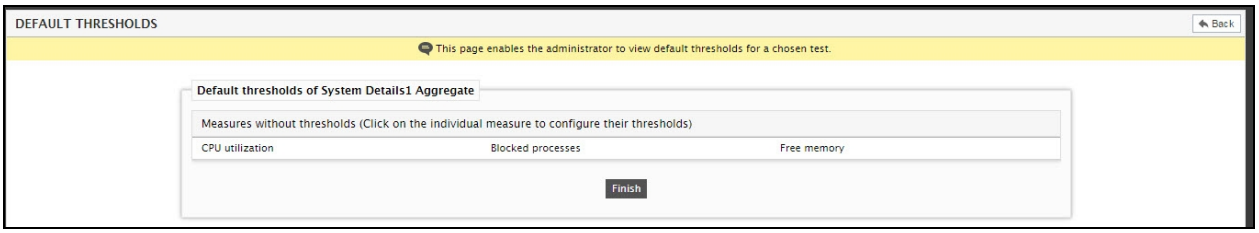

Figure 14.27: Reviewing the measures that will be reported by the new aggregate test that you created

12. By default, no thresholds will govern the measures added to a custom-defined aggregate test. To configure thresholds for these measures, click on a measure name in Figure 14.27. This will open Figure 14.28, using which you can configure maximum and/or minimum thresholds for that aggregate measure.

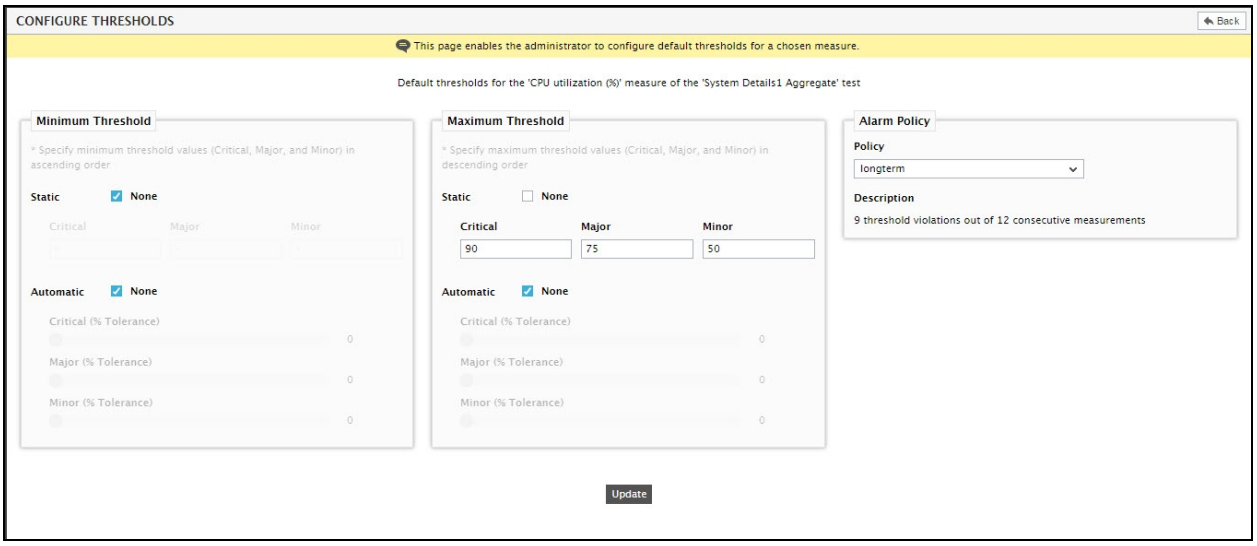

Figure 14.28: Configuring thresholds for an aggregate measure

13. After threshold configuration, click the **Update** button in Figure 14.28 to save the threshold settings. You will now return to Figure 14.29, where you can see which aggregate measures have been configured with thresholds and which ones are awaiting threshold configuration.

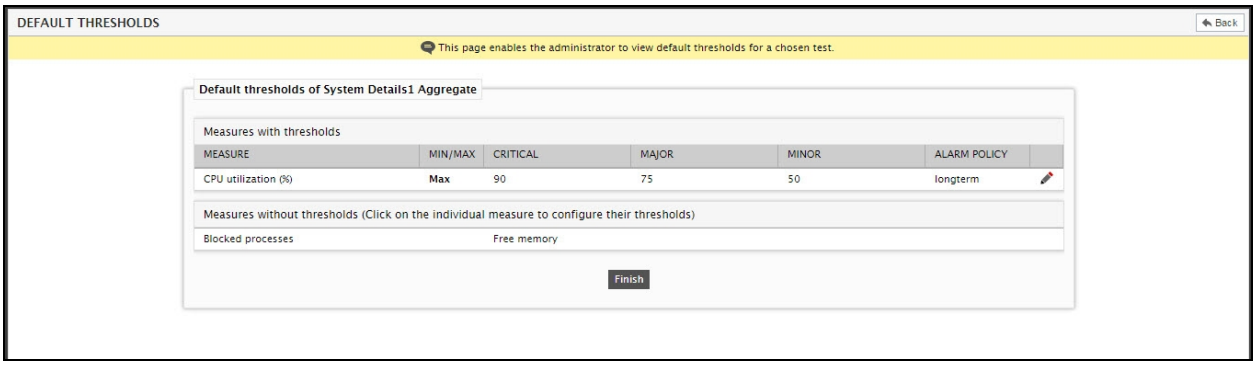

Figure 14.29: A page displaying the aggregate measures for which thresholds have been configured and the ones without thresholds

14. You can click on an aggregate measure in the **Measures without thresholds** section of Figure 14.29 to configure its thresholds. You can even modify the threshold settings of a **Measure with thresholds,** by clicking the icon corresponding to that measure.

15. If you have finished configuring thresholds for all aggregate measures you want, click the **Finish** button in Figure 14.29 to save the changes.

16. By default, as soon as a user-defined aggregate test is added, it will be automatically mapped to all managed aggregate component types to which its base test – i.e., the test chosen from the **Test to be aggregated** list of Figure 13.24 – is mapped. For instance, in the example illustrated by Figure 14.24, the System Details 1 Aggregate test was created from the System Details test. Accordingly, the System Details 1 Aggregate test will be automatically associated with all the 'managed' aggregate component types to which the System Details Aggregate test is mapped. However, if no aggregate component of

such a type is currently managed, then the message depicted by Figure 14.30 will appear, as soon as you click on the **Finish** button in Figure 14.29.

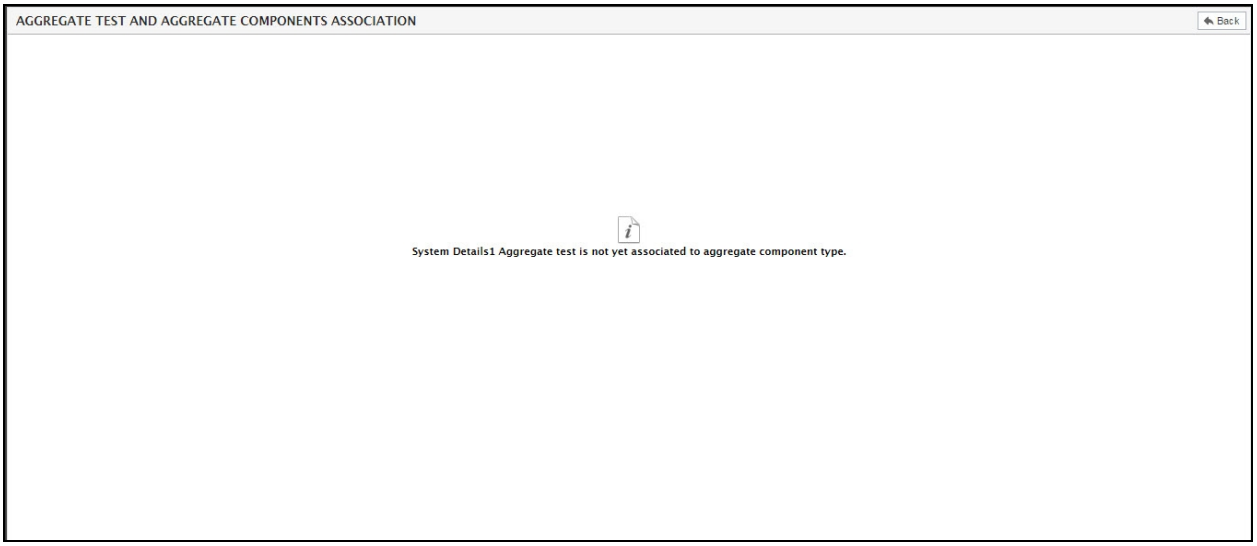

Figure 14.30: A message indicating that the new aggregate test is not mapped to any aggregate component type

17. On the other hand, if you have already managed one/more components of the aggregate component types with which the System Details test is mapped, then, clicking the **Finish** button in Figure 14.29 will reveal the list of component types to which the System Details 1 Aggregate test is automatically mapped.

Figure 14.31: The aggregate component types to which a custom-defined aggregate test is automatically mapped

18. Clicking the **Back** button in Figure 14.31 will lead you to Figure 14.32, where you can see the aggregate test that you created newly.

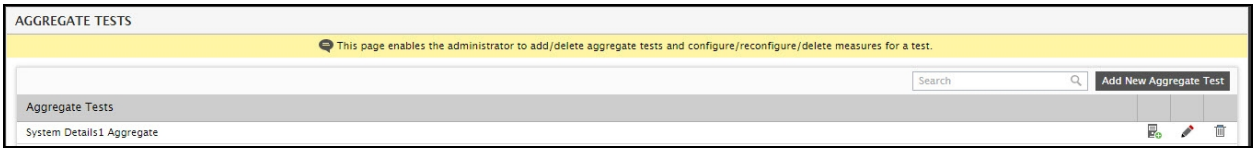

Figure 14.32: The newly added aggregate test listed therein

19. You can view the aggregate component types with which a user-defined aggregate test is currently associated by clicking the icon corresponding to that test in Figure 14.32. You can click the icon against a custom-defined aggregate test to delete it (see Figure 14.32). To modify a test configuration, click the button corresponding to it. This will open Figure 14.33.

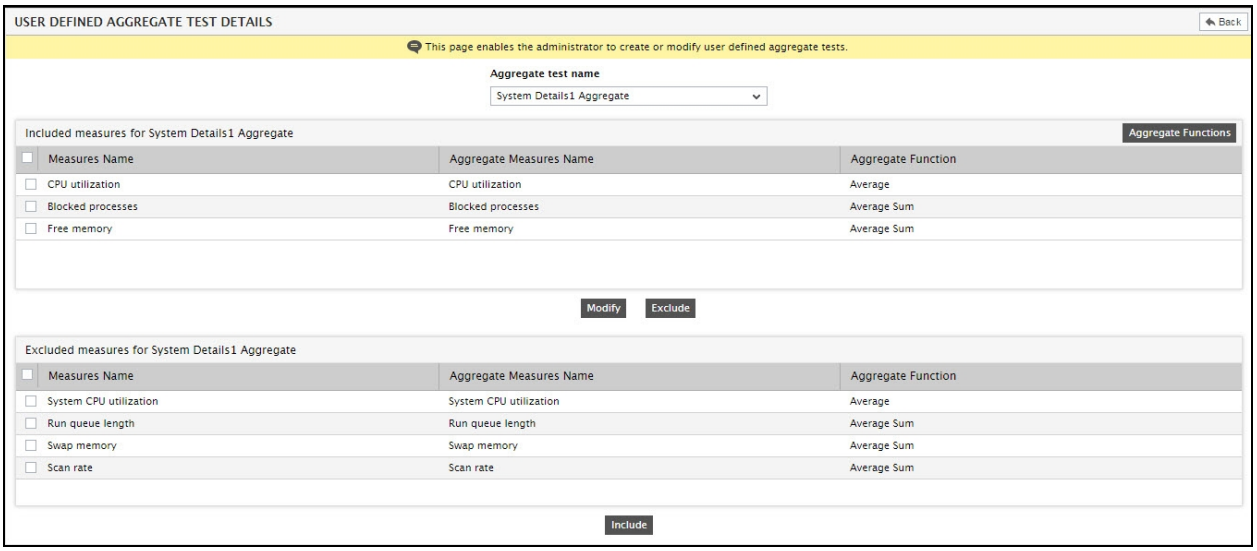

Figure 14.33: Modifying the configuration of a user-defined aggregate test

20. Using Figure 14.33, you can **Include** more measures to the new aggregate test. For this, select the check boxes preceding the measures in the **Excluded measures** section of Figure 14.33, and click the **Include** button below that section. You can exclude any of the measures that were already added to the test. For this, select the check boxes preceding the measures in the **Included measures** section of Figure 14.33, and click the **Exclude** button below that section. You can change the **Aggregate Measure Name** and **Aggregate Function** of any measure in the **Included measures** list. Finally, you can click the **Update** button to save the changes.

14.5 Conditional Aggregation

Aggregate metrics in eG Enterprise help administrators get a farm-wide view (rather than a server by server view of the target infrastructure). In the Section **14.4** topic, we described how mathematical functions like Avg, Sum, Min, Max, etc., can be used for computing the aggregate measure values for new aggregate tests. However, these aggregation functions can hide problem conditions. For instance, say you create an aggregate test to report the average CPU usage of a Windows server farm comprising of 4 Windows servers. If 3 out of the 4 servers register a CPU usage of 40% each, and one server registers 80%, then the average CPU usage for that farm will be 50%. This seemingly low aggregated CPU usage value does not reveal the fact that a single Windows server is seeing more than 80% of CPU resources.

To allow administrators greater flexibility and visibility into the health of the target infrastructure, eG Enterprise uses Conditional Aggregation. This is most useful when administrators only want to know the count or percentage of components that fulfill a defined condition or conditions. For instance, administrators may just want to know how many Windows servers are consuming over 80% of the CPU resources. Conditions are also useful when administrators want aggregate measures to report status values and not aggregated measure values; this will help the administrators determine how the fulfilment of a condition has impacted the health of the aggregate component. For example, an administrator may want the CPU usage measure to report the value Critical, if over 50% of the components in the aggregate are consuming over 80% of their individual CPU resources. The conditional aggregation capability helps in this case.

This conditional aggregation can be performed using a single condition or multiple conditions. Let us see how to perform both using separate illustrated examples.

- Performing Aggregation Using a Single Condition
- Performing Aggregation Using Multiple Conditions

14.5.1 Performing Aggregation Using a Single Condition

The procedure detailed below takes the example a Citrix XenApp farm. Let us see how to define an aggregate test that will report the percentage of XenApp servers in the farm that have CPU utilization over 80%.

1. Select the **Add/Modify Tests** option from the **Aggregates** menu of the **Infrastructure** tile. Figure 14.34 then appears.

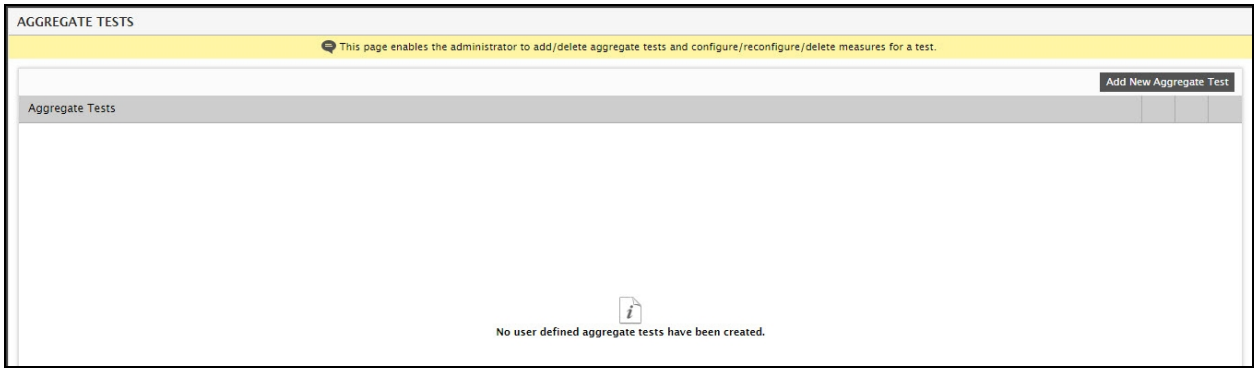

Figure 14.34: Configuring an aggregate test using a single condition

2. If no user-defined aggregate tests pre-exist, then Figure 14.34 will display a message to that effect. To add a new test, click the **Add New Aggregate Test** button in Figure 14.34. This will open Figure 14.35.

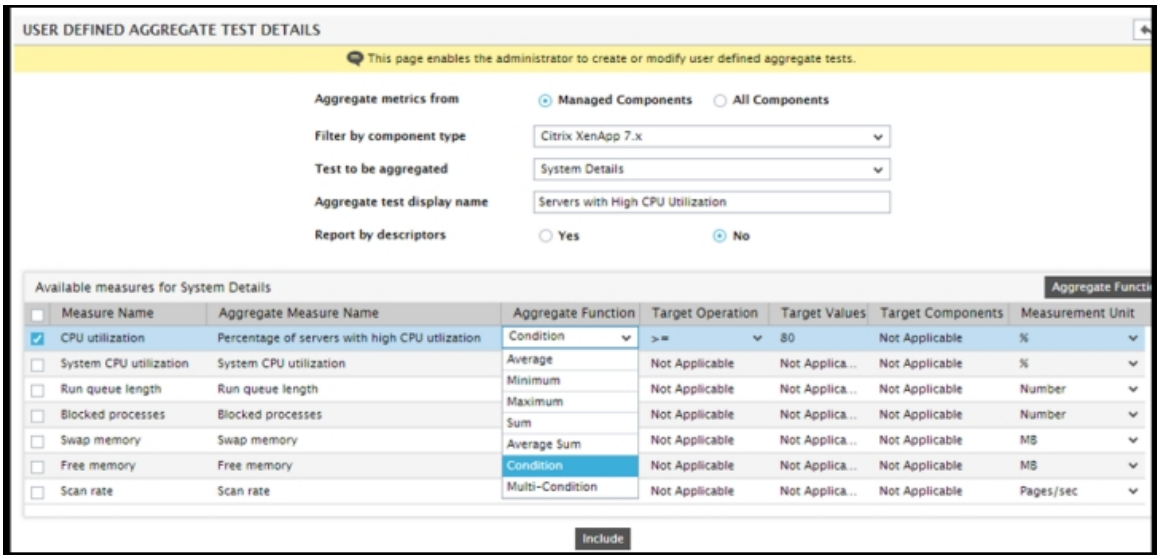

Figure 14.35: Configuring an aggregate test using a single condition

3. Set the **Aggregate** flag in Figure 2 to **Managed Components**.
4. Pick *Citrix XenApp 7.x* as the component type from the **Filter by** component type list.
5. Since the CPU usage measure of interest to us is reported by the System Details test, select *System Details* as the **Test to be aggregated**.
6. Provide a meaningful display name for the new aggregate test in the **Aggregate test display name** text box.
7. As aggregation is not descriptor-based in the case of our example, set the **Report by descriptors** flag to **No**.
8. The **Available measures** section will then display all measures that are reported by the **Test to be aggregated**. Against every measure name, the **Aggregate Measure Name** set by default for that measure and the **Aggregate Function** that is applied by default when aggregating that measure will be displayed.
9. Let us now focus on the *CPU utilization* measure. To aggregate the value of this measure based on a single condition, select the *Condition* option as the **Aggregate Function** for that measure.
10. Once Condition is chosen as the **Aggregate Function**, five new columns get added to Figure 2. They are, namely – Target Operation, Target Values, Target Violation by Descriptors, Target Components, and Measurement Unit.
11. From the **Target Operation** drop-down, select the logical function to define the Conditional Aggregation criteria. Our example seeks to report the percentage of XenApp servers in a farm that are utilizing 80% of the CPU resources or more. To achieve this, select the >= option from the **Target Operation** drop-down. Then, in the **Target Values** text area, specify the value to be checked against the operation chosen from **Target Operation**. Since the condition should check for a value >= 80%, specify 80 in the **Target Values** text area.
12. Since the **Report by Descriptors** flag is set to **No**, you will find an additional **Target Violation**

by Descriptors column. This column offers the following options:

Option	Purpose
Majority	If you select this option for the *CPU utilization* measure in our example, then the measure will report the percentage of only those server in the XenApp farm with a majority of their processors used 80% of the time or more.
All	If you select this option for the *CPU utilization* measure in our example, then the measure will report the percentage of only those servers in the XenApp farm with all their processors used 80% of the time or more.
Any	If you select this option for the *CPU utilization* measure in our example, then the measure will report the percentage of only those servers in the XenApp farm with at least one processor that is used 80% of the time or more.

Since the **Majority** option is what suits our purpose, select that for the *CPU utilization* measure in our example.

13. The **Target Components** specification is not applicable to an aggregate test that is based on a single condition. So, let us move to the **Measurement Unit**.
14. From the **Measurement Unit** drop-down, pick the unit in which the value of this aggregate measure is to be reported. For a single condition, the measurement unit can be Number or % only. As we want the aggregate measure in our example to report the percentage of servers that fulfill a specified condition, select% as the Measurement Unit.
15. Then, select the check box corresponding to the *CPU utilization* measure and click the **Include** button to save the changes.
16. Next, switch to the Monitor interface to check what measure this aggregate test reports. Figure 14.36 depicts the value of the new aggregate measure. As is evident from Figure 3, 25% of the servers in the XenApp farm consume over 80% of CPU. To know which are these XenApp servers, click the icon corresponding to the measure name in Figure 14.36.

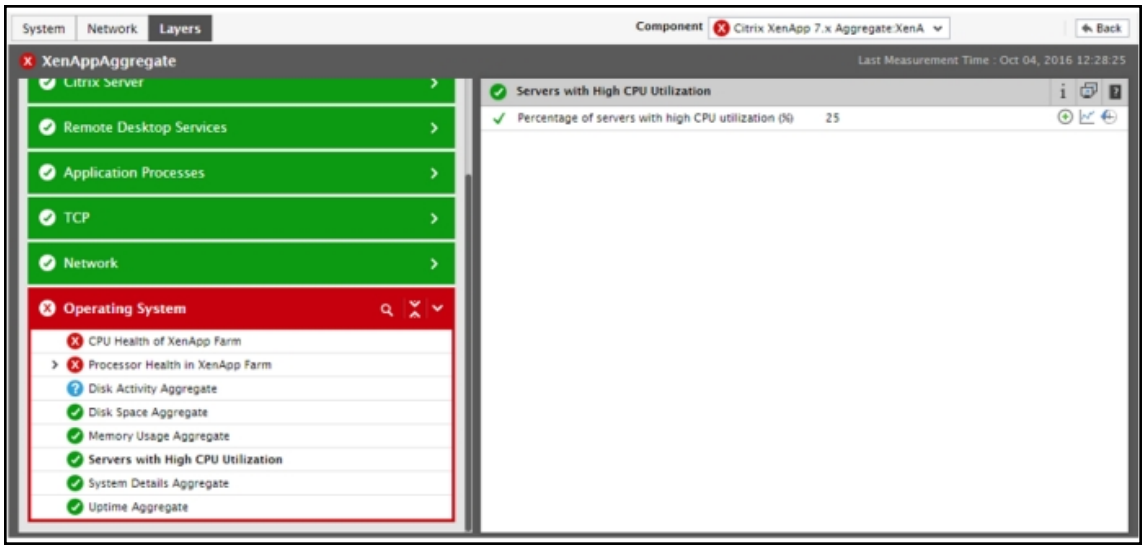

Figure 14.36: Viewing the value of the new aggregate measure

17. Figure 14.37 will then appear.

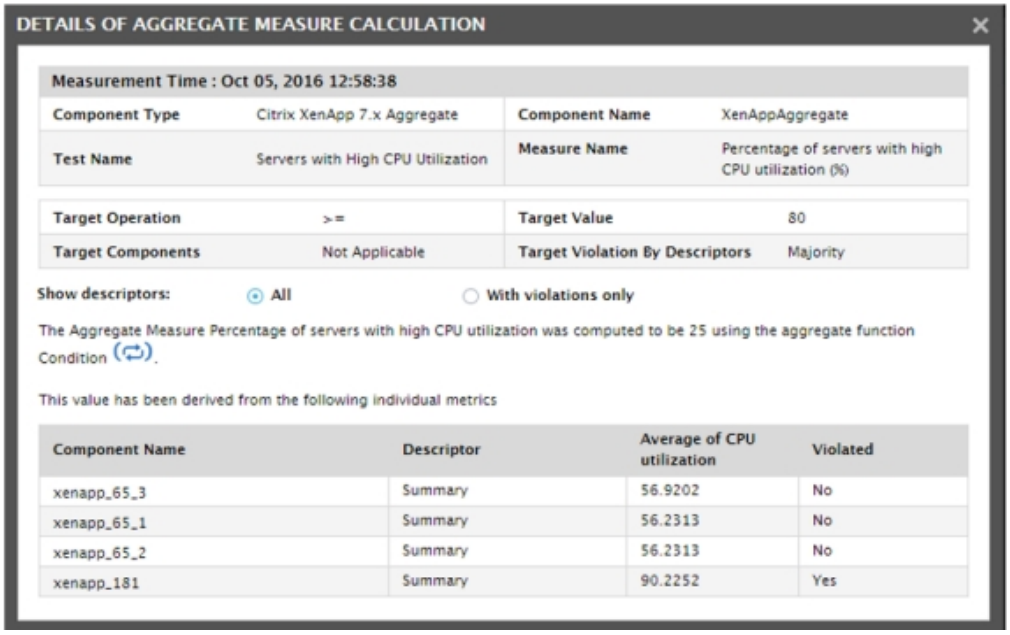

Figure 14.37: How the aggregate measure value was computed

18. Figure 14.37 describes how the value of the measure was computed. In the table at the bottom of Figure 14.37, you can see that the components in the XenApp farm have been listed. Against each component, the descriptors - i.e., the Processors - supporting the corresponding server will be listed. The percent CPU utilization of each processor will also be displayed alongside, with an indication as to whether/not that utilization value has violated the configured **Target Value** condition. In the case of our example, the **System Details** test of each of the XenApp servers in the farm reports metrics for the Summary

descriptor only, and not for the individual processors. Therefore, the aggregate test also looks for violations in the Summary descriptor alone.

19. As per our specification, a server in the farm is considered to have violated the target value of 80%, only if a Majority of its descriptors violates the configured **Target Value**. Since Summary is the only descriptor in our example, any server in the farm that reports a CPU usage summary of over 80% will be in violation of the **Target Value**. If you now look at the values in the Average of CPU utilization column in Figure 14.37, it will be clear to you that the CPU usage summary exceeds 80% in only 1 out of the 4 XenApp servers. In percentage, this count amounts to 25%. This is why, the *Servers with High CPU utilization* measure reports the value 25%.

14.5.2 Performing Aggregation Using Multiple Conditions

The procedure detailed below takes the example of a Citrix XenApp farm. Let us see how to define an aggregate test that will report the following values:

- Raise a Critical alert when 80% of Citrix XenApp servers in the farm consume over 20%;
- Raise a Major alert when 60% of Citrix XenApp servers in the farm consume over 40% of CPU;
- Raise a Minor alert when 30% of Citrix XenApp servers in the farm consume over 60% of CPU;

The steps to achieve this are as follows:

1. Select the **Add/Modify Tests** option from the **Aggregates** menu of the **Infrastructure** tile. Figure 14.38 then appears.

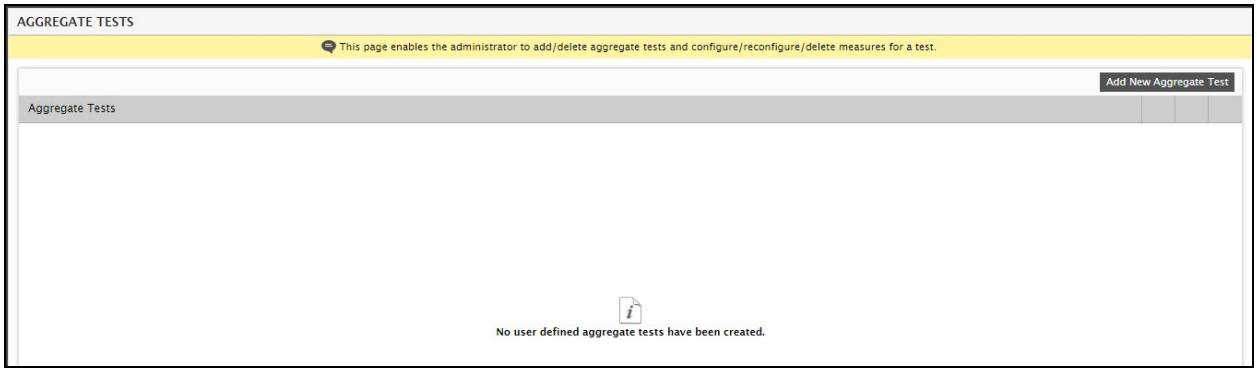

Figure 14.38: Configuring an aggregate test using a single condition

2. If no user-defined aggregate tests pre-exist, then Figure 14.38 will display a message to that effect. To add a new test, click the **Add New Aggregate Test** button in Figure 14.38. This will open Figure 14.39.

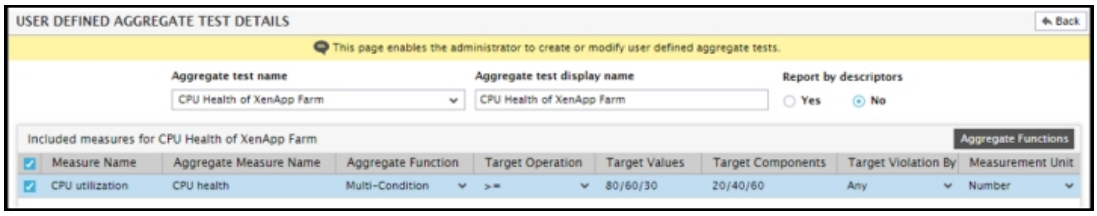

Figure 14.39: Configuring an aggregate test using a single condition

3. Set the **Aggregate** flag in Figure 14.39 to **Managed Components**.
4. Pick *Citrix XenApp 7.x* as the component type from the **Filter by** component type list.
5. Since the CPU usage measure of interest to us is reported by the System Details test, select *System Details* as the **Test to be aggregated**.
6. Provide a meaningful display name for the new aggregate test in the **Aggregate test display name** text box.
7. As aggregation is not descriptor-based in the case of our example, set the **Report by descriptors**flag to **No**.
8. The **Available measures** section will then display all measures that are reported by the **Test to be aggregated**. Against every measure name, the **Aggregate Measure Name** set by default for that measure and the **Aggregate Function** that is applied by default when aggregating that measure will be displayed.
9. Let us now focus on the *CPU utilization* measure. To aggregate the value of this measure based on a single condition, select the *Multi-condition* option as the **Aggregate Function** for that measure.
10. Once Multi-condition is chosen as the **Aggregate Function** and Report by Descriptors is set to No, five new columns get added to Figure 14.39. They are, namely – **Target Operation, Target Values, Target Violation by Descriptors, Target Components**, and **Measurement Unit**.
11. From the **Target Operation** drop-down, select the logical function to define the Conditional Aggregation criteria. For the purpose of our example, select **>=** from this drop-down.
12. Since three conditions need to be set for the purpose of our example, the **Target Values** specification needs to be configured with 3 values - one each for reporting the Critical, Major, and Minor values (respectively) for the new measure.The format for the specification is as follows: *<Critical_target>/<Major_target>/<Minor_target>*. Accordingly, set **Target Values** to *80/60/30* in our example.
13. This specification applies only to multi-conditional aggregation. Here, you need to specify what percentage of components should violate the configured **Target Values** in order to report the corresponding measure. Like **Target Values**, this specification should also be in the format: *Critical/Major/Minor*. For instance, in the case of our example, if more than 80% of components consume more than 20% of CPU, then the measure should report the value *Critical*. So, the first value under **Target Components** should be 80. In the same way, the other two values that should be configured for **Target Components** is 60 and 30. The complete **Target Components** specification therefore is 20/60/30.

    Note:

    - For every **Target Value** configured, a corresponding value should exist in the **Target Components** specification.
    - If a **Target Value** is configured, but no **Target Components** specification corresponds to it, then eG will throw an error message to that effect, when you click the **Include** button in Figure 14.39. Similarly, if a **Target Components** specification does not have a corresponding **Target Value** configuration, then again eG will throw an error message to that effect.
14. Since the **Report by Descriptors** flag is set to **No**, you will find an additional **Target Violation**

by Descriptors column. This column offers the following options:

Option	Purpose
Majority	If you select this option, then the state of the majority of descriptors (for a measure) will be assigned to the component. For example, if 4 out of 5 descriptors of the *CPU utilization* measure violate the critical target value that you have configured, then the state of the corresponding component will also be Critical. However, if no single state has a clear majority, then by default, eG will assign the least state to the component. For instance, if the *CPU utilization* measure of a component reports metrics for 3 descriptors, and each of these descriptors violated the critical, major, and minor target values respectively, then by default, eG will assign the lowest state – Minor – to that component. Similarly, if of the 3 descriptors, one descriptor has not violated any of the target values and is Normal, then the state of the corresponding component will be Normal.
All	If this option is chosen, then eG will first check if all the descriptors for that measure are in the same state for a component. If so, then eG will assign that state to the component. On the other hand, if the descriptors of the measure are in different states, then eG will by default assign the least state to the component. For example, if all 5 descriptors of the *CPU utilization* measure violate the critical target value that you have configured, then the state of the corresponding non-aggregate component will also be Critical. However, if 3 of these descriptors violate the critical target value, 1 descriptor violates the major target value, and the other descriptor violates the minor target value, then eG will assign the lowest state – Minor – to the non-aggregate component. Also note that, if of the 3 descriptors, one descriptor has not violated any of the target values and is Normal, then the state of the corresponding component will also be Normal.
Any	If this option is chosen, then eG will compare each **Target Value** with the real-time value of the non-aggregate measure. This comparison will begin with the critical target value specification. Even if a single descriptor violates this target value, the Critical state is assigned to that component. If no descriptors violate the critical target value, eG then compares the major target value with the measure value. If this value is violated by even a single descriptor, then the state of the measure will be Major. Likewise, if no descriptors violate the critical or major values, eG compares the minor target value with the measure value. If this value is violated by even a single descriptor, then the state of the component will become Minor.

Since the **Any** option is what suits our purpose, select that for the *CPU utilization* measure in our example.

15. From the **Measurement Unit** drop-down, pick the unit in which the value of this aggregate measure is to be reported. For a multi-condition, the measurement unit can only be Number. Therefore, select Number as the Measurement Unit.
16. Then, select the check box corresponding to the *CPU utilization* measure and click he **Include** button to save the changes.
17. Next, switch to the Monitor interface to check what measure this aggregate test reports. Figure 14.40 depicts the value of the new aggregate measure. As is evident from Figure 14.39, the CPU health of the server farm is reported as Critical - this means that over 20% of XenApp servers in the farm are utilizing over 80% of CPU. To know which are these XenApp servers, click the ⊕ icon corresponding to the measure name in Figure 14.40.

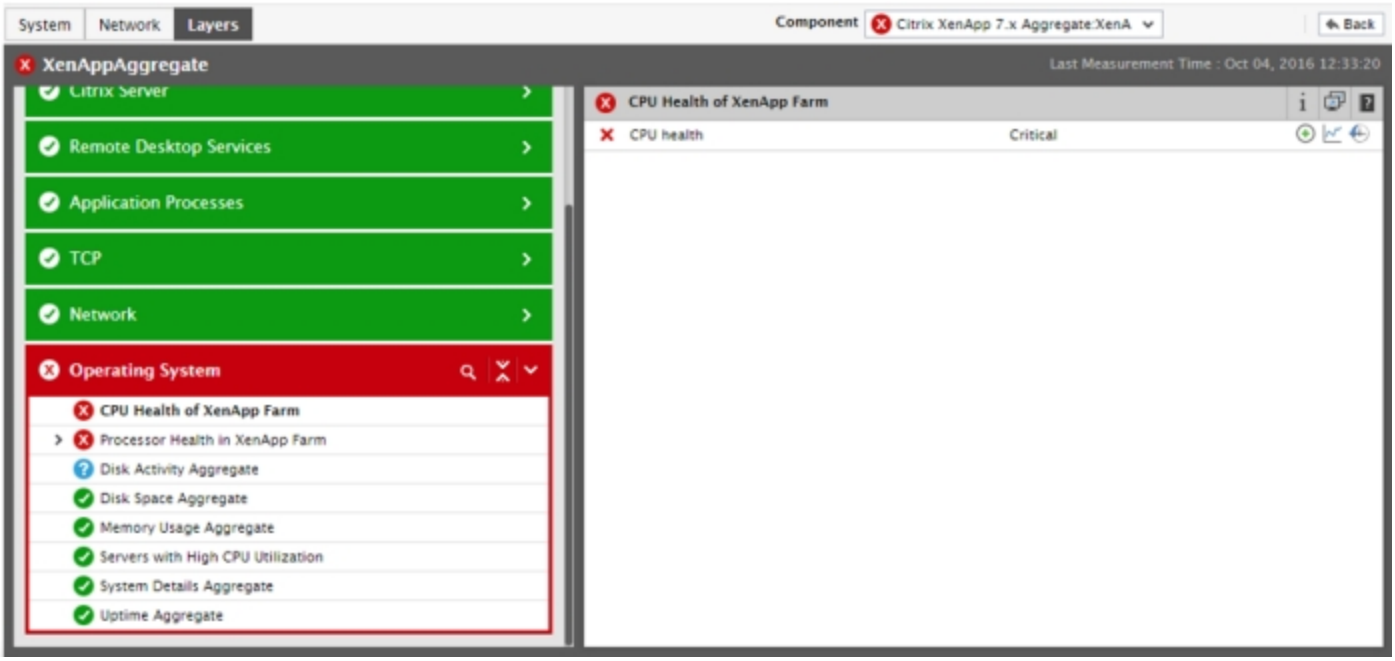

Figure 14.40: Viewing the value of the new aggregate measure

18. Figure 14.41 will then appear.

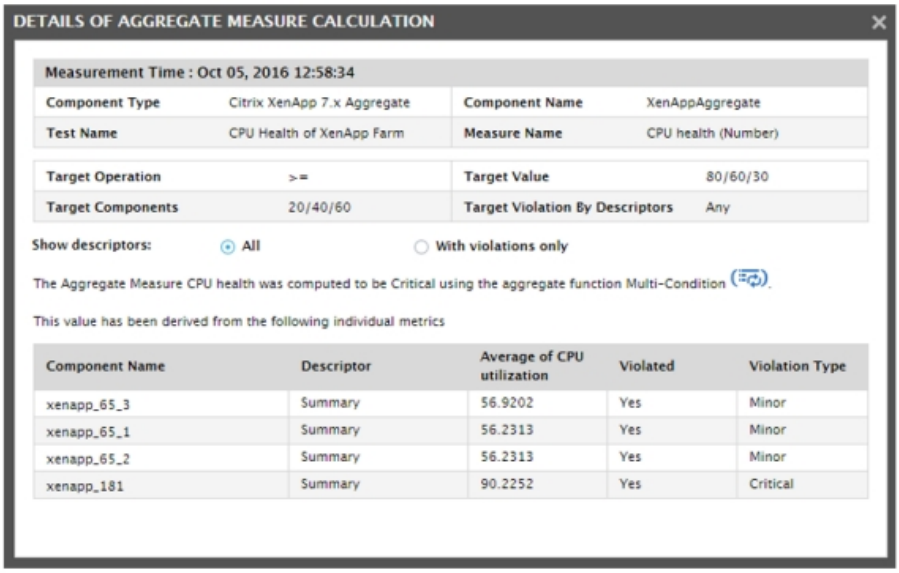

Figure 14.41: How the aggregate measure value was computed

19. Figure 14.41 describes how the value of the measure was computed. In the table at the bottom ofFigure 14.41, you can see that the components in the XenApp farm have been listed. Against each component, the descriptors - i.e., the Processors - supporting the corresponding server will be listed. The percent CPU utilization of each processor will also be displayed alongside, with an indication as to whether/not that utilization value has violated the configured **Target Value** condition. The type of violation - i.e., Critical, Major, or Minor - will also be indicated. In the case of our example, the **System Details** test of each of the XenApp servers in the farm reports metrics for the Summary descriptor only, and not for the individual processors. Therefore, the aggregate test also looks for violations in the Summary descriptor alone.

20. As per our specification, a server in the farm is considered to have violated a target value, if any of its descriptors violates that value. Since Summary is the only descriptor in our example, then a XenApp server will be deemed to have violated a **Target Value** if the CPU utilization value of this descriptor for that XenApp server exceeds any of the **Target Values** configured .

21. If you now look at the values in the Average of CPU utilization column in Figure 14.41, it will be clear to you that the CPU usage summary exceeds 80% in only 1 out of the 4 XenApp servers - i.e., in 25% of the servers. This violates the **Critical Target Value** and **Critical Target Components** specifications. This is why, the measure reports the value *Critical*.

14.6 Configuring Aggregate Tests

Aggregate tests can be configured the same way as non-aggregate tests. Like in the case of non-aggregate tests, you can change the default configuration of an aggregate test or change its configuration for a specific aggregate component alone.

To change the default configuration of an aggregate test, do the following:

1. Select the **Default Configuration** option from the **Tests** menu of the **Agents** tile.
2. When Figure 14.42 appears, select an aggregate component type from the **Component type** list and then the name of the aggregate test to be reconfigured from the **Test name** list.

