Monitoring Using eG Enterprise Suite

Restricted Rights Legend

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

Trademarks

Microsoft Windows, Windows 2008, Windows 2010, Windows 2012, Windows 2016, Windows 7, Windows 8 and Windows 10 are either registered trademarks or trademarks of Microsoft Corporation in United States and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

Copyright

©2016 eG Innovations Inc. All rights reserved.

Table of contents

Table of Figures

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.

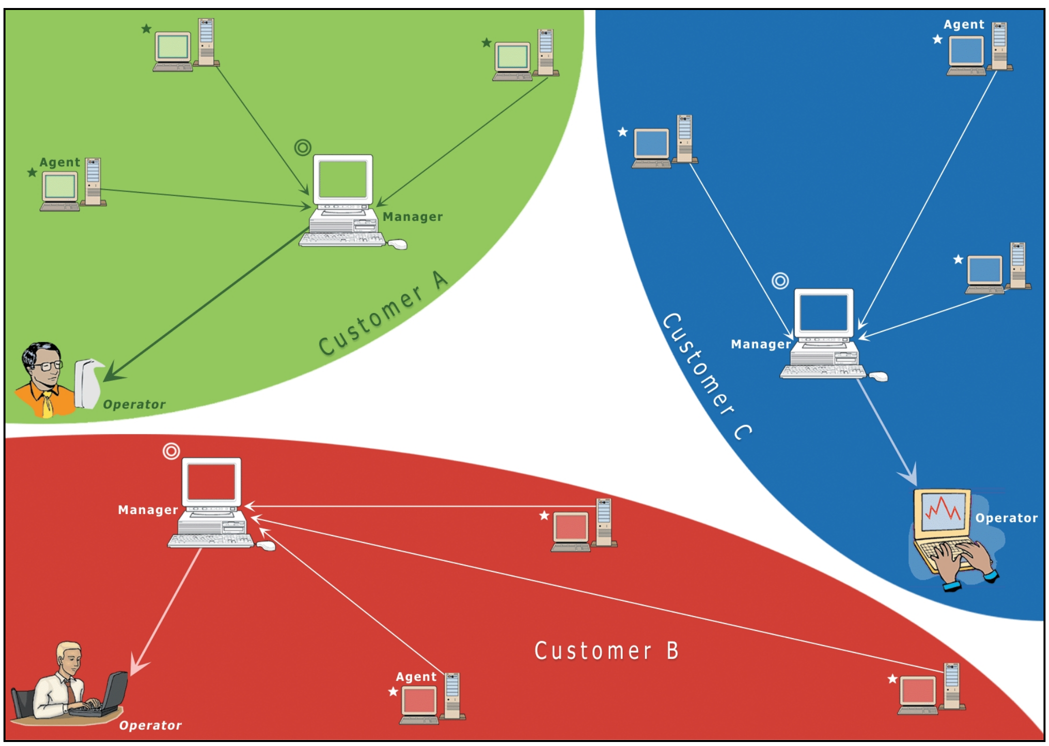

Figure 1.1: Retrofitting existing monitoring solutions for hosted environments

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 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.

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	2003, 2008, XP, Vista, 7, 8, 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

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**
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**
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**
Firewalls	**Check Point Firewall-1, Cisco Pix, Netscreen Firewall, FortiGate Firewall, Microsoft Forefront TMG, Sonic Firewall, WatchGuard Firewall**
Terminal servers	**Citrix MetaFrame 1.8, XP server, 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**
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**

Component Type	Component Brand
SAP	**SAP ABAP server, SAP Internet Transaction server (ITS), SAP Web Application server, MaxDB**
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**
Connection Brokers	**Virtual Desktop Manager, Leostream connection broker, Xen Desktop Broker, VMware 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, InfoBlox**
Other Operating Systems	**Generic SNMP, Generic Netware, AS400 server, OpenVMS server, Mac OS**
Siebel Enterprise	**Siebel Web Server, Siebel Application Server, Siebel Gateway**

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.

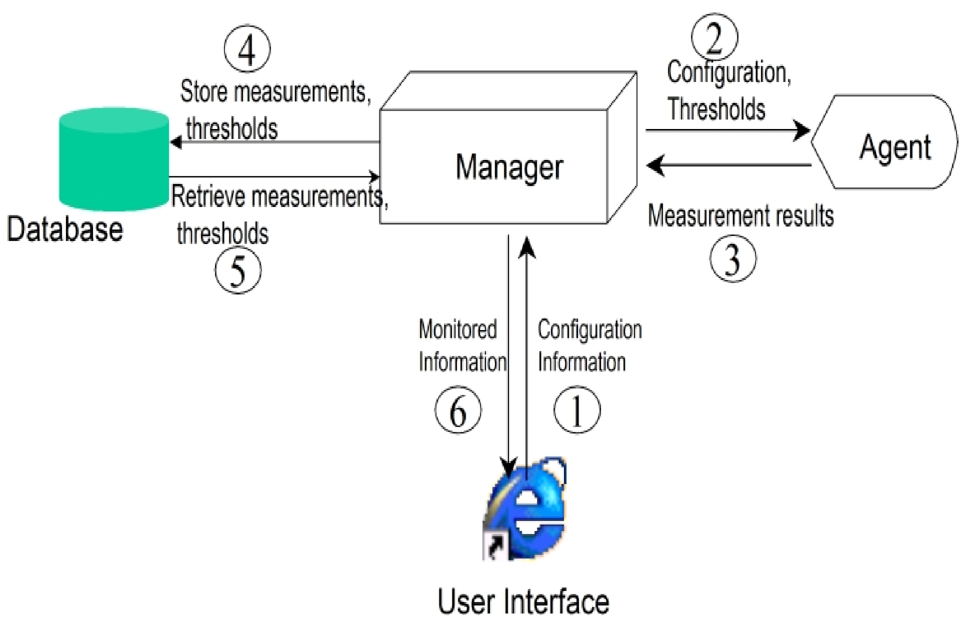

Figure 2.1: Main components of eG Enterprise

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.

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. 2.1 depicts the virtual private manager architecture.

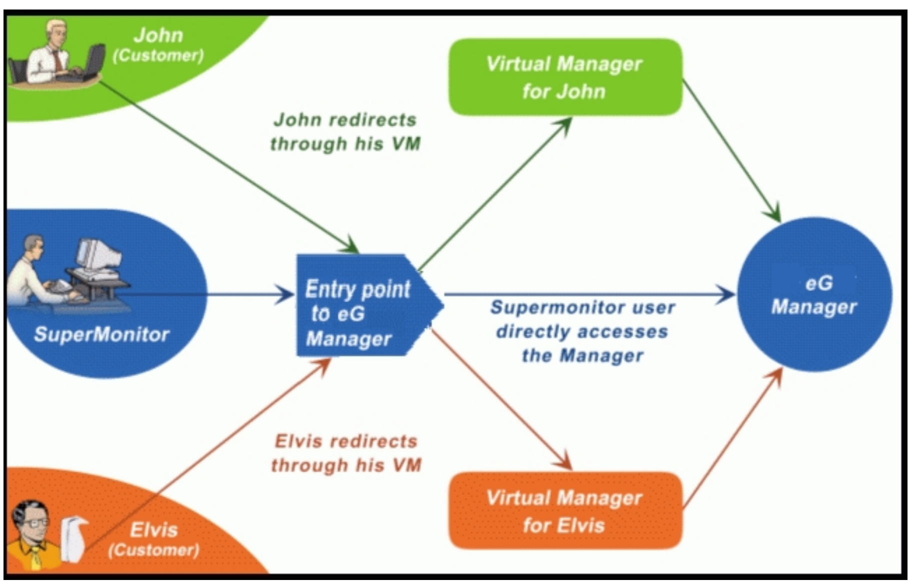

Figure 2.2: Architecture of the virtual, private manager

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*.

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.

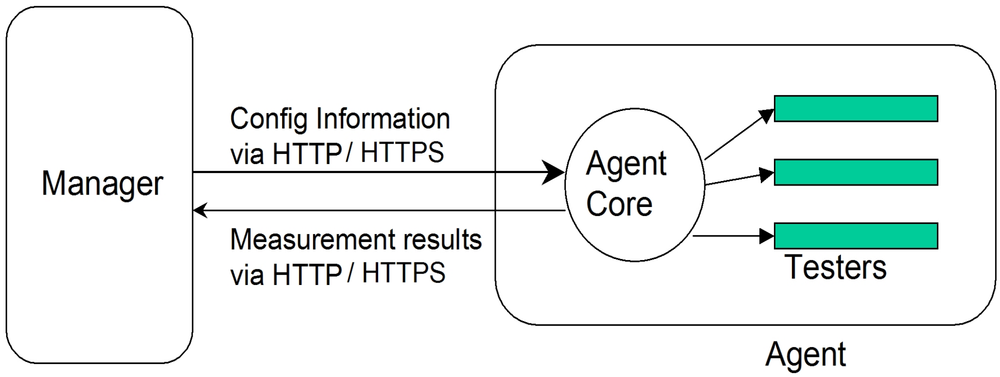

Figure 2.3: The manager-agent communication in the eG architecture

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.

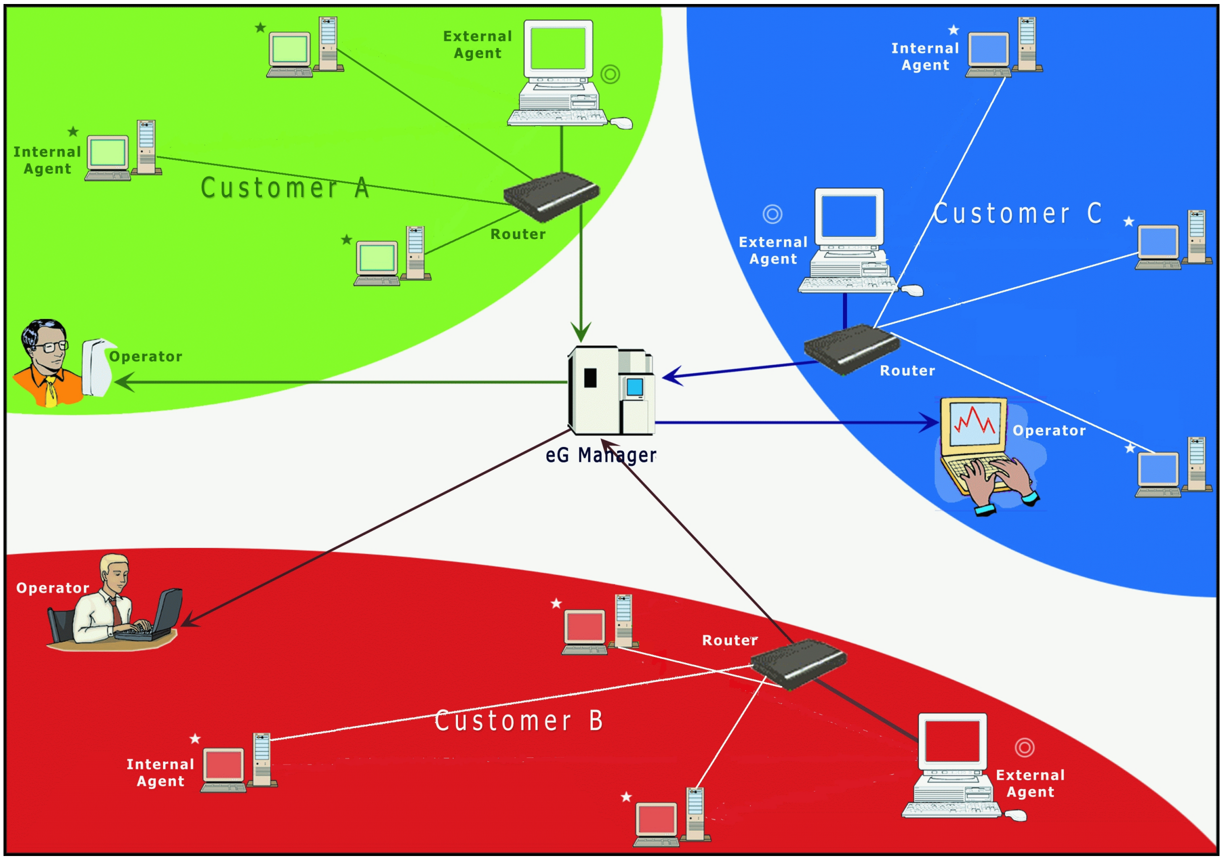

Figure 2.4: Typical deployment architecture of eG Enterprise

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 Asset Management Asset Management section of the Administering the eG Enterprise Suite.

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).

On a Unix manager, follow the steps below to modify the heap memory allocation:

1. Login to the eG manager host.
2. Edit the **catalina.sh** file in the **/opt/egurkha/manager/tomcat/bin** directory.
3. Search for the entry *Xms* in the file. You will then find an entry that reads as follows:

   -Xms256m -Xmx256m

4. 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

5. 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.

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).

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.

   Please refer to *"The eG Installation Guide"* for a detailed description on the above two stages.

3. **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.
4. **Monitoring** using the eG Enterprise system. At this stage, using the user interface, users can monitor various aspects of their IT infrastructure.

Figure 2.5 depicts the various stages involved in deploying eG Enterprise in a target environment.

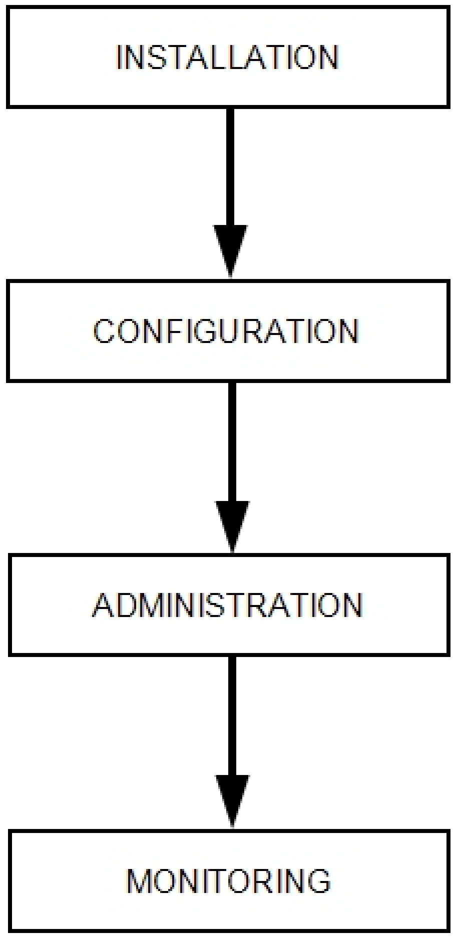

Figure 2.5: Stages involved in deploying eG Enterprise

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

Getting Familiar with Monitor Interface

The eG agents monitor crucial components of an IT infrastructure such as web servers, SSL servers, web application servers, database servers, LDAP servers, DNS servers, payment gateways, WAP gateways, and connections to ISPs, partner sites, and back-end order fulfillment systems. By pseudo-periodically monitoring the entire infrastructure and by comparing the current status of each component with historical data, eG Enterprise detects problems in the IT infrastructure. Using a unique, patented top-to-bottom, end-to-end correlation approach, eG Enterprise can quickly identify and report the root cause of the problem. A graphical representation of the IT infrastructure topology enables a user to quickly drill down to the exact cause of a problem and perform detailed analysis. By enabling speedy diagnosis, eG Enterprise can enable IT infrastructure providers serve their customers better, with minimal downtime.

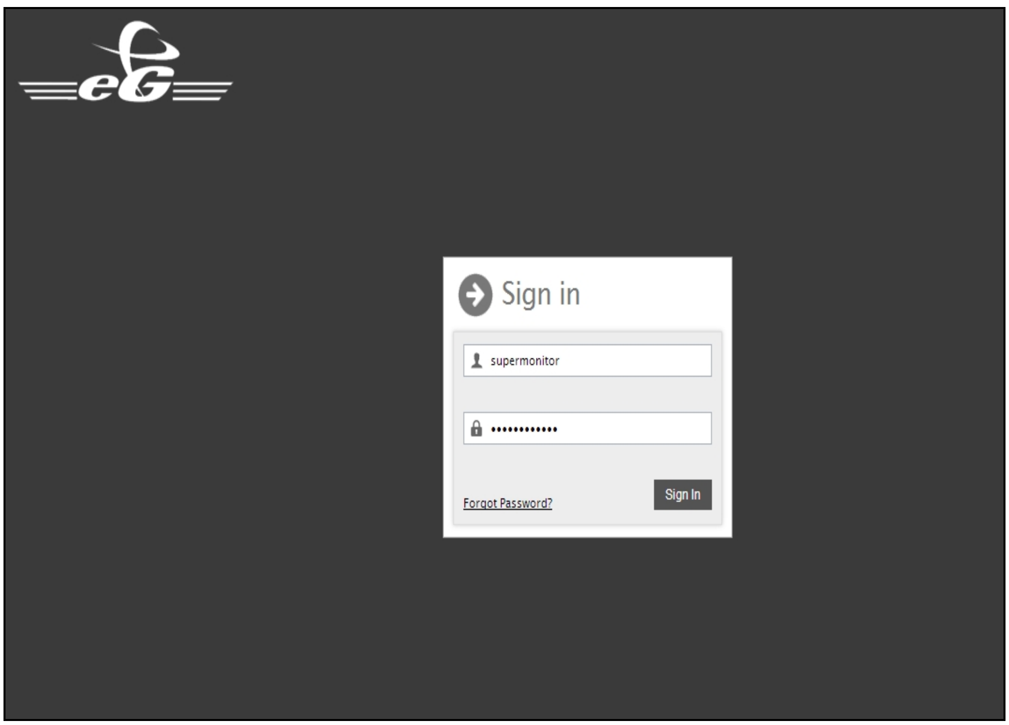

Figure 3.1: Login to eG Enterprise as the supermonitor

This chapter is intended to enable a user to effectively monitor the target IT infrastructure using eG Enterprise. To start monitoring IT infrastructure, a user has to first login to eG Enterprise.

The eG Enterprise system is predefined with a default monitor user account with a login of supermonitor and password supermonitor. A user can also login using any other **Username** also, provided the role assigned to that user name allows him/her monitoring privileges. To know more about user roles and user profile creation in the eG Enterprise system.

If a user forgets his/her login **Password**, he/she can click on the **Forgot Password** link in Figure 3.1. Doing so invokes Figure 3.2 wherein the user would have to provide the **Username** for which the password details are required, and then click the **Get password** button to retrieve the password.

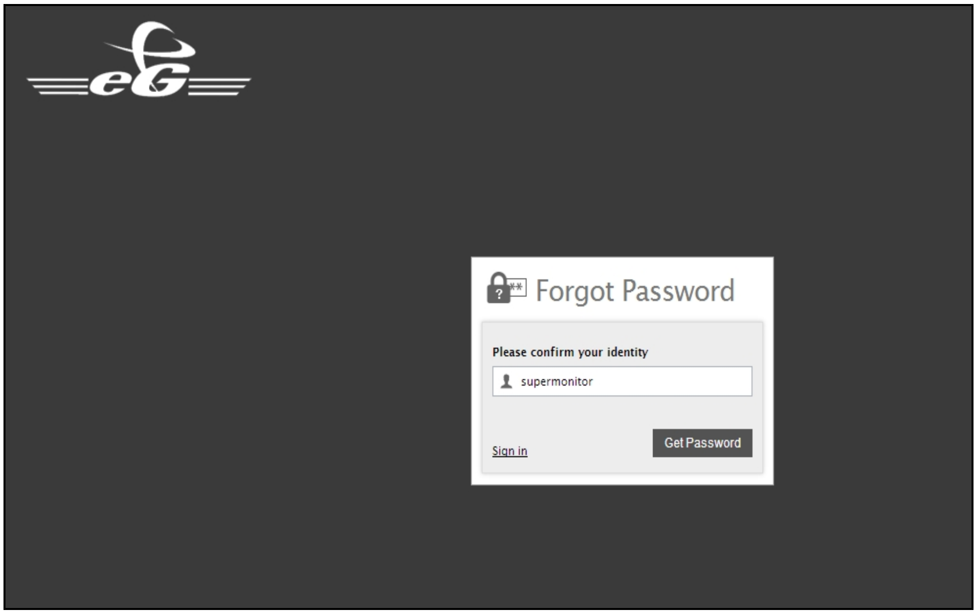

Figure 3.2: Retrieving password

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 3.2 has been configured to receive email alerts, and a valid email ID has been assigned to the user; please refer to 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/SMS SETTINGS** page of the eG administrative interface (see Figure 3.2) a valid mail host and **eG Administrator maild ID** should have been configured. Please refer to Section Figure 3.3 of this manual for more details with regard to the same.

If the user logging in is a domain user or is a user who belongs to an AD group registered with the eG Enterprise system, then the **Username** specified at login should be of the format: *DomainName/Username*, as depicted by Figure 3.3 below:

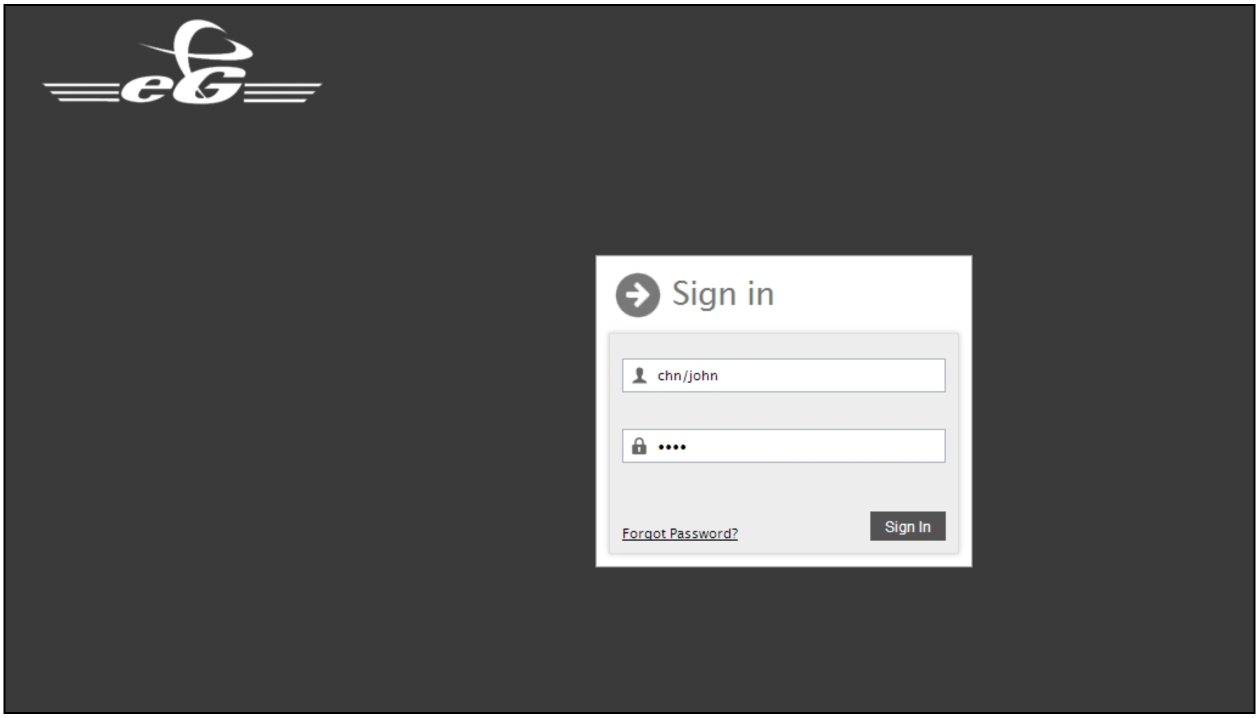

Figure 3.3: A domain user logging in

If a domain user attempting to login to the eG monitoring console belongs to more than one AD group that is created in the eG Enterprise system, then, when the **AUTHENTICATE** button is clicked, Figure 3.4 will appear listing the registered groups to which the user belongs.

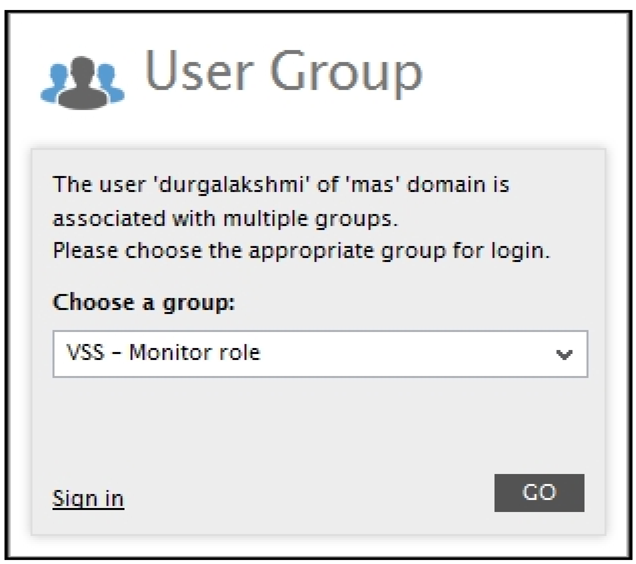

Figure 3.4: Selecting the group to login as

Selecting a group from the **Choose a group** list in Figure 3.4 and clicking the **GO** button soon after enables the user to automatically inherit the access rights and monitoring scope defined for that group.

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=mon&upass=&accessKey=eGm0n1t0r

Users who are assigned the **Admin** role (for eg., the default *admin* user) or the **ServerAdmin** role can access the monitor interface of eG Enterprise directly from the admin interface, without having to log out of the admin interface first. For this, the user will have to select the **MONITOR** tab present next to the **ADMIN** tab in the eG user interface (see Figure 3.5).

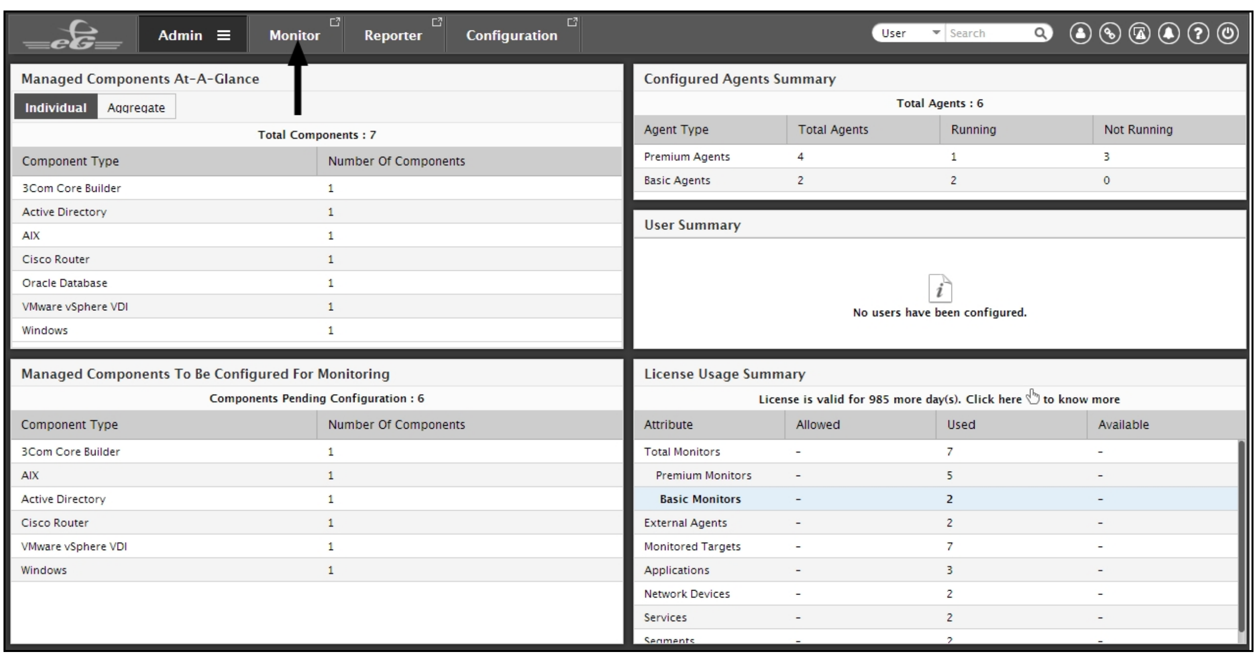

Figure 3.5: Clicking on the MONITOR option to directly connect to the monitor interface

Users who are assigned any other role (**Supermonitor, Monitor, AlarmViewer,** or **SuperAlarmViewer**) will have to connect to the eG manager, and then explicitly login to the monitor interface of eG Enterprise.

To begin with, the *monitor interface* gets to view the **CURRENT ALARMS** window that displays the list of current alarms to the eG Enterprise system (see Figure 3.6), in the order of their priority.

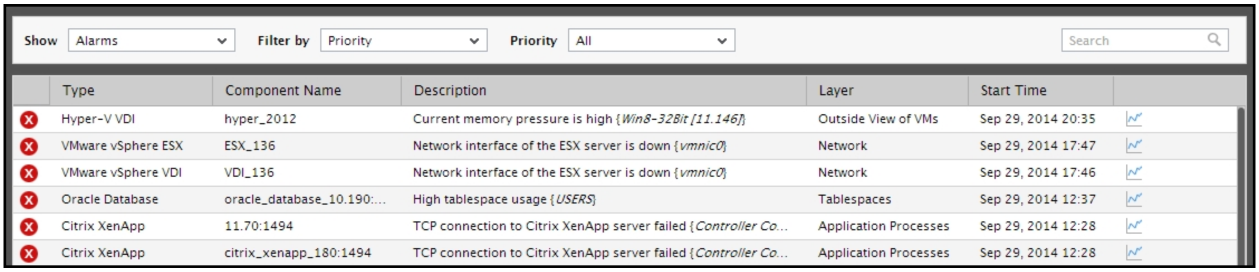

Figure 3.6: The CURRENT ALARMS window displaying the current alarms

Note:

To know more about the current alarms, refer to the Alarms chapter of the guide.

Behind the **ALARMS** window is the **MONITOR HOME** page (see Figure 3.7). If no alarms exist in the environment, then this is the first page that will be displayed in the monitor interface.

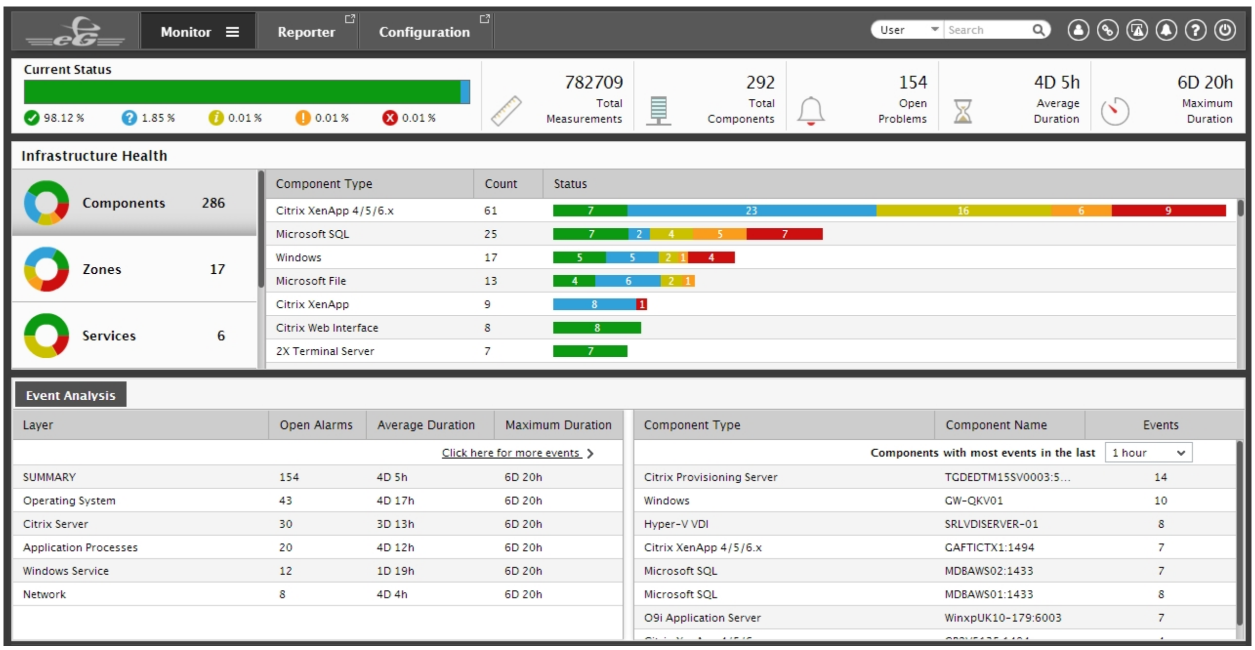

Figure 3.7: The Monitor Home page

The menu at the top of this page permits monitor users to view the status of the monitored elements such as **Zones**, **Services, Segments**, individual **Components, etc.**

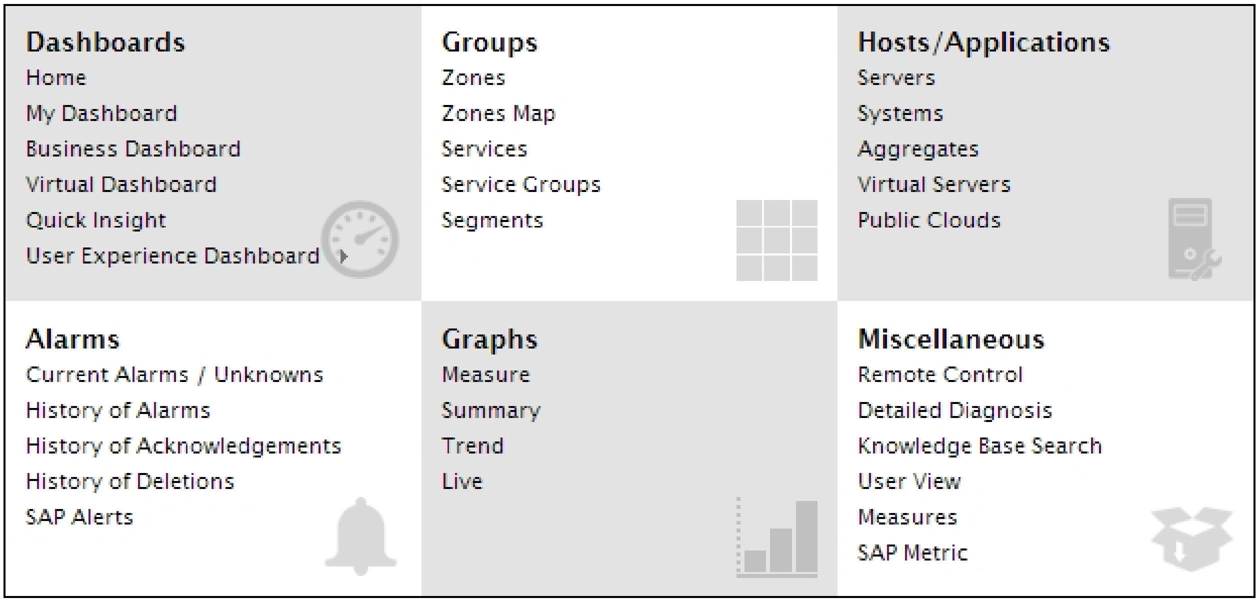

Figure 3.8: The Monitor menu

3.1 Role of a SuperMonitor

As stated already, when a user logs in as the *supermonitor*, he/she is entitled to the privileges of the **Supermonitor** role. This means that a *supermonitor* receives an unrestricted view of the environment.

Note:

The supermonitor as well as other users will not be allowed to login to the eG interface if the instances corresponding to the managed Oracle database servers have not been properly configured by the administrator.

3.1.1 Time Zone Handling in the eG Monitoring Console

By default, the eG manager time zone is displayed in the eG monitoring console. The same time zone applies to metrics collection and alarm generation also. When creating a user profile using the eG admin interface, an administrator can associate a different time zone with the user. When this user logs in, then the time zone assigned to his/her profile automatically applies to all time displays and time-related computations in the eG monitoring console.

Users to the console may sometimes notice discrepancies in how time zones are handled by the console. These caveats have been discussed hereunder:

- Assume that the eG manager's time zone and the current user's time zone are different. If DST (Daylight Savings Time) sets in for the user time zone and the clock is reset to one hour ahead or behind, the user will notice the wrong time being displayed in the console for the next hour. For example, if the user time is actually 2 pm but gets reset to 1 pm owing to DST, the console will show the wrong start time and end time, particularly in graphs, for 1 hour, until the user time under DST becomes 2 pm.

- If DST sets in for the eG manager's time zone, the eG management console will not show the time correctly. The console time may be 1 hour ahead or behind (depending upon whether the clock was reset forward or backward) the actual time in DST. This issue however is resolvable. You can opt for one of the following workarounds:
- Change the JDK of the eG manager to v1.7.0.25 or higher, (OR)
- Use the **tzupdater** tool to update the installed JDK and JRE images with the most recent time zone data to accommodate the US, Asia, Europe, and other locations' DST changes. For this, download the **tzupdater.jar** from the web to any location on the eG manager host. Open the command prompt, go to the copied folder, set the Java path, and run the following command:

  java -jar tzupdater.jar –u
- Finally, restart the manager to get the correct timeline.

4

The Monitor Home Page

Behind the **ALARMS** window is the **MONITOR HOME** page (see Figure 4.1). If no alarms exist in the environment, then this is the first page that will be displayed in the monitor interface.

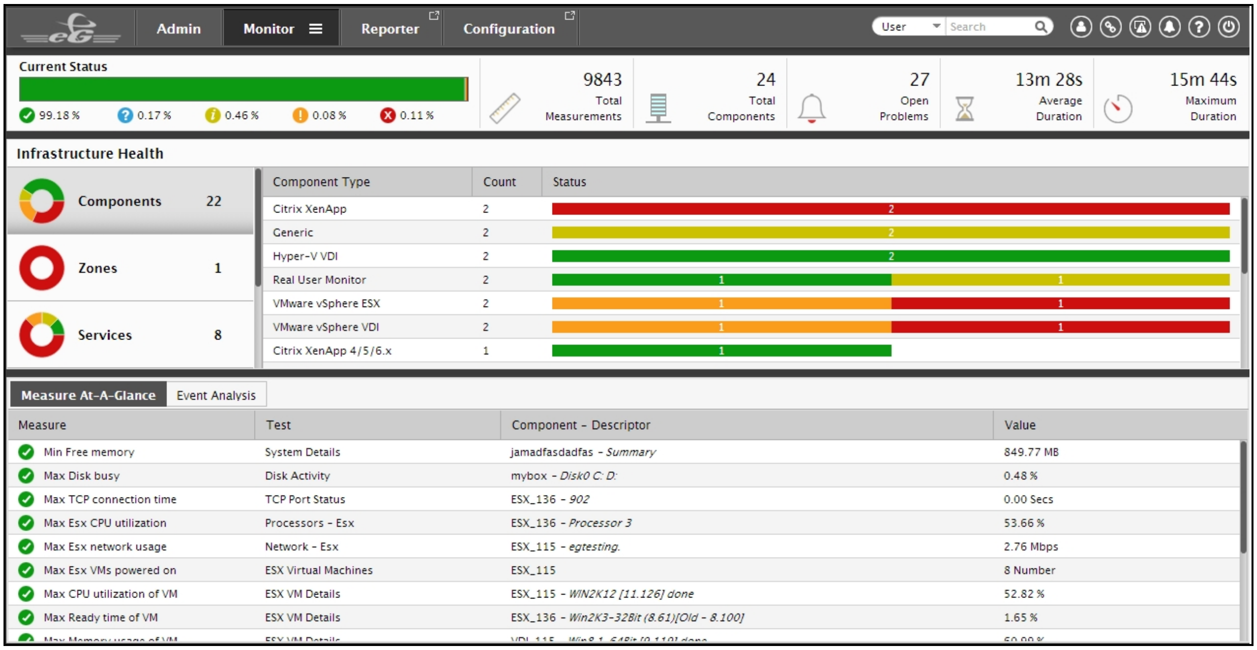

Figure 4.1: The Monitor Home page

This page quickly updates a user with the health of his/her monitored environment. The page reveals the following information:

- The first section is the **Current Status** section that reveals at a glance, the status of the measurements reported to the eG manager. Besides displaying the total number of monitored components and the number of performance metrics collected by the eG agents from these components, this section also reveals the percentage of total measurements that are in the critical, major, minor, normal, and unknown states. Using this information, an accurate assessment of the overall infrastructure performance can be made. Clicking on any of the states will take you to the **Current Alarms** window, where you can view all open alarms of the corresponding priority.
- Below the **Current Status** section, is the section that reveals the **Infrastructure Health**. Since the health of an infrastructure depends entirely upon the performance of each of its key ingredients - namely, the Components, Zones, Services, Service Groups and Segments - this section takes the help of a bar graph to clearly indicate the number of zones, services, service groups, segments, and components that are in the Critical, Major, Minor, Normal, and/or Unknown states.
- The table below explains the color-coding scheme adopted by eG for indicating the states of the

services/segments/components:

Color	State
Red	Critical
Orange	Major
Yellow	Minor
Green	Normal
Blue	Unknown

If you click on a division in the **Zones** doughnut graph, you will be lead to a page that lists the zones which are in that particular state (see Figure 4.2).

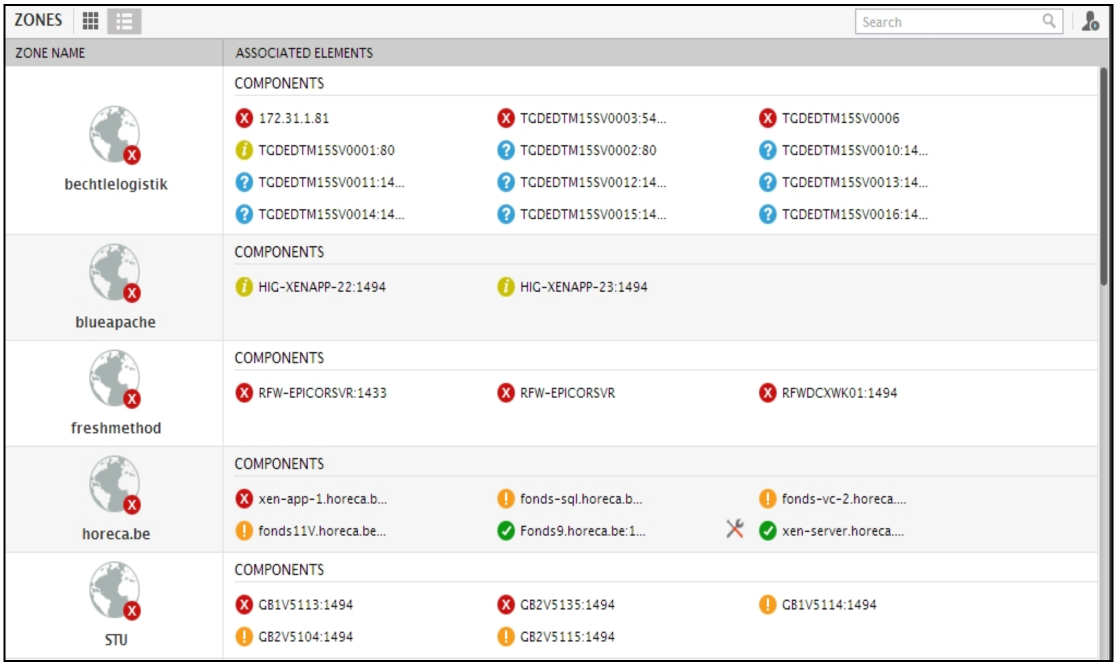

Figure 4.2: The zones in a particular state

Against every zone in this page, the infrastructure elements (which could be subzones/segments/services/components within a zone) that are part of the zone, and the current state of each element will be listed (see Figure 4.2).

Note:

By default, against each zone displayed in the **ZONE LIST** page, the top-10 **Components** included in that zone will be displayed. Typically, to identify the top-10 components, eG Enterprise automatically sorts all the components included in the zone on the basis of their current state, arranges the sorted list in the alphabetical order of the component names, and picks the first 10 components of this list.

If you want more number of components to be displayed against each zone, do the following:

- Login to the eG administrative interface.
- Click on the icon available in the ***Admin*** tab. Then, select the ***Monitor*** option in the ***Settings*** tile.
- Click the **Other Settings** sub-node under the **General** node in the tree-structure in the **MONITOR SETTINGS** page.
- In the right panel, modify the default value 10 that is displayed in the **Components count in segment/service/zone list** text box.
- Click the **Update** button to save the changes.

Clicking on the **Zones** link in the **Infrastructure Health** section reveals the complete list of configured zones and their state.

Note:

While listing the **Components** associated with a zone, if you want the component-types listed instead of the component names in Figure 3.6, then you will have to set the **Display component types in segment/service list** flag in the **MONITOR SETTINGS** page of the eG administrative interface, to **Yes**.

Clicking on an abnormal component in Figure 4.2 will lead you to a page that displays the layer model, tests, and measurements of the component. Clicking on a zone Figure 4.2, reveals a **Zone Dashboard** that summarizes the performance of all the zone elements. For more information on the zone dashboard, refer to Section **8.3** of this manual.

Clicking on a division in the **Service Groups** doughnut graph will lead you to a page that lists the service groups, which are in that particular state.

Clicking on a division in the **Services** doughnut graph in the **Infrastructure Health** section will lead you to a page that lists the services, which are in that particular state (see Figure 4.3). By default, against every service in this page, the components engaged in the delivery of the service will be displayed along with their state.

Note:

By default, against each service displayed in the **SERVICE LIST** page, the top-10 **Components** associated with that service will be displayed. Typically, to identify the top-10 components, eG Enterprise automatically sorts all the service components on the basis of their current state, arranges the sorted list in the alphabetical order of the component names, and picks the first 10 components of this list.

If you want more number of components to be displayed against each service, do the following:

- Login to the eG administrative interface.
- Click on the icon available in the **Admin** tab. Then, select the **Monitor** option in the **Settings** tile.
- Click the **Other Settings** sub-node under the **General** node in the tree-structure in the **MONITOR SETTINGS** page.
- In the right panel, modify the default value 10 that is displayed in the **Components count in segment/service/zone list** text box.
- Click the **Update** button to save the changes.

Note:

By default, only the components associated with a service will be displayed in the **SERVICE LIST** page. If you want the segments associated with the service also to be displayed, then, do the following:

- Login to the eG administrative interface.
- Click on the icon available in the **Admin** tab. Then, select the **Monitor** option in the **Settings** tile.
- Click the **Other Settings** sub-node under the **General** node in the tree-structure in the **MONITOR SETTINGS** page.
- In the right panel, set the **Show segment(s) in service list** flag to **Yes**.
- Click the **Update** button to save the changes.

To make sure that **SERVICE LIST** page does not display the list of components associated with a service, set the **Show component(s) in service list** flag in the **OTHER DISPLAY SETTINGS** panel of the **MONITOR SETTINGS** page to **No**.

If both the segment list and component list are disabled, then the **SERVICE LIST** page will only display a vertical list of services and their current state.

Clicking on the **Services** link in the **Infrastructure Health** section reveals the complete list of configured services and their state.

Note:

While listing the **Components** associated with a service, if you want the component-types listed instead of the component names in Figure 4.3, then you will have to set the **Display component types in segment/service list** flag in the **MONITOR SETTINGS** page of the eG administrative interface, to **Yes**.

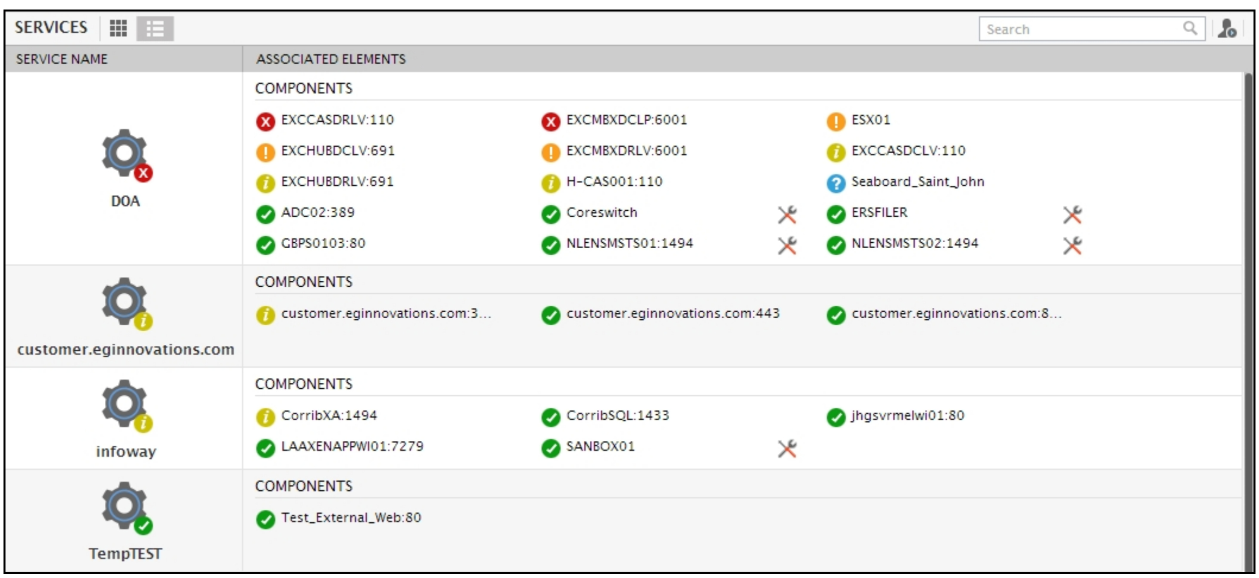

Figure 4.3: Viewing the list of services in a particular state

Clicking on an abnormal component in Figure 4.3 will lead you to a page that displays the layer model, tests, and measurements of the component. Click on a service name in Figure 4.3 to view the topology of the service (in case of a non-web-site service) (see Figure 4.4) or the transactions configured for that service (in case of a web site service). However, if a problem service constitutes more than one problem segment, then clicking on the service name in Figure 4.3 will take you to a page that lists the problem segments, indicates the current state of each segment, and the problem components (if any) in every segment.

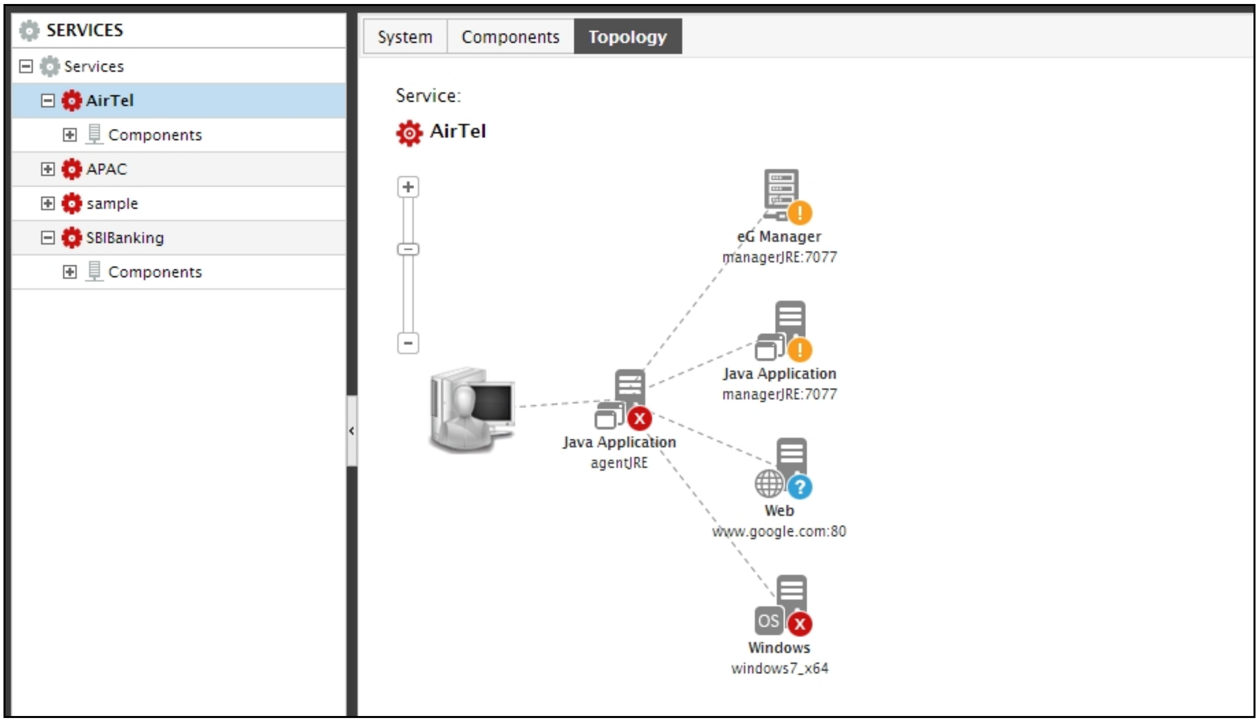

Figure 4.4: The topology of the web site buy.abc.com

Next, to know how well every segment configured in the environment has fared, take a look at the **Segments** doughnut graph in the **Infrastructure Health** section. Clicking on a division in this graph will enable you to view the list of segments in a particular state (see Figure 4.4). Alongside each segment, the IP/hostname of the top-10 (by default) the segment components (i.e., components that are part of the segment) will be displayed along with their state.

Note:

By default, against each segment displayed in the **SEGMENT LIST** page, the top-10 **Components** included in that segment will be displayed. Typically, to identify the top-10 components, eG Enterprise automatically sorts all the components associated with that segment on the basis of their current state, arranges the sorted list in the alphabetical order of the component names, and picks the first 10 components of this list.

If you want more number of components to be displayed against each segment, do the following:

- Login to the eG administrative interface.
- Click on the icon available in the **Admin** tab. Then, select the **Monitor** option in the **Settings** tile.
- Click the **Other Settings** sub-node under the **General** node in the tree-structure in the **MONITOR SETTINGS** page.

- In the right panel, modify the default value 10 that is displayed in the **Components count in segment/service/zone list** text box.
- Click the **Update** button to save the changes.

Note:

While listing the **Components** associated with a segment, if you want the component-types listed instead of the component names, then you will have to set the **Display component types in segment/service list** flag in the **MONITOR SETTINGS** page of the eG administrative interface, to **Yes**.

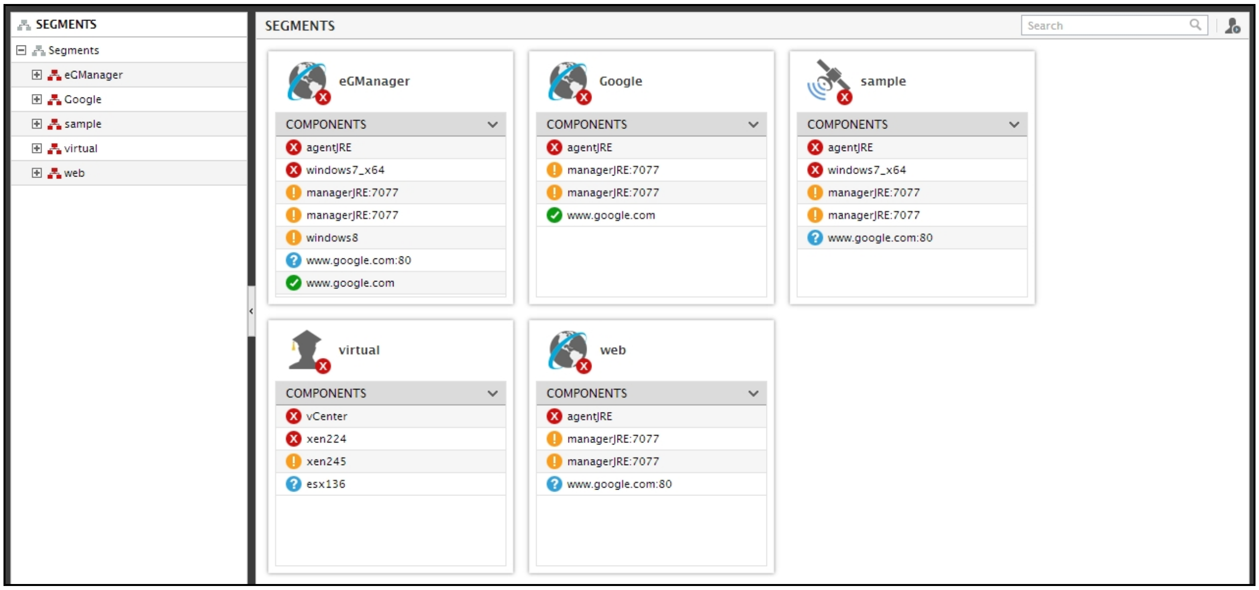

Figure 4.5: Figure 4.5: The list of segments that are in a particular state

Click on a segment name in Figure 4.5 to view the topology of the segment. If you click on a component, then the layer model, tests, and measurements of that component will appear.

Clicking on the **Segments** link in the **Infrastructure Health** section reveals the complete list of configured segments and their state.

Finally, click on a division in the **Components** bar graph in Figure 4.1 for a quick update on components that are currently in a particular state.

Clicking on the **Components** link in the **Infrastructure Health** section reveals the complete list of configured components and their state.

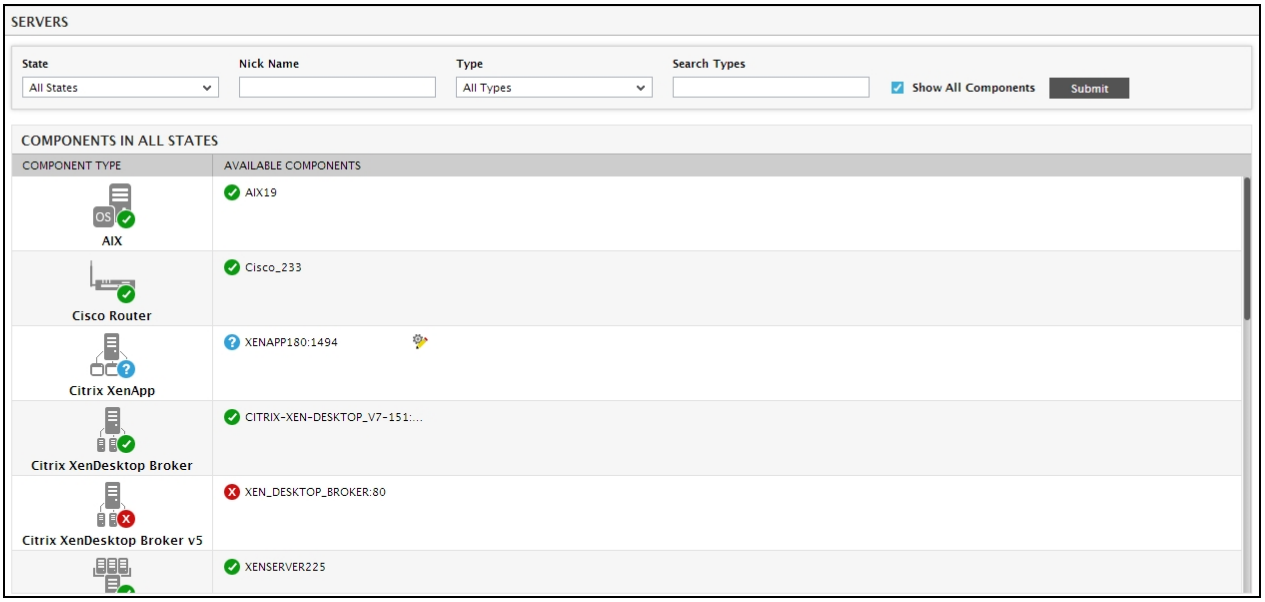

Figure 4.6: The components that are in a particular state

For a closer look at the issues affecting a component's performance, click on the component listing in Figure 4.5. The page depicting the problem layers, tests, and measurements of the chosen component will then appear.

Note:

By default, components that are newly managed will be in the **NORMAL** state. The state will change as the agent starts reporting measures to the eG manager. You can change this default setting to ensure that newly added components are initially in the **UNKNOWN** state, instead of **NORMAL**. To achieve this, set the **NewTestsUnknown** parameter in the **[MISC_ARGS]** section of the **eg_services.ini** (in the **<EG_INSTALL_DIR>\manager\config** directory) to **Yes** (the default value is **No**).

Also, there will be a 30 second time lag between when a component is managed and when its state changes to **NORMAL**.

- Below the **Infrastructure Health** section you will find a **Measures At-A-Glance** section that provides the min/max values of critical measurements updated in real-time. By default, using this section, you can quickly find answers to the following critical performance queries:
  - Which host across the environment is consuming the maximum CPU?
  - Which host in the monitored infrastructure has very little free memory to its credit?
  - Which disk partition on which host is utilized the maximum?
  - Which is the Citrix server that supports the maximum number of active sessions?
  - Across all monitored Citrix servers, which application is CPU-intensive and which Citrix server is it executing on?
  - Which is the web server that services the maximum number of requests over time?

- ◦ Which network interface is using up a lot of bandwidth?
- ◦ On which host are TCP retransmits very high?
- ◦ Where in the target environment is network latency the maximum?
- ◦ TCP connections to which port are taking too long?
- ◦ Which host is currently not available over the network?

You can however, change this default setting to reveal more or less - i.e., you can add to the measure list displayed here, or remove a few measures from the displayed list. To achieve this, do the following:

- ➢ Login to the eG administrative interface as *admin*.
- ➢ Click on the icon available in the **Admin** tab. Then, select the **Monitor** option in the **Settings** tile.
- ➢ Click on the **Measures At-A-Glance** node in the **Monitor Settings** tree.
- ➢ The resulting **MEASURES AT-A-GLANCE CONFIGURATION** page will display the default measure configurations for the **Measures At-A-Glance** section. To know how to manipulate the controls in the page to add more measures or remove a few of the pre-configured measures section.

The min/max values of the configured measures will then be displayed in the **Measures At-A-Glance** section of the Monitor Home page. The first column of the **Measures At-A-Glance** section in Figure 4.6 indicates the current state of the measure (whether Normal/Critical/Major/Minor/Unknown). The **Measure** column is where the configured measures will be displayed. Similarly, each of the configured tests will appear in the **Test** column. Besides, a **Server** column exists, which displays the name of the component and the descriptor, which has currently registered the maximum/minimum value (as the case may be) for every chosen measure. For example, in Figure 3.8, the *tomcat* process executing on the Generic server named *generic*, is the process that is currently consuming a lot of memory. Finally, the current value of the measures for the displayed components will be displayed in the **Value** column.

Note:

The measures will be displayed in the order of the configuration.

Using the information provided by the **Measures At-A-Glance** section, administrators can receive instant status updates on sensitive performance parameters, and can also accurately determine, at a glance, the component on which the parameter is currently experiencing issues (if any), thereby simplifying problem identification.

Clicking on the contents of any cell in this section, will lead you to the layer model of the corresponding component.

Optionally, you can even switch off the **Measures At-A-Glance** section. To do so, you will have to set the **Compute top metrics** parameter in the **MEASURES AT-A-GLANCE CONFIGURATION** page of the eG administrative interface to **No**. By default, this parameter is set to **yes** indicating that, by default, the monitor home page will contain the maximum/minimum computations for measures (i.e., the **Measures At-A-Glance** section). When it is set to **no**, the **Measures At-A-Glance** tab is hidden from the home page, and the **Event Analysis** tab alone appears.

The **Event Analysis** tab page, when clicked, lists the top-5 layers that were most affected by performance issues. Figure 4.7 depicts a sample monitor home page with the **Event Analysis** section.

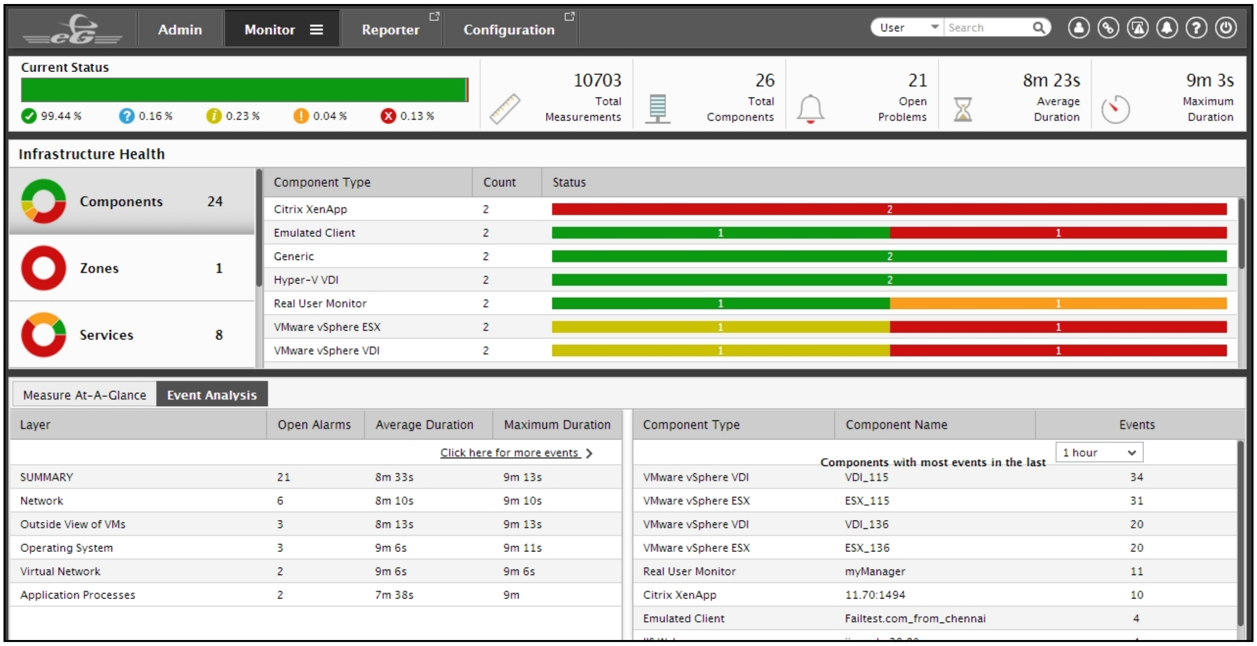

Figure 4.7: A Monitor Home Page with the Event Analysis section

Corresponding to every layer name in the **Event Analysis** section, you will see the number of alarms that are currently open for that layer, the average duration of the open alarms, and the maximum duration for which an alarm had remained open (see Figure 4.7). To the right of the tab page, you will find a list of components that experienced the most number of performance issues during the last one hour (by default). Against every component listing, the corresponding component-type, and the number of events the component encountered during the default period of one hour, is displayed. This information brings to light the most problem- prone components in the environment. Clicking on a component/component- type in this list, reveals the layer model, tests, and the last set of measurements that the eG agent reported for that component. To analyze events across a broader time window in the past, select a different timeline from the **Components with most events in the last** list box here.

The details available in the **Event Analysis** section serve as an effective indicator of the efficiency of the administrative staff in resolving performance issues. To view the complete history of alarms in the environment, click on the **Click here for more events >>** link.

- The **Components At-A-Glance** section comprises of a bar graph depicting the number of components of each type that are being monitored, and their respective states. Clicking on a bar will take you to a page that lists the individual components of the corresponding type and their current state (see Figure 4.8).

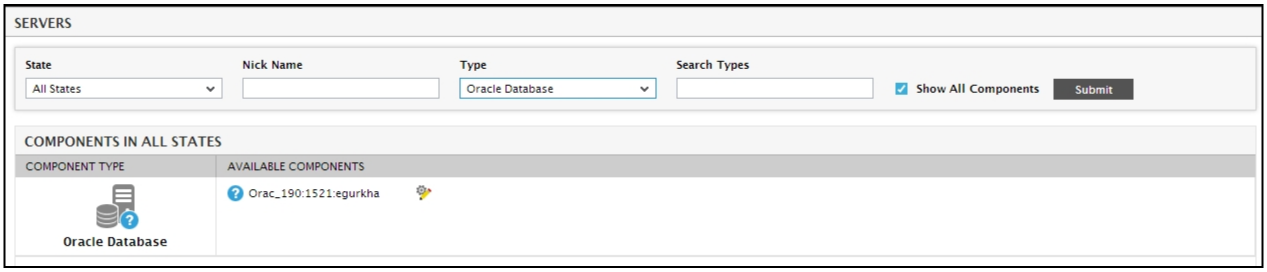

Figure 4.8: Viewing the list of components of a particular type

By clicking on each **Server type** in the **Components At-A-Glance** section, you can view all components of that type and their current state. For instance, clicking on **Generic** in the **Components At-A-Glance** section of Figure 4.8, will reveal all managed components of type **Generic** and their current state.

Note:

By default, in the **Components At-A-Glance** section, the component-types are sorted in the descending order of the total number of monitored components of every type - in other words, in the descending order of the values in the **Count** column of the section. To change the sort order - i.e., to sort the component-types in the ascending order of the contents of the **Count** column - simply click on the down-arrow icon next to **Count**. To sort by a different column, say, the **Server Type** column, simply click on the corresponding column heading. This will instantly sort the contents in the alphabetical order of the names of the displayed server types. You can even override the default sort order, so that the component-types are by default arranged in the alphabetical order of their names, and not on the basis of the **Count**. To achieve this, Click on the icon available in the **Admin** tab. Then, select the **Monitor** option in the **Settings** tile. In the **OTHER DISPLAY SETTINGS** section of the **MONITOR SETTINGS** page, set the **Sort components in dashboards** flag to **By component types**. This ensures that the contents of the **Components At-a-Glance** section are by default sorted in the ascending order of the component-types. Accordingly, the down-arrow icon, by default, appears next to the column heading, **Server Type**.

The menu at the top of this page permits monitor users to view the status of the monitored elements such as **Zones**, **Services, Segments**, and individual **Components** (see Figure 4.8).

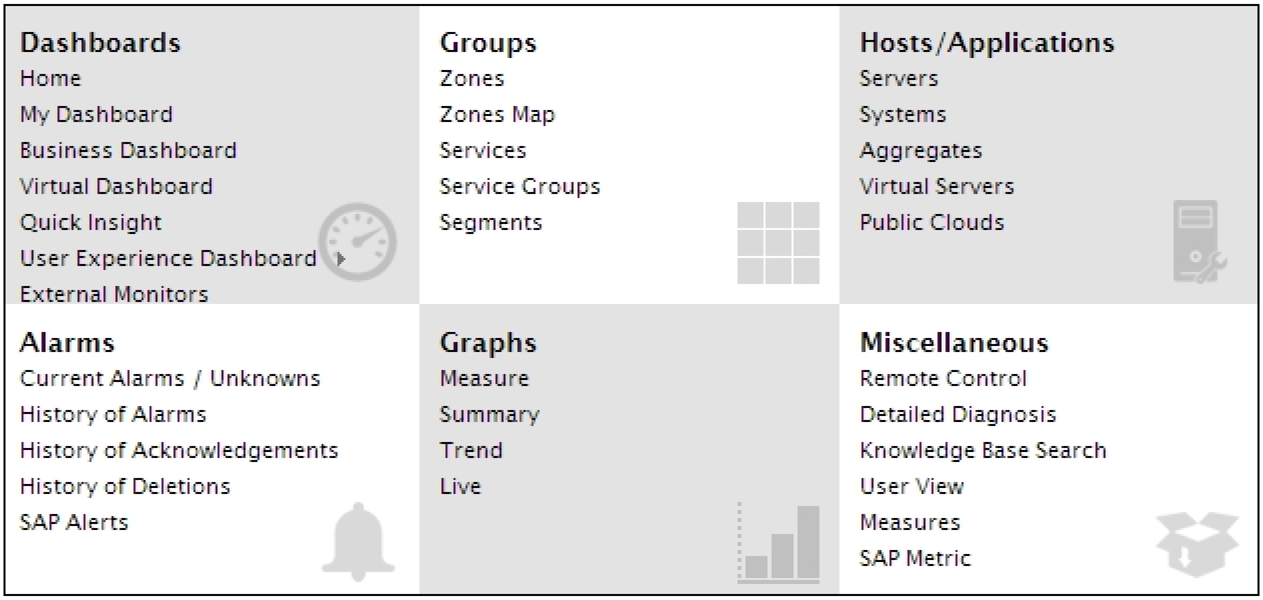

Figure 4.9: The Monitor menu

In addition to this, the menu also facilitates the following:

- Switching to the monitor **Home** page from anywhere in the monitor interface
- Browsing for any measure across the environment and focusing on its current status
- The viewing of the current alarms, and the complete alarm history
- The generation of a wide variety of graphs including, measure, summary, and trend graphs
- The viewing of detailed diagnostic measures and the execution of remote control on agent hosts

Note:

The detailed diagnosis and remote control capabilities will be available only if the license permits them.

The menu at the right top corner of the page (see Figure 4.8) provides for the following:

- Signing out of the monitor interface
- Viewing and modifying the profile of the current monitor user
- Invoking a context-sensitive help topic

5

Alarms

5.1 Alarms

As we have seen in this chapter, eG Enterprise can be configured such that as and when a situation for an alarm arises, the manager automatically generates an email alert to the supermonitor. He/she can choose to view the current set of alarms or the entire history of alarms from the **Alarms** tile in the eG monitor interface.

When the supermonitor selects the **History of Alarms** option, he/she gets to view the entire list of alarms pertaining to the entire infrastructure over a period of time as in Figure 5.1.

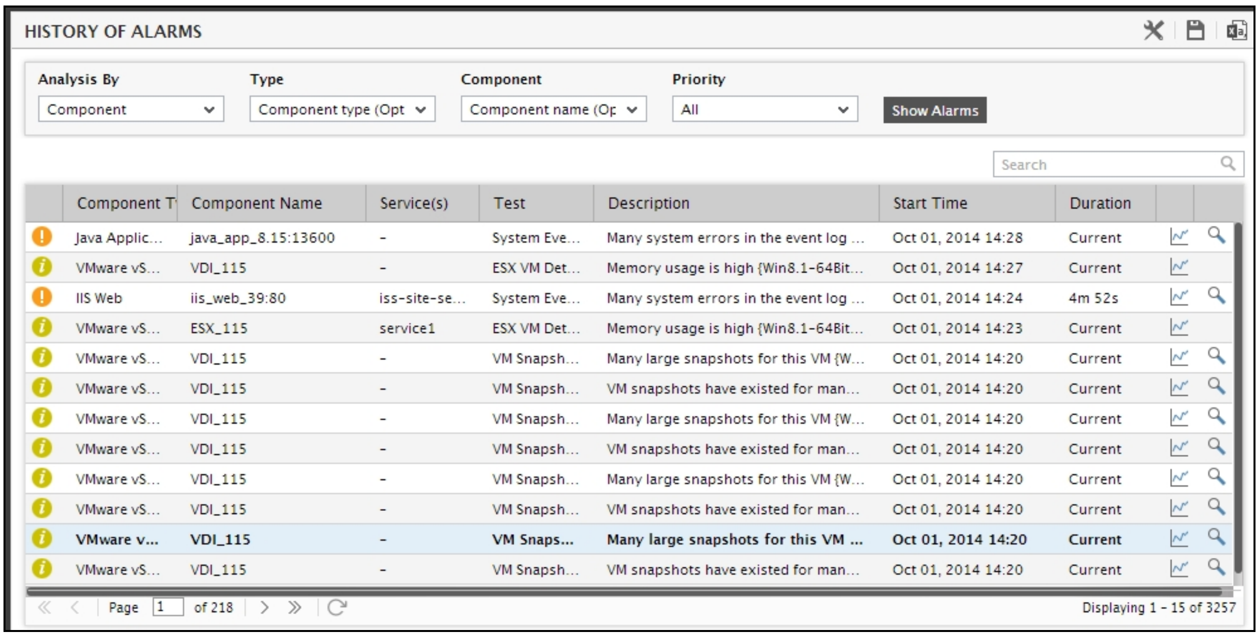

Figure 5.1: History of alarms that the supermonitor can view

- By default, as soon as you access the **HISTORY OF ALARMS** page, you will view the alarm history of all managed components in the environment (see Figure 5.1). If required, you can build filter conditions using this page so that, you can selectively view the alarm history of the following infrastructure elements alone:
  - Any managed component type in the environment;
  - Any managed component;
  - A component of a particular type;
  - A specific segment / service / zone;

- A component-type that is part of a segment / service / zone;
- A component that is part of a segment/service/zone;

The first step towards building these filter conditions is selecting a basis for the filter. This can be achieved by picking an option from the **Analysis By** list. The options available here are as follows:

- **Component** : This is the default selection in the **Analysis By** list. Owing to this default setting, the **HISTORY OF ALARMS** page displays the alarm history of all managed components in the environment, by default. If you proceed with the default selection, then, you will find that the **Component Type** and **Component** lists in Figure 5.1 are populated with all the managed component types and components (respectively) in the environment. If you want to view the alarm history of a particular component-type, pick that type from the **Component Type** list. Likewise, if you want to view the alarm history of a particular managed component, pick the name of that component from the **Component** list. If the **Component** list has too many components to choose from, then, you can condense the list by first picking a **Component Type** ; this will make sure that the **Component** list consists of only those managed components that are of the chosen type. You can then easily pick the component of your choice from the **Component** list.
- **Zone**: Selecting this option from the **Analysis by** list will invoke a **Zone** list. Select a particular zone from this list, if you want to view the history of alarms related to that zone. An **Include Subzone** flag also appears. By setting this flag to **Yes**, you can make sure that the alarm history also includes those alarms that are associated with the sub-zones of the chosen zone.

  Once a **Zone** is selected, the **Component Type** and **Component** lists will be populated with those types and components (respectively) that are part of the selected zone. To view the alarm history of a particular component-type that is part of a zone, pick that type from the **Component Type** list. Similarly, to view the alarm history of a component that is part of a zone, pick that component from the **Component** list. If the **Component** list still has too many components to choose from, then, you can condense the list further by first picking a **Component Type**; this will make sure that the **Component** list consists of those components in the selected zone that are of the chosen type. You can then easily pick the component of your choice from the **Component** list.

  Note that the 'Zone' option will not be available in the 'Analysis By' list if no zones are configured in the environment.
- **Segment**: If this option is chosen from the **Analysis By** list, a **Segment** list will additionally appear. In order to view the alarm history pertaining to a specific segment, pick a segment from the **Segment** list.

  Once a **Segment** is selected, the **Component Type** and **Component** lists will be populated with those types and components (respectively) that are part of the selected segment. To view the alarm history of a particular component-type that is part of a segment, pick that type from the **Component Type** list. Similarly, to view the alarm history of a component that is part of a segment, pick that component from the **Component** list. If the **Component** list still has too many components to choose from, then, you can condense the list further by first picking a **Component Type**; this will make sure that the **Component** list consists of those components in the selected segment that are of the chosen type. You can then easily pick the component of your choice from the **Component** list.

Note that the 'Segment' option will not be available in the 'Analysis By' list if no segments are configured in the environment.

- **Service**: If this option is chosen from the **Analysis By** list, a **Service** list will additionally appear. In order to view the alarm history pertaining to a specific service, pick a service from the **Service** list.

  Once you choose a **Service**, the **Component Type** and **Component** lists in Figure 5.1 will be populated with those types and components (respectively) that are engaged in the delivery of the said service. If you want to view the alarm history of a particular component-type that is part of the selected service offering, then, pick that type from the **Component Type** list. Similarly, if you want to view the alarm history of a component that supports the selected service offering, pick that component from the **Component** list. If the **Component** list still has too many components to choose from, then, you can condense the list further by first picking a **Component Type**; this will make sure that the **Component** list consists of those components in the selected service that are of the chosen type. You can then easily pick the component of your choice from the **Component** list.

  Note that the 'Service' option will not be available in the 'Analysis By' list if no services are configured in the environment.

- Next, to view the alarms that have remained unresolved for a time period that is in excess of a specified duration, select the **greater than** option from the **Duration is** list, enter a value in the adjacent text box, and then select a unit of time from the list box alongside. For example, to view the history of the alarms that have remained unresolved for over 1 hour, select the **greater than** option, enter 1 in the text box alongside, and select **hours** from the list box adjacent to it.
- Similarly, you can view the history of alarms that have remained unresolved for a time period lesser than a specified duration. To achieve this, select the **lesser than** option from the **Duration is** list, specify a value in the adjacent text box, and select a unit of time from the list box.
- You can even choose to view the details of past alarms that are of a particular priority, by selecting that priority from the **Priority** list.

To override the default settings of the **History of Alarms** page, click on the ✕ icon. The settings pop-up window appears:

- For viewing the details of alarms that were generated during a specific time window, select a fixed **Timeline**, or choose **Any** to provide a date/time range.
- If you want to view the alarm history of components with names that embed a specified string, enter the string to search for in the **Component Search** text box.
- You can even search based on alarm description. By providing the whole/part of an alarm description in the **Description search** text box, you can view the details of alarms with descriptions carrying the given search string.
- By default, you cannot view the acknowledgement/deletion history of alarms in the **HISTORY OF ALARMS** page. Accordingly, the **Show acknowledgements** flag is set to **No** by default. To view the

acknowledgement/deletion history of alarms, set this flag to **Yes**.

- Next, select the column by which the alarm history is to be sorted from the **Sort by** list.
- In addition, you can configure the number of event records to be displayed per page of the event history. By default, 15 records are displayed per page. To display more records, select an appropriate value from the **Events per page** list.
- Finally, click the **Show Alarms** button to generate the history of events.
- The details pertaining to every alarm like the start time, duration, name of the component, component-type, test, alarm description, and the service (if any) that is impacted by the issue, are available. Every row of alarm information will be accompanied by a colored indicator, that indicates the corresponding alarm priority. Critical alarms will be of the color red, major alarms will be in orange, and the minor ones come in pink. An alarm with the end time set to current denotes a problem that has still not been fixed.
- Typically whenever an alarm is raised for the problems at the host-level of a component, the **HISTORY OF ALARMS** page automatically sets the **Component type** to **Host system**, even if the component affected is say, an Oracle Database server or a Web server. From this alarm information, users cannot determine the exact **Component type** of the affected component. Moreover, help desk personnel may prefer to view the operating system of the problem host as part of the alarm information displayed in the **HISTORY OF ALARMS** page, as such information will greatly simplify the troubleshooting process. To make sure that the **HISTORY OF ALARMS** page enables help desk to easily understand, interpret, and solve problems affecting a host's performance, you can optionally configure the eG Enterprise system to display the actual **Component type**, **Host system**, or the affected **Operating system** for host-level alarms in the **HISTORY OF ALARMS** page. To enable this capability, do the following:
  - Edit the **eg_ui.ini** flag in the **<EG_INSTALL_DIR>\manager\config** directory
  - In the **[HOST_SYSTEM]** section of this file, set the **Show_HostSystem** flag to any one of the following values mentioned below:
    - Set the **Show_HostSystem** flag to **HostSystem** if you want the component type to be displayed as **Host system** for the host-level alarms;
    - Set the **Show_HostSystem** flag to **CompType** if you want to display the affected component; This is the default setting that is provided;
    - Set the **Show_HostSystem** flag to **OS** if you want to display the operating system of the host;
  - Finally, save the file.

  Note:

  This configuration affects the **CURRENT ALARMS** window, email/SMS alerts, and SNMP traps as well.

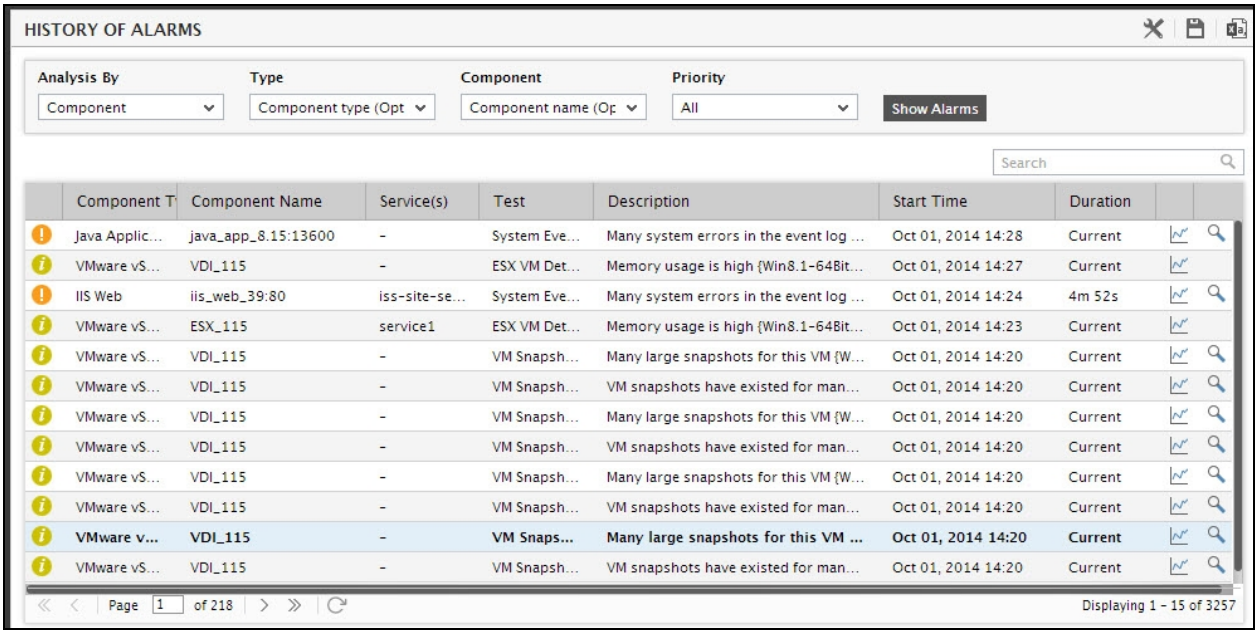

Figure 5.2: The Operating system of the host shown instead of Host System in the Comp Type column

- If you had chosen a particular component name from the **Component** list, then, clicking the **Show Alarms** button will result in an alarm history that pertains to that component alone (see Figure 5.2).

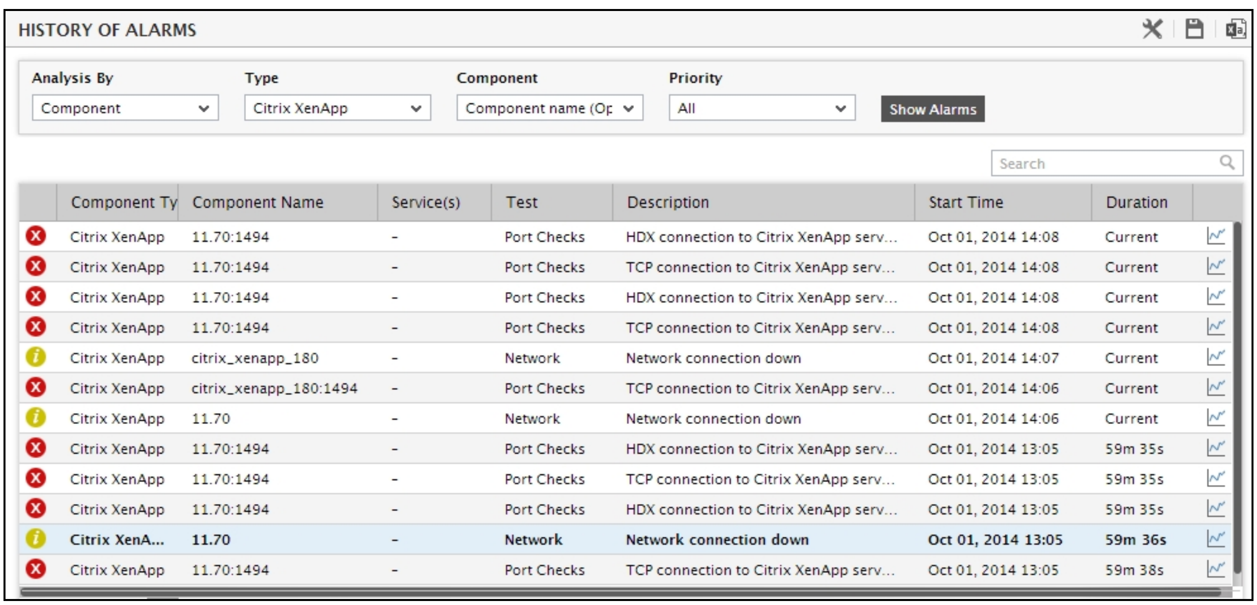

Figure 5.3: Viewing the alarm history of a particular component

- If you had chosen to view only those alarms that have remained unresolved over or within a specified **Duration**, then, clicking the **Show Alarms** button will result in an alarm history that displays only the alarms that fulfill the specified duration condition (see Figure 5.3).

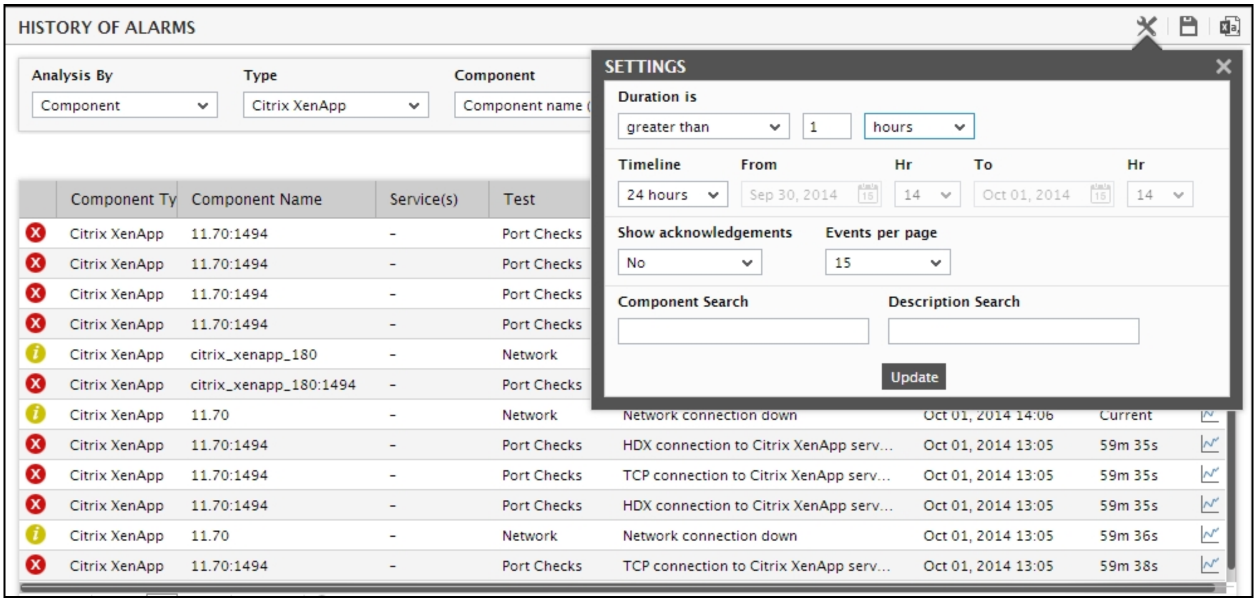

Figure 5.4: History of alarms that have remained open for a duration greater than 1 hour

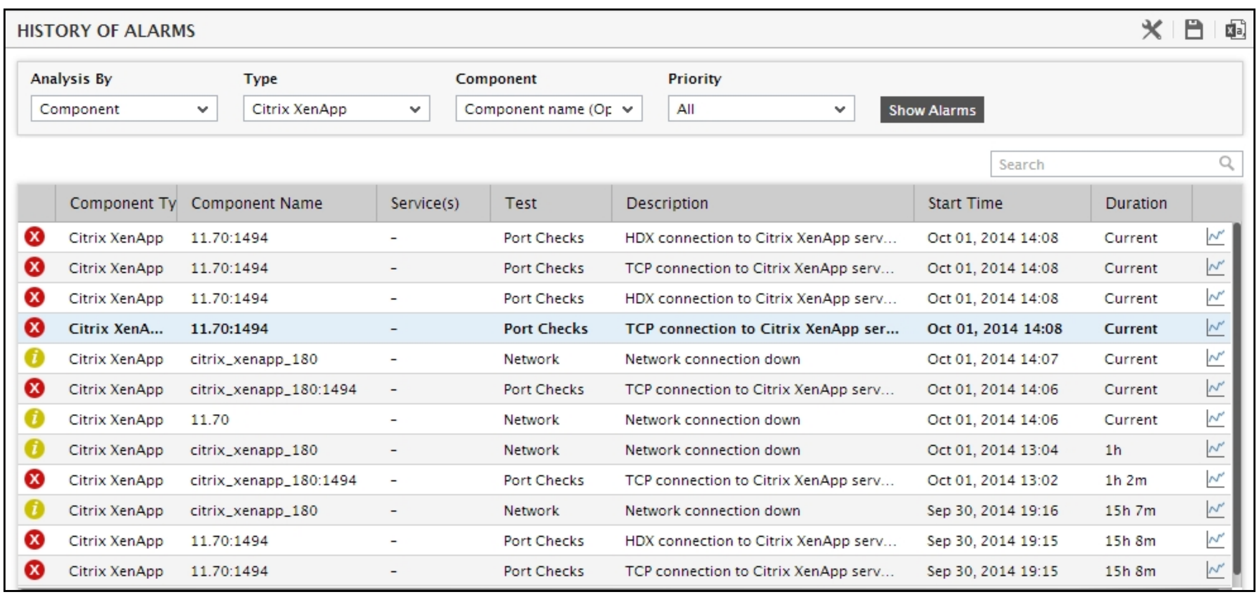

Figure 5.5: The History of Alarms

- Clicking on the **Show Alarms** button after specifying a **Component search** string will display the details of only those alarms that pertain to components with names that embed the specified search string (see Figure 5.4).

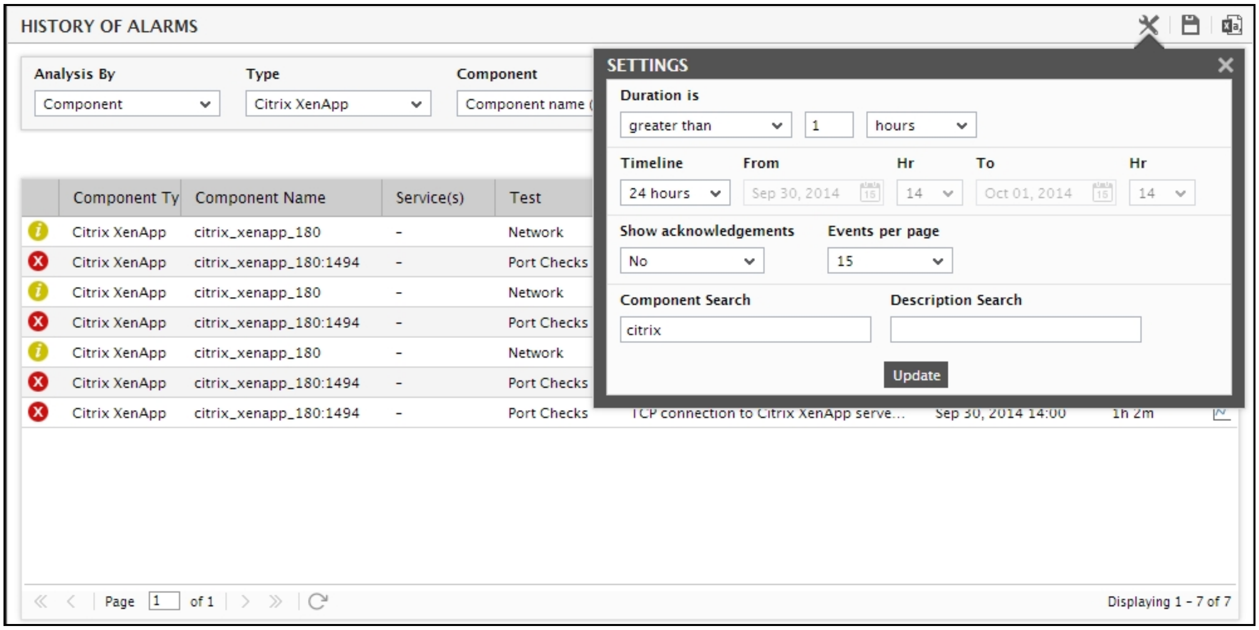

Figure 5.6: Alarm history of components settings page

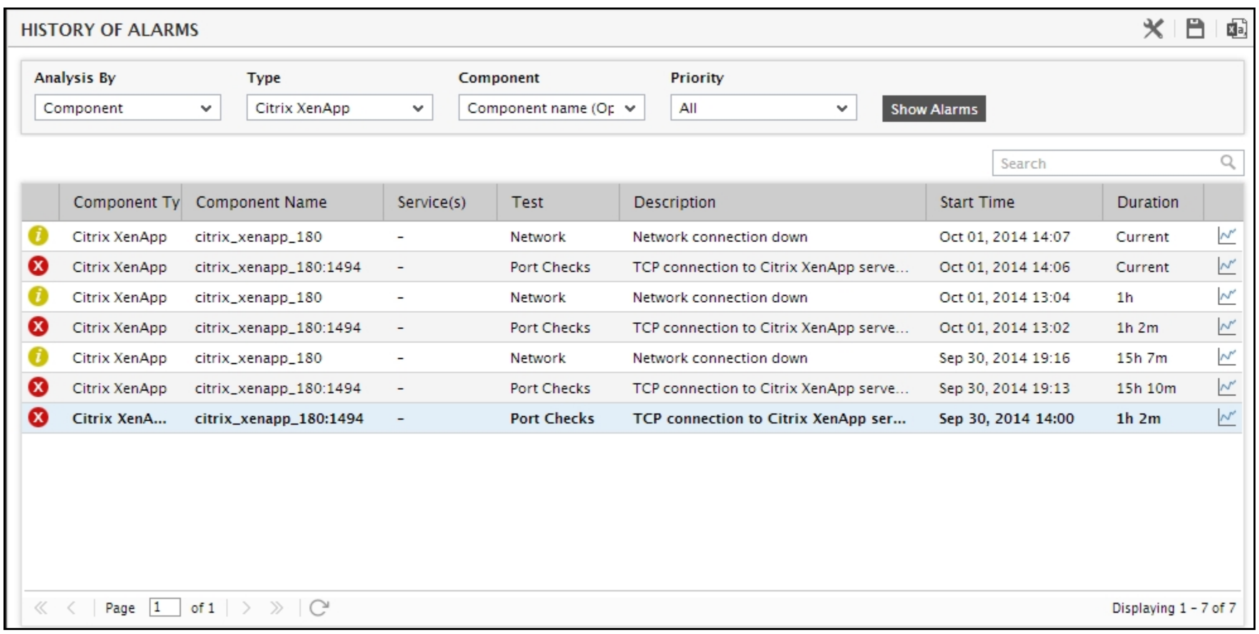

Figure 5.7: The Alarms for a specific component

- In the same way, if you click on the **Show Alarms** button after specifying a **Description search** string, then, only those alarms with descriptions carrying the specified search string will be displayed in the **HISTORY OF ALARMS** page (see Figure 5.6).

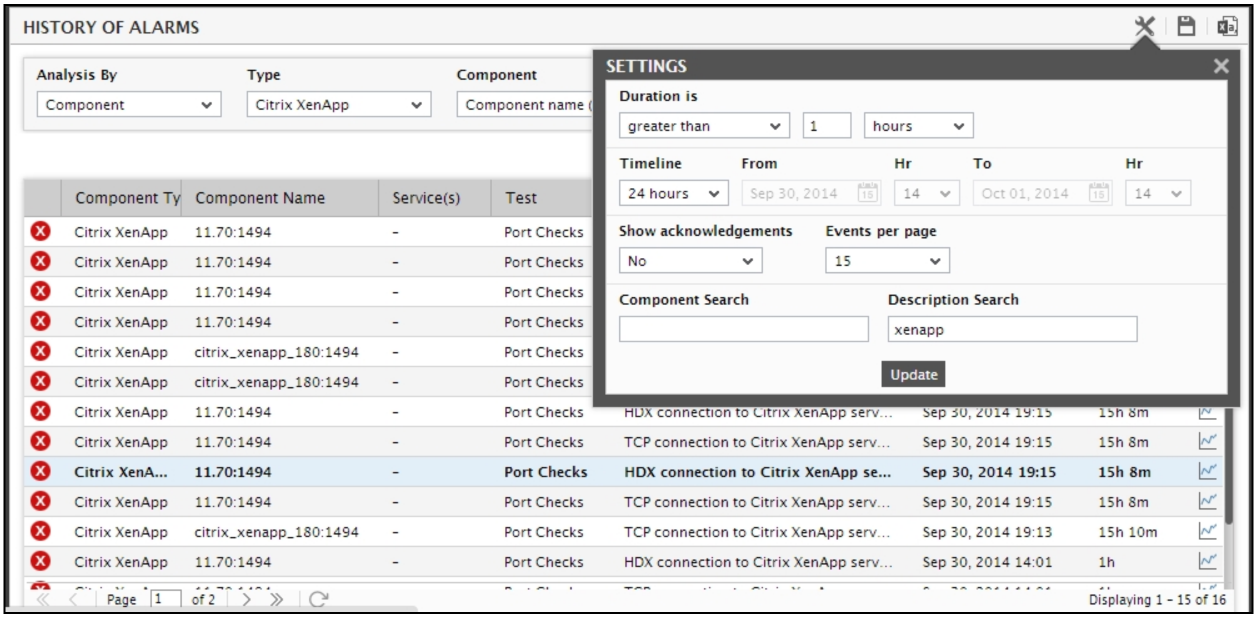

Figure 5.8: Alarm history with alarm descriptions carrying a specified search string

- If the **Show Acknowledgements** flag is set to **Yes**, then, upon clicking the **Show Alarms** button, the acknowledgement/deletion history of alarms will appear as depicted by Figure 5.9 below.

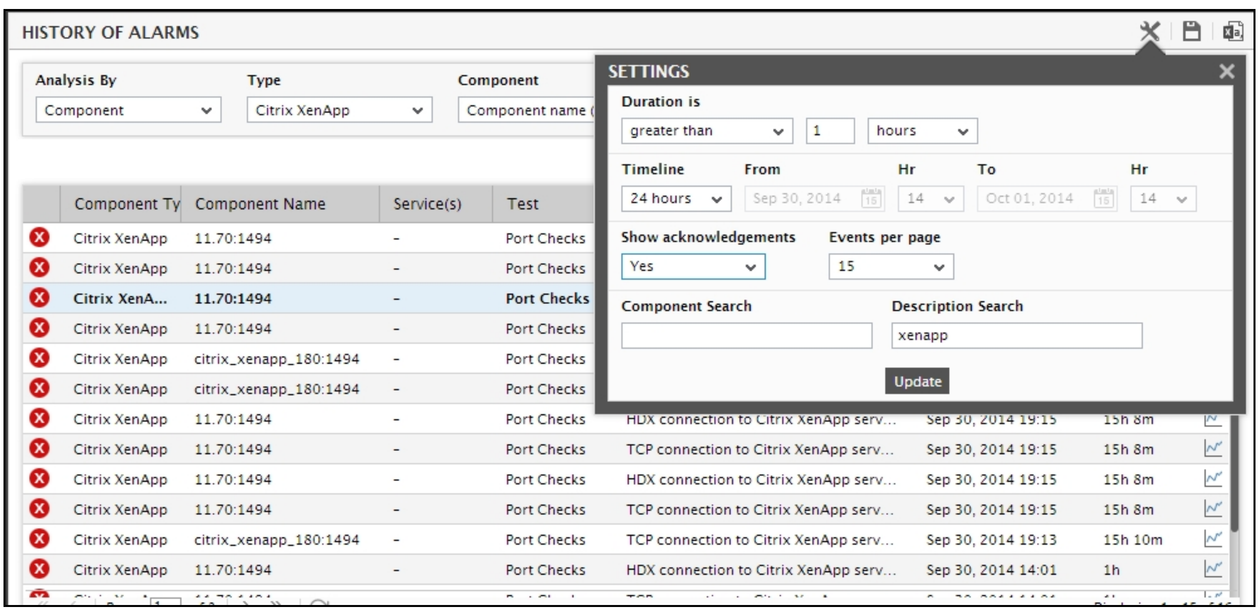

Figure 5.9: Acknowledgment/deletion history displayed in the alarm history page

- By default, the alarm history will not provide information on the users who are responsible for fixing the problems indicated by an alarm - i.e., the users who have been assigned the server/device on which an alarm has been raised. To ensure that every alarm displayed in the **HISTORY OF ALARMS** page is accompanied by this useful user information, do the following:

- Edit the **eg_ui.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
- Set the **Show_Users** flag in the **[ALARM_HISTORY]** section of the file to **true**.
- Save the file.

When this is done, the alarm history will include an additional **User(s)** column, where the names of users who are responsible for fixing the problems indicated by each alarm will be displayed. With this information, the alarm history page will not only enable help desk managers to instantly identify those problems that have remained unresolved for the longest time, but also pin point those help desk personnel who were unsuccessful / had taken a long time to resolve those problems - the efficiency of the help desk staff can thus be ascertained. Moreover, a **User(s)** list will also additionally appear, which will be set to **All** by default. If need be, you can pick a particular user name from this list and click the **Show Alarms** button. Doing so will invoke the history of alarms associated with the chosen user alone (see Figure 5.10).

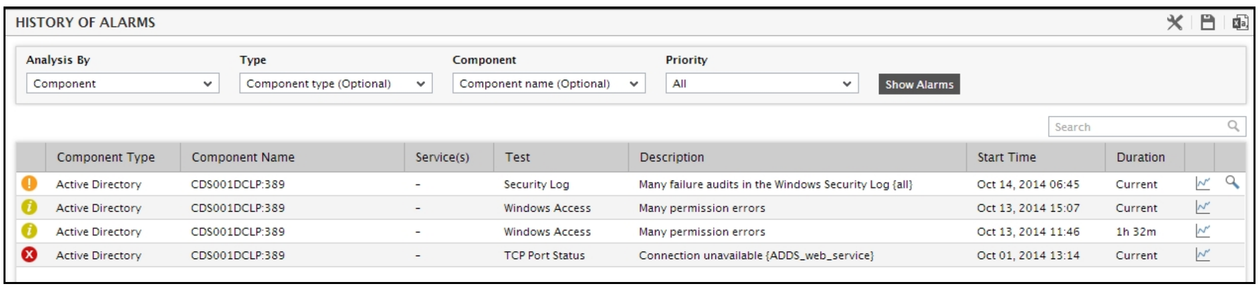

Figure 5.10: Alarm History of a particular user

Note:

Only users with the privilege to monitor all managed components in the environment (eg., users with the **Admin** or **Supermonitor** role) can view user information in the **HISTORY OF ALARMS** page.

- Sometimes, a single alarm raised by the eG manager could have undergone many transitions/changes during the specified `Timeline`. An alarm can change under any of the following circumstances:
  - A change in the alarm priority: This could be a switch to a higher or lower priority.
  - A change in the alarm description: For example, originally, a usage-related alarm may have been raised on disk 'D' of a server. Later, disk 'C' of the same server might have experienced a space crunch, causing another alarm to be raised.
  - A change in the list of impacted services
- Using the **HISTORY OF ALARMS** page, you can now even view the history of transitions experienced by a particular alarm. For this, just click on an alarm in the **HISTORY OF ALARMS** page. If the alarm has not undergone any transitions, then the **Alarm Transitions** window that appears will once again display the details of the alarm that was clicked on. On the other hand, if the alarm had experienced one/more transitions during the given **Timeline**, then the **Alarm transitions** window will provide the details of each transition - such details include, the alarm priority at the time of the transition, the component name, test, and alarm description during the transition, when the transition began (start time), when it ended (end time), and the total duration of the transition (see Figure 5.11).

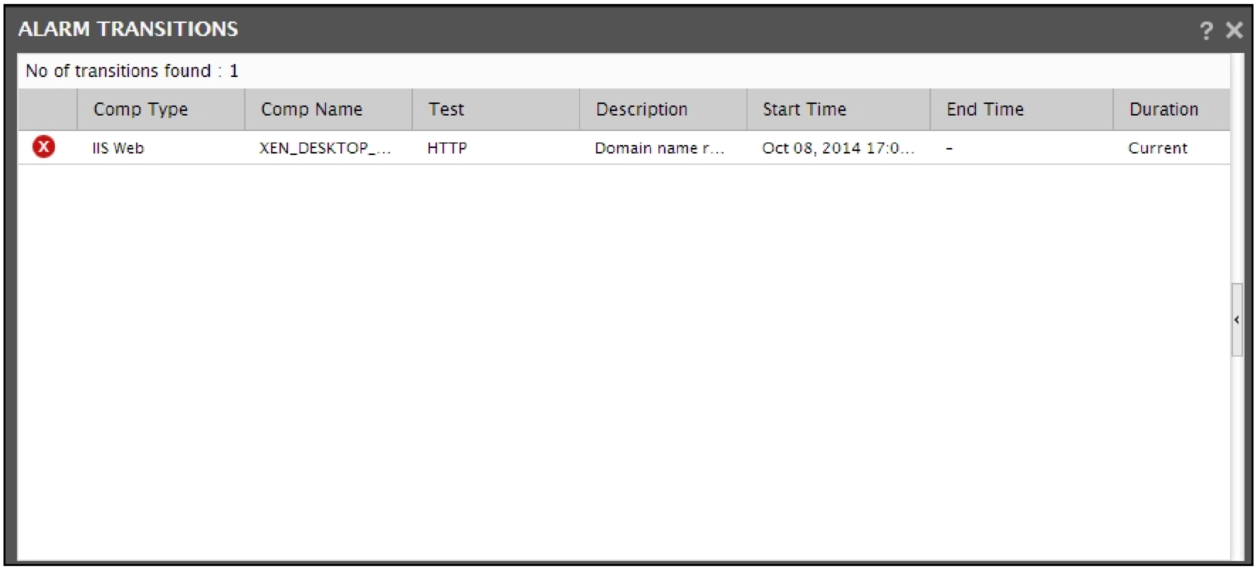

Figure 5.11: Viewing alarm transitions

➢ Using the details provided in the **Alarm transitions** window, you can understand how many transitions have occurred for an alarm in a specified time window, and what they are. To focus only on the state (critical/major/minor) changes that an alarm experienced, click on the left-arrow button to the right of the **Alarm transitions** window. Alternatively, you can click on any of the alarm transitions in this window. This will invoke a distribution pie chart that reveals the percentage of time during the total transition period the alarm has been in the critical, major, and minor states. This reveals how alarm priorities have changed during the entire transition period (see Figure 5.12).

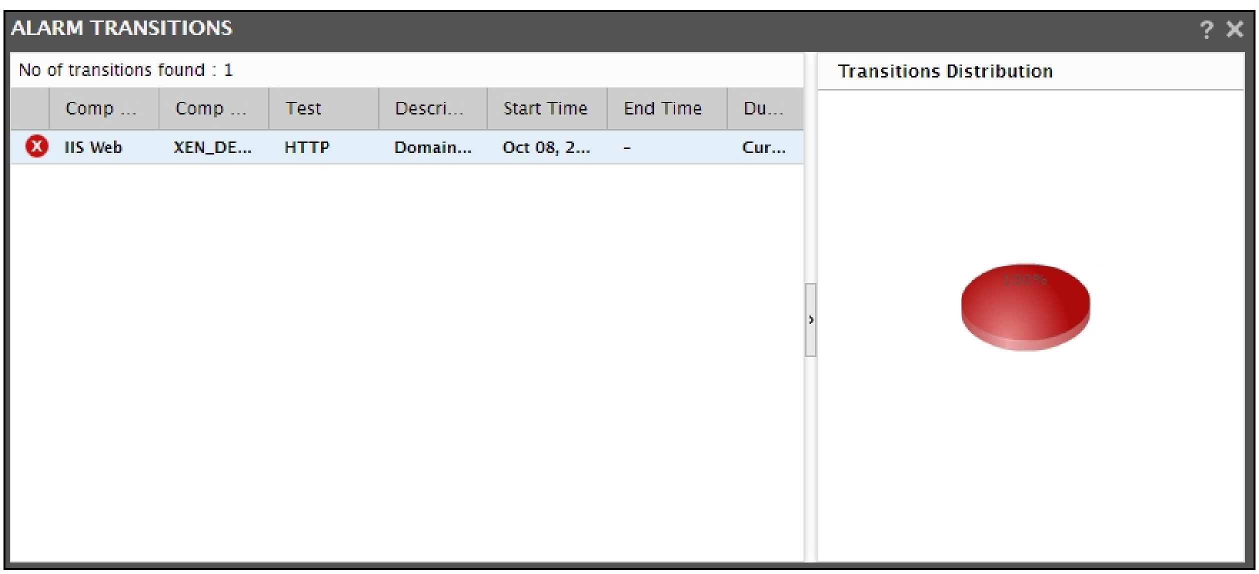

Figure 5.12: Distribution of problems encountered during the transition period

➢ The **HISTORY OF ALARMS** page also comprises of a **GRAPH** icon, which when clicked, allows you to view the graph of the corresponding measure for the last one hour. If the detailed diagnosis capability has been enabled for the eG installation, then problem measures for which detailed diagnosis is available will be accompanied by the **DIAGNOSIS** icon. When this icon is clicked, the detailed diagnosis of the measure will appear, throwing greater light on the problem condition. By default, the graph and detailed diagnosis information will be displayed in the same window as the event history. If you want to view the graph and detailed diagnosis in a separate window, click on the check box preceding the symbol, and then click on the **GRAPH** or **DIAGNOSIS** icons.

- You can save the event history in the CSV format by clicking on the **CSV** button in this page. To save it as a PDF document, click on the **PDF** icon.
- The **NEXT** and **PREVIOUS** buttons, and the hyperlinked page numbers are provided to enable you to easily browse the alarm information that runs across pages.

5.2 The CURRENT ALARMS Window

As stated already, when a user logs in as the *supermonitor*, he/she is entitled to the privileges of the **Supermonitor** role. This means that a *supermonitor* receives an unrestricted view of the environment.

Note:

As discussed already, the monitoring rights of a **Supermonitor** are the same as that of the **Admin** and **ServerAdmin** user roles. Therefore, this chapter will not discuss the **Admin** and **ServerAdmin** roles separately.

To begin with, the *supermonitor* gets to view the **CURRENT ALARMS** window that displays the list of current alarms to the eG Enterprise system (see Figure 5.13), in the order of their priority.

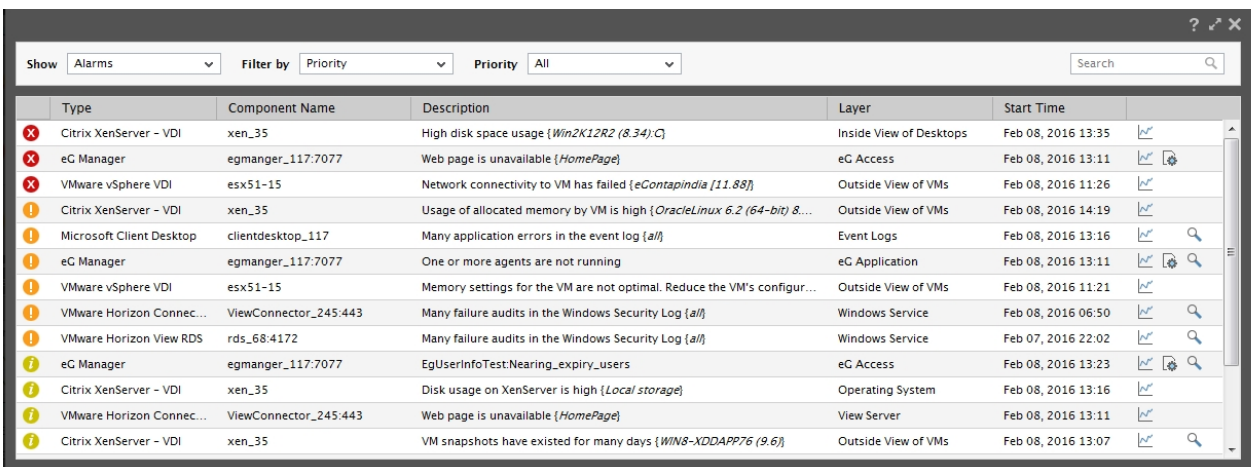

Figure 5.13: The CURRENT ALARMS window displaying the current alarms

Note:

If required, you can make sure that the **CURRENT ALARMS** window does not pop up by default when a user logs into the eG monitoring console. For this, do the following:

- Login to the eG administrative interface.
- Click on the icon available in the **Admin** tab, and select the **Monitor** option from the **Settings** tile.
- Click on the **Alarms** sub-node under the **General** node in the **SETTINGS** tree-structure.

- Set the **Alarm popup** flag that appears in the right panel to **No**. By default, this flag is set to **Yes**.
- Finally, click the **Update** button.

Following conventional management practices, eG Enterprise applies the color-coding scheme mentioned below to indicate alarm priorities:

Color	Priority	Description
Red	Critical	Indicates the existence of a critical issue that requires immediate attention
Orange	Major	Indicates the existence of an issue that could cause serious consequences if not looked into soon.
Yellow	Minor	Indicates the existence of a low priority issue. Proactive alerts fall in this category.

The **CURRENT ALARMS** window (see Figure 5.10) also indicates the problem component-type, IP/host name of the component that has encountered a problem, the layer that has been affected, and the date and time of the problem. Accordingly Figure 5.11 indicates a critical problem with the **Application Processes** layer of the Netscape Application server, and two major problems with the **Web Transactions** and **Web Site** layers of the web server.

Typically, whenever an alarm is raised for problems at the host-level of a component, the **Component type** in the **CURRENT ALARMS** page is automatically set to **Host system**, even if the component affected is a say, Oracle database server or a Web server. The service desk may hence not be able to quickly determine the exact component-type of the affected component from the alarm information. Moreover, help desk personnel may prefer to view the operating system of the problem host as part of the alarm information displayed in the **CURRENT ALARMS** page, as such an information will greatly simplify the troubleshooting process. To make sure that the **CURRENT ALARMS** page enables help desk to easily understand, interpret, and solve problems affecting a host's performance, you can optionally configure the eG Enterprise system to display the actual **Component type**, **Host system**, or the affected **Operating system** for host-level alarms in the **CURRENT ALARMS** page. To enable this capability, do the following:

- Edit the **eg_ui.ini** flag in the **<EG_INSTALL_DIR>\manager\config** directory
- In the **[HOST_SYSTEM]** section of this file, set the **Show_HostSystem** flag to any one of the following values mentioned below:
  - Set the **Show_HostSystem** flag to **HostSystem** if you want the component type to be displayed as **Host system** for the host-level alarms;
  - Set the **Show_HostSystem** flag to **CompType** if you want to display the affected component; This is the default setting that is provided;
  - Set the **Show_HostSystem** flag to **OS** if you want to display the operating system of the host;
- Finally, save the file.

Note:

This configuration affects the **HISTORY OF ALARMS** page, email/SMS alerts, and SNMP traps as well.

To know more about the exact nature of the problem, move your mouse pointer over the alarm displayed in Figure 5.14. Additional alarm information in the form of a brief description of the problem, the test that detected the problem, the test that reported the problem, the host on which the test executed, and the corresponding site name (if any) will be displayed (see Figure 5.14).

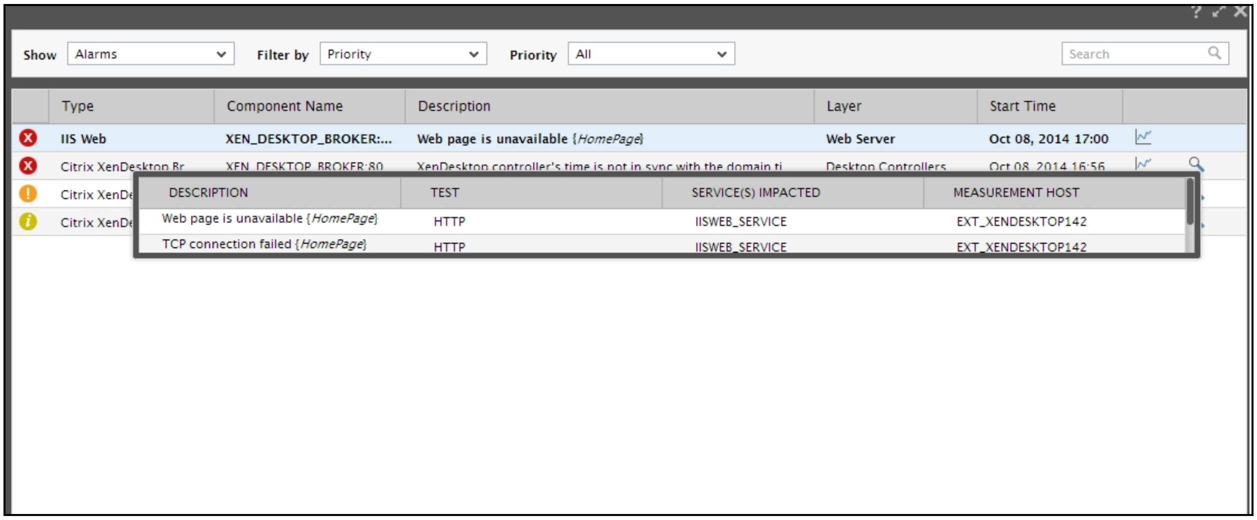

Figure 5.14: Additional alarm information

Note:

The **VALUE** column of the additional alarm details displayed in Figure 5.14 reports the last measure value and unit of the problem measure. The alarms window and email alerts will display the last measure value only if the **Show last measure value in alerts** flag in the **MAIL ALERT CONFIGURATION** section of the **MAIL ALERT PREFERENCES** page is set to **Yes** in the eG admin page.

Besides, a **Graph** icon is available against every alarm. Clicking on this icon invokes a graph of the problem measure for a default period of 1 hour (see Figure 5.15). Using this graph, you can observe the time-of-day variations in the behavior of the problem measure.

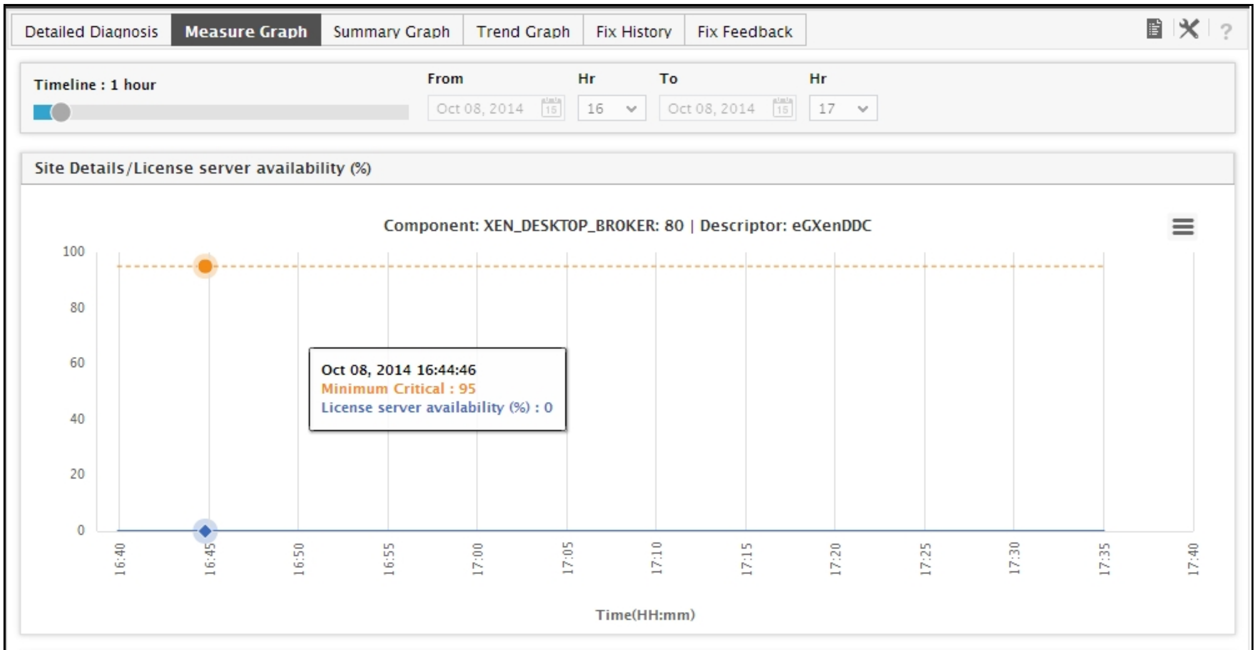

Measurement Name	Avg	Max	Min
License server availability (%)	0	0	0

Figure 5.15: A graph of the problem measure

Note:

To override the default measure graph **Timeline** of 1 hour, do the following:

- Login to the eG administrative interface as *admin*.
- Click on the ▤ icon available in the **Admin** tab, and select the **Monitor** option from the **Settings** tile.
- Select the **Graph** option from the **General Settings** tree Select the Graph option
- Specify a timeline against the **Timeline for graphs** text box in the **GRAPHS** section
- **Update** the changes.

Moreover, to make diagnosis more efficient and accurate, eG Enterprise embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. For example, when the CPU usage of a host reaches the threshold, the agent can be configured to provide more details - e.g., the top 10 process that are consuming more CPU resources. Optionally, this capability can also be configured to periodically generate detailed measures, regardless of the occurrence of problems.

If the detailed diagnostic capability is enabled for the problem measure indicated by the **CURRENT ALARMS** window, then a special **DD** icon will be available against the corresponding alarm (shown in Figure 5.16).

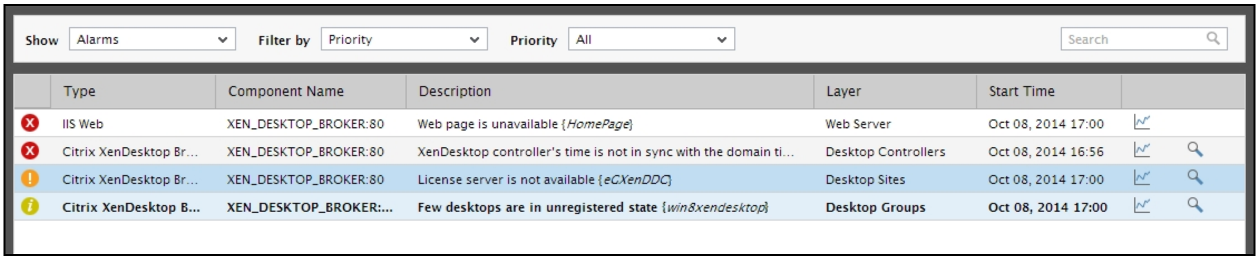

Figure 5.16: The CURRENT ALARMS window displaying the DD icon

Clicking on the icon reveals detailed information pertaining to the problem condition, so that you can quickly and accurately zero-in on the root-cause of the problem. For instance, the **CURRENT ALARMS** window of Figure 5.13 indicates excessive CPU utilization on a host **iis**. Against the 'High CPU utilization' alarm raised on **Processor_0** of this host, you will find a 'magnifying glass' icon. Clicking on this icon will list the top-10 CPU-consuming processes that were executing on **Processor_0** of the host during the last hour (by default), thus enabling you to identify the exact process that is causing the issue (see Figure 5.16). To view the detailed measures related to any other processor supported by the same host, you can pick a different option for analysis from the **Description** list of Figure 5.16 and click the **Submit** button. Similarly, you can change the **Measurement** and the **Timeline** for the detailed diagnosis using Figure 5.17.

Note:

The **DD** icon will not appear in the **CURRENT ALARMS** window under the following circumstances:

- If the eG license does not support the detailed diagnosis capability
- If detailed diagnosis is not enabled for the test that generates the problem measure
- If no detailed measures are available for a problem test

Most administrators will agree that not all performance issues are caused by problems with the internal operations or the external network traffic/connectivity of a component. Sometimes, unplanned/unauthorized/accidental configuration changes can also adversely impact server performance. eG Enterprise optionally provides a dedicated **Configuration Management** module, which enables you to keep track of changes to the configuration of target components and analyze the performance impact of such changes.

Moreover, if the solution captures a configuration change in a component around the same time at which a performance issue was detected with that component, then the **CURRENT ALARMS** window will instantaneously turn your attention to the change by tagging that alarm with symbol, as depicted by Figure 5.17 below:

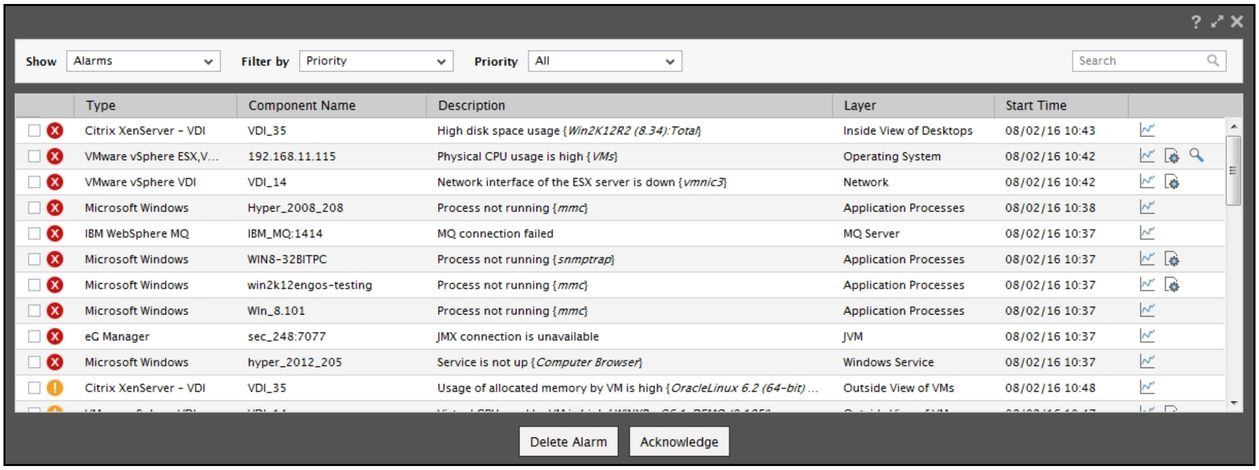

Figure 5.17: Tracking configuration changes in managed components

Clicking on the symbol will open a small pop up window as shown in Figure 5.18 helps you in viewing the recent configuration changes that were made to the problem component from the **CURRENT ALARMS** page itself. This pop-up window will indicate which configuration metric changed, when, what its previous value was, and what it changed to.

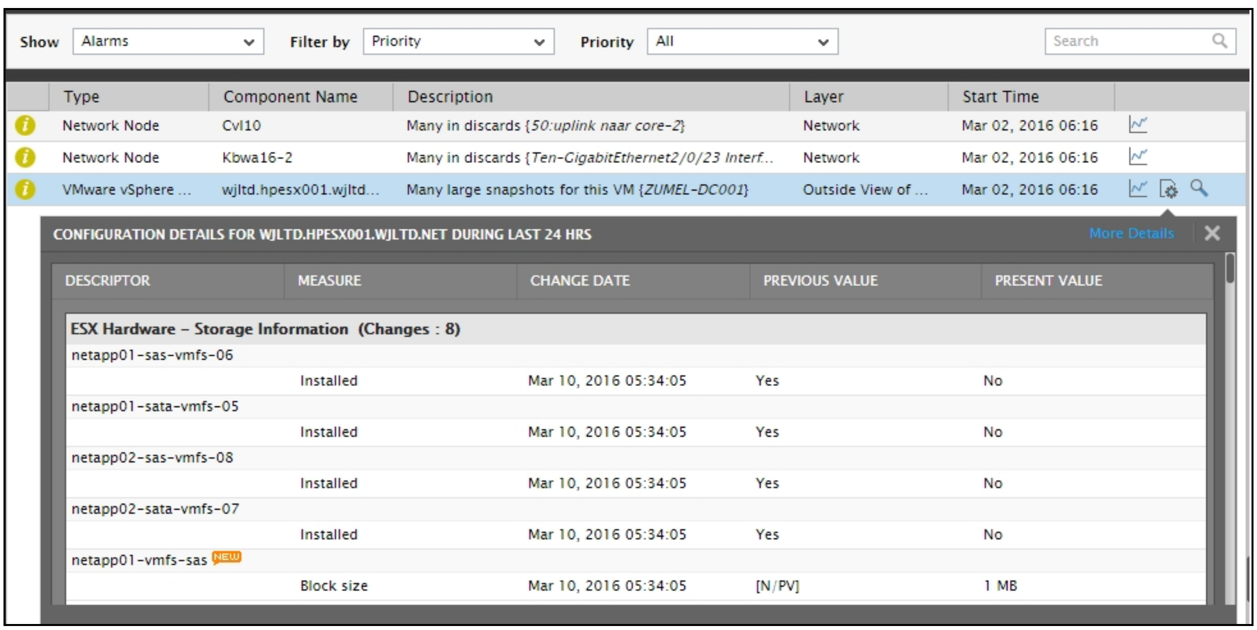

Figure 5.18: Viewing the recent configuration changes to a problem component

Some changes to configuration may not be obvious at first glance. For instance, in the case of configuration metrics with values that run to a few lines, a small change in the middle of a line may go unnoticed, even if both the previous and current values are provided. eG Enterprise zooms into such configuration changes and highlights exactly what has changed, where. Such changes are tagged with the icon at the end (see Figure 5.19).

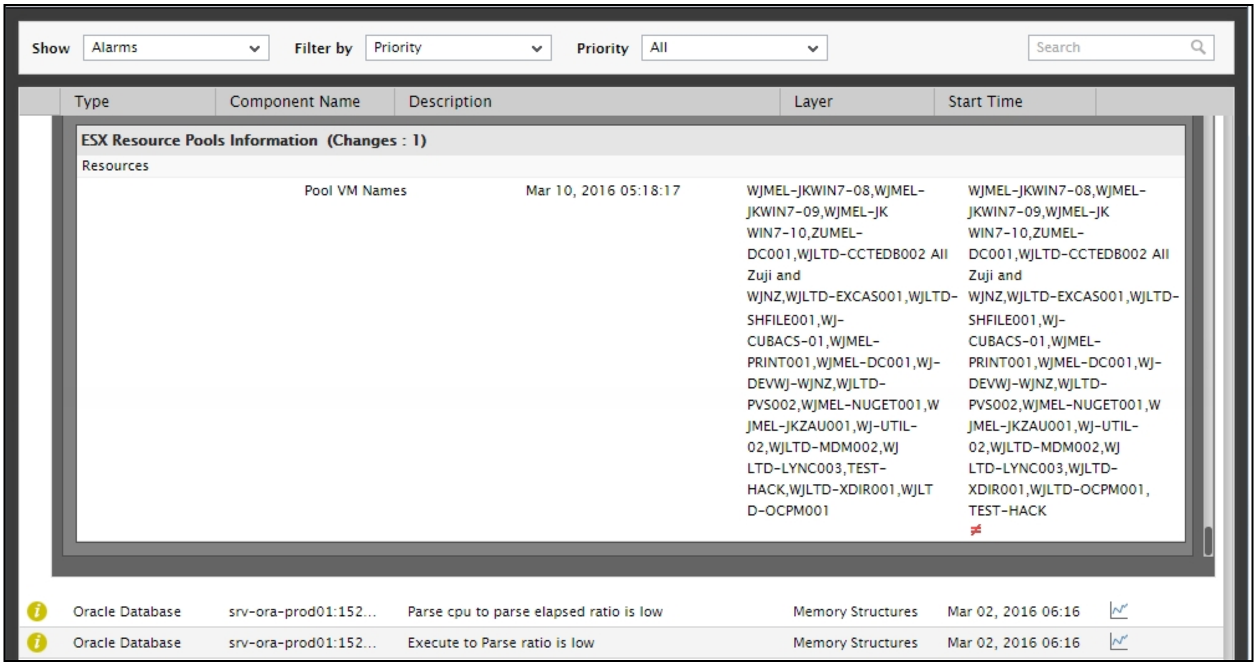

Figure 5.19: An intricate configuration change that has been tagged with a special icon to enable zoom-in

Clicking on this icon will open a **Difference** window (see Figure 5.20). The **Difference** section in this window clearly highlights where changes have taken place. From the **Added** and **Removed** sections, you can instantly figure out what has been newly added to the value of the configuration measure, and/or what has been removed.

Difference

Component	wjltd.hpesx001.wjltd.net:VMware vSphere ESX
Test	ESX Resource Pools Information
Descriptor	Resources
Measure	Pool VM Names
Previous Value	WJMEL-JKWIN7-08,WJMEL-JKWIN7-09,WJMEL-JKWIN7-10,ZUMEL-DC001,WJLTD-CCTEDB002 All Zuji and WJNZ,WJLTD-EXCAS001,WJLTD-SHFILE001,WJ-CUBACS-01,WJMEL-PRINT001,WJMEL-DC001,WJ-DEVWJ-WJNZ,WJLTD-PVS002,WJMEL-NUGET001,WJMEL-JKZAU001,WJ-UTIL-02,WJLTD-MDM002,WJLTD-LYNC003,TEST-HACK,WJLTD-XDIR001,WJLTD-OCPM001
Present Value	WJMEL-JKWIN7-08,WJMEL-JKWIN7-09,WJMEL-JKWIN7-10,ZUMEL-DC001,WJLTD-CCTEDB002 All Zuji and WJNZ,WJLTD-EXCAS001,WJLTD-SHFILE001,WJ-CUBACS-01,WJMEL-PRINT001,WJMEL-DC001,WJ-DEVWJ-WJNZ,WJLTD-PVS002,WJMEL-NUGET001,WJMEL-JKZAU001,WJ-UTIL-02,WJLTD-MDM002,WJLTD-LYNC003,WJLTD-XDIR001,WJLTD-OCPM001,TEST-HACK
Difference	WJMEL-JKWIN7-08,WJMEL-JKWIN7-09,WJMEL-JKWIN7-10,ZUMEL-DC001,WJLTD-CCTEDB002 All Zuji and WJNZ,WJLTD-EXCAS001,WJLTD-SHFILE001,WJ-CUBACS-01,WJMEL-PRINT001,WJMEL-DC001,WJ-DEVWJ-WJNZ,WJLTD-PVS002,WJMEL-NUGET001,WJMEL-JKZAU001,WJ-UTIL-02,WJLTD-MDM002,WJLTD-LYNC003,~~TEST-HACK,~~WJLTD-XDIR001,WJLTD-OCPM001,TEST-HACK
Added	,TEST-HACK
Removed	~~TEST-HACK,~~

Figure 5.20: Zooming into a particular configuration change

By default, a select few metrics alone have been marked by eG for a closer look when configuration changes occur. The value of such measures alone is tagged with the icon. If required, you can override this default list by allowing the 'difference' drill-down for additional metrics. For this, follow the steps below:

1. Edit the **eg_configtests.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
2. In the **[SHOW_DIFFERENCES]** section of the file, you will find some default entries of the following format: *<Internal_name_of_configuration_test>=<Comma-separated list of configuration metrics for which changes need to be highlighted>*
3. If required, you can append more measures of a test to the comma-separated list. You can also add more test-measure combinations to this section. For instance, you can add the following entry to this section: *IpSettings_cf=Description,Ip_add,Subnet,DHCP_server,DNS_server,WINS_server*

   Note that only internal test names and measure names should be specified.
4. Finally, save the file.

Clicking the **More Details** in Figure 5.18 will lead you to Figure 5.21, using which you can instantly figure out the configuration change that may have contributed to the problem at hand.

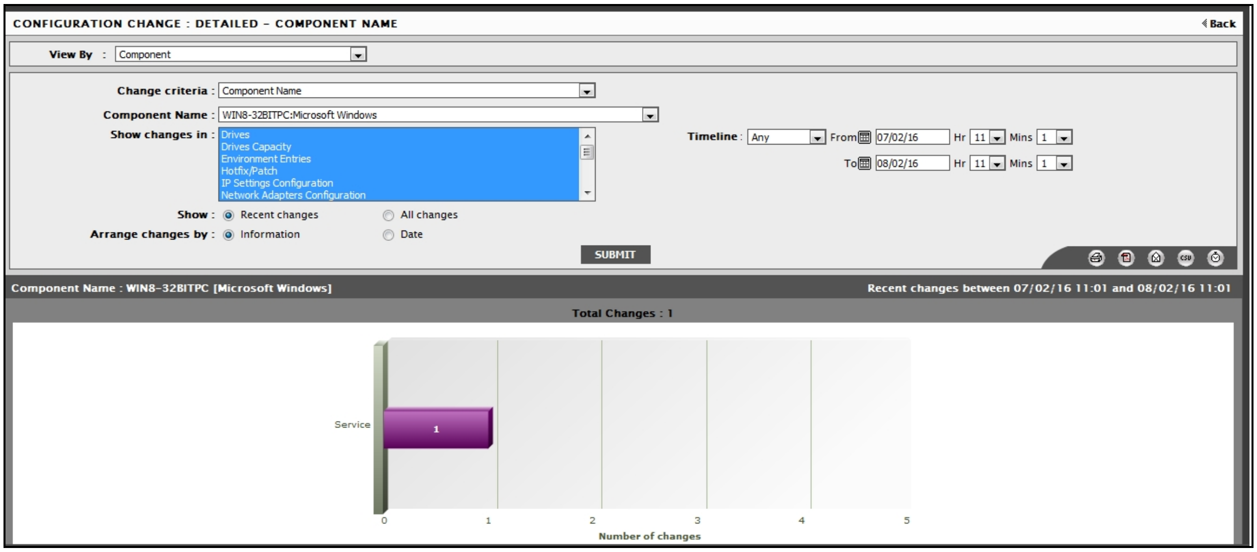

Figure 5.21: Changes in the configuration of a component

By default, the alarm window displays alarms of all priorities. This is indicated by the default selections in the **Filter by** and **Priority** lists. In Figure 5.12, you can see that the *Priority* option is chosen by default from the **Filter by** list, and the *All* option is chosen by default from the **Priority** list.

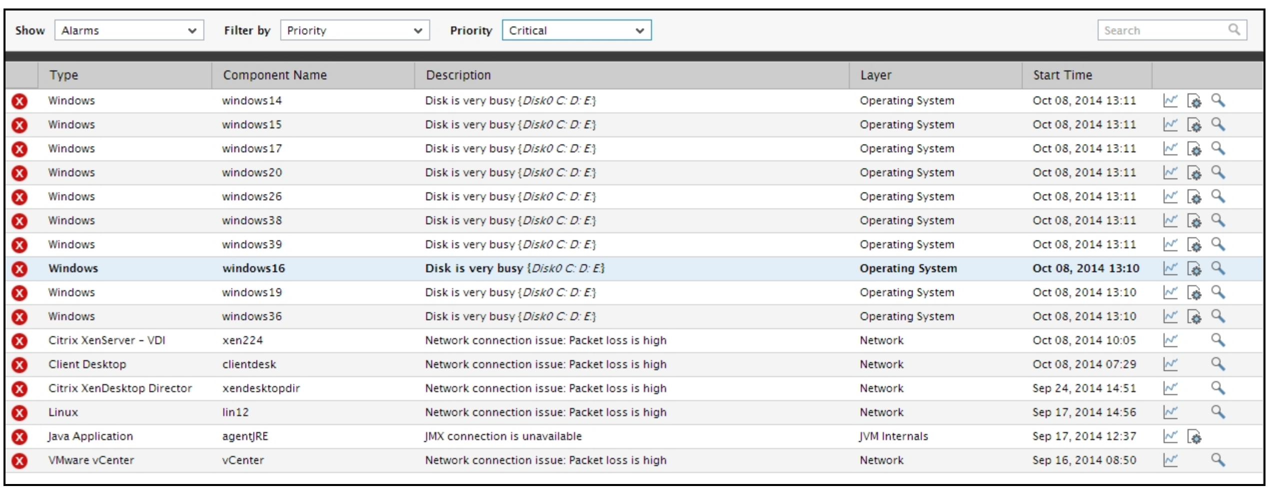

Show: Alarms | Filter by: Priority | Priority: Critical | Search

Type	Component Name	Description	Layer	Start Time
Windows	windows14	Disk is very busy {Disk0 C: D: E}	Operating System	Oct 08, 2014 13:11
Windows	windows15	Disk is very busy {Disk0 C: D: E}	Operating System	Oct 08, 2014 13:11
Windows	windows17	Disk is very busy {Disk0 C: D: E}	Operating System	Oct 08, 2014 13:11
Windows	windows20	Disk is very busy {Disk0 C: D: E}	Operating System	Oct 08, 2014 13:11
Windows	windows26	Disk is very busy {Disk0 C: D: E}	Operating System	Oct 08, 2014 13:11
Windows	windows38	Disk is very busy {Disk0 C: D: E}	Operating System	Oct 08, 2014 13:11
Windows	windows39	Disk is very busy {Disk0 C: D: E}	Operating System	Oct 08, 2014 13:11
Windows	windows16	Disk is very busy {Disk0 C: D: E}	Operating System	Oct 08, 2014 13:10
Windows	windows19	Disk is very busy {Disk0 C: D: E}	Operating System	Oct 08, 2014 13:10
Windows	windows36	Disk is very busy {Disk0 C: D: E}	Operating System	Oct 08, 2014 13:10
Citrix XenServer - VDI	xen224	Network connection issue: Packet loss is high	Network	Oct 08, 2014 10:05
Client Desktop	clientdesk	Network connection issue: Packet loss is high	Network	Oct 08, 2014 07:29
Citrix XenDesktop Director	xendesktopdir	Network connection issue: Packet loss is high	Network	Sep 24, 2014 14:51
Linux	lin12	Network connection issue: Packet loss is high	Network	Sep 17, 2014 14:56
Java Application	agentJRE	JMX connection is unavailable	JVM Internals	Sep 17, 2014 12:37
VMware vCenter	vCenter	Network connection issue: Packet loss is high	Network	Sep 16, 2014 08:50

Figure 5.22: Viewing the critical and major alarms

To view only the critical alarms, select the **Critical** option from the **Priority** list box (see Figure 5.22).

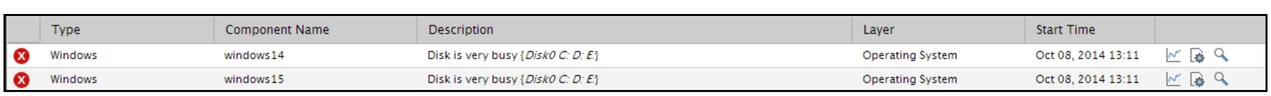

Type	Component Name	Description	Layer	Start Time
Windows	windows14	Disk is very busy {Disk0 C: D: E}	Operating System	Oct 08, 2014 13:11
Windows	windows15	Disk is very busy {Disk0 C: D: E}	Operating System	Oct 08, 2014 13:11

Figure 5.23: Viewing the critical alarms

Likewise, you can view the **Critical & Major** alarms together, or view the **Major** or **Minor** alarms alone by selecting the corresponding options from the **Priority** list.

Besides **Priority**, alarms can also be filtered on the basis of **Component Type**, **Services**, **Segments**, or **Zones**. For instance to view the alarms pertaining to a particular component type alone, pick the *Component Type* option from the **Filter by** list, and then select a component type of your choice from the **Types** list (see Figure 5.22). For instance, Citrix administrators would typically be more concerned with issues pertaining to their mission-critical Citrix XenApp installations. To focus on Citrix-related issues alone, Citrix administrators can filter the alarms list by selecting *Component Type* from the **Filter by** list and then choosing the *Citrix XenApp* option from the **Types** list (see Figure 5.24).

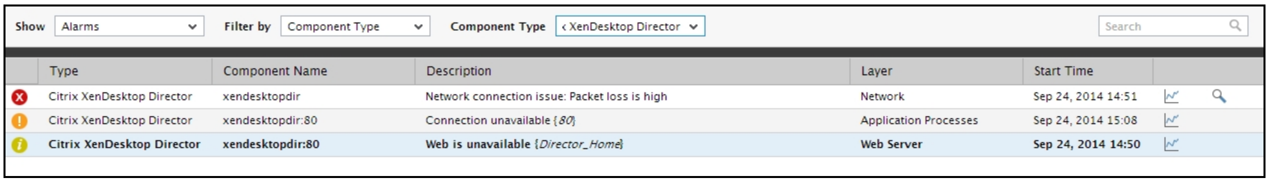

Figure 5.24: Filtering the alarms based on Component Type

Likewise, service managers can filter the alarms list to view only those alarms that are impacting a particular business service's performance. For this, they need to select the *Services* option from the **Filter by** list, and pick a service of interest to them from the **Services** list (see Figure 5.24).

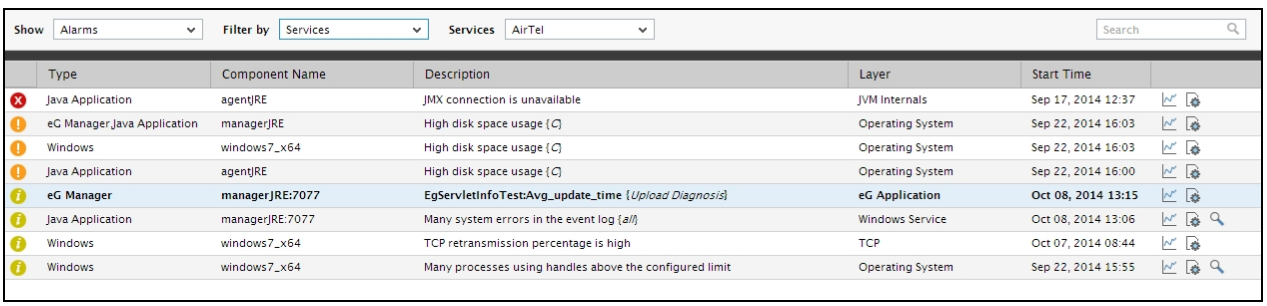

Figure 5.25: Filtering the alarms based on Service name

In the same way, performance degradations experienced by the components in a segment/zone can also be viewed in the **CURRENT ALARMS** window.

You can even run quick searches on the alarm window to locate alarms of interest to you. The default criterion for such searches is component type; this is indicated by the default selection, *Type*, in the **Search by** list in Figure 5.25. To search for alarms related to a component type, specify the whole/part of the name of that component type in the text box next to the **Search by** list and click the 'magnifying glass' icon adjacent to it. All alarms related to component types with names that match the specified search string will then appear in the alarms window (see Figure 5.25).

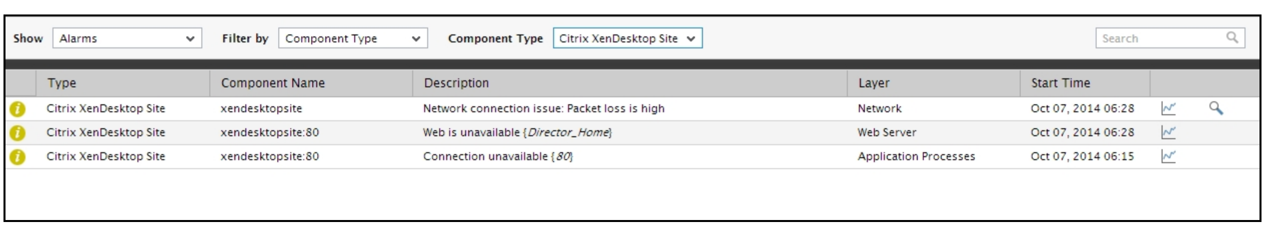

Figure 5.26: Searching for the alarms related to a particular component type

You can even search for alarms related to a **Component**, alarm **Description**, or **Layer**, by picking the desired option from the **Search by** list, and specifying a search string in the text box next to it.

Also, with a single mouse click, you can change the order in which the alarms are sorted in the **CURRENT ALARMS** window. By default, alarms are sorted in the descending order of the **START TIME** of issues. To arrange them in the ascending order of **START TIME**, simply click on the column label - **START TIME** - in Figure 5.26. The current sort order will be depicted by an 'arrow' symbol in the sorted column - while an 'up arrow' symbol signifies the ascending order, the 'down arrow' denotes the descending order. This way, you can quickly arrange the contents of the alarms window in the ascending/descending order of any of the displayed columns.

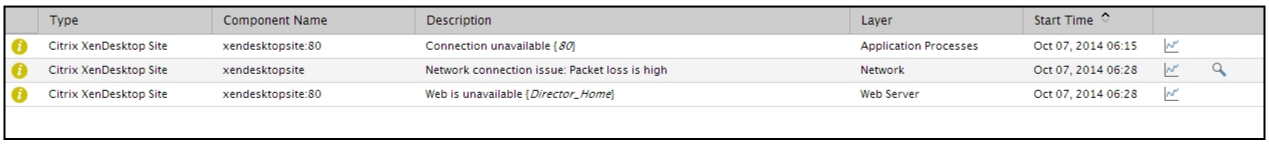

Type	Component Name	Description	Layer	Start Time
Citrix XenDesktop Site	xendesktopsite:80	Connection unavailable {80}	Application Processes	Oct 07, 2014 06:15
Citrix XenDesktop Site	xendesktopsite	Network connection issue: Packet loss is high	Network	Oct 07, 2014 06:28
Citrix XenDesktop Site	xendesktopsite:80	Web is unavailable {Director_Home}	Web Server	Oct 07, 2014 06:28

Figure 5.27: Sorting the alarms in the ascending order of the START TIME of issues

In addition to the above, the option to **DELETE** alarms can be enabled for specific monitor users registered with the eG Enterprise system. While creating/modifying the profile of a user using the eG administrative interface, you can set the **Allow alarm deletion** flag to **Yes** for that user, if you want to grant him/her the right to delete alarms. By default, the alarm deletion capability is disabled for all users (including the *admin* and *supermonitor* users) to the eG monitoring console.

If the capability has been explicitly enabled for a user, say the *supermonitor*, then the **ALARMS** window will display an additional **DELETE** button. To delete an alarm, select the check box corresponding to the alarm, and then click the **DELETE** button.

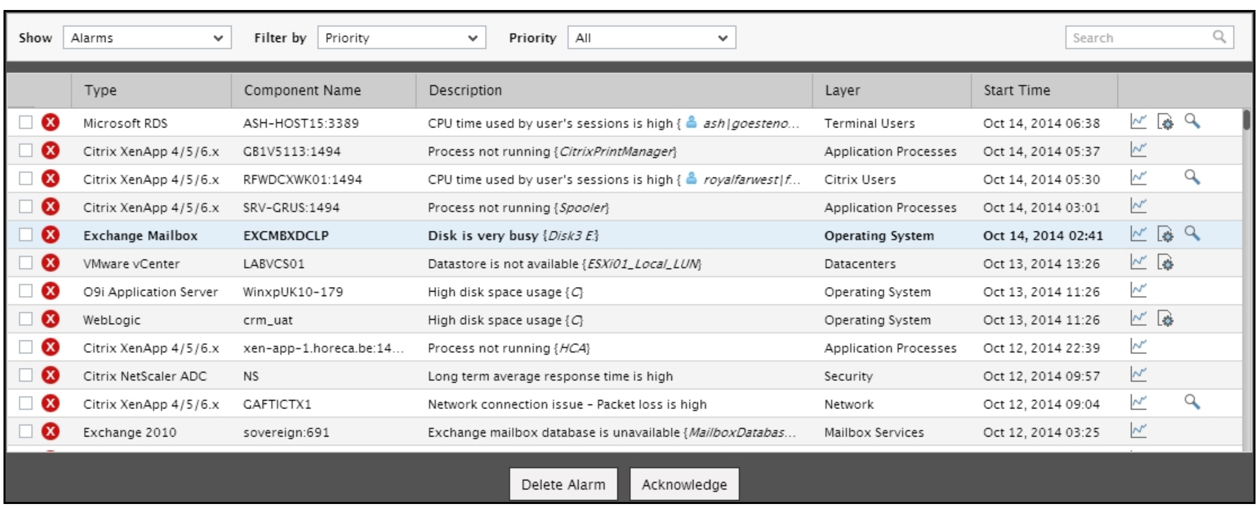

Show Alarms Filter by Priority Priority All Search

Type	Component Name	Description	Layer	Start Time
Microsoft RDS	ASH-HOST15:3389	CPU time used by user's sessions is high { ash|goesteno...	Terminal Users	Oct 14, 2014 06:38
Citrix XenApp 4/5/6.x	GB1V5113:1494	Process not running {CitrixPrintManager}	Application Processes	Oct 14, 2014 05:37
Citrix XenApp 4/5/6.x	RFWDCXWK01:1494	CPU time used by user's sessions is high { royalfarwest|f...	Citrix Users	Oct 14, 2014 05:30
Citrix XenApp 4/5/6.x	SRV-GRUS:1494	Process not running {Spooler}	Application Processes	Oct 14, 2014 03:01
Exchange Mailbox	**EXCMBXDCLP**	**Disk is very busy {Disk3 E}**	**Operating System**	**Oct 14, 2014 02:41**
VMware vCenter	LABVCS01	Datastore is not available {ESXi01_Local_LUN}	Datacenters	Oct 13, 2014 13:26
O9i Application Server	WinxpUK10-179	High disk space usage {C}	Operating System	Oct 13, 2014 11:26
WebLogic	crm_uat	High disk space usage {C}	Operating System	Oct 13, 2014 11:26
Citrix XenApp 4/5/6.x	xen-app-1.horeca.be:14...	Process not running {HCA}	Application Processes	Oct 12, 2014 22:39
Citrix NetScaler ADC	NS	Long term average response time is high	Security	Oct 12, 2014 09:57
Citrix XenApp 4/5/6.x	GAFTICTX1	Network connection issue - Packet loss is high	Network	Oct 12, 2014 09:04
Exchange 2010	sovereign:691	Exchange mailbox database is unavailable {MailboxDatabas...	Mailbox Services	Oct 12, 2014 03:25

Delete Alarm Acknowledge

Figure 5.28: The ALARMS window with the option to ACKNOWLEDGE alarms

Doing so will invoke Figure 5.28, where you can provide a **Reason** for deleting the chosen alarm. You may also decide not to provide any **Reason** for the deletion, if you so desire. Click the **Submit** button to save the reason, and to confirm the deletion of the alarm.

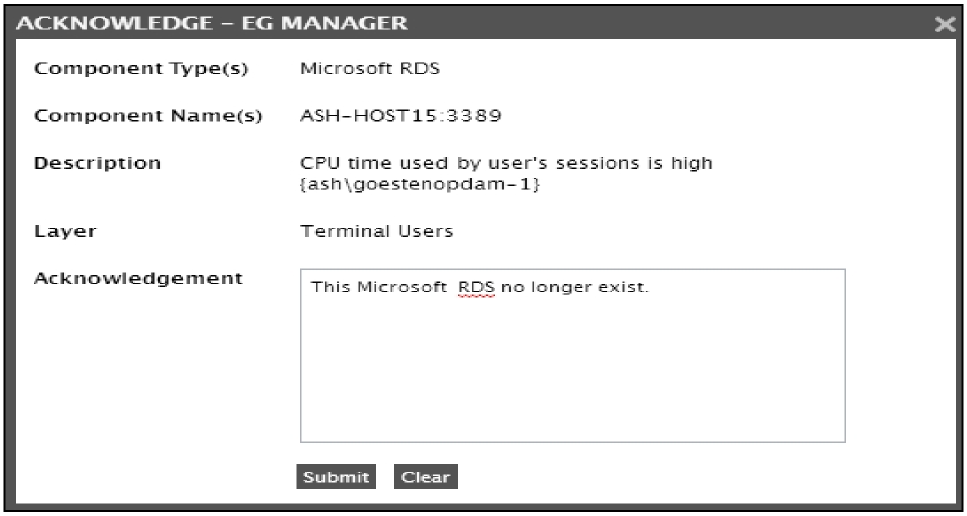

Figure 5.29: Providing a reason for the alarm deletion

Note:

- If required, you can delete multiple alarms simultaneously, by selecting the check boxes corresponding to the alarms and clicking the **Delete** button. A common **Reason** can then be provided for such a deletion.
- Even after an alarm is deleted, the problem measure for which the alarm was originally raised will continue to remain in the abnormal state, until the problem in question is actually resolved. However, the component, layer, and test associated with the problem measure will be in the normal state.

Optionally, an **acknowledgement** can be provided for an alarm displayed in the eG monitor interface (see Figure 5.30). 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 working on 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

In the **CURRENT ALARMS** window of Figure 5.16 select the check box corresponding to the alarm to be acknowledged, and click on the **ACKNOWLEDGE** button therein. This button will appear only if the current user is authorized to either acknowledge an alarm / edit an acknowledgement.

Figure 5.28 will then appear using which the alarm can be acknowledged. To save the acknowledgement, click the **SUBMIT** button in Figure 5.29.

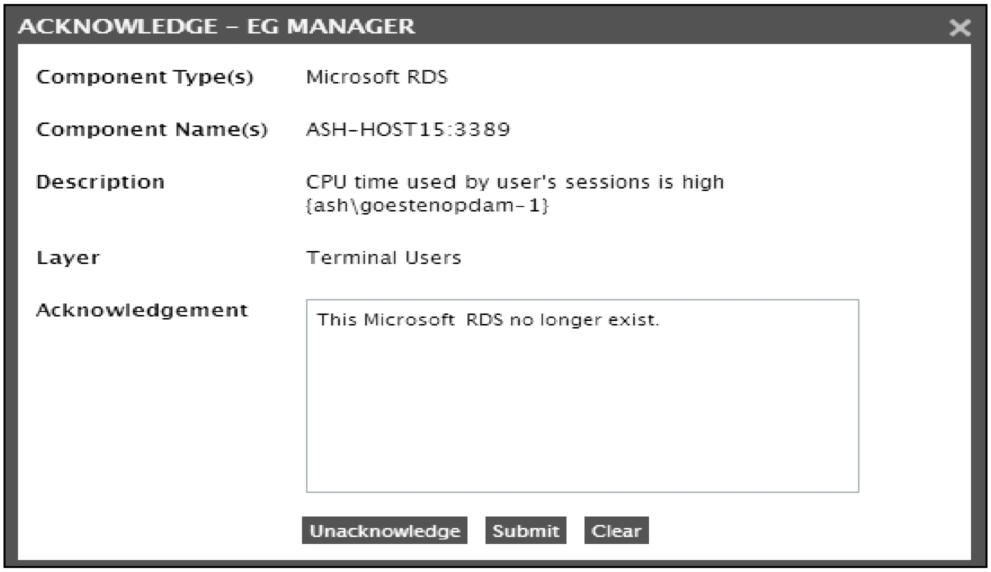

Figure 5.30: Submitting an acknowledgement

Doing so will lead you back to the **CURRENT ALARMS** window, but this time, a symbol indicated by Figure 5.30 will prefix the acknowledged alarm. Moving your mouse pointer over the symbol will reveal the details of the acknowledgement such as its description, and the user who has acknowledged the alarm, and the date and time specifications of the acknowledgement.

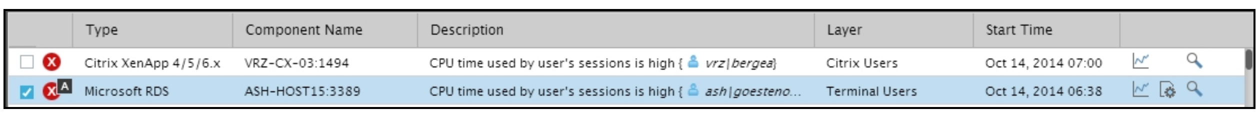

Figure 5.31: Figure 5.31: The acknowledgement description

In large environments, it is but natural that the same set of components are assigned to multiple users for monitoring. In such environments, some/all the users with monitoring rights to a component might want to post their comments for an alarm related to that component. If acknowledgment rights are granted to all these users, then each of them can login to the monitor interface and provide an acknowledgement description for the same alarm using the procedure discussed in Section 1 above.

eG Enterprise maintains a history of the acknowledgement descriptions provided by multiple users with rights to monitor a single component, and lists the entire history the next time one of these users attempts to view the acknowledgement details (see Figure 5.31) in the **CURRENT ALARMS** window. This way, the administrative staff can share the responsibility for resolving critical issues, brainstorm online to identify accurate remedies, and even provide each other with quick updates on problem status.

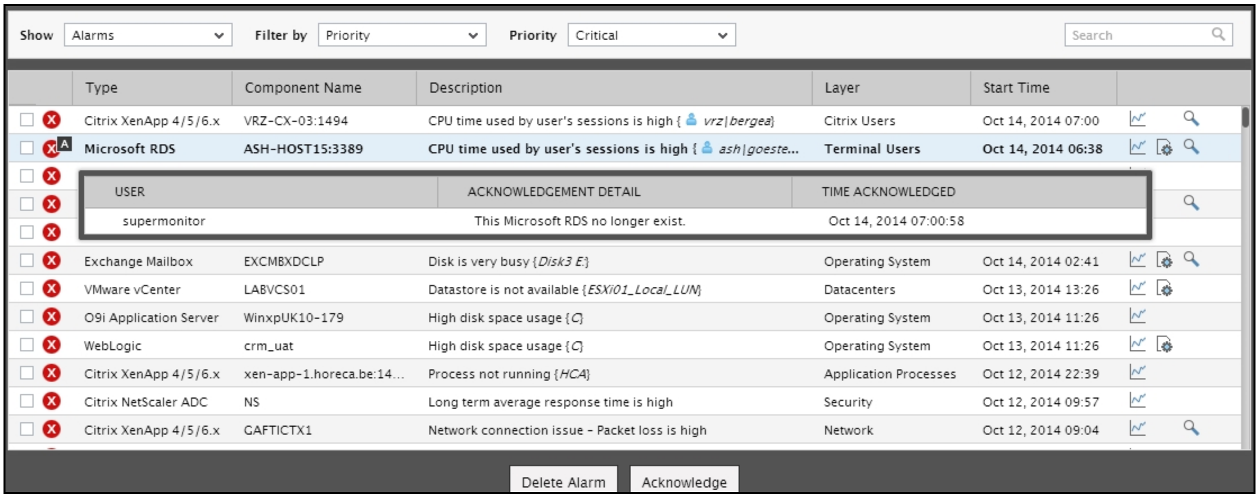

Figure 5.32: Details of the acknowledgement

An alarm can also be unacknowledged, but only by the user who originally submitted the acknowledgement.

For unacknowledging, once again, select the check box corresponding to the acknowledged alarm in the **CURRENT ALARMS** window and click the **ACKNOWLEDGE** button therein. When Figure 5.31 appears, click on the **UNACKNOWLEDGE** button in it. This will make sure that the 'acknowledgement' symbol is removed from **CURRENT ALARMS** window.

Note:

A user can edit/unacknowledge only those acknowledgement descriptions that he/she originally provided.

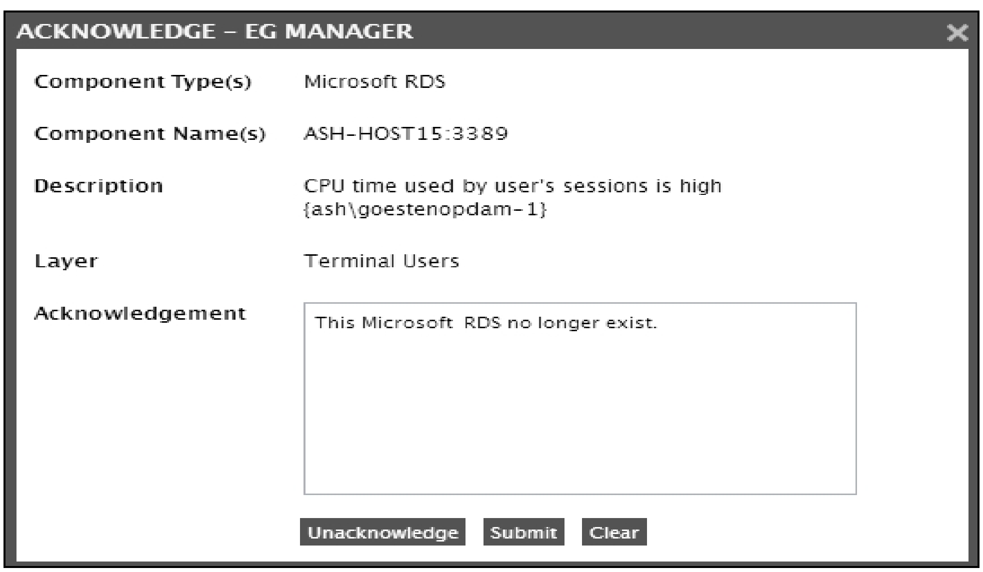

Figure 5.33: Unacknowledging an alarm

The alarm window will refresh periodically to show the latest set of alarms.

Note:

The default refresh period of the alarm page is 60 seconds. This however, can be changed by modifying **How frequently alarm page is refreshed** in the **REFRESH** section of the **MONITOR SETTINGS** page in the eG administrative interface.

Clicking on an alarm will lead monitor users to a page that displays the layer model, tests, and measurements pertaining to the problem component (see Figure 5.33).

Besides open issues, you can also use the **CURRENT ALARMS** window to view the tests that are currently in an **UNKNOWN** state in the environment. A test can switch to an **UNKNOWN** state when the eG Enterprise system is unable to determine the state of one/more metrics that test reports - this could be because of any of the following reasons:

- The test could have been wrongly configured;
- The eG agent executing the test could have suddenly stopped;
- The eG agent may have been unable to pull out metrics from the server;
- The eG agent executing the test may not be able to transmit the metrics collected by that test to the eG manager, owing to say, poor network connectivity;

To enable administrators to receive a heads-up on the **Unknown** tests in the environment, so that issues such as the ones mentioned above can be isolated and resolved, the **Show** list in the left top corner of the **CURRENT ALARMS** window includes a special **Unknowns** option.

Selecting the **Unknowns** option from the **Show** list invokes Figure 5.33, where the **TEST**s in an indeterminate (i.e., **UNKNOWN**) state, the layers to which the tests are mapped, the names of the components they are associated with, and their corresponding component types will be listed.

Show: Unknowns   Filter by: Component Type   Component Type: All   Search

Type	Component Name	Layer	Test
2X Client Gateway	GW1.ash.local:80	Web Server	HTTP
Active Directory	LABDC01:389	AD Service	Global Catalog Search
Cisco ASA	Sparrow	ASA Tunnels	Ike Global Tunnels
Cisco ASA	Sparrow	ASA Hardware	ASA Hardware Status
Cisco ASA	Sparrow	ASA Sessions	ASA Sessions
Cisco ASA	Sparrow	ASA Hardware	ASA Memory Details
Cisco ASA	Sparrow	Network	Network
Cisco ASA	Sparrow	ASA Sessions	ASA Remote Access Sessions
Cisco ASA	Sparrow	ASA Tunnels	Ike Secondary Tunnels
Cisco ASA	Sparrow	ASA Hardware	ASA CPU Details
Cisco ASA	Sparrow	Network	Network Interfaces
Cisco Catalyst Switch	ABAT	Network	Device Uptime
Cisco Catalyst Switch	budgetnet205g	Operating System	Cisco Temperature
Cisco Catalyst Switch	Block-18	Network	Network

Figure 5.34: The components, tests, and metrics in an Unknown state

By moving your mouse pointer over any test in Figure 5.34, you can determine which measure(s) of the test is in the **UNKNOWN** state currently and which **SERVICES** are impacted by this (see Figure 5.35).

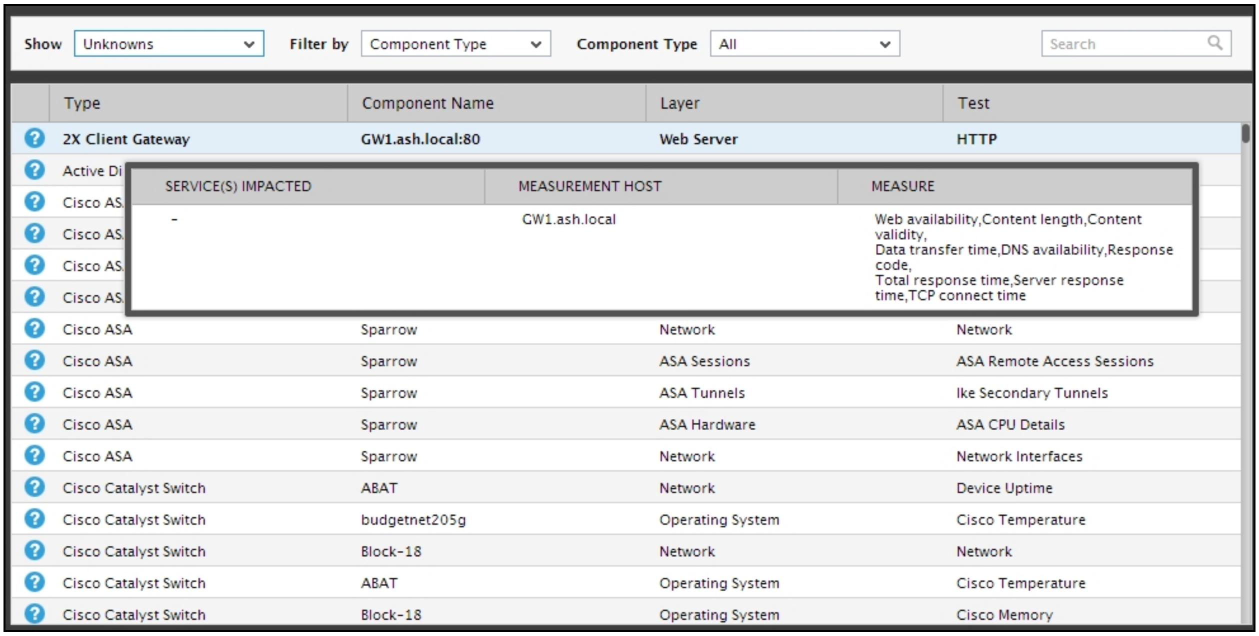

Figure 5.35: Viewing additional information related to the Unknown alerts

By default, the eG Enterprise system displays the **UNKNOWN** tests across all managed component types in the environment in Figure 5.34 and Figure 5.35. This is why, the **Filter by** list is set to **Component Type** by default and the **Types** list is set to **All** by default in Figure 5.36. To view the **Unknowns** related to a specific component type, pick anytype of your choice from the **Types** list. Doing so will list only those **UNKNOWN** tests that are associated with the components of the chosen **Type** (see Figure 5.36).

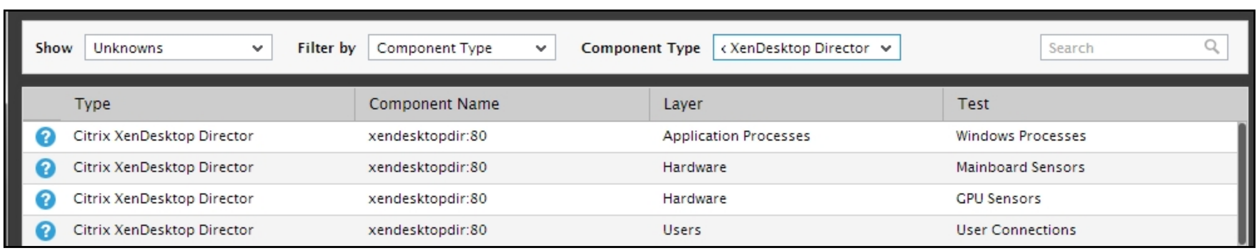

Figure 5.36: Viewing the details of Unknowns related to a specific component type

You can even filter the **Unknowns** list on the basis of service and segment names, so that you can quickly figure out which service/segment components have been affected by **Unknowns** and what tests and metrics are causing this. For instance, to view the **Unknown** tests in a service, select **Services** from the **Filter by** list of Figure 5.36 and pick a service from the **Services** list. The service components with **Unknowns**, their corresponding component types, the layers to which the **UNKNOWN** tests are mapped, and the **UNKNOWN** tests themselves will then be listed, as depicted by Figure 5.36.

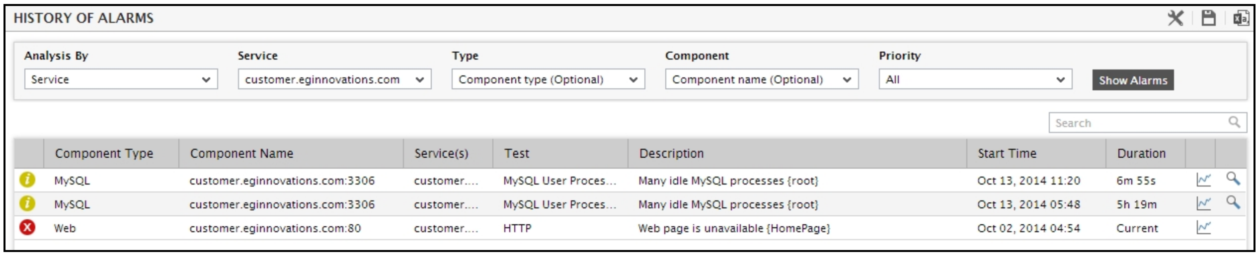

Figure 5.37: Viewing the list of Unknowns in a particular service

Likewise, to view the details of **Unknowns** in a particular segment, pick **Segments** from the **Filter by** list, and select the **Segment** of interest to you. 5.2 will then appear, where you can view the segment components with **Unknowns**, their corresponding component types, the layers to which the **UNKNOWN** tests are mapped, and the **UNKNOWN** tests themselves.

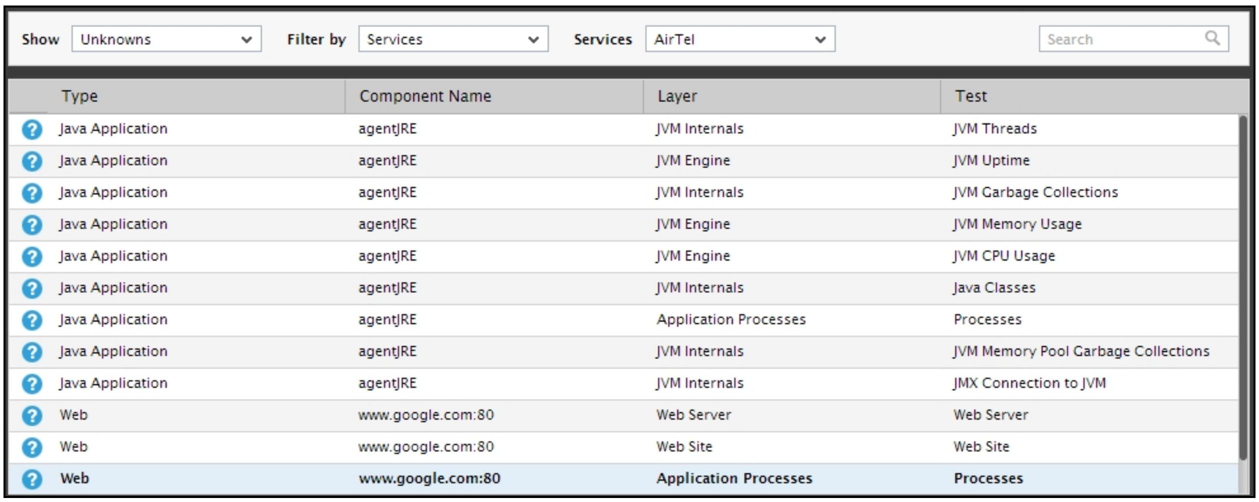

Figure 5.38: Viewing the list of Unknowns in a particular segment

Where there are numerous components with **Unknowns**, you can quickly zoom into the **Unknowns** of specific interest to you using the search utility that the **Unknowns** window embeds. For instance, to view the **Unknowns** related to particular component type, pick the **Type** criterion from the **Search by** list and specify the whole or part of the component type to search for in the text box next to it. Then, click on the 'magnifying glass' icon within. This way, you can instantly access information related to the **Unknowns** in all component types with names that match the specified search string (see Figure 5.38).