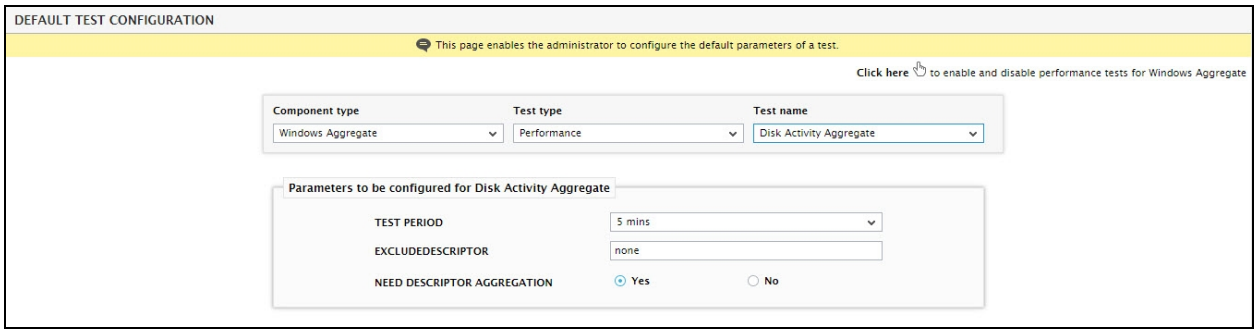

Figure 14.42: Modifying the default configuration of an aggregate test

3. By default, all aggregate tests take the same parameters. These parameters and their default will then be displayed, as depicted by Figure 14.42. Each of these parameters has been discussed below:
   - **TEST PERIOD** – Indicates the frequency at which the aggregate test runs. You can change it by picking a different frequency from the drop-down.

     Note:

     Make sure that the **TEST PERIOD** of the aggregate test is greater than or equal to the test period of the underlying non-aggregate test.
   - **EXCLUDEDESCRIPTOR** – **This flag is applicable to descriptor-based tests alone.** If the test being reconfigured is a descriptor-based test, then provide a comma-separated list of descriptors you want to exclude from aggregation.
   - **NEED DESCRIPTOR AGGREGATION** – **This flag is applicable to descriptor-based tests alone.** If you set this flag to **Yes**, then every measure of this test will report a single value that is aggregated across all descriptors of the test. If you set this flag to **No**, then this test will report an aggregated value per descriptor.
4. Finally, click the **Update** button in Figure 14.42 to save the changes. These changes will apply to all managed aggregate components of the **Component type** chosen.

To change the configuration of an aggregate test for a specific component alone, do the following:

1. Select the **Specific Configuration** option from the **Tests** menu of the **Agents** tile.
2. When Figure 14.43 appears, select an aggregate component type from the **Component type** list, and pick the component of that type for which tests have to be reconfigured from the **Component name** list.

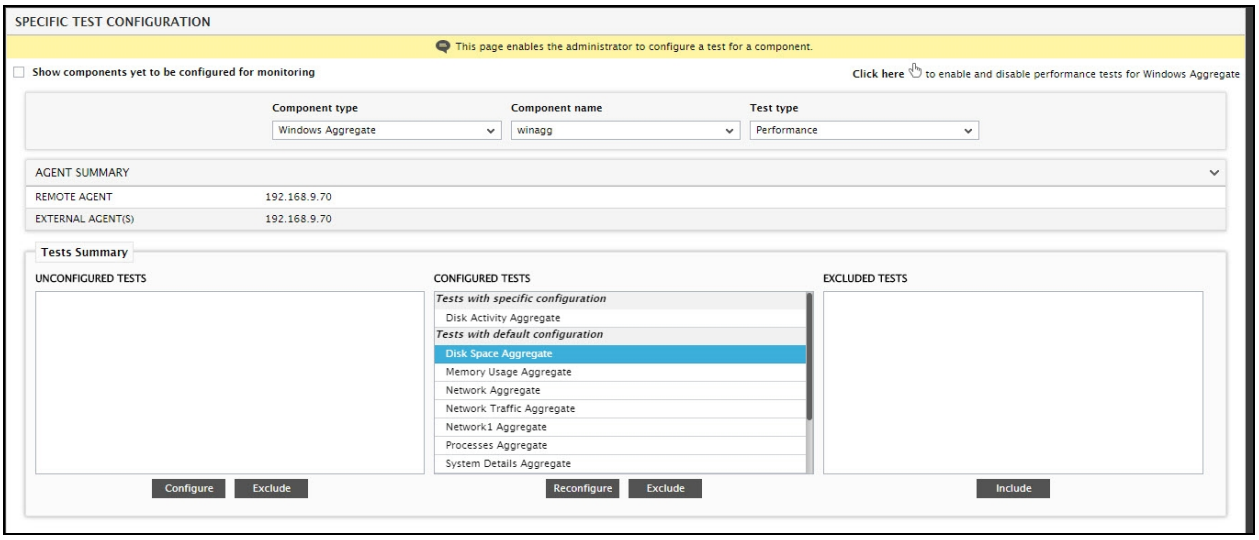

Figure 14.43: Selecting the aggregate component and aggregate test mapped to that component to be reconfigured

3. By default, all aggregate tests are pre-configured. This is why, when a **Component name** is selected, all aggregate tests mapped to that component, automatically appear in the **CONFIGURED TESTS** list of Figure 14.43.
4. To reconfigure a test for a specific component, select the test from the **CONFIGURED TESTS** list and click the **Reconfigure** button.
5. The parameters of the chosen test will then appear, as shown by Figure 14.44. Make changes you deem fit and click the **Update** button to save the changes. These changes will apply only to the component chosen from the **Component name** list of Figure 14.43.

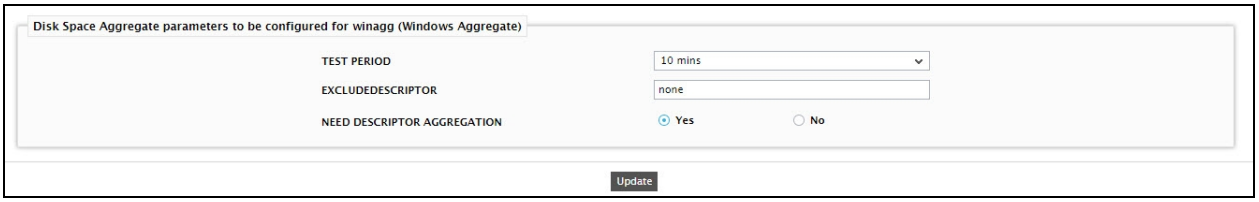

Figure 14.44: Reconfiguring an aggregate test for a specific aggregate component

14.7 Updating Test Period of Aggregate Tests

When one/more 'non-aggregate' components of a type are grouped to create an aggregate component, then the eG Enterprise system automatically computes the default frequency of every aggregate test mapped to that aggregate component on the basis of the default frequencies of the corresponding 'non-aggregate' tests. For example, say two Windows servers – *wina* and *winb* - have been grouped to form an aggregate component, named *winagg*, which is of type *Windows Aggregate*. By default, the **Disk Activity** test of *wina* and *winb* run every 5 minutes. Based on this default setting, eG will automatically configure the *Disk Activity Aggregate* test mapped to *winagg* to run at a 5-minute interval by default. In this case, the *Disk Activity Aggregate* test will be able to compute and report aggregated measures without a glitch. Even if an administrator later changes the frequency of the *Disk Activity Aggregate* test of *winagg* to 10 minutes, the test

execution will not be affected. However, if an administrator configures the *Disk Activity Aggregate* test of *winagg* to run at a frequency less than 5 minutes, or configures the *Disk Activity* test of *wina* or *winb* to run at a frequency higher than 5 minutes, then the *Disk Activity Aggregate* test will not be able to compute aggregate measures correctly.

To enable administrators to rapidly detect incorrect test frequency settings for aggregate tests and adjust these settings on-the-fly, eG Enterprise provides an **UPDATE AGGREGATE TEST PERIOD** page. This web page:

- Keeps track of test frequency changes made by the administrator to aggregate tests and non-aggregate tests related to the aggregate tests;
- Turns the spotlight on those frequency changes that can negatively impact the functioning of aggregate tests;
- Automatically computes and recommends the right frequency setting for the aggregate tests, and;
- Enables administrators to instantly apply these changes to the aggregate tests, so that aggregate test failures can be avoided.

To access the **UPDATE AGGREGATE TEST PERIOD** page and use it effectively, do the following:

1. Select the **Configure Aggregates** option from the **Tests** menu of the **Agents** tile.
2. Figure 14.45 will then appear. To check whether/not the tests mapped to an aggregate component are running in the correct frequency, first, select the **Component type** to which that aggregate component belongs. Then, select the aggregate **Component name**.
3. If one/more tests mapped to the chosen component are running at a frequency below the recommended frequency, then the names of these tests will be displayed as depicted by Figure 14.45.

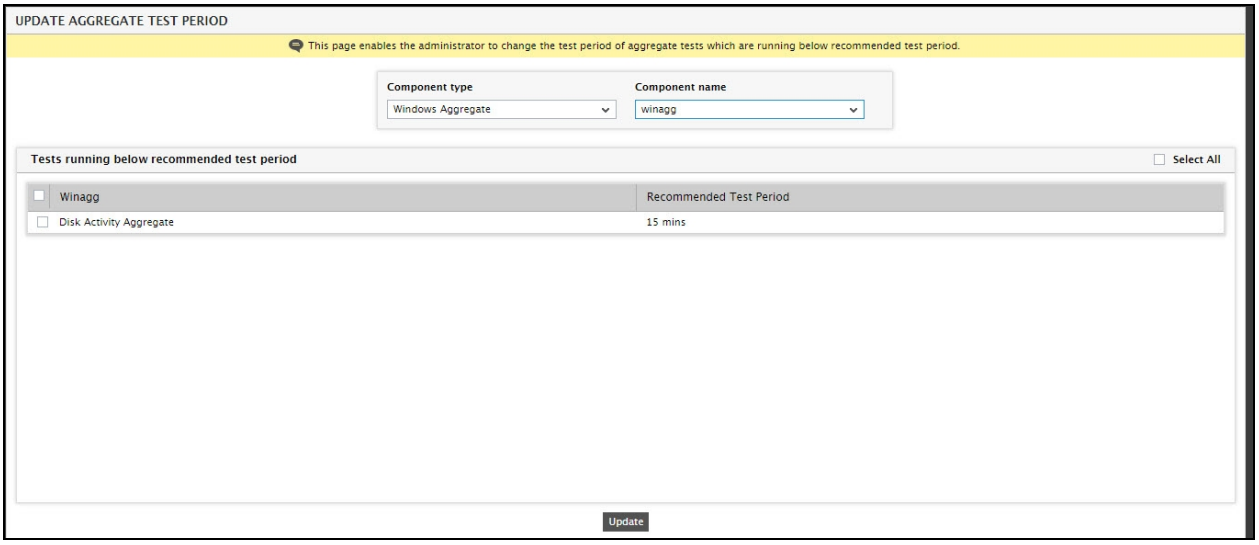

Figure 14.45: Selecting the aggregate component for which aggregate tests may not be running in the right test period

4. Against every such test, the **Recommended Test Period** for that test will be displayed.
5. To instantly apply this recommendation to a test, first select the check box alongside the test name (see

Figure 14.46) and click the **Update** button.

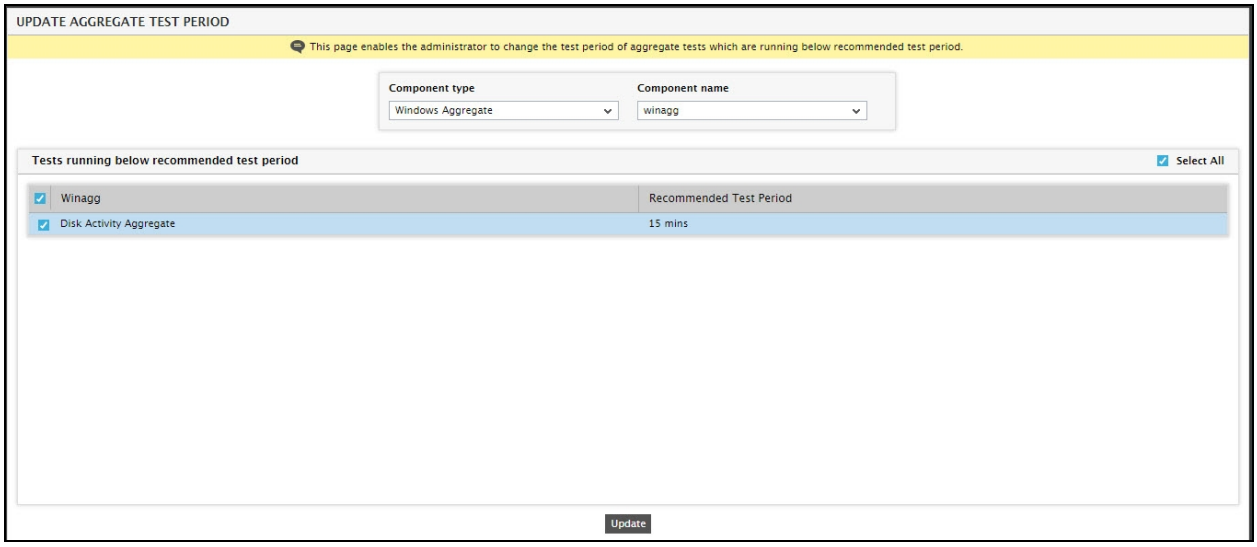

Figure 14.46: Applying recommended test period to a selected aggregate test

6. If the test is successfully updated with the recommended test period, a message depicted by Figure 14.47 will appear.

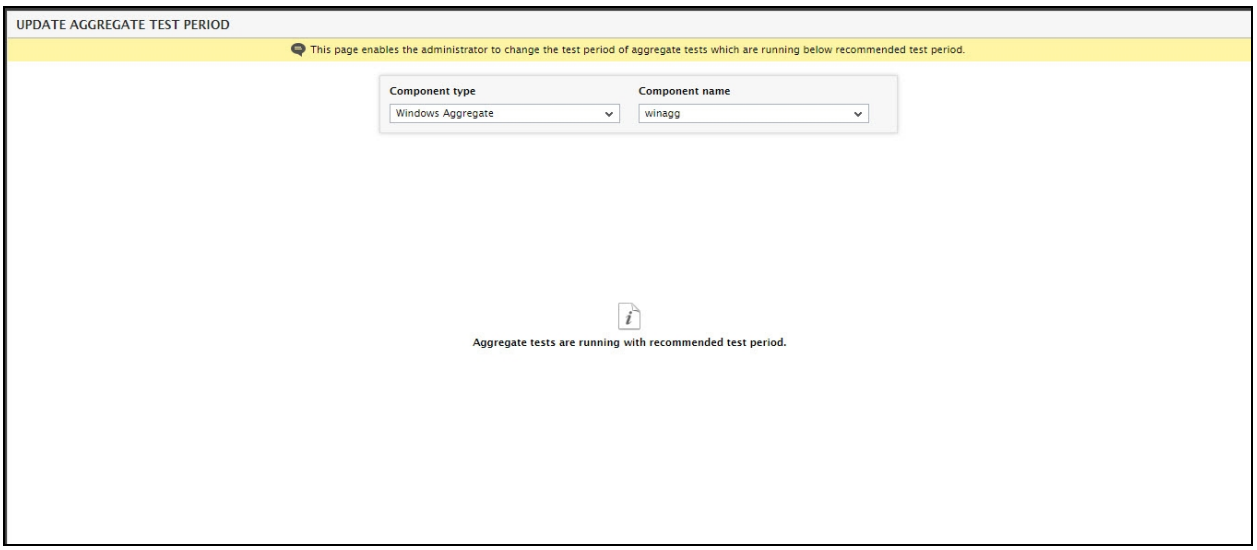

Figure 14.47: A message indicating that the aggregate test has been successfully updated with the recommended test period

15

Manager Redundancy

In the default deployment, eG Enterprise has a single central manager that receives performance metrics from the agents, detects anomalies, and sends email and SMS alerts to administrators regarding a probable problem condition. In this scenario, if for some reason the eG manager is rendered inaccessible, critical issues in the target infrastructure (e.g., failure of a network router, stopping of a key application process) will go unnoticed and therefore, unattended. To ensure high availability of the eG monitoring solution, eG Enterprise offers a redundant manager option wherein a secondary manager can act as an active standby for the primary manager. This capability, together with the ability to deploy redundant external agents in multiple locations, ensures that there is no single point of failure for the monitoring solution.

The Manager Redundancy is a license-controlled feature of eG Enterprise, and is governed by the **Cluster Type** option in the eG license. If the **Cluster Type** license option contains the value *Not Supported*, it indicates that the current installation of eG Enterprise supports a single manager only. If **Cluster Type** is set to *Active-Active*, then it indicates that manager redundancy has been enabled for that eG installation. A cluster can have a single **primary manager** and a single **secondary manager**. An *Active-Active* cluster is one where both the primary and secondary managers can both have agents reporting measures to them during normal operation (see Figure 15.1).

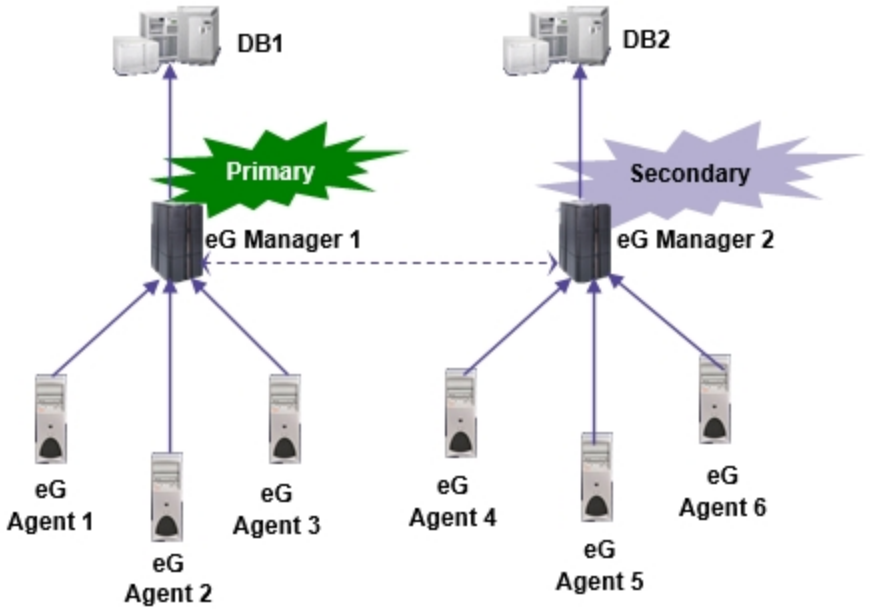

Figure 15.1: An Active-Active manager configuration

15.1 Configuring Manager Redundancy in Unix

In Unix environments, manager redundancy can be configured in one of the following ways:

- While configuring an eG manager
- After configuring the eG manager

15.1.1 Configuring Redundancy during Manager Configuration

When running the **setup_manager** script to set up an eG manager on Unix, the **setup_cluster** script in the **/opt/egurkha/bin** directory is automatically invoked. This script is used to configure redundancy, and when executed, requests for the following inputs:

1. First, the script requests a confirmation to enable the redundant manager capability of eG Enterprise.

   ```
   Would you like to enable eG Manager redundancy y/n [n]? y
   ```

   If **n** is specified, the **setup_cluster** script will automatically terminate.

2. If **y** is specified at step 1, you will then be required to indicate whether SSL has been enabled for the manager being configured.

   ```
   Please indicate if your eG Manager uses SSL y/n :[n] n
   ```

   Press **y** to indicate that the manager has been SSL-enabled, or **n** to deny it.

3. Next, indicate whether Network Address Translation (NAT) is being used for the manager being configured.

   ```
   Please specify if you use Network Address Translation(NAT) y/n :[n] y
   ```

   NAT facilitates multiple managers executing in different domains to communicate with one another. Specify **y** if any of the other managers in the redundant cluster have to use the NAT address of the current manager to communicate with it. If not, enter **n**.

4. If NAT is used (i.e., if **y** is specified at step 3), provide the NAT IP (or hostname), using which the managers interact with each other.

   ```
   Please enter the NAT IP/hostname of this eG Manager:
   ```

5. Similarly, also indicate whether Port Address Translation (PAT) is used. PAT again comes into play only when the managers span multiple domains. In such a case, enter **y**. If not, press **n**.

   ```
   Please specify if you use Port Address Translation (PAT) y/n :[n]y
   ```

6. If **y** is specified against PAT usage, then specify the PAT port number.

   ```
   Please enter the PAT port: 8088
   ```

7. Then, specify **y** if the current eG manager uses a proxy server for communication with other eG managers in the cluster.

   ```
   Please indicate if you would use proxy server for communications y/n :[n] y
   ```

8. If a proxy server is used, you will then have to provide the IP address (or hostname) and port number of the proxy server.

   ```
   Please enter the hostname of the proxy: 192.168.10.60
   ```

   ```
   Please enter the port of the proxy: 80
   ```

9. Indicate whether authentication is required for the proxy server, and if so, proceed to provide the user name and password to be used for the proxy.

```
Do you need authentication for the proxy? y/n [n]: y
```

```
Please enter the username to be used for the proxy: user
```

```
Please enter the password for user :
```

```
Please re-enter the password for user :
```

10. Next, state whether the manager being configured is to be set as the `primary manager`.

```
Is this a primary eG Manager y/n [n]? n
```

11. If the current eG manager is not a primary manager (i.e., if **n** is specified at step 10), then it is a secondary manager. Therefore, proceed to provide the IP address and port number of the primary manager with which this secondary manager communicates.

```
Please enter the hostname of the primary eG Manager: 192.168.10.59
```

```
Please enter the port of the primary eG Manager: 7077
```

Note:

- A target environment can have only one primary manager, and one secondary manager.
- Users with **Admin** privileges can login to the primary and the secondary managers, but he/she can access the eG administrative interface only when he/she logs into the primary manager.
- Monitor users can login to any of the eG managers in a redundant cluster.
- If you need to configure a redundant manager setup in an IPv6 environment, then make sure that the primary and secondary managers in the redundant cluster are configured with their host name and not their IP address.
- When running **setup_cluster** on a secondary manager, make sure that you specify the IP/hostname of the primary manager depending upon how you have configured the cluster in the primary manager. In other words, if when running **setup_cluster** on the primary manager, you have provided the IP address of the primary manager, then make sure that you provide the IP address only when **setup_cluster** prompts you for the details of the primary manager on the secondary manager.

12. Finally, indicate whether the primary manager uses SSL or not by specifying **y** or **n**.

```
Please indicate if your primary eG Manager uses SSL y/n [n]: n
```

Once this is specified, the **setup_cluster** script will exit, and the **setup_manager** script will resume.

Note:

While configuring the two managers in the cluster, ensure that each of the managers uses a separate database. The databases used can be of different types - i.e., Manager A in the cluster can use an Oracle database, and Manager B can use an MS SQL server.

15.1.2 Configuring Redundancy after Manager Configuration

To enable redundancy after installing and configuring the managers, do the following:

1. First, login as the eG user.
2. From the command prompt, move to the **/opt/egurkha/bin** directory, and execute the following command: **./setup_cluster**.
3. Upon execution, the **setup_cluster** script will first request for the location of the Java home directory.

```
Please enter the location of your Java home directory []: /usr/jdk1.3.1_06
```

4. Once the location is specified, setup will request your confirmation to proceed with enabling manager redundancy.

```
Would you like to enable eG Manager redundancy y/n [n]? y
```

5. While specifying **n** at step 4 will terminate the script execution, entering **y** will enable you to proceed with the setup by providing the IP (or hostname) and port number of the manager being configured.

```
Please enter the hostname (or IP address) of this host: 192.168.10.87
```

```
Please enter the port at which this eG Manager listens : 7077
```

Note:

- If an eG manager (primary/secondary) in a cluster supports only an IPv6 address, then its best that you configure redundancy for that manager using its hostname and not its IP address.
- If the eG manager is configured using the hostname, then ensure that cluster setup is also performed using the hostname only. Likewise, if the eG manager is configured using the IP address, then ensure that cluster setup is also performed using the IP address alone.

6. Once the IP and port are provided, steps 2 to 12 of this section will follow.
7. Finally, restart the manager.

15.2 Configuring Redundancy for a Manager on Windows

To enable redundancy for a manager on Windows, the batch file named **setup_cluster.bat** in the **<EG_INSTALL_DIR>\lib** directory needs to be executed. This file when executed, requests the following inputs.

1. Upon execution, the **setup_cluster** batch file will first request your confirmation to proceed with enabling manager redundancy.

```
Would you like to enable eG Manager redundancy y/n [n]? y
```

Specifying **n** here will terminate the script execution. If you enter **y**, steps 2 to 12 of Section **15.1.1** will follow.

15.3 How does Manager Redundancy Work?

To make the eG manager redundant, do the following:

1. After installing the different eG managers that are a part of the cluster, start the managers. To enable the secondary manager to synchronize their time with the primary manager, start the primary manager first and then the secondary manager.

   Note:

   - The primary and secondary managers can exist in different time zones.
   - Every time the secondary manager synchronizes its time with the primary manager, the time is recorded in the **eg_db.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory) of the secondary manager. Typically, the secondary manager will attempt to synchronize itstime with the primary manager every 1 hour. During such attempts, if a primary manager is unavailable, then the secondary manager synchronizes its time with the last value stored in the **eg_db.ini** file.

     However, while attempting to start the secondary manager for the very first time, no time record will be available in the **eg_db.ini** file. During such a situation, if the primary manager is unavailable, then all attempts to start the secondary manager will fail.

2. As an admin user can login to the primary manager only, configure the environment using the primary manager. The secondary manager then downloads the latest configuration details from the primary manager.

   Note:

   Configuration changes subsequently made using the primary manager will immediately be reflected in the secondary manager.

3. Install the eG agents. Refer to the *eG Installation Manual* for a detailed installation and configuration procedure.

4. Typically, the manager to which an agent reports will be set during the time of agent installation itself. To facilitate manager redundancy, this manager-agent association can be overridden using the eG administrative interface. To change the mapping, do the following:

   - Login to the primary manager's administrative interface as *admin*.
   - Next, follow the menu sequence: Agents -> Assign to Manager
   - Figure 15.2 will then appear. By default, the primary manager will appear as the **Manager** (see Figure 15.2). Also, by default, all the agents configured in the environment will appear in the **AGENTS ASSIGNED TO PRIMARY MANAGER** list box of Figure 15.2. In a redundant manager configuration, the manager IP/hostname and port number specified at the time of installing an agent is ideally the IP address and port number of the primary manager. When a component is added for monitoring via the eG admin console, a mapping is created in the eG manager indicating the preferred manager to which this agent should be reporting during normal operation of the cluster. By default, the preferred manager for all agents is the primary manager. The eG admin interface provides a way for an administrator to override this default setting.

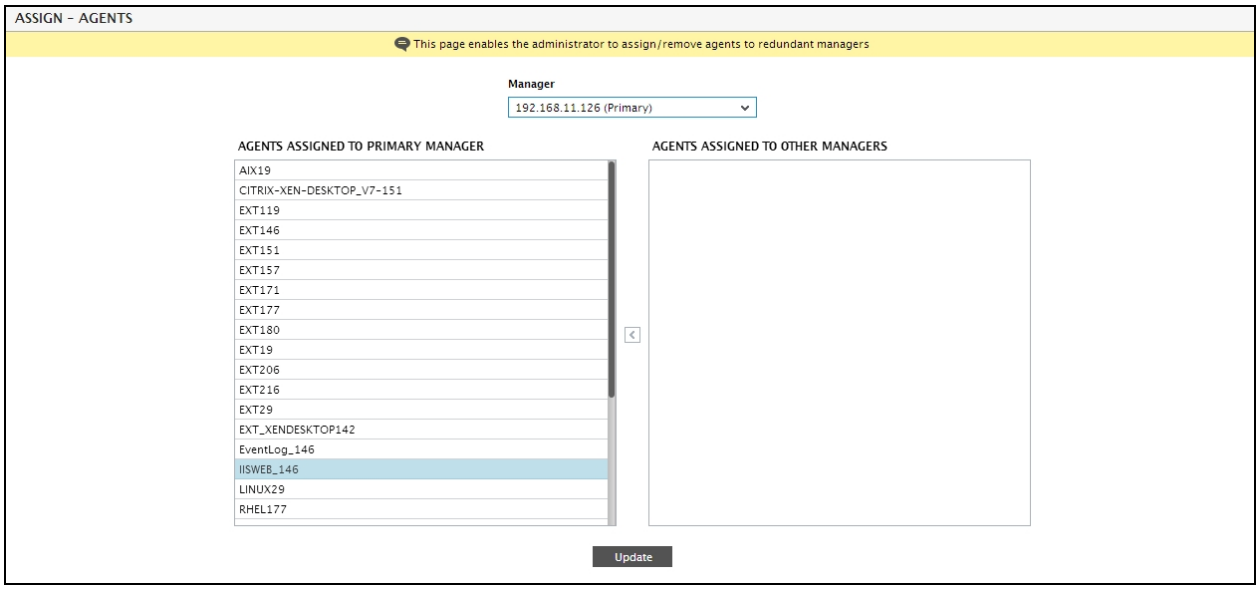

Figure 15.2: All agents configured in the environment appearing in the AGENTS REPORTING TO PRIMARY MANAGER list box

- Note that the eG Enterprise system does not permit users to directly disassociate agents from the primary manager. Instead, you will have to assign an agent to the secondary manager, so that it no longer reports to the primary.
- For this, first associate agents with the secondary manager. To achieve this, first, select a secondary manager from the list box against **Manager** in Figure 15.3.

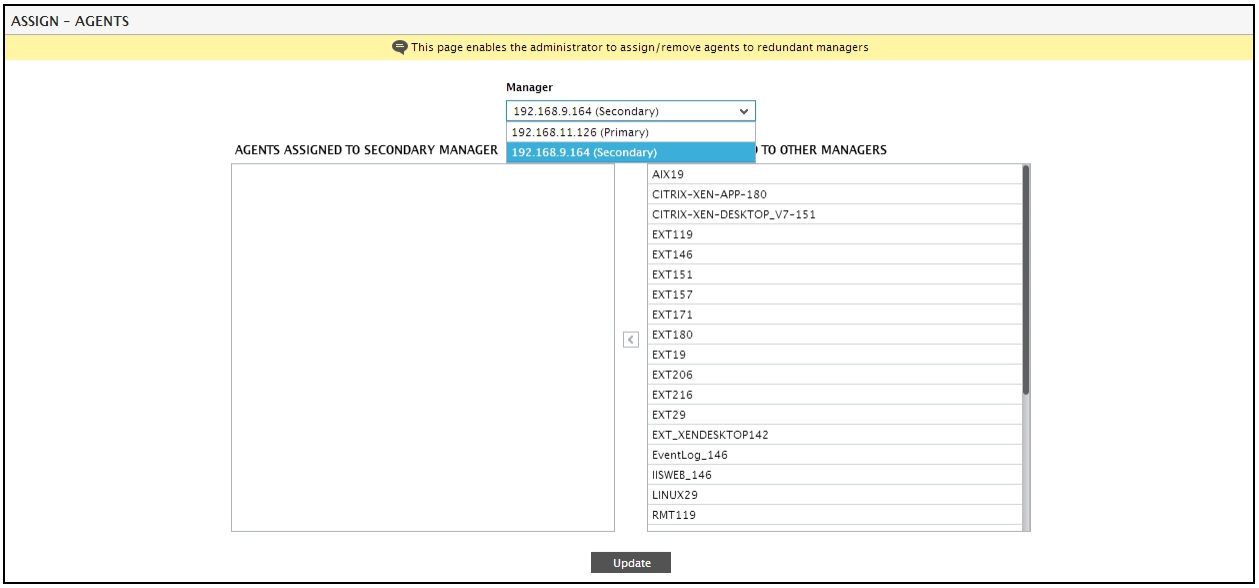

Figure 15.3: Selecting a secondary manager

- By default, the **AGENTS ASSIGNED TO OTHER MANAGERS** list of Figure 15.3 will display all the agents that are currently reporting to the primary manager. To assign agents to the chosen secondary manager, first, select an agent from the **AGENTS ASSIGNED TO OTHER MANAGERS** list (see Figure

15.4).

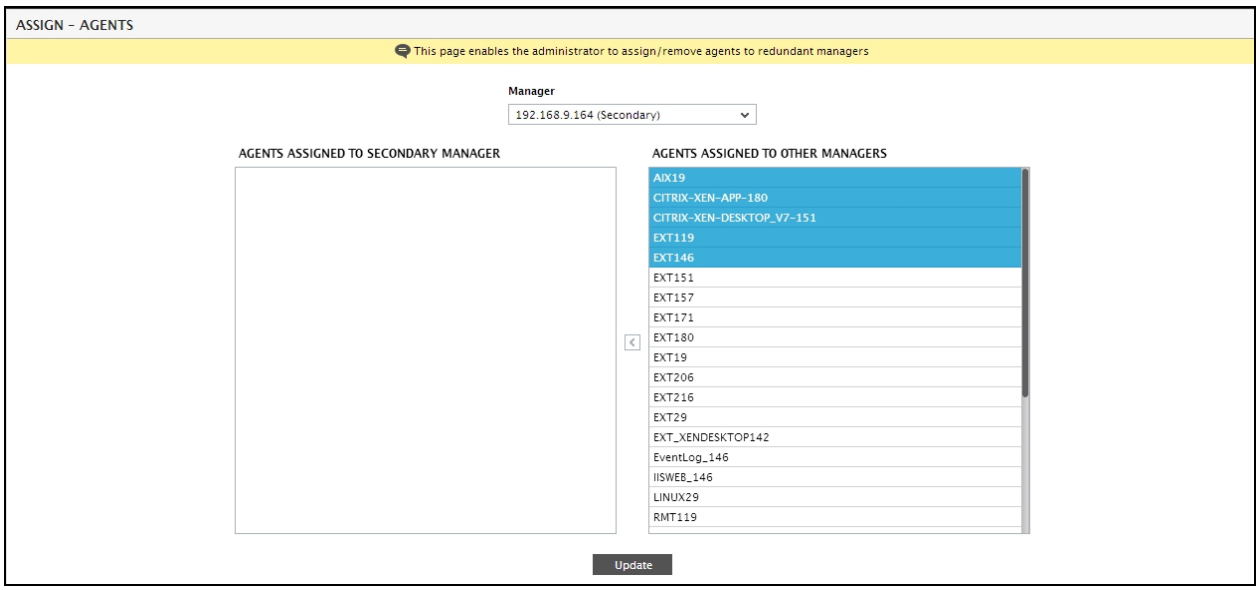

Figure 15.4: Selecting an agent to be associated with the chosen secondary manager

- Next, click on the **<** button in Figure 15.4, so that the selected agent moves to the **AGENTS REPORTING TO SELECTED MANAGER** list (see Figure 15.5).

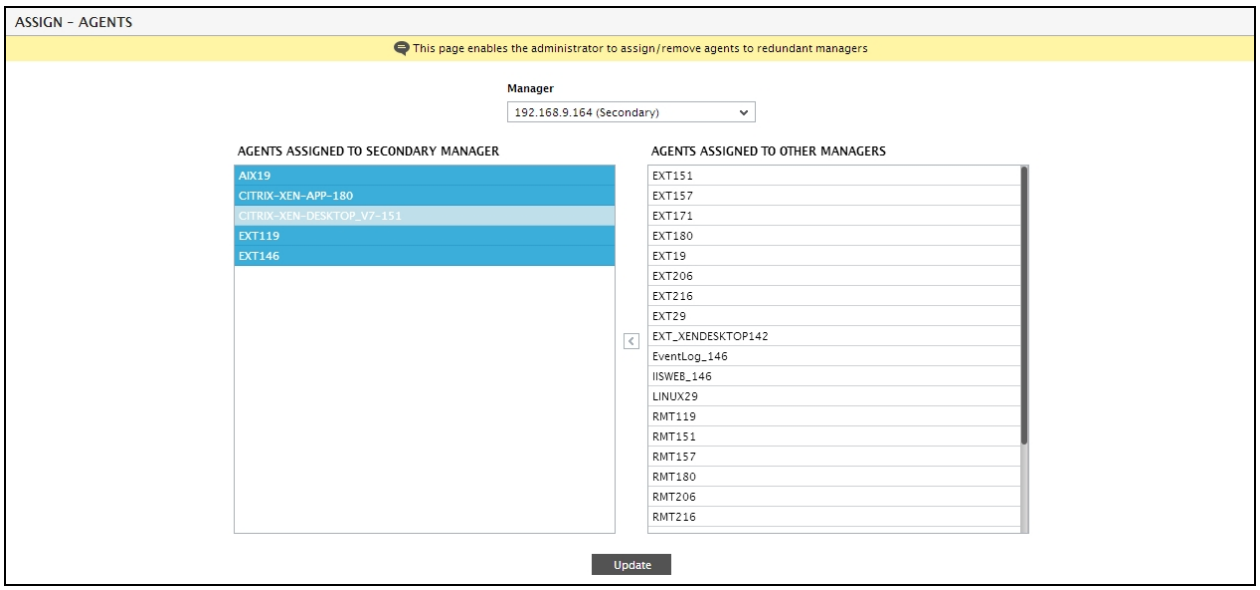

Figure 15.5: Transferring the selection to the AGENTS REPORTING TO SELECTED MANAGER list

- This exercise also ensures that the chosen agent(s) is successfully disassociated from the primary manager.
- Then, start the eG agents.

5. Each of the managers will then receive performance metrics collected by each of the agents assigned to them. Measures reported to a manager will be made available to the other manager as well.

Note:

- As data reported by the agents will be available to all the managers, there will be no difference in the component state and reports displayed by the managers in the cluster.
- The origin of email alerts depends upon the manager-agent association. In other words, a manager receiving measurement data from an agent will generate all the email alerts pertaining to that agent.
- While the printing/mailing of a report can be scheduled using a secondary manager, it is the primary manager that implements the schedule. Therefore, if the primary manager is not available, then the reports will not be printed/mailed as per schedule.

6. If a manager becomes unavailable, agents reporting to it will automatically begin reporting measures to the other available manager (see Figure 15.6).

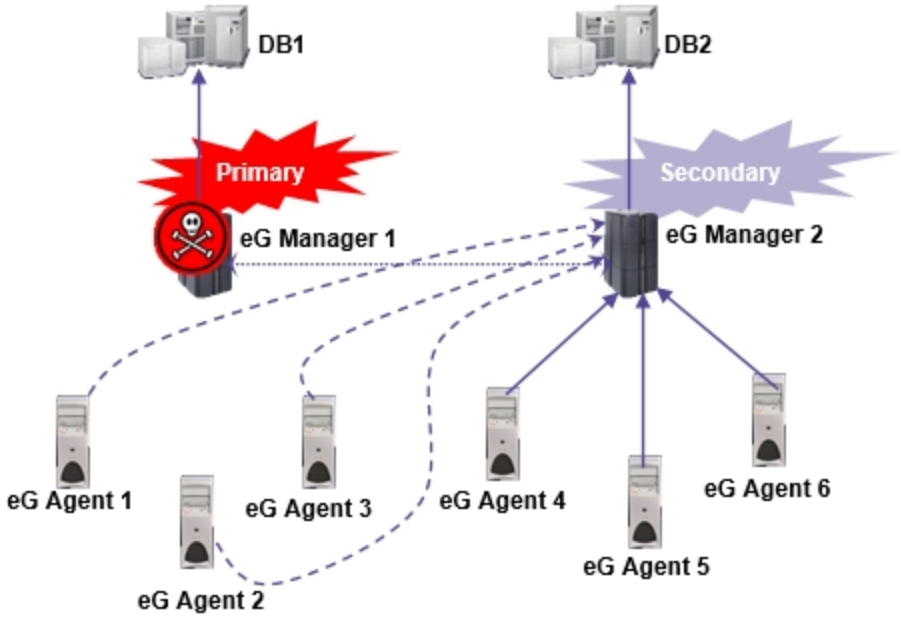

Figure 15.6: A figure depicting what happens when a manager is rendered unavailable

7. When a manager comes back online, the agents which switched loyalties earlier will automatically begin reporting measures to the original manager to which they were assigned.
8. When a manager is offline, other manager receiving data from agents can be configured to store the data that it receives from the agents locally. When the manager comes back up, the other manager will transmit the saved data to the manager that has just come up. This ensures that there is minimal data inconsistency between the different eG managers.

Note:

- The amount of data that can be stored by a manager for transmission to other offline managers is controlled by two configuration settings - **maxStoragePerFile** and **filesPerManager** - that are present in the file **eg_managers.ini** located in the **<EG_INSTALL_DIR>\manager\config** directory.
- The setting **maxStoragePerFile** defines the amount of data (in MB) that can be stored in each temporary file that is used to store data temporarily for transmission to a manager that is offline. An eG manager

can store data in multiple files for transmission to the offline manager. Multiple files are used for storage (rather than a single file) to minimize data read/write operations to memory for transmission to the other manager. The **filesPerManager** setting defines the maximum number of data files per manager that are used for temporary storage of data.

- By default, the **maxStoragePerFile** value is 0, and the **filesPerManager** is 0. This implies that a manager does not save data it receives from agents directly for transmission to another manager that may be offline. If the **maxStoragePerFile** is 10 and the **filesPerManager** is 20, then 200MB of data can be saved for transmission to another manager that is offline.
- When a manager comes back online after a failure/shutdown, the agents that should be reporting to it do not switch back immediately. This is because, while agents are attempting to report new measurements to the original manager, the other manager will attempt to update the original manager with the data that was collected during its non-availability. In order to avoid data congestion at this juncture, the agents will wait until all the data in the **<EG_INSTALL_DIR>\manager\data** folder of the other manager is cleared, and will then begin reporting to their assigned manager.

15.4 How to Remove a Secondary Manager from a Cluster?

To remove a secondary manager from a cluster, do the following:

1. Uninstall the secondary manager.
2. Next, edit the **eg_managers.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory of the primary manager. This file consists of individual sections containing details about each of the managers in the cluster. For example, if 192.168.10.12 is one of the managers in the cluster, then a section headed **[192.168.10.12]** will be present in the **eg_managers.ini** file. Once a secondary manager is uninstalled, remove the section corresponding to that manager from the **eg_managers.ini** file.
3. A section titled **[SECONDARY]** will also exist in the **eg_managers.ini** file. This section lists all the secondary managers in the cluster. Remove the entry corresponding to the uninstalled secondary manager from this list.
4. Then, save the **eg_managers.ini** file and restart all the managers in the cluster.
5. If any agents report measurements to the uninstalled secondary manager, then using the eG administrative interface (Agents -> Assign to Manager menu) assign such agents to another manager.

Note:

To remove a primary manager from a cluster, the secondary manager should be converted into a primary manager.

15.5 How to Convert a Secondary Manager into a Primary Manager?

To convert a secondary manager into a primary manager, do the following:

1. Uninstall the primary manager.
2. Next, identify the secondary manager that should now be set as the primary manager.
3. Then, proceed to edit the **eg_managers.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory of every manager in the cluster. The **eg_managers.ini** file contains a section named **[PRIMARY]**. To convert a primary manager into a secondary manager, first, remove the entry corresponding to the uninstalled primary manager from the **[PRIMARY]** section.
4. Then, from the **[SECONDARY]** section, copy the entry that corresponds to the secondary manager (to be set as the primary manager), to the **[PRIMARY]** section. Later, from the **[SECONDARY]** section, remove the entry that was just copied.
5. Next, save the **eg_managers.ini** file and restart all the managers in the cluster.
6. Agents reporting to the uninstalled primary manager can be reassigned to the new primary manager using the eG administrative interface (Agents -> Assign to Manager menu). Refer to Section **15.3** for further details.

15.6 How to Convert a Primary Manager into a Secondary Manager?

To convert a primary manager into a secondary manager, first set a secondary manager as a primary manager using the procedure explained in Section **15.5**. The only change in the procedure of Section **15.5** would be that instead of removing the original entry in the **[PRIMARY]** section (see step 3) of the **eg_managers.ini** file, copy the entry to the **[SECONDARY]** section.

15.7 How to Convert an Existing Non-Redundant Setup into a Redundant Setup?

To achieve this, the following broad steps will have to be followed:

1. Convert the existing Non-Redundant Manager into a Redundant Manager
2. Add one more redundant manager to the setup.

15.8 How to Convert an Existing Non-Redundant Manager into a Redundant Manager?

In order to achieve this, do the following:

1. Stop the existing eG manager.
2. Execute the **setup_cluster** script in the **<EG_INSTALL_DIR>\lib** directory (on Windows. On Unix, this will be **/opt/egurkha/bin**) of the manager system.
3. Using the script, set the existing manager as the primary or the secondary manager of the redundant setup.
4. Install the eG license that supports a redundant manager setup.
5. Start the manager.

6. If you have set the existing manager as the primary manager, then connect to the primary manager and login to its administrative interface.
7. Open the **MANAGER/UNMANAGE** page using the menu sequence: Components -> Manage.
8. Manage any component using this page and update the management by clicking on the **UPDATE** button. This needs to be done in order to inform the clustered environment of the currently monitored components.
9. Avoid performing configuration changes until the additional managers are configured and added to the redundant setup.
10. If you have set the existing manager as the secondary manager, then do not start this manager until a primary manager is added to the redundant setup and is started.

15.9 How to Add Another Manager to the Redundant Setup?

To achieve this, do the following:

1. Install the new manager and configure it to use the same database server as the old manager or a separate database server.
2. Then, execute **setup_cluster** (on Windows it will be in the **<EG_INSTALL_DIR>\lib** directory; on Unix, it will be in the **/opt/egurkha/bin** directory) on the new manager system.
3. Using the script, set this manager as the primary or the secondary manager of the redundant setup.
4. If you want even the historical data in the old manager's database to be replicated to the database of the new manager, then, follow the steps given below to achieve this:
   - Take a backup of the old manager's database.
   - Restore it to the new manager's database server using the database name assigned to the new manager's database. For example, if you have configured the new manager to use *eg_database* as its database, then you have to to restore the old manager's database it to the new manager's database server as *eg_database*.
5. On the other hand, if you want the old manager to share with the new manager only that data it receives after the redundant cluster is fully configured and started (i.e., only the future data and not the past data), then the above-mentioned backup-restore procedure can be dispensed with.
6. Next, to ensure that the new manager is updated with the details of IC tests configured on the old manager, follow the steps given below:
7. First, from the ini files listed below, search for those entries that are relevant to the IC tests configured for the old manager, and copy them to the same ini files in the new manager. To search, use the name of the IC tests.
   - **eg_agents.ini**
   - **eg_db.ini**
   - **eg_dbase.ini**
   - **eg_specs.ini**
   - **eg_tables.ini**

- **eg_tests.ini**
- **eg_thresholds.ini**
- **eg_udtests.ini**

Typically, all these files will be available in the **/opt/egurkha/manaer/config** directory of the old and new managers.

8. Similarly, search for entries related to IC tests in the **eg_ newtests.ini** file (in the **/opt/egurkha/manager/config/tests** directory), and copy them to the file with the same name in the new manager's system.
9. Copy the test classes defined in the old manager for the new tests added using IC to the new manager system. Typically, the test classes reside in the **/opt/egurkha/manager/config/tests** directory (on Windows, it will be in the **<EG_INSTALL_DIR>\manager\config\tests** directory), and will be named in the format: **<IC_Test_ Name>.class**. Copy the **.class** files from the above-mentioned location in the old manager system, to the same location in the new manager system.
10. Next, in the **/opt/egurkha/bin/database** folder of the old manager, look for sql files that are named after the IC tests configured on the old manager. These files typically contain the queries required for creating the tables for the IC tests. Run each of these queries on the new manager's database as the eG database user, so that the required tables are created therein.
11. To apply the IC-based changes on Unix, execute the **upload** script from the **/opt/egurkha/bin** directory. To do this, switch to the **/opt/egurkha/bin** directory from the command prompt, and type the command: **./upload**. The script will then request you to specify the Java home directory:

```
Please enter the location of your Java home directory :
```

Once the home directory is specified, press the **Enter** key on the keyboard to update the configurations.

12. To apply the IC-based changes on Windows, simply execute the **upload.bat** batch file in the **<EG_INSTALL_ DIR>\lib** directory.
13. Copy the icons/images used by IC tests from the old manager and place them in the appropriate locations (**/opt/egurkha/manager/tomcat/webapps/final/monitor/eg_ images/eg_ layout/eg_ icons/** and (**/opt/egurkha/manager/tomcat/webapps/final/admin/eg_images/**) in the new manager.
14. Install the license that enables the redundant manager capability.
15. Start the new manager.
16. Then, if the new manager is set as the primary manager, connect to it and login to its administrative interface.
17. Open the **MANAGE/UNMANAGE** page using the menu sequence: Components -> Manage.
18. Manage any component using this page and update the management by clicking on the **UPDATE** button. This needs to be done in order to inform the clustered environment of the currently monitored components.
19. If you have set the new manager as the secondary manager, then do not start this manager until a primary manager is added to the redundant setup and is started.

Note:

The secondary manager in a cluster will sync its time with that of the primary manager. Therefore, if the new manager is set as the primary manager of a cluster, and both the old and new managers exist in different time zones, then a data gap or data overlap (as the case may be) is bound to occur.

15.10 Redundancy FAQ

1. **What is an Active/Active cluster?**

   An *Active-Active* cluster is one where both the primary and secondary managers can both have agents reporting measures to them during normal operation.

2. **Do I need to use separate database servers for each of the managers in a cluster?**

   The different managers in a cluster can use the same database server, but the database name assigned for the eG database should be different for all the managers in the cluster.

3. **Do the primary and secondary managers have to be sized the same?**

   Ideally, the same amount of CPU, memory, and disk resources need to allocated to the primary and secondary managers, as the volume of data handled by both the managers will be the same.

4. **Can the managers in a redundant setup use different operating systems?**

   Yes. The eG Redundant Manager setup provides heterogeneous platform support. For instance, you can have one of the managers on a Windows platform, and another on Solaris.

5. **Can the managers in a redundant setup have different double-byte configurations?**

   No. A redundant setup cannot have a primary manager that is double-byte enabled and a secondary manager that is not, or the vice-versa. If so, then the redundant setup will not function properly. It is therefore recommended that all managers in a redundant setup have the same double-byte configuration.

6. **Can only one of the managers in a redundant cluster integrate with the Active Directory server for management of domain users?**

   This is not recommended. For instance, if only the primary manager has been configured to integrate with the Active Directory server and not the secondary manager, then domain users created on the primary manager cannot login to the secondary manager. The vice-versa also holds good. You are therefore advised to ensure that all managers in a redundant setup either integrate with the Active Directory server or do not.

7. **Can AES encryption be enabled for only one of the managers in a cluster?**

   No. You need to make sure that the AES encryption status is uniform across the redundant setup - i.e., you can either enable AES encryption for all managers or not enable AES encryption for all managers. If say, AES encryption is enabled only for the primary manager and not the other managers in the cluster, then the redundant setup will not function properly.

8. **Do I need cluster hardware to support the redundant managers?**

   No. There is no need for any specific cluster hardware to be installed. eG Enterprise handles the cluster functionality at the application-level.

9. **How does a manager in a cluster communicate with the other manager?**

   The primary and secondary managers in a cluster communicate via HTTP/HTTPS. If any of the managers in a cluster is behind a firewall, then the firewall needs to be configured to allow "two-way" communication - i.e., both inbound and outbound traffic from the manager should be allowed.

   The manager port can be configured during installation. 7077 is the default port for all the managers.

10. **How much bandwidth is needed for the communication between redundant managers?**

    Since all the data from one manager will be replicated to the other, the bandwidth requirement for manager-manager communication should be equal to the sum of the data traffic from the agents to the manager.

11. **Can both the managers in the cluster be started simultaneously? If not, why?**

    No. It is recommended that you start the primary manager first, and then the secondary managers. This needs to be done to ensure that all the secondary managers synchronize their time with the primary manager.

12. **What happens if the managers in a cluster are in different time zones?**

    Since the secondary manager synchronizes its time with the primary manager, alerts and measure data will be reported only in the primary manager's time.

13. **What if I change the time zone of the primary manager?**

    If you change the time zone of the primary manager while it is running, then ensure that you restart the primary manager after the timezone change so that, the change takes effect.

14. **Does each manager in the cluster see all the data from the agents?**

    Yes. All the data from all the agents in the environment will be made available to all the managers in the cluster.

15. **Do the agents report to both the managers in a cluster? If not, how does redundancy work?**

    No. At any given point in time, an agent reports to only one manager in a cluster. The measure data reported by an agent to a single manager will be sent by that manager to all other managers in the cluster.

16. **How does an agent know that a manager in the cluster is down and how does the agent react to this?**

    The eG agent periodically reports measurement data to the eG manager for state computation and alarm generation. If the agent is unable to upload metrics to the manager, it is a clear indicator that the manager is down. Under such circumstances, the eG agent will read the secondary manager information configured in the **eg_managers.ini** file, and will try to connect to every secondary manager in the cluster in the same order in which they have been listed in the ini file. The agent will finally start reporting the measure data to the first manager it establishes a connection with.

    The agent will continue reporting measures to the new manager, until its original manager becomes available. Once a manager in the cluster becomes available, every other manager in that cluster will be informed of this change in status. Once the manager to which the agent is currently reporting comes to

know that the original manager is online, it alerts the corresponding eG agent, which will then switch back to the original.

17. **When one manager in the cluster is down, does the other manager in the cluster store data for this manager and then send the data back when the manager comes up?**

    When a manager is offline, the other manager receiving data from agents can be configured to store the data that they receive from the agents locally. When the manager comes back up, the other manager will transmit the saved data to the manager that has just come up. This ensures that there is minimal data inconsistency between the different eG managers.

    The amount of data that can be stored by a manager for transmission to the other offline manager is controlled by two configuration settings - **maxStoragePerFile** and **filesPerManager** - that are present in the file **eg_managers.ini** located in the **<EG_INSTALL_DIR>\manager\config** directory.

    The setting **maxStoragePerFile** defines the amount of data (in MB) that can be stored in each temporary file that is used to store data temporarily for transmission to the other manager that is offline. The **filesPerManager** setting defines the maximum number of data files per manager that are used for temporary storage of data.

    By default, the **maxStoragePerFile** value is 0, and the **filesPerManager** is 0. This implies that a manager does not save data it receives from agents directly for transmission to other managers that may be offline. If the **maxStoragePerFile** is 10 and the **filesPerManager** is 20, then 200MB of data can be saved for transmission to another manager that is offline.

    It is recommended that this storage specification be small.

    Note:

    The **maxStoragePerFile** and **filesPerManager** settings govern how much “measurement data” can be stored by a manager for transmission to other offline manager. It does not however govern “configuration data” - in other words, if configuration changes are made to a manager (using the eG administrative interface) during such time that the other managerin the redundant cluster is offline, then all these changes, regardless of size, will be automatically stored by a manager for transmission.

18. **How does state computation happen in a cluster - is the state of the target infrastructure shared between the managers in the cluster?**

    No. The state of the target infrastructure is not shared between the managers in the cluster. Instead, state computation is performed by the individual managers in the cluster.

19. **When setting up a redundant cluster, do both the managers have to be available first, or can one of the managers be set up first and the second one set up later. If yes, what are the steps involved?**

    The recommended approach is to have both the managers available and configured in the redundant mode. However, if for some reason, both the managers are not available simultaneously, then do the following:

    - Configure the existing manager to function in the redundant mode.
    - Install the eG license that supports a redundant manager setup.
    - If this manager is set as the primary manager, start it and make configuration changes to it.

- Install and configure the other manager to function in a redundant setup.
- Backup the database of the first manager and restore it to the second manager's database server.
- Similarly, copy the config files from the first manager to the second manager.
- Start the primary manager and then the secondary manager.

20. **Can configuration changes (administration privileges) be done on both the managers in a cluster?**

    No. An admin user is authorized to login and make configuration changes to the primary manager only.

21. **Are configuration changes done in one manager automatically propagated to the other manager in the cluster?**

    Yes. The primary manager will automatically transmit the updated configuration information to the other manager in the cluster.

22. **Will both the managers in a cluster send out email/SMS alerts/trouble tickets when a problem is detected?**

    The origin of email/SMS alerts/trouble tickets depends upon the manager-agent association. In other words, a manager receiving measurement data from an agent will generate all the email alerts pertaining to that agent.

23. **How about scheduling of Reports? Does each manager send out scheduled reports?**

    While the printing/mailing of a report can be scheduled using a secondary manager, it is the primary manager that implements the schedule. Therefore, if the primary manager is not available, then the reports will not be printed/mailed as per schedule.

24. **When an existing non-redundant setup is converted into a redundant setup, when do the agents reporting to the managers become cluster aware?**

    Typically, if an agent is started only after the redundant setup is configured and running, then the agent will recognize the existence of a cluster as soon as it polls the manager for configuration information.

    On the other hand, if agents are already running, and the redundant setup is configured only later, then, soon after the redundant cluster is configured, you should restart all the agents using the **Restart all agents** button in the **AGENTS - STATUS** page (Agents -> Status menu sequence) of the eG administrative interface, to ensure that the agent is updated with the changes.

25. **How does agent upgrade happen in a redundant setup?**

    The agent upgrade patch needs to be copied to both the managers in a cluster. The agents then download the patch from the corresponding manager, and get upgraded, automatically.

Caveats

- A manager that locally stores the measure data pertaining to an unavailable manager, will send the data to the other manager when it comes online, but will not send the state computations. This could cause discrepancies in the event history and reporter modules across the managers.

- Settings such as mail server settings should be common to all the managers in the cluster. Moreover, all the managers should have access to the resources specified in the settings.
- Managers which are part of a redundancy setup cannot have the same internal IPs, as each manager identifies the other using the internal IP only.

16

Auto Correction

Increased uptime and lower mean time to repair are critical to ensuring that IT infrastructures deliver a high quality of service to users. Towards this end, eG Enterprise embeds an optional auto-correction capability that enables eG agents to automatically correct problems in the environment, as soon as they occur. With this capability, as and when an abnormal situation is detected, an eG agent can initiate corrective actions automatically to resolve the problem. Automatic correction without the need for manual intervention by IT operations staff reduces service downtime and improves operational efficiency.

By default, the auto-correction capability is available in eG Enterprise for the *Processes running* measure of Processes Test, and the *Service availability* measure of WindowsServices Test. The sections that follow explain how to enable the auto-correction capability of ProcessTest and WinServiceTest.

16.1 Enabling Auto-Correction for Processes Test

eG Enterprise includes a default auto-correction script for the *Processes running* measure of Processes test. When a process that has been configured for monitoring stops, this script automatically executes and starts the process. To enable the auto-correction capability for the Processes test and to associate the default script with it, do the following:

1. When configuring the parameters of Processes test (see Figure 16.1), you will find a **CORRECT** flag that is set to **No** by default.

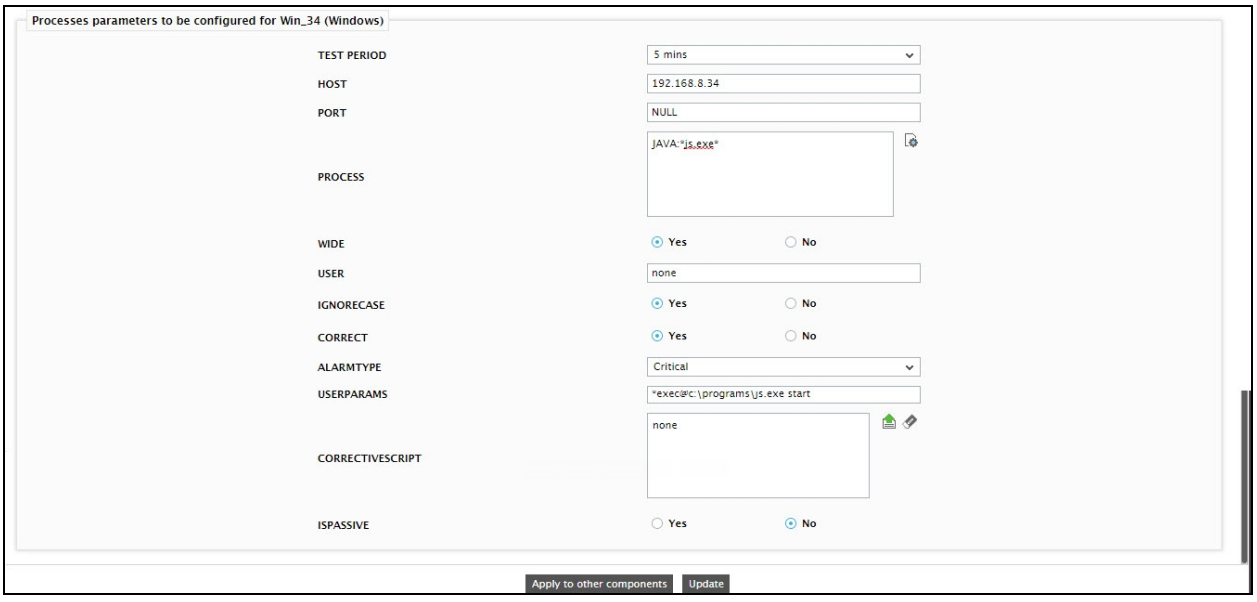

Figure 16.1: Configuring corrective scripts for Processes Test

2. To enable the auto-correction capability, select the **Yes** option against **CORRECT**.
3. Upon selecting the **Yes** option, three new parameters, namely, **ALARMTYPE, USERPARAMS, AND CORRECTIVESCRIPT** will appear. The **ALARMTYPE** parameter indicates when the auto-corrective script should execute. You can set the corrective script to execute when a specific type of alarm is generated, by selecting an option from the **ALARMTYPE** list box. For example, if the **Critical** option is chosen from the **ALARMTYPE** list box, then the corrective script will run only when a critical alarm for the Processes test is generated. Similarly, if the **Critical/Major** option is chosen, then the corrective script will execute only when the eG Enterprise system generates critical or major alarms for the Processes test. In order to ensure that the corrective script executes regardless of the alarm type, select the **Critical/Major/Minor** option.
4. The user-defined parameters that are to be passed to the corrective script are specified in the **USERPARAMS** text box. One of the following formats can be applied to the **USERPARAMS** specification:

   - *exec@processName:command*: In this specification, *processName* is the display name of the process pattern specified against the **PROCESS** parameter, and *command* is the command to be executed by the default script when the process(es) represented by the *processName* stops. For example, assume that the **PROCESS** parameter of ProcessTest has been configured in the following manner: *Apache:*/opt/egurkha/manager/apache/bin/httpd*,Tomcat:*java*tomcat**, where *Apache* and *Tomcat* are the *processNames* or display names of the configured patterns. If auto-correction is enabled for these processes, then the **USERPARAMS** specification can be as follows:
     "exec@Apache:/opt/egurkha/manager/apache/bin/apachectl start,Tomcat: /opt/tomcat/bin/catalina.sh start"
     This indicates that if the processes configured under the *processName* *"Apache"* stop (i.e. **/opt/egurkha/manager/apache/bin/httpd**), then the script will automatically execute the *command* *"/opt/egurkha/manager/apache/bin/apachectl start"* to start the processes. Similarly, if the "*Tomcat"* processes (i.e. **java*tomcat**) stop, the script will execute the *command* *"/opt/tomcat/bin/catalina.sh start"* to start the processes.

   - *"command":* In this specification, *"command"* signifies the command to be executed when any of the processes configured for monitoring, stop. Such a format best suits situations where only a single process has been configured for monitoring, or, a single command is capable of starting all the configured processes. For example, assume that the **PROCESS** parameter has been configured to monitor *IISWebSrv:*inetinfo**. Since only one process requires monitoring, the first format need not be used for configuring the **USERPARAMS**. Therefore, simplify specify the *command*, *"net start World Wide Web Publishing Service"*.
     Note:
     - The **USERPARAMS**specification should be placed within double quotes if this value includes one or more blank spaces (eg.,"*Apache:/opt/egurkha/bin/apachectl start"*).
     - If the Processes test is being configured to run on a Windows system, then in the "command" specification, the process name alone should be enclosed within single quotes. Take the case of the World Wide Web Publishing Service example above. Here, your command specification for Windows should be: *"net start 'World Wide Web Publishing Service'"*.

- Note that if a *processName* configured in the **PROCESS** parameter does not have a corresponding entry in **USERPARAMS** (as discussed in format 1), then the auto-correction capability will not be enabled for these processes.

5. Once the **USERPARAMS** are defined, specify *none* in the **CORRECTIVESCRIPT** text box to use the default auto-correction script.
6. Finally, click the **Update** button.
7. Once this is done, then the default corrective script associated with the *No of processes* measure of the *Processes* test will execute whenever the state of the measure changes to generate an alarm of the **ALARMTYPE** chosen. For instance, if the **ALARMTYPE** chosen is **Critical**, then the corrective script will execute whenever the state of the *No of processes* measure changes from Normal to Critical, Major to Critical, or Minor to Critical.

16.2 Enable Auto-Correction for Windows Services Test

The default script for WindowsServices test executes when the service that the eG agent has been configured to monitor, stops. Upon execution, the script starts the service. To enable the auto-correction capability of the WindowsServices test, do the following:

1. When configuring the Windows Services test, you will find a **CORRECT** parameter that is set to **No** by default (see Figure 16.2).

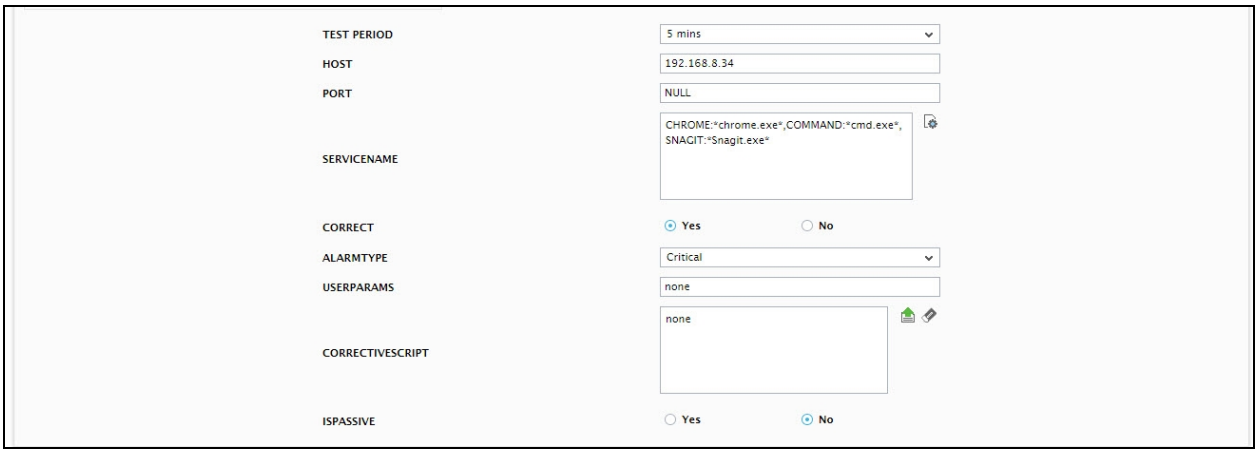

Figure 16.2: Enabling auto-correction for the Windows Services Test

2. To enable auto-correction, set the **CORRECT** parameter to **Yes** and pick an **ALARMTYPE** to indicate when the script execution should begin (see Figure 16.2).
3. The default script for Windows Services test takes no parameters. Therefore, specify “none” against **USERPARAMS** (see Figure 16.2).
4. The **CORRECTIVESCRIPT** text box can also contain none, so that the default script is automatically associated with the test (see Figure 16.2).
5. Finally, click the **Update** button (see Figure 16.2).

6. Once this is done, then the default corrective script associated with the *Service availability* measure of the *Windows Services* test will execute whenever the state of the measure changes to generate an alarm of the **ALARMTYPE** chosen. For instance, if the **ALARMTYPE** chosen is **Critical**, then the corrective script will execute whenever the state of the *Service availability* measure changes from Normal to Critical, Major to Critical, or Minor to Critical.

16.3 Building Custom Auto-Correction Scripts

Administrators can extend the eG Enterprise system's auto-correction capabilities to other tests by writing custom scripts for automatically correcting issues with those tests. You can write a single script that can resolve issues with all the measures of a test, or write individual scripts for each of the measures reported by a test. Typically, these script files can be saved the script ll execute. The extensions supported by Windows environments are: .bat, .vbs, and .exe. Scripts to be executed on Unix commonly use the extension .sh.

The process of script building essentially involves three steps:

- Writing the script for automatically correcting the problems with a test / measure
- Making the option to enable auto-correction, available for a test
- Associating the custom script with the test

16.3.1 Writing a Script

In order to enable you to understand the basic ingredients of a script, the corrective script for the Service Availability measure of Windows Services test, has been provided below.

```
@echo off
rem The following is the format of the parameters
rem 1. <User parameters> (Anything that the user gives during test configuration will be here)
rem 2. -info (constant)
rem 3. <the value of the info>

rem We get all the parameters that the test gets
rem 4. -host (constant)
rem 5. <the monitored target, this could be IP or nick)
rem 6. -port (constant)
rem 7. <port number of the monitored target, if any>
rem 8. -serviceName (constant)
rem 9. <whatever is configured for process>
rem 14 -correct
rem 15 true (This will always be true, otherwise we wouldn't have come here)
```

```
rem 16 -userparams
rem 17 <whatever the user had configured>
rem 18 -correctivescript
rem 19 WinServiceTest_Availability.bat
rem 12 -rptName (constant)
rem 13 <the rpt name>
rem 14 -egMeasHost (constant)
rem 15 <the measurement host>
rem 16 -site (constant)
rem 17 <site name>
rem 18 -egMeasVals (constant)

rem The actual measures and the state follows
rem For WinServiceTestt, the following are the measures and state
rem 19 -Availability (measure name, constant)
rem 20 <state>

set service=%3
net start %service%
```

As is evident from the sample script above, a corrective script should include the following:

- **<User parameters>:** This will store the value provided in the **USERPARAMS** text box in Figure 3.153.
- **-info <value>**: Here **-info** is a constant, and **<value>** is the variable that stores the descriptors (if any) that the test takes. A corrective script will run once for every descriptor that has been enabled for the test, based on state changes.

Besides, the script should extract the values of the test parameters provided in the test configuration page of . In case of the Windows Services Test, the parameters are as follows:

- **-host<>**: The IP of the component for which the test is being configured
- **-port<>**: The port number at which the component listens
- **-serviceName<>**: The names of the services that have been configured for monitoring
- **-correct<>**: Indicates whether corrective action has been enabled for the test or not. This will always be 'true'.
- **-userparams<>**: The parameters that have been configured by the administrator
- **-correctivescript<>**: The name of the corrective script. In the case of the WinServiceTest, to associate the default script with the test, the script name (*WinServiceTest_Availability*) is explicitly mentioned here.

In the test configuration page therefore, leaving the **CORRECTIVESCRIPT** parameter as 'none', will automatically invoke the default corrective script for the test.

- **-rptName:** The reporting name / nick name of the application for which the test is being configured
- **-eGMeasHost<>:** The measurement host
- **-site<>**: The name of site to which the measure pertains
- **-egMeasVals<>**: The constant **egMeasVals** will be followed by the actual measure values and state. For the WinServiceTest, **-Availability** is the measure name constant, which will be followed by the variable, **<state>**, that tracks the state changes.
- Finally, provide the command that will rectify the problems with the test, if any arise. Since the corrective script WinServiceTest should try to start the configured services if they stop, the command to that effect has been specified in the script file.

```
set service=%3
```

```
net start %service%
```

When an eG manager is installed on a host, a folder named **AutoCorrect** is created in the **EG_INSTALL_DIR/bin** directory. The **AutoCorrect** folder consists of some OS-specific folders (eg., folders named "Windows2003", "Linux", etc.). After building the custom script, save it to the OS-specific sub-folder that corresponds to the operating system on which the script is designed to execute.

16.3.2 Ensuring the Availability of the Option to Enable Auto-Correction for a Test

Once a corrective script for a test or a measure is built with the aforesaid contents, you need to make sure that the option to enable / disable the auto-correction capability for the test in question, is made available. In other words, you need to ensure that the test configuration page consists of the **CORRECT** parameter, which when turned on, displays the **ALARMTYPE**, **USERPARAMS**, and **CORRECTIVESCRIPT** parameters. As already stated, this option is by default available for the ProcessTest or WinServiceTest. In order to ensure its availability to any other test, the following parameters should be appended to all occurrences of the corresponding test name in the **eg_specs.ini** (in the **<EG_INSTALL_DIR>\manager\config** directory) file:

```
-correct <true/false> -alarmtype <H/M/L> -userparams none -correctivescript none
```

In the above entry, the text right next to the hyphen ('-') denotes the parameter, and the entry adjacent to the parameter signifies the default value that the parameter will take. For example, the entry '-correct' refers to the **CORRECT** parameter in Figure 16.2. The default setting for this parameter can be true or false. Indicate your choice by specifying either *true* or *false* against '-correct'.

The '-alarmtype' parameter corresponds to the **ALARMTYPE** list box in Figure 16.2, and is followed by the type of alarm (whether High/Medium/Low) that should be selected by default in the **ALARMTYPE** list box. Specifying H, M, or L against '-alarmtype' denotes that a High, Medium, or Low alarm type respectively, will be the default selection in the **ALARMTYPE** list box.

The entry '-userparams', that corresponds to the **USERPARAMS** parameter of Figure 16.2, will contain the value 'none', by default. Similarly, the '-correctivescript' parameter will take the value 'none', by default.

Finally, save the **eg_specs.ini** file to register the changes.

16.3.3 Associating the Corrective Script with the Test

To achieve this, do the following:

1. Open the test configuration page of the test for which auto-correction is to be enabled.
2. Switch auto-correction on by selecting the **TRUE** option against **CORRECT**.
3. Select an **ALARMTYPE**, and if the test takes user parameters, specify the **USERPARAMS** in the same format as discussed for Processes test.
4. Next, upload the **CORRECTIVESCRIPT** to the eG agent, so that the agent executes the script whenever the state of the test or measure changes to generate an alarm of the chosen **ALARMTYPE.** To perform the upload, click on the **Choose** button alongside the **CORRECTIVESCRIPT** text box in Figure 16.2. This will invoke a popup window as depicted by Figure 16.3.

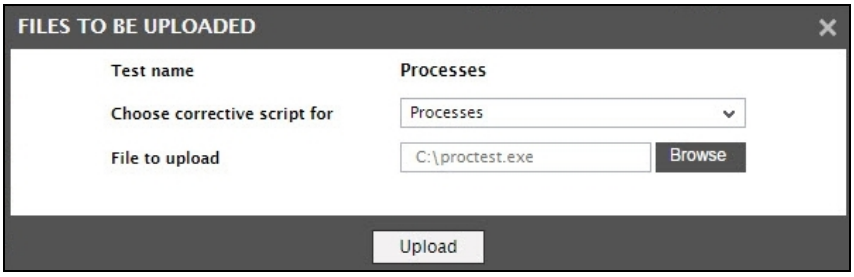

Figure 16.3: Associating the script file to be executed when the measures of the test fail

5. As you can see, Figure 16.3 displays the name of the current test. Using this window, you can associate the script with the test as a whole, or with a particular measure reported by the test. If the script resolves issues pertaining to all the measures mapped to the test, then from the **Choose corrective script for** list, select the test name. Then, in the **File to upload** text box, specify the full path to the script file. You can use the **Browse** button in Figure 16.3 to browse for the file to be uploaded. Finally, click on the **Upload** button in the popup window to upload the file. **This uploading process can be dispensed with if the script has already been uploaded**.
6. Clicking the **Upload** button leads you to the test configuration page. In the **CORRECTIVESCRIPT** text box of the page, you will find that the name of the script file has changed to reflect the corresponding *InternalTestName* - i.e., the name that eG Enterprise maintains internally for the test selected from the *Choose corrective script for* list in Figure 16.3 (see Figure 16.4).

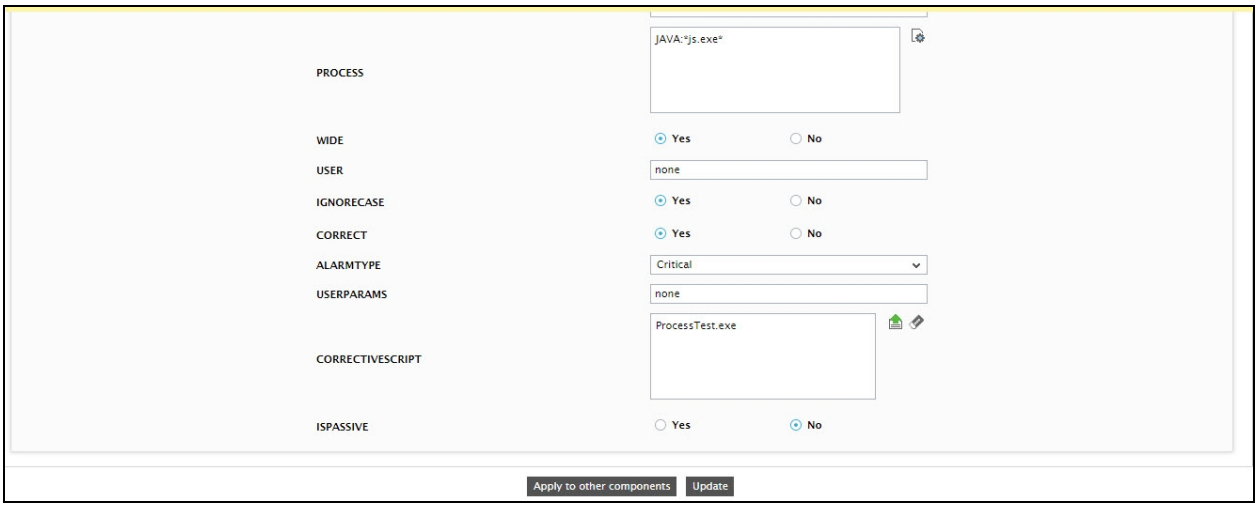

Figure 16.4: The test configuration page displaying the internal test name

7. Click the **Update** button to register the changes to the test configuration.
8. Now, if separate scripts exist for each of the measures of the test, then each of these script files should be uploaded one after another. To ensure this, first click on the **Choose** button in the test configuration page to open the pop-up window. Then, from the **Choose corrective script for** list box in the pop-up window of Figure 16.5, select a *Test:Measure* pair to which you want to associate a corrective script.

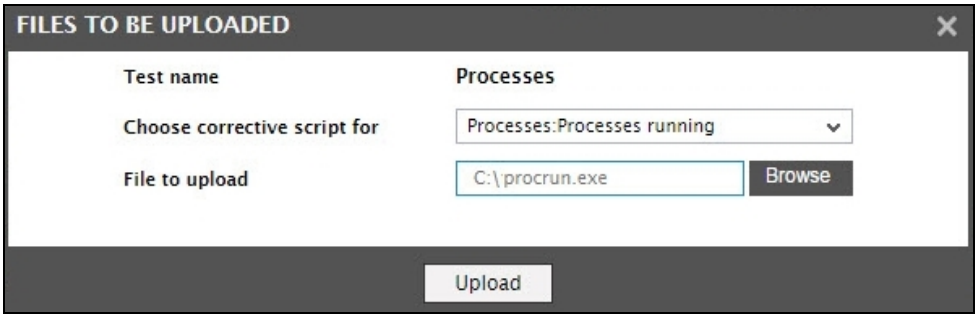

Figure 16.5: Associating a script with a particular measure of the test

9. Then, specify the full path to the corrective script that automatically corrects issues with that measure in the **File to upload** text box, and click the **Upload** button.
10. Once you return to the test configuration page, you will find that the script name in the **CORRECTIVESCRIPT** text box has changed to reflect the *InternalTestName_ InternalMeasureName* - i.e., the name that eG Enterprise maintains internally for the *Test:Measure* pair selected from the **Choose corrective script for** list in Figure 16.6.

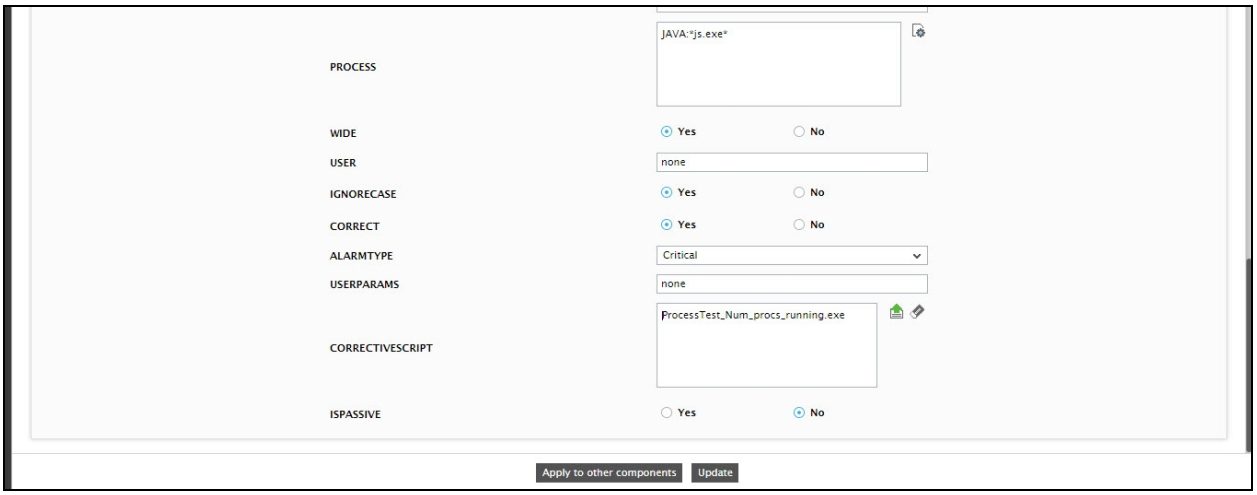

Figure 16.6: The test configuration page of Processes test

11. Finally, click the **Update** button in Figure 16.6.

Once the eG agent executing the test detects a state change in the test or any of its descriptors, it will first look for a script with the *InternalTestName*. Upon locating such a script, it will execute it to bring the state of the test back to Normal. If such a script is not available, then it will look for a script with the *InternalTestName_InternalMeasureName* and executes it.

Note:

- Corrective scripts written for external tests will be executed by the corresponding external agent only.
- The 'Auto Correction' capability is not available for tests run by a remote agent.

17

Audit Logging

An audit log can be best described as a simple log of changes, typically used for tracking temporal information. The eG manager can now be configured to create and maintain audit logs in the eG database, so that all key configuration changes to the eG Enterprise system, which have been effected via the eG user interface, are tracked.

The eG audit logs reveal critical change details such as what has changed, who did the change, and when the change occurred, so that administrators are able to quickly and accurately identify unauthorized accesses/modifications to the eG Enterprise system.

By default, audit logging is disabled. To enable the capability, follow the steps given below:

- Login to the eG administrative interface.
- Select the **Manager** option from the **Settings** tile.
- Select the **Auditing** node from the **MANAGER SETTINGS** panel.
- In the right panel, set the **Enable auditing** flag to **Yes**.
- Click the **Update** button to save the changes.

Subsequent to this, every configuration change that the user makes will be automatically logged in the database. To view the details logged and analyze their implications, eG Enterprise provides an exclusive **Audits** tile in its administrative interface, using which you can generate a variety of **AUDIT LOG REPORTS.**

The following sections deal with each one of these report types.

17.1 Auditing Successful User Logons

To view the details of a chosen user's sessions with the eG Enterprise system, use the **LOGON REPORTS**. This report enables administrators to determine which user(s) was actively using the eG Enterprise system during periods when the target environment was experiencing performance issues or exhibiting a strange behavior. Unauthorized accesses and rogue users can thus be identified quickly.

Moreover, these reports embed a special drill-down feature, which allows you a quick look at the actions performed by a particular user during the period of his/her access. This sheds light on changes effected by the user, which could have caused problems.

1. Login to the eG administrative interface as *admin*.
2. Pick the **Successful Logons** option from the **Audits** tile.
3. Figure 17.1 then appears, providing a wide range of options for report generation.