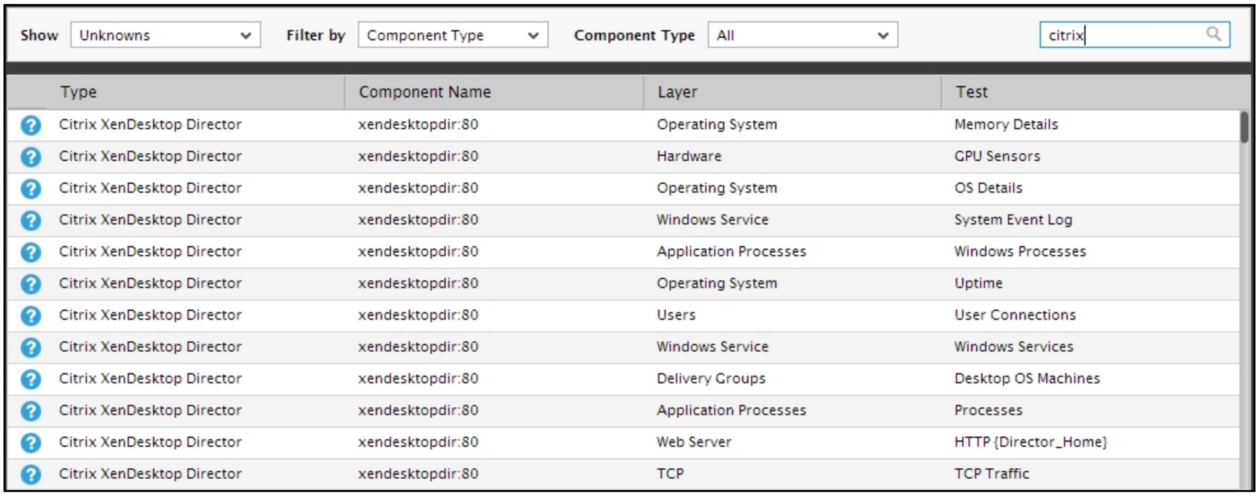

Show: Unknowns | Filter by: Component Type | Component Type: All | citrix

Type	Component Name	Layer	Test
Citrix XenDesktop Director	xendesktopdir:80	Operating System	Memory Details
Citrix XenDesktop Director	xendesktopdir:80	Hardware	GPU Sensors
Citrix XenDesktop Director	xendesktopdir:80	Operating System	OS Details
Citrix XenDesktop Director	xendesktopdir:80	Windows Service	System Event Log
Citrix XenDesktop Director	xendesktopdir:80	Application Processes	Windows Processes
Citrix XenDesktop Director	xendesktopdir:80	Operating System	Uptime
Citrix XenDesktop Director	xendesktopdir:80	Users	User Connections
Citrix XenDesktop Director	xendesktopdir:80	Windows Service	Windows Services
Citrix XenDesktop Director	xendesktopdir:80	Delivery Groups	Desktop OS Machines
Citrix XenDesktop Director	xendesktopdir:80	Application Processes	Processes
Citrix XenDesktop Director	xendesktopdir:80	Web Server	HTTP {Director_Home}
Citrix XenDesktop Director	xendesktopdir:80	TCP	TCP Traffic

Figure 5.39: Searching for the details of Unknowns in a particular component

Besides component **Type**, you can build search conditions on the basis of **Component** name, **Test** name, and **Layer** name by selecting the desired search criterion from the **Search by** list and specifying the corresponding search string in the adjacent text box.

Also, with a single mouse click, you can change the order in which the **Unknowns** are sorted. By default, **Unknowns** are sorted in the ascending order of the contents of the **TYPE** column. This is indicated by the 'up arrow' symbol next to the column label - **TYPE**. To arrange the **Unknowns** in the descending order of **TYPE**, just click on that column label. This will automatically sort the **Unknowns** list in the descending order of component type and will suffix the column heading **TYPE** with a 'down arrow' symbol (see Figure 5.39). To sort the **Unknowns** on the basis of any other column, simply click on that column heading.

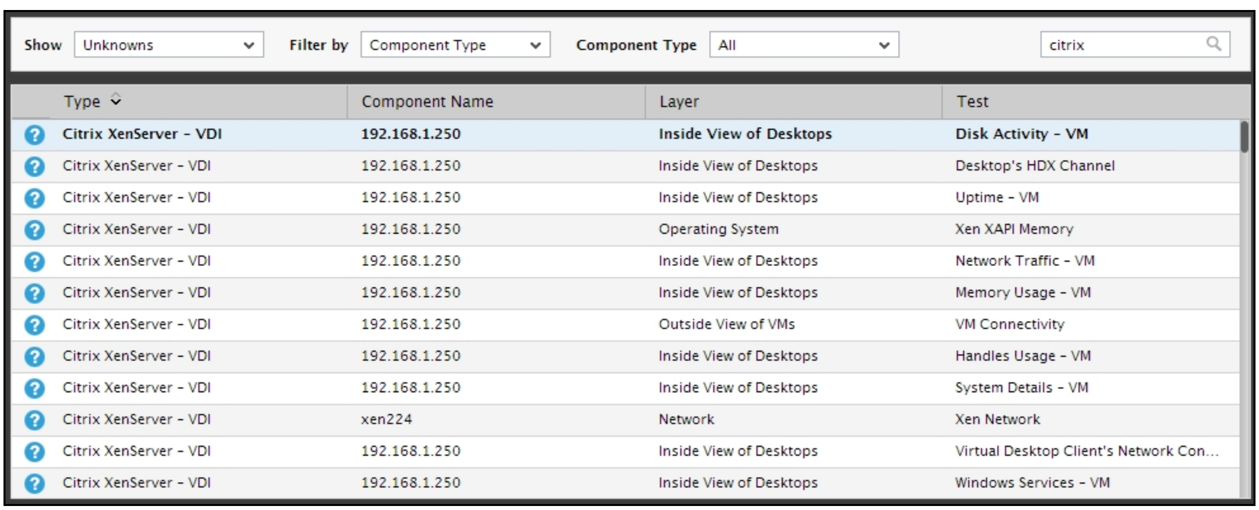

Show: Unknowns | Filter by: Component Type | Component Type: All | citrix

Type	Component Name	Layer	Test
Citrix XenServer - VDI	**192.168.1.250**	**Inside View of Desktops**	**Disk Activity - VM**
Citrix XenServer - VDI	192.168.1.250	Inside View of Desktops	Desktop's HDX Channel
Citrix XenServer - VDI	192.168.1.250	Inside View of Desktops	Uptime - VM
Citrix XenServer - VDI	192.168.1.250	Operating System	Xen XAPI Memory
Citrix XenServer - VDI	192.168.1.250	Inside View of Desktops	Network Traffic - VM
Citrix XenServer - VDI	192.168.1.250	Inside View of Desktops	Memory Usage - VM
Citrix XenServer - VDI	192.168.1.250	Outside View of VMs	VM Connectivity
Citrix XenServer - VDI	192.168.1.250	Inside View of Desktops	Handles Usage - VM
Citrix XenServer - VDI	192.168.1.250	Inside View of Desktops	System Details - VM
Citrix XenServer - VDI	xen224	Network	Xen Network
Citrix XenServer - VDI	192.168.1.250	Inside View of Desktops	Virtual Desktop Client's Network Con...
Citrix XenServer - VDI	192.168.1.250	Inside View of Desktops	Windows Services - VM

Figure 5.40: Unknowns arranged in the descending order of component type

5.3 Acknowledgement History

eG Enterprise stores and displays all the acknowledgement descriptions that are associated with a particular alarm, so that the help desk staff can consult with each other online, exchange views, and thus arrive at effective solutions to persistent performance issues with target components. The eG monitoring interface presents this acknowledgement history to you in various ways. While the **CURRENT ALARMS** window provides you with the complete acknowledgement history of a current issue, the **EVENT HISTORY** page leads you to all acknowledgements associated with a problem that occurred in the recent/distant past.

Sometimes however, to perform better problem diagnosis, you might want to review specific acknowledgment descriptions associated with an alarm and not all of them. For instance, while two users - *john* and *elvis* - may have acknowledged an alarm raised on an Oracle database server, you might only want to view user *john's* acknowledgement description. To facilitate such selective viewing of acknowledgement information, eG Enterprise provides a dedicated **ACKNOWLEDGEMENT HISTORY** page; this page provides a wide variety of filter options with the help of which you can quickly and easily run a search across all alarm acknowledgements, and swiftly locate the acknowledgment information of interest to you.

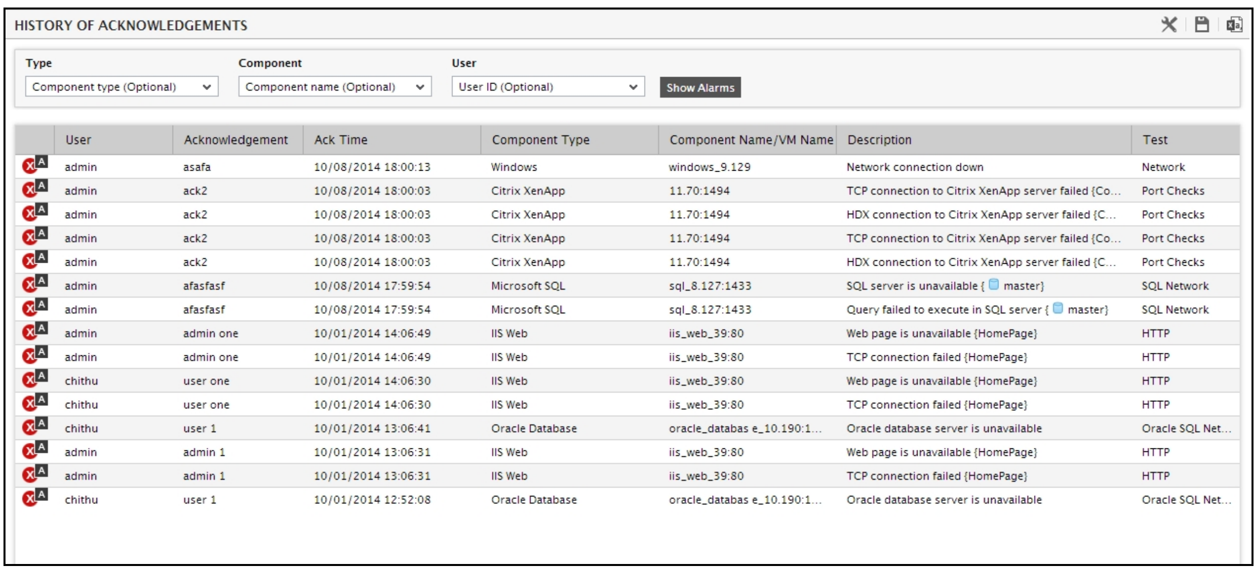

User	Acknowledgement	Ack Time	Component Type	Component Name/VM Name	Description	Test
admin	asafa	10/08/2014 18:00:13	Windows	windows_9.129	Network connection down	Network
admin	ack2	10/08/2014 18:00:03	Citrix XenApp	11.70:1494	TCP connection to Citrix XenApp server failed {Co...	Port Checks
admin	ack2	10/08/2014 18:00:03	Citrix XenApp	11.70:1494	HDX connection to Citrix XenApp server failed {C...	Port Checks
admin	ack2	10/08/2014 18:00:03	Citrix XenApp	11.70:1494	TCP connection to Citrix XenApp server failed {Co...	Port Checks
admin	ack2	10/08/2014 18:00:03	Citrix XenApp	11.70:1494	HDX connection to Citrix XenApp server failed {C...	Port Checks
admin	afasfasf	10/08/2014 17:59:54	Microsoft SQL	sql_8.127:1433	SQL server is unavailable { master}	SQL Network
admin	afasfasf	10/08/2014 17:59:54	Microsoft SQL	sql_8.127:1433	Query failed to execute in SQL server { master}	SQL Network
admin	admin one	10/01/2014 14:06:49	IIS Web	iis_web_39:80	Web page is unavailable {HomePage}	HTTP
admin	admin one	10/01/2014 14:06:49	IIS Web	iis_web_39:80	TCP connection failed {HomePage}	HTTP
chithu	user one	10/01/2014 14:06:30	IIS Web	iis_web_39:80	Web page is unavailable {HomePage}	HTTP
chithu	user one	10/01/2014 14:06:30	IIS Web	iis_web_39:80	TCP connection failed {HomePage}	HTTP
chithu	user 1	10/01/2014 13:06:41	Oracle Database	oracle_databas e_10.190:1...	Oracle database server is unavailable	Oracle SQL Net...
admin	admin 1	10/01/2014 13:06:31	IIS Web	iis_web_39:80	Web page is unavailable {HomePage}	HTTP
admin	admin 1	10/01/2014 13:06:31	IIS Web	iis_web_39:80	TCP connection failed {HomePage}	HTTP
chithu	user 1	10/01/2014 12:52:08	Oracle Database	oracle_databas e_10.190:1...	Oracle database server is unavailable	Oracle SQL Net...

Figure 5.41: Viewing the history of alarm acknowledgements in the ACKNOWLEDGEMENT HISTORY page

To access the **ACKNOWLEDGEMENT HISTORY** page,. When Figure 5.41 appears, you can use any of the filter options discussed below to quickly retrieve the acknowledgement information that you are looking for.

- To view all alarm acknowledgements related to a particular component type, pick a **Type** from the list box in Figure 5.41.
- To view all the acknowledgements associated with a particular component, select an option from the **Component** list in Figure 5.41.
- If you want to view all the acknowledgements submitted by a particular user, then pick an option from the **User** list.
- By default, the **ACKNOWLEDGEMENT HISTORY** page provides the details of all the acknowledgements that were submitted during the last 24 hours. Accordingly, the **Timeline** is set to *24 hours* by default. If need be, you can change this **Timeline**, so that the page displays only those acknowledgements that were submitted during the specified timeline.
- By default, the **ACKNOWLEDGEMENT HISTORY** page displays all acknowledgements that were submitted for all types of alarms, regardless of the problem status. Accordingly, the **Alarm Type** is set to **All** by default. However, if you want to view only those acknowledgements that pertain to the open alarms, pick the **Open** option from the **Alarm Type** list. Similarly, you can pick the **Closed** option from this list to view the acknowledgement information that related to closed issues only.
- You can even indicate how many acknowledgement records need to be provided per page by picking an option from the **Records per page** list.
- Finally, click on the **Show Details** button.
- The resulting report then provides the acknowledgement history that fulfills the specified conditions (seeFigure 5.41).

While you can save the report as a CSV file by clicking on the icon in Figure 5.41, you can save the same as a PDF document by clicking on the icon.

5.4 Deletion History

Typically, in large, multi-user environments, multiple users may be granted the privilege to monitor a single component. In such environments, any of these users can delete an alarm raised on that component without the knowledge of the others, thereby causing confusion. To avoid this confusion, eG Enterprise provides users with the ability to track the deleted alarms.

The **EVENT HISTORY** page for instance has been embedded with the intelligence to indicate whether a past alarm was deleted or not. Sometimes, you might want to view the details of all alarms that were deleted by a particular user. Similarly, you might want to view only the details of those alarms that were deleted during the last 24 hours. Since the **EVENT HISTORY** page allows you to search based on general alarm information alone and not on deleted alarms, this page cannot be used for performing the search operations mentioned above.

To zoom into the details of specific deleted alarms therefore, eG Enterprise offers a dedicated **DELETION HISTORY** page (see Figure 5.42). This page provides a variety of filter options using which you can quickly access the alarm deletion details that you may require.

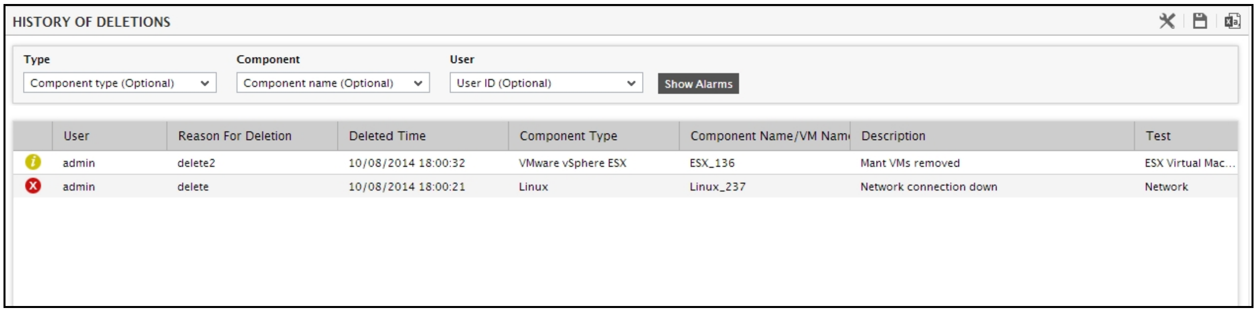

Figure 5.42: The DELETION HISTORY page

To access this page, Click on the icon available in the ***Monitor*** tab. Then, select the ***History of Deletion*** option in the ***Alarms*** tile. When Figure 5.42 appears, you can use any of the filter options discussed below to quickly retrieve the deleted alarm information that you are looking for.

- To view the details of all deleted alarms related to a particular component type, pick a **Type** from the list box in Figure 5.42.
- To view the details of those deleted alarms that pertained to a particular component, select an option from the **Component** list in Figure 5.42.
- If you want to view all the alarms deleted by a particular user, then pick an option from the **User** list.
- By default, the **HISTORY OF DELETIONS** page provides the details of all the alarms that were deleted during the last 24 hours. Accordingly, the **Timeline** is set to *24 hours* by default. If need be, you can change this **Timeline**, so that the page displays only those alarms that were deleted during the specified timeline.
- You can even indicate how many deleted alarm records need to be displayed per page by picking an option from the **Records per page** list.
- Finally, click on the **Show Details** button.
- The resulting report then provides the alarm deletion history that fulfills the specified conditions (see Figure 5.42).

While you can save the report as a CSV file by clicking on the icon in Figure 5.42, you can save the same as a PDF document by clicking on the icon.

6

Monitoring Components

In this chapter we will discuss the following topics:

- Monitoring Server Components
- Monitoring Virtual Components
- Monitoring Aggregate Components

6.1 Monitoring Components

To view the status of the components managed in your infrastructure, select the **Components** option from the **Hosts/Applications** menu in the eG monitor interface. By default, the page that appears lists the key performance metrics of a component that is managed in your infrastructure.

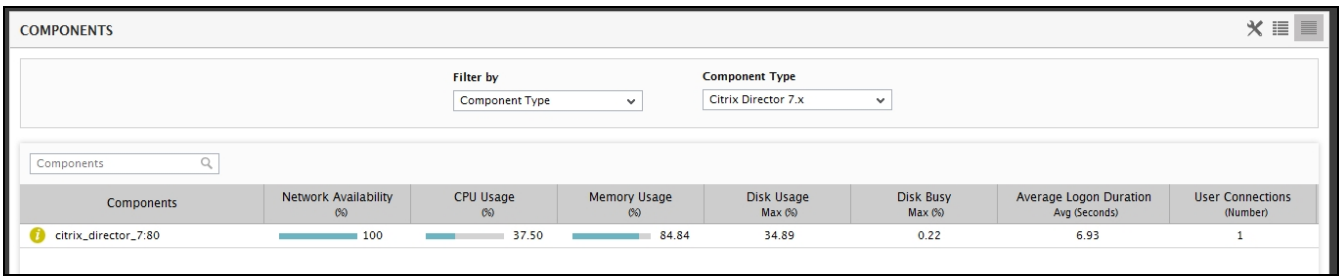

Figure 6.1: The COMPONENTS page displaying the key performance metrics

By default, the **Component Type** option is chosen from the **Filter by** list. If the zones/services/segments are created by the user who is accessing the eG monitor interface, then the components that are part of the zone/segment/service can be filtered using this list. The **Component Type** list box by default, lists all the components managed in the target infrastructure. When the **COMPONENTS** page is accessed for the very first time, the key performance metrics of the component type that comes first in the alphabetical order will be displayed against each component. In Figure 1, the *Citrix Director 7.x* is the first component in alphabetical order displayed in the **Component Type** list box.

By default, eG decides certain metrics to be the key performance metrics and displays the same in the **COMPONENTS** page. If you wish to define the metrics of your choice as the key performance metrics for each component, then you can do so using the icon. Figure 2 will then appear.

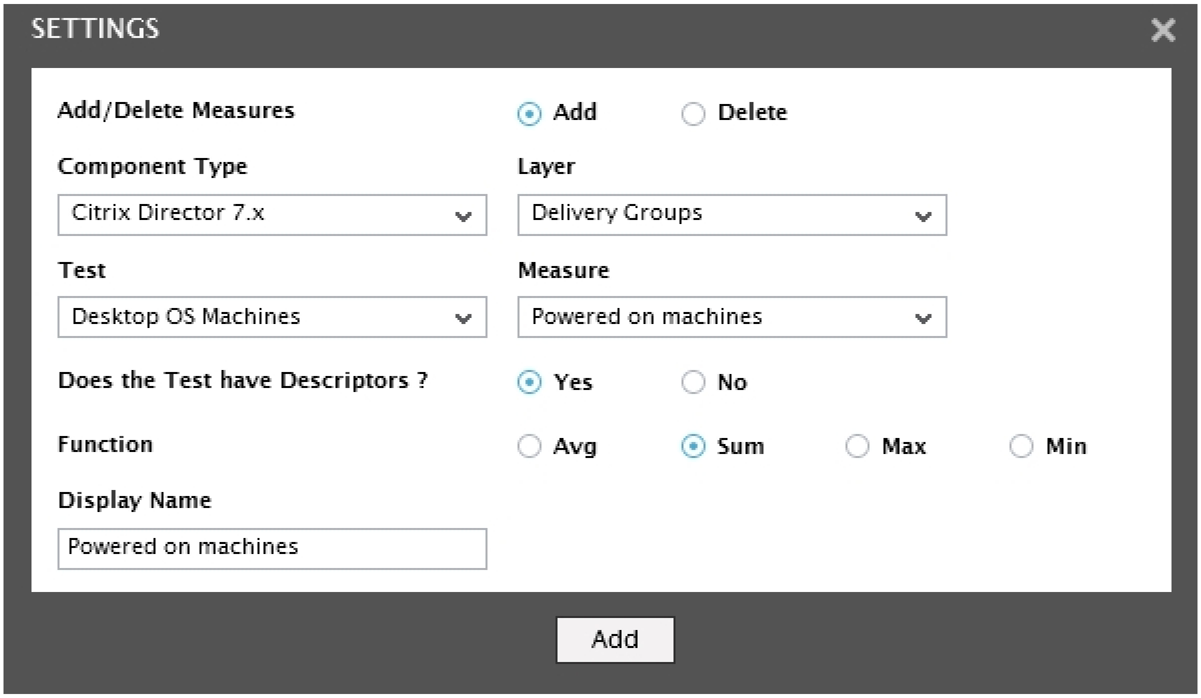

Figure 6.2: The Settings pop up window that appears to add the measure of your choice

Using Figure 2, you can add/delete a measure from being displayed as a key performance metric in the COMPONENTS page. To add/delete a measure, do the following:

- First specify whether you need to add or delete the measures by selecting the **Add** or **Delete** option from the **Add/Delete Measures** flag. Click the **Add** option for adding the measures.
- Then, select the component type from the **Component Type** drop-down list. This will list all the layers associated with the chosen component type in the **Layer** list.
- Now, selecting a particular layer from the **Layer** list will populate the tests associated with the chosen layer in the **Test** list.
- Then, select the test that reports the measure of your choice from the **Test** list.
- Now, all the measures pertaining to the chosen test will be listed in the **Measure** list. Select the measure of your choice from the **Measure** list, and provide a display name for the measure in the **Display Name** text box.
- Next, indicate whether the chosen test is descriptor based or not using the **Does the Test have Descriptors?** flag. If the test has descriptors, then set the **Does the Test have Descriptors?** flag to **Yes**. This will invoke the **Function** flag using which you can aggregate the measure values across all descriptors using a specific aggregate function. By default, the *Avg* flag will be selected against the **Function** flag indicating that the measure value will be displayed as the average of all the descriptors of the chosen test.
- Finally click **Add** button to add the measure. Doing so ensures that the measure will be displayed in the COMPONENTS page for all the components of the chosen Component Type.

- Similarly, you can remove a measure from being displayed from the COMPONENTS page. For this purpose, set the **Add/Delete Measures flag** to **Delete**.(see Figure 3).

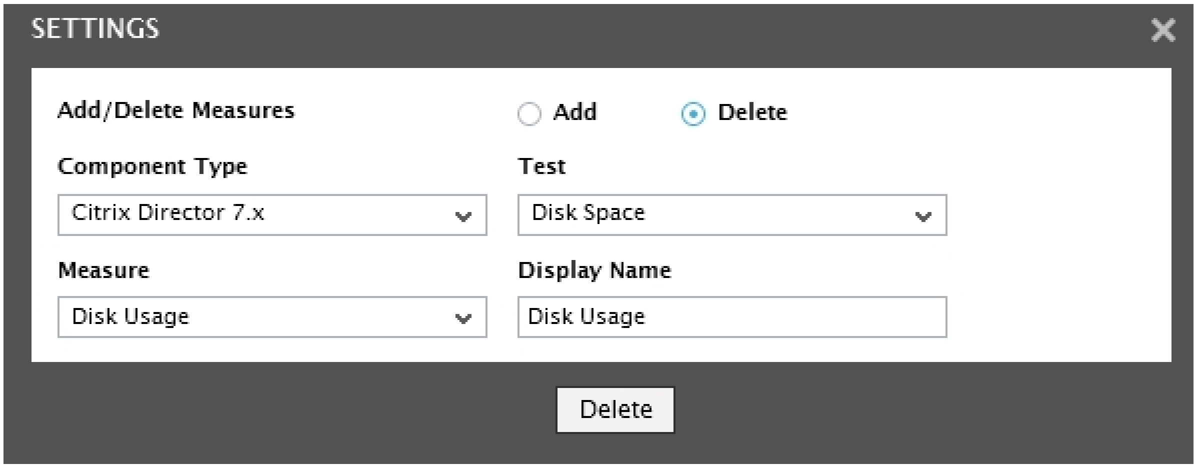

Figure 6.3: Deleting a measure from being displayed in the COMPONENTS page

- Select the component type, test from the **Component Type** and the **Test** lists. Select the measure that is to be deleted from the **Measure** list and click the **Delete** button.
- Clicking on a component will lead the users to a **Layers** tab page that displays the layer model, tests, and measurements pertaining to that component (see Figure 6.4).

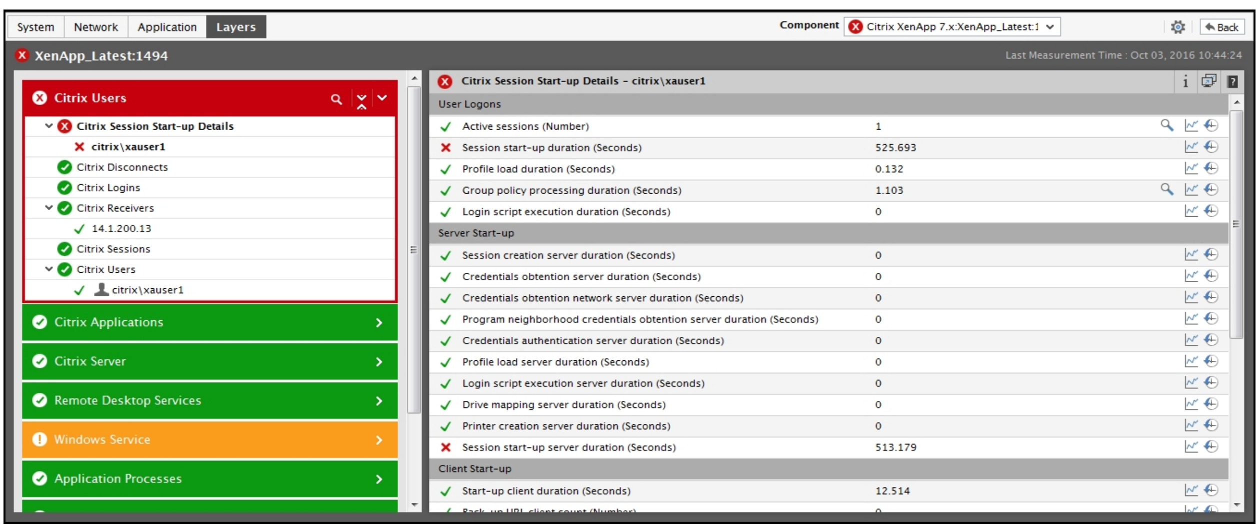

Figure 6.4: The layer model, tests, and measurements pertaining to a Citrix XenApp component

Note:

By default, when you click on a problem component in the eG monitoring console, you will be lead to this page, where you can directly view the problem layer, problem test, and problem measure (as depicted by Figure 6.4). In other words, you will not have to click on a problem layer to view the problem test, and then click on the problem test to view the problem measure. This default setting can be overridden so that, clicking on a problem component reveals only the layer model of that component, instead of revealing the problem layer, problem test, and problem measure. For this purpose, do the following:

  - Login to the eG administrative interface.
  - Click on the icon available in the **Admin** tab. Then, select the **Monitor** option in the **Settings** tile.
  - Click on the **Other Settings** sub-node under the **General** node in the tree-structure of the **MONITOR SETTINGS** page that appears.
  - In the right panel, scroll down to view the **When a problem component is clicked** drop-down list. By default, the **Go to problem measure** option is chosen from this list. This indicates that by default, clicking on a problem component will directly lead you to the problem layer, test, and measure. To go to the layer model of that component instead, just pick the **Go to layer model** option from this list. Finally, click the **Update** button.

  Once this is done, then, everytime you click on a problem component, only the layer model of that component will appear. You will have to manually drill down from the problem layer to view the problem test and measure.

  However, when you click on a component in the Normal state, then, the layer model page will only display the layer model of the component. In this case, you will have to click on a layer to view its tests and then click on a test to view its measures. On the other hand, if a component in the Unknown state is clicked on, then, the layer model page will display the layer and the test in the Unknown state, but will not display the measures.

- Though the **Layers** tab page opens by default, you will find a **System** and **Network** tab page (see Figure 6.4) also available for any component you click from the **COMPONENTS** list. For some components, an additional **Application** tab page will also be available.
- Each of these tab pages are discussed in detail in the sub-sections that follow.

Alternately, if you wish to view a list of all the components managed in your infrastructure, you can do so by clicking the icon. By default, the page that appears displays all the managed components in the environment, regardless of whether or not they are part of a segment, service or zone. Accordingly, the **Show All Components** check box is selected, as depicted by Figure 6.5. To view only the independent components, deselect the **Show All Components** check box (see Figure 6.5).

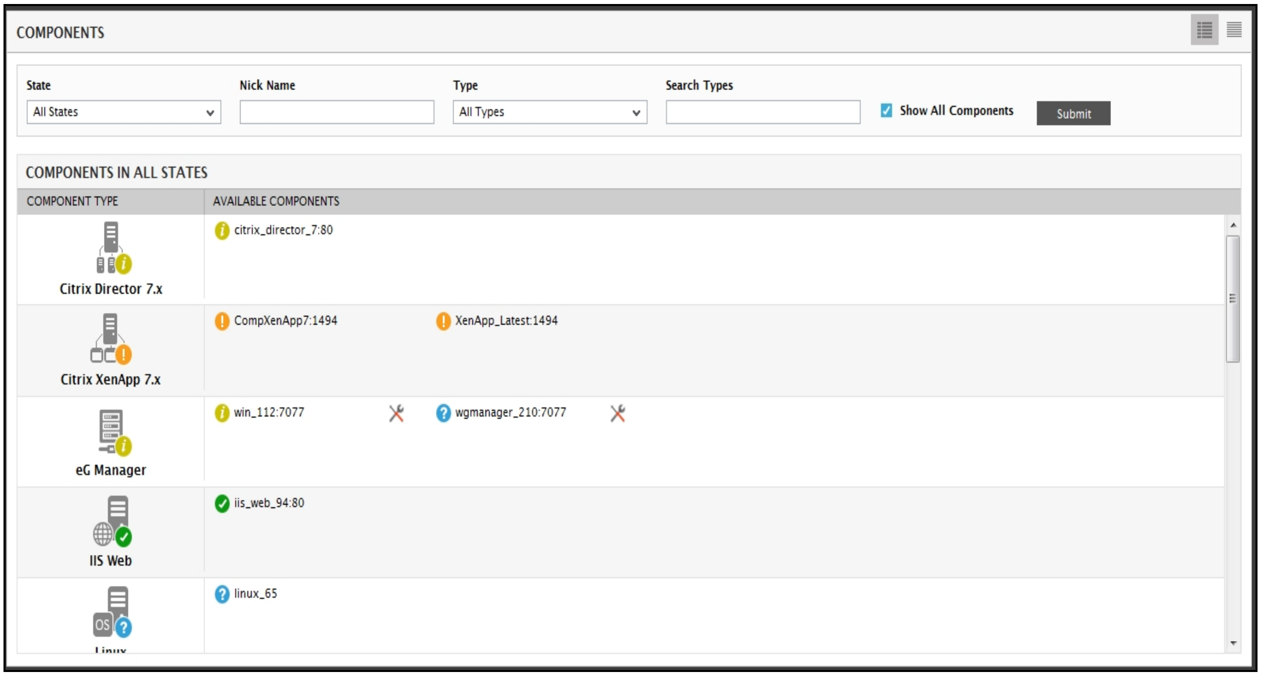

Figure 6.5: All managed components available in the infrastructure

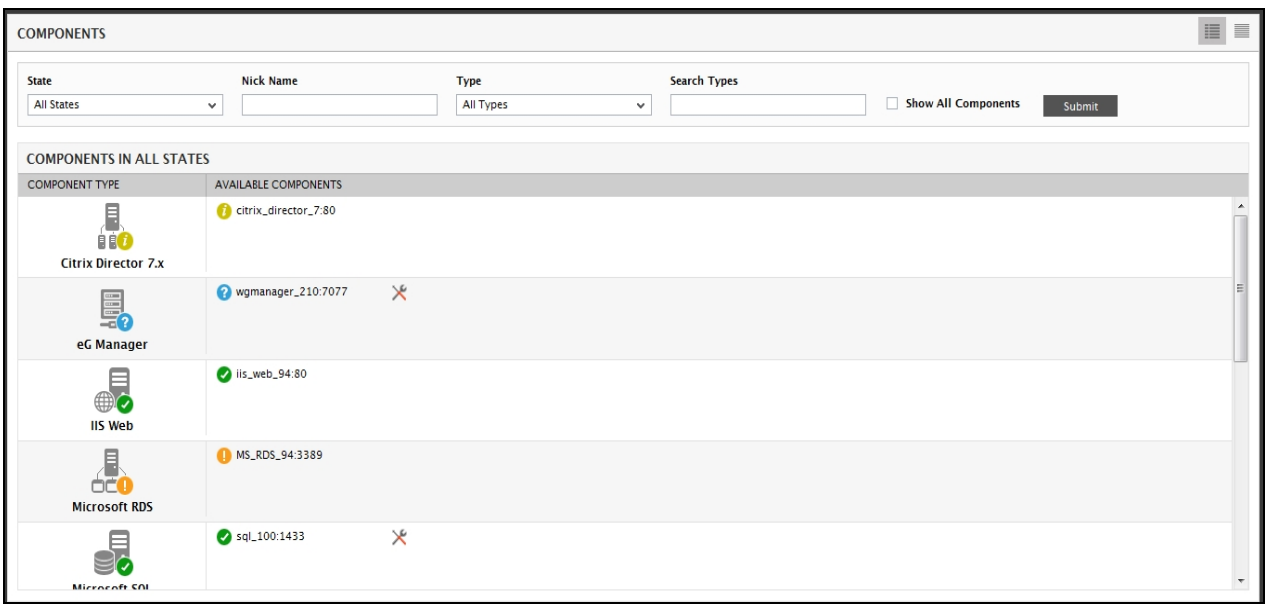

Figure 6.6: Independent components in the environment

By default, the page represents all managed component-types as icons. Corresponding to every component-type icon, the monitored components of that type and their current state will be displayed.

Note:

In target environments where a large number of components are managed, the component-type icons could cramp the component display. To avoid this, you can set the **Show icon for component-type** flag in the **OTHER DISPLAY SETTINGS** section of the **MONITOR SETTINGS** page (Click on the icon available in the **Admin** tab, select the **Monitor** option from the **Settings** tile) to **No**. Doing so removes the icons from the **COMPONENT LIST** page, and replaces them with the names of the corresponding component-types.

The color-coding associated with a component indicates its state: Red indicates critical problems, Green is used to indicate good health, while Blue is used to indicate an unknown state. Also, orange indicates major problems, and amber indicates minor ones.

If one/more maintenance policies are associated with a component, then, during the configured maintenance period, the **COMPONENT LIST** page will display a icon adjacent to that component (see Figure 6.6). This way, eG Enterprise pictorially conveys to administrators that all alerts pertaining to a component will be suppressed as long as the icon appears alongside that component name.

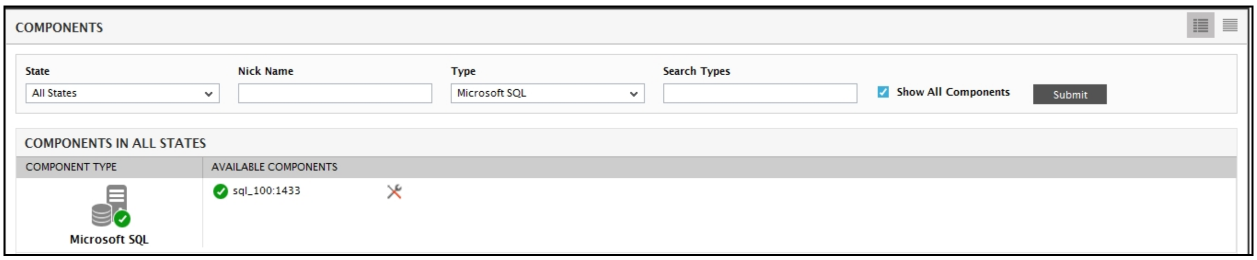

Figure 6.7: A component under maintenance

Likewise, if a component is currently in an abnormal state, and a configuration change has occurred on that component during the same time as its performance abnormality, then, a icon will appear alongside the component name.

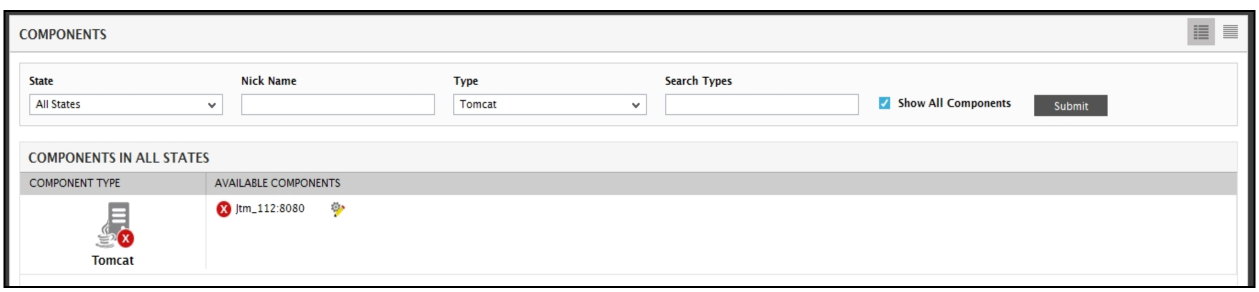

Figure 6.8: A component in an abnormal state owing to a configuration change

Clicking on the icon will invoke a small pop up window which reveals the recent configuration changes. Clicking on the **More Details** in the pop up window will take you to a page revealing the configuration changes that occurred on that component at around the same time at which a performance issue was experienced by the component; using this information, administrators can determine whether/not the performance anomaly can be attributed to the change in configuration.

This page also provides a wide variety of search options using which you can selectively view the status of the component.

- To view the components that are in a particular state, select the state of your choice from the **State** list and click the **Submit** button. All components in that state will then be displayed (see Figure 6.7).

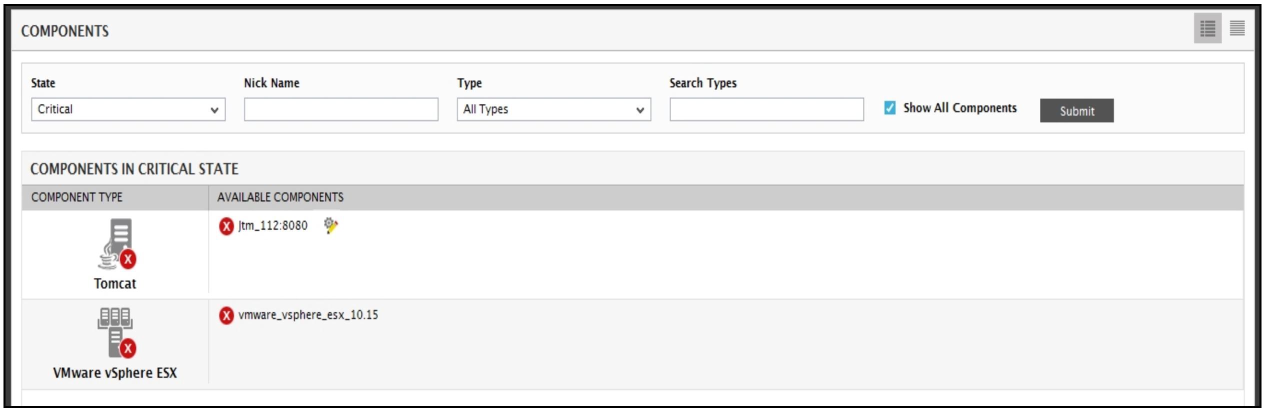

Figure 6.9: Components in a chosen state

- To view the state of components that use a specific nick name, specify the whole/part of the **Nick name** to search for. Multiple nick names can be provided, separated by a white space - for example, *sun_ora web_serv*. Then, click the **Submit** button. All components with nick names that embed the specified string will then appear (see Figure 6.9).

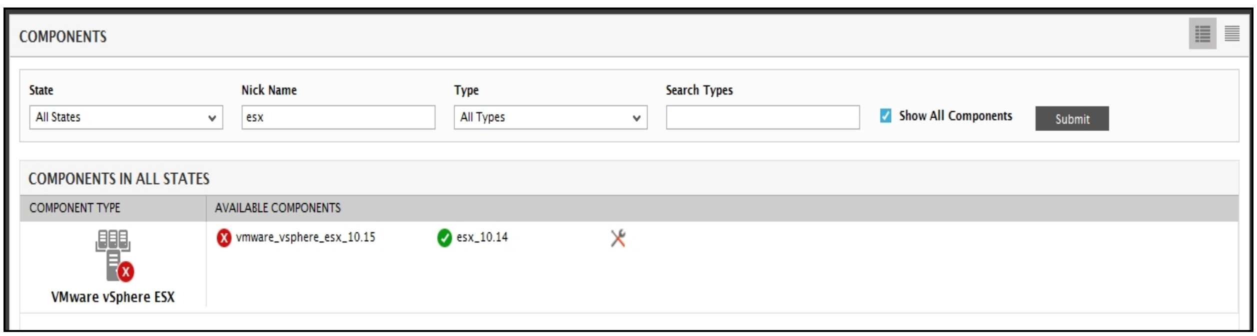

Figure 6.10: Components with a specified nick name

- If you want to view the state of components of a particular type, select the **Type** to search for. You can also enter the whole/part of the component-type to search for in the **Search Types** list box. This will enable you to view the status of those components that are of the types which embed the specified search string. Here again, you can search for multiple component types, by separating every type specification by a white space - for example, *Citrix XenApp 7.x* (see Figure 6.10).

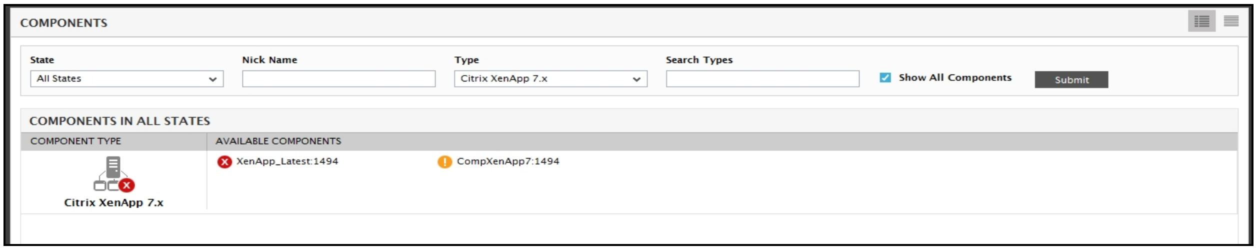

Figure 6.11: Components of a particular type

Note:

If you wish to view the list of components in the monitored environment in the **COMPONENTS** page, then you can do so by setting the **ListView** flag under the **[MONITOR_DASHBOARD]** section of the **<eG_INSTALL_DIR>\manager\config\eg_ui.ini** file to **Yes**.

6.1.1 The Layers Tab Page

Each component managed in eG Enterprise – be it a server, system, or a network device – is represented as a set of hierarchical layers. Each of these models determines what metrics to collect for each application or network device, how to collect these metrics, how to analyze the collected metrics, and what to deduce from the analysis. Since each infrastructure component performs different functions, the layer model for the different applications and network devices are different.

Typically, to view the layer model of a component, you need to click on that component in the **COMPONENTS** page. This will lead you to the **Layers** tab page, where the layer model of that component will be displayed. In Figure 6.12 for instance, you will find that the **Layers** tab page displays the layer model of a Citrix XenApp component.

Each layer of a layer model (also called as monitoring model) is mapped to tests that measure the health of that layer. Clicking on the **>** button alongside a layer name will reveal the tests associated with that layer. For instance, in Figure 6.12, clicking on the **>** button against the layer name **Windows Service**, will reveal the tests mapped to that layer. Clicking on a test will list the measures reported by that test in the panel to the right.

The state of a layer is determined by the state of the tests that are associated with it and the state of the measures that the tests report. For example, the **Windows Service** layer of Figure 6.12 is in the **Minor** state presently.

From the tests list for the **Windows Service** layer, it is evident that a **Minor** problem captured by the **Security Log** test of that layer is responsible for the abnormal state of the layer. Clicking on the **Security Log** test in Figure 6.12 will reveal the measures reported by that test. The measures clearly indicate that audit failures are the reason for the **Minor** problem with the **Windows Service** layer. Since layer state impacts the component state, the Citrix XenApp component too is in a **Minor** state currently.

But, what happens if different layers of a model are in different states? In this case, the component assumes the state of that layer with the most critical problem. For instance, take the case of the Citrix XenApp component in Figure 6.12. As you can see, the **Citrix Applications** layer of the component is in a **Major** state and the **Citrix Users** layer of that component is in the **Minor** state currently. In other words, a **Major** problem has been captured in the **Citrix Applications** layer, and a **Minor** problem has been detected in the **Citrix Users** layer. In this case, since the more critical of the two problems is the **Major** problem, the state of the Citrix XenApp component also becomes **Major**.

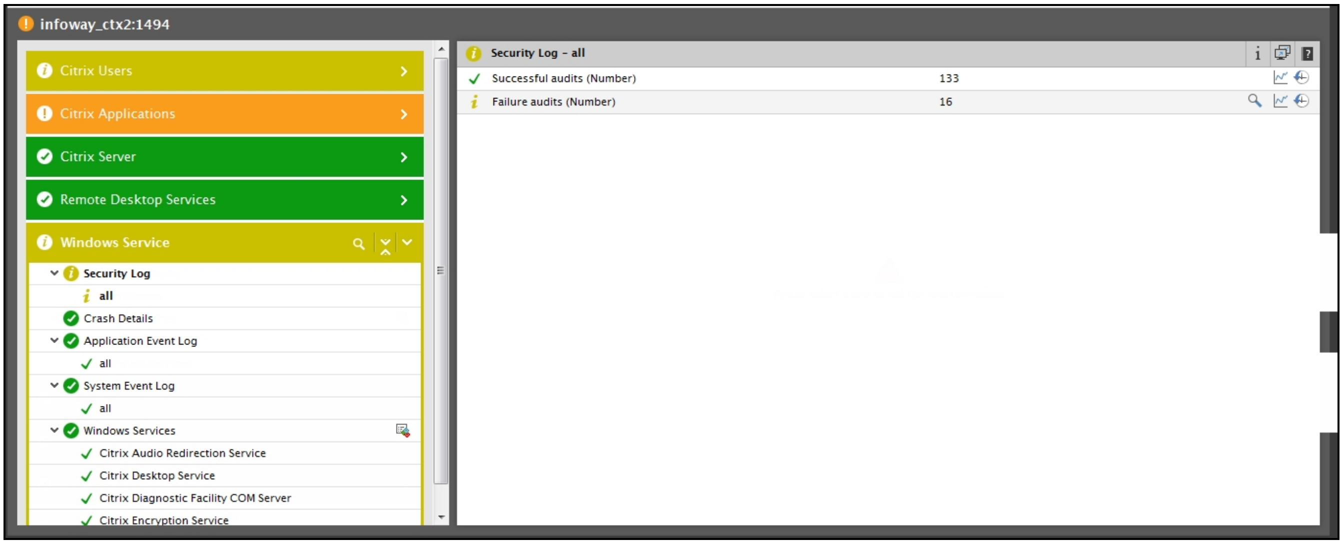

Figure 6.12: State of a component if the layers are in different states

Each pre-defined layer model has been built bottom-up, with the system and network tiers at the bottom and the application tiers at the top. Typically, a problem at the host or network-level – eg., a CPU/memory contention on the host, loss of network connection to the host, etc. – can impact the health of the application layers at the top . eG understands these layer-layer dependencies; this is why, whenever a component encounters more than one issue at the same time, eG automatically correlates issues across the layers of a layer model using a patented top-down correlation algorithm, and intelligently determines the state of the layers. For instance, take the case of the Citrix XenApp component, the layer model of which is depicted by Figure 6.12. As you can see, the **Citrix Applications** layer is in a **Major** state, while the **Citrix Users** layer is in a **Minor** state. To know what is the **Major** problem that is impacting the health of the **Citrix Applications** layer, click the **>** button alongside the layer name in Figure 6.12. Figure 6.13 then appears, revealing that the **Citrix XA Applications** test has captured a **Major** problem with the **Java GUI** application. To know what the problem is, click on **Java GUI** in Figure 6.13. Doing so reveals that the **Java GUI** application is consuming over 30% of the CPU of the XenApp server.

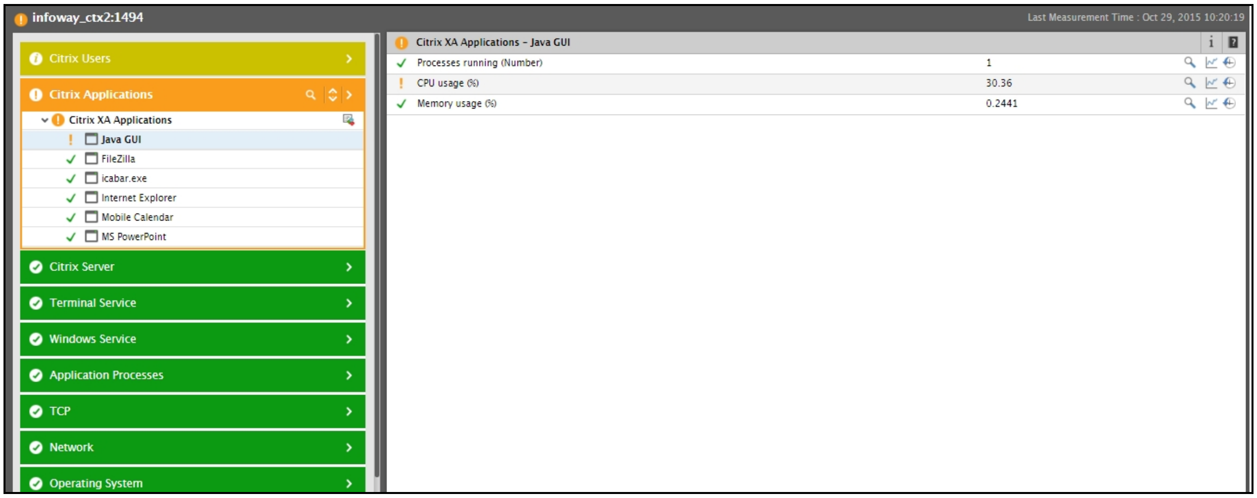

Figure 6.13: Measures reported by the Citrix XA Applications test

To know when this problem occurred, you can instantly invoke the eG alarm that corresponds to this issue from the **Layers** tab page itself. For that, just click on the problematic **CPU usage** measure. Figure 6.14 will appear providing the details of the alarm, along with its **START TIME**.

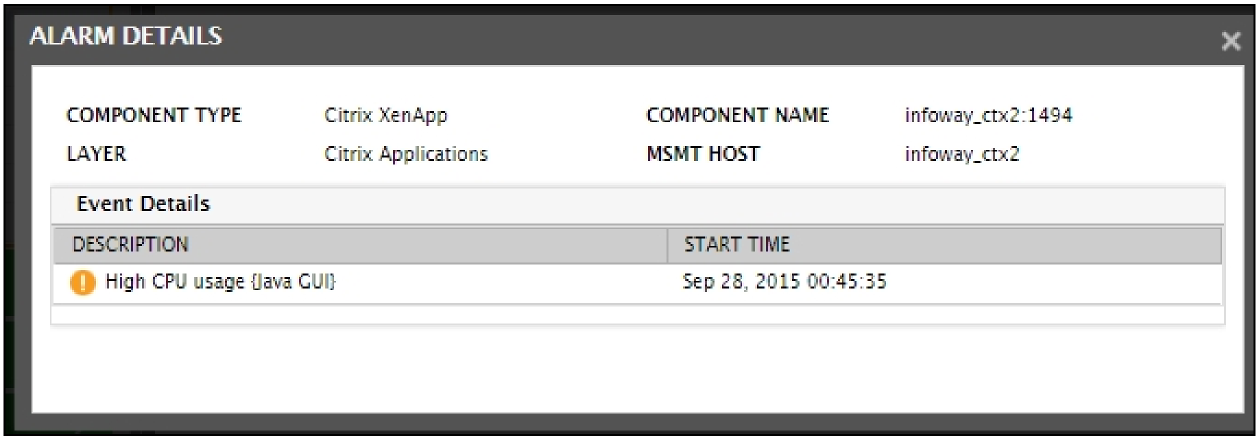

Figure 6.14: Alarm details viewed from the Layers tab page

Excessive CPU usage by an application on XenApp is bound to adversely impact the experience of the users of that application. Is that the reason why the **Citrix Users** layer, which reports user experience issues, is also in an abnormal state? Let's find out. For that, click on the **>** button alongside **Citrix Users** in Figure 6.13.

When Figure 6.15 appears, you can find that the **Citrix User Experience** test is reporting a **Minor** issue with experience of user **kevin**. Clicking on **kevin** reveals that kevin's experience with the XenApp server is very poor, causing a **Minor** alarm to be raised.

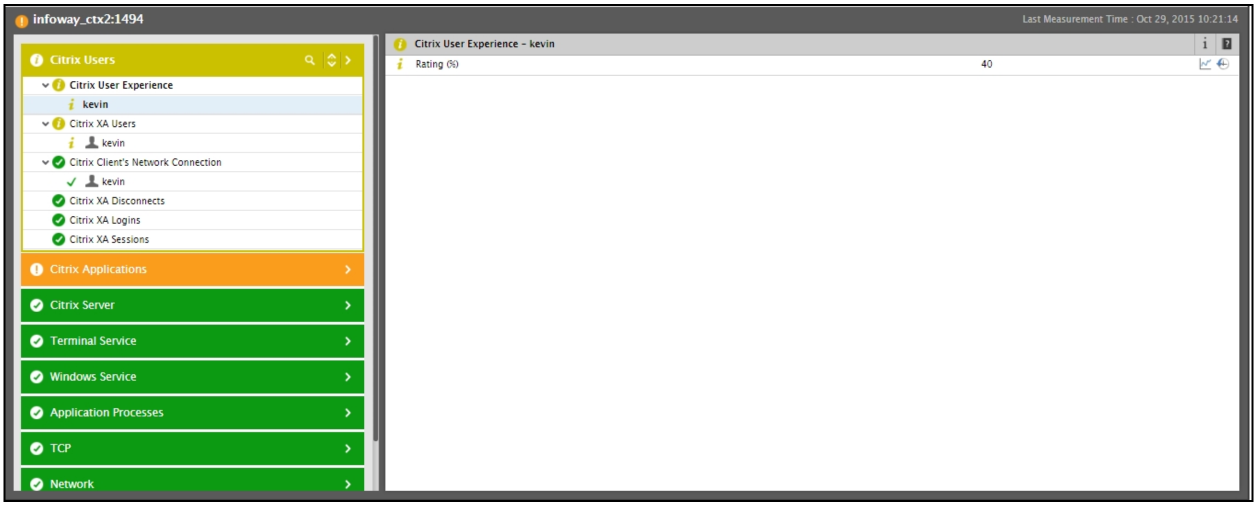

Figure 6.15: Measures reported by the Citrix User Experience test

To know the root-cause of this poor experience, click the **kevin** descriptor of the **Citrix XA Users** test. Figure 6.16 that appears indicates that one/more applications accessed by user **kevin** are consuming over 30% of CPU. This proves beyond doubt that the excessive CPU consumption by the **Java GUI** application has badly affected user **kevin's** experience with the XenApp server. eG's intelligent correlation engine has automatically correlated the **Java GUI** application's high CPU usage issue with user **kevin's** user experience issue and has accurately figured out that the CPU usage issue captured by the **Citrix Applications** layer is the root-cause of the user experience issue reported by the **Citrix Users** layer. Accordingly, eG has automatically assigned the **Major** state to the problem source – i.e., the **Citrix Applications** layer - and has bumped down the state of the problem effect – i.e., **Citrix Users** layer.

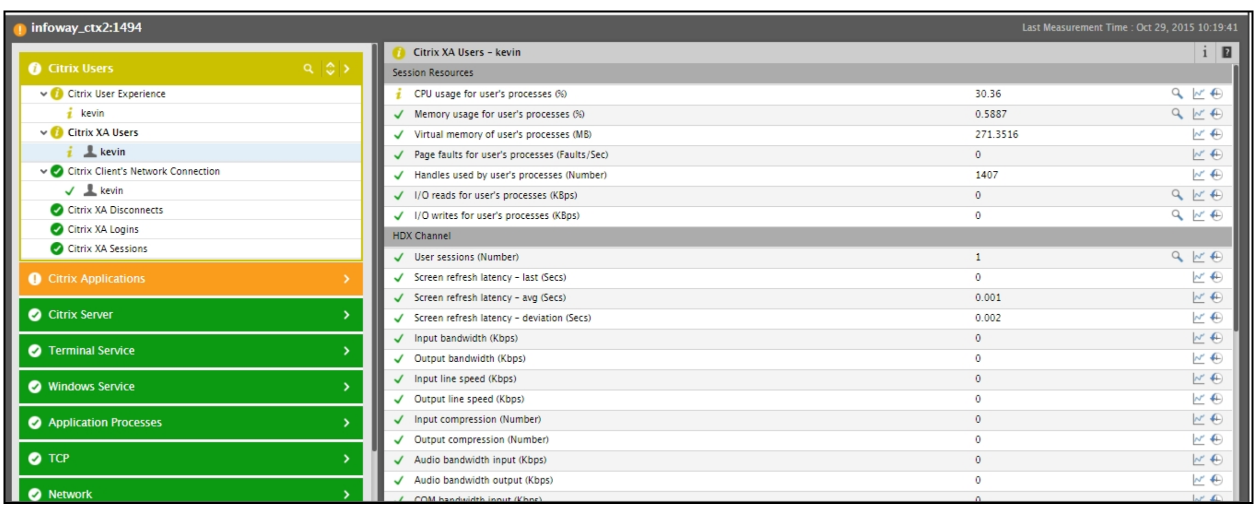

Figure 6.16: Measures reported by the Citrix XA Users test

By default, all the components managed in the infrastructure will be listed in the **Component** list box as shown in Figure 6.17. If too many components are managed in the target environment, then users may find it difficult to scroll down to the component of their choice from this list. A search capability is provided in this list box wherein users can key in a few alphabets from the component of their choice to filter out the results.

eG's layer model representation does more than just report what the problem is; it also enables administrators to figure out "how long the problem had persisted". For this, eG Enterprise allows you to launch a time-of-day graph of the problematic measure from the **Layers** tab page itself! To demonstrate this capability, let us take the example of Figure 6.17, which depicts the layer model of a Citrix XenApp server. As you can see, the **Security Log** test mapped to the **Windows Service** layer of this server has captured multiple security audit failures on the server.

Figure 6.17: The Security Log test reporting multiple security audit failures on a Citrix XenApp server

To view the historical graph of the **Failure audits** measures in Figure 6.17, click the (**GRAPH**) icon alongside that measure name. This will instantly provide you with a measure graph (see Figure 6.18) revealing variations in the count of audit failures during the last hour (by default).

Note:

By default, clicking on the **GRAPH** image against a measure reveals a graph that plots the values registered by that measure during the last hour. This default duration however is configurable. To override the setting, indicate a different duration in the **Timeline for graphs** setting in the **MONITOR SETTINGS** page, which appears when clicking on the icon available in the **Admin** tab and then, selecting the **Monitor** option in the **Settings** tile in the eG administrative interface.

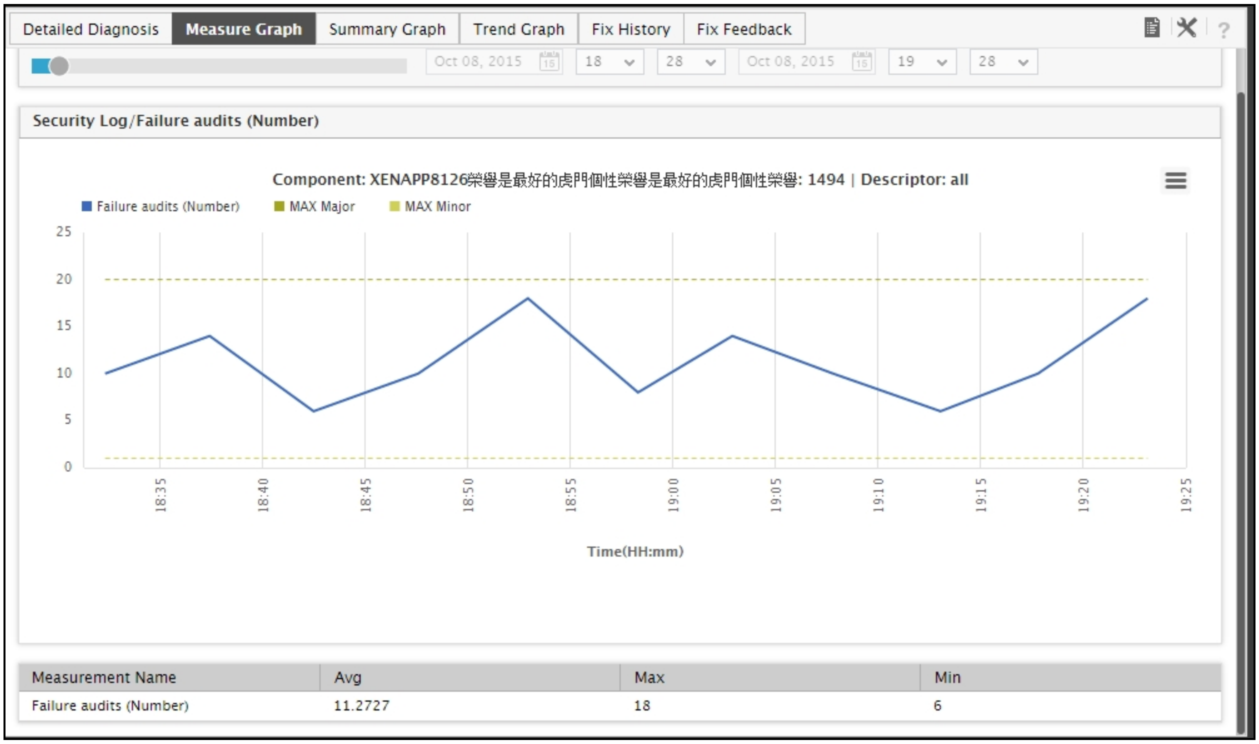

Figure 6.18: A measure graph indicating the number of audit failures that occurred in the last hour

The graph clearly reveals that the **Failure audits** measure violated its **MAX Minor** threshold throughout the last 1 hour. This could also mean that the **Minor** issue that occurred owing to this threshold violation has persisted for an hour or even more! The question now is how to resolve this issue? For that, you will first have to figure out why the issue occurred. This is where, the **Detailed Diagnosis** capability of eG Enterprise comes in handy! This capability, if enabled for a test, provides detailed metrics that will take you a step closer to identifying why a problem occurred. Take the case of the Citrix XenApp component of Figure 6.17 again. As already mentioned, audit failures have been detected on the XenApp component. To know when and why these failures occurred, you can use the **Detailed Diagnostics** provided by the **Security Log** test. To access these additional diagnostics, you need to click on the 'magnifying glass' ( ) icon alongside **Failure audits** in Figure 6.17. This will open Figure 6.19, where the complete description of the failure events will be revealed.

Note:

By default, clicking on the icon against a measure reveals the detailed diagnosis of that measure during the last hour. This default duration however is configurable. To override the setting, do the following:

- Edit the **eg_ui.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory)
- By default, the **DD_DISPLAY** parameter in the **[GRAPHS]** section of the file will display the value 1, indicating that the default detailed diagnosis duration is 1 hour. This value can be changed, if need be.
- After making the necessary changes, save the **eg_ui.ini** file.

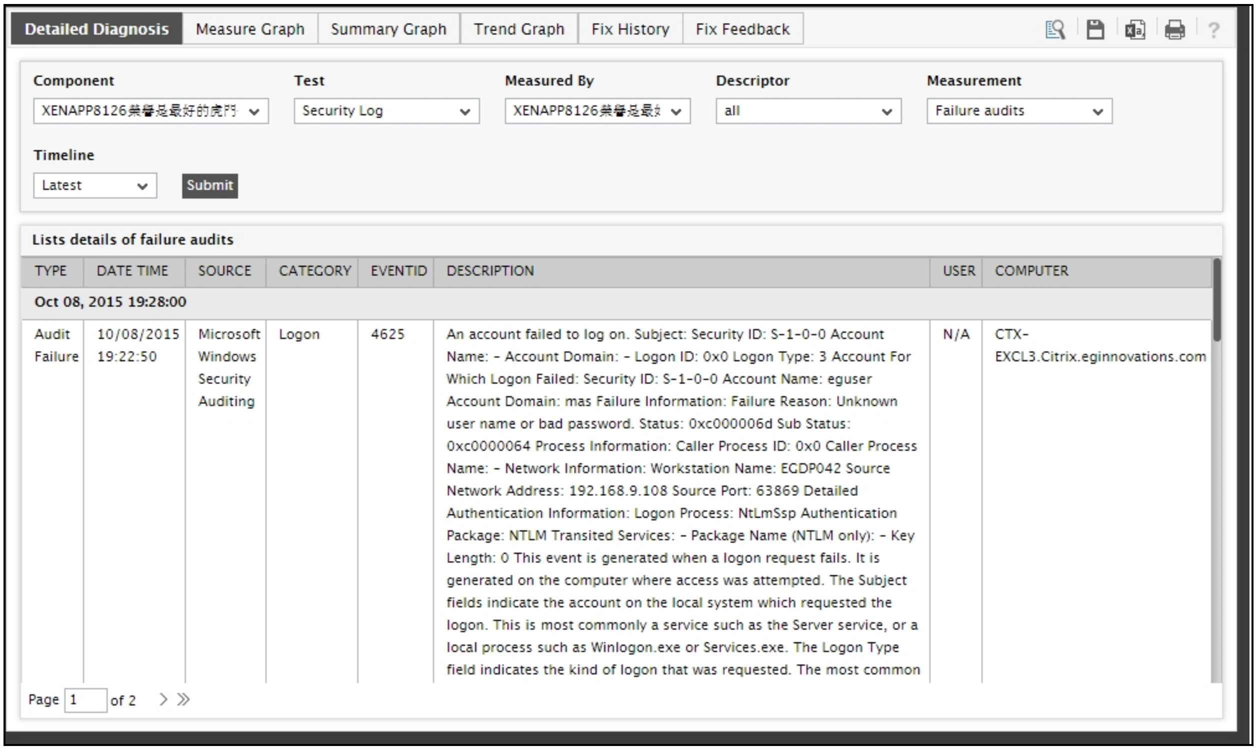

Figure 6.19: The detailed diagnosis of the *Failure audits* measure

From Figure 6.19, it is evident that one of the reasons why a security audit failed on the XenApp server was a login failure experienced by a user on that server. To resolve this issue, you need to figure out which user's login failed and why the failure occurred. The **Detailed Diagnosis** page (see Figure 6.19) provides this information as well! From the **USER** column of Figure 6.19, you can identify the user who experienced the login failure. Armed with the root-cause of the problem, you can quickly resolve the issue!

Sometimes, a configuration change effected in a component could be the root-cause of a performance issue with that component. Take the case of the Oracle database server, the layer model of which is depicted by Figure 6.20. As you can see, a **Critical** issue has been detected in the **Tablespaces** layer of this server – apparently, the **USERS** tablespace has run out of space!

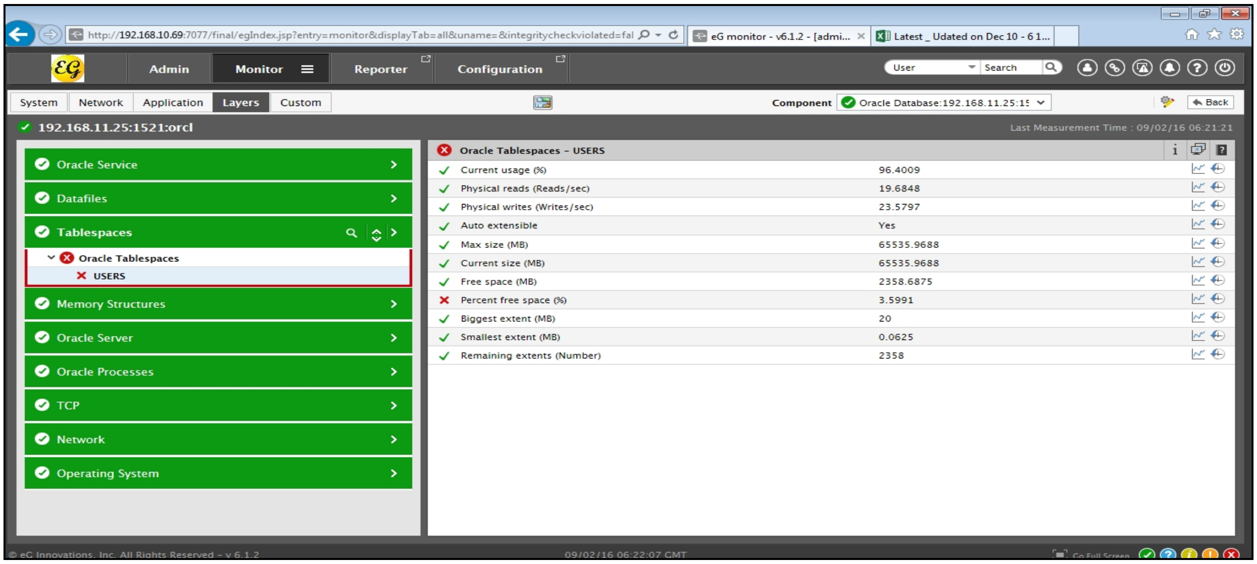

Figure 6.20: A Critical issue with the ORDERS tablespace of an Oracle database server

What could have caused this space drain? The presence of the icon at the right, top corner of the **Layers** tab page (see Figure 6.20) indicates that a configuration change has been made to the Oracle database server at around the same time as the tablespace issue. Could this change have contributed to the problem at hand? eG enables you to answer this question by allowing you to instantly determine what configuration change was made and confirm whether/not that change impacted free space availability on the tablespace. For this, just click on the icon. A pop up window (see Figure 6.21) will appear listing the recent configuration changes in a brief manner.

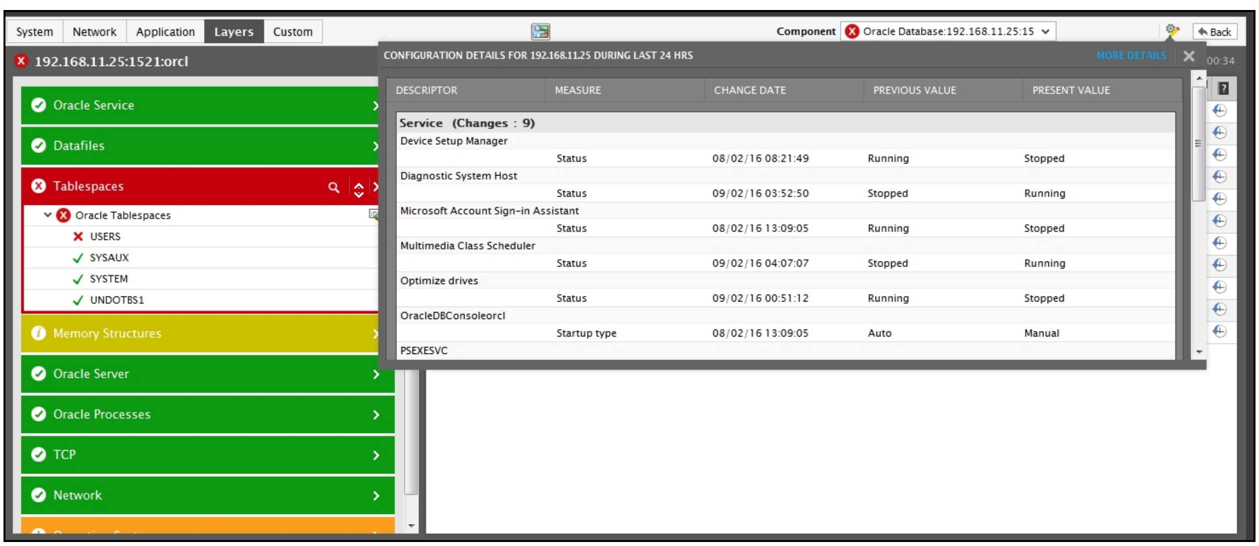

Figure 6.21: The pop up window revealing the recent configuration changes

Clicking the **More Details** in Figure 6.21 will lead you straight into the **Configuration Management** console, where you can view the configuration change made recently on the Oracle database server in question (see Figure 6.22).

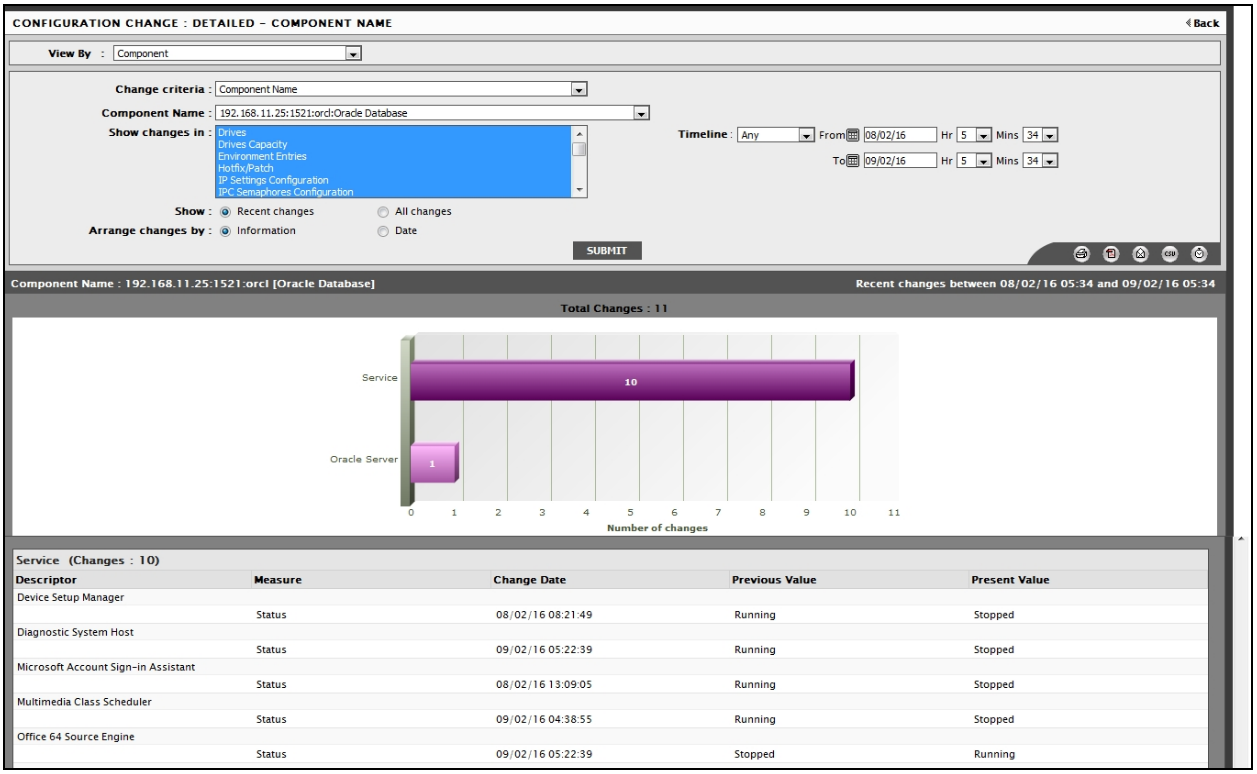

Figure 6.22: The configuration change that was made to the Oracle database server

Figure 6.22 clearly indicates that the Device Setup Manager and Optimize drives service was stopped. In addition, the size of the USERS tablespace has been reduced from 5000000 KB to 200000 KB. Because of this change, the USERS tablespace may not have had enough space to store data, thus resulting in a **Critical** issue in the **Tablespaces** layer. This way, using a quick sequence of mouse clicks, you have deduced what really caused a **Critical** performance issue to occur on a mission-critical server in your environment. We can thus conclude that using just the **Layers** tab page, you can promptly detect problems with a component, accurately pinpoint their root-cause, and in the process, quickly and easily resolve the problems.

But that's not all! Using eG Enterprise, you can also create and maintain a knowledge base of problem fixes. This way, the next time any of your colleagues encounter a similar issue, they can quickly query this knowledge base and easily figure out what remedial action should be taken. For instance, let us once again take the case of the problem – i.e., the security audit failure - illustrated by Figure 6.17. To create a record detailing how this problem is to be fixed, click the **Fix Feedback** tab page in Figure 6.23. Here, provide the **Problem Reason**, describe the **Problem Fix**, and click **Save** to save the details provided.

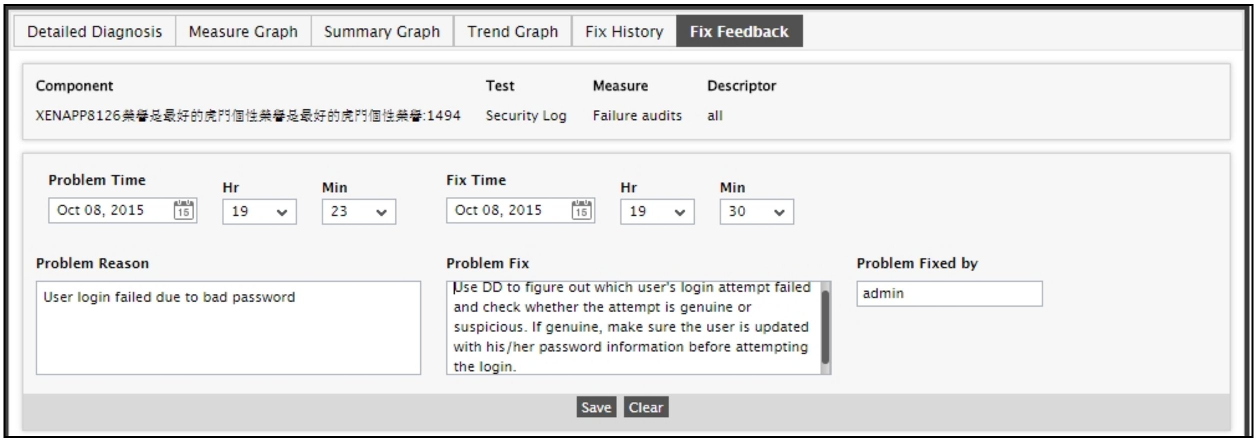

Figure 6.23: Recording the fix for a problem

Later, if the same problem reoccurs on the same server or any other server, all you need to do is click the (**HISTORY**) icon that corresponds to the **Failure audits** measure in Figure 6.17. Figure 6.24 that then appears will display the complete history of all the fixes that have been recorded for the same problem on all components on which the problem was encountered earlier.

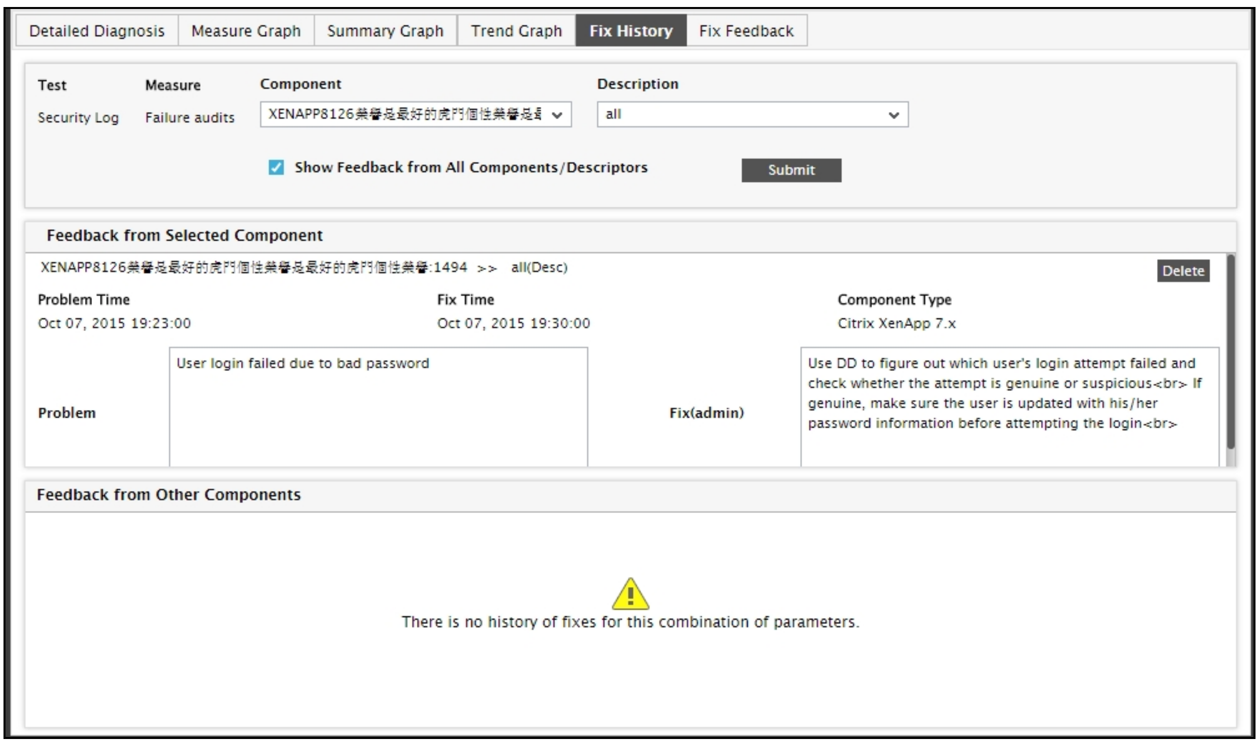

Figure 6.24: History of problem fixes

Note:

By default, monitor users have access to the fix history of problems pertaining to only those components that have been assigned to them. To grant the monitor users access to the fix history of all the monitored components in the environment, do the following:

- Edit the **[FIX_HISTORY]** section of the **eg_ui.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.

The **ALL_COMPONENTS** parameter in this section is set to **FALSE**, by default. This indicates that, by default, monitor users are provided access to only those monitored components that have been associated with them, and not all the components.

By setting this parameter to **TRUE** (as shown below), you can ensure that the fix history of all monitored components is available to the monitor users.

```
[FIX_HISTORY]
ALL_COMPONENTS=TRUE
```

This **Knowledge Management** feature of eG Enterprise greatly saves troubleshooting time and related costs.

As can be inferred from Figure 6.17, the **Security Log** test in our example reports metrics for only one descriptor, which is **all**. Some other tests however may report metrics for multiple descriptors. For example – the **Oracle Tablespaces** test reports usage metrics for each tablespace on an Oracle database server (see Figure 6.20). For such multi-descriptor tests, the **Layers** tab page provides you with the option to quickly compare performance across descriptors. For instance, in the case of the **Oracle Tablespaces** test in Figure 6.20, you can launch a quick comparison of the space usage in all the tablespaces listed as the descriptors of the test in the **Layers** tab page, and thus rapidly identify the the exact tablespace that is running out of space.

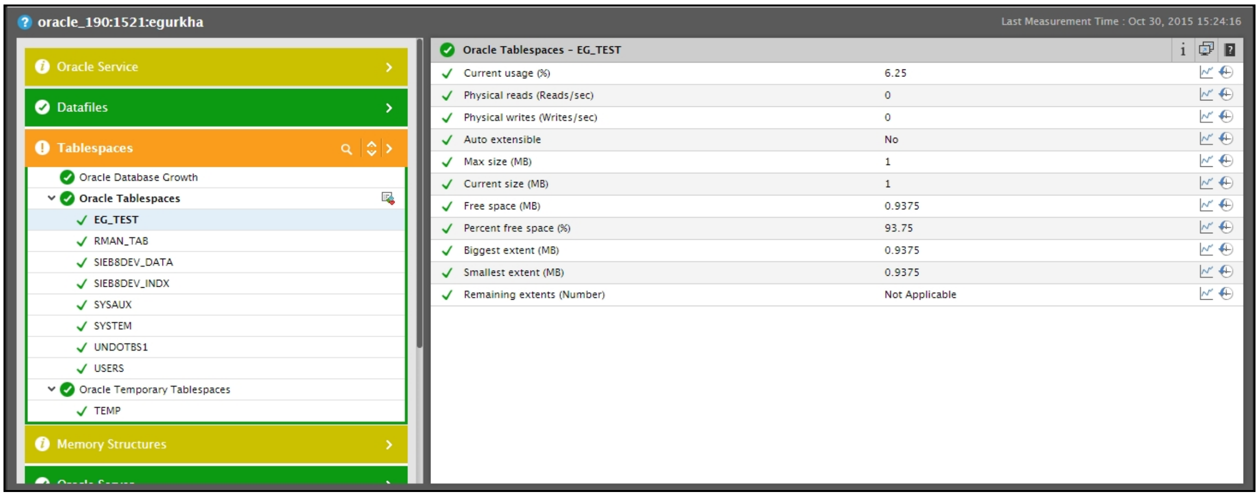

Figure 6.25: The Oracle Tablespaces test reporting metrics for every tablespace

To launch this comparison, click the **(Comparison)** icon alongside the test name in Figure 6.25. This will open Figure 6.26. From the **MEASURE NAMES TO DISPLAY** list, select the measures to be compared. Then, click the **Submit** button in Figure 6.26.

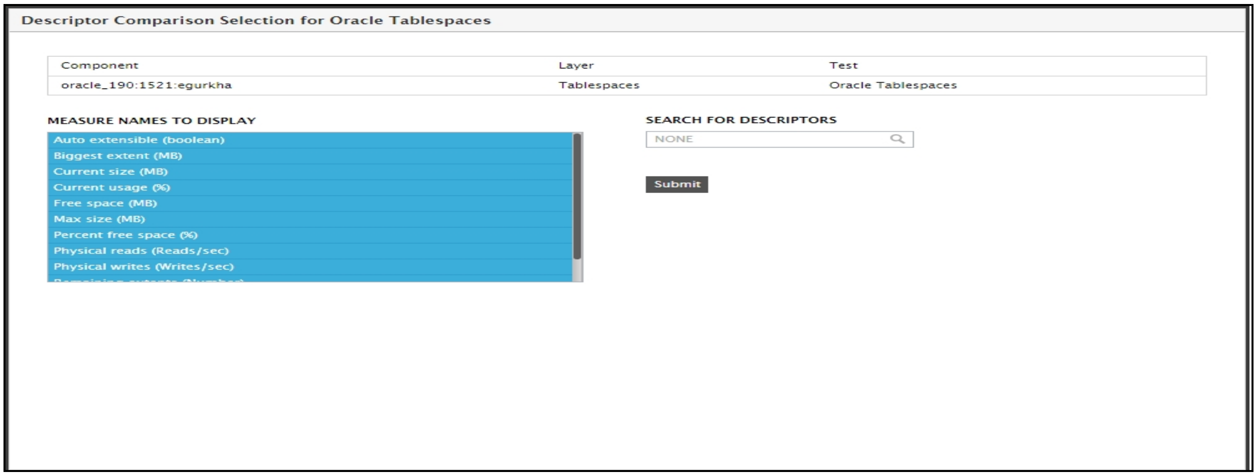

Figure 6.26: Selecting the measures to be compared

Doing so will invoke Figure 6.27, where the chosen metrics will be compared across all the descriptors – i.e., tablespaces, in the case of our example. Using this comparison, you can precisely pinpoint the tablespace that is being over-utilized.

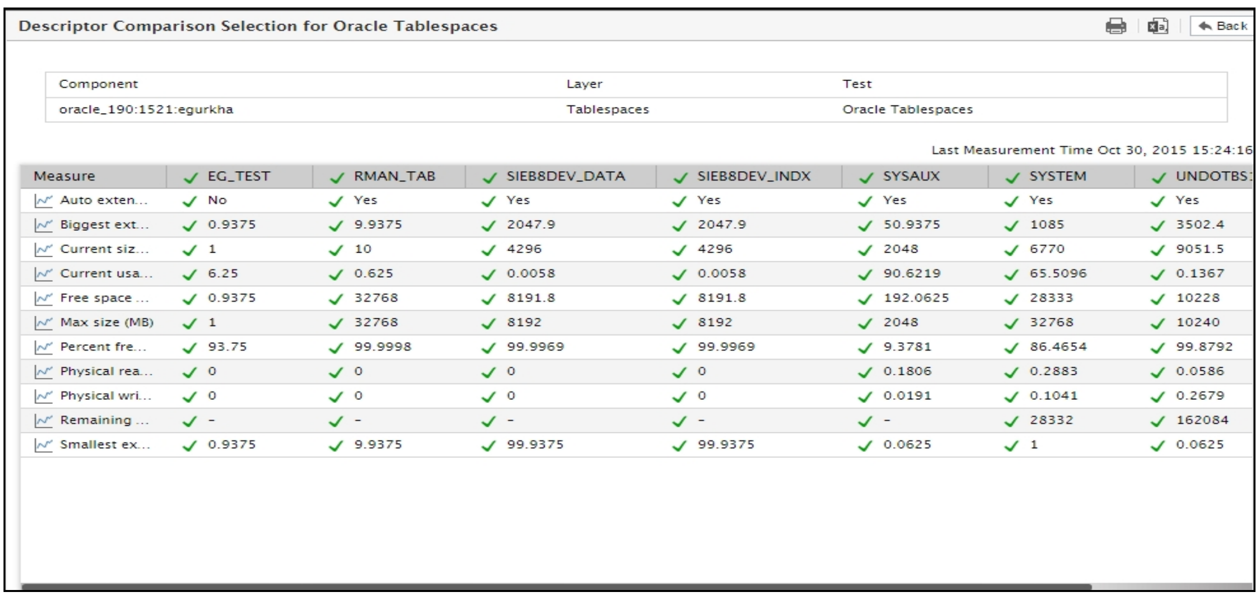

Measure	EG_TEST	RMAN_TAB	SIEB8DEV_DATA	SIEB8DEV_INDX	SYSAUX	SYSTEM	UNDOTBS:
Auto exten...	No	Yes	Yes	Yes	Yes	Yes	Yes
Biggest ext...	0.9375	9.9375	2047.9	2047.9	50.9375	1085	3502.4
Current siz...	1	10	4296	4296	2048	6770	9051.5
Current usa...	6.25	0.625	0.0058	0.0058	90.6219	65.5096	0.1367
Free space ...	0.9375	32768	8191.8	8191.8	192.0625	28333	10228
Max size (MB)	1	32768	8192	8192	2048	32768	10240
Percent fre...	93.75	99.9998	99.9969	99.9969	9.3781	86.4654	99.8792
Physical rea...	0	0	0	0	0.1806	0.2883	0.0586
Physical wri...	0	0	0	0	0.0191	0.1041	0.2679
Remaining ...	-	-	-	-	-	28332	162084
Smallest ex...	0.9375	9.9375	99.9375	99.9375	0.0625	1	0.0625

Figure 6.27: A comparison of the performance results of all descriptors of the Oracle Tablespaces test

If required, you can even compare the performance of specific descriptors alone. For instance, you can compare the value of only those tablespaces that start with **SYS**. For this, simply specify SYS in the **SEARCH FOR DESCRIPTORS** text box (as shown by Figure 6.28), and click the **Submit** button.

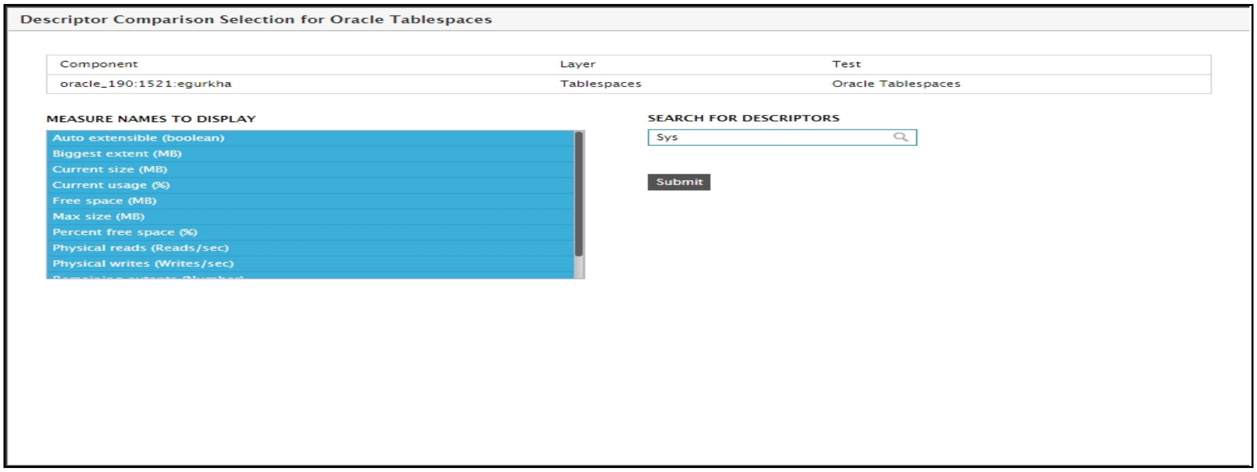

Figure 6.28: Comparing across specific descriptors

Figure 6.29 will then appear, comparing the usage metrics of those descriptors that start with **SYS**.

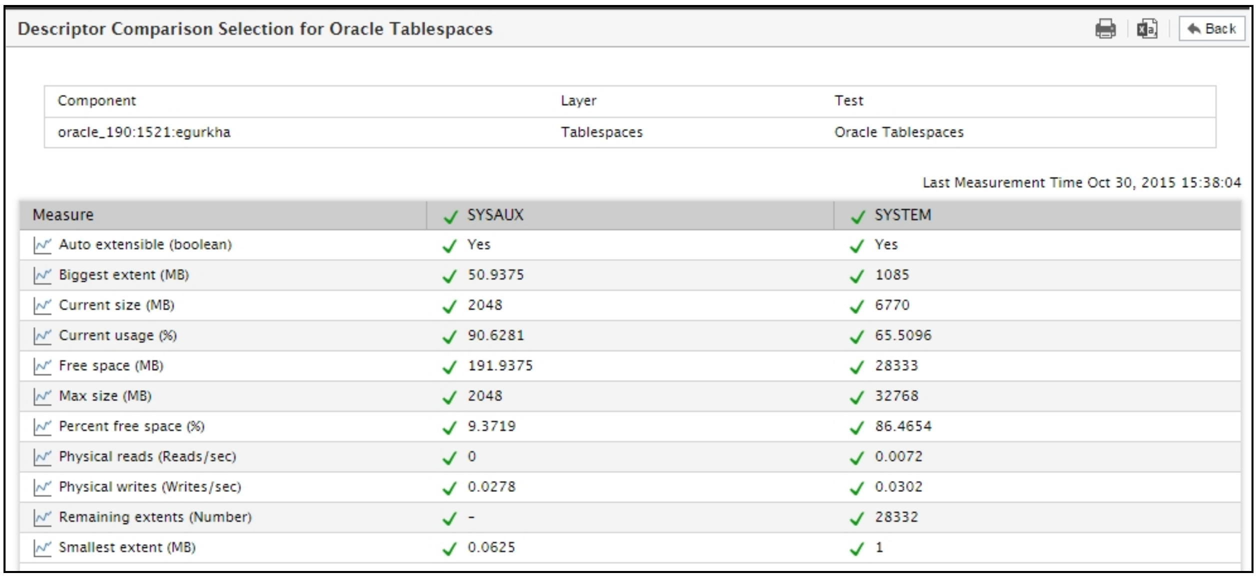

Figure 6.29: A comparison of the usage of all Oracle tablespaces that start with the string SYS

At any given point in time, you can even view a comparison graph of any of the measures chosen for comparison, by clicking the icon alongside the measure name in Figure 6.29. By default, this graph will be generated for a 1 hour timeline for the **TOP-10** descriptors. For instance, the graph in Figure 6.30 compares the variations in the **Current usage** of the top-10 Oracle tablespaces, during the last 1 hour. If required, you can pick a different **Top-N** option from the drop-down in Figure 6.30 and even select a different **Timeline** for the graph.

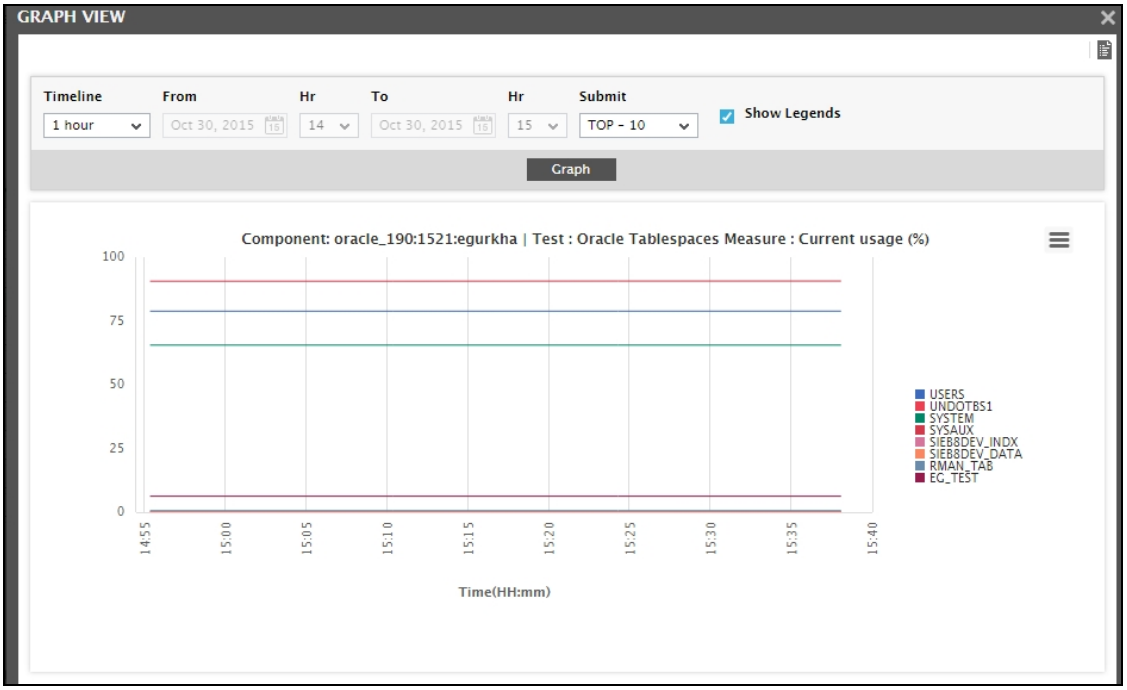

Figure 6.30: A graph comparing the value of a particular measure across descriptors

Let's now go back to the **Security Log** test example. By default, this test displays the details of all security audit failures that occurred, as part of detailed diagnostics. Sometimes, you may want to focus only on those failure events that occurred recently, and not all of them. To ensure this, you will have to reconfigure the **Security Log** test. Typically, to make changes to the configuration of any test, you will have to login to the eG administrative interface, open the **TEST CONFIGURATION** page, select the test to be modified, and then make the changes. To save the time and labor involved in this exercise, eG allows you to launch the **TEST CONFIGURATION** page of a specific test and make necessary changes to it, from the eG monitoring console itself! For example, to reconfigure the **Security Log** test so that the 10 most recent failure events are alone listed as part of detailed diagnostics, first, click the i (**Info**) icon against the test name in Figure 6.8. Doing so will open Figure 6.31, where additional details about the **Security Log** test are displayed. These details include where the test is being executed (i.e., the **Measurement Host**), how frequently the test executes (**Test frequency**), when the test was executed last (**Time since last measure**), and how the test is being executed – i.e., whether internally or remotely (**Test type**).

Figure 6.31: Viewing the details of a test

Note:

As you can see, the **TEST DETAILS**page of Figure 6.31 displays the *IP address:port* of the problem **Component**. If the problem component had been managed using its host/nick name, then the **Component** field in Figure 6.31 will display the *host/nick name:port* of the component. The IP that maps to the displayed host/nick name will appear adjacent to it, within brackets. For example, if the **Component** 192.168.10.28:10818 had been managed using the nick name *Nasserver*, then the **Component** field the **TEST DETAILS** page will display the value: **Nasserver:10818 (192.168.10.28).**

Additionally, the **TEST DETAILS** window of Figure 6.31 also allows you to make changes to the configuration of the **Security Log** test. For this, you will have to click the (**Configure test**) icon alongside the **Test name** in Figure 6.31. This will immediately launch the **TEST CONFIGURATION** page of the **Security Log** test for the **Citrix XenApp server** being monitored (see Figure 6.32). Here, change the value of the **FAILUREEVENTSINDD** parameter to **10** and click the **Update** button to save the changes and return to the **TEST DETAILS** page.

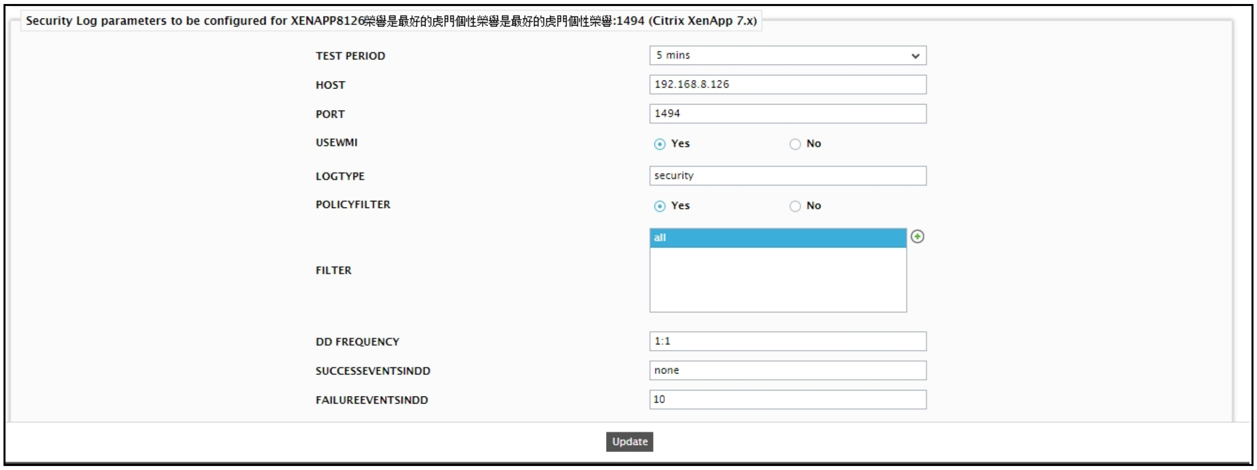

Figure 6.32: Configuring the Security Log test

Likewise, you can also make quick changes to the threshold configuration of the **Security Log** test, without switching to the eG admin interface. For that, click on the (**Configure thresholds**) icon alongside the threshold settings displayed in Figure 6.31, or click on a measure name in the **MEASURES WITHOUT THRESHOLDS** section of Figure 6.31. This will bring up Figure 6.33, where you can change the thresholds of the selected measure.

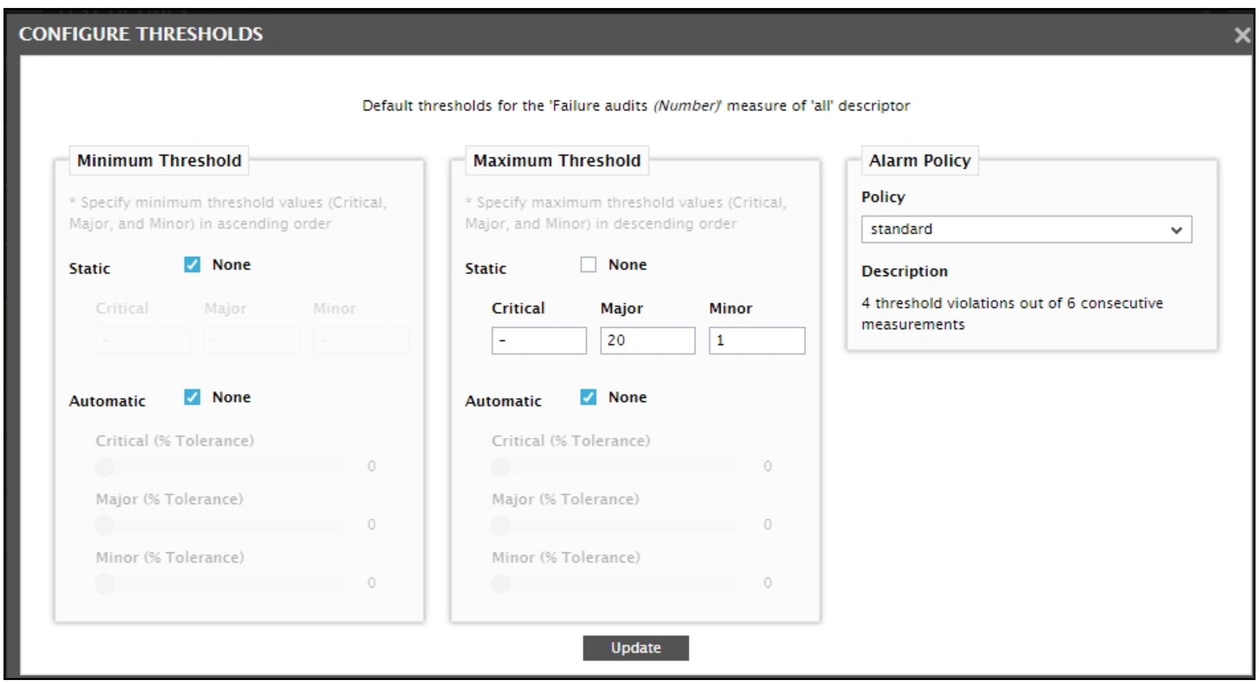

Figure 6.33: Configuring the thresholds of the Security Log test

After making the necessary changes, click the **Update** button in Figure 6.33 to return to the **TEST DETAILS** page.

If a maintenance policy is associated with a test, then, during the configured maintenance period, the corresponding test name in the **Layers** tab page will be accompanied by a icon (as shown in Figure 6.34).

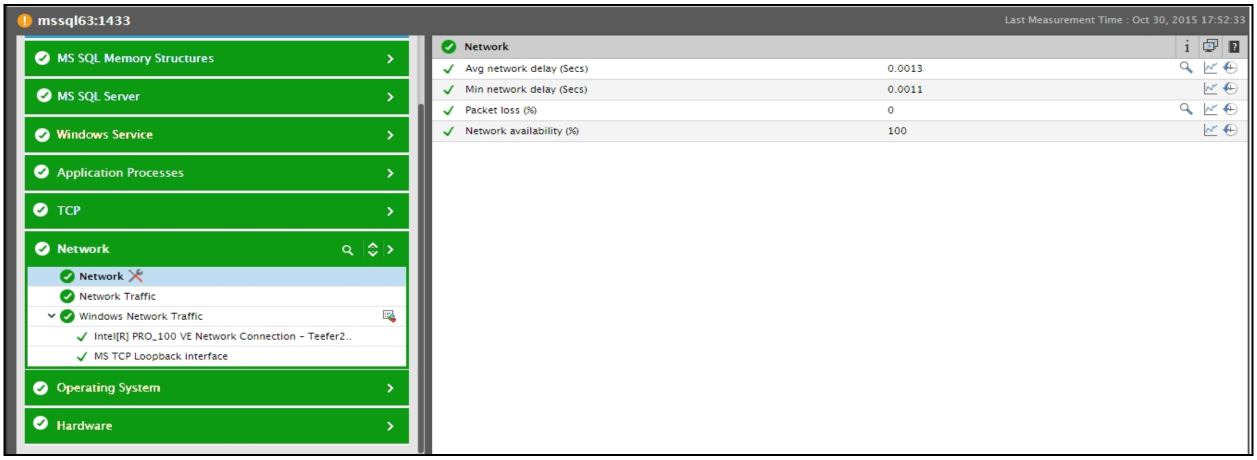

Figure 6.34: A test under maintenance

The icon will continue to appear against the test name until such time that the test is under maintenance. Beyond this period, the icon will disappear. Whenever a test fails to report an expected anomaly, the appearance of the icon against the test name enables administrators to instantly conclude that a maintenance policy prevails during that period, owing to which no alarm was raised by eG Enterprise on that test. For more details about the maintenance policies associated with the test, click on the **(Info)** icon (indicated by Figure 6.34) that appears in the right panel, when the test is clicked. The **TEST DETAILS** page that appears next will include a **Quick Maintenance Details** section (see Figure 6.35), which will display the details of the maintenance policy that is currently active on that test.

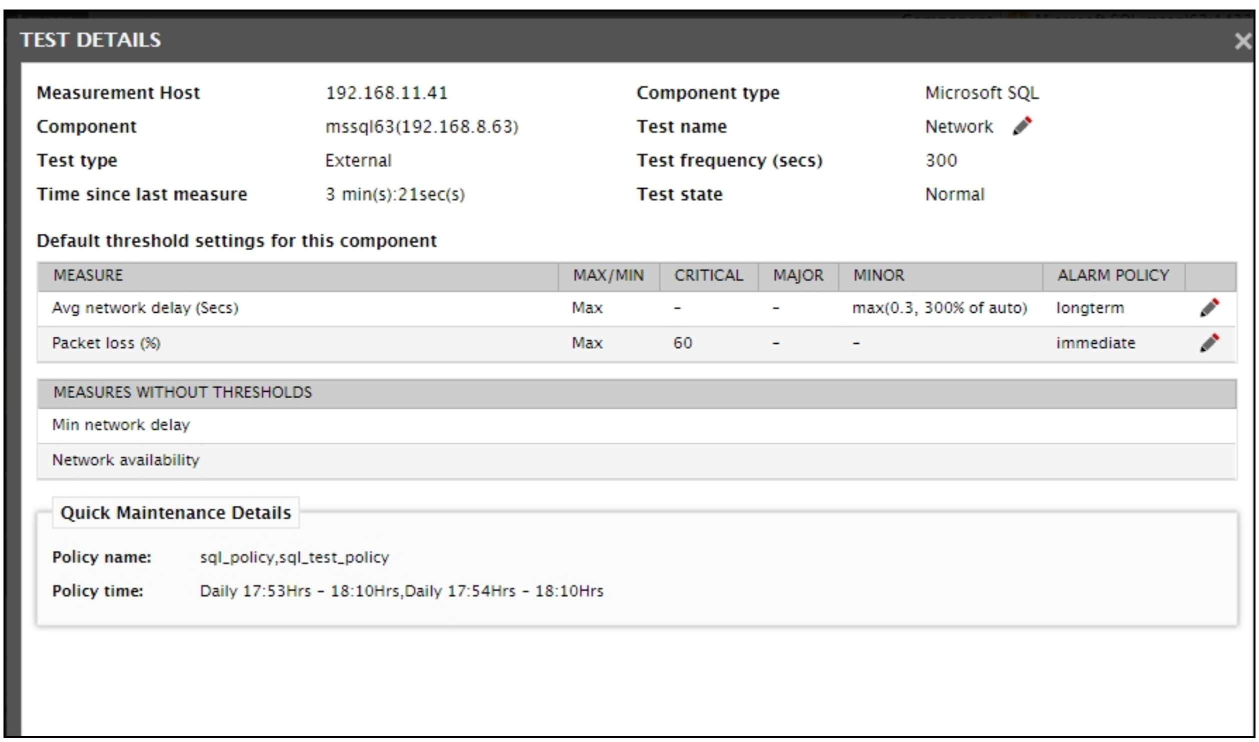

Figure 6.35: Details of the maintenance policy that is currently active on a test

6.1.2 The System Tab Page

Clicking on the **System** tab page in Figure 6.4 will reveal the **System Dashboard** of a component/application.

The **System Dashboard** of an application allows you to focus on the performance of the operating system on which that particular application runs - i.e., the **Operating System** layer of an application. While viewing the layer model of an application using the **Layer Model** tab page, you can, if you so want, instantly switch to its **System Dashboard** for in-depth insights into system performance.

Using the **System Dashboard**, administrators can determine the following:

- The current status of the application host;
- The problems that the host is currently facing, and the type and number of problems it encountered during the last 24 hours;
- The current system configuration (if the eG license enables the **Configuration Management** capability);
- The current state of the critical parameters related to system performance;
- How some of the sensitive performance parameters have performed during the last 1 hour (by default);
- The resource-hungry processors supported by the host, and the disk partitions on the host that are currently experiencing a space crunch.

By default, the **System Dashboard** provides an overview of system performance. Accordingly, the **Overview** option is chosen by default from the **Subsystem** list. Instead of an **Overview**, if you prefer to receive an inside view of system performance - i.e., if you wish to investigate how effectively / otherwise the system in question has been using each of its resources or would prefer to focus on the uptime of the host, you can pick a different option from **Subsystem.**

Note:

To know more about the system dashboard, refer to the Chapter of this document.

6.1.3 The Network Tab Page

Clicking on the **Network** tab page in Figure 6.4 will reveal the **Network Dashboard** of a component/application.

The **Network Dashboard** allows you to zoom into the performance and problems pertaining to the **Network** layer and related layers of a target application/device. Using this dashboard, you can:

- Determine whether/not the application/device currently experiences / has in the past experienced network-related issues;
- Accurately identify the network parameters that are currently failing;
- Understand the current network configuration;
- Analyze network performance over time, study the trends in network connectivity and usage, and accurately deduce problem/performance patterns.
- Identify persistent problems with network health and the network-related layers responsible for the same;

The contents of the **Network Dashboard** and the subsystems it offers for analysis could slightly vary depending upon whether the target is an application or a network device. While the **Network Dashboard** of a host/application enables you to focus on both the network and TCP connections handled by the target, the same for a network device sheds light on the network connectivity of the device and the traffic handled by the network interfaces supported by the device. Accordingly, the **Network Dashboard** of an application/host offers **Network** and **Tcp** as its **Subsystem**s, and that of a network device offers **Network** and **NetworkInterfaces** as its **Subsystem** s. If the target application is a Windows-based one, then the **Subsystems** list will include an additional **WindowsInterfaces** option, which provides the performance information related to the traffic handled and bandwidth used by the network interfaces supported by the system.

The **Network Dashboard** of Hyper-V servers on the other hand, will additionally support a **Hyper-V Switches** and a **Hyper-V Network Adapters** sub-system. Similarly, the **Network Dasboard** of a vSphere/ESX server will include an additional **VirtualNetwork** sub-system.

Note:

To know more about the network dashboard, refer to the Chapter of this document.

6.1.4 The Application Tab Page

In order to ascertain how well an application is/has been performing, analysis of the performance of the **System** and **Network** layers of that application alone might not suffice. A closer look at the health of the **Application Layers** is also necessary, so as to promptly detect instantaneous operational issues with the target application, and also proactively identify persistent problems or a consistent performance degradation experienced by the application. To provide administrators with such in-depth insights into overall application performance and to enable them to accurately isolate the root-cause of any application-level slowdown, eG Enterprise offers the **Application Dashboard**. Each of the critical applications monitored by eG Enterprise is accompanied by an exclusive application dashboard. The contents of the dashboard will therefore primarily vary depending upon the application being monitored.

To view the **Application Dashboard**, click on the **Application** tab page in Figure 6.4.

Like the **System** and **Network** dashboards, the contents of the **Application** dashboard too are further governed by the **Subsystem.** By default, the **Overview** option is chosen from the **Subsystem** list. If need be, you can change this default setting by picking a different option from the **Subsystem** list.

Note:

To know more about the application dashboard, refer to Dashboards of this document.

6.2 Monitoring System Components

While eG Enterprise focuses primarily on monitoring applications, many administrators still prefer to view their infrastructure from a hardware perspective - i.e., as systems they support. The eG monitoring console now provides a “system view”, which represents the overall health of systems in the target infrastructure, with a mapping of the applications that are executing on these systems.

To access the **SYSTEMS LIST** page that provides the ‘system view’, click on the icon available in the **Monitor** tab. Then, select the **Systems** option from the **Hosts/Applications** tile.

The **SYSTEMS LIST** page indicates the current state of those systems/hosts that have been assigned to the current user for monitoring. In case of a 'supermonitor' however, the status of all managed systems in the environment will be available in this page.

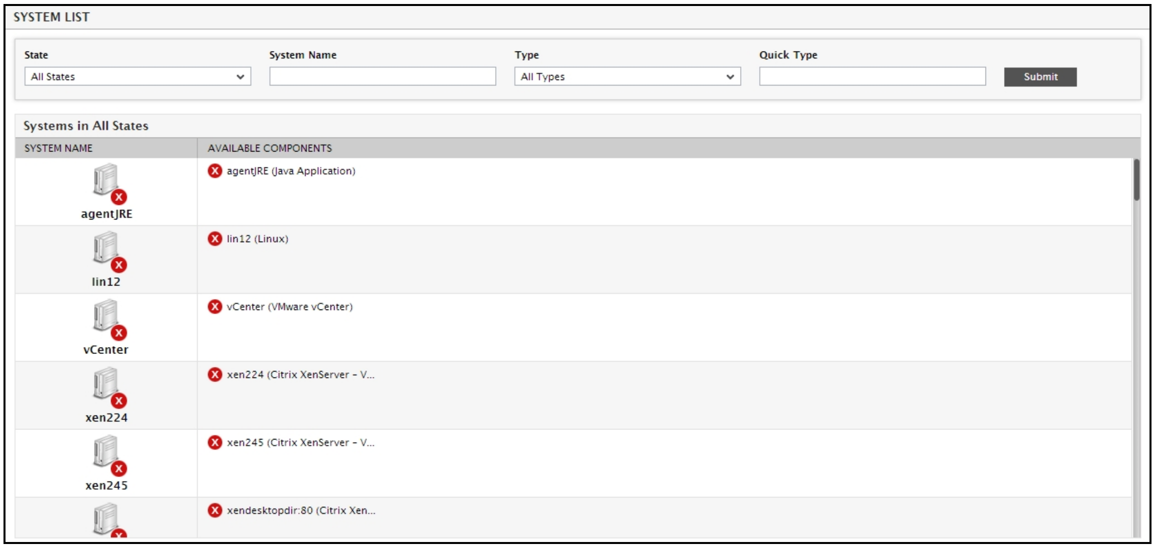

Figure 6.36: Viewing a host-wise list of components and their state

Typically, the state of a system is determined by the state of one/more applications executing on it. Therefore, if a particular system is in an abnormal state, then administrators might want to know which application(s) executing on that system is causing the anomaly. To enable administrators to determine this, against every system listed in this page, eG Enterprise displays the application components deployed on that system along with the component state (see 6.2).

Once the problematic component on a system is identified, you can click on that component to view its layer model, tests, and measurements; this will reveal more details about the exact problem affecting component performance.

If too many systems are being currently monitored, then you can use the search options in this page to search for a particular system(s) and view its current state. To achieve this, follow the steps given below:

- To view the systems that are in a particular state, select the state of your choice from the **State** list. Figure 6.37 incidentally, displays all systems in the **Abnormal State**.
- To view the state of specific systems, specify the whole/part of the **System name** in this page. Multiple system names can be provided, separated by a white space - for example, *agate orac92* (see 6.2). Then, click the **Submit** button.

Figure 6.37: The systems with a given name

- If you want to view the state of systems that are hosting components of a particular type, select the **Type** to search for. You can also enter the whole/part of the component-type to search for in the **Quick Type** list box. This will enable you to view the status of those systems on which component-types that contain the specified search string are executing. Here again, you can search for multiple component types, by separating every type specification by a white space - for example, *Web Oracle*. Finally, click the **Submit** button.

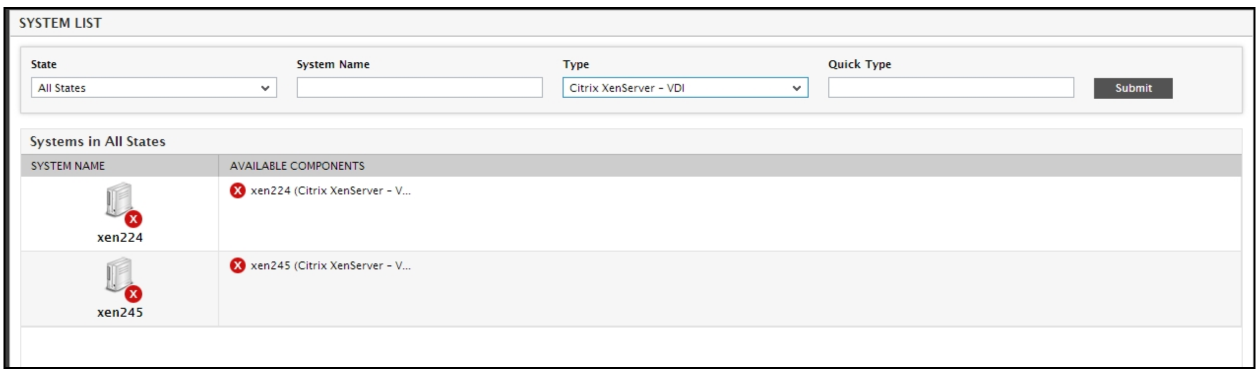

Figure 6.38: The systems of a chosen type

6.3 Monitoring Virtual Servers

If your environment comprises of virtualized components, then Click on the icon available in the ***Monitor*** tab. Then, select the ***Virtual Servers*** option in the ***Hosts/Applications*** tile to view the current state of the VMware® ESX hosts, Solaris virtual servers, and MS Virtual server hosts that have been managed by the eG Enterprise system (see Figure 6.39).

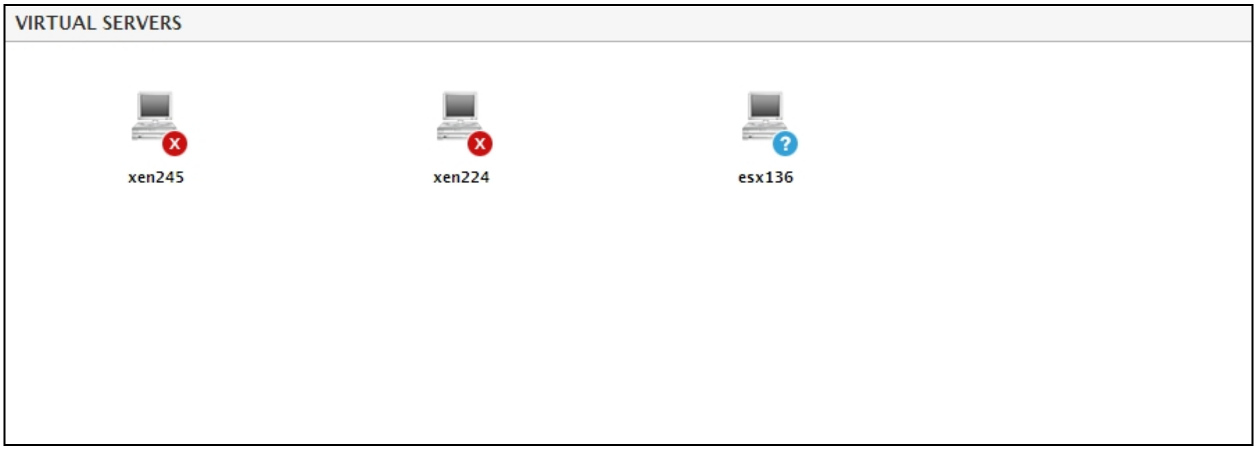

Figure 6.39: The state of virtual hosts

Clicking on any of the virtual servers in Figure 6.38 opens Figure 6.39 that reveals the server applications and the virtual hosts that these applications are running on. The arrows in Figure 6.40 depict the dependencies between the virtual host and the servers running on it.

Figure 6.40: Depicts the server applications that have been deployed on the guest OS' of the ESX server

Sometimes, excessive resource usage at the virtual host-level could cause the server applications on the guest to experience performance degradations. Under such circumstances, the server application to virtual host mapping helps eG Enterprise automatically correlate both the problems and proactively figure out that the issue with the virtual host has impacted the server applications running on it.

6.4 Monitoring Public Clouds

If public clouds have been configured for monitoring in your environment, click on the icon available in the ***Monitor*** tab. Then, select the ***Public Clouds*** option in the ***Hosts/Applications*** tile to invoke a **CLOUD SERVERS** page, where the complete list of monitored cloud servers will appear. For instance, if one/more *AWS-EC2 Regions* are being monitored in your environment, then the **CLOUD SERVERS** page will indicate the names and the current state of the *AWS-EC2 Regions* (see Figure 6.41).

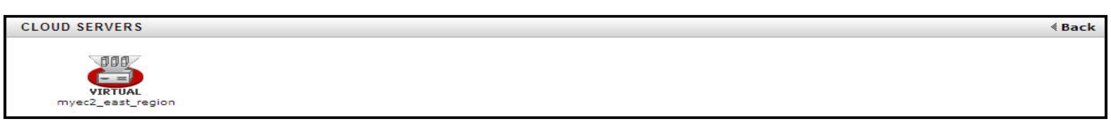

Figure 6.41: List of virtual public clouds that are being monitored

If a cloud server in Figure 6.41 appears to be in an abnormal state, then you can zoom into the root-cause of that abnormality by clicking on the cloud server. This will lead you to Figure 6.41, which reveals the topology of the cloud infrastructure being monitored.

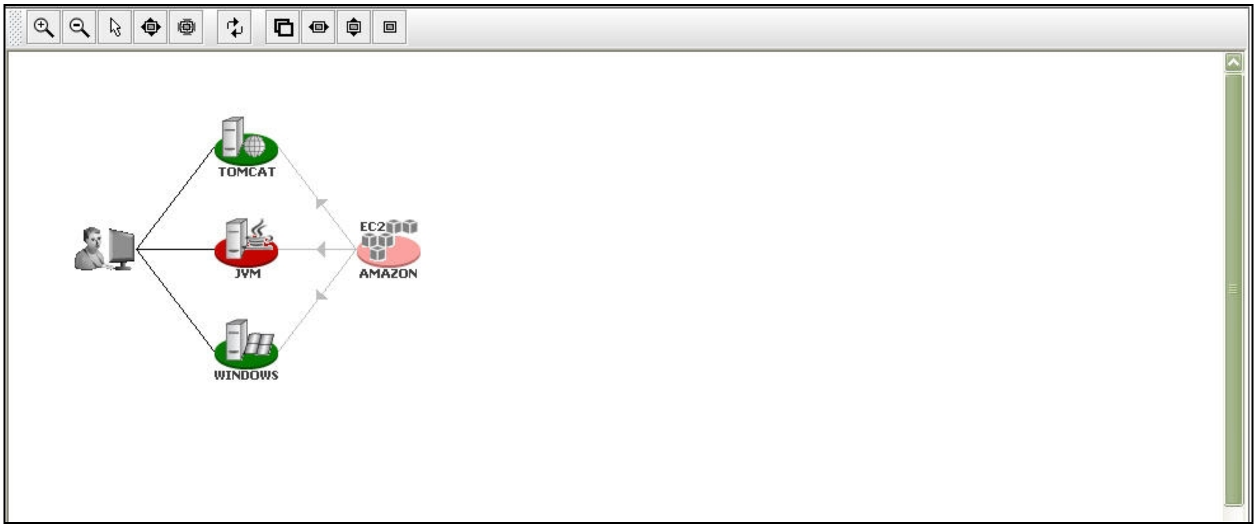

Figure 6.42: Topology of the cloud infrastructure

The eG Enterprise system is capable of automatically discovering the IP address and operating system of the instances launched on a cloud and tracking the powered-on state of these instances. Additionally, the solution also intelligently determines which managed applications are operating on these instances and automatically maps these virtual applications with the cloud server. The topology of Figure 6.42 reveals this server-application mapping and also indicates the current state of the cloud server and the applications operating on each of its instances. A resource deficiency at the cloud server level can affect the resource allocation to the instances, and can in turn cause the performance of the applications executing on the instances to suffer. This implies that a problem with the cloud server can ripple and affect the performance of the applications launched on it. The direction of the arrows in Figure 6.42 depicts this inter-dependency. The state of the cloud service is therefore determined by the state of the cloud server and that of the applications executing on the server. Clicking on the cloud server will lead you to Figure 6.42, where you can view the layer model, tests, and measurements of the cloud server.

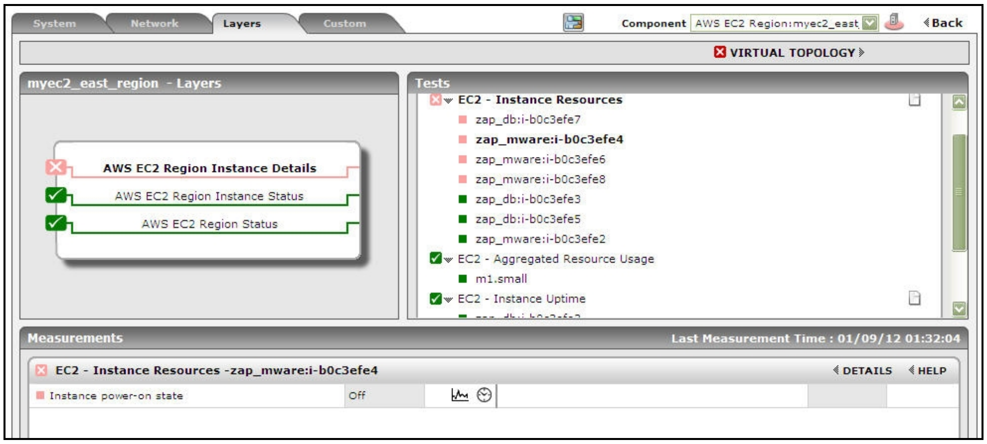

Figure 6.43: The layer model, tests, and measurements of a cloud server

6.5 Monitoring Aggregate Components

eG Enterprise typically monitors every component of a type, separately. However, sometimes, administrators might want to receive an aggregated view of the performance of two/more components of a type. 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*. The eG Enterprise system then automatically aggregates the metrics reported by the components in the group by applying pre-configured aggregate functions on them, and reports these metrics as if they were extracted from the managed *aggregate component*. 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.

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;

If one/more *aggregate components* have been managed in your environment, then the **Aggregates** bar graph in the **Infrastructure Health** section of the **Monitor Home Page** will indicate the number of *aggregate components* that are being monitored and their current state (see Figure 6.44).

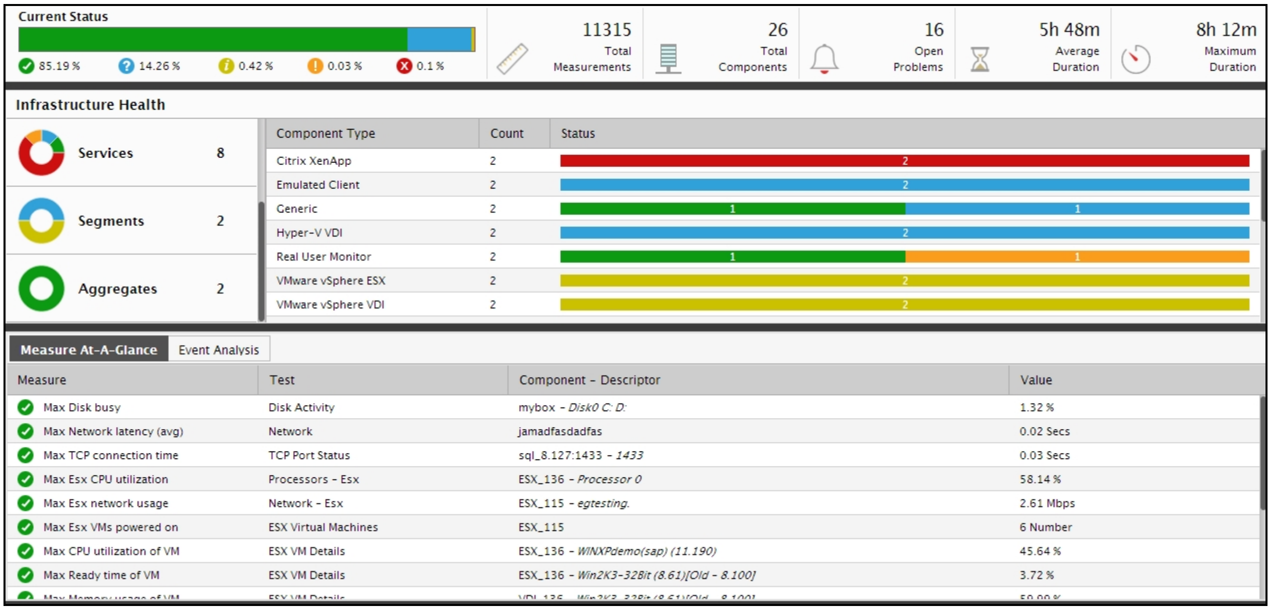

Figure 6.44: The Aggregates bar graph in the Infrastructure Health section

Each division of the **Aggregates** bar graph will indicate the current state of the managed aggregate components and the number of aggregate components in each state.

Clicking on a division in the **Aggregates** bar graph will lead you to the **COMPONENT AGGREGATES** page. This page lists the *aggregate types* that have been managed and the names of *aggregate components* of each type that are currently in the state represented by the division that was clicked on (see Figure 6.44).

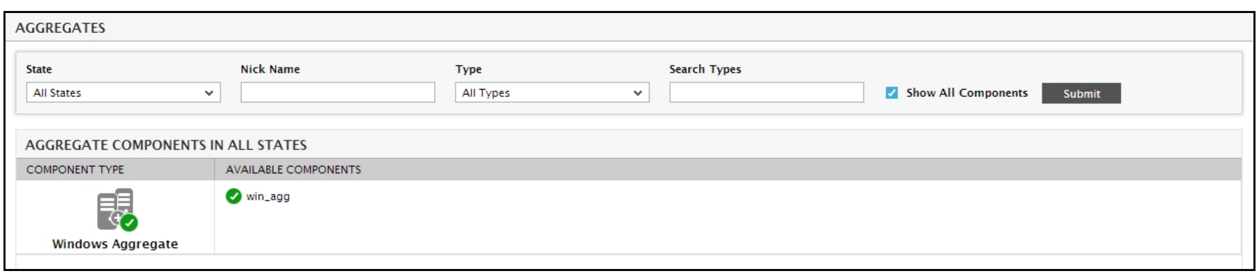

Figure 6.45: The COMPONENT AGGREGATES page

To zoom into the layer model of an aggregate component, click on that component in the **COMPONENT AGGREGATES** page. The **Layer** tab page will then open displaying the layer model of the *aggregate component* (see Figure 6.45).

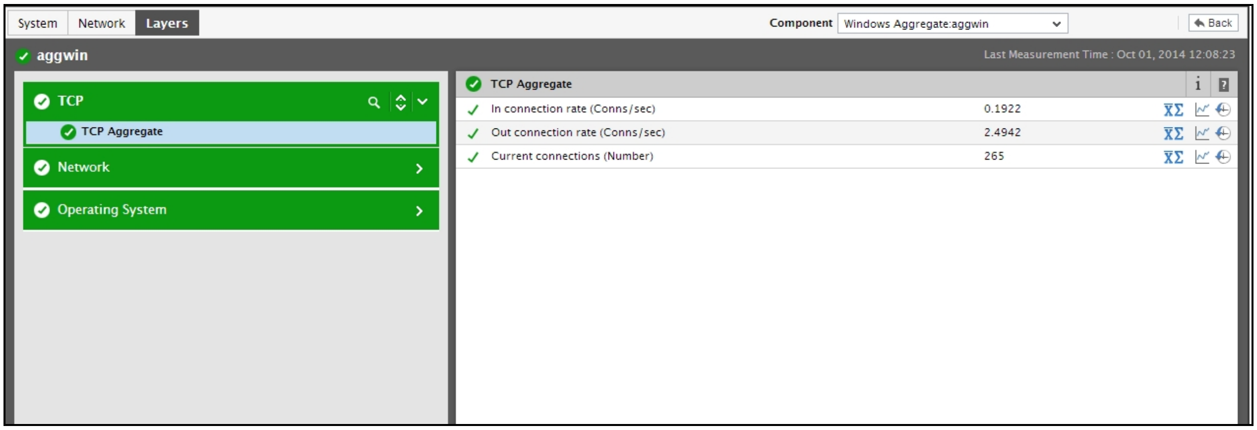

Figure 6.46: Layer model of an aggregate component

Like individual components, aggregate components too are represented in the eG monitoring console using a hierarchical set of layers. Typically, an aggregate component will be represented using the same set of layers that are associated with the corresponding non-aggregate component. For instance, the *Citrix XenApp Aggregate* component will support the same layer model as the *Citrix Xenapp* component.

Each layer of an aggregate component will be mapped to a set of aggregate tests. Some layers will be associated with precanned aggregate tests. Besides these default tests, administrators can also add new aggregate tests and associate them with the *aggregate components*. To know how to add/configure aggregate tests, refer to *Chapter 3* of this document.

While clicking on a layer will list all the aggregate **Tests** mapped to that layer, clicking on a test will open a **Measurements** panel where all the *aggregated metrics* collected by that test will be displayed (see Figure 6.46).

To perform metrics aggregation, the eG Enterprise system applies certain aggregate functions on the measures collected from across all the components of a type that are grouped under a particular *aggregate component*. In case of precanned aggregate tests, these functions are also hard-coded into the eG Enterprise system. In case of a user-configured aggregate test however, the users can indicate what function should be applied on each measure reported by that test. The **Measurements** panel of the **Layer** tab page not only displays the names of the aggregated measures and their aggregated values, but also indicates what function has been applied on the measure to perform aggregation. Statistical symbols have been used to represent functions. These symbols and the functions they represent have been discussed in the table below:

Symbol	Function
Σ	Sum
$\uparrow$	Max
$\downarrow$	Min
$\bar{X}$	Average
$\bar{X}\Sigma$	Avg-Sum

Move your mouse pointer over a symbol in the **Measurements** panel to know what function it represents. Clicking on the symbol that corresponds to a measure in the **Measurements** panel will open Figure 6.46, using which you can understand how the value of that measure was computed. The **Component Name** column in Figure 6.46 lists all the individual components that have been grouped under the *aggregate component* in question. Against every **Component Name**, the **Aggregated Measure** value of that component will be displayed. For instance, for a measure on which the *Avg-Sum* function has been applied, the **Aggregated Measure** column will display the *average* value of that measure for the displayed **Component Name** during the last test frequency. The sum of the *average* values of all **Component Name**s will be displayed as the value of the measure in the **Measurements** panel.

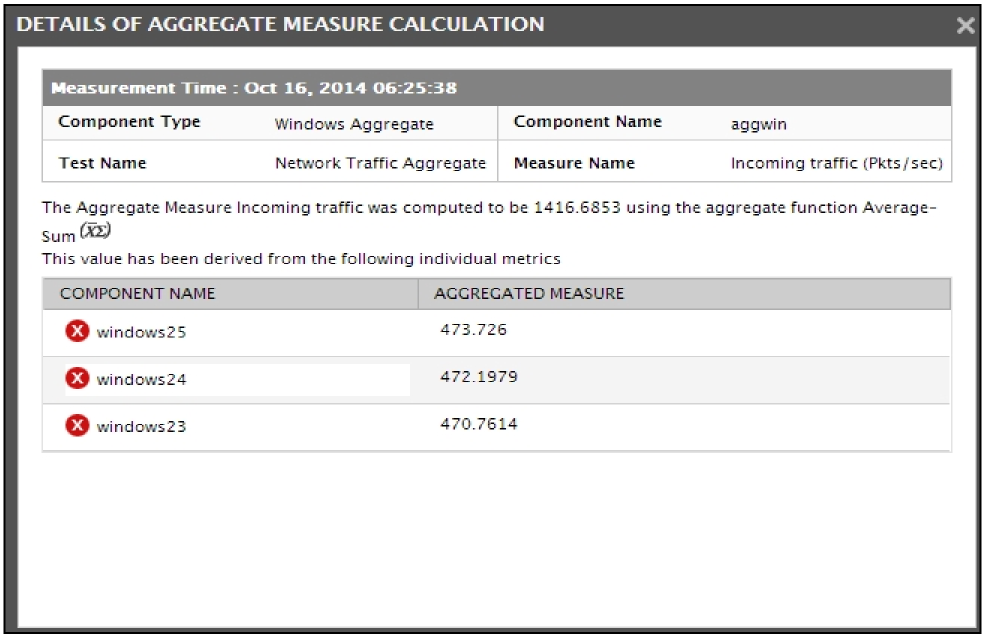

Figure 6.47: How the aggregate measure was calculated?

Likewise, clicking on the **DETAILS** button in the **Measurements** panel will list all the member components that have been included as part of the aggregate component (see Figure 6.46).

Clicking on the **Graph** icon corresponding to an aggregate measure will reveal a **Measure Graph** that plots the time-of-day variations in the aggregated metrics during the last hour (by default) (see Figure 6.48).

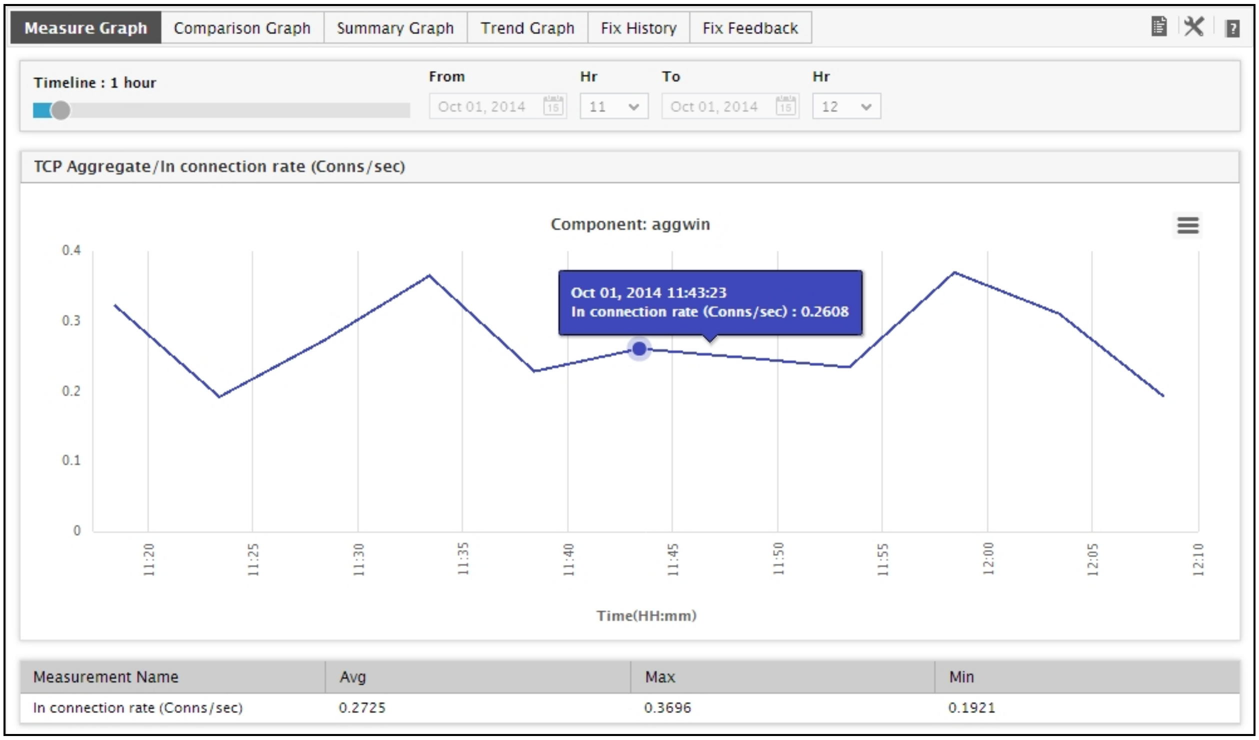

Figure 6.48: Viewing a graph of an aggregate measure

A **Comparison Graph** will also be additionally available for an aggregated measure. To view this graph, you will have to click on the **Comparison Graph** tab page next to the **Measure Graph** tab page (see Figure 6.48). Using this graph, you can compare the historical measures of the target aggregate component with that of each of its member components for a given timeline. If the aggregate measure reports a deviation, then this graph will enable you to isolate the member component that could have caused the deviation. For instance, say that you have managed a *Citrix XenApp Aggregate* component, which is associated with all the Citrix XenApp servers in your Citrix farm. Assume that you have configured this aggregate component to alert you if more than *100* sessions are active on your farm. If this anomaly occurs – i.e., if the number of sessions to the Citrix farm exceeds *100* – then the *Citrix XenApp Aggregate* component will naturally trigger an alert. At this juncture, you can use the **Comparison Graph** to figure out the following:

- Whether this increase in session count was sudden or consistent;
- Whether it is because of a sudden/consistent growth in the number of sessions to a particular Citrix server in the farm; if so, you can even use this graph to point to the problem server.
- Whether it is because the number of sessions to all the servers in the farm has increased over time;

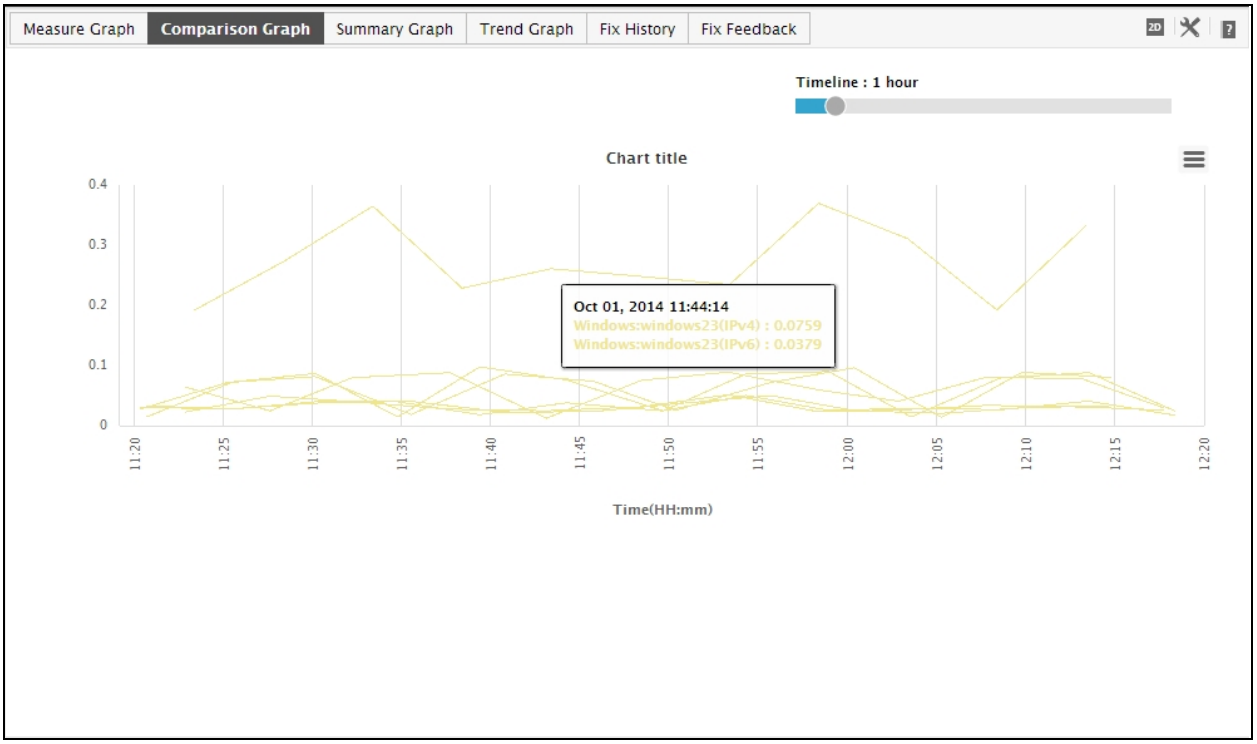

Figure 6.49: The Comparison Graph of an aggregate measure

Moreover, in Figure 6.49, you can see that there is a time lag in the *aggregate graph* line. This is because, aggregate measures are computed one measurement period before the individual components. This is done to ensure that all measures are considered when aggregating metrics.

7

Monitoring Segments

eG Enterprise can also be used to monitor components that are not associated with services. For monitoring components that are a part of the component topology configured, choose the **Segments** option shown in Figure 7.1.

This will result in the display of the different segments configured for the target infrastructure, and their respective states (see Figure 7.1). Each segment will be accompanied by the IP/hostname and state of the segment components associated with that segment.

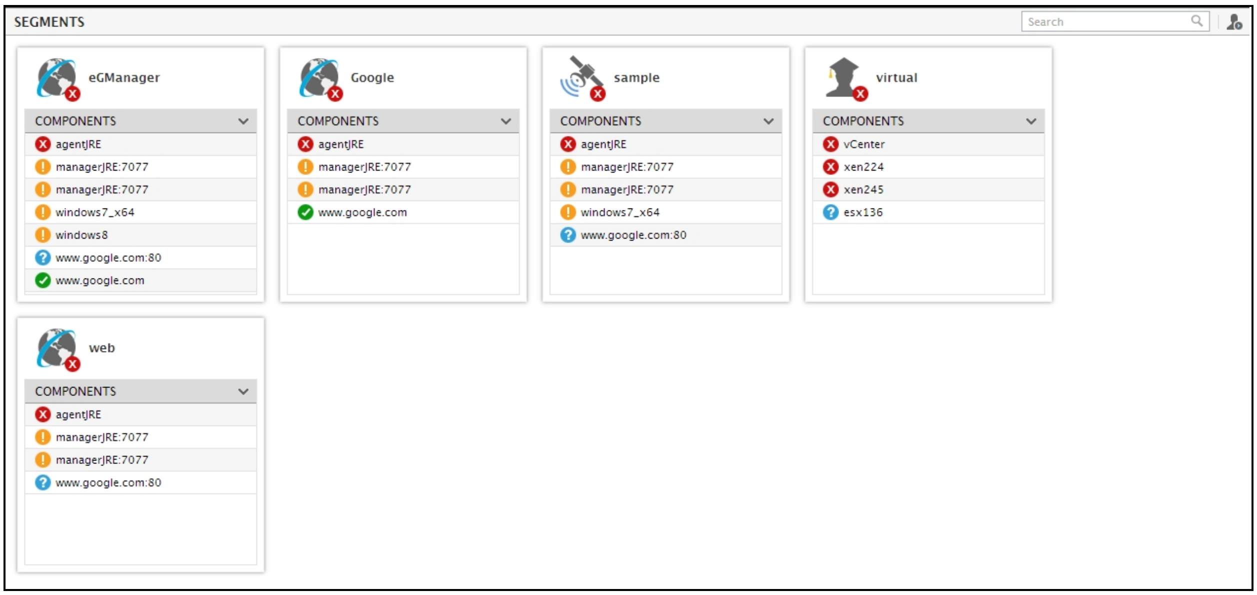

Figure 7.1: Segments configured for the environment

Note:

If you want all the component-types associated with a problem segment to be displayed (in Figure 7.1) instead of the component names, then you will have to set the **Display component types in segment/service list** flag in the **MONITOR SETTINGS** page (Click on the icon available in the **Admin** tab, select the **Monitor** option from the **Settings** tile) in the eG administrative interface to **Yes**. By default, this is set to **No**.

Note:

By default, against each segment displayed in the **SEGMENT LIST** page, the top-10 **Components** included in that segment will be displayed. Typically, to identify the top-10 components, eG Enterprise automatically sorts all the components associated with a segment on the basis of their current state, arranges the sorted list in the alphabetical order of the component names, and picks the first 10 components of this list.

If you want more number of components to be displayed against each segment, do the following:

- Login to the eG administrative interface.
- Click on the icon available in the **Admin** tab, select the **Monitor** option from the **Settings** tile.
- Click the **Other Settings** sub-node under the **General** node in the tree-structure in the **MONITOR SETTINGS** page.
- In the right panel, modify the default value 10 that is displayed in the **Components count in segment/service/zone list** text box.
- Click the **Update** button to save the changes.

If a large number of segments are displayed, you can search for a particular segment, by specifying the whole/part of the segment name in the **Search** text box, and clicking the right-arrow button adjacent to the text box. Doing so will reveal all segments with names that embed the specified search string.

Click on the **Associations** ( )button in Figure 7.2 to view the state of all the fully configured segments in their environment, and the users to which each of them are assigned.

Figure 7.2 also reveals an 'arrow' button to its left. Clicking on this 'arrow' opens a left panel that consists of a tree-structure (see Figure 7.2). Each fully-configured segment will appear as the 'nodes' of this tree - besides the segment name, these 'nodes' will also indicate the current state of the segment.

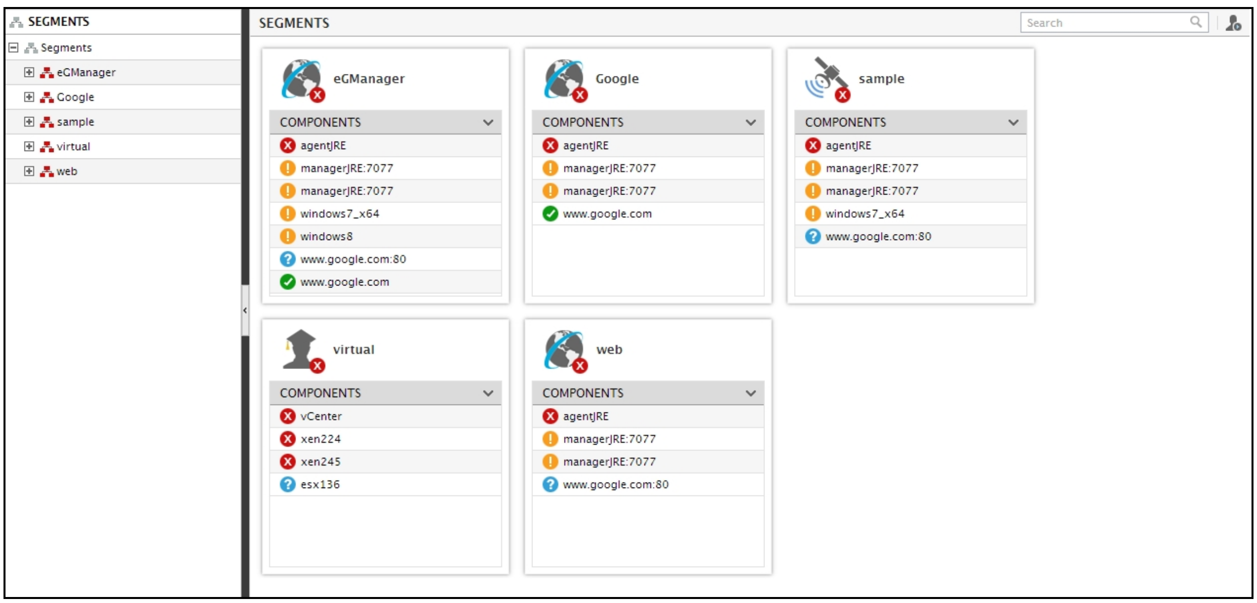

Figure 7.2: Segment Dashboard revealing a tree-structure in its left panel

To view the components that form part of a segment and the state of each of these components, expand the 'node' representing the segment of interest to you. The segment components, along with their current status indicators, will then be listed as the sub-nodes of the 'segment node' as depicted by Figure 4.330. If you click on a 'segment node' in the tree, the right-panel will instantly change to display three tab pages, namely - **Segments**, **Components**, and **Topology**. By default, the **Segments** tab page will open as soon as a 'segment node' in the tree is clicked (see Figure 7.2).

7.1 Systems Tab Page

For every component that is included in the segment, the **Systems** tab page displays the state of each component and the real-time values of key hardware/network/operating system-level measures that have been pre-configured for the components. With the help of this tab page, you can receive instant intimations of hardware failures, network latencies, and other host-level performance ills that the segment components are experiencing currently.

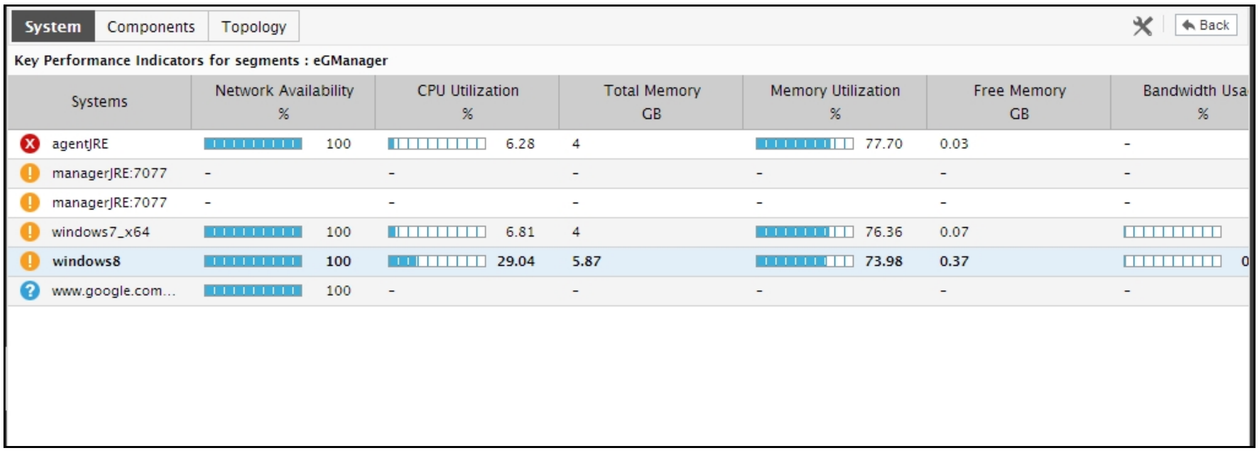

Figure 7.3: The Systems tab page of a Segment node that is clicked on in the tree-structure

If required, you can add more measures to this list or delete one/more existing measures. For this, click on the icon at the right, top corner of Figure 7.3. Figure 7.4 will then appear.

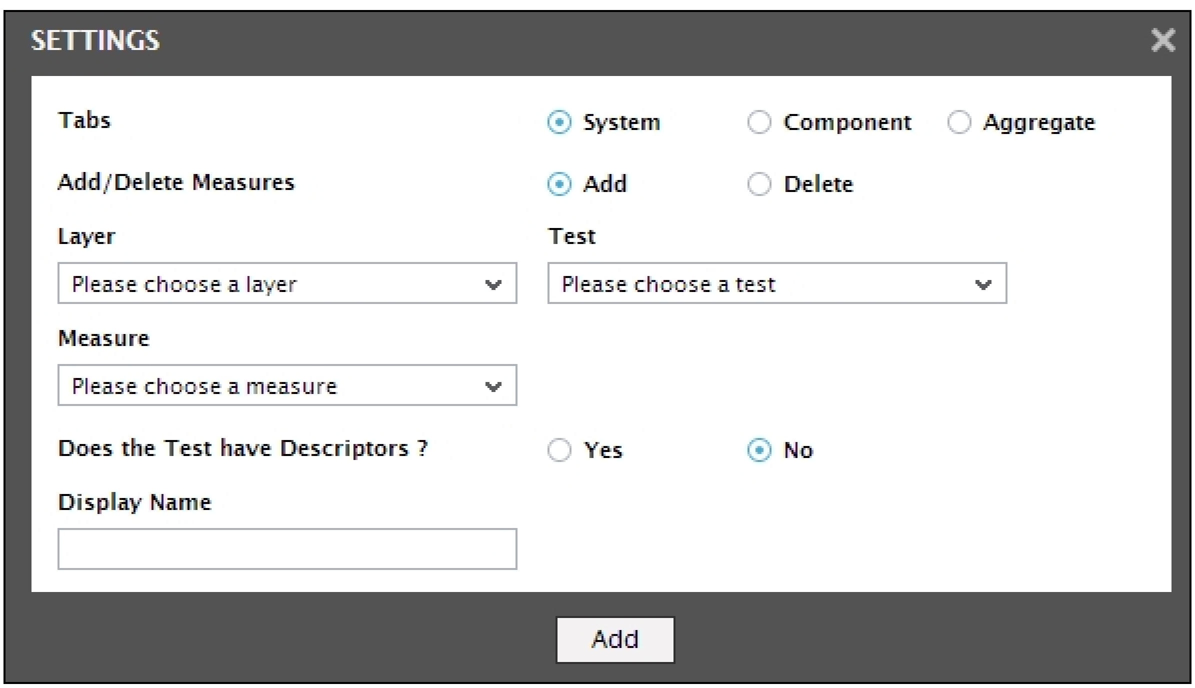

Figure 7.4: Adding a measure to the Systems tab page

To add a measure to the **Systems** tab page using Figure 7.4, do the following:

1. Select the **Systems** option from the **Tabs** section.
2. Set the **Add/Delete Measures** flag to **Add**.
3. The **Layer** list will be populated with all the host-level measures related to all segment components. Pick the **Layer** that reports the measure of interest to you,
4. This will populate the **Test** list with all the tests that execute on the chosen **Layer**. Select the **Test** that reports the measure you need.
5. Then, choose the **Measure** to be added and provide a **Display Name** for the measure.
6. Finally, click the **ADD** button to add the chosen measure to the **Systems** tab page.

To delete a measure from the **Systems** tab page using Figure 7.4, do the following:

1. Select the **System** option from the **Tabs** section.
2. Set the **Add/Delete Measures** flag to **Delete**.
3. Then, from the **Test** list of Figure 7.4, pick the test that reports the measure to be deleted.

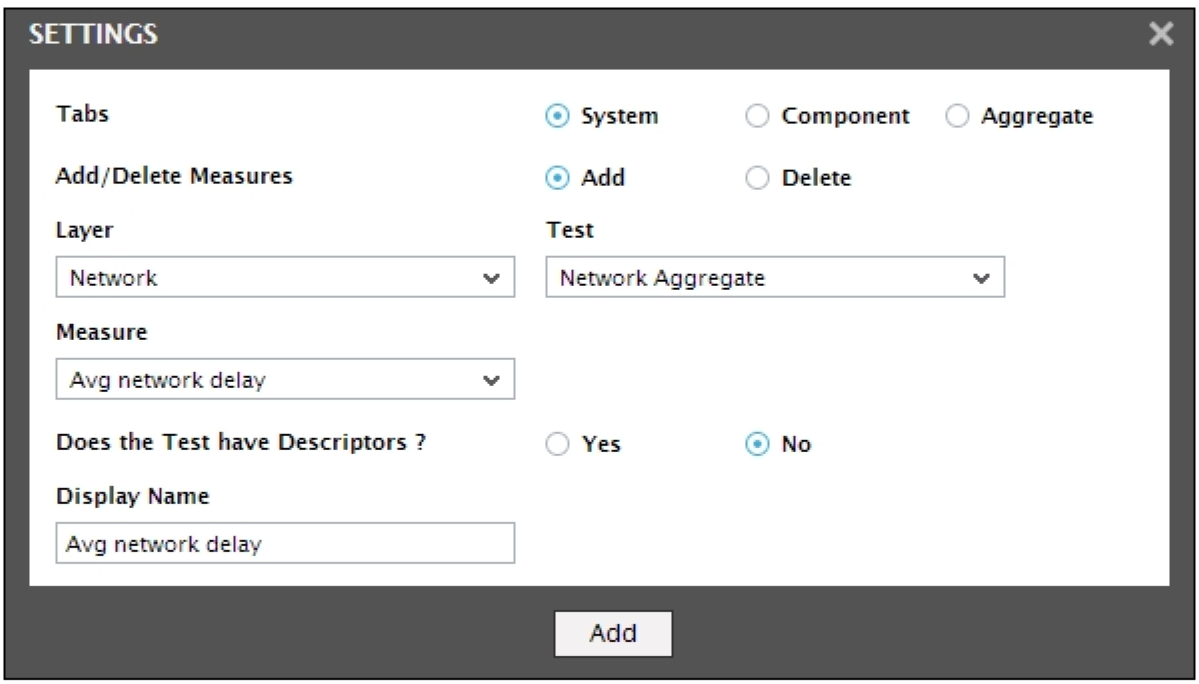

Figure 7.5: Deleting a measure from the Systems tab page

4. Choose the **Measure** to be deleted. The **Display Name** of the selected measure will then be automatically displayed.
5. Click the **DELETE** button to delete the measure.

By default, the contents of the **Systems** tab page are sorted based on the state of the segment components listed therein. If more that one component exists in the same state, then the components of that state will be sorted in alphabetical order. If need be, you can change the sort order. For example, if you wish to sort the components listed in the **Systems** tab page in the descending order of the values of their Disk Usage, just click on the **Disk Usage** label. Doing so, tags the **Disk Usage** label with a down arrow icon – this icon indicates that the **Systems** tab page is currently sorted in the descending order of the total disk space used by each component. To change the sort order to 'ascending', all you need to do is just click again on the **Disk Usage** label or the down arrow icon. Similarly, you can sort the **Systems** tab page based on any column available in the table.

Click on any segment component in the **Systems** tab page to view the layer model of that component.

Note:

- Only a user who is assigned the **Admin** role is allowed to customize the segment dashboard by clicking on the icon.
- While displaying values for descriptor-based measures in the **Systems** tab page, the eG Enterprise system does not display the actual values per descriptor. Instead, the solution computes theaverage or the total sum of values across descriptors and displays it in the corresponding measure column. For instance, for values reported as percentages, the solution computes the average value across descriptors. On the other hand, if the value is reported as a GB or MB, then the total sum of all the descriptor values of the component will be displayed against the component.

7.2 Components Tab Page

If issues at the application-level are hindering segment performance, you can promptly isolate such issues and the segment components affected by them using the **Components** tab page (see Figure 7.6).

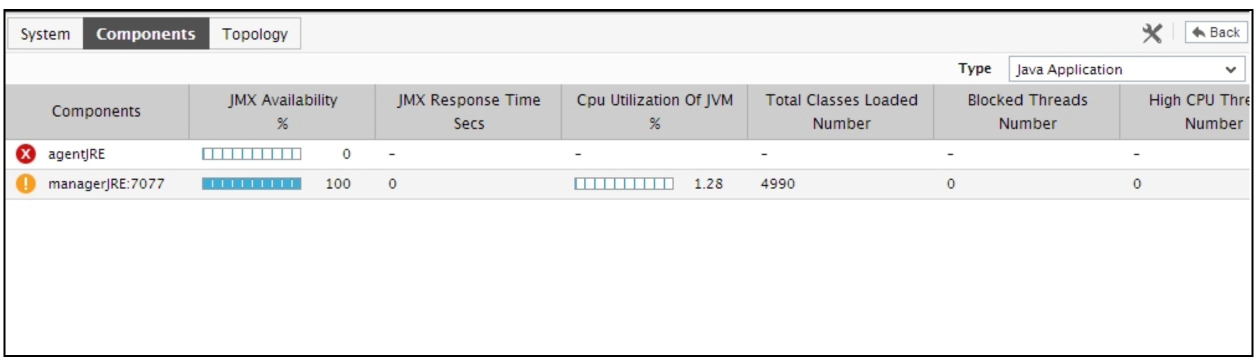

Figure 7.6: The Components tab page of the Segment dashboard

By default, in the **Components** tab page for a particular segment, you will find pre-configured application-level metrics pulled out of that segment component which is currently experiencing problems of the highest priority. This is why, the component type to which such a component belongs is chosen by default from the **Type** list in the **Components** tab page. If need be, you can pick a different **Type** from this list to view the current state of all segment components of that type.

You can even override the default measure list for a component type by configuring additional application-level metrics to be displayed in the **Components** tab page or by removing one/more metrics that pre-exist. For this, click on the icon at the right, top corner of Figure 7.7. Figure 7.7 will appear.

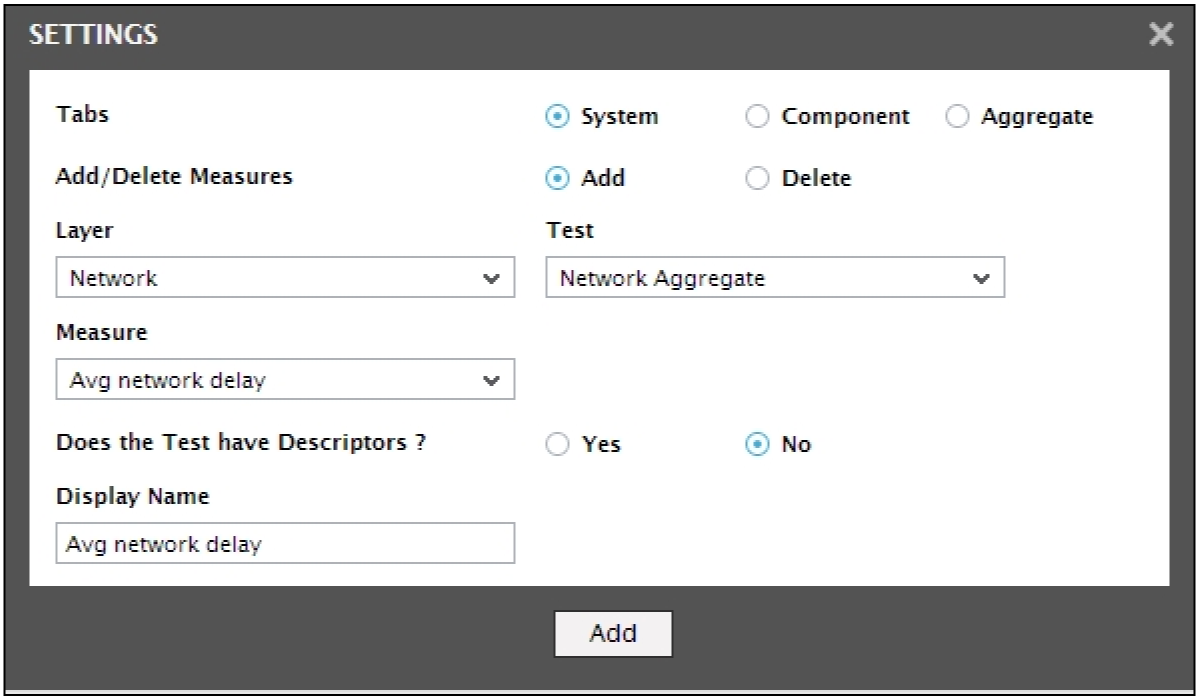

Figure 7.7: Adding a new measure to the Components tab page of the Segment dashboard

To add a measure to the **Components** tab page using Figure 7.7, do the following:

1. Select the **Components** option from the **Tabs** section.
2. Set the **Add/Delete Measures** flag to **Add**.
3. Pick a **Component Type** from the list of component types that are part of all the fully-configured segments in your environment.
4. The **Layer** list will then be populated with all the application-level layers related to all segment components. Pick the **Layer** that reports the measure of interest to you,
5. This will populate the **Test** list with all the tests that execute on the chosen **Layer**. Select the **Test** that reports the measure you need.
6. Then, choose the **Measure** to be added and provide a **Display Name** for the measure.
7. Finally, click the **ADD** button to add the chosen measure to the **Systems** tab page.

To delete a measure from the **Systems** tab page using Figure 7.7, do the following:

1. Select the **Components** option from the **Tabs** section.
2. Set the **Add/Delete Measures** flag to **Delete**.
3. Pick a **Component Type** from the list of component types that are part of all the fully-configured segments in your environment (see Figure 7.8).

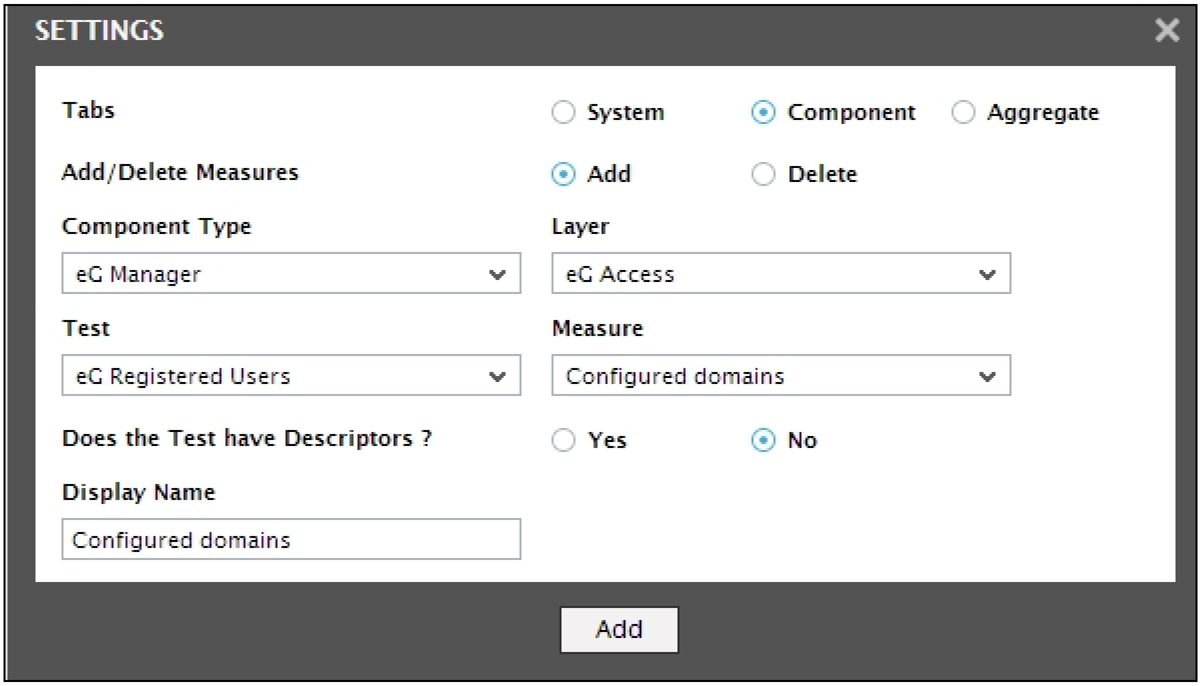

Figure 7.8: Deleting a measure from the Components tab page

4. From the **Test** list. Select the **Test** that reports the measure you need.
5. Choose the **Measure** to be deleted. The **Display Name** of the selected measure will then be automatically displayed.
6. Click the **DELETE** button to delete the measure.

By default, the contents of the **Components** tab page are sorted based on the state of the segment components listed therein. If more that one component of a chosen **Type** exists in the same state, then the components of that state will be sorted in alphabetical order. If need be, you can change the sort order. For example, if you wish to sort the components listed in the **Components** tab page in the descending order of the values of their Response time, just click on the **Response time** label. Doing so, tags the **Response time** label with a down arrow icon – this icon indicates that the **Components** tab page is currently sorted in the descending order of the responsiveness of each component. To change the sort order to 'ascending', all you need to do is just click again on the **Response time** label or the down arrow icon. Similarly, you can sort the **Components** tab page based on any column available in the table.

Click on any segment component in the **Components** tab page to view the layer model of that component.

Note:

- Only a user who is assigned the **Admin** role is allowed to customize the segment dashboard by clicking on the icon.
- While displaying values for descriptor-based measures in the **Components** tab page, the eG Enterprise system does not display the actual values per descriptor. Instead, the solution computes theaverage or the total sum of values across descriptors and displays it in the corresponding measure column. For instance, for values reported as percentages, the solution computes the average value across

descriptors. On the other hand, if the value is reported as a GB or MB, then the total sum of all the descriptor values of the component will be displayed against the component.

7.3 Topology Tab Page

When multiple inter-related components that are part of a segment experience slowdowns, it is often a challenge for administrators to accurately diagnose the 'source' of the slowdown - this is because, the problem in one segment component can affect the performance of other segment components that are dependent on it. To enable administrators to easily understand the inter-relationships between the various components that form a segment, and to help them precisely locate the root-cause of problems with a segment, the segment dashboard provides the **Topology** tab page. If you click on this tab page after selecting a segment node from the tree, Figure 7.9 will appear. This tab page will also appear when you click on a particular segment name in the **SEGMENT LIST** of Figure 7.10.

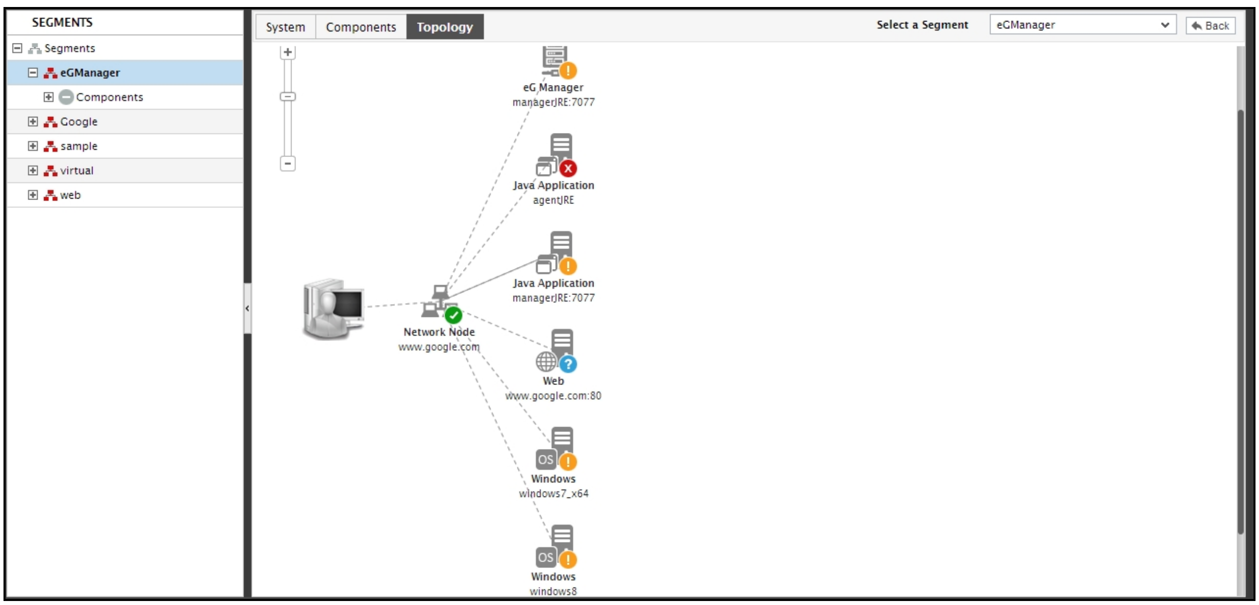

Figure 7.9: The Topology tab page of the Segment dashboard

From the color-coding of the segment components in the topology and the direction of dependencies represented by the 'arrows' in the topology, you will be able to accurately isolate the root-cause of problems in the segment. Clicking on a component in the topology will reveal the layer model of that component.