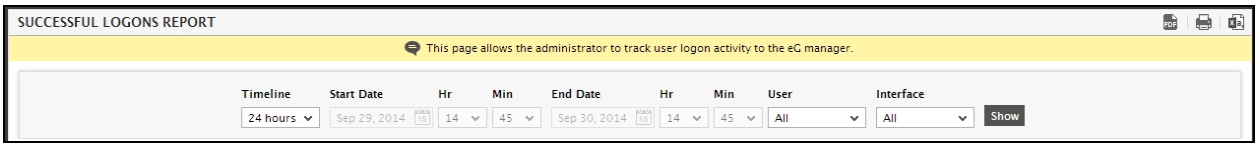

Figure 17.1: Options for generating Successful User Logon reports

4. The default **Timeline** for the report is *24 hours*. You can choose any other fixed period from the **Timeline** list, or select the **Any** option from this list (see Figure 17.2). Choosing the **Any** timeline, allows you to provide a **From** and **To** date and time for report generation. If you change the **Timeline** settings, then make sure that you click the right-arrow button at its end, to register the changes.

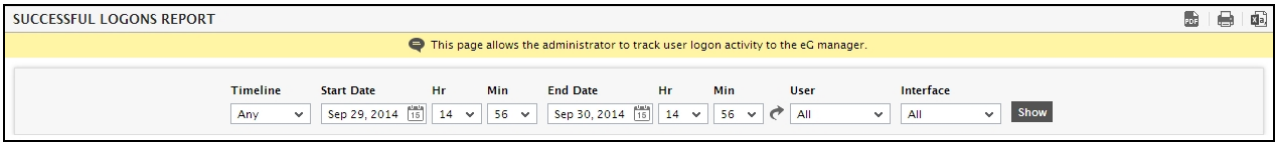

Figure 17.2: Choosing the Any timeline

5. Next, select the **User** whose accesses you want to audit. By default, the **All** option is displayed here, indicating that the report provides the details of the sessions of all users to the eG Enterprise system. However, if only one user has successfully logged into the eG Enterprise system till date, then, by default, that user's name is displayed in the **User** list.
6. Administrators can configure the target environment for monitoring by directly logging into the eG administrative interface or by using the admin command line interface provided by the eG manager. This is why, by default, the audit log not only captures user logins via the web-based eG management console, but also those logins that are performed via the eG Admin Command Line Interface. While generating audit log reports, you have the option of viewing the details of successful logins across both these interfaces, or only those that pertain to a particular interface. To indicate your choice, use the **Interface** drop-down list. The options available in the **Interface** list are as follows:
   - **Web**: Select this option to view the details of successful logins via the web-based eG management console;
   - **Command Line**: Select this option to view the details of successful logins via the admin command line interface;
   - **All**: Select this option to view the details of all successful logins, regardless of interface used.

   If required, you can choose not to maintain audit logs for activities performed via the admin command line interface by setting the **Include activities from the admin command line interface** flag in the **AUDITLOG** section of the **MANAGER SETTINGS** page to **No**. In this case therefore, the **Interface** drop-down list will not appear.
7. Finally, click the **Show** button to generate the report.

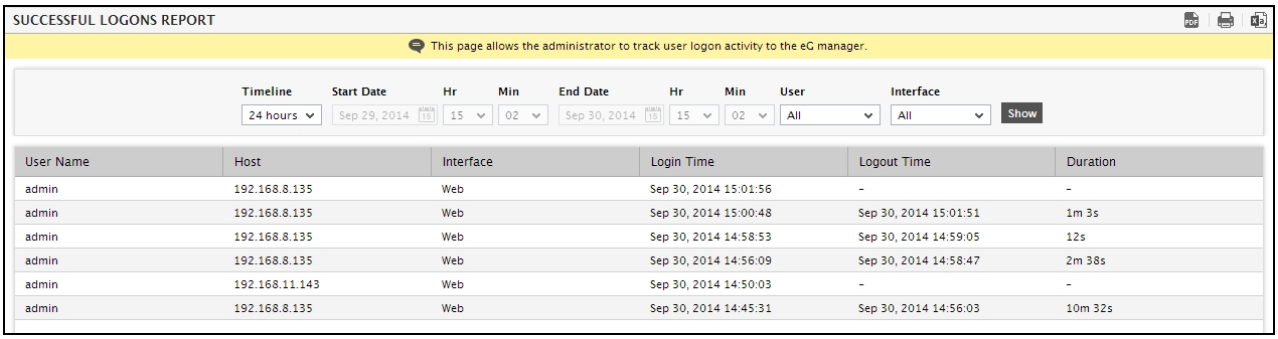

Figure 17.3: Report displaying the details of successful user logons

8. The resulting report provides details of every successful login made by the chosen user(s). These details include (see Figure 17.3):
   - the name of the user
   - the IP address of the host from which the user accessed the eG management console
   - the exact time of login
   - the accurate time of logout
   - the duration of the user access
9. If the report runs across pages, then the hyperlinked page numbers and the **First**, **Next, Prev**, and **Last** links at the bottom of the page will aid navigation.
10. You can print the report by clicking on the **Print** icon at the right, top corner of Figure 17.3, or save the report as a PDF file by clicking on the **Save** icon. You can even save the report as a CSV file by clicking on the **CSV** icon in Figure 17.3.
11. Clicking on a user name in Figure 17.3 leads you to Figure 17.4, which reveals what configuration changes were made by that user during the period of his/her access.

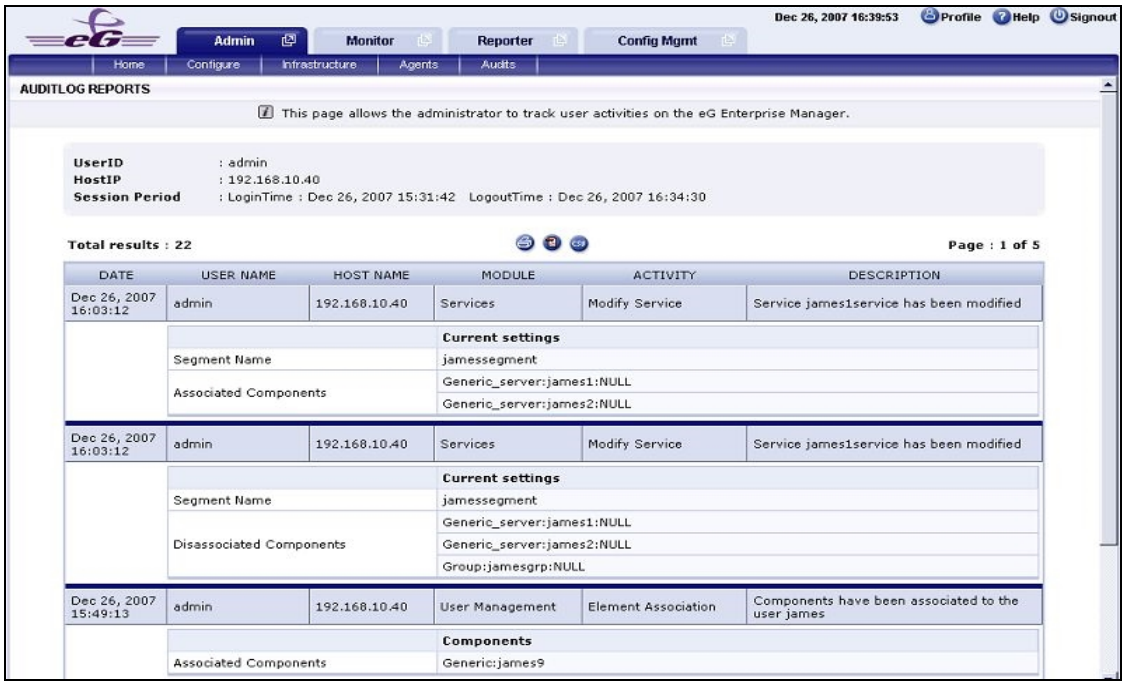

Figure 17.4: Details of changes made by a user

17.2 Auditing Failed Logons

To view the details of user logons to the eG Enterprise system that failed, use the **FAILED LOGON** reports. Using such a report, you can figure out which were the login attempts that failed and why. The reasons can bring to light problems in the network connection that need to be repaired, and even login attempts that are rather 'suspect'.

1. Login to the eG administrative interface as *admin*.
2. Select the **Failed Logons** option from the **Audits** tile.
3. The default **Timeline** for the report is *24 hours*. You can choose any other fixed period from the **Timeline** list, or select the **Any** option from this list. Choosing the **Any** timeline, allows you to provide a **From** and **To** date and time for report generation. If you change the **Timeline** settings, then make sure that you click the right-arrow button at its end, to register the changes.
4. Next, select the **User** whose login attempts you want to audit. By default, the **All** option is displayed here, indicating that the report provides the details of failed login attempts of all users to the eG Enterprise system. However, if only one user had had problems logging in till date, then, by default, that user's name is displayed in the **User** list.
5. Administrators can configure the target environment for monitoring by directly logging into the eG administrative interface or by using the admin command line interface provided by the eG manager. This is why, by default, the audit log not only captures user logins via the web-based eG management console, but also those logins that are performed via the eG Admin Command Line Interface.

While generating audit log reports, you have the option of viewing the details of failed logins across both these interfaces, or only those that pertain to a particular interface. To indicate your choice, use the **Interface** drop-down list. The options available in the **Interface** list are as follows:

- **Web**: Select this option to view the details of login failures that occurred when attempting to login via the web-based eG management console;
- **Command Line**: Select this option to view the details of login failures that occurred when attempting to login via the admin command line interface;
- **All**: Select this option to view the details of all login failures, regardless of interface used to login.

If required, you can choose not to maintain audit logs for activities performed via the admin command line interface by setting the **Include activities from the admin command line interface** flag in the **AUDITLOG** section of the **MANAGER SETTINGS** page to **No**. In this case therefore, the **Interface** drop-down list will not appear.

6. Finally, click the **Show** button to generate the report.

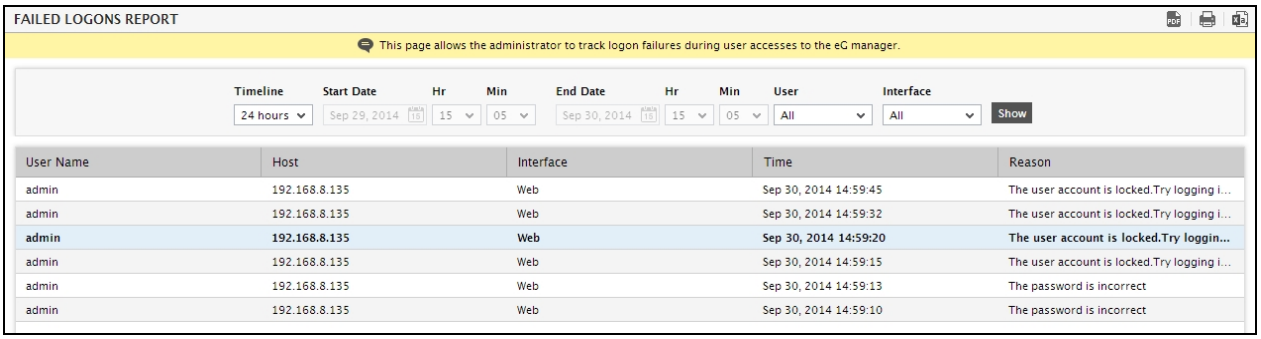

Figure 17.5: Report displaying the details of failed user logons

7. The resulting report provides details of every login made by the chosen user(s) that failed. These details include (see Figure 17.5):
   - the name of the user
   - the IP address of the host from which the user attempted to login to the eG management console
   - the Interface type that was used - whether web or command line
   - the exact time of the login attempt
   - the reason for the login failure
8. You can print the report by clicking on the **Print** icon at the right, top corner of Figure 17.5, or save the report as a PDF file by clicking on the **Save** icon. You can even save the report as a CSV file by clicking on the **CSV** icon in Figure 17.5.

17.3 Auditing Configuration Changes made using the eG Administrative Interface

eG Enterprise provides **AUDITLOG REPORTS** using which you can keep tabs on critical configuration changes made using the eG admin interface, such as password changes, test parameter changes, new server additions, threshold changes etc., which can significantly alter the way the eG Enterprise system performs

monitoring. Sometimes, these configuration changes, if not done properly or if carried out by unauthorized/unqualified personnel, can cause the eG Enterprise system to generate false alerts and perform inaccurate diagnosis.

As these **AUDITLOG REPORTS** reveal what admin settings were modified by which user, along with the details of the original settings, they greatly help administrators in quickly identifying and rectifying errors (if any) in configuration.

To generate these reports, do the following:

1. Login to the eG administrative interface as *admin*.
2. Select the **Admin** option from the **Audits** menu.
3. Figure 17.6 then appears, providing a wide range of options for report generation.

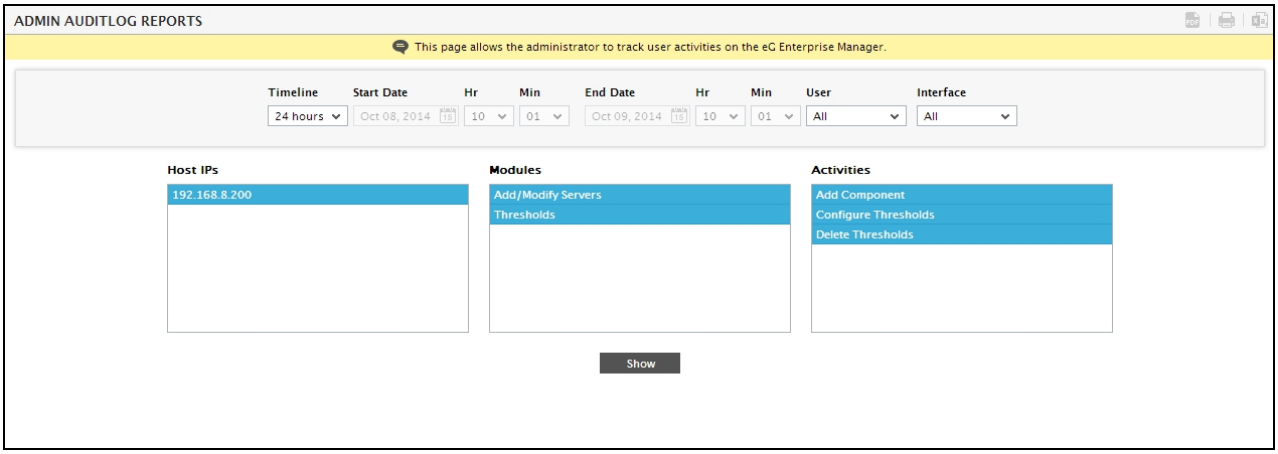

Figure 17.6: Options for generating Admin Audit Log reports

4. The default **Timeline** for the report is *24 hours*. You can choose any other fixed period from the **Timeline** list, or select the **Any** option from this list. Choosing the **Any** timeline, allows you to provide a **From** and **To** date and time for report generation. If you change the **Timeline** settings, then make sure that you click the right-arrow button at its end, to register the changes.
5. Administrators can configure the target environment for monitoring by logging into the eG administrative interface or by using the admin command line interface provided by the eG manager. This is why, by default, the audit log not only captures those configuration changes that are effected via the web-based eG administrative interface, but also logs those activities that are performed via the eG Admin Command Line Interface. While generating audit log reports, you have the option of viewing changes across both these interfaces, or only those changes that pertain to a particular interface. To indicate your choice, use the **Interface** drop-down list. The options available in the **Interface** list are as follows:
   - **Web**: Select this option to view those changes that were effected only via the web interface;
   - **Command Line**: Select this option to view those changes that were effected only via the eG command line interface;
   - **All**: Select this option to view all changes, regardless of interface.

If required, you can choose not to maintain audit logs for activities performed via the admin command line interface by setting the **Include activities from the admin command line interface** flag in the **AUDITLOG** section of the **MANAGER SETTINGS** page to **No**. In this case therefore, the **Interface** drop-down list will not appear.

6. Next, select the **User** whose admin activities you want to audit. By default, the **All** option is displayed here, indicating that the report provides the details of the configuration changes effected by all users to the eG administrative interface. However, if only one user has actively used the eG administrative interface till date, then, by default, that user's name is displayed in the **User** list.
7. The **Host IPs** list displays all the IP addresses from which the chosen user(s) has accessed the eG administrative interface (see Figure 17.6). If you are looking for information on the admin accesses from specific IPs, select those IP addresses alone from the **Host IPs** list.
8. After the selection, the **Modules** list will be populated with those admin modules that the chosen user(s) worked with while accessing the eG admin interface from the selected **Host IPs** (see Figure 17.6). If you want the details of changes that the user made in specific admin modules, select those modules alone from the **Modules** list.
9. Based on the **Modules** selection, the **Activities** list will be populated. While working with the eG admin interface, the selected user(s) might have performed a few specific operations on the chosen **Modules**. eG Enterprise automatically discovers the operations that correspond to the chosen user-host IP-module combination from the audit logs, and populates the **Activities** list with the operations so discovered (see Figure 17.6). If you want the details of specific activities only, select the required options alone from the **Activities** list.
10. Finally, click the **Show** button to generate the report.

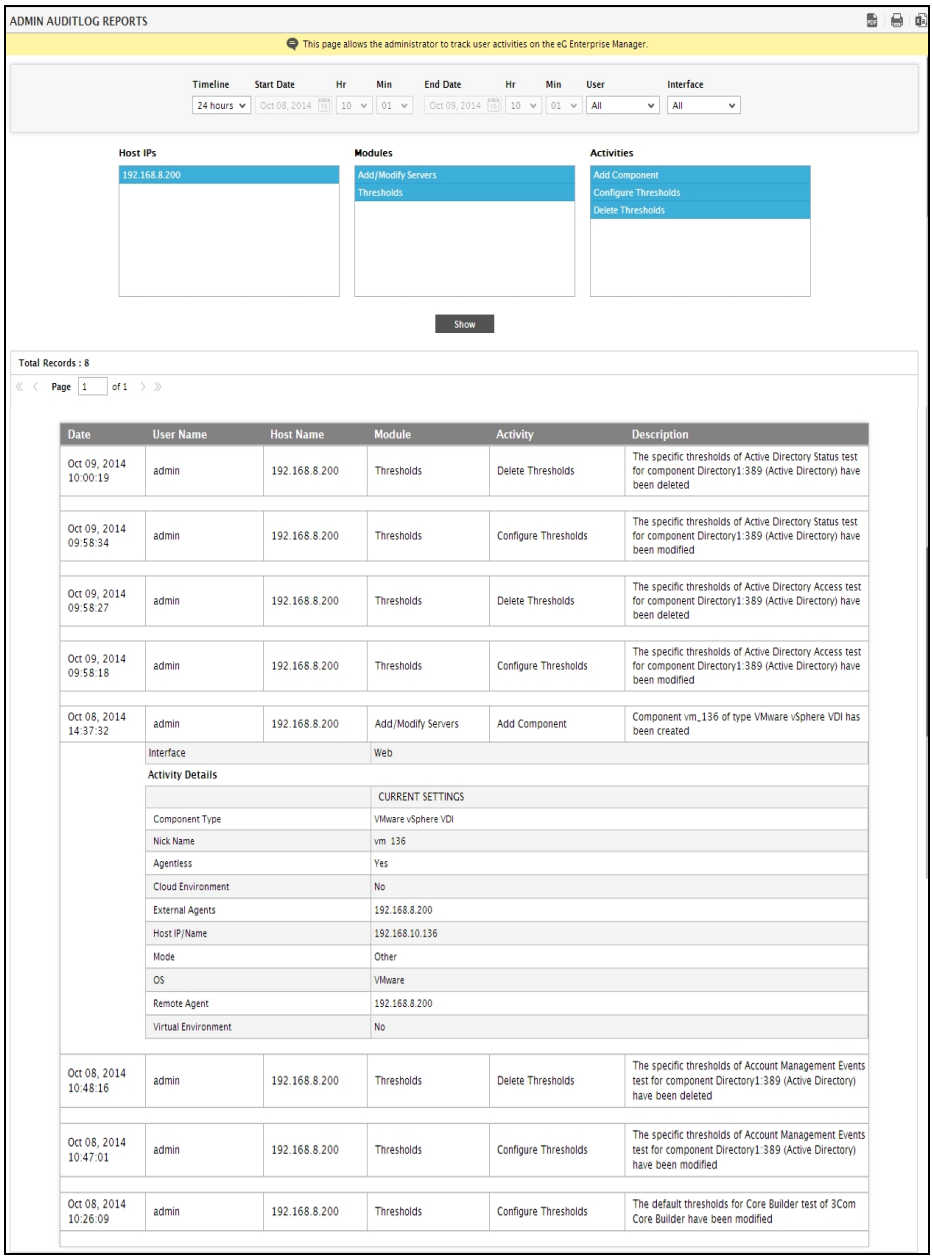

Figure 17.7: An ADMIN AUDITLOG REPORT

11. The resulting report provides the following details (see Figure 17.7):

- the date/time of the change
- the name of the user who made the change
- the IP address of the host from which the user accessed the eG admin interface
- the module that was accessed by the user
- the specific operation/activity that was performed by the user on that module
- the Interface type that was used - whether web or command line

- the detailed description of the change, followed by a snapshot of the settings prior to change, and the settings after the change; if a configuration has been newly introduced (for eg., a server has been newly managed), then only the **Current Settings** will be displayed

  Note:

  By default, every change record that the report displays will be accompanied by the **Current** and **Previous** configuration settings. This can sometimes clutter the report view, making it difficult for you to read and analyze the report. You can therefore hide the both these columns from the report, by setting the **ShowChanges** parameter in the **[AUDIT_LOG_SETTINGS]** section of the **eg_ui.ini** file to **No**.

12. You can print the report by clicking on the **Print** icon at the right, top corner of Figure 17.7 or save the report as a PDF file by clicking on the **Save** icon. You can even save the report as a CSV file by clicking on the **CSV** icon in Figure 17.7.

    Note:

    If the threshold for a measure is changed to -/-/-, then the **Current Settings** column of the **Audit Log Report** will indicate that the threshold has changed to *none*.

17.4 Auditing Configuration Changes made using the eG Monitor Interface

Just like changes made using the eG admin interface, care should also be taken while making changes using the eG monitor interface - eg., while deleting alarms, acknowledging alarms, configuring quick insight/live graph views, etc. Changes that are implemented carelessly can only add to an administrator's confusion, and cause unnecessary delay in problem resolution.

eG Enterprise therefore provides **AUDITLOG REPORTS** that enable administrators to track user activities on the eG monitoring console, and to accurately detect changes wrongly made and the user responsible for the same.

To generate these reports, do the following:

1. Login to the eG administrative interface as *admin*.
2. Select the **Monitor** option from the **Audits** tile.
3. Figure 17.8 then appears, providing a wide range of options for report generation.
4. The default **Timeline** for the report is *24 hours*. You can choose any other fixed period from the **Timeline** list, or select the **Any** option from this list. Choosing the **Any** timeline, allows you to provide a **From** and **To** date and time for report generation. If you change the **Timeline** settings, then make sure that you click the right-arrow button at its end, to register the changes.
5. Next, select the **User** whose monitoring activities you want to audit. By default, the **All** option is displayed here, indicating that the report provides the details of the configuration changes effected by all users to the eG monitoring console. However, if only one user has actively used the eG monitor interface till date, then, by default, that user's name is displayed in the **User** list.

Administrators can configure the target environment for monitoring by logging into the eG administrative interface or by using the admin command line interface provided by the eG manager. This is why, by default, the audit log not only captures those configuration changes that are effected via the web-based eG administrative interface, but also logs those activities that are performed via the eG Admin Command Line Interface. While generating audit log reports, you have the option of viewing the changes across both these interfaces, or only those changes that pertain to a particular interface. To indicate your choice, use the **Interface** drop-down list in Figure 17.8. The options available in the **Interface** list are as follows:

- **Web**: Select this option to view those changes that were effected only via the web interface;
- **Command Line**: Select this option to view those changes that were effeced only via the eG command line interface;
- **All**: Select this option to view all changes, regardless of interface.

If required, you can choose not to maintain audit logs for activities performed via the admin command line interface by setting the **Include activities from the admin command line interface** flag in the **AUDITLOG** section of the **MANAGER SETTINGS** page to **No**. In this case therefore, the **Interface** drop-down list will not appear.

Note:

The eG command line interface can currently be used only for administering the eG manager - i.e., for performing a few administrative tasks such as adding/managing components, configuring external agents/remote agents, assigning agents to secondary manager in a redundant manager setup, etc. Hence, the **Interface** option is currently relevant to the Admin Audit log Reports, and not the Monitor, Reporter, and Configuration Management Audit Log Reports.

6. The **Host IPs** list displays all the IP addresses from which the chosen user(s) has accessed the eG monitor interface (see Figure 17.8). If you are looking for information on the monitor accesses from specific IPs, select those IP addresses alone from the **Host IPs** list.
7. Once the **Host IPs** are chosen, the **Modules** list will be populated with those monitor modules that the chosen user(s) has worked with while accessing the eG monitor interface from the selected **Host IPs** (see Figure 17.8). If you want the details of changes that the user made in specific monitor modules, select those modules alone from the **Modules** list.
8. Based on the **Modules** selection, the **Activities** list will be populated. While working with the eG monitor interface, the selected user(s) might have performed a few specific operations on the chosen **Modules**. eG Enterprise automatically discovers the operations that correspond to the chosen user-host IP-module combination from the audit logs, and populates the **Activities** list with the operations so discovered (see Figure 17.8). If you want the details of specific activities only, select the required options alone from the **Activities** list.
9. Finally, click the **Show** button to generate the report.

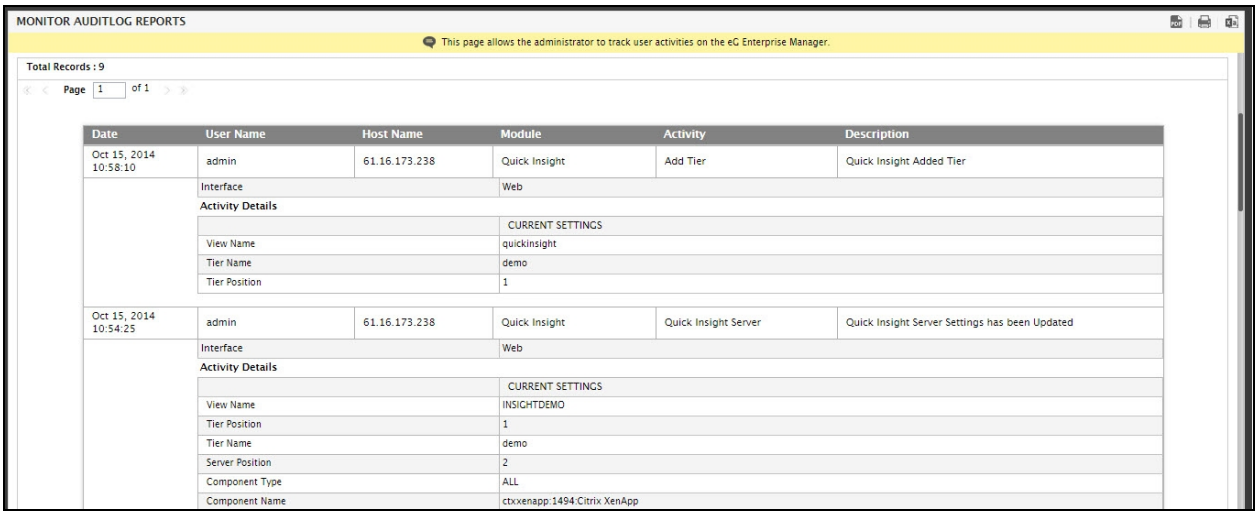

Figure 17.8: Report displaying the details of changes to the eG monitor modules

10. The resulting report provides the following details (see Figure 17.8):
    - The date/time of the change
    - the name of the user who made the change
    - the IP address of the host from which the user accessed the eG monitor interface
    - the module that was accessed by the user
    - the specific operation/activity that was performed by the user on that module
    - the Interface type that was used - whether web or command line
    - the detailed description of the change, followed by a snapshot of the settings prior to change, and the settings after the change; if a configuration has been newly introduced (for eg., quick insight view was newly created), then only the **Current Settings** will be displayed

      Note:

      By default, every change record that the report displays will be accompanied by the **Current** and **Previous** configuration settings. This can sometimes clutter the report view, making it difficult for you to read and analyze the report. You can therefore hide both these columns from the report, by setting the **ShowChanges** parameter in the **[AUDIT_LOG_SETTINGS]** section of the **eg_ui.ini** file to **No**.
11. You can print the report by clicking on the **Print** icon at the right, top corner of Figure 17.8, or save the report as a PDF file by clicking on the **Save** icon. You can even save the report as a CSV file by clicking on the **CSV** icon in Figure 17.8.

17.5 Auditing Configuration Changes made using the eG Reporter Interface

Typically, the key configuration changes that a user can make using the eG Reporter component is to add/modify/remove **FAVORITES** and **SCHEDULE** report configurations. To track the related changes, use the **AUDITLOG REPORTS** that eG Enterprise provides exclusively for eG Reporter.

To generate these reports, do the following:

1. Login to the eG administrative interface as *admin*.
2. Select the **Reports** option from the **Audits** tile
3. Figure 17.6 then appears, providing a wide range of options for report generation.
4. The default **Timeline** for the report is *24 hours*. You can choose any other fixed period from the **Timeline** list, or select the **Any** option from this list. Choosing the **Any** timeline, allows you to provide a **From** and **To** date and time for report generation. If you change the **Timeline** settings, then make sure that you click the right-arrow button at its end, to register the changes.
5. Next, select the **User** whose eG Reporter-related activities you want to audit. By default, the **All** option is displayed here, indicating that the report provides the details of the configuration changes effected by all users to the eG Reporter. However, if only one user has actively used the eG Reporter till date, then, by default, that user's name is displayed in the **User** list.

   Administrators can configure the target environment for monitoring by logging into the eG administrative interface or by using the admin command line interface provided by the eG manager. This is why, by default, the audit log not only captures those configuration changes that are effected via the web-based eG administrative interface, but also logs those activities that are performed via the eG Admin Command Line Interface. While generating audit log reports, you have the option of viewing the changes across both these interfaces, or only those changes that pertain to a particular interface. To indicate your choice, use the **Interface** drop-down list. The options available in the **Interface** list are as follows:

   - **Web**: Select this option to view those changes that were effected only via the web interface;
   - **Command Line**: Select this option to view those changes that were effected only via the eG command line interface;
   - **All**: Select this option to view all changes, regardless of interface.

   If required, you can choose not to maintain audit logs for activities performed via the admin command line interface by setting the **Include activities from the admin command line interface** flag in the **AUDITLOG** section of the **MANAGER SETTINGS** page to **No**. In this case therefore, the **Interface** drop-down list will not appear.

   Note:

   The eG command line interface can currently be used only for administering the eG manager - i.e., for performing a few administrative tasks such as adding/managing components, configuring external agents/remote agents, assigning agents to secondary manager in a redundant manager setup, etc. Hence, the **Interface** option is currently relevant to the Admin Audit log Reports, and not the Monitor, Reporter, and Configuration Management Audit Log Reports.
6. Then, indicate the column by which the **AUDITLOG REPORT** should be sorted. If the **Time** chosen from the **Sort by** list, then the resulting report will be sorted in the descending order of the event time. If any of the other options are chosen from this list box, then the report will be sorted in the ascending order of the values displayed in the chosen column.
7. The **Host IPs** list displays all the IP addresses from which the chosen user(s) has accessed the eG Reporter interface (see Figure 17.9). If you are looking for information on the accesses to eG Reporter from specific IPs, select those IP addresses alone from the **Host IPs** list.

8. Once one/more **Host IPs** are chosen, the **Modules** list will be populated with those Reporter modules that the chosen user(s) has worked with while accessing the eG Reporter interface from the selected **Host IPs** (see Figure 17.9). If you want the details of changes that the user made in specific Reporter modules only, select those modules alone from the **Modules** list.
9. Based on the selections from the **Modules** list, the **Activities** list will be populated. While working with the eG Reporter interface, the selected user(s) might have performed a few specific operations on the chosen **Modules**. eG Enterprise automatically discovers the operations that correspond to the chosen user-host IP-module combination from the audit logs, and populates the **Activities** list with the operations so discovered (see Figure 17.9). If you want the details of specific activities only, select the required options alone from the **Activities** list.
10. Finally, click the **Show** button to generate the report.

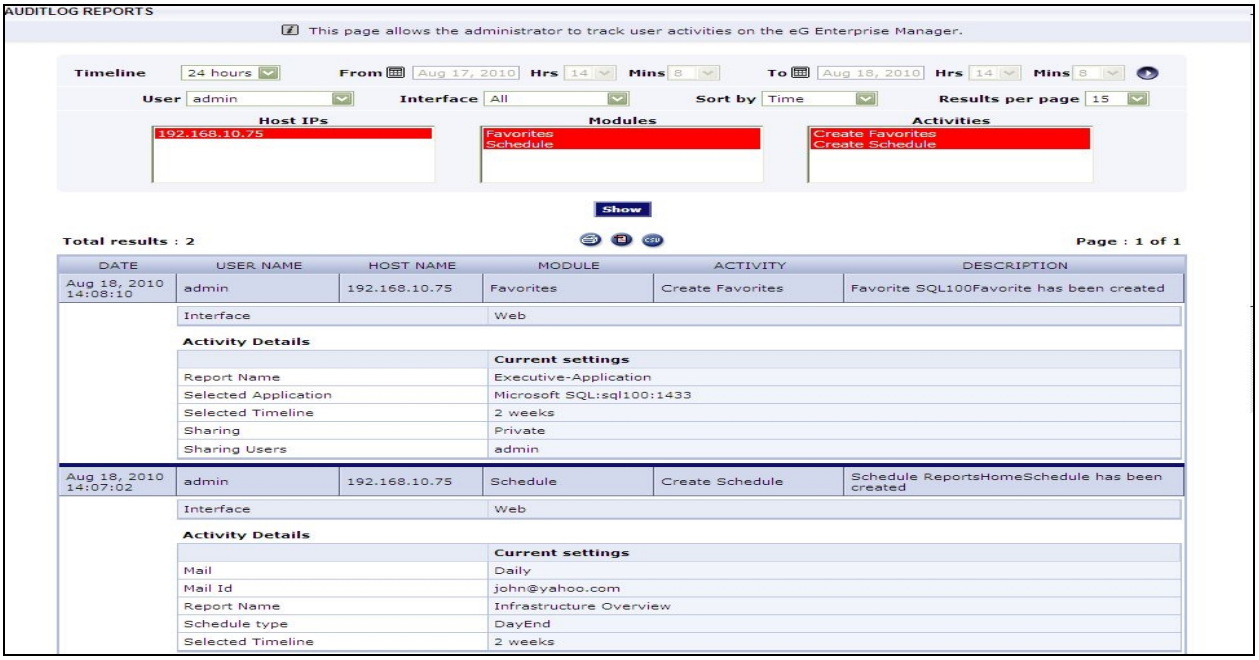

Figure 17.9: Report displaying the details changes made using the eG Reporter interface

11. The resulting report provides the following details (see Figure 17.9):
    - The date/time of the change
    - the name of the user who made the change
    - the IP address of the host from which the user accessed the eG Reporter interface
    - the module that was accessed by the user
    - the specific operation/activity that was performed by the user on that module
    - the Interface type that was used - whether web or command line
    - the detailed description of the change, followed by a snapshot of the settings prior to change, and the settings after the change; if a configuration has been newly introduced (for eg., a new schedule was created), then only the **Current Settings** will be displayed

Note:

By default, the every change record that the report displays will be accompanied by the **Current** and **Previous** configuration settings. This can sometimes clutter the report view, making it difficult for you to read and analyze the report. You can therefore hide both these columns from the report, by setting the **ShowChanges** parameter in the **[AUDIT_LOG_SETTINGS]** section of the **eg_ui.ini** file to **No**.

12. You can print the report by clicking on the **Print** icon at the right, top corner of Figure 17.9, or save the report as a PDF file by clicking on the **Save** icon. You can even save the report as a CSV file by clicking on the **CSV** icon in Figure 17.9.

Note:

In a redundant setup, all the audit log reports discussed above will have an additional **MANAGER NAME** column, which displays the IP or host name of the manager to which a record pertains.

17.6 Auditing the Display Settings Changed Using the eG Configuration Management Interface

Using the **Config** option of the **Audits** menu, you can generate audit log reports that will help you instantly identify whether any changes were made to the dashboard and overall display settings of the eG Configuration Management interface, who made these change, and when.

To generate the **Configuration Management** related audit log reports, do the following:

1. If the **Configuration** option is chosen from the **Audits** tile, Figure 17.10 will appear, using which you can build the report specifications.
2. The default **Timeline** for the report is *24 hours*. You can choose any other fixed period from the **Timeline** list, or select the **Any** option from this list. Choosing the **Any** timeline, allows you to provide a **From** and **To** date and time for report generation. If you change the **Timeline** settings, then make sure that you click the right-arrow button at its end, to register the changes.
3. Next, select the **User** whose Config Management-related activities you want to audit. By default, the **All** option is displayed here, indicating that the report provides the details of the configuration changes effected by all users to the eG Configuration Management interface. However, if only one user has actively used the eG Configuration Management interface till date, then, by default, that user's name is displayed in the **User** list.
4. Administrators can configure the target environment for monitoring by logging into the eG administrative interface or by using the admin command line interface provided by the eG manager. This is why, by default, the audit log not only captures those configuration changes that are effected via the web-based eG administrative interface, but also logs those activities that are performed via the eG Admin Command Line Interface.
5. While generating audit log reports, you have the option of viewing the changes across both these interfaces, or only those changes that pertain to a particular interface. To indicate your choice, use the **Interface** drop-down list. The options available in the **Interface** list are as follows:

- **Web**: Select this option to view those changes that were effected only via the web interface;
- **Command Line**: Select this option to view those changes that were effected only via the eG command line interface;
- **All**: Select this option to view all changes, regardless of interface.

If required, you can choose not to maintain audit logs for activities performed via the admin command line interface by setting the **Include activities from the admin command line interface** flag in the **AUDITLOG** section of the **MANAGER SETTINGS** page to **No**. In this case therefore, the **Interface** drop-down list will not appear.

Note:

The eG command line interface can currently be used only for administering the eG manager - i.e., for performing a few administrative tasks such as adding/managing components, configuring external agents/remote agents, assigning agents to secondary manager in a redundant manager setup, etc. Hence, the **Interface** option is currently relevant to the Admin Audit log Reports, and not the Monitor, Reporter, and Configuration Management Audit Log Reports

6. The **Host IPs** list displays all the IP addresses from which the chosen user(s) has accessed the eG Configuration Management interface. If you are looking for information on the accesses from specific IPs, select those IP addresses alone from the **Host IPs** list.
7. After the **Host IPs** selection, the **Modules** list will be populated with either/all of the following options: **Common Display settings** and **Dashboard settings** (see Figure 17.10). The options displayed depend upon which of the two modules were accessed by the chosen user from the selected **Host IPs**. Select either/both the displayed modules to view the changes made by the user in the respective modules.
8. Based on the selections from the **Modules** list, the **Activities** list will be populated. While working with the eG Configuration Management interface, the selected user(s) might have performed a few specific operations on the chosen **Modules**. eG Enterprise automatically discovers the operations that correspond to the chosen user-host IP-module combination from the audit logs, and populates the **Activities** list with the operations so discovered (see Figure 17.10). If you want the details of specific activities only, select the required options alone from the **Activities** list.
9. Finally, click the **Show** button to generate the report.

DATE	USER NAME	HOST NAME	MODULE	ACTIVITY	DESCRIPTION
Aug 18, 2010 12:12:17	admin	192.168.10.77	Dashboard settings	Dashboard customization	Dashboard settings has been changed

Interface	Web

Activity Details

	Current settings
Machine Distribution	By Operating system
Availability status	SoftwareList
Total count in Availability status	10
Daywise Change Distribution	1 day
Changes At-A-Glance	24 hours
Change Summary	48 hours
View By	Zone

Figure 17.10: Report displaying the details of display settings changed using the eG Configuration Management interface

10. The resulting report provides the following details (see Figure 17.10):
    - The date/time of the change
    - the name of the user who made the change
    - the IP address of the host from which the user accessed the eG Reporter interface
    - the module that was accessed by the user
    - the specific operation/activity that was performed by the user on that module
    - the Interface type that was used - whether web or command line
    - the detailed description of the change, followed by a snapshot of the settings prior to change, and the settings after the change; if a configuration has been newly introduced (for eg., a new schedule was created), then only the **Current Settings** will be displayed

      Note:

      By default, the every change record that the report displays will be accompanied by the **Current** and **Previous** configuration settings. This can sometimes clutter the report view, making it difficult for you to read and analyze the report. You can therefore hide both these columns from the report, by setting the **ShowChanges** parameter in the **[AUDIT_LOG_SETTINGS]** section of the **eg_ui.ini** file to **No**.

11. You can print the report by clicking on the **Print** icon at the right, top corner of Figure 17.10, or save the report as a PDF file by clicking on the **Save** icon. You can even save the report as a CSV file by clicking on the **CSV** icon in Figure 17.10.

    Note:

    In a redundant setup, all the audit log reports discussed above will have an additional **MANAGER NAME** column, which displays the IP or host name of the manager to which a record pertains.

    Note:

    - Concurrent updates to the eG Enterprise configuration could have a malicious effect. So the eG manager tracks admin user sessions. It times out a session after 30 minutes of inactivity. An alert is also displayed whenever the manager deletes multiple simultaneous logins.
    - By default, the title bar of the administrative interface, will display the IP address of the eG manager. However, if you wish to custom define the text on the title bar, then do the following:
      - Open the **eg_services.ini** file in the **<EG_HOME_DIR>/manager/config** directory.
      - Move to the **[MISC_ARGS]** section within and specify a title bar text of your choice against the **ManagerTitle** parameter. Doing so, will ensure that the defined text appears on the title bar of the administrative interface.

18

Integration With Third Party Monitoring Solutions

In some environments, the eG Enterprise Suite may co-exist with one/more third-party monitoring solutions, so that users can avail the benefits of the accurate root-cause identification, intelligent thresholding, and proactive alerting abilities of the eG Enterprise Suite along with the specialized capabilities offered by the other monitoring solutions. Typically, administrators of such environments may prefer not to login to every console - one each for every monitoring solution in use - and receive the performance perspective offered by each solution. This is because, besides being unwieldy, these consoles often present a wealth of performance and problem data, which will first have to be consolidated and then manually correlated and analyzed by the administrators, for any actionable information to emerge. This could take hours, even days!

Instead, administrators would want to login to a single console, and quickly launch the management console of any other monitoring solution from within that console based on need, so that they can receive a clearer picture of performance, and accurately detect bottlenecks and what is causing them. For instance, you could use a Network Management solution to determine which network devices are experiencing performance issues, and can instantly launch the eG monitoring console from within the network management console to identify what is causing the problems. To facilitate such a seamless integration between monitoring solutions, eG Enterprise offers

- You can configure third-party monitoring solutions with quick access to specific pages of the eG monitoring console, so that the eG console can be launched from the third-party console in a click of a button;
- You can configure the eG Enterprise system with URLs of relevant areas of the third-party monitoring console, so that you can quickly launch the third-party console from the eG monitoring console for deeper diagnosis;

This section discusses each of the above-mentioned options in great detail.

18.1 Quickly Launching the eG Monitoring Console from a Third-party Monitoring Console

By default, you cannot launch the eG monitoring console from any other thid-party console. To enable this capability, do the following:

1. Edit the **eg_services.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
2. Set the **quick_lauch** flag in the **[QUICK_LAUNCH]** section of the file to **true** (default: **false**).
3. Finally, save the file.
4. Once this is done, you can proceed to configure the third-party solution with the URLs that allow quick access to specific pages of the eG monitoring console. The URL should be of the following format:

http://<eGManagerIP>:<eGManagerPort>/final/monitor/egQuickLaunch.jsp?user=<UserName>&pass=<Password>&access=<NameofWebPage>&name=<ElementName>

As you can see, the URL supports a **user** and **pass** (i.e., password) parameter, which will have to be configured with the credentials of a user with rights to access the eG monitoring console. The **access** parameter has to be configured with the name of the specific web page (in the eG monitoring console) or a pointer to the specific element (i.e., a component, segment, service, or zone) that is to be launched from the third-party console, and the **name** parameter should be configured with the name of the infrastructure element of interest.

5. For now, users to third-party monitoring consoles can launch only the following pages of the eG monitoring console:
   - The **COMPONENT SERVER** page listing all servers assigned to a specific user;
   - The **COMPONENT SERVER** page listing all components of a specific component type;
   - The **SYSTEMS LIST** page listing all systems assigned to a specific user or a specific system;
   - The **ZONE LIST** page
   - The **Zone Dashboard** of a specific zone
   - The **SEGMENT LIST** page
   - The **SEGMENT TOPOLOGY** page of a specific segment
   - The **SERVICE LIST** page
   - The **SERVICE TOPOLOGY** page of a specific service
   - The **LAYER MODEL** page displaying the layer model of a specific component;
   - The **LAYER MODEL** page displaying the measurements reported by a specific test mapped to a particular layer of a component;
   - The **LAYER MODEL** page displaying the measurements reported by a specific descriptor supported by a test mapped to a particular layer of a component;
   - The **LAYER MODEL** page displaying the tests mapped to a particular layer of a component;
6. Given below are sample URLs for accessing each of the above mentioned pages. These samples conform to the following specifications:
   - eG Manager IP: **192.168.10.64**
   - eG Manager Port: **7077**
   - User name: **admin**
   - Password: **admin**
   - Zone to be accessed: **zone1**
   - Segment to be accessed: **seg_new**
   - Service to be accessed: **myservice**
   - Component to be accessed:

- Internal name of Component type: **Oracle_server**
- Component Name: **oradb:1521:egurkha**
- System Name: **win32**

- To access the **COMPONENT SERVER** page listing all components assigned to user **admin**:

  http://192.168.10.64:7077/final/monitor/egQuickLaunch.jsp?user=admin&pass=admin&access=componentList&name=All

  (OR)

  http://192.168.10.64:7077/final/monitor/egQuickLaunch.jsp?user=admin&pass=admin&access=component&name=All

- To access the **SYSTEMS LIST** page listing all systems monitored by user **admin** and the components managed on each system:

  http://192.168.10.64:7077/final/monitor/egQuickLaunch.jsp?user=admin&pass=admin&access=systemList&name=All

- To view the state of the system, **win32**, in the **SYSTEMS LIST** page and the components managed on that system:

  http://192.168.10.64:7077/final/monitor/egQuickLaunch.jsp?user=admin&pass=admin&access=systemList&name=win32

- To view the current state of all components of type **Oracle** in the **COMPONENT SERVER** page:

  http://192.168.10.64:7077/final/monitor/egQuickLaunch.jsp?user=admin&pass=admin&access=componentList&name=Oracle_server

  Make sure the **name** parameter is configured with the **internal name** of the component type.

- To access the **ZONE LIST** page:

  http://192.168.10.64:7077/final/monitor/egQuickLaunch.jsp?user=admin&pass=admin&access=zoneList&name=All

- To access the **SEGMENT LIST** page:

  http://192.168.10.64:7077/final/monitor/egQuickLaunch.jsp?user=admin&pass=admin&access=segment&name=All

- To access the **SERVICE LIST** page:

  http://192.168.10.64:7077/final/monitor/egQuickLaunch.jsp?user=admin&pass=admin&access=service&name=All

- To access the dashboard of **zone1**:

  http://192.168.10.64:7077/final/monitor/egQuickLaunch.jsp?user=admin&pass=admin&access=zone&name=zone1

- To access the topology page of the segment, **seg_new**:

http://192.168.10.64:7077/final/monitor/egQuickLaunch.jsp?user=admin&pass=admin&access=segment&name=seg_new

- To access the topology page of the service, **myservice**:

  http://192.168.10.64:7077/final/monitor/egQuickLaunch.jsp?user=admin&pass=admin&access=service&name=myservice

- To access the layer model of the Oracle database server, **oradb:1521:egurkha**:

  http://192.168.10.64:7077/final/monitor/egQuickLaunch.jsp?user=admin&pass=admin&access=component&name=Oracle_server:oradb:1521:egurkha

  (OR)

  http://192.168.10.64:7077/final/monitor/egQuickLaunch.jsp?user=admin&pass=admin&access=system&name=Oracle_server:oradb:1521:egurkha

- To access only the host-level layers, tests, and measurements of the Oracle database server, **oradb:1521:egurkha**:

  http://192.168.10.64:7077/final/monitor/egQuickLaunch.jsp?user=admin&pass=admin&access=component&name=Host_system:oradb:1521:egurkha

  (OR)

  http://192.168.10.64:7077/final/monitor/egQuickLaunch.jsp?user=admin&pass=admin&access=system&name=Host_system:oradb:1521:egurkha

  Wherever the **port number** is not applicable, use **NULL**. For instance, to acces the layer model of the **Windows_server**, **win32**:

  http://192.168.10.64:7077/final/monitor/egQuickLaunch.jsp?user=admin&pass=admin&access=component&name=Windows_server:win32:NULL

  (OR)

  http://192.168.10.64:7077/final/monitor/egQuickLaunch.jsp?user=admin&pass=admin&access=system&name=Windows_server:win32:NULL

- To access the metrics reported by the **Oracle SQL Network** test mapped to the **SQL Net** layer of the Oracle database server, **oradb:1521:egurkha**:

  http://192.168.8.64:7077/final/monitor/egQuickLaunch.jsp?user=admin&pass=admin&access=component&name=Oracle_server: oradb:1521:egurkha&layer=SQLNET&test=OraSqlNetTest

  (OR)

  http://192.168.8.64:7077/final/monitor/egQuickLaunch.jsp?user=admin&pass=admin&access=system&name=Oracle_server: oradb:1521:egurkha&layer=SQLNET&test=OraSqlNetTest

- To access the metrics reported by the **Oracle Tablespaces** test for the **ORDERS** tablespace (descriptor) of the Oracle database server, **oradb:1521:egurkha**:

http://192.168.8.64:7077/final/monitor/egQuickLaunch.jsp?user=admin&pass=admin&access=component&name=Oracle_ server: oradb:1521:egurkha&layer=TABLESPACES&test=OraSqlNetTest&descriptor=ORDERS

(OR)

http://192.168.8.64:7077/final/monitor/egQuickLaunch.jsp?user=admin&pass=admin&access=system&name=Oracle_ server: oradb:1521:egurkha&layer=TABLESPACES&test=OraSqlNetTest&descriptor=ORDERS

- To access the tests mapped to the **Memory Structures** layer of the Oracle database server, **oradb:1521:egurkha**:

http://192.168.8.64:7077/final/monitor/egQuickLaunch.jsp?user=admin&pass=admin&access=component&name=Oracle_server: oradb:1521:egurkha&layer=MEMORY_STRUCTURES

(OR)

http://192.168.8.64:7077/final/monitor/egQuickLaunch.jsp?user=admin&pass=admin&access=system&name=Oracle_server: oradb:1521:egurkha&layer= MEMORY_STRUCTURES

18.1.0.1 Frequently Asked Questions

- **What will happen if I try to launch the COMPONENT SERVER page with the credentials of a user who does not have access to the COMPONENT SERVER page?**

  When this URL is launched, the resulting page will indicate that the configured user does not have access to the page. Such a message will also appear while trying to launch a a **SERVICE LIST**, **ZONE LIST**, or **SEGMENT LIST** page with the credentials of a user who does not have access to these pages.

- **What will happen if I try to launch the topology page of the segment, seg1, with the credentials of a user who does not have the rights to monitor seg1?**

  When this URL is launched, the **SEGMENT LIST** page will appear, instead of the topology page of the segment **seg1**.

- **What will happen if I try to launch the topology page of the service, new_service, with the credentials of a user who does not have the rights to monitor new_service?**

  When this URL is launched, the **SERVICE LIST** page will appear, instead of the topology page of the service **new_service**.

  Similarly, if you try to launch the layer model page of a specific component or the dashboard of a specific zone, but the user specified in the URL does not have access to the configured component/zone, then the **COMPONENT SERVER** page or the **ZONE LIST** page (as the case may be) will appear, upon launch.

18.2 Quickly Launching a Third Party Console from the eG Monitoring Console

In environments where multiple monitoring solutions - one each for every component silo - are deployed, you may want to seamlessly integrate the eG monitoring console with a few or all of these solutions, so as to

enable effective investigation and swift identification of the source of problems with the associated components.

Using the eG administrative interface, you can now configure the URLs for launching third-party management consoles from the eG monitoring console. Each such URL is called an **External Monitor**. By associating these monitors with specific component silos, you can enable users to the eG monitoring console to quickly drill down to third-party consoles whenever the associated components experience performance issues, use the expertise of the third-party tools to analyze these issues, and thus accurately diagnose their root-cause.

To configure external monitors, do the following:

1. Select the **External Monitors** option from the **Miscellaneous** tile.
2. If no monitors pre-exist, the message depicted by Figure 18.1 will appear.

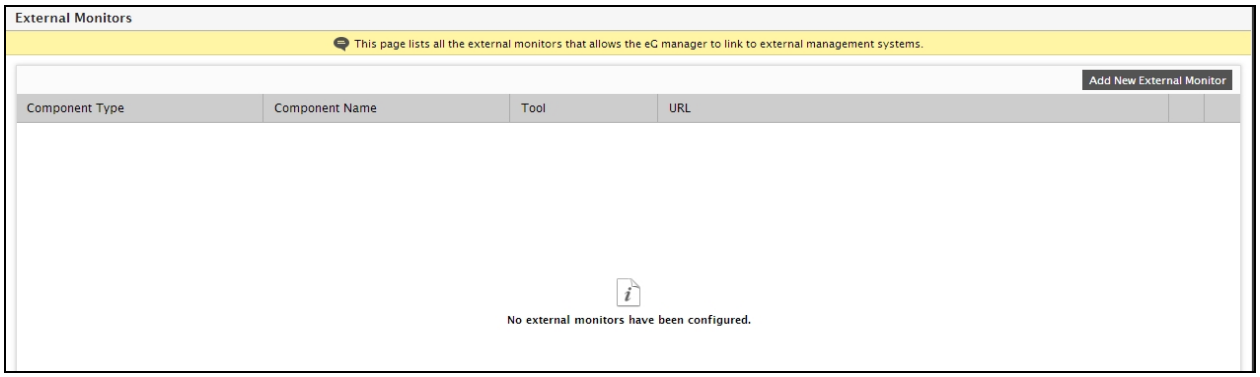

Figure 18.1: No external monitors

3. To add a new monitor, click on the **Add New External Monitor** button in Figure 18.1.
4. Figure 18.2 will then appear, using which you can configure the new external monitor.

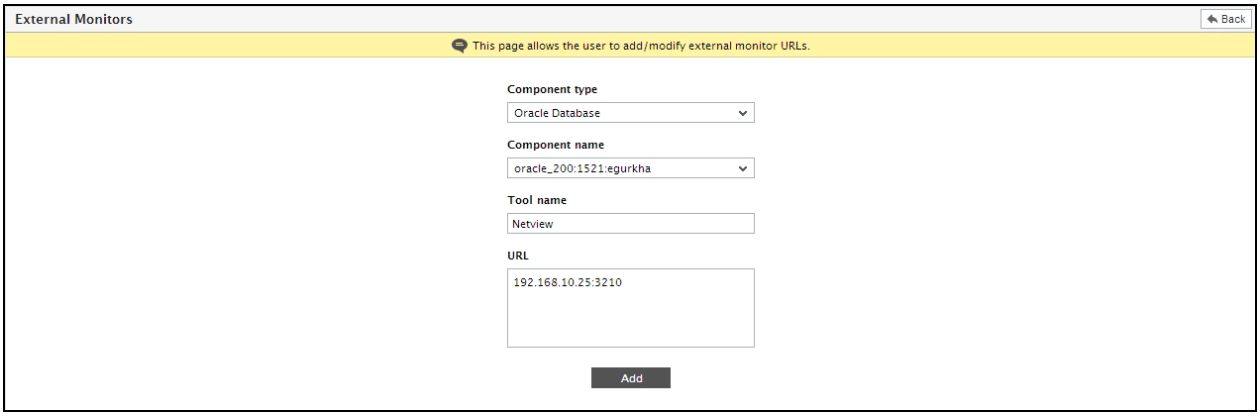

Figure 18.2: Adding a new external monitor

5. Select the **Component type** for which the external monitor is to be configured.
6. The **Component Name** list will then be populated with all components of the chosen type. Pick the component with which the new monitor is to be associated.

7. Specify the name of the third-party monitoring tool that is being integrated with the eG Enterprise Suite in the **Tool name** text box.
8. Specify the **URL** of the web page that should appear when users attempt to launch the third-party monitoring console from the eG monitoring console.
9. Finally, click the **Add** button.
10. The **External Monitor** so added will then be visible in the layer model page (of the eG monitoring console) of the corresponding component. The **Tests** panel of the layer model page of that component will display a link (bearing the configured third-party tool name), which when clicked will instantly launch the monitoring console of the corresponding tool.

19

Quick Links

Based on the roles and responsibilities assigned to them, users to the eG Enterprise system may perform a few operations more frequently than the rest. For instance, in dynamic infrastructures, you would find administrators adding components for monitoring on a daily basis. Likewise, you could have help desk managers who may be interested in checking the alarm history every few days, while performing more detailed event analysis as well using the eG Reporter. Similarly, you could be a VMware administrator, who routinely studies the highs and lows of the virtualized infrastructure using the Virtualization Reports offered by eG Reporter.

Such users may need instant access to those web pages in the eG management console, which they visit often for performing their routine, yet critical operations, so that no time is wasted in following the sequence of menu options offered by the eG Enterprise system or in manually switching from one module to another. To enable this quick, module-independent access, eG Enterprise allows users to configure **Quick Links** to specific web pages that they frequently visit. These links are defined per user and can hence be configured only for those modules a user is permitted to access.

To configure quick links, do the following:

1. The main menu bar of each module of the eG Enterprise system - be it **Admin, Monitor, Reporter,** or **Configuration** - includes a **Quick Links** icon ( ). If the user who logs into a module has already configured quick links for that module, then clicking on the **Quick Links** icon will invoke a window that lists all the links that are enabled for that user for the module that is currently open (by default). If no quick links pre-exist for the module logged in, then a message to that effect will appear in the same window (see Figure 19.1).

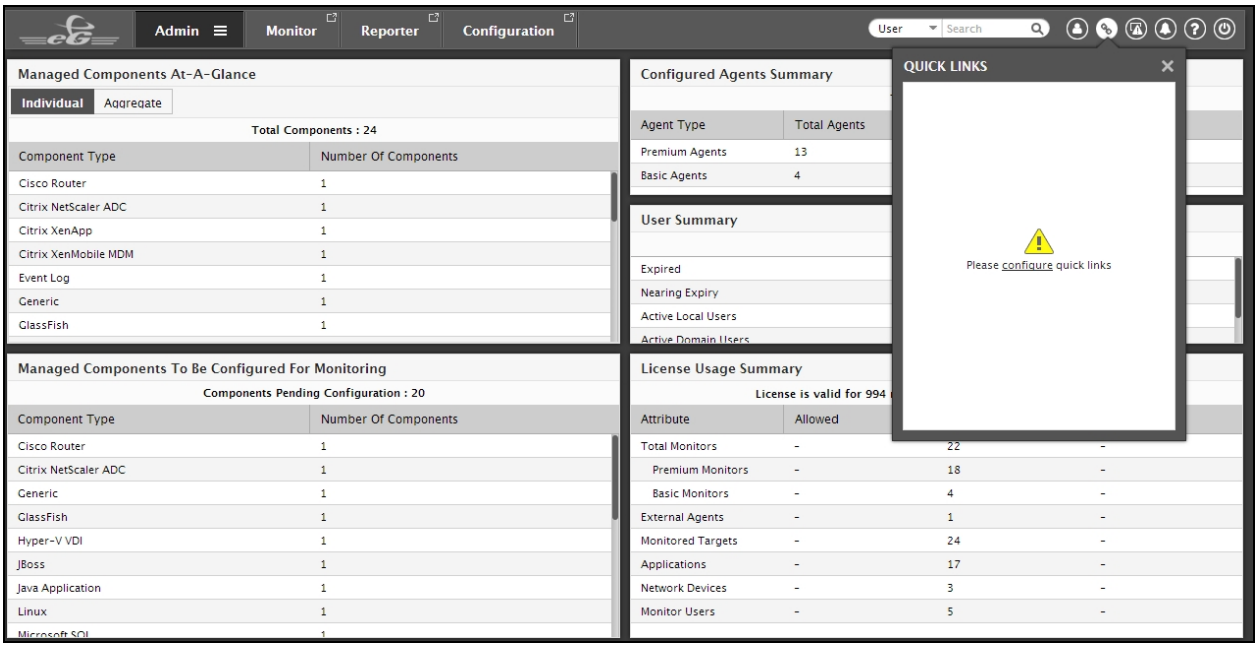

Figure 19.1: A message stating that no quick links pre-exist

To begin configuring quick links, click on the **Please configure quick links** message in Figure 19.1.

This will lead you to Figure 19.2. By default, the module that is currently open will be chosen from the **Module** list in Figure 19.2. To configure quick links for the default module, select the links to enable from the **Disabled Links** list and click **<** button.

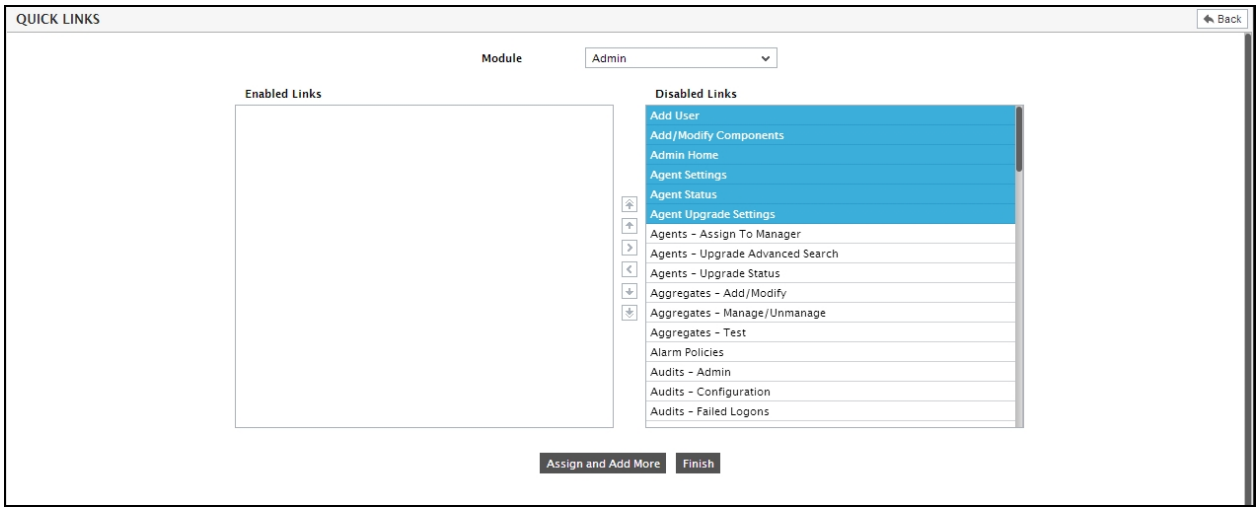

Figure 19.2: Selecting the links to enable

This will transfer your selection to the **Enabled Links** list (see Figure 19.3). To disable one/more enabled links, just select the links from the **Enabled Links** list and click the **>** button.

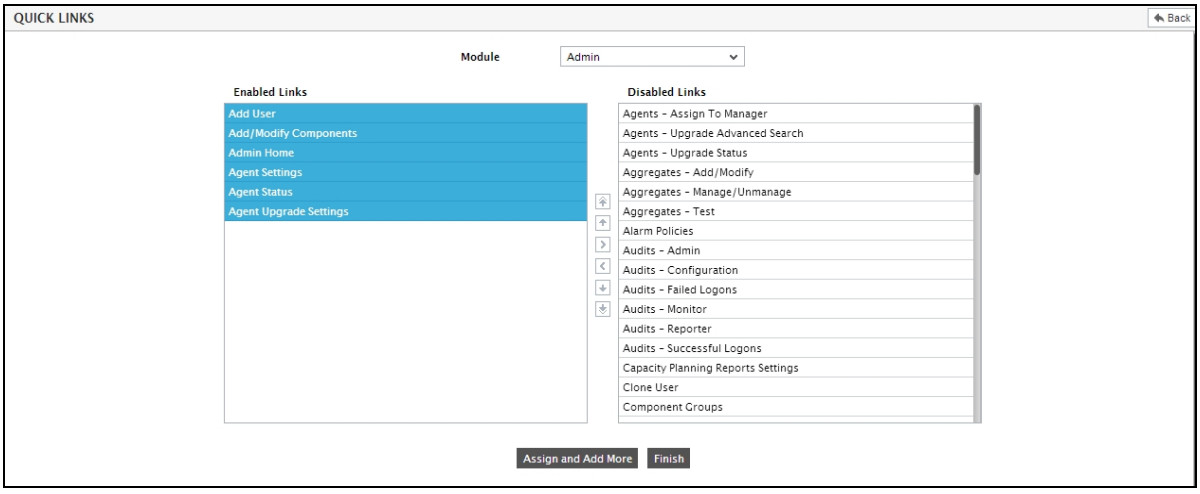

Figure 19.3: Selection transferred to the Enabled Links list

To configure links for other modules, you need to click the **Assign and Add More** button in Figure 19.3 and select a different module from the **Module** list. The list of **Disabled Links** for that module will then be listed. You can then repeat the procedure discussed at step 3 and 4 above to enable links for the chosen module. This way, you can configure quick links for all modules that you have access to.

By default, the quick links in the **Disabled Links** and in the **Enabled Links** lists will be arranged in alphabetical order. If, after moving specific links to the **Enabled Links** list, you wish to change the order in which the links

are displayed for selection, then, you can use the arrow buttons to the left of the **Enabled Links** list for this purpose. For instance, if you want the **Add/Modify Components** link in the **Enabled Links** list to be moved up in the order, then, select the **Add/Modify Components** link from the **Enabled Links** list, and click the **up-arrow** button to the left of the list. This will push the **Add/Modify Components** link above the **Add Users** link.

The other buttons available to the left of the **Enabled Links** list and their purpose have been discussed below:

Button	Purpose
	Moves the selection to the top of the Enabled Links list
	Pushes the selection one down the order
	Moves the selection to the bottom of the Enabled Links list

With the help of these buttons, you can ensure that links to the most frequently accessed pages are readily available at the top of the list, and the less-used links are pushed down the order.

Once you are done enabling links, click the **Finish** button in Figure 19.3 to save the changes and exit the **QUICK LINKS** page.

If you now click on the **Quick Links** icon ( ), Figure 19.4 will appear displaying the quick links enabled for the module that is currently open (by default).

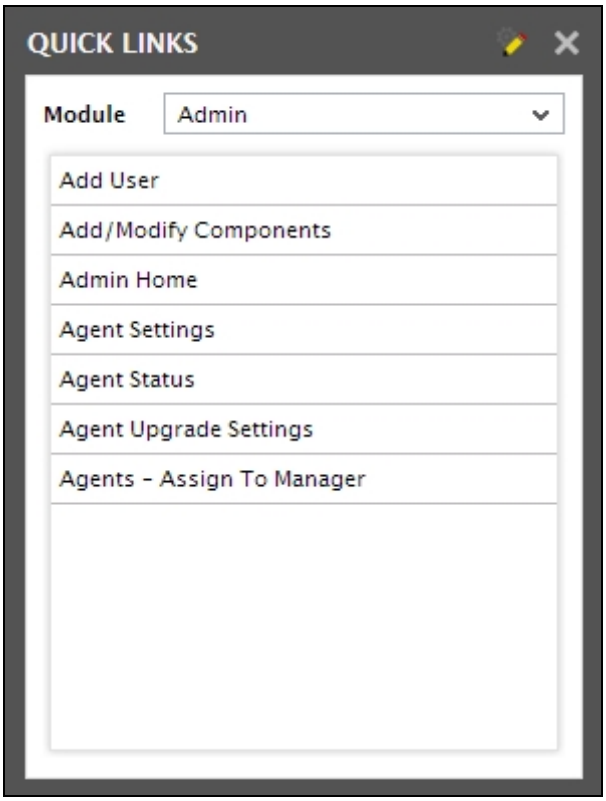

Figure 19.4: The quick links configured for the module that is currently open

Click on any link in Figure 19.4 to immediately switch to the web page of interest to you, without having to follow menu sequences. For instance, clicking on say, **Add/Modify Components** in Figure 19.4 will provide you with instance access to Figure 19.5. In the absence of quick links, you will have to invoke the **Admin** tile menu, browse the **Infrastructure** tile, move your mouse pointer over the **Components** sub-menu, and click the **Add/Modify** option within to access Figure 19.5.

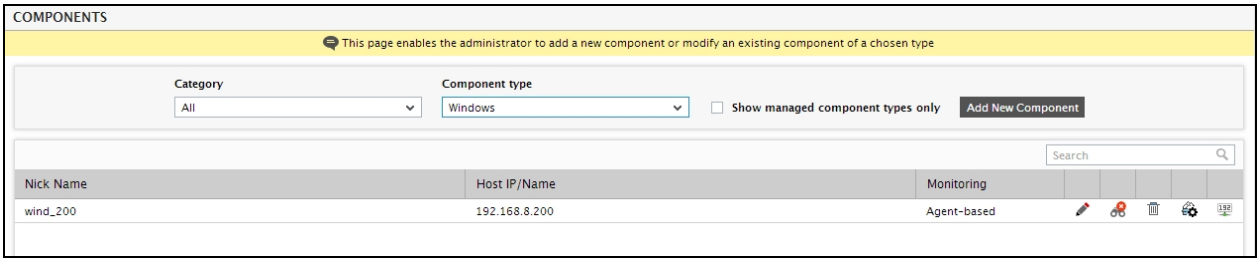

Figure 19.5: A Quick Link instantly opening the Agent Discovery Setting page

You can pick any other **Module** from the **Quick Links** pop-up to view the links configured for that module (see Figure 19.6).

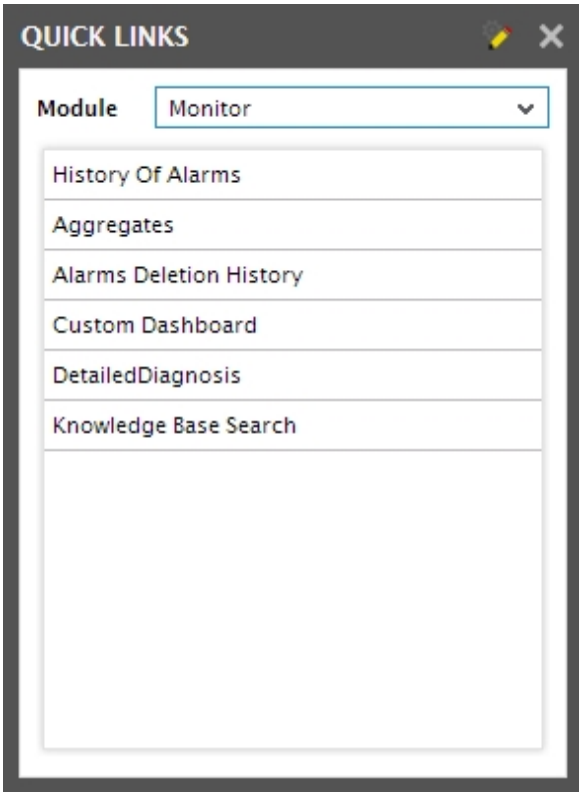

Figure 19.6: Quick links configured for the Monitor module

Clicking on a link in Figure 19.6 will give you 'single-click' access to a web page in a different module, without mandating a manual module switch. For instance, clicking on the **History Of Alarms** link in Figure 19.6 will instantly open Figure 19.7, regardless of which module you are currently logged into.

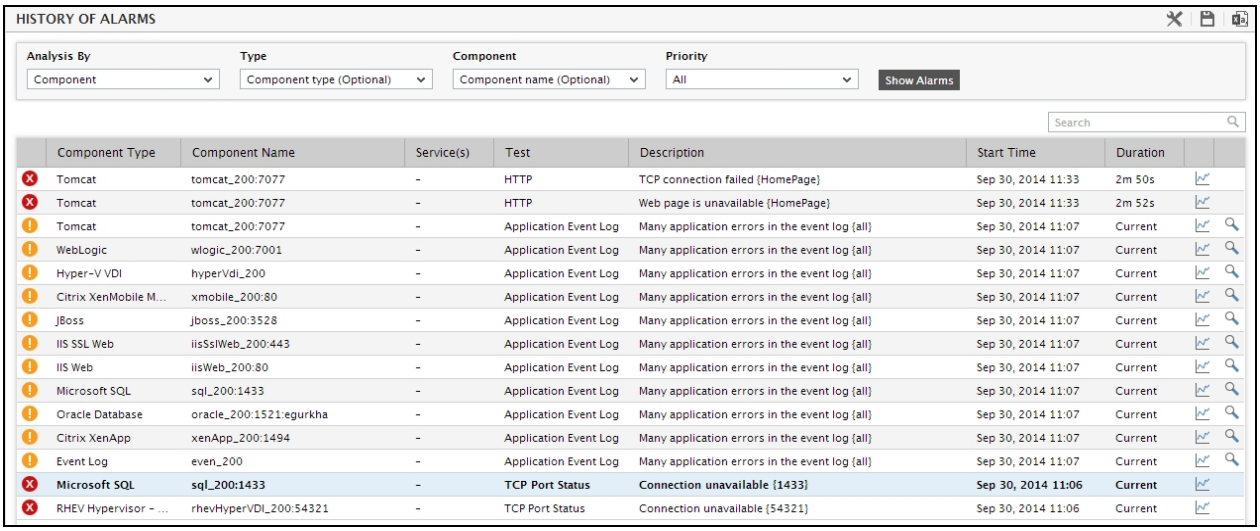

Figure 19.7: A Quick Link providing you with instant access to a web page in a different module

20

Advanced Features

20.1 Importing/Exporting Configuration Across eG Managers

Multiple eG managers are common in many IT infrastructures. Using industry-standard best practices, administrators of such environments may have fine-tuned the configuration (tests, thresholds, data cleanup frequencies, etc.) of one of the eG managers, and may want to apply the same settings to the other managers. To achieve this, administrators may have to painstakingly document the 'ideal' configuration and then manually login to each manager to apply the documented settings one by one. To reduce the manual effort and time involved in this exercise, eG Enterprise enables administrators to quickly apply the configuration of one manager to other managers at one go, using a simple export-import routine. With this capability, administrators can use the eG management interface to quickly export and import the following configurations between managers:

- Default test configurations
- List of enabled/disabled tests
- Default threshold configurations
- Global threshold settings
- Alarm policies
- Data cleanup frequencies
- Fix history

To export the configuration settings of a manager – say, Manager A – do the following:

1. Login to the eG administrative interface of Manager A.
2. Select the **Export/Import Configuration** option from the **Miscellaneous** menu of the **Admin** tile. Figure 20.1 will appear.

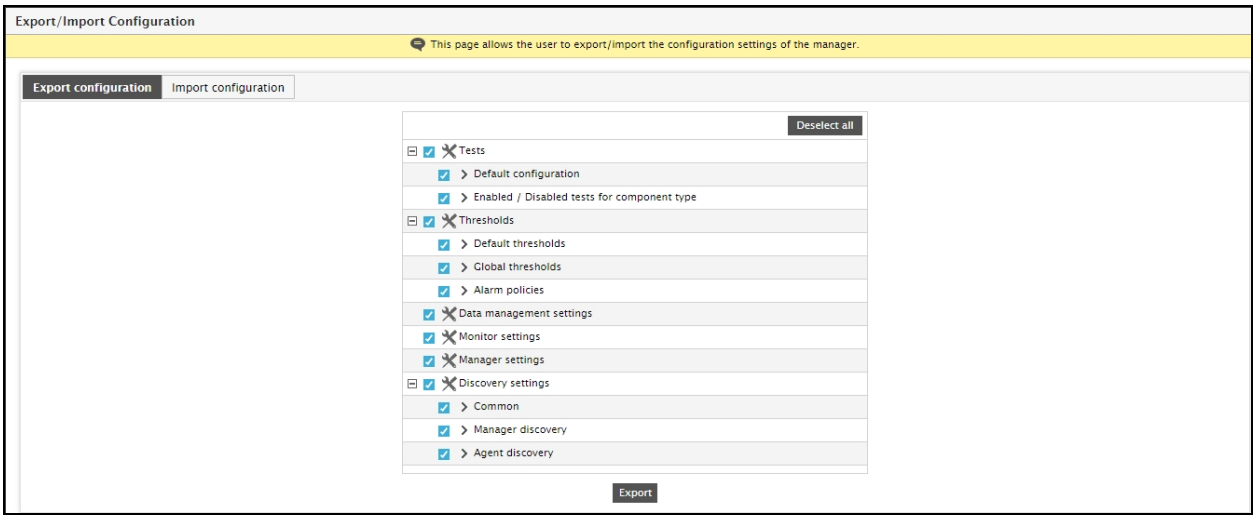

Figure 20.1: Exporting the configuration

3. By default, the **Export configuration** tab page of Figure 20.1 will open. This tab page allows you to choose the configuration settings that you want to export from the current eG manager. As stated earlier, the following configurations are available for export:

   - Default test configurations
   - List of enabled/disabled tests
   - Default threshold configurations
   - Global threshold settings
   - Alarm policies
   - Data cleanup frequencies
   - Manager and monitor settings
   - eG manager and agent discovery settings

   To select a configuration, click on the check box corresponding to it. Since some configurations have been grouped under a head, selecting a group head automatically selects all configurations under that head. For instance, selecting the **Tests** group, automatically selects the **Default configuration** and **Enabled / Disabled tests for component type** options. To deselect all displayed configuration settings, click on the **Deselect all** button. To select all displayed configuration settings, click on the **Select all** button.

4. Once all the required configuration settings have been selected, click on the **Export** button to begin exporting the chosen settings. eG Enterprise exports the configuration settings you select to a zip file, which is by default named in the following format: **<Fully qualified host name of eG manager>_eGconfig_<Date of export>.zip**. This zip file will by default be downloaded to the **default download destination** that you have configured in your browser. Once the export completes successfully, the contents of this zip file will be displayed on-screen for you to take a look (see Figure 20.2).

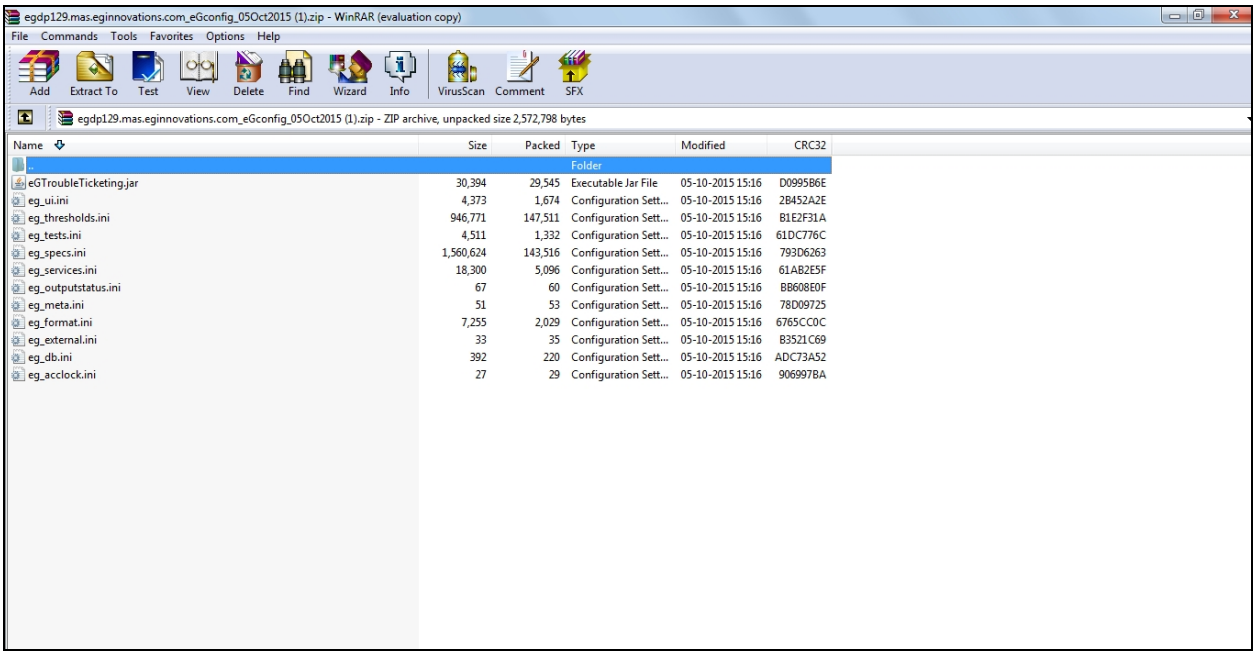

Figure 20.2: Contents of the zip file containing the exported configuration files

To enable another manager – say, Manager B - to import this configuration, follow the steps below:

1. First, copy this zip file from the **default download destination** to any folder on Manager B's host.
2. Then, login to Manager B's administrative interface and select the**Export/Import Configuration** option from the **Miscellaneous** menu of the **Admin** tile.
3. Figure 20.3 will appear. This time, click on the **Import configuration** tab page in Figure 20.3 to open it.

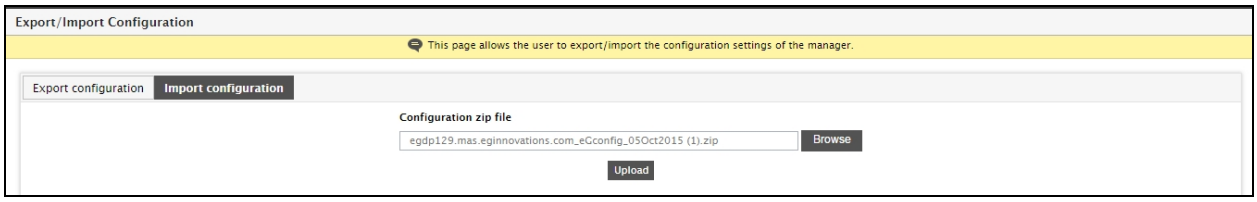

Figure 20.3: Importing the configuration

4. Using the **Choose File** button in Figure 20.3, browse Manager B's host to locate the zip file that contains the configurations to be imported.
5. Then, click on the **Upload** button to upload the file to Manager B.
6. Figure 20.4 will then appear. With the help of Figure 20.4, you can choose the specific configuration settings you want to import. To select a configuration, click on the check box corresponding to it. Since some configurations have been grouped under a head, selecting a group head automatically selects all configurations under that head. For instance, selecting the **Tests** group, automatically selects the **Default configuration** and **Enabled / Disabled tests for component type** options. To deselect all displayed configuration settings, click on the **Deselect** button. To select all displayed configuration settings, click on the **Select all** button.