Large enterprises often have 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 service with tens of servers in a group, the service topology representation can quickly get very cumbersome. To handle such environments, eG Enterprise v 4.1 allows administrators to define server groups that represent a collection of similar servers. By including a server group in a service topology representation, administrators can indicate that all the servers in the group have the same set of dependencies on other parts of the infrastructure. Topology representation using server groups are compact and concise to represent, and simple to understand.

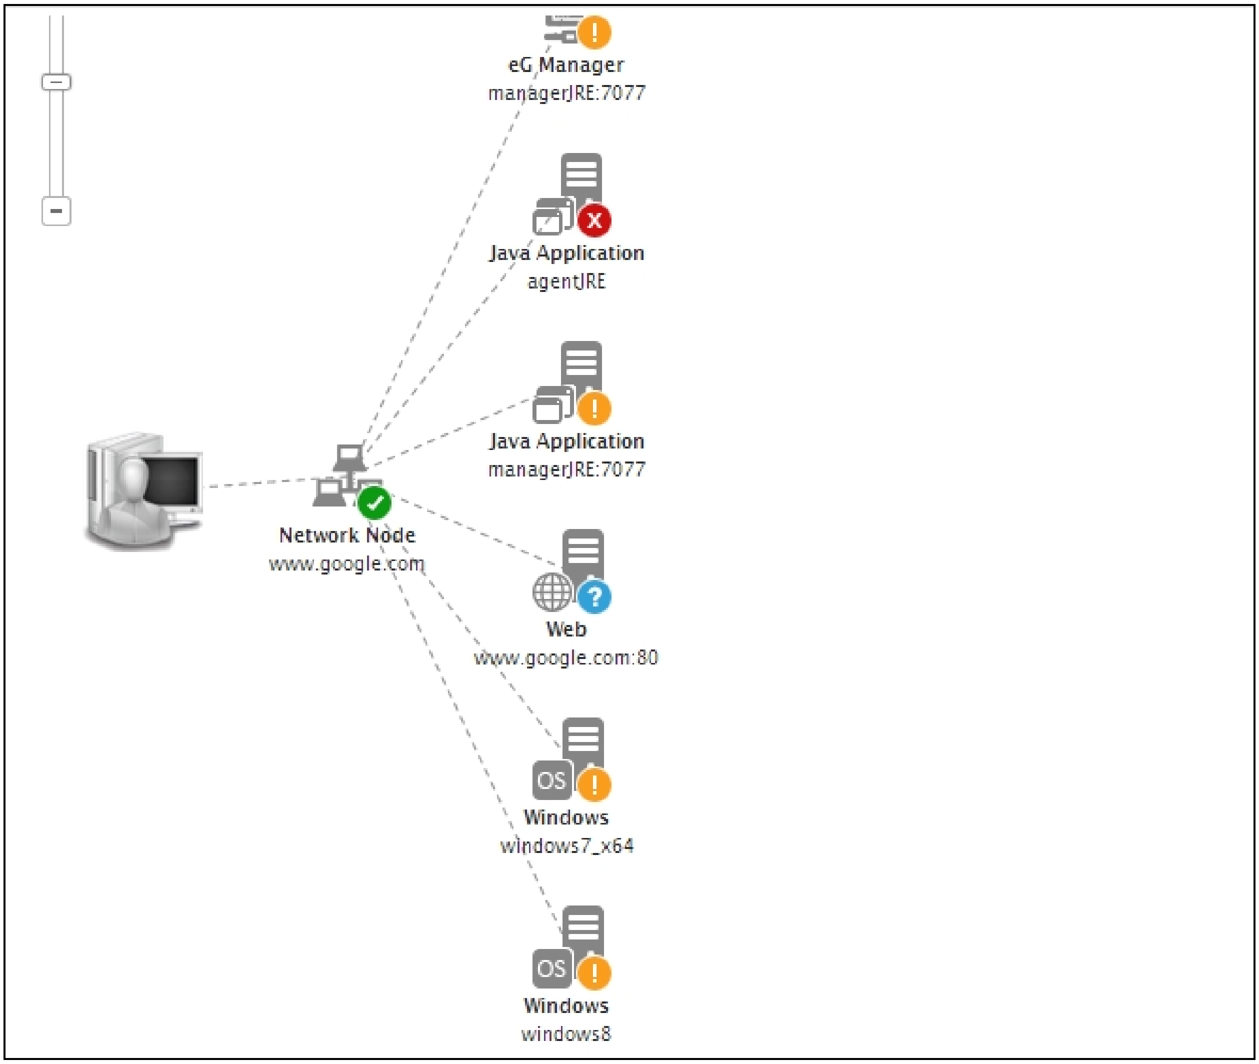

Figure 7.10: The topology of a segment comprising of a group

Clicking on the **GROUP** representation in Figure 7.10 will lead you to Figure 7.10, which displays all the components associated with the group and their current state. If there is an issue with the group, then the **GROUP LIST** will reveal which component in the group is causing the problem.

8

Monitoring Services and Service Groups

8.1 Monitoring Services

A service is a collection of infrastructure components that work together to perform a specific set of functions - e.g., a mobile payment gateway service, an online banking service, a web site, etc. eG Enterprise allows administrators to add one/more services for monitoring. The procedure to configure such services using the eG administrative interface has been described in the *Administering the eG Enterprise Suite* document. This section takes the example of a web site service named *infoway* to explain how eG Enterprise performs 'end-to-end service monitoring'.

Upon logging into the monitor interface, select the **Services** option from the **Groups** tile; this will lead you to the **SERVICES** page that lists the services that you are privileged to monitor and the current state of these services. If you login as *supermonitor*, then all the services that have been configured in your environment will be listed in the **SERVICES** page as depicted by Figure 8.1.

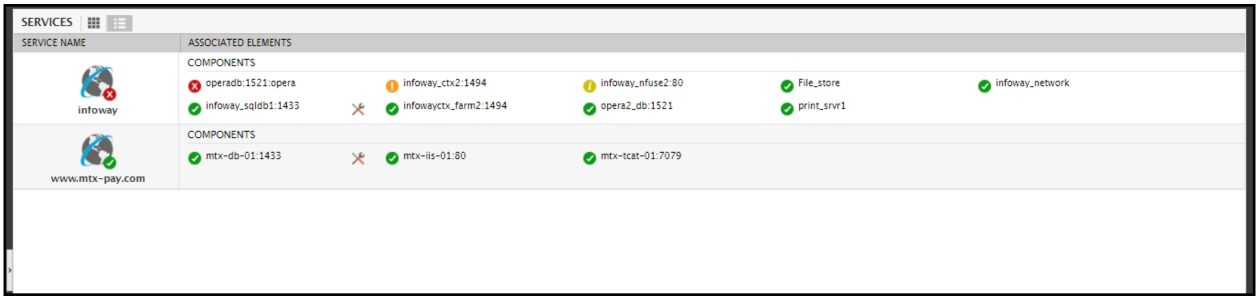

Figure 8.1: Viewing the list of configured services and their states

Note:

By default, against each service displayed in the **SERVICE** page, the top-10 **Components** associated with that service will be displayed. Typically, to identify the top-10 components, eG Enterprise automatically sorts all the components associated with a service on the basis of their current state, arranges the sorted list in the alphabetical order of the component names, and picks the first 10 components of this list.

If you want more number of components to be displayed against each service, do the following:

- Login to the eG administrative interface.
- Click on the icon available in the **Admin** tab. Then, select the **Monitor**option in the **Settings** tile.
- Click the **Other Settings** sub-node under the **General** node in the tree-structure in the **MONITOR SETTINGS** page.

- In the right panel, modify the default value 10 that is displayed in the **Components count in segment/service/zone list** text box.
- Click the **Update** button to save the changes.

Note:

By default, only the components associated with a service will be displayed in the **SERVICE** page. If you want the segments associated with the service also to be displayed, then, do the following:

- Login to the eG administrative interface.
- Click on the icon available in the **Admin** tab**.** Then, select the **Monitor**option in the **Settings** tile**.**
- Click the **Other Settings** sub-node under the **General** node in the tree-structure in the **MONITOR SETTINGS** page.
- In the right panel, set the **Show segment(s) in service list** flag to **Yes**.
- Click the **Update** button to save the changes.

To make sure that **SERVICE LIST** page does not display the list of components associated with a service, set the **Show component(s) in service list** flag in the **OTHER DISPLAY SETTINGS** panel of the **MONITOR SETTINGS** page to **No**.

If both the segment list and component list are disabled, then the **SERVICE LIST** page will only display a vertical list of services and their current state.

If you click on the **>** button to the left of the **SERVICES** list, as indicated by Figure 8.1, a **Services** tree will appear (see Figure 8.2).

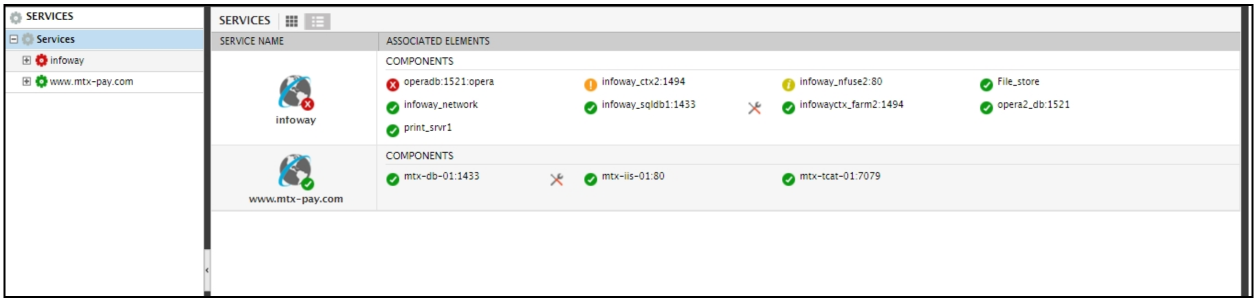

Figure 8.2: The Services tree

Every service configured for monitoring in the eG Enterprise system will appear as the nodes of the **Services** tree. Expanding a service node will list as its sub-nodes all those components that are engaged in the delivery of that service; the state of these components will also be revealed (see Figure 8.3).

Note:

By default, the services listed under the **Services** node will be sorted in the order of the service state. Likewise the components that are listed under each service will also be sorted in the basis of their current state.

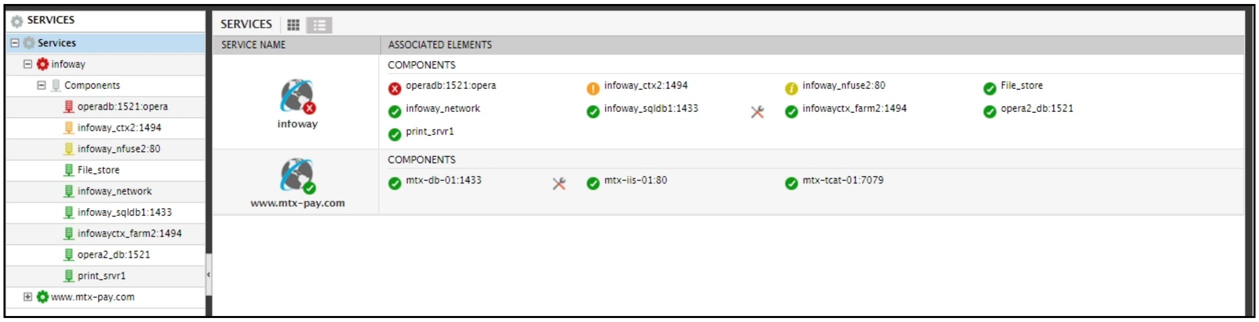

Figure 8.3: Expanding a service node in the Services tree

At any give point in time, you can click on the **<** button indicated by Figure 8.3 to hide the **Services** tree.

Going back to the **SERVICES** list of Figure 8.1, you can see that a service named *infoway* is currently in a **Critical** state. To know what **Critical** problem this service encountered and why, click **infoway** in Figure 8.1.

This opens a **Transactions** tab page (see Figure 8.4). The **Transactions** tab page is applicable only for web site services. Since this tab page appears as soon as the *infoway* service in our example is clicked, it is evident that the *infoway* service is a **web site service**. In this tab page, you can view the current state of the transactions that have been explicitly configured for monitoring for the *infoway* web site service. By closely tracking the requests to, the responsiveness of, and the errors encountered by every transaction to a web site, you can accurately ascertain which transaction is contributing to a slowdown in the *infoway* web site. From Figure 8.4, it is clear that the *UserLogin* and *ApplicationAccess* transactions to the *infoway* site are experiencing **Minor** issues. A look at the transaction metrics reveals that both these transactions are experiencing **Errors**.

If you now look at the **Graphs** section of the tab page, you will find that these **Errors** have persisted for the last hour. To know what is causing these persistent errors, click on either of the transactions – say, the **UserLogin** transaction.

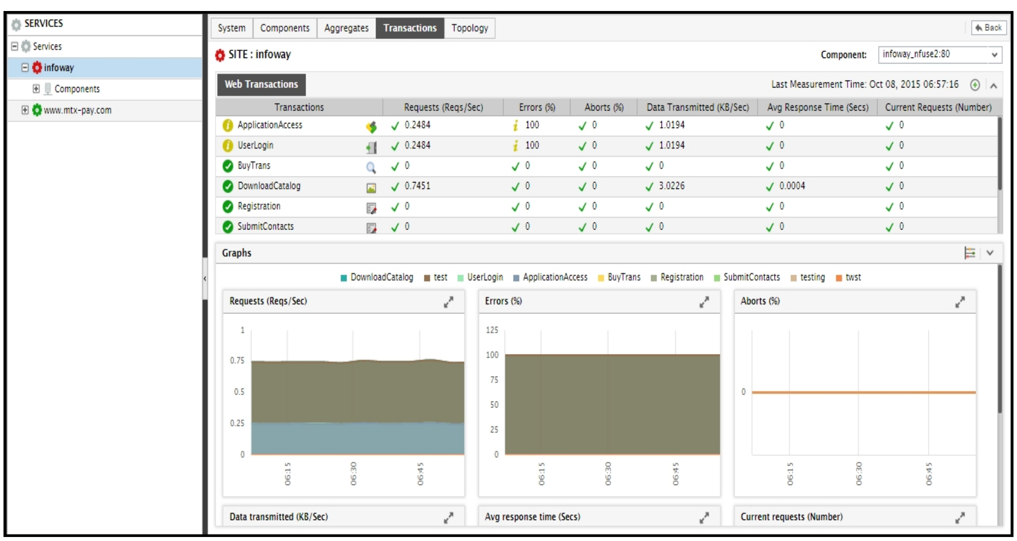

Figure 8.4: The Transactions tab page

This will lead you to the **Topology** tab page of Figure 8.5. This displays the topology of the *infoway* web site, indicating the components engaged in delivering the web site service and the physical/logical relationships that exist between the components.

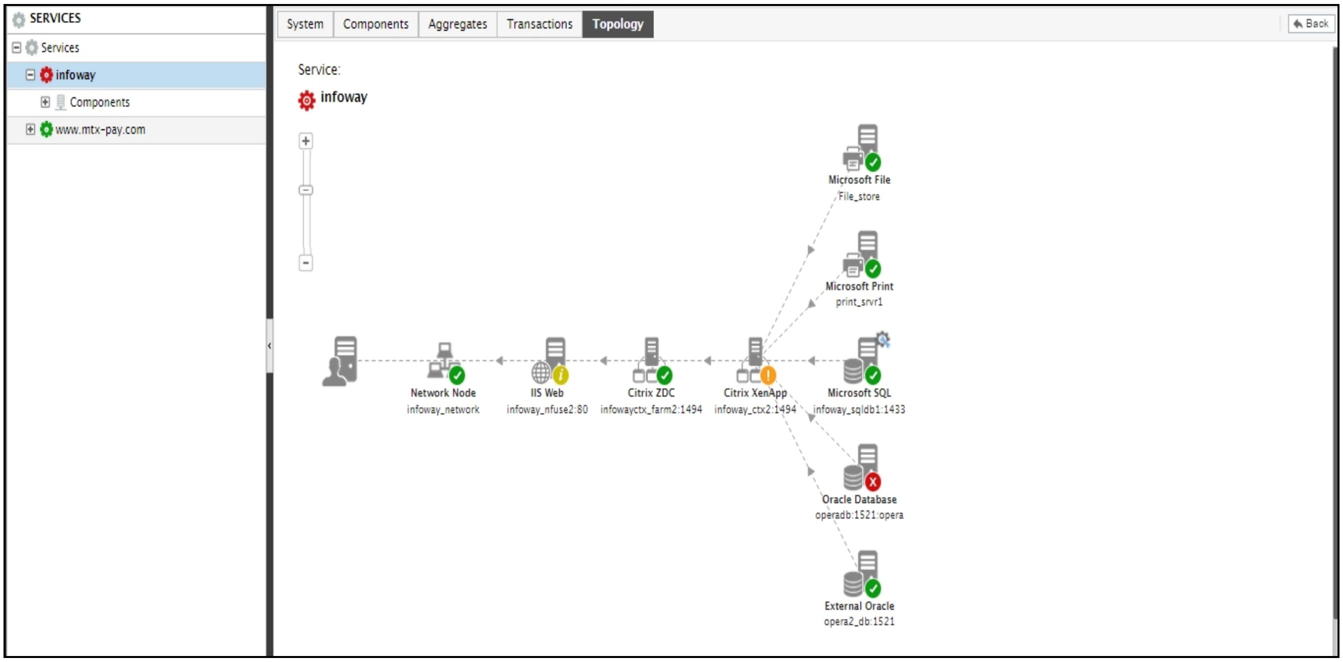

Figure 8.5: The topology of the infoway web site

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.

From the topology view of Figure 8.5, you can easily infer that the **infoway** web site employs a multi-tier architecture. The IIS web server (infoway_nfuse2:80) handles all incoming requests from web clients and forwards them on to a Citrix Zone Data Collector server (infoway_ctx_farm2:1494). The Zone Data Collector server then transmits the request to a Citrix XenApp Server (infoway_ctx2:1494). Back end Oracle and Microsoft SQL databases, a printer, and an Active Directory server are also used in the service delivery.

With the help of the color-coding on the components in the topology diagram, you can figure out that the **Oracle Database Server** is experiencing a **Critical** issue, the **Citrix XenApp** server is suffering a **Major** issue, and the **IIS web server** is having a **Minor** issue. If you look closely at the direction of the arrows used in the topology, you will be able to tell in which direction problems flow. In the case of our example, the direction of the arrows indicate that the **Critical** problem with the **Oracle database server**, has rippled and affected the performance of the **Citrix XenApp** server which depends on it. This problem has also travelled further down the topology to adversely impact the performance of the **IIS web server** that interacts with the XenApp server. eG's patented correlation engine has automatically correlated these issues and has accurately discovered that the root-cause of the problem with the **infoway** web site lies with the **Oracle Database** server. This is why, eG has intelligently assigned the highest problem priority to the problem with the **Oracle Database** server.

Since the performance of the XenApp server and IIS web server suffered as a result of the bottleck at the Oracle database server, eG has smartly downgraded the priority of the problems with the XenApp and IIS servers. This way, eG efficiently differentiates between the cause and effect of service-related issues.

To know what problem with the **Oracle Database** server is delaying the delivery of the **infoway** web site service, first click on the **System** tab page in Figure 8.4. Figure 8.6 will then appear.

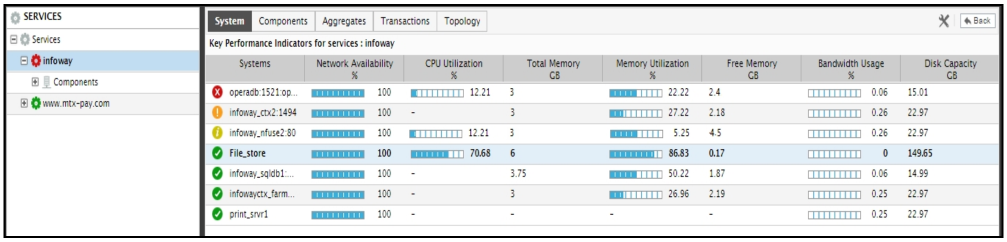

Figure 8.6: The System tab page

Performance issues suffered by an application host can ripple and affect the performance of the application itself, which in turn can delay the delivery of the dependent service(s). For a chosen service, the **System** tab page serves as a central console where you can quickly compare critical host-level metrics captured in real-time from across all components engaged in the delivery of that service. In the case of our example therefore, the **System** tab page reports real-time metrics revealing how the systems supporting the **infoway** web site are performing.

By default, the contents of the **System** tab page are sorted based on the state of the service components listed therein. If more that one component exists in the same state, then the components of that state will be sorted in alphabetical order. If need be, you can change the sort order. For example, if you wish to sort the components listed in the **System** tab page in the descending order of the values of their CPU usage, just click on the **CPU Utilization** label. Doing so, tags the **CPU Utilization** label with a down arrow icon – this icon indicates that the **System** tab page is currently sorted in the descending order of the CPU used by each component. To change the sort order to 'ascending', all you need to do is just click again on the **CPU Utilization** label or the down arrow icon. Similarly, you can sort the contents of the **System** tab page based on any column available in the table.

By default, the CPU, memory, disk space, and network usage of each of the systems engaged in delivering the **infoway** web site service can be tracked using this tab page. You can, if required, override the default measure list in the **System** tab page by adding more critical measures to the list or by removing one/more existing ones from the list. For this, do the following:

- Click on the icon provided near the **Back** button in Figure 8.7. In the **SETTINGS** window that appears (see Figure 8.7), select **System** from the **Tabs** flag.

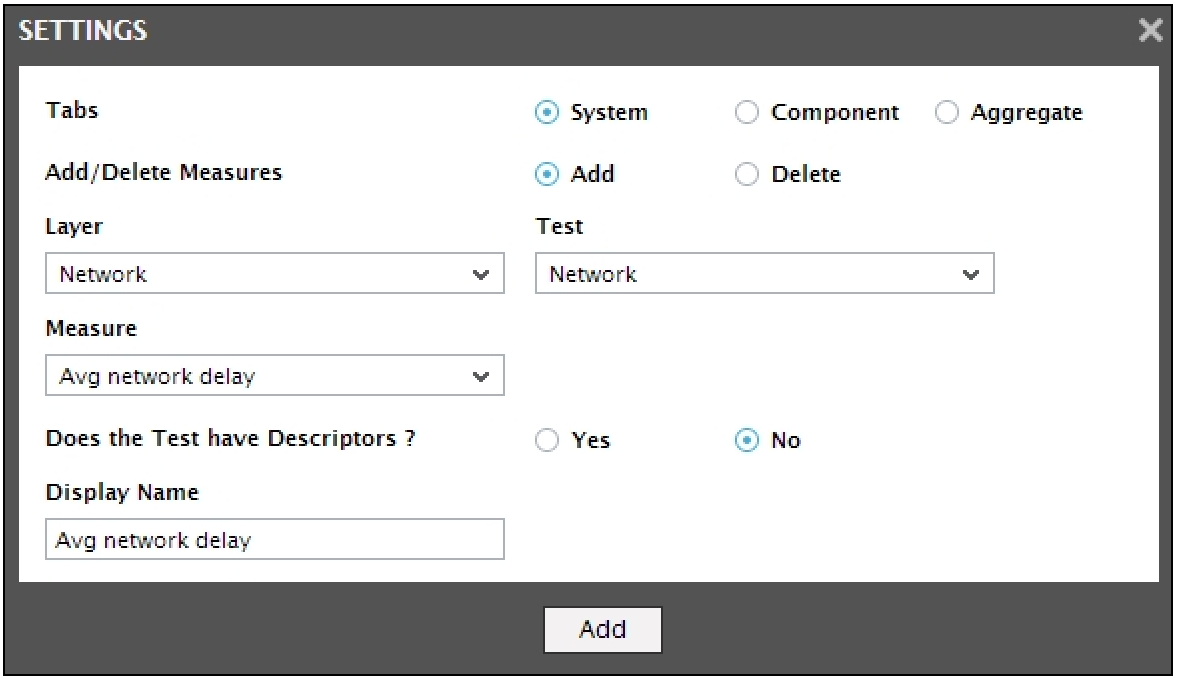

Figure 8.7: Selecting the System flag from the Configuration Settings Window

- To add more metrics to the **System** tab page, first, select the **Add** option from the **Add/Delete Measures** flag.
- Next, select the layer for which you wish to add the test from the **Layer** drop down list. Now, select the **Test** that reports the measure of your choice, pick the measure of your interest from the **Measures** list, provide a **Display** name for the measure, and click the **Add** button to add the chosen measure to the **Systems** tab page.
- If you want to delete one/more measures from this section, then, as soon as you choose the **Delete** option from the **Add/Delete Measures** flag, the **Test** drop down list will be populated with all the existing tests for which measures are displayed. Pick a **Test** and choose the **Measure** to be deleted from the **System** tab page.

  Note:

  While displaying values for descriptor-based measures in the **System** tab page, the eG Enterprise system does not display the actual values per descriptor. Instead, the solution computes theaverage or the total sum of values across descriptors and displays it in the corresponding measure column. For instance, for values reported as percentages, the solution computes the average value across descriptors. On the other hand, if the value is reported as a GB or MB, then the total sum of all the descriptor values of the component will be displayed against the component.

This centralized view of the health of all systems associated with the **infoway** web site helps in quickly determining whether any OS-level issues with the **Oracle Database Server** could be affecting service quality. From the **System** tab page of Figure 8.6, it is obvious that the Oracle database server host is in good health presently. So, could a serious performance snag at the application-level be responsible for the **Critical** issue with the **Oracle Database Server**? To find out, click the **Components** tab page in Figure 8.6. Figure 8.8 will then appear.

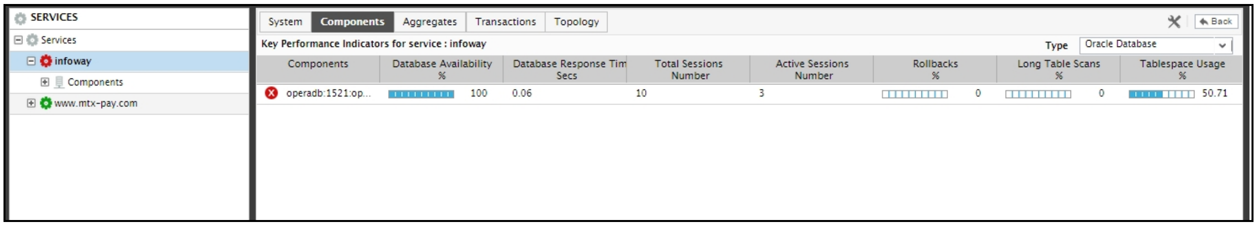

Figure 8.8: The Components tab page

The **Components** tab page provides insights into the performance of the applications that are engaged in service delivery - in other words, the tab page displays the real-time values of the application level metrics collected from each component associated with a service. Using this at-a-glance information, administrators can perform the following with ease:

- Oversee, by a mere glance, how well the components associated with the chosen service are performing;
- Easily analyze and detect abnormalities experienced by the misson-critical applications associated with the chosen service.

At any given point in time, you can view and analyze the application-level metrics related to the service components of a single type alone in the **Components** tab page. Use the **Type** list at the right, top corner of Figure 8.8 to select the component-type of interest of you. Once a component-type is chosen, all components of that type that are associated with the service in question will be listed. For each component of the chosen type, a set of pre-defined application-level metrics will be displayed. If required, you can override this default metrics list by adding more metrics for display in this tab page, or by removing one/more existing metrics. For this, do the following:

- Click on the icon provided near the **Back** button in Figure 8.8. In the **SETTINGS** window that appears (see Figure 8.9), select **Component** from the **Tabs** flag.

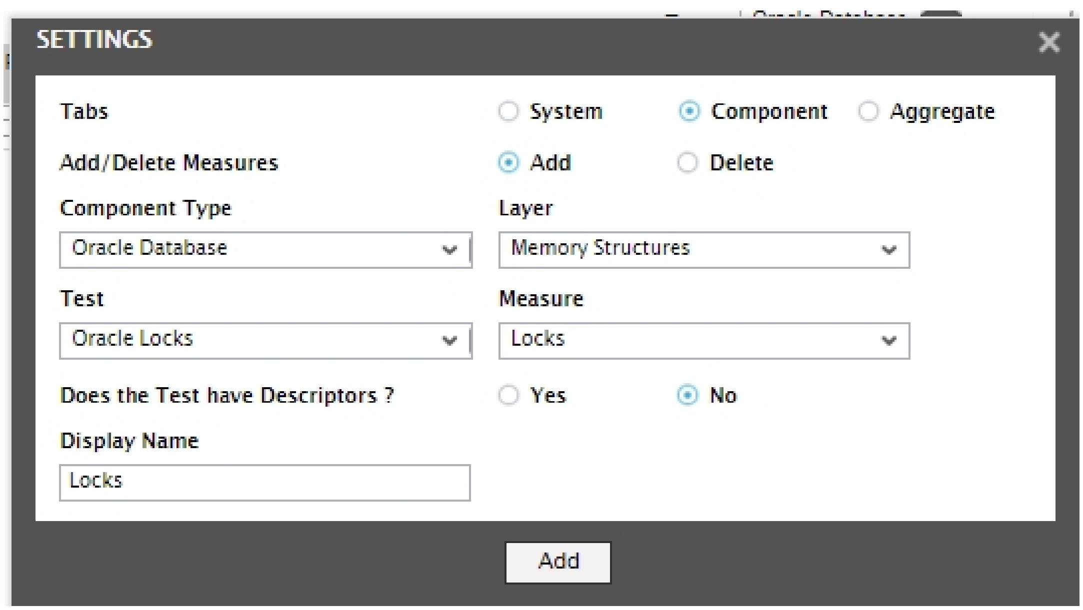

Figure 8.9: Selecting the Component flag from the Settings Window

- To add more metrics to the **Components** tab page, first, select the **Add** option from the **Add/Delete Measures** flag. Then, pick the **Component Type** to which the addition applies.
- Next, select the layer for which you wish to add the test from the **Layer** drop down list. Then, select the **Test** that reports the measure of your choice, pick the measure of interest from the **Measures** list, provide a **Display** name for the measure, and click the **Add** button to add the chosen measure to the **Components** tab page.
- If you want to delete one/more measures from this section, then, as soon as you choose the **Delete** option from the **Add/Delete Measures** flag, the **Test** drop down list will be populated with all the existing tests for which measures are displayed. Pick a test and choose a **Measure** of your interest to delete from the **Components** tab page.

  Note:

  While displaying values for descriptor-based measures in the **Components** tab page, the eG Enterprise system does not display the actual values per descriptor. Instead, the solution computes theaverage or the total sum of values across descriptors and displays it in the corresponding measure column. For instance, for values reported as percentages, the solution computes the average value across descriptors. On the other hand, if the value is reported as a GB or MB, then the total sum of all the descriptor values of the component will be displayed against the component.

In the case of our example by default, **Oracle Database** is chosen as the **Type** in the **Components** tab page (see Figure 8.8). As a result, the **Components** tab page reports the current availability, responsiveness, average tablespace usage, session load, and other key operational metrics of the problematic Oracle database server alone. From these real-time metrics, it is clear that the Oracle database server is available and is responding quickly to requests. Expensive operations such as long table scans and rollbacks are also non-existent on the server. The load on the server also appears minimal. The disconcerting factor however is the **Tablespace Usage**, which is over 50%. Could this be the root-cause of the Critical issue with the Oracle database server? If so, which tablespace is being used excessively? For the answers, click on the Oracle database server component in the **Components** tab page. Figure 8.10 will then appear, revealing the problematic layer, test, and measure of the Oracle database server in our example.

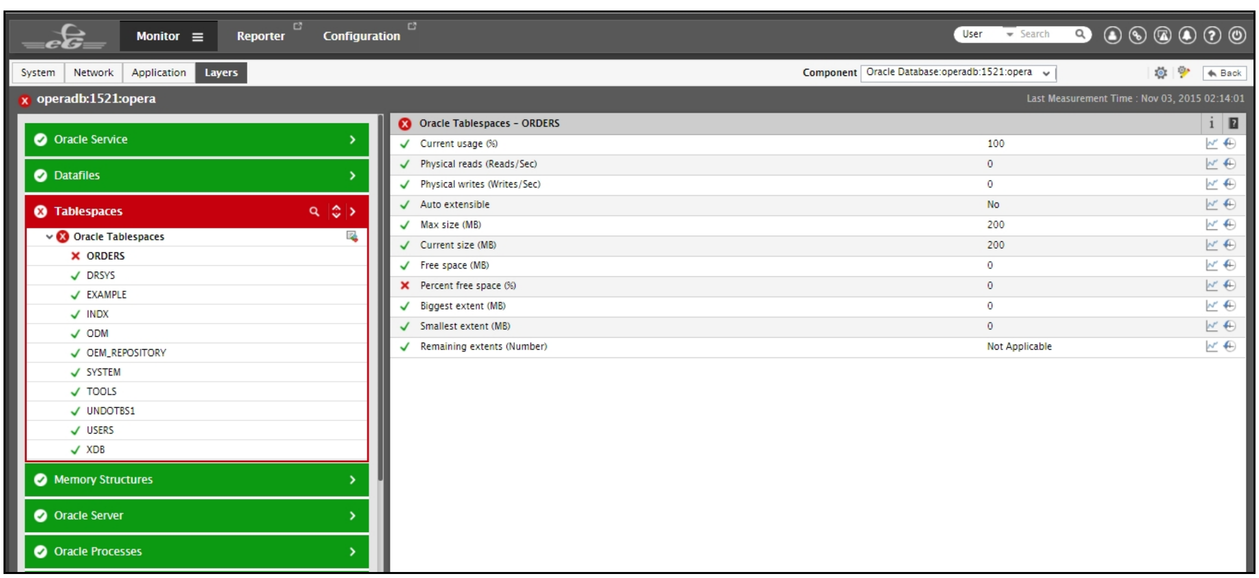

Figure 8.10: The problem layer, test, and measure of the Oracle database server

Figure 8.10 clearly indicates that the **Critical** issue is because the **ORDERS** tablespace of the Oracle database server is being over-utilized, and is hence running out of free space. With that, we can conclude that the lack of free space in the **ORDERS** tablespace is the reason why a **Critical** issue occurred on the Oracle database server.

eG Enterprise saves you the trouble of navigating the **System** and **Component** tab pages to determine what is ailing the Oracle database server. Instead, you can simply click on the Oracle database server component in the **Critical** state in the **Topology** tab page (of Figure 8.5) itself to open Figure 8.10.

The **Topology** tab page (see Figure 8.5) also reveals that it is this **Critical** database server issue that has caused the **Major** issue with the XenApp server, the **Minor** issue with the IIS web server, and has ultimately stalled the delivery of the **infoway** web site service. To know how, let us first figure out what the **Major** problem with the Citrix XenApp server is. For that, switch to the **Topology** tab page and click the **Citrix XenApp** server component in the **Major** state. Figure 8.11 will then appear.

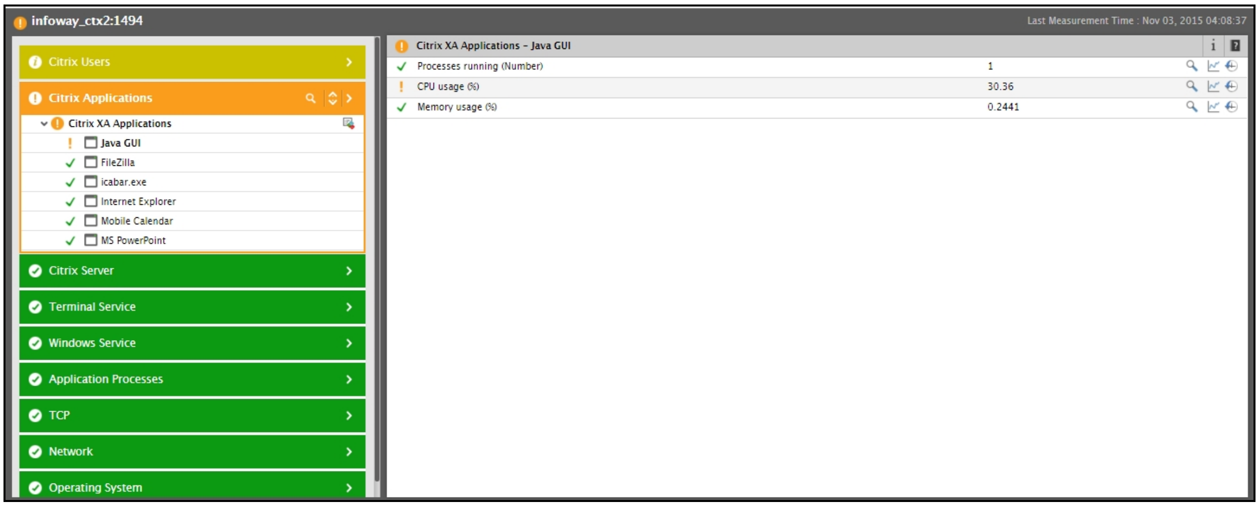

Figure 8.11: The problem layer, test, and measure of the Citrix XenApp server

Figure 8.11 reveals that the **Major** problem is affecting the **Citrix Applications** layer, and the problem is owing to a high CPU usage of the java.exe application executing on the Citrix XenApp server. In this case, it turns out that the **UserLogin** transaction (see Figure 8.4), which registered errors, is being handled by a Java application that is hosted on the XenApp server. When the tablespace is full, the Java application keeps retrying to add a new user record, hence causing a **Major** CPU bottleneck on the XenApp server. Because record insertion failed, the **UserLogin** transaction also failed, resulting in a **Minor** issue on the IIS web server that is hosting the **infoway** web site. With that, we can conclude that the lack of free tablespace is the root-cause of the **Critical** slowdown in the delivery of the **infoway** web site service.

If one/more aggregate components have been created using the components of a service, then an additional **Aggregates** tab page will appear when monitoring that service. In the case of our **infoway** web site service too, you can see the **Aggregates** tab page. Clicking on that tab page will reveal Figure 8.12.

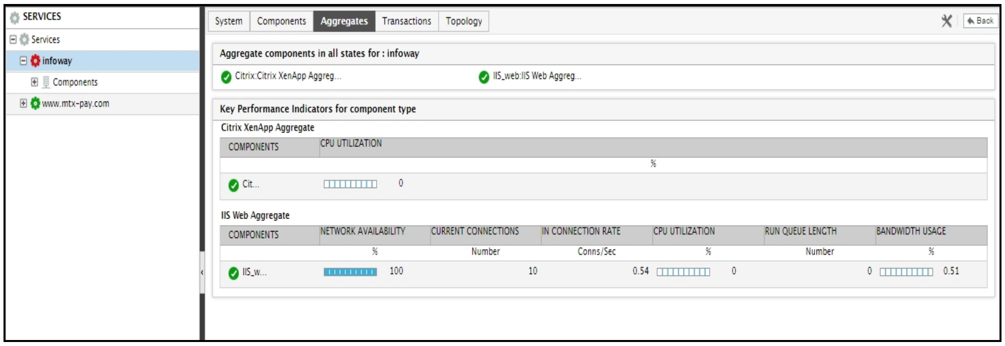

Figure 8.12: The Aggregates tab page

From Figure 8.12, you can infer that the Citrix XenApp server and IIS web server components that are associated with the *infoway* web site are also part of a *Citrix XenApp Aggregate* component and an *IIS Web* aggregate component, respectively. The **Aggregates** tab page also reveals the names of these aggregate components, the current state of these aggregates, and key aggregate metrics reported by each of these components. You can also override the default metrics list by adding more measures to this tab page or removing one/more existing measures. For that, use the **SETTINGS** window that appears when the icon is clicked.

A quick glance at this **Aggregates** tab page can reveal abnormalities related to aggregate components and the metrics responsible for the same. Clicking on an aggregate component here will allow you to zoom into the layer model of that component.

8.2 Monitoring Service Groups

Once you are in the eG monitoring interface, you will find that the **Infrastructure Health** section of the eG monitoring console displays the number of configured service groups and the current state of these groups. Clicking on the **Service Group** label here will lead you to Figure 8.13, which lists all the configured service groups and their current state. Against every service group, the services included in that group and the current state of each of the services will be available so that, you can instantly identify the services that are in responsible for the abnormal state of the service group. Clicking on a service here will allow you to preview the topology of that service, using which you can identify the root-cause of the problems associated with that service.

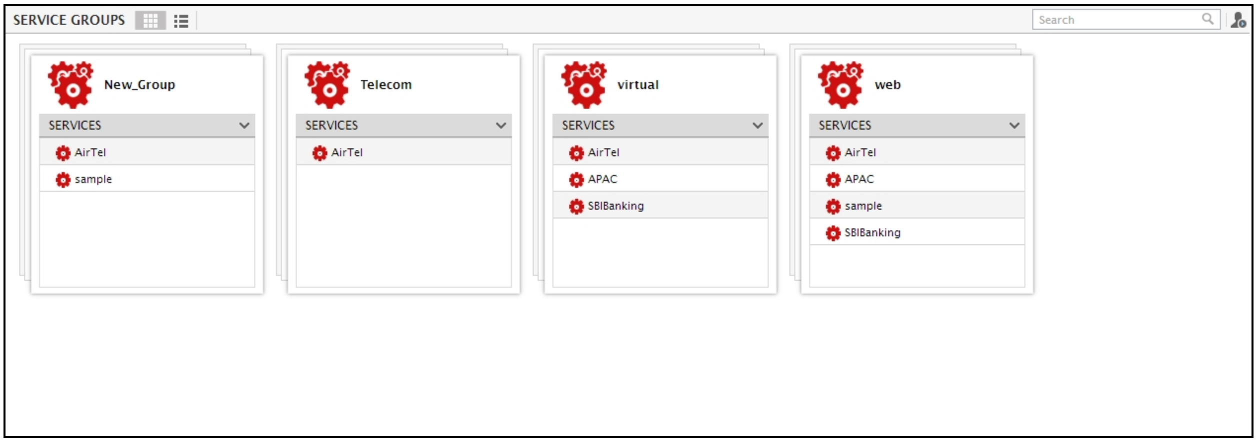

Figure 8.13: List of configured service groups and their current state

8.3 Service Health

When monitoring the performance of their mission-critical business services, service managers may want to know what percentage of service components are currently in an abnormal state. This knowledge will not only enable them to quickly identify the services that are experiencing performance degradations, but will also help them swiftly assess the extent of the damage! The optional **SERVICE HEALTH** page of the eG monitoring console imparts this knowledge to service managers. By default, this page is not available in the eG monitoring console. To enable this page, do the following:

- Edit the **eg_ui.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory).
- Set the **enable_service_health** parameter in the file to **true**. By default, this is set to **false**.
- Finally, save the file.

Once this is done, then, you would be able to access the **SERVICE HEALTH** page by clicking the icon amenu sequence in the eG monitoring console.

This uses a wide variety of chart types, namely - gauge chart, pie chart, HLinear gauge chart, and column chart - to indicate the state of components engaged in the delivery of each managed business service.

By default, the chart that corresponds to a service denotes the following:

- the percentage of service components that are currently in a **Critical** state, indicated by the color Red;
- the collective percentage of components that are in the non-critical states (i.e., Normal, Major,Minor, and Unknown), indicated by the color Green;

This default setting is governed by the **show_states** parameter in the **[SERVICE_HEALTH]** section of the

eg_ui.ini file. By default, the **show_states** parameter is set as follows:

```
[SERVICE_HEALTH]
show_states=HIGH,GOOD
```

Since the **show_states** parameter is set to **HIGH** and **GOOD** by default, each service chart indicates the percentage of service components in the **CRITICAL** state and **GOOD** states (by default) only. While **HIGH** denotes **CRITICAL**, the **GOOD** here is the collective term used to denote all non-critical states, such as Normal, Unknown, Major, and Minor.

You can override this default setting so that the service charts indicate more or less number of 'distinct' states. For instance, to make sure that each of the service charts 'distinctly denote' the percentage of components in the Critical, Major, Minor, Normal, and Unknown states, your **show_states** parameter setting should be changed as follows:

```
show_states=HIGH,INTERMEDIATE,LOW,UNKNOWN,GOOD
```

Also under each service chart, the list of service components in the abnormal state (if any) is displayed. While the chart itself can alert you to any slowdown that is currently experienced by the service, a quick look below the chart will help you instantly identify those service components that could be contributing to the service slowdown. By default, this service component list will display the top-5 abnormal components involved in delivering that service. To override this default setting, do the following:

- Edit the **eg_ui.ini** file.
- Against the **show_top** parameter in the **[SERVICE_HEALTH]** section, specify any number of your choice. By default, this is set to 5.
- Finally, save the file.

Moreover, as mentioned already, the **SERVICE HEALTH** page uses any or all of the following types of charts to indicate the current state of configured services:

- Angular Gauge chart
- HLinear Gauge chart
- Pie chart
- Column chart

If you choose to use all the above-mentioned chart types in the **SERVICE HEALTH** page, then set the **show_charts** parameter in the **[SERVICE_HEALTH]** section of the **eg_ui.ini** file to **all**, as indicated below:

```
show_charts=all
```

If the **show_charts** parameter is set to **all**, you then need to indicate the chart types to be used and the sequence in which they are to be used. For this, provide a comma-separated list of chart types against the **chart_list** parameter in the **[SERVICE_HEALTH]** section of the **eg_ui.ini** file:

```
chart_list=AngularGauge:bottom,HLinearGauge:bottom,Pie:right,Column:right
```

At runtime, the charts will be displayed in the same sequence in which they appear in the comma-separated **chart_list** above.

Each entry against the **chart_list** parameter should be of the following format:

```
chart_list=<charttype:position_of_legend>,....
```

For instance, if an *Angular Gauge* chart is one of the charts you wish to use in the **SERVICE HEALTH** page, and you want the legend of the chart to appear below the chart, then your **chart_list** specification will be:

```
chart_list=AngularGauge:bottom
```

Sometimes, you may not want to use a variety of charts in the **SERVICE HEALTH** page, and instead use a single chart type across the entire interface. In this case, you will need to alter the **show_charts** parameter as indicated below:

```
show_charts=<charttype:position_of_legend>
```

For instance, if you want the health of all services to be depicted using only pie charts, and you want the legend to be displayed to the right of every pie chart, your **show_charts** specification will be:

```
show_charts=Pie:right
```

Clicking on the service chart of a non-web-site service in the SERVICE HEALTH page here will lead you to the **Topology** tab page of the **Service Dashboard**, where you can view the state of all components engaged in the delivery of that service and the inter-dependencies they share. Clicking on the service chart of a web site service on the other hand, will list the transactions configured for that web site, the current state of each transaction, and the metrics that are collected for every transaction.

9

Monitoring 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 know how to configure zones, refer to Chapter 2 of this manual.

Any number of zones can be configured using eG Enterprise. To quickly determine the state of the configured zones, you can login to the eG monitor interface and just move your mouse pointer over the **Zones** option in Figure 9.1. Then, pick the **List** option from the **Zones** menu.

Choosing the **List** option opens Figure 9.2, where those zones that have been assigned to you will be listed, along with their states; if you login as supermonitor, then all the fully-configured zones will be displayed in Figure 9.2.

Note:

If you have not configured locations (using the eG map interface) for all the managed zones, then clicking on the **Zones**option in Figure 9.1 itself will invoke the zone list of Figure 9.1.

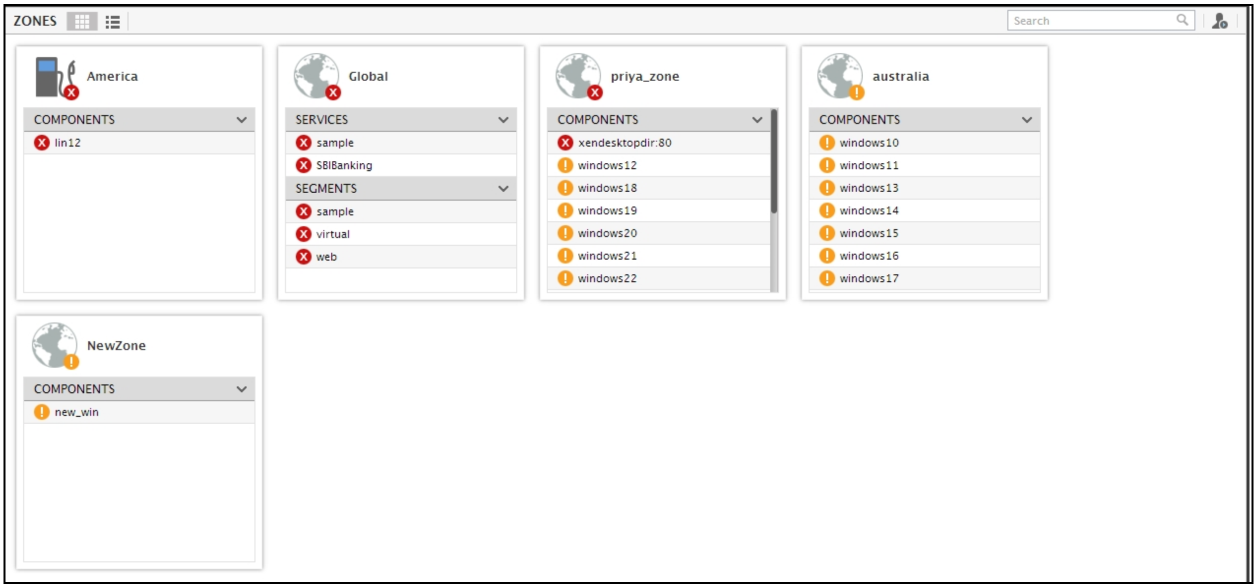

Figure 9.1: The state of all the zones being monitored

For each configured zone, the **ZONE LIST** page displays the state of elements (segments/components/services/other zones) within that zone. This way, you can quickly identify those infrastructure elements that are responsible for problems with a zone.

Note:

By default, against each zone displayed in the **ZONE LIST** page, the top-10 **Components** included in that zone will be displayed. Typically, to identify the top-10 components, eG Enterprise automatically sorts all the components included in the zone on the basis of their current state, arranges the sorted list in the alphabetical order of the component names, and picks the first 10 components of this list.

If you want more number of components to be displayed against each zone, do the following:

- Login to the eG administrative interface.
- Click on the icon available in the **Admin** tab**.** Then, select the **Monitor** option from the **Settings** tile**.**
- Click the **Other Settings** sub-node under the **General** node in the tree-structure in the **MONITOR SETTINGS** page.
- In the right panel, modify the default value 10 that is displayed in the **Components count in segment/service/zone list** text box.
- Click the **Update** button to save the changes.

Clicking on the right-arrow button alongside the **ZONE LIST** page will reveal a tree-structure in the left panel (see Figure 9.1). This tree-structure consists of a **Zones** node, which displays all the configured zones as its sub-nodes. Besides the names of the zones, these sub-nodes also indicate the current state of each zone.

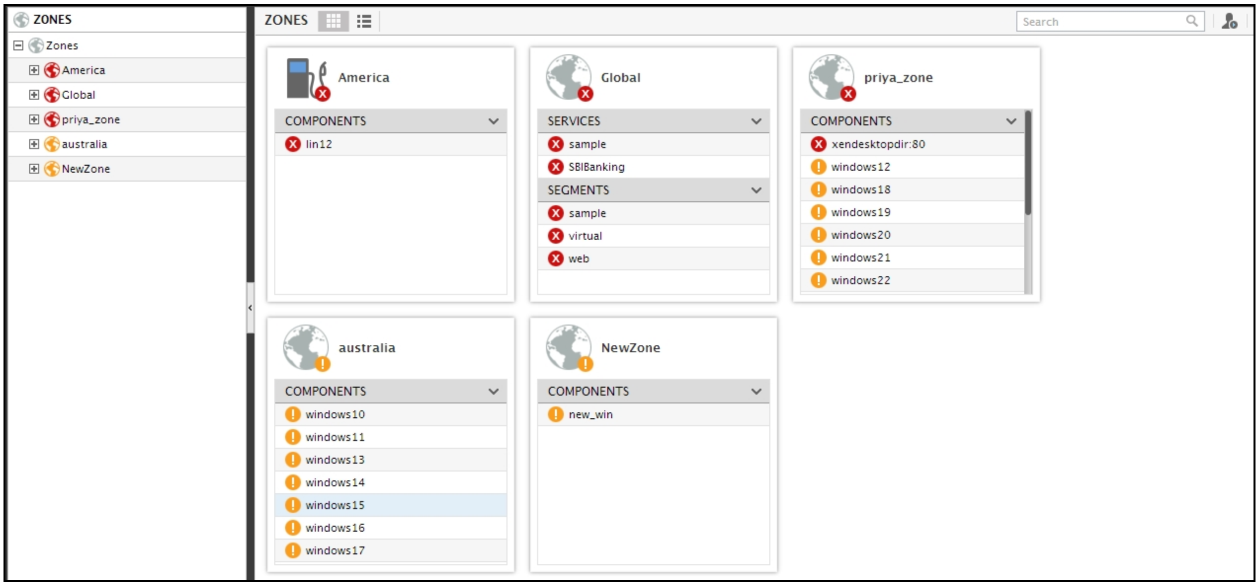

Figure 9.2: The Zone Dashboard

If you expand a node representing a zone in the tree, you will find sub-nodes representing each type of infrastructure element that has been included in the zone. For instance, if one/more services and segments have been added to a zone, then expanding that zone's node will reveal a **Services** and a **Segments** sub-node. To know which segments and services have been added to the zone, expand the **Segments** and **Services** sub-nodes, respectively.

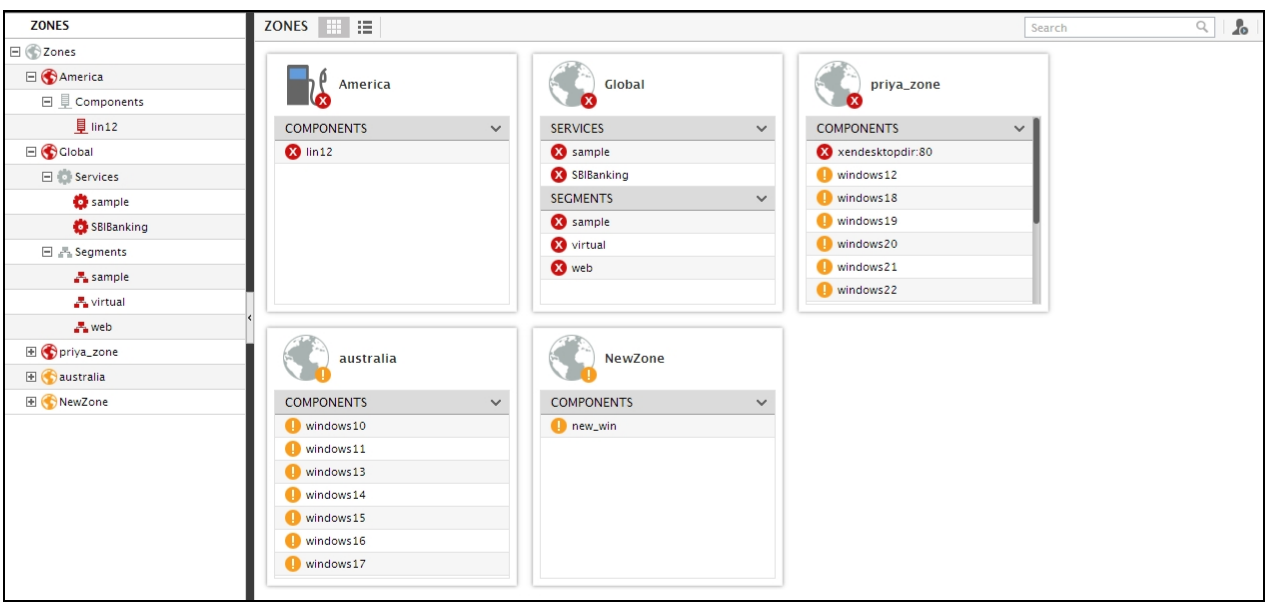

Figure 9.3: Expanding the node representing a particular zone

Likewise, if a zone consists of sub-zones and independent components, then, expanding such a zone's node will reveal sub-nodes named **Zones** and **Components**, respectively. Here again, you can expand the **Zones** and **Components** sub-nodes to figure out which other zones and components have been included in the zone. If the zone includes aggregate components, then the node representing that zone will host an **Aggregates** sub-node, which when expanded, will list all the aggregate components that are part of that zone and their current state.

Click on a particular segment in the tree to view the **Topology** of that segment, using which you can determine the root-cause of performance issues in the segment. In the same way, you can click on a particular service in the tree to view the **Topology** of that service and accurately diagnose the source of any slowdowns that may be experienced by that service.

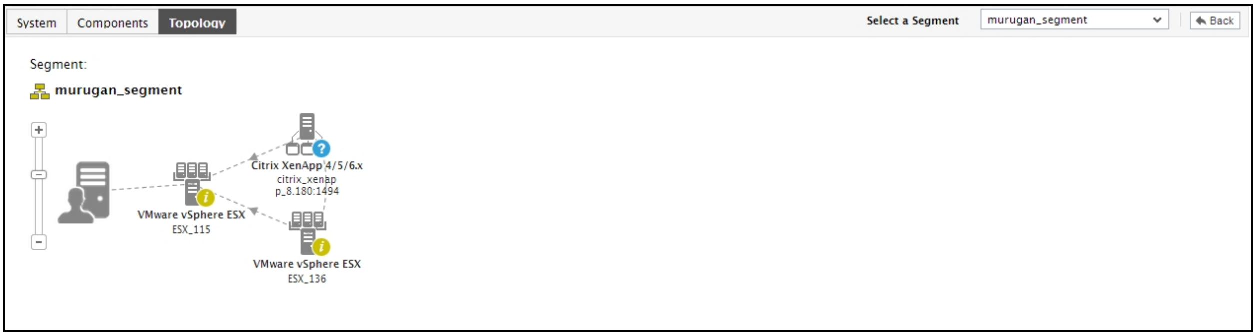

Figure 9.4: The Segment Topology tab page that appears when a segment that is part of a zone is clicked in the Zones tree

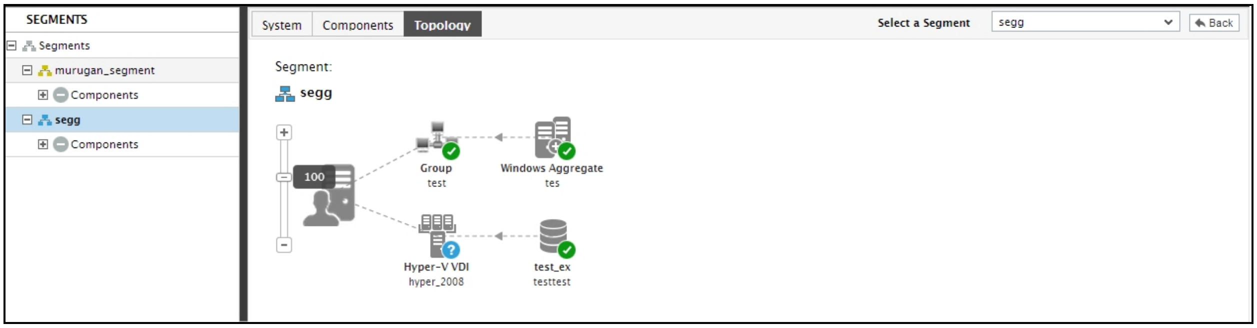

Figure 9.5: The Service Topology tab page that appears when a specific service that is part of a zone is clicked in the Zones tree

When you click on a node representing a (parent) zone in the tree - i.e., if you click on any of the sub-nodes of the global **Zones** node in the tree - the right panel will change to display three tab pages, namely - the **Systems**, **Components**, and **Details** tab pages. The sections that follow will discuss each of these tab pages in detail.

9.1 The Systems Tab Page

By default, selecting a zone node from the tree opens the **Systems** tab page. This tab page serves as a single, central interface that allows administrators to ascertain, from just a glance, the current operating system-level health of every component that is part of the chosen zone. The values of key host-level parameters are captured in real-time from all the components that belong to the chosen zone and are displayed here, so that administrators are instantly alerted to issues related to the network connection/traffic, TCP connectivity, and resource usage of every component.

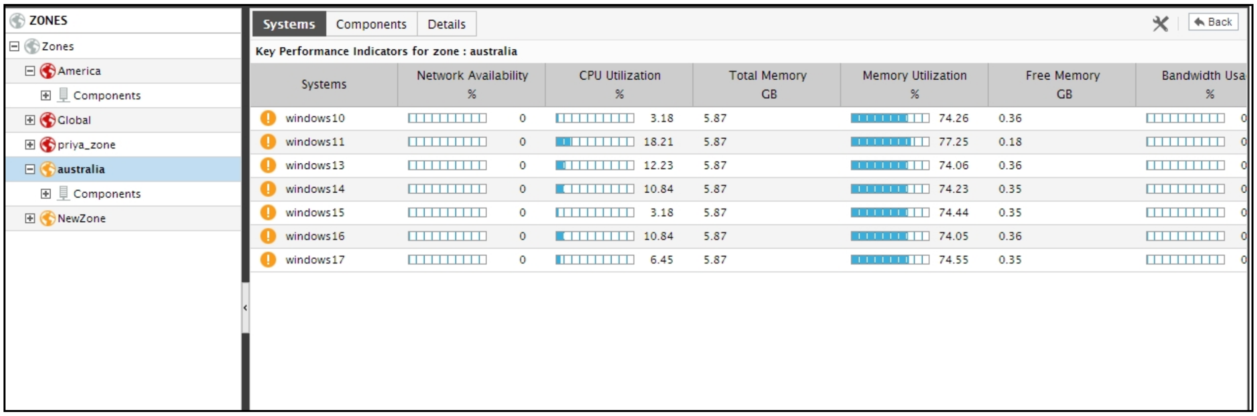

Figure 9.6: The Systems tab page revealing operating system-level health of every component that is part of a zone

Clicking on a **System** in the right panel of Figure 9.6 will lead you to the layer model page, where you can view the exact layer where a problem has occurred, the test that reported the problem, and the problematic measure.

By default, the contents of the **Systems** tab page are sorted based on the state of the zone components listed therein. If more that one component exists in the same state, then the components of that state will be sorted in alphabetical order. If need be, you can change the sort order. For example, if you wish to sort the components listed in the **Systems** tab page in the descending order of the values of their Disk Usage, just click on the **Disk Usage** label. Doing so, tags the **Disk Usage** label with a down arrow icon – this icon indicates that the **Systems** tab page is currently sorted in the descending order of the total disk space used by each component. To change the sort order to 'ascending', all you need to do is just click again on the **Disk Usage** label or the down arrow icon. Similarly, you can sort the **Systems** tab page based on any column available in the table.

If you are an **Admin** user, then, you can override the default measure list in the **Systems** tab page by adding more critical measures to the list or by removing one/more existing ones from the list. For this, do the following:

- Click on the  icon provided near the **Back** button in Figure 9.7. In the **Settings** window that appears (see Figure 9.6), select **Systems** from the **Tabs** flag.

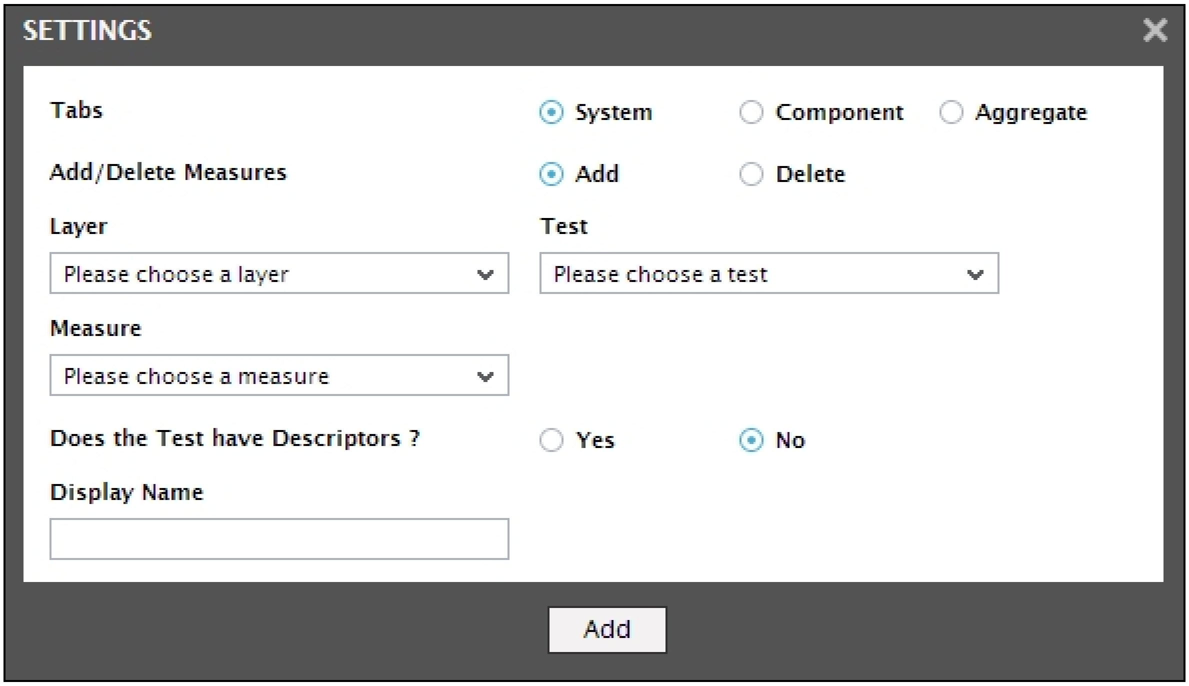

Figure 9.7: Adding a measure to the Systems tab page of the Zone dashboard

- To add more metrics to the **Systems** tab page, first, select the **Add** option from the **Add/Delete Measures** flag.
- Next, select the layer for which you wish to add the test from the **Layer** drop down list. Now, select the **Test** that reports the measure of your choice, pick the measure of your interest from the **Measures** list, provide a **Display** name for the measure, and click the **Add** button to add the chosen measure to the **Systems** tab page.
- If you want to delete one/more measures from this section, then, as soon as you choose the **Delete** option from the **Add/Delete Measures** flag, the **Test** drop down list will be populated with all the existing tests for which measures are displayed. Pick a **Test** and choose a **Measure** of your interest to delete from the **Systems** tab page.

  Note:

  - Only a user who is assigned the **Admin** role is allowed to customize the zone dashboard by clicking on the icon.
  - While displaying values for descriptor-based measures in the **Systems** tab page, the eG Enterprise system does not display the actual values per descriptor. Instead, the solution computes theaverage or the total sum of values across descriptors and displays it in the corresponding measure column. For instance, for values reported as percentages, the solution computes the average value across descriptors. On the other hand, if the value is reported as a GB or MB, then the total sum of all the descriptor values of the component will be displayed against the component.

9.2 The Components Tab Page

The **Components** tab page provides insights into the performance of the applications that are part of the chosen zone. For each application that has been added to a zone, users can configure key application-level metrics that are to be captured in real-time and displayed in the **Components** tab page. This way, users can ensure that they receive a heads-up on common, yet critical operational issues encountered by an application that is part of a zone, without having to go to the layer model page of that application for this purpose.

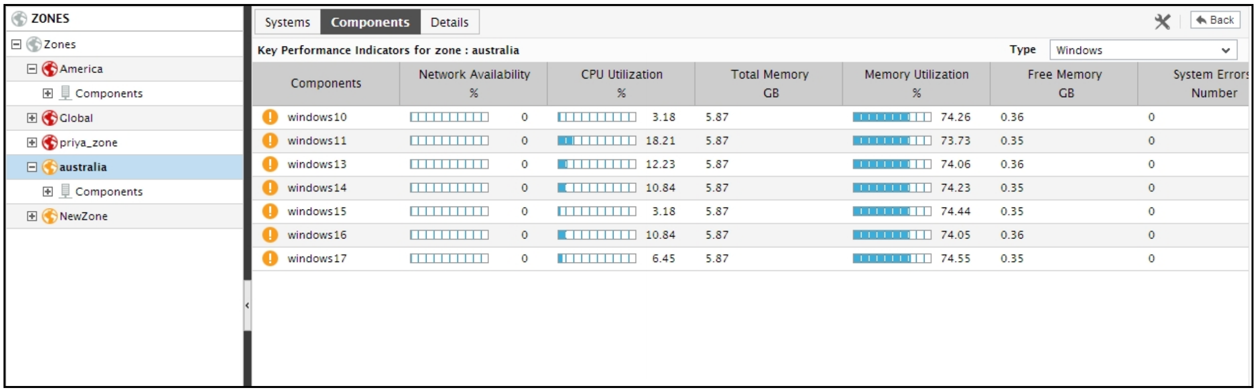

Figure 9.8: The Components tab page of the Zone Dashboard

Since metrics are configured per application, the application level metrics displayed in this tab page will differ based on the type of the component. A **Type** drop-down list as shown in Figure 9.8 will be populated with all the component types associated with the chosen zone. You can pick the component type of interest to you from the **Type** drop-down list to view the user-configured application level metrics of all zone components of that type.

Note:

The **Type** drop down list will be sorted based on the current state of the zone components of each type.

By default, the components listed in the **Components** tab page will be sorted in the order of their state - starting from the critical to the normal. If more that one component exists in the same state for the chosen component type, then the components of that state will be sorted in alphabetical order. If need be, you can change the sort order based on the application level metrics that are displayed against each component. For example, if you wish to view the sort the *Oracle Database* server list in the **Components** tab page in the descending order of the number of **Table Space usage**, just click on the **Table Space usage** label. Doing so, tags the **Table Space usage** label with a down arrow icon – this icon indicates that the **Components** tab page is currently sorted in the descending order of the table space usage of each zone component of type *Oracle Database*. To change the sort order to 'ascending', all you need to do is just click again on the **Table Space usage** label or the down arrow icon. Similarly, you can sort the **Components** tab page based on any column available in the table.

Clicking on a component here will lead you to the layer model page of that component, where the problem layer, test, and measures are revealed.

On the other hand, if no metrics have been configured for the **Type** chosen, then a message to that effect will appear.

To modify the measure-list associated with a component type, do the following:

- Click on the icon provided near the **Back** button in Figure 9.8. In the **Settings** window that appears (see Figure 9.8), select **Components** from the **Tabs** flag.

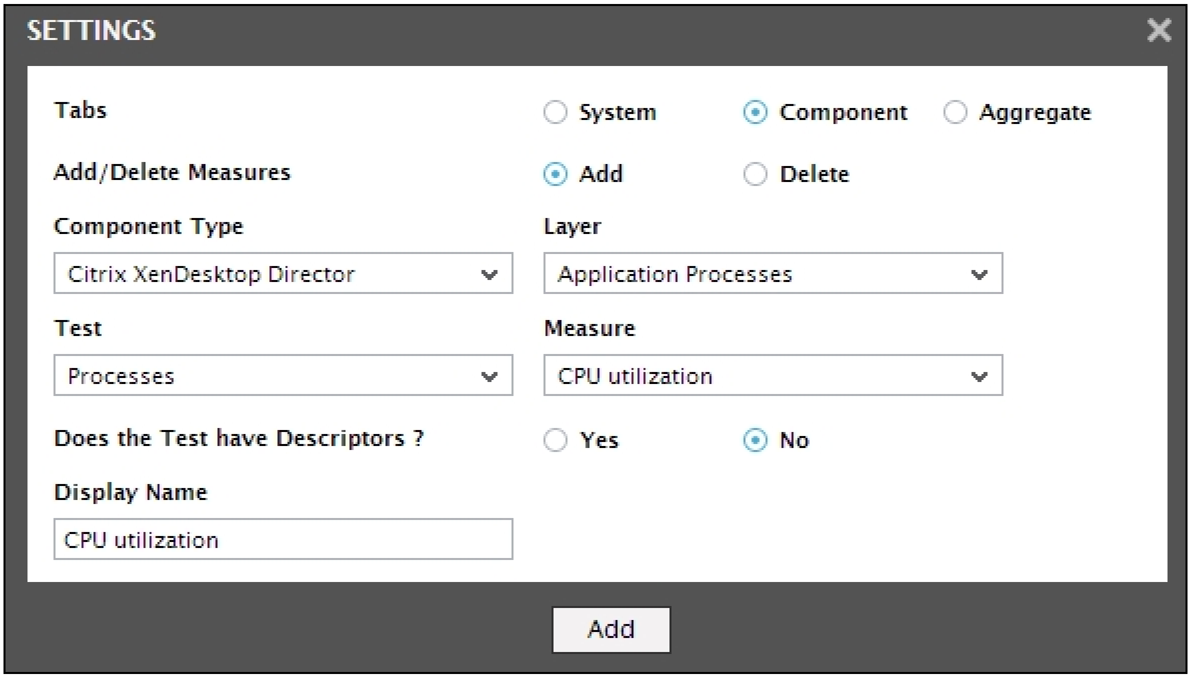

Figure 9.9: Selecting the Components flag from the Configuration Settings Window

- To add more metrics to the **Components** tab page, first, select the **Add** option from the **Add/Delete** Measures flag. Then, pick the **Component Type** to which the addition applies.
- Next, select the layer for which you wish to add the test from the **Layer** drop down list. Then, select the **Test** that reports the measure of your choice, pick the measure of interest from the **Measures** list, provide a **Display** name for the measure, and click the **Add** button to add the chosen measure to the **Components** tab page.
- If you want to delete one/more measures from this section, then, as soon as you choose the **Delete** option from the **Add/Delete Measures** flag, the **Test** drop down list will be populated with all the existing tests for which measures are displayed. Pick a test and choose a **Measure** of your interest to delete from the **Components** tab page.

Note:

- Only a user who is assigned the **Admin** role is allowed to customize the zone dashboard by clicking on the icon.
- While displaying values for descriptor-based measures in the **Components** tab page, the eG Enterprise system does not display the actual values per descriptor. Instead, the solution computes theaverage or the total sum of values across descriptors and displays it in the corresponding measure column. For instance, for values reported as percentages, the solution computes the average value across descriptors. On the other hand, if the value is reported as a GB or MB, then the total sum of all the descriptor values of the component will be displayed against the component.

9.3 The Details Tab Page

The **Details** tab page, when clicked, provides a quick overview of the performance of a chosen zone.

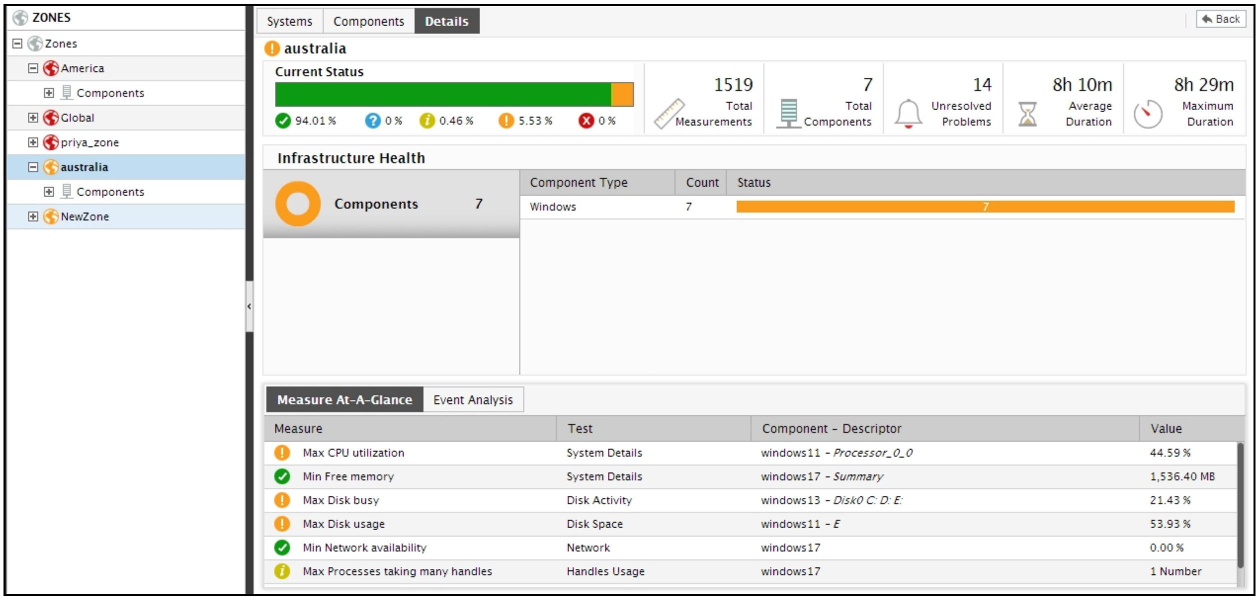

Figure 9.10: A Zone Dashboard

Just like the Monitor Dashboard, the **Details** tab page too comprises of four panels, each of which sheds light on a critical performance aspect of the chosen zone. The **Current Status** panel displays the total number of measurements that eG Enterprise has collected from the zone elements, and also indicates the percentage of measurements that are in abnormal, unknown, and normal states. A count of currently unresolved issues at the zone-level is also available here. This panel thus provides an overview of the health of the zone. Clicking on any of the states here will take you to the **Current Alarms** window, where you can view all open alarms of the corresponding priority.

A zone can contain a wide variety of infrastructure elements starting with independent components to segments, services, and even other zones. The **Infrastructure Health** section of Figure 9.10 therefore, graphically represents the different categories of infrastructure elements that a zone contains, and how well each category is currently performing. The **Sub Zones** bar graph for instance indicates the number of zones that have been added to the zone being monitored, and the current state of these subzones. You can zoom into individual subzone performance, by clicking on a division in the bar graph; the subzones which are in that particular state will then appear.

Clicking on the **Sub Zones** link in the Infrastructure Health section also invokes the **ZONE LIST** page, but in this case, the page displays all sub-zones that are part of the parent zone, regardless of state.

Either way, by default, the **ZONE LIST** page that appears displays the following sub-zones: direct sub-zones of the original zone, and zones (if any) that are included in the direct sub-zones. For instance, assume that 3 zones - zone A, zone B, and zone C - have been configured. While zone B has been directly assigned to zone A, zone C has been added to zone B. Now, while viewing the dashboard of zone A, if say, the **Sub Zones** link in the **Infrastructure Health** section is clicked, then the resulting **ZONE LIST** page will list the following by default:

- *zone B* which is the direct sub-zone of zone A;
- *zone C* which is added to zone B

In the same way, the **Components** bar graph represents the number and state of the components that are part of the zone. Clicking on a division in the **Components** bar lists the components in that particular state. Instead, if you click on the **Components** label in the **Infrastructure Health** section, you can view the complete list of components associated with the chosen zone, regardless of state.

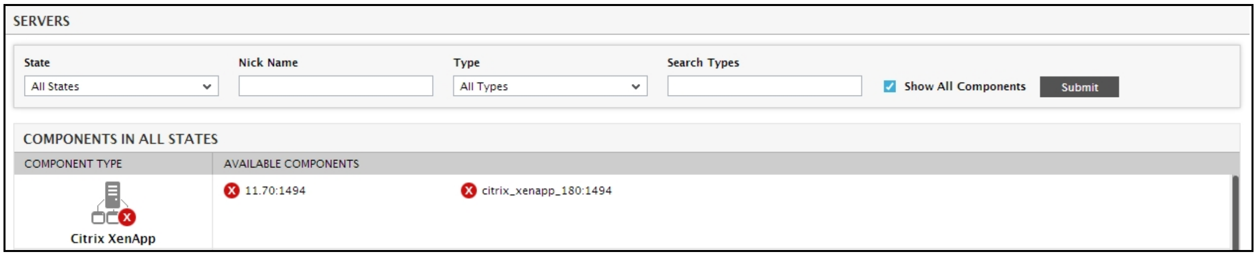

Figure 9.11: A page displaying the zone components in a CRITICAL state

Likewise, the **Services** and **Segments** bars in the **Infrastructure Health** section indicate the number and state of services and segments (if any) that are part of the said zone. While clicking on any division in the **Services** graph provides you with a list of services in that particular state, segments of a specific state will be displayed when you click on the corresponding division in the **Segments** graph. Alternatively, you can click on the **Services** or **Segments** label (as the case may be) to view all the segments/services (as the case may be) included in the zone, regardless of their state.

If the **Measures At-A-Glance** section is enabled by Clicking on the icon available in the **Admin** tab. Then, select the **Monitor** option in the **Settings** tile in the eG administrative interface, and if measures have been configured to be displayed in this section, then the **Details** tab page will display a **Measures At-A-Glance** section. The **Measures At-A-Glance** section (not shown in Figure 9.11) provides the min/max values of critical performance data collected in real-time from the zone being monitored. A quick look at this panel will instantly reveal significant deviations in zone performance. Click on any of the measures in this section to view the layer model, tests, and measurements pertaining to the corresponding component.

Alongside the **Measures At-A-Glance** tab is an **Event Analysis** tab, that primarily lists the top-5 layers at the zone-level, which were most affected by performance issues. Corresponding to every layer name in this section (see Figure 9.11), you will see the number of alarms that are currently open for that layer, the average duration of the open alarms, and the maximum duration for which an alarm had remained open. If you have a dedicated troubleshooting cell for the zone, then this information will serve as an effective indicator of the efficiency of the cell in resolving performance issues pertaining to the zone. To view the complete history of alarms in the environment, click on the **Click here for more events >>** link.

Besides a layer-wise event analysis, this section also enables a component-wise review of events that occurred during the last hour (by default). By choosing a different duration from the **Components with most events in the last** list, you can view the zone components that experienced performance degradations during the chosen duration, and the number of problem events each such component encountered. This sheds light on the most problem-prone components in the zone. Clicking on a component or component-type in this section, will lead you to the layer model of the corresponding component, revealing the current status of the component layers.

The **Components At-a-Glance** section comprises of a bar graph depicting the number of components of each type that available in the monitored zone, and their respective states. Clicking on a bar will take you to a page that lists the individual components of the corresponding type and state. To view the complete list of zone components of a particular type, just click on the corresponding component-type in the **Components At-A-Glance** section.

Note:

By default, in the **Components At-A-Glance** section, the component-types are sorted in the descending order of the total number of monitored components of every type - in other words, in the descending order of the values in the **Count** column of the section. To change the sort order - i.e., to sort the component-types in the ascending order of the contents of the **Count** column - simply click on the down-arrow icon next to **Count**. To sort by a different column, say, the **Server Type** column, simply click on the corresponding column heading. This will instantly sort the contents in the alphabetical order of the names of the displayed server types. You can even override the default sort order, so that the component-types are by default arranged in the alphabetical order of their names, and not on the basis of the **Count**. To achieve this, first switch to the eG administrative interface, and then Click on the icon available in the **Admin** tab. Then, select the **Monitor** option in the **Settings** tile. In the **OTHER DISPLAY SETTINGS** section of the **MONITOR SETTINGS** page, set the **Sort components in dashboards** flag to **By component types**. This ensures that the contents of the **Components At-a-Glance** section are by default sorted in the ascending order of the component-types. Accordingly, the down-arrow icon, by default, appears next to the column heading, **Server Type**.

10

Dashboards

10.1 The System Dashboard

The **System Dashboard** of an application allows you to focus on the performance of the operating system on which that particular application runs - i.e., the **Operating System** layer of an application. While viewing the layer model of an application using the **Layer Model** tab page, you can, if you so want, instantly switch to its **System Dashboard** for in-depth insights into system performance; for this, click on the **System** tab page.

Figure 10.1 depicts a sample **System Dashboard** of a Java Application server.

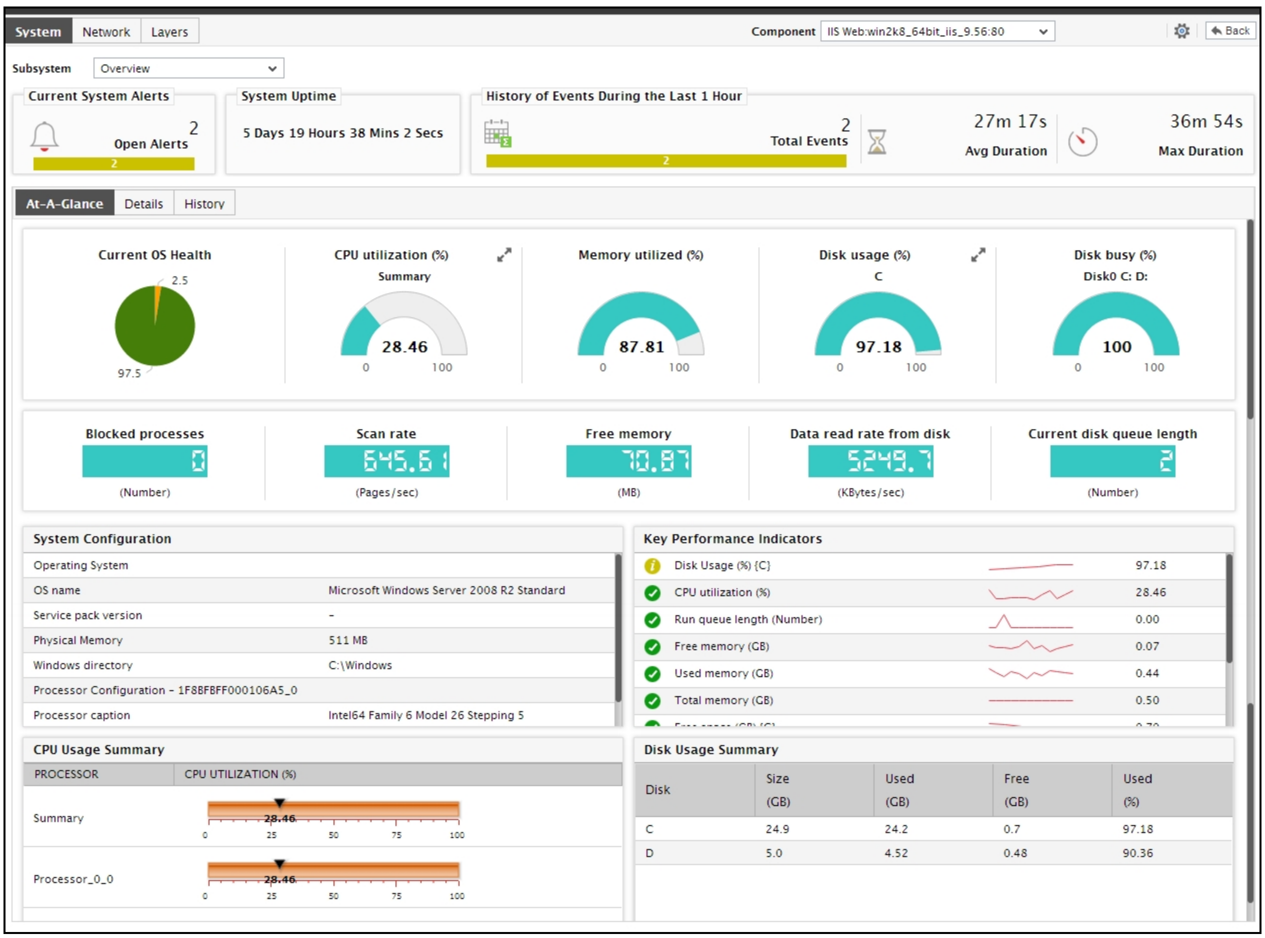

Figure 10.1: The System Dashboard of a Java Application

Using the **System Dashboard**, administrators can determine the following:

- The current status of the application host;
- The problems that the host is currently facing, and the type and number of problems it encountered during the last 24 hours;
- The current system configuration (if the eG license enables the **Configuration Management** capability);
- The current state of the critical parameters related to system performance;
- How some of the sensitive performance parameters have performed during the last 1 hour (by default);
- The resource-hungry processors supported by the host, and the disk partitions on the host that are currently experiencing a space crunch.

By default, the **System Dashboard** provides an overview of system performance. Accordingly, the **Overview** option is chosen by default from the **Subsystem** list of Figure 10.1. Instead of an **Overview**, if you prefer to receive an inside view of system performance - i.e., if you wish to investigate how effectively / otherwise the system in question has been using each of its resources or would prefer to focus on the uptime of the host, you can pick a different option from **Subsystem**. The sections below discuss each of these options in great detail.

10.1.1 Overview

In the **Overview** mode, the **System Dashboard** reveals the following:

1. The **Current System Alerts** section indicates the number of unresolved issues at the host-level, and also reveals how these issues have been distributed based on priority - i.e., the number of current issues of each priority. By clicking on an alarm priority, you can view the details of current alarms of that priority (see Figure 10.2). This way, you not only determine how problem-prone your operating system is, but also figure out the number and type of current problems at the system-level.

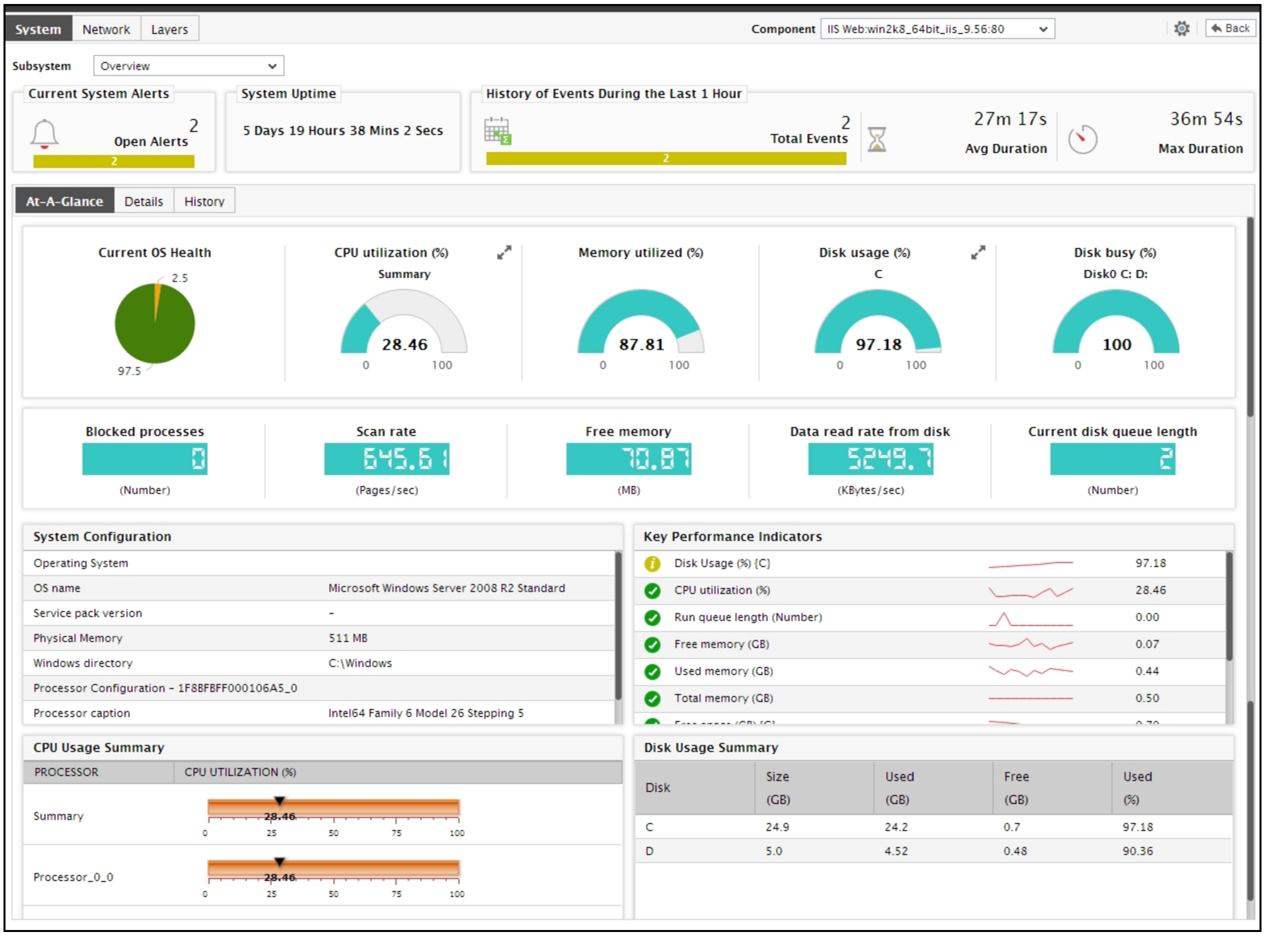

Figure 10.2: Figure 10.2: The details of current alarms raised at the system-level

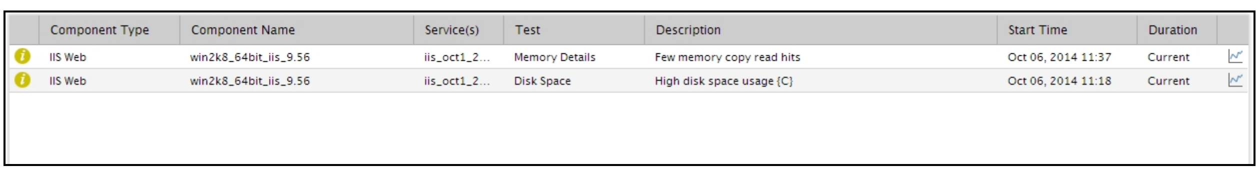

	Component Type	Component Name	Service(s)	Test	Description	Start Time	Duration
	IIS Web	win2k8_64bit_iis_9.56	iis_oct1_2...	Memory Details	Few memory copy read hits	Oct 06, 2014 11:37	Current
	IIS Web	win2k8_64bit_iis_9.56	iis_oct1_2...	Disk Space	High disk space usage {C}	Oct 06, 2014 11:18	Current

Figure 10.3: Figure 10.3: The details of problems of a particular priority

2. If too many alarms are displayed in Figure 10.2, you can use the **Search** text boxes placed at the end of this alarm window to perform quick searches based on the **Description**, **Layer**, and **StartTime** columns to locate the specific alarm(s) of interest to you. For instance, to look for an alarm with a specific description, specify the whole/part of this description in the text box below the **Description** column in Figure 10.3. Doing so will automatically display the details of only that alarm containing the specified description.

3. If you click on any alarm in Figure 10.3, an **Alarm Details** section will be introduced in the **Alarms** window itself, providing additional details of the alarm clicked on. These details include the **Site** affected by the problem for which the alarm was raised, the test that reported the problem, and the **Last Measure** value of the problem measure.

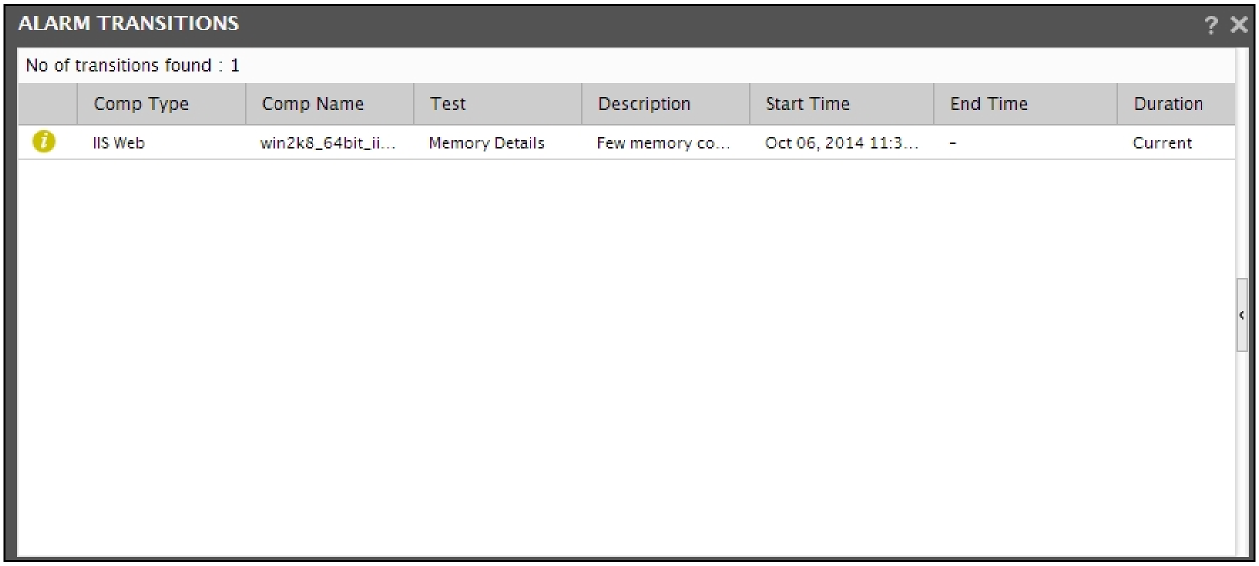

Figure 10.4: Figure 10.4: Additional alarm details

4. To figure out the type of problems that occurred the maximum on the system during the last 1 hour, refer to the **History of Events** section in Figure 10.4. This section provides a bar graph that reveals the number of problems of each priority that the application host experienced during the last 1 hour. Clicking on a bar will lead you to the **HISTORY OF ALARMS** page (see Figure 10.4) that displays the complete list of problem events that occurred at the system-level in the last 1 hour. This information provides administrators with quick and effective insights into recurring problems, and enables them to deduce problem patterns.

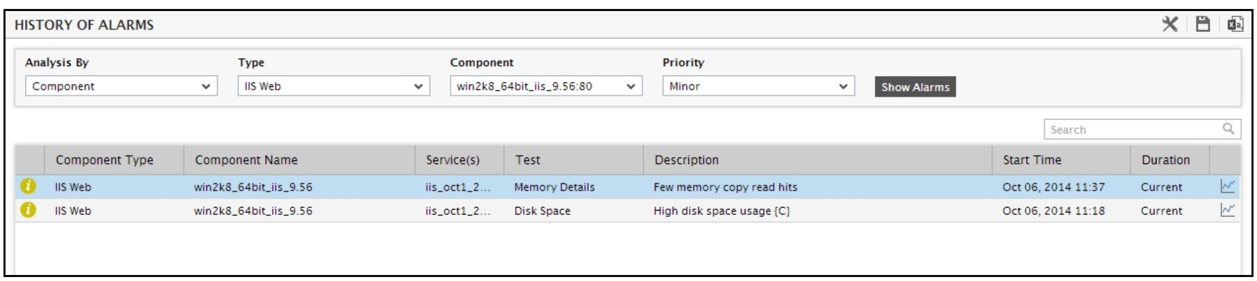

Figure 10.5: Figure 10.5: History of alarms in the last 24 hours

5. Once back in the dashboard, you will find an **At-A-Glance** tab page that enables you to determine, at a glance, the current state of the system. This tab page begins with a **Current OS Health** section, which provides a pie chart revealing how problem-prone the system currently is - in other words, it indicates the service level that has been achieved by the system currently. Clicking on a slice will lead you to the **EVENT HISTORY** page again (see Figure 10.4).

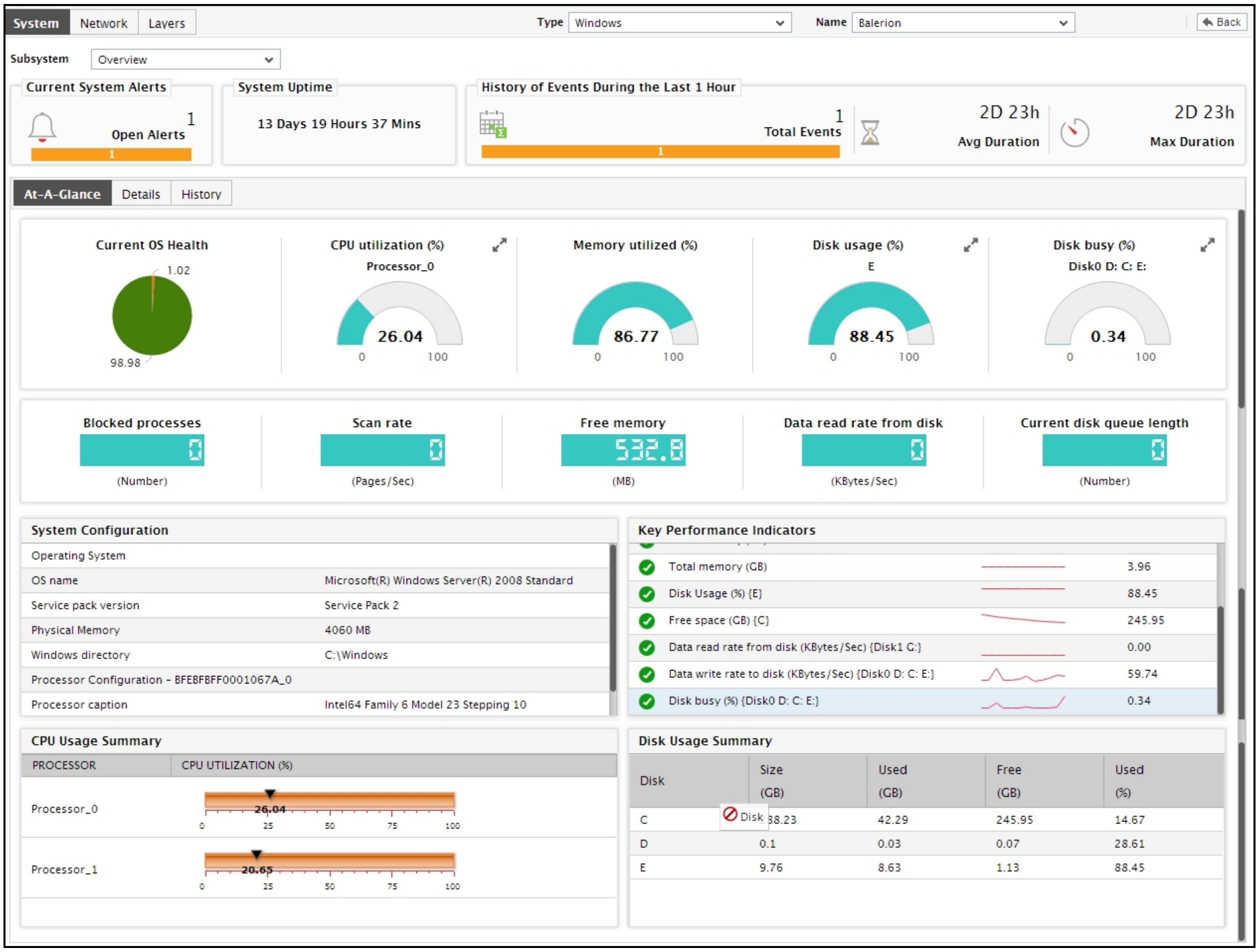

Figure 10.6: Figure 10.6: The At-A-Glance tab page

6. This pie chart will be followed by a series of pre-configured host-level measures and their current values, with the help of which unhealthy metrics can be instantly detected and impact analysis easily performed. While measures that report percentage values are typically represented using a dial chart, other measures are reported using digital displays. Since all these values are rounded-off to two decimal places (by default), you are advised to move your mouse pointer over these “approximations”, so that the “actuals” can be viewed as a tool tip.

   Note:

   If you have configured one/more measures of a descriptor-based test to be displayed as a dial chart, then, in real-time, the descriptor that is in an abnormal state or is currently reporting the maximum value for that measure will be represented in the dial chart. You can view the dial charts pertaining to the other descriptors, by clicking on the **More** button that appears alongside a dial chart in the dashboard.

7. Also, to enable administrators to instantly and accurately detect deviations from the norm, the dial charts, by default, indicate the threshold settings of a measure along with the real-time values reported by that measure. If multi-level thresholds are set for a measure, then each such threshold will be indicated using the conventional color-codes (*Red* for *Critical*, *Orange* for *Major*, and *yellow* for *Minor*) used across the eG monitoring and reporting consoles. By default, the dial charts display the Maximum thresholds alone. If a

measure is associated with Minimum thresholds only, then the dial chart will display the minimum thresholds settings instead. Thanks to the threshold representations in the dial charts, administrators can easily identify when and what type of thresholds were violated.

8. Let us now return to the dashboard. If you click on any dial/digital graph in the dashboard, you will be directly lead to the layer model page, where the exact layer-test-measure combination that corresponds to the dial/digital graph will be displayed.

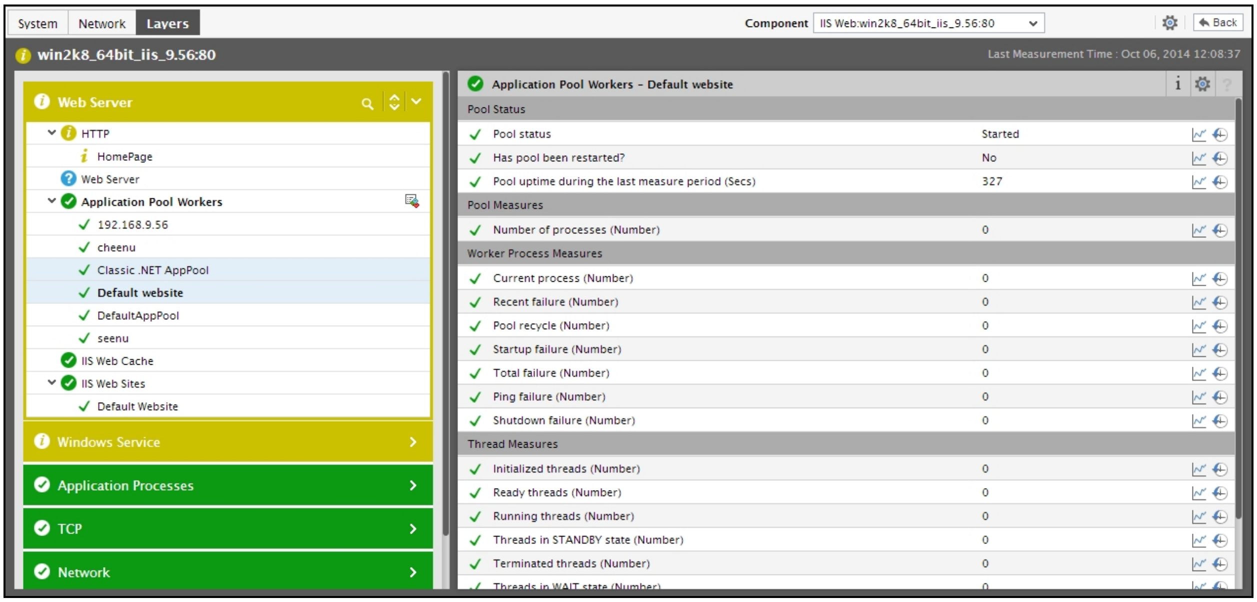

Figure 10.7: Figure 10.7: The layer model, test, and measure that appear when a dial graph is clicked on

9. Now that we are done exploring the dial and digital graphs, let us proceed to focus on the other sections of the dashboard. A quick look at the current **System Configuration** in Figure 10.5 helps determine whether a change in system configuration can make the host less vulnerable to performance issues. **Note that the System Configuration will appear only if your eG license allows Configuration Management**; if not, then this section will display a bar chart indicating the current status of the **Operating System** layer of the host, in terms of the percentage of time the host has been in normal/critical/major/minor/unknown states.

10. Let us now focus on the dashboard again. A list of critical system-related measures and their current state is provided under the head **Key Performance Indicators** in the dashboard, so that administrators can swiftly determine if the eG agent has detected any abnormalities with any of the factors that significantly influence system performance. This way, remedial measures can be immediately initiated. Clicking on a measure here will lead you the **Layer Model** tab page displaying the monitoring model of the target application, and the value reported by the measure that was clicked on (see Figure 10.6).

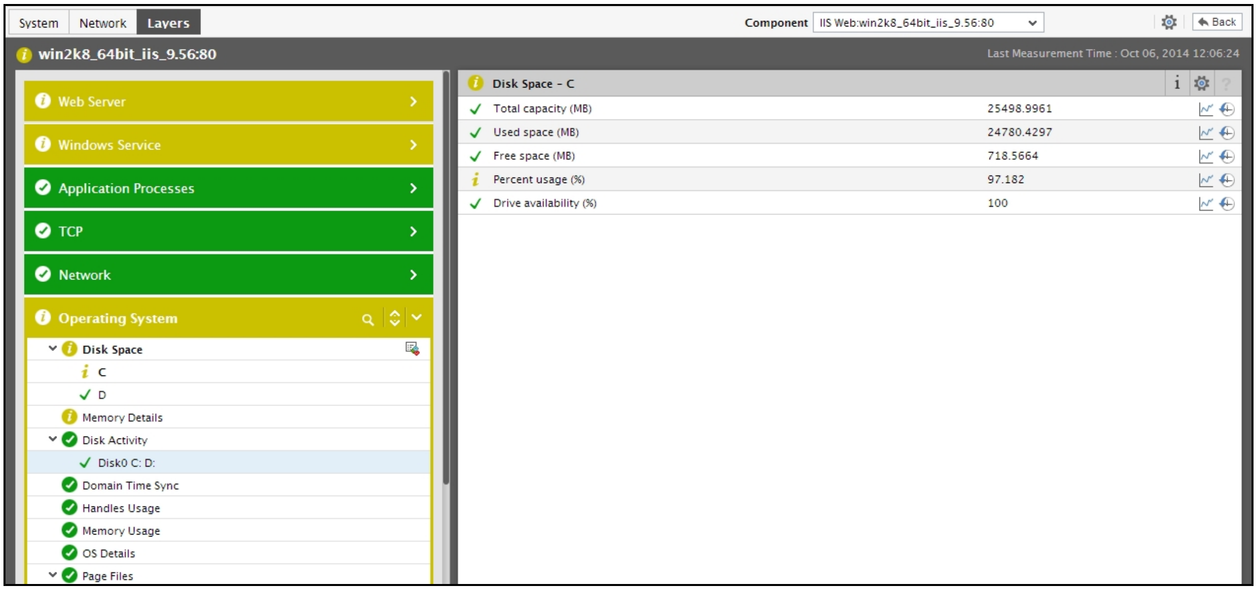

Figure 10.8: Figure 10.8: The layer model tab page that appears when a key performance indicator is clicked

11. Moreover, corresponding to each of the core measures displayed in the **Key Performance Indicators** section, a miniature graph will be available, which provides a quick look at the variations in that measure during the last 1 hour (by default). By observing these variations more closely and clearly - say, for even longer time periods - you can rapidly detect disturbing performance trends and proactively isolate potential problems. To achieve this, click on the miniature graph to expand it. Figure 10.7 then appears, displaying a zoomed out graph. Using Figure 10.10, you can alter the **Timeline** of the graph, and also change the default dimension of the graph from **3D** to **2D**. In addition, if the time-of-day values reported by multiple descriptors are plotted in the graph, you can choose to focus on the historical performance of the best/worst descriptors alone by picking a **TOP-N** or **LAST-N** option from the **Show** list that appears in the expanded graph (not shown in Figure 10.7). For instance, if the graph tracks the usage of all the disk partitions on the host over time, then, you can pick the **TOP-3** option from the **Show** list to make sure that the graph plots the historical values of only those 3 disk partitions, which are being used the maximum.

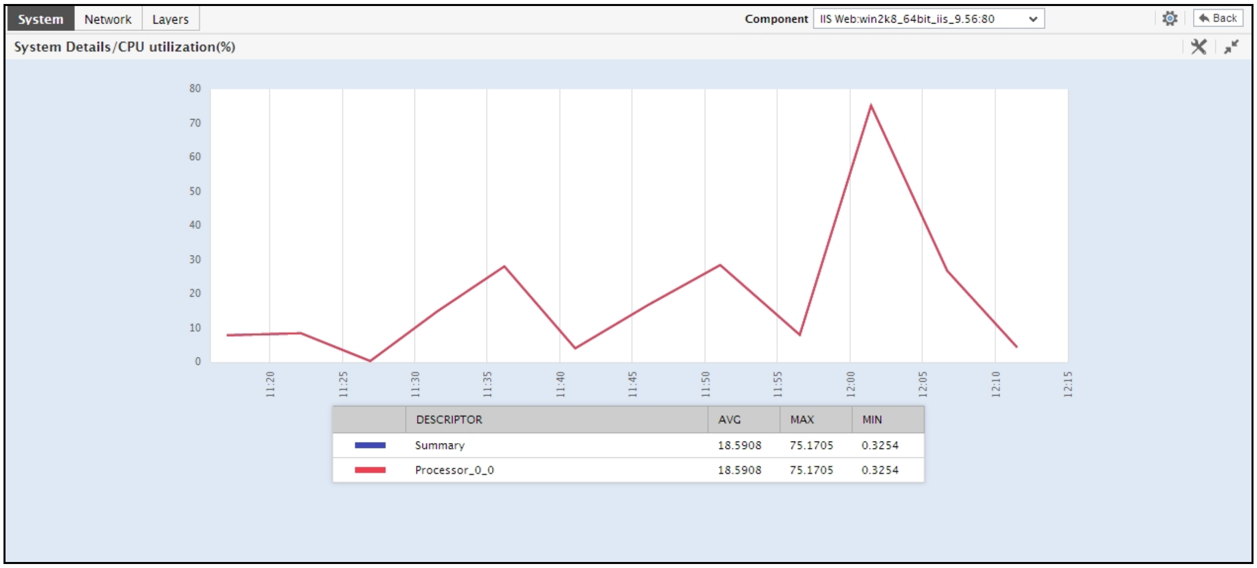

Figure 10.9: Figure 10.9: The expanded graph that appears upon clicking the miniature graph in the Key Performance Indicators section

12. Beneath the **Key Performance Indicators** section, you will find a **CPU Usage Summary**; from this summary, you can quickly understand how well all the processors supported by the system are currently utilizing the host's CPU resources, and also accurately identify those processors that are eroding these critical resources.

13. Similarly, the **Disk Usage Summary** will reveal the current capacity and usage of each of the disk partitions on the host, so that you can swiftly isolate disk partitions that are running out of space.

14. Thus, with the help of tabulated usage statistics, the **At-A-Glance** tab page turns the spot light on resource-intensive processors and disk partitions on the host. Alternatively, if you prefer an interface that provides a graphical comparison of resource usage across processors and disk partitions (as the case may be), combined with quick insights into the root-cause of usage excesses (if any), then, you can switch to the **Details** tab page instead. For this purpose, click on the **Details** tab page in Figure 10.9. Figure 10.10 will then appear.

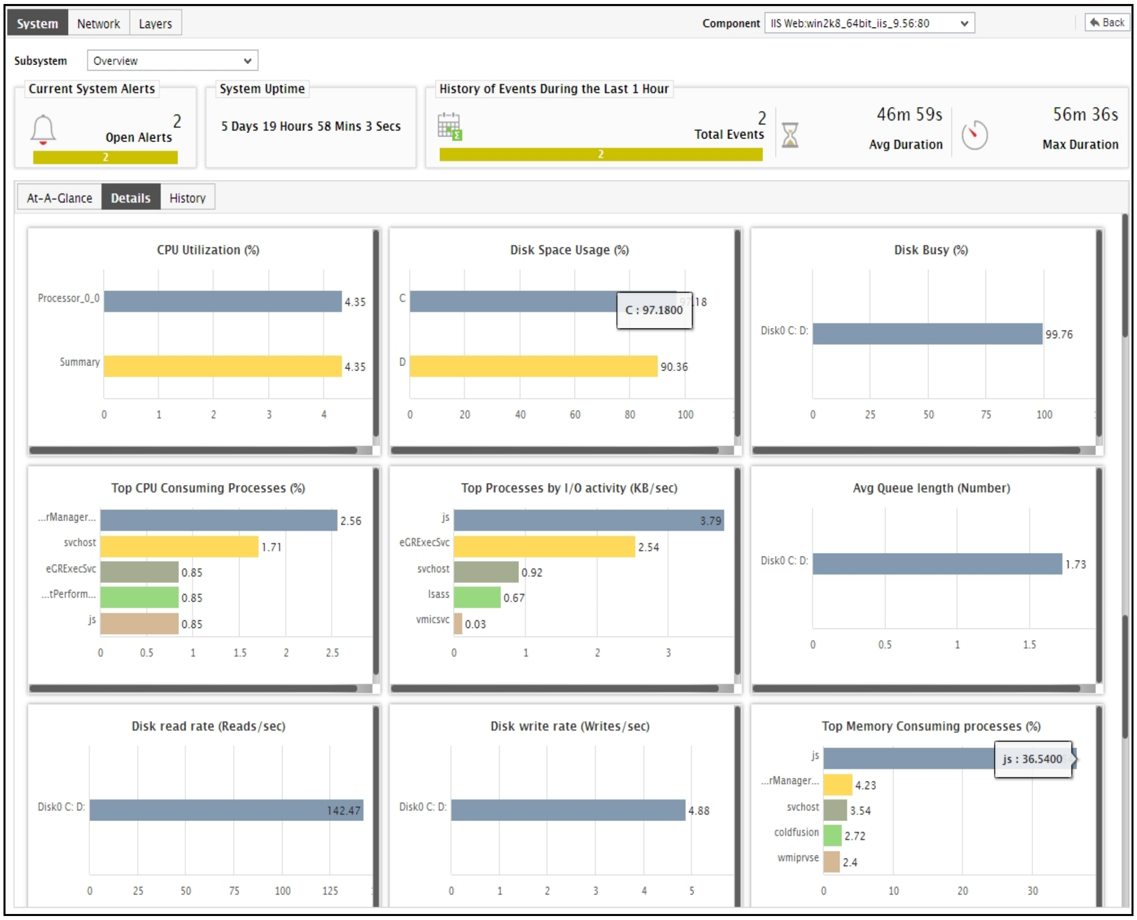

Figure 10.10: Figure 10.10: The Details tab page of the Overview subsystem

15. As you can see, Figure 10.10 provides a set of pre-defined bar charts, each of which focuses on the current usage of a key resource (disk/CPU/memory). Using these default graphs, you can easily and accurately determine the following:

   - Which processor is currently utilizing the maximum CPU resources? Which process currently executing on the host is causing this resource-drain?
   - Which disk partitions are left with limited free space?
   - Which disk partitions are the busiest in terms of the rate of I/O requests they handle? Which processes currently executing on the host are causing high disk I/O?
   - Which processes on the host are consuming memory excessively?

16. These bar charts, in fact, can also be configured to aid effective *postmortem analysis* of resource usage. For instance, you can use one of these bar charts to find out which process caused the memory usage on the host to increase during a time period in the past. For this purpose, click on the corresponding bar chart in Figure 10.11. The graph will then zoom out as depicted by Figure 10.11.

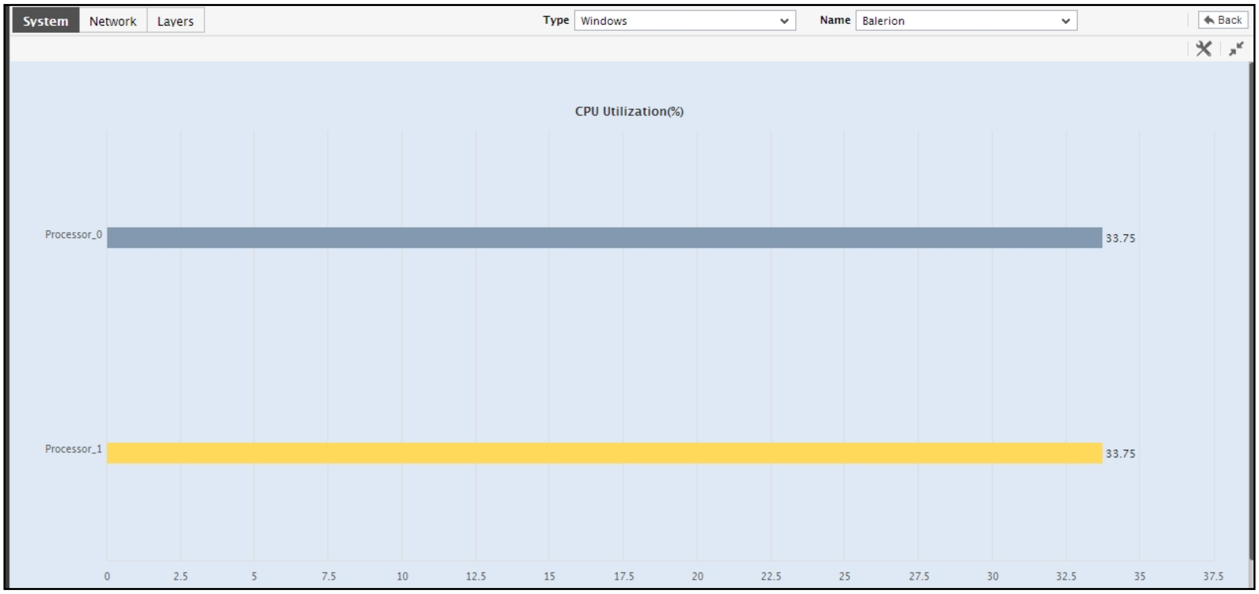

Figure 10.11: Figure 10.11: An enlarged bar chart in the Details tab page

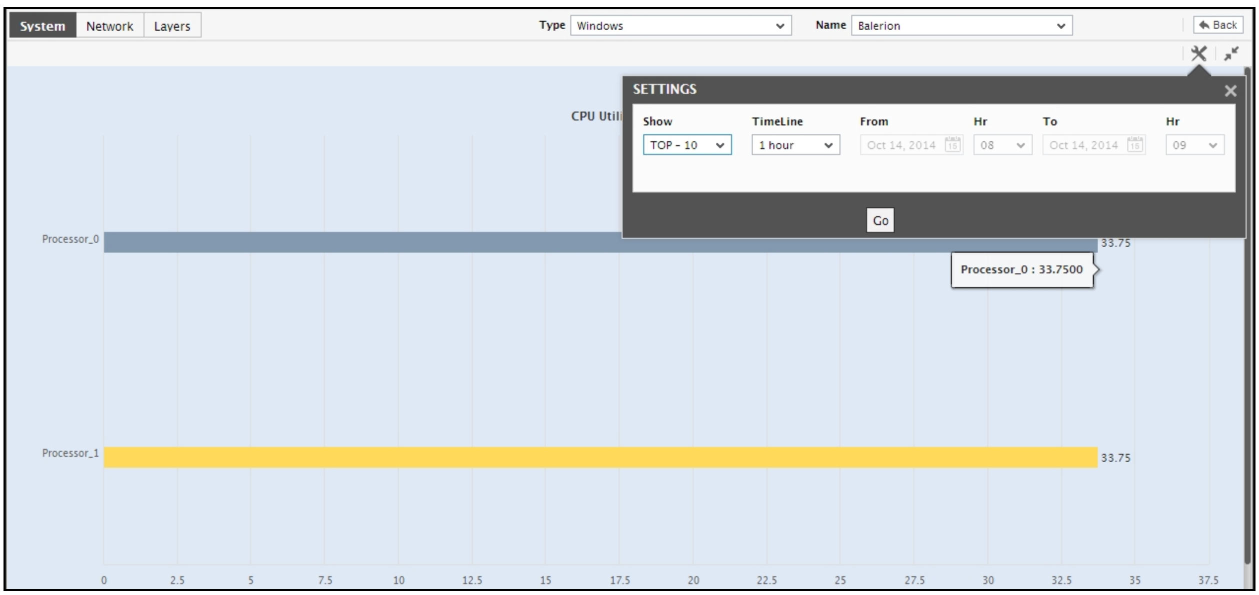

Figure 10.12: Figure 10.12: The Top-N graph for a past period

17. By default, the resulting graph will display the top-10 processes that executed on the host during the specified **Timeline**, in the descending order of memory usage. Accordingly, the **TOP-10** option is chosen by default from the **Show** list. To view only a limited number of processes in the graph, pick a different **TOP-N** or **LAST-N** option from the **Show** list.

18. This way, with the help of the bar charts, you can quickly get to the source of any resource contention at the host. However, to engage in an elaborate historical analysis of the behavior of the host-level measures and isolate probable problems, the measure/summary/trend graphs offered by the **History** tab page will be more useful. To switch to this tab page, click on **History** in Figure 10.13. Figure 10.14 will then appear.

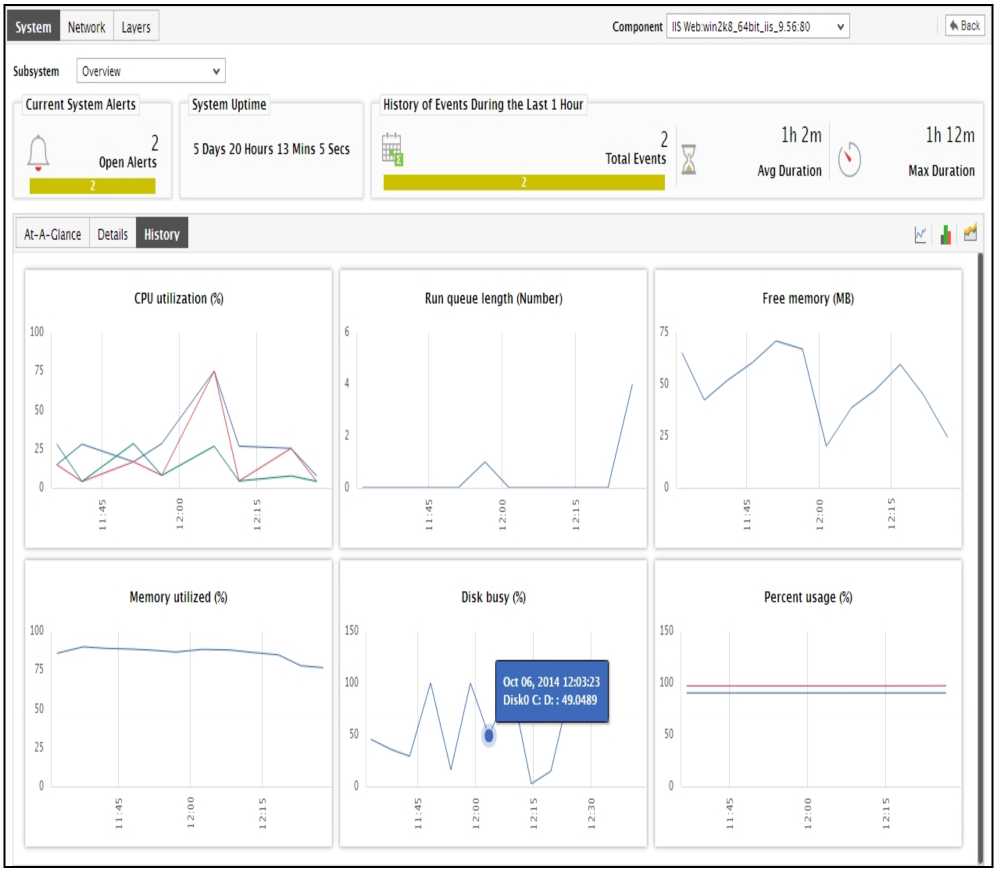

Figure 10.13: Figure 10.13: The History tab page of the Overview subsystem

19. By default, the **History** tab page provides measure graphs for each of the key host-level measures, using which you can efficiently track the changes in the performance of the measures during the last 24 hours (by default). These graphs help determine when a measure, which is currently in an abnormal state, began exhibiting performance inconsistencies.

20. Some measure graphs in the **History** tab page may plot values for multiple descriptors; such graphs will appear very cluttered, making analysis a nightmare! To view such measure graphs clearly, you will have to first enlarge the graph by clicking on it. The graph will then zoom out as depicted by Figure 10.14.

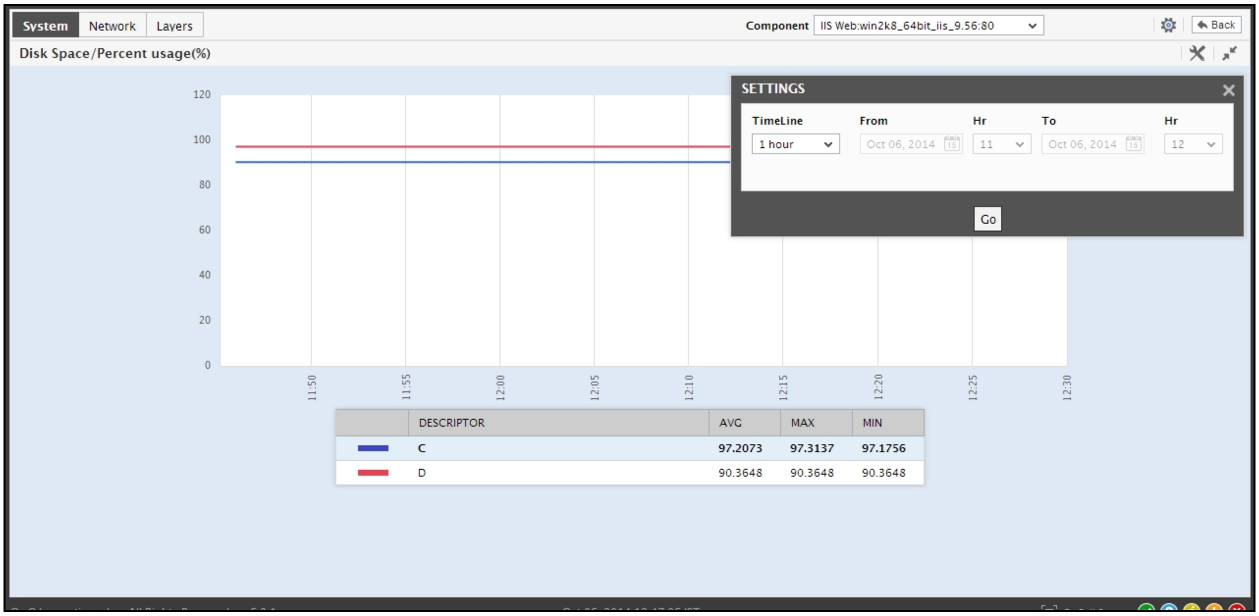

Figure 10.14: Figure 10.14: Expanding a measure graph in the History tab page

21. If need be, you can change the **Timeline** of the enlarged graph, choose to view only a few of the descriptors in the graph by picking a **TOP-N** or **LAST-N** option from the **Show** list.

22. You can change the graph timeline by clicking on the **Timeline** link in Figure 10.15. You can even expand the graph by clicking on it, and then alter its **Timeline**.

23. In addition to the timeline and dimension, the enlarged summary graph also allows you to change its **Duration**. By default, the **Duration** is set to **Hourly**, indicating that the summary graphs plot only the hourly summaries by default. If required, you can change the **Duration** of the summary graph in the enlarged mode so that, you can perform daily or monthly summary analysis.

24. Similarly, to observe and understand the past trends in the performance of the host and to predict future measure behavior, click on the icon at the right, top corner of the **History** tab page. Figure 10.16 will appear revealing trend graphs for the host-level metrics. By default, these graphs plot the maximum and minimum values registered by a measure during the default period of 1 day.

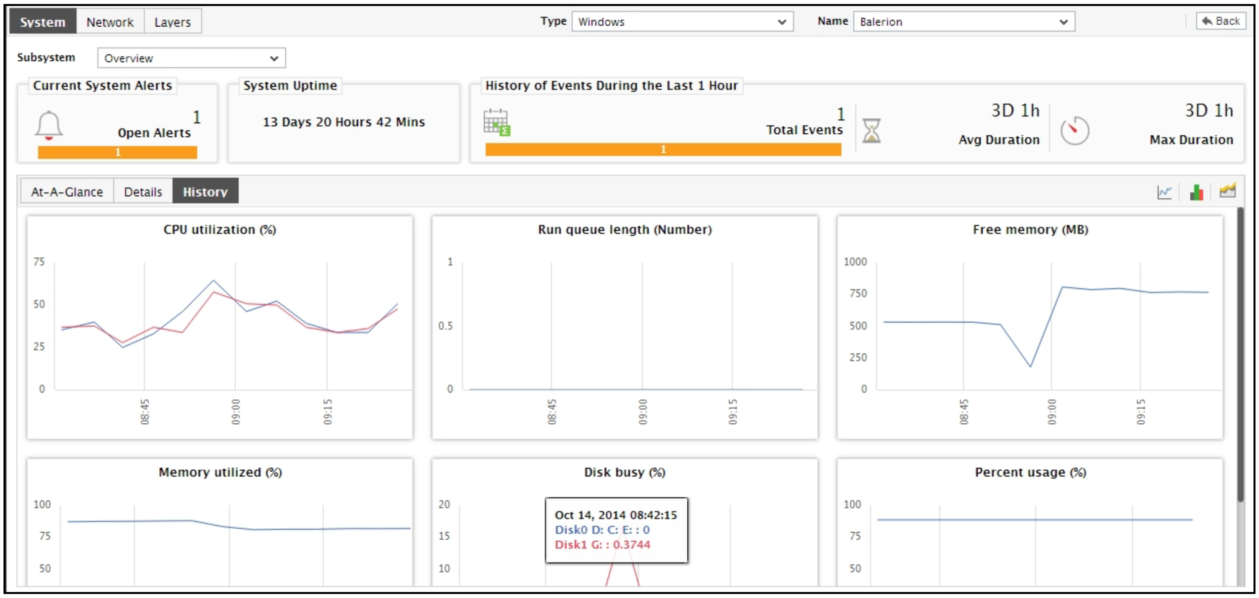

Figure 10.15: Figure 10.15: Trend graphs in the History tab page

25. If need be, you can change the graph timeline by clicking on the **Timeline** link in Figure 10.15. You can even expand the graph by clicking on it.

26. Doing so will invoke Figure 10.16, where you can view the enlarged graph. By default, only hourly trend values are plotted in a trend graph. If need be, you can change the **Duration** of the trend graph in the enlarged mode, so that you can perform daily or monthly trend analysis. Likewise, you can change the graph **Timeline** using Figure 10.16.

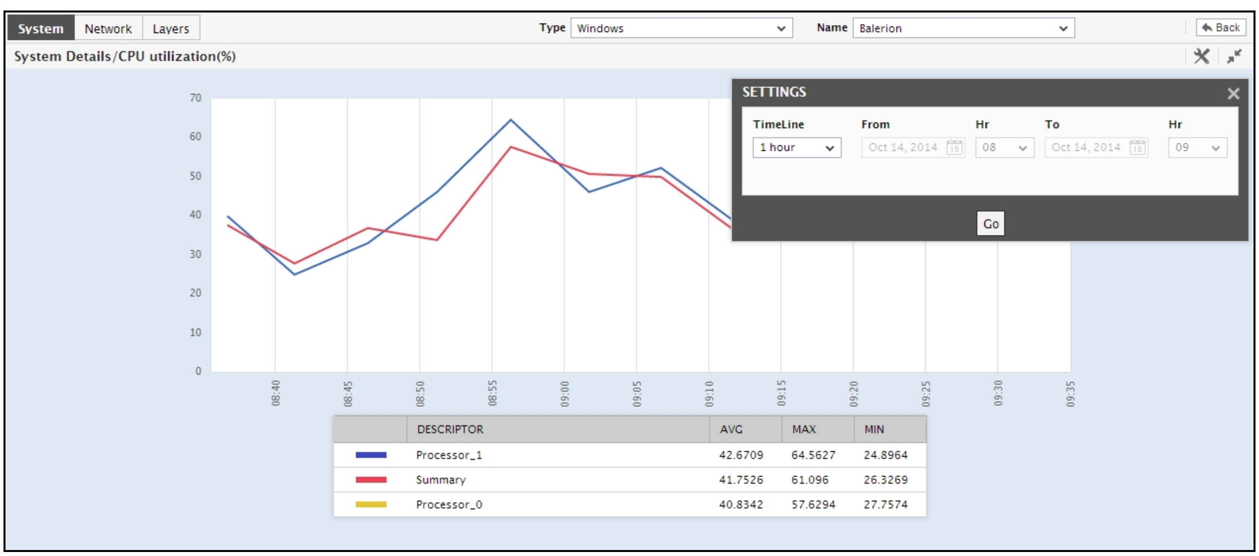

Figure 10.16: Figure 10.16: Changing the Duration of the trend graph

27. Also, by default, the trend graph displays the **AVG**, **MAX** and MIN value below the graph.

Note:

In case of descriptor-based tests, the **Summary** and **Trend** graphs displayed in the **History** tab page typically plot the values for a single descriptor alone. To view the graph for another descriptor, pick a descriptor from the drop-down list made available above the corresponding summary/trend graph. For instance, in Figure 10.18 above, note that the trend graph for **CPU utilization** plots the CPU usage trends of *Processor_1* only. If you want to view the CPU usage trend graph for *Processor_0* instead, pick the *Processor_0* option from the drop-down list adjacent to the graph title, **CPU utilization (%)**.

28. At any point in time, you can switch to the measure graphs by clicking on the button.

29. Typically, the **History** tab page displays measure, summary, and trend graphs for a default set of measures.

10.1.2 CPU

You can also use the **System** dashboard to instantly identify CPU bottlenecks and the processes responsible for the same. For this, select the **CPU** option from the **Subsystem** list. Figure 10.17 will then appear.

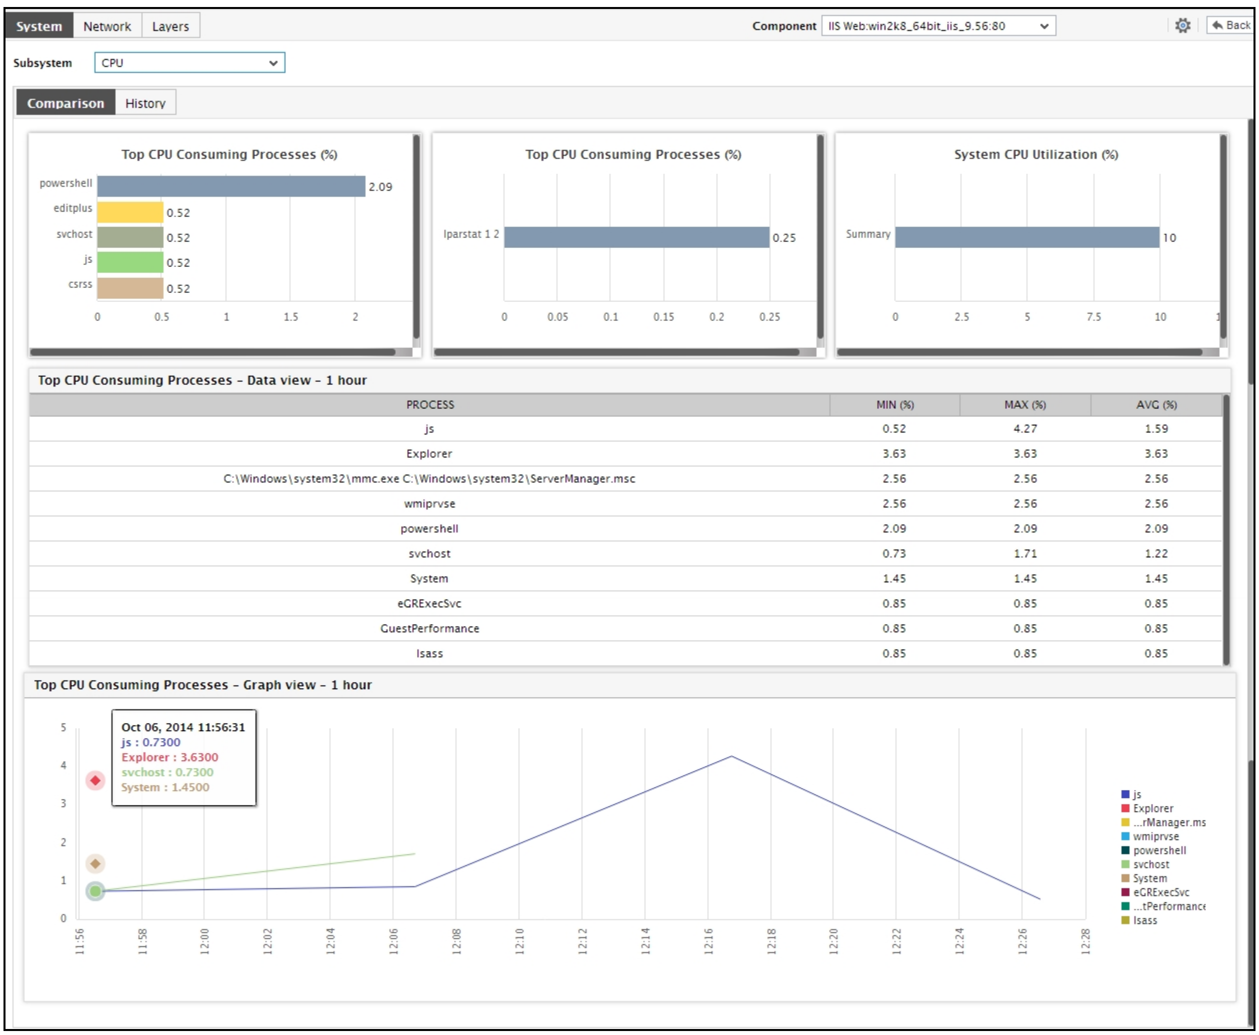

Figure 10.17: The System Dashboard for the CPU subsystem

1. The first section of the dashboard helps determine whether/not the host is currently facing any CPU-related issues - dial charts and digital number displays provided by this section enables administrators to figure out whether or not all the key CPU health indicators are currently operating well-within their thresholds limit. The dial charts do not indicate the threshold settings for measures by default.

2. Let us now return to the **CPU** dashboard. If you click on a dial chart or a digital graph in the **CPU** dashboard, Figure 10.18 will appear displaying the exact measure that is represented by the dial/digital chart, and the layer and test to which that measure is mapped.

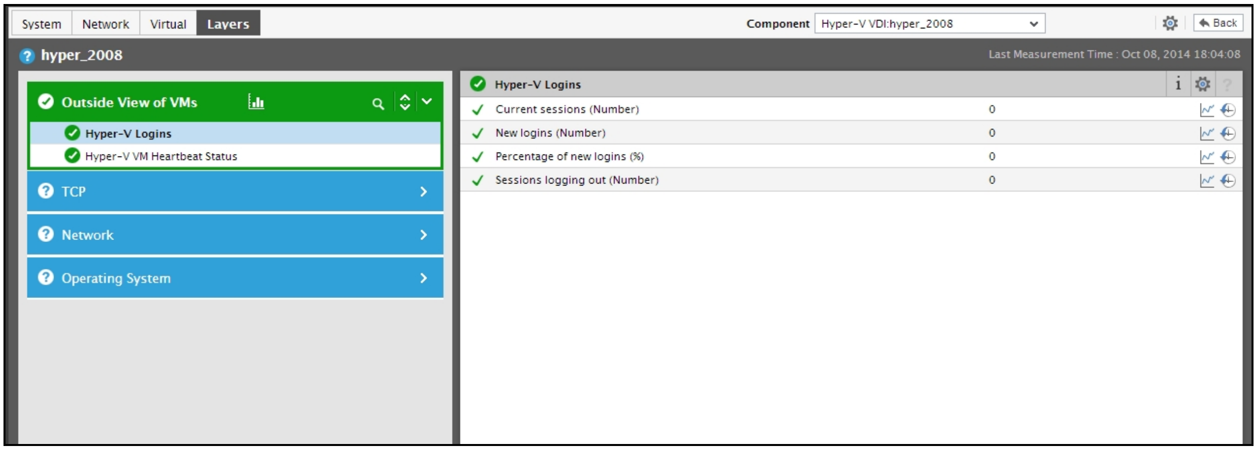

Figure 10.18: Clicking on a dial chart/digital chart to view the corresponding layer model, test, and measure

3. Below the dial and digital charts, you will find the **Comparison** tab page, which provides a default set of bar charts comparing the current CPU usage of the processors supported by the host and the processes executing on the host. Using these default bar graphs, you can accurately identify those processors that are excessively utilizing the CPU resources of the host, and those processes that are responsible for this CPU erosion.

4. Clicking on any of the bar charts in the **Comparison** tab page will expand the chart; for instance, to clearly view the top CPU consuming processes on the host, you can enlarge the **Top CPU Consuming Processes** bar chart as depicted by Figure 10.19 below.

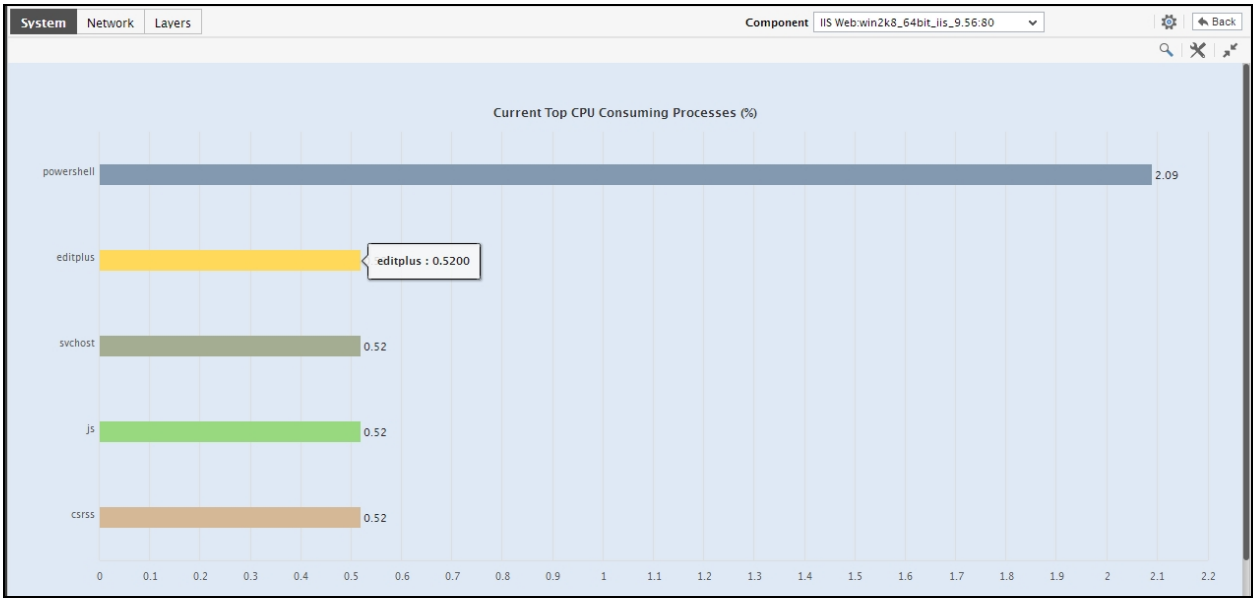

Figure 10.19: Expanding a Top-N bar chart in the Comparison tab page of the CPU subsystem dashboard

5. By default, the enlarged bar graph in Figure 10.20 displays **TOP-10** the processes currently executing on the host and the percentage CPU utilized by each. If need be, you can customize the enlarged bar graph, so that it displays only a few processes - say, only the top-5 processes in terms of CPU usage. For this, simply select the **TOP-5** option from the **Show** list in Figure 10.21.
6. Moreover, besides the current CPU usage, you can also compare the historical CPU usage of processes by clicking on the **Compare History** link in Figure 10.22. Doing so will allow you to alter the **Timeline** of the bar graph, and enable you to zero-in on that process that was devouring the CPU resources on the host during a particular time period in the past.

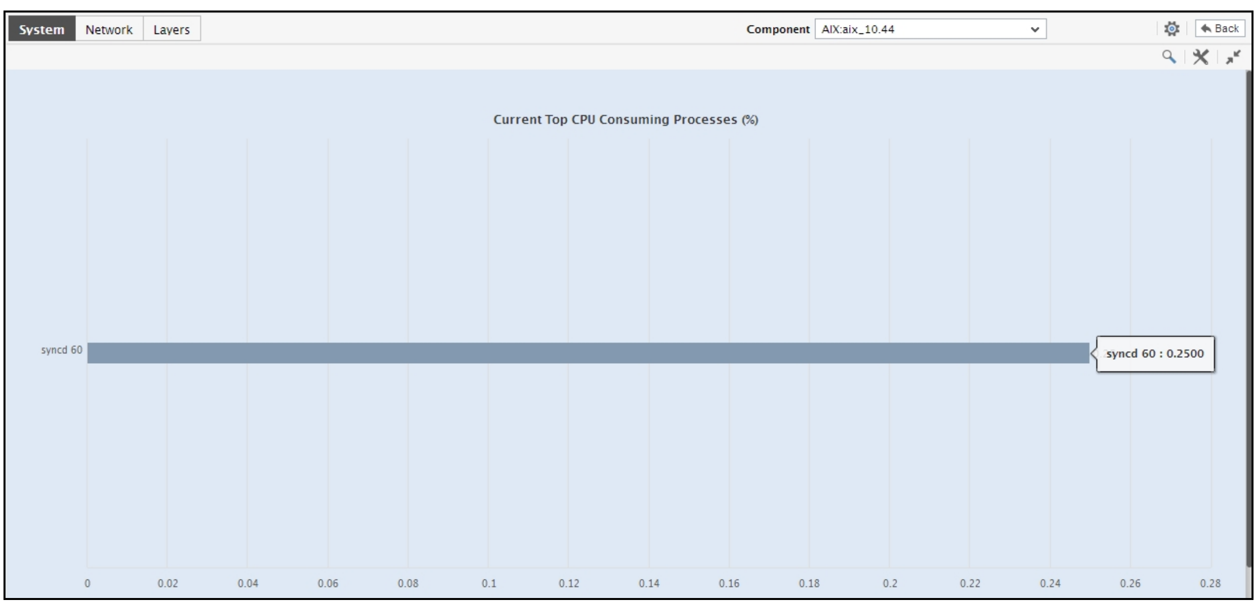

Figure 10.20: Comparing the historical CPU usage of processes

7. You can also click on the icon in Figure 10.20 to invoke the **DETAILED DIAGNOSIS** page. Since this page provides the **PROCESS ID** of the top CPU consumers on the host, it enables you to easily locate the CPU-intensive processes and initiate remedial measures.

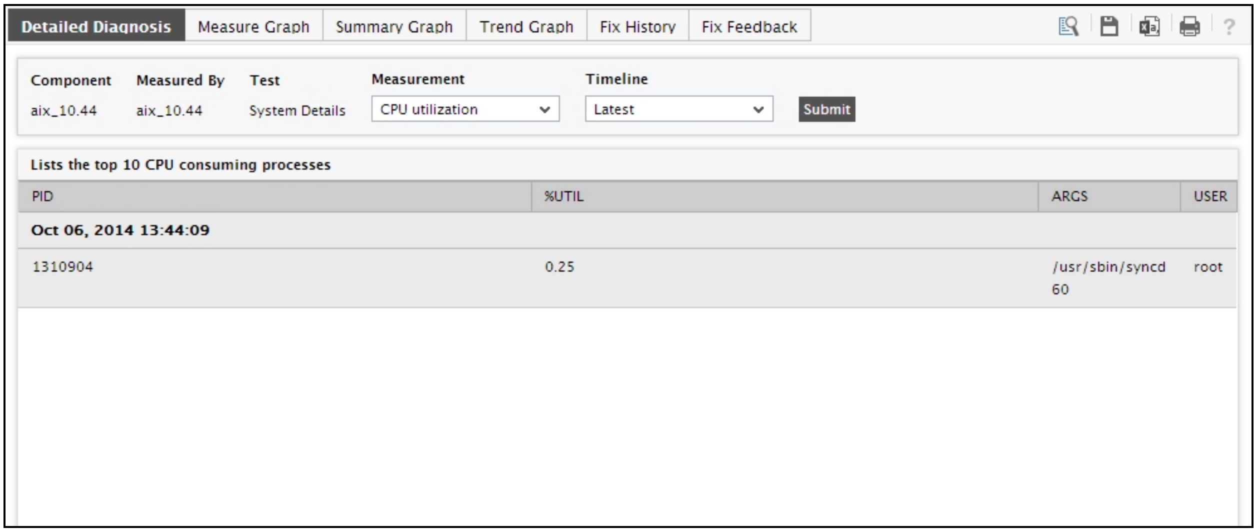

Figure 10.21: The detailed diagnosis of the top CPU-consuming processes

8. In addition to these bar charts, this tab page also provides you with a **CPU Usage Summary by Top Processes** table. While the charts focus on current CPU usage by default, the table reveals the top CPU consumers on the host during the last 1 hour (by default). Besides, the table also indicates how high and how low the CPU usage of each process has scaled during the same hour. This way, you can understand whether the high CPU usage of a process was just a sudden spike or a consistent phenomenon.

9. Besides the above, the tab page also embeds a **CPU Usage History by Top Processes** chart. This chart graphically compares the CPU usage of the processes that were executing on the host during the last 1 hour (by default), and enables the quick and easy identification of the process that is the top CPU consumer. Clicking on this chart enlarges it enabling you to perform the comparative analysis more effectively.

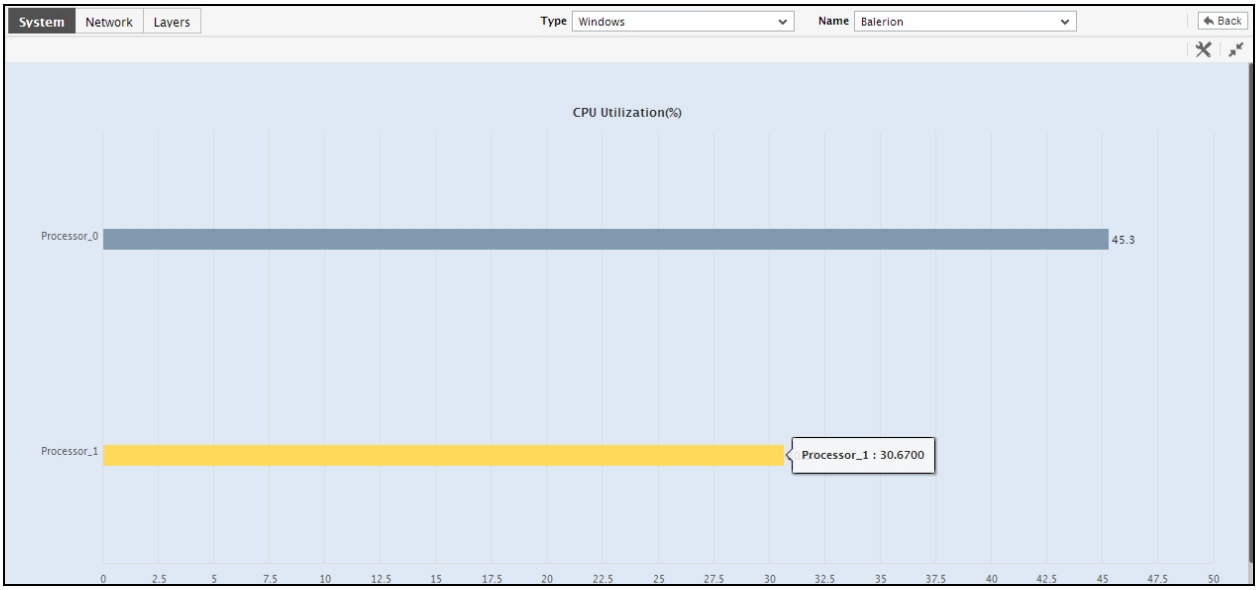

Figure 10.22: The CPU utilization graph enlarged

10. You can alter the **Timeline** of the comparison, change its dimension, and even invoke the detailed diagnosis page (see Figure 10.24) to determine the **Process ID** of the leading CPU consumer.

11. To shift your focus from current performance to historical performance, click on the **History** tab page. Figure 10.25 will then appear.

Figure 10.23: The History tab page of the CPU dashboard

12. By default, the **History** tab page displays time-of-day measure graphs revealing how the host has been using the CPU resources of the host over the last 24 hours. Using these graphs, you can effortlessly figure out when exactly a CPU contention (if any) crept into the host. If required, you can click on the **Timeline** link in Figure 10.26 to go further back in time and analyze the CPU usage.

13. Back in the dashboard, you can enlarge a graph in the **History** tab page by clicking on that graph.

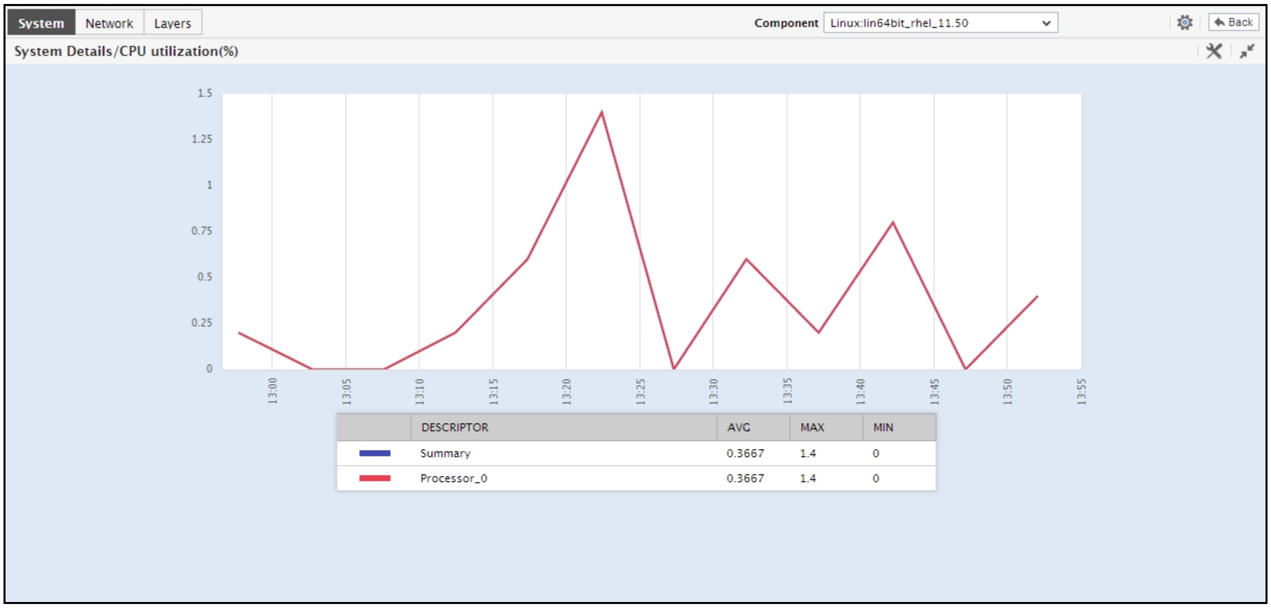

Figure 10.24: An enlarged CPU usage measure graph

14. Here again, you can modify the graph **Timeline**, or change its dimension from **3D** to **2D**. In case of graphs that plot values for multiple descriptors, you can selectively compare the CPU usage of a few descriptors alone by selecting a **TOP-N** or **LAST-N** option from the **Show** list.

15. By clicking on the icon in the **History** tab page you can convert the measure graphs into summary graphs; these summary graphs, by default, reveal the percentage of time in the last 24 hours during which the CPU usage of the host has remained optimal. Besides indicating whether the overall CPU usage of the host has been within the acceptable limits, it also indicates how often during a day these usage levels have been compromised. You can click on a summary graph to enlarge it; furthermore, you can change the **Timeline** of the graph in the enlarged mode.

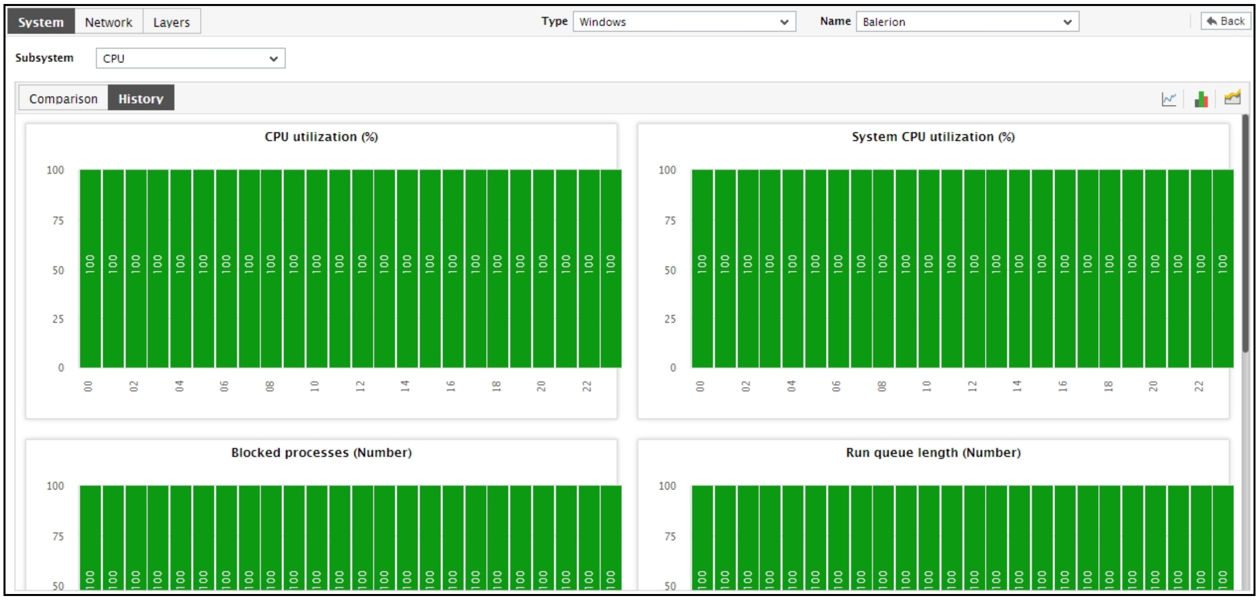

Figure 10.25: Summary graphs on CPU usage displayed in the History tab page of the CPU dashboard

16. You can even expand the summary graph by clicking on it, and then alter its **Timeline**.
17. Similarly, you can click on the icon in the **History** tab page to view trend graphs revealing the maximum and minimum CPU used by the host during the last 24 hours (by default). This enables you to better analyze the past trends in CPU usage, and foresee the future trends.

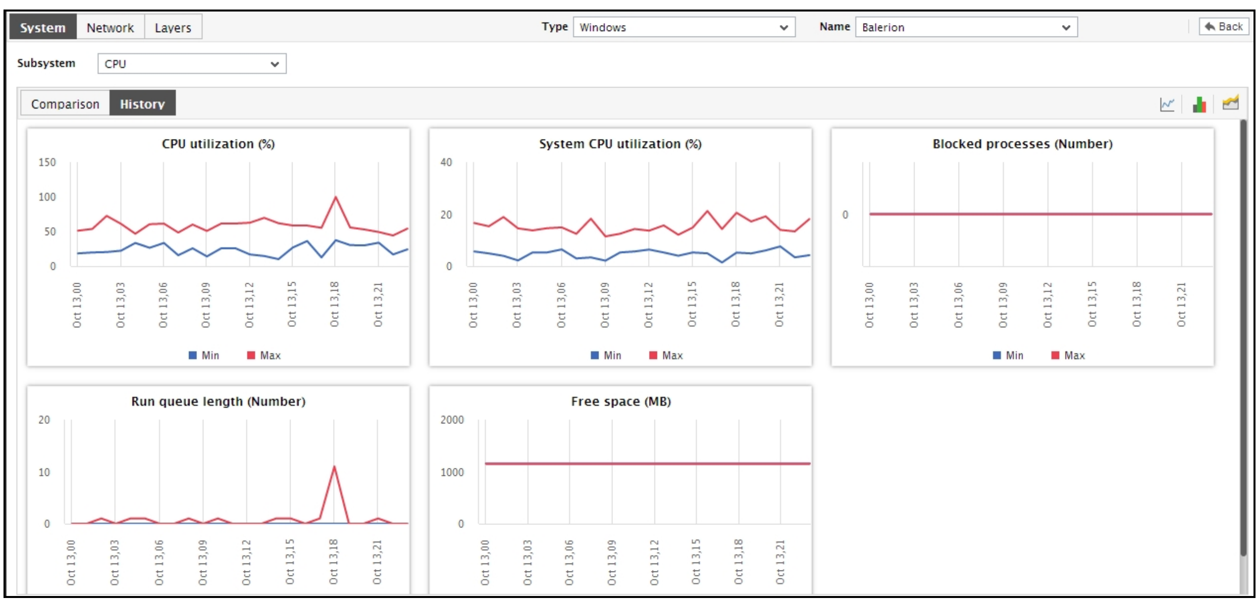

Figure 10.26: The Trend graphs on CPU usage displayed in the History tab page of the CPU dashboard

18. You can even expand the trend graph by clicking on it, and then alter its **Timeline**. By default, the trend graph plots the minimum and maximum values of a measure during the given timeline. In the enlarged

mode, you can change the **Graph** type so that the average values or sum of trend values are plotted in the trend graphs instead.

Note:

In case of descriptor-based tests, the **Summary** and **Trend** graphs displayed in the **History** tab page typically plot the values for a single descriptor alone. To view the graph for another descriptor, pick a descriptor from the drop-down list made available above the corresponding summary/trend graph.

19. At any point in time, you can switch to the measure graphs by clicking on the button.

10.1.3 Disk

To take a closer look at how a host uses the disk space available to it, and to promptly detect probable usage excesses, select the **Disk** option from the **Subsystem** list. Figure 10.27 will then appear.

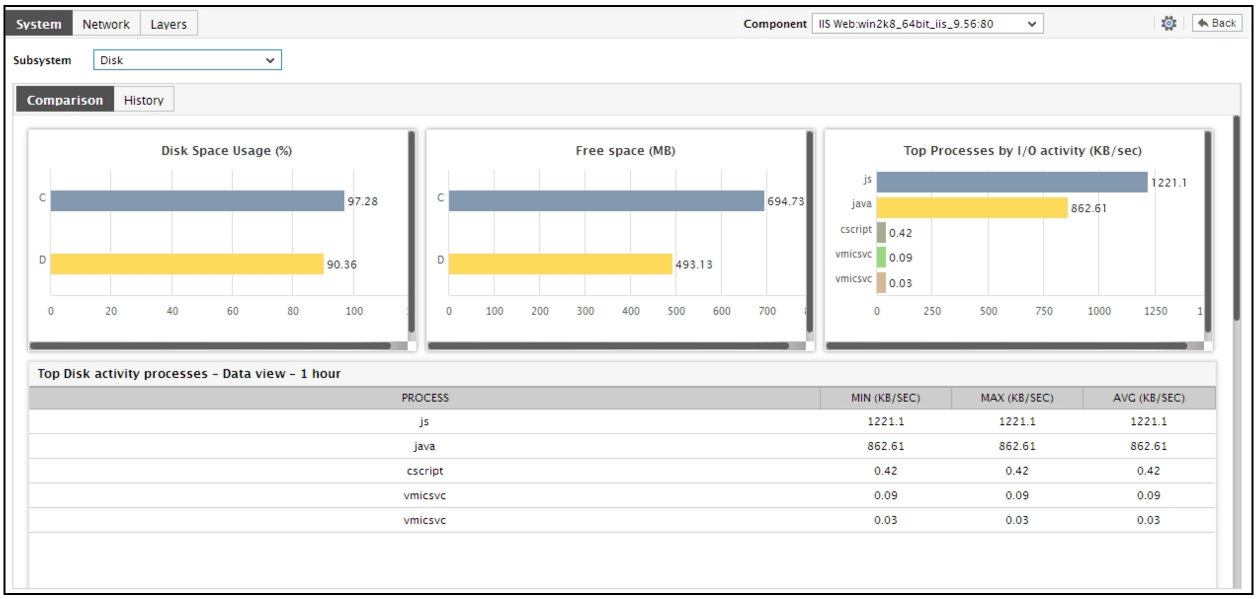

Figure 10.27: The dashboard of the Disk Subsystem

This dashboard provides the following:

1. A quick look at the pie charts displayed in the **Disk Usage Summary** section, can provide you with a clear idea of how well each disk partition on the host is being used currently; the partition that is currently running out of space can be swiftly identified from this section.
2. In addition to the above, the **Comparison** tab page of Figure 10.27 provides a default collection of bar charts. These default charts not only enable you to compare space usage across disk partitions, but also help you evaluate the level of activity on each disk. This way, you can accurately isolate those disks that are currently experiencing a space crunch and also those that are very busy processing requests. If these bar charts reveal that a particular disk is currently experiencing high I/O activity, you can use the **Top Processes by I/O activity** bar chart to zero-in on the process on the host that is responsible for generating this I/O.

3. To view a bar chart in the **Comparison** tab page more clearly, click on it; this enlarges the chart as depicted by Figure 10.28 below.

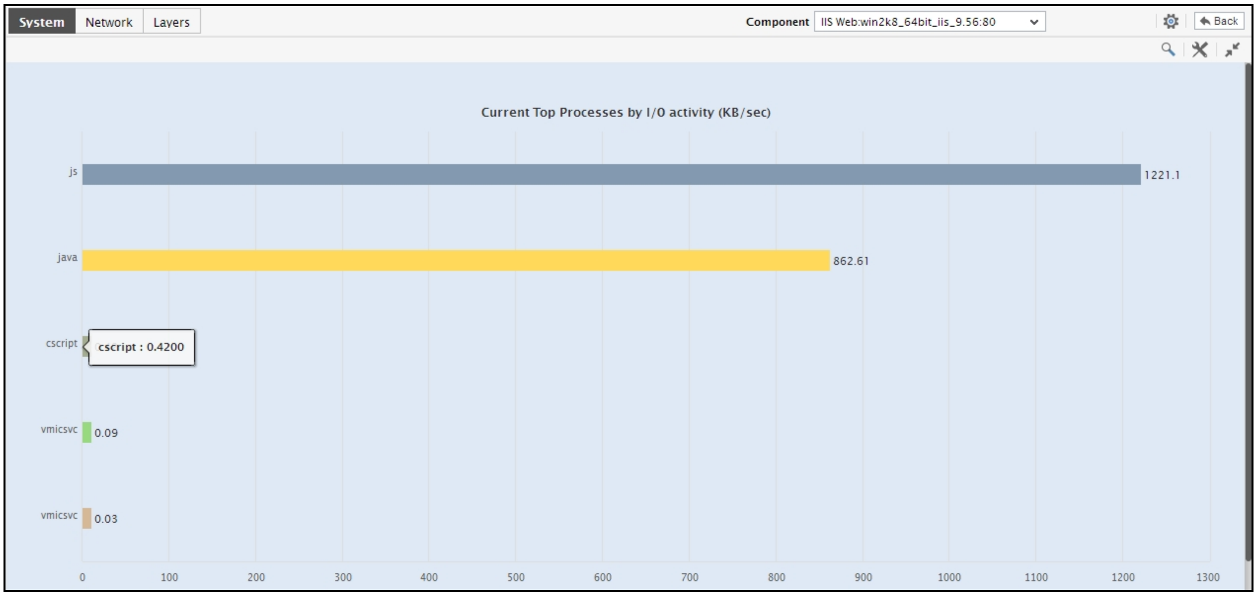

Figure 10.28: The expanded bar chart comparing the disk I/O activity generated by processes on the host

4. By default, the comparison bar charts compare the current disk usage or I/O. If need be, you can alter the timeline of a bar chart, so as to compare the status of disks at some point of time in the past. This enables you to investigate past problems better.

5. If your expanded bar chart appears cluttered owing to a large number of disk partitions/processes, you can easily filter out the 'not-very-important' disk partitions / processes from the chart. Besides enhancing the readability of your bar charts, this enables you to focus on selected descriptors alone. For instance, in the bar chart, you can choose to view only the top-5 processes on the basis of the level of disk I/O activity. To achieve this, simply select the **TOP-5** option from the **Show** list (see Figure 10.29).

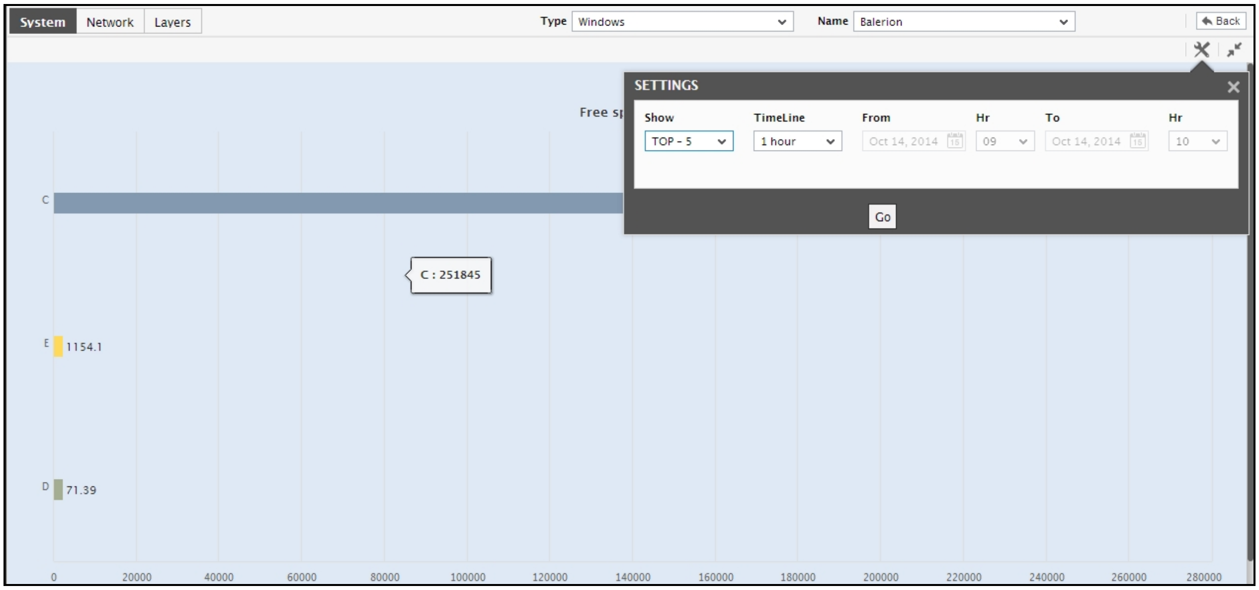

Figure 10.29: Configuring the expanded bar chart to display the top-5 processes with the highest I/O activity

6. Besides, to view more details about the I/O-intensive processes on the host, you can simply invoke the detailed diagnosis page from the enlarged **Top Processes by I/O activity** bar chart.
7. After analyzing current disk performance thoroughly, if you want to engage in an in-depth analysis of the historical disk usage metrics, switch to the **History** tab page by clicking on it.

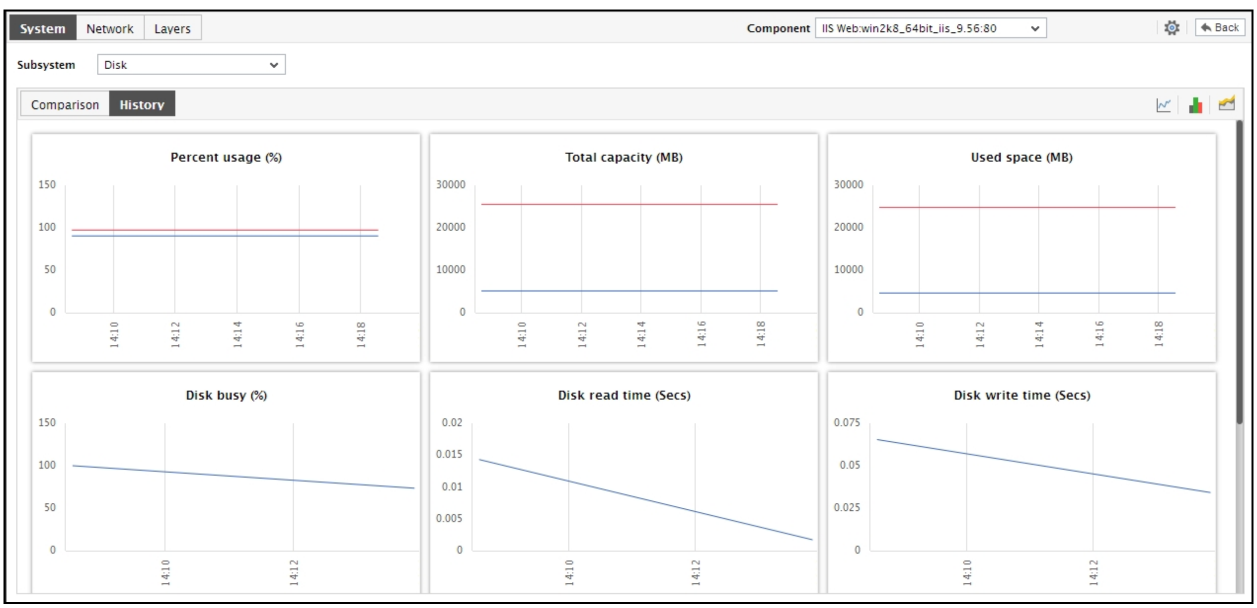

Figure 10.30: The History tab page of the Disk dashboard

8. By default, the **History** tab page provides measure graphs revealing the time-of-day variations in disk usage and disk I/O during the last 24 hours (by default). By carefully studying these measure graphs, you can accurately identify which disk experienced excessive usage / high I/O at what time during the last hour.
9. Like the bar charts, these measure graphs also enlarge when clicked (see Figure 10.30).

Figure 10.31: An enlarged measure graph in the History tab page of the Disk Dashboard

10. Here again, the graph **Timeline** can be altered. Moreover, you can reduce the number of disk partitions for which usage values are plotted in a single graph, by picking a top-n or last-n option from the **Show** list.

11. To assess the overall health of the disk partitions and to perform efficient service level audits on disk usage, you can convert the measure graphs into summary graphs, on-the-fly. For this purpose, click on the icon at the right, top corner of Figure 10.31. 10.1.3 will then appear displaying summary graphs for a pre-configured list of measures.