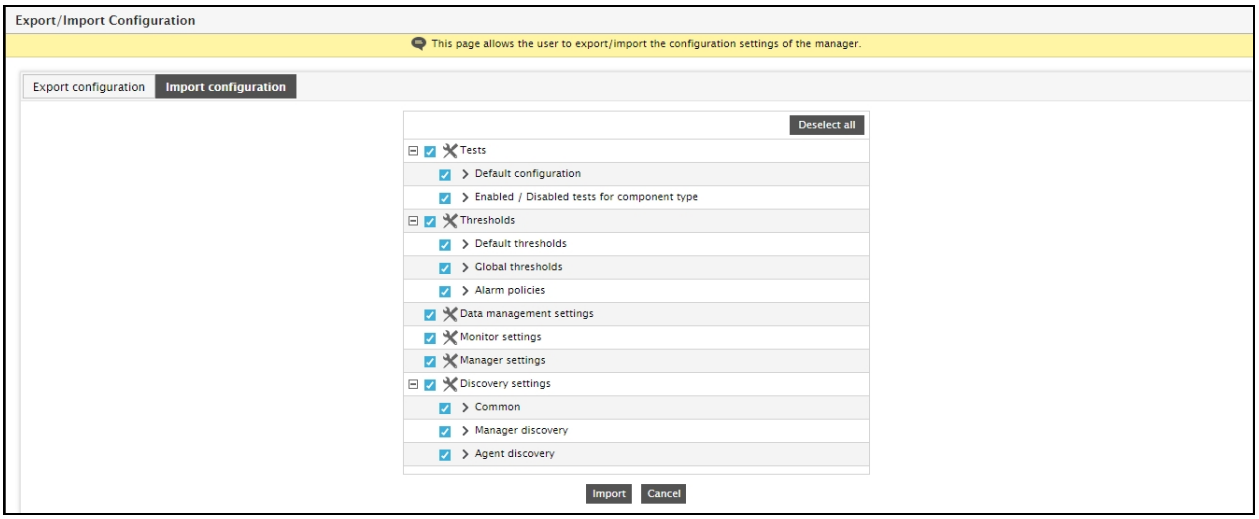

Figure 20.4: Selecting the configuration settings to be imported

7. To begin the import of the selected settings, click on the **Import** button in Figure 20.4.
8. If import is successful, a message to that effect will appear.

Note:

- The default configuration of host-level tests in Manager A will overwrite the default configuration of the same tests in Manager B post the import, regardless of the component-type to which the tests are mapped.
- On the other hand, changes made to an application-level test's default configuration in Manager A, can be viewed in Manager B post the import, only if one/more applications of the same type are managed in Manager B.
- In case of default threshold configurations, post the import, all default threshold specifications related to host-level tests in Manager A wil overwrite the default threshold settings of the same tests in Manager B, regardless of component-type. However, the default threshold settings of application-level tests in Manager A will be applied to Manager B post the import, only if one/more applications of the same type are managed in Manager B.
- Alarm policies and data management settings configured in Manager A will overwrite the existing configurations in Manager B after a successful import. On the other hand, the fix history maintained by Manager A will be appended to the fix history that pre-exists in Manager B, post the import. Accordingly, if you access the **Knowledge Base Search** page in the eG monitoring console of Manager B and click the **Submit** button, the complete history of fixes maintained by that manager and the fix history imported from Manager A will be displayed therein. This is why, you will find the **Knowledge Base Search** page display the details of even those fixes that pertain to components that are not managed by Manager B.

20.2 Importing/Exporting user-defined component types

In many IT infrastructures, as an industry-standard best practice, administrators may have managed a brand new component type using the Integration Console capability offered by the eG Enterprise Suite, and may want to manage the same component type across the servers in their environment. To achieve this, administrators may have to painstakingly document the 'ideal' configuration and then manually login to each

server, add the IC component type one by one. To reduce the manual effort and time involved in this exercise, eG Enterprise enables administrators to quickly apply the configuration of a component type managed using the Integration Console to other servers at one go, using a simple export-import routine.

To export the configuration settings of a component type that is managed using the Integration Console capability, do the following:

1. Login to the eG administrative interface.
2. Follow the menu sequence: *Admin -> Miscellaneous -> Export/Import Configuration -> Integration Console -> Component*. Figure 20.5 will then appear.

Figure 20.5: Exporting the configuration of a user-defined component type

3. By default, the **Export configuration** tab page of Figure 20.5 will open. Using this tab page you can export the configuration settings of a component type listed in this page.

   To select a configuration, click on the check box corresponding to the component type. To select all displayed component types, click on the check box against the **Component Type**. To search for a component type of your choice, you can use the **Search** text box. If you wish to deselect a component type, then, you can uncheck the check box against the component type. To deselect all the component types, uncheck the check box against the **Component Type**.

4. Then, click on the **Export** button (see Figure 20.5) to begin exporting the settings of the chosen component types. eG Enterprise exports the configuration settings you select to a zip file, which is by default named in the following format: *<Fully qualified host name>_eGICComponent_<Date of export>.zip*. This zip file will by default be downloaded to the default download destination that you have configured in your browser.

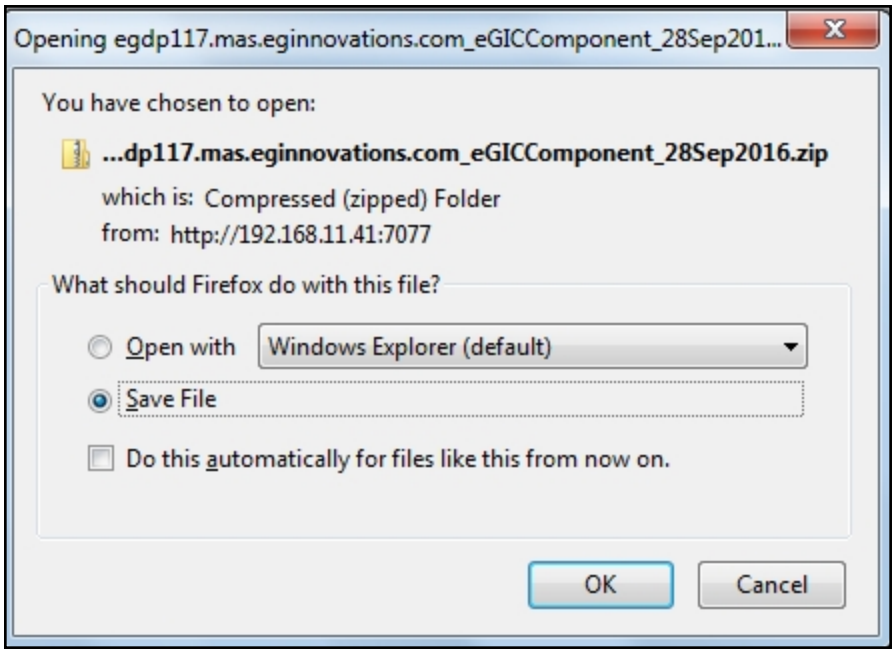

Figure 20.6: Saving the exported configuration

Once the export completes successfully, the contents of this zip file will be displayed on-screen for you to take a look as shown in Figure 20.7.

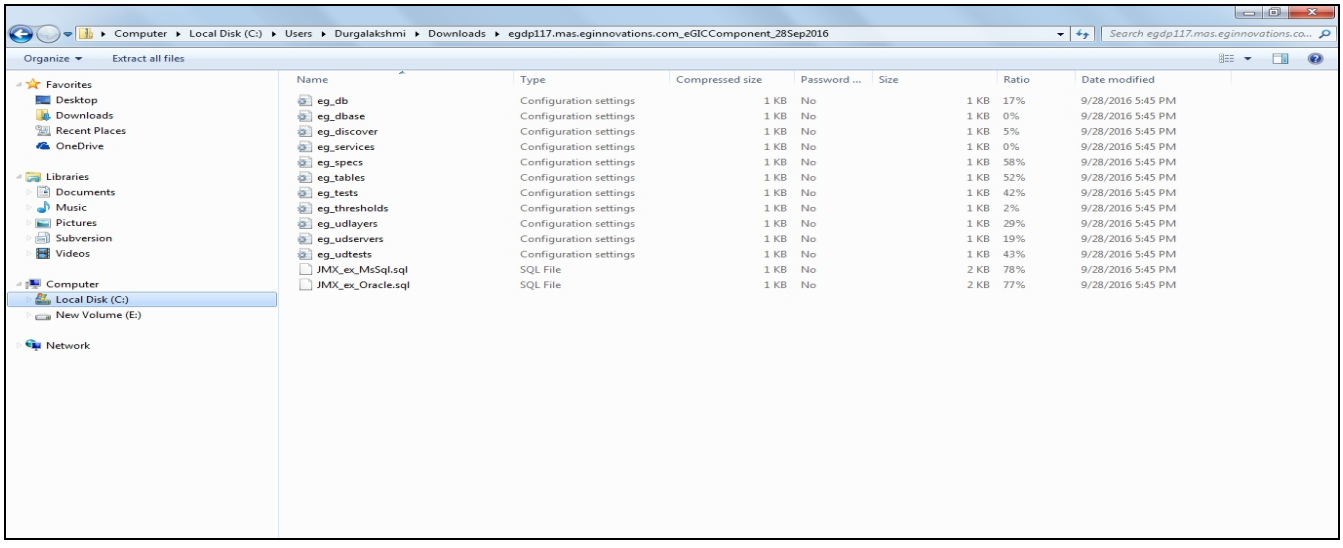

Figure 20.7: Contents of the zip file containing the exported configuration files

To enable an eG manager say, Manager B to import the configuration that was exported from a manager say Manager A, follow the steps below:

1. First, copy this zip file from the **default download destination** of the Manager A to any folder on Manager B's host.

2. Then, login to Manager B's administrative interface and select the **Component** option from the **Export/Import Configuration** option that is available under the **Miscellaneous** menu of the **Admin** tile.
3. Figure 20.8 will appear. This time, click on the **Import configuration** tab page in Figure 20.8 to open it.

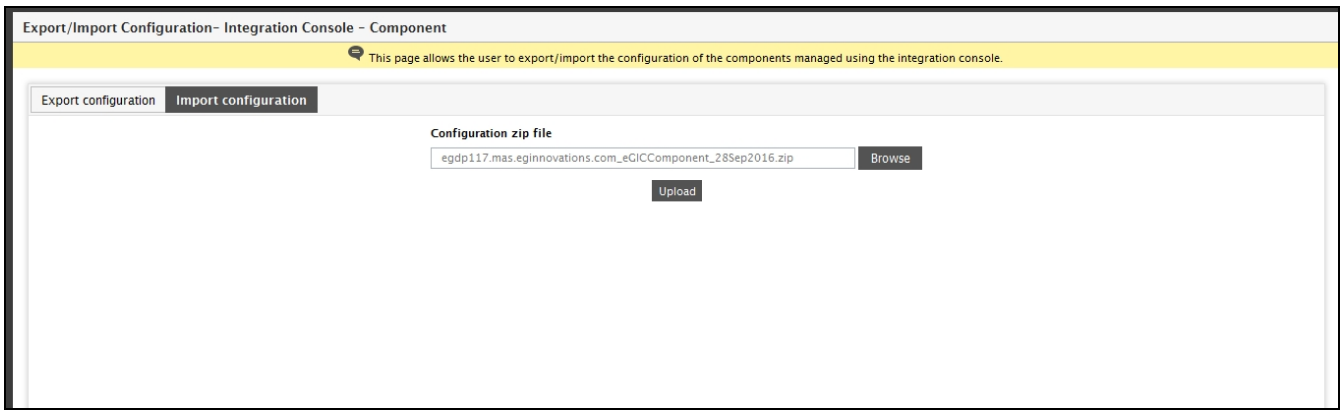

Figure 20.8: Importing the configuration

4. Using the **Browse** button in Figure 20.8, browse the Manager B's host to locate the zip file that contains the configurations to be imported.
5. Then, click on the **Upload** button in Figure 20.8 to upload the file to the server.
6. Figure 20.9 will then appear. With the help of Figure 20.9, you can choose the component types that you want to import. To select a component type, click on the check box corresponding to it. To select all displayed component types, click on the check box against the **Component Type**. To deselect a particular component type, simply uncheck the check box corresponding to the component type in this page. To search for a component of your choice, you can use the **Search** option.

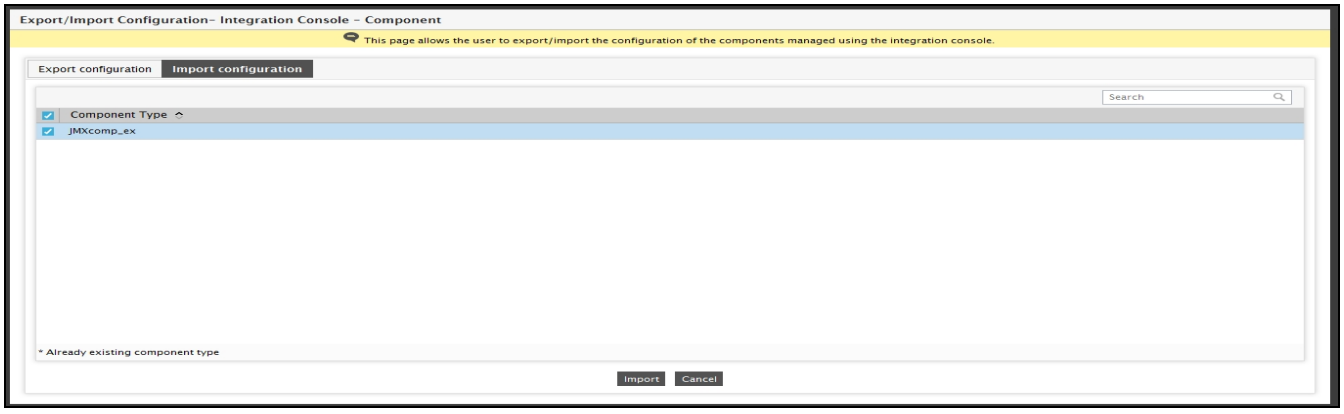

Figure 20.9: Selecting the configuration settings to be imported

7. To begin the import of the selected component types, click on the **Import** button in Figure 20.9.
8. If import is successful, a message to that effect will appear.

Note:

While a user-defined component type is being exported/imported, the configuration of the tests corresponding to that component type will be retained and imported to the manager in a multi-manager environment.

20.3 Importing/Exporting user-defined Tests

In many IT infrastructures, as an industry-standard best practice, administrators may have added a new test using the Integration Console capability offered by the eG Enterprise Suite, and may want to add the same test across the servers in their environment. To achieve this, administrators may have to painstakingly document the 'ideal' configuration and then manually login to each server, add the test one by one. To reduce the manual effort and time involved in this exercise, eG Enterprise enables administrators to quickly apply the configuration of a test added using the Integration Console to other servers at one go, using a simple export-import routine.

To export the configuration settings of a user-defined test, do the following:

1. Login to the eG administrative interface.
2. Follow the menu sequence: *Admin -> Miscellaneous -> Export/Import Configuration -> Integration Console -> Test*. Figure 20.10 will then appear.

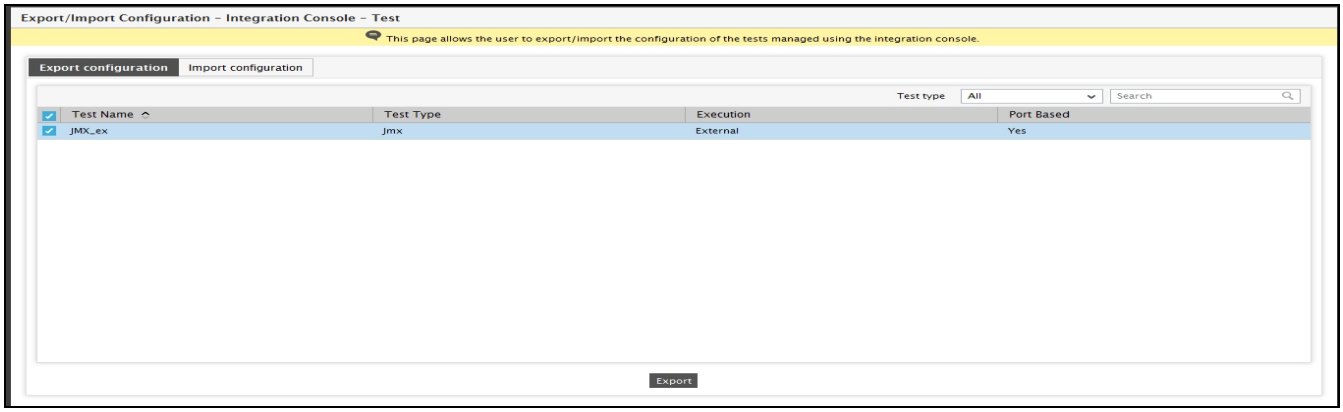

Figure 20.10: Exporting the configuration of a user-defined component type

3. By default, the **Export configuration** tab page of Figure 20.10 will open. Using this tab page you can export the configuration settings of the user-defined tests listed in this page.

   To select a test, click on the check box corresponding to the test i.e., Test A_ex. To select all displayed tests, click on the check box against the **Test Name**. To search for a test of your choice, you can use the **Search** text box. If you wish to deselect a test, then you can uncheck the check box against the test. To deselect all the tests, uncheck the check box against the **Test Name**.

   If you wish to filter the tests based on the category that was chosen while adding the test, then you have to select an option from the **Test type** list. By default, *All* option will be chosen from this list indicating that all the tests will be listed in this page, by default.

4. Then, click on the **Export** button to begin exporting the settings of the chosen component types. eG Enterprise exports the configuration settings you select to a zip file, which is by default named in the

following format: **<Fully qualified host name>_eGICTest_<Date of export>.zip**. This zip file will by default be downloaded to the **default download destination** that you have configured in your browser. Once the export completes successfully, the contents of this zip file will be displayed on-screen for you to take a look.

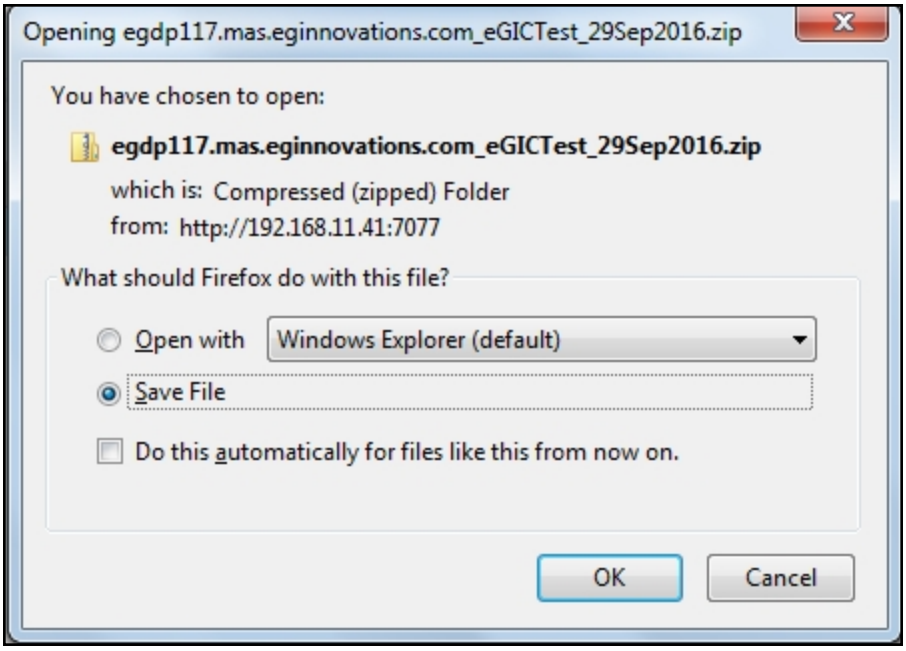

Figure 20.11: Saving the exported configuration

Once the export completes successfully, the contents of this zip file will be displayed on-screen for you to take a look as shown in Figure 20.12.

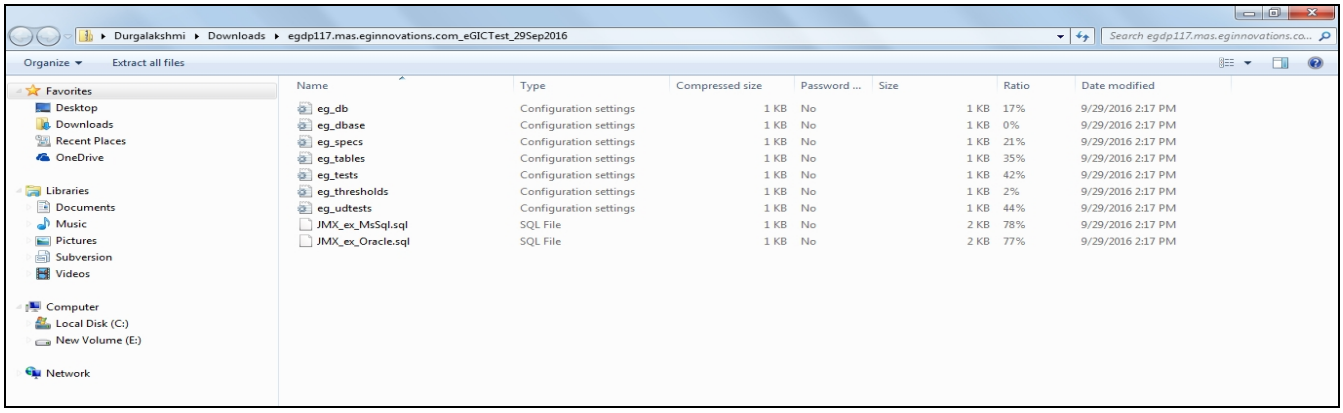

Figure 20.12: Contents of the zip file containing the exported configuration files

To enable an eG manager say, Manager B to import the configuration that was exported from a manager say Manager A, follow the steps below:

1. First, copy this zip file from the **default download destination** of the Manager A to any folder on Manager B's host.

2. Then, login to Manager B's administrative interface and select the **Test** option from the **Export/Import Configuration** option that is available under the **Miscellaneous** menu of the **Admin** tile.
3. Figure 20.13 will appear. This time, click on the **Import configuration** tab page in Figure 20.13 to open it.

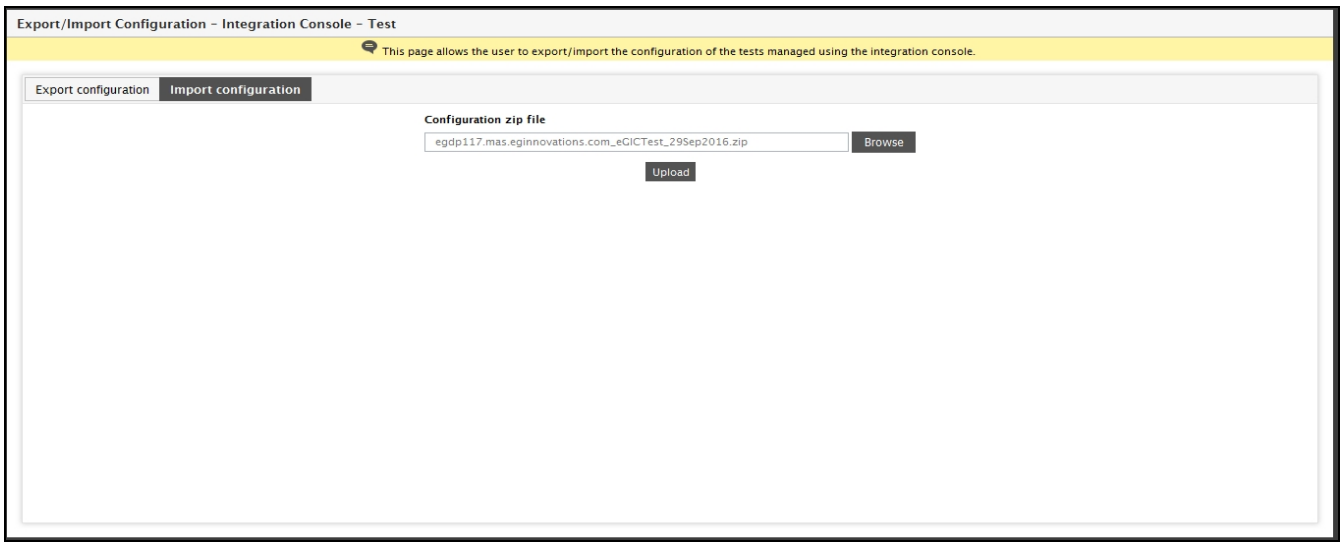

Figure 20.13: Importing the configuration

4. Using the **Browse** button in Figure 20.13, browse the Manager B's host to locate the zip file that contains the configurations to be imported.
5. Then, click on the **Upload** button in Figure 20.13 to upload the file to the server.
6. Figure 20.14 will then appear. With the help of this page, you can choose the tests that you want to import. To select a test, click on the check box corresponding to it. To select all displayed tests, click on the check box against the **Test Name**. To deselect a particular test, simply uncheck the check box corresponding to the test in this page. To search for a test of your choice, you can use the **Search** option.

   If you wish to filter the tests based on the category that was chosen while adding the test, then you have to select an option from the **Test type** list. By default, *All* option will be chosen from this list indicating that all the tests will be listed in this page, by default.

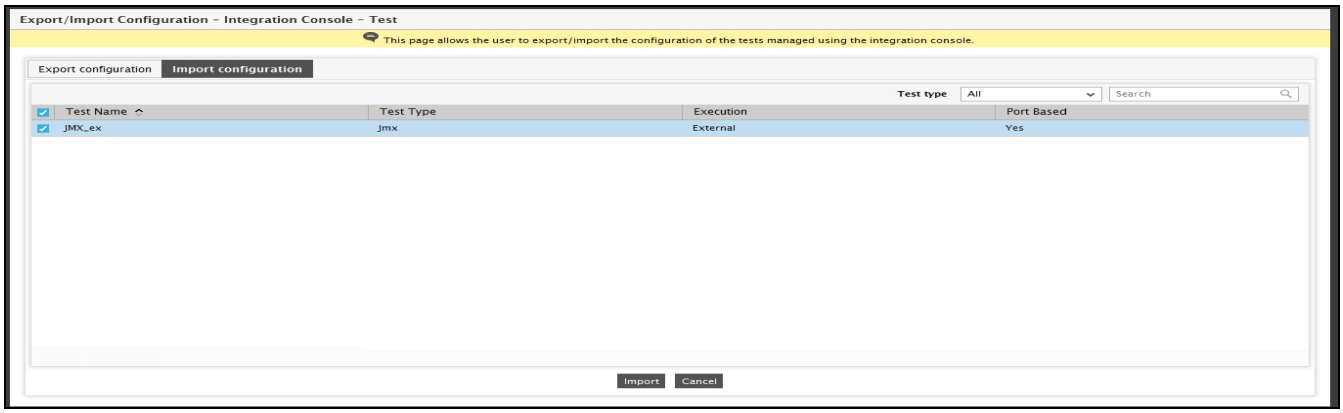

Figure 20.14: Selecting the configuration settings to be imported

7. To begin the import of the selected tests, click on the **Import** button in Figure 20.14.
8. If import is successful, a message to that effect will appear.

20.4 Configuring Remote Control Commands

Monitoring solutions often provide the ability to alert an administrator over email, pager, SMS, etc., when problems occur. In response to an alert, the administrator has to perform domain-specific detailed analysis of the problem, often by running different commands on the target system. In this process, the administrator has to figure out and initiate corrective measures. Most monitoring solutions provide remote problem alerting capability, but the ability to **remotely login in a secure manner and perform detailed analysis and troubleshooting** is not available. To allow true anytime, anywhere management capability, such remote control of the target IT infrastructure must be possible using a web browser.

eG's Remote Control Action capability allows an administrator to remotely and securely access any monitored server in an IT infrastructure and to execute remote commands in order to perform detailed analysis of problems and to initiate corrective actions against them.

By default, eG Enterprise provides a set of pre-defined remote commands, for use during remote problem correction or diagnosis. If need be, administrators can override this default list using the **REMOTE CONTROL COMMAND** page. To access this page, click the **Settings** menu option in the **Agents** tile, and select the **Remote Control** node in the **AGENT SETTINGS** panel. This will open Figure 20.15.

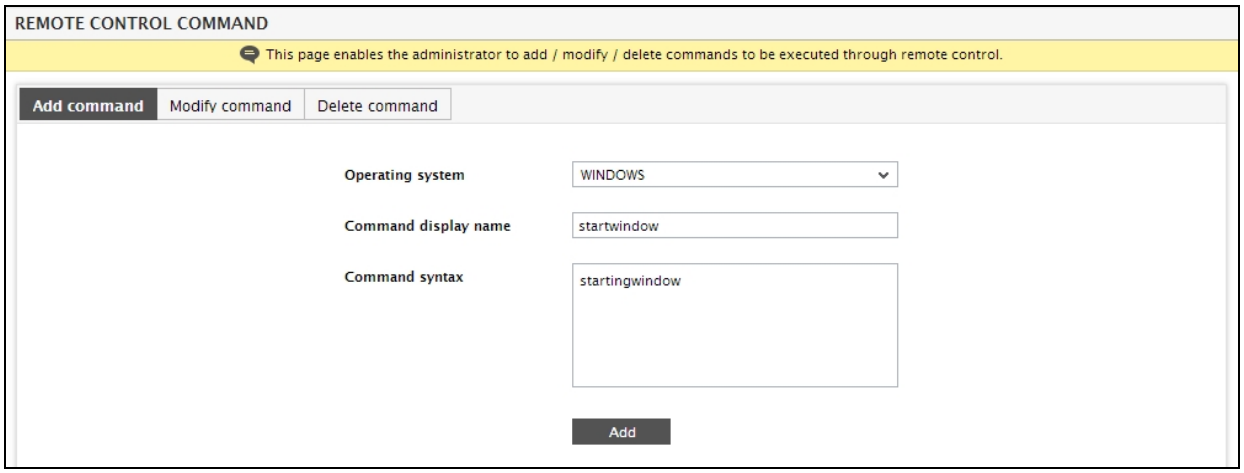

Figure 20.15: Adding a command

Using the **REMOTE CONTROL COMMAND** page, one can add new commands, modify existing ones, or even delete one/more commands. To add a command, follow the steps given below:

1. Click the **Add command** tab page in Figure 20.15.
2. Next, choose the **Operating system** on which the remote command will function.
3. Then, provide a unique display name for the command against **Command display name**.
4. Finally, provide the complete **Command syntax** and click the **Add** button to register the command.

To modify a command, do the following:

1. Select the **Modify command** tab page as depicted by Figure 20.16.
2. Select the **Operating system** to which the command applies.
3. The **Existing commands** list box displays all existing commands that pertain to the chosen **Operating system**. From this list, select the command to be modified.
4. Make necessary changes to the **Command syntax**, and **Update** the changes.

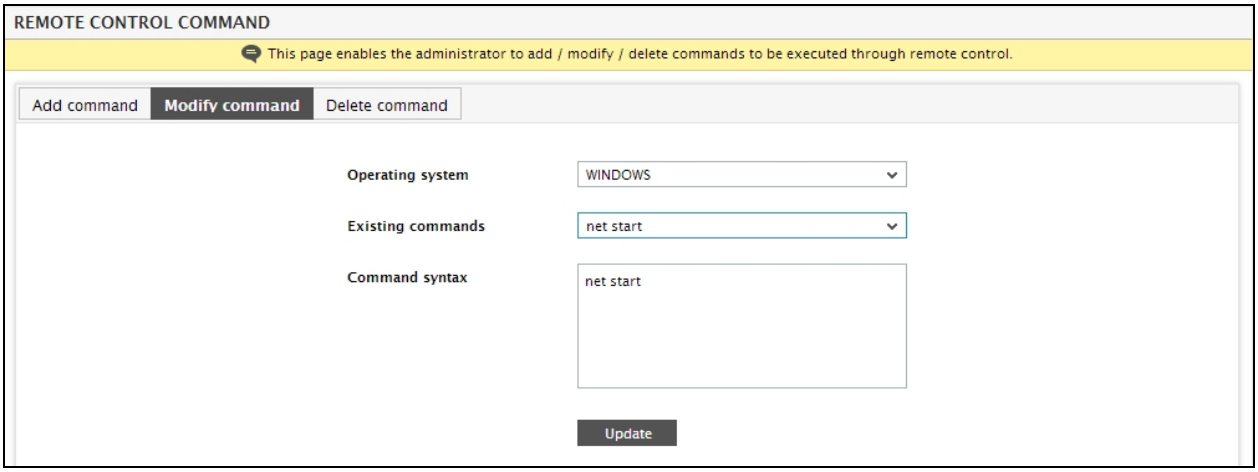

Figure 20.16: Modifying a command

To delete a remote command, do the following:

1. Select the **Delete command** tab page as depicted by Figure 20.17.
2. Select the **Operating system** to which the command to be deleted applies.
3. The **Existing commands** list box displays all existing commands that pertain to the chosen **Operating system**. From this list, select the command(s) to be deleted.
4. Finally, click the **Delete** button.

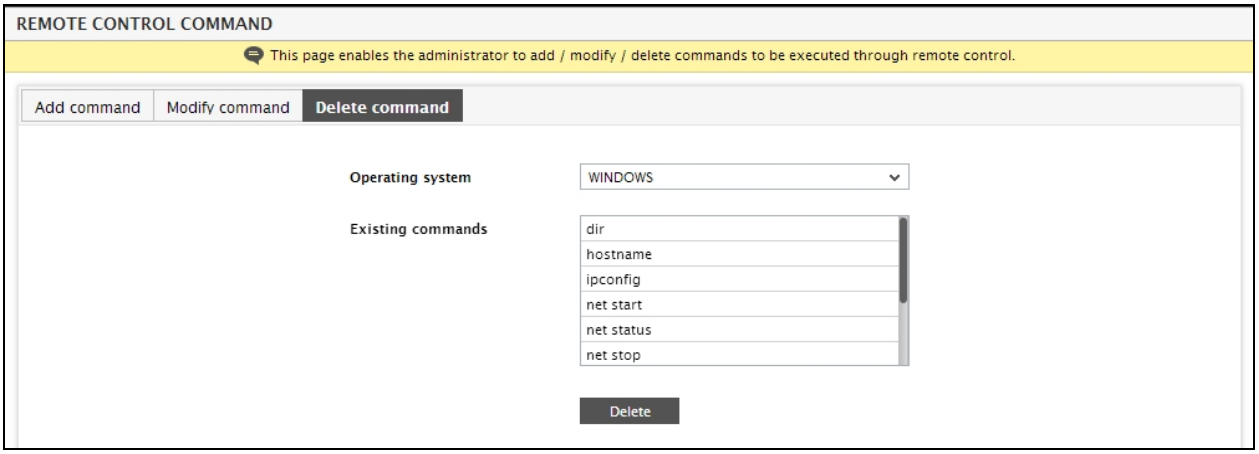

Figure 20.17: Deleting a command

20.5 Viewing VM Statistics

eG remote agents that perform 'In-N-Out' monitoring of hypervisors (such as VMware vSphere/ESX servers, Citrix XenServers, etc.) and the VMs configured on them capture critical VM-related statistics and errors that occur when monitoring the VMs to a file named **eg_vm.ini**; this file is automatically created on the agent host. When troubleshooting the failure of the remote agent to obtain the 'inside view' of VMs, this file serves as the primary source of problem information to the administrator, as it provides effective pointers to the root-cause of such failures. Typically, administrators have to physically login to the agent host to view the contents of the **eg_vm.ini** file. Where multiple hypervisors are being monitored - each with a dedicated remote agent - administrators will have to access the **eg_vm.ini** file of each agent to initiate investigations. To save administrators the time and trouble involved in this exercise, the eG administrative console now embeds a comprehensive interface, which provides a hypervisor-independent view of the following:

- The physical servers managed by a chosen remote agent;
- The details of VMs running on each physical server;
- The powered-on state of every VM;
- Whether the 'inside view' tests failed on any VM, and if so, the reason for the same

In other words, the dedicated interface allows administrators a sneak peek at the contents of the **eg_vm.ini** file for a chosen remote agent.

To access this interface, do the following:

1. Select the **Remote Agents** menu option from the **Agents** tile.
2. Click the **VM Statistics** button in the **REMOTE AGENT CONFIGURATION** page.
3. When Figure 20.18 appears, select a remote agent from the **Agent** list.
4. Upon selecting a remote agent, the contents of the **eg_vm.ini** file associated with that agent will then be displayed as depicted by Figure 20.18 below.

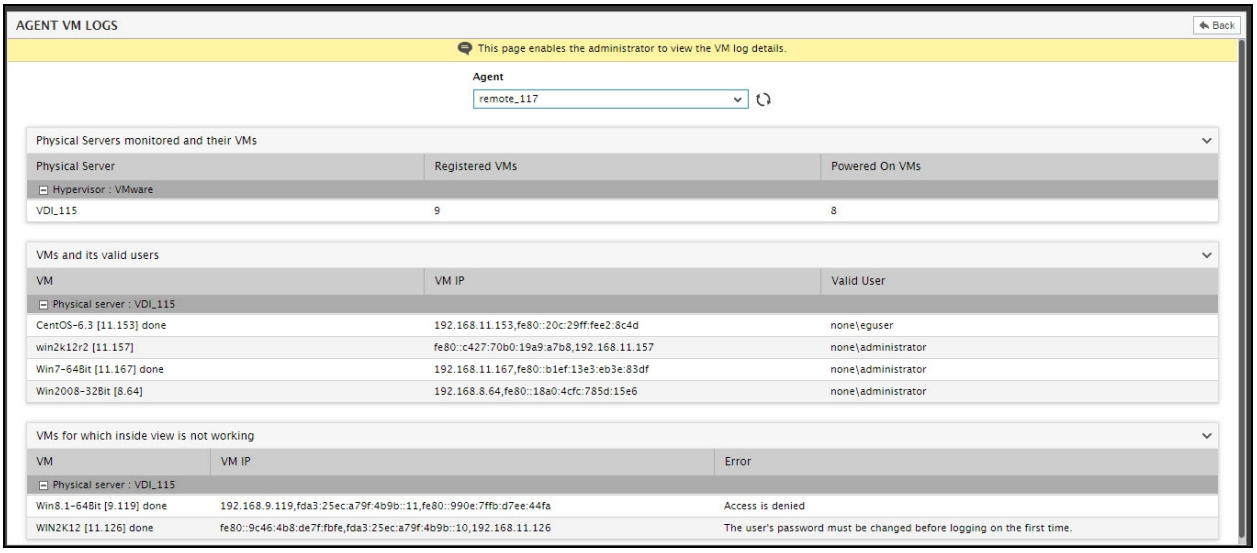

Figure 20.18: Viewing VM Statistics

5. The **Physical Servers monitored and their VMs** section of Figure 20.18 reveals which hypervisors are being monitored by the chosen remote agent. For each hypervisor, this section additionally reports the number of VMs that are currently registered with that hypervisor and the number of VMs that are currently powered on. From this section, administrators can quickly infer how many VMs on which hypervisor are powered-off currently.
6. Next, the **VMs and its valid users** section of Figure 20.18, provides a hypervisor-wise break-up of VMs. In addition, the section also reveals the IP address of every VM and the valid user to that VM. This information will be useful when configuring the eG agent to collect the "inside view" of VMs.
7. Finally, Figure 20.18 includes a **VMs for which inside view is not working** section. This section lists the VMs from which the eG agent was unable to obtain inside view metrics and the reason for the same. This will greatly help administrators troubleshoot the failure of the eG agent to collect inside view metrics.
8. At any point in time, you can click on the **Refresh** button next to the **Agent** list to refresh this page and ensure that the details displayed therein are up-to-date.

20.6 Configuring Auto Upgrade of eG Agents

Upgrades (or patches) to the eG agents add new features and enhancements to the eG product suite. Manual installation of the agent upgrades involves a lot of time, labor and cost, especially in environments comprising of hundreds of agents spanning multiple locations. In order to simplify the process of deploying the agents, eG Enterprise offers the auto upgrade capability. By default, this capability is disabled for all agents. Once it is

enabled, then, the next time the agents check the manager for the existence of an upgrade, the manager will send the upgrade (if any) to the agents. The agent will then install the upgrade automatically.

Figure 20.19 and Figure 20.20 depict how the auto upgrade capability is enabled for an agent. Using these pages administrators can perform the following tasks:

- Enable the auto upgrade capability for specific agents or all of them, as required
- Specify the frequency with which the agents will check the manager for upgrades
- Select the agent that needs to be upgraded immediately

This page (see Figure 20.19) appears when the **Settings** option of the **Upgrade** menu in the **Agents** tile is clicked.