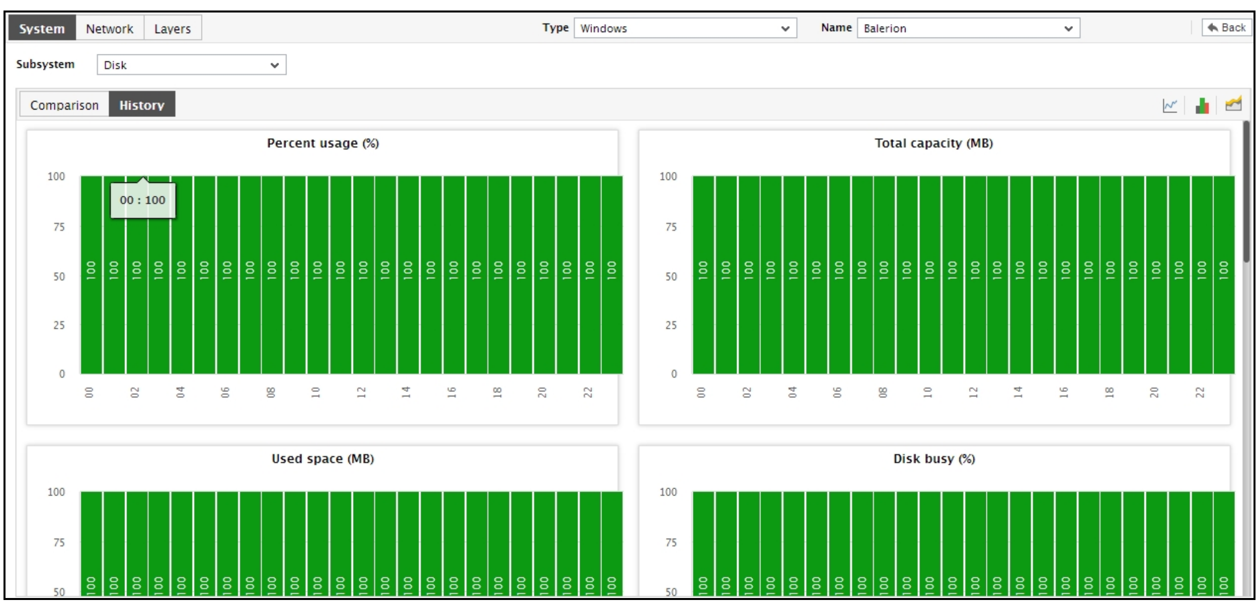

Figure 10.32: Summary graphs of disk usage

12. You can even expand the summary graph by clicking on it, and then alter its **Timeline** and its dimension (**2D / 3D**).

13. Similarly, to analyze past trends in disk usage and accordingly plan the future disk capacity of the host, you can view the historical trend graphs in the **History** tab page, instead of the measure/summary graphs. For this, click on the icon.

14. You can even expand the trend graph by clicking on it, and then alter its **Timeline**. By default, the trend graph plots the minimum and maximum values of a measure during the given timeline. In the enlarged mode, you can view the **Graph** type so that the average values or sum of trend values are plotted in the trend graphs instead.

    Note:

    In case of descriptor-based tests, the **Summary** and **Trend** graphs displayed in the **History** tab page typically plot the values for a single descriptor alone. To view the graph for another descriptor, pick a descriptor from the drop-down list made available above the corresponding summary/trend graph.

15. At any point in time, you can switch to the measure graphs by clicking on the button.

10.1.4 Memory

Apart from CPU and disk usage, you can pick the **Memory** option from the **Subsystem** list to determine how the system has been using its memory resources. Figure 10.33 will then appear.

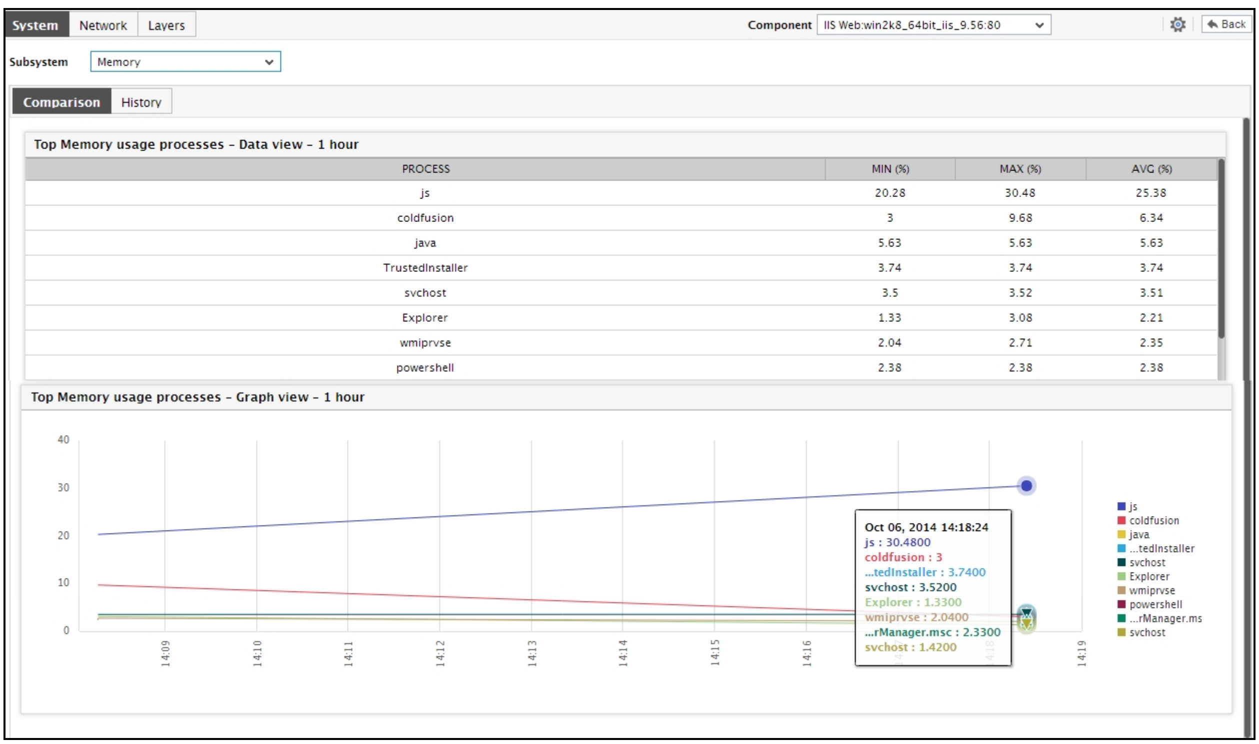

Figure 10.33: Figure 10.33: The System Dashboard for the Memory subsystem

1. Since the dial charts and digital number displays in Figure 10.34 report the current state of a default set of memory usage-related metrics, continuous review of these metrics can provide you with a heads-up on any usage irregularities that may have surfaced recently.

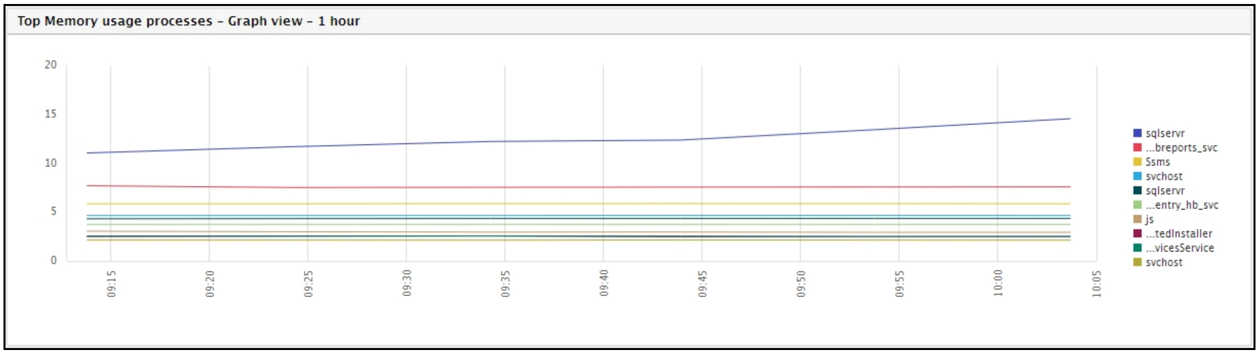

Figure 10.34: Figure 10.34: The enlarged memory usage graph

2. While a comparative analysis of current/1-hour performance can reveal a memory contention that occurred recently, deeper historical analysis is necessary for deducing usage trends and problem patterns. To perform this historical analysis, use the **History** tab page. This tab page displays a number of graphs, plotted for a default period of 24 hours, that reveal the time-of-day variations in the memory usage of the host.

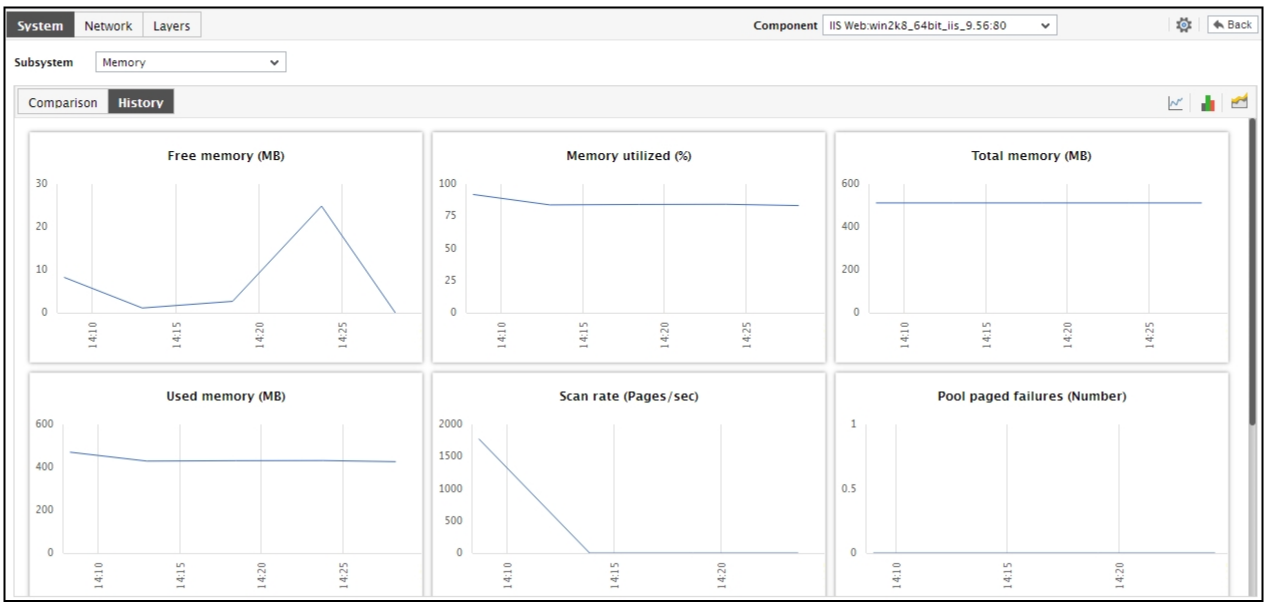

Figure 10.35: Figure 10.35: The History tab page of the Memory subsystem

3. Instead of changing the default timeline, you can change the timeline of the measure graphs on-the-fly too by clicking on the **Timeline** link in Figure 10.35, and can even enlarge a graph by simply clicking on it. The enlarged graph will appear as depicted by Figure 10.36.

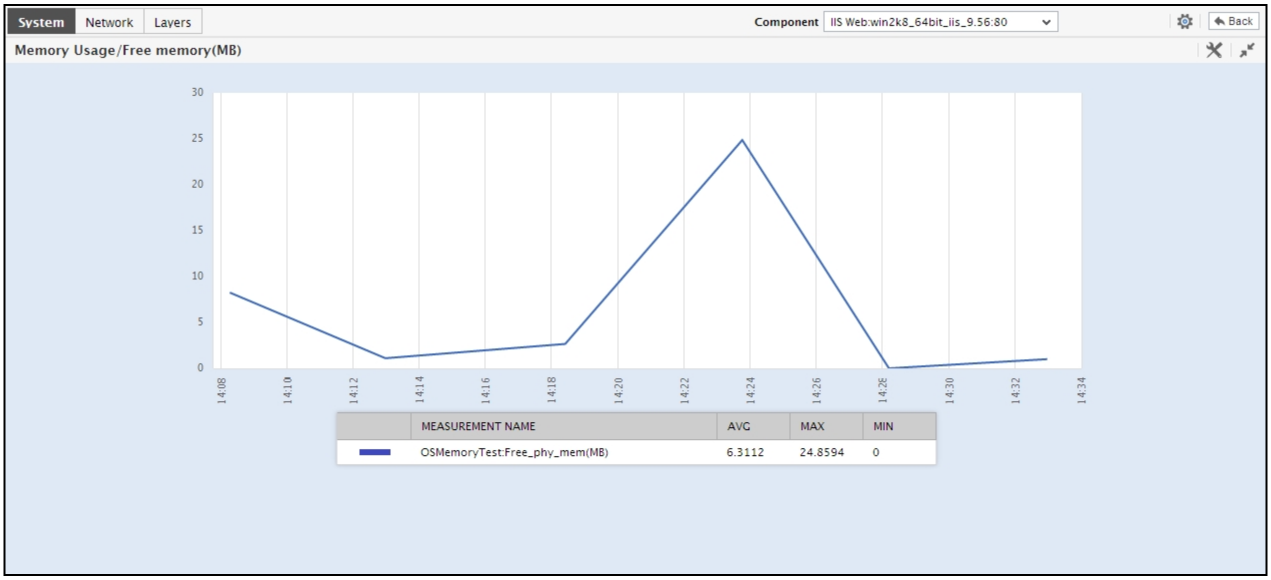

Figure 10.36: Figure 10.36: An enlarged Memory usage measure graph

4. Here again, you can modify the **Timeline** of the graph.

5. If required, you can configure the **History** tab page to display **Summary** graphs instead of the measure graphs. To achieve this, click on the icon at the right, top corner of the **History** tab page of Figure 10.37. Figure 10.38 will then appear revealing summary graphs for the critical memory usage measures. From these graphs, you can easily determine the percentage of time (during the last day by default) the system has experienced memory-related issues. Besides indicating how memory-efficient the system was during the default period, these graphs also enable you to guage the efficiency of the administrative staff in resolving issues that might have surfaced.

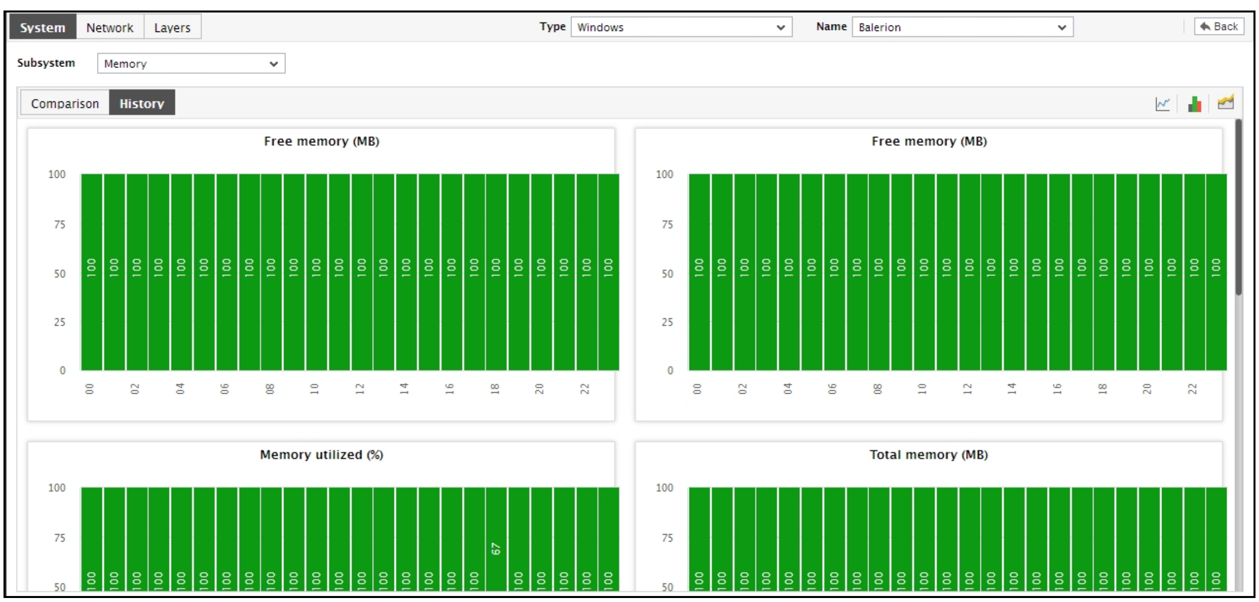

Figure 10.37: Figure 10.37: The History tab page displaying Summary graphs on memory usage

6. You can even expand the summary graph by clicking on it, and then alter its **Timeline** and its dimension (**2D / 3D**).
7. Similarly, with a mere click of a button, you can have the **History** tab page display trend graphs instead of summary/measure graphs. For this, just click on the icon at the right, top corner of Figure 10.39.

Figure 10.38: Figure 10.38: The History tab page displaying trend graphs on memory usage

8. You can even expand the trend graph by clicking on it, and then alter its **Timeline**. By default, the trend graph plots the minimum and maximum values of a measure during the given timeline. In the enlarged mode, you can change the **Graph** type so that the average values or sum of trend values are plotted in the trend graphs instead.

   Note:

   In case of descriptor-based tests, the **Summary** and **Trend** graphs displayed in the **History** tab page typically plot the values for a single descriptor alone. To view the graph for another descriptor, pick a descriptor from the drop-down list made available above the corresponding summary/trend graph.

9. At any point in time, you can switch to the measure graphs by clicking on the button.

10.1.5 Uptime

To view the uptime details of the system, you can select the **Uptime** option from the **Subsystem** list. Figure 10.39 then appears.

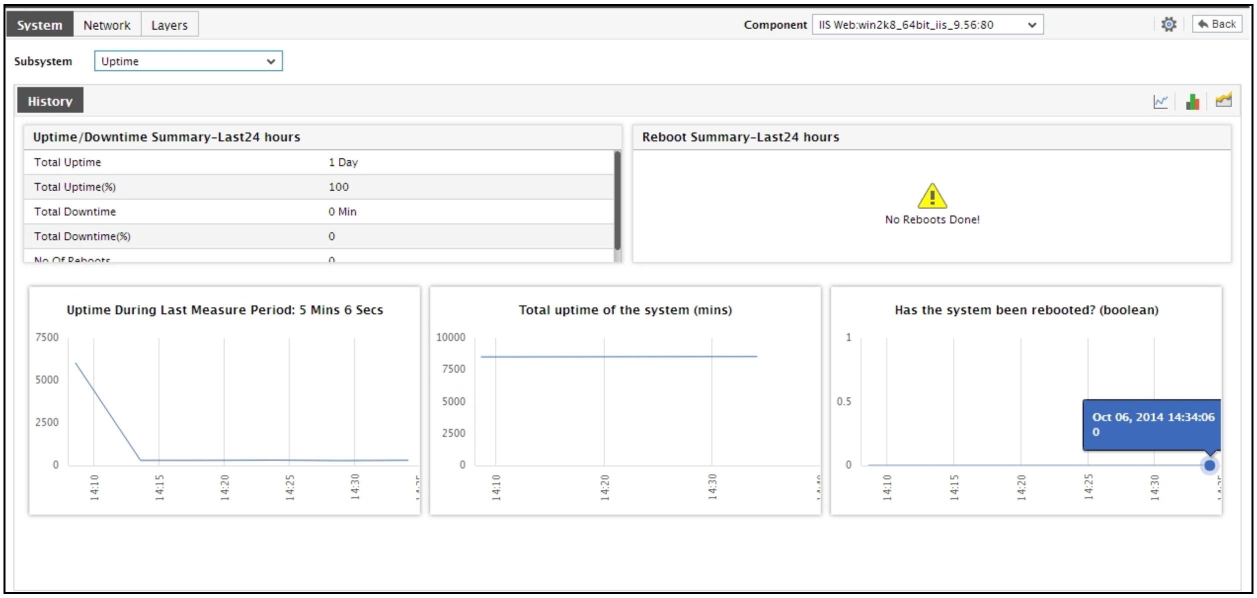

Figure 10.39: The System Uptime Dashboard

This dashboard reveals the following:

1. The first section of the dashboard reveals the total time for which the system has been up and running since it was last rebooted. The **Uptime/DownTime Summary** section provides a quick summary of the availability of the system during the last 24 hours (by default) - the details include: the total duration for which the system has been up and running in the last 24 hours, the percentage uptime, the total duration (in the last 24 hours) for which the system was down, the percentage of downtime, and number of reboots during the last 24 hours. Using these details, you can determine whether the agreed uptime levels for the system were met or not, and if not, how much is the system falling short of its desired performance levels.
2. You can also infer whether the system experienced any reboots during the last 24 hours (by default). To know more about each reboot, refer to the **Reboot Summary** section. For every reboot that occurred in the last 24 hours (by default), this section reveals when the system was shutdown, when the reboot occurred, and how long did the system remain down until it was rebooted. This clearly indicates the frequency of the reboots, and helps determine whether such reboots were scheduled or unexpected.
3. This will be followed by a default set of graphs indicating how long during the last 24 hours (by default) the system has been up, and whether the reboots scheduled for the system have occurred during the last hour or not.
4. To view the measure graphs clearly, click on a graph of interest to you to enlarge it (see Figure 10.40); you can alter the **Timeline** in the enlarged graph. Using these graphs, breaks in the availability of the system and failure of reboot schedules can be accurately identified and investigated.

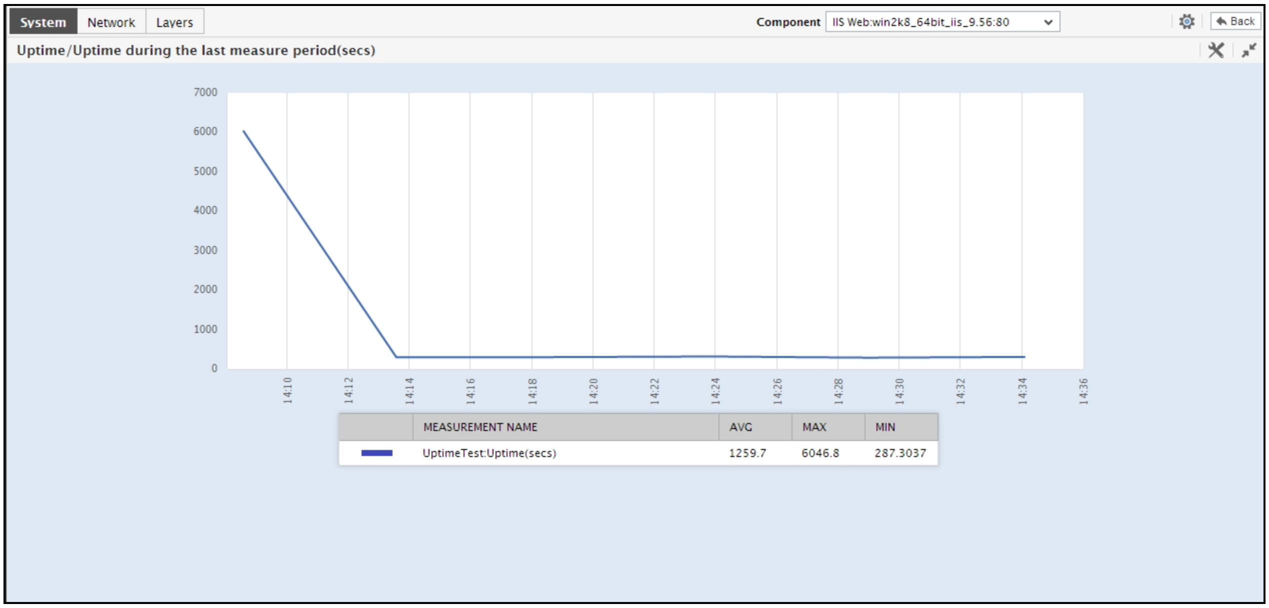

Figure 10.40: The enlarged Uptime graph

5. Click on the icon at the right, top corner of the **History** section to view summary graphs using which you can effectively perform service level audits on a host, based on the duration of their availability. Determine the percentage of time for which the host was operational during the last day (by default), and also be notified of reboots that might have occurred on the host during the default timeline. If required, you can click on the **Timeline** link to alter the graph timeline.

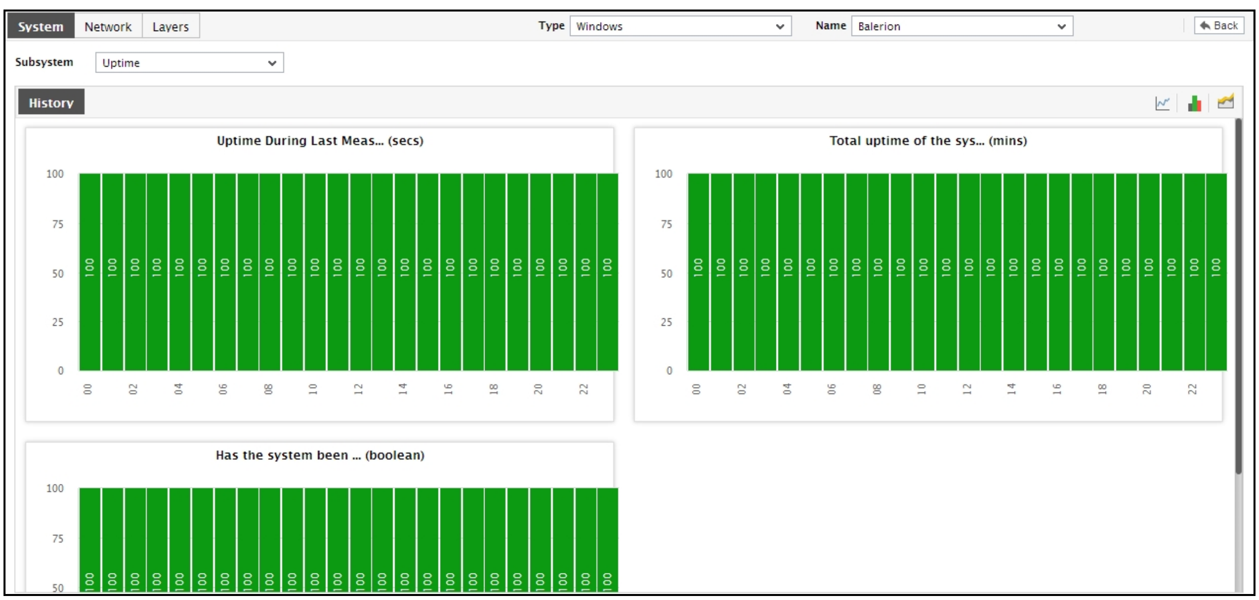

Figure 10.41: The Summary graphs in the System Uptime Dashboard

6. Clicking on the icon in the **History** section will display trend graphs on system uptime; these trends reveal when during the last 24 hours (by default) uptime was the lowest, and when reboots failed (see Figure 10.42). If required, you can click on the graph to expand it and alter its **Timeline**.

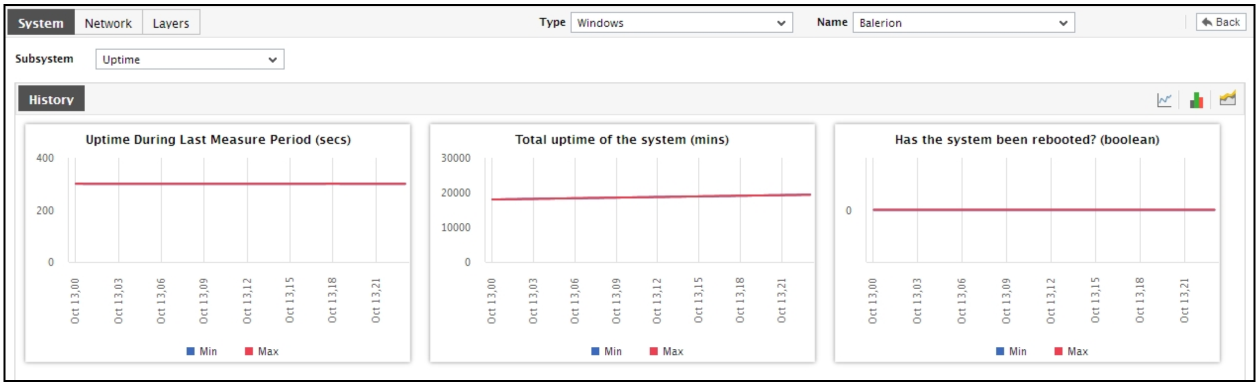

Figure 10.42: The System Uptime Dashboard with Trend graphs

7. You can even expand the trend graph by clicking on it, and then alter its **Timeline**. By default, the trend graph plots the minimum and maximum values of a measure during the given timeline. In the enlarged mode, you can change the **Graph** type so that the average values or sum of trend values are plotted in the trend graphs instead.

8. At any point in time, you can switch to the measure graphs by clicking on the button.

10.2 The Network Dashboard

Clicking on the **Network** tab page in Figure 10.43 will reveal the **Network Dashboard**, which allows you to zoom into the performance and problems pertaining to the **Network** layer and related layers of a target application/device. Using this dashboard, you can:

- Determine whether/not the application/device currently experiences / has in the past experienced network-related issues;
- Accurately identify the network parameters that are currently failing;
- Understand the current network configuration;
- Analyze network performance over time, study the trends in network connectivity and usage, and accurately deduce problem/performance patterns.
- Identify persistent problems with network health and the network-related layers responsible for the same;

The contents of the **Network Dashboard** and the subsystems it offers for analysis could slightly vary depending upon whether the target is an application or a network device. While the **Network Dashboard** of a host/application enables you to focus on both the network and TCP connections handled by the target, the same for a network device sheds light on the network connectivity of the device and the traffic handled by the network interfaces supported by the device. Accordingly, the **Network Dashboard** of an application/host offers **Network** and **Tcp** as its **Subsystems**, and that of a network device offers **Network** and **NetworkInterfaces** as its **Subsystems**. If the target application is a Windows-based one, then the **Subsystems** list will include an additional **WindowsInterfaces** option, which provides the performance information related to the traffic handled and bandwidth used by the network interfaces supported by the system.

The **Network Dashboard** of Hyper-V servers on the other hand, will additionally support a **Hyper-V Switches** and a **Hyper-V Network Adapters** sub-system. Similarly, the **Network Dasboard** of a vSphere/ESX server will include an additional **VirtualNetwork** sub-system.

The sections that follow will discuss each of these subsystems at length.

10.2.1 Overview

By default, the **Overview** option is chosen from the **Subsystem** list of Figure 10.43. As the name suggests, the **Overview** dashboard provides an *all-round* view of the network health of a target application/device.

Figure 10.43 depicts the **Overview Dashboard** for a network device.

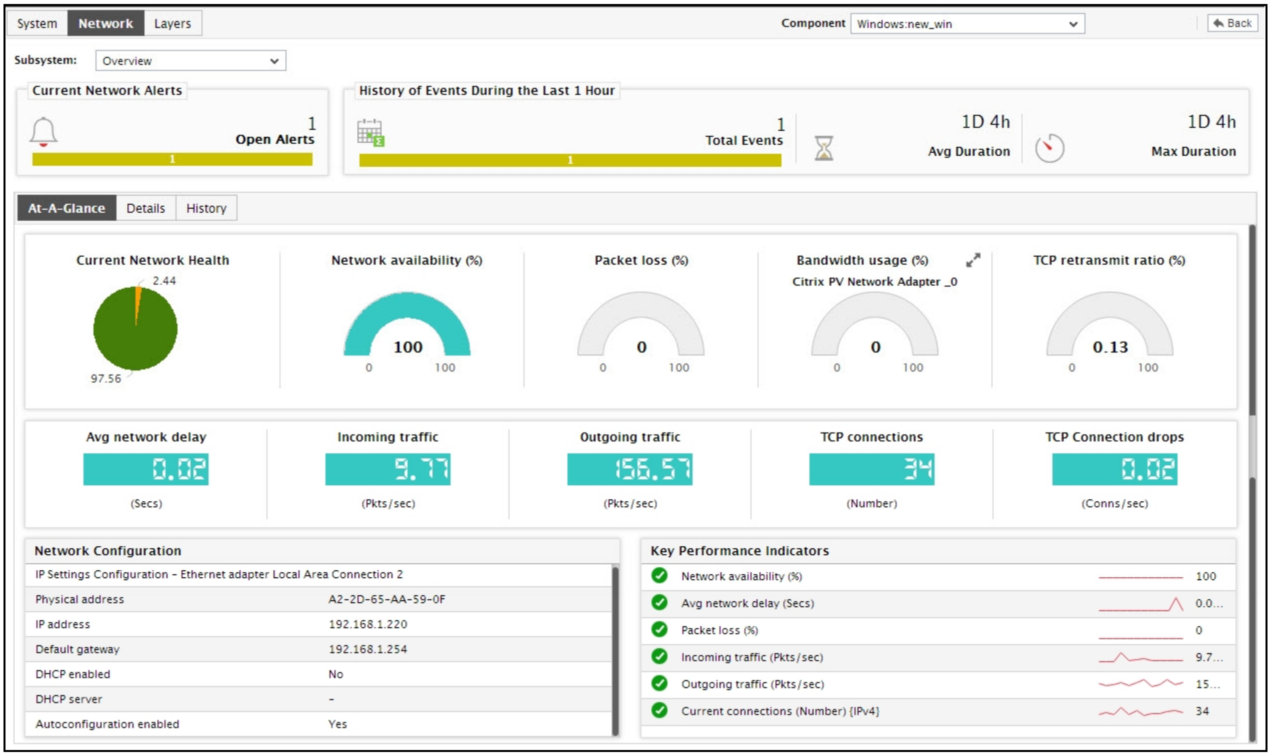

Figure 10.43: The Network Dashboard of a network device

For a monitored host, the **Overview** dashboard reveals the following:

1. Using the **Current Network Alerts** section of Figure 10.43, you can instantly understand how many problems of what severity are currently affecting the health of the **Network** layer of the device. Clicking on a non-zero number here will open Figure 10.43, which will list the network-related alarms of the priority clicked on.

2. To determine the quality of the network links leading to the network device during the last 1 hour, view the **History of Events** section; this section displays a bar graph that provides a quick look at the number and priority of network problems faced by the target application during the last 1 hour. In addition, the average and maximum duration for which these issues remained open will also be available in this secton; this enables you to assess the efficiency of your administrative staff.

3. For more details about the historical problems, click on any of the bars in the bar graph. The **HISTORY OF**

ALARMS page of Figure 10.44 will then appear, listing all the network-related events that were captured by the eG agent on the target network device, during the last 24 hours (by default). This information enables you to understand what type of problems were faced by the target device during the default timeline, and how long each problem remained.

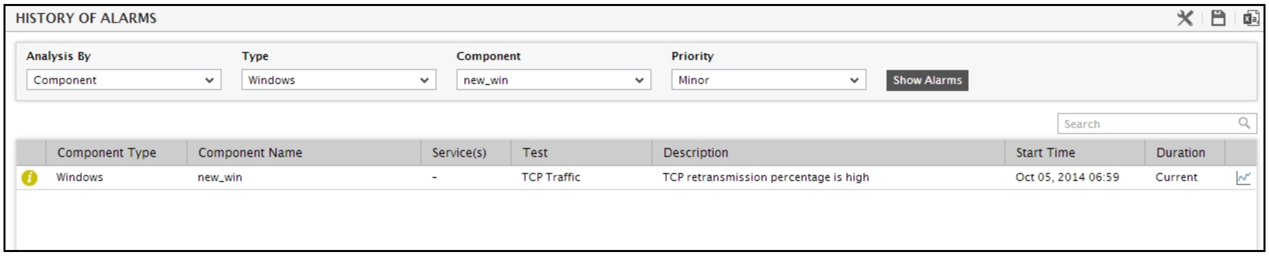

Figure 10.44: The History of Alarms page

4. For a network device, this tab page reveals, at a single glance, the current network health, network latencies (if any), and the status of the network interfaces supported by the device. Besides enabling you to accurately detect sudden changes in the state of the device, this section also helps you nail the root-cause of this state change; moreover, you can even identify network interfaces that are performing poorly, with the help of this section.

5. You can use the **Current Network Status** section in the **At-A-Glance** tab page to know how problem-prone your network is; the pie chart here indicates the percentage of current network-related measures that are in varying states of activity. Clicking on a slice in this pie chart will once again lead you to the **Event History** page.

6. This will be followed by dial charts and digital displays that will update you with the status of a pre-configured list of metrics; these metrics reveal the following:

   - Whether the network device is currently available over the network or not;
   - Has the device experienced any loss of data packets? If so, to what extent?
   - How long has the device been up? Was the device rebooted anytime after it was started?
   - Was there any delay in connecting to the device? How significant is the delay? - in the event of a high network latency, you can use the **Routers by HopDelay** table to figure out where - i.e., at which hop - the maximum delay occurred.
   - What is the maximum rate at which data is currently transmitted and received by the network interfaces supported by the device?
   - What is the maximum speed at which the network interfaces currently operate?

   Note:

   If you have configured one/more measures of a descriptor-based test to be displayed as a dial chart, then, in real-time, the descriptor that is in an abnormal state or is currently reporting the maximum value for that measure will be represented in the dial chart. You can view the dial charts pertaining to the other descriptors, by clicking on the **More** button that appears alongside a dial chart in the dashboard.

7. Clicking on a dial chart/digital display in the dashboard will lead you to the layer model page that will reveal

the exact layer and test that reports the measure represented by the dial chart/digital display.

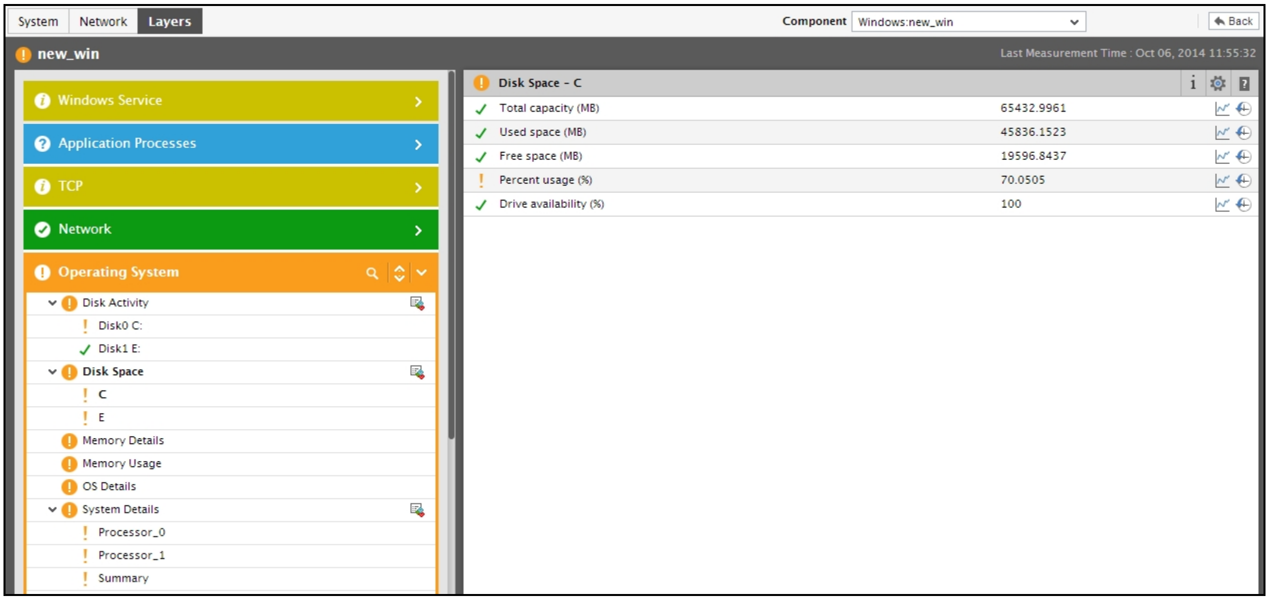

Figure 10.45: The page that appears when a dial/digital chart in the Network Overview Dashboard is clicked

8. If users are experiencing delays while connecting to the component via the network, then you can use the **Routers by Hop Delay (ms)** section to nail the root-cause of such a network latency. This section details the hop-by-hop connectivity and delay, and thus help isolate the exact hop at which a network delay occurred.

   Note:

   If you have configured more than two dial charts for the **Network Overview** dashboard using the **Dashboard Settings** window of Figure 4.140, then the **Routers by Hop delay (ms)** section will not appear in the **At-A-Glance** tab page.

9. The **Network Configuration** section displays the current network configuration, and enables you to ascertain whether a change in configuration could have contributed to the network problems at hand.

10. By closely monitoring the status of these metrics, you can be instantly updated with critical network issues. Clicking on a measure listed in the **Key Performance Indicators** section will lead you to the **Layer Model** tab page of the application (see Figure 10.45), which reveals the exact layer and test that reported the measure.

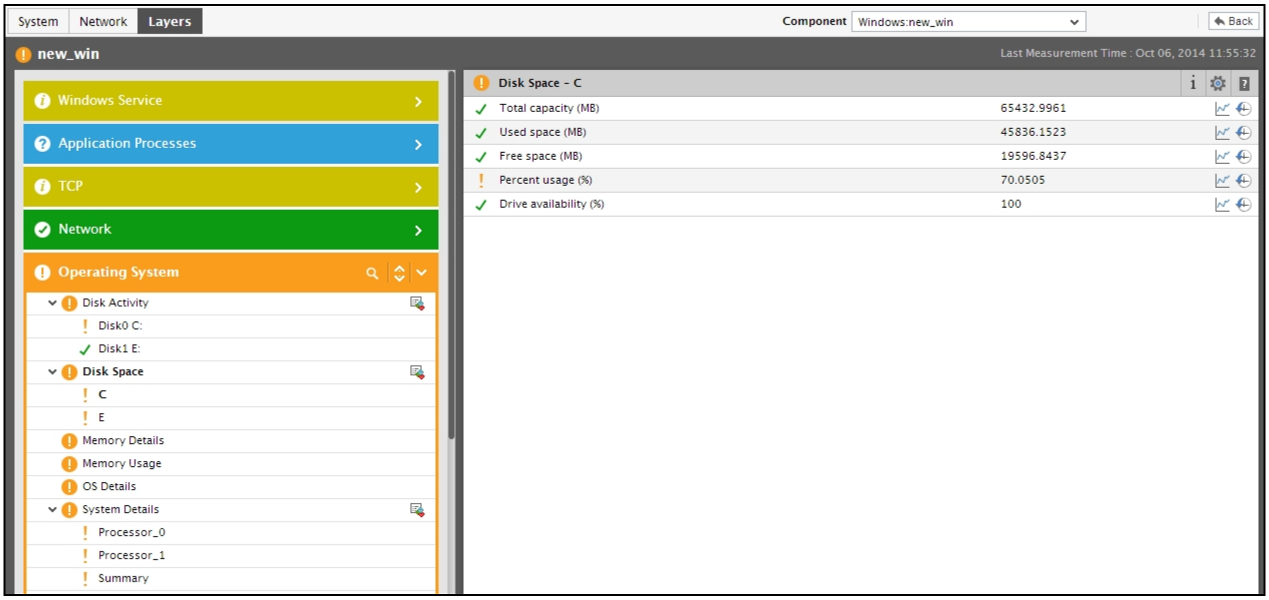

Figure 10.46: Clicking on a key performance indicator in the Network Overview Dashboard

11. This way, from a single console, you can receive rapid, real-time updates on the overall network status of your network devices, and instantly identify the root-cause of such network-related anomalies. In order to zoom into the performance of network interfaces in particular, switch to the **Details** tab page by clicking on it. Figure 10.47 will then appear.

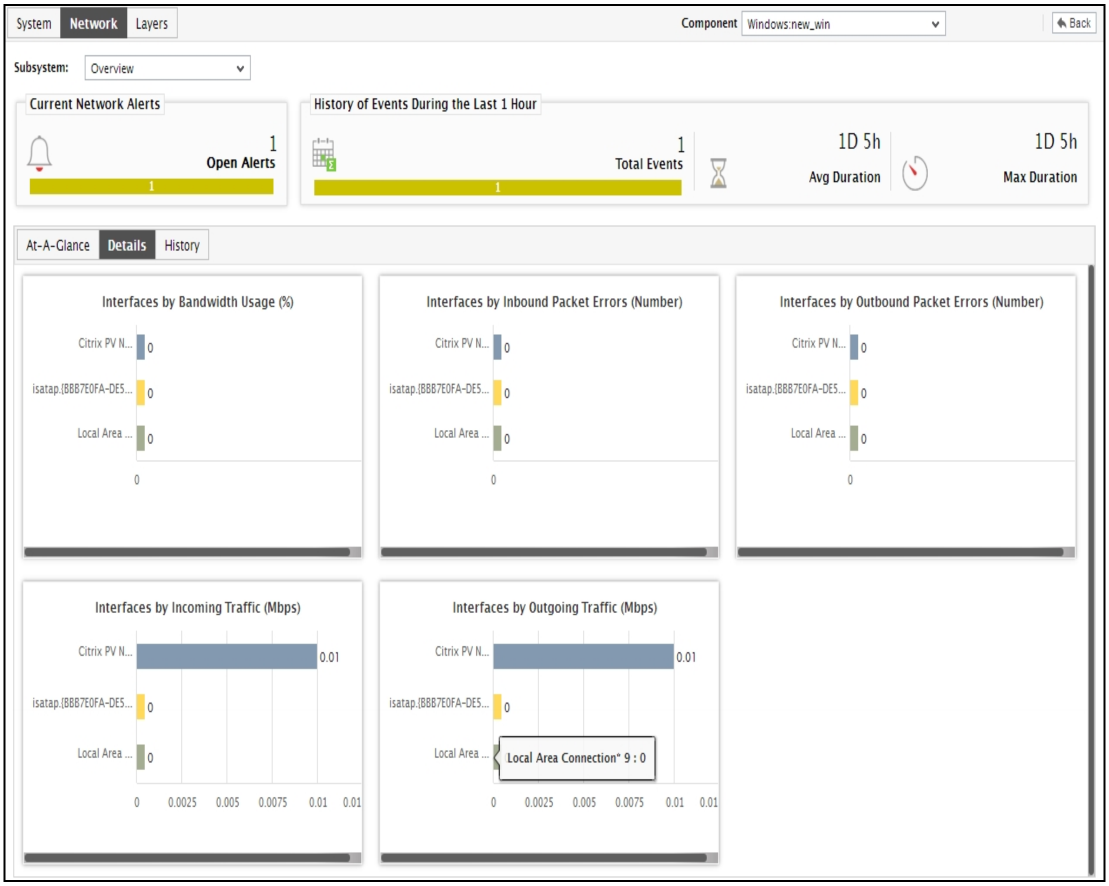

Figure 10.47: The Details tab page of the Network Dashboard of a network device

12. The **Details** tab page provides a default set of bar charts, each of which lists the top-10 network interfaces in a particular performance arena. Using these default bar charts, you can instantly figure out the following:

    - Which network interface is currently available?
    - Which network interface is operating with the maximum speed?
    - Which network interface is utilizing the maximum bandwith?
    - The incoming and outgoing traffic to which network interface is very high currently?

13. In addition to the default/user-configured bar graphs, the **Details** tab page includes a **Routers by HopDelay** bar chart that graphically depicts the hop-by-hop delay, so that you can easily identify the hop at which significant delay has occurred.

14. By default, all the bar charts in the **Details** tab page pertain to the current period. However, sometimes, you might want to go back in time, so as to investigate a past problem and identify its source. For this purpose, click on the corresponding bar chart in the **Details** tab page. Doing so will invoke Figure 10.47.

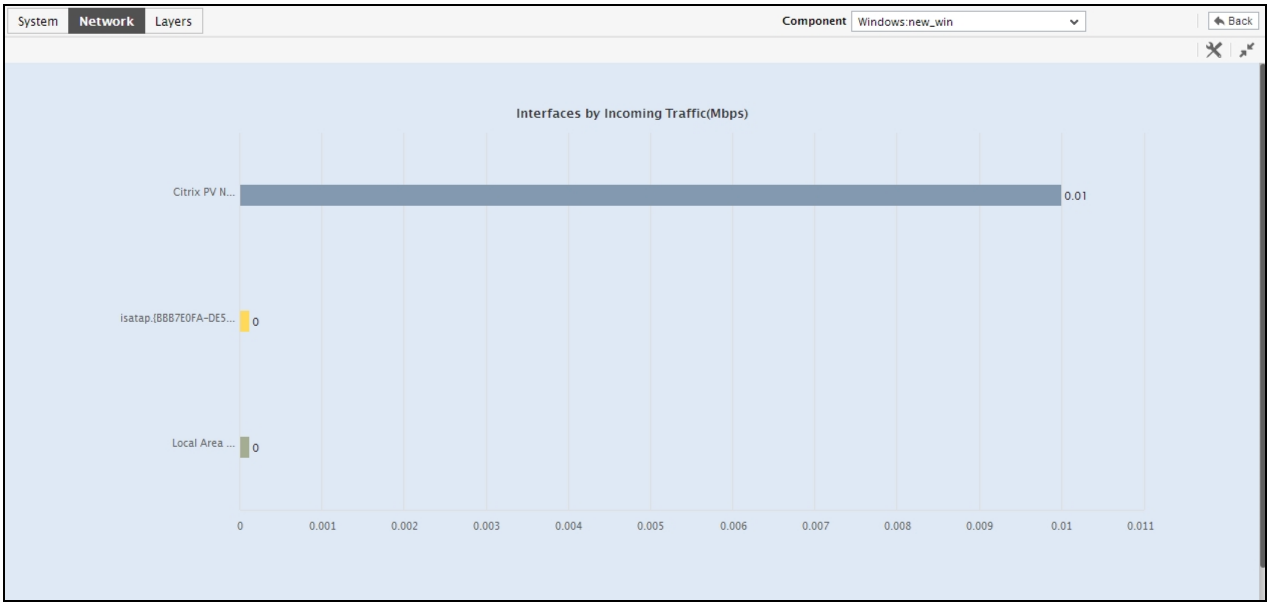

Figure 10.48: The enlarged top-10 bar chart in the Details tab page of a Network Dashboard

15. The enlarged bar chart, by default, lists all the network interfaces that are supported by the target server/network device. To view the performance of a specific set of interfaces alone, select a **TOP-N** or **LAST-N** option from the **Show** list in the settings option.
16. Though the **Details** tab page enables you to focus on both the current and past performance of network interfaces, to proactively detect any type of potential anomaly, the past performance of the network in its entirety needs to be analyzed. For this purpose, switch to the **History** tab page by clicking on it.

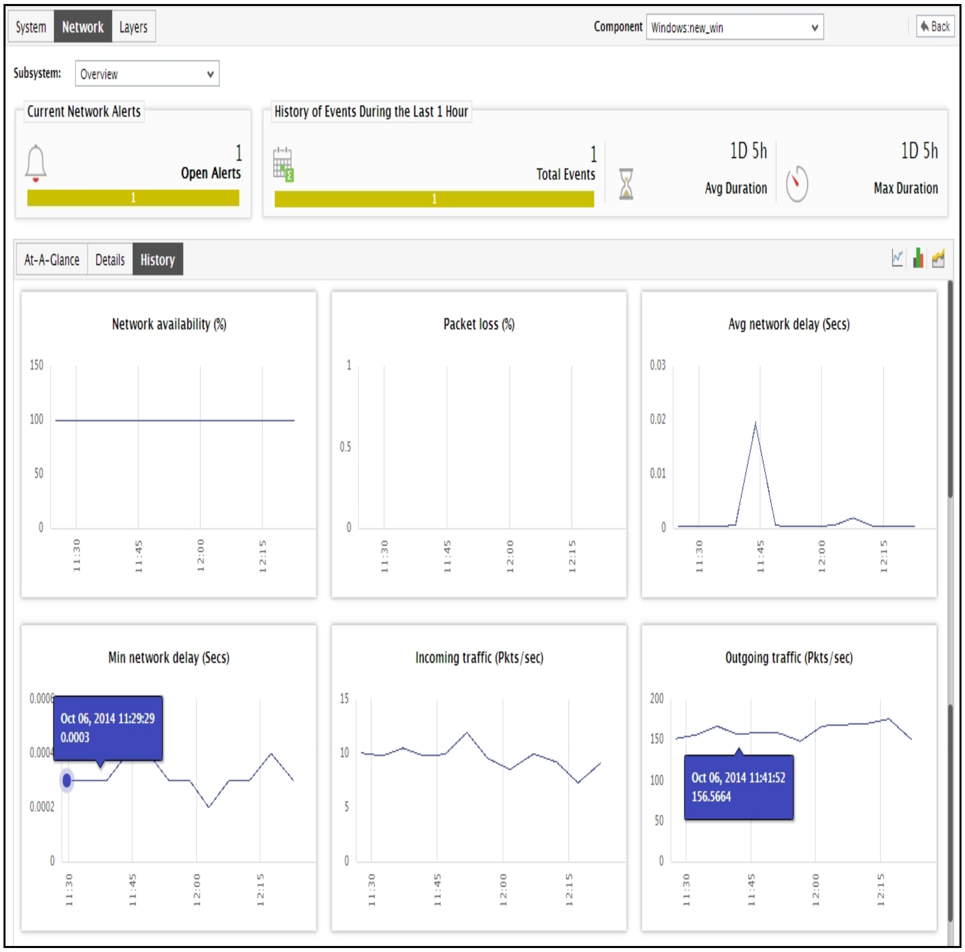

Figure 10.49: The History tab page of the Network dashboard of a network device

17. Figure 10.49 that appears reveals measure graphs for a pre-configured set of performance metrics; these graphs, which are typically plotted for a default period of 24 hours, indicate the following:

    - Were there any breaks in the availability of the network connection to the device during the last 24 hours?
    - Were any data packets lost in transit during the last 24 hours?
    - Was any significant delay detected in connecting to the device during the last 24 hours?
    - How was the data traffic to and from the device during the last 24 hours?
    - Were any TCP connections dropped by the target during the default timeline?
    - How busy was the target, in terms of TCP connections, during the default time period?

    You can click on the graph to expand it and view its contents clearly.

18. If need be, you can view **Summary** graphs in the **History** tab page to figure out the percentage of time (during the last day by default) the device experienced network-related issues. To achieve this, click on the icon at the right, top corner of Figure 10.49. Figure 10.50 will then appear displaying summary graphs plotted for a default period of 24 hours.

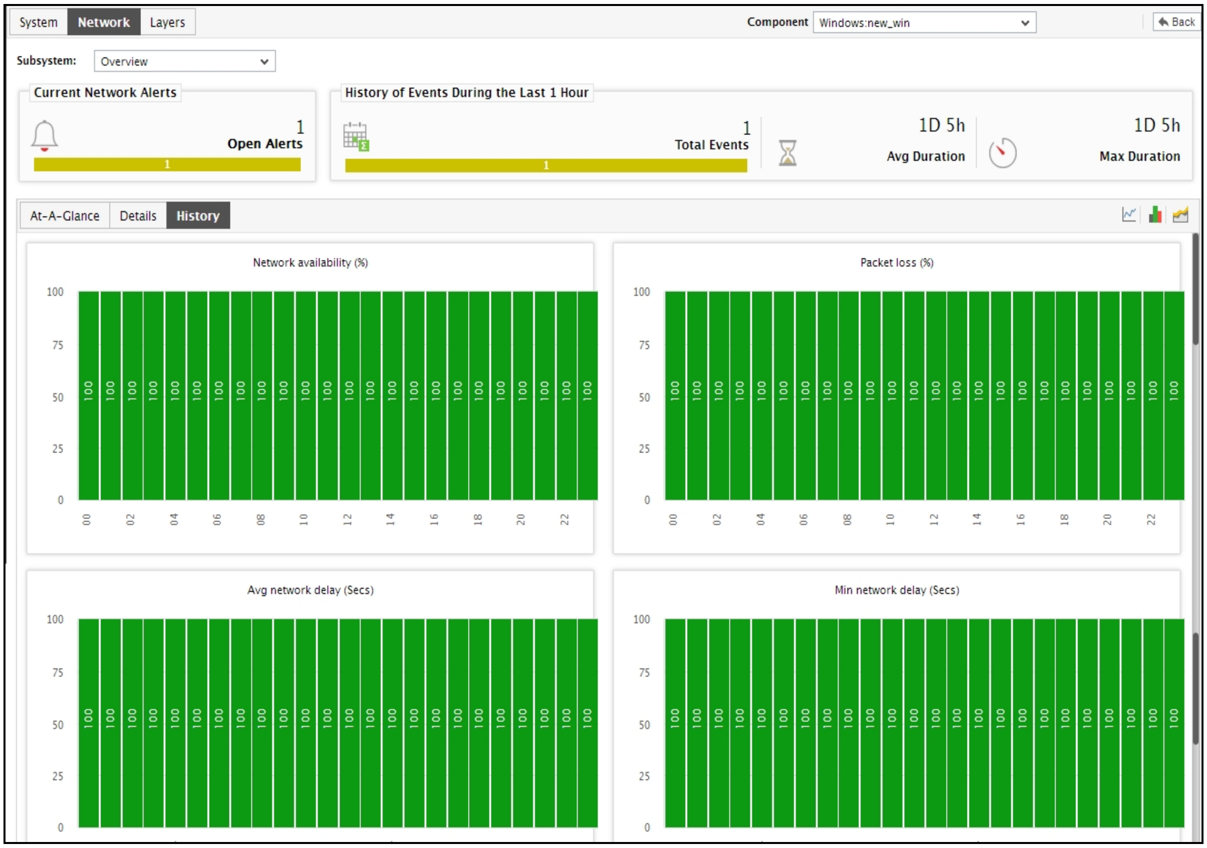

Figure 10.50: The History tab page displaying summary graphs on network performance

19. Similarly, you can choose to analyze the past trends in network performance using the **History** tab page, instead of time-of-day variations or performance summaries, by clicking on the icon at the right, top corner of Figure 10.50. This will invoke Figure 10.51, which will reveal trend graphs that capture the minimum and maximum values that pre-configured measures registered during the default period of 24 hours.

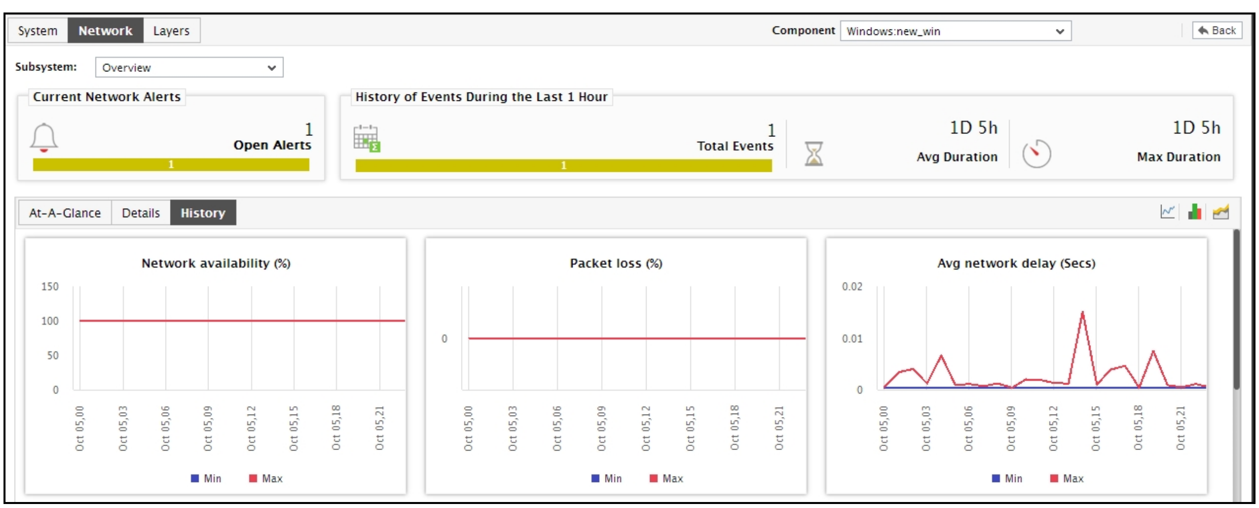

Figure 10.51: Figure 10.51: The History tab page displaying trend graphs on network performance

20. You can even expand the trend graph by clicking on it (see Figure 10.52).

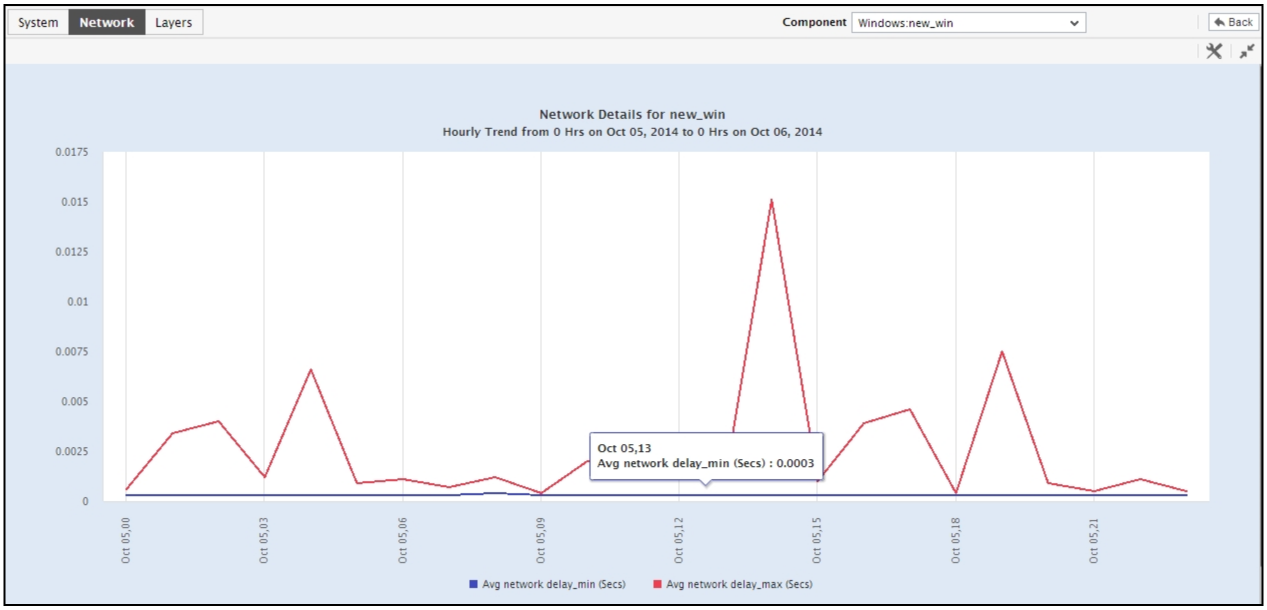

Figure 10.52: An enlarged Network

21. In the enlarged mode, you can even alter the **Timeline** of trend graphs. By default, the trend graph plots the minimum and maximum values of a measure during the given timeline. In the enlarged mode, you can change the **Graph** type so that the average values or sum of trend values are plotted in the trend graphs instead.

    Note:

    In case of descriptor-based tests, the **Summary** and **Trend** graphs displayed in the **History** tab page typically plot the values for a single descriptor alone. To view the graph for another descriptor, pick a descriptor from the drop-down list made available above the corresponding summary/trend graph.

22. At any point in time, you can switch to the measure graphs by clicking on the button.

10.2.2 Network

If you want to strictly focus on the current and past issues related to the availability of the network connection to the device/server, and network latencies / packet losses experienced by the device/server, pick the **Network** option from the **Subsystem** list. Figure 10.52 will then appear.

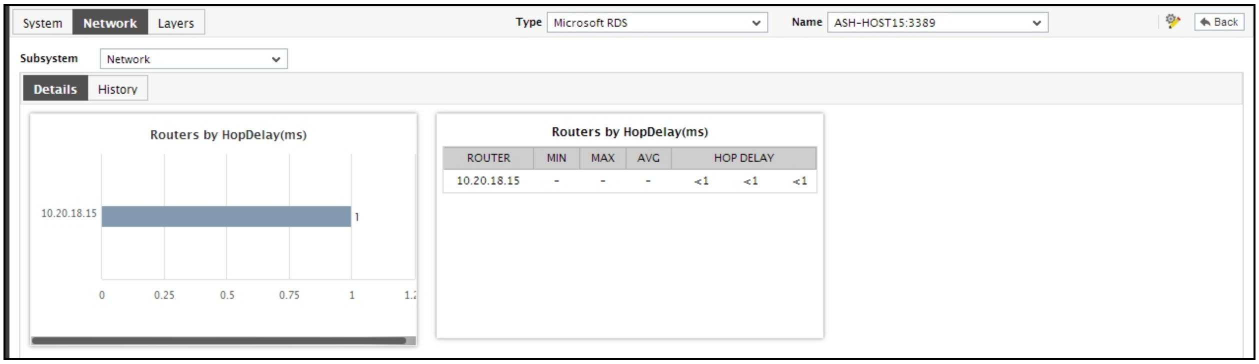

Figure 10.53: The dashboard of the Network subsystem

For a network device, this dashboard reveals the following:

1. With the help of the digital displays available in Figure 10.52, which report the current network latency values for the device, you can proactively determine network-level slowdowns (if any). While digital displays are available for a default set of measures, you can configure additional digital displays or remove any of the existing displays by following the steps discussed below:
2. Below the digital graphs is the **Details** tab page. If the target network device experiences a high network latency, then, using the **Routers by Hop Delay** bar graph and table available in the **Details** tab page, you can view the hop-by-hop connectivity and delay, zero-in" on the exact hop at which the delay has occurred, probe into the root-cause of the delay, and resolve the issue, so as to optimize network performance. You can expand the bar graph by clicking on it.
3. The **History** tab page, on the other hand, displays measure graphs by default, which indicate the time-of-day variations in the availability and overall network health of the device during the last 24 hours. You can click on a graph to enlarge it and view it clearly. You can even click on **Timeline** to modify the graph timeline, and thus analyze network performance over a longer time period.

Figure 10.54: History tab page displaying measure graphs in the dashboard of the Network sub-system

4. To view the percentage of time during the last 24 hours for which the network performance was affected by issues, click on the icon at the right, top corner of Figure 10.54.

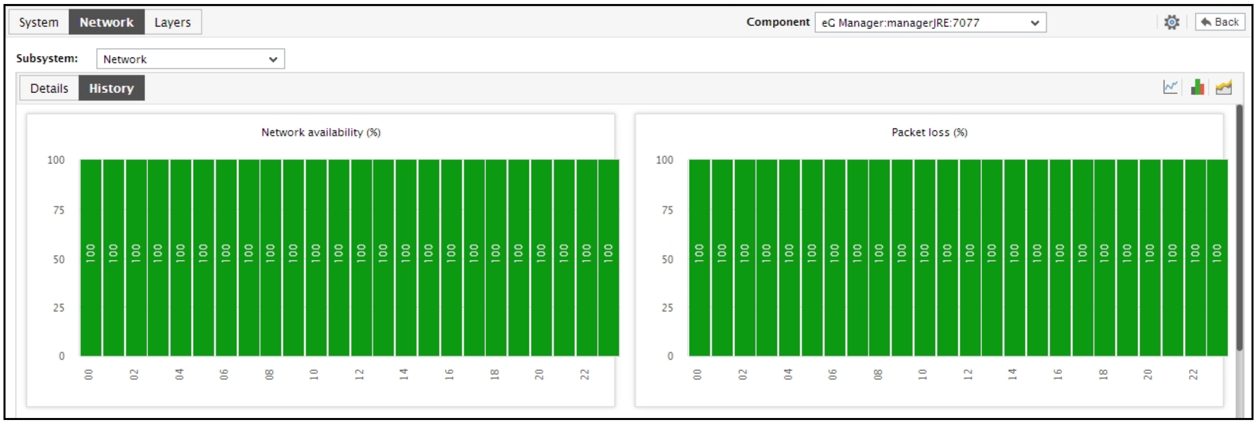

Figure 10.55: The Summary graphs in the dashboard of the Network subsystem

5. Using the graphs in Figure 10.54, you can effectively perform service level audits and detect when and what type of network issues caused the agreed-upon service levels to be compromised.
6. Similarly, click on the icon at the right, top corner of the **History** tab page in Figure 10.54 to view and analyze the past trends in network performance. By default, the trend graphs will pertain to the last 24 hours.

Figure 10.56: The Trend graphs in the dashboard of the Network subsystem

7. Using these trend graphs, you can determine when network performance peaked and when it hit rock bottom - this way, you can easily infer how network performance has varied during the last 24 hours, and thus receive a heads-up on potential network anomalies.
8. To change the timeline of a single graph, click on the graph to enlarge it.
9. Also, an enlarged summary/trend graph allows you to alter the graph **Duration** - i.e., view the daily or weekly summary/trend information, instead of the default hourly data in the graphs.
10. Also, by default, the trend graphs in the **History** tab page plot the minimum and maximum values of a measure during the given timeline. In enlarged trend graphs, you can change the **Graph** type so that the average values or sum of trend values are plotted in the trend graphs instead.

    Note:

    In case of descriptor-based tests, the **Summary** and **Trend** graphs displayed in the **History** tab page typically plot the values for a single descriptor alone. To view the graph for another descriptor, pick a descriptor from the drop-down list made available above the corresponding summary/trend graph.
11. At any point in time, you can switch to the measure graphs by clicking on the  button.

As stated earlier, the **Network Subsystem Dashboard** of a network device provides measure graphs that track the availability of the device over the network, network latencies, and packet loss experienced by the device during communication. For a server/application however, the same **Network Subsystem Dashboard** will additionally display the following:

- Measure graphs tracking the incoming and outgoing traffic of the target server during the last 24 hours (by default); using these graphs, you can determine when during the last 24 hours the network activity to and from the server was very high.

- A bar graph depicting the network delay experienced at every hop; using this bar graph, you can accurately isolate the hop at which maximum network latency has occurred.
- A table depicting the hop-by-hop delay.

10.2.3 Tcp

To zoom into the health of the TCP connections to and from a server, select **Tcp** from the **Subsystem** list. Doing so invokes Figure 10.57.

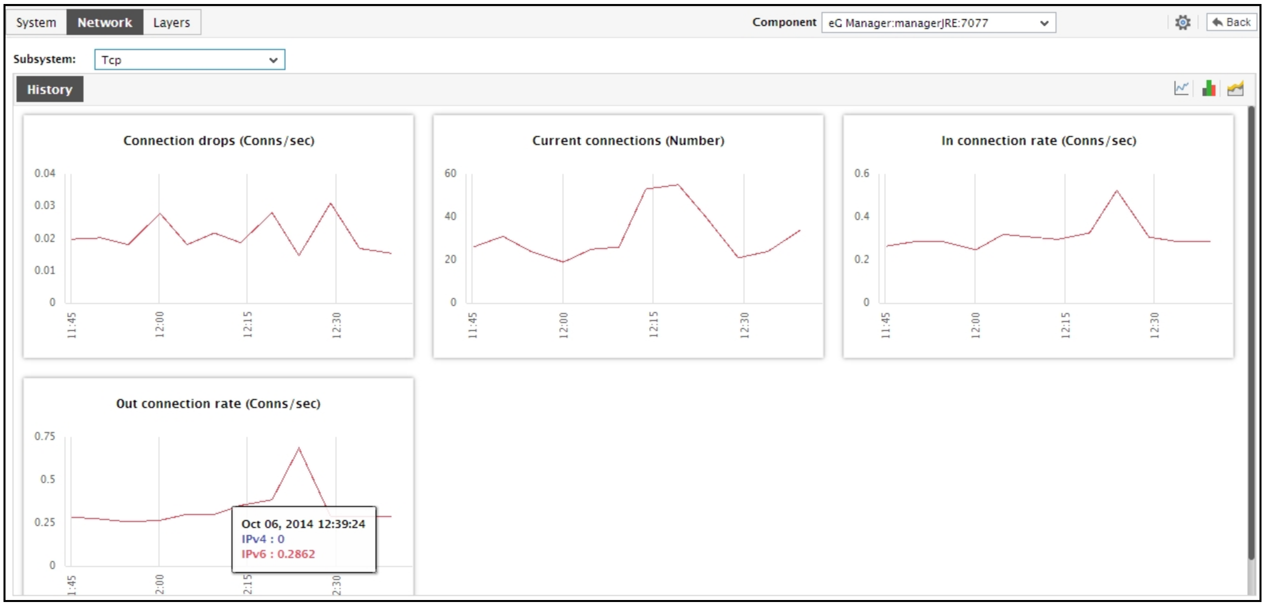

Figure 10.57: The Dashboard of the Tcp Subsystem

For shifting from the present to the past and analyzing past TCP performance, use the **History** tab page inFigure 10.58. By default, this tab page displays measure graphs revealing the time-of-day variations in current TCP connections, dropped TCP connections, and incoming and outgoing TCP connections, during the last 24 hours (by default). Using these graphs, you can quickly figure out the following:

- When during the last 24 hours was the TCP load on the target host unusually high? When was the load uncharacteristically low?
- Is there any increase in the number of TCP connections that were dropped? Is this increase sudden or consistent? If the TCP connection drops have increased steadily during the last 24 hours, when did this pattern begin?

Clicking on a graph in Figure 10.57 will lead you to the enlarged graph as shown in Figure 10.58.

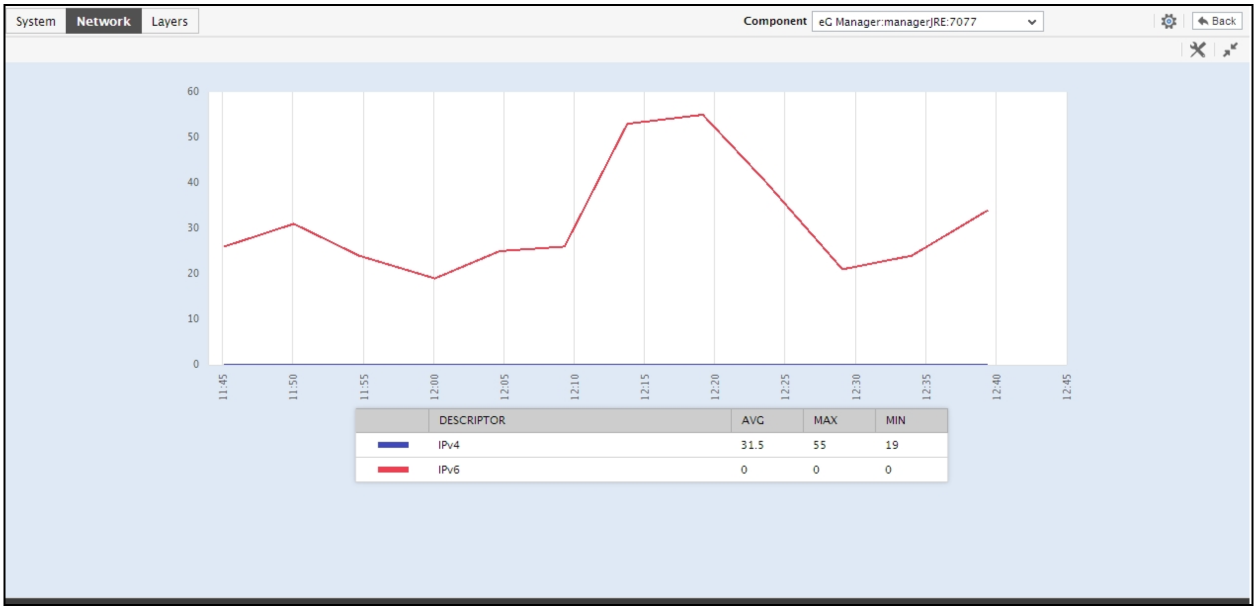

Figure 10.58: The enlarged graph that appears when a graph is clicked

If you want to perform service level audits on the TCP connection handling ability of a target host, then, you can configure the **History** tab page to display summary graphs of the TCP-related metrics, instead of the default measure graphs. For this, just click on the icon at the right, top corner of the **History** tab page. Figure 10.59 will then appear.

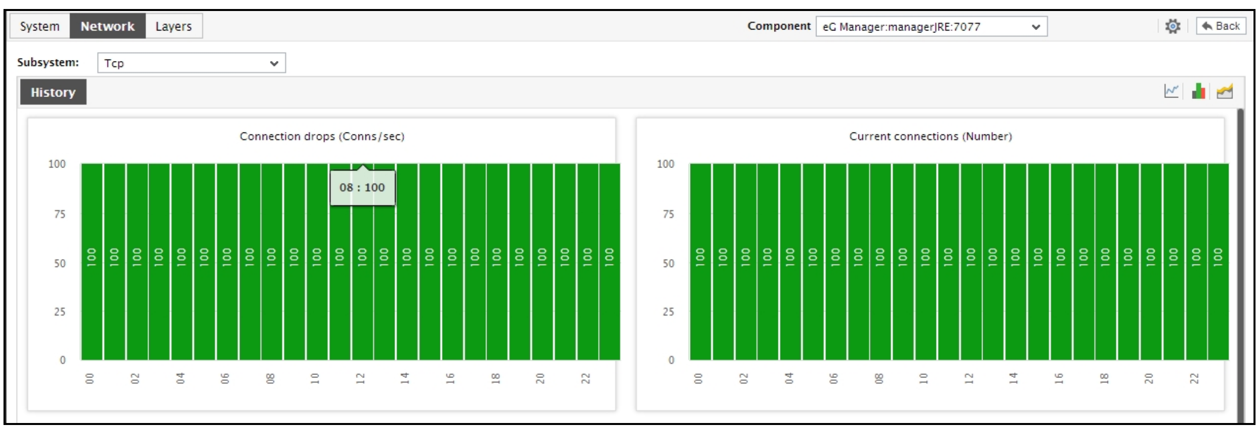

Figure 10.59: Summary graphs in the dashboard of the Tcp Subsystem

The summary graphs denote the percentage of time (during the last 24 hours) for which each TCP-related measure had been in an abnormal state. This way, problem-prone performance metrics can be isolated and the reasons for their poor performance can be investigated.

Similarly, to study the past trends in TCP-related statistics and predict the future behavior of the target host, you can have the **History** tab page display trend graphs of the TCP-related metrics. For this, click on the icon. Figure 10.60 will then appear.

Figure 10.60: Trend graphs displayed in the dashboard of the Tcp Subsystem

Using the trend graphs, you can figure out when (during the last 24 hours) the overall TCP activity and connection drops had peaked on the target host and when it was abnormally low

You can change the **Timeline** of the measure/summary/trend graphs on-the-fly by clicking on the **settings** option on the enlarged graph.

Also, an enlarged summary/trend graph allows you to alter the graph **Duration** - i.e., view the daily or monthly summary/trend information, instead of the default hourly data in the graphs.

Moreover, by default, the trend graphs in the **History** tab page plot the minimum and maximum values of a measure during the given timeline. In enlarged trend graphs, you can change the **Graph** type so that the average values or sum of trend values are plotted in the trend graphs instead.

Note:

In case of descriptor-based tests, the **Summary** and **Trend** graphs displayed in the **History** tab page typically plot the values for a single descriptor alone. To view the graph for another descriptor, pick a descriptor from the drop-down list made available above the corresponding summary/trend graph.

At any point in time, you can switch to the measure graphs by clicking on the button.

10.2.4 NetworkInterfaces

If you want to focus only on the speed, bandwidth usage, and traffic handled by the network interfaces supported by a network device, pick the **NetworkInterfaces** option from the **Subsystem** list in your **Network Dashboard**. When this is done, Figure 10.61 will appear.

Figure 10.61: The NetworkInterfaces dashboard

1. Like the other sub-systems, the dashboard for the **NetworkInterfaces** sub-system too begins with dial graphs and digital displays. These dial and digital graphs are configured for a default set of measures - typically, these measures will be key indicators of the performance of network interfaces. If one/more of these measures are currently abnormal for a particular interface, then, these dial and digital graphs will give you a heads-up on the anomaly, and also indicate which network interface is experiencing the issue. On the other hand, if more than one network interface is experiencing performance issues or if all interfaces are operating normally, then the eG Enterprise system randomly picks a network interface and presents the default metrics extracted from that interface in the dial and digital format. If required, you can configure dial and digital displays for additional measures, or can even remove the graphs that pre-exist by deleting the corresponding measures; for this, follow the steps discussed hereunder:

2. Let us now return to the dashboard. If you click on any dial/digital graph in the dashboard, you will be directly lead to the layer model page, which will display the exact measure represented by the dial/digital graph and the layer-test combination that reports the said measure.

3. The **Comparison** tab page below displays a series of bar charts depicting the top network interfaces in various performance areas. These bar charts enable you to quickly and accurately identify the following:

   - Which network interface is operating with the maximum speed?
   - Which network interface is consuming bandwidth excessively?
   - Which network interface is unavailable currently?
   - Which network interface is receiving data at a very high rate?
   - Which network interface is sending out data at a very high rate?

4. By default, these bar charts depict the top-10 network interfaces only. To view this graph clearly, click on the corresponding bar chart. Doing so enlarges the bar chart as depicted by Figure 10.62.

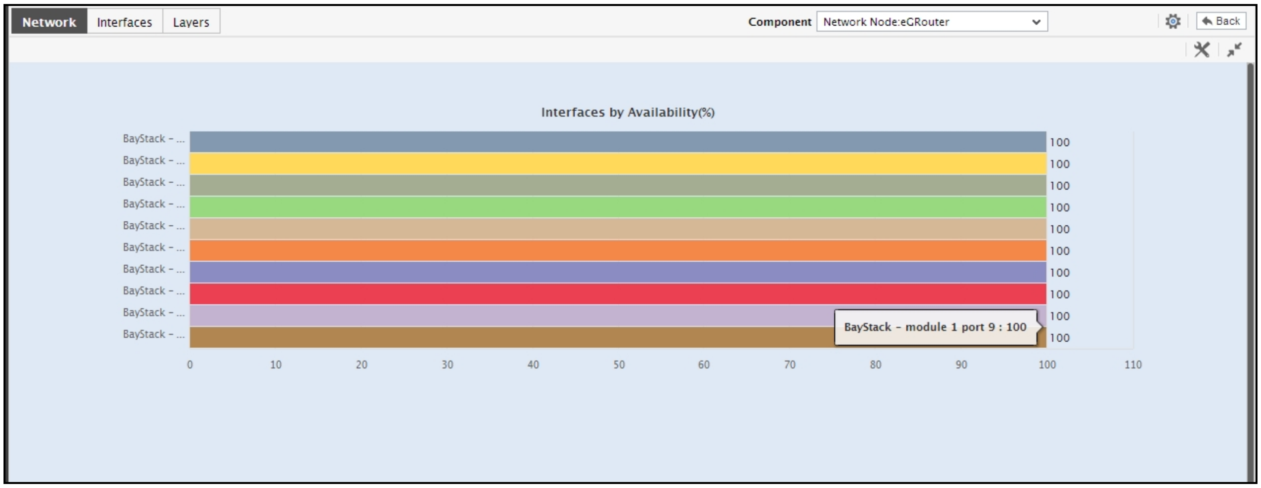

Figure 10.62: A bar chart in the NetworkInterfaces dashboard that has been enlarged

5. The enlarged bar chart too, by default, displays the **TOP-10** network interfaces in the chosen performance area. For instance, if the **Interfaces by Bandwidth usage** chart is enlarged, then the magnified chart will display the **TOP-10** network interfaces in terms of current bandwidth usage. Moreover, since these network interfaces are arranged in the descending order of the bandwidth usage, you can instantly identify the top bandwidth consumer. To view all the network interfaces supported by the device and the bandwidth usage of each, select **ALL** from the **Show** list. You can choose to view only a few of the top/poor players in a performance area, by picking a **TOP-N** or **LAST-N** option from the **Show** list.

6. Also, note that the bar chart displays only those network interfaces that are currently supported by the device. However, sometimes, to investigate past issues, you might want to determine the top performers or the poor performers during a time period in the past. For this, click on **Compare History** in Figure 10.63. This will invoke a **Timeline** field, using which you can alter the **Timeline** of the enlarged bar chart.

7. This way, the **Comparison** tab page enables you to effectively compare the current performance of the network interfaces on the basis of various performance parameters, so that the erring interfaces are swiftly identified. While updates on the current state of the network interfaces can introduce you to operational errors that crept in suddenly, to identify potential anomalies, the knowledge of both current and past performance is essential. The **History** tab page, when clicked, provides a series of measure graphs (by default) that allow you to analyze the past performance of the network interfaces. By default, these graphs are plotted for a period of 24 hours. You can however change the timeline of the graphs on-the-fly by clicking on the **Timeline** link at the right, top corner of the **History** tab page.

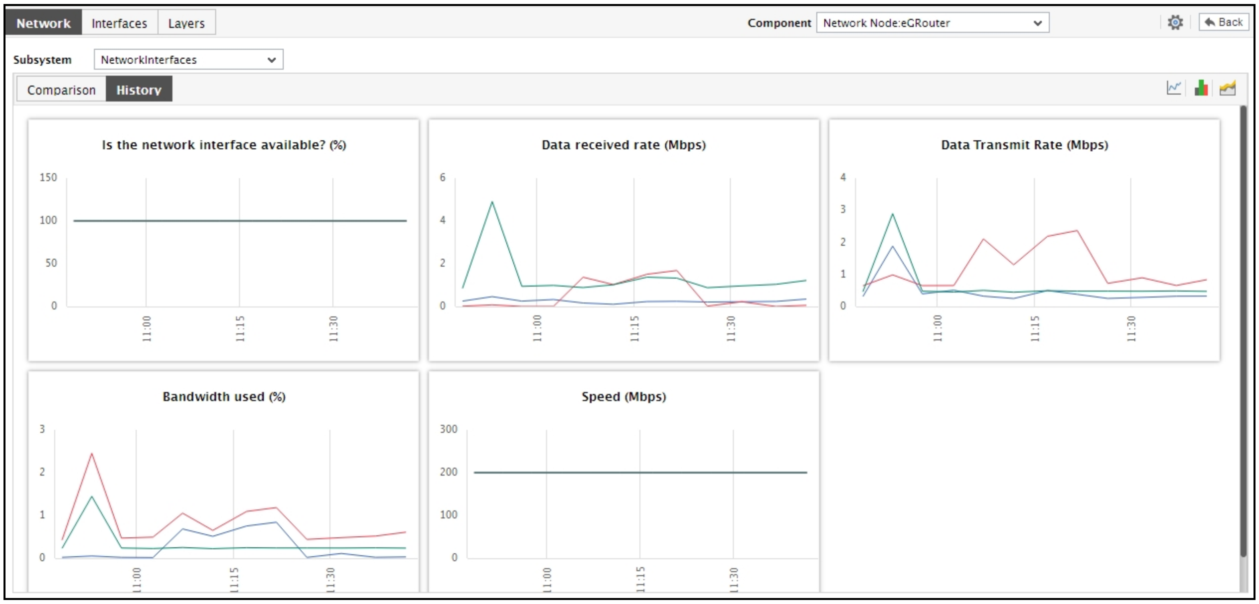

Figure 10.63: The History tab page of the NetworkInterfaces dashboard

Using these measure graphs, you can determine the following:

- Were there consistent/intermittent breaks in the availability of any of the network interfaces during the specified time period?
- Did data traffic to/from any network interface increase significantly during the designated period?
- Were any irregularities noticed in load balancing across the network interfaces during the designated period?
- Was bandwidth usage optimal during the specified period? Which network interface contributed to an increase in bandwidth usage?
- Were all network interfaces operating in normal speed during the said period? Did any network interface experience significant slowdowns? If so, which one?

8. You can enlarge a measure graph by clicking on it, and thus view it more clearly (see Figure 10.64).

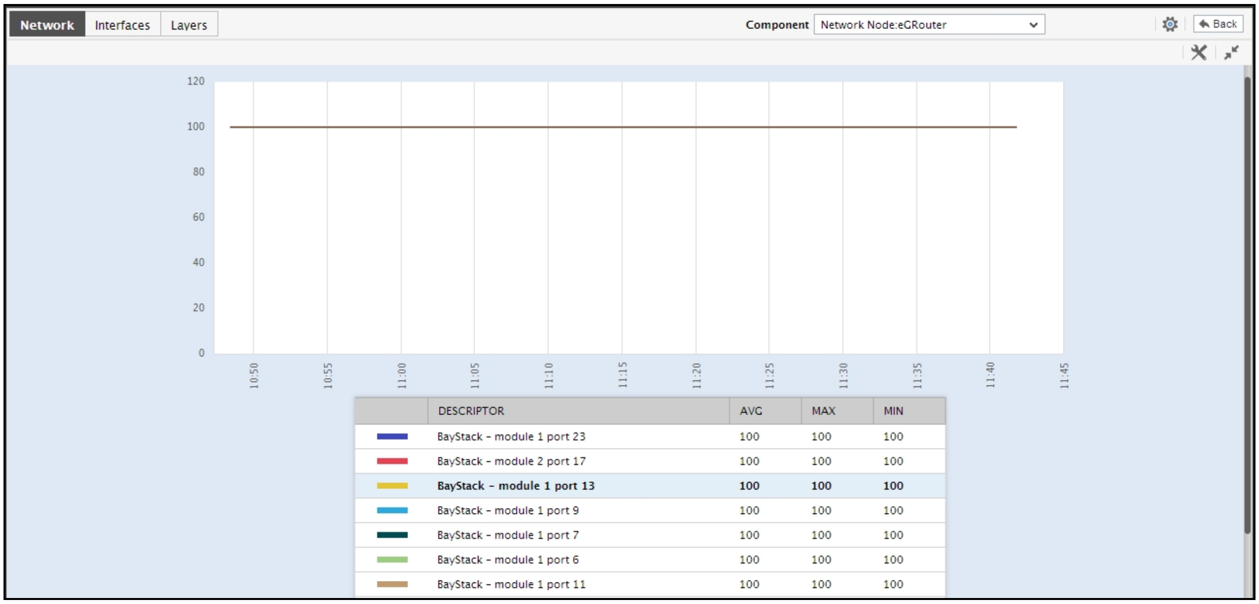

Figure 10.64: An enlarged measure graph in the History tab page of the NetworkInterfaces dashboard

9. Like the comparison graphs, the enlarged measure graphs also, by default, plot the values of the **TOP-10** network interfaces supported by the device. Accordingly, the **TOP-10** option is by default chosen from the **Show** list. To zoom into the performance of only a few of the top players / weak players in that performance area, pick a **TOP-N** or **LAST-N** option from the **Show** list.

10. Besides enabling you to identify the best/worst performers in a chosen performance arena, the enlarged graph also enables you to assess performance of network interfaces across broader time periods - for this, you will have to select a different **Timeline** for the enlarged graph. Similarly, you can also change the dimension (**3D** / **2D**) of a graph in its enlarged mode.

10.2.5 WindowsInterfaces

To analyze the speed, bandwidth usage, and traffic handled by the network interfaces supported by a target server/application over time, pick the **WindowInterfaces** option from the **Subsystem** list. When this is done, Figure 10.65 will appear.

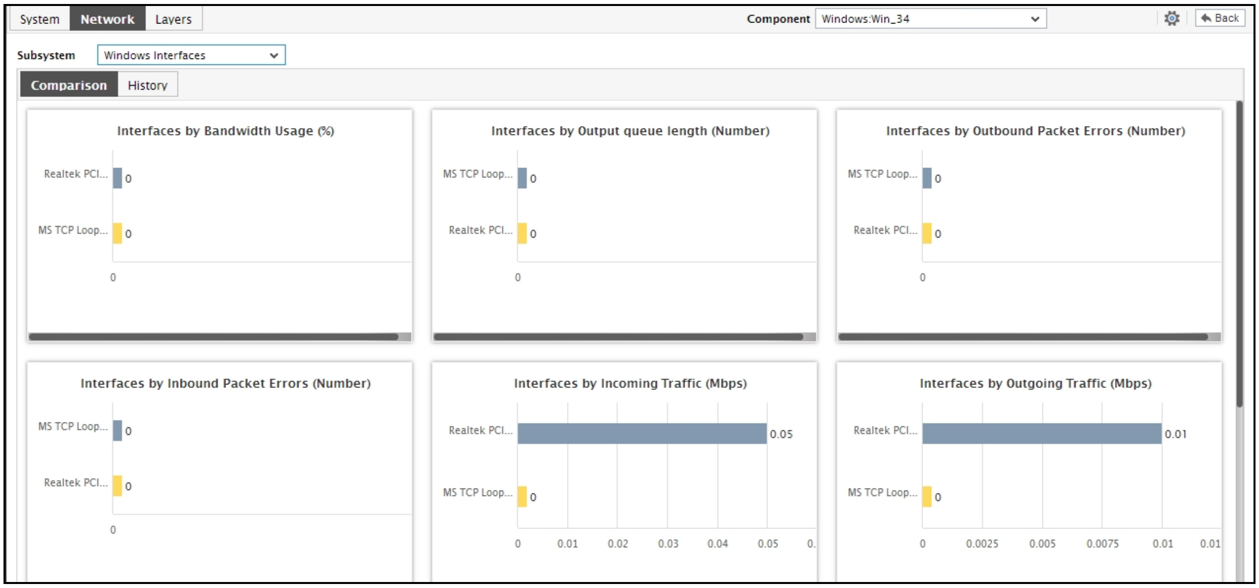

Figure 10.65: The WindowInterfaces dashboard

1. Using the dial graphs and digital displays available in this dashboard, you can receive quick updates on current anomalies pertaining to the network interfaces supported by the target Windows-based component. Besides knowing what went wrong, you will also be able to identify which network interface has been affected, with the help of the dial/digital graphs displayed here.

2. Below the dial/digital displays, you will find the **Comparison** tab page, which provides a default collection of comparison bar charts, each of which compares the performance of the network interfaces in a specific performing sphere. Using these graphs, you can not only isolate performance anomalies, but also identify those network interfaces that are contributing to such anomalies.

3. Let us now refocus on the **Comparison** tab page. To view a comparison graph clearly, click on it; this will enlarge the graph as depicted by Figure 10.66.

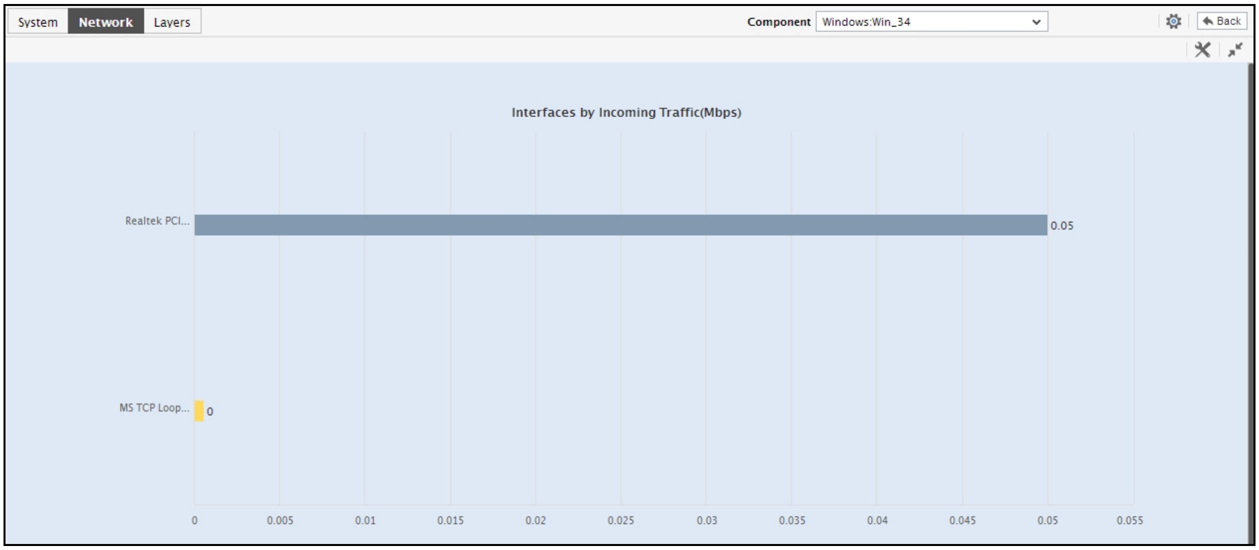

Figure 10.66: An enlarged comparison graph in the WindowsInterfaces dashboard

4. By default, the enlarged comparison graph will reveal only the top-10 network interfaces in a specific performance area. You can, if need be, view all the network interfaces supported by the target component, or only a few best/worst players in a performance area, in the enlarged mode. For this, select the relevant option from the **Show** list in the **Settings** option.

5. For deeper insights into the historical performance of the network interfaces, use the **History** tab page. By default, the **History** tab page of Figure 10.67 displays measure graphs depicting the time-of-day variations in the performance of the network interfaces supported by the target host. By default, these measure graphs are plotted for the last 24 hours. To analyze performance over a wider time range, click the **Timeline** link at the right, top corner of the **History** tab page to change the graph timeline.

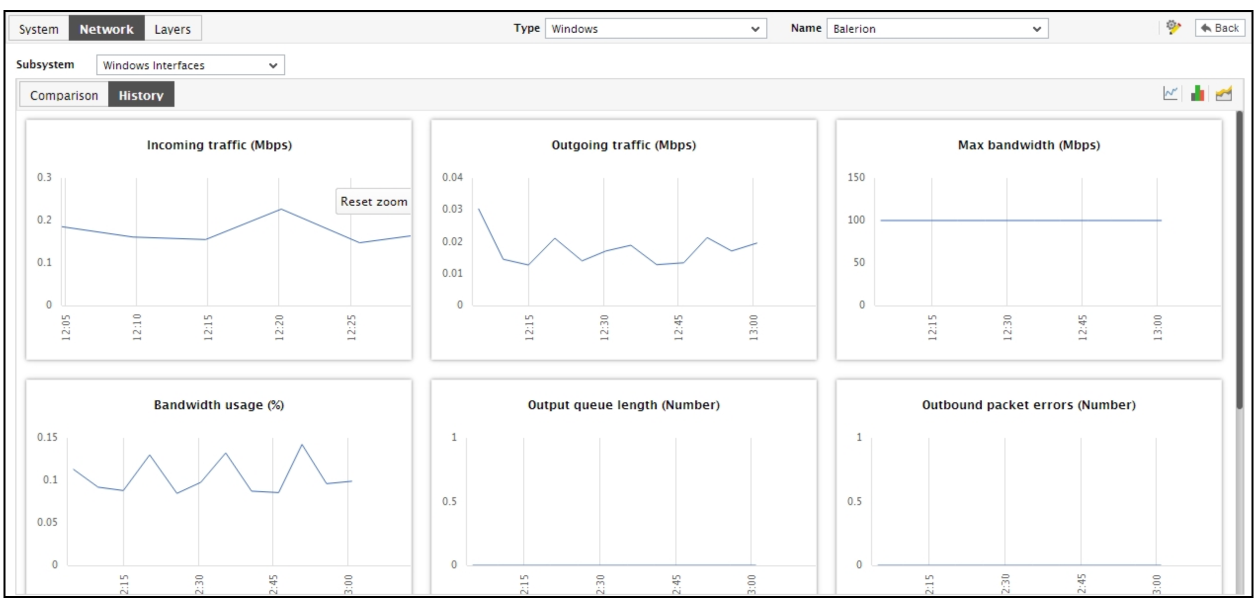

Figure 10.67: The History tab page of the WindowsInterfaces dashboard

6. You can enlarge a measure graph by clicking on it, and thus view it more clearly (see Figure 10.67).

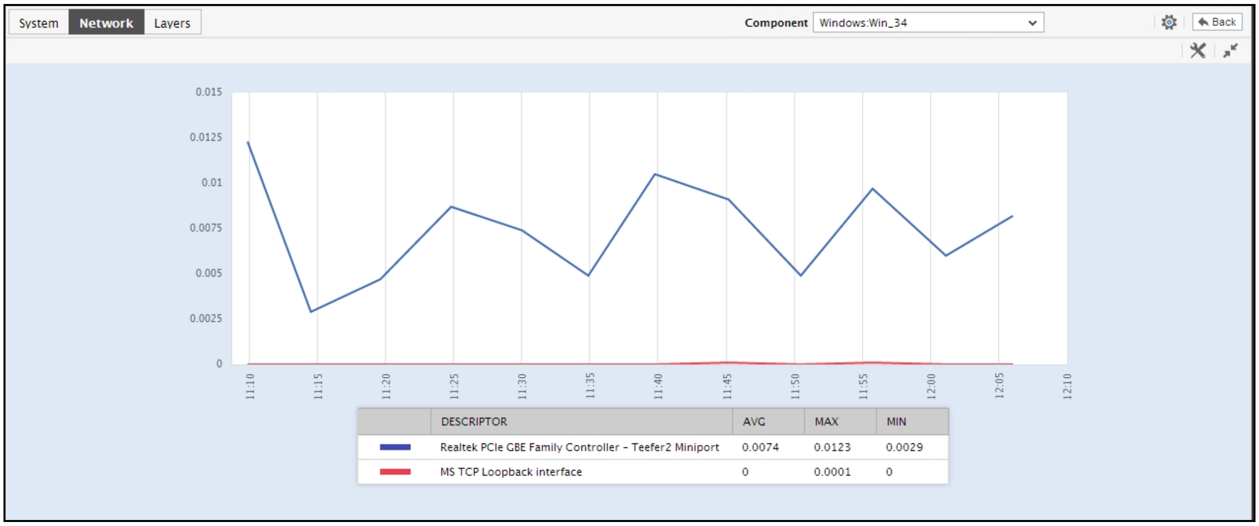

Figure 10.68: An enlarged measure graph in the History tab page of the WindowsInterfaces dashboard

7. Like the comparison graphs, the enlarged measure graphs also, by default, plot the values of the **TOP-10** network interfaces supported by the device. Accordingly, the **TOP-10** option is by default chosen from the **Show** list. To zoom into the performance of only a few of the top players / weak players in that performance area, pick a **TOP-N** or **LAST-N** option from the **Show** list of the **Settings** option.

8. Besides enabling you to identify the best/worst performers in a chosen performance arena, the enlarged graph also enables you to assess performance of network interfaces across broader time periods - for this, you will have to select a different **Timeline** for the enlarged graph.

9. To view the percentage of time during the last 24 hours for which a network interface was affected by issues, click on the icon at the right, top corner of Figure 10.68.

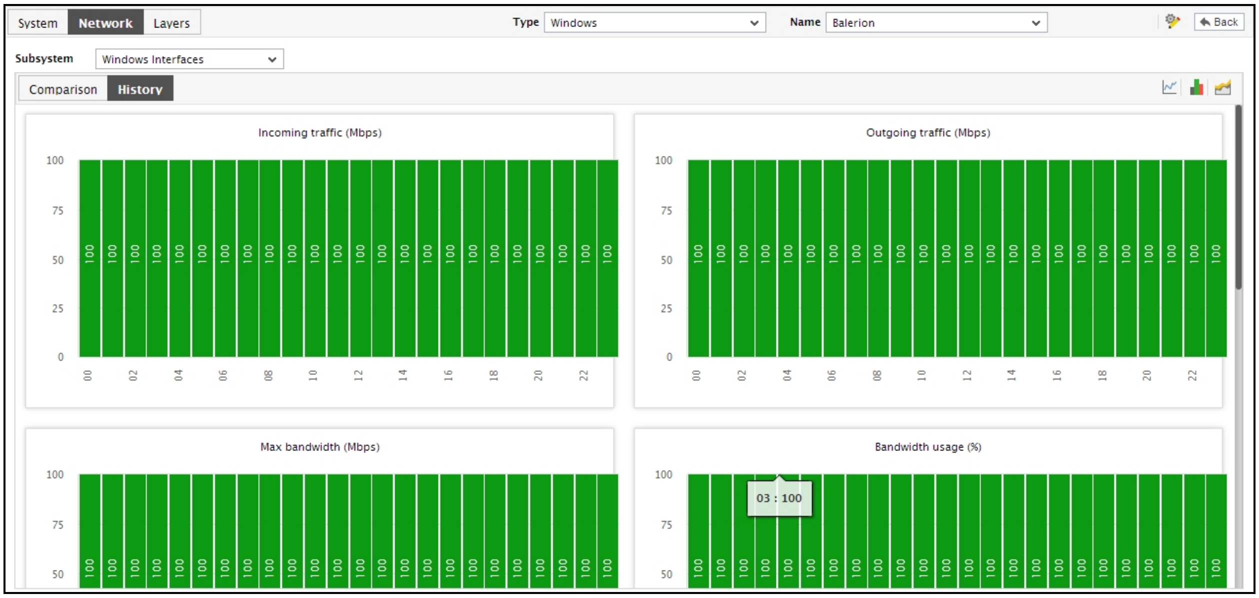

Figure 10.69: Summary graphs in the History tab page of the WindowsInterfaces dashboard

10. Using the graphs in Figure 10.69, you can effectively perform service level audits and detect when and what type of network issues caused the agreed-upon service levels to be compromised.

11. Similarly, click on the  icon at the right, top corner of the **History** tab page in Figure 10.70 to view and analyze the past trends in network interface performance. By default, the trend graphs will pertain to the last 24 hours.

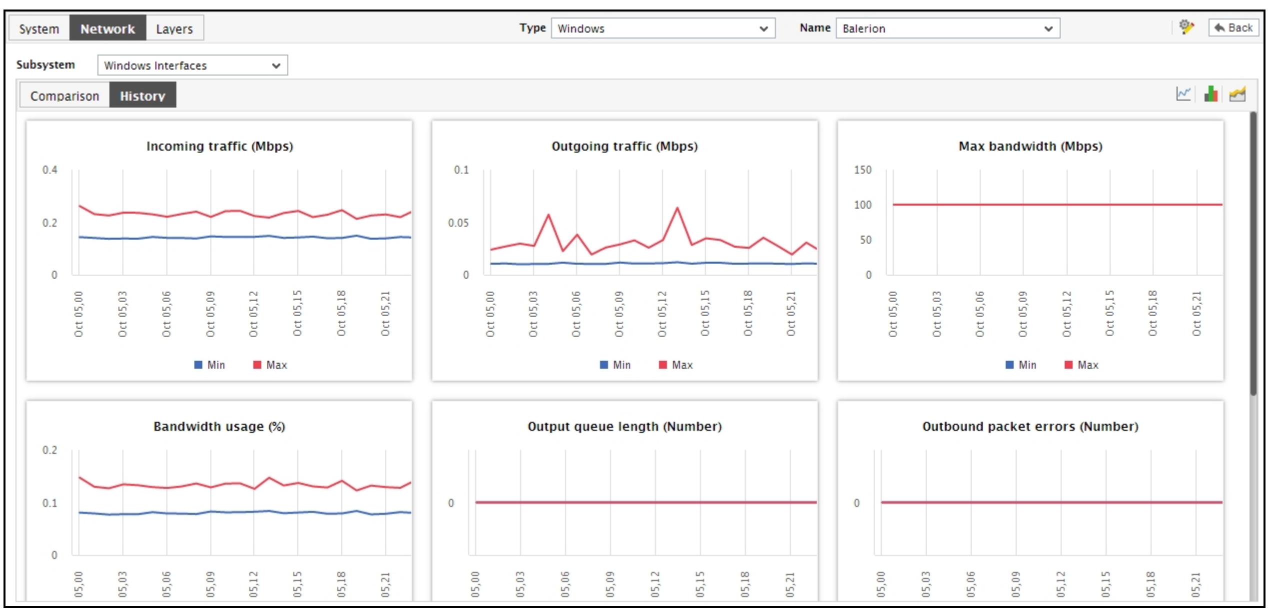

Figure 10.70: The Trend graphs in the dashboard of the Network subsystem

12. Using these trend graphs, you can determine when the performance of network interfaces peaked and when it hit rock bottom - this way, you can easily infer how network interface health has varied during the last 24 hours, and thus receive a heads-up on potential interface-related anomalies.

13. Besides the above, you can instantly change the timeline of the measure/summary/trend graphs in the **History** tab page by clicking on the **Timeline** link at the right, top corner of the tab page. To change the timeline of a single graph on the other hand, click on the graph to enlarge it, and then proceed to change its timeline. In the enlarged mode, you can even change the graph dimension (**3D / 2D**).
14. Also, an enlarged summary/trend graph allows you to alter the graph **Duration** - i.e., view the daily or monthly summary/trend information, instead of the default hourly data in the graphs.
15. Moreover, by default, the trend graphs in the **History** tab page plot the minimum and maximum values of a measure during the given timeline. In enlarged trend graphs, you can change the **Graph** type so that the average values or sum of trend values are plotted in the trend graphs instead.

    Note:

    In case of descriptor-based tests, the **Summary** and **Trend** graphs displayed in the **History** tab page typically plot the values for a single descriptor alone. To view the graph for another descriptor, pick a descriptor from the drop-down list made available above the corresponding summary/trend graph.

16. At any point in time, you can switch to the measure graphs by clicking on the button.

10.2.6 Uptime

To view the uptime details of a network device, you can select the **Uptime** option from the **Subsystem** list. Figure 10.71 then appears.

Figure 10.71: The Network Device Uptime Dashboard

This dashboard reveals the following:

1. The **Device Uptime** section of the dashboard reveals the total time for which the network device has been up and running since it was last rebooted. The **Uptime/DownTime Summary** section provides a quick summary of the availability of the device during the last 24 hours (by default) - the details include: the total duration

for which the device has been up and running in the last 24 hours, the percentage uptime, the total duration (in the last 24 hours) for which the device was down, the percentage of downtime, and number of reboots during the last hour. Using these details, you can determine whether the agreed uptime levels for the device were met or not, and if not, how much is the device falling short of its desired performance levels.

2. You can also infer whether the device experienced any reboots during the last 24 hours. To know more about each reboot, refer to the **Reboot Summary** section. For every reboot that occurred in the last 24 hours (by default), this section reveals when the device was shutdown, when the reboot occurred, and how long did the device remain down until it was rebooted. This clearly indicates the frequency of the reboots, and helps determine whether such reboots were scheduled or unexpected.
3. The default **Measure** graphs in the **History** section indicate how long during the last 24 hours (by default) the device has been up, and whether the reboots scheduled for the device have occurred during the last hour or not.
4. To view the measure graphs clearly, click on a graph of interest - this will enlarge that graph. Using these graphs, breaks in the availability of the device and failure of reboot schedules can be accurately identified and investigated. In addition, you can modify the dimension of the measure graphs from the default **3D** to **2D**.
5. Click on the icon at the right, top corner of the **History** section to view summary graphs using which you can effectively perform service level audits on a device, based on the duration of their availability. Determine the percentage of time for which the device was operational during the last day (by default), and also be notified of reboots that might have occurred on the device during the default timeline. If required, you can click on the **Timeline** link alter the graph timeline.
6. Clicking on the icon in the **History** section will display trend graphs on device uptime; these trends reveal when during the last 24 hours uptime was the lowest, and when reboots failed. If required, you can click on the graph to expand it and alter its **Timeline**.
7. Besides the above, you can instantly change the timeline of the measure/summary/trend graphs in the **History** tab page and that of the **Uptime/Downtime Summary** and **Reboot Summary**. To change the timeline of a single graph on the other hand, click on the graph to enlarge it, and then proceed to change its timeline.
8. Also, an enlarged summary/trend graph allows you to alter the graph **Duration** - i.e., view the daily or monthly summary/trend information, instead of the default hourly data.
9. Moreover, by default, the trend graphs in the **History** tab page plot the minimum and maximum values of a measure during the given timeline. In enlarged trend graphs, you can change the **Graph** type so that the average values or sum of trend values are plotted in the trend graphs instead.
10. At any point in time, you can switch to the measure graphs by clicking on the button.

10.2.7 Virtual Network

In order to provide in-depth insights into the network traffic handled by the virtual networks on a vSphere/ESX server, the **Network Dashboard** of such a server supports a specialized **Virtual Network** sub-system. Figure 10.72 depicts the **Virtual Network** dashboard that appears when the **Virtual Network** option is chosen from the **Subsystem** list.

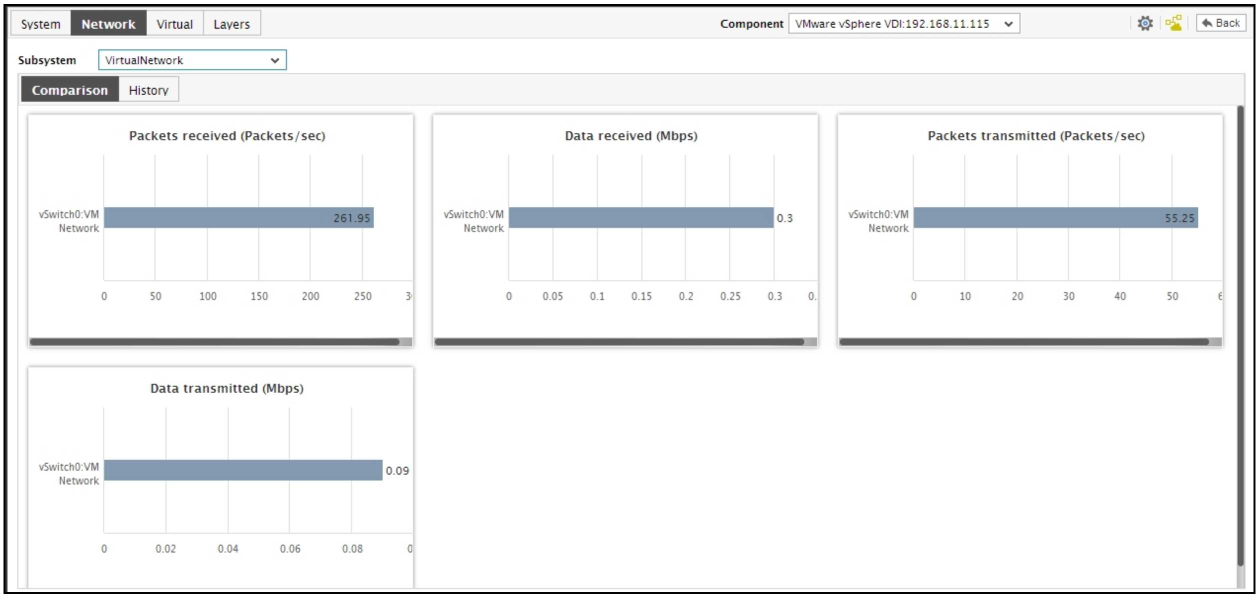

Figure 10.72: The Virtual Network Dashboard

1. To the right of the **Subsystem** list in the **Virtual Network** dashboard, you will find a series of digital graphs displaying the current values for a pre-configured set of measures related to the traffic handled by the virtual networks. Network traffic congestions and the number of VMs affected by such a congestion can be instantly detected with the help of these default digital displays
2. Below the digital graphs, you will find the **Comparison** tab page. This tab page provides comparison bar graphs, which, by default, graphically compare the data and packet traffic handled by the vSwitches servicing the virtual networks. Using these default comparison graphs, you can accurately identify which vSwitches are experiencing heavy traffic.
3. If a comparison graph appears to be very cluttered, then, you can view the graph clearly by clicking on it - this will enlarge the graph. In the enlarged mode, you can clearly view the top-10 (by default) virtual switches in a specific performance area, and thus identify that switch that has performed well/badly in that area. If need be, you can pick a different **TOP-N** or **LAST-N** option from the **Show** list in the enlarged mode so that, you can focus on the performance of more or a less number of vSwitches.
4. However, to perform a more elaborate historical analysis of the performance of virtual networks, you need to switch to the **History** tab page. This tab page displays a default set of measure graphs that track the variations in the traffic to/from each of the vSwitches during the last 24 hours (by default).

Figure 10.73: The History tab page in the VirtualNetwork Dashboard

5. To instantly change the timeline of the measure graphs, click on the graph; this will enlarge the graph. Click on the **Settings** option and change the **Timeline** for the graph.

Figure 10.74: The enlarged history graph in the History tab page of the VirtualNetwork dashboard

6. Besides the timeline, by default every measure graph will plot the values for the top-10 vSwitches on a vSphere/ESX server. If required, you can choose to view the historical performance of a more or a less number of vSwitches in the graph, by picking a different **TOP-N** or **LAST-N** option from the **Show** list in the enlarged mode.

7. Instead of the time-of-day variations, if you prefer to view a quick summary of virtual network performance across vSwitches so that, service level slippages are brought to light, click on the icon at the right, top corner of the **History** tab page. Summary graphs revealing the service level achievements of the pre-configured measures during the last 24 hours (by default), will then appear.

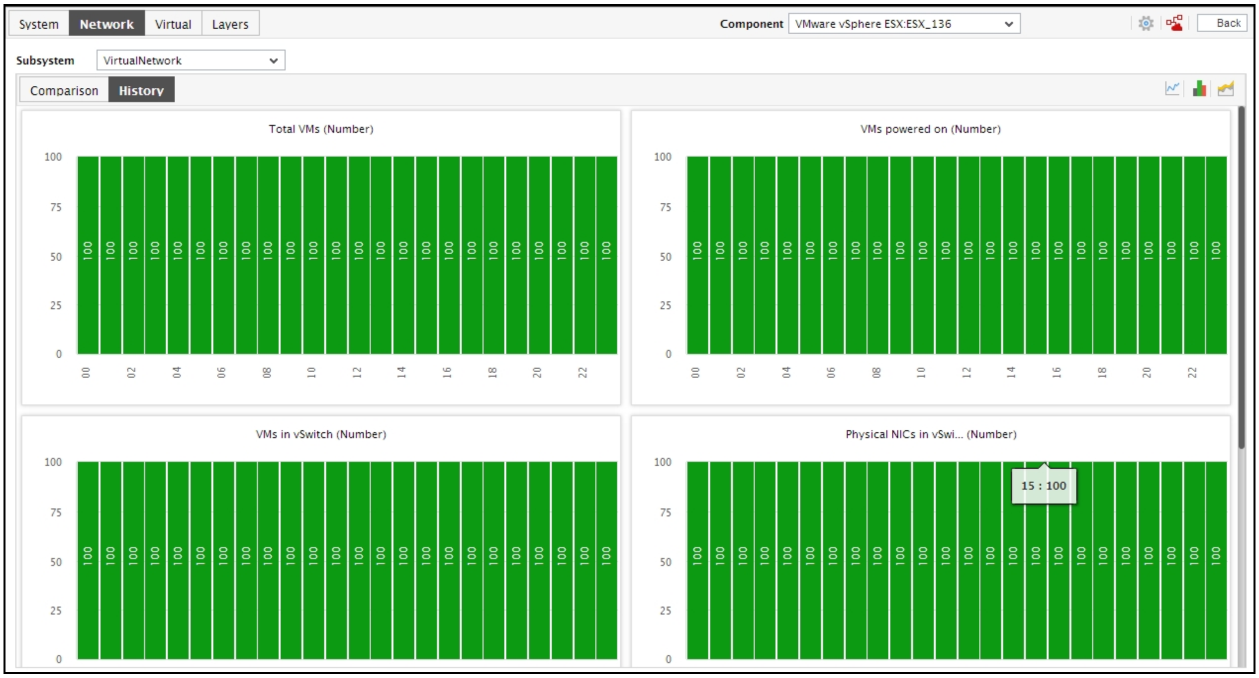

Figure 10.75: Summary graphs displayed in the History tab page of the VirtualNetwork dashboard

8. If you want to alter the timeline of a single graph alone, click on the graph; this will enlarge the graph.
9. Besides the timeline, an enlarged summary graph allows you to alter the graph **Duration** - i.e., view the daily or monthly summary information, instead of the default hourly data.
10. For analyzing past trends in virtual network performance, understanding traffic patterns, and isolating bottlenecks to smooth transmission, click on the icon at the right, top corner of the tab page. Doing so will invoke a series of trend graphs for pre-configured measures, which are plotted for a default duration of 24 hours.

Figure 10.76: Trend graphs displayed in the History tab page of the VirtualNetwork dashboard

11. If you want to alter the timeline of a single graph alone, click on the graph; this will enlarge the graph. Click on the **Settings** option and change the **Timeline**.

12. Besides the timeline, an enlarged trend graph allows you to alter the graph **Duration** - i.e., view the daily or monthly trend information, instead of the default hourly data.

13. Moreover, by default, the trend graphs in the **History** tab page plot the minimum and maximum values of a measure during the given timeline. In enlarged trend graphs, you can view the **Graph type** so that the average values or sum of trend values are plotted in the trend graphs instead.

    Note:

    In case of descriptor-based tests, the **Summary** and **Trend** graphs displayed in the **History** tab page typically plot the values for a single descriptor alone. To view the graph for another descriptor, pick a descriptor from the drop-down list made available above the corresponding summary/trend graph.

14. At any point in time, you can switch to the measure graphs by clicking on the button.

10.2.8 Hyper-V Network Adapters

While monitoring *Hyper-V* and *Hyper-V VDI* environments, the **Network Dashboard** comes embedded with two additional sub-systems, namely - **Hyper-V Network Adapters** and **Hyper-V Switches**.

The **Hyper-V Network Adapters** dashboard provides you with a quick look at the current and historical performance of the network adapters used by a Hyper-V host. Using this dashboard, you can be proactively alerted to real and potential irregularities in load-balancing across the adapters, and identify those adapters that may be experiencing unusually heavy load.

Figure 10.77: Figure 10.77: The Hyper-V Network Adapters Dashboard

1. The dashboard begins with digital graphs for a default set of metrics, which indicate the data/packet load currently handled by the network adapters. The eG Enterprise system automatically compares the current value of a pre-configured metric across all adapters, picks the maximum value, and displays this high watermark in the digital format in the dashboard. A look at these digital graphs will hence instantly reveal whether any adapter is currently experiencing unhealthy load patterns. In the event of an abnormality, you can simply move your mouse pointer over the problem digital graph to know which network adapter is responsible for the anomaly. Alternatively, click on the digital graph, so that you can view the layer, test, descriptor (i.e., network adapter), and measure that has deviated from the norm.

2. The **Comparison** tab page, as the name suggests, displays comparison bar charts for pre-configured measures. With the help of these bar charts, you can effortlessly determine on which network adapter load is abnormally high. To view a graph clearly, click on it. Figure 10.78 will then appear.

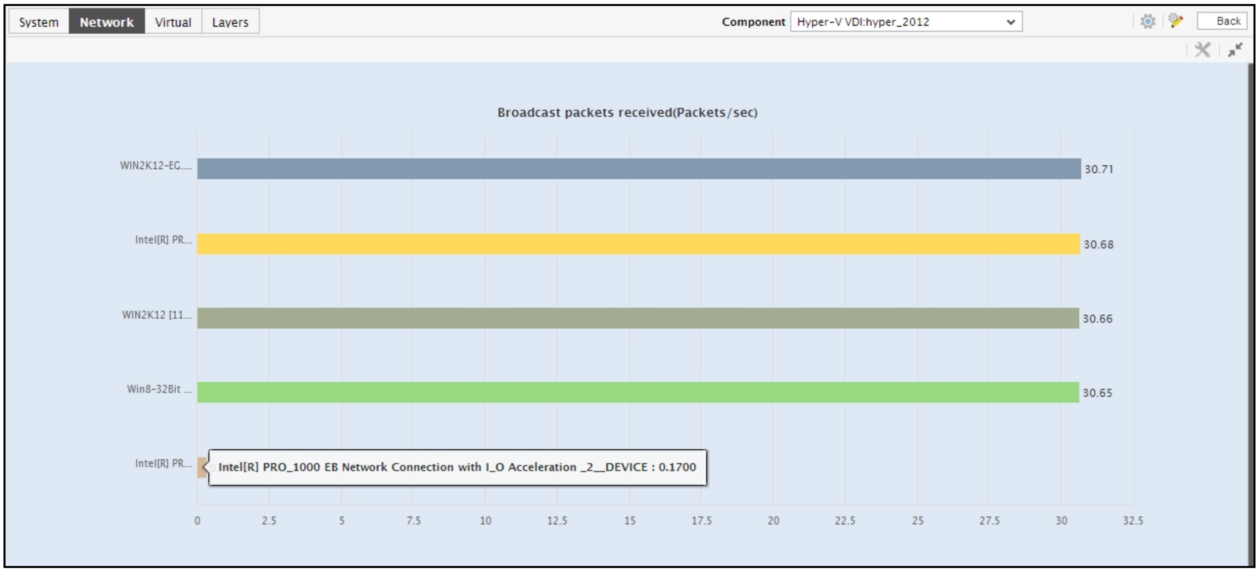

Figure 10.78: Figure 10.78: An enlarged comparison graph in the Hyper-V Network Adapters dashboard

3. In the enlarged mode too, by default, the graph will only reveal the top-10 network adapters in that performing sphere. You can choose a different **TOP-N** or **LAST-N** option from the **Show** list in Figure 10.78 for viewing a few other best/worst performers.
4. For receiving greater visibility into the historical performance of the network adapters, use the **History** tab page. Clicking on this tab page will reveal measure graphs for a default set of load-related metrics, each of which will indicate how well the corresponding measure has performed across adapters during the last 24 hours (by default).
5. In the event of an abnormality, you can use these historical measure graphs to ascertain whether the issue occurred suddenly or is only the climax of a consistent performance deterioration over time.
6. To view a measure graph clearly, click on it.
7. You can change the **Timeline** of the graph in the enlarged mode, or pick a different **TOP-N** or **LAST-N** option from the **Show** list to view the historical performance of a few other best/worst network adapters.
8. To view the percentage of time during the last 24 hours for which a network adapter was affected by issues, click on the icon at the right, top corner.

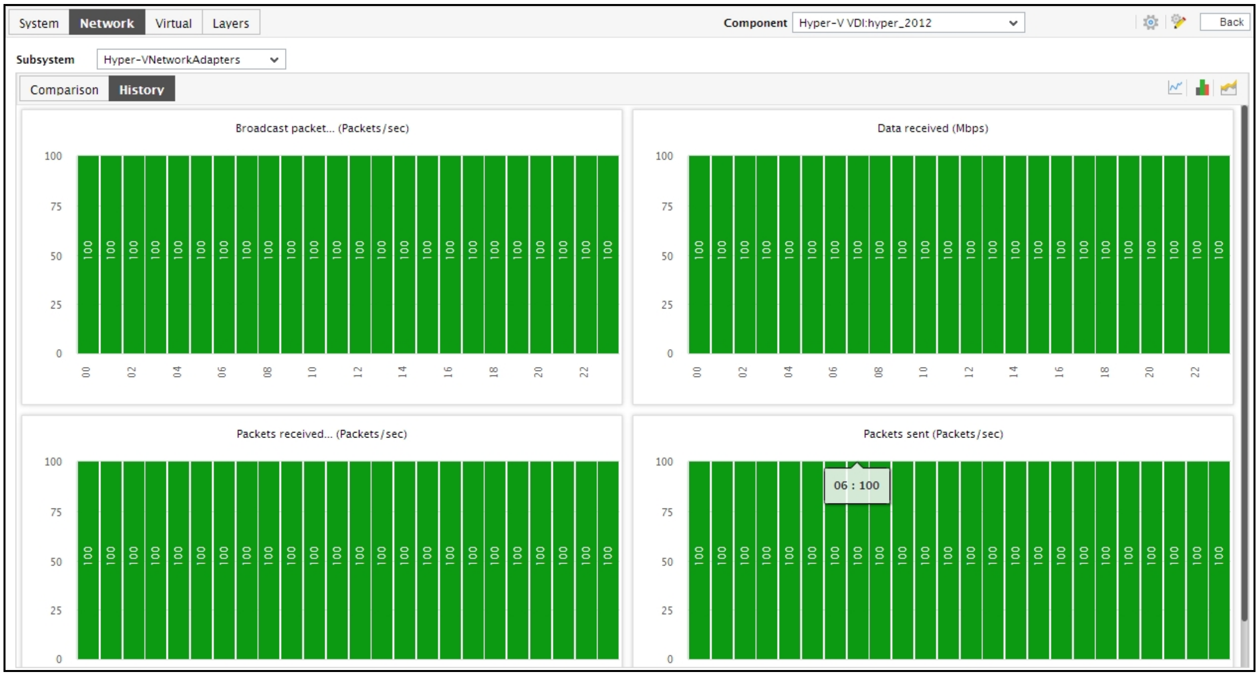

Figure 10.79: Figure 10.79: Summary graphs in the History tab page of the Hyper-V Network Adapters dashboard

9. Using the graphs in Figure 10.80, you can effectively perform service level audits and detect when and what type of network issues caused the agreed-upon service levels to be compromised.

10. Similarly, click on the icon at the right, top corner of the **History** tab page in Figure 10.80 to view and analyze the past trends in network adapter performance. By default, the trend graphs will pertain to the last 24 hours.

    You can click on a summary graph to enlarge it.

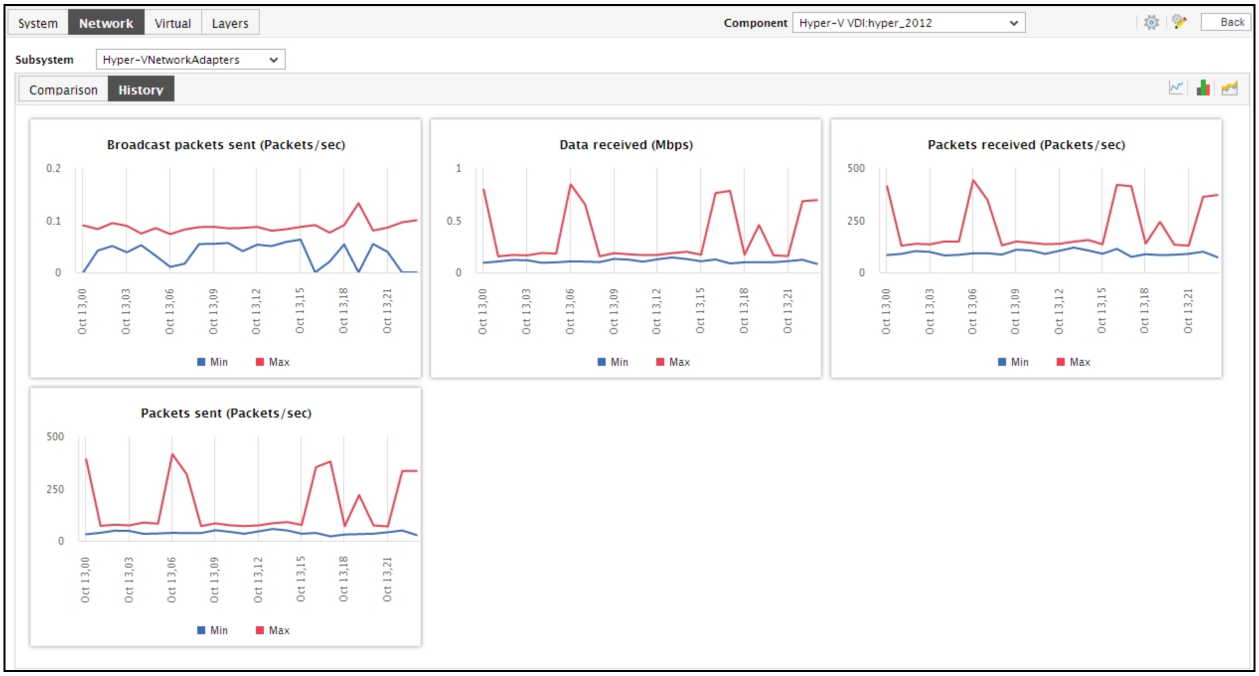

Figure 10.80: Figure 10.80: The Trend graphs in the Hyper-V Network Adapters dashboard

11. Using these trend graphs, you can determine when the performance of a network adapter peaked and when it hit rock bottom - this way, you can easily infer how a network adapters load patterns have varied during the last 24 hours, and thus receive a heads-up on potential adapter-related anomalies.

12. You can click on a trend graph to enlarge it. An enlarged summary/trend graph allows you to alter the graph **Duration** - i.e., view the daily or monthly summary/trend information, instead of the default hourly data in the graphs.

13. Moreover, by default, the trend graphs in the **History** tab page plot the minimum and maximum values of a measure during the given timeline. In enlarged trend graphs, you can change the **Graph** type so that the average values or sum of trend values are plotted in the trend graphs instead.

    Note:

    In case of descriptor-based tests, the **Summary** and **Trend** graphs displayed in the **History** tab page typically plot the values for a single descriptor alone. To view the graph for another descriptor, pick a descriptor from the drop-down list made available above the corresponding summary/trend graph.

14. At any point in time, you can switch to the measure graphs by clicking on the button.

15. If need be, you can instantly change the timeline of the measure/summary/trend graphs in the **History** tab page by clicking on the **Timeline** link at the right, top corner of the tab page.

10.2.9 Hyper-V Switches

1. Selecting the **Hyper-v switches** option from the **Subsystem** list will invoke a **Hyper-v switches** dashboard that serves as a single console from which you can observe the network traffic flowing into and out of the virtual switches associated with a Hyper-V server, identify abnormalities in traffic patterns, isolate virtual

switches that are experiencing heavy network traffic, and review historical performance of the switches to ascertain whether such traffic-related anomalies occurred sporadically or regularly.

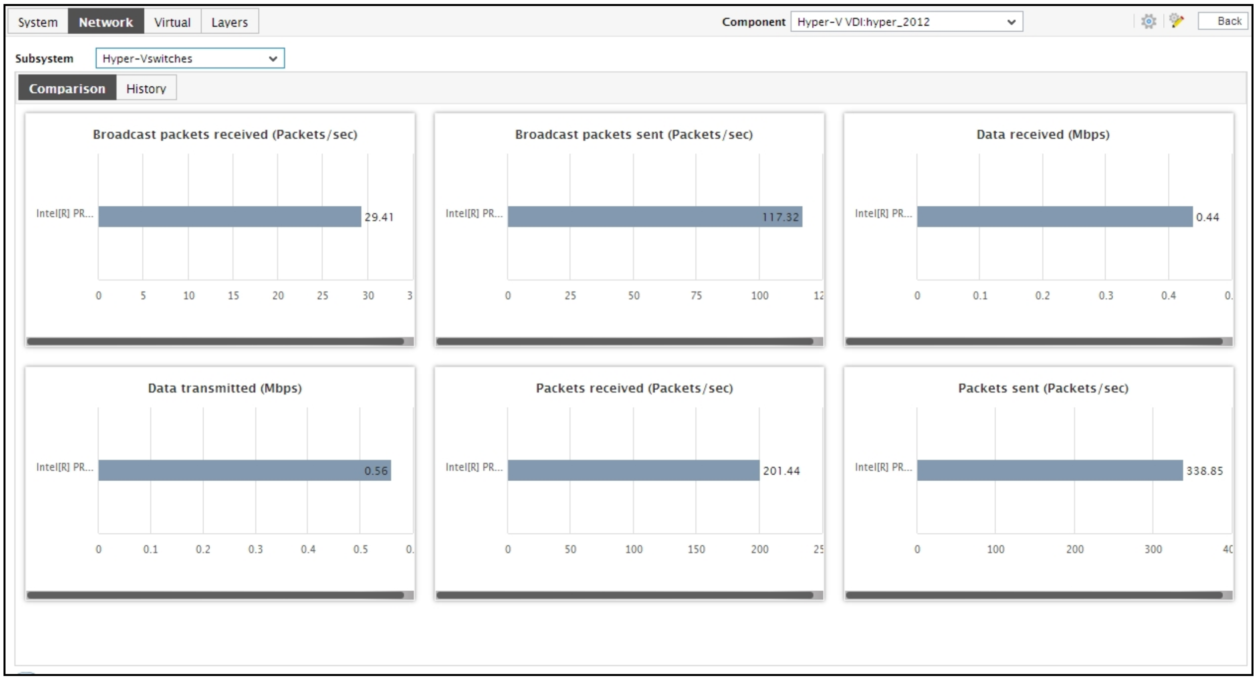

Figure 10.81: The Hyper-V Switches dashboard

2. The **Comparison** tab page, as the name suggests, displays comparison bar charts for pre-configured measures. With the help of these bar charts, you can effortlessly point to the virtual switch that is experiencing heavy traffic. To view a graph clearly, click on it. This will enlarge the graph. In the enlarged mode too, by default, the graph will only reveal the top-10 virtual switches in that performing sphere. You can choose a different **TOP-N** or **LAST-N** option from the **Show** list in the enlarged graph for viewing a few other best/worst performers.

3. For receiving greater visibility into the historical performance of the virtual switches, use the **History** tab page. Clicking on this tab page will reveal measure graphs for a default set of traffic-related metrics, each of which will indicate how well the corresponding measure has performed across switches during the last 24 hours (by default).

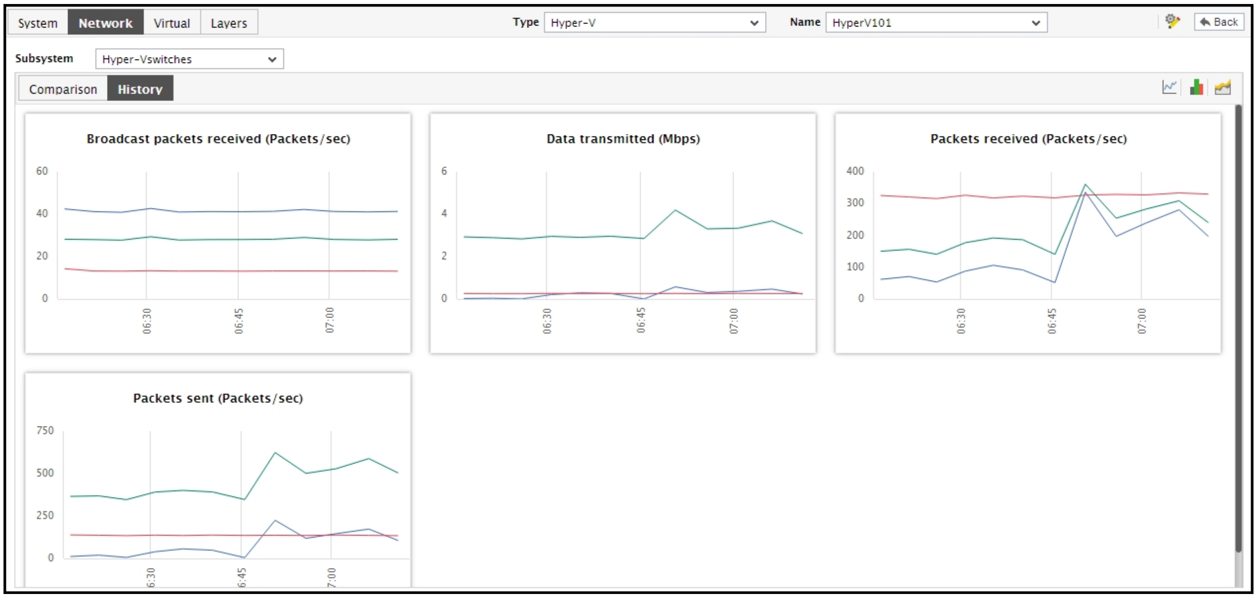

Figure 10.82: The History tab page of the Hyper-V Switches dashboard

4. In the event of an abnormality, you can use these historical measure graphs to ascertain whether the issue occurred suddenly or is only the climax of a consistent performance deterioration over time.

5. To view a measure graph clearly, click on it. The graph will then enlarge as depicted by Figure 10.79.

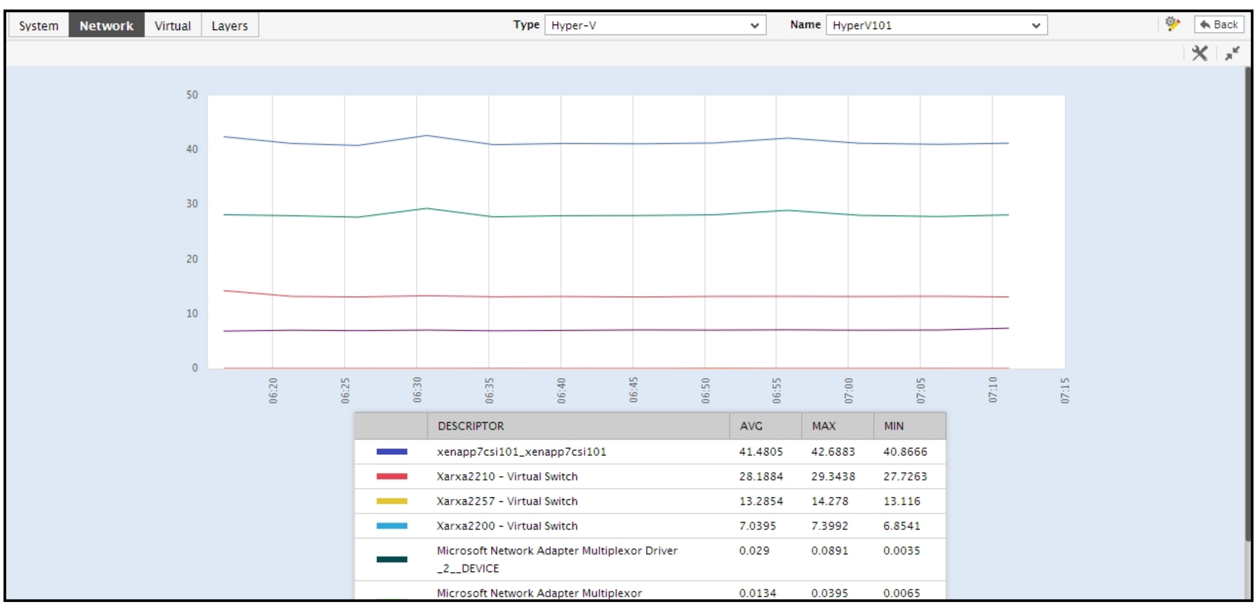

Figure 10.83: An enlarged measure graph in the History tab page of the Hyper-V Network Adapters dashboard

6. You can change the **Timeline** of the graph in the enlarged mode, or pick a different **TOP-N** or **LAST-N** option from the **Show** list to view the historical performance of a few other best/worst virtual switches.

7. To view the percentage of time during the last 24 hours for which a virtual switch was affected by issues, click on the icon at the right, top corner of Figure 10.80.

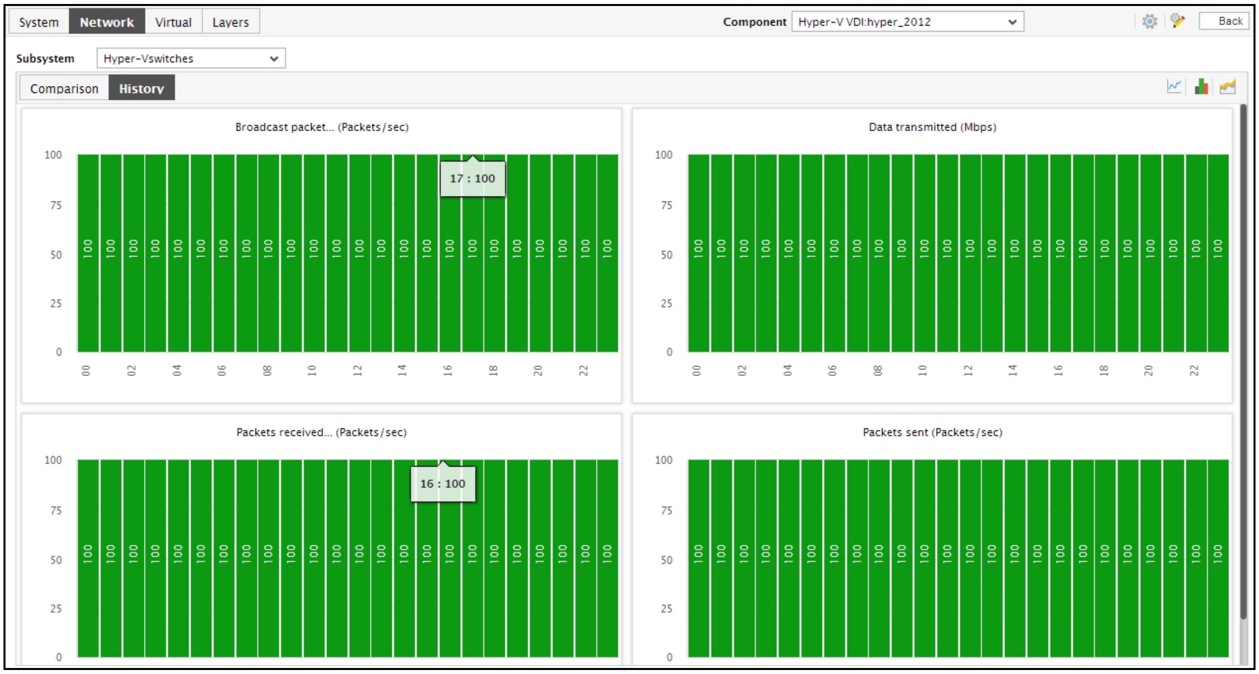

Figure 10.84: Summary graphs in the History tab page of the Hyper-V Switches dashboard

8. Using the graphs in Figure 10.81, you can effectively perform service level audits and detect when and what type of network issues caused the agreed-upon service levels to be compromised.
9. Similarly, click on the icon at the right, top corner of the **History** tab page in Figure 10.82 to view and analyze the past trends in virtual switch performance. By default, the trend graphs will pertain to the last 24 hours.

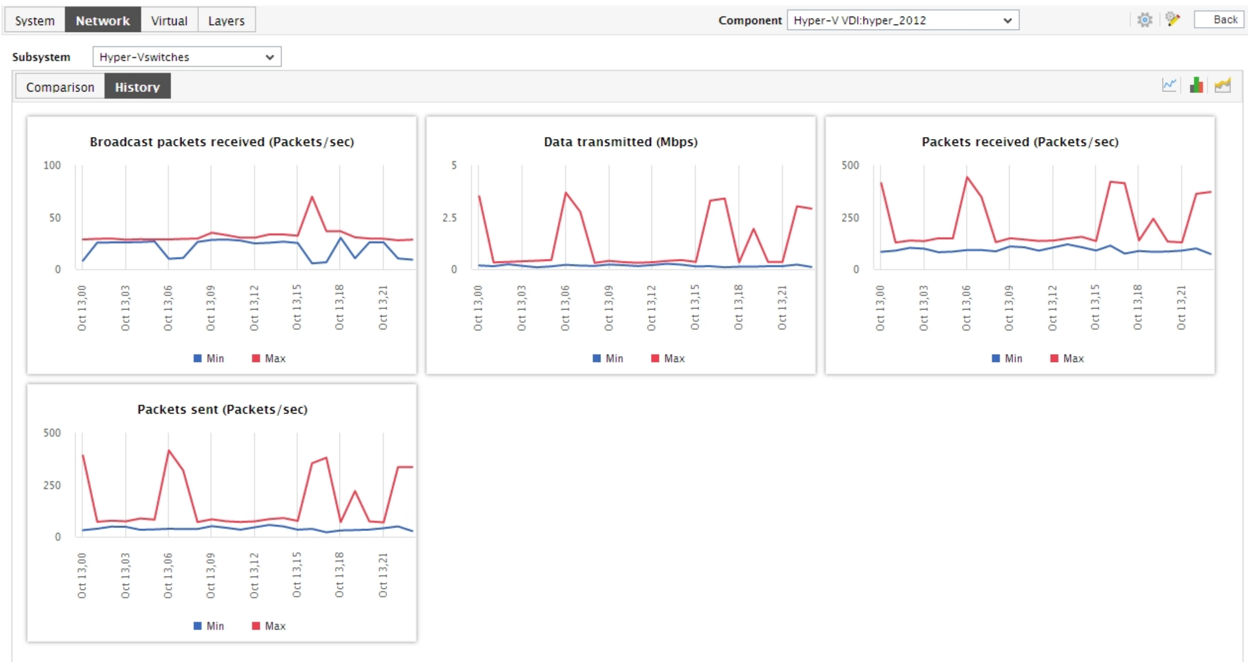

Figure 10.85: The Trend graphs in the Hyper-V Switches dashboard

10. Using these trend graphs, you can determine when the performance of a virtual switch peaked and when it hit rock bottom - this way, you can easily infer how a virtual switch's load patterns have varied during the last 24 hours, and thus receive a heads-up on potential switch-related anomalies.

11. You can click on a trend graph to enlarge it. An enlarged summary/trend graph allows you to alter the graph **Duration** - i.e., view the daily or monthly summary/trend information, instead of the default hourly data in the graphs.

12. Moreover, by default, the trend graphs in the **History** tab page plot the minimum and maximum values of a measure during the given timeline. In enlarged trend graphs, you can view the **Graph** type, the average values or sum of trend values are plotted below the graphs.

    Note:

    In case of descriptor-based tests, the **Summary** and **Trend** graphs displayed in the **History** tab page typically plot the values for a single descriptor alone. To view the graph for another descriptor, pick a descriptor from the drop-down list made available above the corresponding summary/trend graph.

13. At any point in time, you can switch to the measure graphs by clicking on the button.

14. If need be, you can instantly change the timeline of the measure/summary/trend graphs in the **History** tab page by clicking on the **Timeline** link at the right, top corner of the tab page.

15. Typically, the **History** tab page displays measure graphs for a default set of measures.

10.3 The Application Dashboard

In order to ascertain how well an application is/has been performing, analysis of the performance of the **System** and **Network** layers of that application alone might not suffice. A closer look at the health of the **Application**

Layers is also necessary, so as to promptly detect instantaneous operational issues with the target application, and also proactively identify persistent problems or a consistent performance degradation experienced by the application. To provide administrators with such in-depth insights into overall application performance and to enable them to accurately isolate the root-cause of any application-level slowdown, eG Enterprise offers the **Application Dashboard**. Each of the critical applications monitored by eG Enterprise is accompanied by an exclusive application dashboard. The contents of the dashboard will therefore primarily vary depending upon the application being monitored. Figure 10.86 for instance depicts the **Application Dashboard** of a **Java application**.

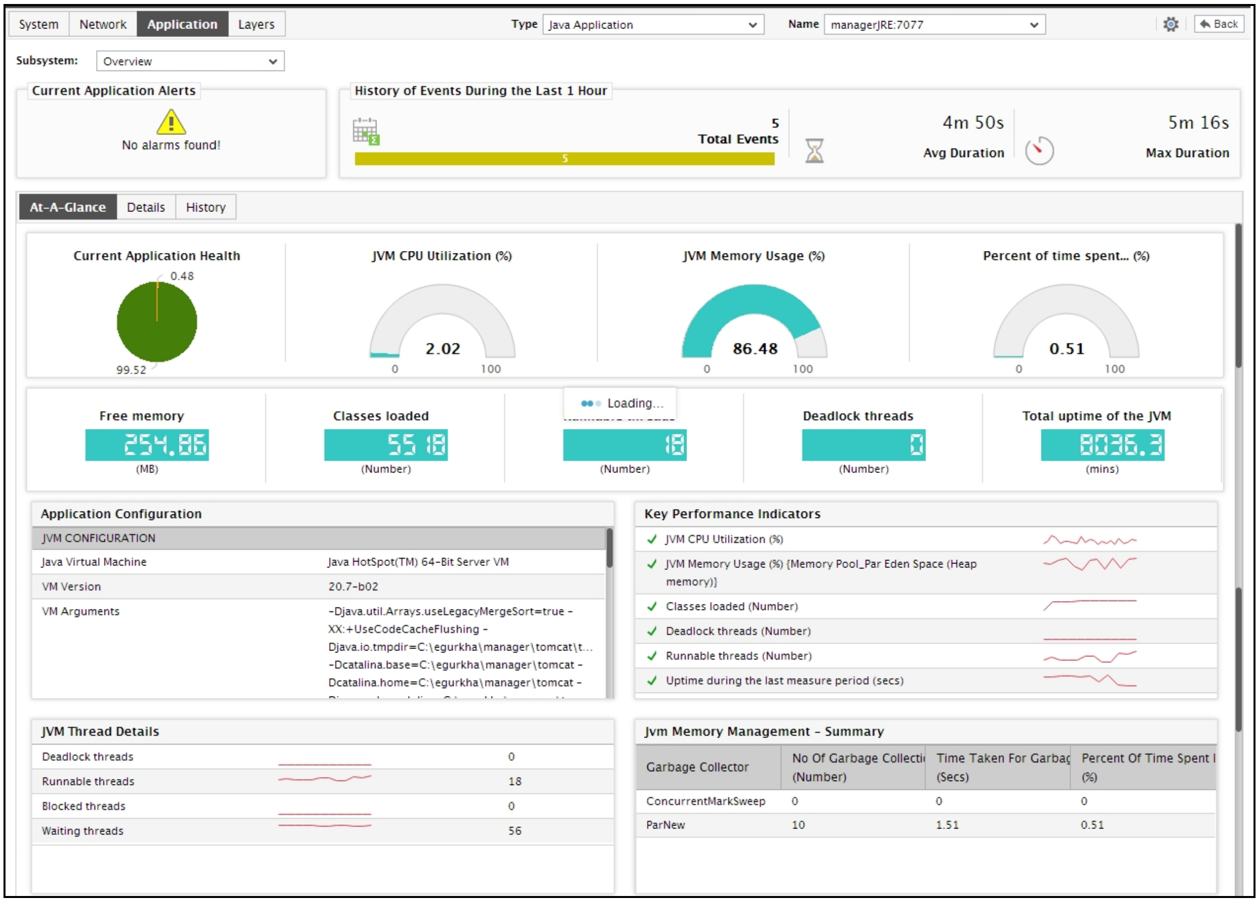

Figure 10.86: The Application Dashboard

In addition, like the **System** and **Network** dashboards, the contents of the **Application** dashboard too are further governed by the **Subsystem** chosen from Figure 10.86. By default, the **Overview** option is chosen from the **Subsystem** list. If need be, you can change this default setting by picking a different option from the **Subsystem** list. The sections that follow will discuss each of the **Subsystems** offered by the sample **Java application dashboard** shown in Figure 10.86 above.

10.3.1 Overview

As its name suggests, the **Overview** dashboard of a Java application provides an all-round view of the health of the Java application being monitored, and helps administrators pinpoint the problem areas. Using this dashboard therefore, you can determine the following quickly and easily:

- Has the application encountered any issue currently? If so, what is the issue and how critical is it?
- How problem-prone has the application been during the last 24 hours? Which application layer has been badly hit?
- Has the administrative staff been able to resolve all past issues? On an average, how long do the administrative personnel take to resolve an issue?
- Are all the key performance parameters of the application operating normally?
- Is the JVM (of the application) utilizing CPU optimally or is the current CPU usage of the JVM very high? Did the CPU usage increase suddenly or gradually - i.e., over a period of time?
- How many threads are currently live on the JVM? Which of these threads is currently consuming high CPU?
- Have any JVM threads been blocked? Which thread is it?
- Are any threads deadlocked?
- Which is the busiest garbage collector on the JVM?
- Which garbage collector is taking too long to collect garbage?
- Which JVM process is currently consuming CPU and memory excessively?
- Which memory pool on the JVM is utilizing too much memory?

The contents of the **Overview Dashboard** have been elaborated on hereunder:

1. The **Current Application Alerts** section of Figure 10.86 reveals the number and type of issues currently affecting the performance of the Java application being monitored.
2. While the list of current issues faced by the application serves as a good indicator of the current state of the application, to know how healthy/otherwise the application has been over time, a look at the problem history of the application is essential. Therefore, the dashboard provides the **History of Events** section; this section presents a bar chart, where every bar indicates the number of problems of a particular severity, which was experienced by the Java application during the last 1 hour (by default). Clicking on a bar here will lead you to Figure 10.87, which provides a detailed history of problems of that priority. Alongside the bar chart, you will also find a table displaying the average and maximum duration for problem resolution; this table helps you determine the efficiency of your administrative staff.

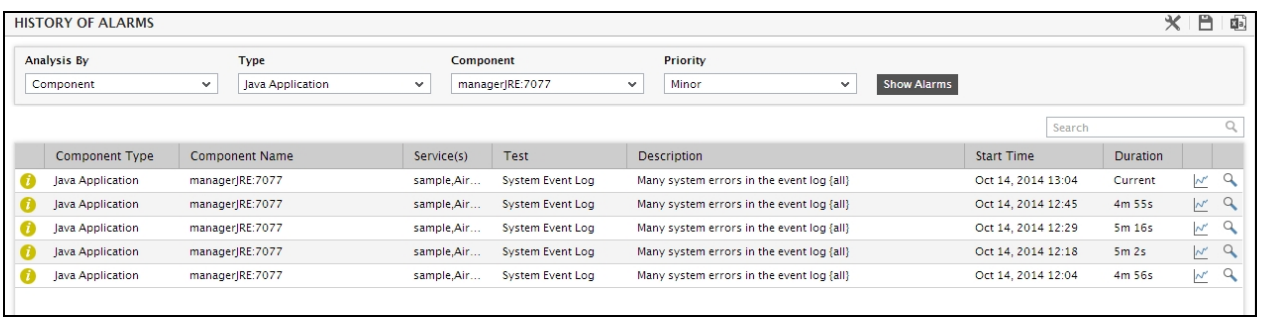
HISTORY OF ALARMS

Analysis By	Type	Component	Priority
Component	Java Application	managerJRE:7077	Minor

Show Alarms

Search

Component Type	Component Name	Service(s)	Test	Description	Start Time	Duration
Java Application	managerJRE:7077	sample,Air...	System Event Log	Many system errors in the event log {all}	Oct 14, 2014 13:04	Current
Java Application	managerJRE:7077	sample,Air...	System Event Log	Many system errors in the event log {all}	Oct 14, 2014 12:45	4m 55s
Java Application	managerJRE:7077	sample,Air...	System Event Log	Many system errors in the event log {all}	Oct 14, 2014 12:29	5m 16s
Java Application	managerJRE:7077	sample,Air...	System Event Log	Many system errors in the event log {all}	Oct 14, 2014 12:18	5m 2s
Java Application	managerJRE:7077	sample,Air...	System Event Log	Many system errors in the event log {all}	Oct 14, 2014 12:04	4m 56s

Figure 10.87: The problem history of the target application

3. Back in the dashboard, you will find that the **History of Events** section is followed by an **At-A-Glance** section; this section, using pie charts, digital displays and gauge charts, reveals, at a single glance, the current status of some of the critical metrics and key components of the Java application. For instance, the **Current Application Health** pie chart indicates the current health of the application by representing the number of application-related metrics that are in various states.

4. The dial and digital graphs that follow provide you with quick updates on the status of a pre-configured set of resource usage-related metrics pertaining to the JVM. If required, you can configure the dial graphs to display the threshold values of the corresponding measures along with their actual values, so that deviations can be easily detected.

   Note:

   If you have configured one/more measures of a descriptor-based test to be displayed as a dial chart, then, in real-time, the descriptor that is in an abnormal state or is currently reporting the maximum value for that measure will be represented in the dial chart. You can view the dial charts pertaining to the other descriptors, by clicking on the **More** button that appears alongside a dial chart in the dashboard.

5. Clicking on a dial/digital graph will lead you to the layer model page of the Java Application; this page will display the exact layer-test combination that reports the measure represented by the dial/digitial graph.

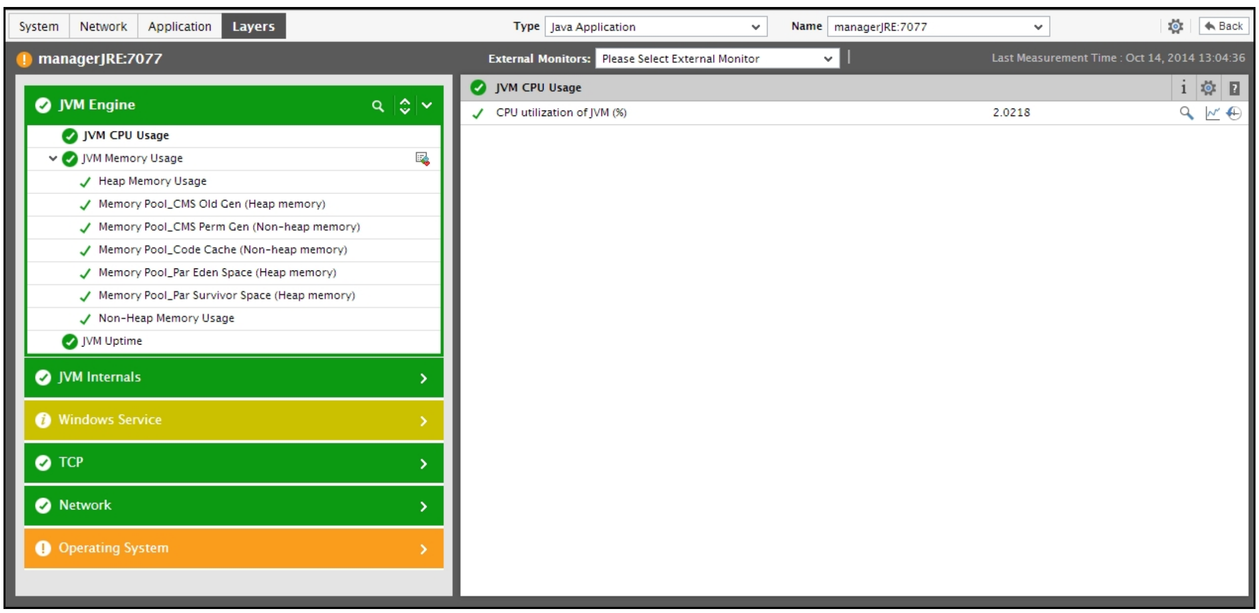

Figure 10.88: The page that appears when the dial/digital graph in the Overview dashboard of the Java Application is clicked

6. If your eG license enables the **Configuration Management** capability, then, an **Application Configuration** section will appear here providing the basic configuration of the application

7. Next to this section, you will find a pre-configured list of **Key Performance Indicators** of the Java application. Besides indicating the current state of and current value reported by a default collection of critical metrics, this section also reveals 'miniature' graphs of each metric, so that you can instantly study how that measure has behaved during the last 1 hour (by default) and thus determine whether the change in state of

the measure was triggered by a sudden dip in performance or a consistent one. Clicking on a measure here will lead you to Figure 10.90, which displays the layer and test that reports the measure.

You can, if required, override the default measure list in the **Key Performance Indicators** section by adding more critical measures to the list or by removing one/more existing ones from the list. Clicking on a 'miniature' graph that corresponds to a key performance indicator will enlarge the graph, so that you can view and analyze the measure behavior more clearly, and can also alter the **Timeline** and, if need be.

Figure 10.89: Enlarging the Key Performance Indicator graph

8. This way, the first few sections of the **At-A-Glance** tab page help understand what issues are currently affecting application health, and when they actually originated. To diagnose the root-cause of these issues however, you would have to take help from the remaining sections of the **At-A-Glance** tab page. For instance, the **Key Performance Indicators** section may indicate a sudden/steady increase in the CPU usage by a JVM. However, to determine whether the rise in CPU usage was a result of one/more high CPU threads executing on the JVM or a couple of resource-intensive Java processes, you need to focus on the **JVM Thread Details** section, **JVM Memory Management - Summary** section, and the **Application Process - Summary** section in the dashboard. The **JVM Thread Details** section for starters reveals the number of JVM threads that are in varying states of activity. With the help of this section therefore, you can quickly figure out whether there are currently any:

   - Threads that are blocking each other (deadlocked threads);
   - Threads that are being blocked by other threads;
   - Threads that are waiting for other threads to release a block;
   - Threads that are consuming high CPU resources, etc.

   Say, you notice that too many threads are currently in a BLOCKED state. Immediately, you might want to know whether this is a sudden occurrence, or a situation that has become worse over time. To enable you to determine this, every thread count displayed in the **JVM Thread Details** section is accompanied by a 'miniature' graph, which tracks the changes in the corresponding thread count during the last 1 hour (by default). To enlarge the graph, click on it; this will invoke Figure 10.89. The enlarged graph allows you to change the **Timeline** for analysis, and also the graph dimension.

Figure 10.90: The enlarged thread count graph

9. To zoom into a particular thread-type and analyze its resource usage, click on the thread type in the **JVM Thread Details** section. For instance, to gain deeper insights into the performance and resource usage of the runnable threads, click on **Runnable Threads** in the **JVM Thread Details** section. Figure 10.90 will then appear, where a list of threads of the chosen type will be displayed, starting with the most CPU-intensive thread. To enable in-depth analysis of the resource usage of a thread, a pie chart depicting the percentage of time the thread used CPU and the percentage of time it was idle, is provided. If a thread is observed to have used CPU excessively, then, you can study the **stack trace** information available alongside the pie chart to zero-in on the exact line of code that the thread was executing when its CPU usage spiked.

10. The **JVM Memory Management - Summary** section reveals how well the JVM manages its memory resources by measuring and reporting the effectiveness of its garbage collection activity. For every garbage collector, this section reveals the number of garbage collections initiated by the collector, the time taken for the garbage collection, and the percentage of time spent on garbage collection. From this information, you can infer which garbage collector is spending too much time and resources on garbage collection. By default, the garbage collector list provided by this section is sorted in the alphabetical order of the names of the collectors. If need be, you can change the sort order so that the garbage collectors are arranged in, say, the descending order of values displayed in the **Time taken for garbage collection** column - this column displays the time taken by each collector to perform garbage collection. To achieve this, simply click on the column heading **Time taken for garbage collection**. Doing so tags the **Time taken for garbage collection** label with a **down arrow** icon - this icon indicates that the **JVM Memory Management** table is currently sorted in the descending order of the time taken by each garbage collector for collecting garbage. To change the sort order to ‘ascending’, all you need to do is just click again on the **Time taken for garbage collection** label or the **down arrow** icon. Similarly, you can sort the table based on any column available in it.

11. The **Application Process - Summary** section, on the other hand, traces the CPU and memory usage of each of the Java processes currently executing on the JVM of the target application, and thus leads you to the resource-intensive processes. By default, the process list provided by this section is sorted in the alphabetical order of the process names. If need be, you can change the sort order so that the processes are arranged in, say, the descending order of values displayed in the **Instances** column - this column

displays the number of instances of each process that is in execution currently. To achieve this, simply click on the column heading - **Instances**. Doing so tags the **Instances** label with a **down arrow** icon - this icon indicates that the process list is currently sorted in the descending order of the instance count. To change the sort order to 'ascending', all you need to do is just click again on the **Instances** label or the **down arrow** icon. Similarly, you can sort the process list based on any column available in the **Application Process - Summary** section.

12. While the **At-A-Glance** tab page reveals the current state of the JVM threads and the overall resource usage of the JVM, to perform additional diagnosis on problem conditions highlighted by the **At-A-Glance** tab page and to accurately pinpoint their root-cause, you need to switch to the **Details** tab page by clicking on it. For instance, the **At-A-Glance** tab page may indicate the number of threads that are currently blocked, but to know which thread has been blocked for the longest time, you will have to use the **Details** tab page.

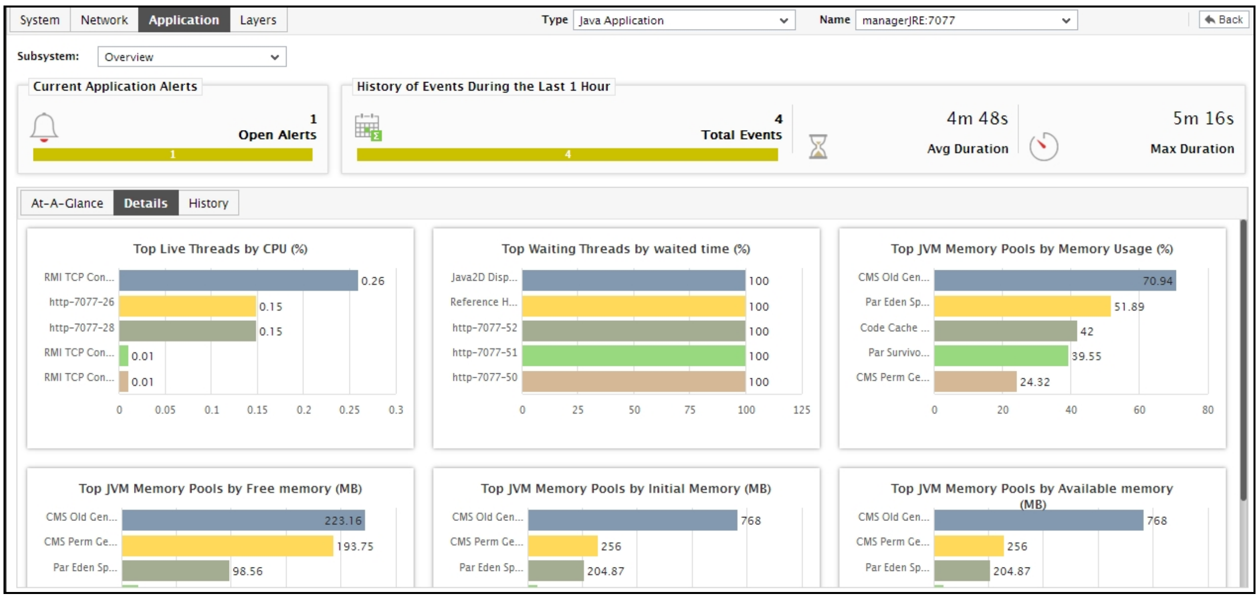

Figure 10.91: The Details tab page of the Application Overview Dashboard

13. The **Details** tab page comprises of a default set of comparison bar graphs using which you can accurately determine the following:

    - How many threads are currently executing on the JVM? Which is the most CPU-intensive thread?
    - Have any threads been blocked? If so, which thread has been blocked for the maximum duration?
    - Are any threads in the WAITING state? If so, which thread has been waiting for the longest time?
    - Which memory pool on the JVM is consuming memory excessively?

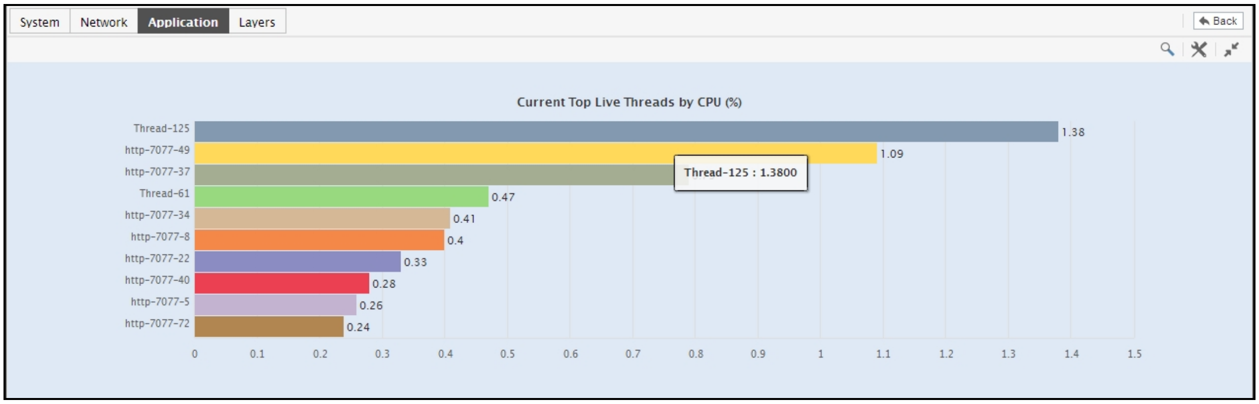

Figure 10.92: The expanded top-n graph in the Details tab page of the Application Overview Dashboard

14. Though the enlarged graph lists all the memory pools or threads (as the case may be) by default, you can customize the enlarged graph to display the details of only a few of the best/worst- performing threads/memory pools by picking a **TOP-N** or **LAST-N** option from the **Show** list in Figure 10.91.

15. Another default aspect of the enlarged graph is that it pertains to the current period only. Sometimes however, you might want to know what occurred during a point of time in the past; for instance, while trying to understand the reason behind a sudden spike in memory usage on a particular day last week, you might want to first determine which memory pool is guilty of abnormal memory consumption on the same day.

16. Where detailed diagnosis is applicable, you can quickly view the detailed measures that correspond to a comparison graph by clicking on the icon at the right, top corner of the enlarged graph. This will invoke Figure 10.93, using which you can arrive at the root-cause of a problem.

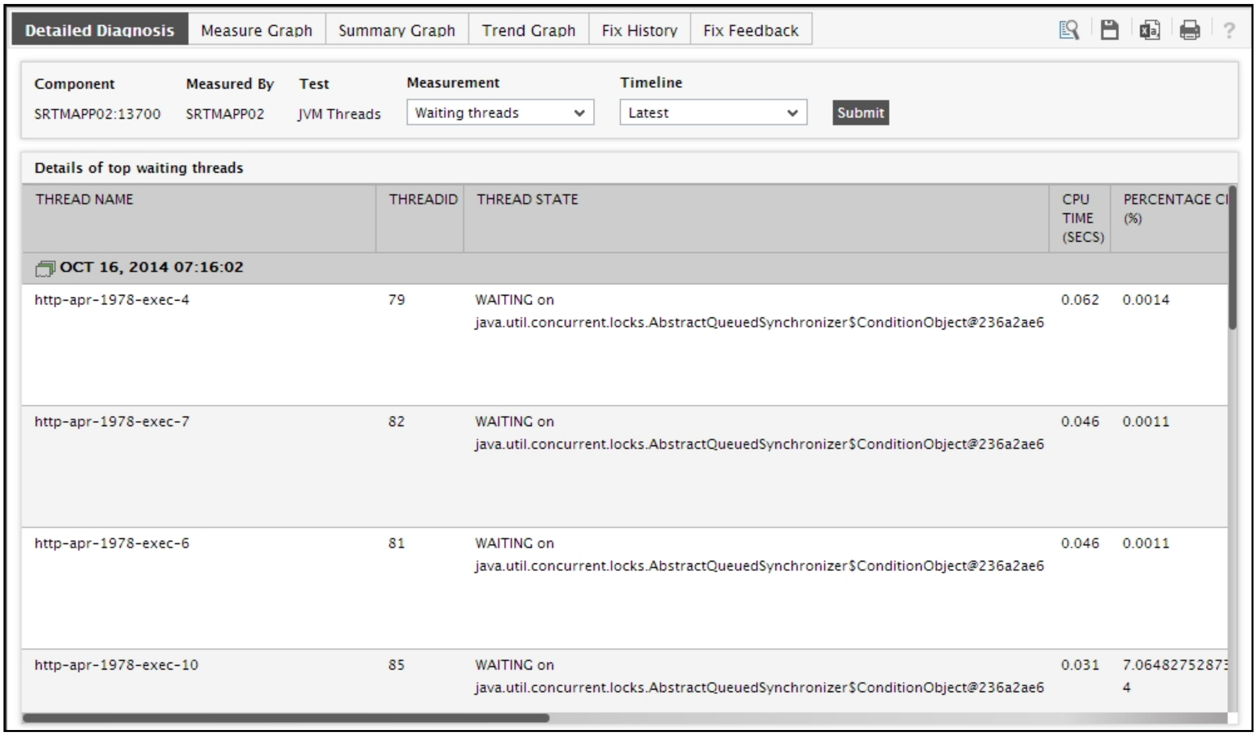

Figure 10.93: The detailed diagnosis that appears when the DD icon in the enlarged comparison bar graph is clicked

17. For detailed time-of-day / trend analysis of the historical performance of a Java application, use the **History** tab page. By default, this tab page (see Figure 10.93) provides time-of-day graphs of critical measures extracted from the target Java application, using which you can understand how performance has varied during the default period of 24 hours. In the event of a problem, these graphs will help you determine whether the problem occurred suddenly or grew with time.

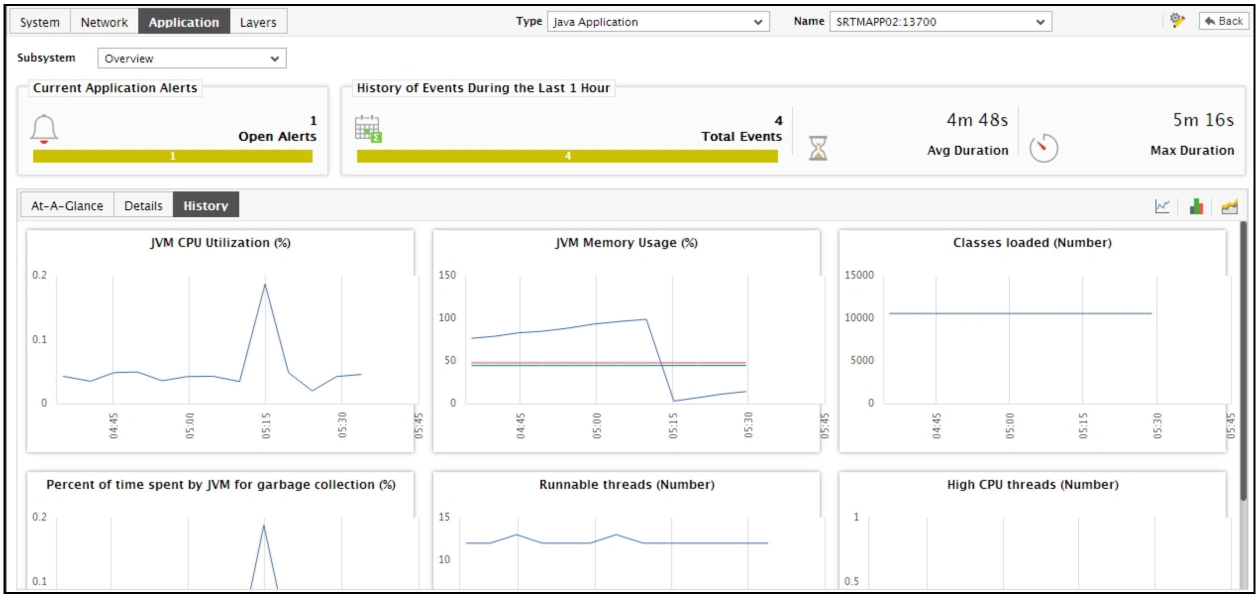

Figure 10.94: Time-of-day measure graphs displayed in the History tab page of the Application Overview Dashboard

18. You can click on any of the graphs to enlarge it, and can change the **Timeline** of that graph in the enlarged mode.

Figure 10.95: An enlarged measure graph of a Java Application

19. In case of tests that support descriptors, the enlarged graph will, by default, plot the values for the **TOP-10** descriptors alone. To configure the graph to plot the values of more or less number of descriptors, select a different **TOP-N** / **LAST-N** option from the **Show** list in Figure 10.93.

20. If you want to quickly perform service level audits on the Java application, then summary graphs may be more appropriate than the default measure graphs. For instance, a summary graph might come in handy if you want to determine the percentage of time during the last 24 hours the Java application consumed excessive CPU. Using such a graph, you can determine whether the CPU usage levels guaranteed by the Java application were met or not, and if not, how frequently did the application falter in this regard. To

invoke such summary graphs, click on the icon at the right, top corner of the **History** tab page. Figure 10.96 will then appear.

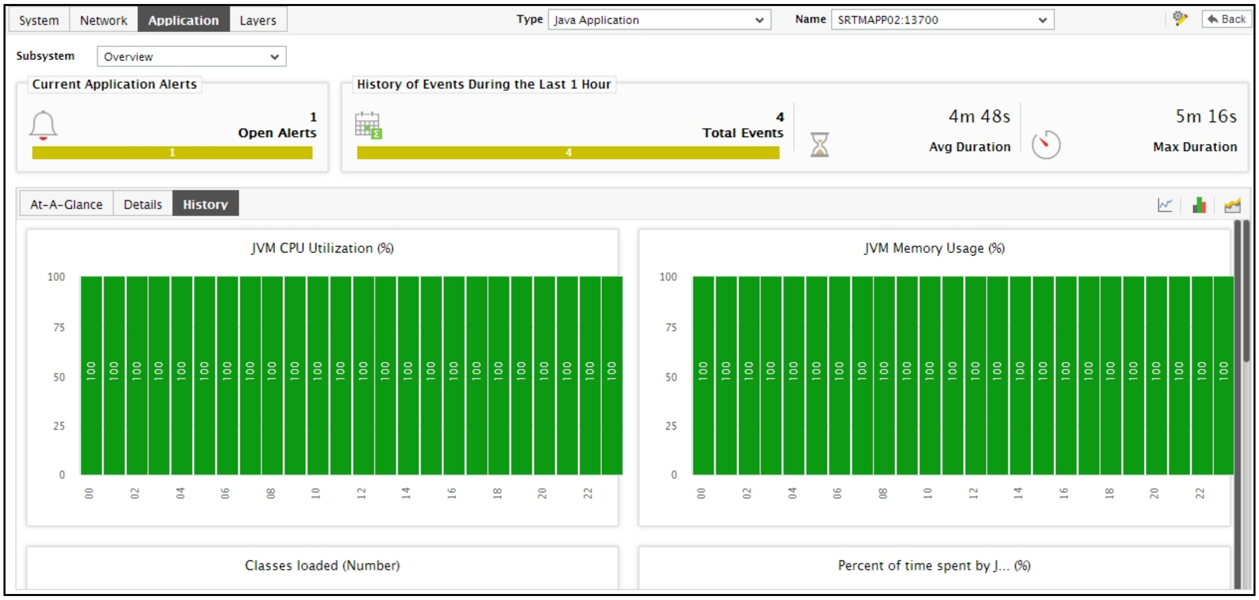

Figure 10.96: Summary graphs displayed in the History tab page of the Application Overview Dashboard

21. You can alter the timeline of all the summary graphs at one shot by clicking the **Timeline** link at the right, top corner of the **History** tab page of Figure 10.97.