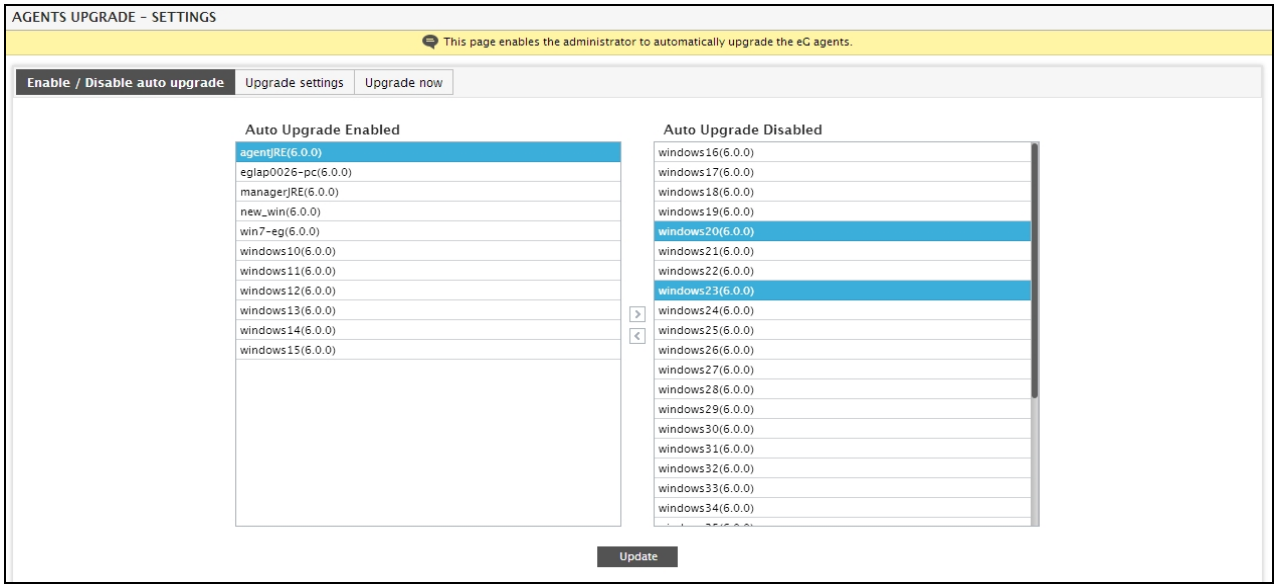

Figure 20.19: Selecting the agent for which the auto upgrading capability is to be enabled

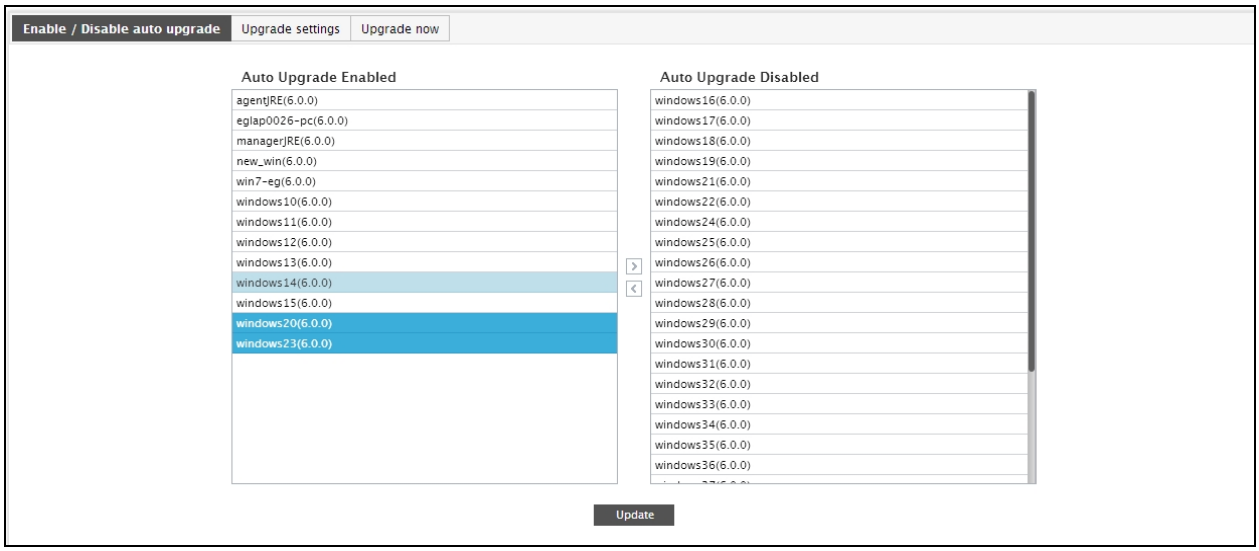

Figure 20.20: Enabling the auto upgrading capability for an agent

To enable the auto upgrade capability for specific agents, do the following:

1. From the **Auto Upgrade Disabled** list in the **Enable/Disable auto upgrade** tab page (see Figure 20.19), select the agent(s) for which the auto upgrade capability is to be enabled.
2. Then, click the **<** button to transfer the selection to the **Auto Upgrade Enabled** list (see 20.6).
3. To disable this capability later, select the agent(s) from the **Auto Upgrade Enabled** list, click the **>** button, and transfer the selection back to the **Auto Upgrade Disabled** list.

To specify the frequency with which the agent should check the manager for upgrades, do the following:

1. Click on the **Upgrade settings** tab page as depicted by Figure 20.21.
2. Then, from the **How often agents should check for auto upgrade package** list box (see), select the time interval at which the agents (for which auto upgrade has been enabled) need to check the manager for upgrades.

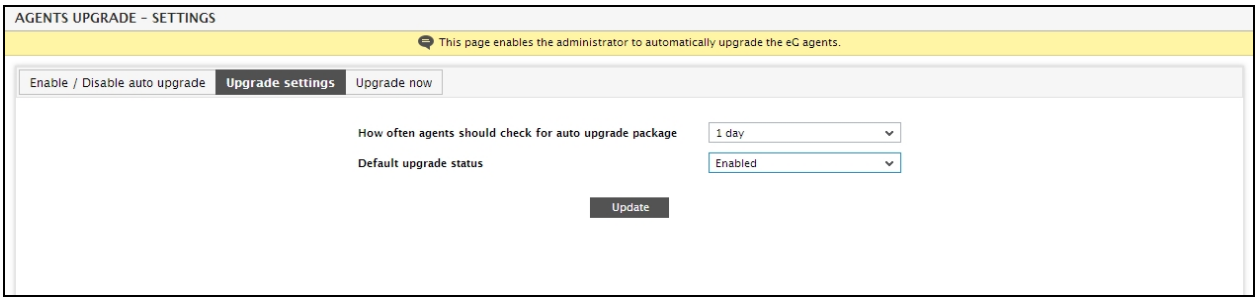

Figure 20.21: Specifying the upgrade interval

3. By selecting the **Enabled** or **Disabled** option from the **Default upgrade status** list box, administrators can indicate whether auto upgrade is, by default, enabled/disabled for new agents to the eG Enterprise system

(see Figure 20.21).

4. Finally, click the **Update** button (see Figure 20.21).

To upgrade agents within the next 15 minutes, do the following:

1. Click on the **Upgrade now** tab page as shown by Figure 20.22.
2. All the agents for which auto-upgrade has been enabled will then appear in the **Auto Upgradeable Agents** list in Figure 20.22.
3. From this list, select the agents which need to be auto-upgraded within the next 15 minutes, and click the **<** button (see Figure 20.22).

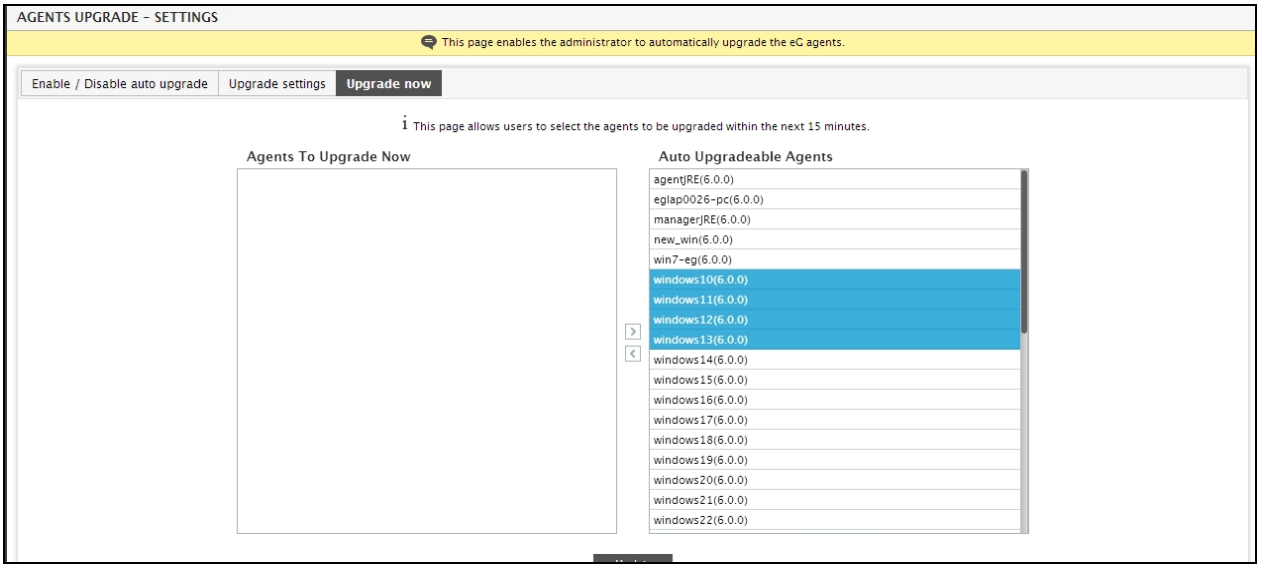

Figure 20.22: Selecting the agents to be upgraded now

4. The selected agent will then be transferred to the **Agents To Upgrade Now** list (see Figure 20.23).

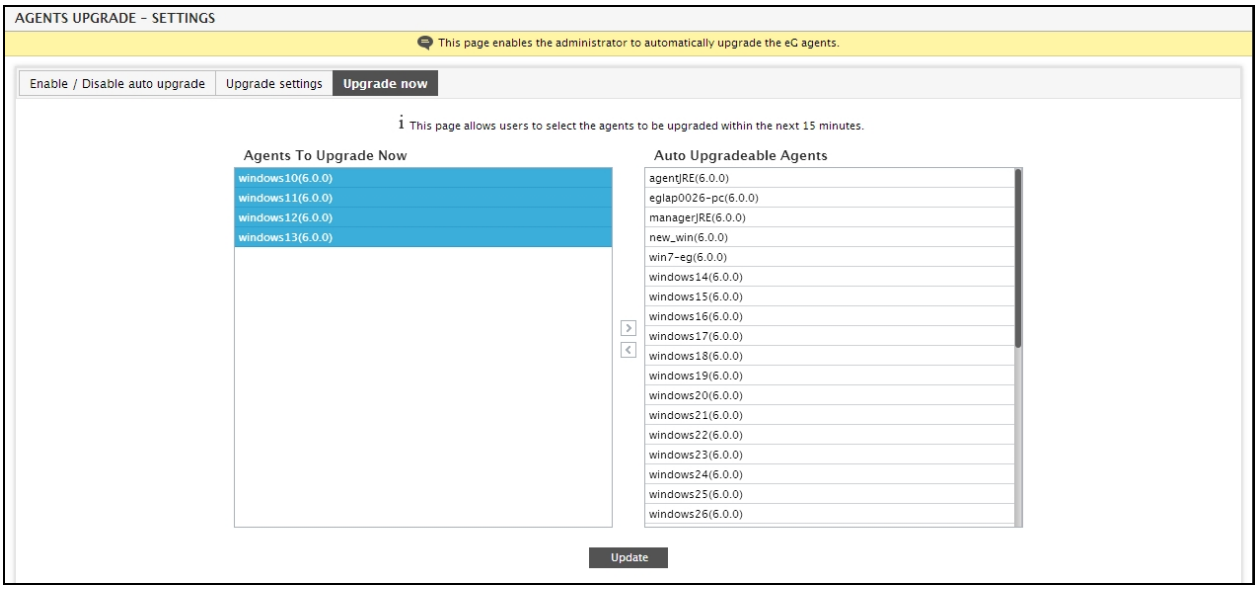

Figure 20.23: The agents for which Upgrade Now has been enabled

5. To disable the **Upgrade Now** capability later, you can select the agents from the **Agents To Upgrade Now** list and click the **>** button.

   Note:

   The **Upgrade now** capability can only be enabled for a maximum of 10 agents, simultaneously.

20.7 Advanced Search Options

In environments comprising of a large number of components, it is often very difficult for administrators to remember which agent manages each of the monitored components, and whether upgrade/upgrade now has been enabled for those agents or not. eG Enterprise therefore, provides a single interface using which administrators can search for and view details of agents based on a given component name, component type, and/or IP address(es). In addition, this interface permits the display of agents based on upgrade status (i.e., whether auto- upgrade/Upgrade now has been enabled/not), and allows administrators to instantly enable/disable upgrade for one/more of the listed agents.

To access this user interface, select the **Advanced Search** menu option from the **Upgrade** menu in the **Agents** tile.

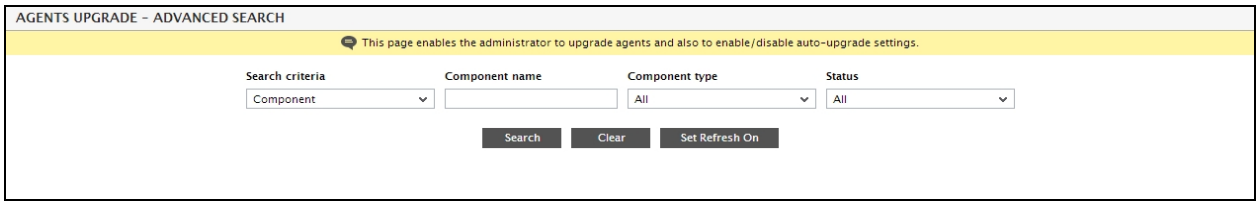

Figure 20.24: The ADVANCED SEARCH page displaying the filter criteria

1. Figure 20.24 appears next.
2. First, pick a **Search criteria**.
3. If you want to view the details of agents monitoring specific components, select **Component** from **Search criteria** and specify a search string in the **Component name** text box.
4. Then, click the **Search** button in Figure 20.25. The result set that appears (see Figure 20.25) displays the details of agents that are monitoring those components which have names that embed the specified string.

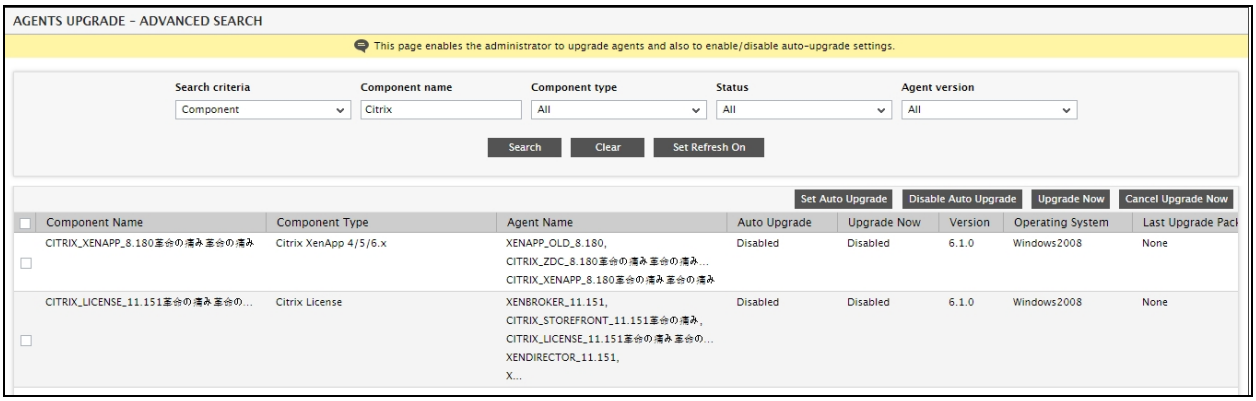

Figure 20.25: Searching based on Component name

5. In the same way, to view the details of all the agents that are monitoring components of a particular type, specify the **Component type**, and click the **Search** button to retrieve the results (Figure 20.26).

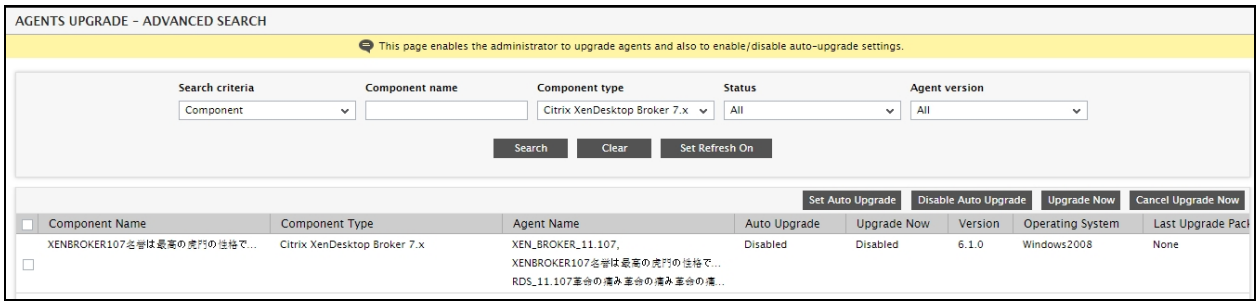

Figure 20.26: Searching based on Component type

6. Similarly, you can also view the details of agents of a particular status, by selecting the desired option from the **Status** list. You can thus choose to view information pertaining to agents for which auto upgrade is disabled/enabled, or upgrade now is disabled/enabled (see Figure 20.27).

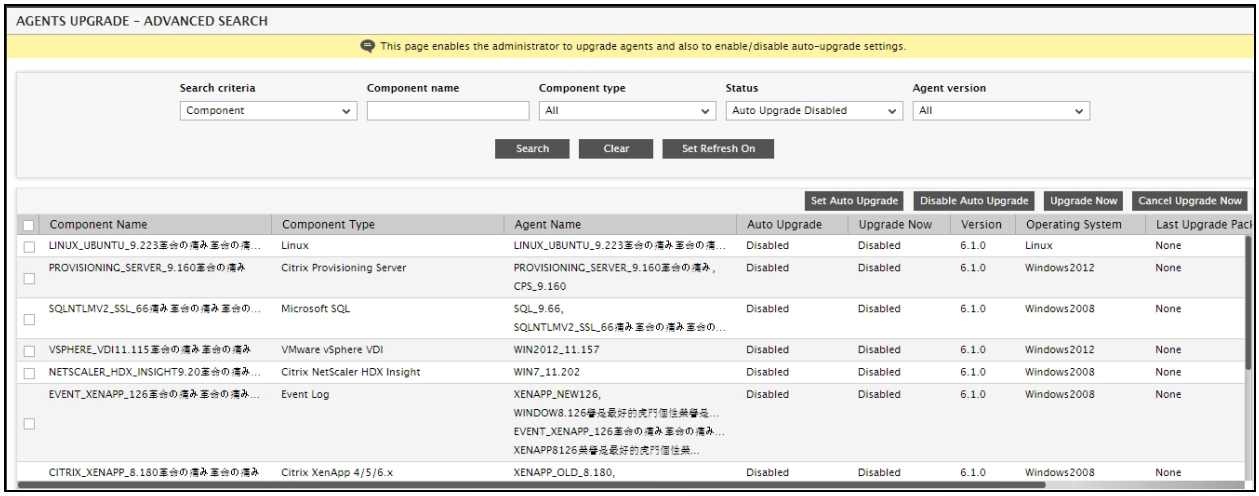

Figure 20.27: Searching based on Status

7. Where multiple versions of eG agents co-exist, you can use this page to identify those agents that are yet to

be upgraded to the latest version. For this, pick an agent version from the **Agent version** drop-down and click the **Search** button. Figure 20.28 will then appear, listing the details of all agents that are of the chosen version.

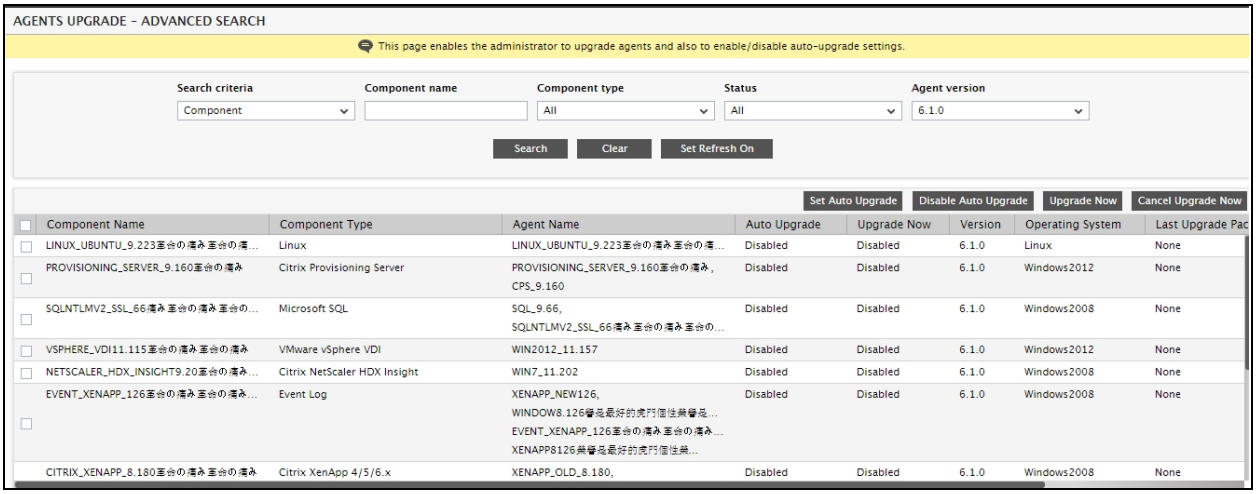

Figure 20.28: Searching based on Agent version

8. Alternatively, a combination of search criteria can also be specified as indicated Figure 20.29.

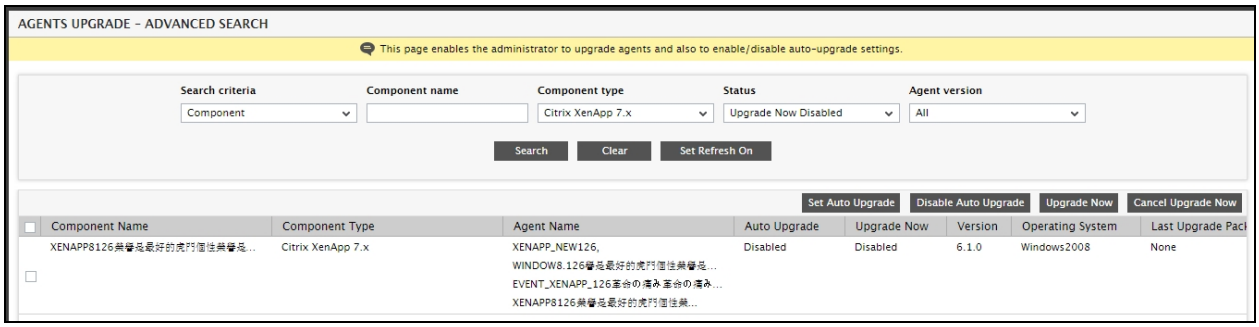

Figure 20.29: Searching based on Component Type and Status

9. Also, instead of filtering your agent-view on the basis of a specific component name, type, or agent status, you can simply provide an IP address or range of IP addresses for which agent information is required. Clicking on the **Search** button then, will display the details of all agents with host/nick names that are associated with the given IP address(es) (see Figure 20.30 and Figure 20.31).

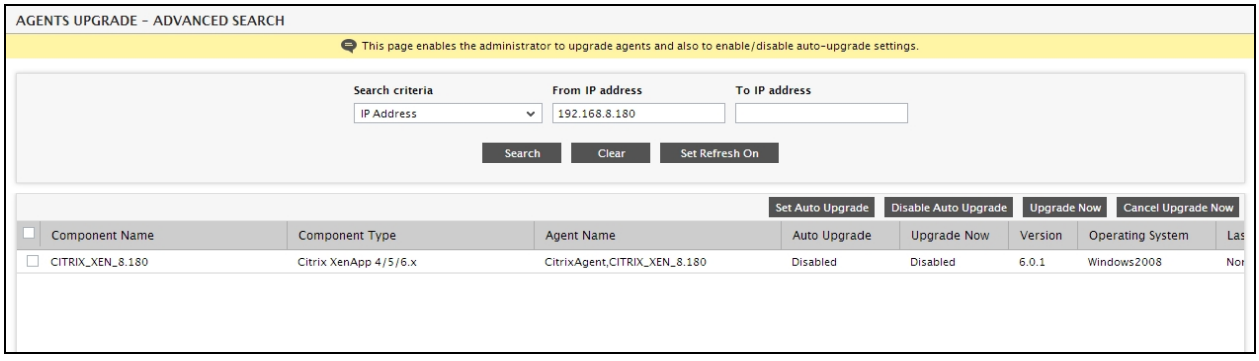

Figure 20.30: Searching based on a single IP address

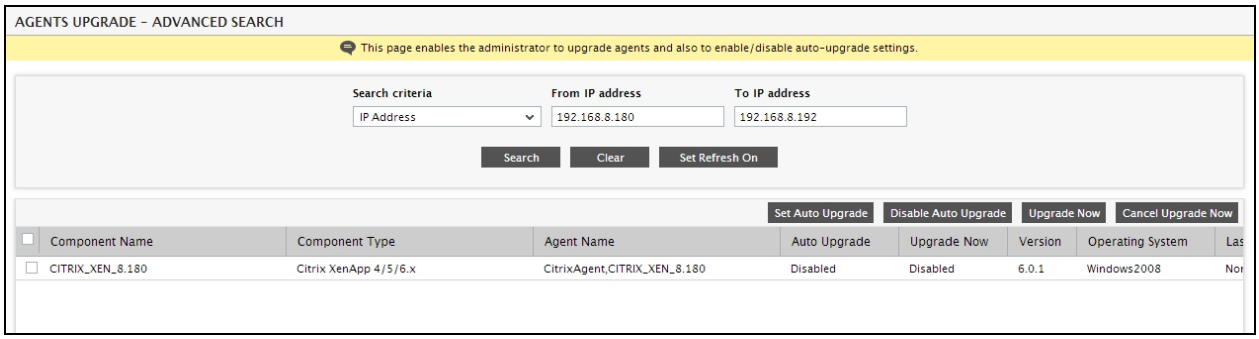

Figure 20.31: Searching based on a range of IP addresses

10. However, regardless of the search criteria specified, the following information is typically retrieved and displayed in the **ADVANCED SEARCH** page:

    - **Component Name**
    - **Component Type**
    - **Agent Name** - All the nick names that map to the IP address of the displayed **Component Name**
    - Whether **Auto Upgrade** and **Upgrade Now** have been enabled for the agent or not
    - The current **Version** of the eG agent
    - The **Operating System** on which the eG agent functions
    - The ID of the **Last Upgraded Package**
    - **Last Upgraded Time** - The time at which the agent was last upgraded

11. By default, the agent information displayed is sorted in the ascending order of the contents of the **Agent name** column. To sort the agent details on any other column, click on the corresponding column heading. For instance, to sort the agent details on the basis of the contents of the **Component type** column, just click on the column label, Component type. This will automatically sort the agent information in the ascending order of the component types displayed in the **Component type** column. An 'up arrow' mark will appear against the Component type column label to indicate the same. Clicking on the 'up arrow' once again will switch the sort order from ascending to descending and will change the 'up arrow' to the 'down arrow' (see Figure 20.33).

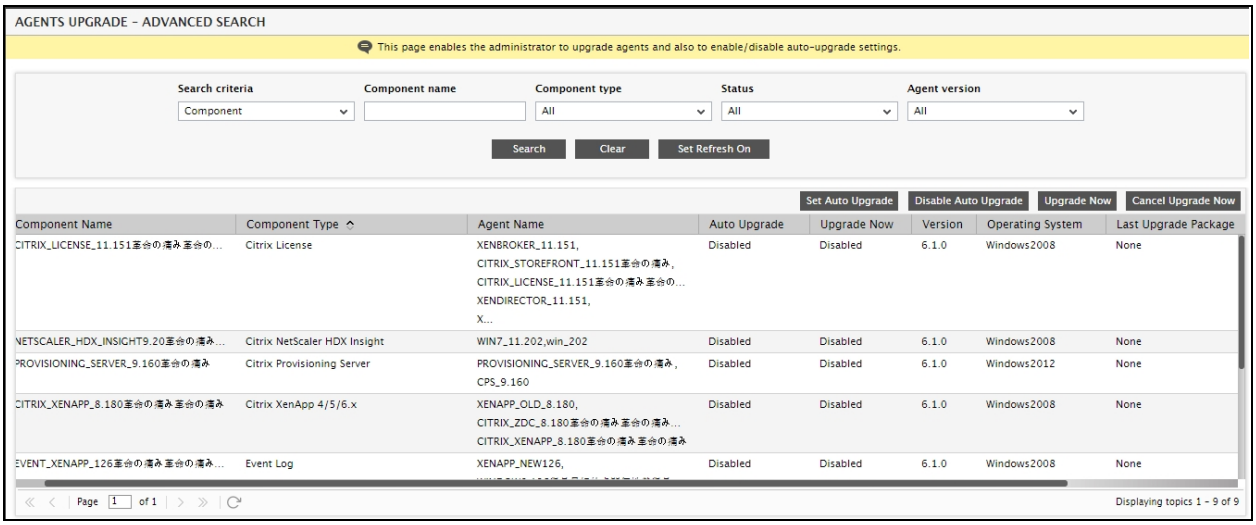

Figure 20.32: Sorting in the ascending order of component types

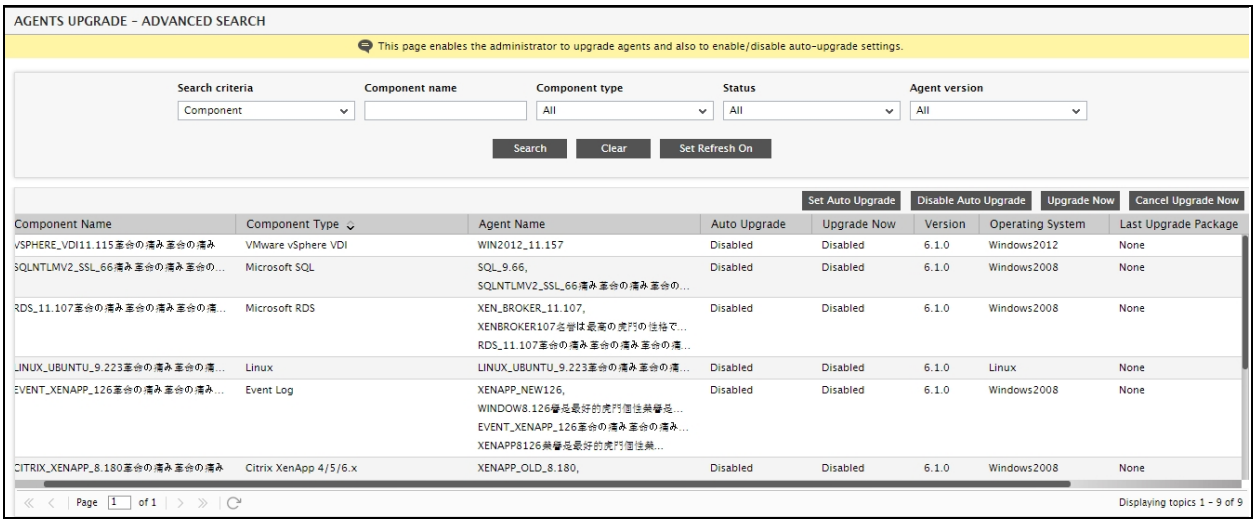

Figure 20.33: Sorting in the descending order of component types

12. At any point in time, you can clear the displayed information by clicking on the **Clear** button in Figure 20.33.
13. As stated earlier, the **ADVANCED SEARCH** page not only provides agent information, but also allows you to enable/disable auto-upgrade or the 'upgrade now' capabilities of the agents. To enable the auto-upgrade capability of multiple eG agents simultaneously, click on the check boxes that prefix every row of information (in the **ADVANCED SEARCH** page) related to these agents as depicted by Figure 20.34, and click the **Set Auto Upgrade** button.

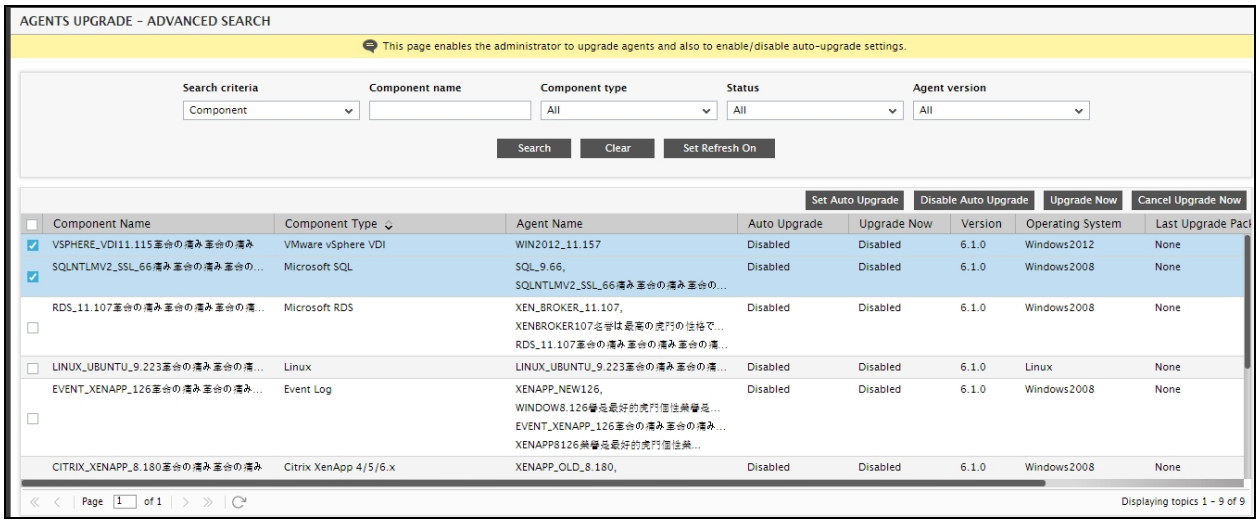

Figure 20.34: Selecting the agents for which auto-upgrade is to be enabled

14. If the auto-upgrade capability was enabled successfully for the chosen agents, then the **Auto Upgrade** column of Figure 20.35 that appears next, will indicate the same.

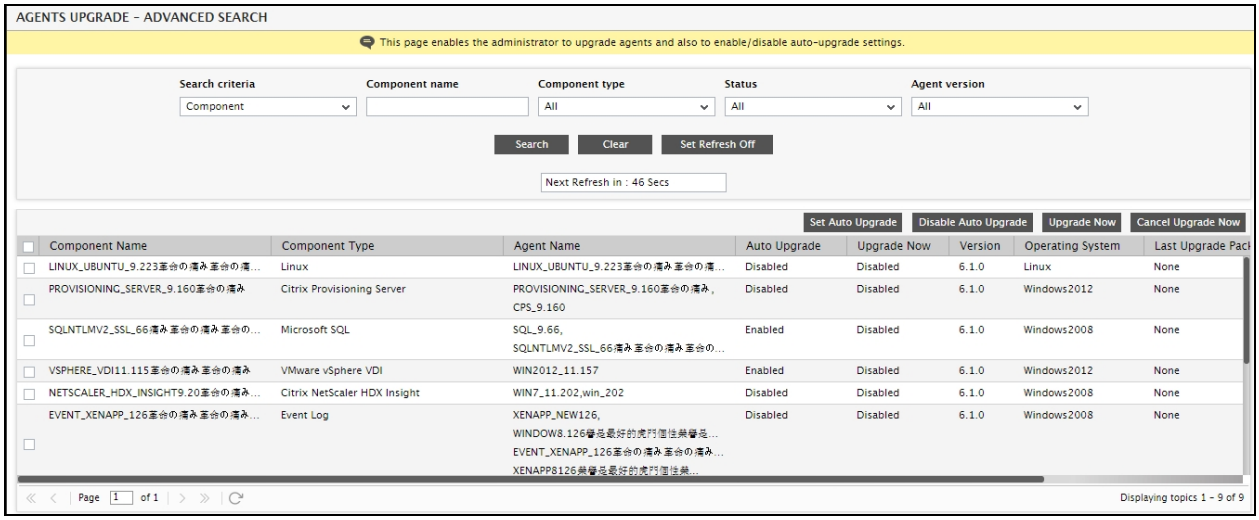

Figure 20.35: Enabling the Auto Upgrade capability

15. Similarly, to disable the auto-upgrade capability, select the check boxes prefixing the corresponding agent details, and click the **Disable Auto Upgrade** button. The **Auto Upgrade** column will then indicate whether the auto-upgrade capability of those agents was successfully disabled or not.

16. Likewise, you can enable/disable the 'Upgrade now' capability of agents by first selecting the check boxes corresponding to the agent information, and clicking the **Upgrade Now** or **Cancel Upgrade Now** buttons (as the case may be). Figure 20.36 and Figure 20.37 indicate the procedure for upgrading a few chosen agents, now (i.e., within the next 15 minutes).

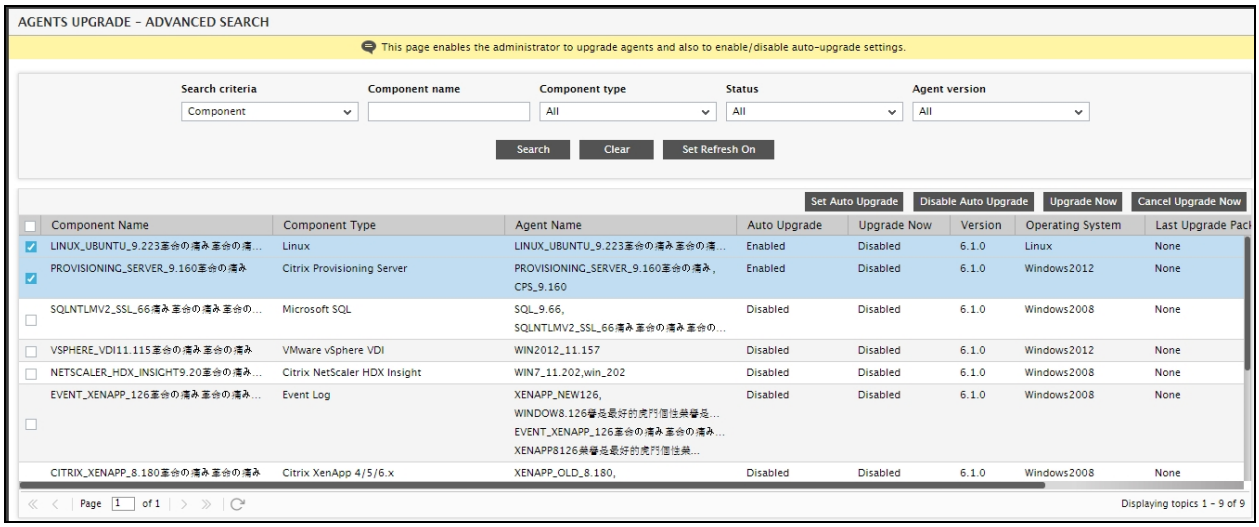

Figure 20.36: Selecting the agents to be upgraded now

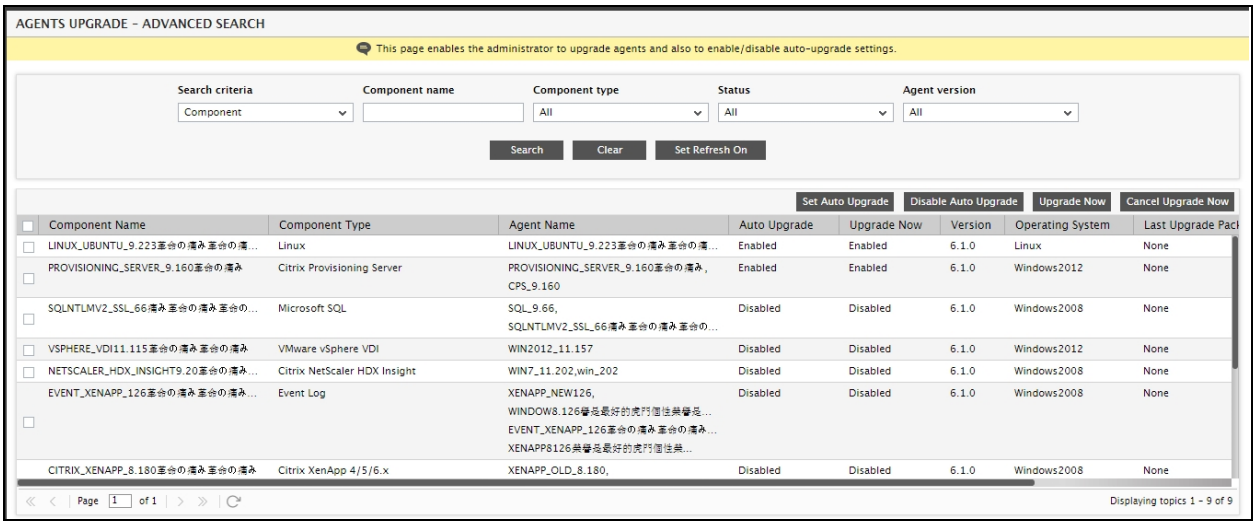

Figure 20.37: Enabling the Upgrade Now capability of selected agents

Note:

- The **Upgrade Now** capability can be enabled only for those agents for which the **Auto Upgrade** capability has been enabled.
- If the **Auto Upgrade** capability is disabled for an agent, the **Upgrade Now** capability of that agent (if previously enabled) will also be automatically disabled.

17. By default, the **ADVANCED SEARCH** page does not refresh automatically. Clicking on the **Set Refresh On** button in Figure 20.38 allows the page to automatically refresh according to a pre-configured refresh period, and also enables administrators to track how long it would be before the next reload occurs (see Figure 20.38).

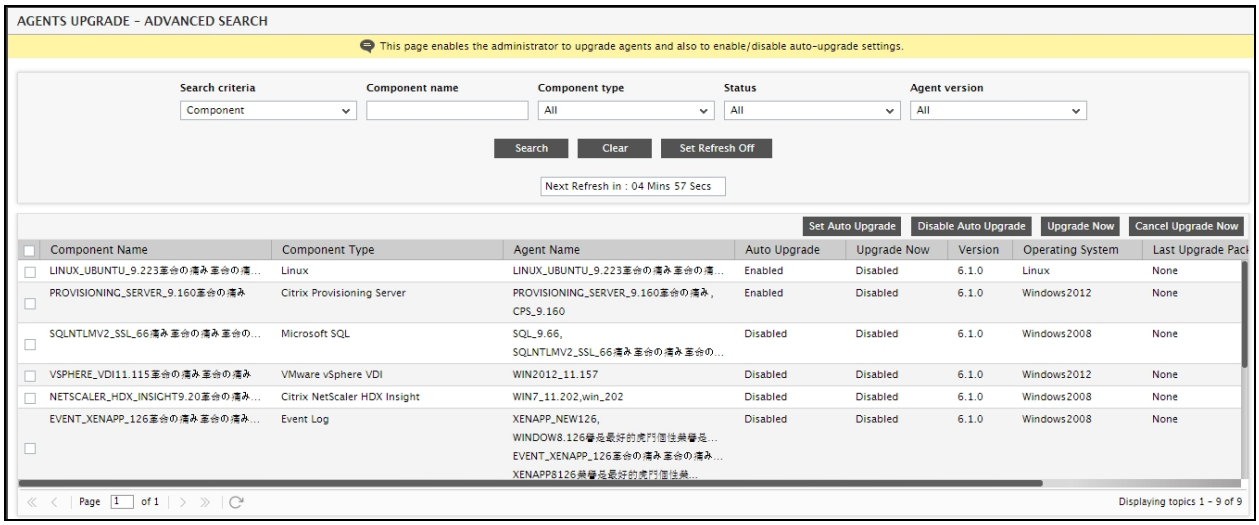

Figure 20.38: Setting refresh on and tracking time to refresh

18. Once the **Set Refresh On** button is clicked, the **ADVANCED SEARCH** page refreshes every 5 minutes (i.e., 300 seconds), by default. You can however, modify the refresh period by editing the **eg_ui.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory. The **AutoUpgrade** parameter in the **[REFRESH]** section of this file is set to 300 (seconds) by default. If need be, this default setting can be overridden. To disable the automatic refresh capability of this page, click on the **Set Refresh Off** button in Figure 20.38.