22. To change the timeline of a particular graph, click on it; this will enlarge the graph as depicted by Figure 10.98. In the enlarged mode, you can alter the **Timeline** of the graph. Also, though the graph plots hourly summary values by default, you can pick a different **Duration** for the graph in the enlarged mode, so that daily/monthly performance summaries can be analyzed.

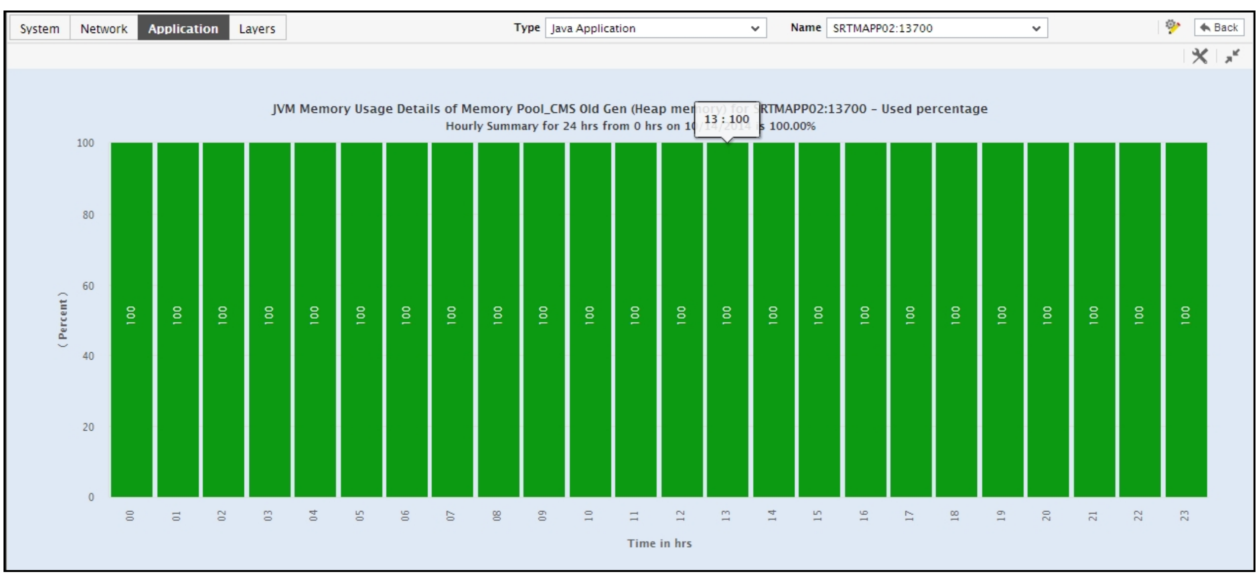

Figure 10.97: An enlarged summary graph of the Java Application

23. To perform effective analysis of the past trends in performance, and to accurately predict future measure behavior, click on the icon at the right, top corner of the **History** tab page. These trend graphs typically show how well and how badly a measure has performed every hour during the last 24 hours (by default). For instance, the CPU usage trend graph of a Java application will help you figure out the maximum and minimum percentage of CPU that was consumed by the application every hour during the last 24 hours. If the gap between the minimum and maximum values is marginal, you can conclude that CPU usage has been more or less constant during the designated period; this implies that CPU usage has neither increased nor decreased steeply during the said timeline. On the other hand, a wide gap between the maximum and minimum values is indicative of erratic usage of CPU, and may necessitate further investigation. By carefully studying the trend graph, you can even determine the points of time at which CPU usage had been abnormally high during the stated timeline, and this knowledge can greatly aid further diagnosis.

Figure 10.98: Trend graphs displayed in the History tab page of the Application Overview Dashboard

24. To analyze trends over a broader time scale, click on the **Timeline** link at the right, top corner of the **History** tab page, and edit the **Timeline** of the trend graphs. Clicking on any of the miniature graphs in this tab page will enlarge that graph, so that you can view the plotted data more clearly and even change its **Timeline**.

25. Besides the timeline, you can even change the **Duration** of the trend graph in the enlarged mode. By default, **Hourly** trends are plotted in the trend graph. By picking a different option from the **Duration** list, you can ensure that **Daily** or **Monthly** trends are plotted in the graph instead.

26. Also, by default, the trend graph only plots the minimum and maximum values registered by a measure. Accordingly, the **Graph** type is set to **Min/Max** in the enlarged mode. If need be, you can change the **Graph** type to **Avg**, so that the average trend values of a measure are plotted for the given **Timeline**. For instance, if an average trend graph is plotted for the *Live threads* measure, then the resulting graph will enable administrators to ascertain how many threads, on an average, were executing in the JVM during a specified timeline; such a graph serves as a good indicator of the growth in the workload of the JVM over

time.

Figure 10.99: Viewing a trend graph that plots average values of a measure for a Java application

27. Likewise, you can also choose **Sum** as the **Graph** type to view a trend graph that plots the sum of the values of a chosen measure for a specified timeline. For instance, if you plot a 'sum of trends' graph for the measure that reports the CPU usage of the JVM, then, the resulting graph will enable you to analyze, on an hourly/daily/monthly basis (depending upon the **Duration** chosen), how CPU usage of the JVM has varied.

    Note:

    In case of descriptor-based tests, the **Summary** and **Trend** graphs displayed in the **History** tab page typically plot the values for a single descriptor alone. To view the graph for another descriptor, pick a descriptor from the drop-down list made available above the corresponding summary/trend graph.

28. At any point in time, you can switch to the measure graphs by clicking on the  button.

10.3.2 JVM Memory

If you want to assess how efficiently the JVM uses the memory resources available to it, and thus promptly detect memory-intensive pools, select the **JVM Memory** option from the **Subsystem** list.

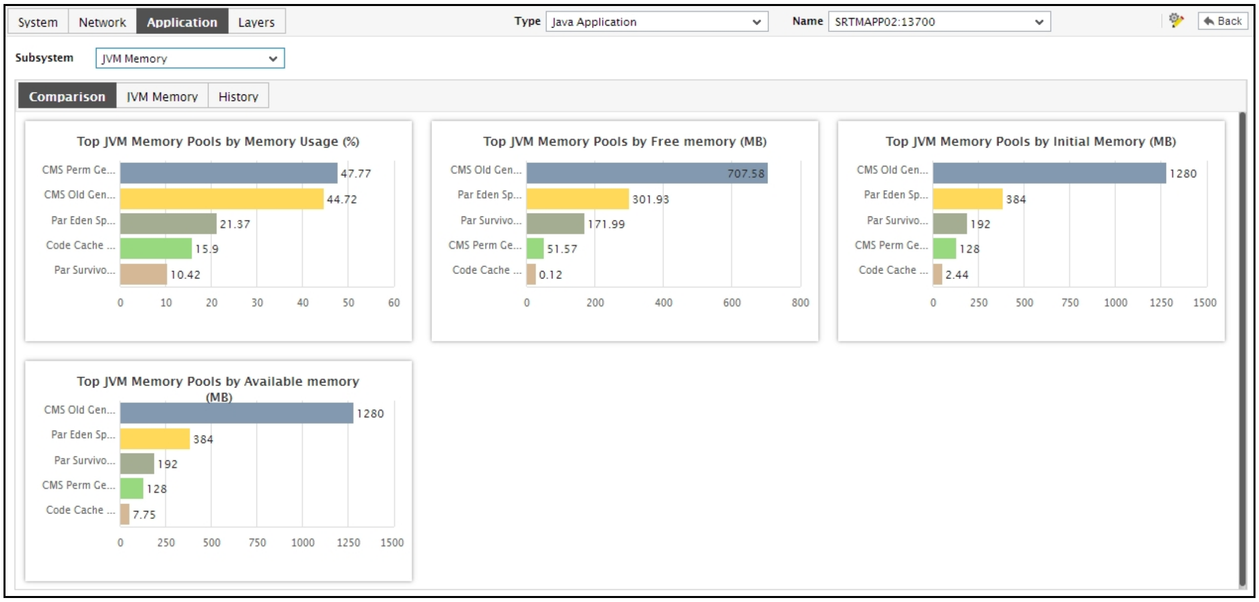

Figure 10.100: The JVM Memory Dashboard

The contents of the **JVM Memory** dashboard that then appears (see Figure 10.100) are as follows:

1. The dashboard begins with a **JVM Memory Usage - Summary** section, which enables you to visually track the percentage of memory used by each of the memory pools on the JVM. Pools that are currently running short of memory resources can be instantly identified using the usage chart provided by this section.

2. The **Comparison** tab page that follows the **JVM Memory Usage - Summary** section, provides a series of top-10 charts, using which you can quickly isolate those memory pools that are leading the lot in the following default performance areas: overall memory consumption, amount of free memory, committed/available memory, and initial memory.

3. If an application slowdown can be attributed to the lack of adequate memory resources, then these top-10 bar charts can aid you in swiftly nailing the exact memory pool that could be serving as the source of this memory contention.

4. Typically, these bar charts depict the current usage data. Sometimes however, you might want to detect which memory pool was over-utilizing memory at some point of time in the past. In such a case, you will have to click on the corresponding graph in the **Comparison** tab page to enlarge it. In the enlarged mode, you can click on the **Compare History** link, so that you can alter the graph **Timeline**, and view which memory pool was the leading memory consumer during the specified timeline.

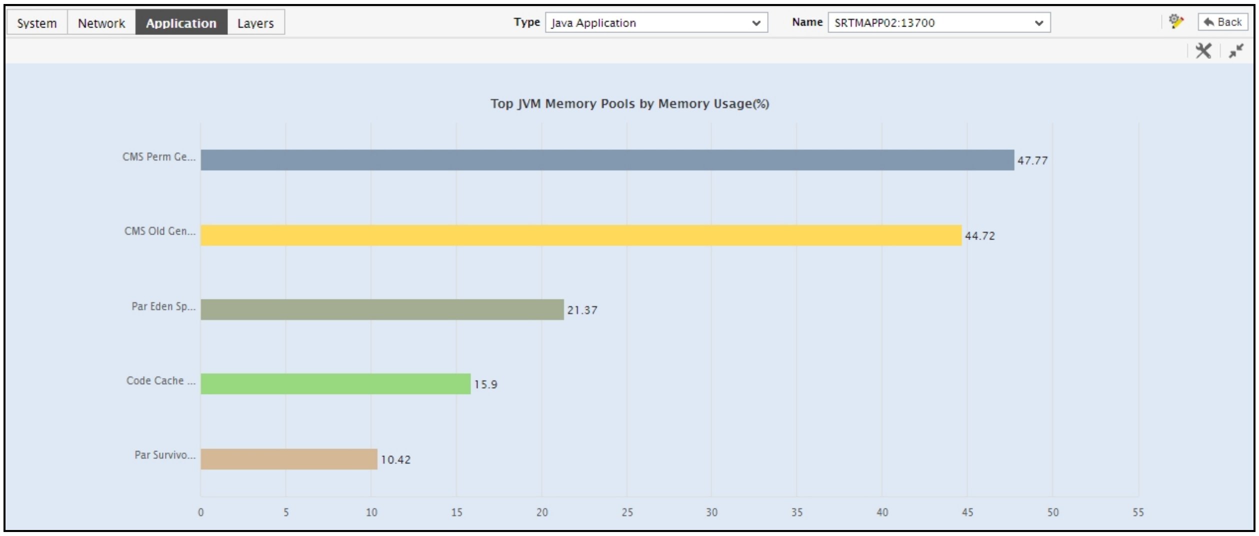

Figure 10.101: The enlarged memory usage graph

5. Also, though the enlarged graph displays the **TOP-10** memory pools on the JVM by default, you can choose to view a more or a less number of memory pools by picking a **TOP-N** or **LAST-N** option from the **Show** list in the enlarged graph.
6. In contrast to the **Comparison** tab page, which, by default, reports the current memory usage levels of individual memory pools, the **History** tab page displays measure graphs that depict how each memory pool has been using the JVM memory over time. In the event of a memory contention, this time-bound analysis will help you easily differentiate between a sudden spike in memory usage and a consistent rise in the same.

Figure 10.102: The History tab page of the JVM Memory Dashboard

7. To alter the timeline for a single graph, just click on that graph - this will enlarge the graph. You can change

the **Timeline** of the graph in the enlarged mode.

Figure 10.103: An enlarged measure graph in the Comparison tab page of the JVM Memory dashboard

8. Besides changing the timeline, you can also change the number of best/worst performers whose performance results are to be plotted in the graph. By default, the enlarged graph reveals the variations in the performance of the **TOP-10** memory pools. If need be, you can pick a different **TOP-N** or **LAST-N** option from the **Show** list in the enlarged graph.

9. Instead of these measure graphs, you can, if required, view summary graphs of the memory-related measures in the History tab page. For this, click on the icon at the right, top corner of the History tab page. Summary graphs help you figure out the percentage of time during the last 24 hours (by default) the Java application was hogged by memory-related issues. While monitoring mission-critical applications that are governed by rigid service level agreements, summary graphs will help you determine whether the guaranteed memory usage levels were fulfilled or not, and if not, how often did the usage levels slip.

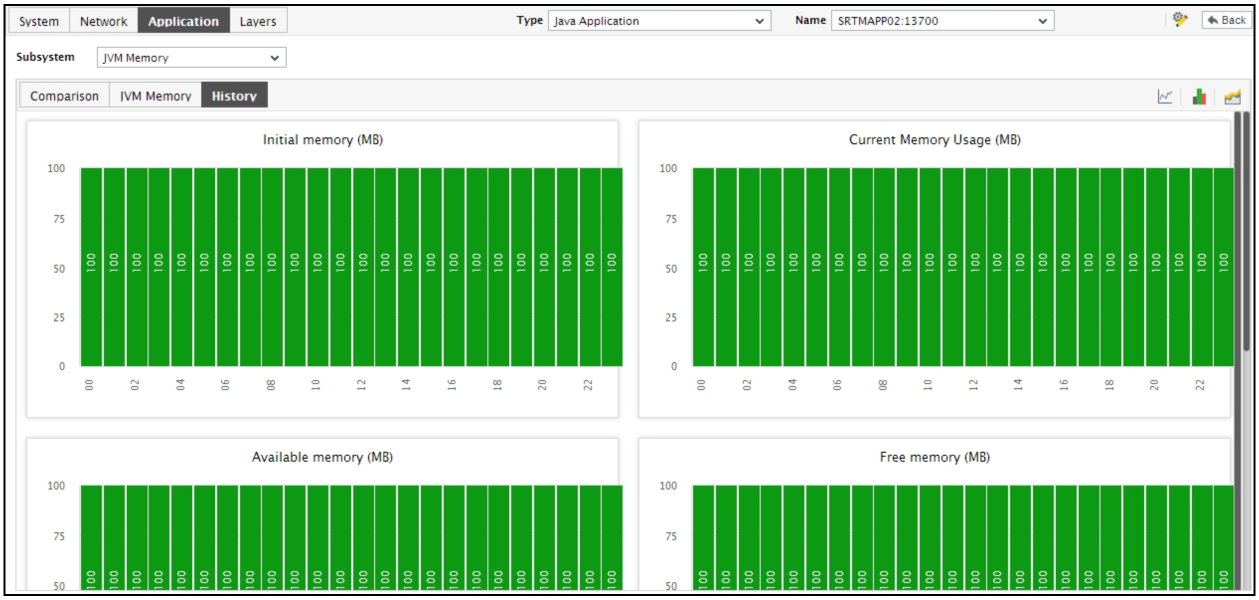

Figure 10.104: Summary graphs displayed in the History tab page of the JVM Memory Dashboard

10. Here again, you can change the **Timeline** of all the summary graphs by clicking on the **Timeline** link in Figure 10.105, or click on a graph, enlarge it, and change its **Timeline** in the enlarged mode. Also, though the graph plots hourly summary values by default, you can pick a different **Duration** for the graph in the enlarged mode, so that daily/monthly performance summaries can be analyzed.

11. You can click on the  icon at the right, top corner of the **History** tab page to view trend graphs of the memory usage-related measures. By default, these trend graphs plot the maximum and minimum memory usage values for every hour of the last 24 hours (by default).

12. Using these trend graphs, you can understand the variations in the memory usage of each pool during the last 24 hours (by default), deduce the future usage trends, and accordingly recommend changes to the memory pool size.

Figure 10.105: Trend graphs displayed in the History tab page of the JVM Memory Dashboard

13. Click on a graph, enlarge it, and change its **Timeline** in the enlarged mode. Also, though the graph plots hourly trend values by default, you can pick a different **Duration** for the graph in the enlarged mode, so that daily/monthly performance trends can be analyzed.

14. Also, by default, the trend graph only plots the minimum and maximum values registered by a measure. Accordingly, the **Graph** type is set to **Min/Max** in the enlarged mode. If need be, you can view the **Graph** type of the measure for the given **Timeline**. Such a graph will enable you to assess whether the memory resources were utilized effectively or not, over time.

    Note:

    In case of descriptor-based tests, the **Summary** and **Trend** graphs displayed in the **History** tab page typically plot the values for a single descriptor alone. To view the graph for another descriptor, pick a descriptor from the drop-down list made available above the corresponding summary/trend graph.

15. At any point in time, you can switch to the measure graphs by clicking on the button.

16. Typically, the **History** tab page displays measure, summary, and trend graphs for a default set of measures.

10.3.3 JVM Thread

If you require an integrated dashboard for analyzing the present/past performance and problem information pertaining to the threads executing on the JVM, so that you can efficiently and accurately diagnose the root-cause of the thread-related abnormalities, select the **JVM Thread** option from the **Subsystem** list. Using this single, central dashboard, you can ascertain the following quickly and easily:

- Are any threads currently executing on the JVM? Which of these threads are consuming CPU excessively and why? Has the CPU consumption of the thread been high always or is this a sudden occurrence?

- Are any threads blocked currently? Have they been blocked for too long a time? Why did it happen?

- Are any threads in a deadlock? If so, what caused the deadlock?
- Have any threads been waiting for too long a time for other threads to release an object? If so, how long, and what caused the waiting?

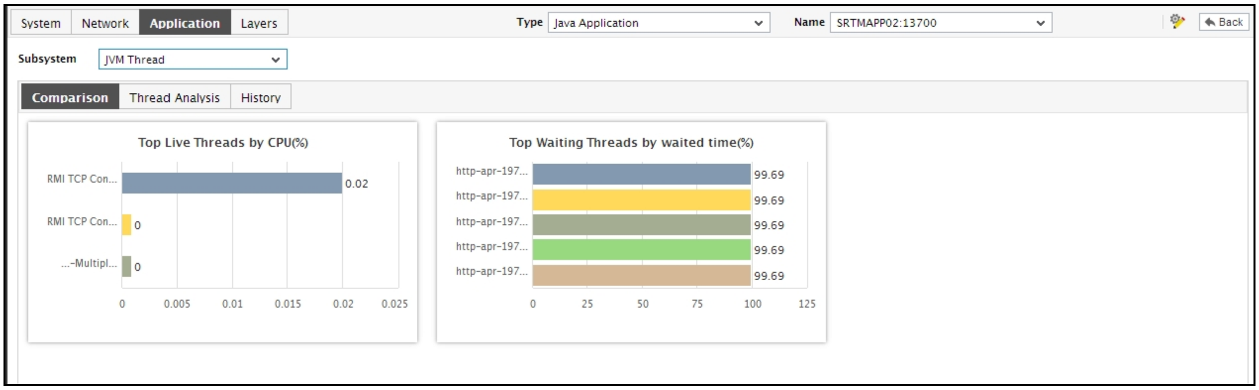

Figure 10.106: The JVM Thread Dashboard

The contents of this dashboard are discussed hereunder:

1. The **Thread Analysis** section, by default, displays the **stack trace** of each of the threads that are currently running on the JVM. Accordingly, the default selection in the **Analysis By** list is **Runnable threads**. To view the stack trace of those threads that are in a different state (say, **Blocked**, **Waiting**, **Timed waiting**, etc.), you will have to pick a different option from the **Analysis By** list.

2. A **stack trace** (also called **stack backtrace** or **stack traceback**) is a report of the active stack frames instantiated by the execution of a program. It is commonly used to determine what threads are currently active in the JVM, and which threads are in each of the different states - i.e., alive, blocked, waiting, timed waiting, etc.

   Typically, when a Java application begins exhibiting erratic resource usage patterns, it often takes administrators hours, even days to figure out what is causing this anomaly - could it be owing to one/more resource-intensive threads being executed by the application? If so, what is causing the thread to erode resources? Is it an inefficient piece of code? In which case, which line of code could be the most likely cause for the spike in resource usage? To be able to answer these questions accurately, administrators need to know the complete list of threads that the application executes, view the **stack trace** of each thread, analyze each stack trace in a top-down manner, and trace where the problem originated.

   The **JVM Thread** dashboard simplifies this seemingly laborious procedure by not only alerting administrators instantly to excessive resource usage by a thread, but also by providing the administrator with quick and easy access to the **stack trace** information of that thread; with the help of stack trace, administrators can effortlessly drill down to the exact line of code that requires optimization.

3. Regardless of the **Analysis By** option chosen, the thread list in the **Thread Analysis** section is sorted in the descending order of the percentage CPU time of the threads. We can thus conclude that the first thread for which stack trace is provided in this section is the top consumer of CPU. In the event of abnormally high CPU usage by this thread, you can use the stack trace to determine which line of code executed by this thread was causing the CPU usage to soar.

4. You will have to scroll down the **Thread Analysis** section to view the stack trace of the other threads. Alternatively, you can click on the icon next to the **Analysis By** list to invoke the **Thread Analysis** window (see Figure 10.107) using which you can quickly review the stack trace of each of the top CPU consumers.

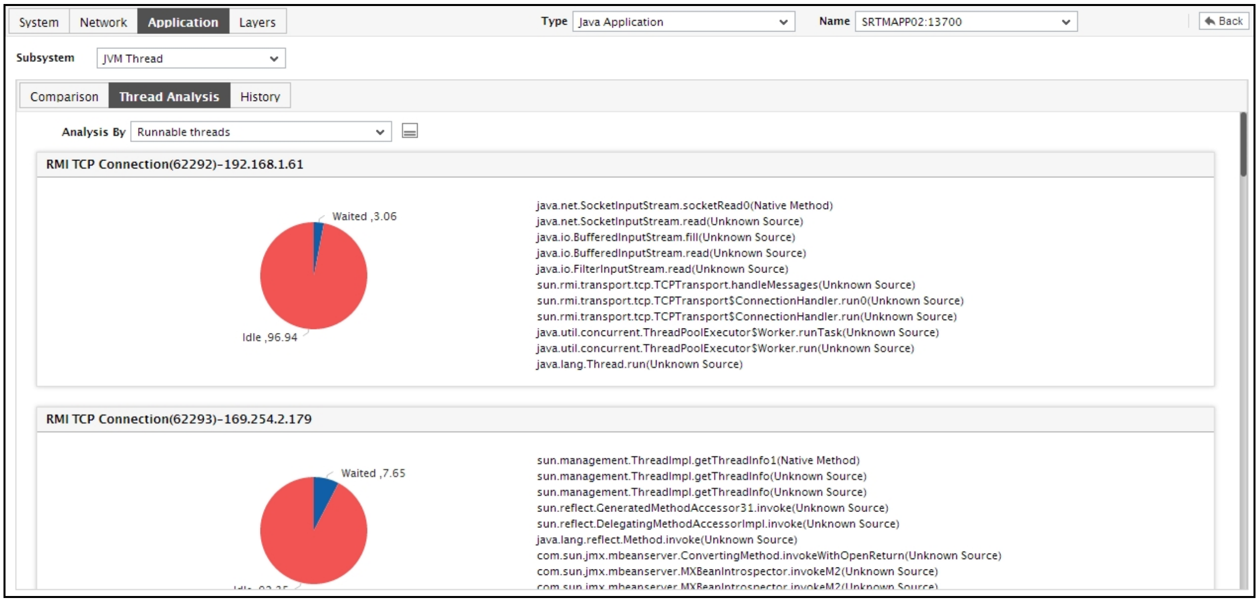

Figure 10.107: The Thread Analysis window

5. Below the **Thread Analysis** section is the **Comparison** tab page that displays a series of top-10 charts. These charts, by default, aid the quick and accurate identification of the thread that is currently consuming the maximum CPU, the thread that has been blocked for the longest time, and the thread that has been in waiting for the longest time. You can override this default setting by including comparison graphs for more measures in the **Comparison** tab page, or by removing one/more existing graphs from this tab page.

6. To view a comparison graph more clearly, click on it; this enlarges the comparison graph as depicted by Figure 10.108.

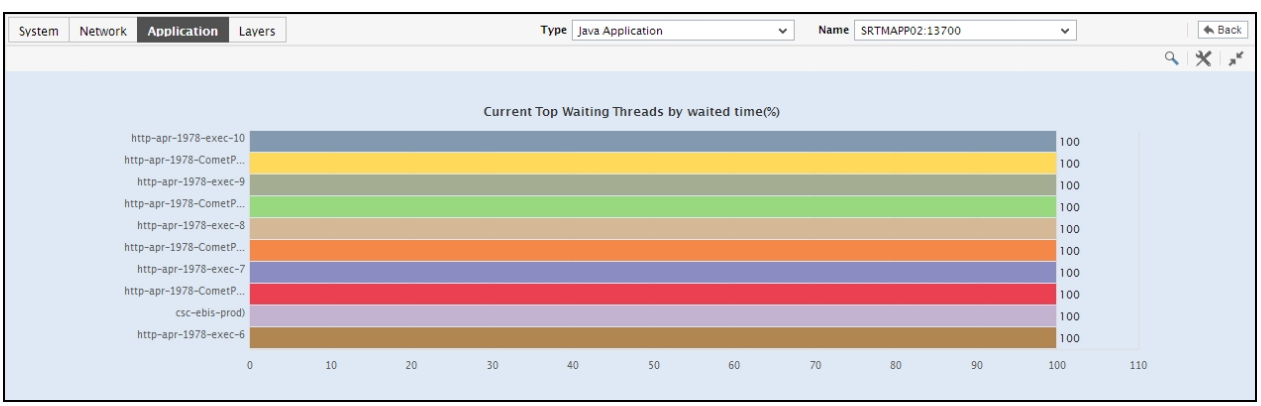

Figure 10.108: An enlarged top-n bar graph in the JVM Thread dashboard of the Java Application

7. In the enlarged mode, you can pick a different **TOP-N** or **LAST-N** option from the **Show** list to focus on a more or a less number of JVM threads. Also, to perform a post-mortem on issues that occurred in the past and to zero-in on threads that may have contributed to this past problem, click on the **Compare History** link in Figure 10.109 and provide a **Timeline** for this comparison. A comparison bar graph indicating the top-10 (by default) threads in a specific performance area during the specified timeline, will then appear.
8. For historically analyzing the state of the JVM threads, click on the **History** tab page. This tab page displays time-of-day graphs for all the thread-related measures for a default duration of 24 hours.

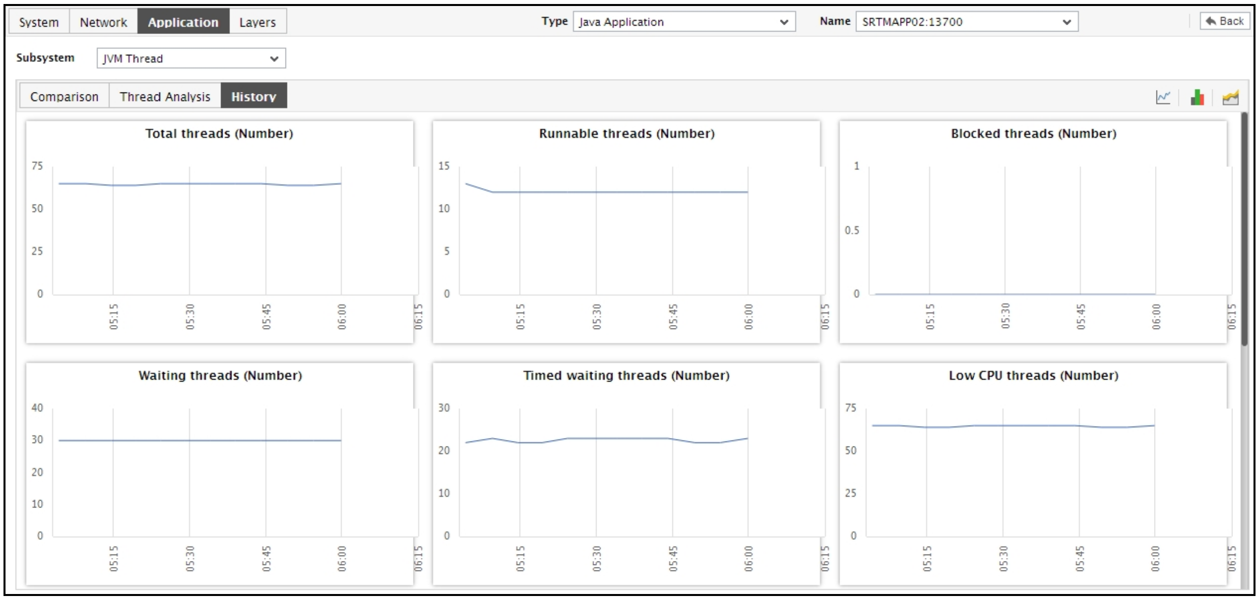

Figure 10.109: The History tab page of the JVM Thread Dashboard

9. Say, you suddenly notice that the number of blocked threads has increased; in such a case, you can use these measure graphs to figure out when during the last 24 hours the block occurred. If required, you can even look beyond the last 24 hours - i.e., you can find out whether the anomaly originated much earlier. For this, you just need to click on the graph of interest to you. This will enlarge the graph; in the enlarged mode, you can alter the graph **Timeline**, so that the performance of that measure can be analyzed over a broader time window.

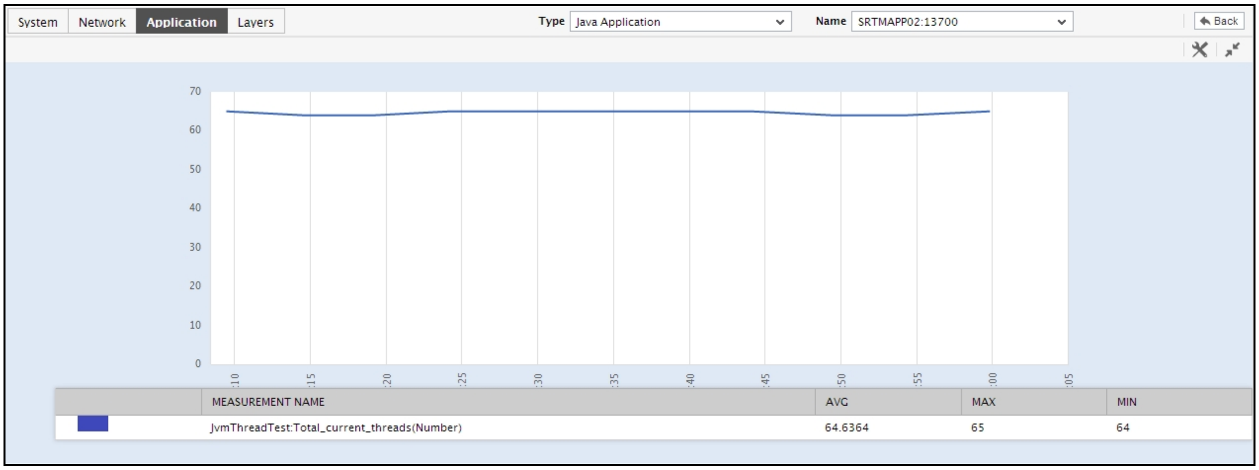

Figure 10.110: An enlarged measure graph in the History tab page of the JVM Thread dashboard

10. To view summary graphs on thread state instead of the default measure graphs, just click on the icon at the right, top corner of the **History** tab page. Figure 10.110 will then appear. The summary graphs of Figure 10.111 reveal the percentage of time during the last 24 hours (by default) the Java application has been affected by thread-related issues, and the type of issues (whether critical/major/minor) the application was experiencing. These graphs help determine whether the assured service levels were delivered or not. For instance, if the Java application assures its users of zero deadlock situations, then the summary graph will reveal whether the Java application did encounter any deadlock situation, and if so for how long - this way, while performing service level audits on the application, the auditor can understand whether the application delivered on its promise or has slipped. If service level slippages are detected, then the summary graphs will also reveal the extent of the slip - i.e., the percentage of time for which the desired service levels were not delivered.

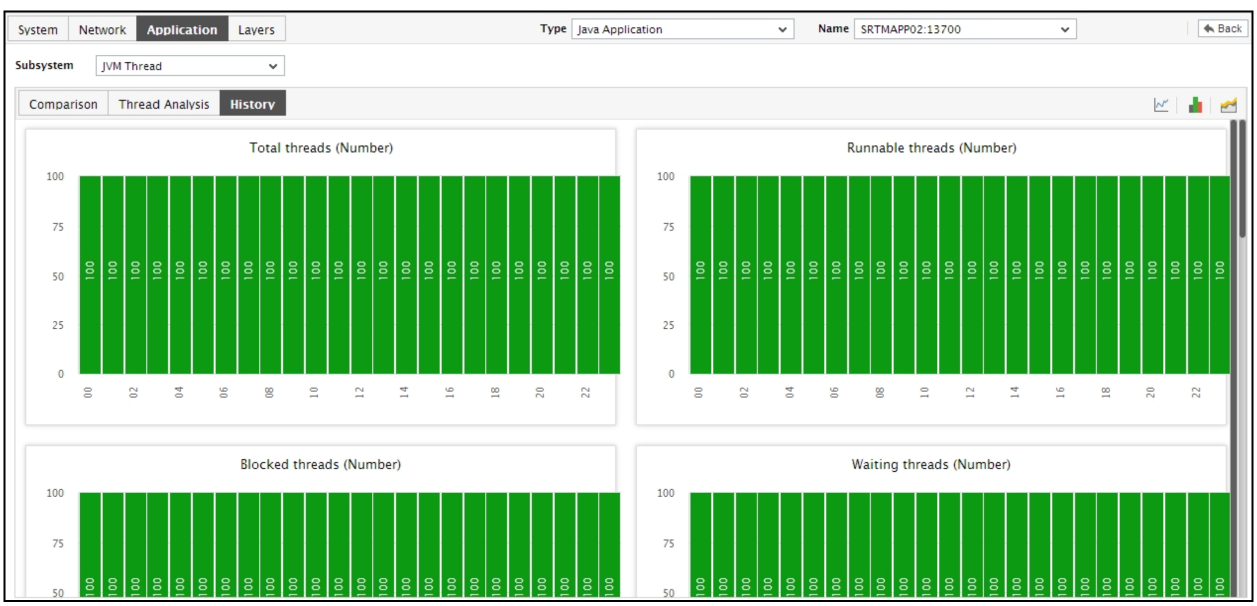

Figure 10.111: Summary graphs displayed in the History tab page of the JVM Thread Dashboard

11. Use the **Timeline** link at the right, top corner of the tab page to change the timeline of all the summary graphs at one shot. For altering the timeline of a single graph, click on it; this will enlarge the graph. In the enlarged mode, you can change the **Timeline** of the summary graph and modify the dimension (3D/2D) of the graph. Also, by default, hourly summaries are plotted in the summary graph; you can configure these graphs to plot daily/monthly summaries instead by picking the relevant option from the **Duration** list in the enlarged mode.
12. If you want to view the past trends in thread performance, click on the icon at the right, top corner of the **History** tab page. Figure 10.112 will then appear. Using the trend graphs displayed inFigure 10.112, you can better assess the current capacity of your application and can accordingly plan its future capacity. By default, these trend graphs plot the maximum and minimum values registered by every thread-related measure during every hour of the last 24 hours. From this data, you can clearly figure out when during the last 24 hours the application performance has peaked and when it has been below-normal.

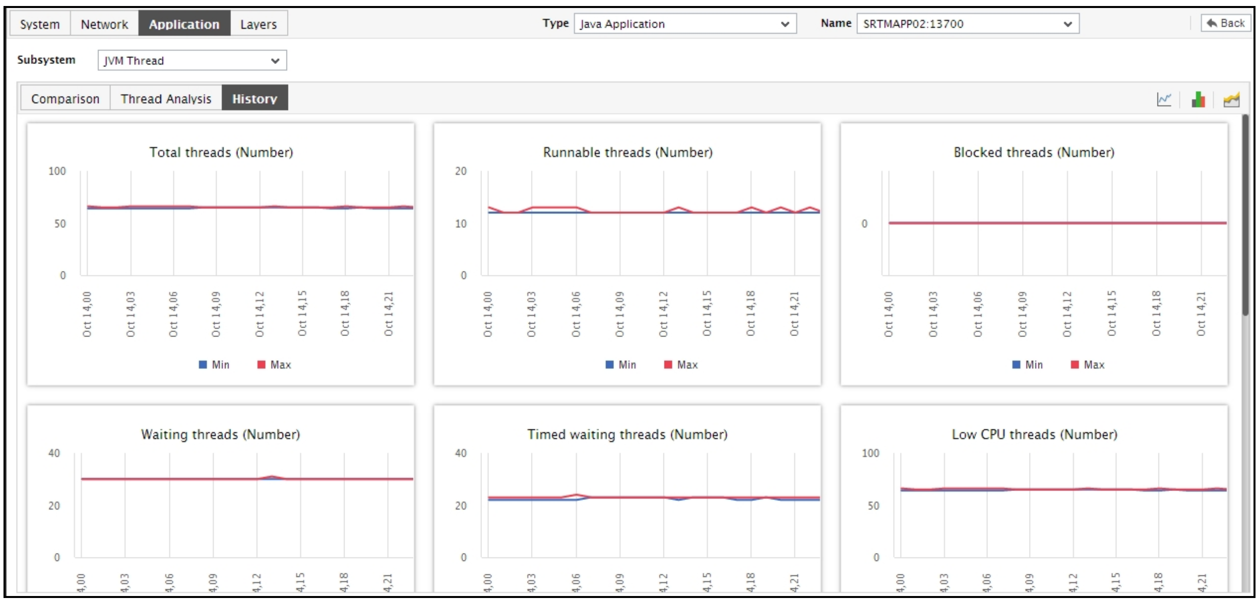

Figure 10.112: Trend graphs displayed in the History tab page of the JVM Thread Dashboard

13. Use the **Timeline** link at the right, top corner of the tab page to change the timeline of all the trend graphs at one shot. For altering the timeline of a single graph, click on it; this will enlarge the graph. In the enlarged mode, you can change the **Timeline** of the trend graph and modify the dimension (3D/2D) of the graph. Also, by default, hourly trends are plotted in the trend graph; you can configure these graphs to plot daily/monthly trend values instead by picking the relevant option from the **Duration** list in the enlarged mode. Moreover, by default, the trend graphs plot only the minimum and maximum values registered by a measure during the specified timeline - this graph will enable you to isolate those times at which performance of that measure had peaked and the times it had fared poorly. For instance, using the default trend graph for the *Blocked threads* measure, you can clearly identify when too many threads were blocked and when blocked threads were minimum. If need be, you can view the **Avg** option from the **Graph type** list in the enlarged mode to make sure that the trend graph plots the average trend values for the specified timeline - in the case of the above example, such a graph will help you understand how the blocked threads count has varied during the set timeline. Alternatively, you can select the **Sum** option from the **Graph type**

list to have the trend graph plot the sum of trends for the specified timeline.

Note:

In case of descriptor-based tests, the **Summary** and **Trend** graphs displayed in the **History** tab page typically plot the values for a single descriptor alone. To view the graph for another descriptor, pick a descriptor from the drop-down list made available above the corresponding summary/trend graph.

14. At any point in time, you can switch to the measure graphs by clicking on the button.

10.3.4 JVM Classes

Select the **JVM Classes** option from the **Subsystem** list to know how efficiently the class loader used by the Java application is and has been loading/unloading classes onto memory. Upon selection, 10.3.4 will appear.

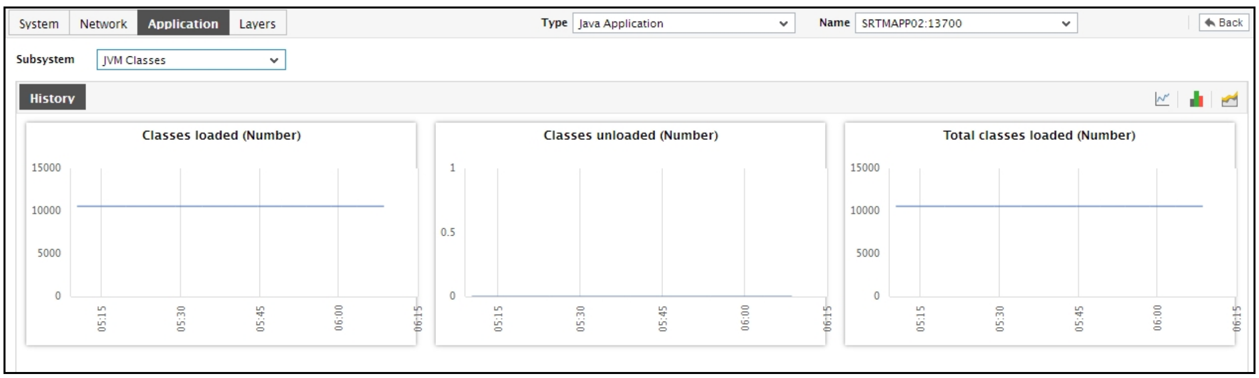

Figure 10.113: The JVM Classes Dashboard

The contents of this dashboard are as follows:

1. The **JVM Classes** section indicates the current health of the class loader by providing a pie chart that graphically depicts the number of classes currently loaded and unloaded by the application.

2. The **History** tab page below, by default, provides a series of measure graphs that reveal how the class loader has been performing over the default duration of the last 24 hours. If the number of classes loaded/unloaded dramatically decreases, it could indicate that the class loader is experiencing issues with loading/unloading. In such a case, a look at these measure graphs will help you figure out when exactly the bottleneck surfaced - did it happen suddenly or is it a condition that has become worse with time?

3. If need be, you can even alter the timeline of all these measure graphs so that you can analyze performance across days and weeks; for this, simply click the **Timeline** link at the right, top corner of the **History** tab page and change the timeline for the graphs using the calendar that pops out. To change the timeline of a single graph alone, simply click on that graph to enlarge it, and then modify the **Timeline** of the graph in the enlarged mode.

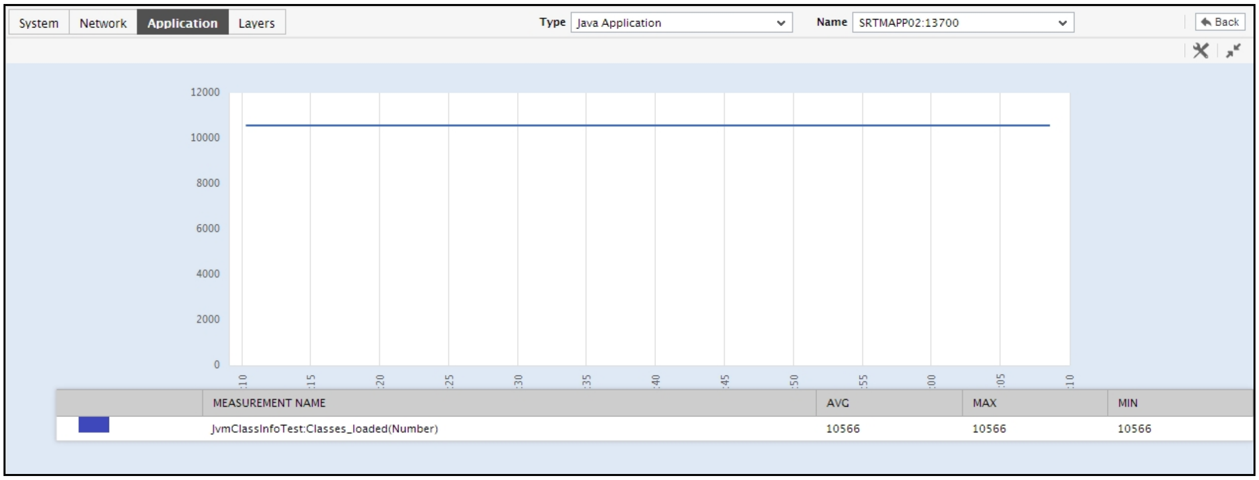

Figure 10.114: An enlarged measure graph in the History tab page of the JVM Classes dashboard

4. To determine the service level achievements / slippages of the class loader, you need to view summary graphs of the measures and not the default measure graphs. For this, just click on the icon at the right, top corner of the **History** tab page. Figure 10.114 then appears.

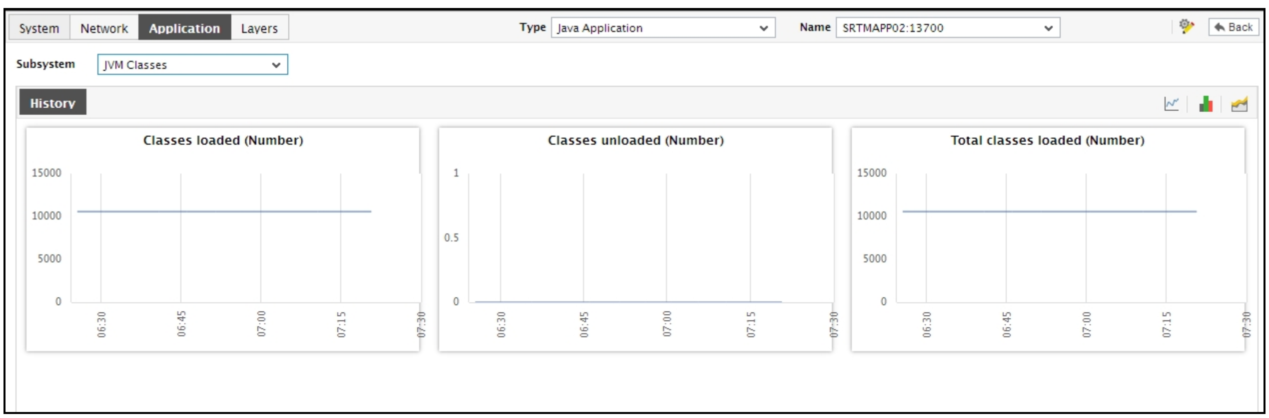

Figure 10.115: Summary graphs displayed in the JVM Classes Dashboard

The summary graphs displayed in Figure 10.115 reveal the percentage of time the Java application experienced problems in loading/unloading classes on to the memory. Besides revealing the efficiency of your administrative staff in recognizing bottlenecks and mitigating them, these summary graphs also indicate whether the class loader has been able to maintain the assured performance levels during the default duration of 24 hours.

5. In case of the summary graphs too, you can change the **Timeline** of all graphs by clicking on the **Timeline** link at the right, top corner of the **History** tab page. To alter the timeline of a single graph, here again, you will have to click on that graph, enlarge it, and modify the timeline. Also, by default, hourly summaries are plotted in the summary graph; you can configure these graphs to plot daily/monthly summaries instead by picking the relevant option from the **Duration** list in the enlarged mode.

6. To analyze past trends in the loading/unloading of classes, click on the icon at the right, top corner of the **History** tab page. Figure 10.116 will then appear.

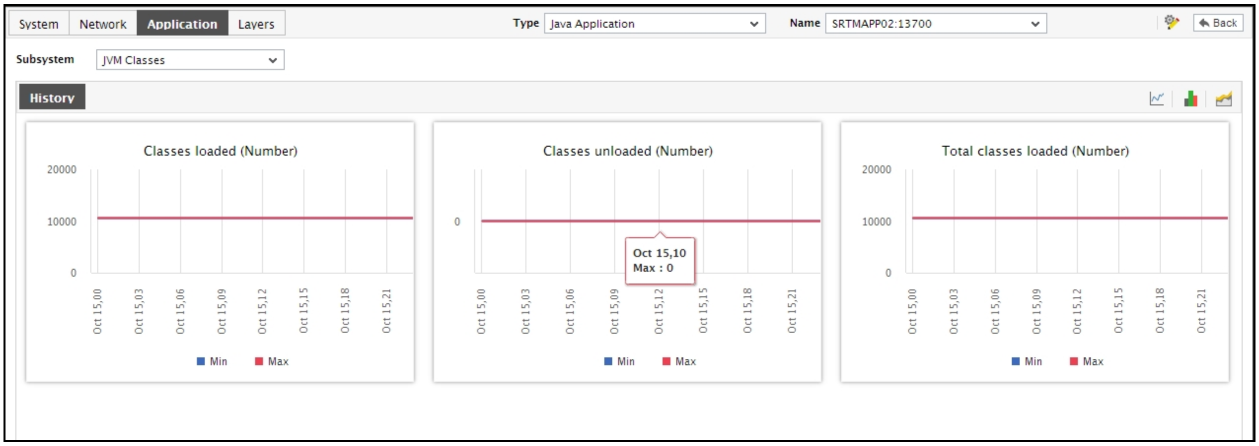

Figure 10.116: Trend graphs displayed in the JVM Classes Dashboard

7. These trend graphs, by default, plot the minimum and maximum values that every measure registered during each hour of the last 24 hours (by default). Using such graphs, you can accurately point to the time windows in which the class loader was actively loading/unloading classes, and the times at which there was a lull. By carefully observing these past trends, you can effectively gauge the workload that the application has been imposing on the class loader, predict future workloads accordingly, and suggest measures to enhance the efficiency of the loader

8. In addition, when a trend graph is enlarged, it is not just the **Timeline** that you can modify. The **Duration** of the graph can also be altered. By default, trend graphs reveal only the hourly trends in performance. By picking the relevant option from the **Duration** list, you can ensure that the trend graph in question plots daily/monthly trend values instead. Also, in the enlarged mode, the **Graph type** can also be modified. Since the default **Graph type** is **Min/Max**, the trend graph, by default, reveals the minimum and maximum values registered by a measure. If need be, you can select the **Avg** or **Sum** option from the **Graph type** list to plot average trend values of a measure or sum of trends (as the case may be) in the graph.

   Note:

   In case of descriptor-based tests, the **Summary** and **Trend** graphs displayed in the **History** tab page typically plot the values for a single descriptor alone. To view the graph for another descriptor, pick a descriptor from the drop-down list made available above the corresponding summary/trend graph.

9. At any point in time, you can switch to the measure graphs by clicking on the  button.

10.3.5 JVM Uptime

To investigate issues with availability and uptime of the JVM, select **JVM Uptime** as the **Subsystem**. Figure 10.117 will then appear.

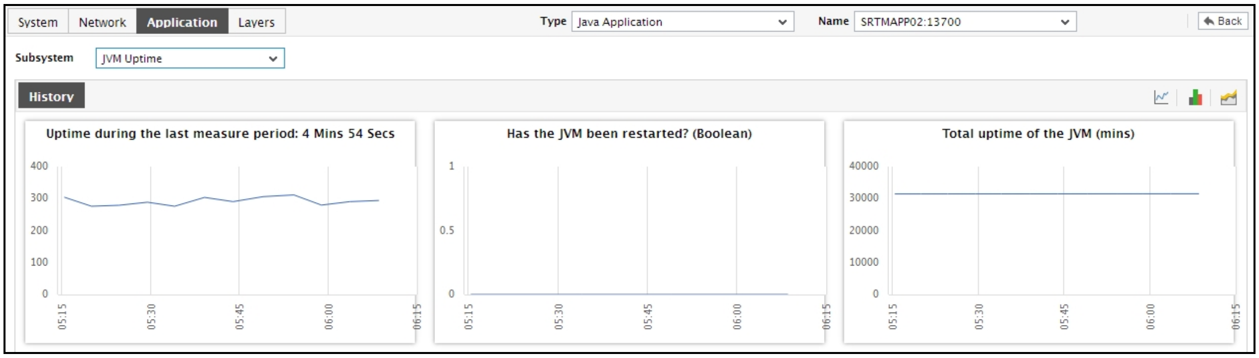

Figure 10.117: The JVM Uptime Dashboard

The contents of the uptime dashboard are as follows:

1. From the **JVM Uptime** section, you can determine the total duration for which the JVM has been up and running since its startup. If the JVM was started two days ago, but the uptime indicated by this section spans only a day, it is a clear indication that the JVM was unavailable for a while in-between; a possible reason for this could be a scheduled/unscheduled reboot of the Java application.
2. A careful study of this graph over time periods longer than 24 hours, can reveal intermittent breaks (if any) in JVM availability and failure of scheduled JVM reboots (if any). To ensure that all graphs plot values for longer time periods, click on the **Timeline** link at the right, top corner of the **History** tab page, and then change the timeline using the calendar that pops out. To modify the timeline for a particular graph alone, click on the graph to enlarge it, and alter the timeline in the enlarged mode. Besides the timeline, you can even change the graph dimension (**3D** / **2D**) in the enlarged mode.
3. Sometimes, you might have to periodically determine the percentage of time for which certain critical Java applications have been up and running, so that you know whether/not the application has been able to maintain the desired uptime levels. To run such uptime checks, summary graphs of the uptime measures are useful. To view summary graphs in the **History** tab page, click on the icon at the right, top corner of the **History** tab page. Figure 10.118 will then appear.

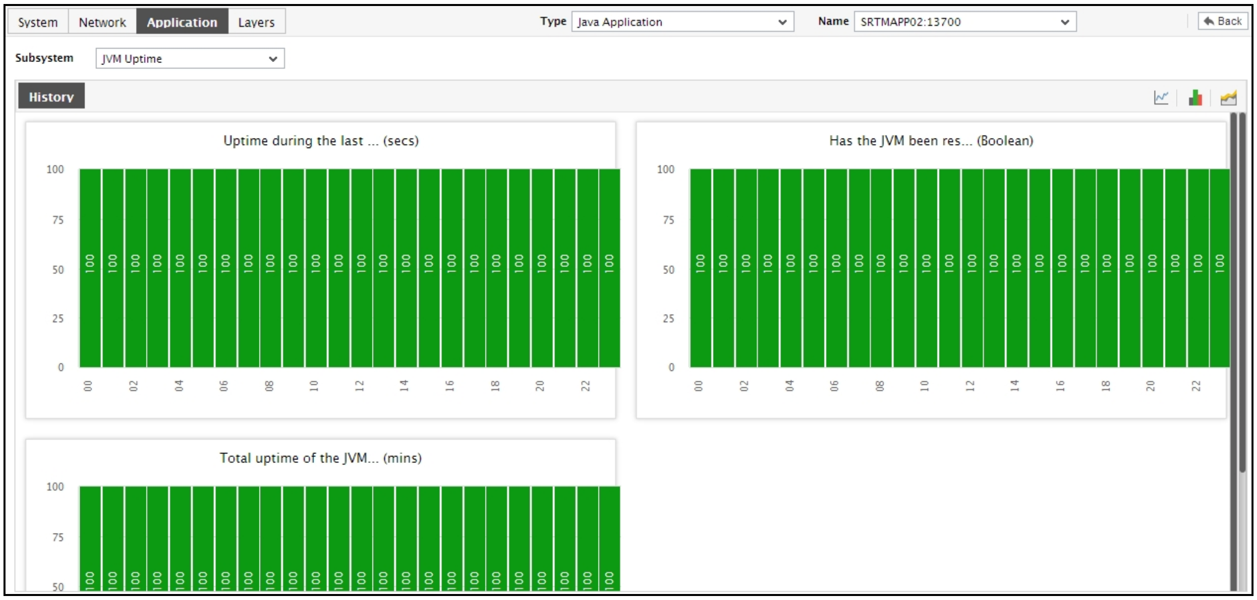

Figure 10.118: Summary graphs displayed in the JVM Uptime Dashboard

4. To perform the summary analysis over a broader time window, click on the **Timeline** link at the right, top corner of the **History** tab page and change the timeline; this will alter the timeline for all the graphs. To change the timeline of a particular graph alone, click on the graph to enlarge it, and then alter its timeline. Also, by default, hourly summaries are plotted in the summary graph; you can configure these graphs to plot daily/monthly summaries instead by picking the relevant option from the **Duration** list in the enlarged mode.

5. Similarly, you can analyze uptime trends by viewing trend graphs in the **History** tab page. For this, click on the icon at the right, top corner of the tab page. Figure 10.118 will then appear.

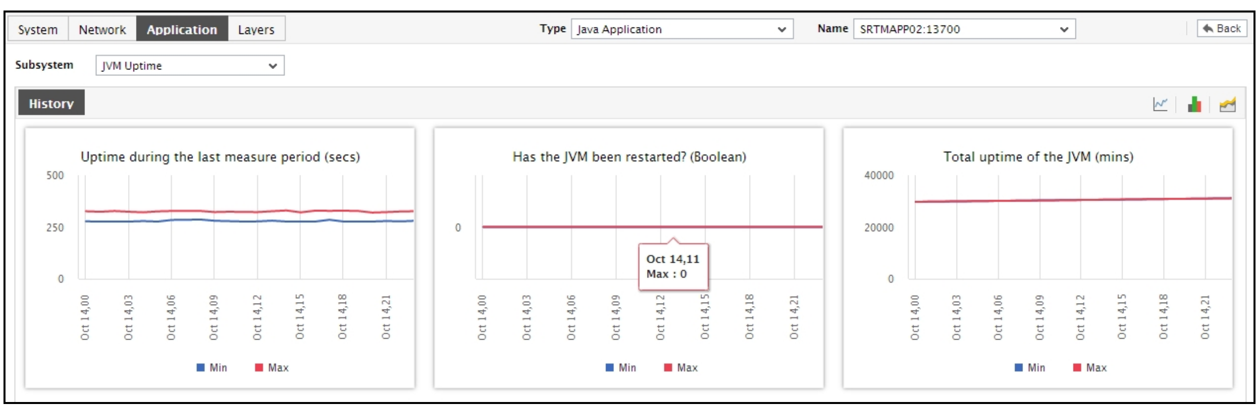

Figure 10.119: Trend graphs displayed in the JVM Uptime Dashboard

6. These trend graphs, by default, plot the minimum and maximum values registered by every uptime-related measure during every hour for the last 24 hours. Using these graphs, you can ascertain when during the last 24 hours uptime was very high, and when it was low.

7. To change the timeline of a particular graph, click on the graph to enlarge it, and then alter its timeline. The

Graph type can be viewed in the enlarged mode. By default, the graph **Duration** is **Hourly**, indicating that trend graphs plot hourly trend values by default. To ensure that these graphs plot the daily/monthly trend values instead, select the relevant option from the **Duration** list. Similarly, as already mentioned, trend graphs plot only the minimum and maximum values registered by a measure during the specified timeline. Accordingly, the **Graph type** is set to **Min/Max** by default in the enlarged mode. If you want the trend graph to plot the average trend values instead, set the **Graph type** to **Avg**. On the other hand, to configure the trend graph to plot the sum of trends set the **Graph type** to **Sum**.

Note:

In case of descriptor-based tests, the **Summary** and **Trend** graphs displayed in the **History** tab page typically plot the values for a single descriptor alone. To view the graph for another descriptor, pick a descriptor from the drop-down list made available above the corresponding summary/trend graph.

8. At any point in time, you can switch to the measure graphs by clicking on the button.

10.4 Business Dashboard

IT executives often require a high-level view of the performance of their mission-critical business services. The **Business Dashboard** of v6 provides this view in a form that is easy to comprehend and analyze. This dashboard quickly compares service demand with resource consumption and service quality to enable IT executives swiftly determine where service performance is most likely bottlenecked – at the demand level? resource consumption level? or user experience level? Moreover, it allows IT executives to rapidly triage performance issues tier-wise, so that they can accurately isolate the tier where the problem originated.

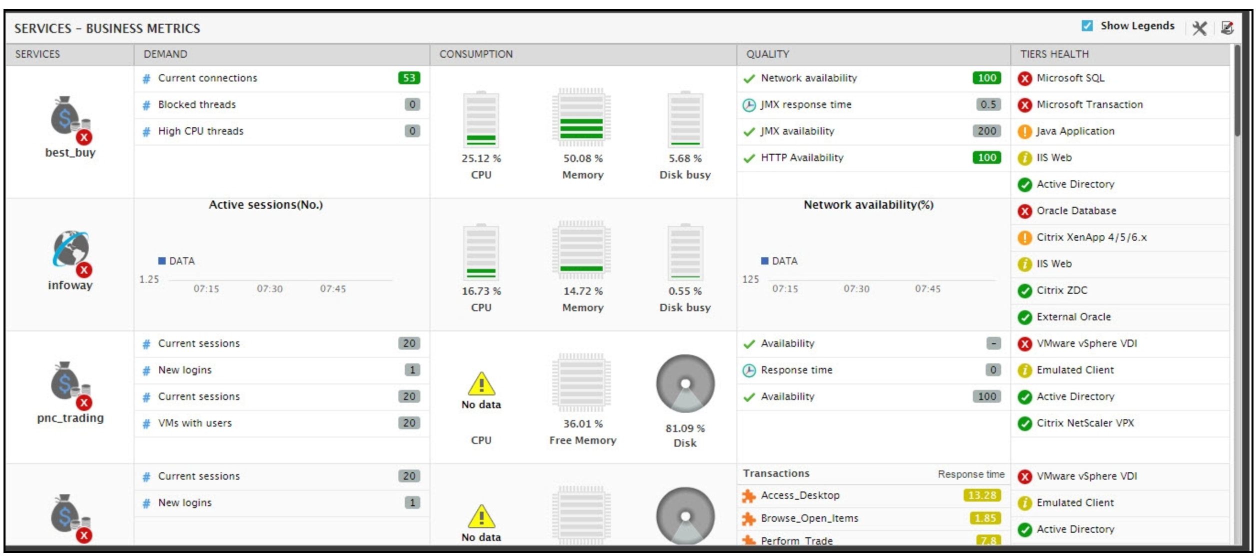

Figure 10.120: Figure 10.120: The Business Dashboard

You can customize the Business Dashboard to include demand, consumption, and quality indicators of your choice. For this, do the following:

Click the icon to the right, top corner of the Business dashboard (as indicated by Figure 10.1). Figure 10.2 will then appear.

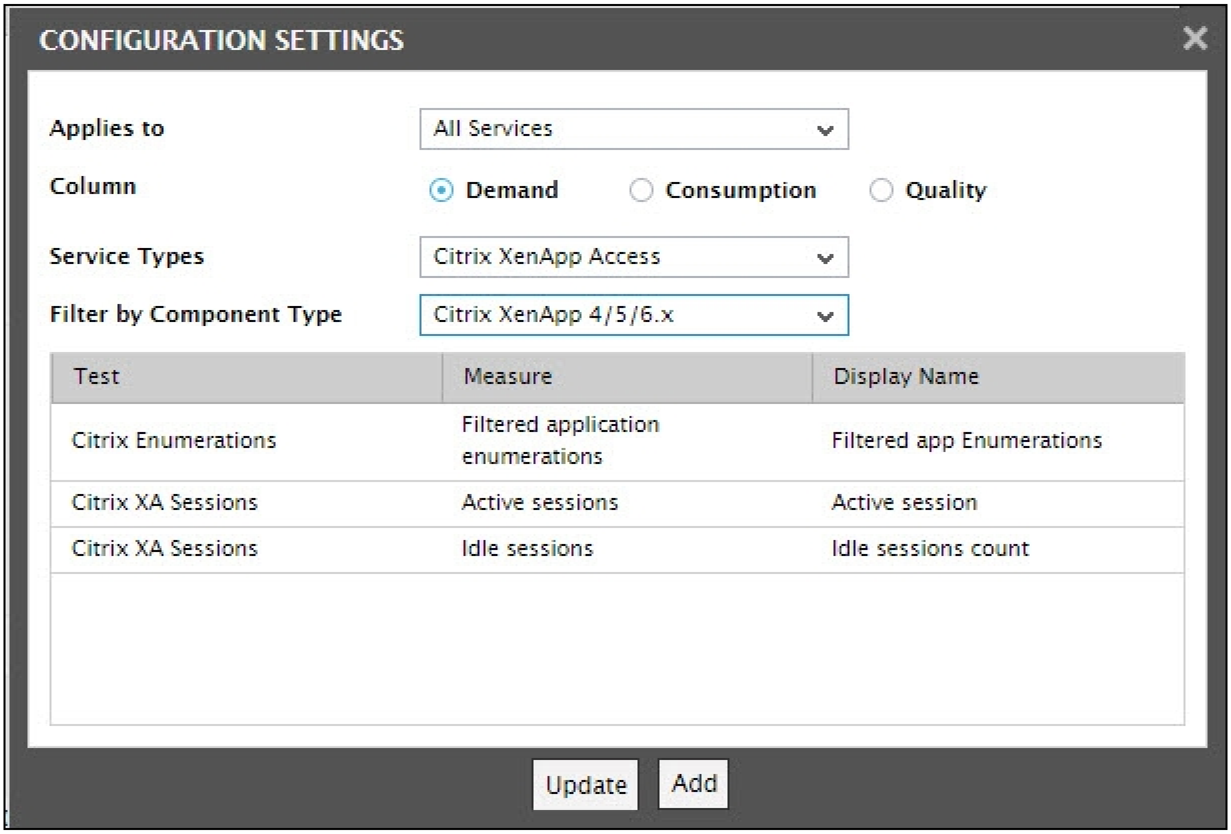

Figure 10.121: Figure 10.121: Configuring the business dashboard

1. In the **CONFIGURATION SETTINGS** window, first select the service to which the changes apply from the **Applies tp** drop-down. The steps that follow will vary depending upon your choice of service. By default, the **All Services** option will be selected, indicating that the settings being configured apply to all services by default. However, if you want to make changes to a particular service's settings alone, select the service of interest from the **Applies to** list.

2. If you go with the default selection (i.e., the **All Services** option) of the **Applies to** list, then follow the steps below to configure your dashboard:

   - First, indicate which **Column** of the dashboard you intend to alter – **Demand, Consumption,** or **Quality**.
   - Then, pick a specific service type from the **Service Types** list. This implies that the changes you make subsequently will apply to all services of that type.
   - Next, select a component type from the **Filter by Component Type** list. This list will display all component types that are engaged in delivering all the managed services in the environment.
   - Once a component type is chosen from this list, then the **CONFIGURATION WINDOW** will display a table listing all measures related to that component type that have been pre-configured (if any) for display in the selected **Column** for all services of the chosen **Service Type**. If no such pre-configuration exists, then the table will be empty.
   - You can either modify/delete any of the listed measures or add more measures to the chosen **Column**.

- To add a measure, first click the **Add** button in Figure 10.2.
- When Figure 10.3 appears, select a **Layer** from all layers that are key to the performance of the **Service Type** selected. For instance, if **Citrix XenApp Access** is the **Service Type** chosen, then the **Layer** list will include all layers of the Citrix XenApp component that is imperative to a service of that type.

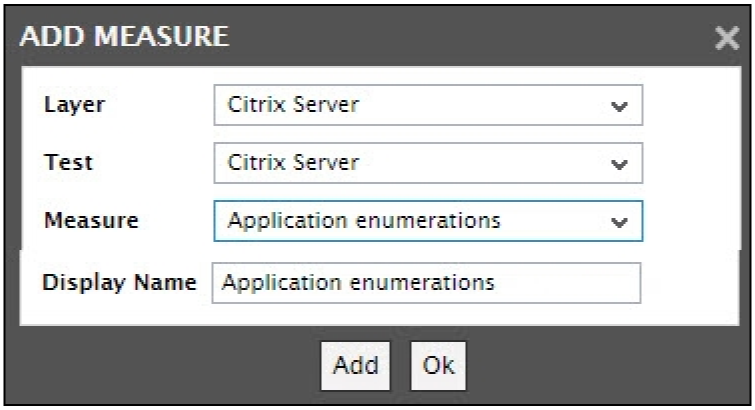

Figure 10.122: Figure 10.122: Adding a measure

  - Pick a **Test** mapped to the chosen **Layer**.
  - Select the **Measure** that you want displayed in the chosen **Column** of the dashboard.
  - If the chosen **Test** is a descriptor-based test, then, before adding any measure reported by that test, you should aggregate the measure value across all descriptors using a specific aggregate function. The result of this aggregation is then displayed in the dashboard. For this, pick a **Function**.

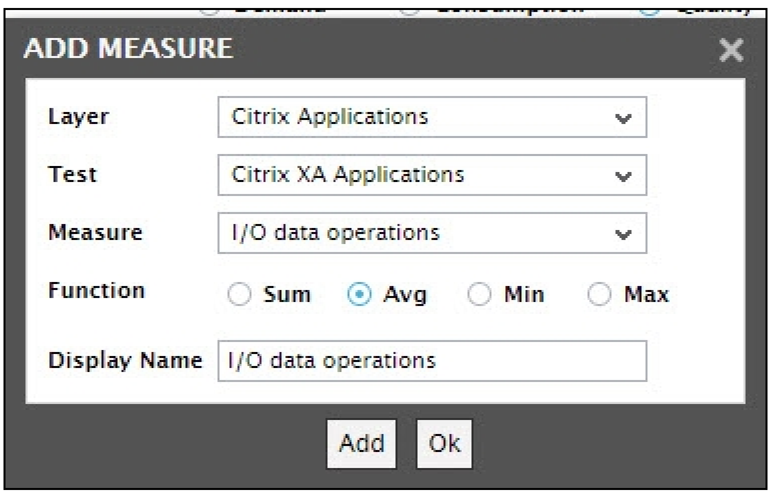

Figure 10.123: Figure 10.123: Adding a measure that supports descriptors

  - Next, provide a **Display Name** for the measure.
  - Click the **Add** button to add the measure. Similarly, you can add multiple measures to a column.
  - Once you are done adding, click **OK** to exit the **ADD MEASURE** pop-up.
  - The measure so added will then be appended to the measure list in the **CONFIGURATION SETTINGS**

window.

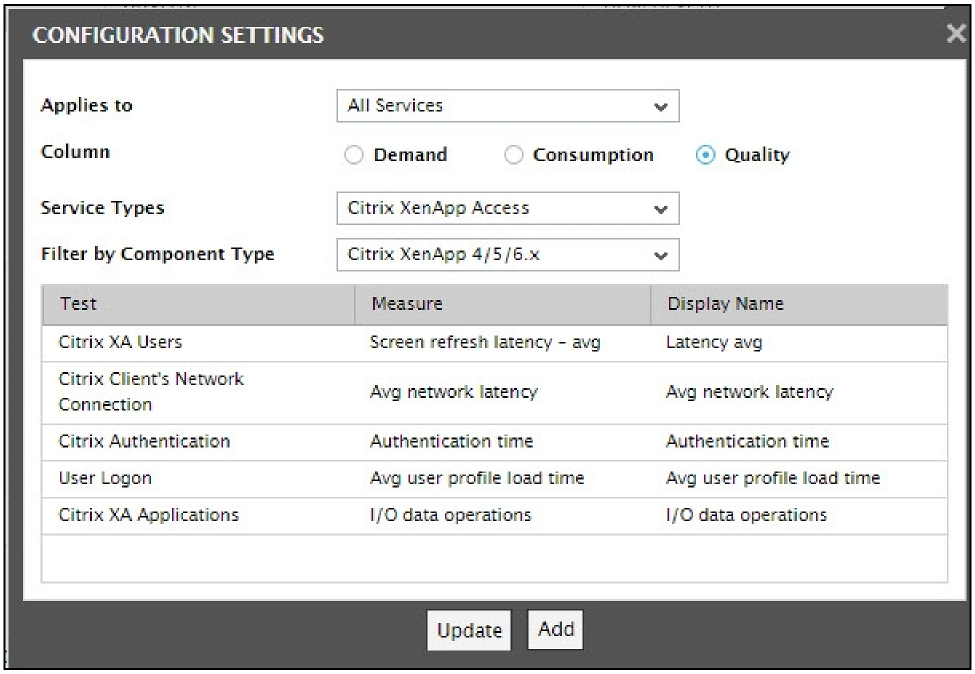

Figure 10.124: Figure 10.124: Added measure displayed in the configuration window

- To modify any of the listed measures, simply select the row representing that measure in Figure 10.5 and click the **Modify** button. When Figure 10.6 appears, you can change the **Display Name** of the chosen measure. In case of descriptor-based measures, you have the additional option of picking a different aggregate **Function** for the measure from Figure 10.6. Finally, click the **OK** button in Figure 10.6.

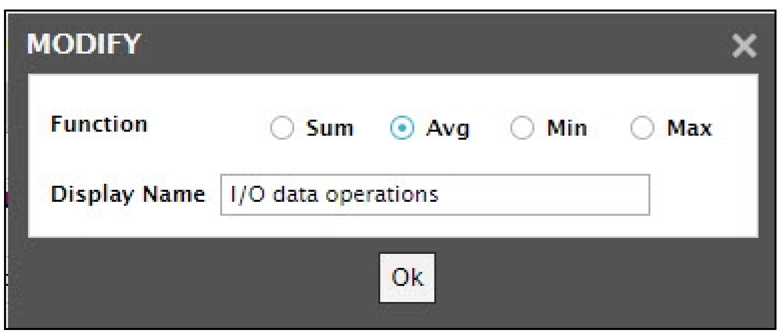

Figure 10.125: Figure 10.125: Modifying a measure that applies to All Services

- To delete a pre-existing measure, select the row representing that measure in Figure 10.5 and click the **Delete** button.

3. Now, if a specific service is chosen from the **Applies to** list, then follow the steps below to configure the dashboard:
   - First, indicate which **Column** of the dashboard you intend to alter – **Demand, Consumption,** or **Quality**.
   - Upon selecting a **Column**, the **CONFIGURATION SETINGS** window will automatically display the

measures related to the chosen service that have already been configured for display in that column. If no pre-configuration exists, then no measures will be displayed.

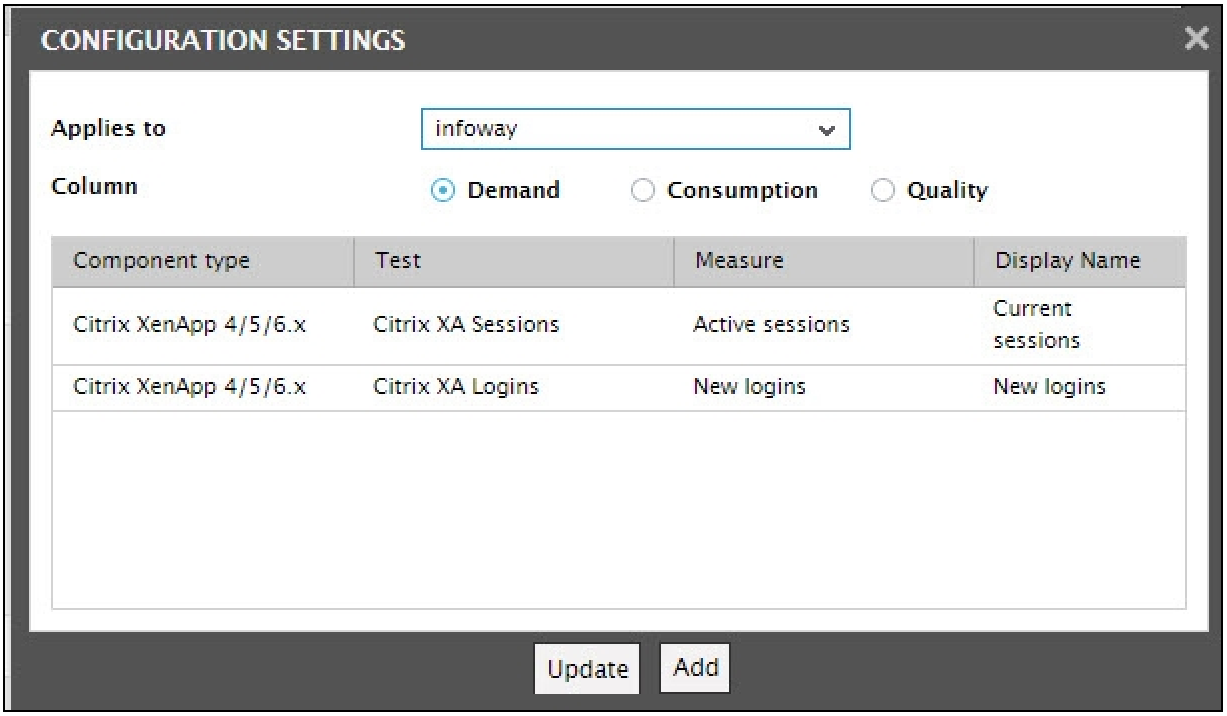

Figure 10.126: Figure 10.126: Configuring the dashboard settings for a specific service

- You can either modify/delete any of the listed measures or add more measures to the chosen **Column**. To add a measure, click the **Add** button in Figure 10.7.
- Figure 10.8 will then appear. Of the **Comp Types** that are part of the chosen service, select the component type from which measures are to be added to the dashboard.

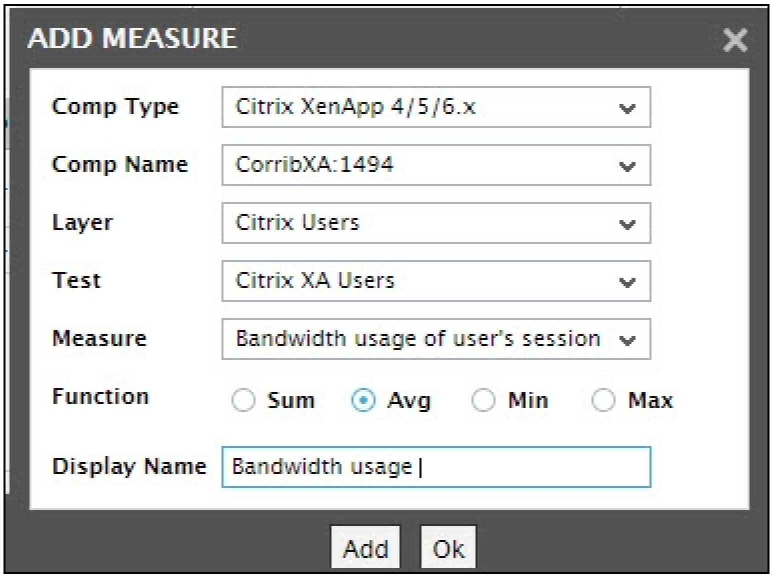

Figure 10.127: Figure 10.127: Adding a measure that applies to a specific service in the business dashboard

- From the **Comp Name** list, select a service component of the chosen type.
- Pick the **Layer** and the **Test** to which the measure to be added is mapped.
- Select the **Measure** to be added.
- If the chosen **Test** is a descriptor-based test, then choose the **Function** using which the chosen measure's values are to be aggregated across all descriptors of the test.
- Provide a **Display Name** for the measure.
- Click the **Add** button to add the measure. Similarly, you can add multiple measures to a column.
- Once you are done adding, click **OK** to exit the **ADD MEASURE** pop-up.
- The measure so added will then be appended to the measure list in the **CONFIGURATION SETTINGS** window.

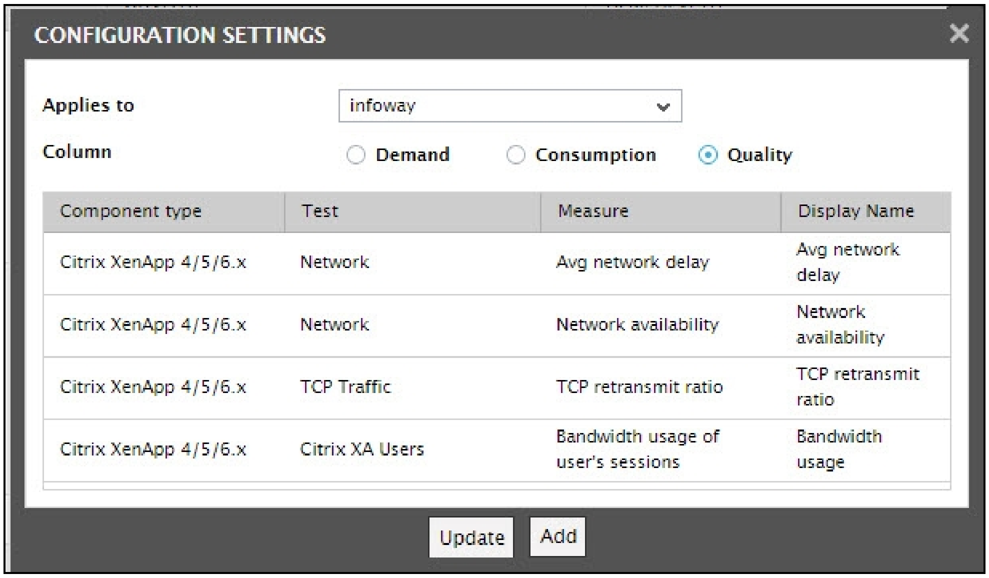

Figure 10.128: Figure 10.128: Added measure displayed in the configuration window

- To modify any of the listed measures, simply select the row representing that measure in Figure 10.9 and click the **Modify** button. When Figure 10.10 appears, you can change the **Display Name** of the chosen measure. In case of descriptor-based measures, you have the additional option of picking a different aggregate **Function** for the measure from Figure 10.10. Finally, click the **OK** button in Figure 10.10.

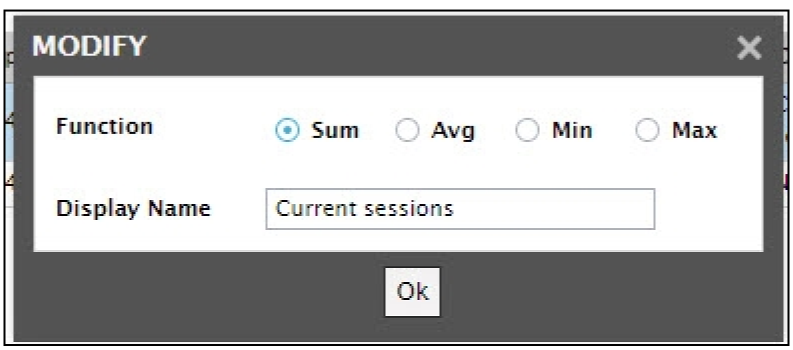

Figure 10.129: Figure 10.129: Modifying a measure that applies to a specific service

- To delete a pre-existing measure, select the row representing that measure in Figure 10.9 and click the **Delete** button.

You can also hide those services that you deem inconsequential from the business dashboard, so that you can focus on those services that are critical to you from a performance perspective. The capability to 'hide' services is not available by default. To enable this capability, do the following:

1. Edit the **eg_serviceperfdash.ini** in the **<EG_INSTALL_DIR>\manager\config** directory.
2. In the **[SERVICES_USER_CONFIG]** section, set the **admin:enableServiceSkipReorder** parameter to **true** (default is false).
3. Save the file.

Once this is done, then, to show/hide services, follow the steps below:

1. Click the icon to the right, top corner of the business dashboard to open the **CONFIGURATION SETTINGS** window.
2. You will find a new **What would you like to do?** drop-down in Figure 10.130 that then appears.

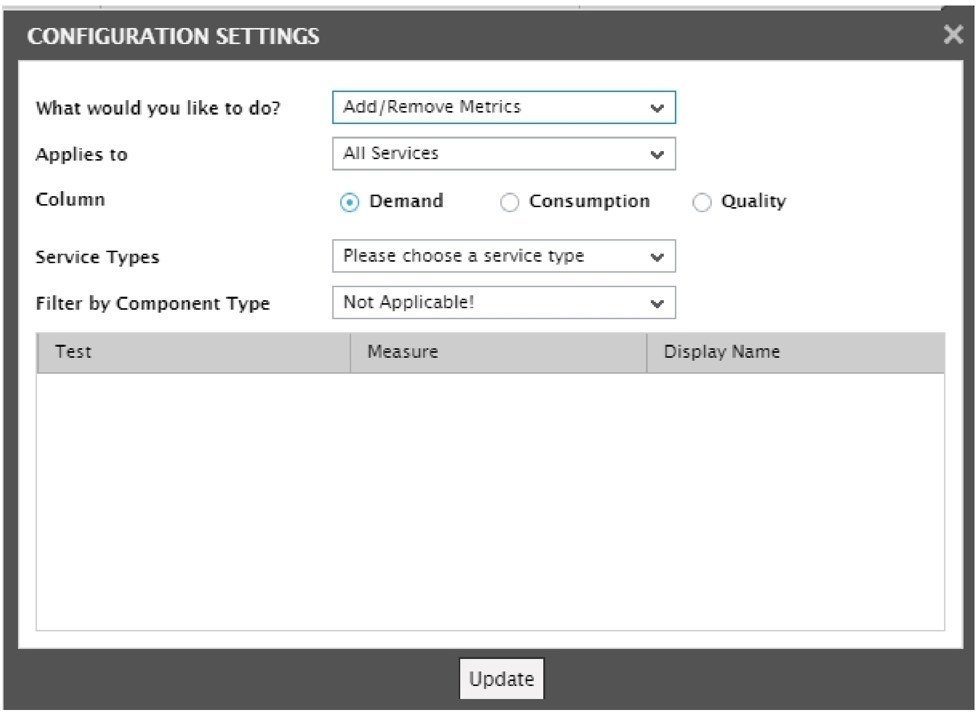

Figure 10.130: Figure 10.130: The Configuration Settings window with the option to show/hide services from the dashboard

3. Select the **Show/Hide Services** option from the **What would you like to do?** drop-down.

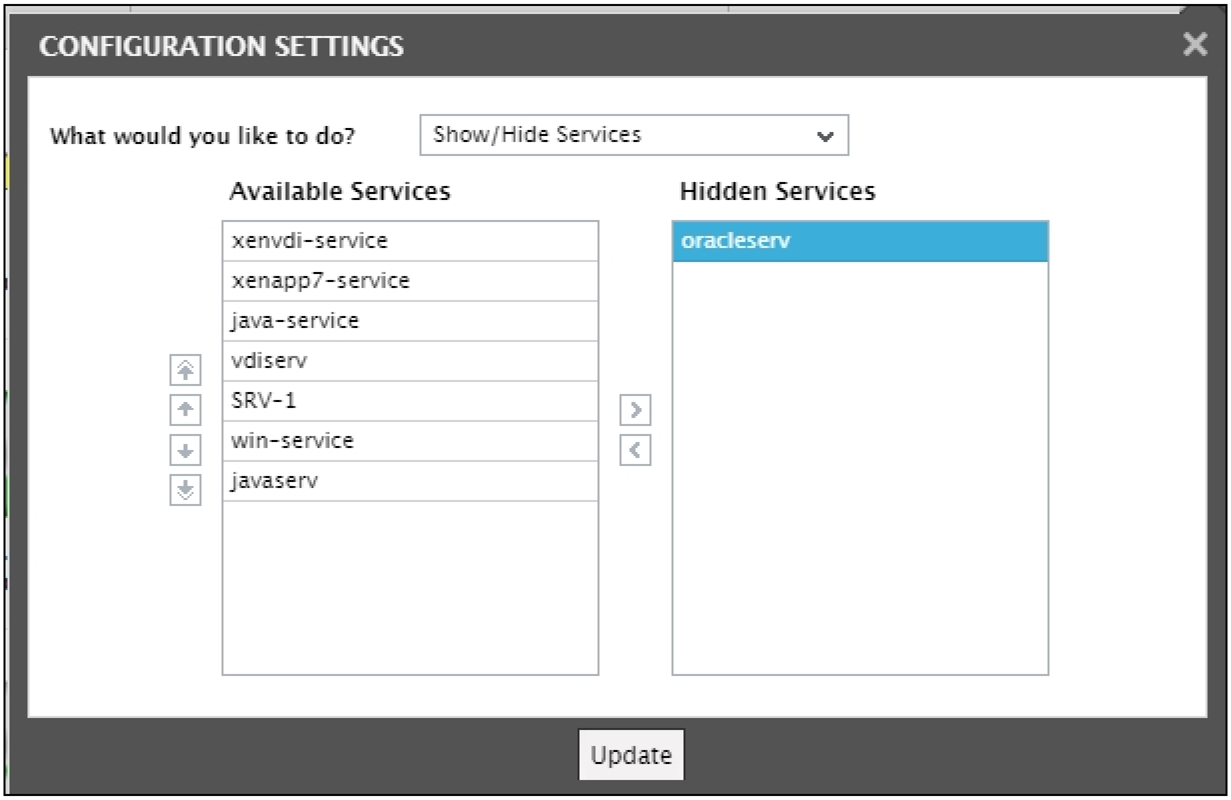

Figure 10.131: Figure 10.131: Hiding services from the dashboard

4. To hide a service, select it from the **Available Services** list in Figure 10.130 and then click the **>** button. This will transfer the selection to the **Hidden Services** list (see Figure 10.131). Likewise, you can add multiple services to the **Hidden Services** list.

5. You can also rearrange the order in which services are to be listed in the dashboard. This way, you can have your mission-critical services on top, and the less critical ones at the bottom of the dashboard. To achieve this, select a service from the **Available Services** list and click any of the following buttons, which are available alongside. The table below lists the buttons and indicates what will happen if a button is clicked.

Button	Purpose
	Click on this button to move the selection a level higher.
	Click on this button to move the selection a level lower
	Click on this button to move the selection to the top of the list
	Click on this button to move the selection to the bottom of the list

6. Finally, click the **Update** button to save the changes.

7. Moreover, super-users to the eG Enterprise system (i.e., users with **Admin** or **Supermonitor** rights) can share the changes they make to the business dashboard with other users, so that some/all users receive a

consistent view of the demand, consumption, and quality of the service. For this purpose, the super-user needs to follow the steps detailed below:

8. Once changes are made to one/more service displays in the business dashboard, click the icon at the right, top corner of Figure 10.120.

9. Figure 10.130 will then appear. From the **Applies to** drop-down, select a service. The changes made by the super-user to the business dashboard settings of this service will then be shared with other users. To share changes effected on all services, pick the **All Services** option.

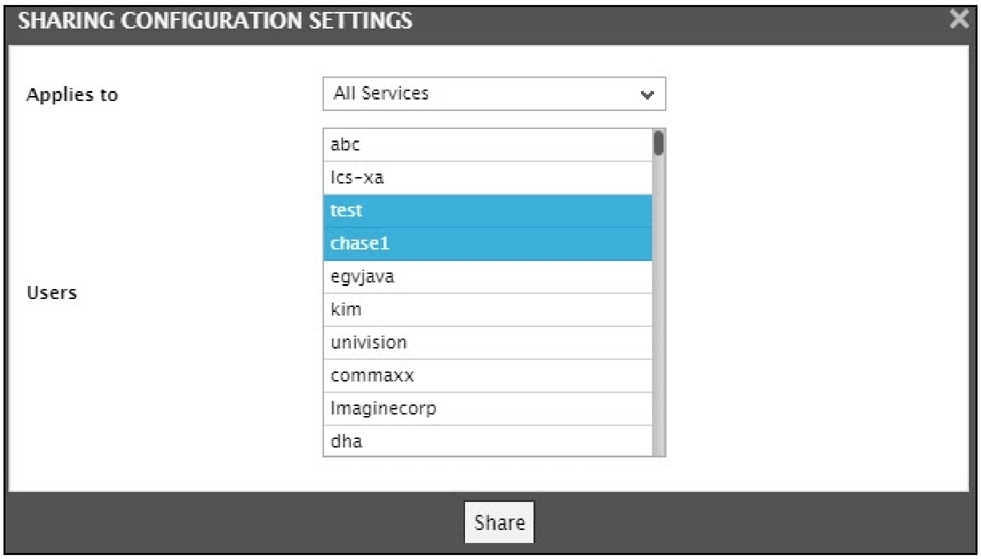

Figure 10.132: Selecting the users with whom the dashboard changes are to be shared

10. Then, from the **Users** list, select the users with whom the dashboard changes of the chosen service are to be shared.

11. Finally, click the **Share** button to share.

10.5 User Experience Dashboard

One of the biggest challenges that Citrix/virtual/terminal desktop administrators have is that they often have to spend time troubleshooting problems that may be caused in other parts of the infrastructure that they do not control. For instance, a slowdown in the home network that a user is connecting from can impact the user experience when accessing a Citrix service. eG Enterprise includes a **User Experience Dashboard** that makes it possible for end-users themselves to view the performance metrics related to their access to the Citrix/VDI/Terminal server infrastructure. This way, end users can easily determine when they see a slowdown, is the problem being caused by connectivity to the Citrix/VDI/Terminal server 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.

The self-service capability that the end-user performance dashboard provides results in fewer complaints and trouble calls to the Citrix/virtual desktop helpdesk. As a result, support costs are lower, end-users are less frustrated and the Citrix/virtual desktop deployment can proceed to successful completion.

End-users do not have to login to the eG monitoring console to access the dashboard. Administrators can publish the dashboard for public viewing . By accessing the URL, **<eG Manager IP:Port>/final/monitor/endUserDashboard.jsp? username=<name of the user>**, an end-user can get to see the performance of his / her Citrix or virtual desktop session or terminal server session.

Citrix/Virtual desktop/Terminal server administrators can also use the same dashboard to handle user complaints. When a specific user calls, they can view the performance dashboard for that user and determine what action needs to be taken to resolve the issue. This industry first end-user performance dashboard for Citrix/virtual desktop/terminal server infrastructures greatly simplifies the day to day operation of a Citrix/virtual desktop infrastructure. To access this dashboard, a Citrix/VDI/Terminal server administrator registered with the eG Enterprise system should first invoke the **Monitor** menu in the eG monitoring console, browse the **Dashboards** menu, move the mouse pointer over the **User Experience Dashboard** group, and pick the **VDI** or **XenApp** or **Terminal** option depending upon which user experience dashboard they want – the one for VDI users? or the one for XenApp users? or the one for the users logged into the terminal servers?(seeFigure 10.133)

Note:

The **VDI**, **XenApp** and **Terminal** options will be available only if the VDI, XenApp and Microsoft RDS components are managed in the monitored environment. If only one of the two is available – say, if only XenApp servers are managed in the environment - then, the **User Experience Dashboard** menu option will not display any sub-options. Instead, clicking on the **User Experience Dashboard** option will automatically lead administrators to the Overview dashboard of the XenApp environment.

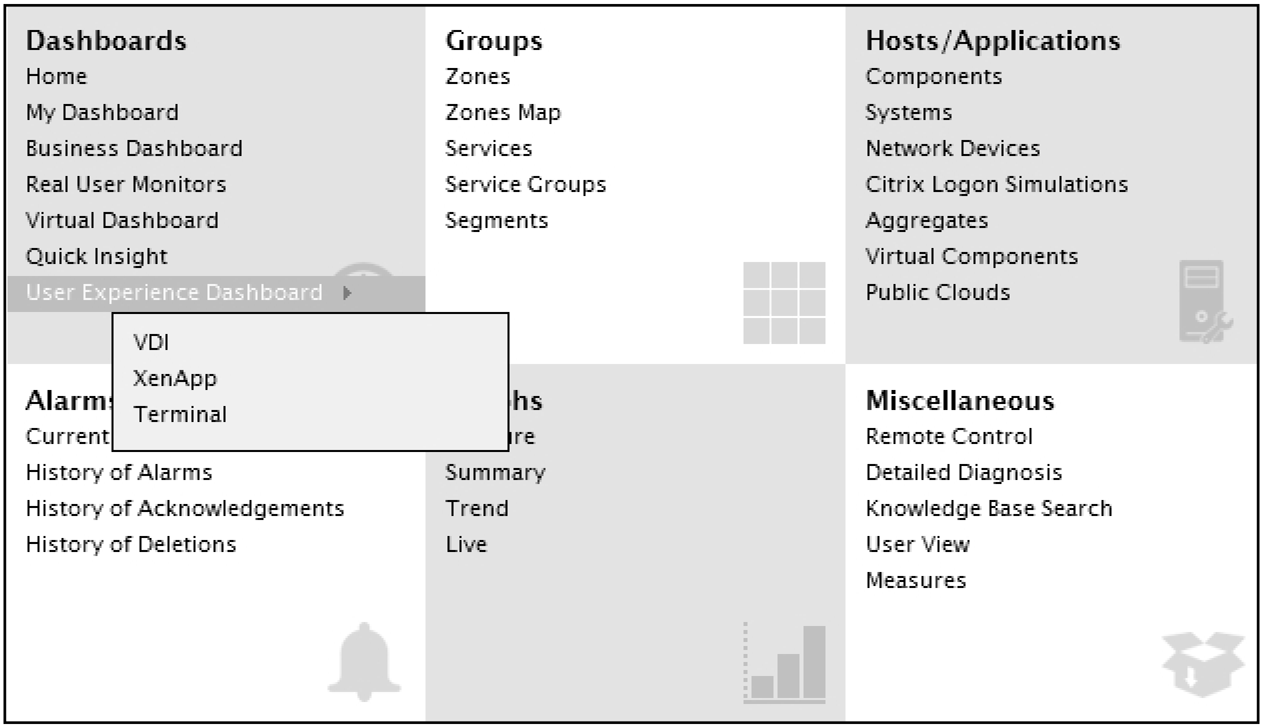

Figure 10.133: Accessing the User Experience dashboard

If the **XenApp** option is chosen, then the XenApp administrator will first be lead to an interface that reveals at-a-glance, which user is currently logged into which VM/desktop across the virtual infrastructure, the current state of each VM/desktop, and an overview of each user's experience with his/her VM/desktop. A quick look at this dashboard will point administrators to those users who are currently experiencing performance issues with their VMs/desktops and what the problem areas could be.

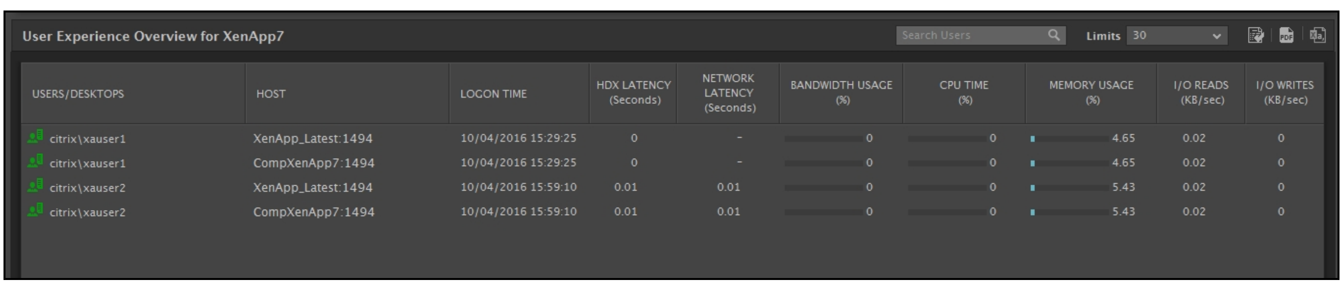

Figure 10.134: The User Experience Dashboard Overview that can be viewed by the administrators

Note:

The User Experience Dashboard Overview is available only for Citrix/VDI/Microsoft RDS administrators only and not for the end-users.

By default, this interface will list a maximum of 30 virtual desktops in the monitored environment, arranged in the order of their current state and then in the ascending order of their names. An administrator can however increase/decrease this default limit and can even add/delete the desktop health metrics that are displayed per desktop. To increase/decrease this default limit and to add/delete the desktop health metrics, administrators should navigate to the MONITOR SETTINGS page of the eG administrative interface. To know more about changing the default settings for the User Experience Dashboard, refer to Configuring Monitor Settings of the Administering the eG Enterprise Suite document. Administrators can override the default number of desktops displayed in the User Experience Overview Dashboard by selecting an appropriate number from the Limits list.

To search for a user in the User Experience Overview dashboard, administrators can use the Search Users text box. Administrators can print the User Experience Overview dashboard using the icon and export the User Experience Overview Dashboard in Microsoft Excel format using the icon.

For deeper analysis of user experience with a desktop, the administrator can click on any desktop in the Overview dashboard of Figure 10.134. This will invoke the **User Experience Dashboard** for the user who is currently logged into that desktop.


User Experience Dashboard for 'citrix\xauser1'
Search Users
Back
Session - Login Details
Currently Logged In
Yes
Last Login
10/05/2016 09:48:06
Last accessed
10/05/2016 17:00:39
Server
CompXenApp7:1..
Client IP
192.168.8.148
Client Version
14.1.200.13
Client Name
eglap0025-pc
Idle Time
9 mins
User sessions
1
Alerts for this user
Recent 10
Component
Description
StartTime
Duration
CompXenApp7:1494 Session start-up server duration is high 10/05/2016 14:11:06 4m 6s
XenApp_Latest:1494 Session start-up server duration is high 10/05/2016 14:11:06 5m 46s
XenApp_Latest:1494 Session start-up duration is high 10/05/2016 14:11:02 5m 50s
CompXenApp7:1494 Session start-up duration is high 10/05/2016 14:11:02 4m 10s
CompXenApp7:1494 Session start-up server duration is high 10/03/2016 10:46:37 5m 36s
CompXenApp7:1494 Session start-up duration is high 10/03/2016 10:46:32 5m 41s
XenApp_Latest:1494 Session start-up server duration is high 10/03/2016 10:44:52 6m 41s
XenApp_Latest:1494 Session start-up duration is high 10/03/2016 10:44:50 6m 43s
CompXenApp7:1494 Session start-up server duration is high 09/30/2016 10:38:56 7m 41s
User Experience - Latencies
HDX Screen Refresh Latency (Seconds)
0.003
Session Consumption - CPU/Memory Usage
Memory Usage (%)
6.0562
CPU Usage (%)
0.8
I/O reads/writes of user
IO Reads (KB/sec)
0
IO Writes (KB/sec)
0
User Session - Top CPU processes (%)
User Session - Top Memory processes (%)
Session - Profile Details
Profile Size (MB)
2.57
Number of files (Number)
94
Large files (Number)
0
User Session - Top I/O write processes (KBps)
User Session - Top I/O read processes (KBps)
ctxgfx : 0
Session Details
24 hours
Server
Last Accessed Time
XenApp_Latest:1494 10/05/2016 17:01:59
CompXenApp7:1494 10/05/2016 17:00:39
HDX Channel - Input Bandwidth
Audio (Kbps)
Printer (Kbps)
Total (Kbps)
HDX Channel - Output Bandwidth
Line Speed
Output (Kbps)
17235
Input (Kbps)
Memory Usage
Memory usage for user's processes (%)
6.06
test
Time since last activity (Minutes)
10
Virutal memory of user
Virtual memory of user's processes (MB)
915.2
Idle time
Total idle time in session (Minutes)
177.73
Oct 05, 2016 16:06:05
Total idle time in session (Minutes): 126.6000
Handles Usage
Handles used by user's processes (Number)
1708
CPU Percent - DD

Figure 10.135: The User Experience Dashboard

Note:

An end-user can just hit the URL, **<eG Manager IP:Port>/final/monitor/endUserDashboard.jsp?username=<name of the user>**, to instantly open this dashboard and review the impact of their operations on the health of their desktop.

By default, the **User Experience Dashboard** displays a wide variety of panels that reveal the following for a desktop user:

- ICA session details of the user, indicating whether/not the user is engaged in any bandwidth-intensive communication over the ICA channel, and if so, what type of communication it is;
- Recent alerts related to the user;
- How the user is consuming the network, CPU, memory, disk space, and disk I/O resources of the desktop;
- The processes run by the user on the desktop, which are the leading resource (CPU/memory/disk I/O) consumers
- When the user last accessed the desktop;
- TCP connections to and from the user's desktop

From these default panels, administrators will be able to quickly pinpoint the exact source of a user's poor experience with his/her virtual desktop. If required, administrators can hide those default panels that they deem unnecessary, replace them with new panels with newer, meaningful metrics, or can even modify the default panels slightly to include the performance information of interest to them. To perform this customization, administrators must click on the icon available in the right top corner of Figure 3. The sections to be followed will discuss in details on the following:

- Add a panel
- Modify an existing panel
- Delete custom panels
- Show/Hide default panels

10.5.1 Adding a widget

To add a new widget to the **User Experience** dashboard, the user must first click on the  icon. Figure 1 will then appear.

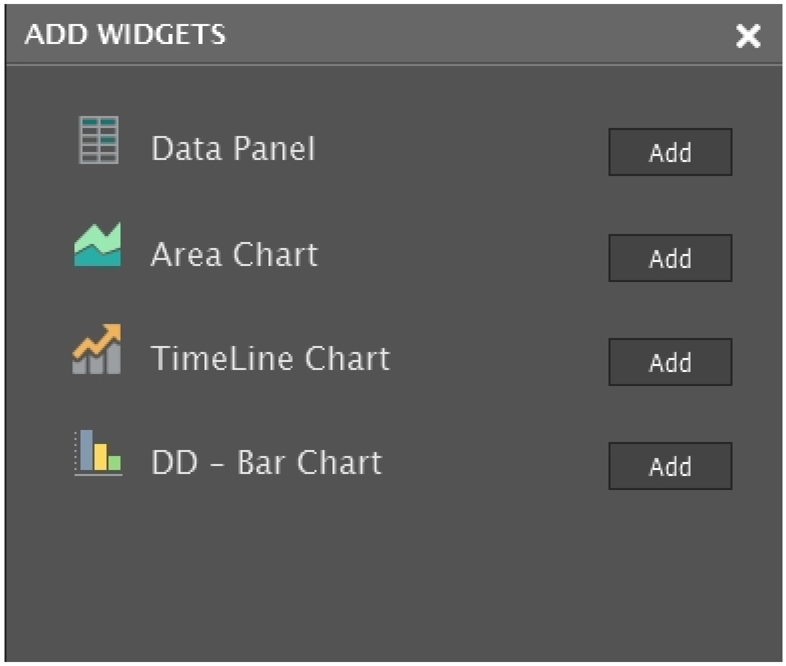

Figure 10.136: Adding a widget to the User Experience dashboard

The User Experience Dashboard supports four type of widgets, namely:

- The **Data** widget: When included in the dashboard, this will display the current state and value of a selected measure, with a link to detailed metrics (if available).
- The **Area Chart** widget: When included in the dashboard, this will display a historical area chart indicating the trends in the behavior of a chosen measure over a given timeline. This chart is easy to read and helps in rapidly detecting abnormal performance trends.
- The **TimeLine Chart** widget: This widget will display a line graph of a chosen measure over a given timeline. Using this graph, time-of-day variations in measure values can be observed closely and sporadic/consistent aberrations in performance can be quickly identified.
- The **DD – Bar Chart** widget: This widget will graphically represent the detailed diagnostics of a chosen measure. This way, administrators can quickly compare performance across various parameters and can accurately isolate the root-cause of performance degradations.

To add a **Data** widget, do the following:

1. Click the **Add** button in Figure 1. Figure 2 will then appear.

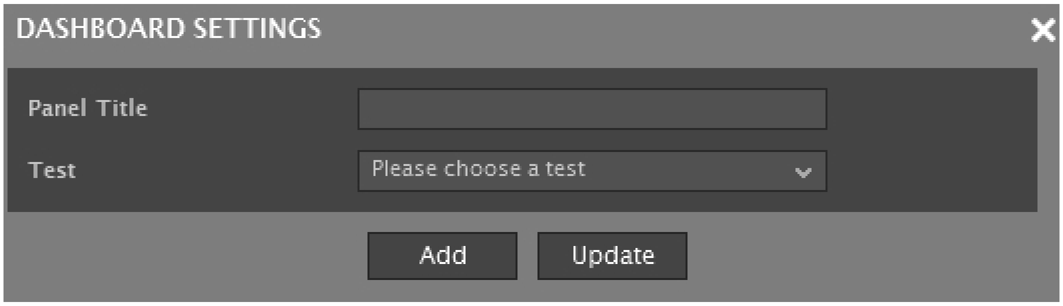

Figure 10.137: Th pop up window that appears when the Add button against the Data Widget is clicked

2. In Figure 2, specify the **Panel Title.**

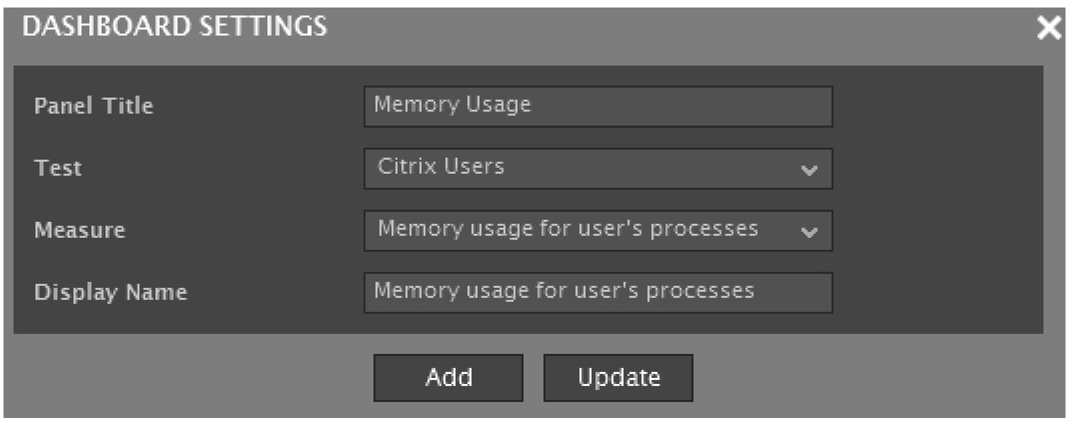

Figure 10.138: Adding a Data widget

3. Next, select the **Test** that reports the measure to be displayed in the Data widget.
4. Pick the **Measure.**
5. Provide a **Display Name** for the measure.
6. Click **Add** to add the measure to the Data widget. Similarly, multiple measures can be added to the same Data widget.
7. Once you are done adding, click the **Update** button to apply the configuration changes to the dashboard.
8. Figure 4 displays a Data widget that has been added to the dashboard.

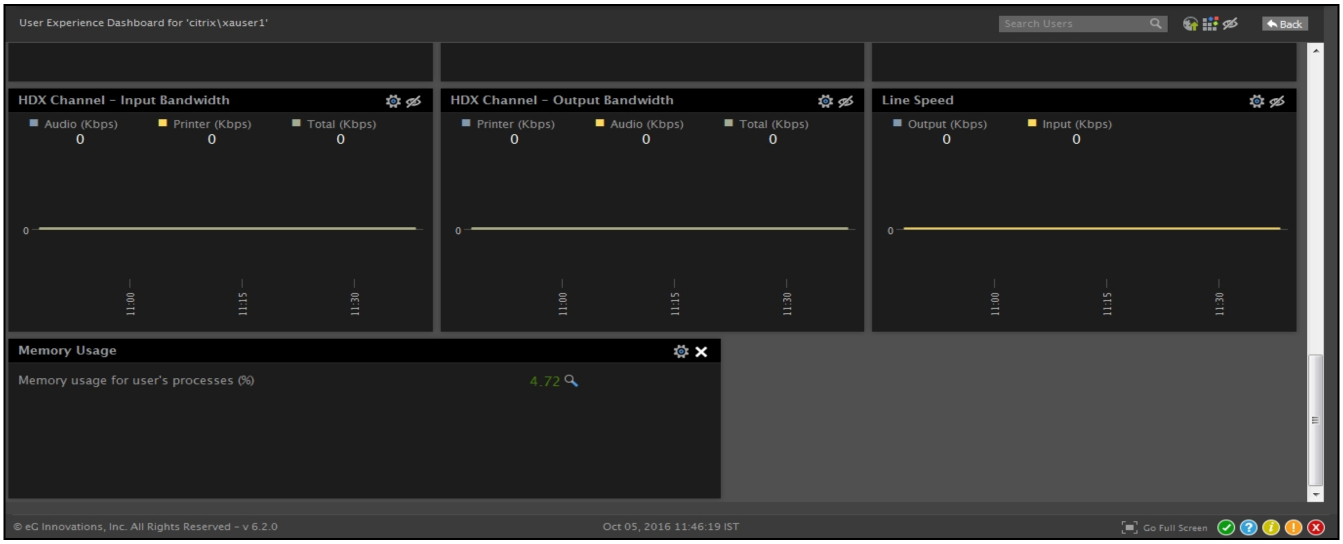

Figure 10.139: User Experience Dashboard with a Data widget

9. To add an **Area Chart**, do the following:
10. Click the Add button against the Area widget in Figure 1.
11. In Figure 5 that appears, specify the **Panel Title** (see Figure 5).

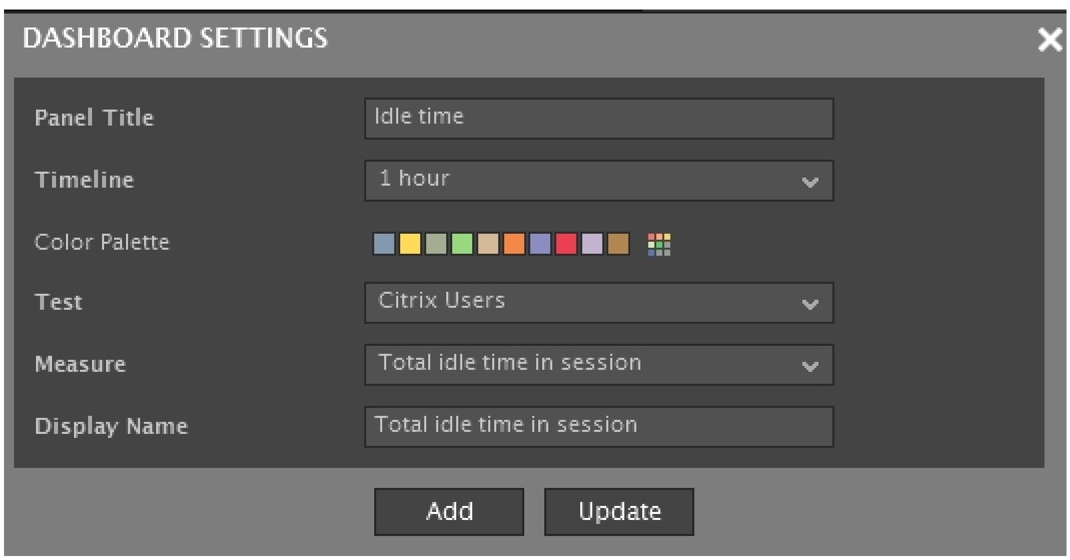

Figure 10.140: Adding an Area widget

12. Next, select a **Timeline** for the Area chart.
13. Then, pick a **Color Palette** for the Area chart by clicking the  icon.
14. Then, select the **Test** that reports the measure for which the Area chart is to be generated.
15. Pick the **Measure** for the chart.
16. Provide a **Display Name** for the measure.

17. Click **Add** to add the widget and **Update** to apply the configuration changes to the dashboard.
18. Figure 6 displays a Data widget that has been added to the dashboard.

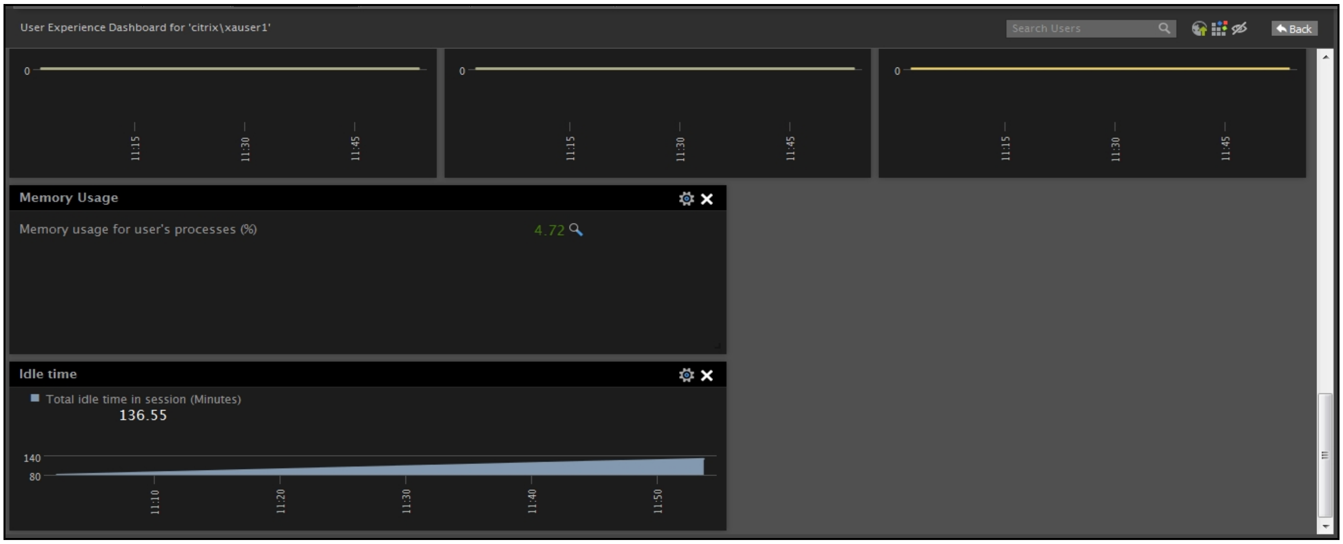

Figure 10.141: A User Experience Dashboard with an Area widget

19. To add a **TimeLine Chart**, do the following:
20. Click the Add button against the TimeLine Chart widget in Figure 1.
21. Then, specify the **Panel Title** (see Figure 7).

Figure 10.142: Adding a Line widget

22. Next, select a **Timeline** for the graph.
23. Then, pick a **Color Palette** for the Area chart by clicking the button.
24. Then, select the **Test** that reports the measure for which the Line graph is to be generated.

25. Pick the **Measure** to be plotted in the graph.
26. Provide a **Display Name** for the measure.
27. Click **Add** to add the widget and **Update** to apply the configuration changes to the dashboard.
28. Figure 8 displays a TimeLine Chart widget that has been added to the dashboard.

Figure 10.143: A User Experience Dashboard with a Line widget

To add a **DD – Bar Chart** widget, do the following:

1. Click the Add button against the **DD – Bar Chart** widget in Figure 1. Figure 9 will then appear.
2. Then, specify the **Panel Title** (see Figure 9).

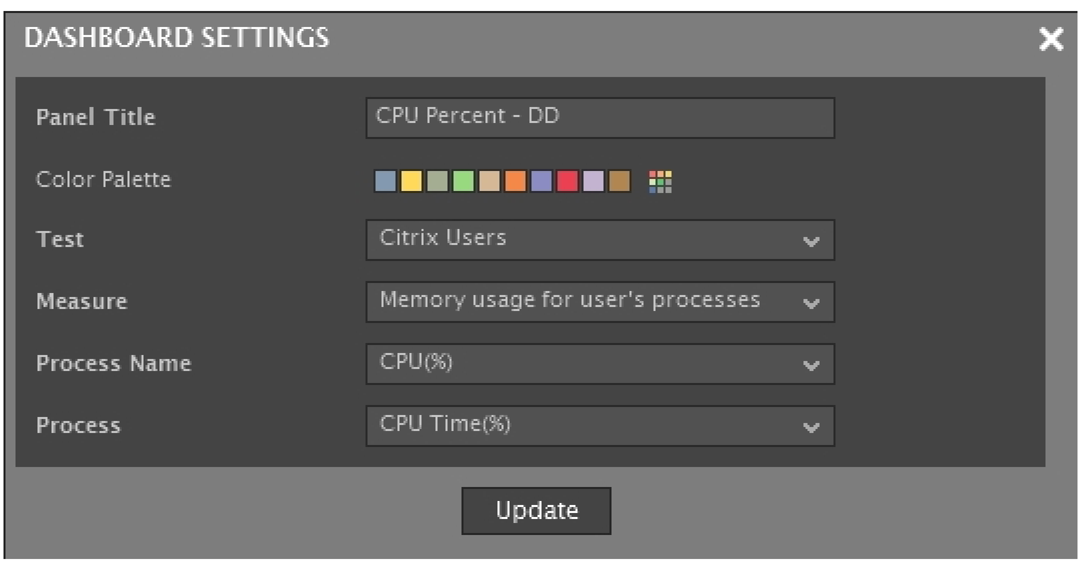

Figure 10.144: Adding a DD – Bar Chart widget

3. Pick a **Color Palette** for the bar chart by clicking the icon.

4. Then, select the **Test** that reports the detailed metrics for which the bar chart is to be generated.
5. Pick the **Measure**.
6. From **Process Name** and **Process** lists, select the DD columns that report the values to be used as the X and Y axis of the bar graph.
7. Click **Update** to apply the configuration changes to the dashboard.
8. Figure 10 displays a Line widget that has been added to the dashboard.

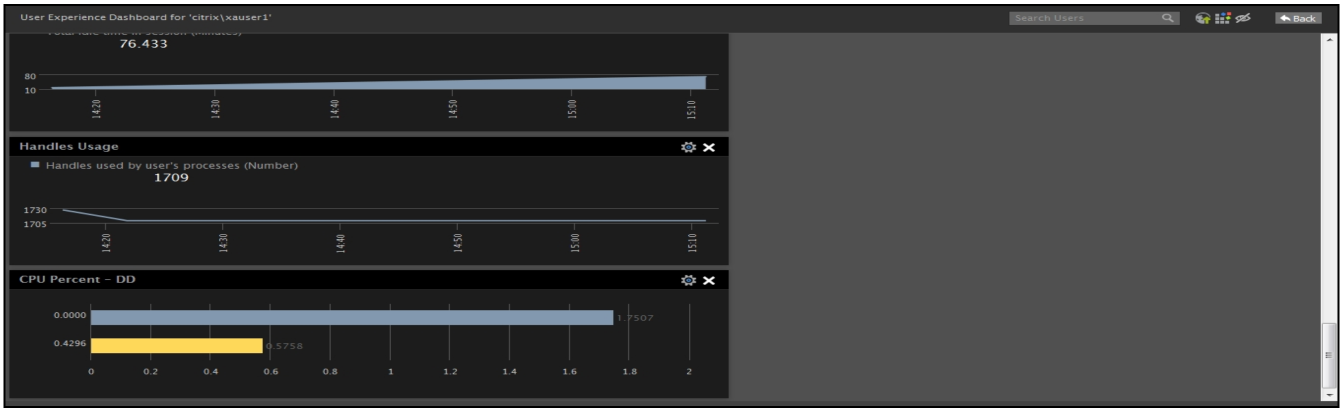

Figure 10.145: A User Experience Dashboard with a DD – Bar Chart widget

Note:

- The position of any widget (default/custom) can be changed in the dashboard using a simple click-and-drag routine.

10.5.2 Modifying an Existing Panel

To modify an existing panel in the User Experience Dashboard, click the icon available in the top right corner of the panel. Any default or custom panel can be picked for editing. However, **the procedure for making changes to a panel will differ from one panel to another**.

For instance, if you choose the **Consumption-Disk Usage** panel for modification, Figure 2 will appear.

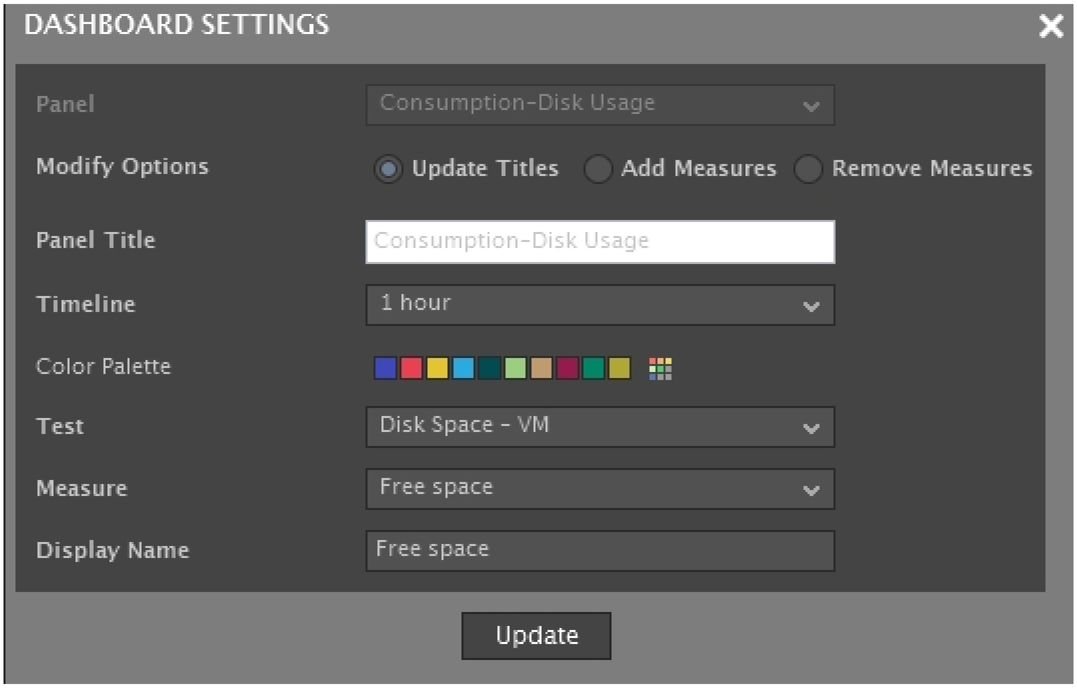

Figure 10.146: Modifying a panel in the User Experience Dashboard.

From Figure 1, it can be inferred that you can modify the Panel Title, Add measures to the widget and remove measures from the widget. You can also change the timeline of the graph. and change the color palette of the graph, if applicable.

Now, take the case of the **Consumption – Disk Usage** panel (See Figure 1). By default, the Modify Options flag will be set to Update Titles. You can change the tite of the panel using the panel Title text box and click the Update button. If this panel is chosen for modification, the **Panel Title** can be changed (see Figure 1) by selecting the **Update Titles** option. Alongside, you can also select a **Test** and **Measure** that have already been configured for display in the **Consumption – Disk Usage** panel, and change the **Display Name** of that measure (see Figure 1).

In the same way, to add new measures to an existing panel, you can do so by setting the Modify Options flag to Add Measures.

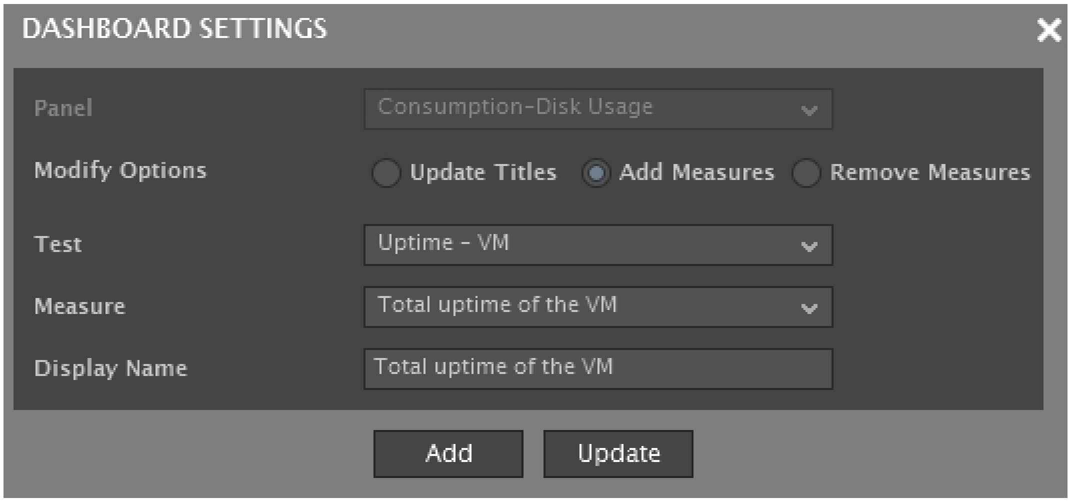

Figure 10.147: Adding new measures to the Consumption – Network panel

Similarly, to removea measure from the Consumption - Disk Usage panel, set the Modify Option flag to Remove Measures.

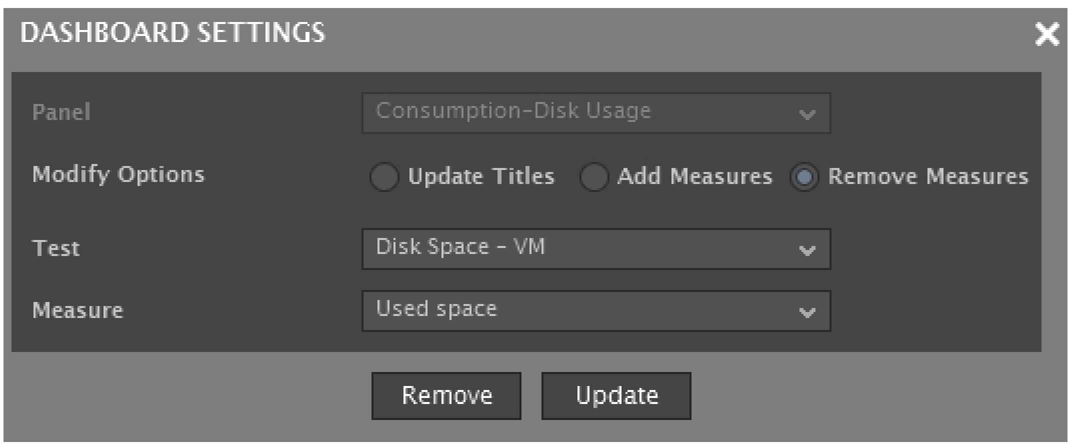

Figure 10.148: Deleting an existing measure from the Consumption – Network panel

The table lists all the default panels of the User Experience Dashboard and indicates what can be changed in each panel.

Default Panel	What can be changed
ICA Session Data	• **Panel Title** can be updated. • **Measure Display Name** can be changed. • New measures can be added. • Existing measures can be deleted.
Top CPU processes of user	• **Panel Title** can be updated. • The **Color palette** can be changed.

Default Panel	What can be changed
Alerts for this user	• **Panel Title** can be updated. • The **No of records** can be changed – i.e., the number of recent alerts to be displayed for a user can be set.
Top memory processes of user	• **Panel Title** can be updated. • The **Color palette** can be changed.
Top processes by disk busy	• **Panel Title** can be updated. • The **Color palette** can be changed.
Consumption – CPU/Memory Usage	• **Panel Title** can be updated. • **Measure Display Name** can be changed. • The **Color palette** can be changed. • The **Timeline** can be altered. • New measures can be added. • Existing measures can be deleted.
Consumption – Disk I/O	• **Panel Title** can be updated. • **Measure Display Name** can be changed. • New measures can be added. • Existing measures can be deleted.
Consumption – Disk Usage	• **Panel Title** can be updated. • **Measure Display Name** can be changed. • The **Color palette** can be changed. • The **Timeline** can be altered. • New measures can be added. • Existing measures can be deleted.
Demand – TCP Connections	• **Panel Title** can be updated. • **Measure Display Name** can be changed. • The **Color palette** can be changed. • The **Timeline** can be altered. • New measures can be added. • Existing measures can be deleted.

After making changes to a panel, remember to click the **Update** button to save the changes.

10.5.3 Deleting a Custom Panel

To delete a custom panel, simply click on the **X** button that appears on the top right corner of the panel (see Figure 1).

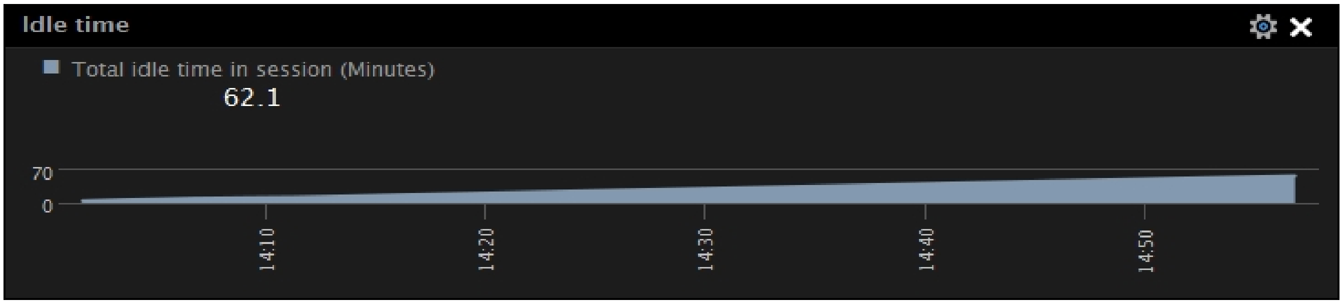

Figure 10.149:

Note:

- A default panel in the User Experience Dashboard cannot be permanently deleted. You can only hide a default panel.
- Multiple panels cannot be deleted at a single shot.

10.5.4 Show/Hide Default Panels

You can hide two/more default panels and unhide them later using this procedure.

Click on the icon at the right top corner of the User Experience Dashboard. Figure 10.150 will then appear.

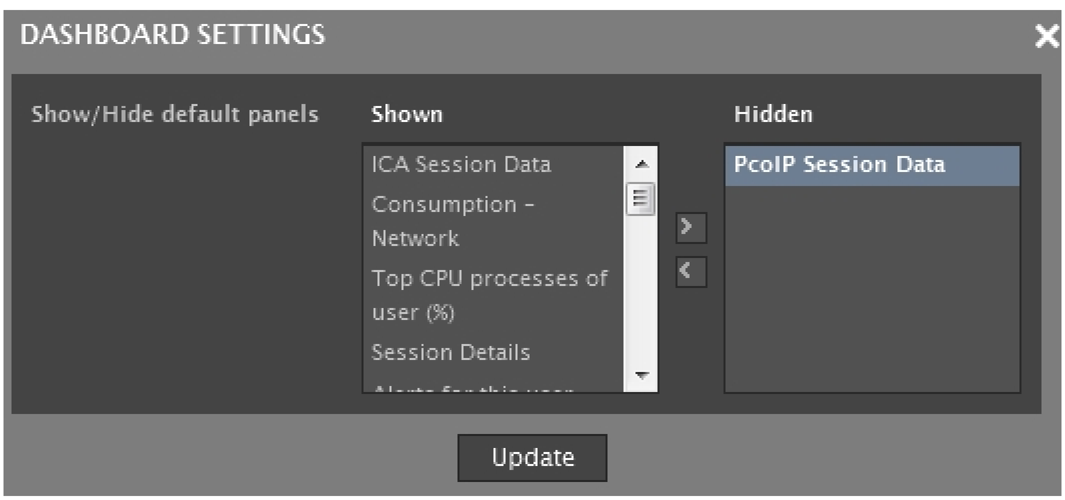

Figure 10.150: Hiding/unhiding the default widgets

All default panels that are currently visible in the dashboard will be listed in the **Shown** list box. The default panels that are currently hidden will be displayed in the **Hidden** list box. To hide some visible panels, select them from the **Shown** list and click the **>** button in Figure 10.150, so that the selection moves to the **Hidden** list.

To unhide a few hidden panels, select them from the **Hidden** list and click the **<** button, so that they move to the **Shown** list.

Finally, click the **Update** button.

Note:

- A default panel cannot be permanently deleted. You can only hide a default panel.
- The procedure discussed in this section is ideal for hiding multiple default panels at one shot. To hide a single default panel, just click the button at the right top corner of that panel.

10.6 Virtual Dashboard

Large virtualized environments are typically characterized by hundreds of virtual hosts, with tens of VMs configured on each host. In such environments, if a mission-critical application executing on a VM experiences a slowdown, the key challenge for the administrator is to determine what caused the slowdown - is it because the VM is under-sized? is it owing to a resource contention at the virtual host level? Is it because of other resource-hungry VMs? or is it owing to resource-intensive processes spawned by the target application itself? Also, as VM migration would be rampant in such environments, administrators are required to know when, why, and to which virtual host a VM was migrated, and even assess the impact of such a shift on the source and the destination hosts.

eG Enterprise provides dedicated monitoring models for a variety of virtualization platforms to enable administrators to address the performance concerns cited above. By analyzing the resource usage of VMs from inside and outside, and using an intelligent correlation engine to differentiate problem symptoms from sources, these models accurately point to the root-cause of slowdowns experienced by a virtual application. In addition, these models are also capable of tracking the movement of VMs from one host to another and alerting administrators to the same. However, the key limitation of these models is that the aforesaid capabilities are 'hidden' inside the layers, tests, and measures offered by the models. To bring these capabilities to light, administrators would have to invest a considerable amount of time and mouse clicks!

A quicker, 'single glance' alternative to this is the **Virtual Dashboard**. The Virtual Dashboard provides administrators of virtualized environments the wherewithal to accurately diagnose the cause for slowdowns experienced by a virtual application, with minimal effort and time! This dashboard collates critical resource usage data from the host operating system and from 'inside' and 'outside' all the VMs on a virtual server, presents them in graphical and tabulated formats in a single interface in the eG monitoring console, and thus aids instant and effective performance analysis. In the event of a problem situation therefore, administrators can use the **Virtual Dashboard** to rapidly find answers to the following questions:

- Is the host experiencing a resource crunch? If so, which resource is being drained - CPU/memory/disk? Which processes executing on the host could be causing the resource bottleneck?
- How are the VMs using the physical resources of the host? Is any VM consuming resources excessively? If so, which VM is it?
- Were sufficient resources allocated to the VMs? How are the VMs using the resources allocated to them?
- Are resource-intensive processes executing on any VM? If so, which VM is being impacted, and what are the rogue processes?

- Did any VM migrate during the last hour (by default)? If so, which VM is it? Which server was it migrated to? How is the VM handling the physical and allocated resources in the destination host?

In addition to current problems, the dashboard also sheds light on the probable causes for issues that occurred in the past, thus paving the way for effective post-mortem analysis.

To view the virtual dashboard, click on the **Virtual** tab page that is available for a target virtual server. Doing so will invoke Figure 10.151.

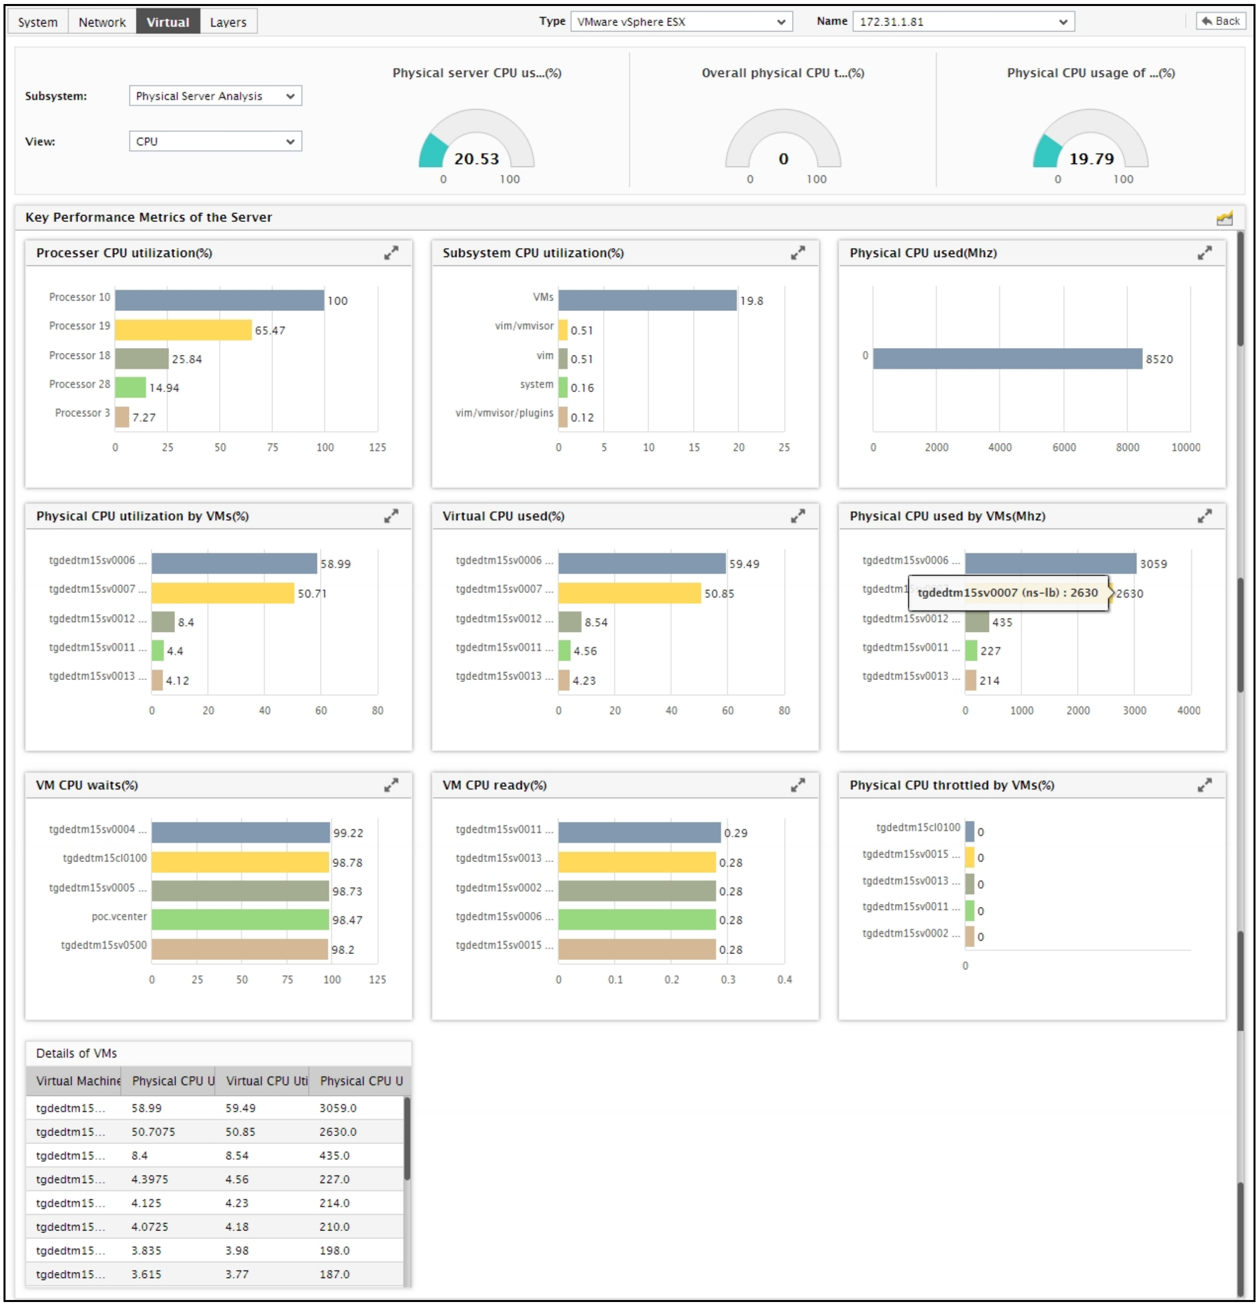

Figure 10.151: A Virtual Dashboard

The **Virtual Dashboard** groups the data collected into subsystems, and presents the data in the interface based on your choice of subsystem. The subsystems supported by this interface are as follows:

- Physical Server Analysis
- Outside view of VMs
- Inside view of VMs

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

10.6.1 Physical Server Analysis - CPU

By default, the **Virtual Dashboard** allows you to focus on how the host operating system and the VMs are utilizing the physical CPU resources of the target host. Accordingly, **Physical Server Analysis** is chosen as the **Subsystem**, and **CPU** is chosen from the **View** list, by default (see Figure 10.152). The contents of this default **CPU** dashboard have been discussed below:

1. The **CPU** dashboard begins with dial charts that will frequently update you with the current state and values for critical CPU usage metrics related to the host operating system. You can promptly detect sudden spikes in CPU usage by the host using these dial charts.

   Note:

   You can configure what CPU usage statistics need to be represented using dial charts by following the steps discussed below:

   - Edit the **eg_dashboard.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
   - In the **[CPUAnalysis]** section of the file, you will find a parameter of the following format for each of the virtualization platforms that is being monitored by eG Enterprise:

     <InternalComponentType>_GAUGE

     This parameter is typically set to a comma-separated list of measures for which dial charts are to be displayed in the dashboard.

     By default, this parameter is set to a pre-configured list of CPU-related metrics extracted from the corresponding **<InternalComponentType>**.

   - You can override this default setting by adding more measures to the comma-separated list, or by removing one/more existing measures. To add another measure, append an entry of the following format to the comma-separated list that follows the **<InternalComponentType>_GAUGE** parameter:

     <InternalTest>:<InternalMeasure>:<DisplayName>

     For instance, to add a dial chart for the **Physical CPU Usage** measure reported by the **CPU - ESX** test of the **VMware vSphere ESX** component, your entry should be: **VmEsx_i_server_GAUGE=. . . . . . .,EsxCpuSummaryTest:Usage:Physical server CPU usage**
     To know the internal component type, test, and measure names, refer to page of this document.

   - Finally, save the file.

Note that dial charts can be used to represent only those measures that report values in percentage.

2. Clicking on a dial chart will lead you to the **Layer** tab page, where you can view the layer and test that reported the measure represented by that dial chart (see Figure 10.153).

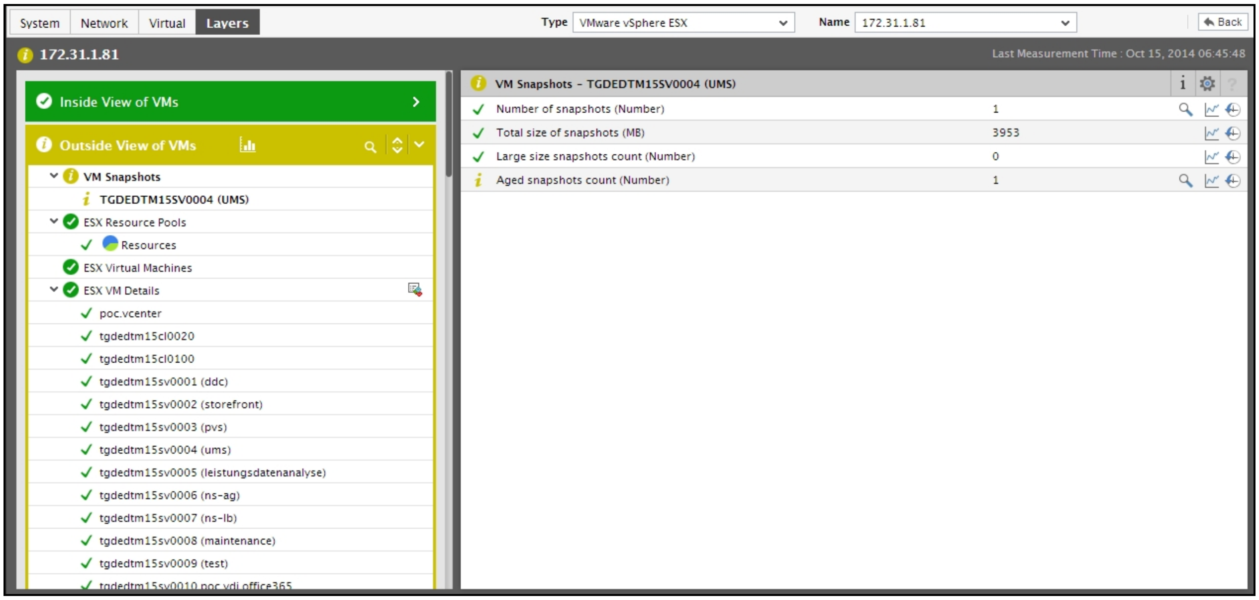

Figure 10.152: Figure 10.155: Clicking on a dial chart in the Physical Server CPU Analysis dashboard

3. Below the dial charts, you will find a default collection of comparison bar charts that typically compare CPU usage across the host's processors, the host's subsystems, the physical CPU usage across VMs, and the virtual CPU usage across VMs. Using these graphs, resource-intensive processors, subsystems, and VMs can be accurately identified.

   Note:

   You can configure what CPU usage statistics need to be represented using the comparison bar charts, by following the steps discussed below:

   - Edit the **eg_dashboard.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
   - In the **[CPUAnalysis]** section of the file, you will find a parameter of the following format for each of the virtualization platforms that is being monitored by eG Enterprise:

     <InternalComponentType>

     This parameter is typically set to a comma-separated list of measures for which comparison bar charts are to be displayed in the dashboard.

     By default, this parameter is set to a pre-configured list of CPU-related metrics extracted from the corresponding **<InternalComponentType>**.
   - You can override this default setting by adding more measures to the comma-separated list, or by removing one/more existing measures. To add another measure, append an entry of the following format to the comma-separated list that follows the **<InternalComponentType>** parameter:

<InternalTest>:<InternalMeasure>:<DisplayName>

For instance, to add a comparison graph for the **VM CPU Extra** measure reported by the **EsxGuestDetails** test of the **VMware vSphere ESX** component, your entry should be: **VmEsx_i_server =. . . . . . .,EsxGuestTest:Cpu_extra:VM CPU Extra**

To know the internal component type, test, and measure names, refer to page of this document.

- Finally, save the file.

It is recommended that you configure comparison graphs for only those measures that support descriptors.

4. If a comparison graph appears too cluttered, you can view that graph more clearly by enlarging it. To do so, click on that graph. The graph will then appear as depicted below:

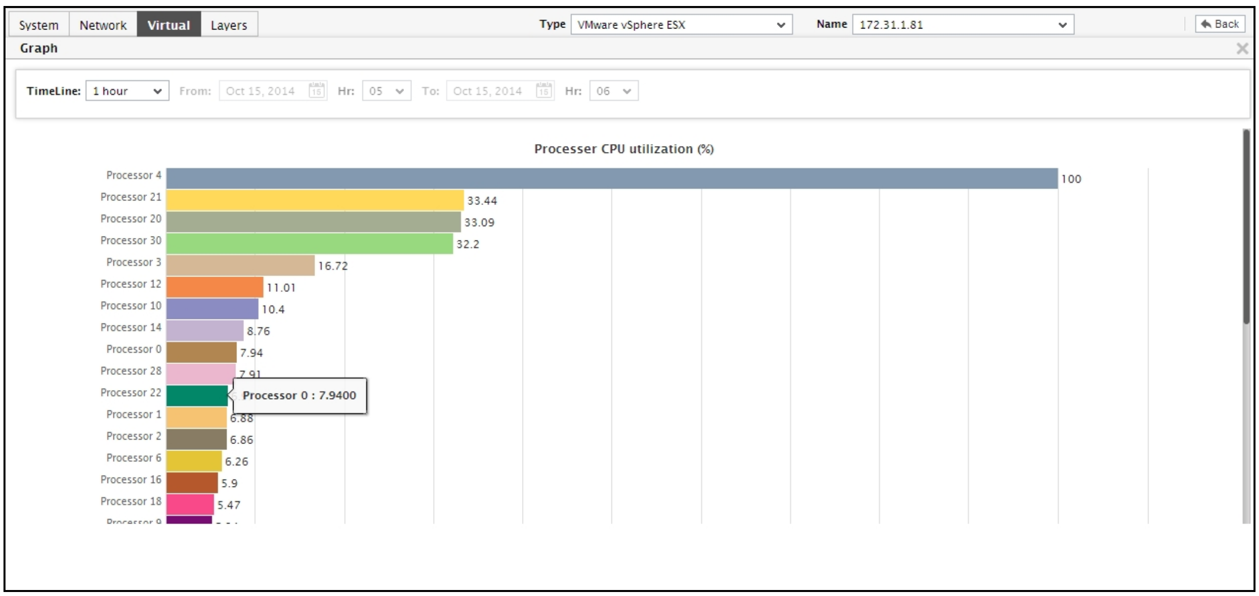

Figure 10.153: Figure 10.156: Enlarging a comparison bar chart in the Physical Server CPU Analysis dashboard

5. By default, in the enlarged mode, the comparison bar chart will display all the compared elements (i.e., descriptors). This is why, the **Show** list is set to **ALL** by default. If you want the graph to display only a few best or a few worst players in a particular area, then, pick a **TOP-N** or **LAST-N** option from the **Show** list. To collapse the enlarged graph, click on it again.

6. Back in the dashboard, you will find that a **Details of VMs** table follows the comparison bar charts. By default, this table compares the physical and virtual CPU usage of each of the VMs so that, you can clearly identify the following:

   - The VM causing a severe dent in the physical resources of the host;
   - The VM using the allocated CPU resources excessively

   By default, this table is sorted in the descending order of the **Physical CPU utilization** column. To sort the table in the ascending order of the same column, click on the down arrow button that appears adjacent to the column heading, **Physical CPU utilization (%)**. To sort the table on the basis of the values of another column, click on the title of the corresponding column.

Note:

You can configure additional measure columns to the **Details of VMs** table by following the procedure discussed below:

- Edit the **eg_dashboard.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
- In the **[CPUAnalysis]** section of the file, you will find a parameter of the following format for each of the virtualization platforms that is being monitored by eG Enterprise:

  <InternalComponentType>_DATA

  This parameter is typically set to a comma-separated list of measures that will appear as columns in the **Details of VMs** table.

  By default, this parameter is set to a pre-configured list of CPU usage statistics extracted from each VM on the corresponding **<InternalComponentType>**.
- You can override this default setting by adding more measures to the comma-separated list, or by removing one/more existing measures. To add another measure, append an entry of the following format to the comma-separated list that follows the **<InternalComponentType>_DATA** parameter:

  <InternalTest>:<InternalMeasure>

  For instance, to add a column for the **VM CPU Extra** measure reported by the **EsxGuestDetails** test of the **VMware vSphere ESX** component, your entry should be: **VmEsx_i_server_DATA =. . . . . . .,EsxGuestTest:Cpu_extra**

  To know the internal component type, test, and measure names, refer to page of this document.
- Finally, save the file.

7. To historically analyze the physical CPU usage of the host and the VMs, you can view historical measure graphs in your dashboard, instead of the default comparison bar charts. To do this, click on the icon at the right, top corner of your dashboard. This will result in the display of measure graphs, plotted for a default duration of 1 hour, for each of the measures for which comparison bar charts were originally displayed.

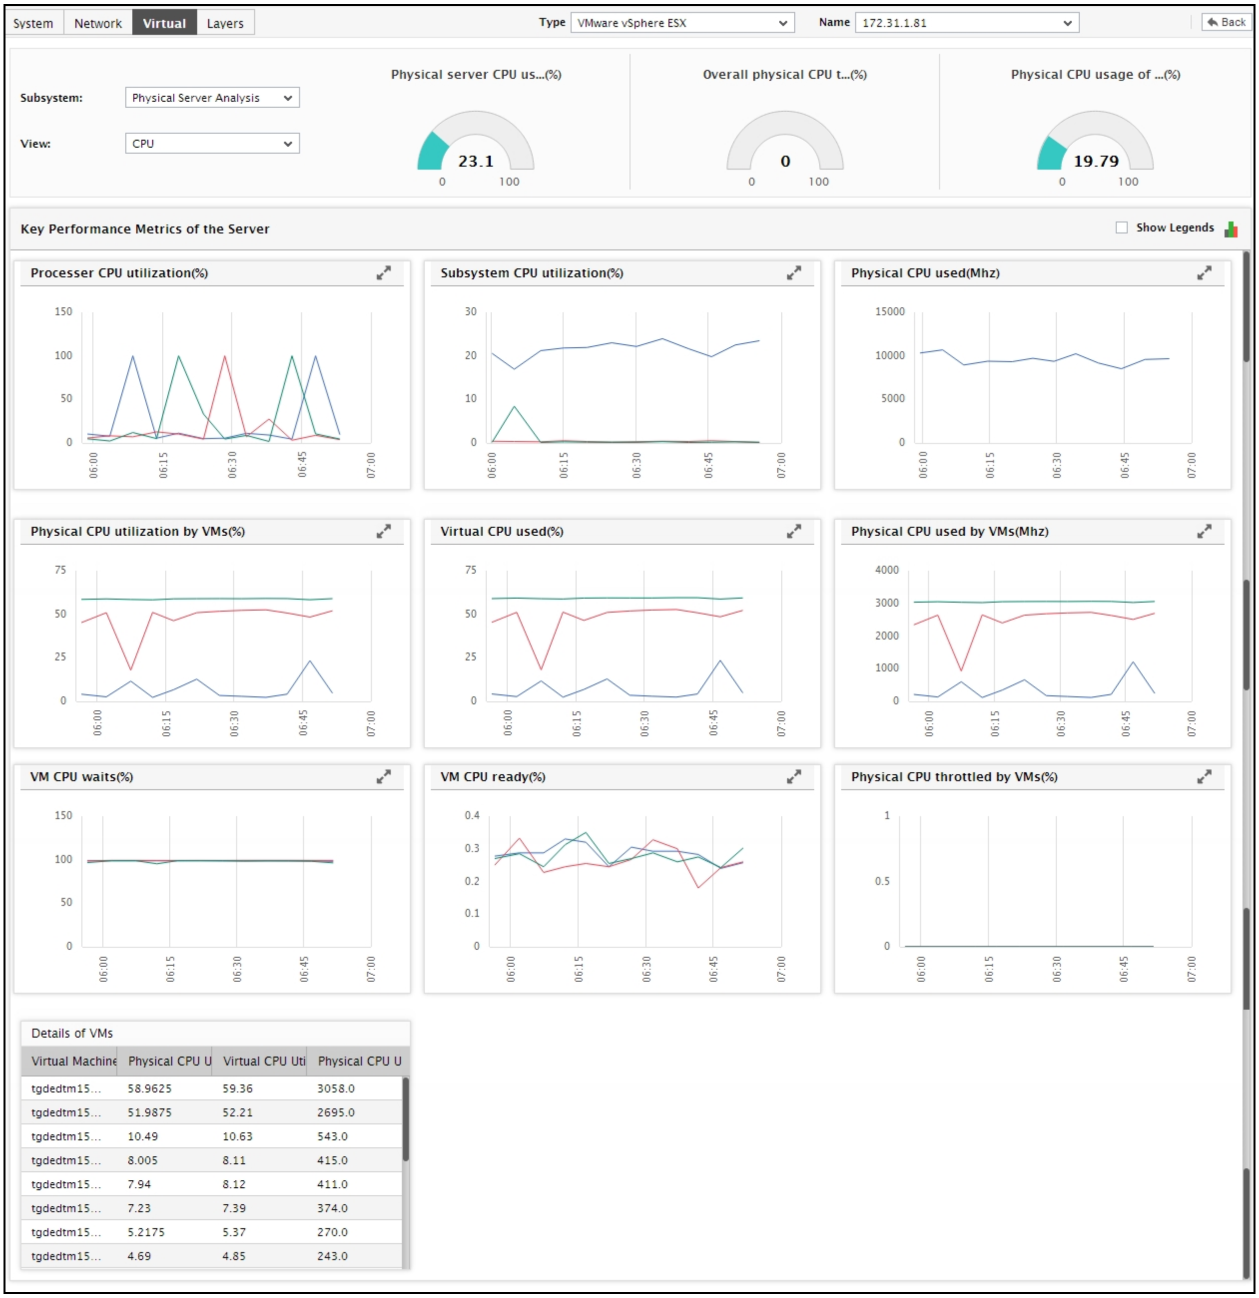

Figure 10.154: Figure 10.157: Viewing historical measure graphs in the Physical Server CPU Analysis dashboard

8. In the event of a contention for CPU resources, you can use these measure graphs to understand how CPU usage has varied during the last hour, and when the CPU contention actually began.
9. To view a measure graph more clearly, click on it. This will enlarge the graph (see Figure 10.155).

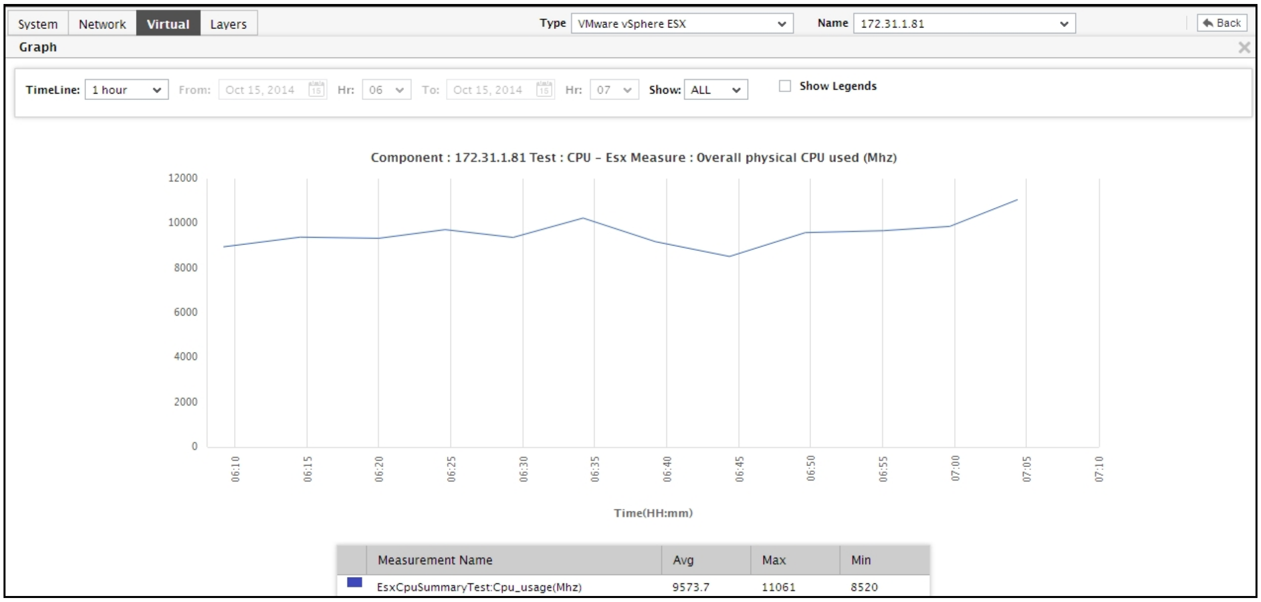

Figure 10.155: Figure 10.158: Enlarging a measure graph in the Physical Server CPU Analysis dashboard

10. In the enlarged mode, you can change the **Timeline** of the graph. You can also pick a **TOP-N** or **LAST-N**option from the **Show** list to analyze the time-of-day variations in the performance of a few best/worst players in the chosen performance arena. To return to the dashboard, click on the enlarged graph.

10.6.2 Physical Server Analysis - Memory

To quickly analyze physical memory usage by the virtual host and its VMs, and to precisely point to the VMs that are eroding the memory resources of the host, select **Memory** from the **View** list. Figure 10.156 will then appear.

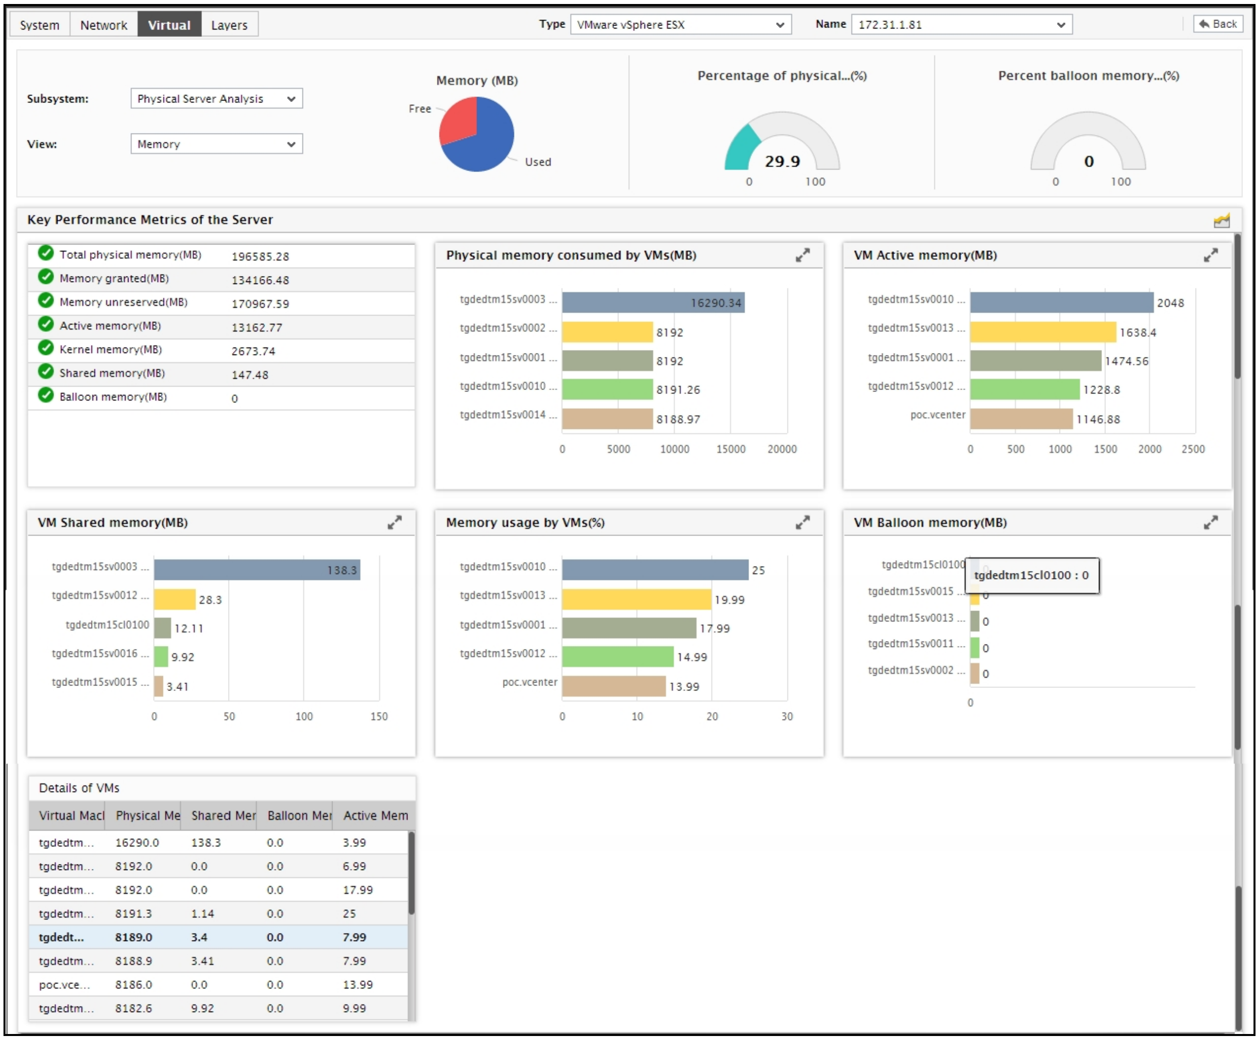

Figure 10.156: The Physical Server Memory Analysis dashboard

The contents of Figure 10.156 have been detailed below:

1. A glance at the pie chart in the **Memory** dashboard will indicate to you whether a memory crunch exists on the host or not.

   Note:

   By default, the pie chart will denote the size of **Free memory** and **Used Memory** on the target virtual host. If you want the pie chart to represent more measures, do the following:

   - Edit the **eg_dashboard.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
   - In the **[MemoryAnalysis]** section of the file, you will find a parameter of the following format for each of the virtualization platforms that is being monitored by eG Enterprise:

     <InternalComponentType>_PIE

This parameter is typically set to a comma-separated list of measures that will be represented in the pie chart in the **Memory** dashboard.

By default, this parameter is set to the **Free memory** and **Used memory** measures.

- You can override this default setting by adding more measures to the comma-separated list, or by removing one/more existing measures. To add another measure, append an entry of the following format to the comma-separated list that follows the **<InternalComponentType>_PIE** parameter:

  <InternalTest>:<InternalMeasure>:<DisplayName>

  For instance, to make sure that the pie chart also represents the value of the **Memory granted** measure reported by the **Memory - ESX** test of the **VMware vSphere ESX** component, your entry should be: **VmEsx_i_server_PIE=. . . . . . .,EsxMemoryTest:Memory_granted:Memory granted**

  To know the internal component type, test, and measure names, refer to page of this document.

- Finally, save the file.

  Note that dial charts can be used to represent only those measures that report values in percentage.

2. The dial charts in the dashboard indicate the current state and values for critical memory usage metrics related to the host operating system. You can promptly detect sudden spikes in host's memory usage using these dial charts.
3. Clicking on a dial chart will lead you to the **Layer** tab page, where you can view the layer and test that reported the measure represented by that dial chart.
4. Below the dial charts, you will find a table that indicates the current state and the current values for a default collection of critical memory usage-related metrics. This table allows you to focus on important usage metrics so that, you can instantly detect any change in the state of any such a metric, and immediately initial remedial steps.

   Note:

   If need be, you can alter the list of metrics that are displayed in this table by following the steps below:

   - Edit the **eg_dashboard.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
   - In the **[MemoryAnalysis]** section of the file, you will find a parameter of the following format for each of the virtualization platforms that is being monitored by eG Enterprise:

     <InternalComponentType>_Memory_DATA

     This parameter is typically set to a comma-separated list of measures that need to be included in the table of metrics in the **Memory** dashboard.

     By default, this parameter is set to a pre-configured list of memory usage-related metrics extracted from the corresponding **<InternalComponentType>**.

   - You can override this default setting by adding more measures to the comma-separated list, or by removing one/more existing measures. To add another measure, append an entry of the following format to the comma-separated list that follows the **<InternalComponentType>_ Memory_ DATA** parameter:

<InternalTest>:<InternalMeasure>

For instance, to add the **Free physical memory** measure reported by the **Memory - ESX** test of the **VMware vSphere ESX** component, your entry should be: **VmEsx_ i_ server_ Memory_ DATA =. . . . . . .,EsxMemoryTest:Machine_free_memory**

To know the internal component type, test, and measure names, refer to page of this document.

- Finally, save the file.

5. Next to the table, you will find a default collection of comparison bar charts that typically compare physical memory usage across VMs. Using these graphs, memory-intensive VMs can be isolated.

   Note:

   You can configure what memory usage statistics need to be represented using the comparison bar charts, by following the steps discussed below:

   - Edit the **eg_dashboard.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
   - In the **[MemoryAnalysis]** section of the file, you will find a parameter of the following format for each of the virtualization platforms that is being monitored by eG Enterprise:

     <InternalComponentType>

     This parameter is typically set to a comma-separated list of measures for which comparison bar charts are to be displayed in the dashboard.

     By default, this parameter is set to a pre-configured list of memory usage-related metrics extracted from the corresponding **<InternalComponentType>.**

   - You can override this default setting by adding more measures to the comma-separated list, or by removing one/more existing measures. To add another measure, append an entry of the following format to the comma-separated list that follows the **<InternalComponentType>** parameter:

     <InternalTest>:<InternalMeasure>:<DisplayName>

     For instance, to add a comparison graph for the **Current swap memory** measure reported by the **EsxGuestDetails** test of the **VMware vSphere ESX** component, your entry should be: **VmEsx_i_server =. . . . . . .,EsxGuestTest:Swapped_memory:Swap memory**

     To know the internal component type, test, and measure names, refer to page of this document.

   - Finally, save the file.

     It is recommended that you configure comparison graphs for only those measures that support descriptors.

6. If a comparison graph appears too cluttered, you can view that graph more clearly by enlarging it. To do so, click on that graph. Figure 10.157 will then appear as depicted below:

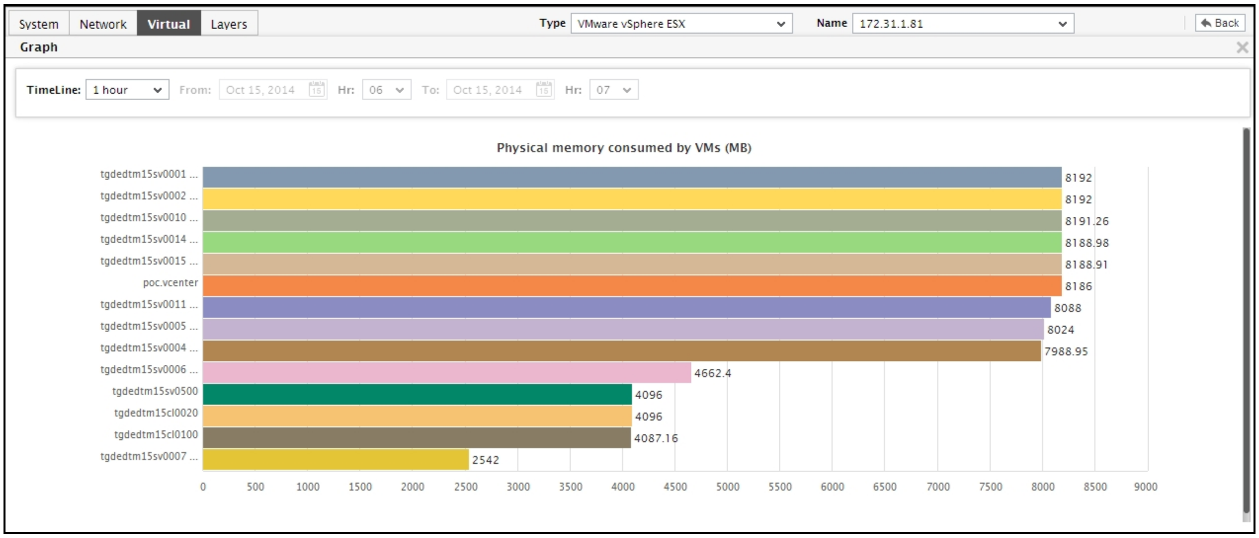

Figure 10.157: Enlarging a comparison bar chart in the Physical Server Memory Analysis dashboard

7. By default, in the enlarged mode, the comparison bar chart will display all the compared elements (i.e., descriptors). This is why, the **Show** list is set to **ALL** by default. If you want the graph to display only a few best or a few worst players in a particular area, then, pick a **TOP-N** or **LAST-N** option from the **Show** list.

8. Back in the dashboard, you will find that a **Details of VMs** table follows the comparison bar charts. By default, this table compares the physical and virtual memory usage of each of the VMs so that, you can clearly identify the following:

   - The VM causing a severe dent in the physical memory resources of the host;
   - The VM using the allocated memory resources excessively

   By default, this table is sorted in thedescending order of the **Physical memory consumed** column. To sort the table in the ascending order of the same column, click on the down arrow button that appears adjacent to the column heading, **Physical memory consumed (GB)**. To sort the table on the basis of the values of another column, click on the title of the corresponding column.

   Note:

   You can configure additional measure columns for the **Details of VMs** table by following the procedure discussed below:

   - Edit the **eg_dashboard.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
   - In the **[MemoryAnalysis]** section of the file, you will find a parameter of the following format for each of the virtualization platforms that is being monitored by eG Enterprise:

     <InternalComponentType>_DATA

     This parameter is typically set to a comma-separated list of measures that will appear as columns in the **Details of VMs** table.

     By default, this parameter is set to a pre-configured list of memory usage statistics extracted from each VM on the corresponding **<InternalComponentType>**.

- You can override this default setting by adding more measures to the comma-separated list, or by removing one/more existing measures. To add another measure, append an entry of the following format to the comma-separated list that follows the **<InternalComponentType>_DATA** parameter:

  <InternalTest>:<InternalMeasure>

  For instance, to add a column for the **Current swap memory** measure reported by the **EsxGuestDetails** test of the **VMware vSphere ESX** component, your entry should be: **VmEsx_i_server_DATA =. . . . . . .,EsxGuestTest:Swapped_memory**

  To know the internal component type, test, and measure names, refer to page of this document.

- Finally, save the file.

9. To historically analyze the physical memory usage of the host and the VMs, you can view historical measure graphs in your dashboard, instead of the default comparison bar charts. To do this, click on the icon at the right, top corner of your dashboard. This will result in the display of measure graphs, plotted for a default duration of 1 hour, for each of the measures for which comparison bar charts were originally displayed.