20.8 Determining the Status of the eG Agents

The eG manager is able to determine and report the operational status of all the eG agents in the target environment. The sections that follow will discuss how to view this status information.

If you select the **Agent Status** option from the **Agents** tile, you will be lead to Figure 20.39 , which will provide status information for agents based on the agent types.

To obtain the status of the eG agents of a particular type, the administrator has to first select the type of agent (whether basic, premium, external, or remote) from the **Agent type** list box.

The IP address / host names of the agents of the selected type will then be displayed. A symbol against each agent indicates that the agent has been deployed. A symbol appears against each agent implying that the agent has not been deployed. While the symbol indicates that the agent is running currently, the **Status** column indicates that the eG agent is not running.

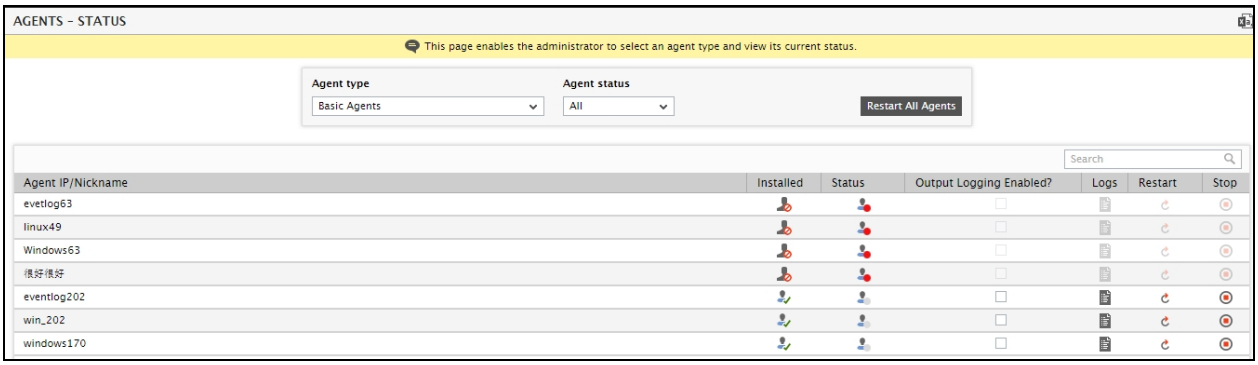

Figure 20.39: Status information for agents

Also, using the **Search** text box, you can find out the status of a particular agent. To know the status of a particular agent, just specify the IP address / host name of that agent in the **Search** text box, and then click the 'magnifying glass' icon next to it. The status of the specified agent will then appear. If the exact IP address / host name of the agent is not known, then a string or a character that features in the IP / host name of the agent can be provided in this text box (see 20.8). Multiple search conditions can be specified as a comma-separated list.

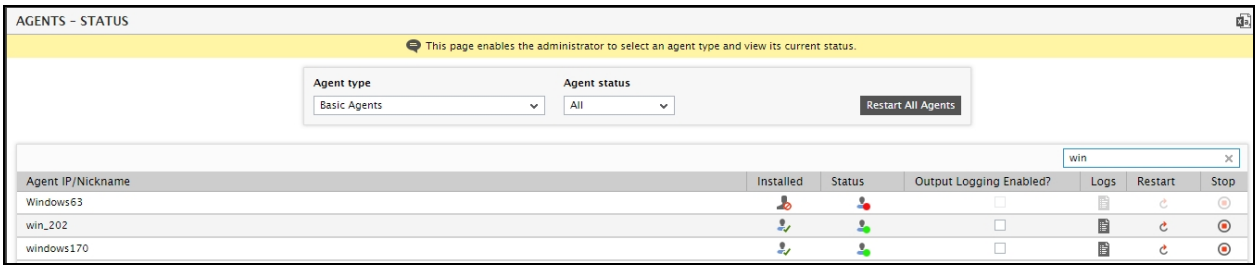

Figure 20.40: Searching for agent status

To know the agents that are currently in a particular state, simply select an **Agent status** (which can be Running/Not Running/All). The default selection here is *All*.

You can even remotely initiate an agent-restart, by simply clicking on the **Restart** icon that corresponds to an agent. To restart all agents, click on the **Restart All Agents** button in Figure 20.40. Doing so immediately sends out restart requests to all the agents that are currently running and reporting metrics to the eG manager. If an agent is not running currently, then the eG Enterprise system sends out the restart request soon after that agent starts running.

You can also stop an eG agent from the eG management console itself, using this page. Click on the **Stop** icon corresponding to an eG agent in Figure 20.39, to stop that agent. Doing so immediately sends out a stop request to that agent. The next time that agent sends metrics to the manager, it reads the stop request and serves it.

If an administrator needs to be alerted upon login, about agents that are not reporting measures to the manager, then do the following:

1. Open the **eg_services.ini** file in the **<EG_HOME_DIR>/manager/config** directory.
2. In the **[MISC_ARGS]** section, set the **AlertAgentsNotRunning** flag to **Yes** (default is **No**).
3. Once this is done, the next time the administrator logs into the admin interface, a message listing the agents that are not running will be displayed.

   Note:

   An eG agent can be configured to run specific tests once a day or once every few hours. You can configure the eG manager to exclude tests that are infrequently run when it determines whether an agent is running or not. To do this, modify the value of **NotReportingCutoffFactor** in the **[MISC_ARGS]** section of the **eg_services.ini** file. By default, tests running with measure period of greater than 20 minutes are not considered by the eG manager for determining if an agent is running or not.

4. Also, by default, output logging is disabled for the eG agents configured in an environment. The eG Enterprise system allows you to enable output and error logging for a specific agent from the eG administrative interface itself, thereby saving you the trouble of running the **debugon.bat** file to achieve the same. When output logging is enabled, an **agentout.log** file is created in the **<EG_INSTALL_DIR>\agent\logs** directory to which details of the tests run and measures reported by that agent to the manager are recorded. To enable output logging for an agent, set the **Output Logging Enabled?** flag for that agent to **ON**. When you attempt to enable output logging, a message box shown by Figure 20.41 will appear, requesting your confirmation to enable output logging for that agent.

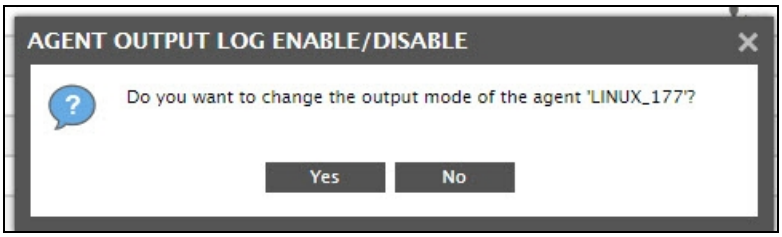

Figure 20.41: A message box requesting your confirmation to enable output logging

5. Click on the **OK** button in the message box to enable output logging or the **Cancel** button to disable it. You can then click on the **LOGS** icon that corresponds to an agent in to view both error logs and output logs related to that agent. Clicking on the **LOGS** icon corresponding to that agent will lead you to Figure 20.42, where the contents of the **error_log** of the corresponding agent can be viewed by default.

   Note:

   If you have turned on output logging for an eG agent using the **AGENTS – STATUS** page, then you should not turn off output logging for that eG agent by manually running the **debugon.bat** file. Likewise, if you have turned on output logging for an eG agent by running the **debugon.bat** file, then you should not turn it off using the **AGENTS – STATUS** page.

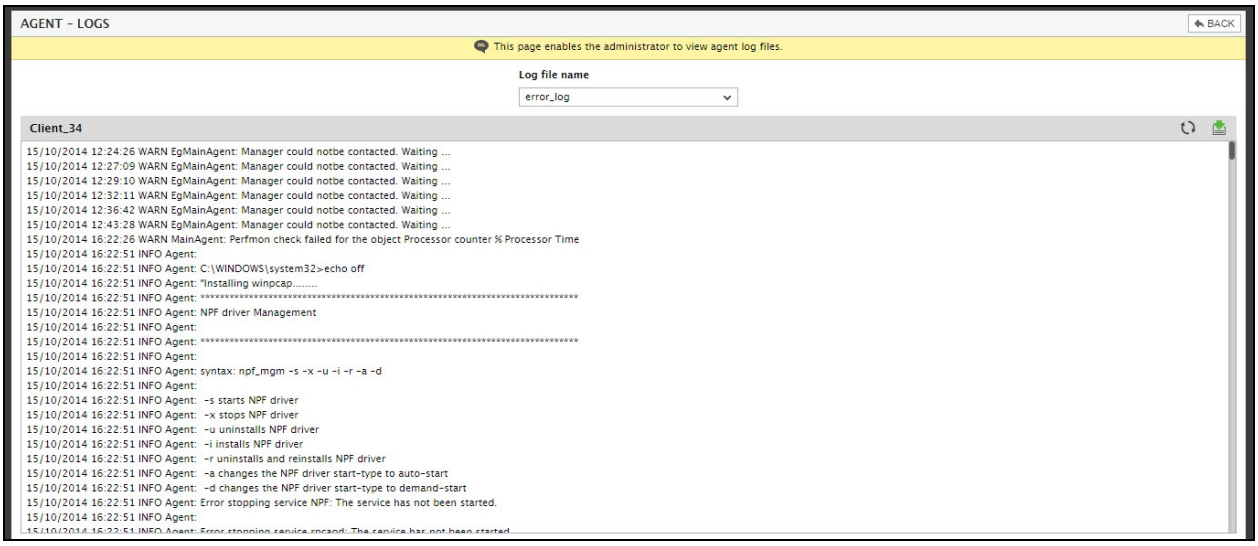

Figure 20.42: Viewing the error_log of an agent

6. You can pick any log file from the **Log file name** list to view its contents (see Figure 20.43).

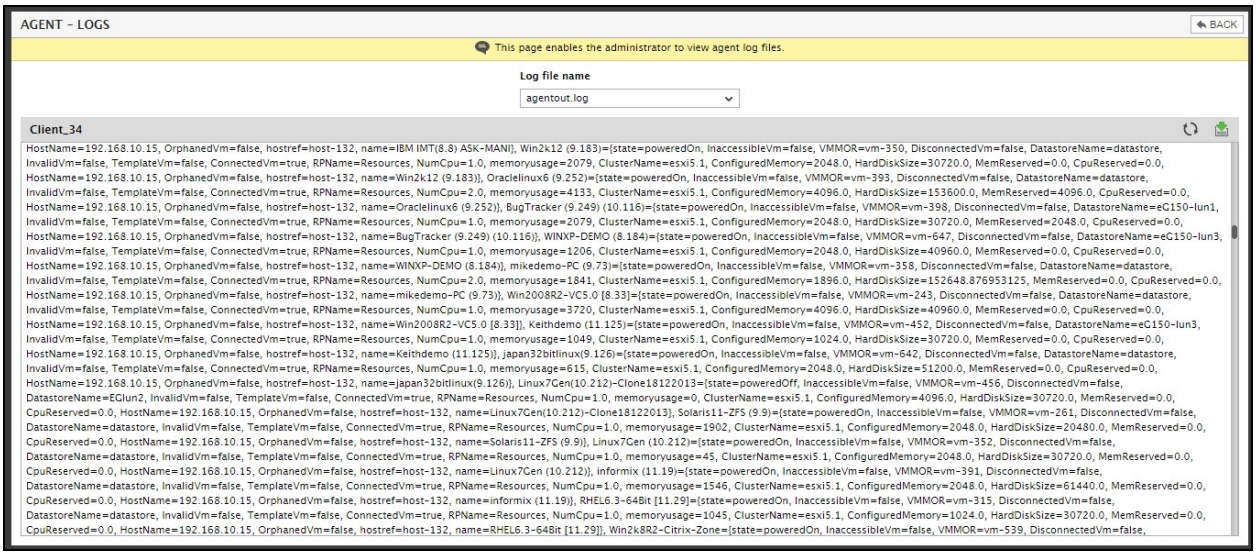

Figure 20.43: Viewing a different log file

7. At any given point in time, you can click the **Refresh** button at the right, top corner of the area where the log file contents are displayed to refresh the contents of the log file. This way, you can make sure that the log file you are viewing is up to date.
8. You can also click on the **Download** button next to the **Refresh** button to download the chosen log file.
9. Clicking on an agent displayed in Figure 20.39 will lead the users to an **Agent Information** page (see), which provides some agent-related information. This includes:
   - The **Agent IP/Nickname**
   - An indicator as to whether the auto upgrade capability has been **Enabled** for that agent, or **Disabled**

- The ID of the last upgraded package (if any) (if no upgrading has occurred, then this will be 'None")
- The date and time at which the agent was last upgraded
- The **HostName** of the agent
- The operating system on which the agent is executing
- The current version of the agent
- The date and time at which the agent last updated the manager with configuration changes
- A **Reset** button
- A **Restart Agent** button

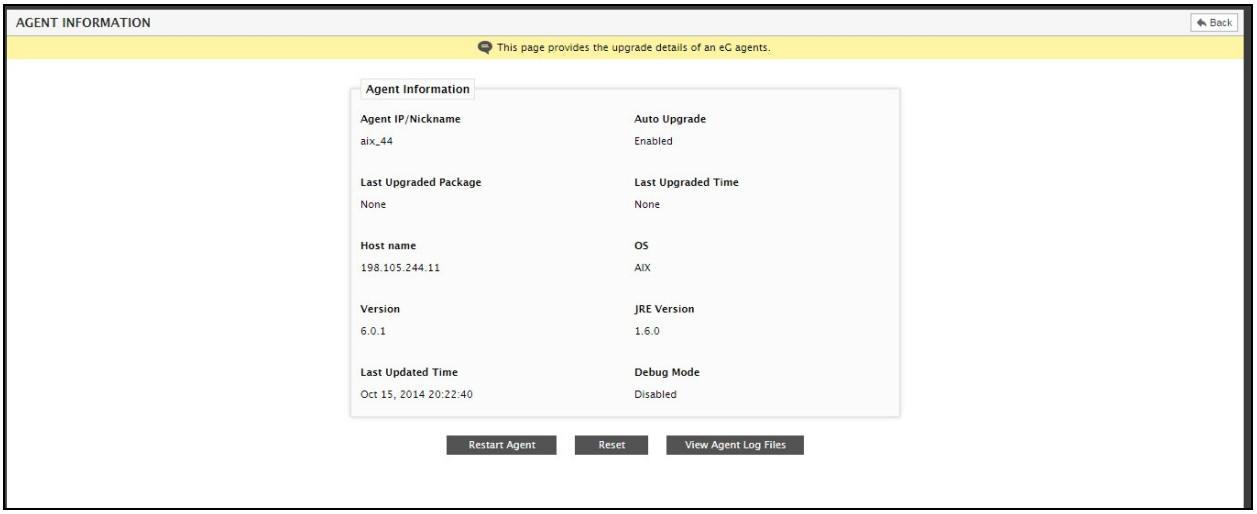

Figure 20.44: A page displaying the upgrade information of an agent

10. Once an agent is upgraded, information regarding the upgraded package will be registered with the manager. Figure 20.44 provides that information. Now, the next time the agent requests for an upgrade, the manager checks whether any newer upgrades are available. If any such upgrade is found, it sends the same to the agent. If for some reason the information pertaining to the last upgrade has to be cleared from the agent's upgrade history, then click on the **Reset** button. This ensures that the details of the last upgrade are lost, and helps the agent download the last upgrade once again from the manager. To restart the agent, click on the **Restart Agent** button in 20.8. To view agent logs, click on the **View Agent Log Files** button in 20.8.
11. Moreover, if the **Agent type** chosen from Figure 20.39 is **External Agents** or **Remote Agents**, then you will also be able to view the count of hosts assigned (if any) to each external/remote agent. For this, you will have to click on the '+' button that pre-fixes an agent (see Figure 20.45). Beneath the assigned host count, you can see that Figure 19.35 also reveals which specific hosts have been assigned to that agent. From a single glance therefore, you can precisely identify the external/remote agents that are been actively utilized, and those that are not.

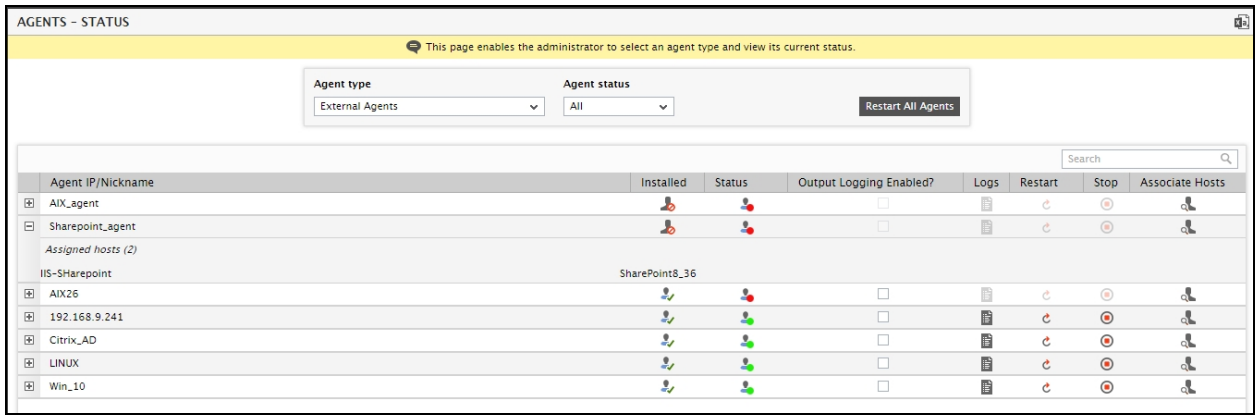

Figure 20.45: Viewing the status of external agents

12. To add more hosts to an external/remote agent, click on the button corresponding to an agent in Figure 20.45. This will open Figure 20.46, using which you can assign more hosts to the agent or disassociate existing hosts from it.

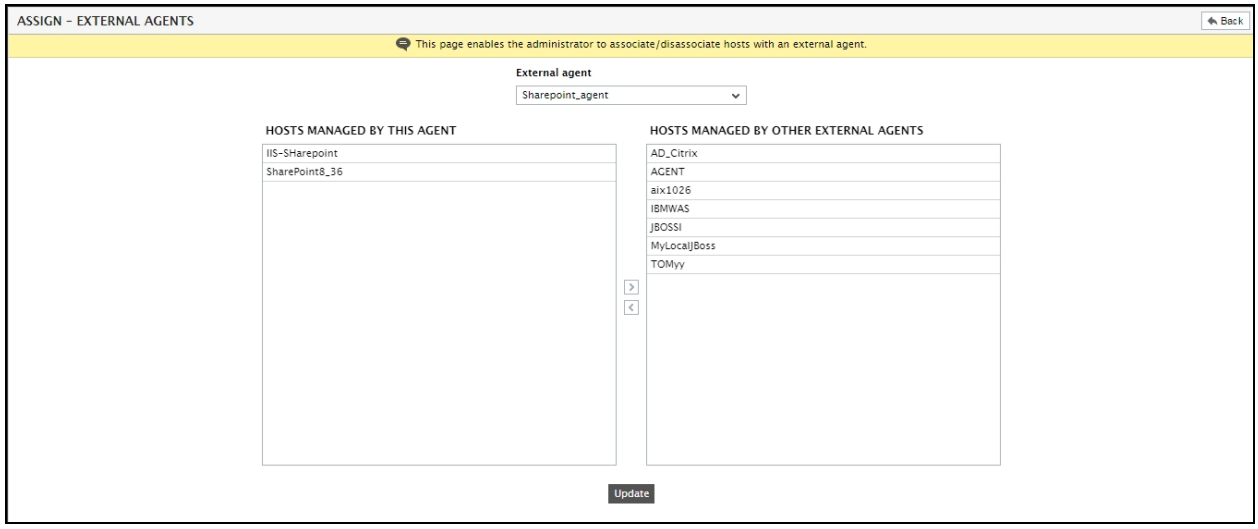

Figure 20.46: Associating/Disassociating hosts from an external agent

20.9 Viewing the Upgrade Status

Clicking on the **Status** option in the **Upgrade** menu of the **Agents** tile will open Figure 20.47 that reveals the following information indicating the upgrade status of every agent reporting to a manager:

- The IP/hostname of the agent
- The unique package id of the last upgraded package of the agent
- The time of upgrade
- Whether upgrade is currently disabled or enabled for the agent
- The operating system on which the agent executes

- The current version of the agent
- The version of JRE used by the agent

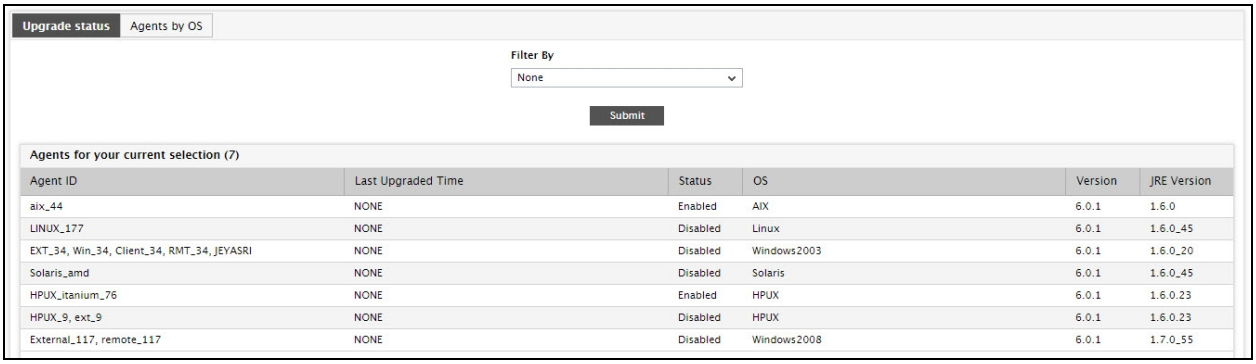

Figure 20.47: Viewing the upgrade status of all agents

To view the upgrade status selectively, choose a **Filter By** option. By default, **None** (see Figure 20.47) will be selected in this list box. Besides this, the list box offers the following filtering options:

- To view the upgrade information pertaining to agents of a particular version (see Figure 20.48), select the **Version** option from the Filter By list box. From the **Filter Criteria** list box that appears next, select a particular version number, and finally, click the **Submit** button.

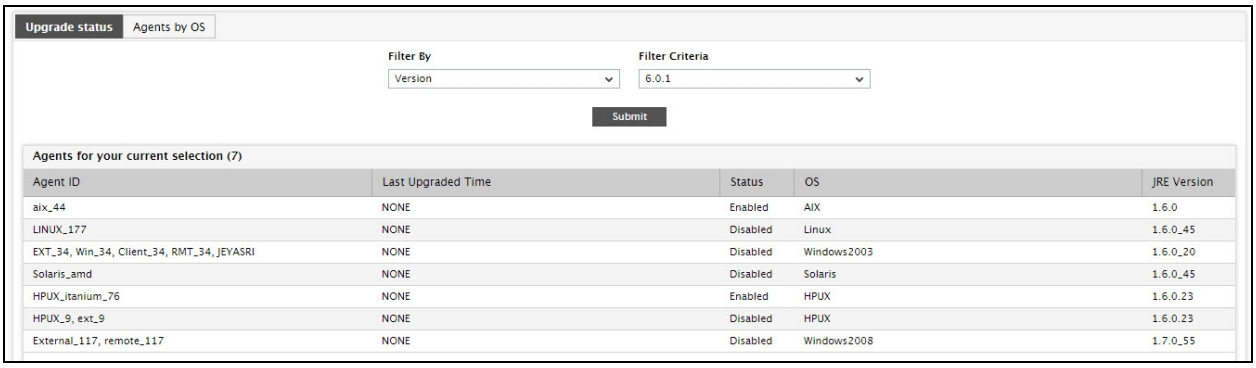

Figure 20.48: Viewing the upgrade status of agents of a specific version

- To view the upgrade information of agents executing on a specific operating system (see Figure 20.49), select the **Operating system** option from the Filter By list box. From the **Filter Criteria** list box that appears next, select a particular operating system, and finally, click the **Submit** button.

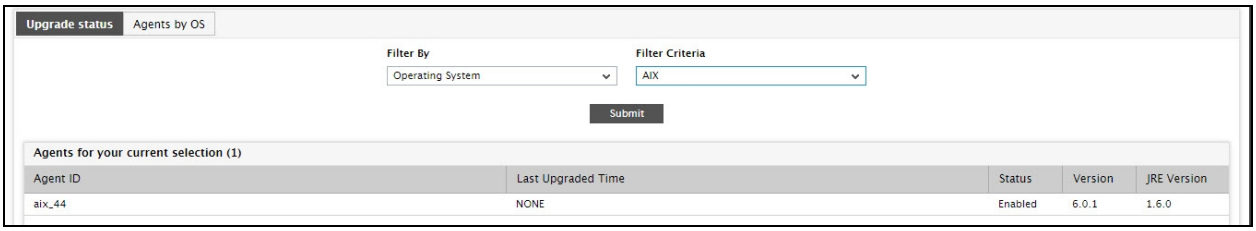

Figure 20.49: Viewing the upgrade status of agents executing on a particular operating system

- To view the agent upgrade status based on the upgrade setting (i.e. whether enabled/disabled) (see Figure 20.50), select the **Upgrade setting** option from the **Filter By** list box. From the **Filter Criteria** list box that appears next, select either **Enabled** or **Disabled**, and finally, click the **Submit** button.

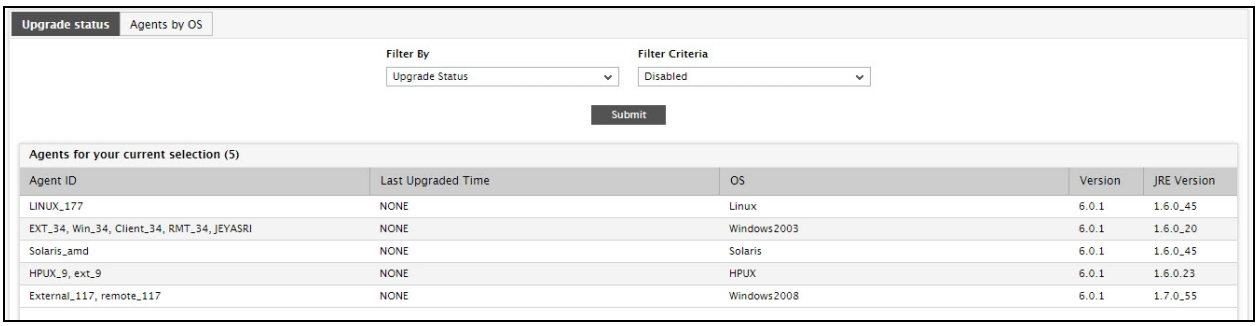

Figure 20.50: Viewing the upgrade status of agents with a specific upgrade setting

- To view the agent upgrade status based on the JRE version, select the **JRE version** option from the **Filter By** list box (see Figure 20.51). From the **Filter Criteria** list box that appears next, select the JRE version to search for, and click the **Submit** button.

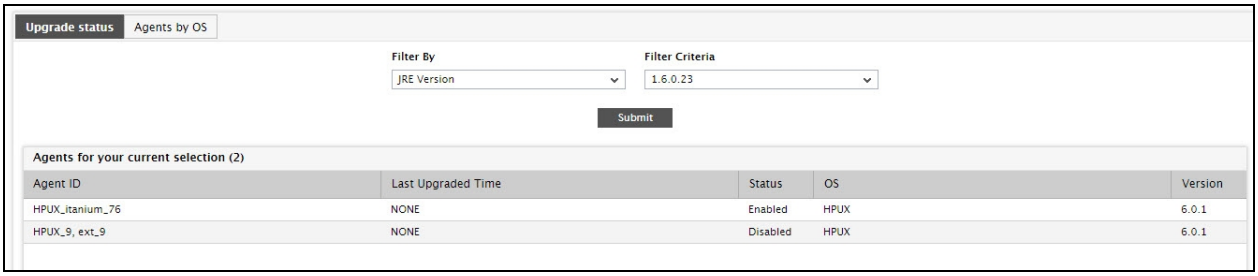

Figure 20.51: Viewing the upgrade status of agents with a specific JRE version

- If you click on the **Agent by OS** tab page in Figure 20.51 you will also view a brief summary of the number of agents executing on every OS (see Figure 20.52).

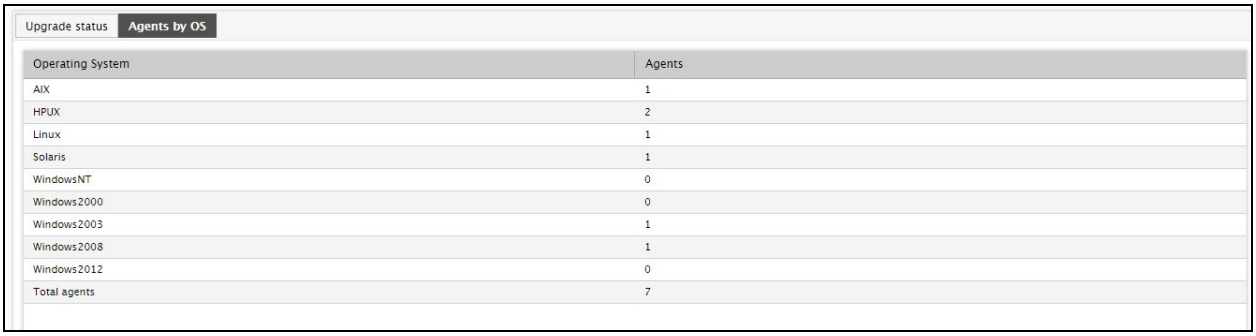

Figure 20.52: Agent summary by OS

20.10 Viewing the eG Manager Logs

When attempting to setup the eG manager, a number of log files get created in the **<EG_INSTALL_DIR>\manager\logs** directory to which error/warning/general status messages related to the installation,

configuration, control, and operations of the eG manager are logged. When faced with slowdowns/errors, a look at these logs may provide you with useful pointers/hints to the source of the problem. To enable you to view these manager logs online and troubleshoot issues on-the-fly, the eG administrative interface provides a **MANAGER - LOGS** page. To view these logs, do the following:

1. Select the **Manager Logs** menu option from the **Miscellaneous** tile.
2. Figure 20.53 will appear displaying the contents of the **checkmgr_log** file in the **manager/logs** directory by default.

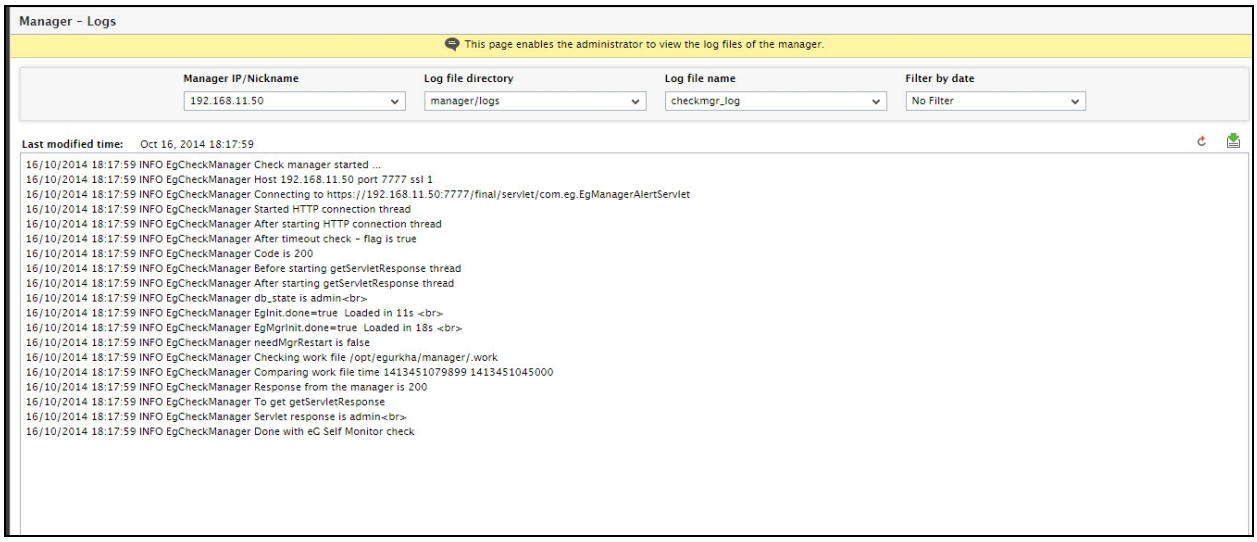

Figure 20.53: Viewing the contents of the checkmgr_log'

3. You can even view the contents of a log file on a particular date by selecting a **Filter by date** (see Figure 20.54).

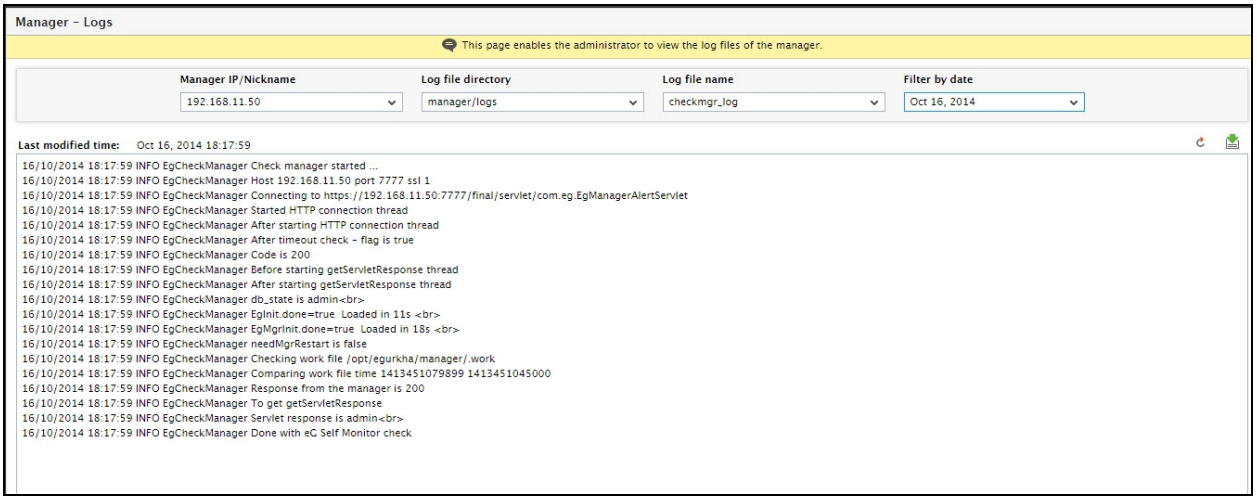

Figure 20.54: Viewing the contents of a log file on a particular date

4. To view the contents of a different log file in the same directory, pick a different option from the **Log file**

name list.

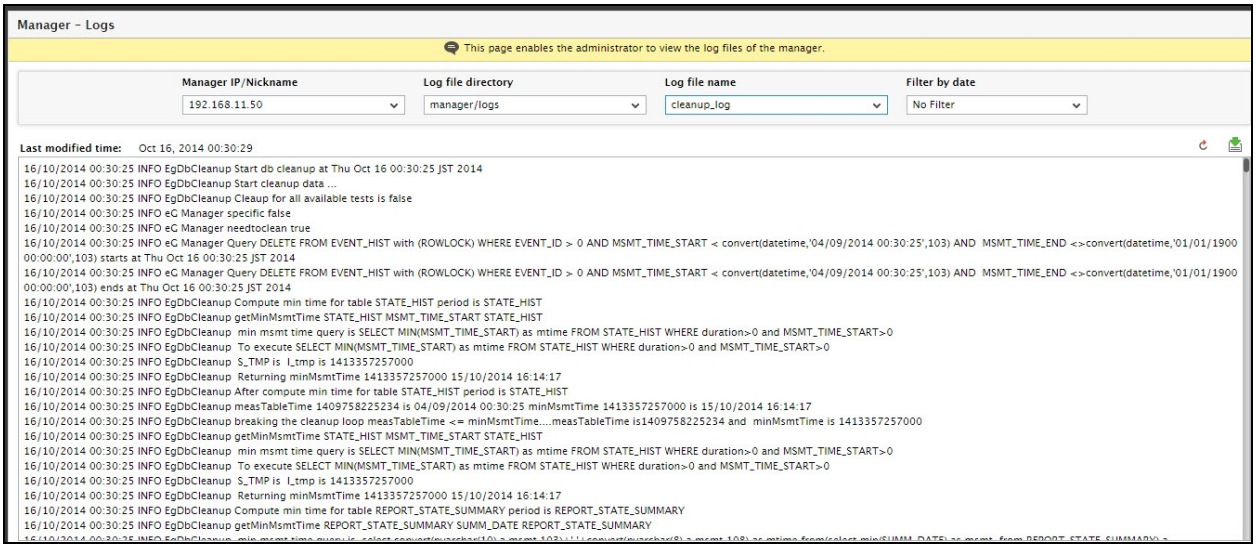

Figure 20.55: Selecting a different log file from the same directory

5. Like the **manager/logs** directory, you can even view the contents of log files in the **tomcat/logs** directory using the **MANAGER - LOGS** page. For this, select the **tomcat/logs** file from the **Log file directory** list and then pick a log file of your choice from the **Log file name** list. Doing so will display the contents of the chosen log file as depicted by Figure 20.56.

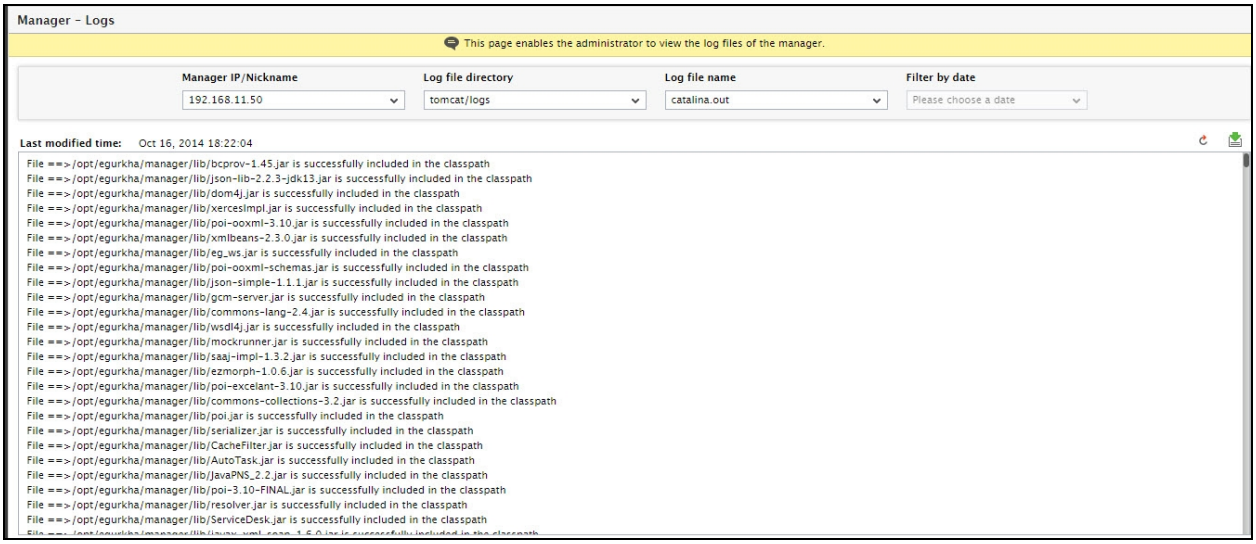

Figure 20.56: Viewing the contents of a log file in the tomcat/logs directory

6. At any given point in time, to refresh the contents of the log file, you can click the **Refresh** button at the right, top corner of the panel where the log file contents are displayed. This way, you can make sure that the log file you are viewing is up to date.

Note:

In case of a redundant setup, by default, the **MANAGER - LOGS** page displays the contents of the primary manager's log files only. To view the secondary manager's log file as well, select the IP address of the secondary manager from the **Manager IP/Nickname** list in Figure 20.56.

21

Conclusion

eG Enterprise of products has been specially designed keeping in mind the unique requirements of IT infrastructure operators. For more information on the eG family of products, please visit our web site at www.eginnovations.com.

This document has described the administration, and usage of eG Enterprise that enables IT infrastructure operators monitor their web infrastructure efficiently and effectively. It has gone a long way in clarifying concepts in various aspects of using eG Enterprise.

For more details regarding the eG architecture and the details of the metrics collected by the eG agents, please refer to the following documents:

- The eG Installation Guide
- Monitoring Using the eG Enterprise Suite
- The eG Reporter

We recognize that the success of any product depends on its ability to address real customer needs, and are eager to hear from you regarding requests for enhancements to the products, suggestions for modifications to the product, and feedback regarding what works and what does not. Please provide all your inputs as well as any bug reports via email to support@eginnovations.com.