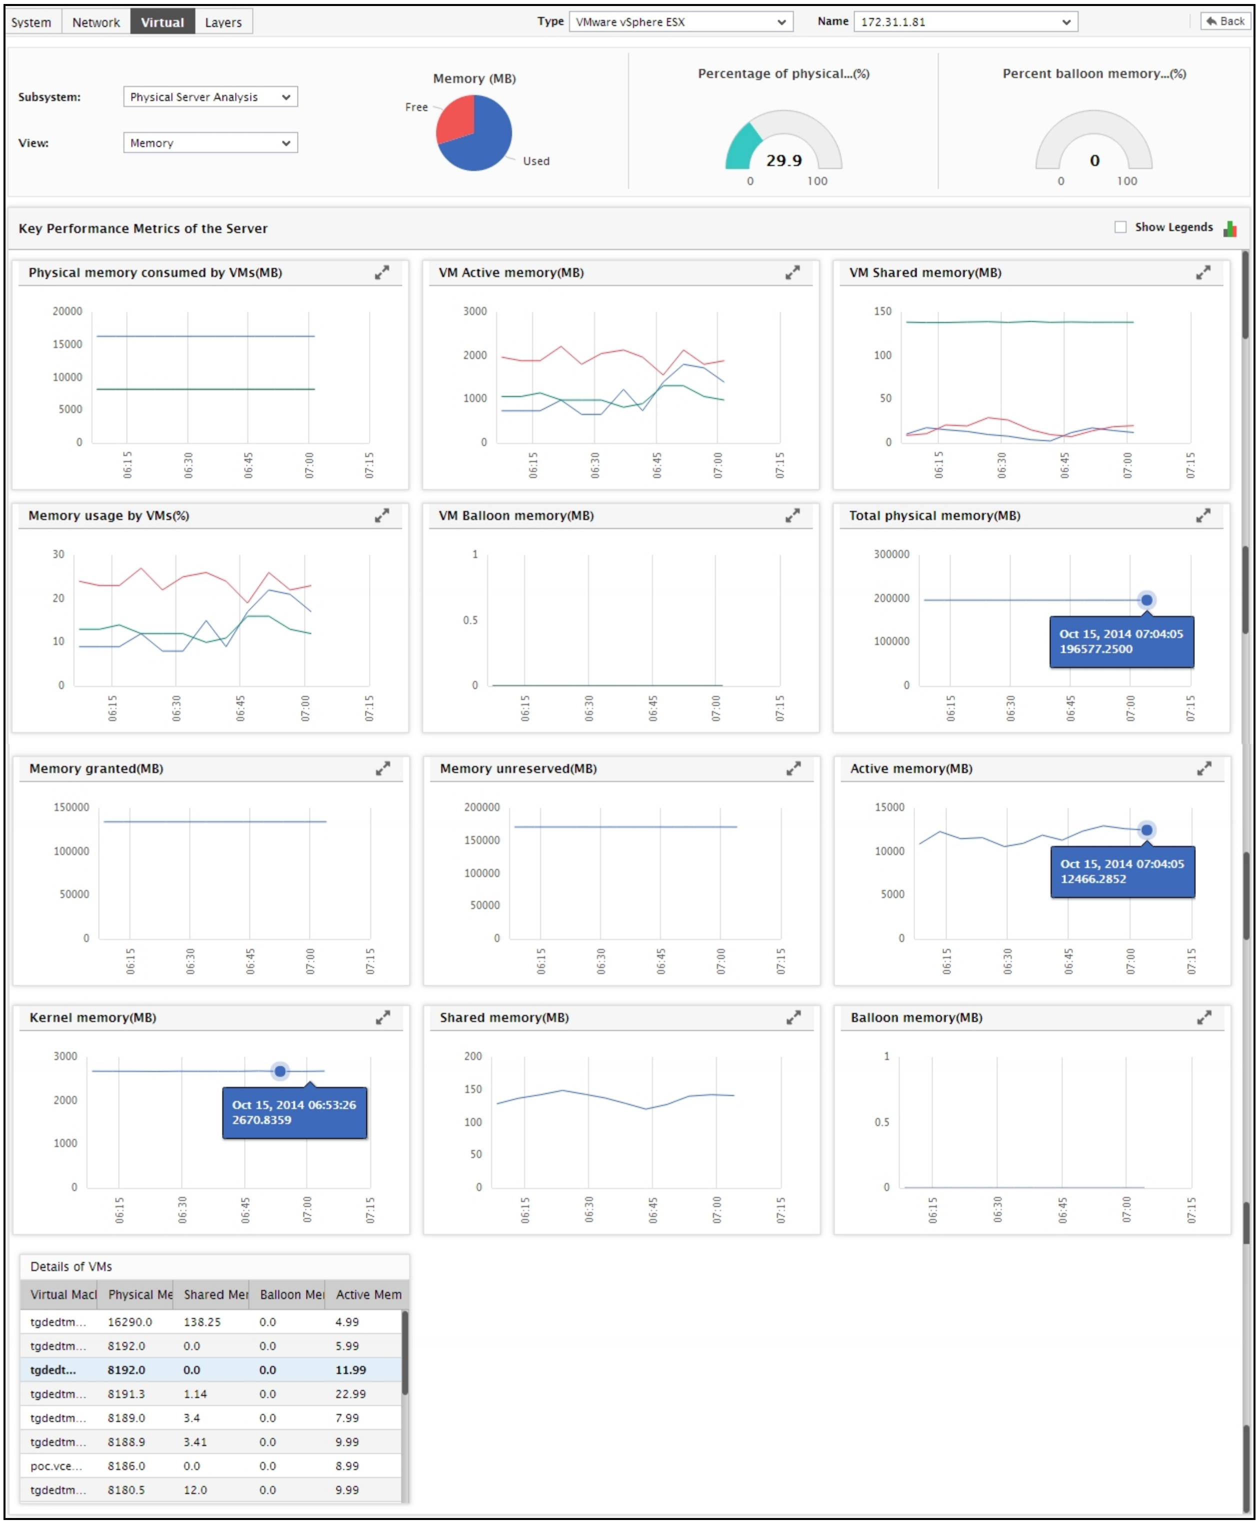

Figure 10.158: Viewing historical measure graphs in the Physical Server Memory Analysis dashboard

10. In the event of a contention for memory resources, you can use these measure graphs to understand how memory was used during the last hour, and when the memory erosion actually began.

11. To view a measure graph more clearly, click on it. This will enlarge the graph (see Figure 10.159).

Figure 10.159: Enlarging a measure graph in the Physical Server Memory Analysis dashboard

12. In the enlarged mode, you can change the **Timeline** of the graph. You can also pick a **TOP-N** or **LAST-N** option from the **Show** list to analyze the time-of-day variations in the performance of a few best/worst players in the chosen performance arena. To return to the dashboard, click on the enlarged graph.

10.6.3 Physical Server Analysis - Disk

To quickly analyze disk activity and disk space usage by the virtual host and its VMs, and to precisely point to the VMs that are consuming too much disk space, select **Disk** from the **View** list. Figure 10.160 will then appear.

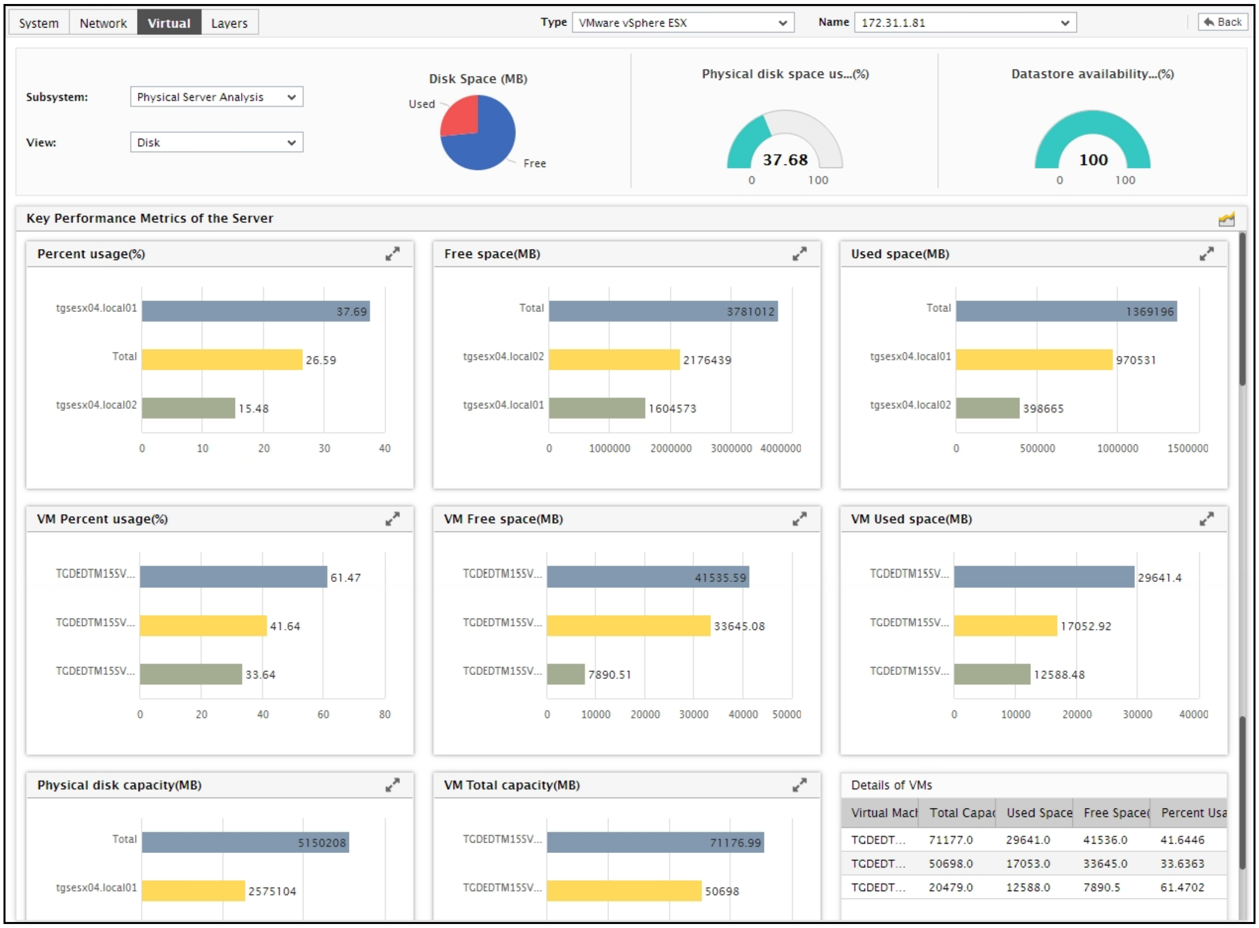

Figure 10.160: The Physical Server Memory Disk dashboard

The contents of Figure 10.160 have been detailed below:

1. A glance at the pie chart in the **Disk** dashboard will indicate to you whether a space crunch exists on the host or not.
   Note:
   By default, the pie chart will denote the **Free space** and **Used space** on the target virtual host. If you want the pie chart to represent more measures, do the following:

   - Edit the **eg_dashboard.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
   - In the **[DiskAnalysis]** section of the file, you will find a parameter of the following format for each of the virtualization platforms that is being monitored by eG Enterprise:

     <InternalComponentType>_PIE

     This parameter is typically set to a comma-separated list of measures that will be represented in the pie chart in the **Disk** dashboard.

     By default, this parameter is set to the **Free space** and **Used space** measures.
   - You can override this default setting by adding more measures to the comma-separated list,

or by removing one/more existing measures. To add another measure, append an entry of the following format to the comma-separated list that follows the **<InternalComponentType>_PIE** parameter:

<InternalTest>:<InternalMeasure>:<DisplayName>

For instance, to make sure that the pie chart also represents the value of the **Physical disk capacity** measure reported by the **Datastores - Esx** test of the **VMware vSphere ESX** component, your entry should be: **VmEsx_i_server_PIE=. . . . . .,EsxDiskUsageTest:Total_capacity:Physical disk capacity**

To know the internal component type, test, and measure names, refer to page of this document.

2. The dial charts in the dashboard indicate the current state and values for critical space usage metrics related to the host operating system. You can promptly detect which disk partition / storage adapter on the host is currently facing a space crunch, using these dial charts.

   Note:

   You can configure what disk space usage statistics need to be represented using dial charts by following the steps discussed below:

   - Edit the **eg_dashboard.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
   - In the **[DiskAnalysis]** section of the file, you will find a parameter of the following format for each of the virtualization platforms that is being monitored by eG Enterprise:

     <InternalComponentType>_GAUGE

     This parameter is typically set to a comma-separated list of measures for which dial charts are to be displayed in the dashboard.

     By default, this parameter is set to a pre-configured list of space usage-related metrics extracted from the corresponding **<InternalComponentType>.**
   - You can override this default setting by adding more measures to the comma-separated list, or by removing one/more existing measures. To add another measure, append an entry of the following format to the comma-separated list that follows the **<InternalComponentType>_GAUGE** parameter:

     <InternalTest>:<InternalMeasure>:<DisplayName>

     For instance, to add a dial chart for the **Percent usage** measure reported by the **DiskSpace - Console** test of the **VMware vSphere ESX** component, your entry should be: **VmEsx_i_server_GAUGE=. . . . . .,CSLDiskSpaceTest:Percent usage:Percent usage**
   - To know the internal component type, test, and measure names, refer to page of this document.
   - Finally, save the file.

     Note that dial charts can be used to represent only those measures that report values in percentage.

3. Clicking on a dial chart will lead you to the **Layer** tab page, where you can view the layer and test that reported the measure represented by that dial chart.

4. Below the dial charts, you will find a default collection of comparison bar charts that reveal how disk space is utilized by the disk partitions on the host, and how the VMs use the disk space allocated to them. Using these graphs, disk partitions and VMs that use their disk space resources excessively, can be clearly identified.

   Note:

   You can configure what space usage statistics need to be represented using the comparison bar charts, by following the steps discussed below:

   - Edit the **eg_dashboard.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
   - In the **[DiskAnalysis]** section of the file, you will find a parameter of the following format for each of the virtualization platforms that is being monitored by eG Enterprise:

     <InternalComponentType>

     This parameter is typically set to a comma-separated list of measures for which comparison bar charts are to be displayed in the dashboard.

     By default, this parameter is set to a pre-configured list of disk space usage-related metrics extracted from the corresponding **<InternalComponentType>**.
   - You can override this default setting by adding more measures to the comma-separated list, or by removing one/more existing measures. To add another measure, append an entry of the following format to the comma-separated list that follows the **<InternalComponentType>** parameter:

     <InternalTest>:<InternalMeasure>:<DisplayName>

     For instance, to add a comparison graph for the **Percent usage** measure reported by the **DiskSpace - Console** test of the **vSphere/ESX(i)** component, your entry should be: **VmEsx_i_server =. . . . . .,CSLDiskSpaceTest:Percent_usage:Percent usage**

     To know the internal component type, test, and measure names, refer to page of this document.
   - Finally, save the file.

     It is recommended that you configure comparison graphs for only those measures that support descriptors.

5. If a comparison graph appears too cluttered, you can view that graph more clearly by enlarging it. To do so, click on that graph. The graph will then appear as depicted below:

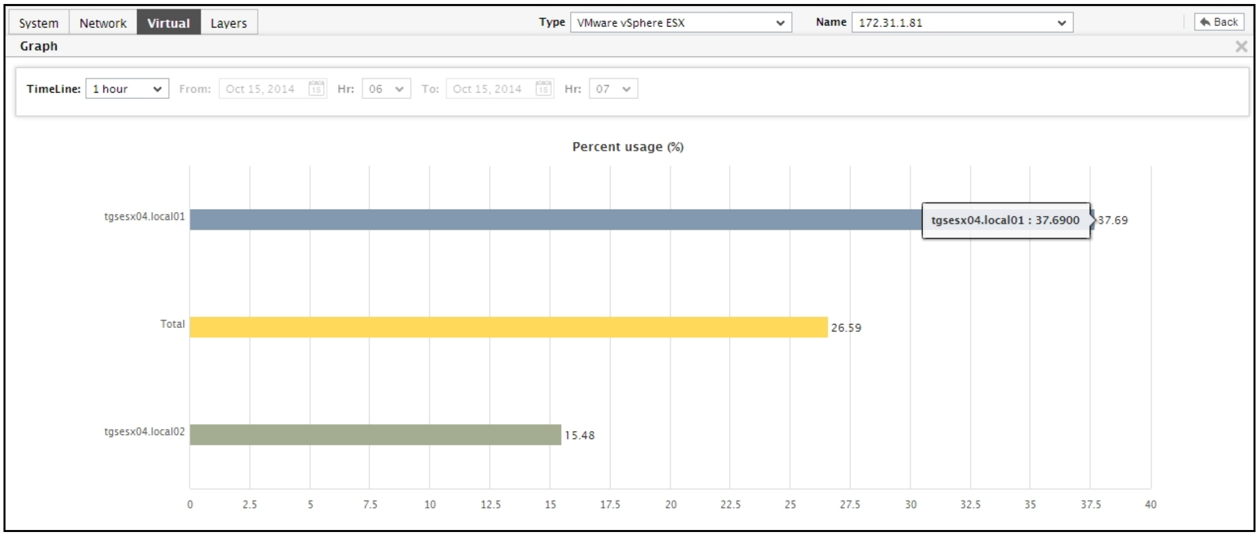

Figure 10.161: Enlarging a comparison bar chart in the Physical Server Disk Analysis dashboard

6. By default, in the enlarged mode, the comparison bar chart will display all the compared elements (i.e., descriptors). This is why, the **Show** list is set to **ALL** by default. If you want the graph to display only a few best or a few worst players in a particular area, then, pick a **TOP-N** or **LAST-N** option from the **Show** list.

7. Back in the dashboard, you will find that a **Details of VMs** table follows the comparison bar charts. By default, this table compares the space usage of the VMs configured on the target host so that, you can clearly identify the VM that may soon run out of space

   By default, this table is sorted in the descending order of the **Total capacity** column. To sort the table in the ascending order of the same column, click on the down arrow button that appears adjacent to the column heading, **Total capacity (GB)**. To sort the table on the basis of the values of another column, click on the title of the corresponding column.

   Note:

   You can configure additional measure columns for the **Details of VMs** table by following the procedure discussed below:

   - Edit the **eg_dashboard.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
   - In the **[DiskAnalysis]** section of the file, you will find a parameter of the following format for each of the virtualization platforms that is being monitored by eG Enterprise:

     <InternalComponentType>_DATA

     This parameter is typically set to a comma-separated list of measures that will appear as columns in the **Details of VMs** table.

     By default, this parameter is set to a pre-configured list of space usage statistics extracted from each VM on the corresponding **<InternalComponentType>**.

   - You can override this default setting by adding more measures to the comma-separated list, or by removing one/more existing measures. To add another measure, append an entry of the following format to the comma-separated list that follows the **<InternalComponentType>_DATA** parameter:

<InternalTest>:<InternalMeasure>

For instance, to add a column for the **Disk reads** measure reported by the **EsxGuestDetails** test of the **vSphere/ESX(i)** component, your entry should be: **VmEsx_i_server_DATA =. . . . . . .,EsxGuestTest:Reads**

To know the internal component type, test, and measure names, refer to page of this document.

- Finally, save the file.

8. To historically analyze the disk space usage of the host and the VMs, you can view historical measure graphs in your dashboard, instead of the default comparison bar charts. To do this, click on the icon at the right, top corner of your dashboard. This will result in the display of measure graphs, plotted for a default duration of 1 hour, for each of the measures for which comparison bar charts were originally displayed.

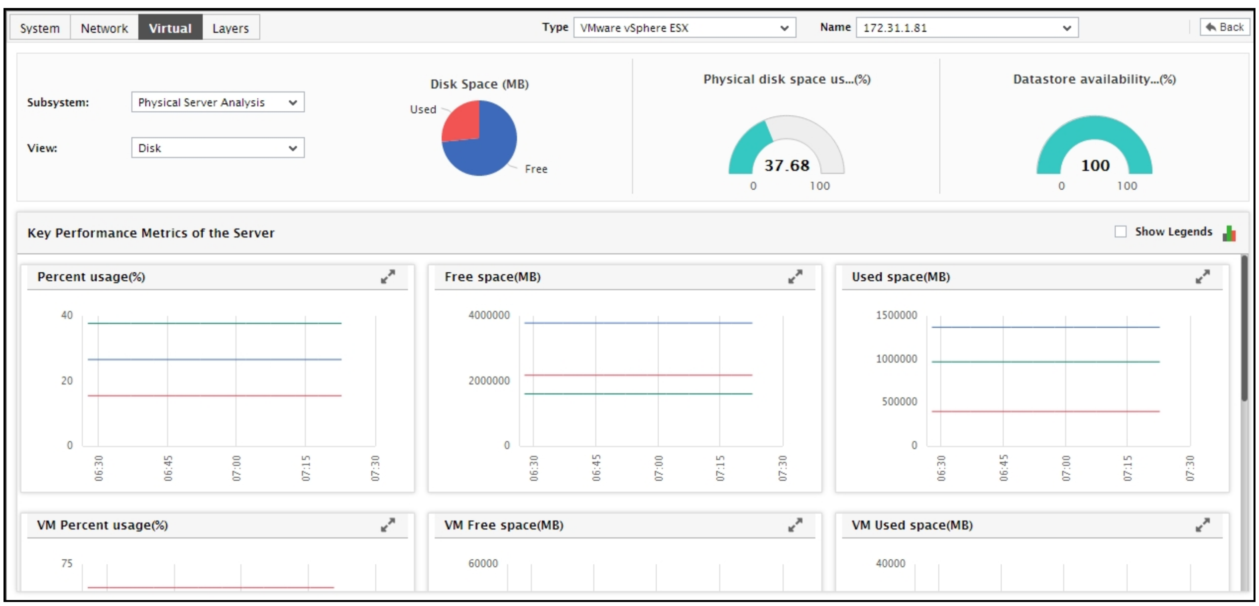

Figure 10.162: Viewing historical measure graphs in the Physical Server Memory Analysis dashboard

9. In the event of a space crunch, you can use these measure graphs to understand how disk space was used by the host and the VMs during the last hour, and when the space erosion actually began.

10. To view a measure graph more clearly, click on it. This will enlarge the graph (see Figure 10.163).

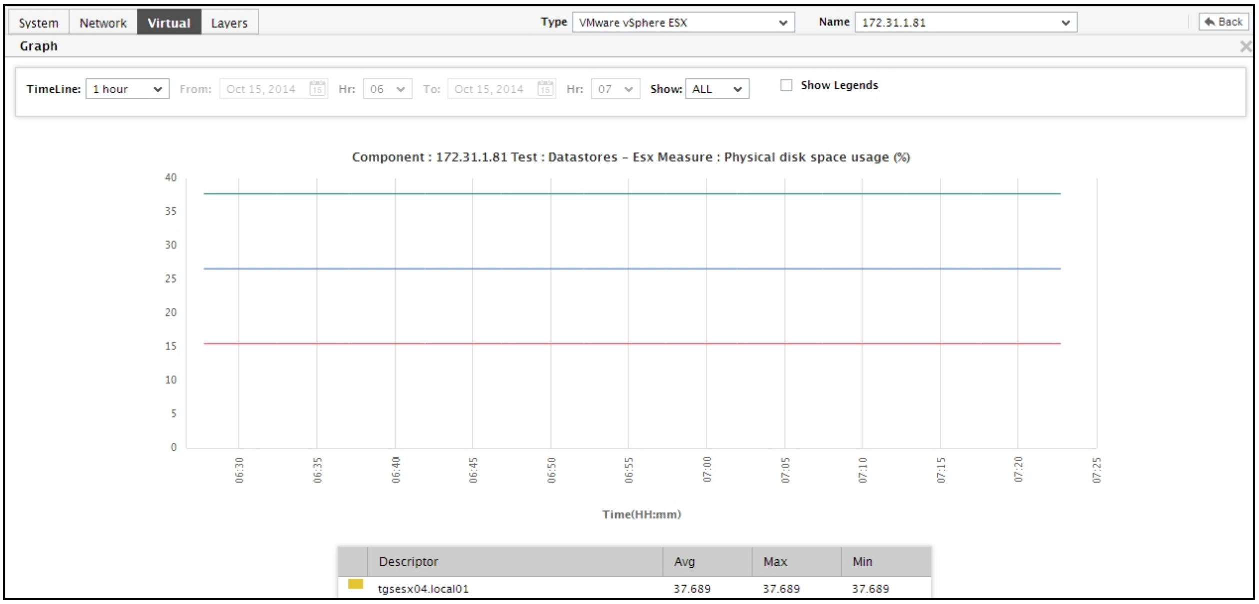

Figure 10.163: Enlarging a measure graph in the Physical Server Memory Analysis dashboard

11. In the enlarged mode, you can change the **Timeline** of the graph. You can also pick a **TOP-N** or **LAST-N** option from the **Show** list to analyze the time-of-day variations in the performance of a few best/worst players in the chosen performance arena. To return to the dashboard, click on the enlarged graph.

10.6.4 Outside View Analysis

Knowing how the VMs on a virtual host are utilizing the physical resources of the host will enable administrators to isolate resource-deficiencies at the host level, and to accurately identify resource-intensive VMs. The **Outside View** dashboard provided by eG Enterprise facilitates such an analysis for each VM configured on a virtual host. To view this dashboard, select the **Outside view of VMs** option from the **Subsystem** list, and pick the VM for which the outside view is to be analyzed, from the **VM** list.

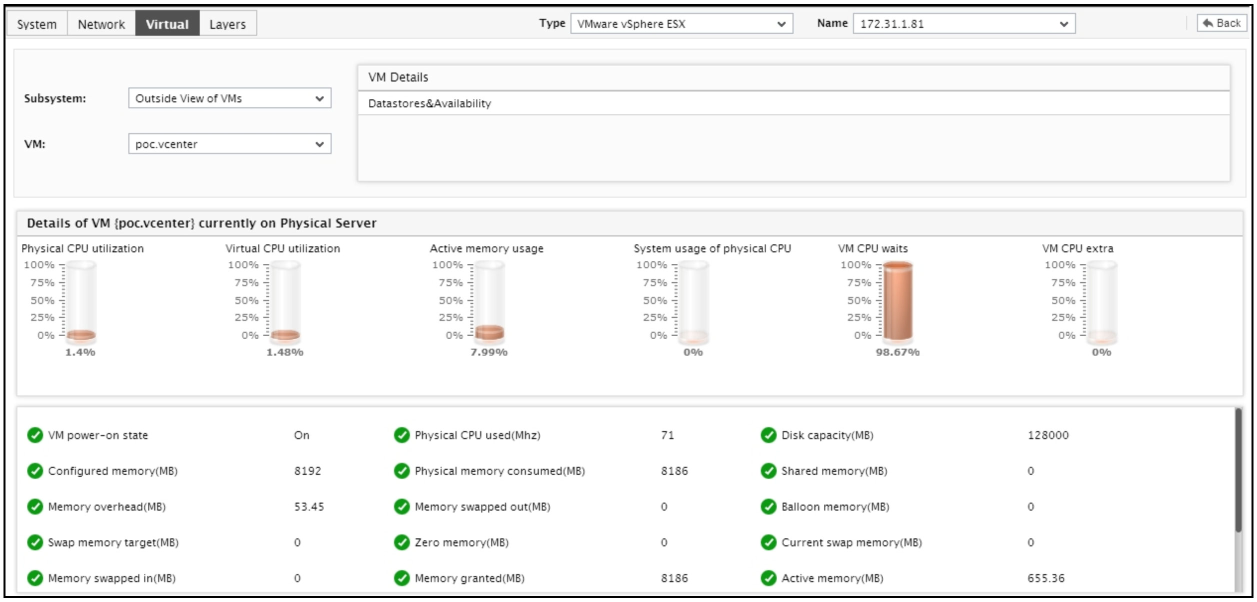

Figure 10.164: The Outside view of VMs dashboard

Doing so will invoke the following details (see Figure 10.164):

1. The **VM Details** section displays the IP address of the chosen VM, the operating system on which the VM runs, the datastores used by the VM and the availability of each datastore.
2. Below this section you will find a series of bar charts, each reporting the current value of a pre-configured outside view metric. From these bar charts, you can quickly figure out how that VM is utilizing the physical CPU, memory, and disk resources of the host.

   Note:

   You can configure what outside view statistics need to be represented using bar charts by following the steps discussed below:

   - Edit the **eg_dashboard.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
   - In the **[OutsideViewAnalysis]** section of the file, you will find a parameter of the following format for each of the virtualization platforms that is being monitored by eG Enterprise:

     <InternalComponentType>_BAR

     This parameter is typically set to a comma-separated list of measures for which bar charts are to be displayed in the dashboard.

     By default, this parameter is set to a pre-configured list of outside view metrics extracted from the corresponding **<InternalComponentType>.**
   - You can override this default setting by adding more measures to the comma-separated list, or by removing one/more existing measures. To add another measure, append an entry of the following format to the comma-separated list that follows the **<InternalComponentType>_BAR** parameter:

     <InternalTest>:<InternalMeasure>:<DisplayName>

For instance, to add a bar chart for the **Memory swap in rate** measure reported by the **EsxGuestDetails test** test of the **vSphere/ESX(i)** component, your entry should be: **VmEsx_i_server_BAR=. . . . . . .,EsxGuestTest:Swap_In_Rate:Swap rate**

To know the internal component type, test, and measure names, refer to Page of this document.

- Finally, save the file.

3. Below the bar charts, you will find a metrics table, where pre-configured outside view metrics are displayed, and the current value and state of each metric is reported. This way, you can quickly spot any sudden change in the state or values of critical statistics.

Note:

You can configure what outside view metrics are to be displayed in the metrics table, by following the steps below:

- Edit the **eg_dashboard.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
- In the **[OutsideViewAnalysis]** section of the file, you will find a parameter of the following format for each of the virtualization platforms that is being monitored by eG Enterprise:

  <InternalComponentType>_data

  This parameter is typically set to a comma-separated list of measures that need to appear in the metrics table in the outside view dashboard.

  By default, this parameter is set to a pre-configured list of outside view metrics extracted from the corresponding **<InternalComponentType>**.
- You can override this default setting by adding more measures to the comma-separated list, or by removing one/more existing measures. To add another measure, append an entry of the following format to the comma-separated list that follows the **<InternalComponentType> _DATA** parameter:

  <InternalTest>:<InternalMeasure

  For instance, to add the **Memory swap in rate** measure reported by the **EsxGuestDetails test** test of the **vSphere/ESX(i)** component to the metrics table, your entry should be:: **VmEsx_i_server_DATA =. . . . . . ., EsxGuestTest:Swap_In_Rate:Swap rate**

  To know the internal component type, test, and measure names, refer to page of this document.
- Finally, save the file.

10.6.5 Inside View Analysis

Though the outside view can point you to the resource-intensive VM on a virtual host, you would have to zoom into the problem VM to figure out what is causing the resource contention. For this purpose, you require an *inside view* of a VM. Using the **inside view** dashboard provided by eG Enterprise, you can perform an effective analysis of the internal workings of a particular VM.

To access this dashboard, pick the **Inside view of VMs** option from the **Subsystem** list in the **Virtual** tab page, and then pick the **VM** of interest to you.

10.6.6

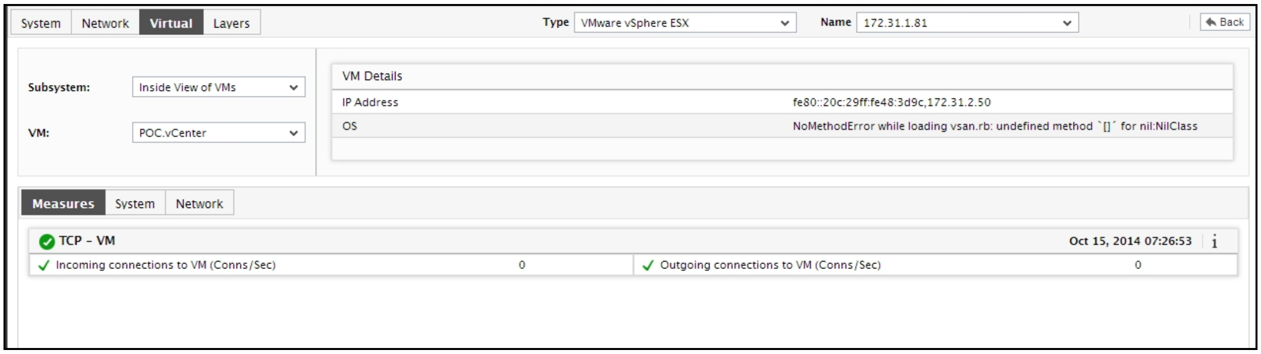

Figure 10.165: Figure 10.168: The Inside view of VMs dashboard

Doing so will display all the performance metrics extracted from within the chosen guest, in real-time. You are thus enabled to cross-correlate across the various metrics, and quickly detect the root-cause of current/probable disturbances to the internal health of a guest.

To view the time-of-day variations in a measure, you can view its graph by clicking on that measure. By default, the graph in Figure 10.166 is plotted for a period of 1 hour. If need be, you can alter the graph **Timeline**, so that you can effectively analyze performance patterns across longer/shorter time periods.

If the VM had been migrated/Vmotioned to another ESX server during the specified **Timeline**, then the resulting graph will depict how that measure has performed in both the ESX servers. This way, you can better understand the impact of a VM migration across ESX servers, and decide on resource allocations and subsequent migrations accordingly.

10.7 The VM Dashboard

In large virtualized infrastructures today, a number of business-critical applications are being deployed on tens of virtual machines configured on a variety of virtualization platforms (e.g., VMware ESX, Citrix XenServer, Solaris LDoms, etc.), so as to optimize space and resource usage. In such environments, excessive resource usage by a single VM or a resource pool on a virtual host can cause a huge dent in the resources available to other VMs, thereby affecting the performance of the applications executing on those VMs! To ensure a high uptime for their key applications, administrators need to track in real-time the resource usage across VMs and physical servers regardless of the underlying virtualization platform, quickly detect abnormal usage patterns, accurately identify the VM(s) responsible for the same, and promptly initiate corrective action. Likewise, in environments where centralized management tools such as VMware vCenter are used, administrators also need to keep tabs on the availability and overall health of these tools, so as to ensure that performance degradations that the tool experiences does not impact the performance of the ESX servers it manages.

eG Enterprise provides a single, central **VM Dashboard** that provides an integrated interface from where administrators can compare resource usage across physical servers and VMs on each physical server, and provides them with real-time insights into the health of the physical servers, the status of the VMs, and how each VM currently uses the allocated and physical resources available to it.

Using this dashboard, administrators can:

- Understand, from a single look, the composition of the managed virtualized environment - in other words, using the dashboard, you can easily view and instantly figure out the complete hierarchical structure of your virtualized environment; for instance, for a VMware environment, the dashboard will help you identify the number of vCenter servers that are being managed, the number of datacenters configured on each vCenter, the clusters (if any) that exist within the datacenters, the ESX servers being managed by every datacenter, the resource pools on each ESX server, and the VMs that each resource pool comprises of.
- Detect at a glance, excessive resource usage by any VM, cluster, resource pool, or a physical server from across the environment, regardless of the virtualization technology in use; quickly diagnose the root-cause of the resource drain with the help of efficient drill down features;
- Accurately identify resource-intensive VMs/resource pools/clusters/physical servers;
- Promptly detect unavailable datastores;
- Instantly spot a powered-off VM anywhere in the environment;
- Know which VM is currently operating on which physical server, and thus keep tabs on VMotion/XenMotion (as the case may be) activities;
- View a consolidated list of issues currently encountered by physical servers and virtual machines across the environment and also per virtual component, so as to ease the troubleshooting efforts of a dedicated help-desk;
- Quickly identify and troubleshoot issues with vCenter servers (if any) in the environment;

To access the **VM Dashboard**, click on the icon available in the **Monitor** tab. Then, select the **Virtual Dashboard** option in the **Dashboards** tile. Figure 10.166 then appears:

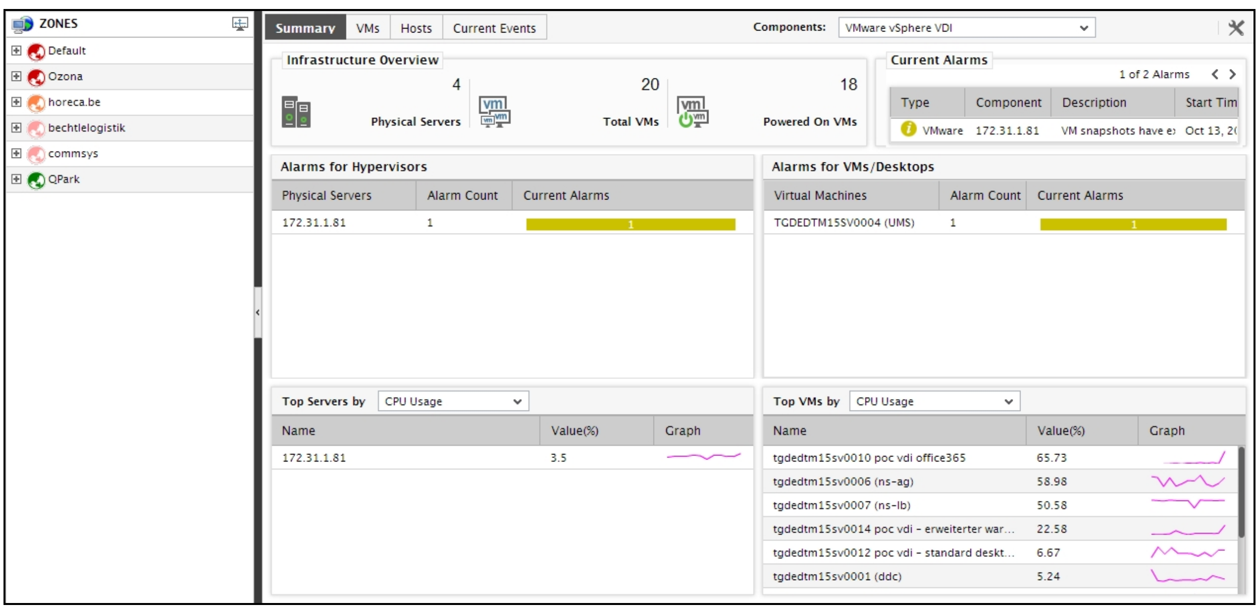

Figure 10.166: Figure 10.169: The VM Dashboard

The **VM Dashboard** of Figure 10.166 comprises of two panels. The left panel boasts of a tree-structure, comprising of a default global node named **Zones**. All the zones/farms in the target environment that have been configured with one/more virtual hosts (eg., VMware ESX servers, Citrix XenServers, etc.) or VMware vCenter servers, will be the sub-nodes of the **Zones** node. If you expand a particular zone node in the tree, you will find that the virtual component-types that have been added to that zone appear as its sub-nodes. If you expand any node under the zone node that corresponds to a particular virtual server-type (eg., *VMware vSphere ESX*, *VMware vSphere VDI*, *Citrix XenServer*, etc.), then such a tree will typically comprise of virtual hosts of that type that are included in the zone. Expanding the virtual host node will reveal the VMs executing on that host and resource pools (if any) configured on that host.

In addition to sub-nodes representing a virtual server-type, if a *VMware vCenter* component has been added to the zone, then a *VMware vCenter* sub-node will also appear under the corresponding zone node. If you expand the node representing the *VMware vCenter* component-type under a zone, then all the managed vCenter servers will appear as its sub-nodes. Expanding a particular vCenter server node will reveal the folders (if any) configured on that vCenter server; similarly, expanding a particular folder node will reveal the datacenters (if any) that are being managed by that vCenter server. To view the virtual hosts/clusters configured within a datacenter, you will have to expand the corresponding datacenter node. While the cluster tree will contain virtual hosts within the cluster as its sub-nodes, expanding a virtual host tree will reveal the VMs and resource pools that are executing on that virtual host. You can also view the virtual host tree by expanding any node that corresponds to a virtual host type under the a particular zone node. The state of a node in the tree depends upon the current state of its sub-nodes. Independent virtual hosts/vCenter servers that are not part of any existing zone will be automatically grouped under the **Default** zone tree (see Figure 10.167).

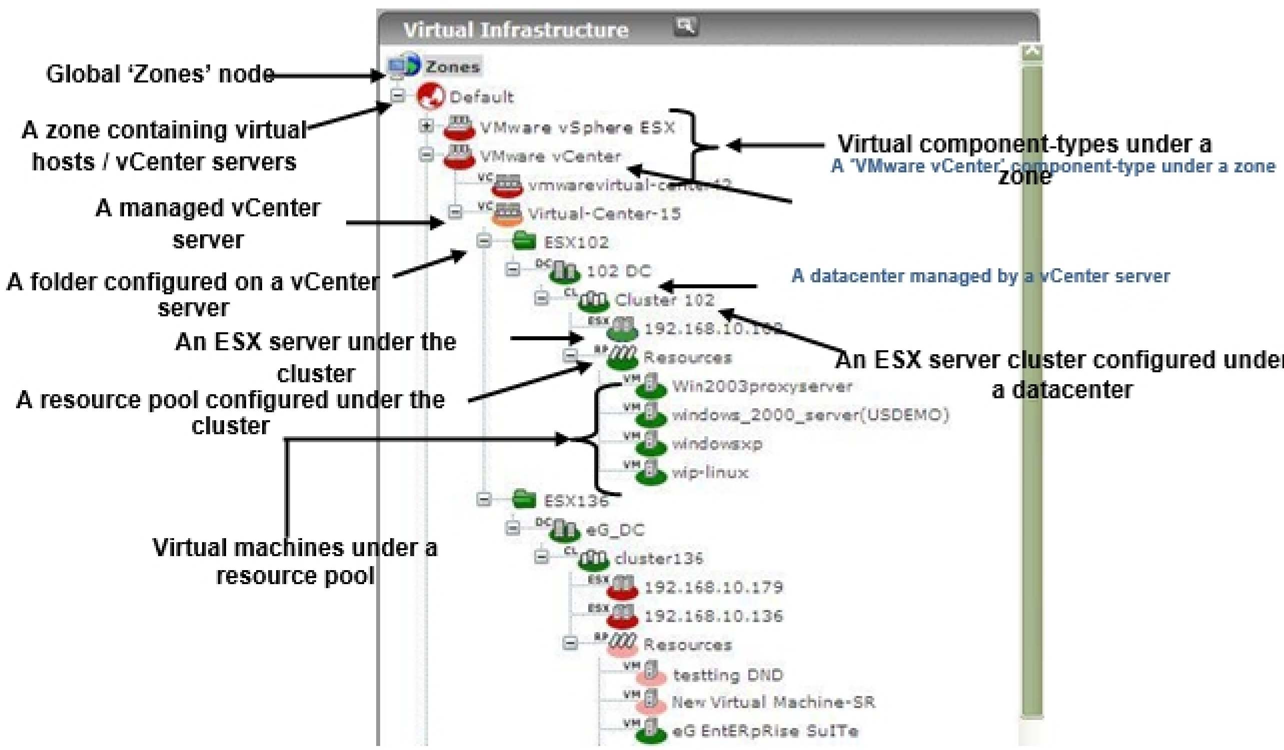

Figure 10.167: Figure 10.170: The Virtual Infrastructure tree

By default, the tree-structure lists all virtual hosts, clusters, and virtual machines. Accordingly, when you click on the button at the right, top corner of the left panel, both the **Hosts And Clusters** and **Show Virtual Machines** options will be chosen by default from the **Display Settings** section (see Figure 4.84) that appear. To view only the VMs in the tree, select the **Show Virtual Machines** check box, but deselect the **Hosts And Clusters** check box. To hide VMs from the tree-view, select the **Hosts And Clusters** check box, but leave the **Show Virtual Machines** check box deselected.

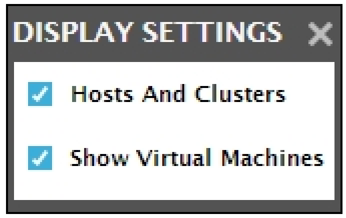

Figure 10.168: Figure 10.171: Display settings for the tree-view

The right panel (see Figure 10.168) is a context-sensitive panel, the contents of which will vary according to the node chosen from the left panel. By default, this panel provides 4 tab pages - a **Summary** tab page that provides a quick summary of performance and problem information pertaining to the node chosen from the tree, a **VMs** tab page that provides current status updates related to virtual machines, a **Hosts** tab page that displays virtual host-specific metrics gathered in real-time, and a **Current Events** tab page that lists the problems currently experienced by virtual hosts and guests. By default, all the tab pages provide information pertaining to managed *VMware vSphere ESX servers* in the environment. Accordingly, the **VMware vSphere ESX** option is chosen from the drop-down list in the top, right corner of the right panel (indicated by Figure 10.168). You can view details for a different virtual component type, by selecting a different option from this drop-down list.

Note:

As stated earlier, by default, the tab pages in the right panel provide details related to *VMware vSphere ESX* servers. This default setting can be overridden in the following manner:

- Click on the button at the right, top corner of the right panel. Figure 10.169 then appears.

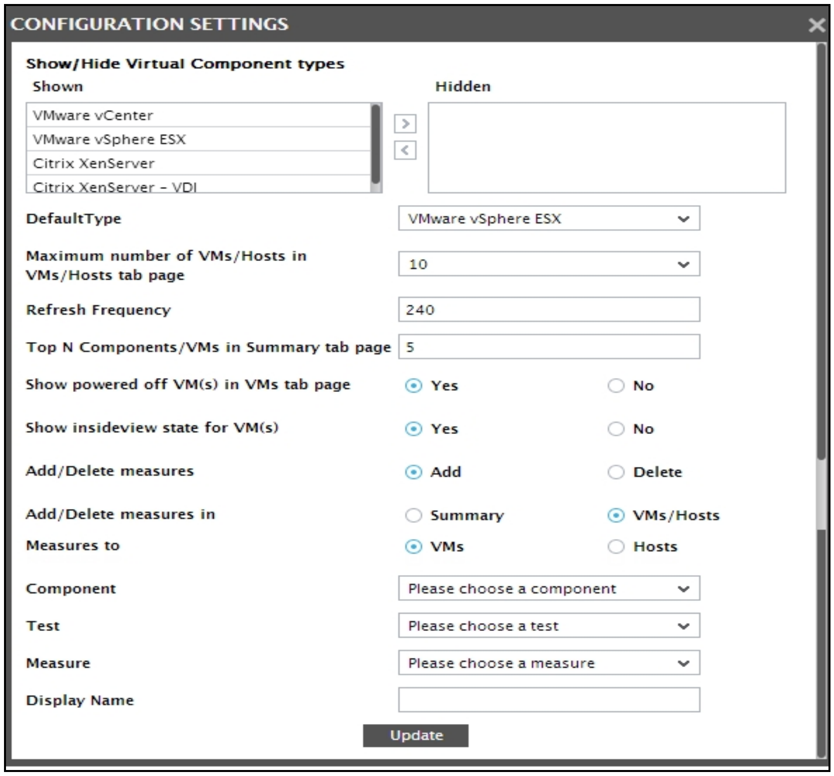

Figure 10.169: Figure 10.172: The Configuration Settings page

- Where vSphere servers are managed, by default, the *VMware vSphere ESX* option is chosen from the **DefaultType** list in Figure 10.169, indicating that the details in the right panel pertain to the managed vSphere components by default. To change the default setting, you will have to select a different option from the **DefaultType** list as indicated by Figure 10.169.
- Then, click the **Update** button in Figure 10.169.

Similarly, every tab page, by default, displays the details of the top 10 VMs/virtual hosts only. This is indicated by the default value **10** chosen from the **Limit to** list at the right, top corner of each tab page. To view more (or less) number of records in the tab pages, you will have to select a different value from the **Limit to** list. Alternatively, you can even override the default value **10**, so that the tab pages display more or a less number of records by default. To achieve this, click the button in the right panel, and set the **Maximum number of VMs/Hosts in VMs/Hosts tab page** parameter in Figure 10.169 that appears next to any value other than 10.

The sections to come take you on a guided tour of each of the tab pages.

10.7.1 The Summary Tab Page

The **Summary** tab page serves as a single, central interface that combines ‘problem and performance information’ related to virtual infrastructures. Using this tab page, administrators can perform the following with ease:

- Oversee, by a mere glance, the composition of and the all-round performance of the virtual infrastructure as a whole or of the virtual infrastructure component chosen from the tree;
- View a consolidated list of current alarms pertaining to the node chosen from the tree, and instantly identify problem-prone virtual infrastructure elements;
- Receive real-time updates of the resource usage of physical servers and virtual machines, and instantly identify the hosts/guests experiencing a resource contention;
- Easily analyze and accurately detect disconcerting trends in the resource usage of the physical servers and virtual machines.

This section elaborately discusses the contents and usage of this tab page.

1. If the global **Zones** node is chosen from the tree, the **Summary** tab page in the right panel will, by default, provide a quick overview of the composition and performance of the monitored *VMware vSphere ESX* servers spread across all the managed zones in the environment (see Figure 10.170). To view the performance summary of a different virtualization platform, select a different option from the drop-down list in the right, top corner of the **Summary** tab page. By default, *VMware vSphere ESX* is chosen from this list.

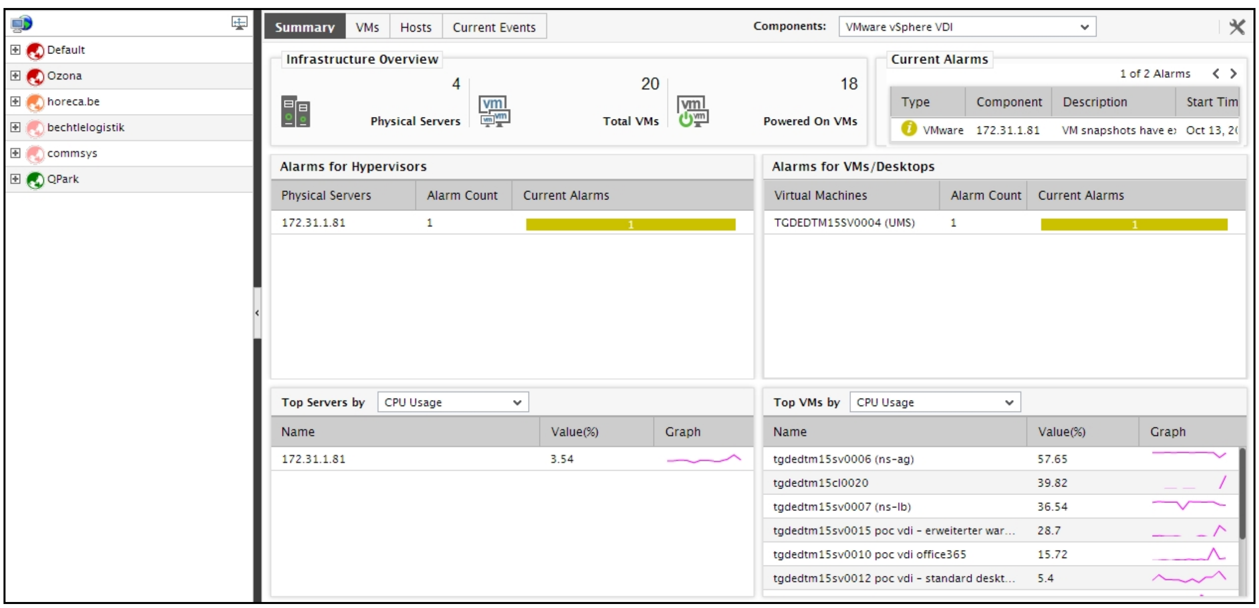

Figure 10.170: The Summary tab page if the global Zones node is clicked

2. For the default *VMware vSphere ESX* infrastructure, this tab page provides an **Infrastructure Overview** section that briefly discusses the key ingredients of the virtual infrastructure - i.e., the number of managed physical servers in the target virtual environment, the total number of VMs on the physical servers, the number of VMs that are currently powered on, and the number of alarms currently open on the infrastructure elements. For more details about the target environment, simply move your mouse pointer over every value displayed in the **Infrastructure Overview** section. For instance, to know the names of physical servers that are being managed in the environment, simply move the mouse pointer over the value corresponding to **Physical Servers** in the **Infrastructure Overview** section. The **Summary** tab page will then change to display a pop-up that lists the names of the physical *VMware vSphere ESX* servers managed in the environment (see 10.7.1). This way, you can view the names of virtual machines executing on the

physical servers, the names of powered on VMs, and also the list of current alarms pertaining to the environment. Besides helping you identify VMs that are powered-off currently, the **Infrastructure Overview** also enables you determine the number and nature of the unresolved problems in the environment.

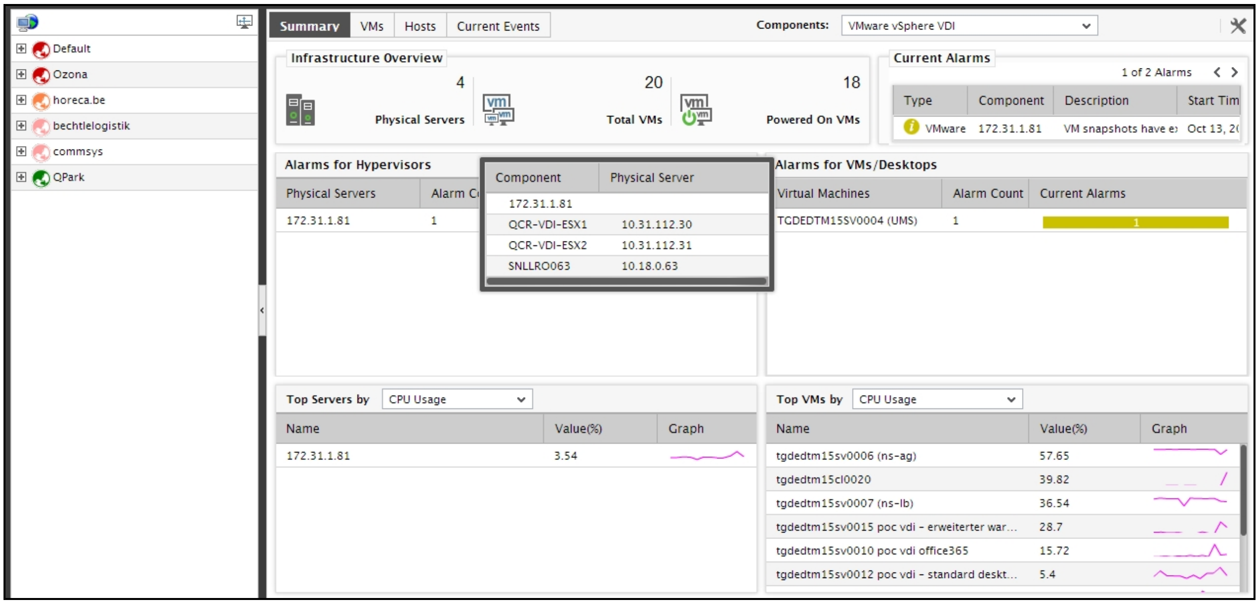

Figure 10.171: The names of physical VDI servers displayed in a pop-up

3. Now that you know the names of the physical servers, you might want to analyze the current resource usage of each of these servers to ascertain whether they are experiencing any resource shortages or not. For that, click on the **Physical Servers** label in the **Infrastructure Overview** section of Figure 10.173. This leads you straight to the **Hosts** tab page of the global **Zones** node, which displays the physical servers, their current state, and also the resource usage metrics pertaining to each server (see Figure 10.173).

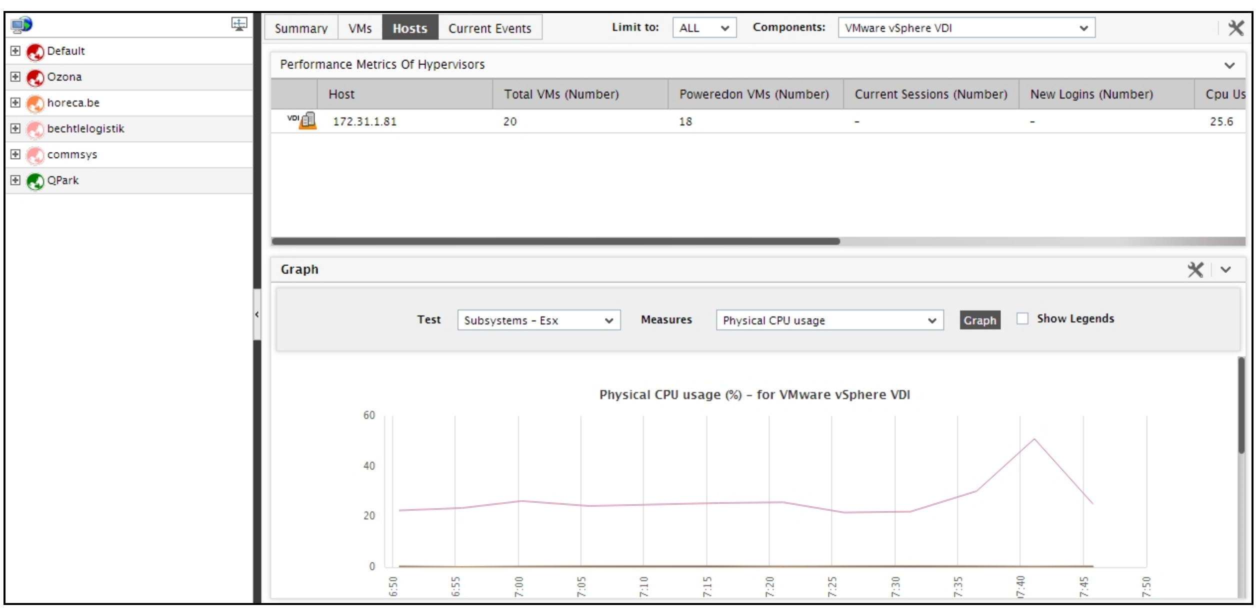

Figure 10.172: The Hosts tab page appears when the ‘Physical servers’ label in the Summary tab page is clicked

4. Similarly, you can click on the **Total VMs** label in the **Infrastructure Overview** section of 10.7.1 to switch to the **VMs** tab page of the **Zones** node, focus on the performance of the individual VMs, and identify the VM that could be consuming resources excessively (see 10.7.1).

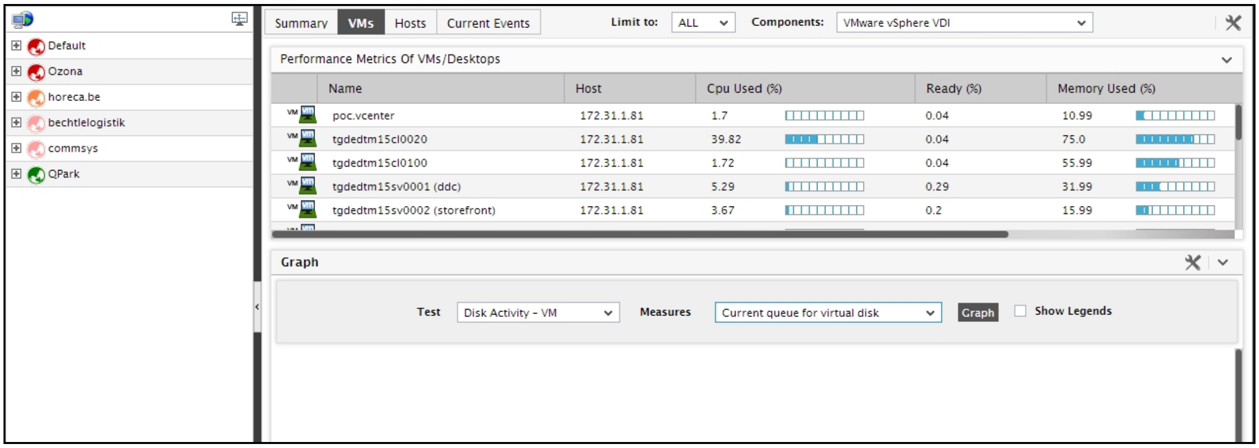

Figure 10.173: The VMs tab page that appears when the ‘VMs’ label in the Summary tab page is clicked

5. While clicking on the **Powered On VMs** label in the **Infrastructure Overview** section takes you to the **VMs** tab page and allows you to analyze the resource usage of powered-on VMs alone (see 10.7.1), clicking on the **Current Alarms** label leads you straight to the **Current Events** tab page, where you can view the complete list of problems the target virtualized environment (see Figure 10.173) is currently experiencing.

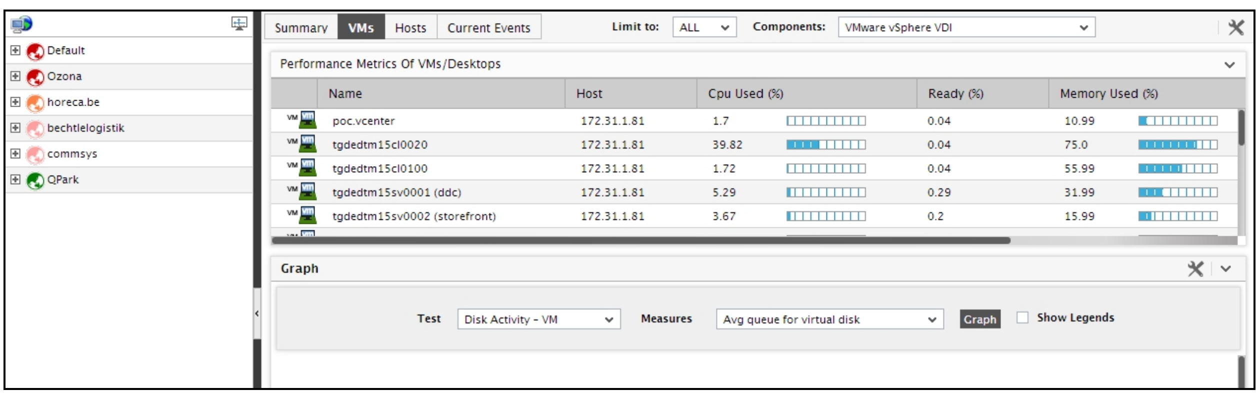

Figure 10.174: The VMs tab page listing only powered-on VMs

Figure 10.175: The Current Events tab page listing all the current problems related to the virtualized infrastructure

6. If you want to focus on each problem closely, then, you can use the **Current Alarms** section adjacent to the **Infrastructure Overview** section in the **Summary** tab page. Use the arrow buttons (<< and >>) above an alarm to navigate to the next alarm. Move your mouse pointer over an alarm to know which test has reported the problem.
7. The **Physical Servers** section below the **Infrastructure Overview** section enables you to determine how problem-prone the physical servers in your environment are, by revealing the number of critical, major, and minor issues that are currently unresolved for each physical server. By moving your mouse pointer over an alarm priority corresponding to a physical server, you can view the details of current alarms of that priority (see Figure 10.175).

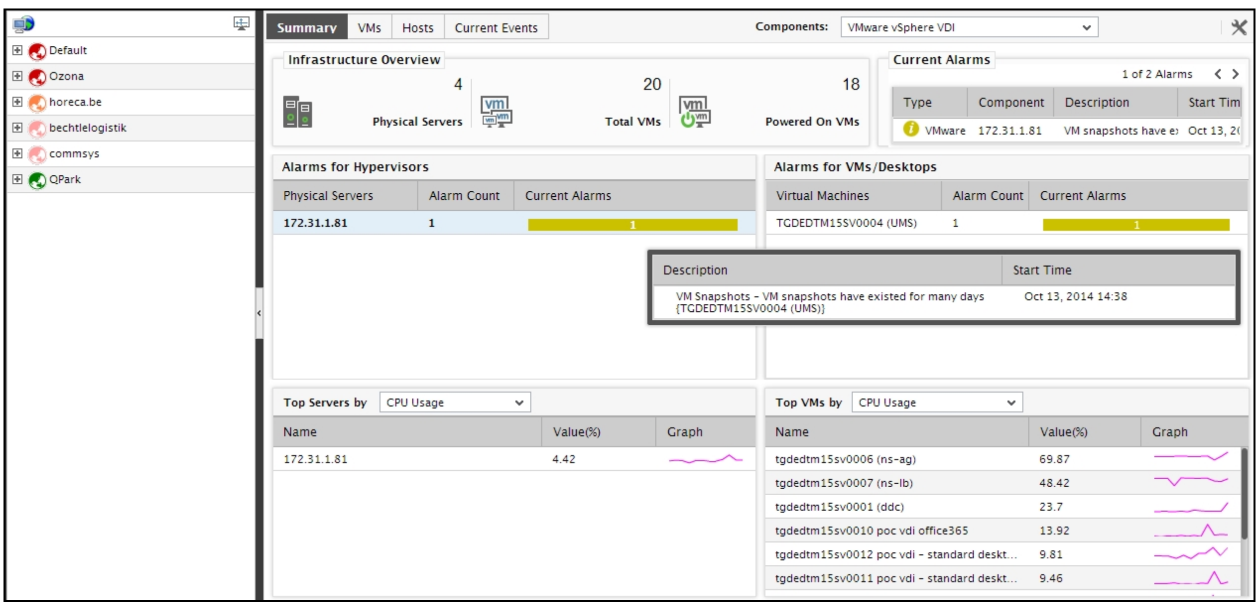

Figure 10.176: Moving the mouse pointer over an alarm priority that corresponds to a physical server

8. By clicking on a physical server in the **Alarms for Hypervisor** section, you can zoom into the layer model of that server; this will indicate all the layers that have been affected by problems. From the color-coding of the layers, you can easily infer from which layer the problem originated (see Figure 10.176). Click on that layer to view the problem tests, and then, click on a problem test to view the measures that have reported anomalies (see Figure 10.176).

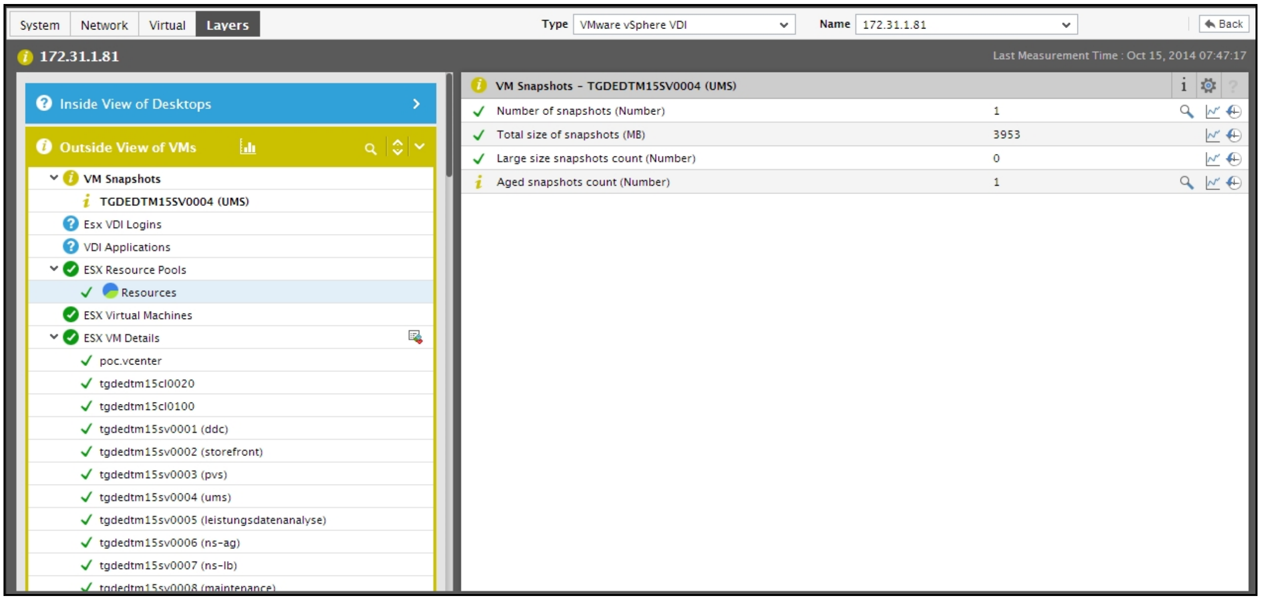

Figure 10.177: The layer model of a physical server indicating the layer where the root-cause of a problem lies

9. The **Alarms for Hypervisor** section enables you to determine how problem-prone the virtual machines in your environment are, by revealing the number of critical, major, and minor issues that currently remain unresolved for each virtual machine. By moving your mouse pointer over an alarm priority corresponding to a virtual machine, you can view the details of current alarms of that priority (see Figure 10.177).

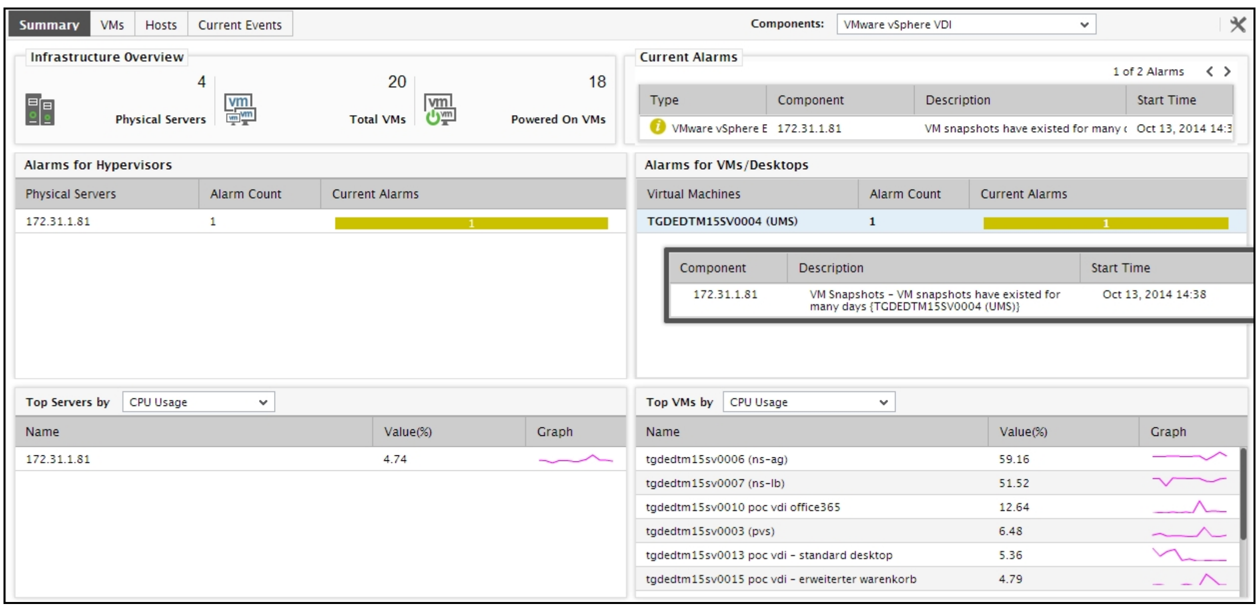

Figure 10.178: The alarms of a particular priority corresponding to a virtual machine

10. To know the exact layer where the problem occurred and the test that reported the problem, click on any virtual machine in the **Virtual Machines** section. Figure 10.179 will then appear indicating the same.

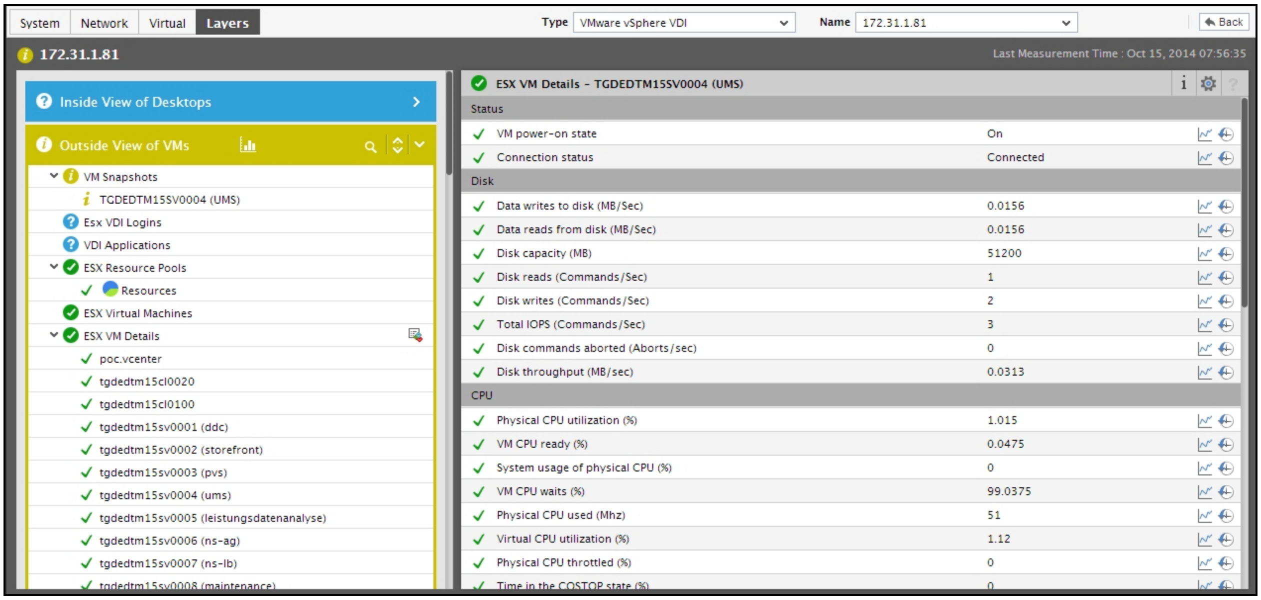

Figure 10.179: The layer model of the physical server on which the chosen VM executes, revealing the exact layer and test that reported the problem clicked on

11. In large virtualized environments comprising of a multitude of virtual hosts that are configured with tens of VMs, it is often difficult for administrators to instantly and accurately locate resource-intensive physical servers and/or virtual machines. Similarly, identifying the physical servers with too many VMs is also a herculean task. In order to ease the pain of the administrators, the **Summary** tab page provides two sections - one each for physical servers and VMs - which can be configured to list the top physical servers and VMs (respectively) in the environment, in terms of resource consumption. Also, by default, both sections will reveal the top consumers of physical CPU resources, starting with the leading consumer. Accordingly, the *Physical CPU utilization* measure is the default selection in both the **Top Servers by** and the **Top VMs by** lists. To view the top consumers of another resource, select a different measure from these lists.

    In the **Top Servers by** list, in addition to the *Physical CPU utilization* measure, the following measures are available for selection by default: *Used physical memory*, *Used space*, *Registered guests*, and *VM power-on state*. If need be, you can override this default setting, so that new measures can be added to the list and one/more existing measures can be removed from the list. To do this, follow the steps given below:

    - Click on the button at the right, top corner of the **Summary** tab page.
    - Figure 10.178 then appears. To add a new measure to the **Top Servers by** list in the **Summary** tab page, first select the **Add** option from the **Add/Delete measures** section, and then pick the **Summary** option from the **Add/Delete measures in** section in Figure 10.178.

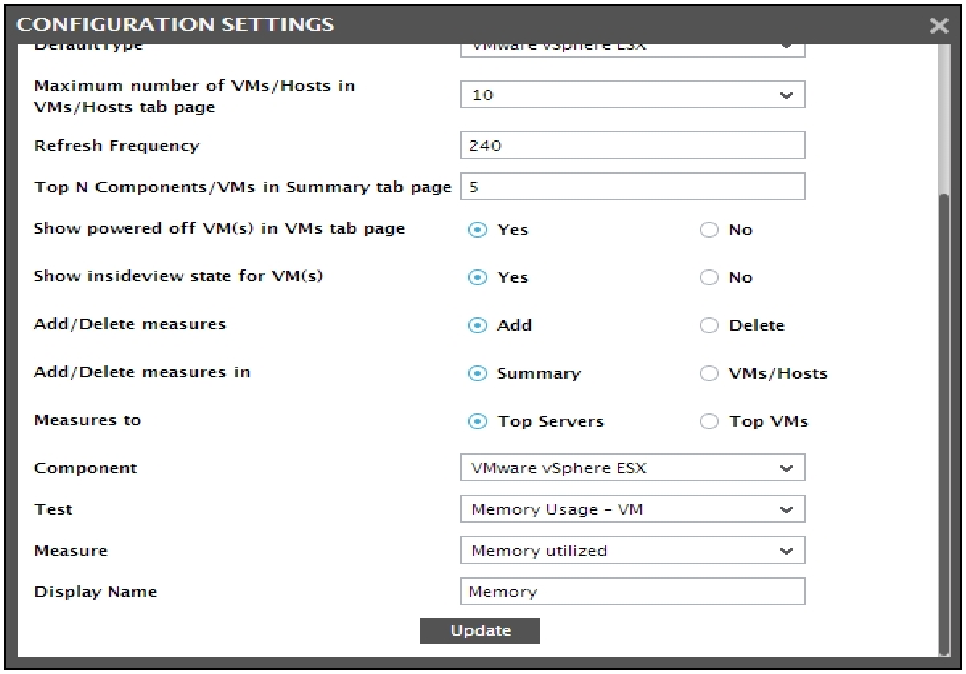

Figure 10.180: Adding a new measure to the Top Servers by list

- Next, set the **Measures to** flag to **Top Servers**.
- Then, pick the **Component** to which the new measure pertains.
- Select the **Test** that reports the measure of interest.
- Select the **Measure** to be added, and provide a **Display Name** for the measure.
- To add the measure, click the **Update** button. Doing so ensures that the **Display Name** specified in Figure 10.179 appears as an option in the **Top Servers by** list in the **Summary** tab page.
- Similarly, you can remove a measure from the **Top Servers by** list. For this purpose, set the **Add/Delete** flag to **Delete**, select the **Component** to which the measure to be deleted pertains, select the **Test** reporting the measure, select the **Measure** to be deleted, and finally, click the **Update** button.

In the **Top VMs by** list on the other hand, in addition to the *Physical CPU utilization* measure, the following measures are available for selection by default: *Memory usage*, *Disk capacity*, *and Percent disk usage*. If need be, you can override this default setting, so that new measures can be added to the list and one/more existing measures can be removed from the list. To do this, follow the steps given below:

- Click on the button at the right, top corner of the **Summary** tab page.
- Figure 10.180 then appears. To add a new measure to the **Top VMs by** list in the **Summary** tab page, first select the **Add** option from the **Add/Delete measures** section, and then pick the **Summary** option from the **Add/Delete measures in** section in Figure 10.181.

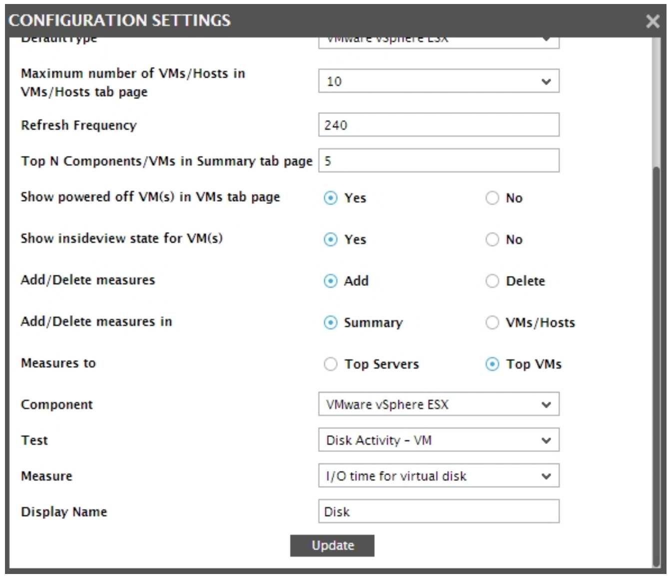

Figure 10.181: Adding a new measure to the Top VMs by list

- Next, set the **Measures to** flag to **Top VMs.**
- Then, pick the **Component** to which the new measure pertains.
- Select the **Test** that reports the measure of interest.
- Select the **Measure** to be added, and provide a **Display Name** for the measure.
- To add the measure, click the **Update** button. Doing so ensures that the **Display Name** specified in Figure 10.181 appears as an option in the **Top VMs by** list in the **Summary** tab page.
- Similarly, you can remove a measure from the **Top VMs by** list. For this purpose, set the **Add/Delete** flag to **Delete**, select the **Component** to which the measure to be deleted pertains, select the **Test** reporting the measure, select the **Measure** to be deleted, and finally, click the **Update** button.

12. Also, by default, both lists will display the top-5 resource consumers only. This default setting can be overridden by following the steps given below:

    - Click on the  button at the right, top corner of the **Summary** tab page.
    - Figure 10.182 then appears. By default, the value **5** is displayed in the **Top N Component/VM in Summary tab page** text box, indicating that the **Summary** tab page displays the top-5 resource consumers, by default.

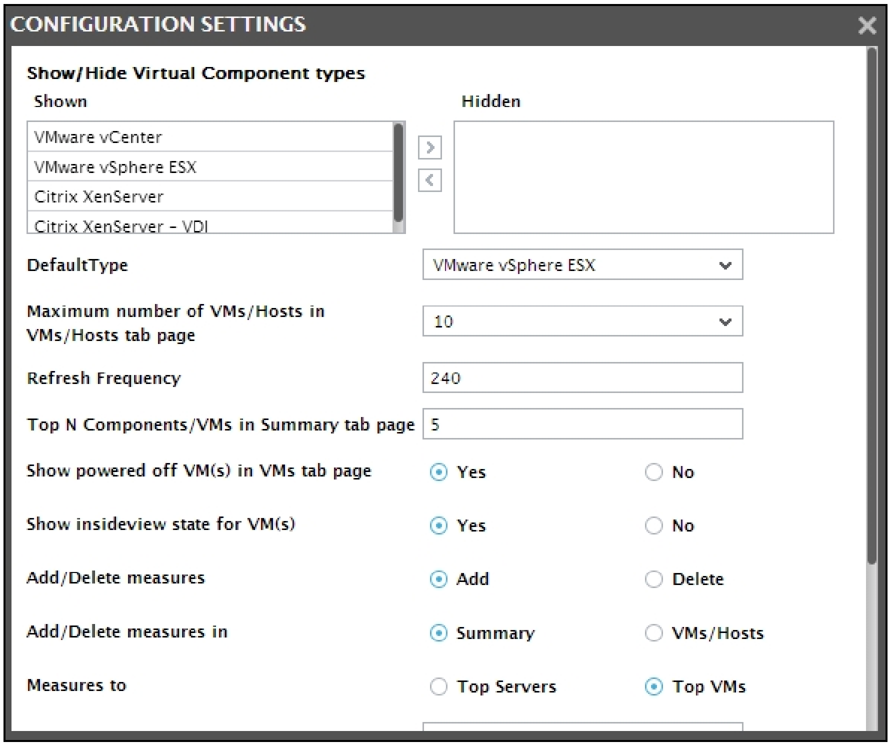

Figure 10.182: The Dashboard configuration page

- Override this default setting by specifying a different number in the **Top N Component/VM in Summary tab page** text box.
- Then, click the **Update** button in Figure 10.183.

13. Let us now focus on the **Top Servers by** section alone. Against every physical server displayed in the **Top Servers by** section, the percentage of the chosen resource currently utilized by each physical server will be displayed, followed by a miniature graph tracking the usage of that resource over a period of time. If you click on a physical server in this section - say, the server that is the leading consumer of physical CPU resources - Figure 10.182 will appear revealing the layer model of that server. From the layer model, you can navigate to the test and the measure reporting the physical CPU usage of the server, perform further analysis, and accurately identify which processor supported by the server has contributed to the excessive resource usage.

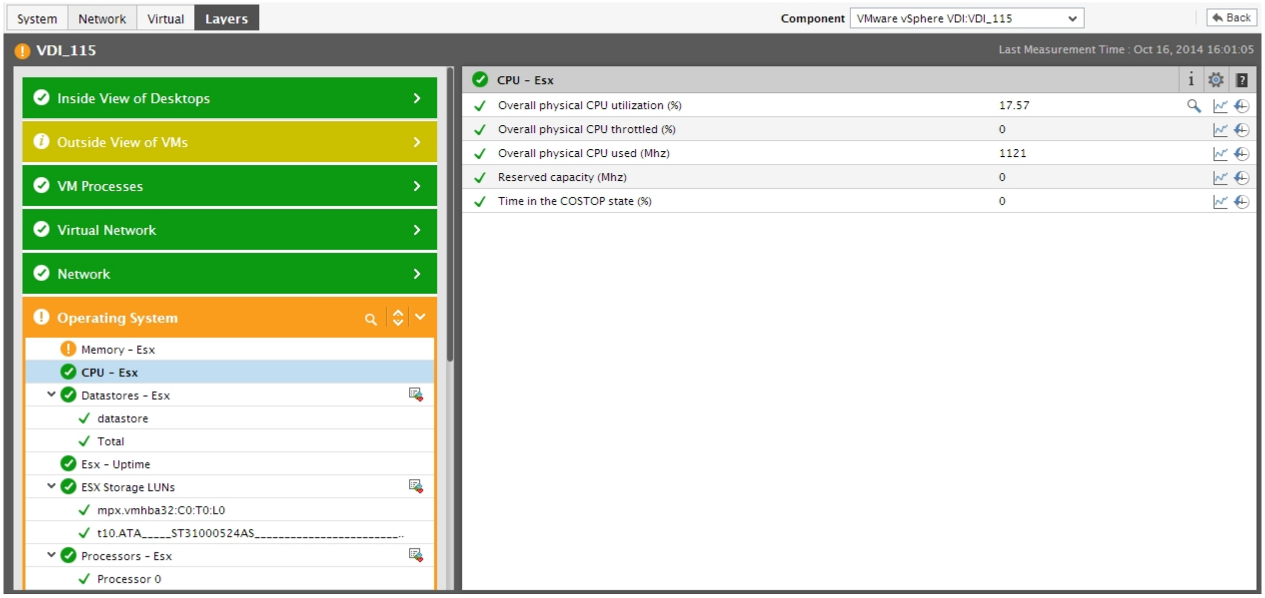

Figure 10.183: The layer model of a CPU-intensive physical server revealing the processor that is excessively consuming physical CPU

14. If you click on the miniature graph that corresponds to a physical server, the graph will expand as depicted by Figure 10.183. By default, the expanded graph tracks the variations in the measure selected from the **Top Servers by** list, during the last 1 hour. With the help of this graph, you can effortlessly observe how the physical server has been using the chosen resource over the last 1 hour by default. You can change this default period by choosing a different **Timeline** for the graph. This analysis will enable you to effectively study usage trends, and accurately detect the exact time at which the physical server began experiencing spikes in resource usage.

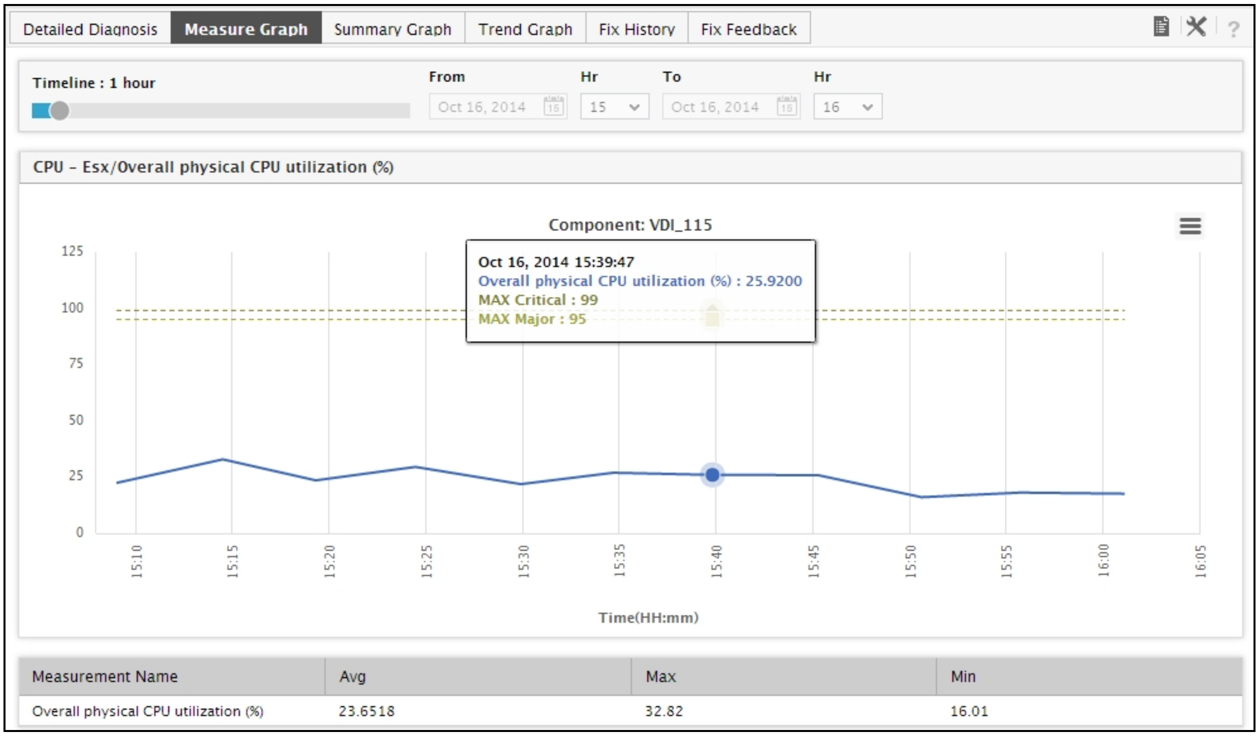

Figure 10.184: The expanded resource usage graph of a physical server

15. Let us now shift our focus to the **Top VMs by** section in the **Summary** tab page. Against every virtual machine displayed in the **Top VMs by** section, the percentage of the chosen resource currently utilized by each VM will be displayed, followed by a miniature graph tracking the usage of that resource over a period of time. If you click on a VM in this section - say, the VM that is the leading consumer of physical CPU resources - Figure 10.185 will appear revealing the layer model of the physical server on which the VM is executing. By default, the test that reports the physical CPU usage of the VM in question will be selected in the layer model, and all the measures reported by that test for the chosen VM will also be displayed. Using these metrics, you can effectively assess the overall resource usage of that VM.

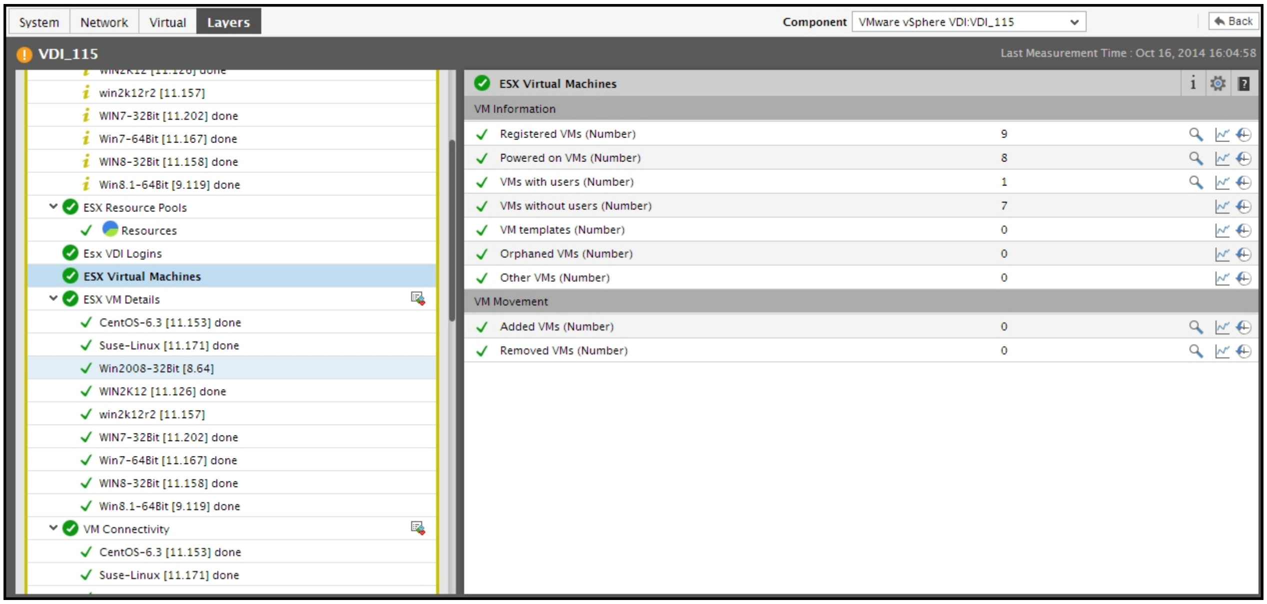

Figure 10.185: The layer model of a CPU-intensive VM revealing the resource usage of the VM

16. If you click on the miniature graph that corresponds to a VM in the **Top VMs by** section, the graph will expand as depicted by Figure 10.184. By default, the expanded graph tracks the variations in the measure selected from the **Top VMs by** list, during the last 1 hour. With the help of this graph, you can effortlessly observe how the VM has been using the chosen resource over the last 1 hour by default. You can change this default period by choosing a different **Timeline** for the graph. This analysis will enable you to effectively study usage trends, and accurately detect the exact time at which the VM began exhibiting unhealthy resource usage trends.

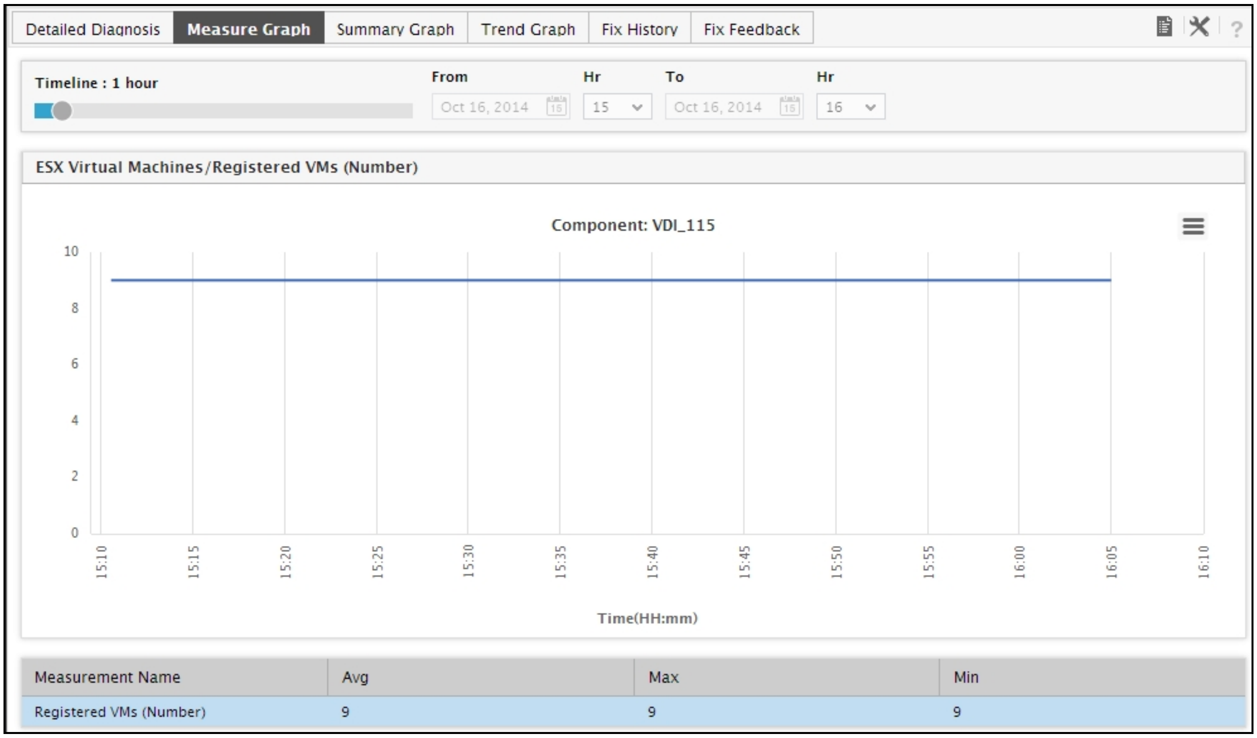

Figure 10.186: The expanded resource usage graph of a VM

17. If the node representing a particular zone is chosen from the tree-structure in the left panel, the contents of the **Summary** tab page will change accordingly. Besides revealing the number of physical servers, VMs, and powered-on VMs in the chosen zone, the **Infrastructure Overview** section also presents a macro view of the health of the zone by indicating the number of unresolved problems in the zone. The **Current Alarms** section will enable you to view these unresolved problems one after another. To know which physical servers in the zone are responsible for these problems, you can take the help of the **Physical Servers** section, which lists the names of the servers along with the number and severity (critical/major/minor) of problems (if any) each server is associated with. Similarly, the **Virtual Machines** section, in addition to displaying the names of VMs that are executing on the physical servers included in the zone, also reveals the problematic VMs by indicating the number and severity of problems (if any) that each VM is currently experiencing. Besides the above, the **Summary** tab page for a particular zone will also enable you accurately identify the resource-intensive physical servers and VMs in a zone. The **Top Servers by** section of this tab page displays the top-5 (by default) resource-hungry physical servers; this enables you to quickly identify the server in the zone that is most resource-intensive. The **Top VMs by** section displays the top-5 (by default) resource-hungry VMs, and thus enables you to identify the most resource-intensive VM in the zone.

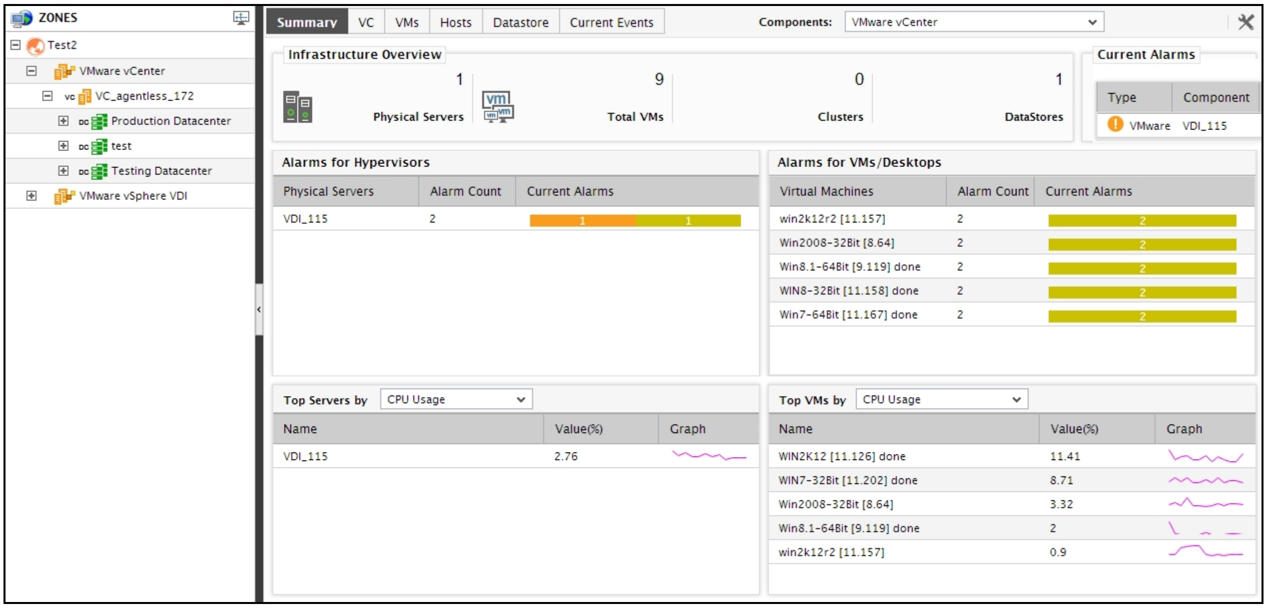

Figure 10.187: The Summary tab page if a particular zone is chosen from the tree-structure in the left panel

18. If a particular virtual server-type (eg., *VMware vSphere ESX*, *Citrix XenServer*, etc.) under a zone is chosen from the tree-structure, then the contents of the **Summary** tab page will change as depicted by Figure 10.185.

19. The **Infrastructure Overview** section in this case will provide a quick performance summary of only those servers in the zone that are of the type chosen from the tree-structure. In other words, if the *VMware vSphere ESX* node is chosen from the tree-structure, then the **Infrastructure Overview** section will display the number of ESX servers in the zone, the number of VMs and powered-on VMs executing on the ESX servers, and the current alarms related to these ESX servers. You can move your mouse pointer over any number displayed in this section, to view the corresponding details.

    To pay individual attention to each alarm, view them one after another in the **Current Alarms** section. You would now want to know which ESX servers in the zone are the most problematic, and also determine whether any of the VMs on ESX servers have contributed to these problems. For this purpose, you can use the **Physical Servers** and **Virtual Machines** sections, which list the problem virtual hosts and guests in the zone, and also indicate the number and type of problems currently encountered by each of the displayed hosts and guests. By moving your mouse pointer over an alarm priority corresponding to a physical server/VM, you can take a quick look at the alarms of that priority that are currently open on that physical server/VM.

    Also, you can receive real-time updates on the resource utilization levels of these hosts and guests from the **Top Servers by** and **Top VMs by** sections; these sections, by default, display the top-5 physical servers and VMs in terms of *Physical CPU utilization*. To view the toppers in a different performance realm, you can select a different measure from the list box available in both the sections. You can even add a new measure to the list or remove one/more of the existing measures using the procedure discussed in page 289 and page 290 of this document. Similarly, to view more number of physical servers and VMs in this section, you can change the default value 5 to a different number using the procedure discussed in page 291 of this document. Using the information provided by these sections, you can determine the current resource usage of each displayed physical server and VM, and view graphs that can enable you to effectively assess the resource usage trends of these components over a broader period of time. Moreover, resource-intensive hosts and guests can be rapidly identified, sporadic/consistent surges in resource utilization by these hosts and guests can be promptly detected, and any potential resource contention can be diagnosed before its too late, and averted.

20. If a zone consists of a VMware vCenter server, then the node that corresponds to this zone in the tree-structure, when expanded, will also reveal a *VMware vCenter* sub-node. If this sub-node is clicked, the contents of the **Summary** tab page in the right panel will change as depicted by Figure 10.186.

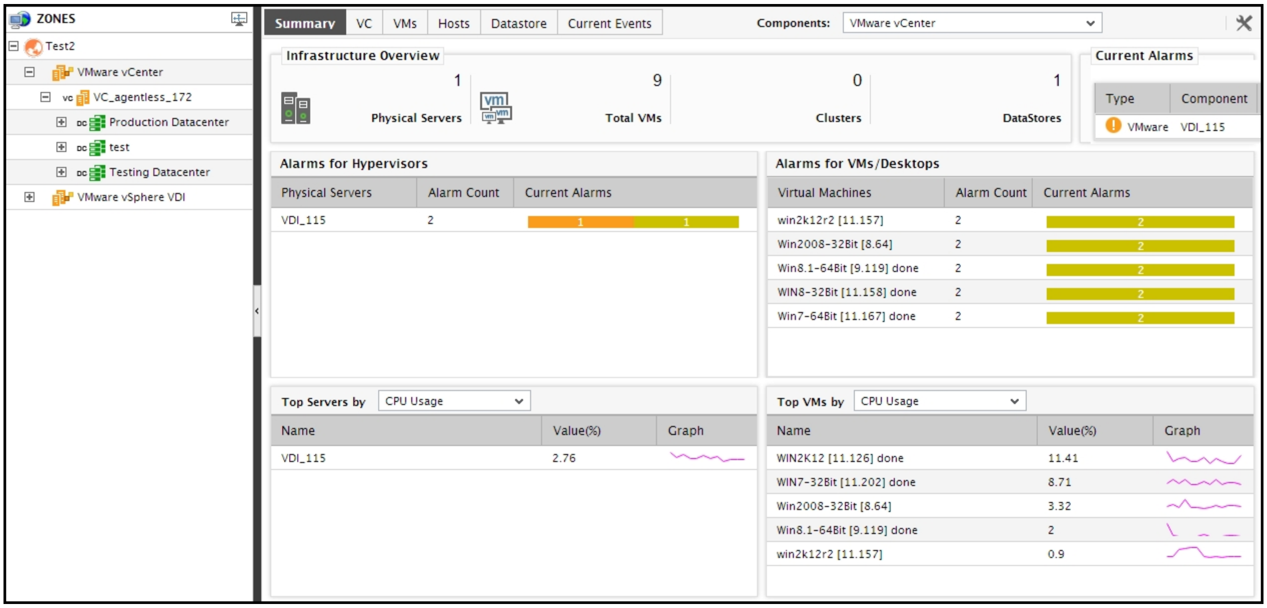

Figure 10.188: The Summary tab page if the 'VMware vCenter' node is clicked

21. The **Infrastructure Overview** section of Figure 10.188 reveals the number of physical servers and VMs managed by all vCenter servers in a zone. Move your mouse pointer over each of these numbers to know the names of the ESX servers and VMs, and also that of the vCenter server managing them. Also, the number of unresolved issues related to these physical servers and VMs, and the number of performance degradations currently experienced by the managed vCenter servers themselves, will be added and displayed as the total number of **Current Alarms** in this section; this will provide you with a fair idea of how healthy the virtualized environment managed by vCenter is. To view the complete list of current alarms, move your mouse pointer over the number of current alarms. You can even focus on every performance issue individually by browsing the alarms, one after another, using the **Current Alarms** section in the **Summary** tab page.

    To know which physical sever and which VM has contributed the maximum to the problems list, use the **Physical Servers** and **Virtual Machines** sections; these sections display the problem-prone ESX servers and VMs (respectively) across vCenter servers in a zone, and indicate how many problems of what severity are currently affecting each of the ESX servers and VMs. Move your mouse pointer over a problem severity corresponding to an ESX server or VM to view the details of the related alarms.

    Besides, you can quickly identify the most resource-hungry ESX servers and VMs across vCenter servers in a zone, using the **Top Servers by** and **Top VMs by** sections in the **Summary** tab page. By default, these sections display the top-5 physical servers and VMs in terms of *Physical CPU utilization*. To view the toppers in the usage of a different resource, you can select a different measure from the list box available in both the sections. You can even add a new measure to the list or remove one/more of the existing measures using the procedure discussed in page 289 and page 290 of this document. Similarly, to view more number of physical servers and VMs in this section, you can change the default value 5 to a different number using the procedure discussed in page 291 of this document. Using the information provided by these sections, you can determine the current resource usage of each displayed physical server and VM, and thus identify the physical server or VM that is consuming resources excessively. Also, by clicking on the miniature graph alongside a physical server or VM, you can expand the graph and effectively analyze the ups and downs in resource usage of the corresponding physical server or VM over time. This way, you can accurately determine whether the increase in resource usage (if any) occurred suddenly, or whether an upward trend in resource usage began earlier on.

22. If you click on a particular vCenter server under the *VMware vCenter* node, then the resulting **Summary** tab page will provide an overview of the performance of that vCenter server alone (see Figure 10.189).

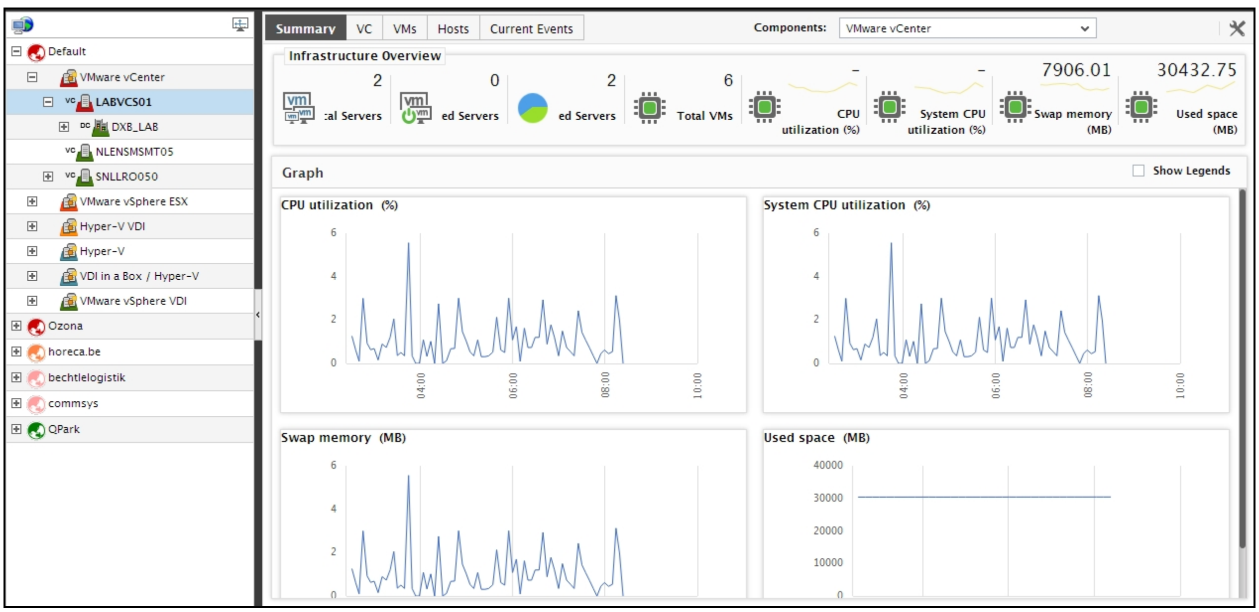

Figure 10.189: The Summary tab page if the node representing a particular vCenter server is clicked in the tree

23. Know how many ESX servers and VMs are managed by the chosen vCenter server using the **Infrastructure Overview** section of the **Summary** tab page. Also, determine how healthy the vCenter server and the virtualized environment it manages is by viewing the number of **Current Alarms** in the **Infrastructure Overview** section. Move your mouse pointer over any number in this section to view the corresponding details.

24. Adjacent to the **Infrastructure Overview** section, you will find a **vCenter Health** section that reports in real-time, the availability and responsiveness of the vCenter server, the load on the server in the terms of current sessions to vCenter, and the current license usage of the server. You can proactively detect the non-availability or a slowdown of the vCenter server, a server overload, or excessive license usage by the server using the metrics reported by this section. Against every value displayed, a miniature graph is available, tracking the time-of-day variations in the values of the corresponding measures. Expand the

graph by clicking on it. The expanded graph, by default, reveals how the corresponding measure has performed during the last 1 hour. You can plot the graph for a broader period by choosing a different **Timeline** for the graph. Using this graph, you can easily analyze the performance of the vCenter over time.

25. The **Graphs** section in the **Summary** tab page provides graphs that enable you to assess the CPU, disk, and memory usage of the vCenter server, over a default period of 6 hours. Resource usage trends can be accurately deduced from these graphs, and probable resource crunches can be proactively detected and averted. To zoom into a particular graph, click on it. The graph then expands (see Figure 10.190), so that you can study it clearly and make sound inferences. You can even change the **Timeline** of the graph, so that you can generate resource usage graphs for longer time periods, and perform more effective analysis.

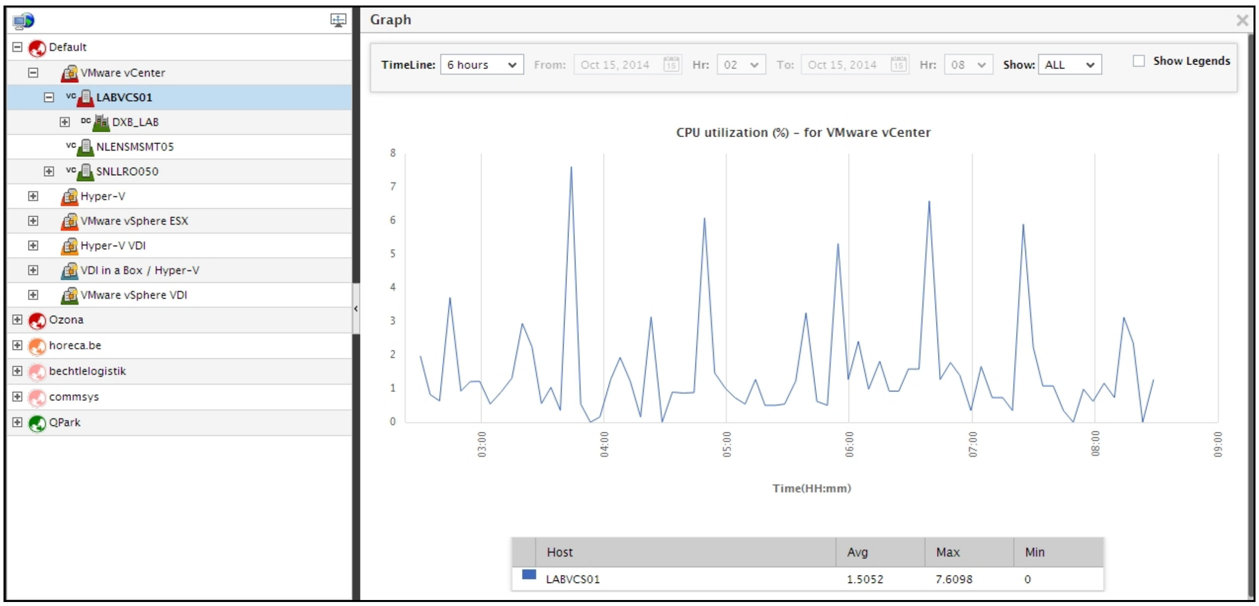

Figure 10.190: Expanding a graph in the Graphs section of the Summary tab page

26. If vCenter manages the target virtualized environment as **folders**, then expanding a vCenter server node in the tree will reveal the folders managed by that vCenter. If you click on the folder node in the tree, the **Summary** tab page will change to provide an overview of the composition and current state of that folder (see Figure 10.191).

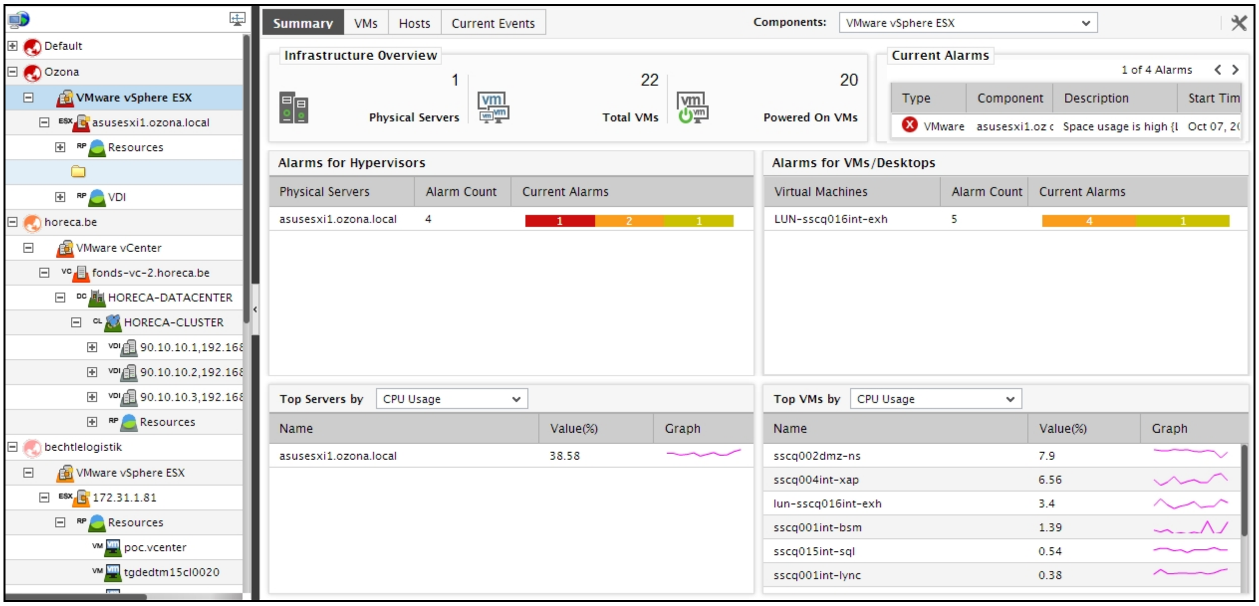

Figure 10.191: The Summary tab page if a folder in the tree is clicked

27. The **Infrastructure Overview** section of Figure 10.192 displays the number of physical servers, VMs, datastores, and datacenters managed by the chosen folder. To know the names of these elements, move your mouse pointer over the corresponding element count. In addition, the section also displays the number of **Current Alarms** related to the elements managed by the folder, and thus enables you to quickly assess the overall health of folder. If you want to take a close look at each of the current alarms, use the **Current Alarms** section, and browse the alarms one after another.

    In order to figure out which physical servers and VMs are responsible for these alarms, use the **Physical Servers** and **Virtual Machines** sections. These sections list the problem-prone physical servers and virtual machines (as the case may be) in the folder, along with the number and type of problems that each physical server/virtual machine is currently experiencing. Move your mouse pointer over an alarm priority corresponding to a physical server/VM to view the details of alarms of that priority that are currently open on that physical server/VM.

    To track the resource usage of physical servers and VMs within a folder and to identify those physical servers and VMs that are consuming resources excessively, you can use the **Top Servers by** and **Top VMs by** sections. By default, both these sections list the top-5 physical servers and VMs (as the case may be) in terms of *Physical CPU utilization*. To identify the top-5 consumers of a different resource, you can select a different measure from the list box available in both the sections. You can even add more measures to these list boxes for selection using the procedure discussed in page 289 and page 290 of this document. Likewise, you can view more number of leading resource consumers in this section by changing the default value 5 to another number using the procedure discussed in page 291 of this document.

28. If a folder node exists, then expanding this node in the tree, will reveal the datacenters included in that folder. If no folders exist, then you would have to expand the vCenter server node in the tree to view the datacenter sub-nodes. If you click on the node representing a datacenter in the tree-structure, then the contents of the **Summary** tab page will change to provide an overview of the composition and health of that datacenter (see Figure 10.192).

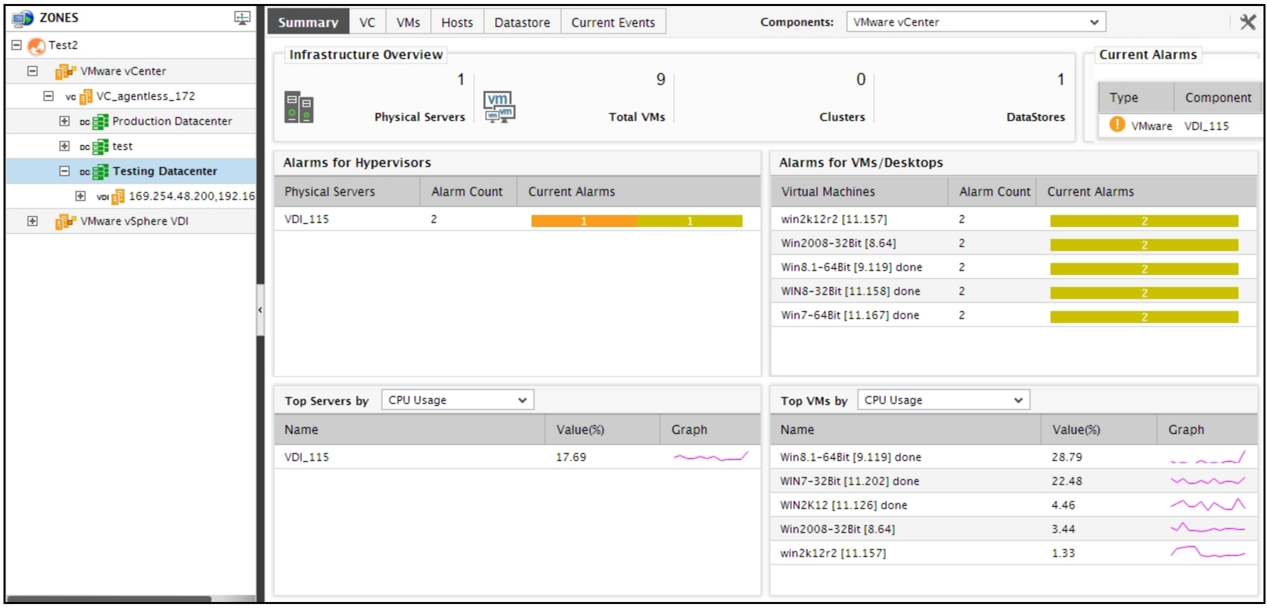

Figure 10.192: The Summary tab page of a datacenter chosen from the tree-structure

29. The **Infrastructure Overview** section of Figure 10.192 displays the number of physical servers, VMs, and datastores managed by the chosen datacenter. To know the names of these elements, move your mouse pointer over the corresponding element count. In addition, the section also displays the number of **Current Alarms** related to the elements managed by the datacenter, and thus enables you to quickly assess the overall health of the datacenter. If you want to take a close look at each of the current alarms, use the **Current Alarms** section, and browse the alarms one after another.

    In order to figure out which physical servers and VMs are responsible for these alarms, use the **Physical Servers** and **Virtual Machines** sections. These sections list the problem-prone physical servers and virtual machines (as the case may be) in the datacenter, along with the number and type of problems that each physical server/virtual machine is currently experiencing. Move your mouse pointer over an alarm priority corresponding to a physical server/VM to view the details of alarms of that priority that are currently open on that physical server/VM.

    To track the resource usage of physical servers and VMs within a datacenter and to identify those physical servers and VMs that are consuming resources excessively, you can use the **Top Servers by** and **Top VMs by** sections. By default, both these sections list the top-5 physical servers and VMs (as the case may be) in terms of *Physical CPU utilization*. To identify the top-5 consumers of a different resource, you can select a different measure from the list box available in both the sections. You can even add more measures to these list boxes for selection using the procedure discussed in page 289 and page 290 of this document. Likewise, you can view more number of leading resource consumers in this section by changing the default value 5 to another number using the procedure discussed in page 291 of this document.

30. To receive an overview of the performance of a cluster within a datacenter, you will have to click on the cluster sub-node under the datacenter node in the tree-structure. The **Summary** tab page will then change as depicted by Figure 10.193.

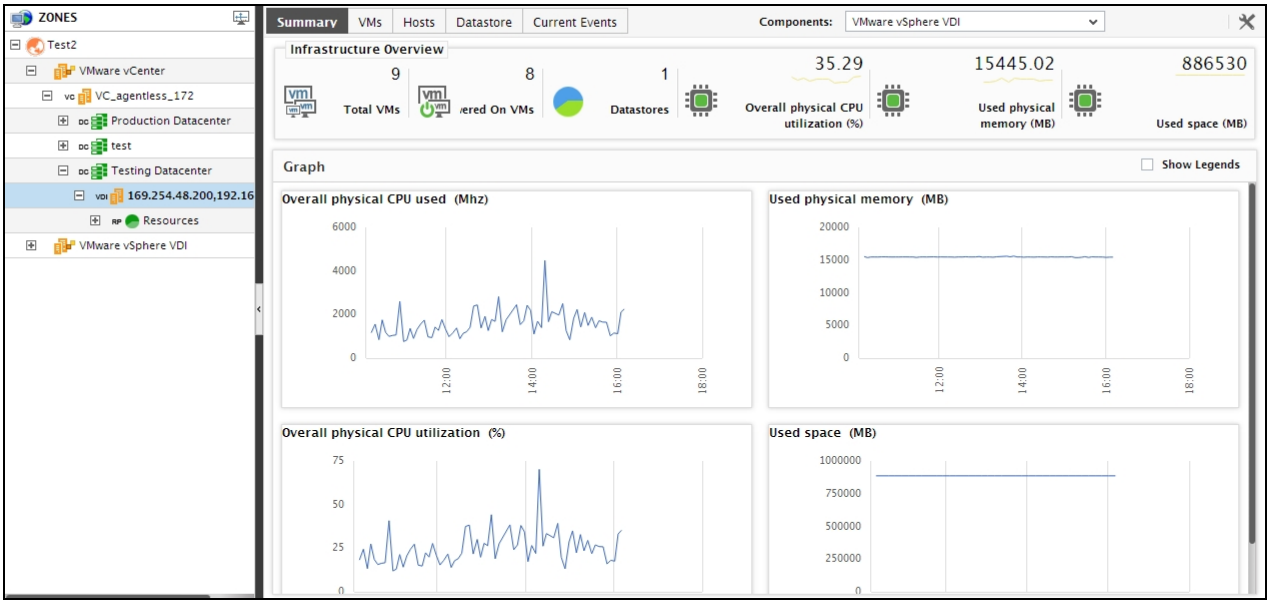

Figure 10.193: The Summary tab page of the cluster sub-node

31. For a particular cluster, the **Summary** tab page will provide a quick summary of the performance of that cluster, so as to enable you to gauge how healthy the cluster is. The **Infrastructure Overview** section of the **Summary** tab page displays the number of physical servers, VMs, powered-on VMs, and resource pools that the chosen cluster consists of. In addition, the **Current Alarms** information in this section displays the number of issues related to the cluster that are still unresolved; this serves as an effective indicator of the overall health of the cluster. Move your mouse pointer over any number in this section to view the corresponding details.

    To be alerted in real-time to abnormalities in resource usage by the cluster, you can use the **Resources** section. This section reports how well the cluster is currently using the physical memory and CPU resources available to it; alongside every usage value displayed in this section, you will find a miniature graph. Click on the graph to expand it, and view the time-of-day variations in resource usage during the last 1 hour (by default). You can even change the **Timeline** of the graph for analyzing usage patterns over a longer period of time. This graph helps you quickly detect disturbing trends in the usage of physical resources, and initiate corrective actions at the earliest.

    The **Graphs** section in the **Summary** tab page displays a series of graphs depicts how the cluster has been using the physical memory and CPU resources available to it over a default period of 6 hours. You can magnify any of these graphs by clicking on it. You can even change the **Timeline** of the expanded graph, to facilitate more effective analysis of the resource usage. Using these graphs, you can determine how resource-intensive the cluster has been.

32. To figure out, from a mere glance, the current state of an ESX host included in a cluster, click on the sub-node representing an ESX host under the cluster node. The **Summary** tab page will then change as depicted by Figure 10.193.

33. For an ESX host, the **Infrastructure Overview** section of Figure 10.195 reveals the number of VMs configured on the ESX host, the number of powered-on VMs, and the number of datastores used by the

ESX host. By moving your mouse pointer over any of these numbers, you can view the corresponding details - i.e., the names of VMs/datastores, as the case may be. This way, you can rapidly identify the VMs that are currently powered-off. Moreover, the section also displays the number of **Current Alarms** related to the ESX host chosen from the tree, and thus enables you to quickly judge the current health of the ESX host. For complete details of the current alarms, move your mouse pointer over the number of **Current Alarms.**

The **Resources** section of Figure 10.193 reports the percentage/amount of physical CPU, memory, and disk resources that the ESX host chosen is currently using. Sudden spikes in resource consumption by the ESX host can be promptly detected by closely observing the change in the resource usage levels reported by this section. Every value displayed in this section is accompanied by a miniature graph, which when clicked, expands to reveal how well the corresponding resource usage metric has performed during the last 1 hour (by default). To analyze the behavior of the said measure over a longer period of time, you can change the **Timeline** of the expanded graph. In the event of excessive resource usage by the ESX host, you can use this graph to figure out when exactly the upward trend in resource consumption began, and then, proceed to investigate the reasons for the same.

In addition to the **Resources** section, a **Graphs** section is also available in the **Summary** tab page. This section provides a series of graphs, which track the physical CPU, memory, and disk resources used by the ESX host during the last 6 hours (by default). Click on a graph in this section to expand it. Using the expanded graph, you can even change the **Timeline** of the graph, so that you can observe usage patterns over longer periods of time.

34. If resource pools are configured on a cluster, then each resource pool will appear as a sub-node of that cluster node. When a resource pool sub-node is clicked, the **Summary** tab will change as depicted by Figure 10.196.

Figure 10.194: The Summary tab page of a resource pool

35. The **Infrastructure Overview** section of Figure 10.194 reveals the number of VMs, powered-on VMs, and current alarms in the resource pool clicked on. Move your mouse pointer over any number in this section to

view the corresponding details. This information enables you to swiftly identify powered-off VMs and also assess the overall health of the resource pool.

The **Resources** section reports the availability and usage of critical resources allocated to the resource pool, in real-time. Any sudden increase in usage of a resource can be promptly detected using the values reported by this section. Adjacent to every value, a miniature graph is provided. To track the usage of a resource over time, click on the miniature graph. The graph expands to reveal how well the resource pool used the corresponding resource during the last 1 hour (by default). You can even change the **Timeline** of the graph to understand resource usage trends over longer time periods.

The **Graphs** section of the tab page provides time-of-day graphs that reveal the variations in the physical CPU and memory usage by the resource pool during the last 6 hours (by default). Analysis of the resource consumption of the pool over time reveals whether the pool has been using resources optimally or inefficiently. Clicking on a graph expands it. You can change the **Timeline** to analyze resource usage over broader time periods.

10.7.2 The VMs Tab Page

The **VMs** tab page provides VM-centric information such as the name of the discovered VMs, the physical server on which each VM executes, and the metrics indicating how every VM uses its allocated and physical resources. Besides revealing resource-hungry VMs, this tab page also brings to light improper resource allocations to VMs.

This tab page, by default, provides the details of those VMs that are executing on the managed *VMware vSphere ESX* servers in your environment. Similarly, the details of the top-10 VMs alone will be displayed in this tab page, by default. These default settings can however be overridden by following steps given below:

- Click on the button at the right, top corner of the **VMs** tab page.
- Figure 10.195 then appears.

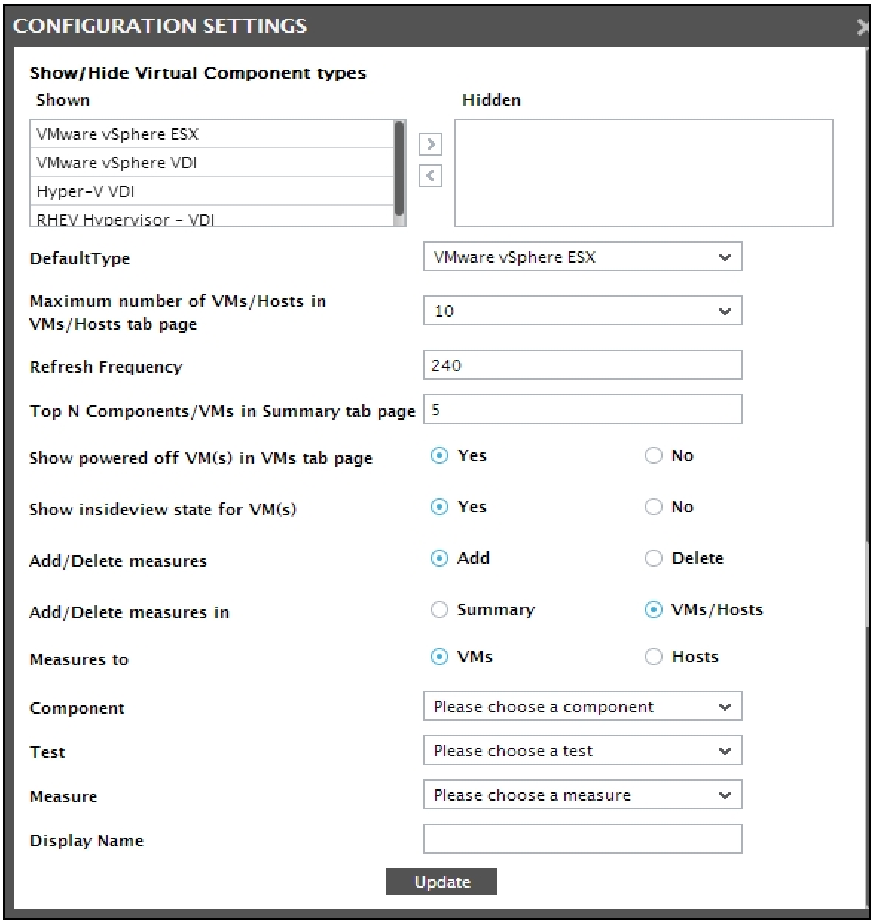

Figure 10.195: Changing the default dashboard configuration

1. By default, the **Maximum number of VMs/Hosts in VMs/Hosts in tab pages** parameter is set to 10, indicating that the top-10 VMs are by default listed in the **VMs** tab page. You can override this default setting by changing the value of this parameter in Figure 10.195.
2. Similarly, you will find that the **DefaultType** list is set to *VMware vSphere ESX* in Figure 10.195. This indicates that, by default, only those VMs that are executing on managed *VMware vSphere ESX* servers in the environment will be listed in the **VMs** tab page. To view details pertaining to the VMs on another virtualized component-type by default, select a different option from the **DefaultType** list.
3. Finally, click the **Update** button in Figure 10.195.

The VMs listed in this tab page change according to the node chosen from the tree-structure in the left panel. This section brings out these differences.

1. If the global **Zones** node is selected in the left panel, then the **VMs** tab page in the right panel will list the top 10 (by default) virtual machines that the eG agent auto-discovers from across all the managed virtual hosts of the chosen type (see Figure 10.196).

   Note:

By default, the VMs displayed in the **VMs** tab page are sorted in the order of their state - i.e., the powered-off VMs will top the list, followed by the powered-on VMs. Sometimes, administrators may want to hide the details of powered-off VMs from the **VMs** tab page, and instead view the resource usage metrics of the powered-on VMs alone. To enable this, follow the steps discussed below:

- Edit the **eg_ui.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
- In the **[VMDASHBOARD_DISPLAY]** section of the file, you will find that the **showPoweredOffVMs** flag is set to **Yes** by default. This indicates that, by default, the **VMs** tab page will display powered-off VMs as well. Set this flag to **No** if you want the **VMs** tab page to only display powered-on VMs.
- Finally, save the file.

2. If state is the same across VMs, then the VMs are arranged in the order of their names. Against every VM listing, the state of the VM (whether powered on or off), the physical server on which the VM executes and the current resource usage metrics pertaining to each VM are displayed, so that administrators will be able to accurately identify powered-off VMs and resource-intensive VMs across all the managed physical servers (of the chosen type) in the environment, from just a quick glance. If required, you can sort the VM listing on the basis of any of the resource usage metrics. To change the sort order on-the-fly, click on the column head that represents the usage metric you want to sort on - for instace, to sort based on the CPU usage of the VMs, click on the column heading **CPU used**.
3. The resource usage metrics that accompany each VM displayed in this tab page are pre-configured in the eG Enterprise system.
4. If need be, you can alter this default measure list, so that more useful measures are displayed per VM or one/more unnecessary measures are removed from the display. To effect this change, follow the steps given below:

- Click on the button at the right, top corner of the **VMs** tab page.
- Figure 10.196 then appears. First, select the **Add** option from the **Add/Delete** section.
- Then, to add a new measure to the **VMs** tab page, select the **VM/Hosts** option from the **Add/Delete measures in** section in Figure 10.196.

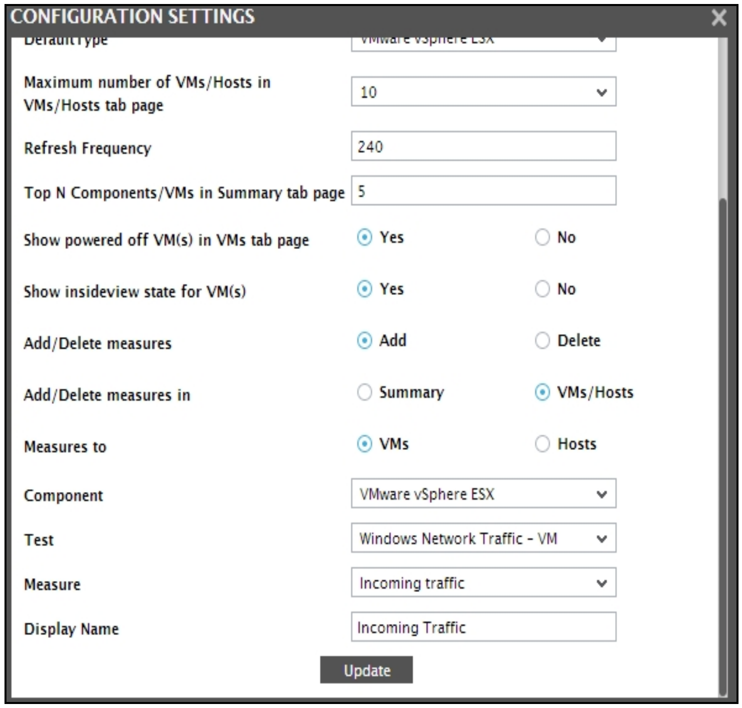

Figure 10.196: Changing the default dashboard configuration

- Then, select **VMs** from the **Measures to** section.
- Then, pick the virtualized **Component** type for which a measure is to be added to the **VMs** tab page.
- Select the **Test** that reports the measure.
- Choose the **Measure**.
- Provide a **Display Name** for the measure.
- Finally, click the **Update** button in Figure 10.196 to save the changes.
- To remove an existing measure from the **VMs** tab page, select the **Delete** option from **Add/Delete**, pick the **VMs/Hosts** option from the **Add/Delete measures in** section, set the **Measures to** flag to **VMs**, select the **Component** type, pick the **Test** that reports the measure, pick the **Measure**, and click the **Update** button.

5. To zoom into the performance of a "powered on" VM, simply click on the VM **Name** in the right panel. This will invoke Figure 10.197 displaying all the performance metrics extracted from that VM in real-time. You are thus enabled to cross-correlate across the various metrics, and quickly detect the root-cause of current/probable disturbances to the internal health of a VM.

   Note:

   If you click on the **Name** of a powered off VM in the **VMs** tab page, then Figure 10.197 will not appear. Instead, you will be lead to the layer model page of the physical server on which the VM executes, which will allow you to verify whether the VM is indeed powered off or not.

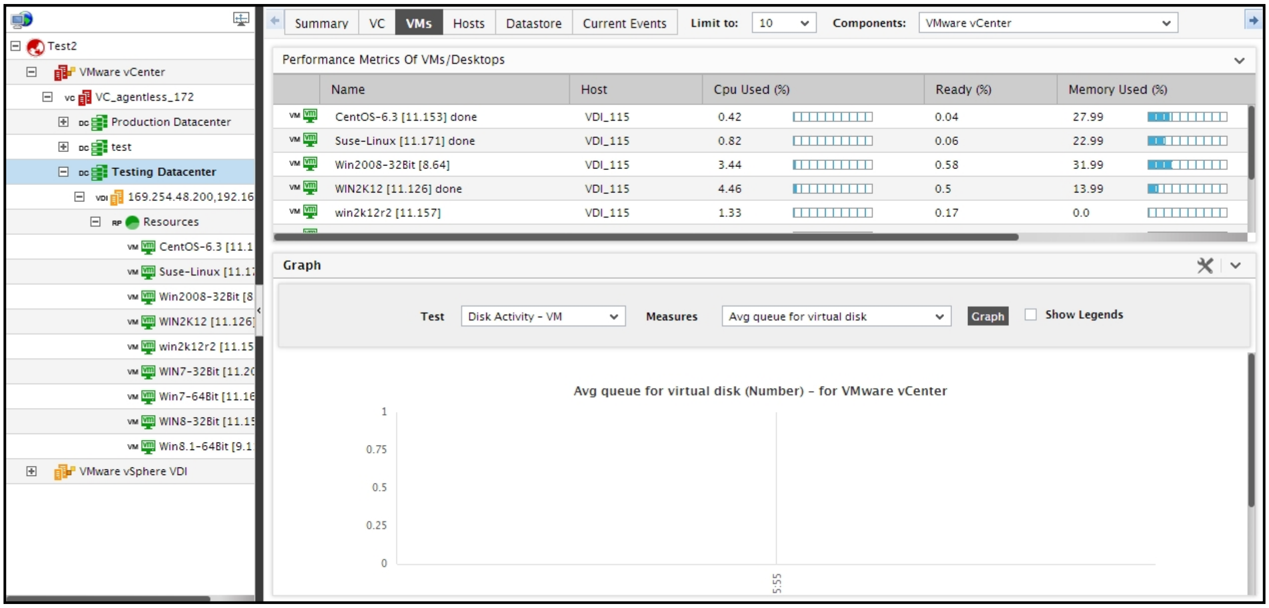

Figure 10.197: The VMs tab page if the global Zones node is clicked

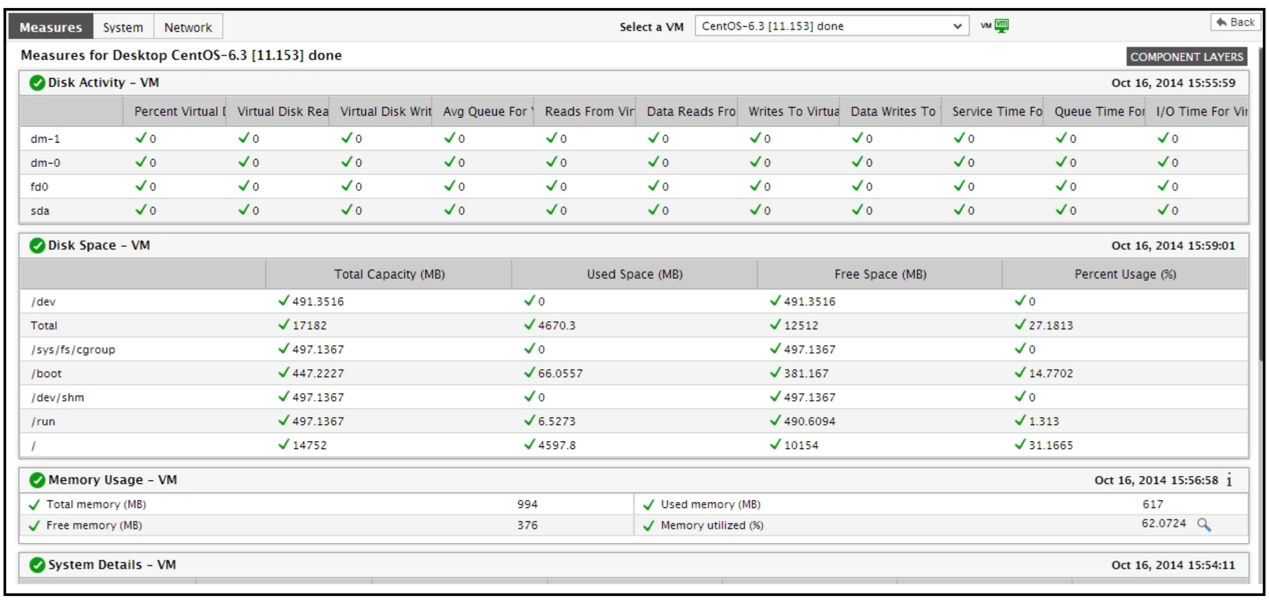

Figure 10.198: Zooming into the performance of a VM from the VM dashboard

6. Below the VMs list, you will find a **Graph** section (see Figure 10.198). This section will be available in every tab page in the right panel of the VM dashboard. For the **VMs** tab, the **Test** list in this section will be populated with all those tests that report metrics for each VM on the virtual hosts of the chosen type. These tests will be sorted in the order of the test names, and the top test in the sorted test list will be selected by default in the **Test** list box. The **Measures** list naturally, will be populated with those metrics that the default **Test** reports. Since this list too is sorted in the order of the measure names, the top measure in the sorted list will be chosen by default in the **Measures** list. Accordingly, for the global **Zones** node in the left panel, the graph that appears in this section traces the variations in the default measure across all VMs (on

virtual hosts of the chosen type) during the default timeline of 1 hour. Using this graph, administrators can compare the performance of a particular measure across VMs, and accurately identify those VMs that are weak in a chosen performance arena. If need be, you can plot a comparison graph for a different **Test-Measure** pair, for a different timeline. To change the timeline, click on **Timeline**; the window depicted by Figure 10.198 will then pop out, allowing you to change the date and time.

7. A legend is provided at the end of every graph clearly indicating which VM is represented using which color in the graph. This is accompanied by the **Avg**, **Max**, and **Min** values that the chosen **Measure** has recorded for every VM during the chosen **Timeline**. To view the legend, you can scroll down the **Graph** section (see Figure 10.199):

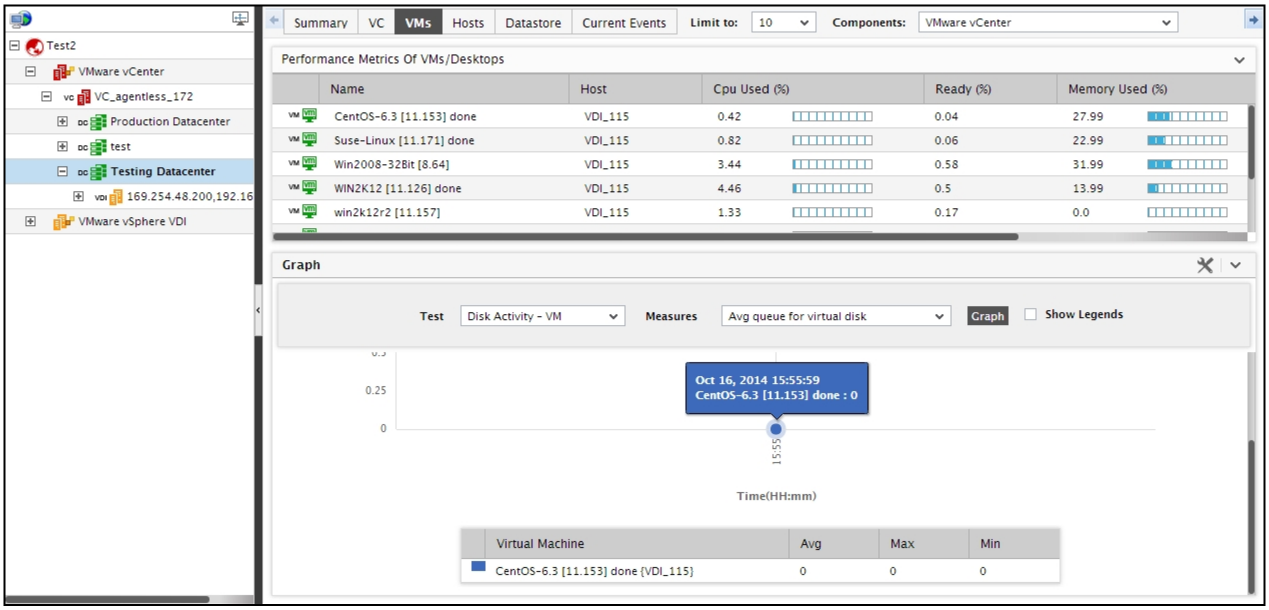

Figure 10.199: Scrolling down the graph to view the legend

8. To view the graph more clearly, you can enlarge it by clicking the button indicated by Figure 10.199 above. The graph then zooms as depicted by Figure 10.200.

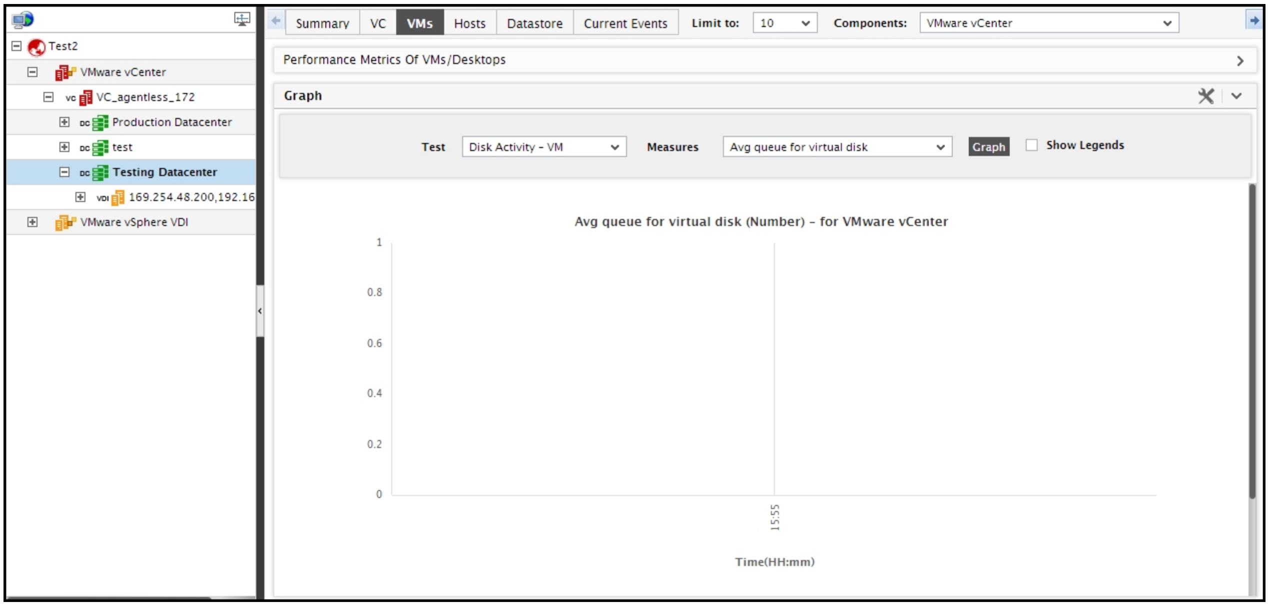

Figure 10.200: Zooming the graph

9. You can even hide the **Graph** section by clicking on the button below the VMs list. The dashboard will then look as depicted by Figure 10.200.

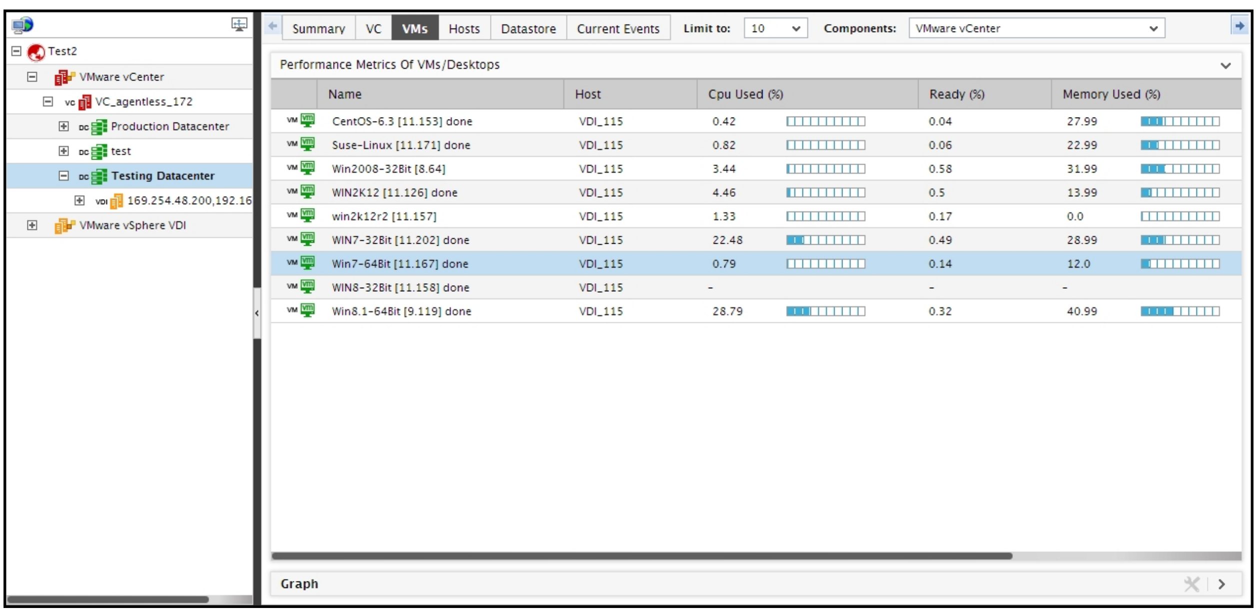

Figure 10.201: The VM dashboard after the graph is hidden

10. You can then click on the button indicated by Figure 10.201 to restore the **Graph** section. Similarly, you can hide the tree-structure by clicking on the button next to the tree (see Figure 10.202). This ensures that the right panel expands and fills the vacuum created by the tree (see Figure 10.202).

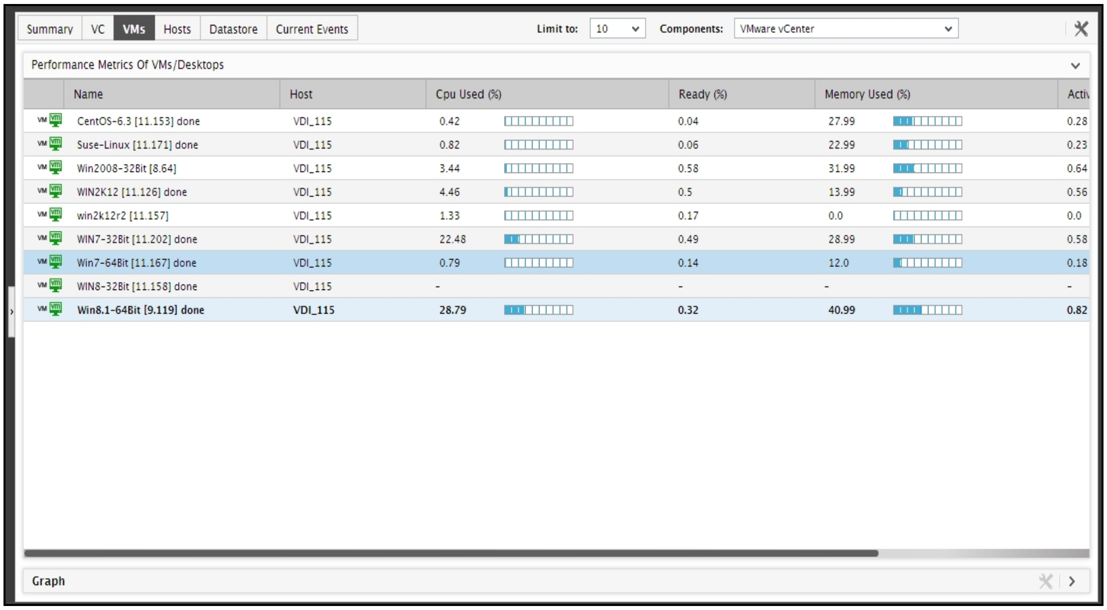

Figure 10.202: The VM dashboard without the tree

11. To restore the tree, click on the button indicated by Figure 10.203.
12. Now, let us see what happens if a particular zone is chosen from the tree-structure in the left panel. When this is done, the **VMs** list in the right panel will change to display the state and resource usage metrics related to the top-10 (by default) VMs that are executing on those physical servers (of the chosen type) that are included in the zone that is clicked on. This information helps administrators analyze how the performance of one/more VMs in a zone impact the performance of the zone as a whole (see Figure 10.204).
13. If you then drill down a particular zone in the tree, you will be able to view the virtual component-types that form part of the zone, and their current state. If you click on a particular component-type in the tree, the VMs tab page in the right panel will allow you to view the state and usage metrics pertaining to the top-10 (by default) VMs executing on the virtual hosts of that type that are included in the corresponding zone (see Figure 10.205).

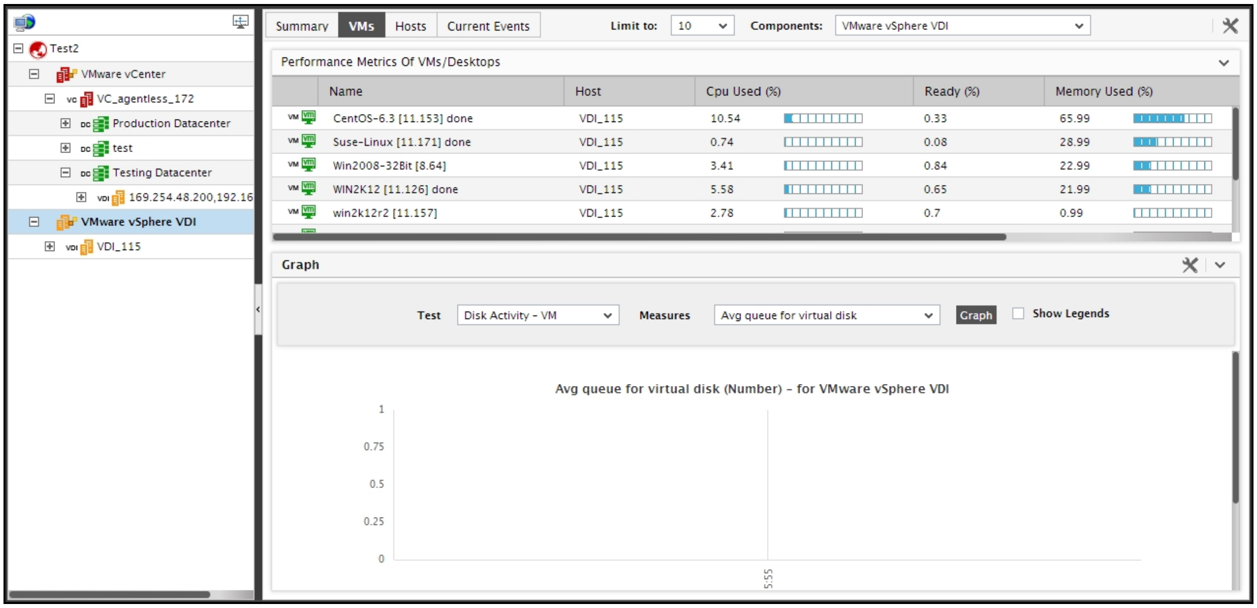

Figure 10.203: The VMs tab page for a virtual component type in the tree

14. If the vCenter servers in your environment are being monitored as part of a zone, then expanding that zone's node in the tree would reveal the *VMware vCenter* component-type. When this component-type is clicked on, the **VMs** tab page will change to display the state and resource usage metrics related to the top-10 (by default) VMs that are executing on the virtual hosts managed by all the vCenter servers included in that zone.

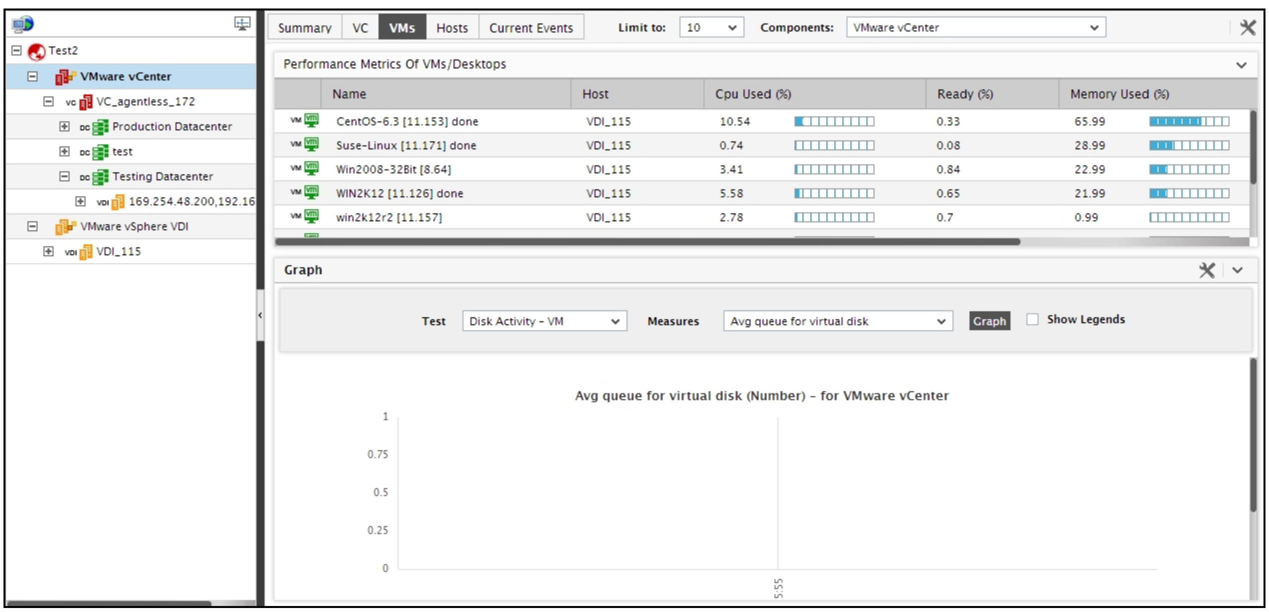

Figure 10.204: The VMs tab page if the 'VMware vCenter' node is clicked

15. To know which vCenter servers have been added to a zone, just expand the *VMware vCenter* sub-node under the zone node in the tree. This will reveal the name and the current state of the vCenter servers in that zone. Clicking on a particular vCenter server in the tree will provide the complete details of the top-10

VMs (by default) executing on the virtual hosts managed by that vCenter component.

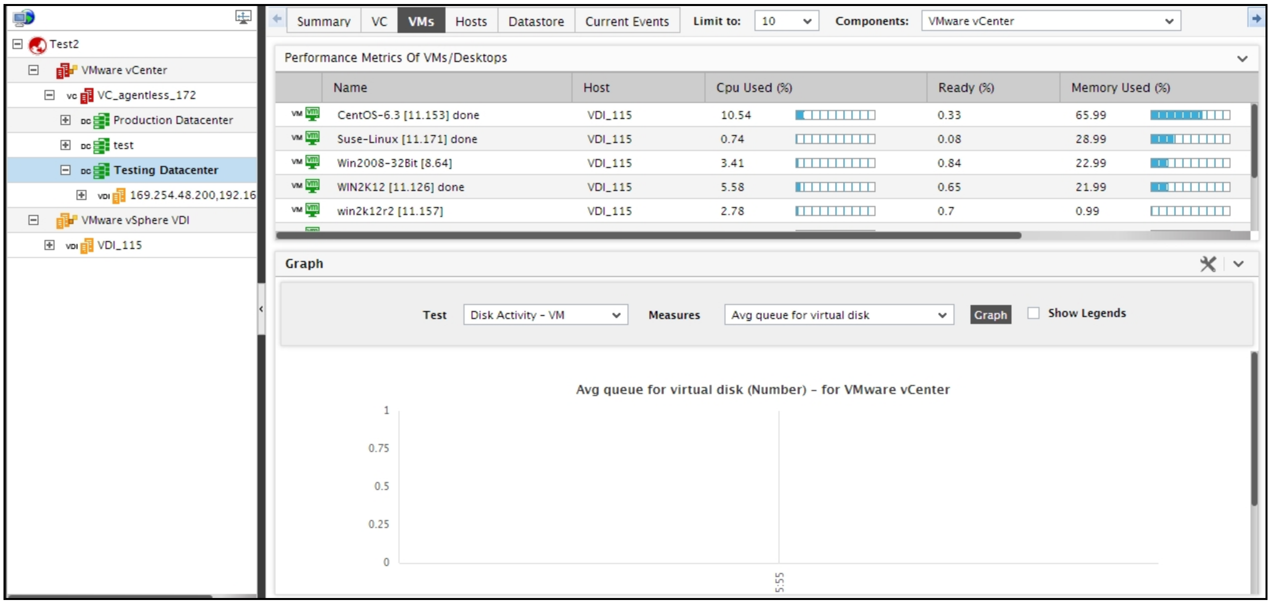

Figure 10.205: The VMs tab page if a particular vCenter server is chosen from the tree

16. If a vCenter server manages the virtualized environment as **folders**, then expanding the vCenter server node in the tree will reveal sub-nodes representing the folders configured on that vCenter server. Click on a folder to view the resource usage metrics related to the VMs executing on the ESX servers included in that folder (see Figure 10.205).

17. Expanding the node representing a vCenter server will reveal the datacenters that have been configured on that vCenter. If the datacenters are included in a folder, then expanding the folder sub-node under the vCenter server node will reveal the datacenters. Click on a datacenter to view the names of the VMs executing on the hosts that reside within that datacenter, and the resource usage of each VM (see Figure 10.206).

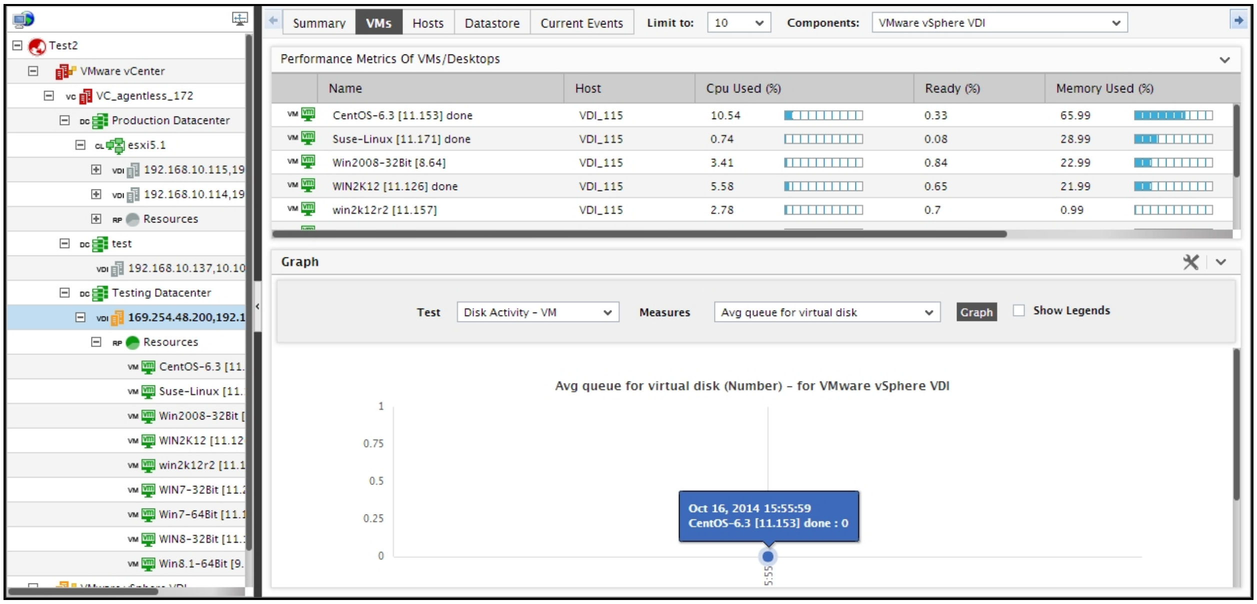

Figure 10.206: The VMs tab page if a datacenter configured on vCenter is clicked

18. Typically, when you expand datacenter in the tree, all the physical ESX servers that are being managed by that vCenter server will appear.

    Note:

    While the vCenter server tree will list even those ESX servers that are not monitored by eG Enterprise, the tree will not indicate the current state of such servers; also, clicking on any such server will not display corresponding performance information in the tab pages in the right panel.
    However, many vCenter installations manage clusters of ESX servers. If such clusters have been configured on any monitored vCenter server, then, in the **Virtual Infrastructure** tree, these clusters will appear as sub-nodes of that datacenter node. If you click on a cluster sub-node in the tree, the **VMs** tab page will reveal the state and performance information pertaining to the top-10 VMs (by default) that are executing on the ESX servers that are part of the cluster clicked on.

19. To view the individual virtual hosts that are part of a zone, do any of the following:

    - Expand the nodes representing the virtual component-types in the zone;
    - If the zone consists of components of type *VCenter*, expand the node representing the monitored vCenter server in your environment;
    - If datacenters are configured on a monitored vCenter server, expand a datacenter sub-node under the vCenter server node;
    - If clusters are configured within a datacenter, expand the cluster sub-node;

20. If you then click on a virtual host in the tree, the **VMs** tab page will change to display the state and measures extracted from the top-10 (by default) virtual machines that are executing on the chosen virtual host alone (see Figure 10.207).

21. Similarly, if you expand the virtual host node in the tree-structure, you can view the name and state of the VMs that are executing on that virtual host. If you now click on a VM in the tree, the metrics extracted from

that VM alone will be displayed in the **VMs** tab page (see Figure 10.207).

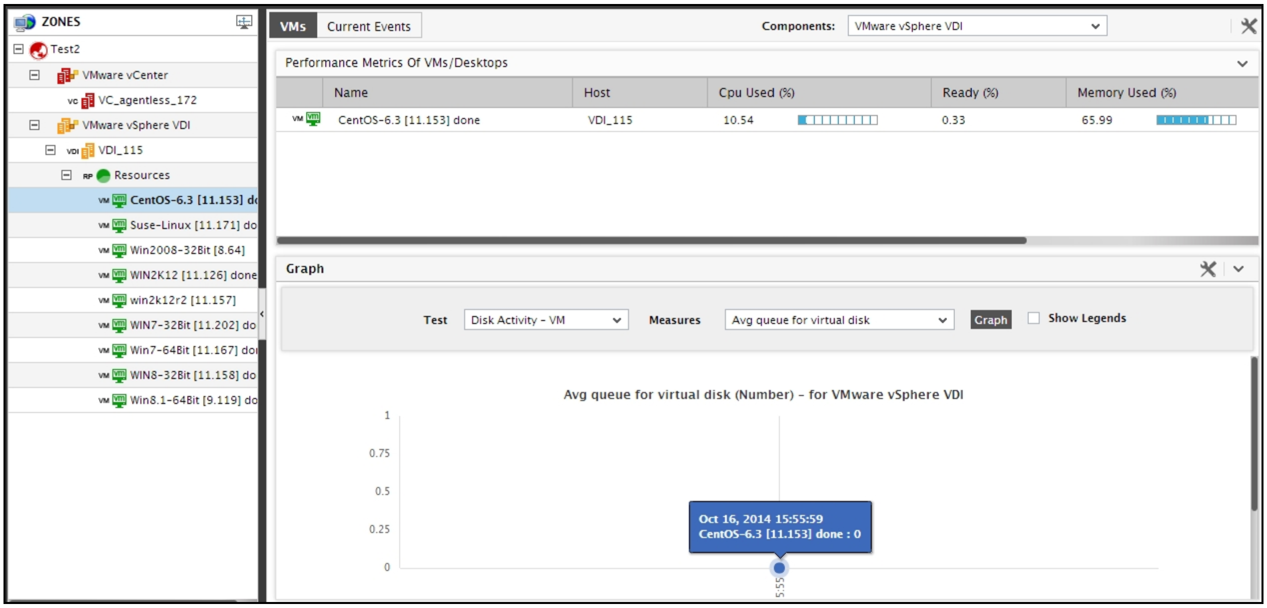

Figure 10.207: The VMs tab page for a particular VM

22. If resource pools are configured on a virtual host, then the resource pools also will appear as the sub-nodes of the virtual host-node. Clicking on a resource pool in the tree will reveal the current state and resource usage metrics related to the top-10 (by default) VMs present in that resource pool, in the **VMs** tab page (see Figure 10.207).

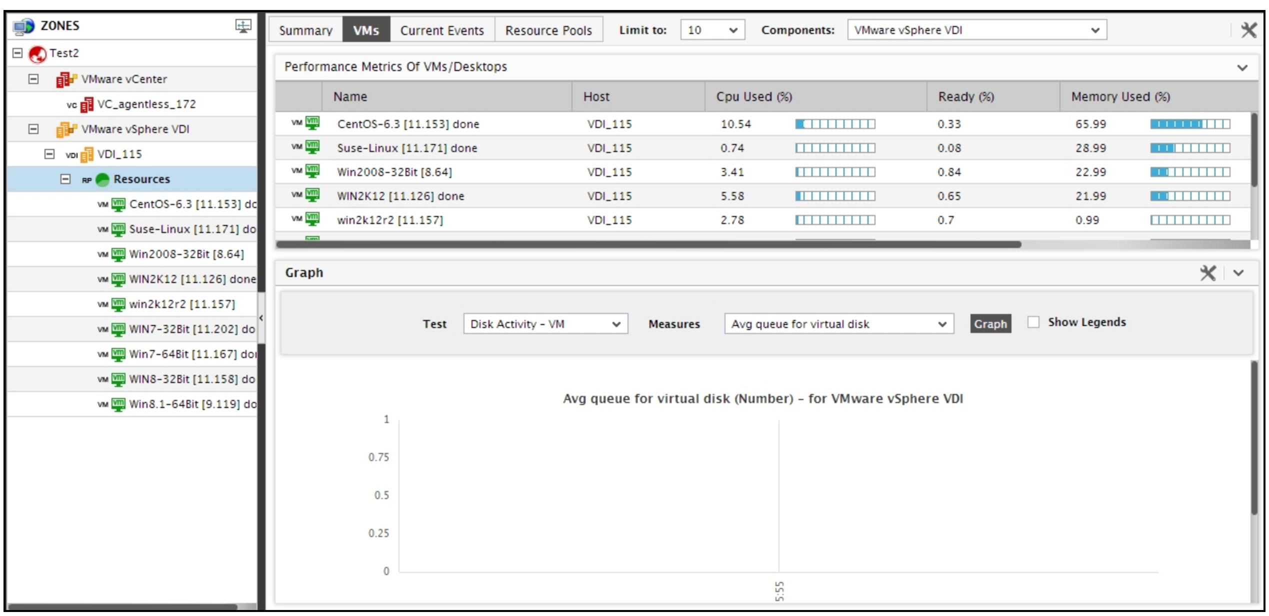

Figure 10.208: The VMs tab page listing the VMs executing within a chosen resource pool

10.7.3 The VC Tab Page

Promptly detect issues with vCenter and swiftly diagnose the root-cause of these issues, using the **VC** tab page offered by the VM dashboard. This tab page lists one/all vCenter servers in the environment (depending upon the node chosen from the tree), and reports the number of ESX servers managed, the resource usage, and the overall health of each vCenter server.

This section discusses how the contents of this tab page change with the node chosen from the tree.

1. If the global **Zones** node is chosen from the tree, then the **VC** tab page will display the names of all vCenter servers included in all configured zones, and will report a default set of metrics per vCenter server, revealing the availability and resource usage of each server. Resource-intensive/unavailable vCenter servers in the target environment can thus be instantly identified.

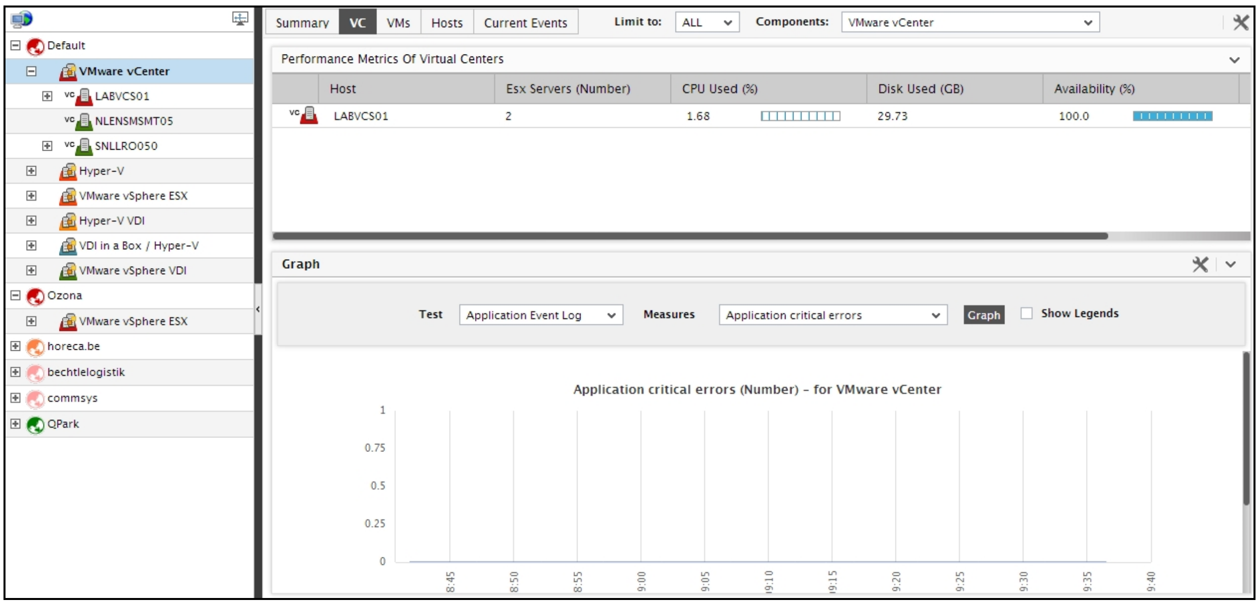

Figure 10.209: The VC tab page if the global Zones node is clicked

2. The default measure list that accompanies every vCenter listed in the **VC** tab page can be modified. For this purpose, follow the steps given below:

- Click on the button at the right, top corner of the tab page. Figure 10.210 will appear.

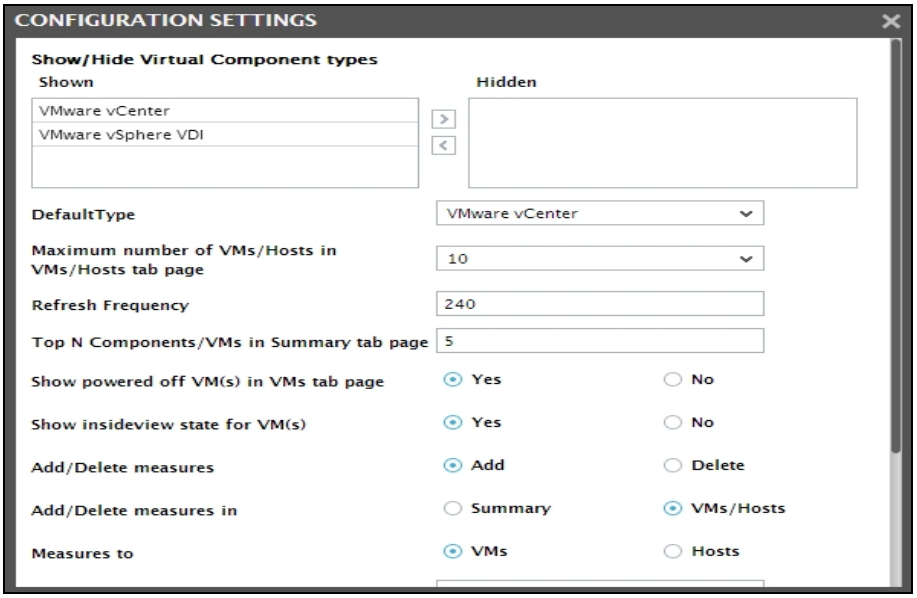

Figure 10.210: Adding a new measure to the VC tab page

- Select **Add** from **Add/Delete measures** section to add a new measure to the **VC** tab page.
- Select **VMs/Hosts** from the **Add/Delete measures in** section in Figure 10.210.
- Set the **Measures to** flag to **VMs**.
- Then, set *VMware vCenter* as the **Component** type. Select the **Test** that reports the measure to be added, and then pick the **Measure**. Provide a **Display Name** for the new measure, and finally, click the **Update** button.
- To remove a measure that pre-exists in the **VC** tab page, select **Delete** from **Add/Delete**, select **VM/Hosts** from the **Add/Delete measures in** section, set the **Measures to** flag to **VM**, select the **Component** type, pick the **Test**, choose the **Measure**, and click the **Update** button.

3. In the event of a slowdown in the performance of a vCenter server, you can click on the name of that server in the **VC** tab page; this will lead you straight to the layer model of the vCenter server, which reveals the exact layer where the problem has originated. Click on the problem layer to view the problem test, and click on the problem test to view the measures (see Figure 10.211).

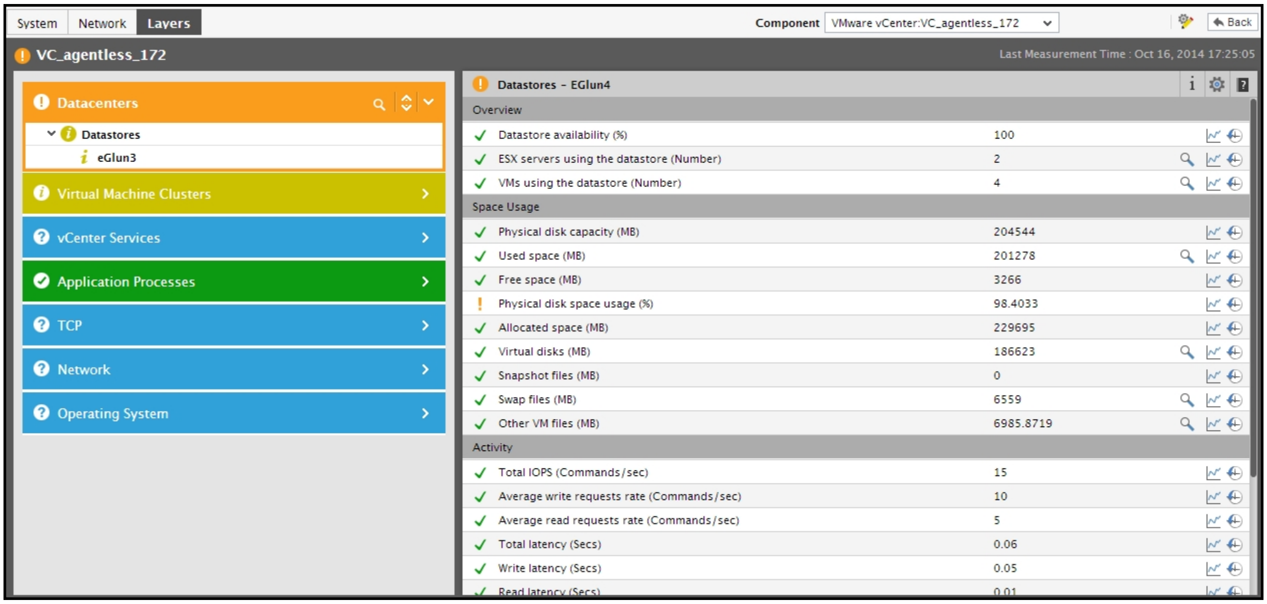

Figure 10.211: The layer model, tests, and measures pertaining to the vCenter server

4. The **VC** tab page also embeds a time-of-day graph, which, by default, reveals the variations in the number of VMs in each resource pool configured within each cluster on vCenter, during the last 1 hour. Accordingly, **ClusterResourcePools** is chosen by default from the **Test** list, the default **Measure** is **VMs in pool**, and the default **Timeline** is 1 hour. You can select a different **Test** and **Measure** combination for the graph, and also define a different **Timeline** for the graph, if required. Using this graph, you can efficiently analyze performance trends and proactively detect performance issues.

5. To view the zones configured with one/more virtual component-types, expand the global **Zones** node in the tree. The individual zones will appear as sub-nodes of the global **Zones** node. Selecting a particular zone will reveal the performance details related to each vCenter server within that zone in the **VC** tab page.

6. Alternatively, you can select the *VMware vCenter* sub-node under a zone node to see how all the vCenter servers in a zone are performing (see Figure 10.212).

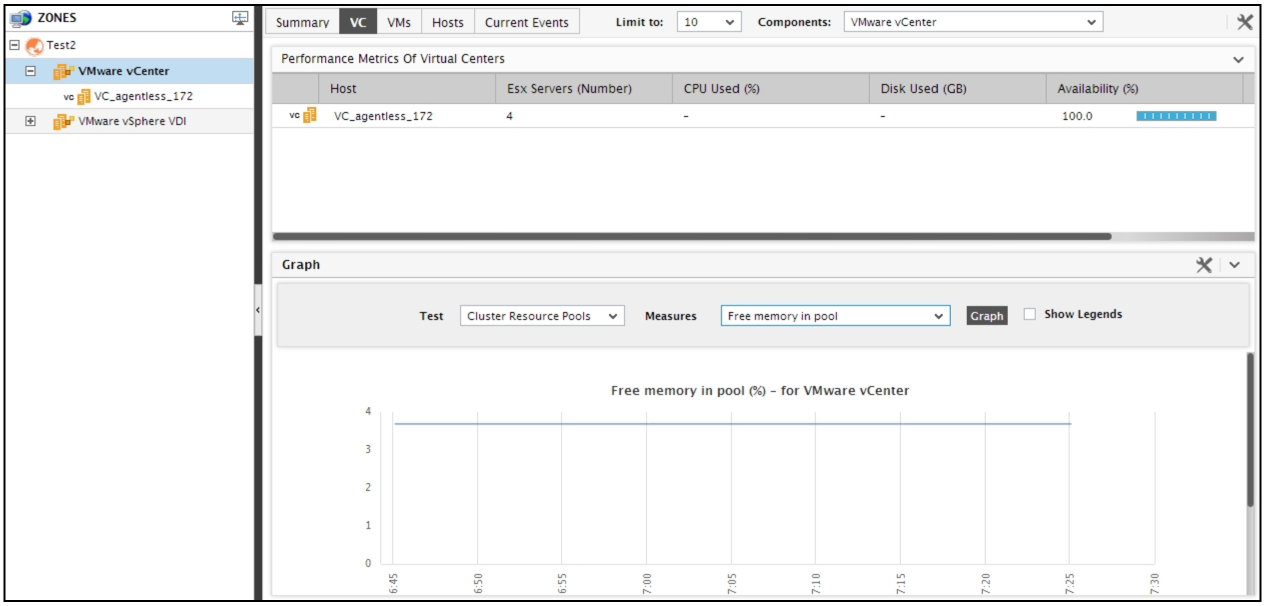

Figure 10.212: The VC tab page if the ‘VMware vCenter’ node is chosen from the tree

7. If you want to check whether a particular vCenter server within a zone is available or not, and if available, how well its using the resources available to it, select the node representing a vCenter server under the *VMware vCenter* node in the tree structure (see Figure 10.213).

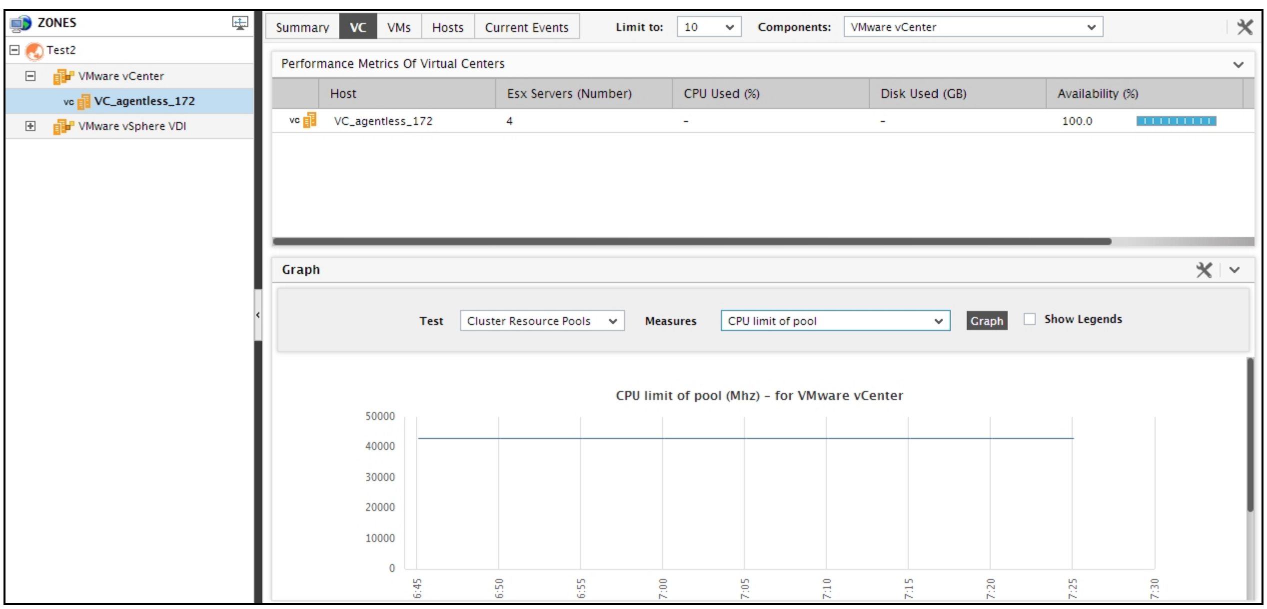

Figure 10.213: The VC tab page for a particular vCenter server chosen from the tree

10.7.4 The Hosts Tab Page

The **Hosts** tab, as mentioned already, provides insights into the current state and the extent to which resources are currently used by the managed virtual hosts in the environment. This tab page sheds light on resource-intensive hosts, and embeds efficient drill downs to discover the underlying cause for the high resource consumption of a host.

As stated earlier, this tab page, by default, provides the details of those managed virtual hosts that are of the type *VMware vSphere ESX*. Similarly, the details of the top-10 virtual hosts alone will be displayed in this tab page, by default. These default settings can however be overridden using the procedure already discussed.

Here again, the hosts displayed depend upon the node chosen from the tree-structure in the left panel. This section explains how the contents of this tab page change with context.

1. If the global **Zones** node is selected in the left panel, then the **Hosts** tab page in the right panel will list the top 10 (by default) virtual hosts that the eG agent auto-discovers from across all the managed virtual hosts of the chosen type (see Figure 10.214). This list is typically sorted by the current state of the hosts. If state is the same across hosts, then the hosts are arranged in the order of their host names. Against every virtual host, the state of the host, the total number of VMs configured on the host, the number of VMs powered-on, and a default set of metrics indicating the extent to which resources are currently utilized by the host, will be displayed. This default measure list can be modified by adding new measures to be displayed in the **Hosts** tab page or by removing one/more existing measures from the tab page. To achieve this, follow the steps given below:
2. Besides revealing the VM load on each virtual host, this tab page also enables administrators to instantly figure out the following:
   - Is any host experiencing performance issues currently?
   - Is any VM on this host currently powered off?
   - Are there any resource-hungry virtual hosts in the environment? If so, which ones are they?

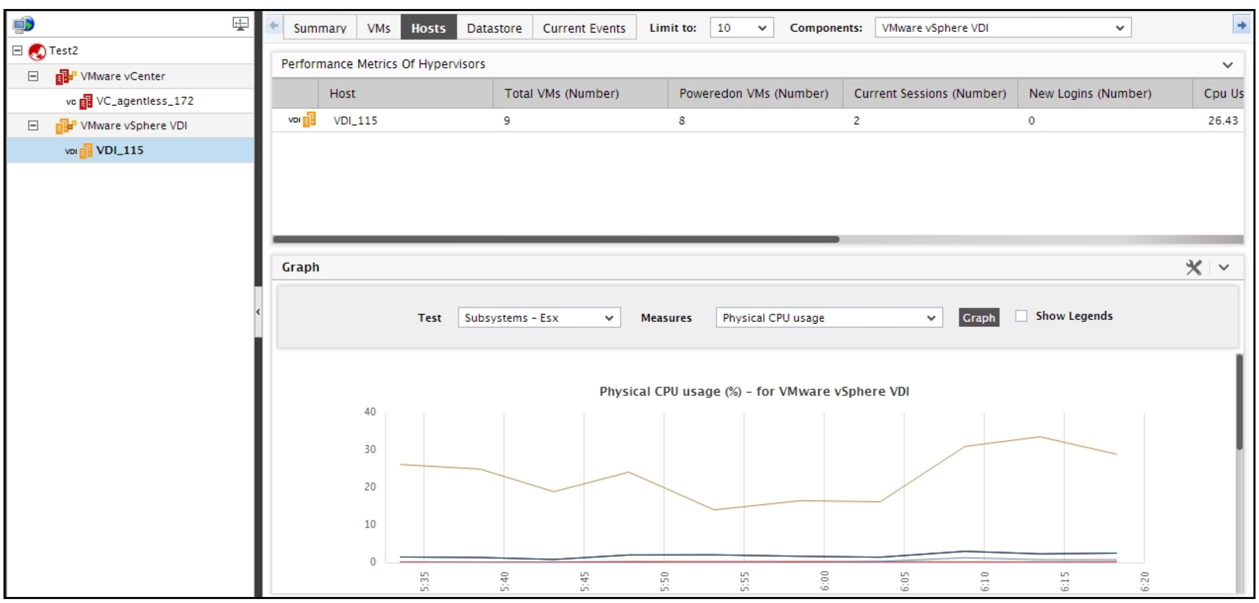

Figure 10.214: Zooming into the performance of a VM from the VM dashboard

3. Below the host list, you will find a **Graph** section (see Figure 10.214). This section will be available in every tab page in the right panel of the VM dashboard. The **Test** list in this section will be populated with all the

tests related to the virtual hosts of the chosen type. **Note that VM-related tests will not be available for selection in this list**. These tests will be sorted in the order of the test names, and the top test in the sorted test list will be selected by default in the **Test** list box. The **Measures** list naturally, will be populated with those metrics that the default **Test** reports. Since this list too is sorted in the order of the measure names, the top measure in the sorted list will be chosen by default in the **Measures** list. Accordingly, for the global **Zones** node in the left panel, the graph that appears in this section traces the variations in the default measure across the top 10 (by default) virtual hosts (of the chosen type) in the environment during the default timeline of 1 hour. Using this graph, administrators can compare the performance of a particular measure across virtual hosts, and accurately identify those virtual hosts that are poor performers. If need be, you can plot a comparison graph for a different **Test-Measure** pair, for a different timeline. To change the timeline, click on **Timeline**; the window depicted by Figure 10.214 will then pop out, allowing you to change the date and time.

4. A legend is provided at the end of every graph clearly indicating which virtual host is represented using which color in the graph. This is accompanied by the **Avg**, **Max**, and **Min** values that the chosen **Measure** has recorded for every virtual host during the chosen **Timeline**. To view the legend, you can scroll down the **Graph** section. If need be, you can hide/unhide the **Graph** section or the **Tree** in the left panel using the buttons provided in the **VM Dashboard**. These buttons have already been described in Section **10.7.1** of this document.

   Note:

   By default, while plotting a graph for a descriptor-based test across virtual hosts, eG Enterprise aggregates the measure values across all descriptors for a host, and plots only a single value for each virtual host. Accordingly, the **AggregateGraphs** flag in the **Configuration Settings** window that appears when the button is clicked, is set to **Yes,** by default. Sometimes, administrators might want the graph to plots values per descriptor. In such a case, set the **AggregateGraphs** flag to **No**.

5. Now, let us see what happens if a particular zone is chosen from the tree-structure in the left panel. When this is done, the **Hosts** list in the right panel will change to display the state and resource usage metrics related to the top 10 virtual hosts that are operating within the zone that is clicked on. This information helps administrators analyze how the performance of one/more virtual hosts in a zone impact the performance of the zone as a whole.

6. If you then drill down a particular zone in the tree, you will be able to view the virtual component-types that form part of the zone, and their current state. If you click on a particular component-type in the tree, the **Hosts** tab page in the right panel will allow you to view the state and usage metrics pertaining to the top-10 (by default) virtual hosts of that type that are included in the corresponding zone (see Figure 10.214).

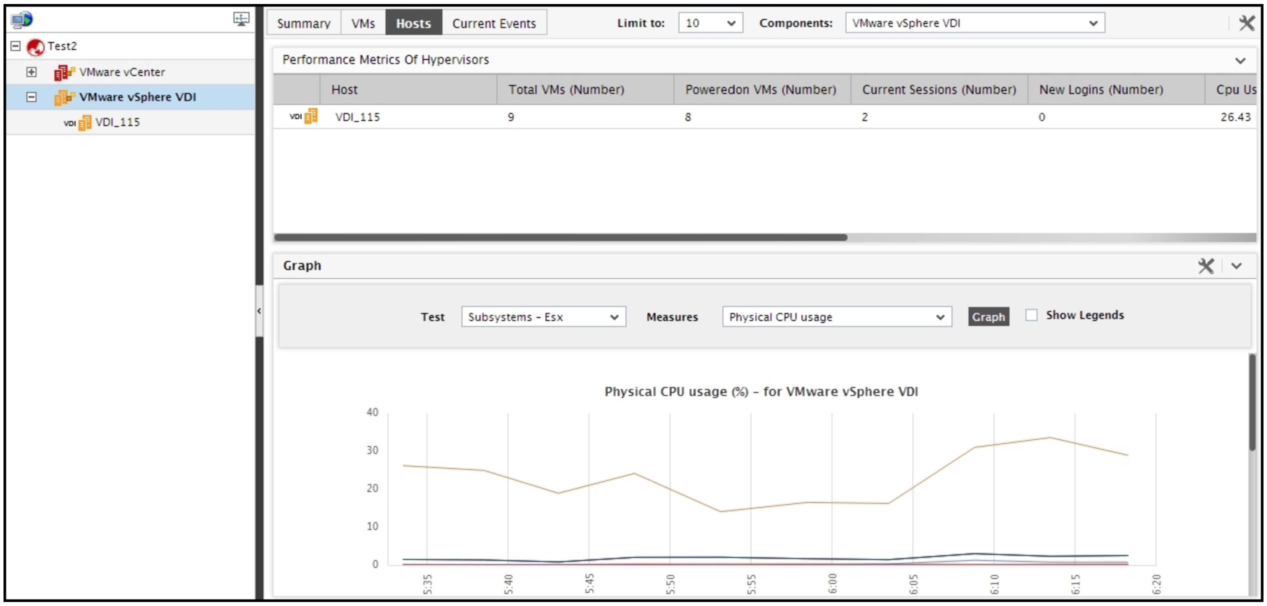

Figure 10.215: The Hosts tab page for a virtual component type in the tree

7. If the vCenter servers in your environment are being monitored as part of a zone, then expanding that zone's node in the tree would reveal the *VMware vCenter* component-type. When this component-type is clicked on, the **Hosts** tab page will change to display the state and resource usage metrics related to the top-10 (by default) virtual hosts managed by all vCenter servers included in that zone.

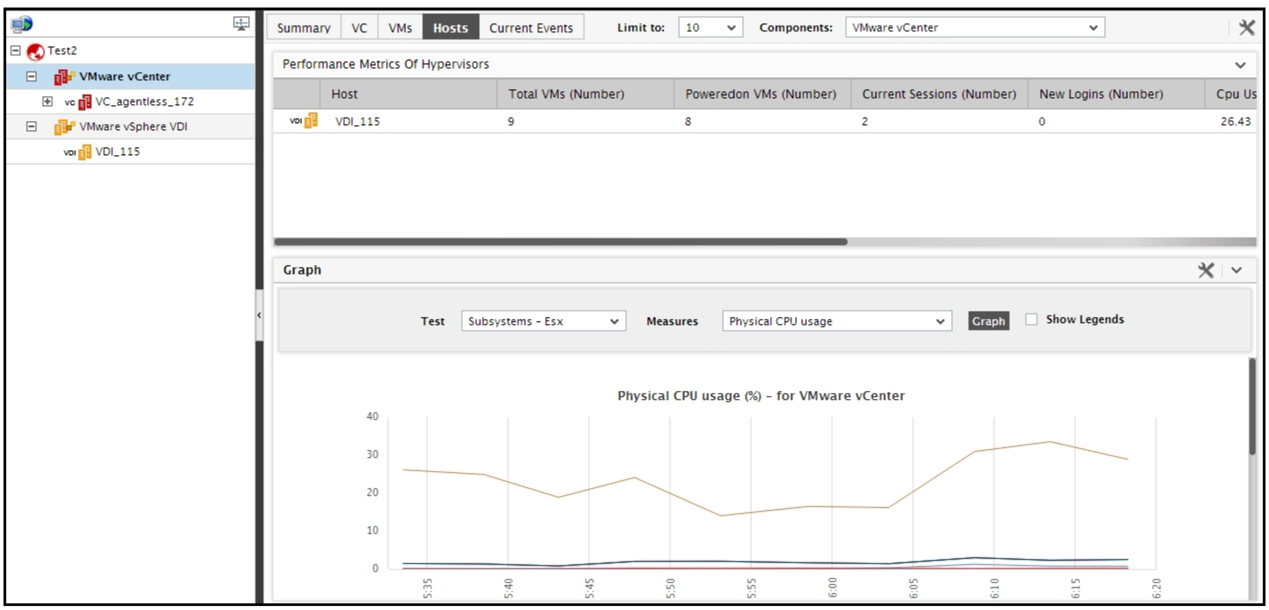

Figure 10.216: The Hosts tab page if the 'VMware vCenter' node is chosen from the tree

8. To know which vCenter servers have been added to a zone, just expand the *VMware vCenter* sub-node under the zone node in the tree. This will reveal the name and the current state of the vCenter servers in that zone. Clicking on a particular vCenter server in the tree will provide the complete details of the top-10 hosts (by default) managed by that vCenter component.

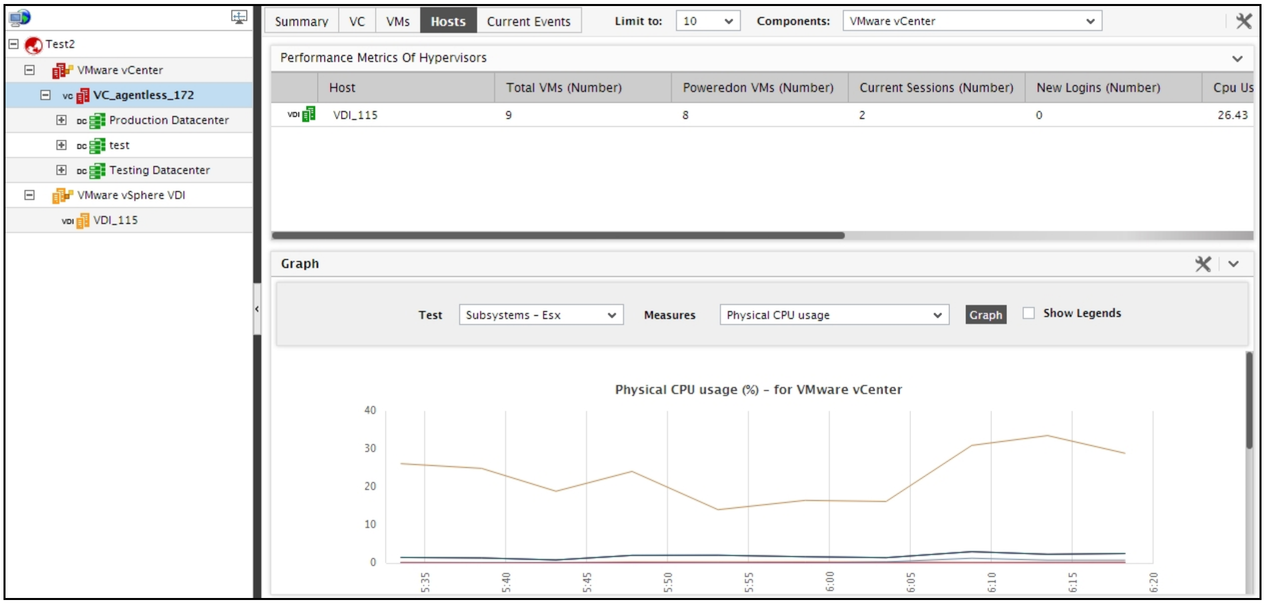

Figure 10.217: The Hosts tab page if a particular vCenter server is chosen from the tree

9. Typically, when you expand a vCenter server node in the tree, the physical ESX servers that are being managed by that vCenter server will appear.

Note:

While the vCenter server tree will list even those ESX servers that are not monitored by eG Enterprise, the tree will not indicate the current state of such servers; also, clicking on any such server will not display corresponding performance information in the tab pages in the right panel.

However, if folders have been configured on the vCenter server, then these folders will appear as sub-nodes of the vCenter server node. If you click on a folder, then the **Hosts** tab page will indicate how well the top-10 hosts (by default) within the folder are performing.

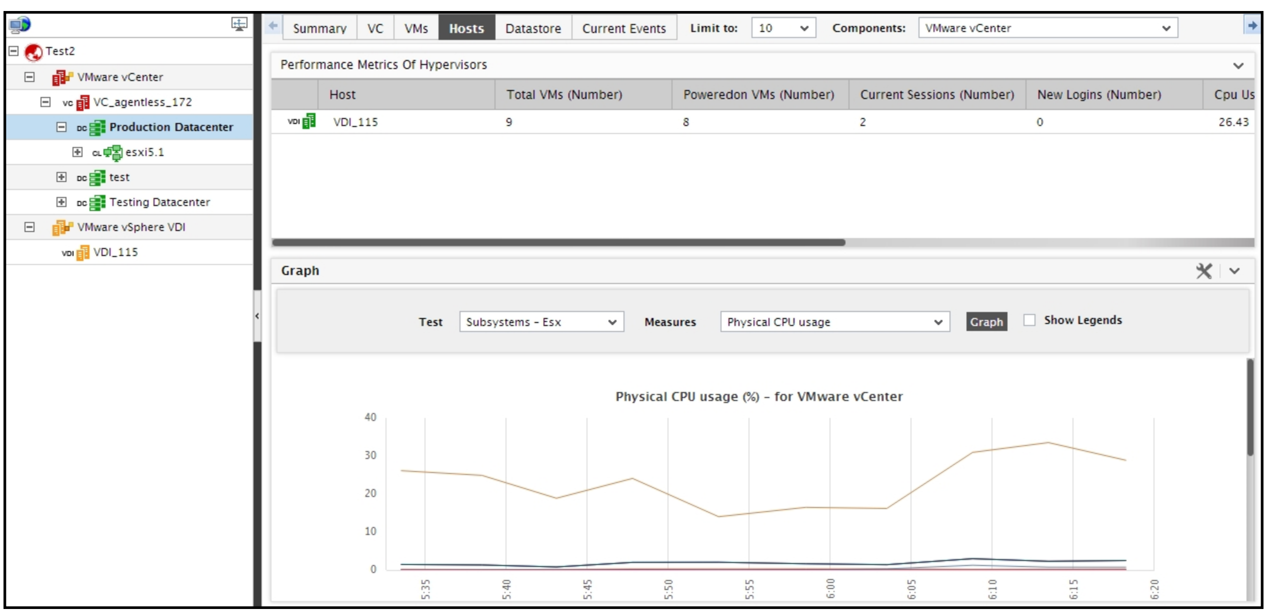

Figure 10.218: The Hosts tab page for a folder selected from the tree

10. If folders exist within a vCenter server, then expanding the folder node reveals sub-nodes representing the datacenters that have been configured on the vCenter server. If no folders exist, then expanding the vCenter server node will reveal the datacenter sub-nodes. If you click on a datacenter, then the **Hosts** tab page will indicate how well the top-10 hosts (by default) within that datacenter are performing (see Figure 10.219).

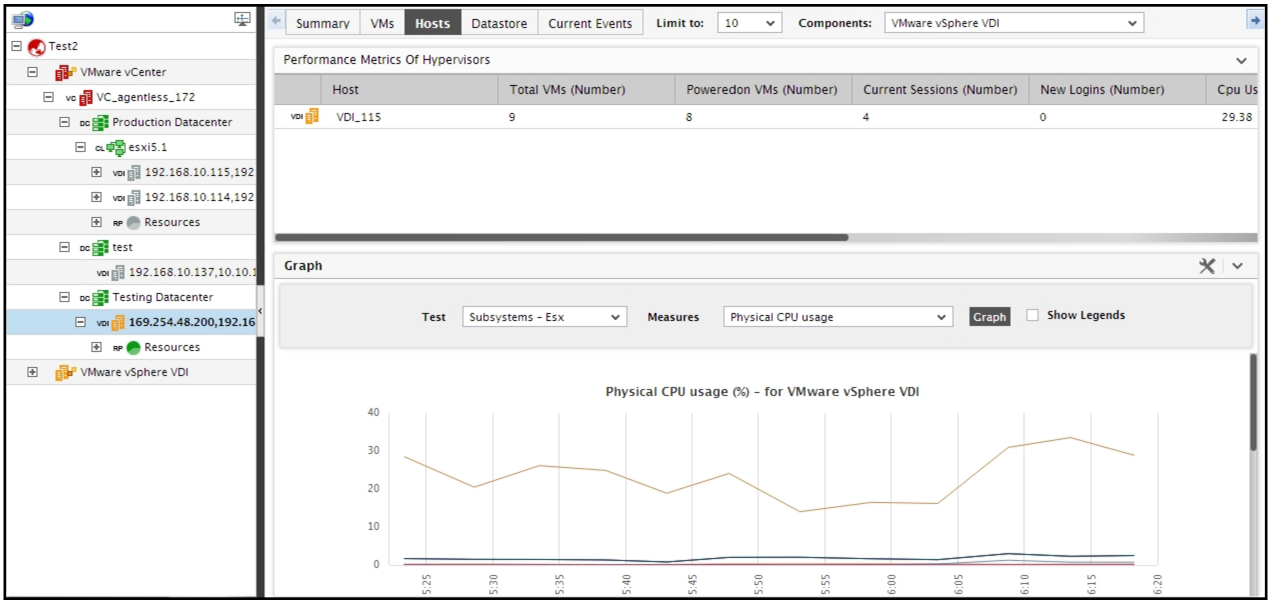

Figure 10.219: The Hosts tab page for a datacenter chosen from the tree

11. Many vCenter installations manage clusters of ESX servers. If such clusters have been configured on any monitored vCenter server, then, in the **Virtual Infrastructure** tree, these clusters will appear as sub-nodes of the datacenter node. If you click on a cluster sub-node in the tree, the **Hosts** tab page will reveal the state and performance information pertaining to the top-10 hosts (by default) that are part of the cluster clicked on.

12. To view the individual virtual hosts that are part of a zone, do any of the following:

    - Expand the nodes representing the virtual component-types in the zone;
    - If the zone consists of components of type *VMware vCenter*, expand the node representing the monitored vCenter server in your environment;
    - If datacenters are configured on a monitored vCenter server, expand a datacenter sub-node under the vCenter server node;
    - If clusters are configured within a datacenter, expand the cluster sub-node;

    If you then click on a virtual host in the tree, the **Hosts** tab page will change to display the state and measures extracted from the chosen virtual host alone (see Figure 10.220).

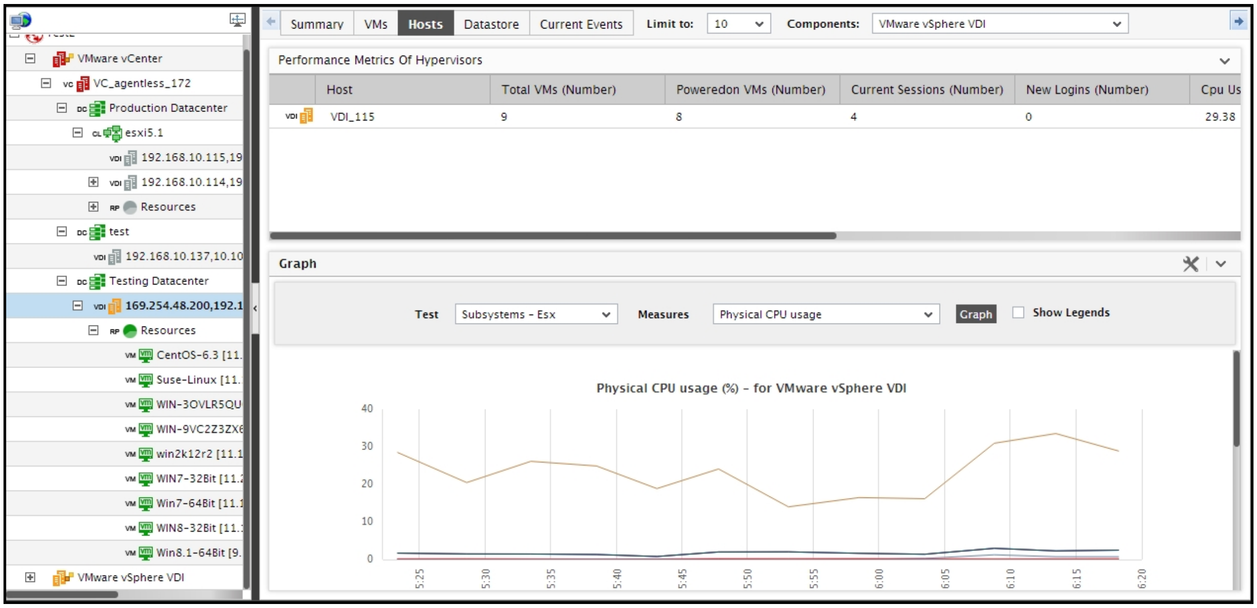

Figure 10.220: The VMs tab page for a particular virtual host selected from the tree

13. If a virtual host in the **Hosts** tab page is found to be in a critical state, then to zoom into the problems affecting the health of that virtual host, simply click on it. Figure 10.221 then appears revealing the layer model of the virtual host, and clearly indicating the problem layer. While clicking on the problem layer will reveal the problem test, a click away from the problem test is the problem measure, which sheds light on the root-cause of the problem with the virtual host (see Figure 10.221). To return to the VM dashboard, just click on the **Back to Virtual Dashboard** link at the right, top corner of the layer model page.

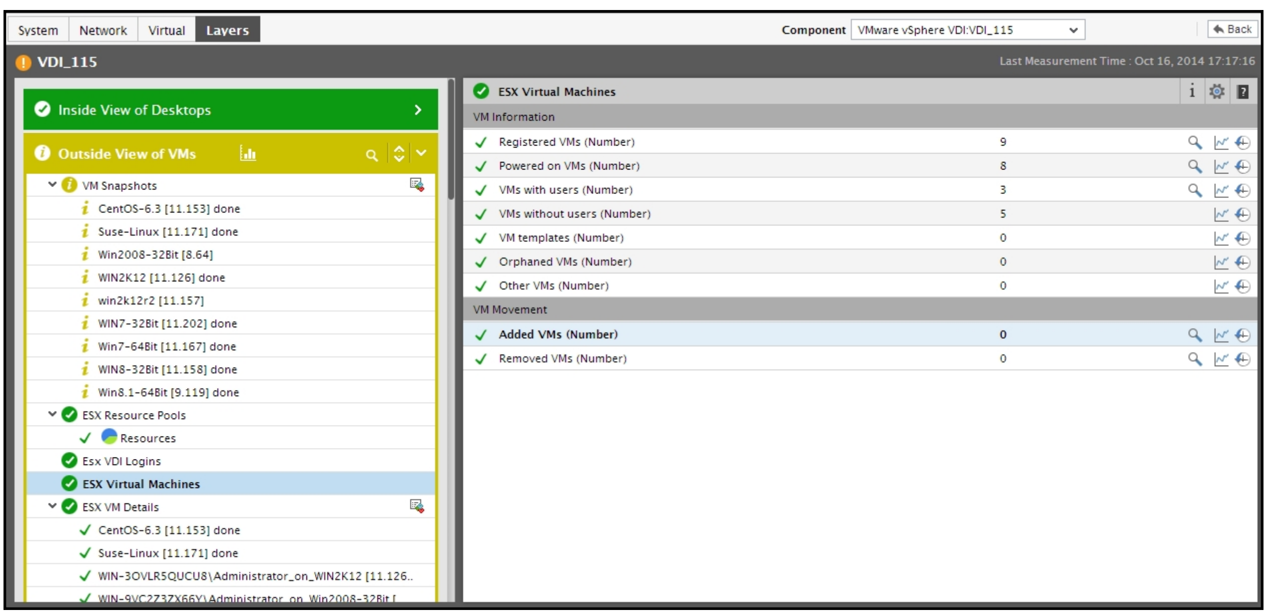

Figure 10.221: The layer model of the virtual host

10.7.5 The Datastores Tab Page

The **Datastores** tab page lists the datastores used by a particular ESX host or one/more ESX hosts within a datacenter, and reports their availability and usage, so that unavailable datastores and those that are currently running out of space can be accurately identified, and the hosts affected by this anomaly can be easily isolated.

This section discusses how the contents of this tab page change according to the node chosen from the tree-structure.

1. If the node representing a particular folder is chosen from the tree, then the **Datastores** tab page will display the top-10 (by default) datastores used by the ESX hosts managed by that folder. Against every datastore displayed here, the availability and space usage metrics pertaining to that datastore will be provided, along with the number of VMs and ESX servers that have been using the datastore. Using this information, you can swiftly isolate unavailable or over-utilized datastores, and the number of ESX servers and VMs that have been impacted by the anomaly (see Figure 10.222).

2. If the node representing a particular datacenter is chosen from the tree, then the **Datastores** tab page will display the top-10 (by default) datastores used by the ESX hosts managed by that datacenter. Against every datastore displayed here, the availability and space usage metrics pertaining to that datastore will be provided, along with the number of VMs and ESX servers that have been using the datastore. Using this information, you can swiftly isolate unavailable or over-utilized datastores, and the number of ESX servers and VMs that have been impacted by the anomaly.

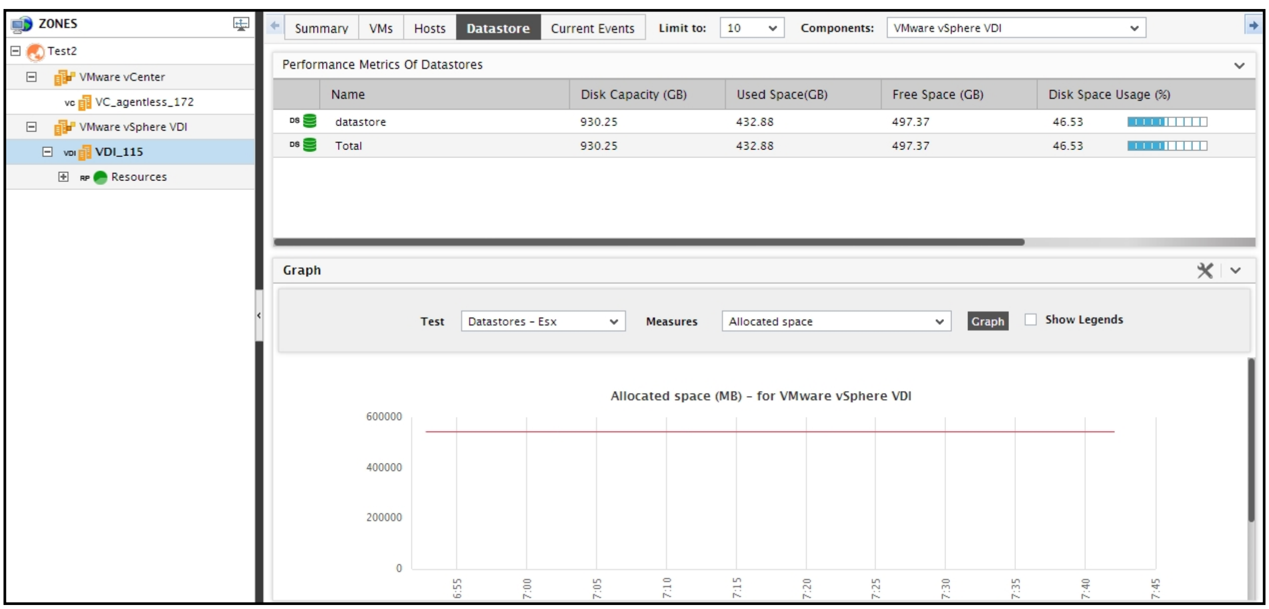

Figure 10.222: The Datastores tab if a particular datacenter is chosen from the tree

3. In addition to listing datastores, the **Datastores** tab page provides a time-of-day graph, which, by default, reveals the number of ESX servers that have been using each datastore during the last 1 hour. Accordingly, the default selection in the **Measure** list is *ESX servers using the datastore*, and the default **Timeline** is 1 hour. Typically, the **Measure** list contains all the measures reported by the default **Test**, which is the *Datastores* test. The options available in the **Measure** list box are sorted in the ascending order of the

measure names, and, by default, the first measure in the sorted list will be displayed as the default **Measure**. The default graph clearly indicates how workload on a datastore has varied during the last 1 hour. You can plot the graph for a different measure or a timeline by choosing a different option from the **Measure** list, and by altering the timeline for the graph by clicking the right-arrow button that prefixes **Timeline**.

4. If you click on a datastore in the **Datastores** tab of Figure 10.222, Figure 10.223 will appear leading you straight to the layer model of the *VMware vCenter* server, which manages that datastore. In the event of the non-availability or excessive usage of a datastore, you can use this model to instantly identify the layer that has been affected by the problem with the datastore. Click on the problem layer to view the problem test, and then, click on the problem test to view the problem measure(s). This way, you can easily understand the nature of the problem with the datastore, and how it has impacted the state of the vCenter server.

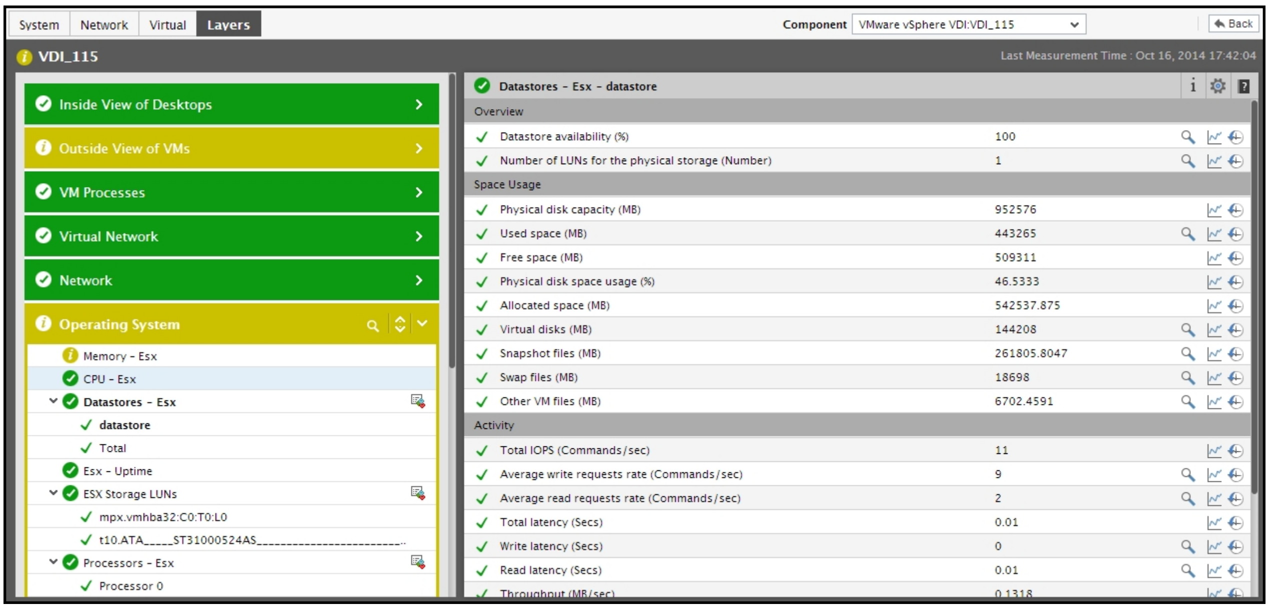

Figure 10.223: The layer model of the vCenter server that appears when a datastore is clicked on in the Datastores tab page

5. Expanding the datacenter node in the tree will enable you to view the ESX hosts that are managed by that datacenter. If a sub-node representing an ESX host is clicked in the tree, then the **Datastores** tab page will only display those datastores that are currently in use by the ESX host and the VMs on it.

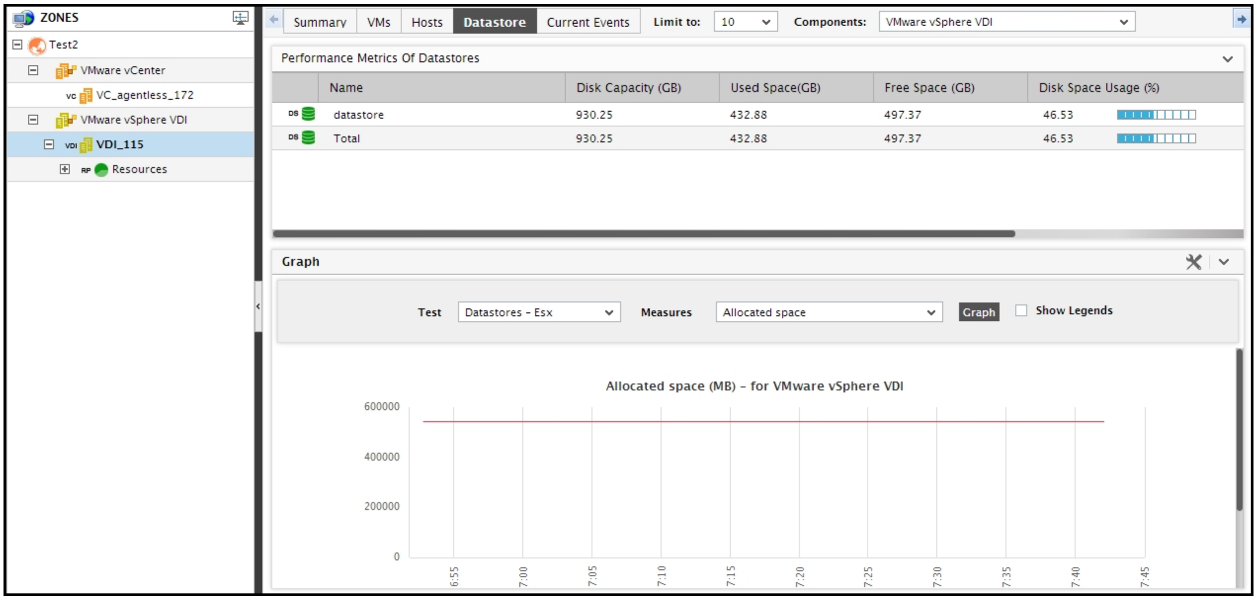

Figure 10.224: The Datastores tab if a particular ESX host is selected from the tree

6. Here again, a time-of-day graph is available, but this graph, by default, reveals the physical disk capacity of each datastore during the last 1 hour. Accordingly, the default selection in the **Measure** list is *Physical disk capacity*, and the default **Timeline** is 1 hour. Typically, the **Measure** list contains all the measures reported by the default **Test**, which is the *Datastores-Esx* test. The options available in the **Measure** list box are sorted in the ascending order of the measure names, and, by default, the first measure in the sorted list will be displayed as the default **Measure**. You can plot the graph for a different measure or a timeline by choosing a different option from the **Measure** list, and by altering the timeline for the graph by clicking the right-arrow button that prefixes **Timeline**.

10.7.6 The Current Events Tab Page

The **Current Events** tab page lists the problems that are currently affecting the performance of managed virtual hosts and virtual machines executing on them. This tab page enables administrators to focus on issues related to their virtualized environment alone, without being distracted by the “non-virtual” issues.

As stated earlier, this tab page, by default, lists the alarms pertaining to those virtual hosts that are of type *VMware vSphere ESX*.

Like the other tab pages, the **Current Events** tab page too changes with respect to the node chosen from the tree-structure. To see how, read on.

1. If the global **Zones** node is selected in the left panel, then the **Current Events** tab page in the right panel will list the problems that are adversely impacting the performance of the virtual hosts of the chosen type across the environment (see Figure 10.225). This list is typically sorted by event priority. If event priority is the same across events, then the events are arranged in the order of the names of the problem hosts. The details provided here include the name of the problem component, a brief description of the problem, and the time at which the problem was reported. This information enables administrators to understand how problem-prone their virtualized environment is, and also provides them with pointers to the root-cause of

the problem.

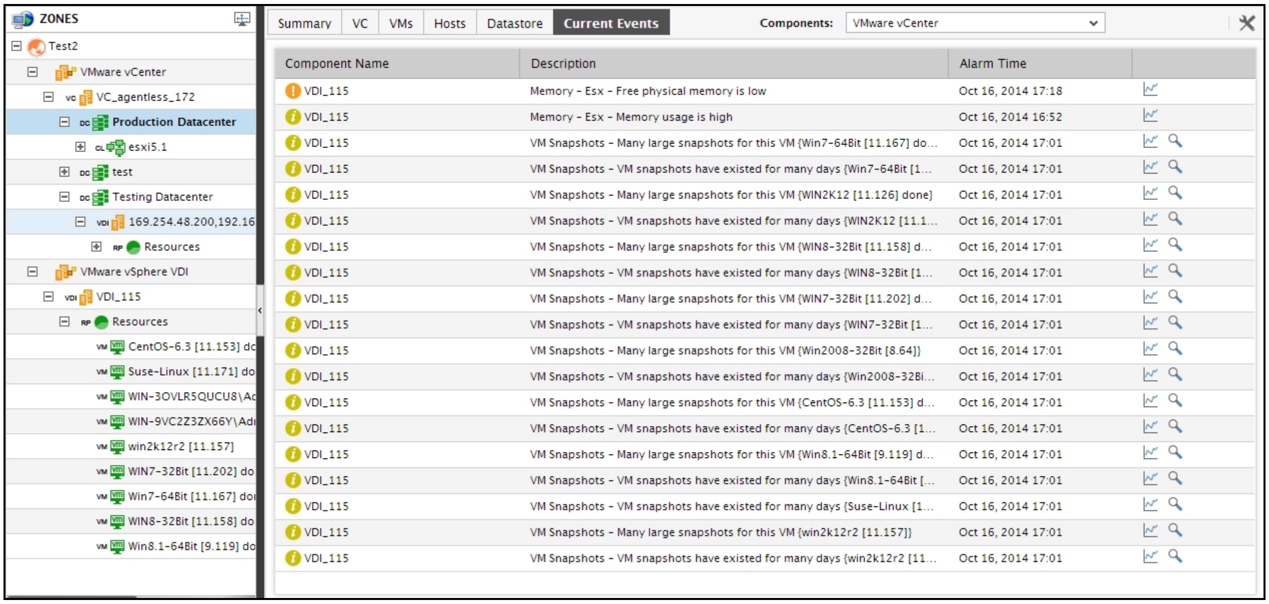

Figure 10.225: The Current Events tab for the global Zones node

2. Now, let us see what happens if a particular zone is chosen from the tree-structure in the left panel. When this is done, the **Current Events** list in the right panel will change to display the problems related to those virtual hosts that are included in the zone that is clicked on. This information helps administrators identify those problems that are affecting the performance of a particular zone.

3. If you then drill down a particular zone in the tree, you will be able to view the virtual component-types that form part of the zone, and their current state. If you click on a particular component-type in the tree, the **Current Events** tab page in the right panel will allow you to view the problems related to virtual hosts of that type that are included in the corresponding zone (see Figure 10.226).

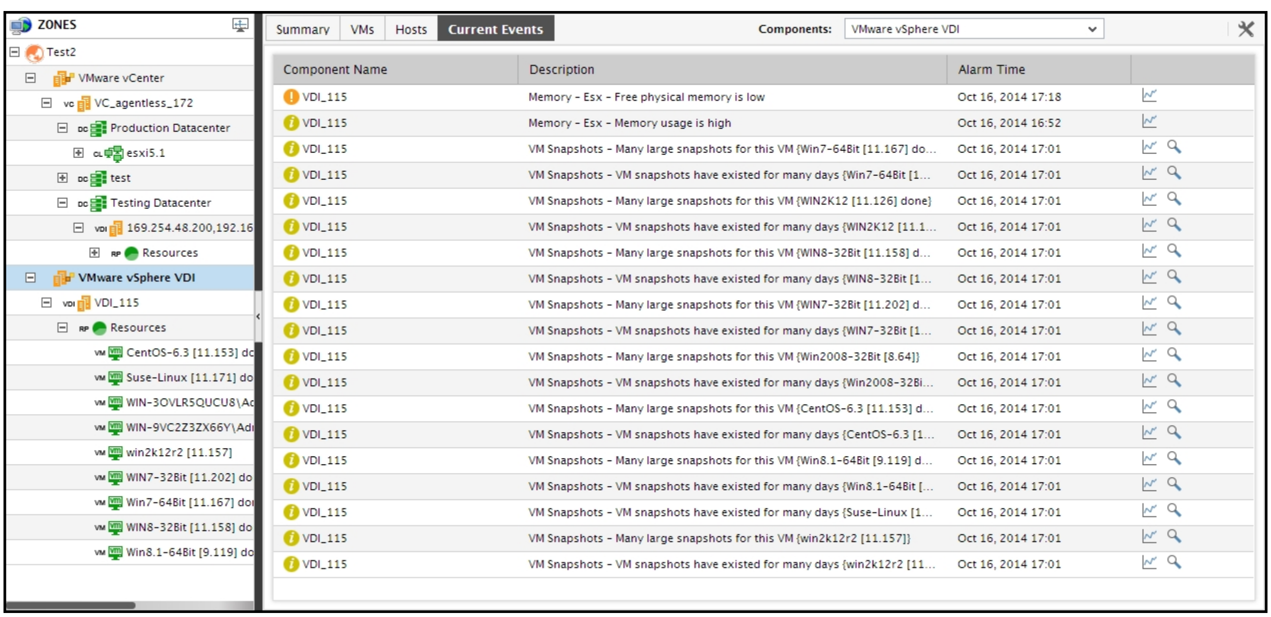

Figure 10.226: The Events tab page for a virtual component type in the tree

4. If the vCenter servers in your environment are being monitored as part of a zone, then expanding that zone's node in the tree would reveal the *VMware vCenter* component-type. When this component-type is clicked on, the **Current Events** tab page will change to display the problem events pertaining all the vCenter servers included in that zone, and those that correspond to the ESX servers managed and VMs by every vCenter server in the zone.

5. To know which vCenter servers have been added to a zone, just expand the *VMware vCenter* sub-node under the zone node in the tree. This will reveal the name and the current state of the vCenter servers in that zone. Clicking on a particular vCenter server in the tree will display problem events pertaining to that vCenter component and those that relate to ESX servers and VMs managed by that vCenter, in the **Current Events** tab page.

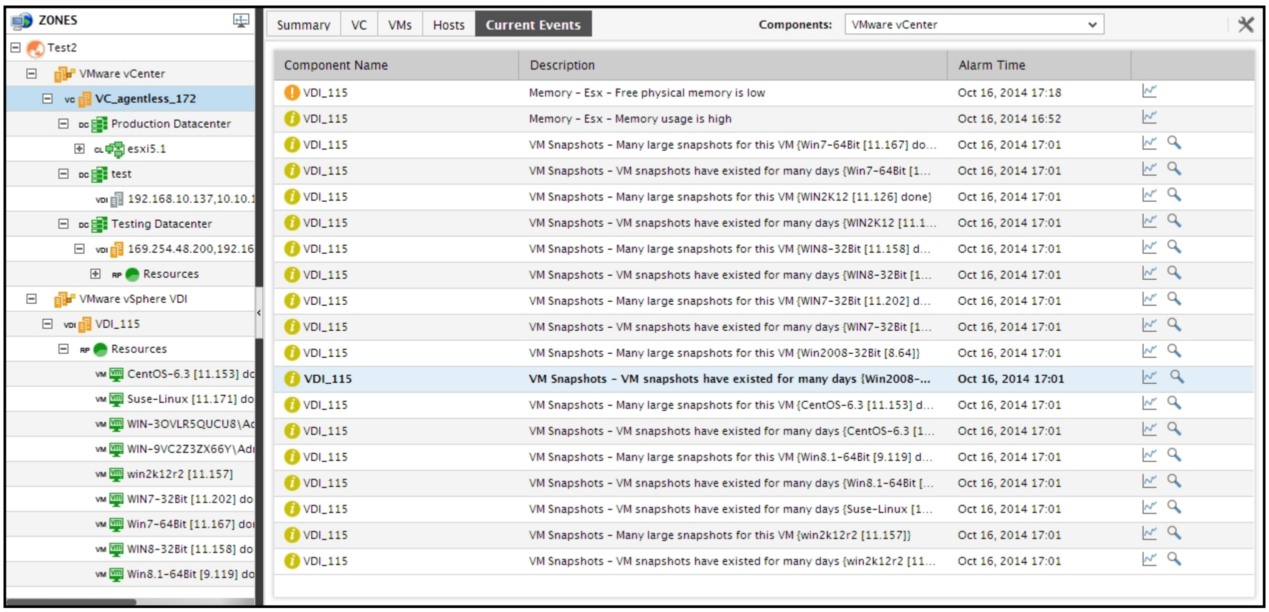

Figure 10.227: The Events tab page for a vCenter server chosen from the tree

6. Typically, when you expand a vCenter server node in the tree, the physical ESX servers that are being managed by that vCenter server will appear.

   Note:

   While the vCenter server tree will list even those ESX servers that are not monitored by eG Enterprise, the tree will not indicate the current state of such servers; also, clicking on any such server will not display corresponding performance information in the tab pages in the right panel.

7. However, in some environments, folders may be configured on vCenter server, where every folder could contain one/more datacenters, ESX servers, clusters, datastores, and VMs. In such cases, expanding the vCenter server node will reveal the folders configured on that vCenter server as sub-nodes. Clicking on a folder sub-node will display the complete list of problems currently affecting the performance of datacenters, clusters, ESX servers, and VMs that are included in that folder in the **Current Events** tab page (see Figure 10.227).

8. Expanding the folder node will reveal the datacenters within that folder. If no folders exist, then expanding

a vCenter server node in the tree, will display the datacenters that have been configured within a vCenter. To know what problems are affecting the performance of a particular datacenter currently, click on the node representing a datacenter in the tree. Figure 10.226 then appears.

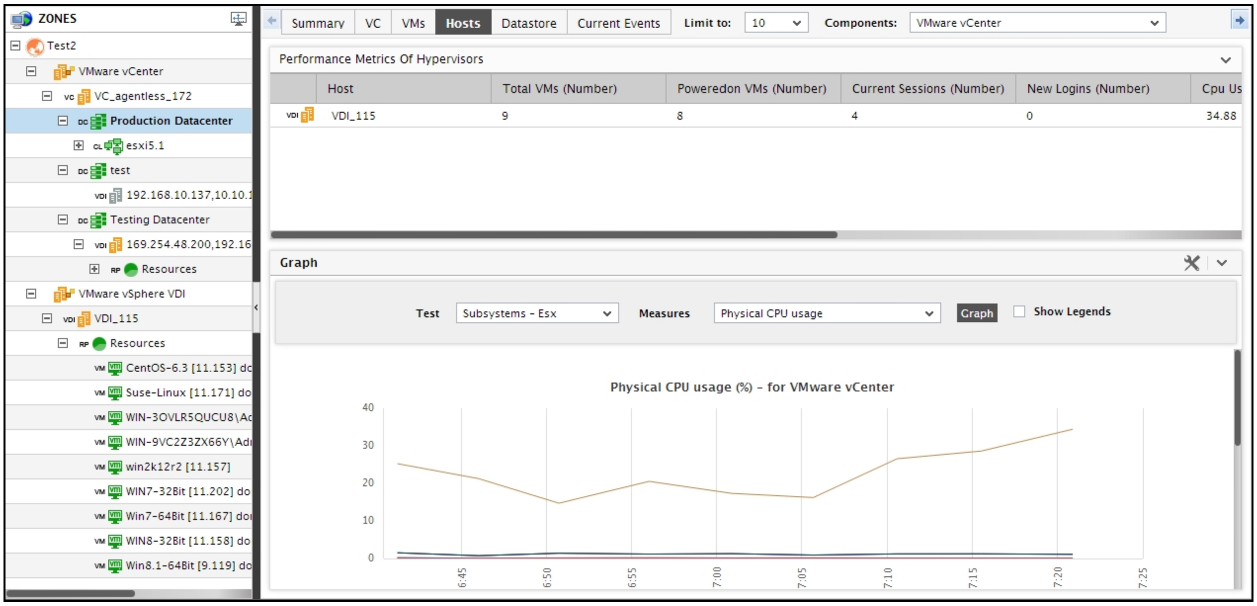

Figure 10.228: The Events tab page for a datacenter chosen from the tree

9. Every datacenter, in turn, may manage clusters of ESX servers. If such clusters have been configured on any monitored vCenter server, then, in the **Virtual Infrastructure** tree, these clusters will appear as sub-nodes of that datacenter node. If you click on a cluster sub-node in the tree, the **Current Events** tab page will reveal the problem events pertaining to all managed ESX servers included in the cluster.

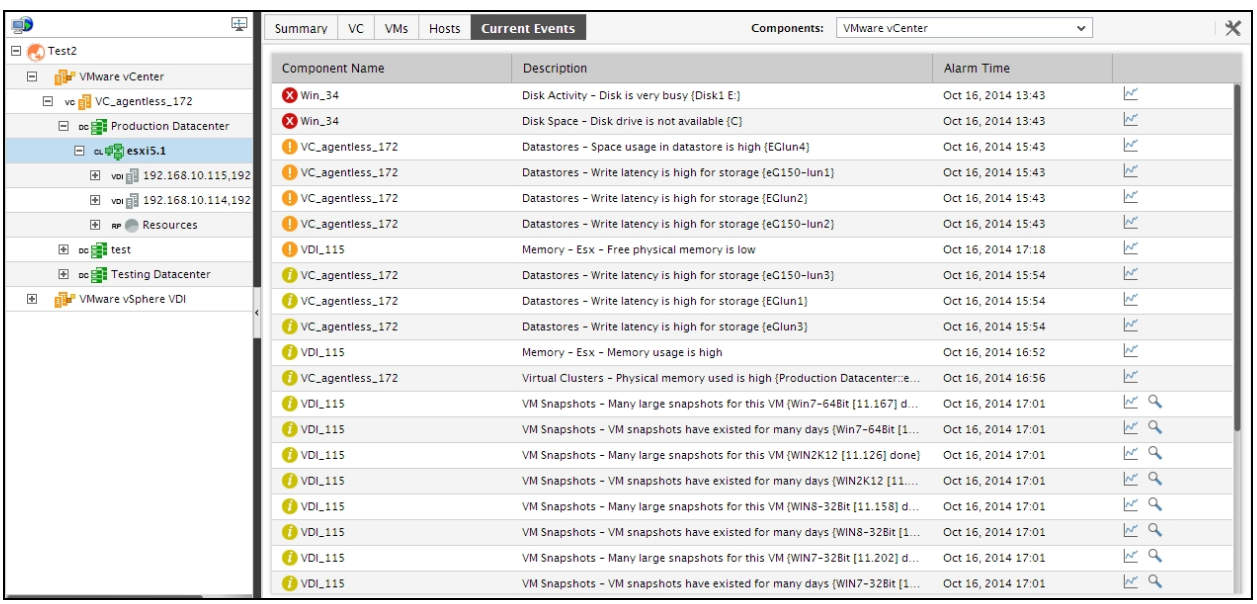

Figure 10.229: The Current Events tab page for a cluster

10. To view the individual virtual hosts that are part of a zone, do any of the following:

- Expand the nodes representing the virtual component-types in the zone;
- If the zone consists of components of type *VMware vCenter*, expand the node representing the monitored vCenter server in your environment;
- If datacenters are configured on a monitored vCenter server, expand a datacenter sub-node under the vCenter server node;
- If clusters are configured within a datacenter, expand the cluster sub-node;

If you then click on a virtual host in the tree, the **Current Events** tab page will change to display problems affecting that virtual host (see Figure 10.227).

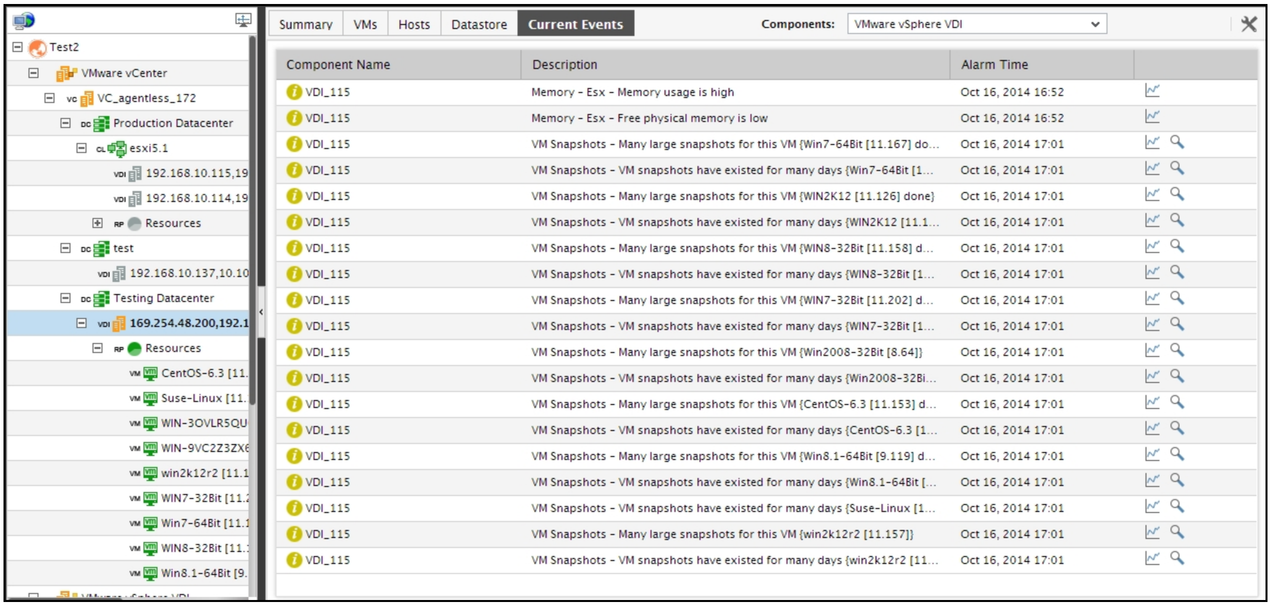

Figure 10.230: The Events tab page for a particular virtual host selected from the tree

11. Similarly, if you expand the virtual host node in the tree-structure, you can view the name and state of those VMs that are executing on that virtual host. If you now click on a VM in the tree, the problem affecting that VM alone will be displayed in this tab page.
12. If resource pools are configured on a virtual host, then the resource pools also will appear as the sub-nodes of the virtual host-node. If you click on a resource pool in the tree, the corresponding **Current Events** tab page will list the current problems affecting the performance of one/more virtual machines included in the resource pool.
13. To perform additional diagnostics on one of the problems listed in this tab page, you can click on the corresponding **Graph** icon. A graph of the problem measure for the last 1 hour (by default) will then appear revealing when exactly the problem occurred.
14. Clicking on an event listed in this tab page will lead you to the layer model page of the problem component, using which you can quickly determine the problem layer, test, and measure.

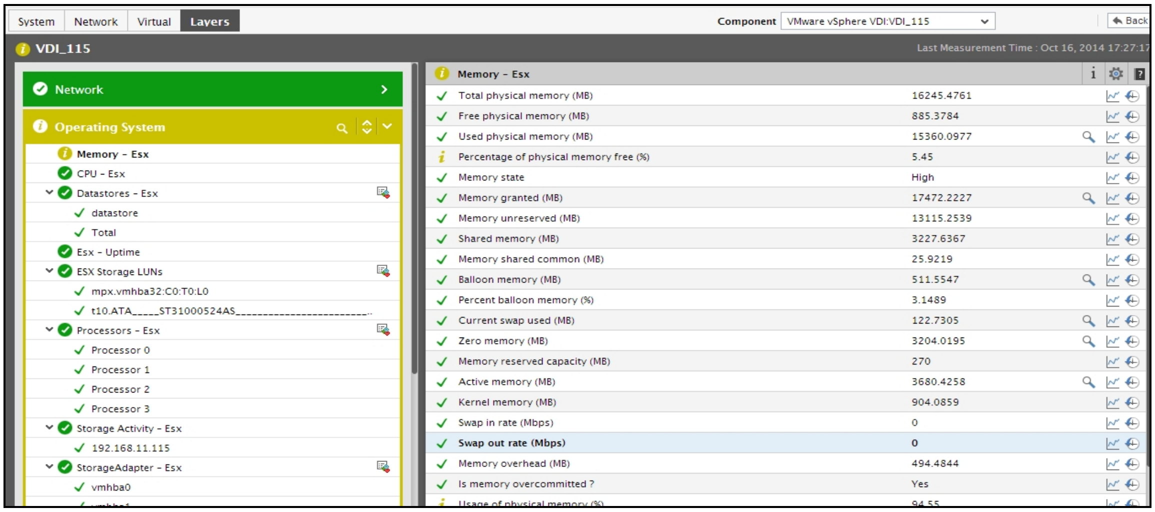

Figure 10.231: The problem layer, test, and measure of the problem component

10.7.7 The Resource Pools Tab Page

If resource pools have been configured on a virtual host, then the eG agent auto-discovers these pools and displays them as sub-nodes of a virtual host-node. The **Resource Pools** tab page appears only when a resource pool in the tree is clicked on (see Figure 10.232). This tab page reveals the current configuration of the chosen resource pool, which includes the number of virtual machines on the resource pool, the number of running virtual machines, and the number of child resource pools. Besides the configuration, administrators can also use this tab page to determine the current state of each of the virtual machines and child resource pools under the chosen resource pool, and simultaneously analyze the resource usage by the pool.

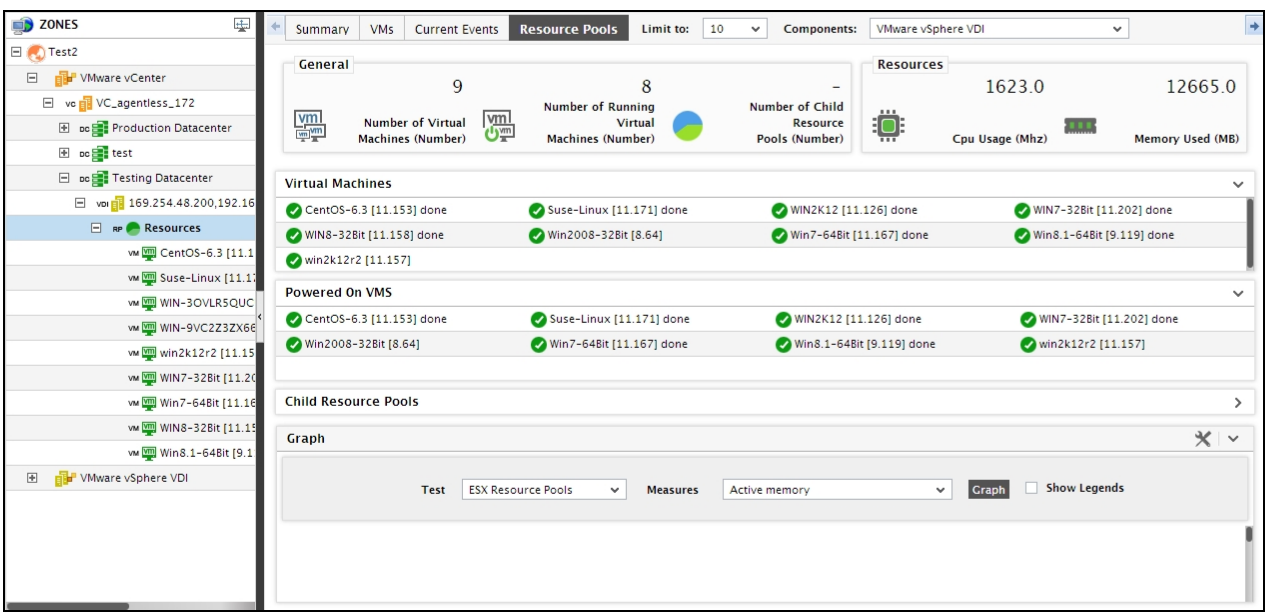

Figure 10.232: The Resource Pools tab page

Clicking on a VM listed against the **Virtual Machines** section of a resource pool leads administrator directly to the layer model page of the virtual host on which that resource pool is configured, and automatically displays metrics that provide an “outside view” of that VM’s performance (see Figure 10.233). These metrics help administrators in understanding the impact the VM has on the physical resources of the virtual host.

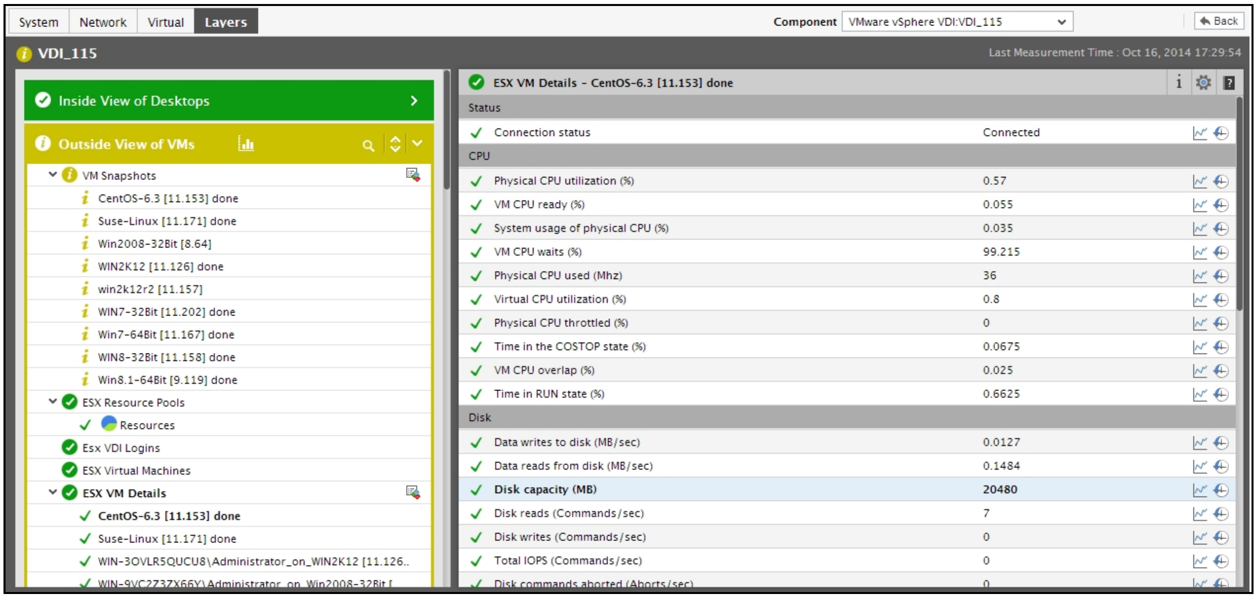

Figure 10.233: The outside view of a VM under the resource pool

Similarly, clicking on a particular child resource pool displayed against the **Child Resource Pools** section of Figure 10.233, takes the administrator to the layer model page of the corresponding virtual host, thereby granting him a sneak peek at the resource usage metrics of that child resource pool, and enabling him to analyze how resource-intensive the child is (see Figure 10.233).

10.8 Application-Specific Custom Dashboard

Though the contents of the system/network/application dashboards are customizable, the right to customization rests with the **Admin** and **Supermonitor** users alone, and not all users to the eG monitoring console. This means that if any monitor user logs into the console, he/she will only be allowed to use the pre-defined dashboards. These dashboards, as we know, will focus on only those metrics that an **Admin** or **Supermonitor** has configured - a normal monitor user can neither customize the layout nor alter the contents of such dashboards.

Also, note that by default, dashboards are available only for those applications that are supported out-of-the-box by the eG Enterprise Suite. For in-house/legacy applications that may have been integrated into the eG Enterprise Suite using the Integration Console plugin, ready-to-use dashboards are not available.

Therefore, to enable every user with monitoring rights to personalize his/her dashboard experience, the eG monitoring console allows the creation of **Custom Dashboards**. These dashboards can be designed for both existing applications and legacy applications. This capability empowers users to control what data is to be displayed in the dashboards and how to present it (whether to use dial charts or digital displays or comparison graphs or tables). This way, users can see what they want to see in the dashboards.

Using the eG monitoring console, a user can build a custom dashboard for a specific host/application, or can design a dashboard to visually compare performance across multiple applications. The sections that follow will discuss both these custom dashboard options.

To build custom dashboard for a particular application, do the following:

1. By default, clicking on a component in the **COMPONENT LIST** page will lead you to the **Layer** tab page of that component, where the layer model, tests, and measurements pertaining to that component will be displayed.
2. If a custom dashboard is enabled for that component, then a **Custom** tab page will appear next to the **Layer** tab page. If custom dashboard templates pre-exist for the component in question, then, upon clicking on the **Custom** tab page, the eG Enterprise system will automatically sort these dashboard templates in the ascending order of their names, and display the first template in the sorted list in the **Custom** tab page. However, if no dashboard templates pre-exist for the said component, then a message to the effect will appear, as depicted by Figure 10.234 below:

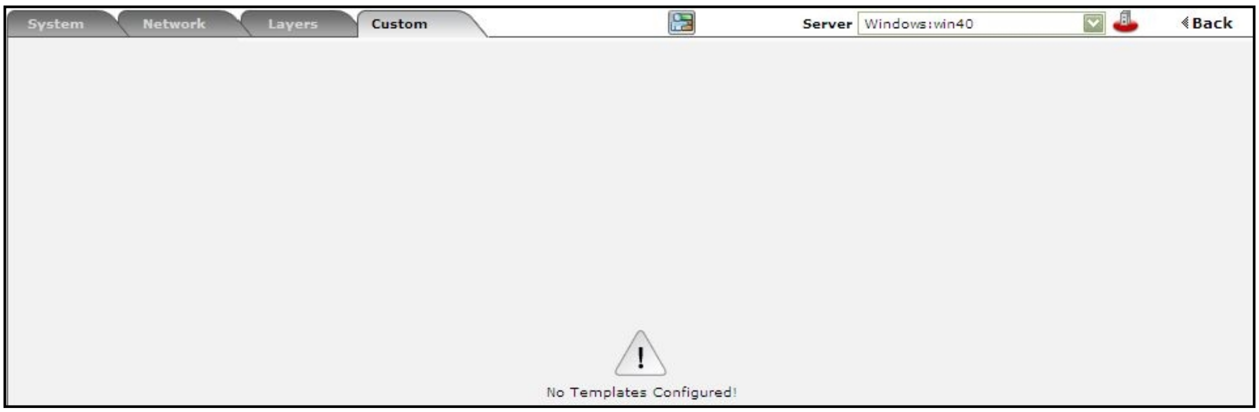

Figure 10.234: The message that appears when no dashboard templates have been configured for a component

3. To create a new dashboard, click on the icon at the top of Figure 10.234. Figure 10.235 will then appear displaying a **Personal Templates** section with the complete list of custom dashboard templates that pre-exist, regardless of the component with which they are associated.

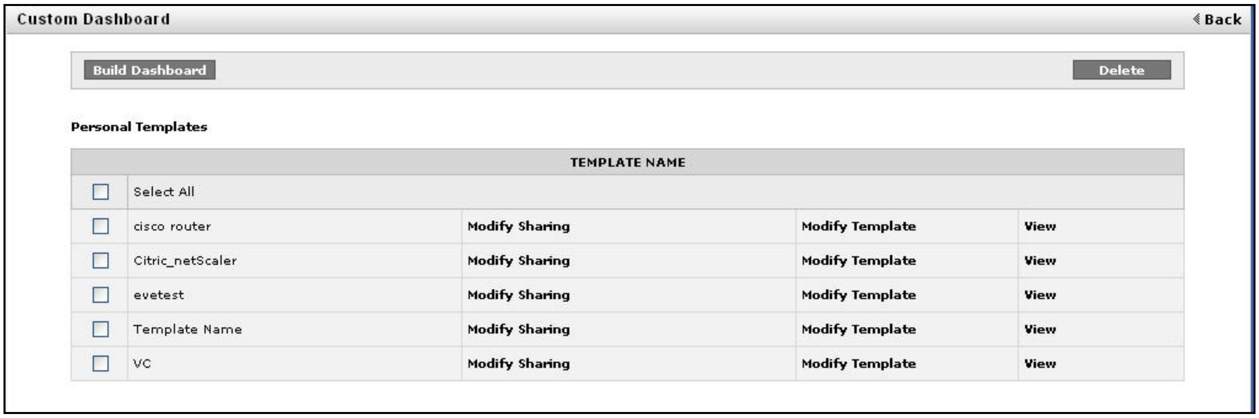

Figure 10.235: List of dashboard templates

4. To modify a dashboard template, click on the **Modify Template** link corresponding to it, and to view a dashboard template, click on **View**. If you want to delete any of the listed templates, select the check box that pre-fixes the template and click the **Delete** button. To remove all templates, select the **Select All** check box, and then click the **Delete** button. If required, you can even modify how a dashboard template is to be shared with the other users registered with the eG Enterprise system. For this, you will have to click on the **Modify Sharing** link that corresponds to the template.

5. If other users have shared one/more dashboards that they created with you, then a separate **Templates shared by other Users** section will appear below the **Personal Templates** section (as depicted by Figure 10.236). This section will list all those dashboard templates that other users have shared with you, along with the names of the users who shared the dashboard (in the **SHARED BY** column). Shared dashboards can only be viewed (not modified or deleted). Clicking on the **View** link against a dashboard template in this section will allow you to view a shared dashboard.

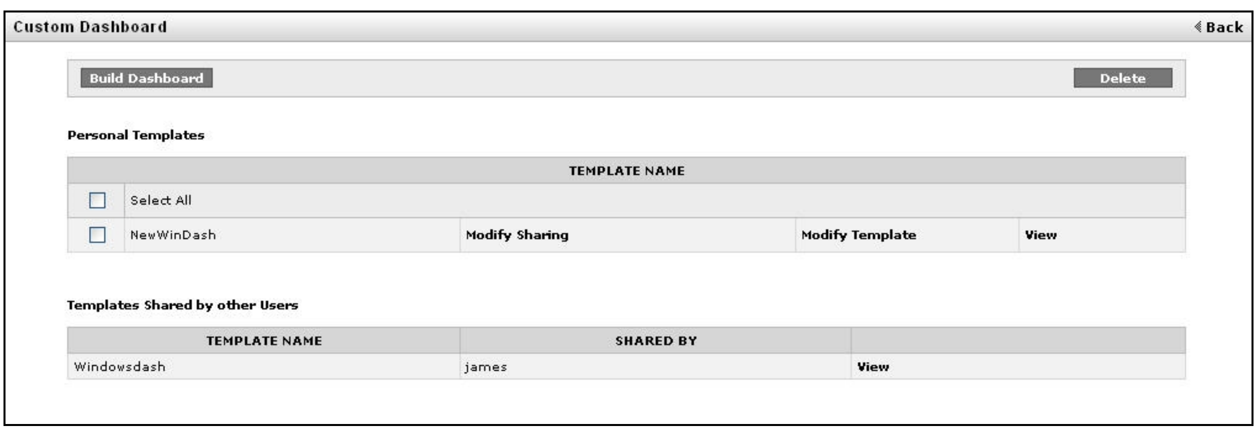

Figure 10.236: Viewing the list of templates shared by other users

6. To build a new dashboard, click on the **Build Dashboard** button. Provide a name for the new dashboard template in the **Dashboard Name** text box in Figure 10.236. Then, from the **Application** list, pick the component for which the new dashboard is being configured. By default, the component in question will be selected from the **Application** list. Then, by selecting an option from the **Sharing** list, indicate how you want to share the dashboard with other users to the eG Enterprise system. By default, the **Private** option is

chosen from this list. This indicates that the user building the dashboard is alone to authorized to view/modify/delete the dashboard. Selecting the **Public** option allows all users to the eG Enterprise system to view (not modify/delete) the dashboard that is being created. To share the dashboard with specific users to the eG Enterprise system, select the **Share** option from the **Sharing** list. Then, from the **Available Users** list that then appears (see Figure 10.238), select the users with whom you want to share the dashboard, and click the **Grant** button. This will transfer the selection to the **Selected Users** list. To revoke the share, select one/more users from the **Selected Users** list and click the **Revoke** button.

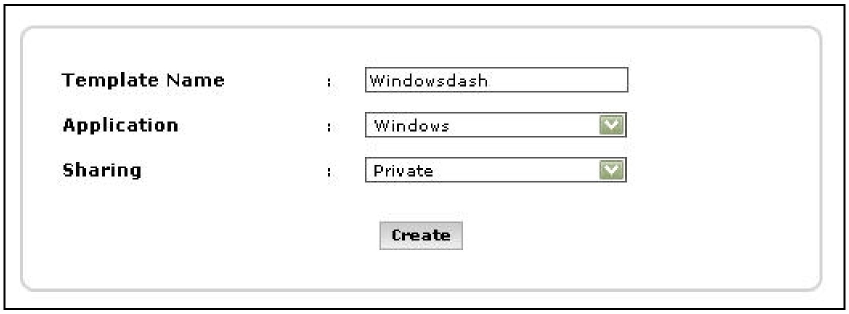

Figure 10.237: Creating a new dashboard template

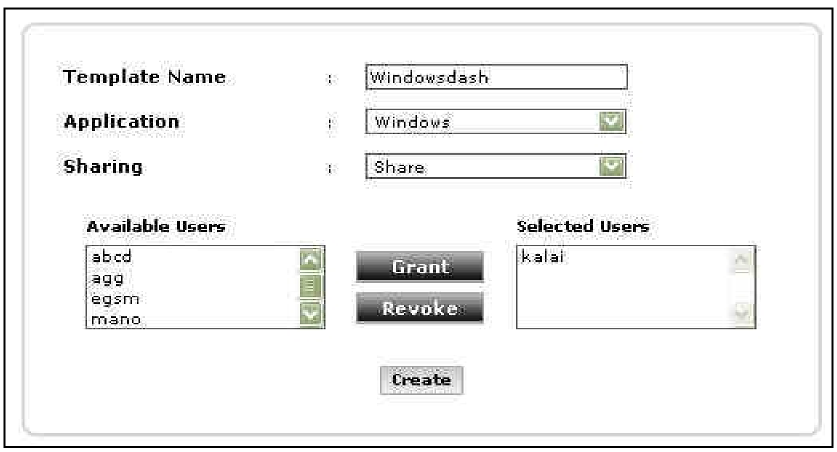

Figure 10.238: Sharing the dashboard template with specific users

Note:

Users with whom you share your dashboard template - i.e., users in the **Selected Users** list of Figure 10.238 - will only be allowed to view the dashboard, and not modify/delete it.

7. Finally, click the **Create** button to add the new dashboard template.
8. Doing so will invoke Figure 10.239.

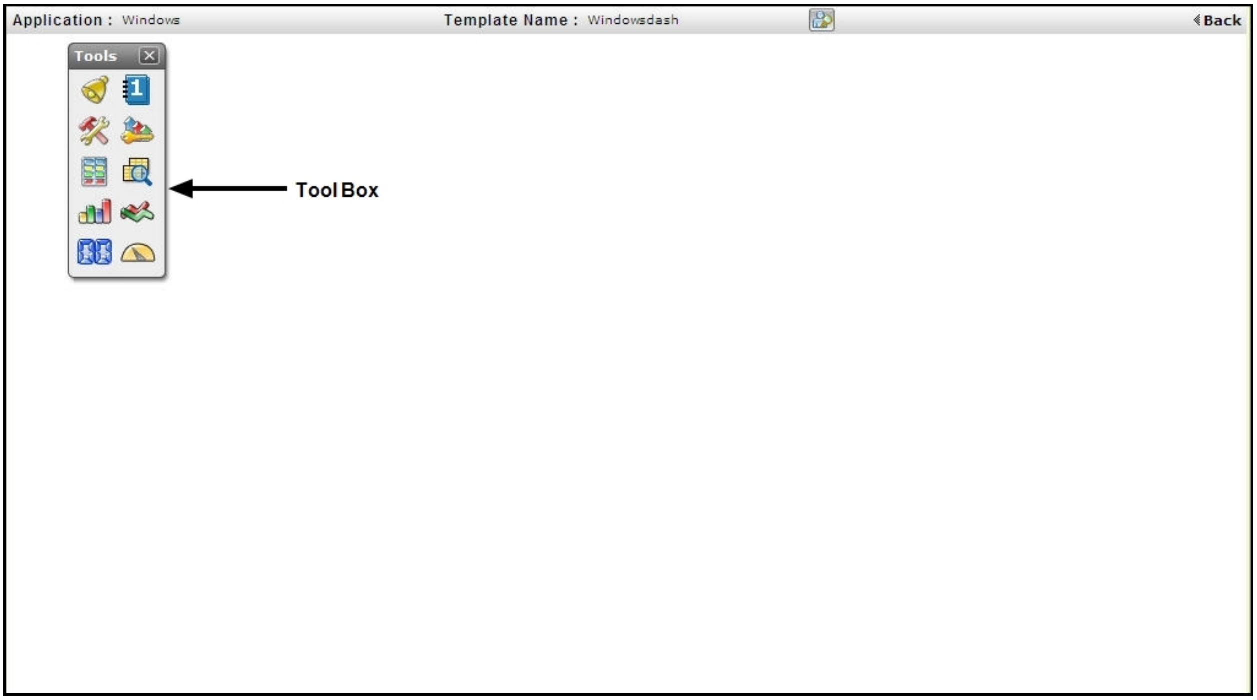

Figure 10.239: The Design window of the custom dashboard

9. Using Figure 10.239, you can design 'look and feel' of the new dashboard. To help you in building your dashboard, a tool box is provided to you, as indicated by Figure 10.239. The tools available in this tool box have been detailed in the table below:

Tool	Tool Name
	Current Alerts
	Event History
	Configurations
	Performance Indicators
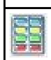	Comparison Table
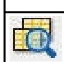	DD Comparison Table
	Comparison Graph
	Timeline Chart
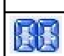	Digital Chart
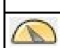	Dial Chart

10. If you want your dashboard to display a quick summary of the problems that the target component is

currently experiencing, then, click on the tool in the tool box. This will insert a **Current Alerts** section in your custom dashboard as depicted by 10.8.

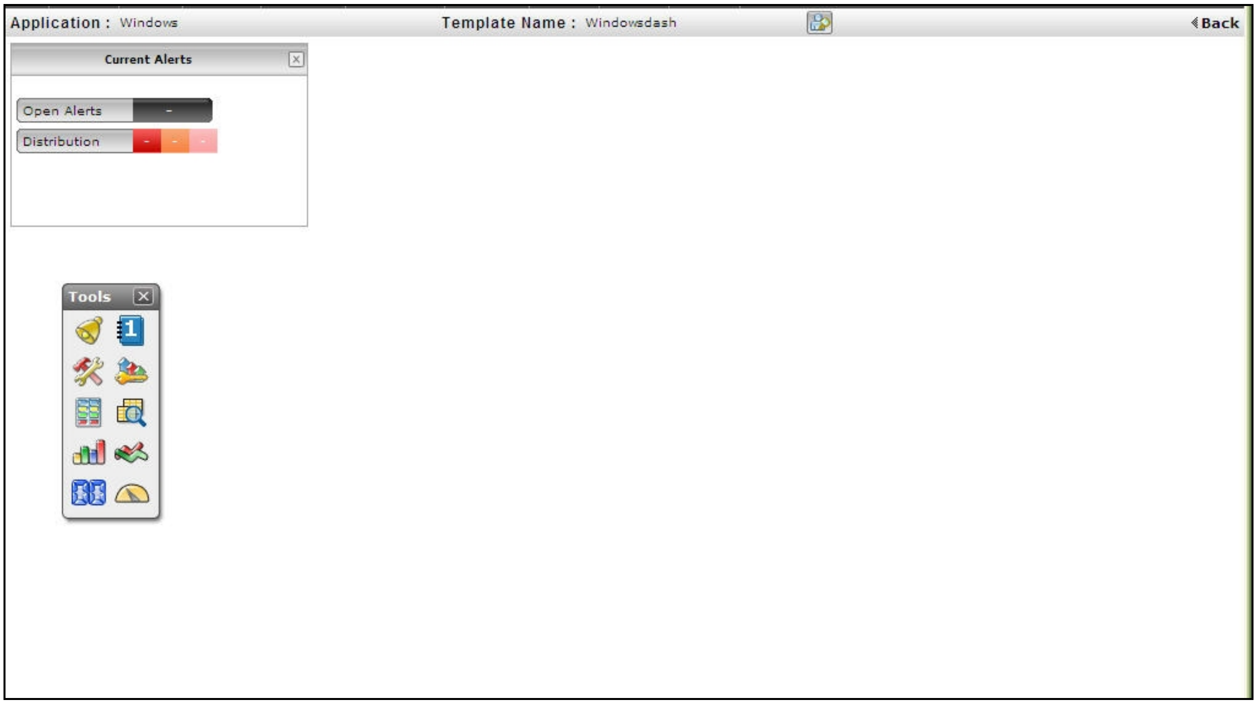

Figure 10.240: Inserting a Current Alerts section in the dashboard

11. In real-time, the **Current Alerts** section will display the number of problem events currently associated with the target component, and will also reveal how these problems are distributed based on problem priority.

12. If you want to include an event history in your dashboard, then, click on the tool in the tool box. A bar graph will be inserted in your dashboard, as depicted by 10.8.

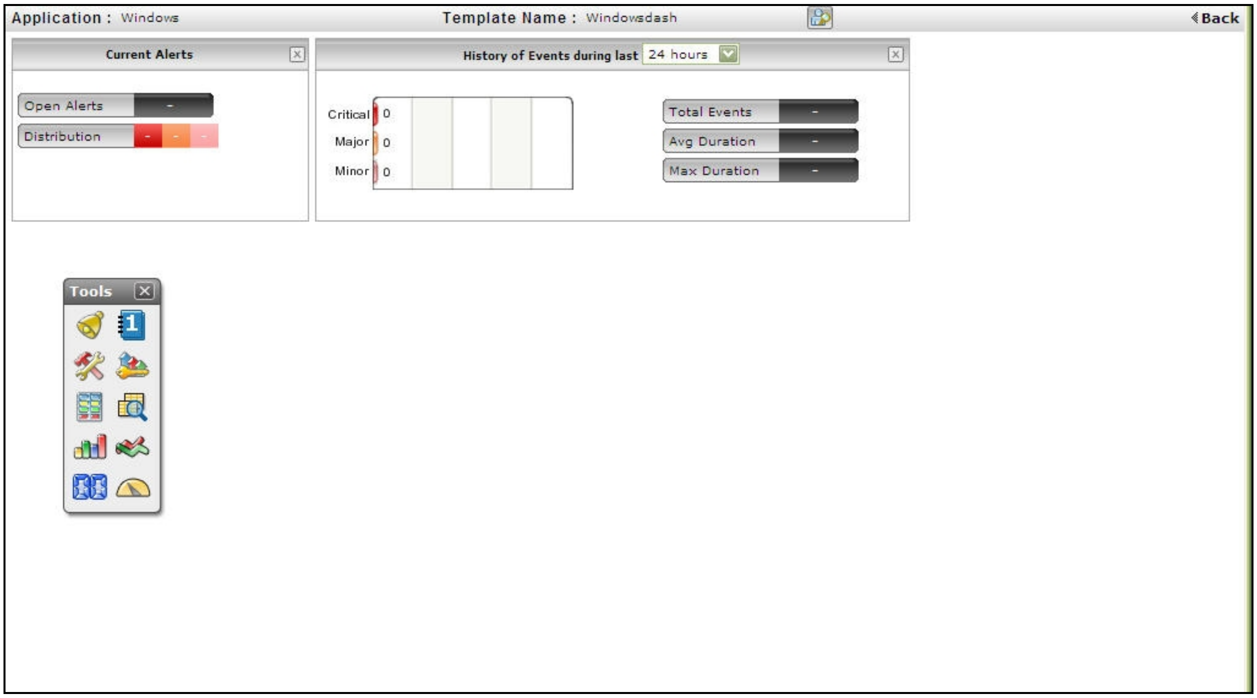

Figure 10.241: Including an Event History in your dashboard

13. At runtime, the event history bar graph will reveal the **History of Events** experienced by the target component during the last 24 hours (by default). If need be, you can pick another duration from the **History of Events during last** list in 10.8, so as to override the default duration of 24 hours.

14. If you want to view the basic configuration of the target component in the dashboard, then, click on the  tool in the tool box. This will add a **Configuration** section to the dashboard, as depicted by Figure 10.242.

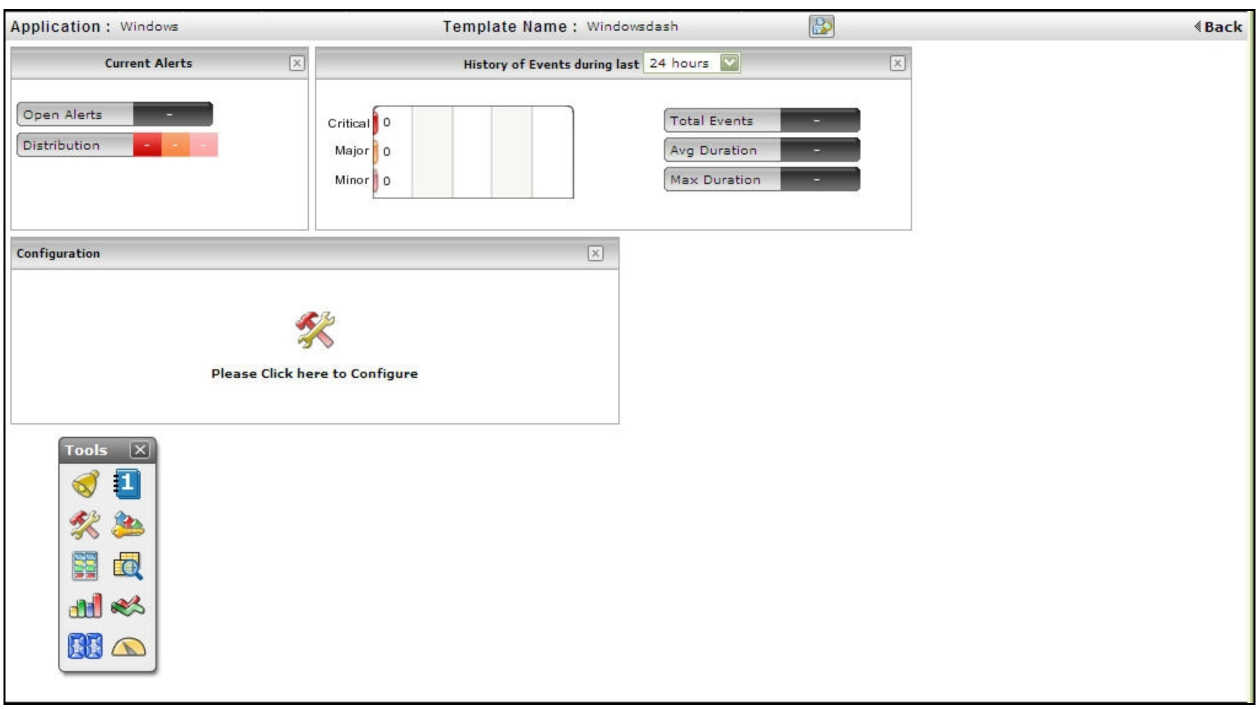

Figure 10.242: Adding a Configuration section to the custom dashboard

15. You can even choose the type of configuration information that you want displayed in the **Configuration** section. For this purpose, click on the **Please Click here to Configure** icon within that section in Figure 10.242. This will invoke a **Properties** window as depicted by Figure 10.243. From the **Tests** list in the **Properties** pop-up, select the test that reports the configuration metrics of interest to you. This will populate the **Measures** list with the measures reported by the chosen test. Now, from the **Measures** list, pick the measures to be added to the **Configuration** section. You can add multiple measures reported by multiple tests to this section. To remove any measure from the section, simply deselect the check box alongside the measure in the **Configuration** section.

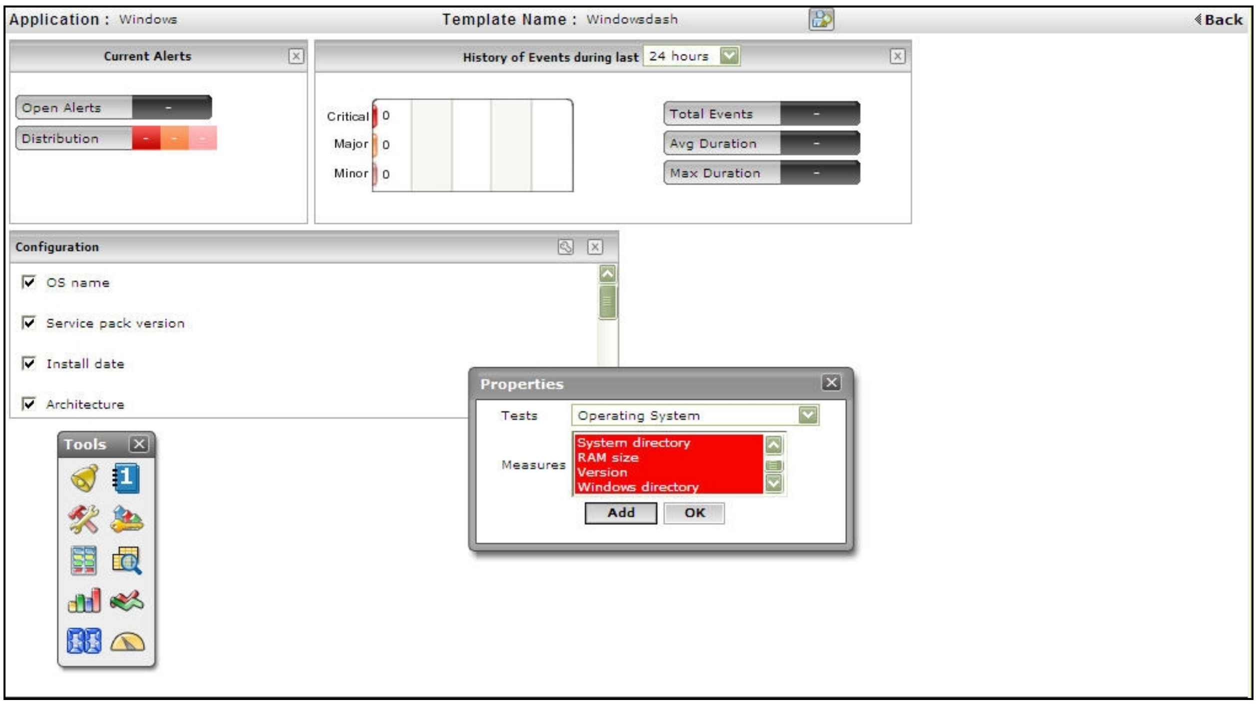

Figure 10.243: Configuring the configuration metrics to be displayed in the Configuration section

16. To focus on specific metrics reported by a component so that you are promptly alerted to any change in the state of those measures, you can mark such metrics as key performance indicators and include them in the **Key Performance Indicators** section in your dashboard. To add this section to your dashboard, click on the  tool in the tool box. Figure 10.244 will then appear.

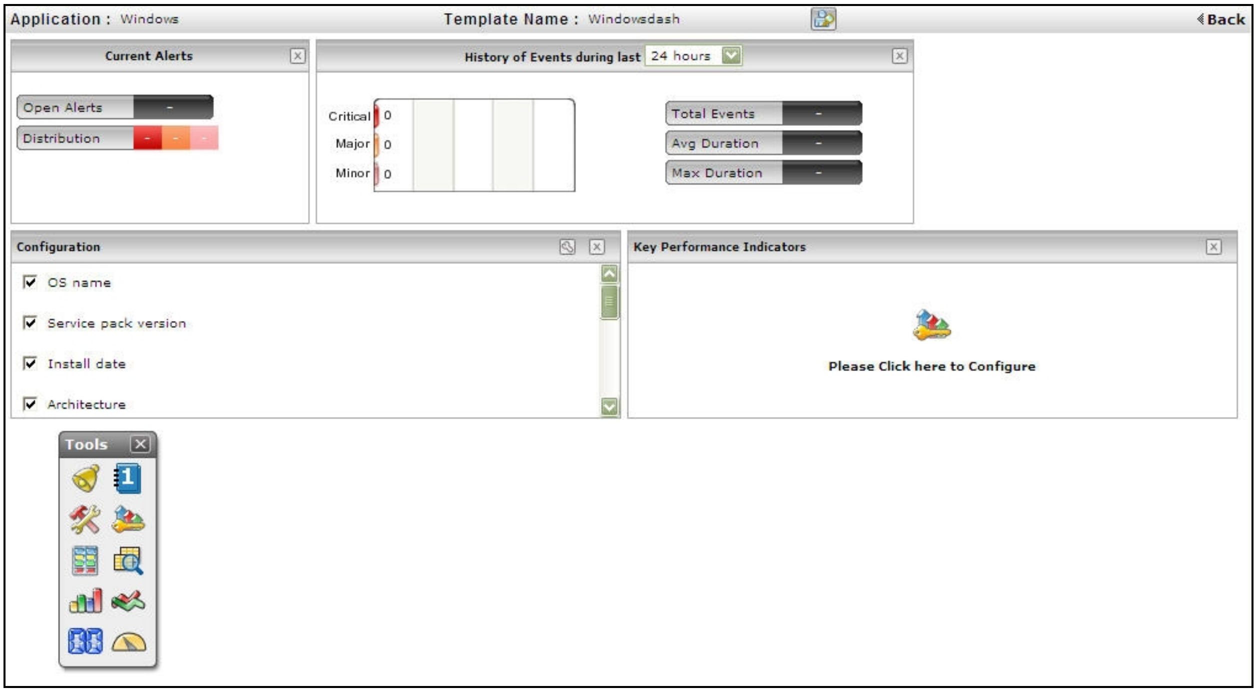

Figure 10.244: Inserting the Key Performance Indicators section in the dashboard

17. To configure the metrics to be displayed in the **Key Performance Indicators** section, click on the **Please Click here to Configure** icon within that section in Figure 10.245.

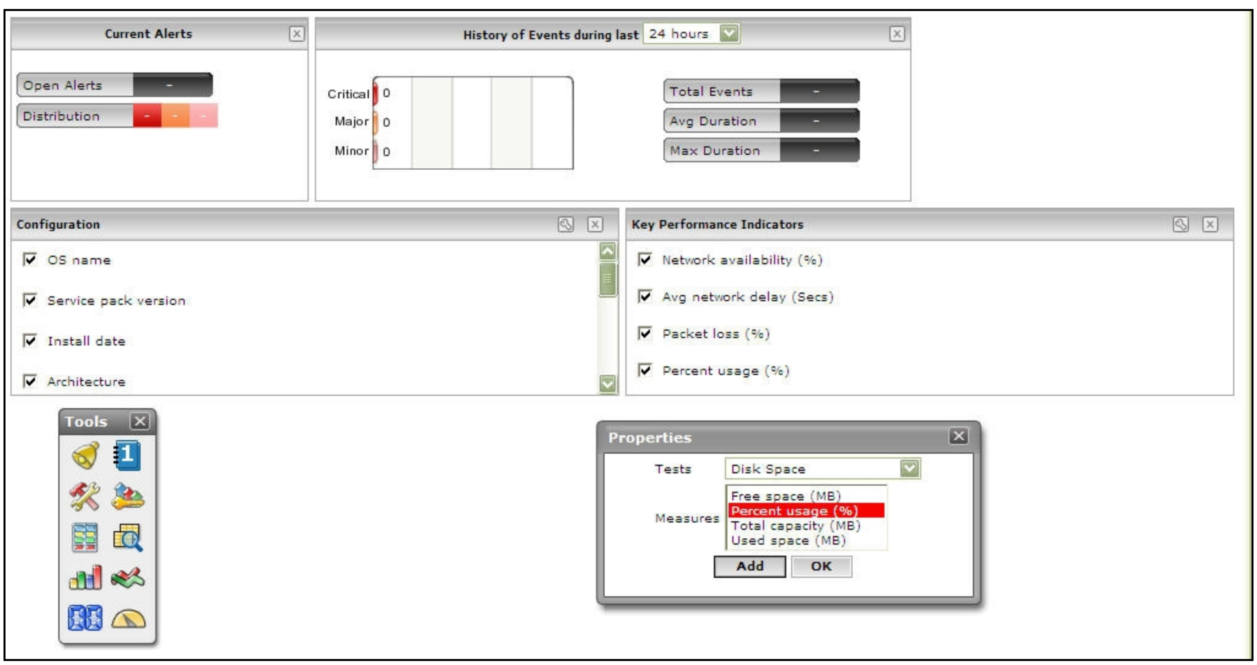

Figure 10.245: Configuring the metrics to be included in the Key Performance Indicators section

18. A **Properties** pop-up then appears (see Figure 10.245). From the **Tests** list in the **Properties** pop-up, select the test that reports the critical metrics. This will populate the **Measures** list with the measures reported by the chosen test. Now, from the **Measures** list, pick the measures to be added to the **Key Performance Indicators**

section. You can add multiple measures reported by multiple tests to the **Key Performance Indicators** section. To remove any measure from the section, simply deselect the check box alongside the measure in the **Key Performance Indicators** section.

19. Sometimes, you may want to quickly compare the performance of one/more measures across a set of descriptors, so that potential bottlenecks and the descriptors responsible for the same can be isolated. For instance, you may want to compare disk space usage across disk partitions to identify the partition that may soon run out of space. To facilitate this comparative analysis, you may want to include a comparison table in your dashboard. To do so, simply click on the  tool in the tool box. A new section will be inserted in your dashboard layout, as depicted by Figure 10.246.

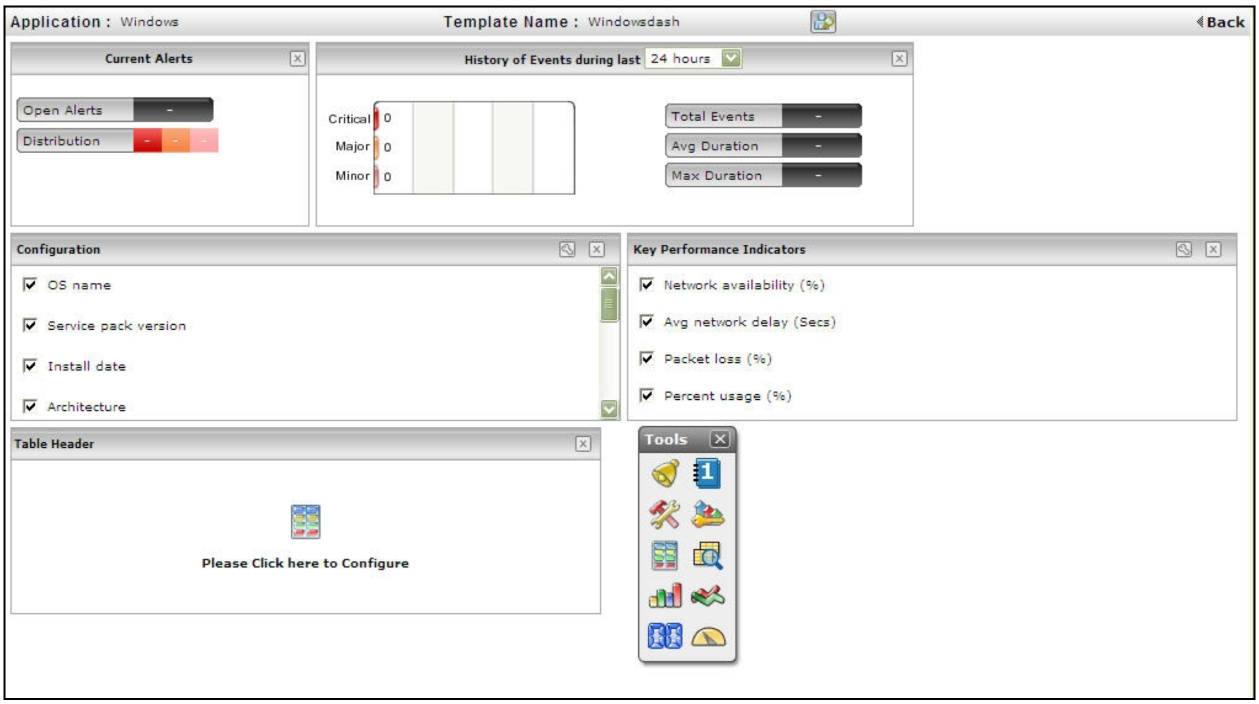

Figure 10.246: Adding a Comparison Table to your dashboard

20. To indicate which measures need to be compared across descriptors, click on the **Please Click here to Configure** icon within the **Table Header** section in Figure 10.246. A **Properties** window will then pop-up as shown by Figure 10.247. From the **Tests** list in the **Properties** pop-up, pick a descriptor-based test, and then pick one/more measures reported by that test from the **Measures** list. **Note that a single comparison table can be associated with a single test only**. Then, click the **OK** button in the **Properties** pop-up. Upon clicking **OK**, the name of the test will become the title of your comparison table section, and the chosen measures will be listed therein. To remove a measure from the comparison table, simply deselect the check box alongside that measure in the comparison table section in your dashboard.

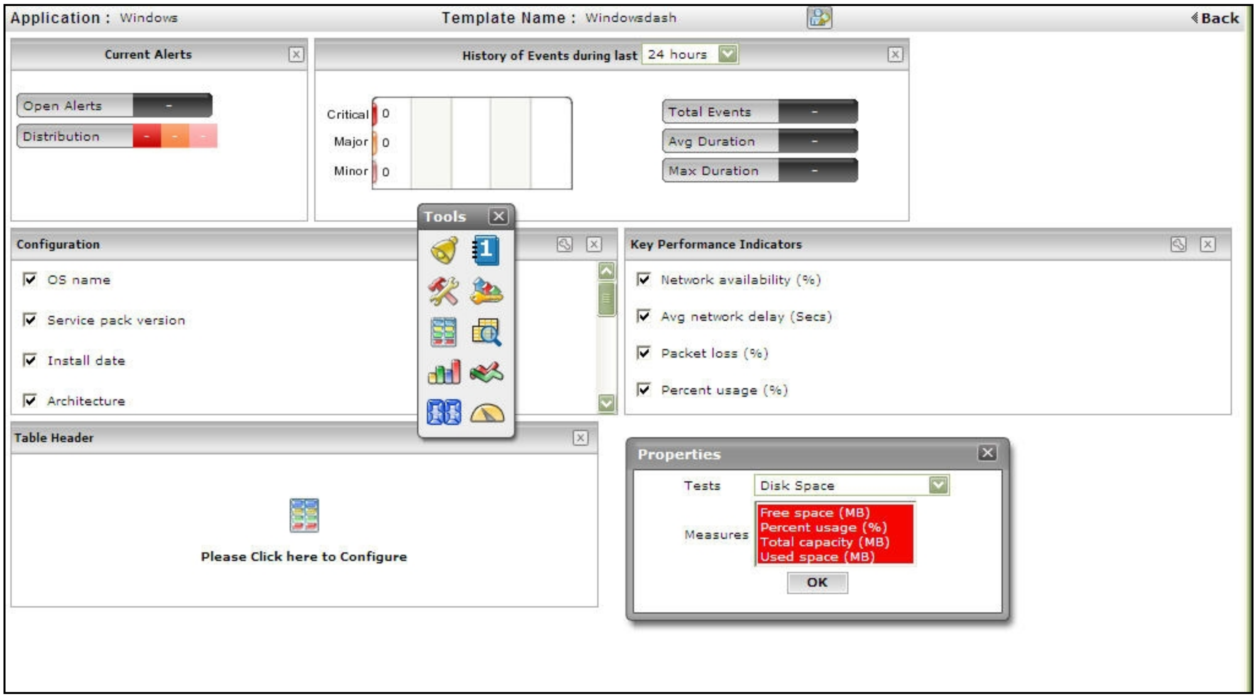

Figure 10.247: Configuring the measures to be compared in the comparison table

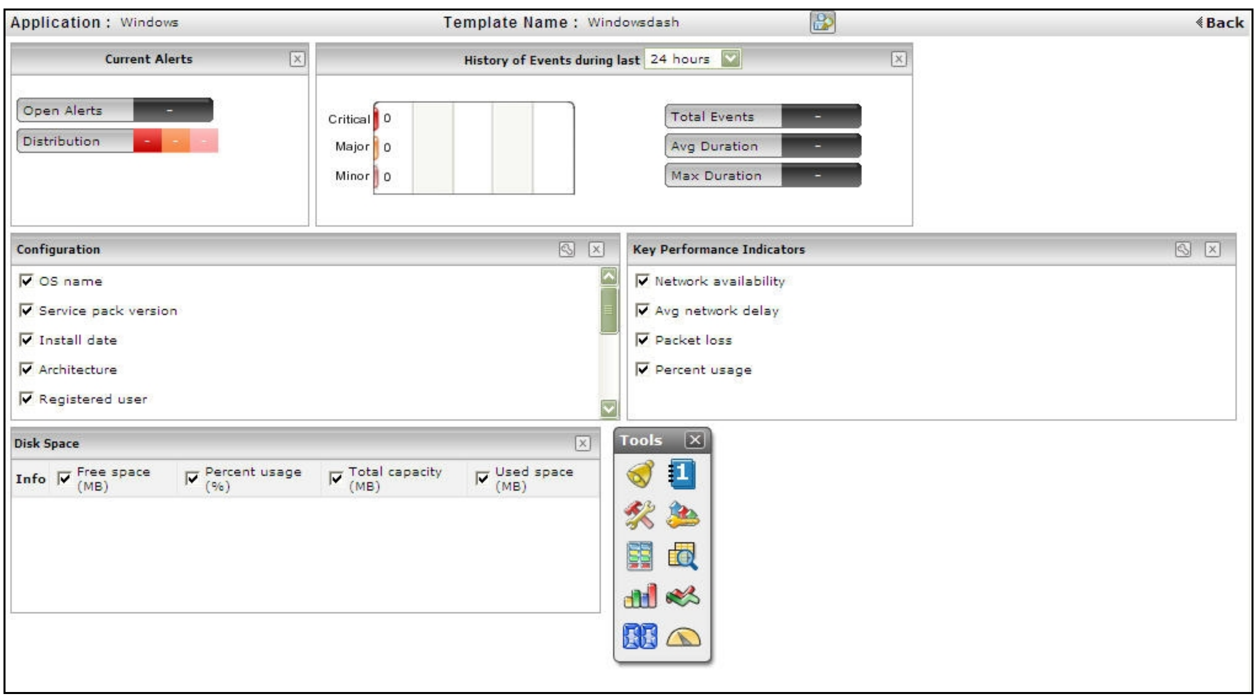

Figure 10.248: Selected measures appearing in the 'Comparison Table' placeholder

21. Sometimes, you may want to comparatively analyze the information provided in the detailed diagnosis of a test in order to understand the problem reported by that test better and to accurately isolate the source of the problem. For instance, the detailed diagnosis of the *Free memory* measure of the **System Details** test of a

host provides the PID (process ID) of the top memory-consuming processes on that host and the percentage of memory consumed by each process. In the event of a memory contention on the host, you can use the detailed diagnosis to compare the memory usage of the processes and identify that process which is responsible for the memory drain. If this comparison is available in a custom dashboard, you can get to the root-cause of the memory drain much quickly. To display such useful detailed diagnosis information in your custom dashboard, insert a **DD Comparison Table** section in it. For this, click on the icon in the tool box.

22. This will insert a **DD Comparison Table Header** in your dashboard.

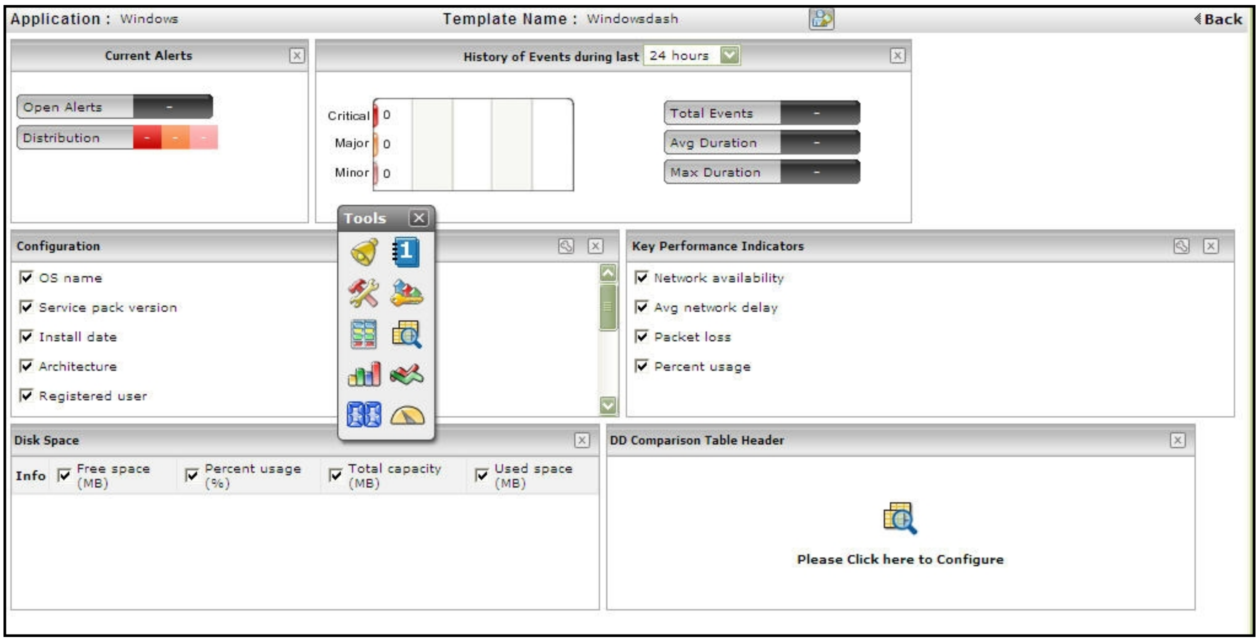

Figure 10.249: Inserting a DD Comparison Table section in the dashboard

23. Once you click on the **Please Click here to Configure** link in Figure 10.249, a **Properties** dialog box will appear. In the **Tests** drop-down list here, the tests for which the 'detailed diagnosis' capability is enabled will be listed. Once you select a test from this list, the measures for which detailed diagnosis is available will be populated in the **Measures** drop-down list. **Note that you can select only one test and one measure pertaining to that test for viewing the detailed diagnosis in a single DD Comparison table. You may add more DD Comparison tables if you wish to view the detailed diagnosis for more measures.** Upon choosing a measure, a **DD Columns** list will appear. Here, select the columns that you wish to view in the dashboard and click the **OK** button. At runtime, the **DD Comparison Table** section will display chosen **DD columns** in a tabular format.

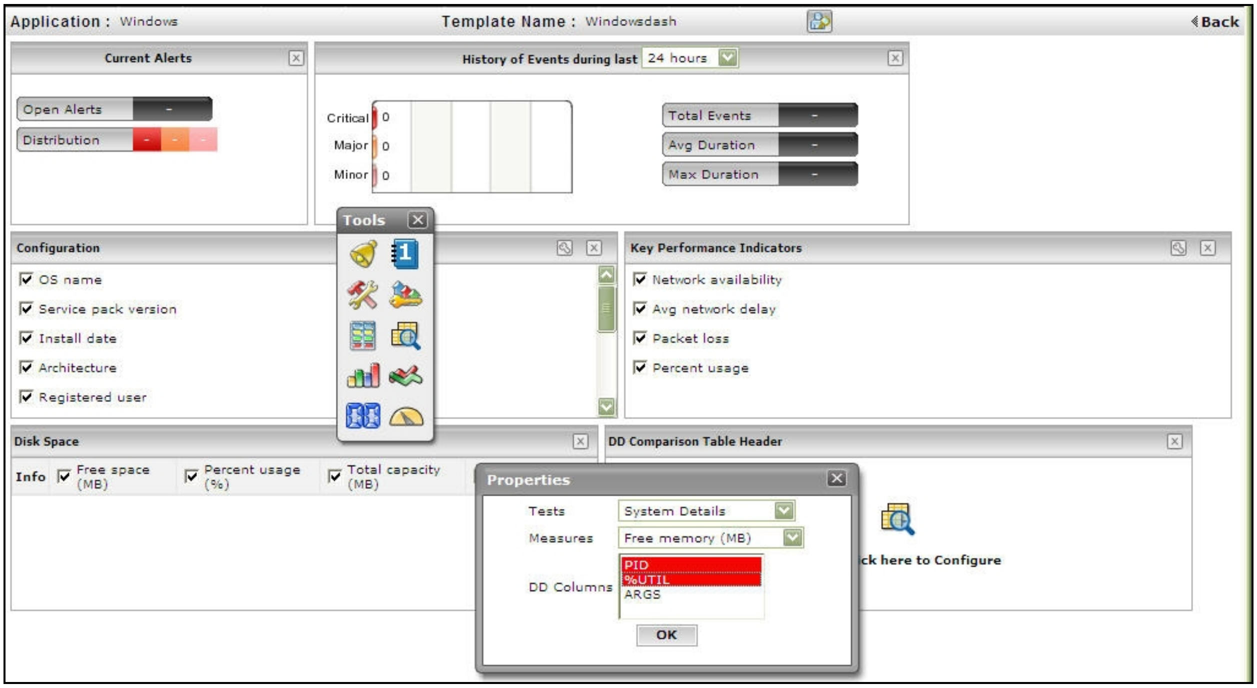

Figure 10.250: Selecting the DD columns to be included in the DD Comparison Table section

24. If you want to graphically compare the values of specific measures across their descriptors, then, you can insert a **Comparison Graph** in your dashboard. Such a graph will enable you to accurately identify the best/worst performers in a chosen performance area. To insert this graph in your dashboard, click on the tool in your tool box. A section depicted by Figure 10.251 will then be included in your dashboard design.

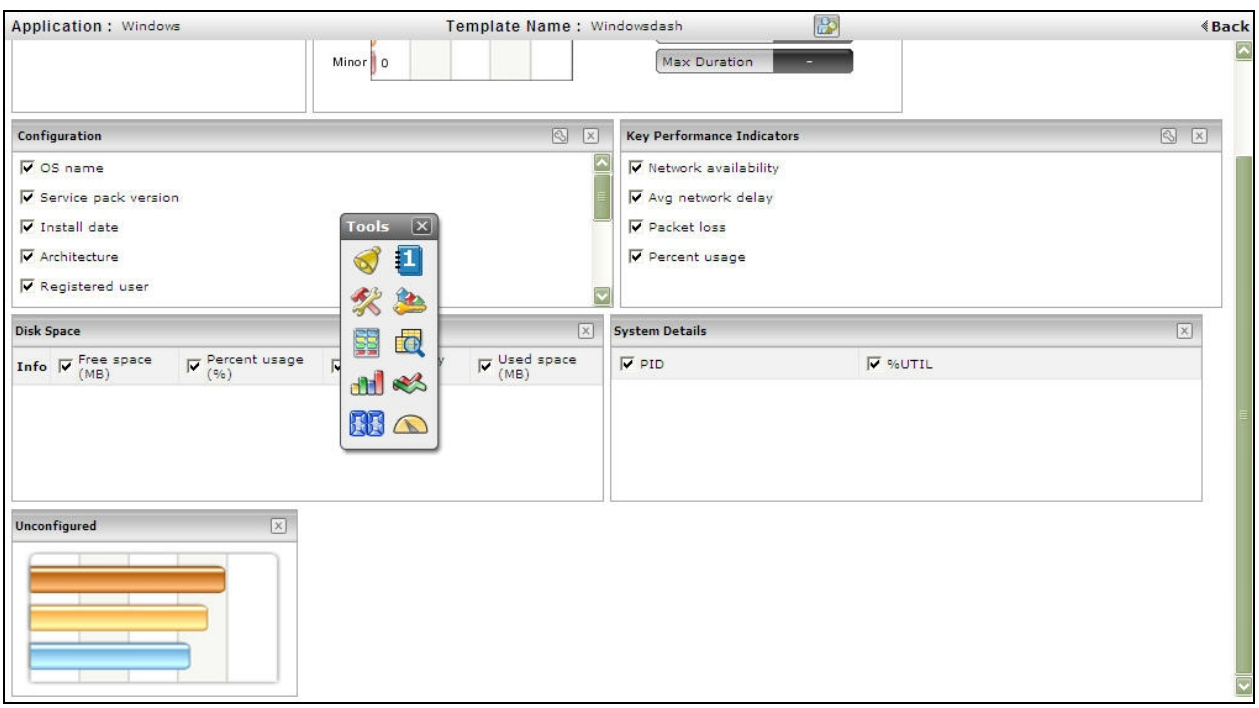

Figure 10.251: Including a comparison graph in your dashboard

25. To configure the measure to be compared using the comparison graph, click on the graph titled **Unconfigured** in Figure 10.251. A **Properties** window will then pop-up as shown by Figure 10.252. From the **Tests** list in the **Properties** pop-up, pick a descriptor-based test, and then pick a measure reported by that test from the **Measures** list. Then, click the **OK** button. Upon clicking **OK**, the name of the chosen measure will appear as the title of your comparison graph section (see Figure 10.253).

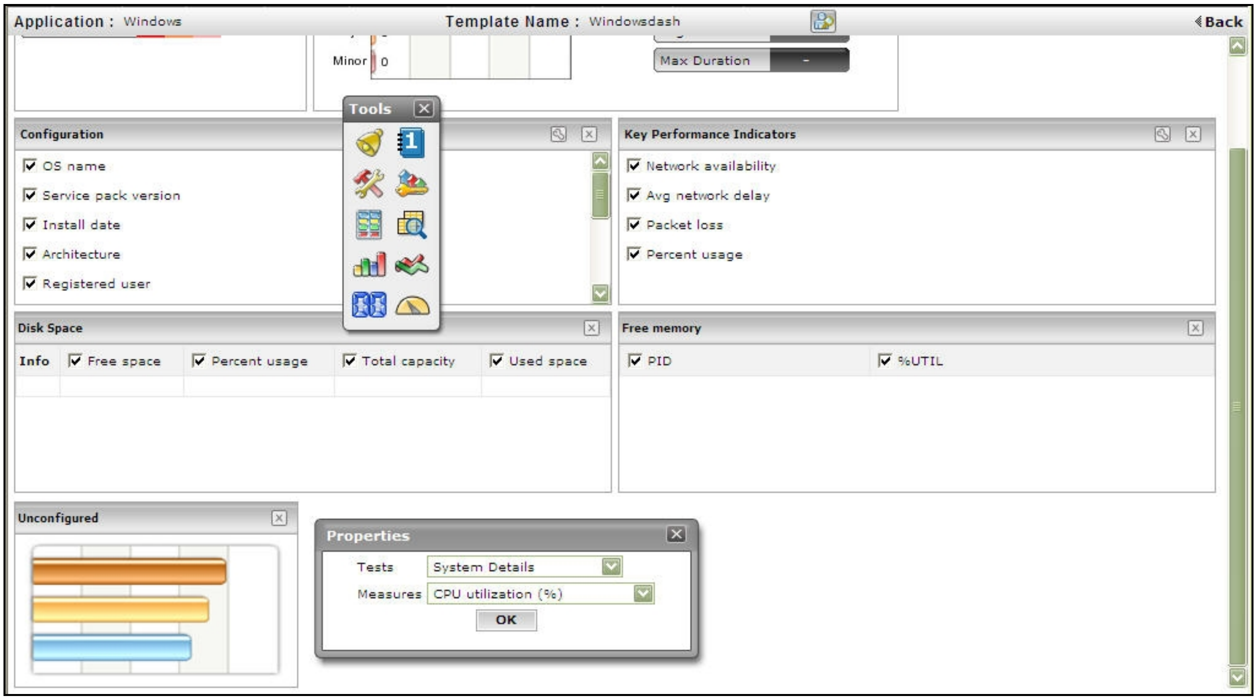

Figure 10.252: Configuring the measure for the comparison graph

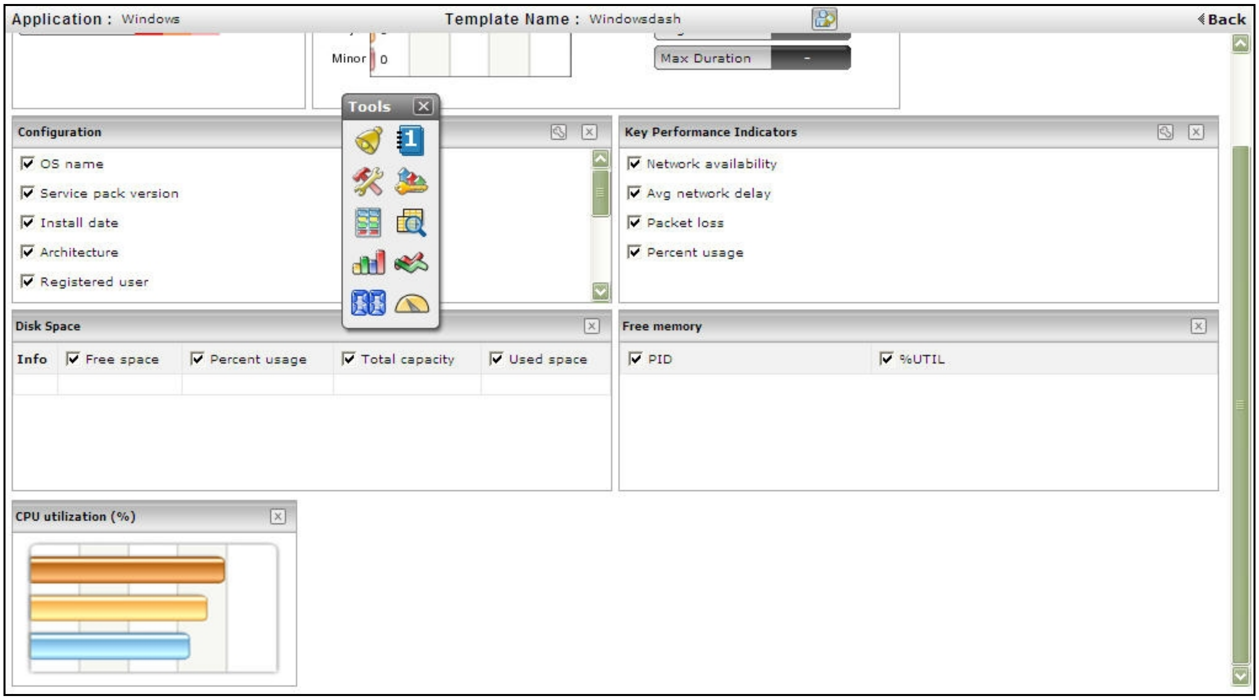

Figure 10.253: The selected measure appearing as the title of the Comparison Graph section

26. To historically analyze performance of a measure, you can include a **Timeline Graph** in your dashboard. For this purpose, click on the tool in the tool box. A measure graph depicted by Figure 10.254 will then get appended to your dashboard as shown by Figure 10.254.

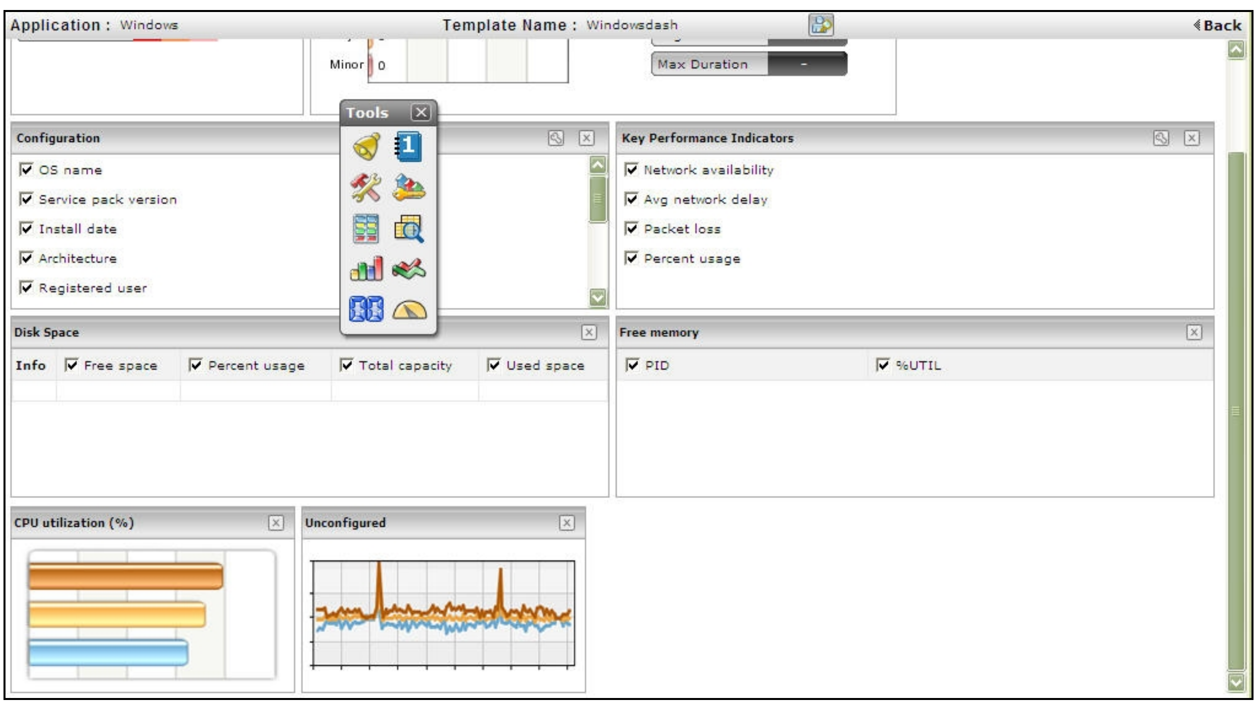

Figure 10.254: Including a timeline graph in the dashboard

27. To configure the measure for which the history graph is to be generated, click on the graph titled **Unconfigured** in Figure 10.254. A **Properties** window will then pop-up as shown by Figure 10.255. From the **Tests** list in the **Properties** pop-up, pick a test, and then pick a measure reported by that test from the **Measures** list. Then, click the **OK** button (see Figure 10.255). Doing so will set the name of the chosen measure as the title of your history graph section.

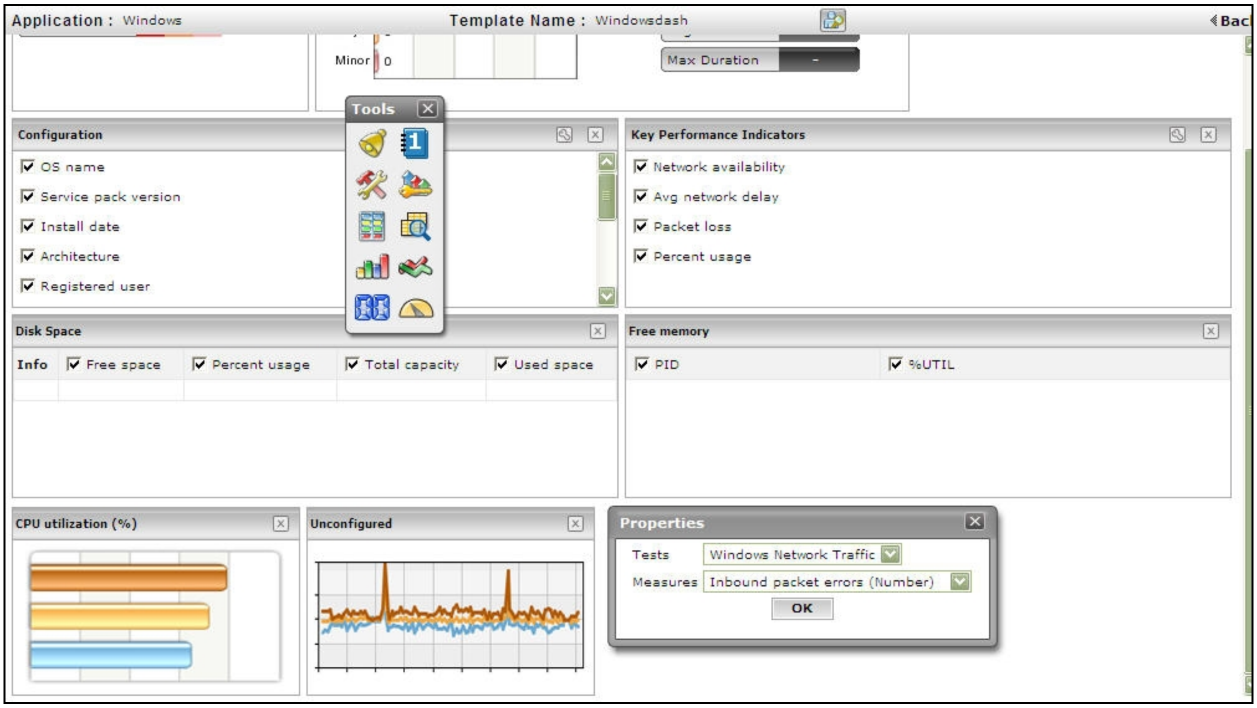

Figure 10.255: Configuring a measure for the history graph in the dashboard

28. To receive quick updates on the current status and values of critical measures, you can include digital displays and/or dial charts in your dashboard for each such measure. To add a digital display to your dashboard, click on the tool in the tool box. A digital graph titled **Unconfigured** will then appear in the custom dashboard (see Figure 10.256). To associate the graph with a particular measure, click on the digital graph in Figure 10.256.

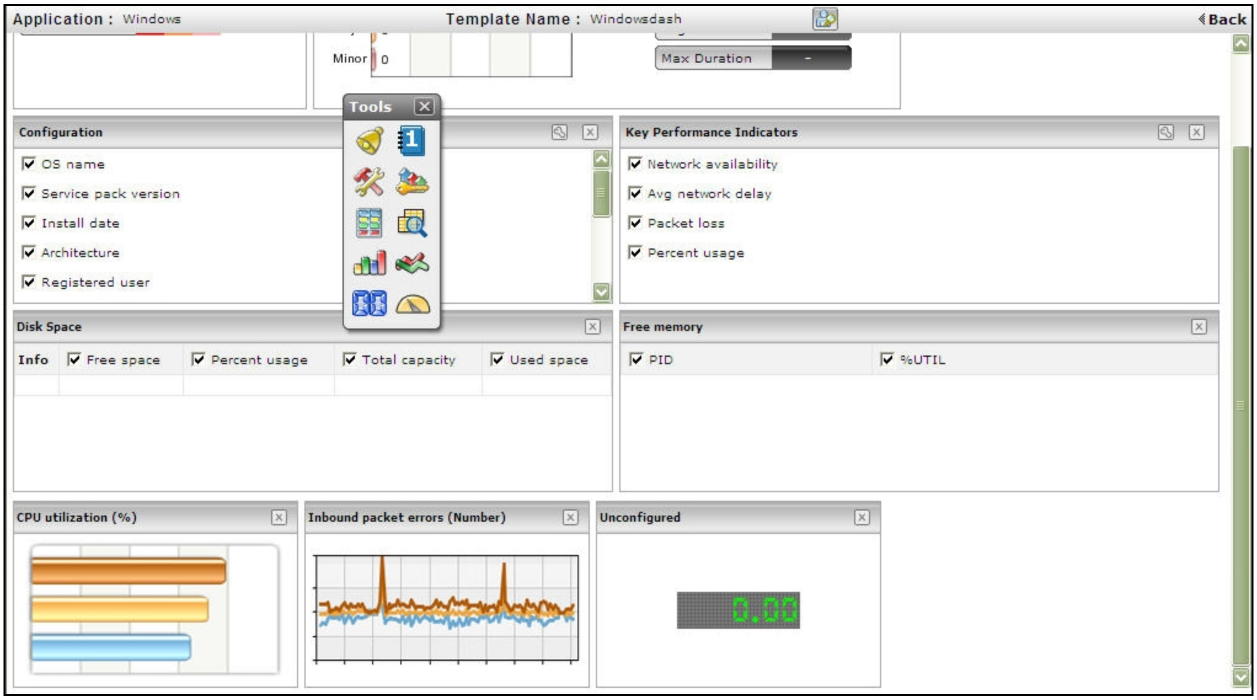

Figure 10.256: Adding a digital graph to the dashboard

29. A **Properties** window will then pop-up as shown by Figure 10.257. From the **Tests** list in the **Properties** pop-up, pick a test, and then pick a measure reported by that test from the **Measures** list. Then, click the **OK** button. Upon picking **OK**, you will find that the title of the digital graph changes to reflect the measure name.

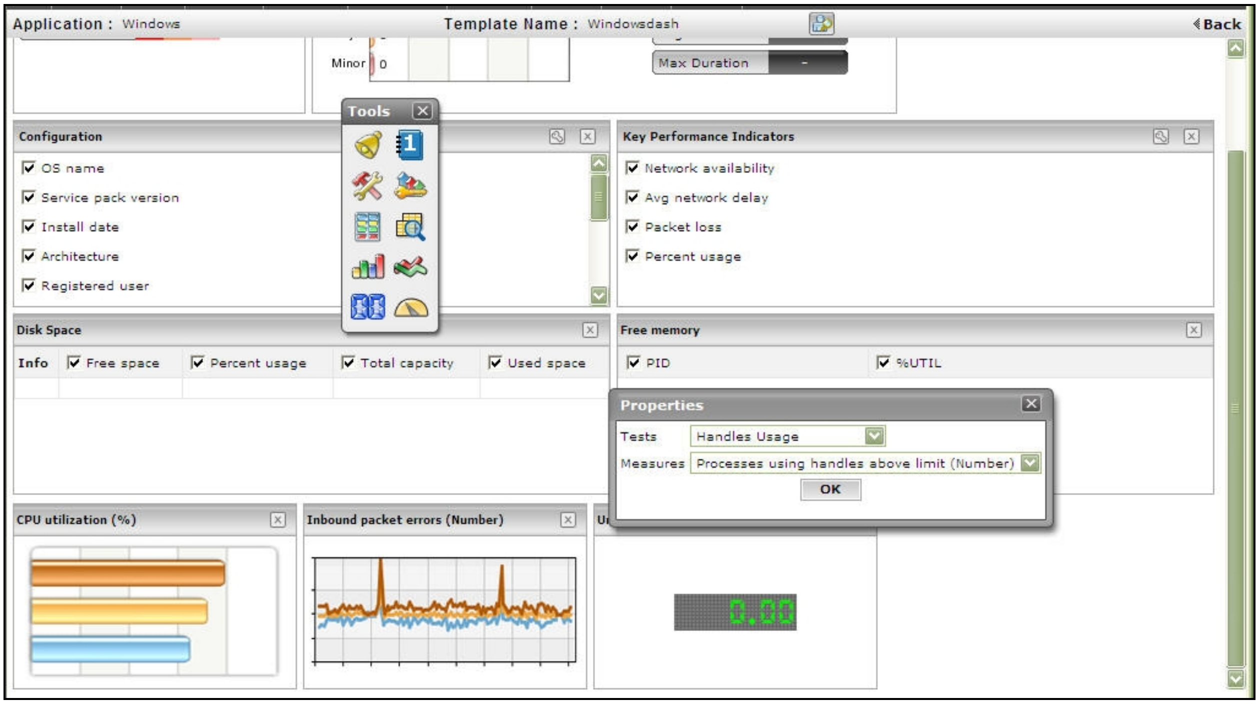

Figure 10.257: Configuring the digital graph in the custom dashboard

30. To insert a dial chart, click on the tool in the tool box. A section titled **Unconfigured** will then appear, as shown by Figure 10.258. For configuring the measure to be represented by the dial chart, click on the **Unconfigured** section in Figure 10.258.

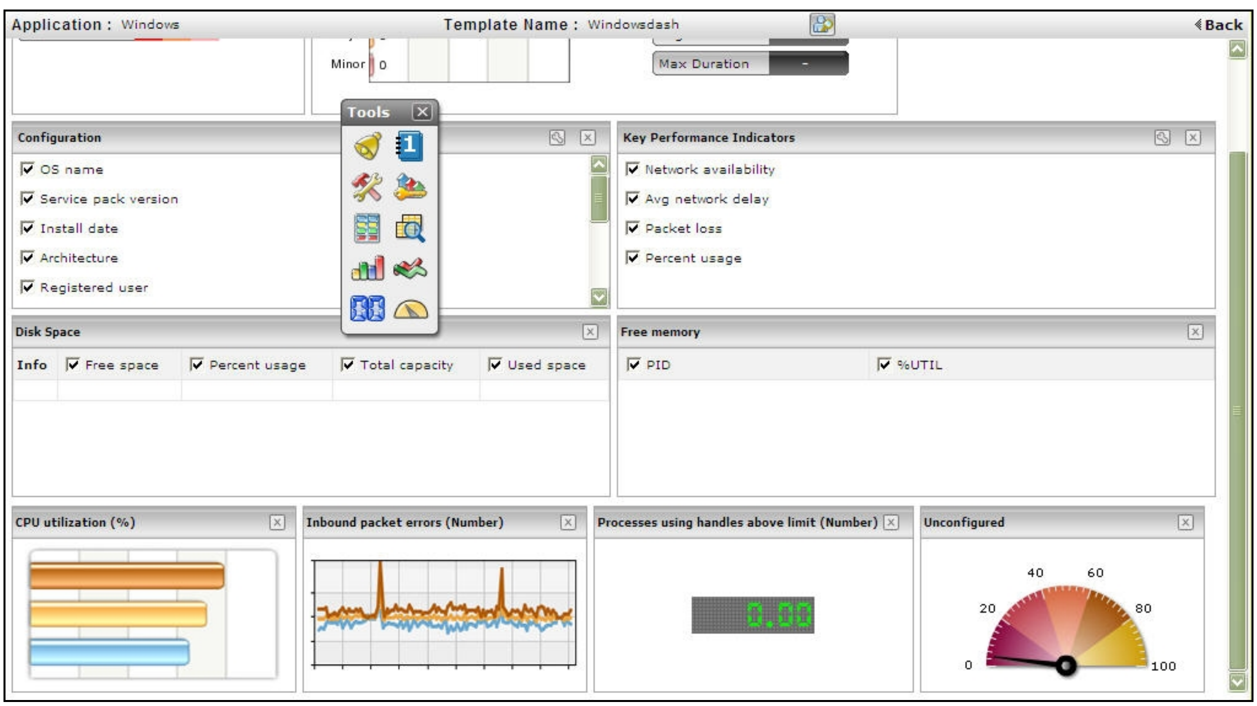

Figure 10.258: Inserting a dial chart into the custom dashboard

31. A **Properties** window will then pop-up as shown by Figure 10.259. From the **Tests** list in the **Properties** pop-

up, pick a test, and then pick a measure reported by that test from the **Measures** list. Then, click the **OK** button. Upon clicking **OK**, you will find that the title of the dial chart changes to reflect the measure name.

Note:

- Dial charts can be configured only for those measures that report percentage values.
- If the **Show threshold** flag is enabled, then, only the **Absolute Thresholds** configured for the chosen measure will be displayed in the dial chart.

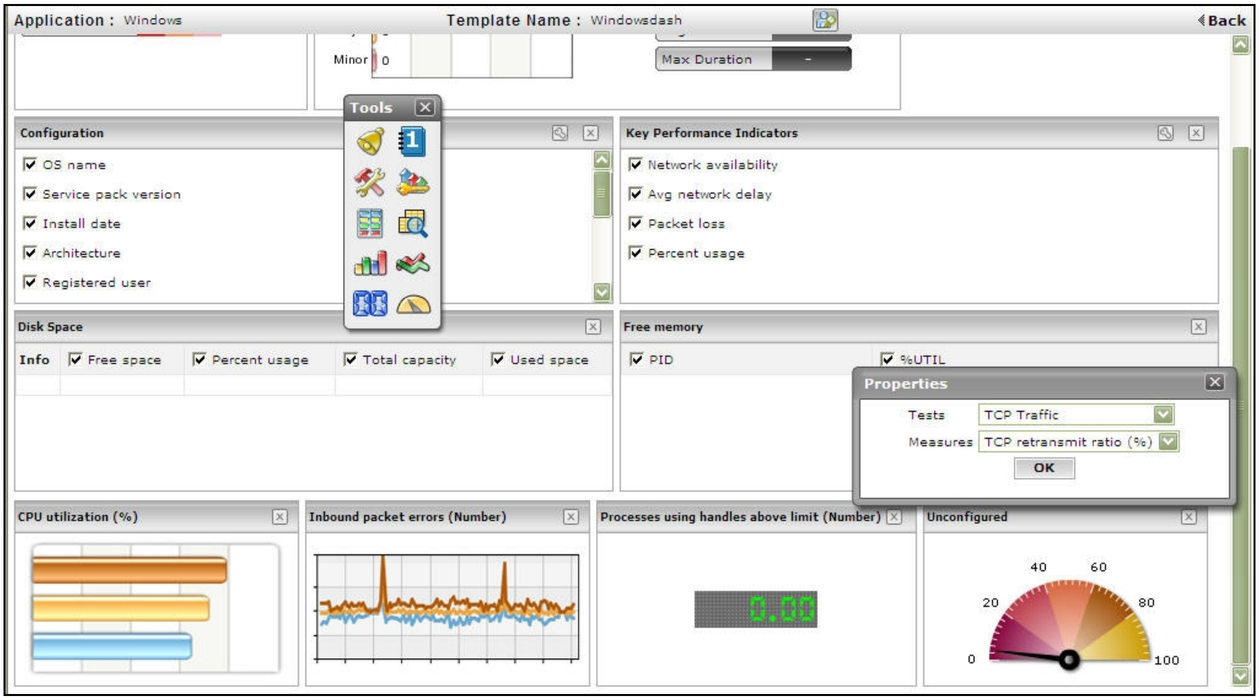

Figure 10.259: Configuring a measure for the dial chart

32. At any time during dashboard design, if you close the tool box by clicking on the **X** button at its right, top corner, you can restore the tool box by clicking on the icon that appears at the top of the custom dashboard.

33. Once you are done designing the dashboard, you can preview it by clicking on the at the top of Figure 10.259. Figure 10.260 and Figure 10.261 depict a sample preview.

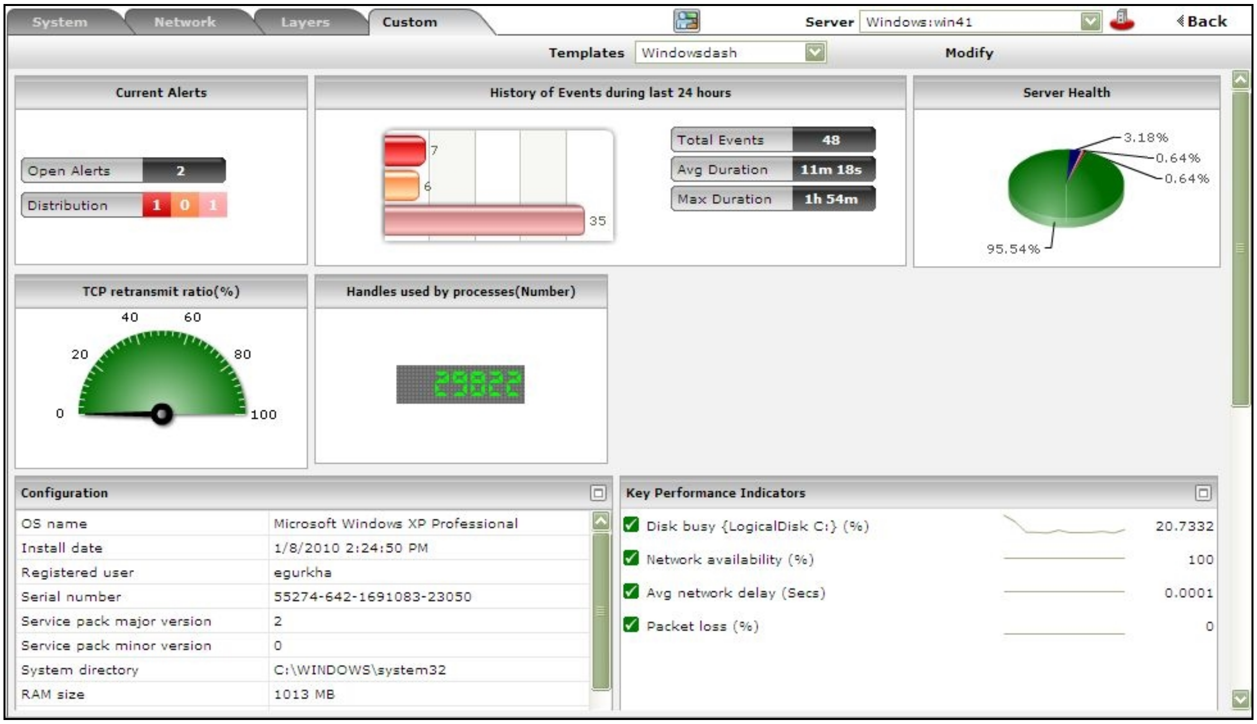

Figure 10.260: A sample preview of the custom dashboard

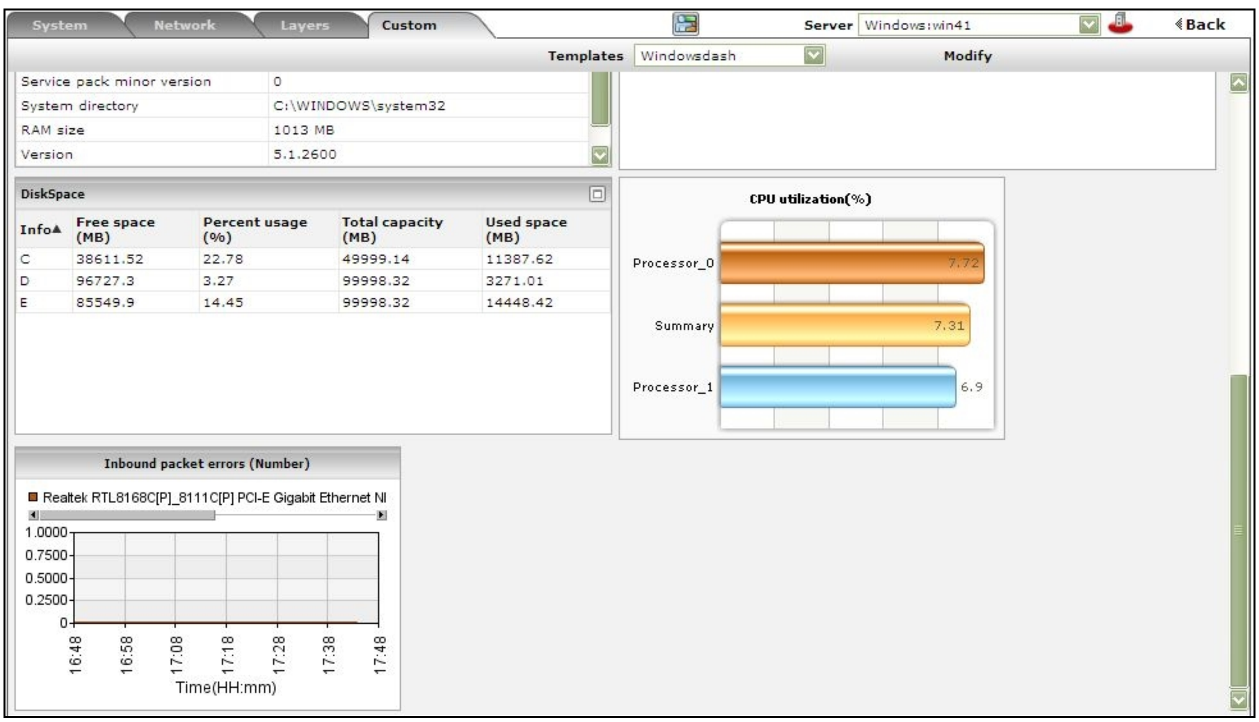

Figure 10.261: A sample preview of the custom dashboard

34. To return to the design mode of the dashboard being viewed, click on the **Modify** button next to the **Templates** list in Figure 10.235. To modify another dashboard, simply select the dashboard name from the **Templates** list and click the **Modify** button next to it.

10.9 Application-Independent My Dashboard

With eG Enterprise, you have the option to build a custom dashboard, using which you can graphically compare current and historical performance of multiple applications of interest to you. This dashboard not only allows you to choose the components to focus on, but also pick the measures and even descriptors that should be featured in it. By default, this dashboard capability is enabled. To build an application-independent custom dashboard therefore, click on the ☰ icon available in the **Monitor** tab and then select the **My Dashboard** option from the **Dashboards** section. To disable this capability, you need to do the following:

- Edit the **[MULTISERVER_CUSTOM_DASHBOARD]** section of the **eg_ui.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
- By default, the **ShowCustomDashboard** flag is set to **Yes**. To disable the capability to view / modify / build new custom dashboards that are application-independent, switch off this flag by setting it to **No**.
- Finally, save the file.

Note:

- By default, the user who creates a dashboard is the only one who is authorized to view / modify / delete that dashboard.
- Users can create custom dashboards using only those components that are assigned to them for monitoring.

Follow the steps given below to design a My Dashboard:

1. Select **My Dashboard** from the **Dashboards** section that appears upon clicking the ☰ icon available against the **Monitor** tab.
2. If no dashboards pre-exist, then "No Dashboards configured! Click here to Add Dashboard" message will appear (see Figure 10.262).

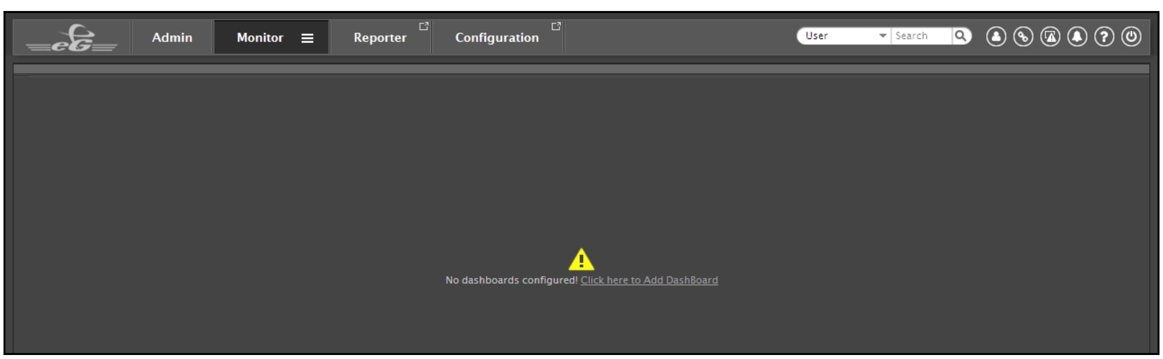

Figure 10.262: The message that appears when no dashboard pre-exists

If multiple dashboards pre-exists, then the dashboard that was created soon after clicking the Click here to Add Dashboard message will be set as the default dashboard.

3. If any dashboards pre-exists for the components that are being monitored, then, upon clicking icon, the **LIST OF MY DASHBOARDS** page as shown in Figure 10.263 will appear listing all the pre-existing dashboards. The eG Enterprise system automatically sorts these dashboards in the ascending order of their names while listing.

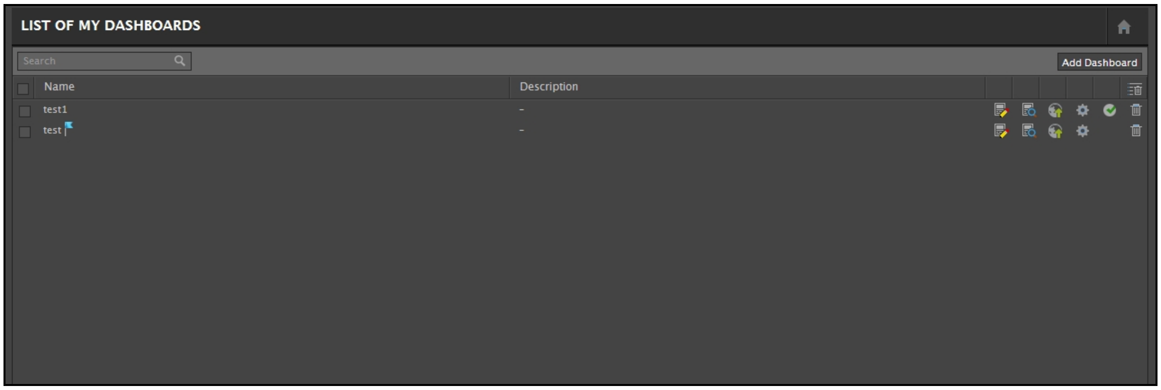

Figure 10.263: The List of My Dashboards page

4. If other users to the eG Enterprise system have shared one/more dashboards that they created with you, then such dashboards will be listed in the **Dashboards shared by other users** section. Besides listing the names of the shared dashboards, this section will also display (in the **Shared by** column) who shared each of the listed dashboards with you. Shared dashboards can only be viewed (not modified or deleted). To view a shared dashboard, click on the **Name** of the dashboard, or click on the icon that corresponds to each dashboard.

5. To create a new dashboard, click on the **Add Dashboard** button in the **LIST OF MY DASHBOARDS** page (see Figure 10.263). This will lead you to Figure 10.264 where you are required to enter the credentials for creating a new dashboard.

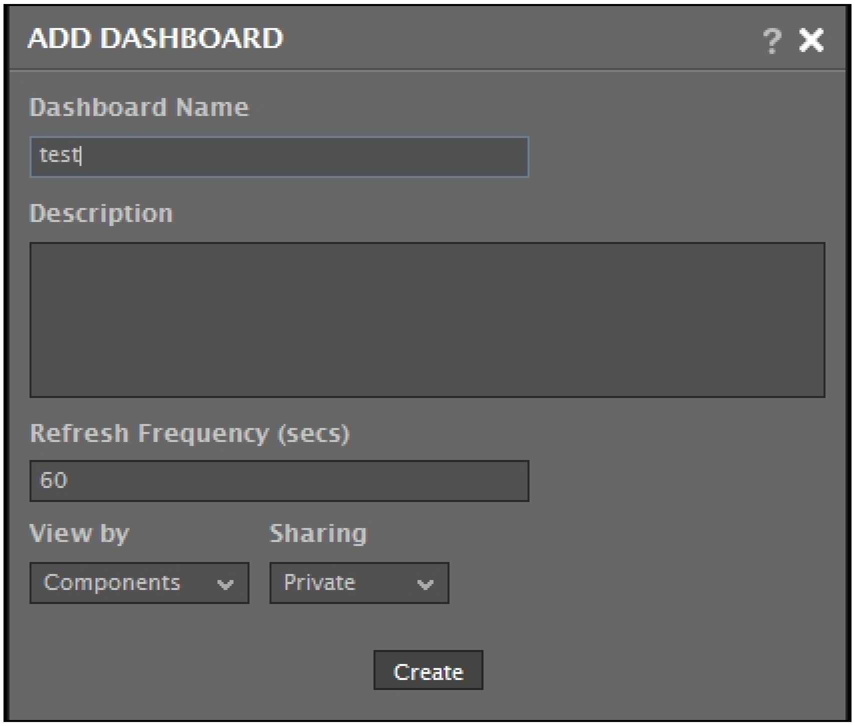

Figure 10.264: Adding a new dashboard

6. In Figure 10.264, specify the following:

- **Dashboard Name**: Provide a name for the new dashboard.
- **Description**: Here, provide a brief description of the new dashboard.
- **Refresh frequency :** Specify the time interval after which the dashboard needs to be refreshed. By default, this is **60** seconds.
- **View by** : To categorize the dashboard, or to view only the components available in the **Zone/Segment/Services**, select the option under the **View by** dropdown. For example, selecting the **Service** option will list only the components that are part of a service. To view all the components in the environment, click on the **Components** option. By default, the **Components** option will be chosen in the **View by** list.
- **Sharing**: Here, indicate whether or not you want to share the new dashboard with other users, and if so, what should be the level of sharing. For this, pick any of the following options from the **Sharing** drop-down list. By default, the **Private** option is chosen from this list.
  - **Private**: Indicates that the user building the dashboard is alone authorized to view/modify/delete the dashboard. This is the default option.
  - **Public**: Allows all users to the eG Enterprise system to view (not modify/delete) the dashboard that is being created.
  - **Share:** You can select this option to share the dashboard with specific users to the eG Enterprise system. From the **Available Users** list that then appears (see Figure 5), select the users with whom you want to share the dashboard, and click the **>** button. This will transfer the selection to the

Selected Users list. To revoke the share, select one/more users from the **Selected Users** list and click the **<** button. You can also provide read/write access to the users with whom you have shared the dashboard by checking the **Allow R/W sharing** checkbox. The **Available users for R/Wsharing** list will then be populated with the users with whom you have shared the dashboard. Select the users to whom you wish to provide the read/write access from this list.

- Finally, click the **Create** button to create the new dashboard.

7. To modify a pre-existing dashboard, click on the icon that corresponds to each dashboard in the **LIST OF MY DASHBOARDS** page or by clicking the icon while editing the dashboard. A **MODIFY DASHBOARD** pop up window will then appear as shown in Figure 10.265 using which you can change the name of the dashboard, the frequency with which the dashboard is to be refreshed etc.

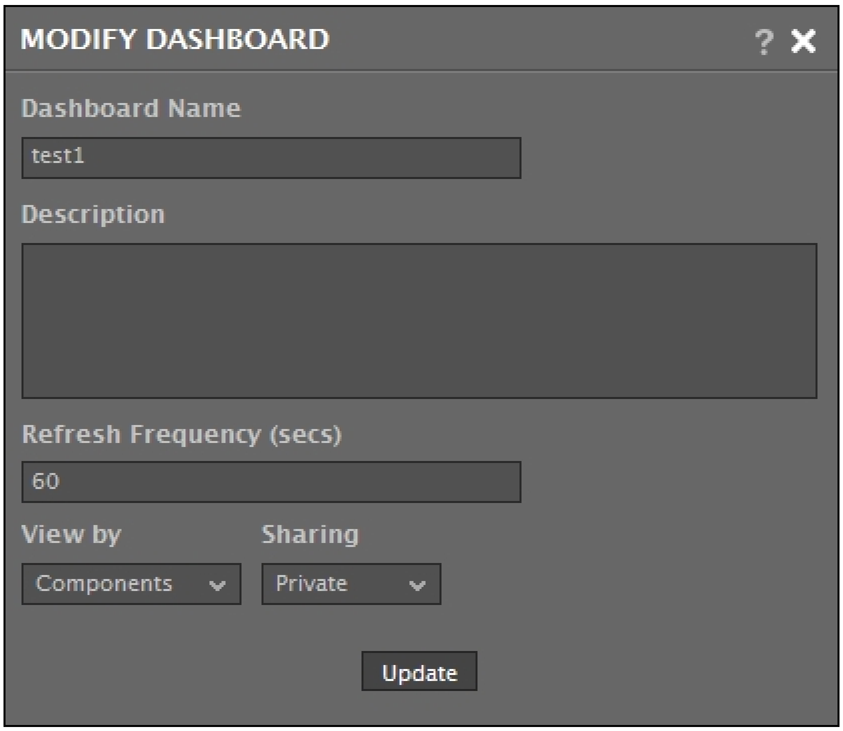

Figure 10.265: Modifying an existing dashboard

If required, you can even modify how a dashboard is to be shared with other users registered with the eG Enterprise suite. For this, you will have to pick an option from the **Sharing** drop down list (see Figure 10.266). Finally, clicking the **Update** button will help you apply the necessary changes made.

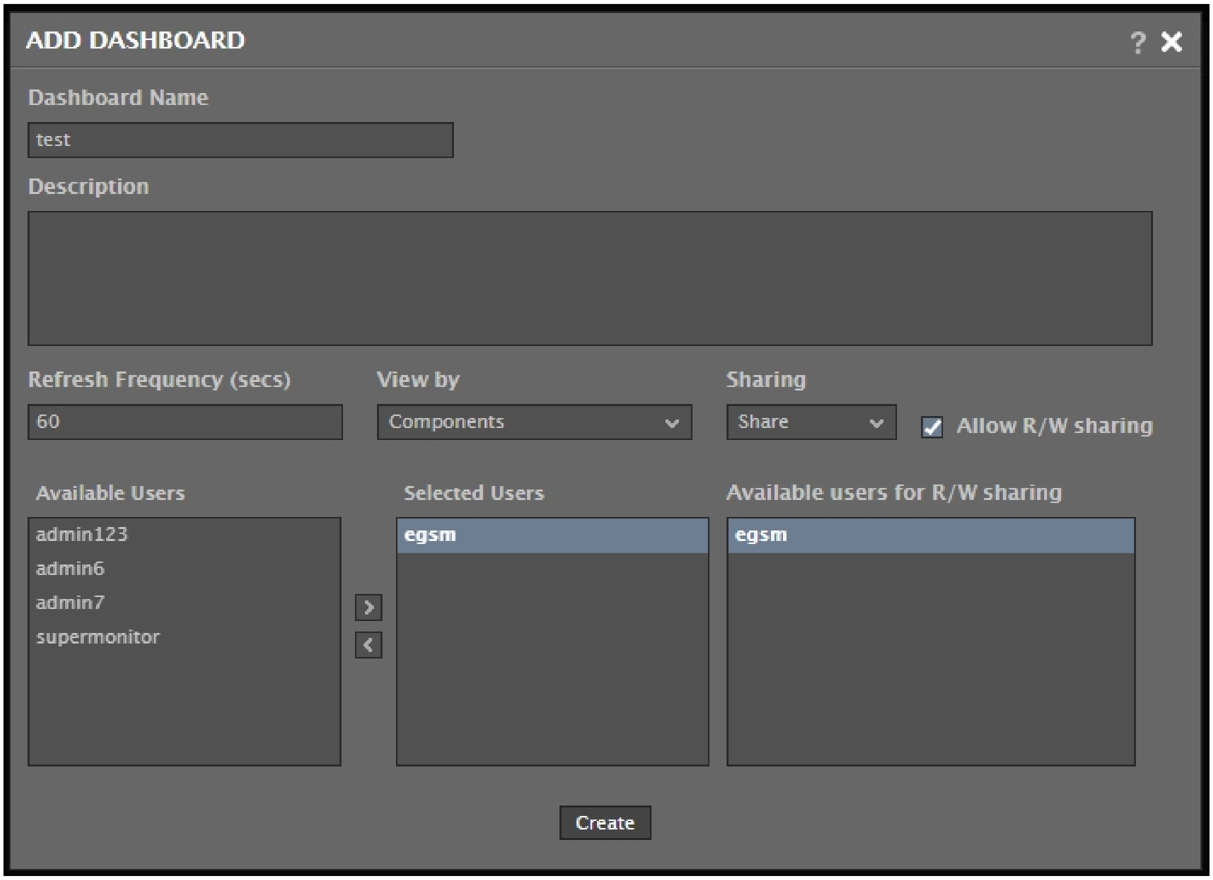

Figure 10.266: Sharing the dashboard with the users of your choice

8. To view any of the existing dashboards, click on the corresponding to each dashboard in the **LIST OF MY DASHBOARDS** page (see Figure 10.263). If you want to delete any of the listed dashboards, click the icon corresponding to each dashboard (see Figure 10.263). To remove all dashboards, select the check box that pre-fixes each dashboard, and then click the icon (see Figure 10.263). To search for a dashboard from the list of pre-existing dashboards, enter the name of the dashboard in the **Search** text box and click the icon within (see Figure 10.263).
9. To publish any of the existing dashboards through Microsoft SharePoint to other users in your organization, click on the icon corresponding to each dashboard (see Figure 10.263). Figure 10.267 will then appear listing the **Name** and the **Publish Url** of the dashboard that is to be published.

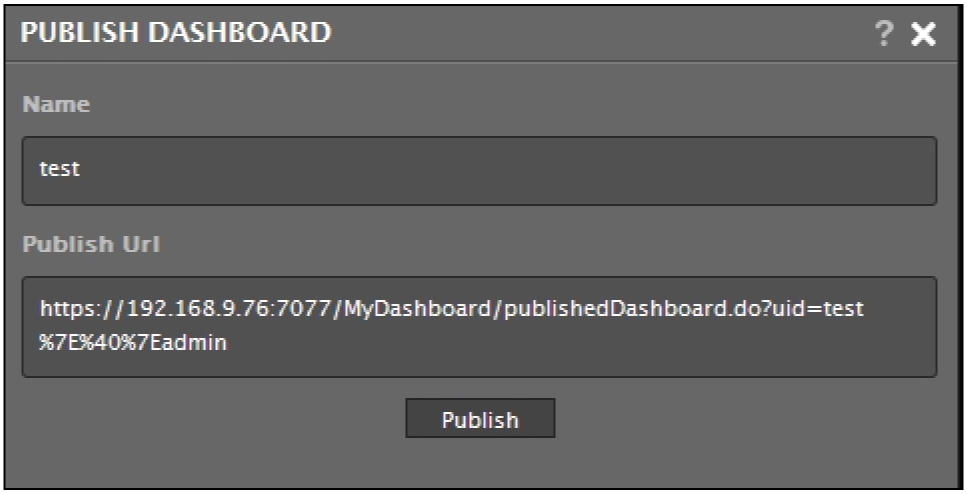

Figure 10.267: The Publish Dashboard pop up window

Clicking the **Publish** button will help you publish the chosen dashboard. Once the dashboard is published, the icon will appear along with the **Name** of the dashboard as shown in Figure 10.268.

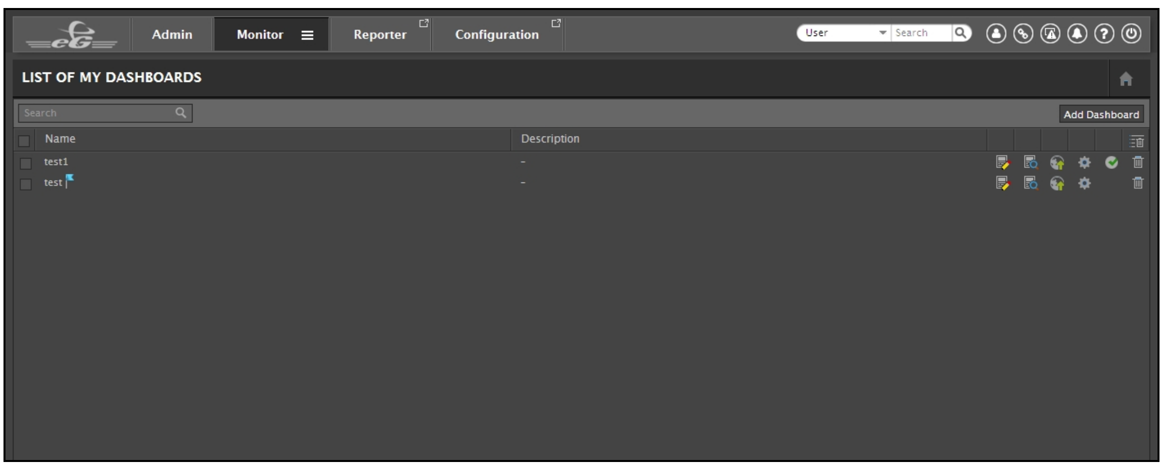

Figure 10.268: An icon appearing next to the published dashboard

You can hover over this icon to view the Publish url at any point of time. Likewise, you can unpublish a published dashboard by clicking the corresponding to each published dashboard. Figure 10.269 will then appear. Clicking the **Unpublish** button in the **UNPUBLISH DASHBOARD** pop up window will unpublish the dashboard.

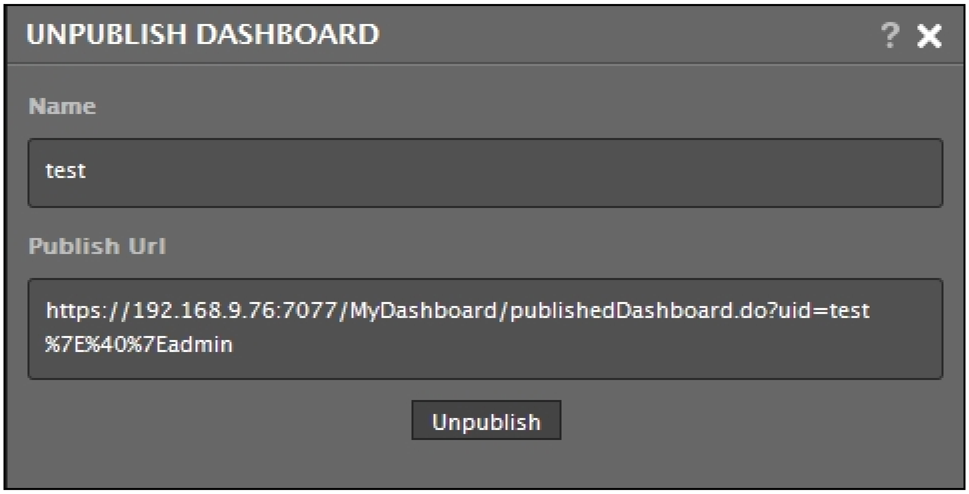

Figure 10.269: Unpublishing the dashboard

10. To edit a pre-existing dashboard, click on the icon corresponding to each dashboard in the **LIST OF MY DASHBOARDS** page (see Figure 10.263). Figure 10.270 will then appear.

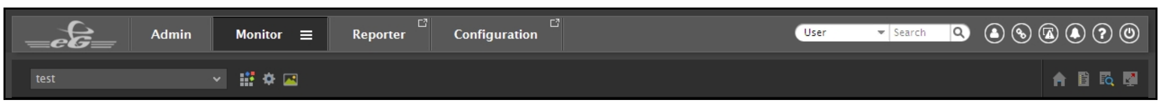

Figure 10.270: Icons that are used for designing the dashboard

11. A set of icons will appear as shown in Figure 10.270 using which you can design the dashboard of interest to you.

12. To design a dashboard, you need to click on the icon. Figure 10.271 will then appear using which you can customize the dashboard of your interest. Section Section **10.9.1** will explain in detail about how to design a dashboard.

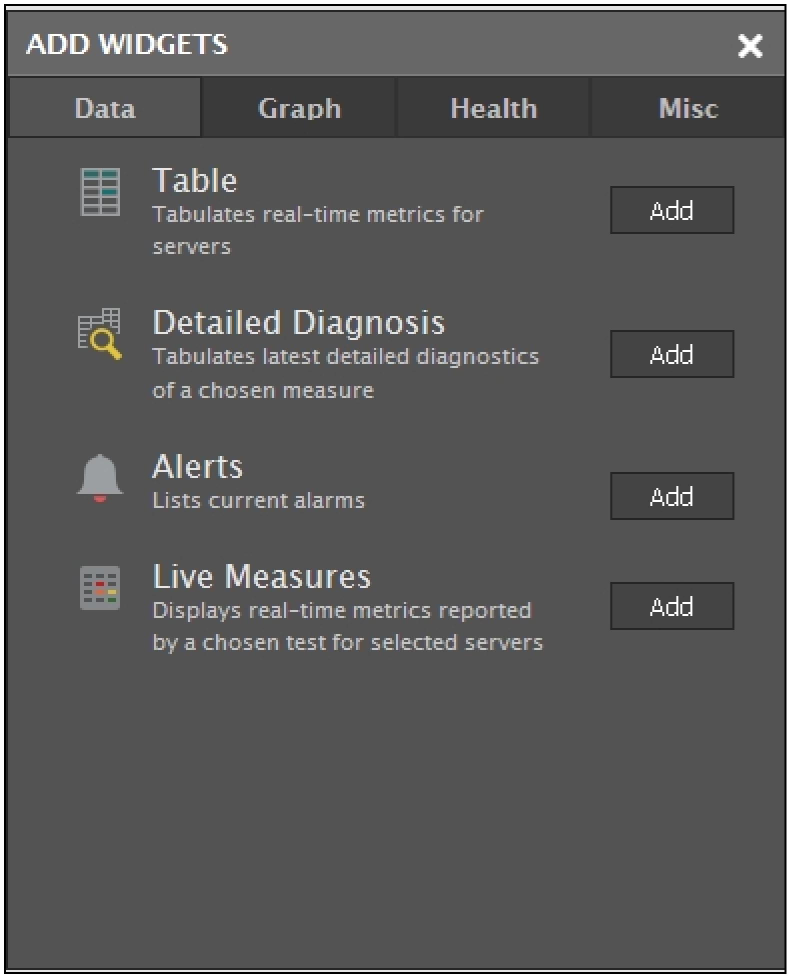

Figure 10.271: The Add Widgets pop up window used for building the dashboard

13. To modify the dashboard on the fly, click on the icon. The **MODIFY DASHBOARD** pop up window (see Figure 10.265) will then appear allowing you to modify the settings of the dashboard that is currently being built.

14. You can even personalize the background of your dashboard by clicking the icon. Figure 10.272 window will then appear listing a set of background themes.

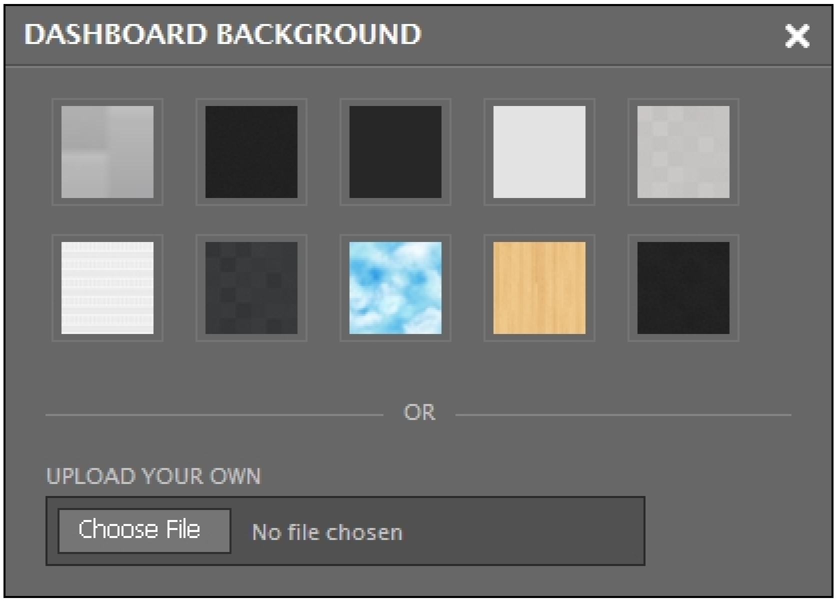

Figure 10.272: Changing the background of the dashboard

Clicking on any background theme will change the dashboard background accordingly. You can also import any image of your interest as a dashboard's background to reflect your interest or mood by simply uploading an image using the **Choose file** option available in the **UPLOAD YOUR OWN** section.

15. If the dashboard that you are currently building is not the default dashboard, then an additional icon will appear which when clicked will set that particular dashboard as the default dashboard (see Figure 10.273).

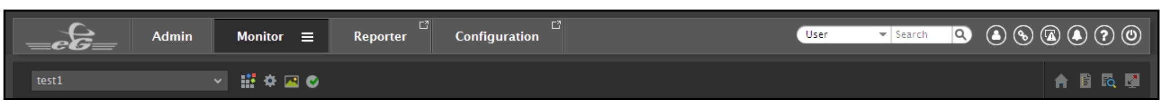

Figure 10.273: Changing the default dashboard

16. If you wish to open the dashboard in a separate window, then you can do so by clicking the icon (seeFigure 10.273).

10.9.1 Designing a My Dashboard

To design a **My dashboard** of interest to you, click on the icon in Figure 10.270. Figure 10.271 will then appear. Using the widgets mentioned in Figure 10.271, you can design the look and feel of the new dashboard.

By default, the **ADD WIDGETS** pop up window consists of the following four tabs:

- Data
- Graph
- Health
- Misc

Each tab in this window consists of a group of widgets that are used in building the dashboard. Let us now discuss in detail about each tab and the widgets that are currently available:

10.9.1.1 The Data Tab

If you wish to compare the values of metrics for a wide range of components or if you wish to compare the detailed diagnosis of various measures or if you wish to view the list of alarms for a chosen measure, then you can use the widgets from the **Data** tab. The forthcoming sections will explain in detail about each widget available in the **Data** tab.

1. Table

   If you want to compare the values of metrics collected from a wide array of components in a tabular column, then, you can use the **Table** widget. To add a **Table** widget, click on the **Add** button against the **Table** option in the **Data** tab. A blank **Table** widget will be appended to the dashboard as shown in Figure 10.274.

Figure 10.274: The Table widget

Clicking on the  icon in this widget will invoke Figure 10.275 using which you can configure the metrics that are to be displayed in this widget.

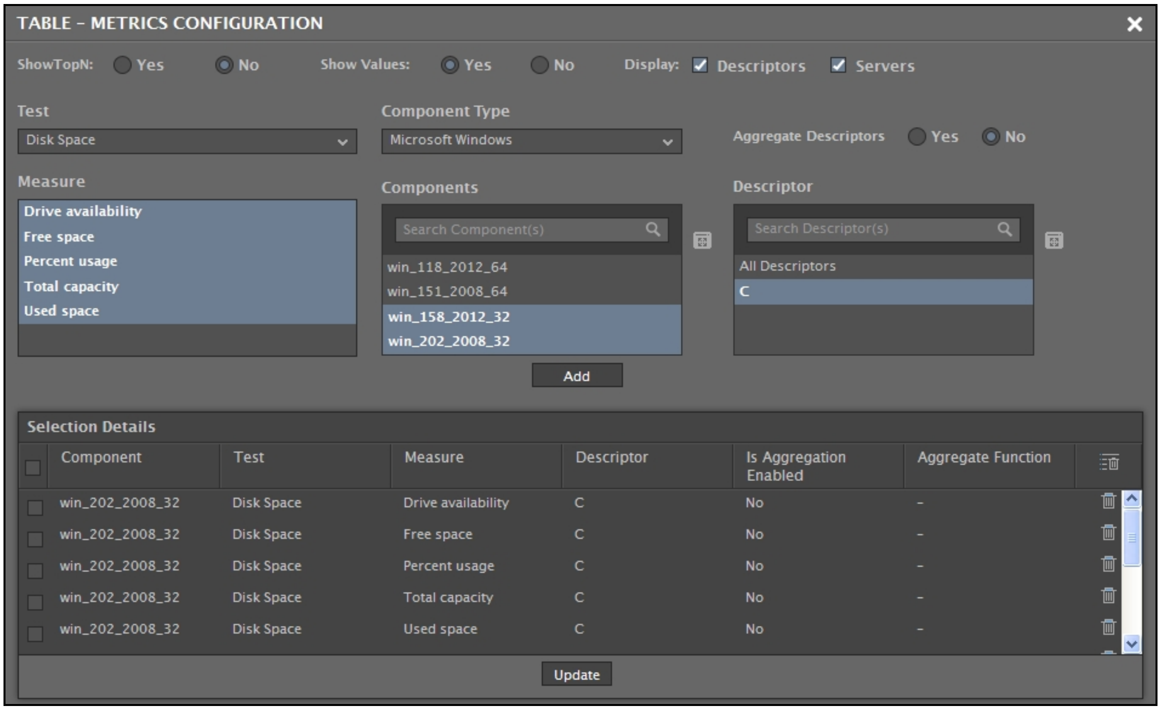

Figure 10.275: Configuring metrics that are to be shown in the Table widget

Let us now discuss on how to configure metrics in the Table widget using Figure 10.275:

- Select a **Test**, **Component Type**, **Component** and **Measure** for which you wish to configure the **Table** widget. Once you have chosen the **Component**, the descriptors of the chosen component will be displayed in the **Descriptor** section. Select the descriptor of your choice from this list. If the **Components** list consists of too many components, then viewing all the components and selecting the ones you need for configuring the widget could require endless scrolling. To avoid this, you can click the  icon next to the **Component** list. Figure 10.276 will then appear using which you can view almost all the descriptors in a single interface and **Select** the ones that you wish to view in the **Table** widget.

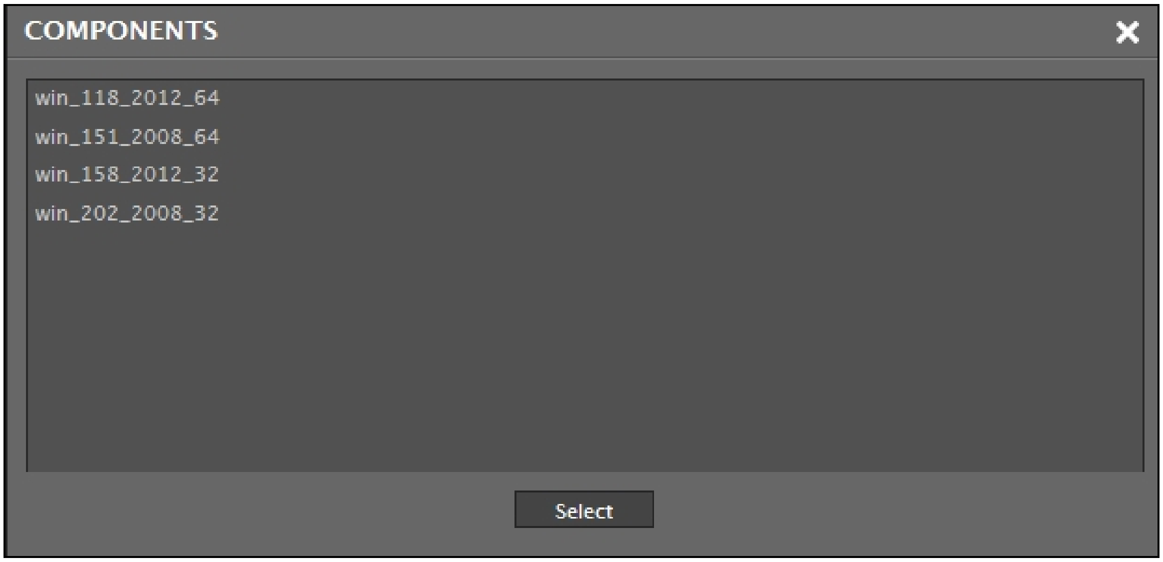

Figure 10.276: The COMPONENTS pop up window

- You can even search for a component of your search using the **Search Component(s)** field in the **Components** list.
- If the **Descriptor** list consists of too many descriptors, then viewing all the descriptors and selecting the ones you need for configuring the widget could require endless scrolling. To avoid this, you can click the  icon next to the **Descriptor** list. Figure 10.276 will then appear using which you can view almost all the descriptors in a single interface and **Select** the ones that you wish to view in the **Table** widget.

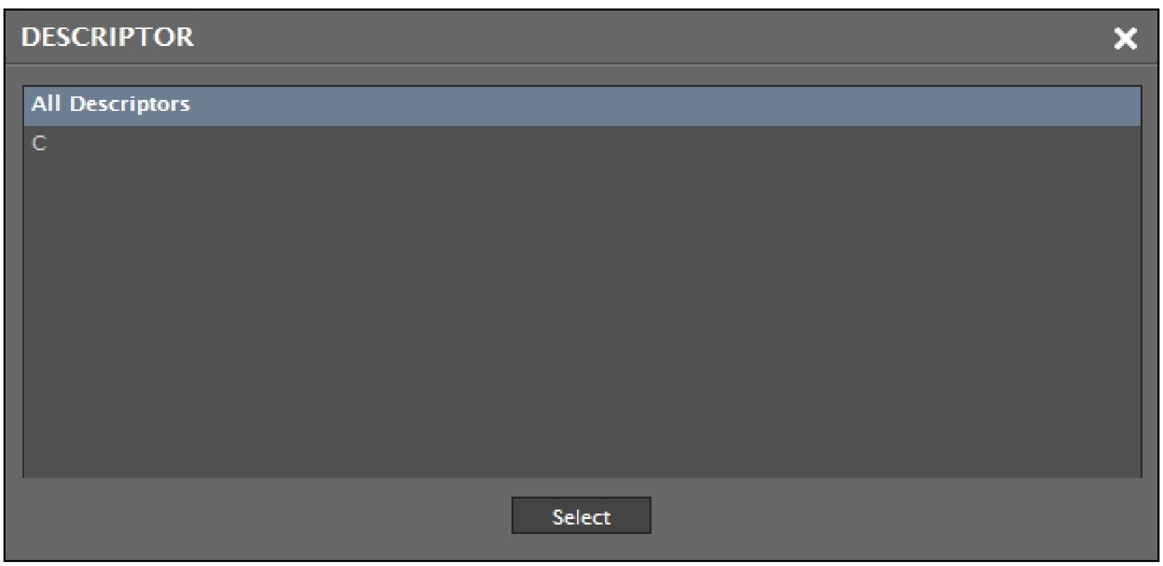

Figure 10.277: The DESCRIPTOR pop up window

- You can even search for a descriptor of your search using the **Search Descriptor(s)** field in the **Descriptor** list.
- If descriptors exist for a chosen component, an **Aggregate Descriptors** flag will appear as shown in

Figure 10.278. By default, this flag is set to **No**. If this flag is set to **Yes**, then an **Aggregate Function** drop down list will appear with the functions using which you can aggregate the chosen measures and display the same in the **Table** widget.

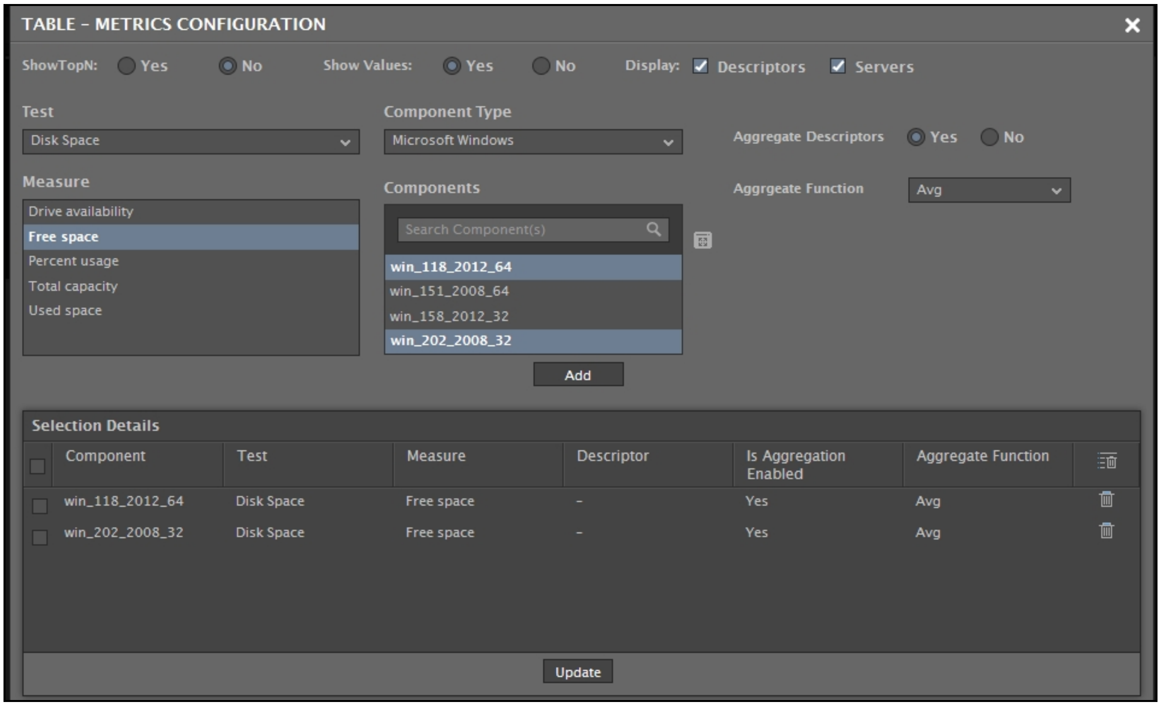

Figure 10.278: The Aggregate Descriptors flag that appears when a descriptor is chosen

- If you wish to see the Top N values of the chosen measures in the **Table** widget, then set the **Show TopN** flag (seeFigure 10.275) to **Yes**. By default, this flag is set to **No**.
- By default, the **Show Values** flag (see Figure 10.275) is set to **Yes** indicating that the real time values of the chosen measures are displayed in the **Table** widget. If this flag is set to **No**, then the state of the chosen measure during the last measurement period for each chosen component is displayed as against the real time values.
- By default, the **Servers** option is checked against the **Legend Display** field in Figure 10.275. If you wish to view descriptor-wise measure values, then you can check the **Descriptors** option against the **Legend Display** field.
- Clicking the **Add** button in Figure 10.275 will display the selections that you have made in the **Selection Details** section (see Figure 10.275).
- If you wish to delete a selection, then you can do so clicking the icon. If you wish to delete multiple selections at a single shot, then select the check box against the selections and click the icon.
- Clicking the **Update** button in Figure 10.275 will populate the **Table** widget with the values of the chosen measures in a tabular format as shown in Figure 10.279.

Figure 10.279: The Table widget that is designed as per your choice

- If you have set the **Show TopN** flag to **Yes**, then an additional **Weighted Average** flag will appear as shown in Figure 10.280.

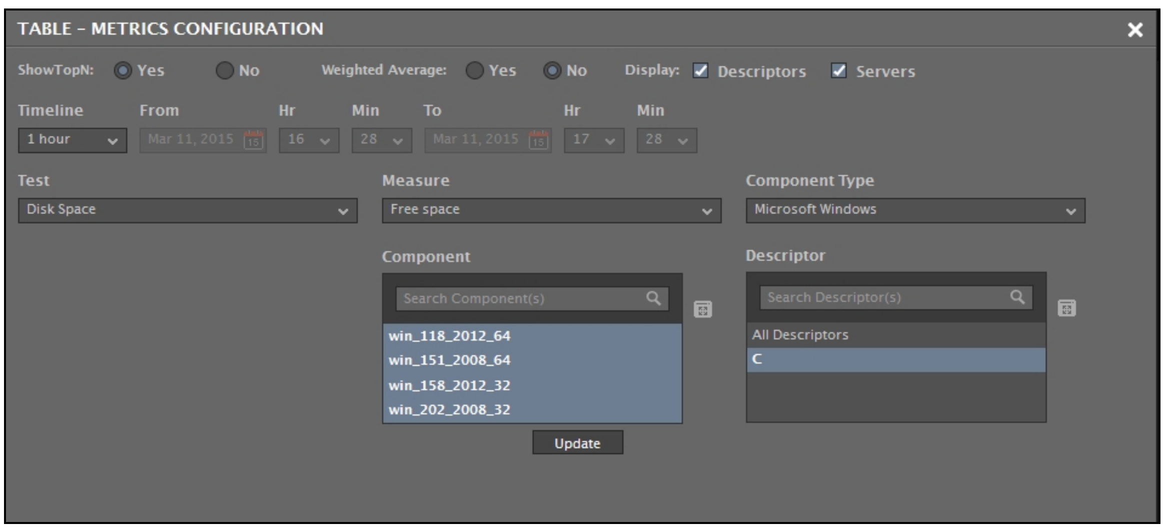

Figure 10.280: When the Show TopN flag is set to Yes while designing a Table widget

Using the **Weighted Average** flag, you can indicate how the **Avg** value is to be computed for a chosen measure. The status of the **Weighted Average** flag is relevant only if the **Test** chosen is a descriptor-based test, and the descriptors are dynamic in nature. For example, the `Net flows` test auto-discovers the net flows that are currently active on a Cisco router. For each net flow (i.e., source IP - destination IP pair) so discovered, the test reports a variety of statistics. Unlike descriptors such as disk partitions or processors that rarely change, the net flows are dynamic descriptors, which may change often; in other words, a net flow that is active on the network now, may not communicate at all during the next measurement period. The Net flows test will neither report metrics for the inactive descriptors nor display it in the eG monitoring console. This is why, net flows are considered 'dynamic descriptors'.

- By default, when the **Show TopN** flag to **Yes**, the `Avg` value of the chosen measure is compared across all selected descriptors. Since the **Weighted Average** flag is set to **No** by default, this `Avg` is computed as the ratio of the sum total of the measure values reported by a descriptor during the given timeline and the total number of times the test ran during the same timeline. In case of dynamic descriptors however, the `Avg` values so computed may not reveal the 'true picture of performance'. This is because, the test may not discover or report metrics for dynamic descriptors throughout a given timeline. For example, take the case of the `Net flows` test. Say, the test auto-discovered two net flows - namely, 'A' and 'B' - during its first measurement period; both net flows registered a Data in Flow of 3 MB and 8 MB respectively. Assume that the second time the Net flows test ran, it captured 5 MB of data in net flow 'A'. Net flow 'B' however was inactive during the second measurement period, and hence, was not discovered at all. If the Weighted Average flag is set to No by default, then, the TOP-N report will plot the Avg value of 4 MB (3+5=8/2=4) for net flow 'A' and 4 MB (8/2=4) again for net flow 'B'. If you notice, unlike net flow 'A', where 8 MB of data was transacted over a period of time, in case of net flow 'B', 8 MB of data was transacted at one shot! Logically therefore, net flow 'B' has to be ranked above net flow 'A' in terms of data in flow. However, since the default Avg value computation does not clearly bring out this difference, both net flow 'A' and 'B' are treated at par! This is why, in case of dynamic descriptors, you may want to set the **Weighted Average** flag to **Yes**. In this case, the eG Enterprise system expresses `Avg` as the ratio of the sum total of the measure values reported by a descriptor during a given timeline and the 'total number of times that descriptor was active' during the same timeline. This implies that if the **Weighted Average** flag is set to **Yes** in the net flow example above, the `Avg` value for net flow 'A' will continue to be 4 (3+5=8/2=4), but the same for net flow 'B' will be 8 MB (8/1=8). In the **Table** widget report therefore, net flow 'B' will be placed above net flow 'A', thereby accurately pointing you to the top communicators over the network.

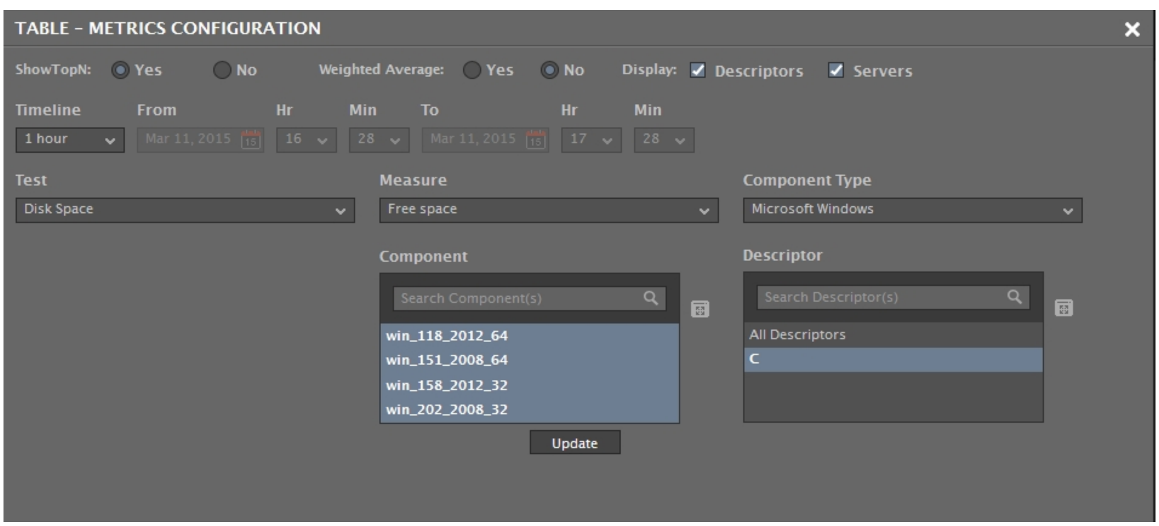

Figure 10.281: Setting the ShowTopN and Weighted Average flags to Yes

- If you have set the **Show TopN** flag to **Yes**, chosen a **Test** and **Measure** to be displayed in the **Table** widget as shown in Figure 10.281, then an additional bar graph will appear against each chosen component in the **Table** widget as shown in Figure 10.282.

Figure 10.282: The Table widget when both the ShowTopN and Weighted Average flags are set to Yes

- If you wish to edit your selection, then you can click on the icon that appears in the right top corner of the **Table Header** widget once you have configured the widget.

2. Detailed Diagnosis

   If you wish to periodically monitor the detailed diagnosis of a particular measure in your environment, then you can use the **Detailed Diagnosis** widget. Clicking the **Add** button against the **Detailed Diagnosis** option in the **Data** tab of Figure 10.271 will append the **Detailed Diagnosis** widget to the dashboard see Figure 10.283.

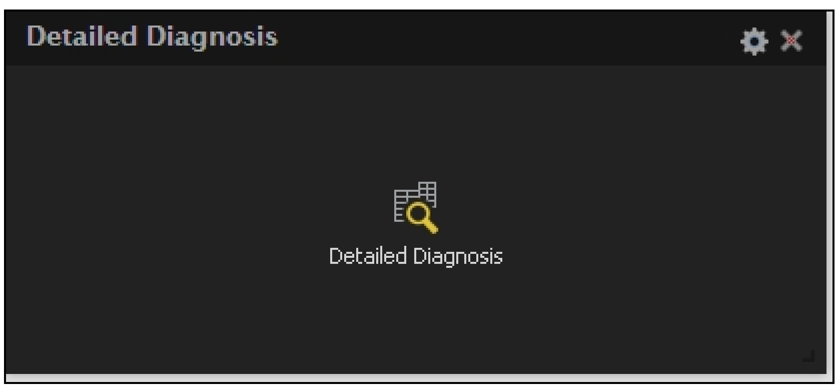

Figure 10.283: The Detailed Diagnosis widget

Now, clicking the  icon in the widget will lead you to Figure 10.284.

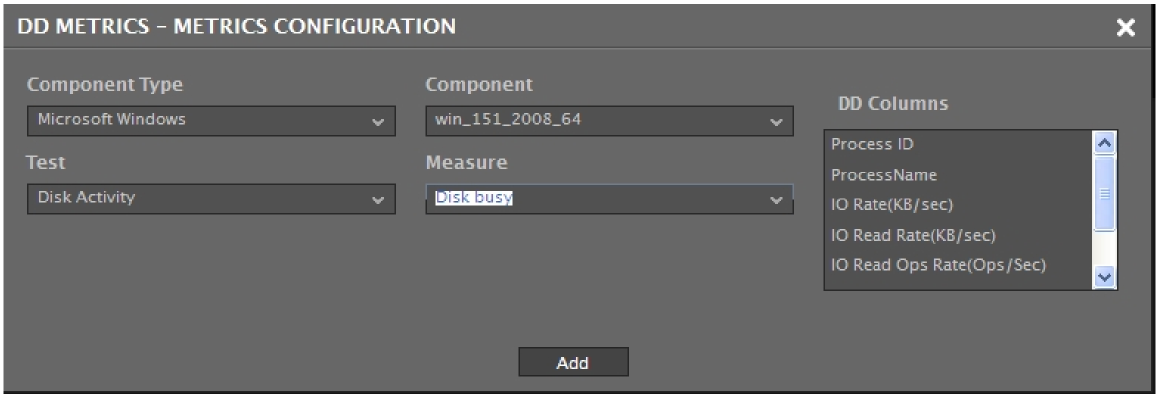

Figure 10.284: Configuring the Detailed Diagnosis widget

Select the **Component Type, Component**, **Test** and **Measure** for which you wish to view the detailed diagnosis in Figure 10.284. The **DD Columns** will then be populated with the attributes that are available for the chosen measure. You can either select all the attributes or select the ones that are of interest to you. Clicking the **Add** button in Figure 10.284 will list your selection in the **Detailed Diagnosis** widget as shown in Figure 10.285.

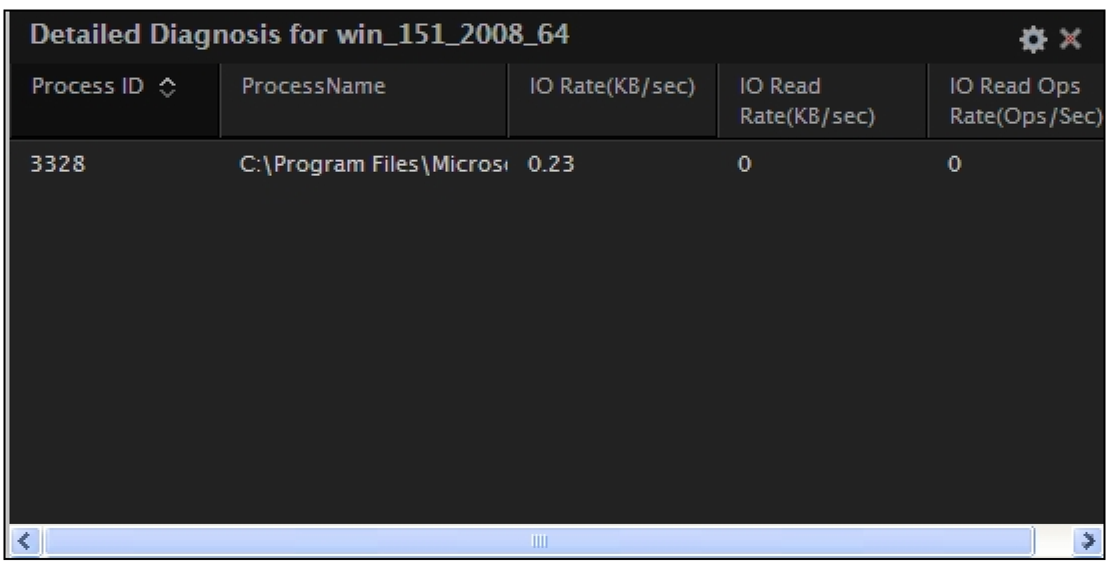

Figure 10.285: The Detailed Diagnosis widget designed based on your choice

Note:

The detailed diagnosis of a single measure can alone be viewed in the **Detailed Diagnosis** widget. If you wish to view the detailed diagnosis for multiple measures, then you need to include multiple **Detailed Diagnosis** widgets in your dashboard.

3. Alerts

To view the list of currently open problems in your environment, you can use the **Alarms** widget. To include an **Alarms** widget in the dashboard, click on the **Add** button against the **Alerts** option of the **Data** tab (see Figure 10.271). The **Alarms** widget will then be appended to the dashboard with all the alarms that are currently open in your environment as shown in Figure 10.286.

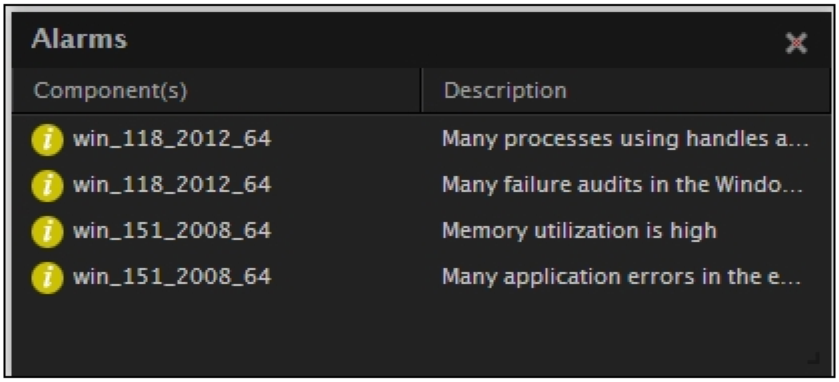

Figure 10.286: The Alerts widget that shows the current alarms in your environment

Note:

- You cannot edit the **Alarms** widget.
- The **Alarms** widget can be added only once in the dashboard.
- The displayed alarms will be based on the **View by** option that is chosen while adding the dashboard. If you have associated a service with the dashboard, then the alarms pertaining to that service alone will be displayed in the **Alarms** widget.

4. Live Measures

If you wish to monitor the real time metrics reported by certain measures, then you can use the **Live Measures** widget. Clicking the **Add** button against the **Live Measures** option in the **Data** tab as shown in Figure 10.271 will append the **Live Measures** widget to your dashboard (see Figure 10.287).

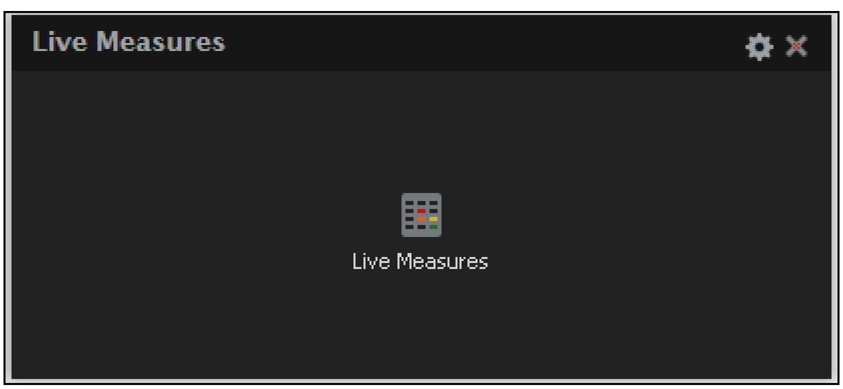

Figure 10.287: The Live Measures widget

Clicking the  icon in the widget will lead you to Figure 10.288 where you will be required to configure the measure of your choice.

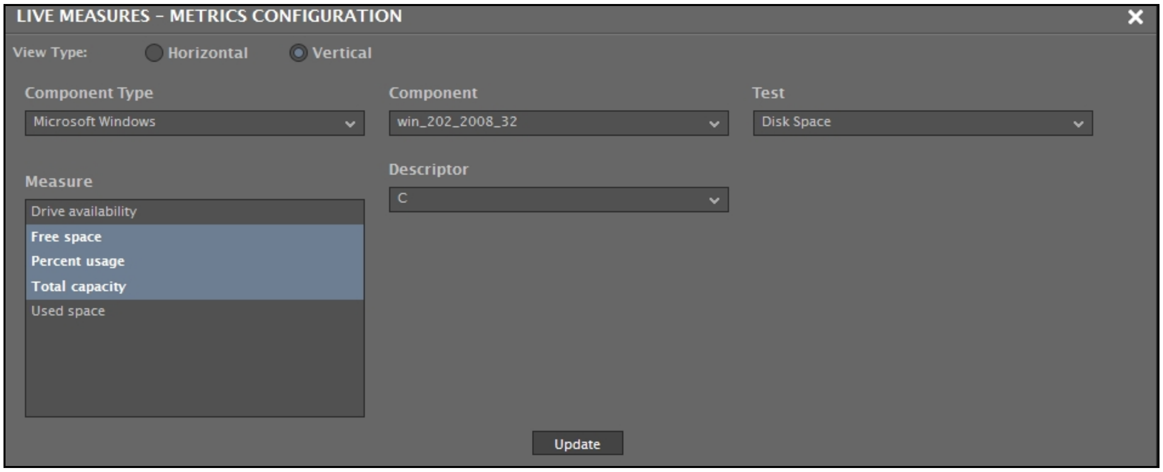

Figure 10.288: Configuring the metrics to be displayed in the Live Measures widget

Select the **Component Type**, **Component**, **Test**, **Descriptor** and **Measure** foe which you wish to view the real time metrics in the **Live Measures** widget and click the **Update** button in Figure 10.288. Figure 10.290 will then appear.

Figure 10.289: The Live Measures widget that is designed as per your choice

By default, **Vertical** option will be chosen against the **View Type** field (see Figure 10.288) indicating that the metrics of the chosen measure will be displayed vertically along with the current state of the measure (see Figure 10.290).

If you choose **Horizontal** option against the **View Type** field, then metrics will be displayed in a column manner as shown in Figure 10.290. Here, the state of the measure cannot be viewed.

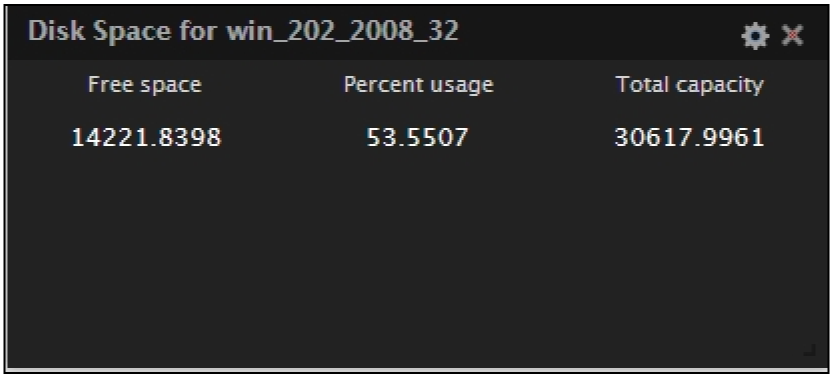

Figure 10.290: The Live Measures widget that appears when Horizontal option is chosen

Note:

The **Live Measures** widget can be configured to view the measures of a single **Test** alone. If you wish to configure **Live Measures** for multiple tests, then you may need to configure a **Live Measures** widget for each chosen **Test**.

10.9.1.2 The Graph Tab

To view a graphical representation of the measures, compare the measures in real-time, you can use the **Graph** tab (see Figure 10.291).

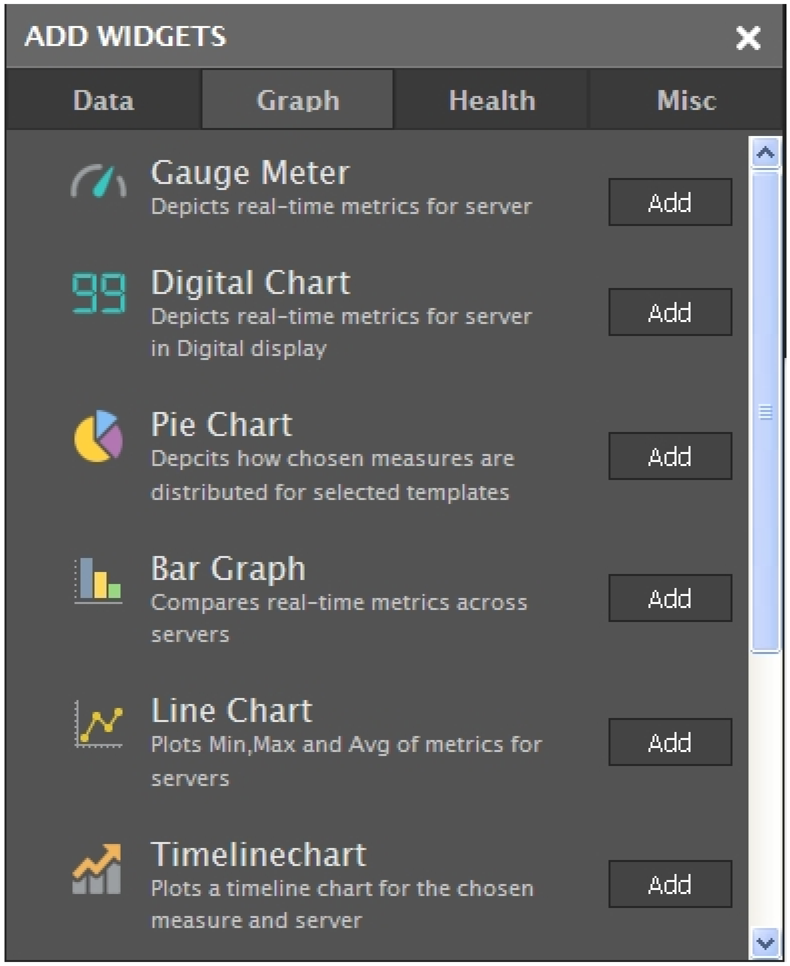

Figure 10.291: The Graph tab in the Add Widgets window

Each of the widgets of the Graph tab is discussed in detail below:

1. Gauge Meter

   Often administrators may want to keep track on certain mission critical metrics round the clock. To view such real-time metrics, you can use the **Gauge Meter** widget. To add a **Gauge Meter** widget to the dashboard, just click the **`Add`** button against the **`Gauge Meter`** in the **`Graph`** tab (see Figure 10.291). The **Gauge Meter** widget will then be appended to the dashboard as shown in Figure 10.292.

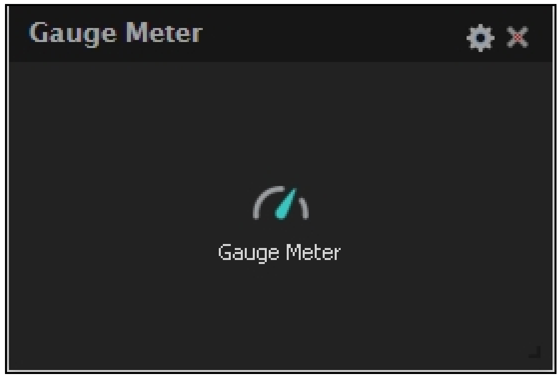

Figure 10.292: The Gauge Meter widget

Clicking the  icon in the widget will lead you to Figure 10.293 using which you can configure the metrics for the widget.

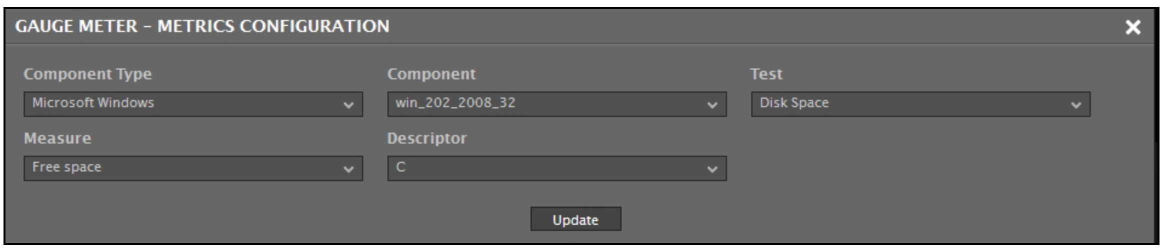

Figure 10.293: Configuring the Gauge Meter

Selecting a **Component Type**, **Component**, **Test**, **Measure** and **Descriptor** in Figure 10.293 and clicking the **Update** button will display the **Gauge Meter** for the chosen measure.

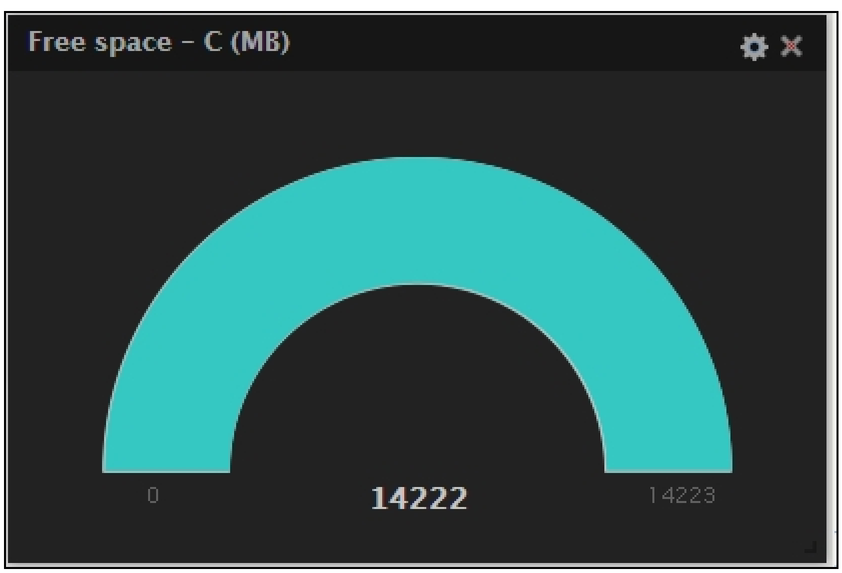

Figure 10.294: The Gauge Meter widget that is configured as per your choice

Using this **Gauge Meter** widget you can figure out the maximum value of the measure and the current value of the metric in real time. This widget also helps you to figure out the load/utilization of the chosen measure in real time.

Note:

If the measurement unit of the measure plotted in the Gauge Meter is Percent, then the maximum value of the Gauge Meter is 100. If the measure has a measurement unit other than Percent, then the maximum value is picked by analyzing the values reported by the measure over a timeperiod. If you wish to change this timeperiod, then you can do so by editing the **GaugeMaxTimeline** parameter available in the **[DASHBOARD_ SETTINGS]** section of the **<eG_ INSTALL_ DIR>\manager\config\eg_ customdashboard.ini** file. By default, the **GaugeMaxTimeline** parameter is set to 24 hours.

By default, the **Gauge Meter** is plotted using the default **HighCharts** plugin. If you wish to use a different plugin, then you can do so by setting the **EnableHighChartsforGauge** flag available in the **[DASHBOARD_ SETTINGS]** section of the **<eG_INSTALL_DIR>\manager\config\eg_customdashboard.ini** file to **false**.

2. Digital Chart

The **Digital Chart** helps administrators to track the fluctuations of a chosen measure over a period of time. To configure a **Digital Chart**, just click the **Add** button against the **Digital Chart** widget in the **Graph** tab (seeFigure 10.291). The **Digital Chart** will now be appended to the dashboard as shown in Figure 10.295.

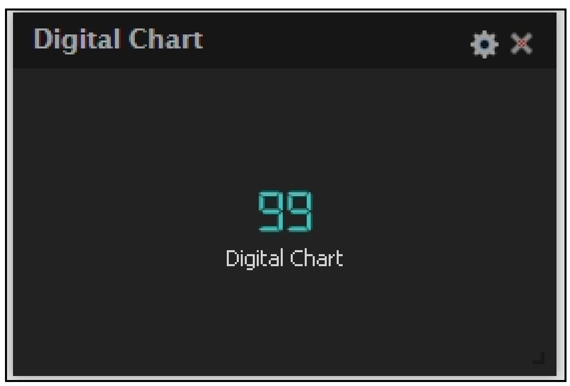

Figure 10.295: The Digital Chart widget

To configure the measure of your choice, just click the  icon in the widget which leads you to Figure 10.296.

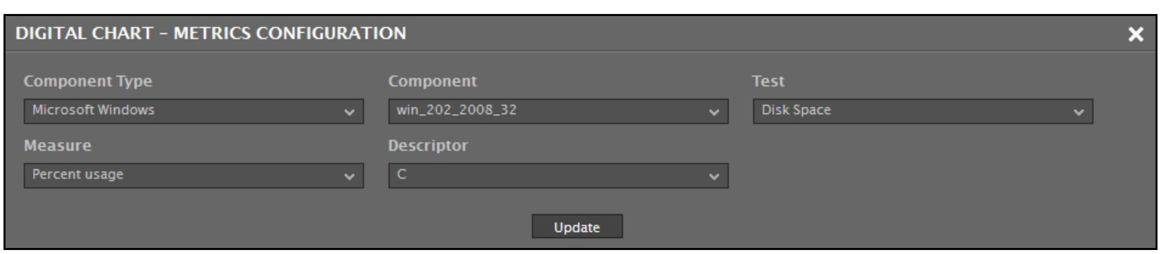

Figure 10.296: Configuring the Digital Chart widget

Selecting a **Component Type**, **Component**, **Test**, **Measure** and **Descriptor** in Figure 10.296 and clicking the **`Update`** button will display the **Digital Chart** for the chosen measure as shown in Figure 10.297.

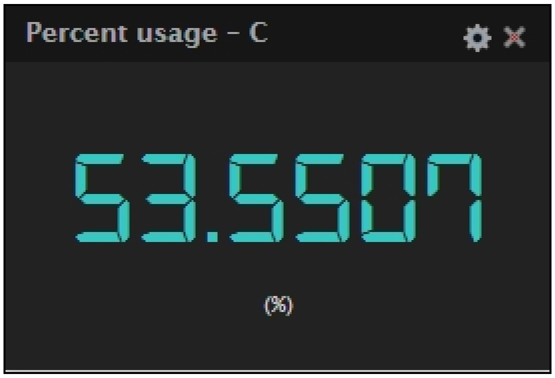

Figure 10.297: The Digital Chart widget that is configured as per your selection

3. Pie Chart

   To view the distribution of metrics for a chosen component based on a pre-defined template, you can use the **Pie Chart** widget. Clicking the **Add** button against the **Pie Chart** in the **Graph** tab will append the **Pie Chart** widget to the dashboard as shown in Figure 10.298.

Figure 10.298: The Pie Chart widget

Clicking the  icon will invoke Figure 10.299.

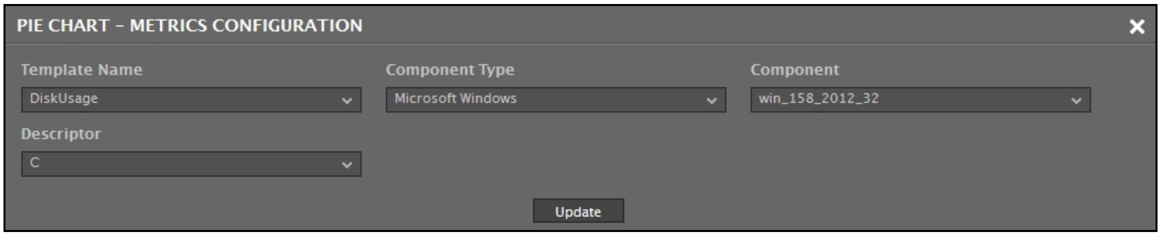

Figure 10.299: Configuring the Pie Chart widget

Select a **Template Name**, **Component Type**, **Component** and **Descriptor** of your choice from Figure 10.299 and click the **Update** button to view the distribution of metrics in the **Pie Chart** widget.

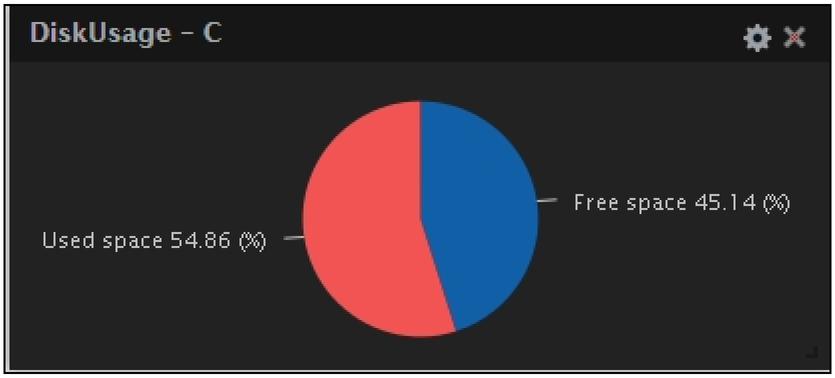

Figure 10.300: The Pie Chart metrics displaying the distribution of metrics

By default, the **Template Name** list will be populated with the templates that have been pre-defined in the **[PIE_ CHART_ TEMPLATES]** section of the **<eG_ INSTALL_ DIR>\manager\config\eg_ customdashboard.ini** file. If you wish to add a template on your own, you can do so by editing the **[PIE_ CHART_TEMPLATES]** section. You can append a new template in the following format:

Name of the Template#TestName:Measure=TestName:Measure 1,TestName:Measure 2,.....Testname:Measure n where, Measure is a cumulative sum of Measure 1, Measure 2,.....Measure n.

Note:

For a pie chart to be plotted effectively, the measurement units of all the measures appended to a template should be the same.

4. Bar Graph

If you want to compare the values of the configured measure-descriptor combinations during the last measurement period, you can use the **Bar Graph** widget. To append a bar graph to your dashboard, click the **Add** button against the **Bar Graph** in the **Graph** tab (see Figure 10.291). Figure 10.301 will then appear in your dashboard.

Figure 10.301: The Bar Graph widget

You may add the measures and components using Figure 10.302 that appears when you click the 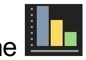 icon in the **Bar Graph** widget. Figure 10.302 will then appear.

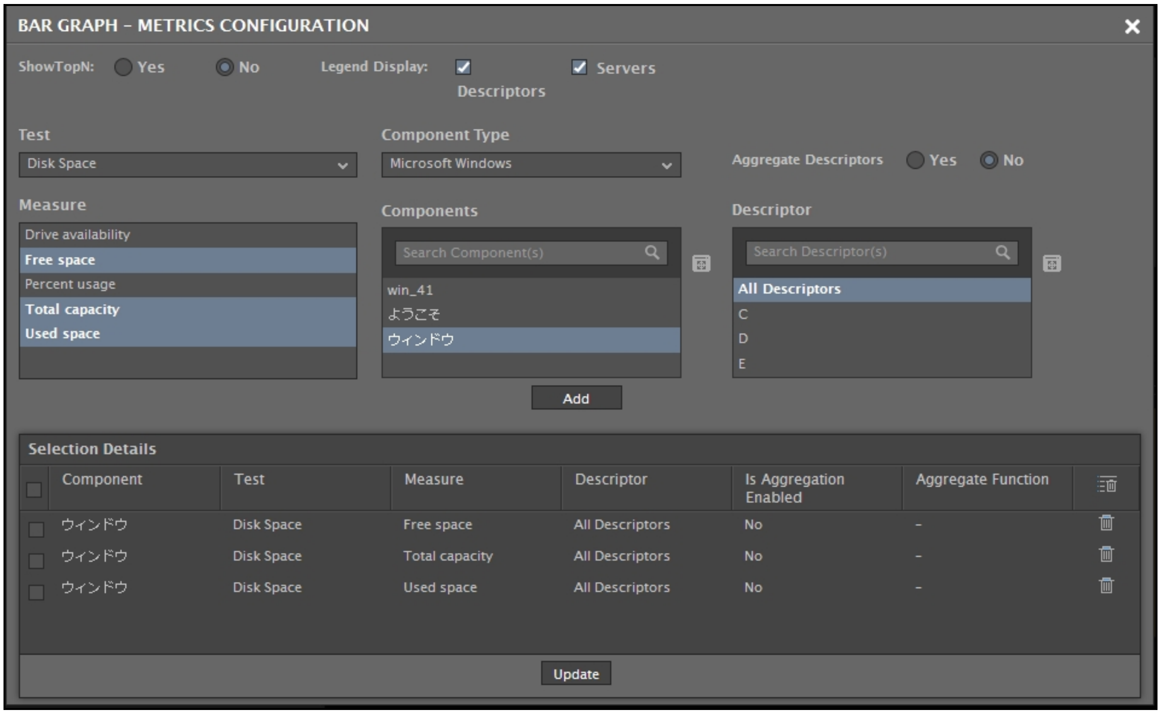

Figure 10.302: Configuring the bar graph

Choosing the options of your choice in Figure 10.302 and clicking the **Update** button will invoke Figure 10.303. The Configuration of a Bar Graph widget is similar to that of the table widget. Refer to Page 10 of this document to know more about configuring a bar graph widget.

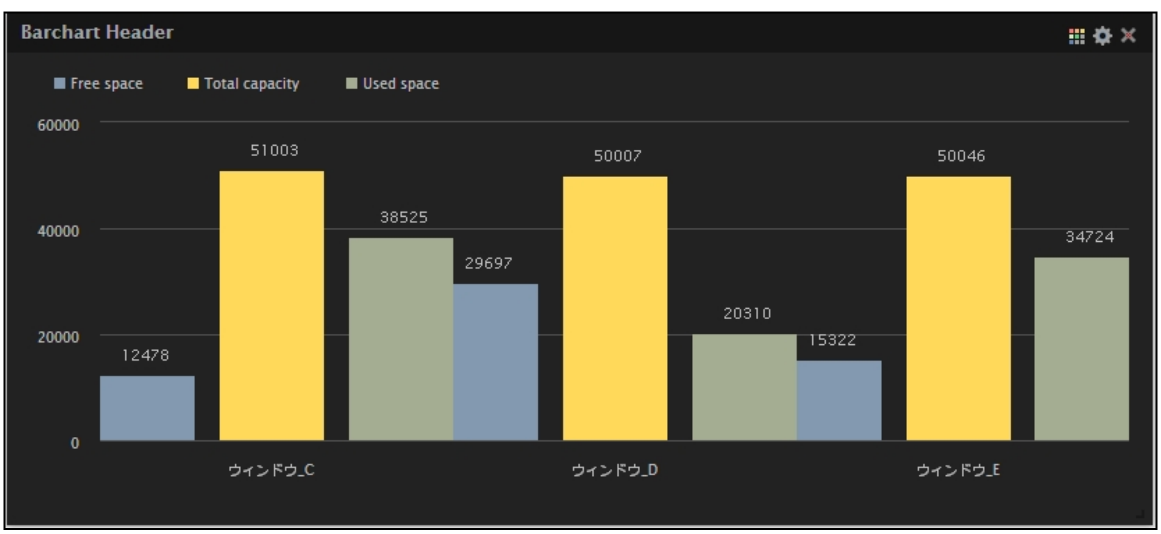

Figure 10.303: The Bar Graph widget that is configured as per your selection

If the **ShowTopN** flag is set to **Yes** as shown in Figure 10.304, then the Bar Graph widget will appear as shown in Figure 10.305.

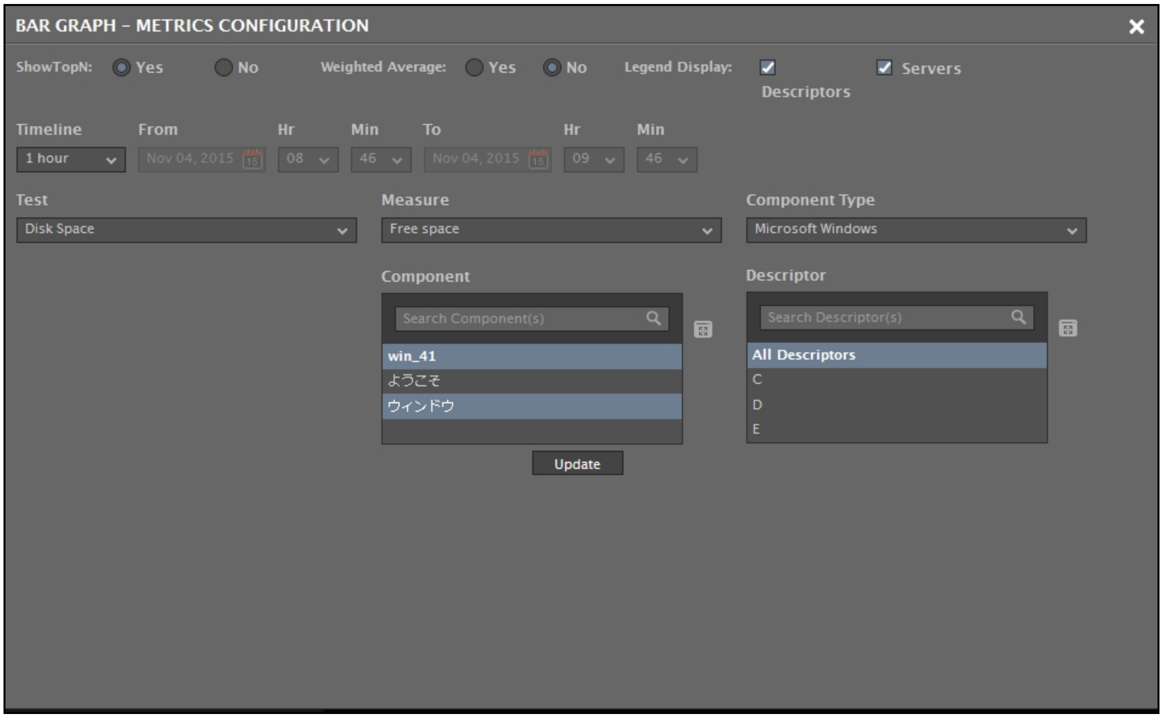

Figure 10.304: Configuring a Bar Graph when the ShowTopN flag is set to Yes

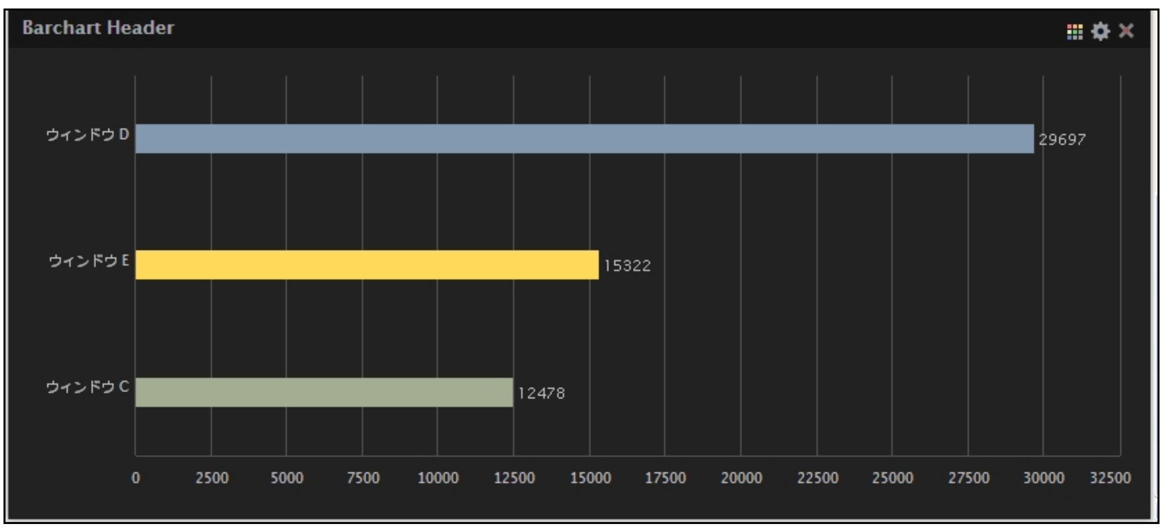

Figure 10.305: The Bar graph widget when the ShowTopN flag is set to Yes

If such a bar graph is generated for closely related measures, it will enable you to accurately diagnose the root-cause of performance problems in an infrastructure.

If you wish to change the colors used in the bar graph, simply click the  icon and choose the color palette of your interest. You can even define your own set of colors in the **[GRAPH_PALETTES]** section of the **<eG_INSTALL_DIR>\manager\config\eg_customdashboard.ini** file.

5. Line Chart

   By analyzing the Minimum, Maximum and Average values of a chosen measure over a chosen time period, you can deduce past performance trends, predict future trends, and accordingly draw capacity plans for the future. To deduce such trends, you can use the **Line chart** widget. Clicking the **Add** button against the **Line chart** widget in the **Graph** tab as shown in Figure 10.291 will append the **Line chart** to the dashboard as shown in Figure 10.306.

Figure 10.306: The Line Chart widget

To add the measures for which you wish to view the line chart, click on the icon in Figure 10.306. Figure 10.307 will then appear.

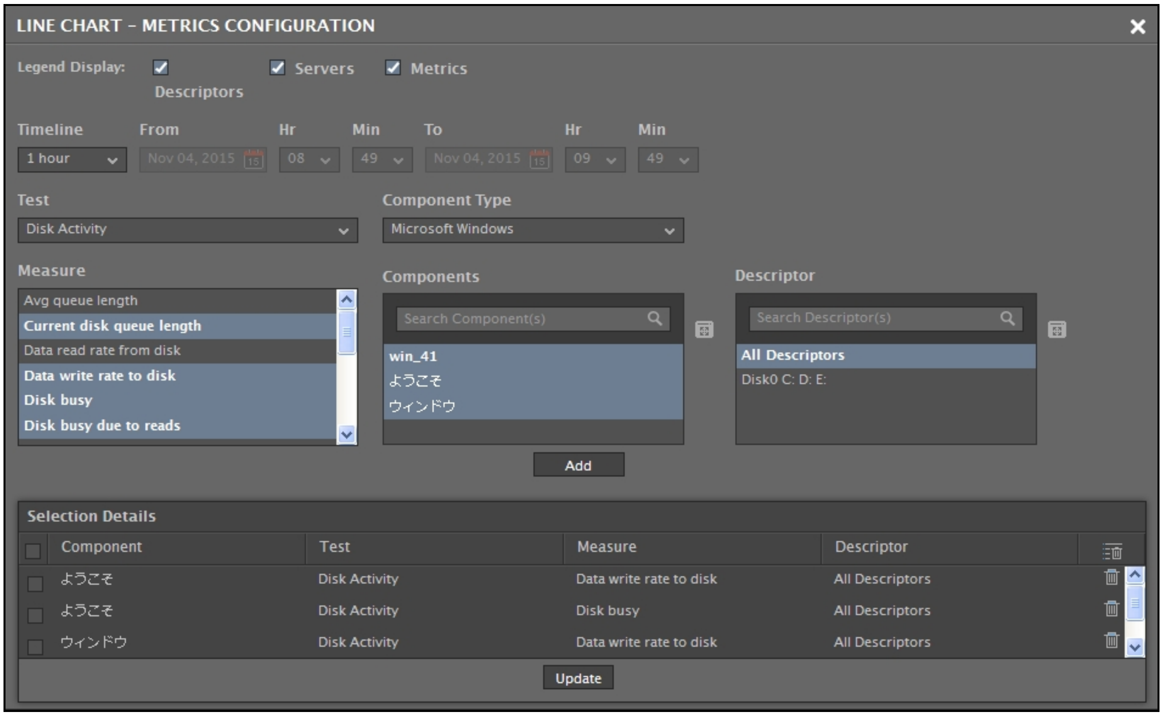

Figure 10.307: Configuring the Line Chart widget

If you wish to change the colors used in the line chart, simply click the icon and choose the color palette of your interest. You can even define your own set of colors in the **[GRAPH_PALETTES]** section of the **<eG_INSTALL_DIR>\manager\config\eg_customdashboard.ini** file.

Let us now discuss on how to configure metrics using Figure 10.307:

- Select a **Test**, **Component Type**, **Component** and **Measure** for which you wish to configure the **Line Chart**. If **the Component** list consists of too many components, then viewing all the components and selecting the ones you need for configuring the widget could require endless scrolling. To avoid this, you can click the icon next to the **Component** list. The **COMPONENTS** pop up window will then appear using which you can view almost all the components in a single interface and **Select** the ones that you wish to view in the **Line Chart** widget. You can even search for a component of your choice using the **Search Component(s)** field in the **Component** list. Once you have selected the **Component**, the descriptors of the chosen component will be displayed in the **Descriptor** section. Select the descriptor of your choice from this list. If the **Descriptor** list consists of too many descriptors, then viewing all the descriptors and selecting the ones you need for configuring the widget could require endless scrolling. To avoid this, you can click the icon next to the **Descriptor** list. The **DESCRIPTOR** pop up window will then appear using which you can view almost all the descriptors in a single interface and **Select** the ones that you wish to view in the **Line Chart**. You can even search for a descriptor of your choice using the **Search Descriptor(s)** field in the **Descriptor** list.

- By default, the **Servers** option is checked against the **Legend Display** field. If you wish to view descriptor-wise measure values, then you can check the **Descriptors** option against the **Legend Display** field. If you wish to view the measure-wise values then you can check the **Metrics** option against this field.
- You can even pick a **Timeline** for which the **Line Chart** should be configured. If you wish to specify the time for which the **Line Chart** should be configured, then you can do so using the **From** and **To** fields.
- Clicking the **Add** button will display the selections that you have made in the **Selection Details** section.
- If you wish to delete a selection, then you can do so clicking the icon. If you wish to delete multiple selections at a single shot, then select the check box against the selections and click the icon.
- Clicking the **Update** button will display the **Line Chart** for the chosen measures as shown in Figure 10.308.

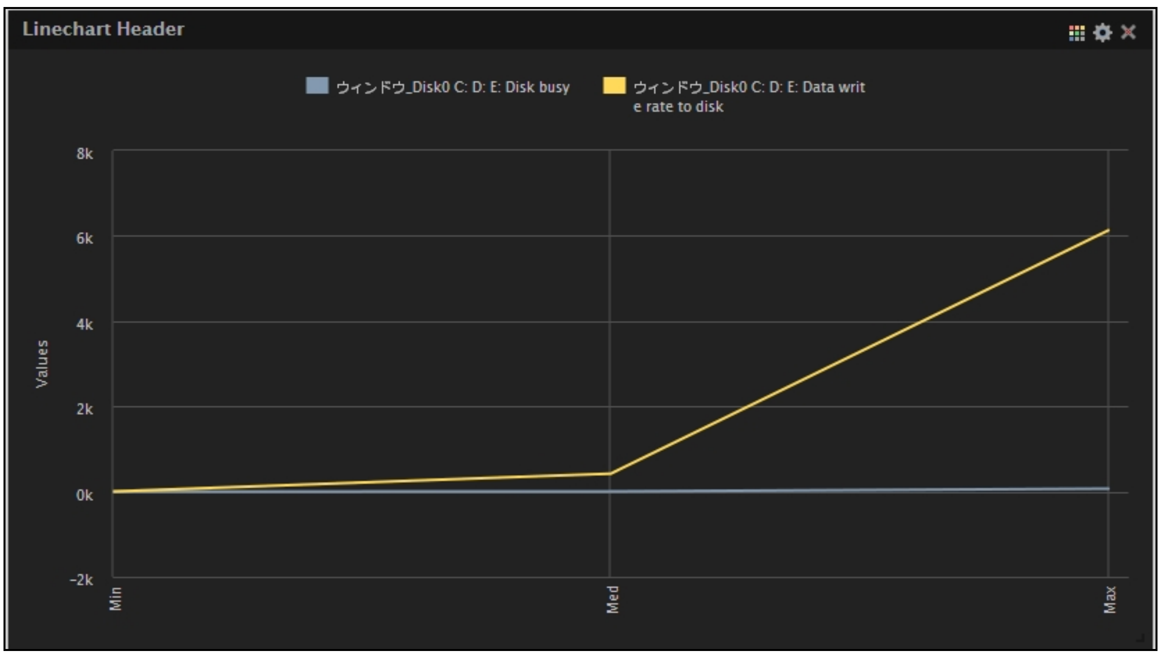

Figure 10.308: The Line Chart widget that is configured based on your selection

6. Timeline Chart

To historically analyze the performance of the chosen measure, you can include a **Timeline Chart** in your dashboard. Clicking the **Add** button against the **Timeline Chart** widget in the **Graph** tab of Figure 10.291 will append the **Timeline Chart** to the dashboard as shown in Figure 10.309.

Figure 10.309: The Timeline Chart widget

To add the measures for which you wish to view the timeline chart, click on the  icon. Figure 10.310 will then appear.

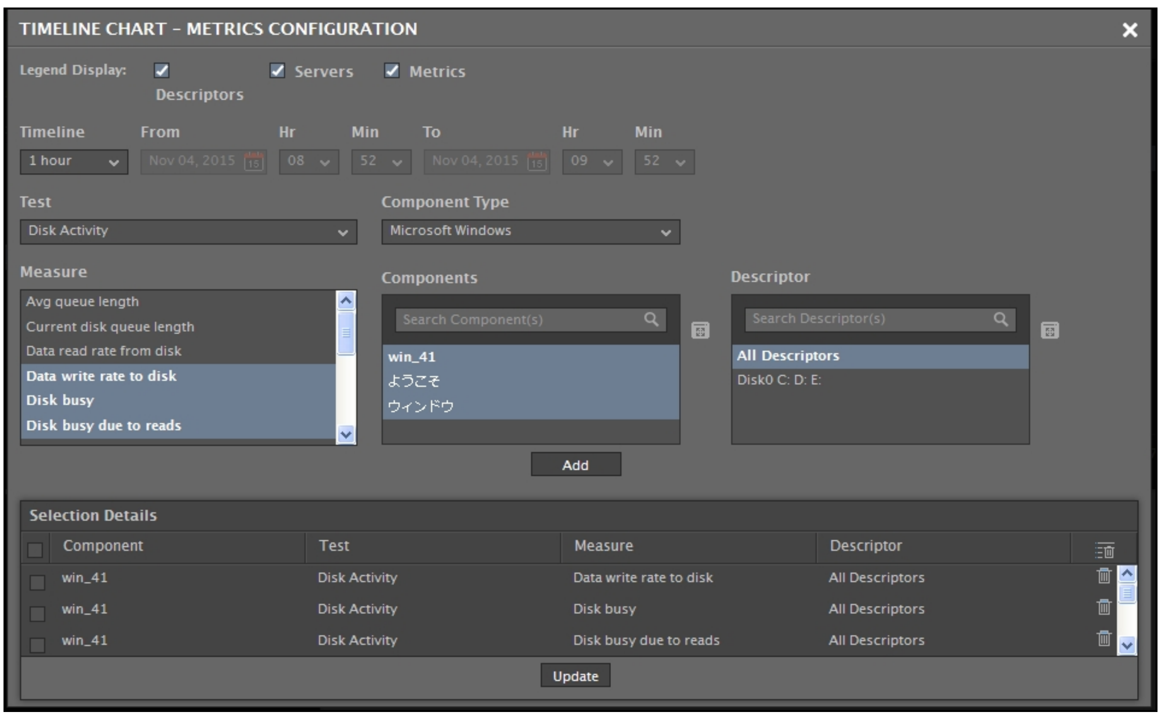

Figure 10.310: Configuring the Timeline Chart with the metrics of your choice

In Figure 10.310, do the following:

- Select a **Test**, **Component Type**, **Component** and **Measure** for which you wish to configure the **Timeline Chart** widget. If the **Components** list consists of too many components, then viewing all the components and selecting the ones you need for configuring the widget could require endless scrolling. To avoid this, you can click the  icon next to the **Components** list. The **COMPONENTS** pop up window will then appear using which you can view almost all the components in a single interface and **Select** the ones that you wish to view in the **Timeline Chart** widget. You can even search for a component of your choice using the **Search Component(s)** field in the **Component** list. Once you have selected the **Component**, the descriptors of the chosen component will be displayed in the **Descriptor** section. Select the descriptor of your choice from this list. If the **Descriptor** list consists of too many descriptors, then viewing all the descriptors and selecting the ones you need for configuring the widget could require endless scrolling. To avoid this, you can click the icon next to the **Descriptor** list. The **DESCRIPTOR** pop up window will then appear using which you can view almost all the descriptors in a single interface and **Select** the ones that you wish to view in the **Timeline Chart**. You can even search for a descriptor of your choice using the **Search Descriptor(s)** field in the **Descriptor** list.

- By default, the **Servers** option is checked against the **Legend Display** field. If you wish to view descriptor-wise measure values, then you can check the **Descriptors** option against the **Legend Display** field. If you wish to view the measure-wise values, then you can check the **Metrics** option against this field.
- You can even pick a **Timeline** for which the **Timeline Chart** should be configured. If you wish to specify the time for which the **Timeline Chart** should be configured, then you can do so using the **From** and **To** fields.
- Clicking the **Add** button will display the selections that you have made in the **Selection Details** section.
- If you wish to delete a selection, then you can do so clicking the icon. If you wish to delete multiple selections at a single shot, then select the check box against the selections and click the icon.
- Clicking the **Update** button in Figure 10.310 will display the **Timeline Chart** for the chosen measures as shown in Figure 10.311.

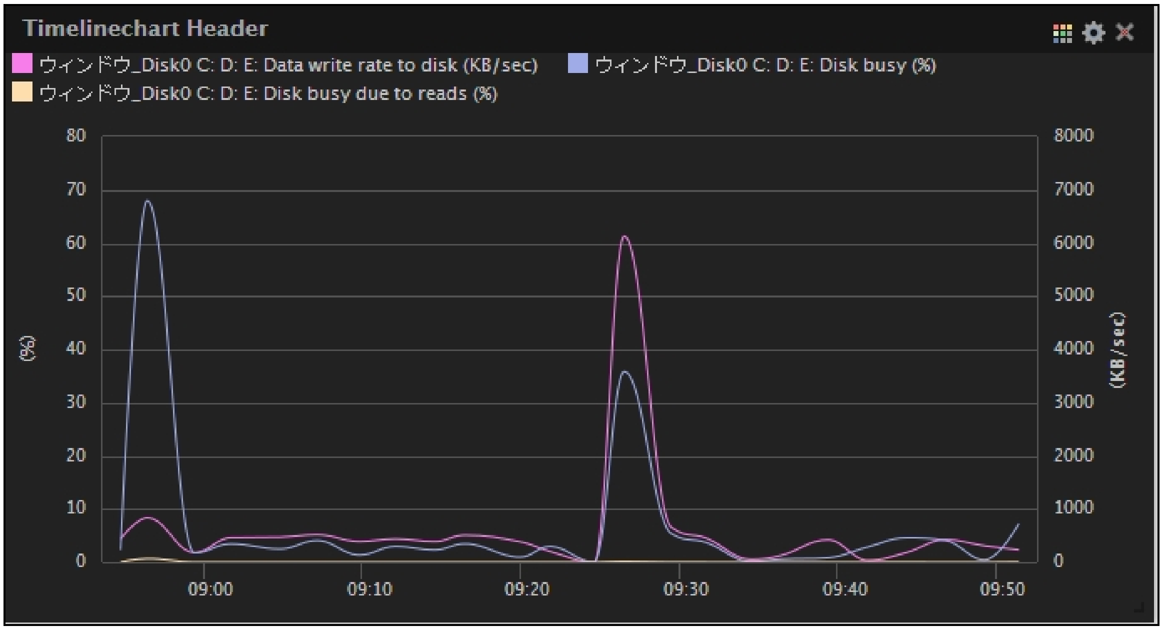

Figure 10.311: The Timeline Chart widget configured based on the measures of your choice

If you wish to change the colors used in the **Timeline Chart**, simply click the icon and choose the color palette of your interest. You can even define your own set of colors in the **[GRAPH_PALETTES]** section of the **<eG_INSTALL_DIR>\manager\config\eg_customdashboard.ini** file.

7. Area Chart

To view the area-wise distribution of metrics in a single graph, you can use the **Area Chart**. For example, if you wish to plot a chart for the *Incoming traffic* and *Outgoing traffic* of the *Windows Network Traffic* test and identify how well the network traffic is utilized over a timeperiod, then you can plot an **Area Chart**. To add an **Area Chart** to your dashboard, click the **Add** button against the **Area Chart** option in the **Graph** tab of Figure 10.291. Figure 10.312 will then appear.

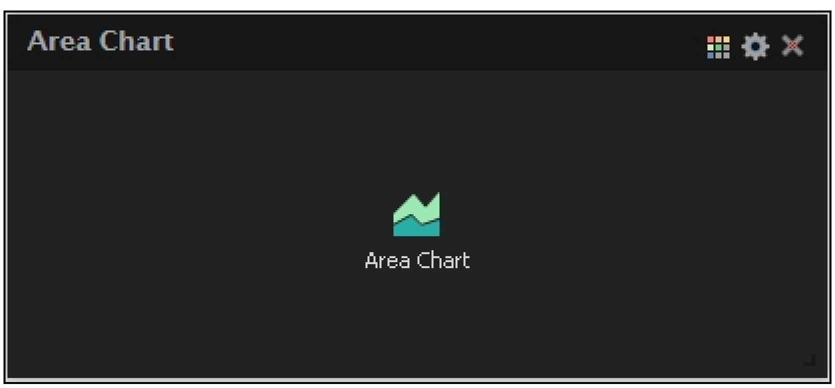

Figure 10.312: The Area Chart widget

Once the **Area Chart** widget is appended to the dashboard, click on the icon. Figure 10.313 will then appear.

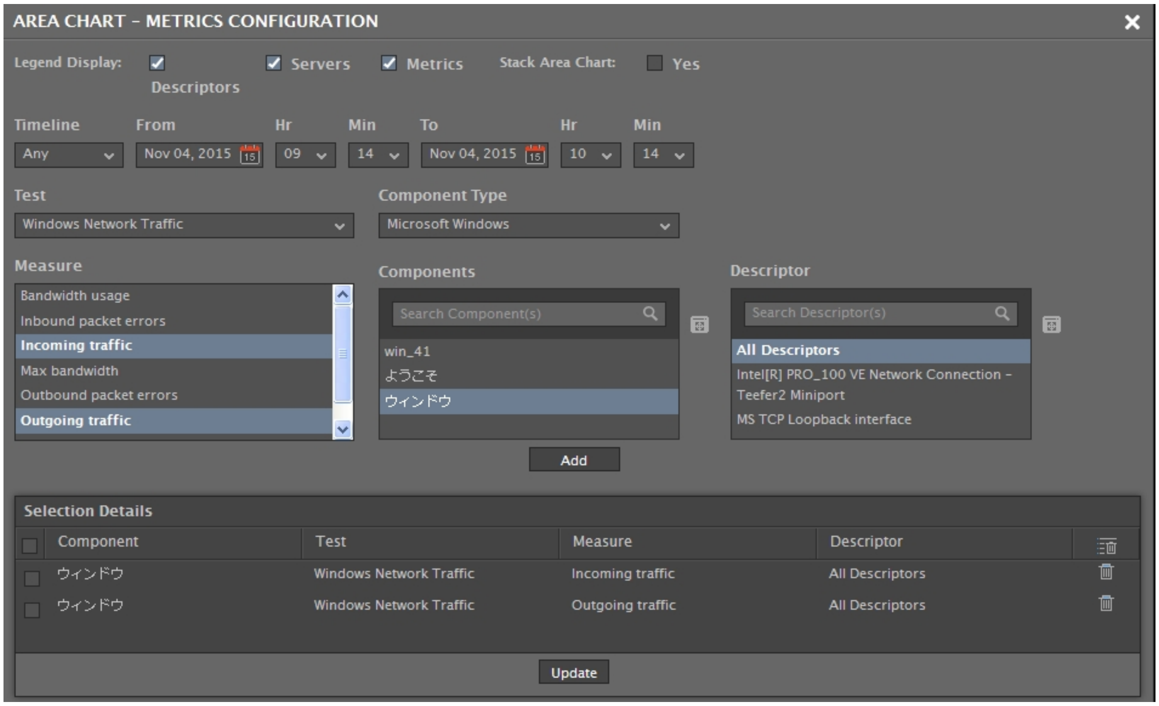

Figure 10.313: Configuring the Area Chart

Configure the metrics in Figure 10.313 as specified below:

- Select a **Test**, **Component Type**, **Component** and **Measure** for which you wish to configure the **Area Chart**. If the **Components** list consists of too many components, then viewing all the components and selecting the ones you need for configuring the widget could require endless scrolling. To avoid this, you can click the icon next to the **Components** list. The **COMPONENTS** pop up window will then appear using which you can view almost all the components in a single interface and **Select** the ones that you wish to view in the **Area Chart** widget. You can even search for a component of your choice using the **Search Component(s)** field in the **Components** list. Once you have selected the **Component**, the descriptors of the chosen component will be displayed in the **Descriptor** section. Select the descriptor of your choice from this list. If the **Descriptor** list consists of too many descriptors, then viewing all the descriptors and selecting the ones you need for configuring the widget could require endless scrolling. To avoid this, you can click the icon next to the **Descriptor** list. The **DESCRIPTOR** pop up window will then appear using which you can view almost all the descriptors in a single interface and **Select** the ones that you wish to view in the **Area Chart**. You can even search for a descriptor of your choice using the **Search Descriptor(s)** field in the **Descriptor** list.
- By default, the **Servers** option is checked against the **Legend Display** field. If you wish to view descriptor-wise measure values, then you can check the **Descriptors** option against the **Legend**

Display field. If you wish to view the measure-wise values, then you can check the **Metrics** option against this field.

- You can even pick a **Timeline** for which the **Area Chart** should be configured. If you wish to specify the time for which the **Area Chart** should be configured, then you can do so using the **From** and **To** fields.
- Clicking the **Add** button will display the selections that you have made in the **Selection Details** section.
- If you wish to delete a selection, then you can do so clicking the icon. If you wish to delete multiple selections at a single shot, then select the check box against the selections and click the icon.
- Clicking the **Update** button will display the **Area chart** for the chosen measures as shown in Figure 10.314.
- By default, the area chart will be displayed with an overlap of one metric over the other i.e., the measure with the maximum value will be plotted first and the rest of the measures will be plotted within the area of the measure with the maximum value.

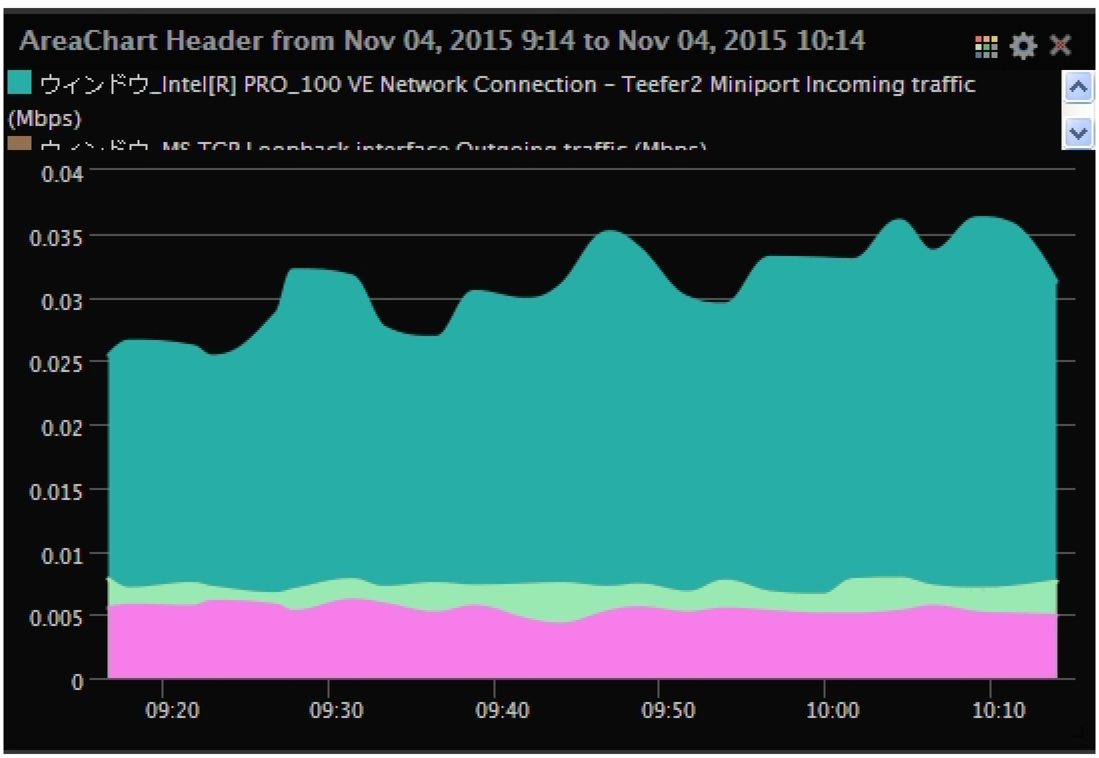

Figure 10.314: The Area Chart widget that is configured based on your selection

- If you wish to view the exact area distribution of measures without overlap, you can check the **Stack Area Chart** check box in Figure 10.313. For example, if you wish to view the area-wise distribution of the *Incoming Traffic* and *Outgoing Traffic* metrics of the *Windows Network Traffic* test using the **Area Chart**, then upon adding those metrics, the **Area Chart** will appear as shown in Figure 10.315 with the minimum value being plotted in the graph first and the value higher than the minimum value plotted over the minimum value. **Note that only the metrics with the same measurement unit should be added if the Stack Area Chart is checked .**

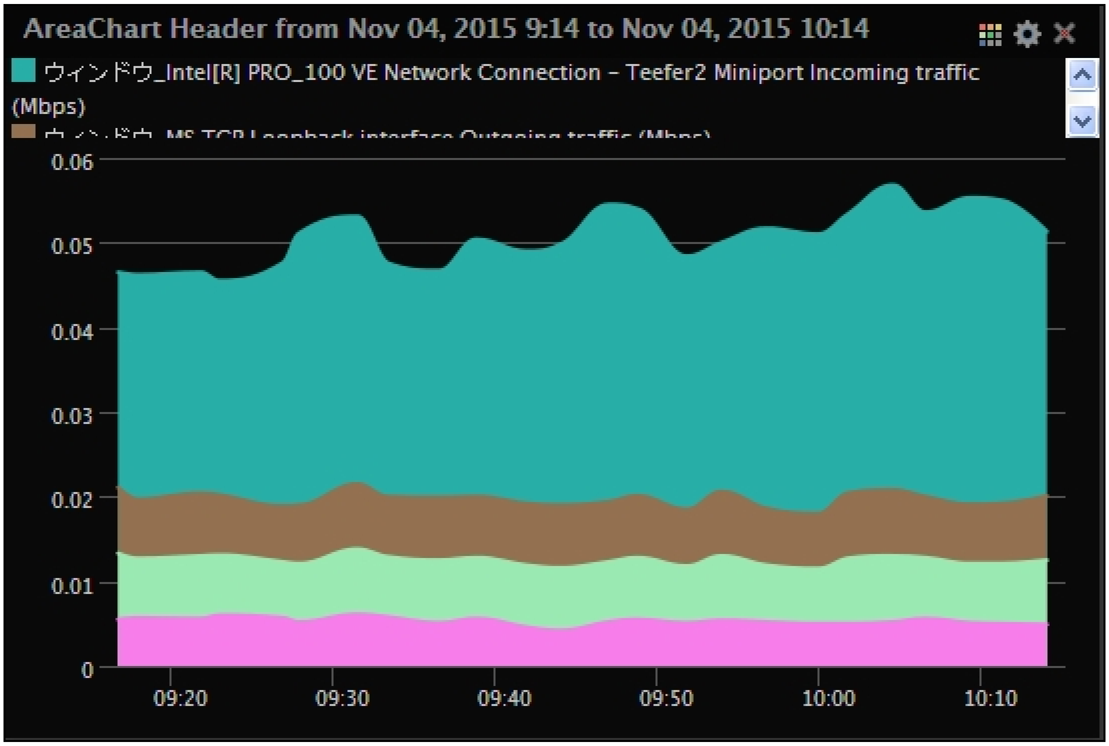

Figure 10.315: The Area Chart when the Stack Area Chart check box is checked

If you wish to change the colors used in the **Timeline Chart**, simply click the icon and choose the color palette of your interest. You can even define your own set of colors in the **[GRAPH_PALETTES]** section of the **<eG_INSTALL_DIR>\manager\config\eg_customdashboard.ini** file.

8. Combination Chart

If you wish to view a graphical representation of a set of metrics with different measurement units, then you can use the **Combination Chart** widget. This **Combination Chart** combines the **Line Chart**, **Bar Graph** and **Area Charts** and provides you with a single graph. To add a **Combination Chart** to your dashboard, click the **Add** button against the **Combination Chart** option in the **Graph** tab. Figure 10.316 will then appear.

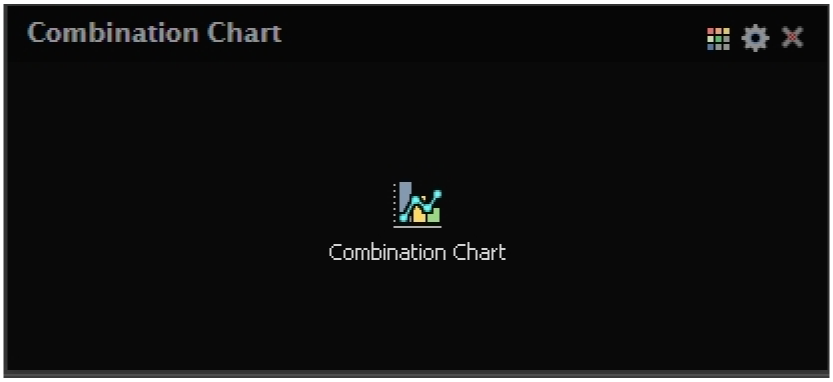

Figure 10.316: The Combination Chart widget

Once the **Combination Chart** widget is appended to the dashboard, click on the  icon. Figure 10.317 will then appear.

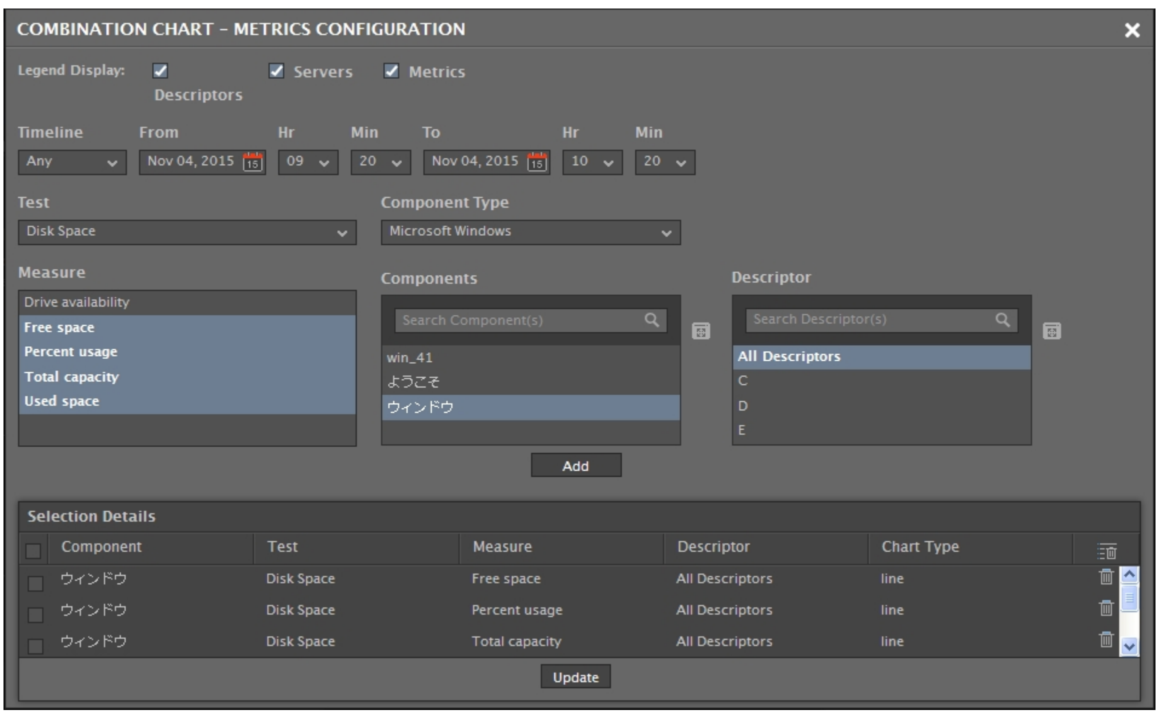

Figure 10.317: Configuring the Combination Chart

In Figure 10.317, you can do the following:

- Select a **Test**, **Component Type**, **Component** and **Measure** for which you wish to configure the **Combination Chart**. If the **Components** list consists of too many components, then viewing all the components and selecting the ones you need for configuring the widget could require endless scrolling. To avoid this, you can click the icon next to the **Components** list. The **COMPONENTS** pop up window will then appear using which you can view almost all the components in a single interface and **Select** the ones that you wish to view in the **Combination Chart** widget. You can even search for a component of your choice using the **Search Component(s)** field in the **Components** list. Once you have selected the **Components**, the descriptors of the chosen component will be displayed in the **Descriptor** section. Select the descriptor of your choice from this list. If the **Descriptor** list consists of too many descriptors, then viewing all the descriptors and selecting the ones you need for configuring the widget could require endless scrolling. To avoid this, you can click the icon next to the **Descriptor** list. The **DESCRIPTOR** pop up window will then appear using which you can view almost all the descriptors in a single interface and **Select** the ones that you wish to view in the **Combination Chart**. You can even search for a descriptor of your choice using the **Search Descriptor(s)** field in the **Descriptor** list.
- By default, the **Servers** option is checked against the **Legend Display** field. If you wish to view descriptor-wise measure values, then you can check the **Descriptors** option against the **Legend Display** field. If you wish to view the measure-wise values, then you can check the **Metrics** option against this field.
- You can even pick a **Timeline** for which the **Combination Chart** should be configured. If you wish to specify the time for which the **Combination Chart** should be configured, then you can do so using the **From** and **To** fields.
- Clicking the **Add** button will display the selections that you have made in the **Selection Details** section.
- By default, the graph will be plotted as a line chart and therefore all your selections will be listed in the **Selection Details** section with the **Chart Type** column indicating **line** as the default graph type. If you wish to view certain metrics of your selection as a bar graph or an area chart, then simply double clicking the **line** option under the **Chart Type** column against the chosen measure will invoke a drop down list. From this list, select **column** or **Area** according to your wish.
- If you wish to delete a selection, then you can do so clicking the icon. If you wish to delete multiple selections at a single shot, then select the check box against the selections and click the icon.
- Clicking the **Update** button will display the **Combination chart** for the chosen measures as shown in Figure 10.318.

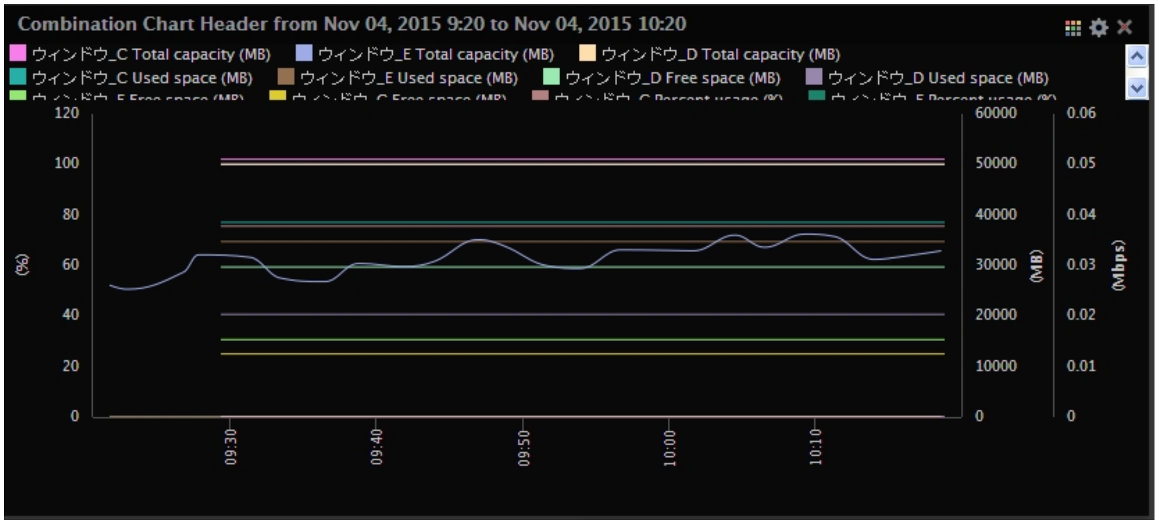

Figure 10.318: The Combination Chart that is displayed based on the metrics of your choice

➢ If the metrics of your choice are with different measurement units, then each type of graph will have a separate Y axis for your reference (see Figure 10.318). If you plot the *Used space*, *Total capacity*, *Free Space* metrics of the *Disk Space* test and the *Incoming Traffic* of the *Network Traffic* test, then the graph will be plotted with 3 different Y axis. The Y axis for the *Total Capacity* measure will be in represented in `MB`, the Y axis of the *Percent usage* measure will be represented in *Percent* and the Y axis of the *Incoming traffic* will be represented in Mbps.

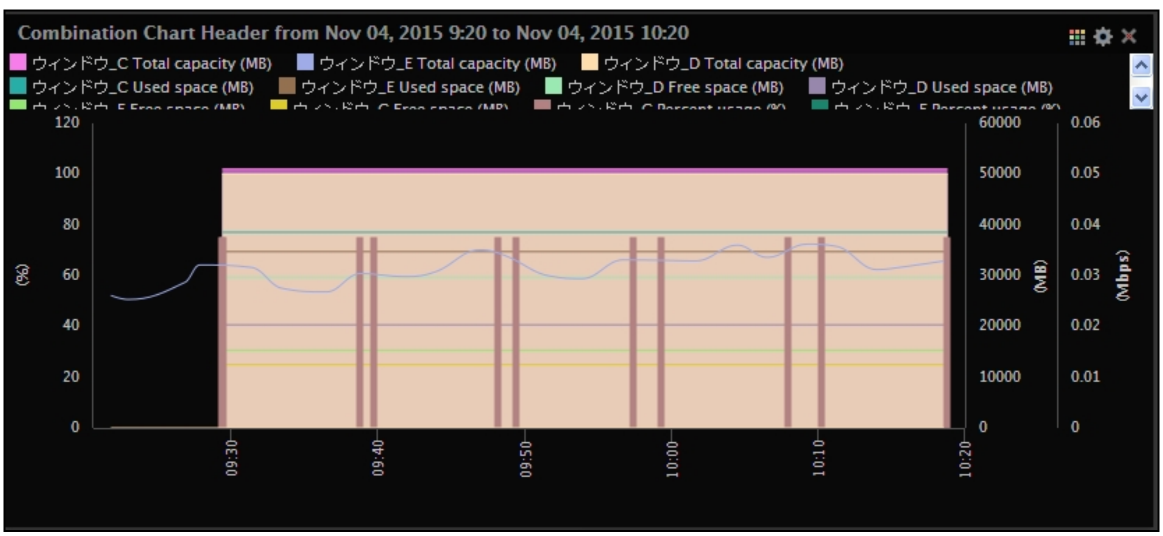

Figure 10.319: The Combination Chart when different chart types are chosen

If you wish to change the colors used in the **Combination Chart**, simply click the icon and choose the color palette of your interest. You can even define your own set of colors in the **[GRAPH_PALETTES]** section of the **<eG_INSTALL_DIR>\manager\config\eg_customdashboard.ini** file.

9. Heat Map Chart

If you wish to view the distribution of metrics possessing the same measurement unit, then you can use the **Heat Map Chart**. To add a **Heat Map Chart** to your dashboard, click the **Add** button against the **Heat Map Chart** option in the **Graph** tab. Figure 10.320 will then appear.

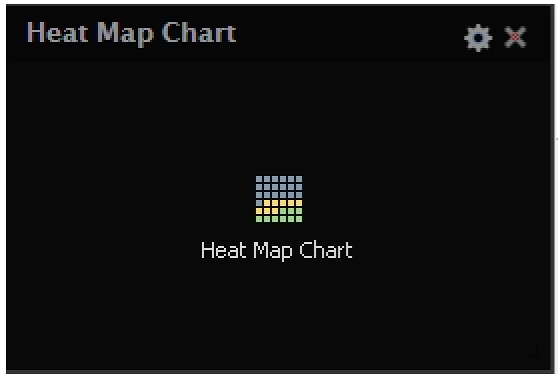

Figure 10.320: The Heat Map Chart widget

Once the **Heat Map Chart** widget is appended to the dashboard, click on the 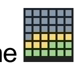 icon. Figure 10.321 will then appear.

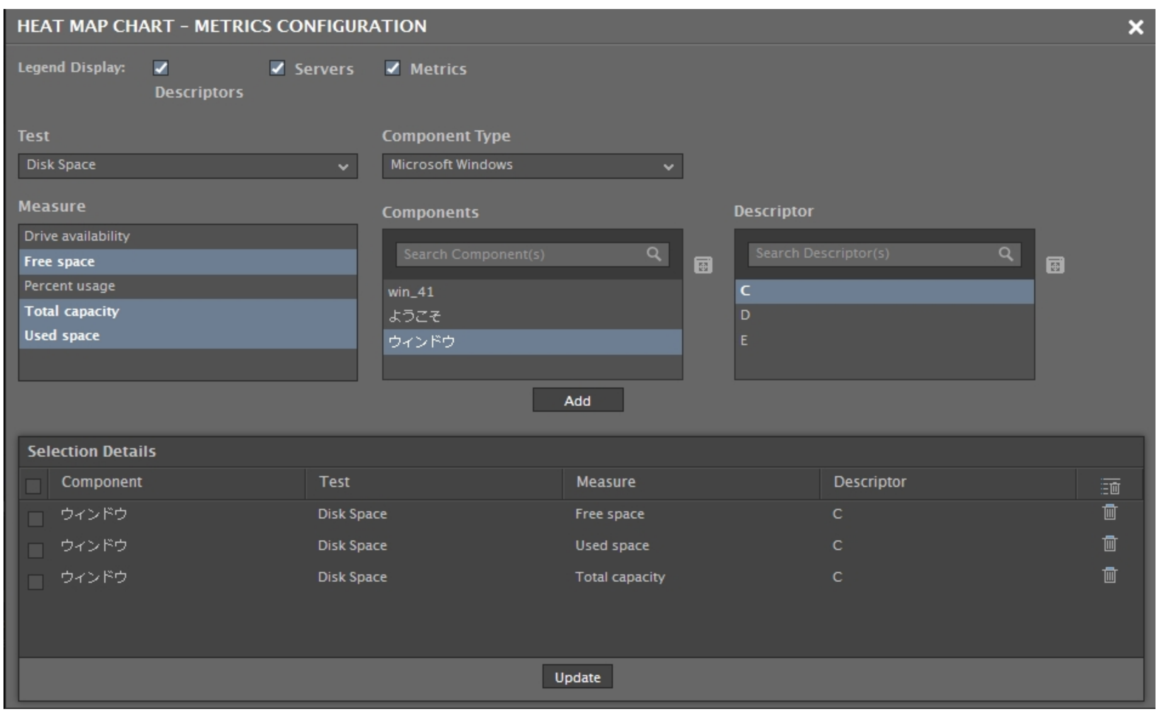

Figure 10.321: Configuring the Heat Map Chart

Specify the following in Figure 10.321:

- Select a **Test**, **Component Type**, **Component** and **Measure** for which you wish to configure the **Heat Map Chart**. If the **Components** list consists of too many components, then viewing all the components and selecting the ones you need for configuring the widget could require endless scrolling. To avoid this, you can click the icon next to the **Components** list. The **COMPONENTS** pop up window will then appear using which you can view almost all the components in a single interface and **Select** the ones that you wish to view in the **Heat Map Chart** widget. You can even search for a component of your choice using the **Search Component(s)** field in the **Components** list. Once you have selected the **Component**, the descriptors of the chosen component will be displayed in the **Descriptor** section. Select the descriptor of your choice from this list. If the **Descriptor** list consists of too many descriptors, then viewing all the components and selecting the ones you need for configuring the widget could require endless scrolling. To avoid this, you can click the icon next to the **Descriptor** list. The **DESCRIPTOR** pop up window will then appear using which you can view almost all the descriptors in a single interface and **Select** the ones that you wish to view in the **Heat Map Chart**. You can even search for a descriptor of your choice using the **Search Descriptor(s)** field in the **Descriptor** list.

- By default, the **Servers** option is checked against the **Legend Display** field. If you wish to view descriptor-wise measure values, then you can check the **Descriptors** option against the **Legend Display** field. If you wish to view the measure-wise values, then you can check the **Metrics** option against this field.
- You can even pick a **Timeline** for which the **Heat Map Chart** should be configured. If you wish to specify the time for which the **Heat Map Chart** should be configured, then you can do so using the **From** and **To** fields.
- Clicking the **Add** button will display the selections that you have made in the **Selection Details** section.
- If you wish to delete a selection, then you can do so clicking the icon. If you wish to delete multiple selections at a single shot, then select the check box against the selections and click the icon.
- Clicking the **Update** button will display the **Heat Map chart** for the chosen measures as shown in Figure 10.322. By default, the overall representation of the **Heat Map Chart** will be in percentage values, irrespective of the measurement unit of the chosen measures.

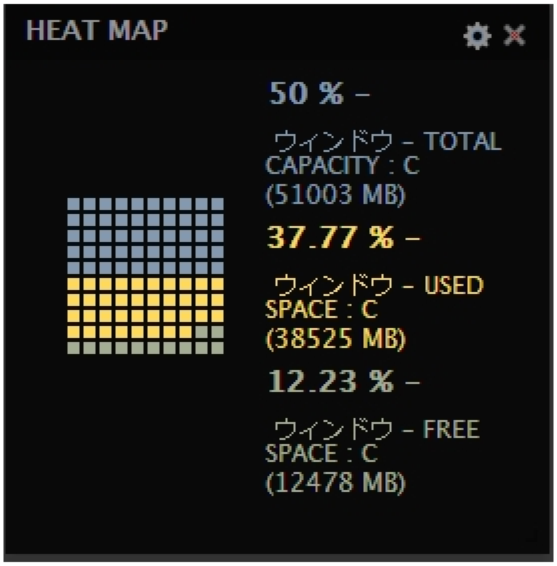

Figure 10.322: The Heat Map Chart that is configured based on the metrics of your choice

Note:

- The distribution of the Heat Map Chart will be relevant only for measures with the same measurement unit.
- If you configure more than 5 measures, the Heat Map Chart may not be evenly distributed.
- Unlike other graphs, when you expand the Heat Map Chart widget, the actual **Heat Map Chart** cannot be expanded to fit the widget.

10.9.1.3 The Health tab

This tab (see Figure 10.323) allows you to add Tier Health, Service Health and Network Health to the dashboard.

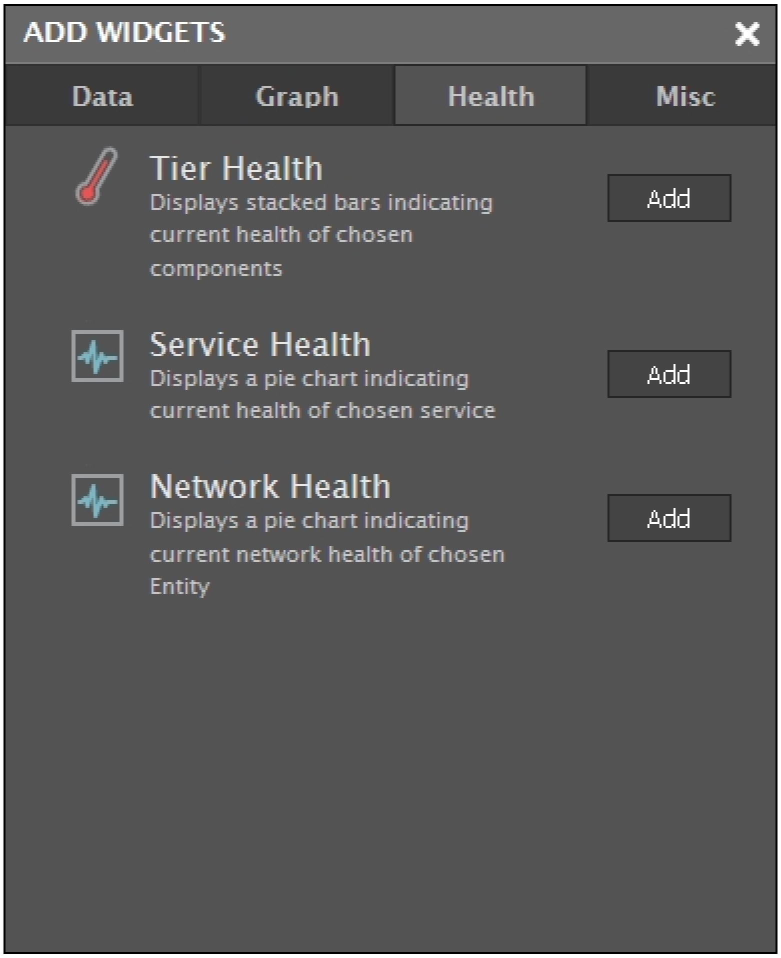

Figure 10.323: The Health Tab of the ADD WIDGETS window

1. Tier Health

   In order to view the number of components and their current status for each **Component Type** in your infrastructure, you can use the **Tier Health** widget. Once you click the **Add** button against this widget in Figure 10.323, the **Tier Health** widget will appear in your dashboard displaying all the component types in your infrastructure (see Figure 10.324).

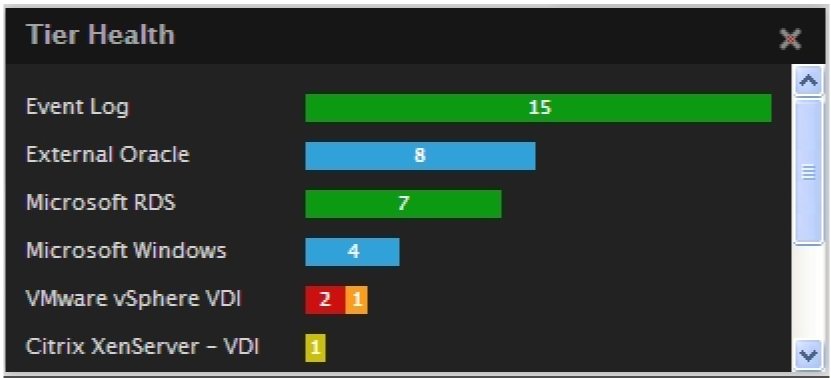

Figure 10.324: The Tier Health widget

   Against each component type, a bar graph will appear listing the number of components of each component type in Critical, Major, Minor, Normal, and/or Unknown states. Clicking on each section of the bar graph will lead you to the **SERVERS** page (see Figure 10.326) where you will be able to view the components that are currently in the chosen state.

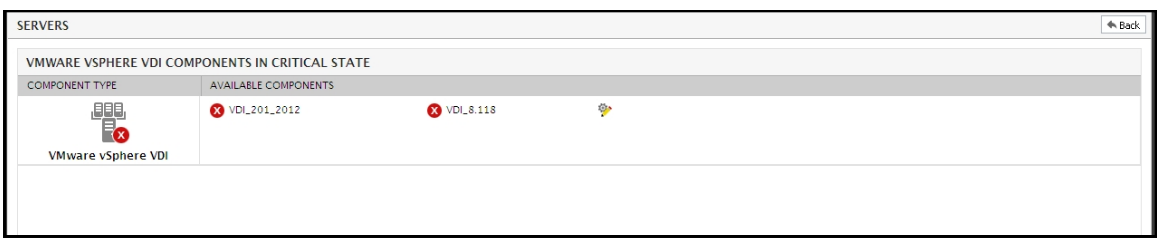

Figure 10.325: The SERVERS page that appears when a bar is clicked from the Tier health widget

   Note:

   - You cannot edit the **Tier Health** widget.
   - The **Tier Health** widget can be added only once in the dashboard.

2. Service Health

If you wish to view the state wise distribution of your infrastructure in your dashboard, then use the **Service Health** widget. The Service Health widget will appear once you click the **Add** button against the **Service Health** widget in the **Health** tab (see Figure 10.326).

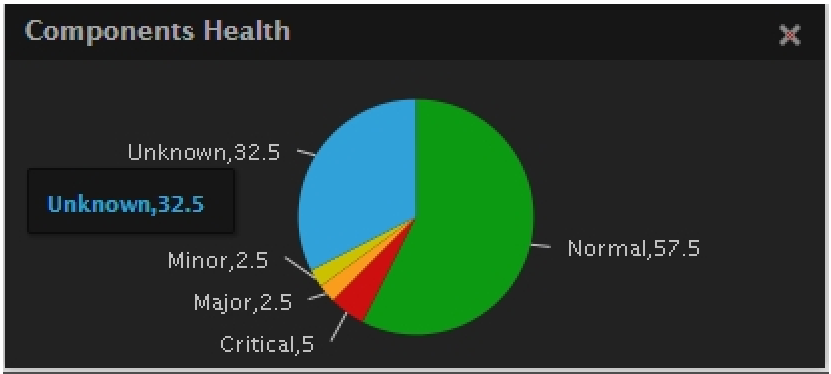

Figure 10.326: The Service Health widget

The pie chart will then appear with the state wise distribution of your infrastructure based on the option that you have chosen from the **View by** list in the **ADD DASHBOARD** window while adding the dashboard. If you have chosen **Components** from the **View by** list, then the **Service Health** widget will display a state wise distribution of all the components in your infrastructure. Clicking on a slice of the pie chart will lead you to the **SERVERS** page (see Figure 10.327) where you will be able to view the components that are currently in the chosen state.

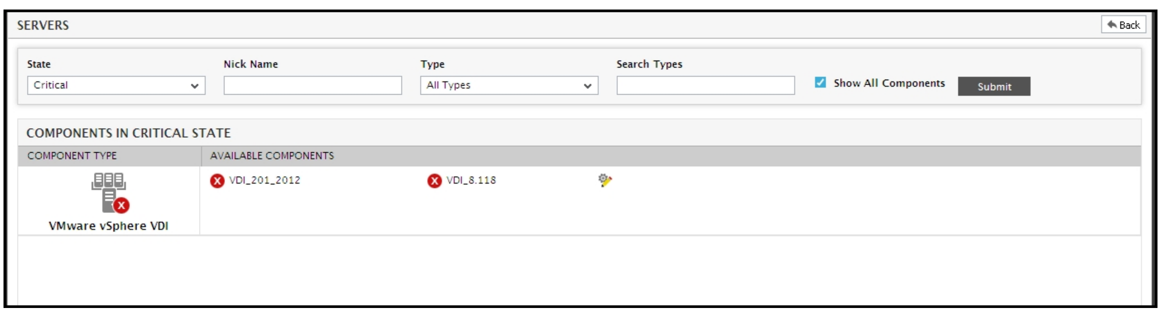

Figure 10.327: The SERVERS page that appears upon clicking a slice of the Service Health widget

Note:

- You cannot edit the **Service Health** widget.
- The **Service Health** widget can be added only once in the dashboard.

3. Network Health

If you wish to view the state-wise network health of the components in your infrastructure in the dashboard, then use the **Network Health** widget. The Network Health widget will appear once you click the **Add** button against the **Network Health** widget in the **Health** tab (see Figure 10.328).

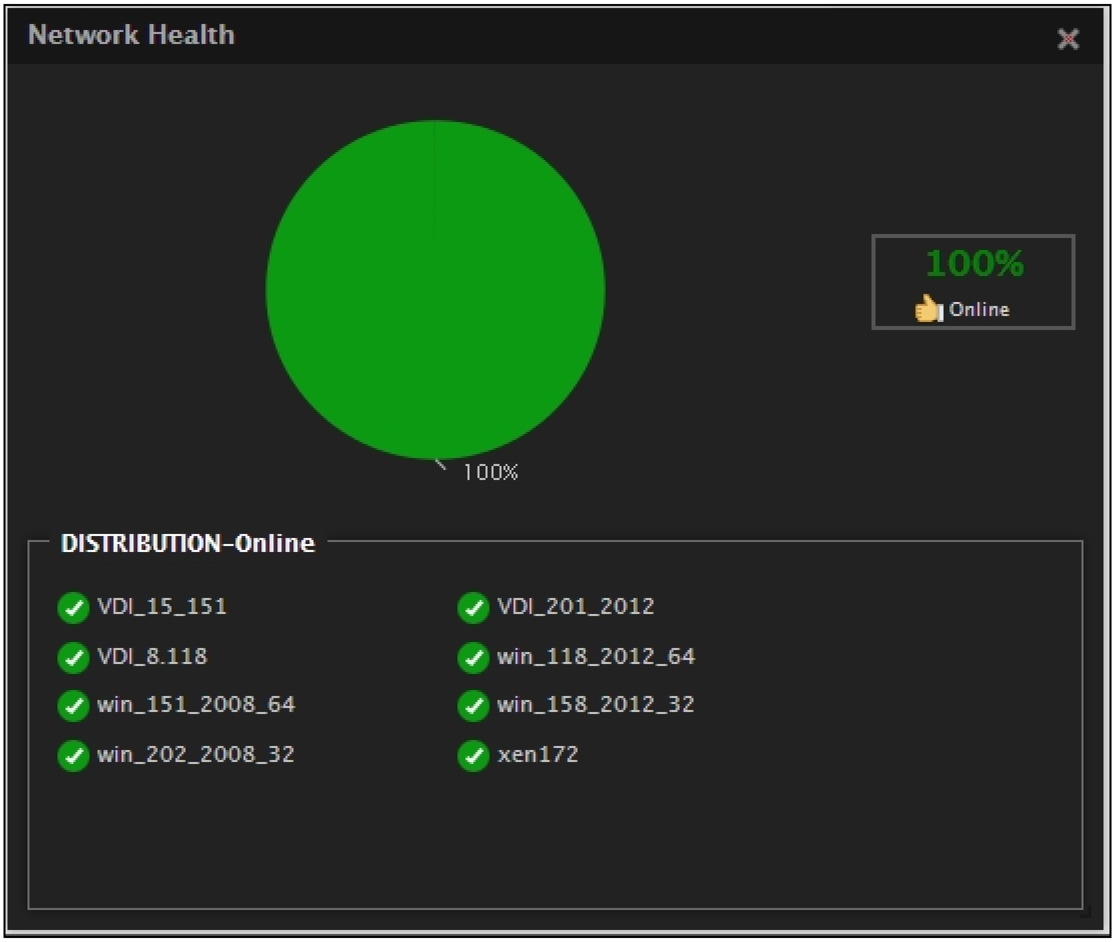

Figure 10.328: The Network Health widget

The pie chart will then appear with the state wise distribution of your infrastructure based on the option that you have chosen from the **View by** list in the **ADD DASHBOARD** window while adding the dashboard. If you have chosen **Components** from the **View by** list, then the **Network Health** widget will display a state wise distribution of all the components in your infrastructure. Clicking on a slice of the pie chart will lead you to the **SERVERS** page where you will be able to view the components that are currently in the chosen state. If you wish to view the network health of an individual component, then you can do so by clicking the component in Figure 10.328. This will lead you to the layer model of the chosen component where you can view all the network related details of the chosen component.

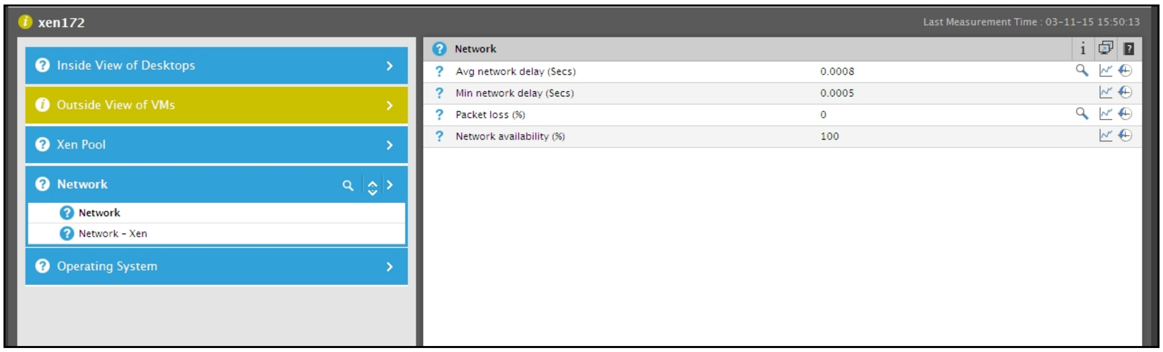

Figure 10.329: The layer model page that appears upon clicking a component in the Network Health widget

Note:

- You cannot edit the **Network Health** widget.
- The **Network Health** widget can be added only once in the dashboard

10.9.1.4 The Misc tab

If you wish to stay updated with the latest happening around the world from within your dashboard, you can use the **Misc** tab. You can configure weather widget and the News of your interest using the widgets available in this tab. **For this tab to show real time metrics, you would require an internet connection.**

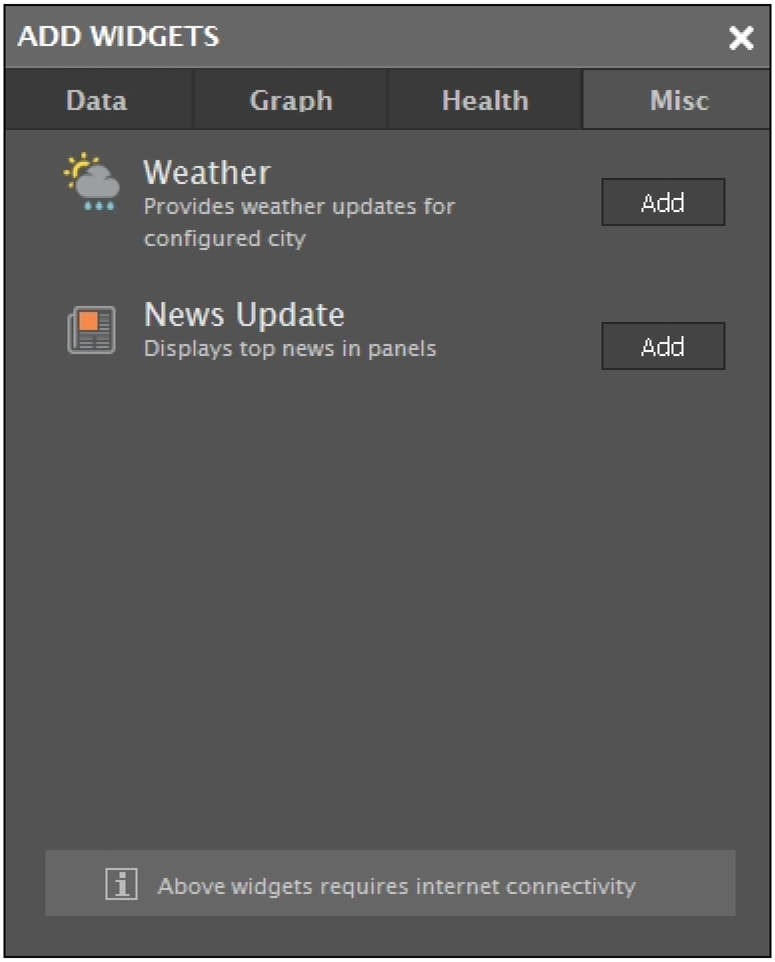

Figure 10.330: The Misc tab

1. Weather

   If you wish to view the weather forecast for the city of your choice, click on the **Add** button against the **Weather** widget in the **Misc** tab. Doing so, the **Weather** widget will appear in your dashboard as shown in Figure 10.331.

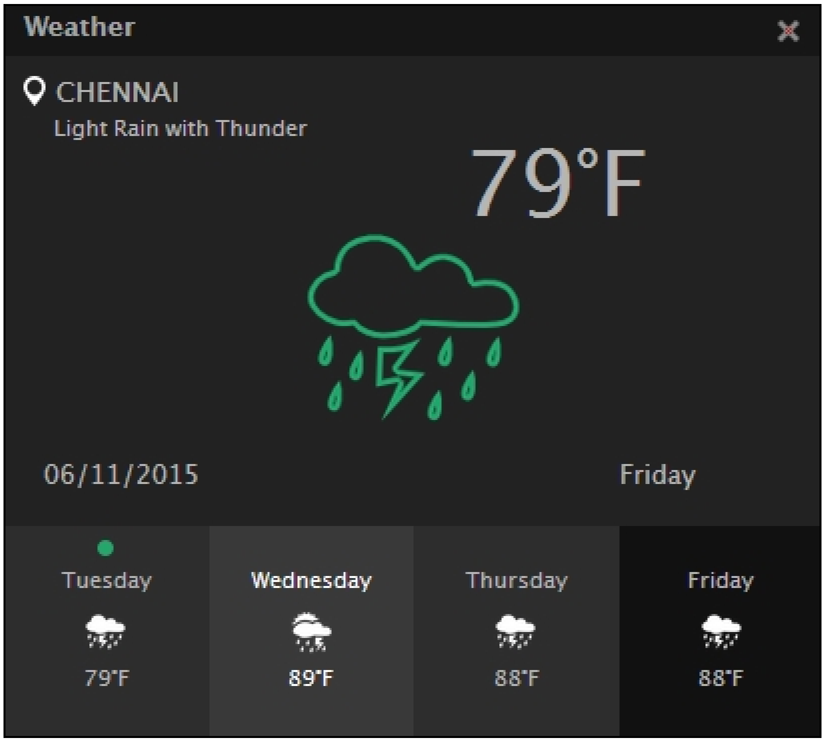

Figure 10.331: The Weather widget that appears in your dashboard

By default, the weather forecast for **Chennai** city will be displayed. To receive the weather forecast for the city of your choice, you can provide the name of the city against the **CurrentLocation** flag in the **[DASHBOARD_ SETTINGS]** section which is available in the **<EG_ INSTALL_ DIR>\manager\config\eg_ customdashboard.ini** file. If the city of your choice is from a country other than **INDIA**, then you have to set the **CurrentCountry** flag in the **[DASHBOARD_SETTINGS]** section which is available in the **<EG_INSTALL_ DIR>\manager\config\eg_customdashboard.ini** file to the country to which the city of your choice belongs to. For example, if you wish to view the weather report of New York, then you have to set the following:

```
CurrentLocation=NewYork
```

```
CurrentCountry=US|United States of America
```

Note:

- The **Weather** widget can be added only once in a dashboard.
- If the country mentioned in the **CurrentCountry** flag does not match with the city specified in the **CurrentLocation** flag, then the **Weather** widget will not be displayed.

2. News Update

If you wish to view the latest trending news in your dashboard, click on the **Add** button against the **News Update** widget in the **Misc** tab. The **News - India** widget will then appear as shown in Figure 10.332 listing the top trending news in India under various categories.

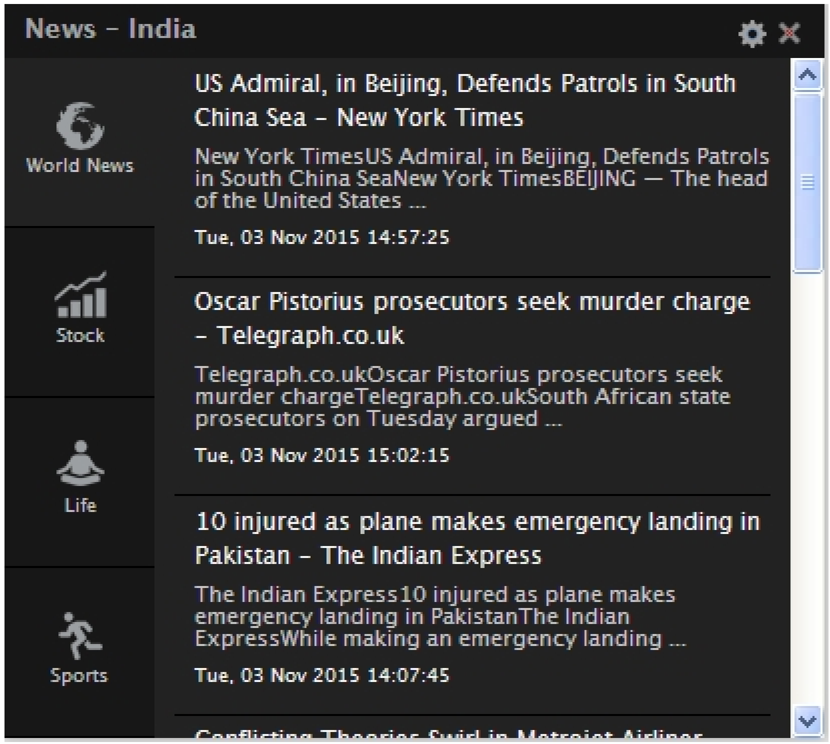

Figure 10.332: The News widget that appears in the dashboard

To view the news that is currently trending in the country of your choice, click on the  icon. Figure 10.333 will then appear prompting you to choose the **Country** of your choice.

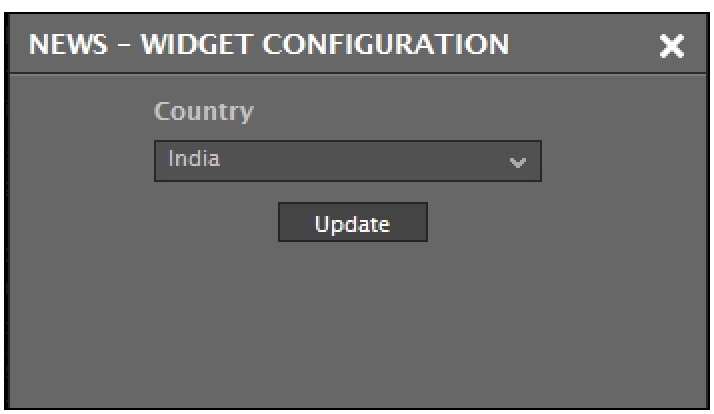

Figure 10.333: Selecting the country of your choice for which you wish to view the trending news

Clicking the **Update** button in Figure 10.333 after entering the **Country** of your choice will display the news that is trending currently in the country of your choice in the **News** widget.

Note:

- The **News** widget can be added only once in a dashboard.

11

Quick Insight

Often administrators may wish to track the values of certain key metrics in a single dashboard, so that they can proactively determine when the IT infrastructure may need attention. The metrics to be tracked may differ from one administrator to another. The eG Quick Insight offers an easy way for administrators to quickly track key metrics at each of the infrastructure tiers. With the help of this option, users can define infrastructure tiers to be monitored, add critical components for monitoring within each tier, and associate key metrics requiring closer observation with every component. Besides providing a holistic view of the environment at a single glance, this option enables users to focus on the more sensitive and critical components in the environment, and keep a close watch over their performance.

To build a custom view for monitoring, the following broad steps need to be followed:

Create a new view and define the layout of the view

- Add infrastructure tiers to the view
- Add servers to a tier
- Associate metrics with the servers

The sections to come will discuss each of these steps in detail.

11.1 Creating a New View and Defining its Layout

1. To achieve this click on the icon available in the **Monitor** tab. Then, select the **Quick Insight** option in the **Dashboards** tile in the eG monitor interface.

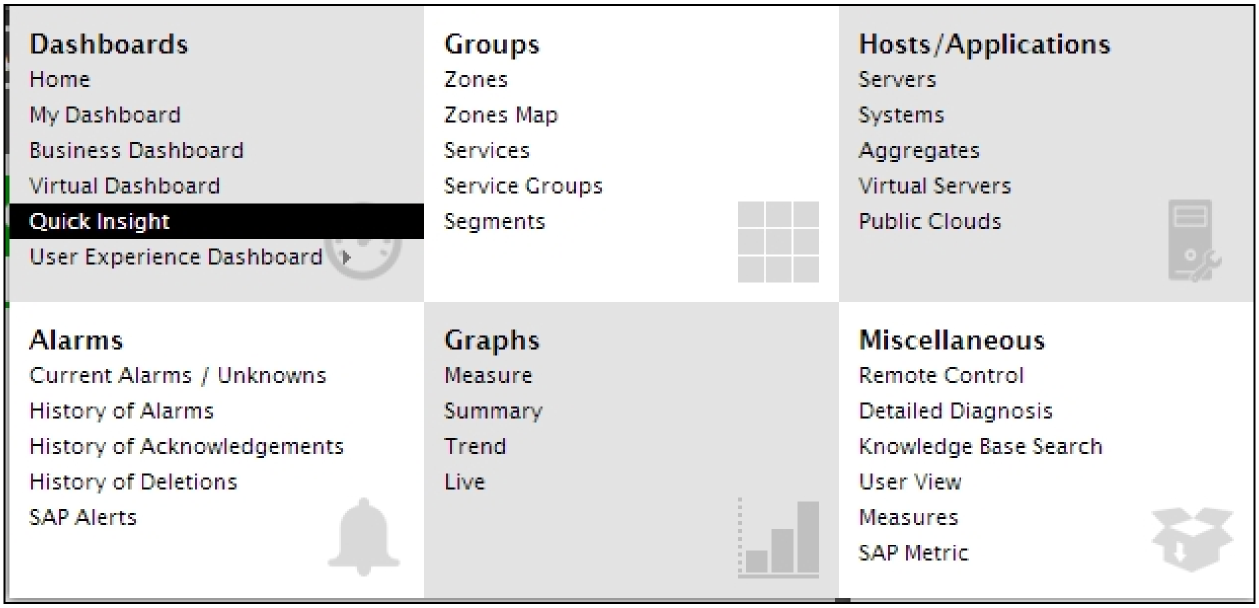

Figure 11.1: The Options menu

2. If custom views pre-exist, then upon selecting the **Quick Insight** option from the **Dashboard** tile, the eG Enterprise system will display the measures that have been configured for that custom view which has been set as the default view. If only one custom view has been defined so far, the eG Enterprise system will automatically set this view as the default view and display the state and the values of the measures configured for it. On the other hand, if no custom views have been defined yet, then you will be led straight to the **CONFIGURE LAYOUT** page of Figure 11.2.

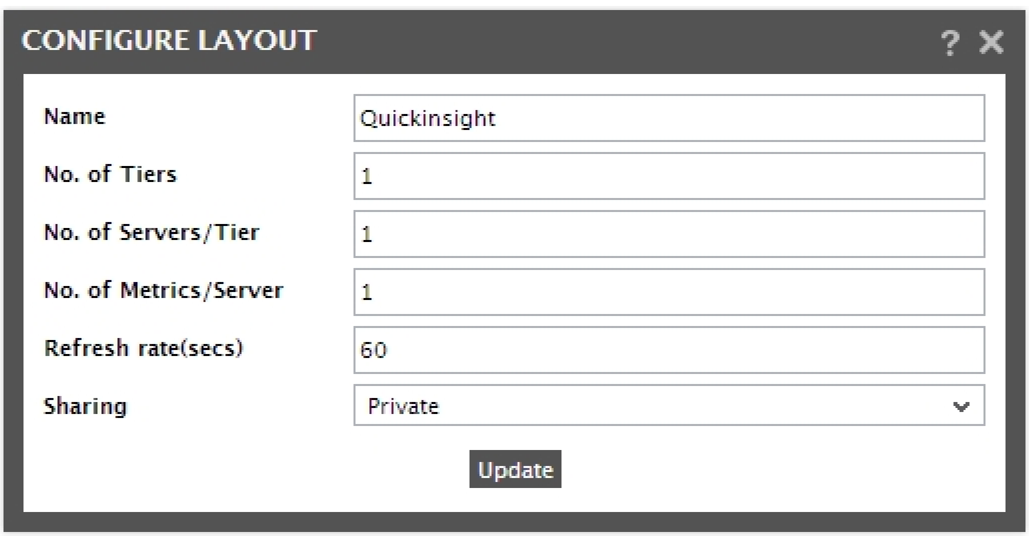

Figure 11.2: Creating a new view and defining its layout

3. In Figure 11.2, provide a **Name** for the new view.
4. In order to create a new Quick Insight view, you need to specify the following:

- **Tiers**: A custom view is divided into broad sections known as tiers. A tier can be based on geography or it can match the different infrastructure tiers (for e.g., the database tier, the web server tier), or it can even be a service / segment. While configuring the layout of your view using Figure 11.2, specify the **No. of Tiers** that your view should constitute. A maximum of 15 tiers can be configured.
- **Servers**: Each tier comprises of one or more **servers**. Typically, the critical servers in the monitored infrastructure are added to a tier. In the **No. of Servers/Tier** text box of Figure 11.2, specify the number of servers that you intend adding to a tier. A maximum of 15 servers can be added.
- **Metrics**: These are the key performance statistics that need to be extracted from every component in a tier. You can configure the number of measures that your layout should support by providing a number in the **No. of Metrics/Server** text box of Figure 11.2. Here again, the maximum limit of 15 prevails.
- **Refresh rate**: Specify the screen refresh time in seconds.
- **Sharing**: To facilitate collaboration and knowledge sharing among administrators, eG now allows users the flexibility to share useful quick insight views they create with other users. For sharing a view, you need to first set a share type by selecting the desired option from the **Sharing** list box in Figure 11.2. The options are as follows:
  - **Private**: If you set a view as **Private**, then all other users to the eG Enterprise system will be denied access to that view. Only the creator of the view will be able to access the view.
  - **Public**: If you set the view as **Public**, then only users with the following rights will have access to that custom view:
- Users with access to all the managed components in the environment
- User with access to one/more components that are included in the view being shared

  Note:

  - If a view has been granted **Public** access, then users with the privilege to monitor only a few components in that view will be able to take a look at only those statistics in the view that pertain to the components in his/her monitoring scope.
  - If a view has been allowed **Public** access, then users whose monitoring scope does not include any of the components in that view will not be allowed access to the view.
  - If a view that does not include any components/metrics has been allowed **Public** access, then such a view will not be available to any other user.
- **Share:** On the other hand, if you choose the **Share** option from the **Sharing** list, then you can pick and choose the users who need to be granted access to that view. To map users to a view, do the following:

5. Select the users to be granted access from the **Available Users** list box.
6. Then, click on **Grant** to grant the selected users access to the view. Doing so shifts the selection to the **Selected Users** list.
7. To revoke the access permission granted to a user, you can select the user from the **Selected Users** list, and click the **Revoke** button.
8. Finally, click on **Update** to save the changes.

Note:

- A user without access to **Quick Insight** cannot view any of the views shared with him/her.
- A view can be modified / deleted only by the user who creates it. The shared users on the other hand, can only view the Quick Insight data, and cannot modify/delete the view.
- If a user has been granted **Share** access to a view, but he/she has the right to monitor only a few components that are included in the view, then such a user will be able to take a look at only those statistics in the shared view that pertain to the components in his/her monitoring scope.
- If a user has been allowed **Share** access to a view, but he/she does not have the privilege to monitor any component included in the view, then such a view, despite being shared, will not be available to that user.
- If a user has been allowed **Share** access to a view to which no servers/metrics have been assigned, then such a view will not be available to that user upon login.

9. Finally, specify how frequently the configured measures have to be refreshed, in the **Refresh Rate** text box of Figure 11.3, and click the **Update** button to save the changes. Upon updating, Figure 11.3will appear displaying the layout that was just designed (see Figure 11.4).

Figure 11.3: Designing a view

10. To view the complete list of custom views, click on the icon in the Quick Insight menu at the right top corner of Figure 11.5. The existing custom views (including the newly configured view) will appear (see Figure 11.3).

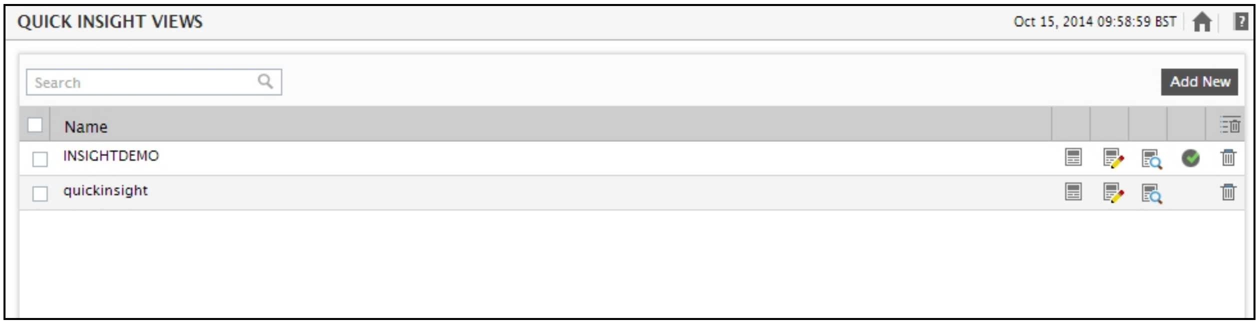

Figure 11.4: Viewing the list of custom views

11. To add a new view, click the **Add New** button in Figure 11.4. To set any of the listed views as the default view, click the icon corresponding to the view name in Figure 11.5. To delete a view, select the check box corresponding to the view name in Figure 11.5, and click the icon. Choose **Select all** and then click the icon, if all the listed views need to be deleted.

12. If other users have shared their views with you, then those views will be listed separately in the **Views Shared by other Users** section as depicted in the Figure 11.4. Against every shared view, this section also displays the name of the view creator - i.e., the user who originally created the view and shared it with the current user.

Figure 11.5: Viewing the Shared views

13. To configure the tiers/servers/measures to be associated with a view, click on the icon against a view in Figure 11.5. This will lead you back to Figure 11.5 where such configurations can be performed.

14. To view the state and the values of the measures that have been configured under a view, click the **View** ( ) icon against the view name. Figure 11.6 will then appear.

15. Clicking on the **Layout** ( ) icon against a view, will lead you to Figure 11.6, using which the layout of the

view can be modified.

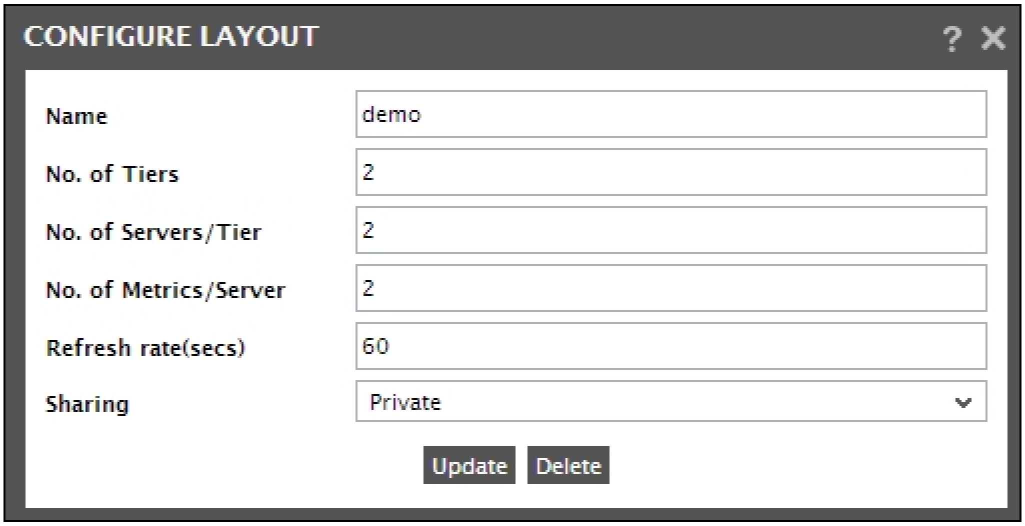

Figure 11.6: Modifying a layout

16. In the Modify mode, all the layout specifications except the **Name** of the view, can be modified. Additionally, a **Delete** button appears in the modify mode (see Figure 11.6), clicking on which deletes the displayed layout.

11.1.1 Creating a Tier

Now that the layout has been set, proceed to add a new tier. To achieve this, do the following:

1. When Figure 11.7 appears, provide a **Tier name** and click on the **UPDATE** button to register the changes.

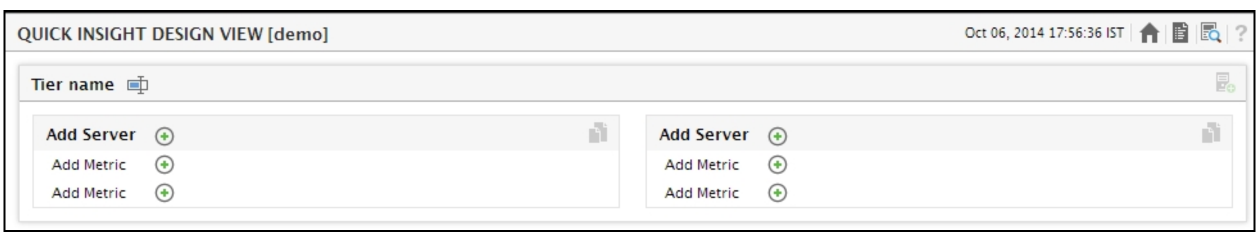

Figure 11.7: Configuring a tier

2. The **DESIGN VIEW** will now reflect the name of the tier(see Figure 11.8).

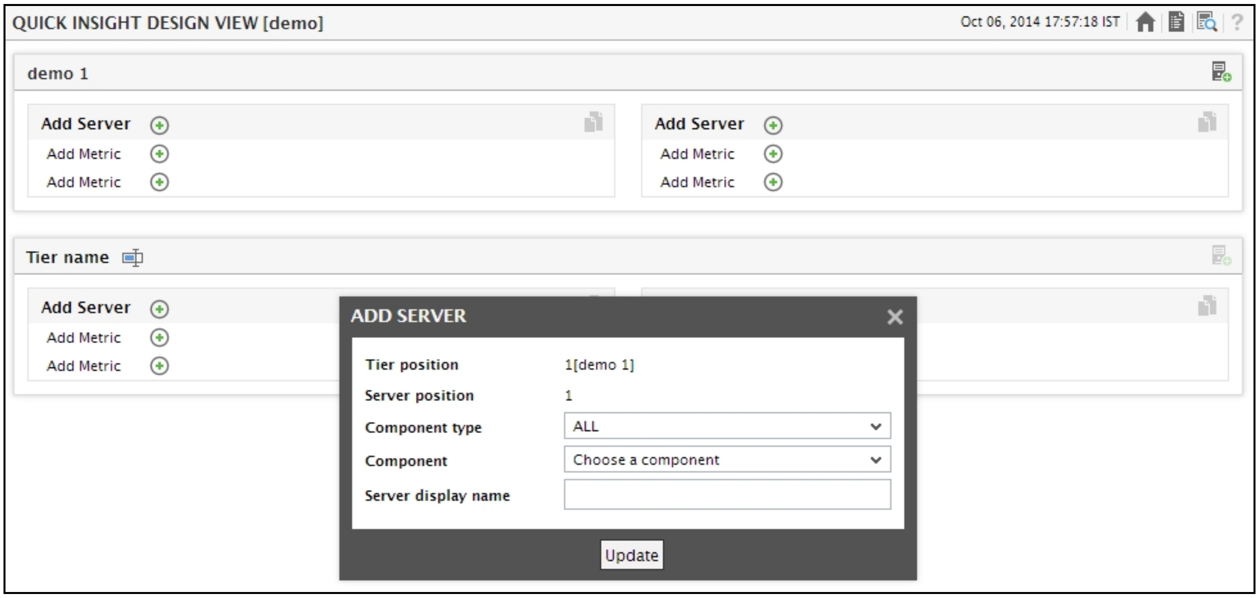

Figure 11.8: The newly added tier appearing in the layout

3. To add a new tier, click on the **Tier Name**. Once the addition is updated, the name of the new tier will be visible in this page. You can also modify or delete the tier name. Similarly, clicking on the **`[Add server]`** link here, will allow you to add a new server to a tier.

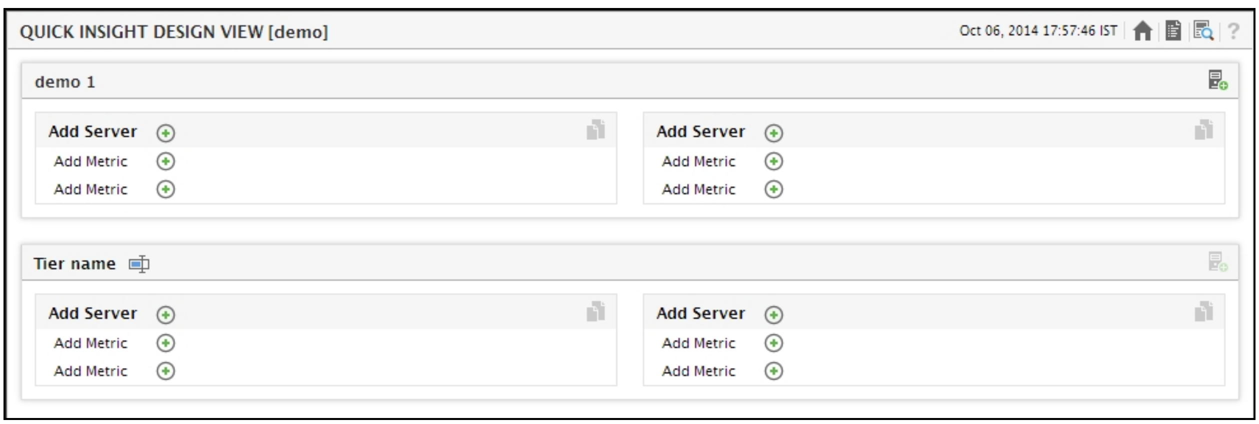

Figure 11.9: Configuring multiple tiers

4. To modify a tier, click on the name of the tier in Figure 11.9. Figure 11.10 will then appear displaying the tier name. Change the name if required and update the changes by clicking on the **Update** button. The displayed tier can also be deleted using the **Delete** button in Figure 11.9

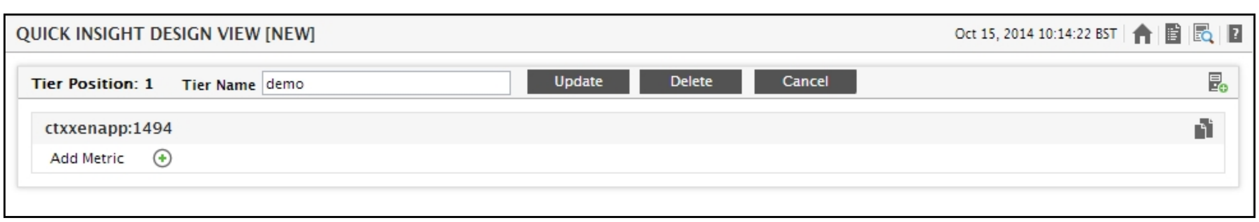

Figure 11.10: Modifying a tier name

11.1.2 Adding a Component to a Tier

Next, proceed to add a component to a tier, by doing the following:

1. Click on the **[Add server]** link under a tier (see Figure 11.10).
2. Figure 11.11 will then appear.

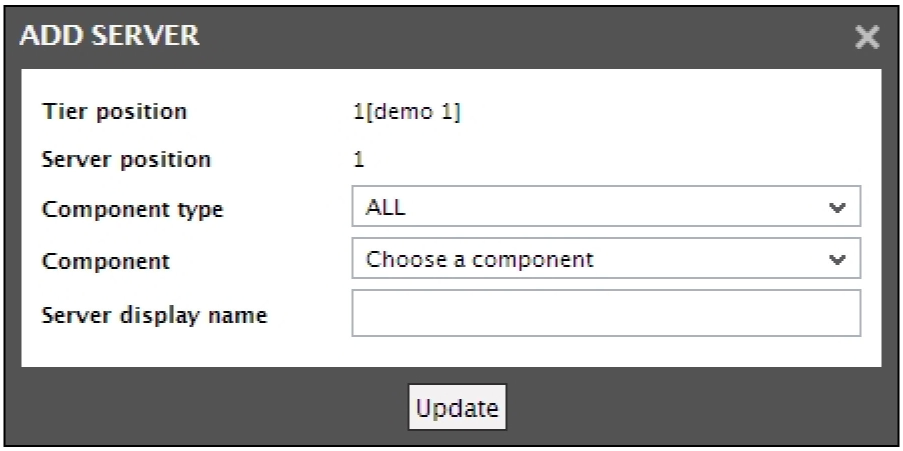

Figure 11.11: Adding a server

3. Select a **Component type** to be added, and from the list of components of the chosen type, select the **Component** to be added.
4. Provide a display name for the component in the **Server dislay name** text box, and click the **Update** button to register the changes.
5. Figure 11.11 will then appear displaying the newly added server name in the place of the **[Add server]** link that was earlier clicked on.

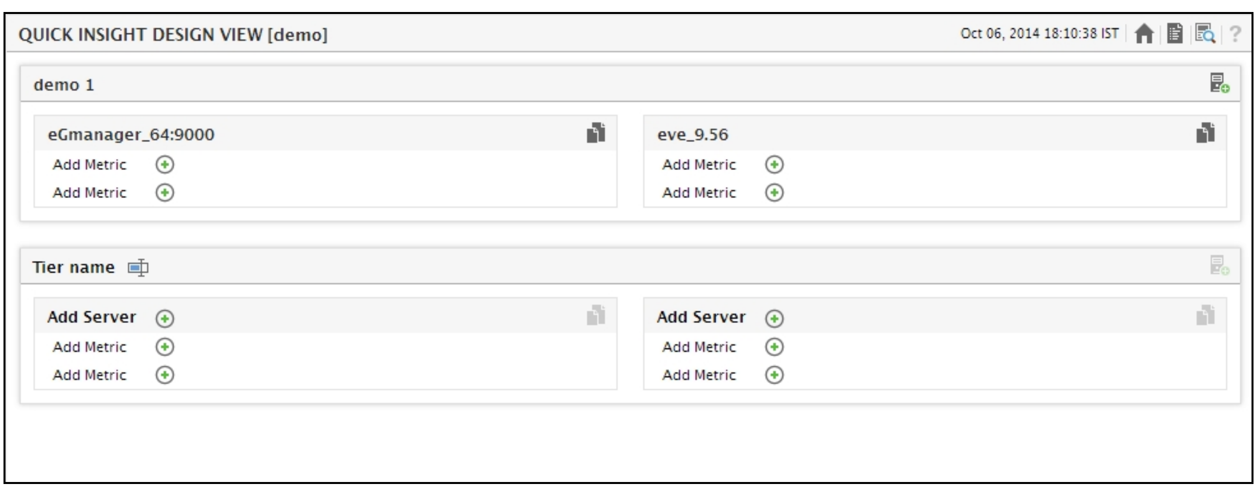

Figure 11.12: The name of the new server appearing in the layout

6. The added server will then appear in this page. To modify the server details, click on the server name here and then make changes to it in the **Configure Server** page that will subsequently appear. To save users the trouble of adding the servers individually, this page houses the icon, using which all the servers that need to be added to a tier can be added at one shot. The button appears next to the tier name in this page.

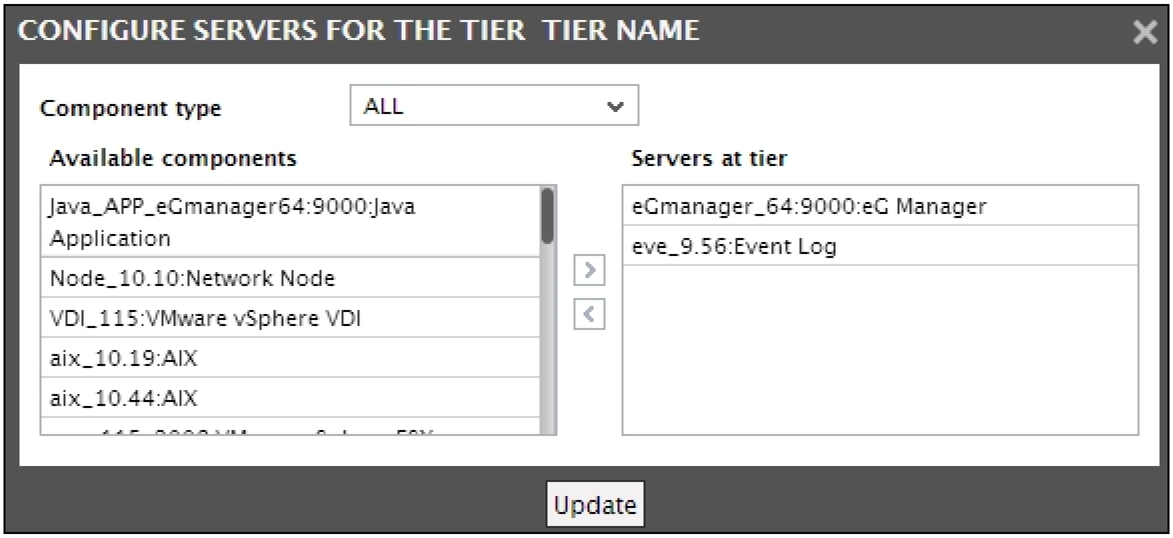

Figure 11.13: A page displaying all the component that are available for association with a tier

7. Figure 11.13 also houses a **Servers at tier** list box that lists multiple **[ADD SERVER]** slots, equal to the number of servers per tier configured in the layout.

8. To assign a component to a particular slot, first select a **Component type** as depicted by Figure 11.14. Then, from the list of components of the chosen type, select the component that is to occupy a particular slot (see Figure 11.14).

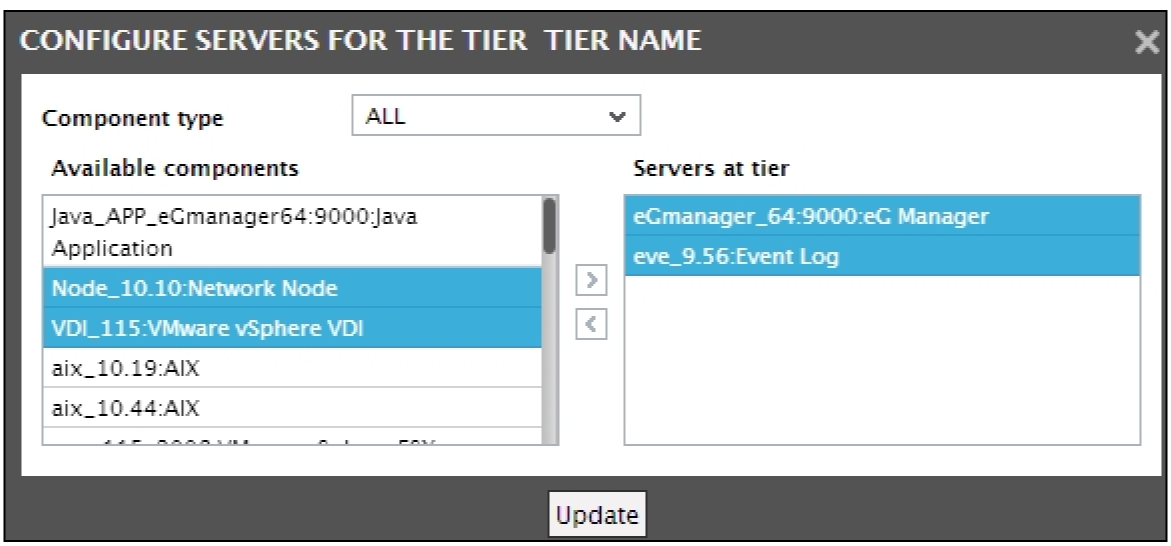

Figure 11.14: Associating a component with a slot

9. Next, with the component selected, select the **[ADD SERVER]** option (from the **Servers at tier** list) that corresponds to the slot that is to be occupied by the chosen component (see Figure 11.14). For example, if the chosen component needs to be assigned the second position in the tier, then select the second **[ADD SERVER]** option in Figure 4. Then, click the **>** button in Figure 11.14.
10. Figure 11.14 will then appear indicating that the selected component(s) has replaced the **[ADD SERVER]** slot (s) that was chosen (see Figure 11.14).
11. Similarly, you can assign components of different types to occupy the **[ADD SERVER]** slots.
12. Finally, click the **Update** button in Figure 11.14.
13. Upon clicking the **Update** button in Figure 11.14, you will return to the **DESIGN VIEW** (see Figure 11.15), which will reflect the recent changes in the server configuration.

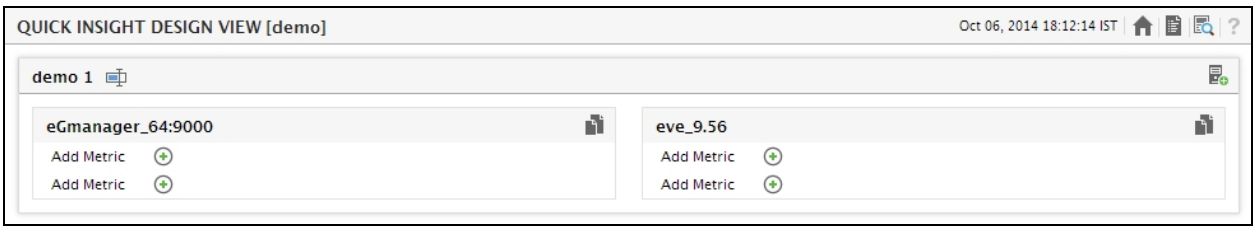

Figure 11.15: Adding multiple servers to a tier at one shot

14. Using Figure 11.15 that appears, change the **Server name** of the server that was clicked on, and click the **Update** button therein. In the modify mode, a **Delete** button appears, which when clicked on removes the server configuration from the database (see Figure 11.15).
15. In the same way, the display names for the other servers configured for a tier can also be modified. Figure 11.15 depicts the **DESIGN VIEW** after the display name change.
16. Similarly, servers can be configured for all the tiers in a custom view (see Figure 11.15).

11.1.3 Associating Metrics with the Components in a Tier

Next, proceed to associate every component in the tier with performance metrics. To do so, click on the **[Add metric]** link in Figure 11.16.

1. Figure 11.16 will then appear wherein the details of the measurement will have to be specified.

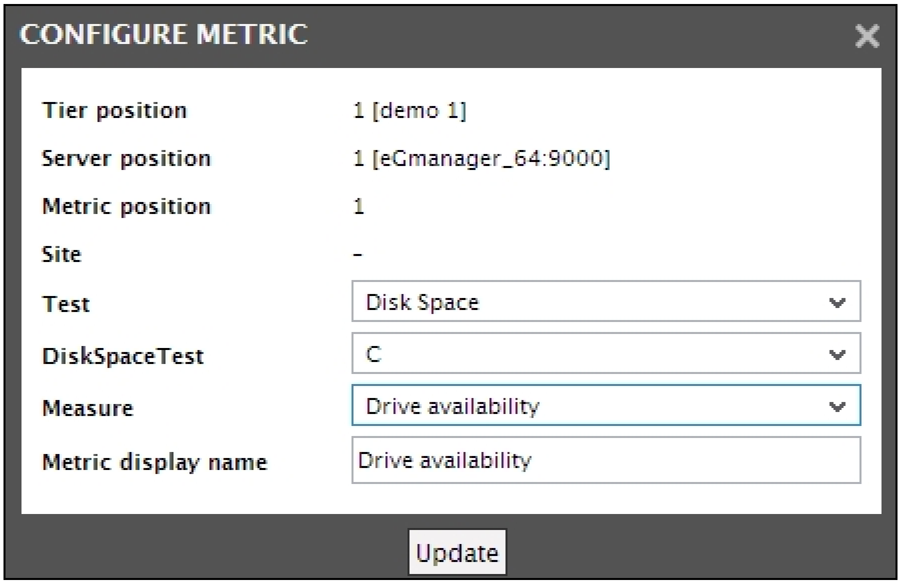

Figure 11.16: Configuring a measure

2. If the measurement to be added pertains to a particular site, select the **Site**.
3. Select a **Test**, and if the test takes descriptors, pick a **Descriptor**.
4. Then, select a **Measure**, and provide a display name for the measure in the **Metric name** text box.
5. Finally, click the **Update** button.
6. The newly added measure will now appear in the **DESIGN VIEW** (see Figure 11.17).

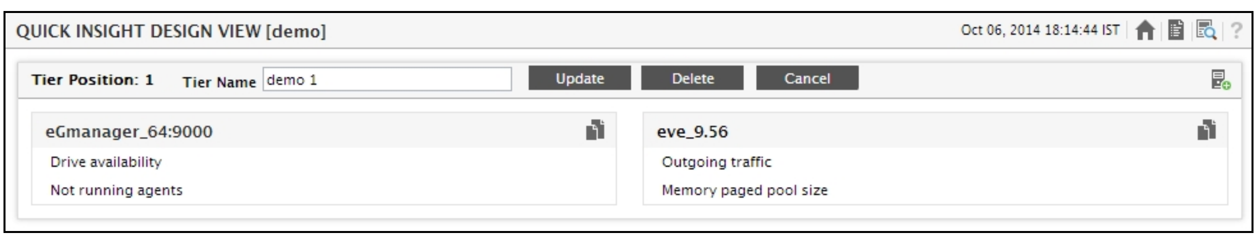

Figure 11.17: Associating a measure with a tier component

7. Similarly, multiple measures can be associated with a single component (see Figure 11.18).
8. If the measures configured for one component in a tier need to be associated with a few / all the other components in the tier, then, click on the icon next to the server name (for which at least one measure has been configured) in Figure 11.17. Figure 11.18 will then appear. The **Available metrics** list in Figure 11.18 will display the measures that have been configured for the displayed **Server name.** From this list, select the measures to be copied to the other components in the tier. Then, from the **Available servers** list that lists all the other components in the tier, select those components to which the chosen measures need to be copied. Finally, click the **Update** button in Figure 11.19.

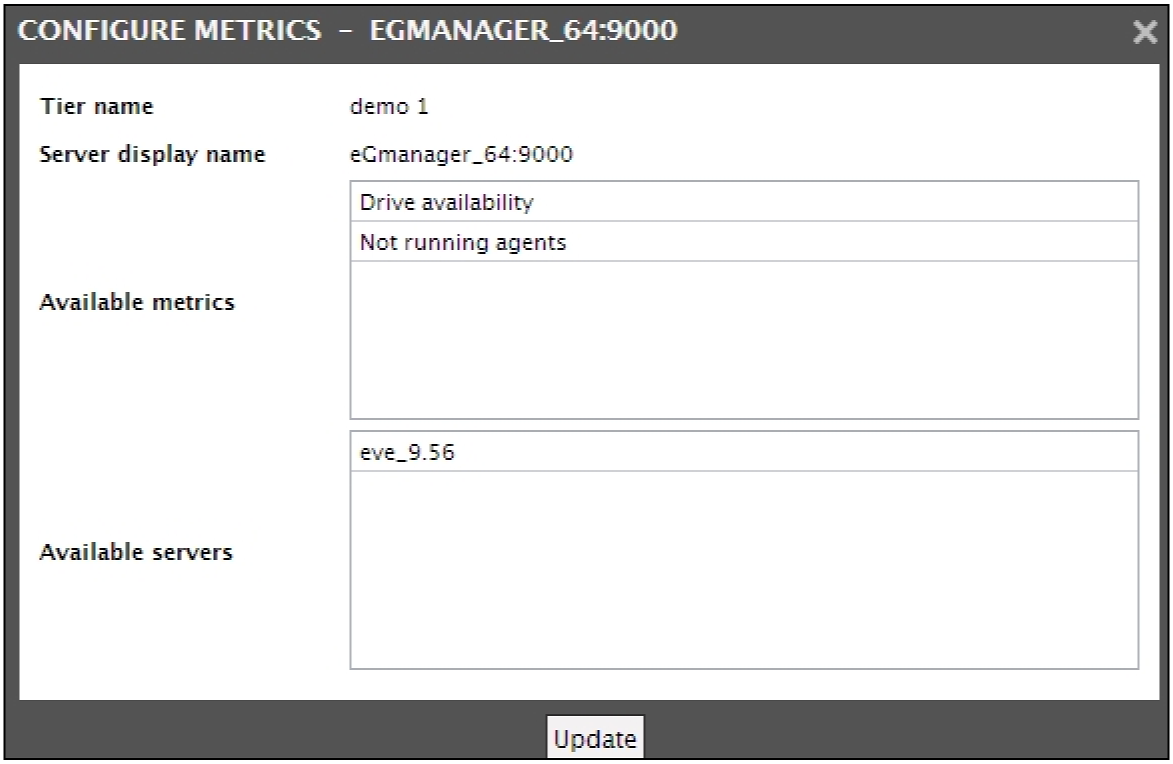

Figure 11.18: Applying measures to a few / all the other components in a tier

9. The changes will be reflected in the **DESIGN VIEW** as depicted by Figure 11.18.

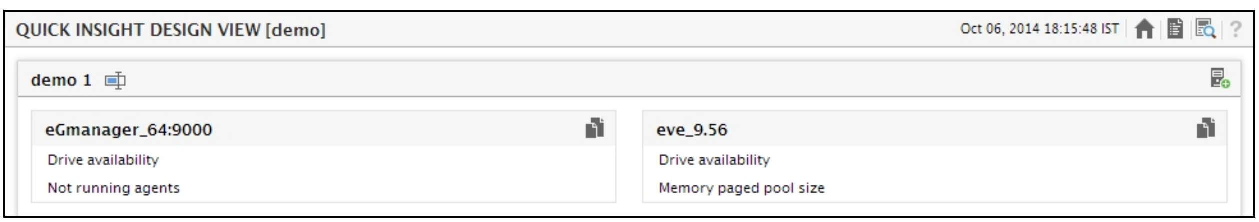

Figure 11.19: Copying measures to other components in a tier

10. Similarly, measures can be associated with the components in all other tiers of the custom view.
11. Once the layout is designed, begin monitoring the behavior of the associated metrics by clicking the icon at the right top corner of Figure 11.19

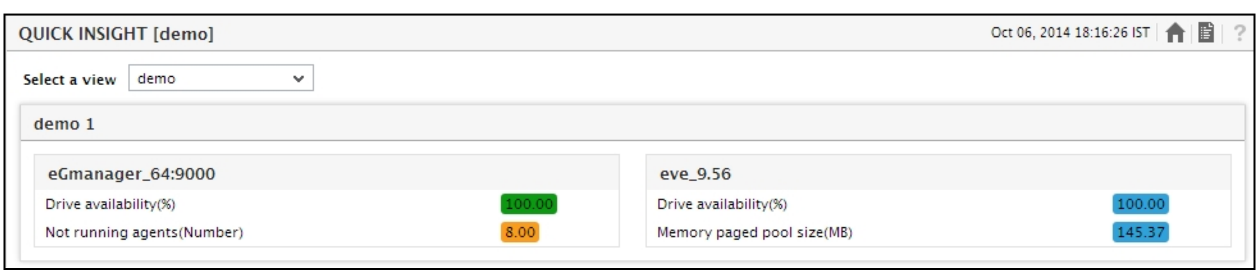

Figure 11.20: Monitoring the measures

12. eG will then begin displaying the measure values in the custom view that was designed (see Figure 11.20). The measures will get updated based on the **Refresh Rate** set using Figure 11.20. While the **NORMAL** measure values will be displayed in green, the **CRITICAL** values will appear in red, **MAJOR** ones in orange, **MINOR** values in yellow, and **UNKNOWN** values in blue.
13. The **Drive Availability** measure of the **eGmanager** server in Figure 11.20 displays **100**, indicating the non-availability of the web server. To know more about this issue, click on the **Drive Availability** measure in Figure 11.20.
14. Figure 11.20 will appear displaying the layer model of the **eGmanager** web server indicating the exact layer of the problem, the tests that detected the problem, and the measures that were rendered abnormal.

12

Graphs

So far in this chapter we have illustrated eG Enterprise's capability to lead the supermonitor to the root cause of the problems in his/her infrastructure. In more complex scenarios than the one depicted, manual analysis may become necessary.

eG Enterprise includes a variety of graphing capabilities for manual diagnosis. eG Enterprise supports the following graph types:

- Measure
- Summary
- Trend

Figure 12.1 shows the measurement graph that is obtained via the **Measure** option under the **GRAPHS** menu. This graph shows the instantaneous values of the measures reported by the agents for the TcpTest for a web server. A measure graph is used by the supermonitor to plot the instantaneous value of any of the measurements made by eG Enterprise with time of day. The measure graph is accessed via the **Measure** option under *GRAPHS* in the menu across the top of this page. Follow the steps given below to plot a measure graph.

1. The service for which the graph is to be plotted is first selected from the **Service** list box. This step is optional. If no specific service is chosen, the **Component** list box is populated with all the components being currently managed by the eG manager. Alternatively, if a specific service is chosen, the **Component** list only includes components that are related to the service under consideration.
2. Next, a specific component associated with this service can be selected from the **Component** list box. If there are too many components in the list to choose from, you can narrow your search further by using the **Search** text box. Specify the whole/part of the component name/type to search for in this text box, and click the right-arrow button next to it. The **Component** list will then be populated with all component names and/or component types that embed the specified search string (see Figure 12.1). Select the component of your choice from this list.
3. All the layers that map to the selected component will be available in the **Layer** list box and the supermonitor can choose any layer.
4. The tests corresponding to the layer under consideration form the options of the **Test** list box. The required test can be chosen from this box.
5. If the test is executed by multiple **Measurement Hosts**, then you can select a particular **Msmt Host** for graph generation.
6. The measurements reported by this test appear in the **Measurements** list box, one or more of which can be picked.
7. To override the default timeline, click on the  icon. The settings page appears, the period for which the variations of the selected measurements have to be analyzed can be specified using the **Timeline** list. You

can either provide a fixed time line such as 1 hour, 2 days, etc., or select the **Any** option from the list to provide a **From** and **To** date/time for graph generation.

8. The **Graph** button can be utilized to view the graph. The *D***ata** button will enable the supermonitor to view the table of data corresponding to the graph as in Figure 12.1. The user can print the graph using the *P***rint** button. Similarly, the graph data can also be saved as a PNG image using the **Save** button in Figure 12.1.

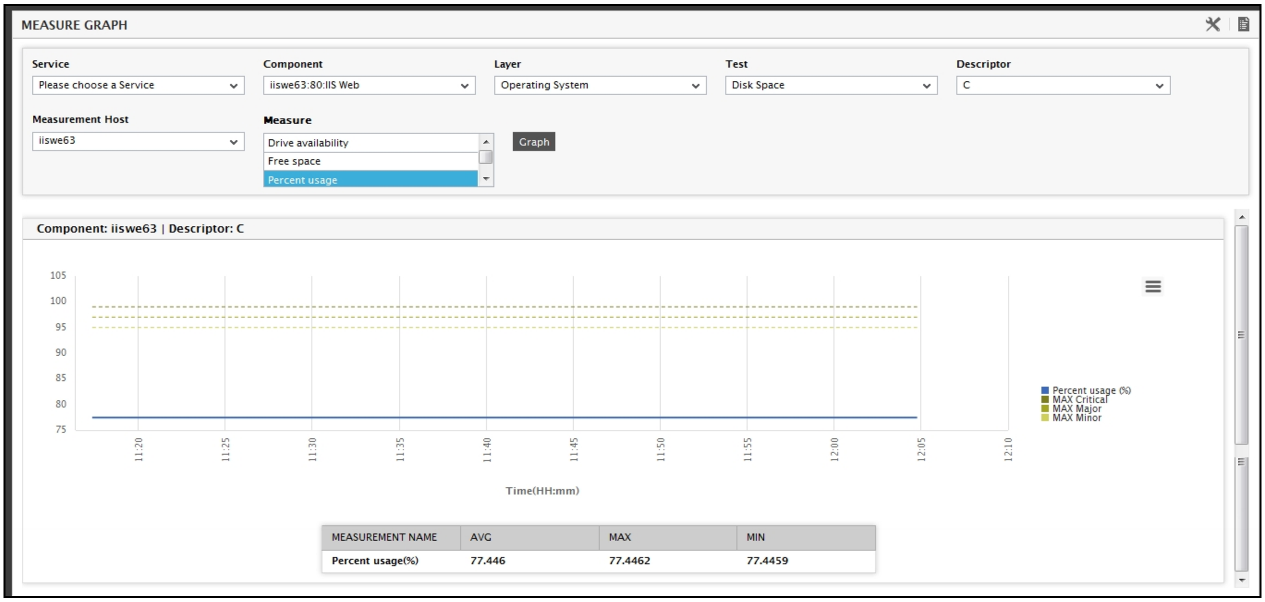

Figure 12.1: Figure 12.1: eG Enterprise’s measurement graph

In Figure 12.1, the blue line denotes the actual variations in the measurement values. The lines in various shades of gold and green, indicate the different levels of thresholds (high, medium, and low) that have been assigned to the measure. An alarm is raised when the actual measurement crosses any of the multiple threshold values that have been set. The alarm priority though, will differ according to the nature of the threshold that was violated. For example if a MIN CRITICAL or MAX CRITICAL threshold value is violated, then a critical priority alarm will be generated by the eG Enterprise system.

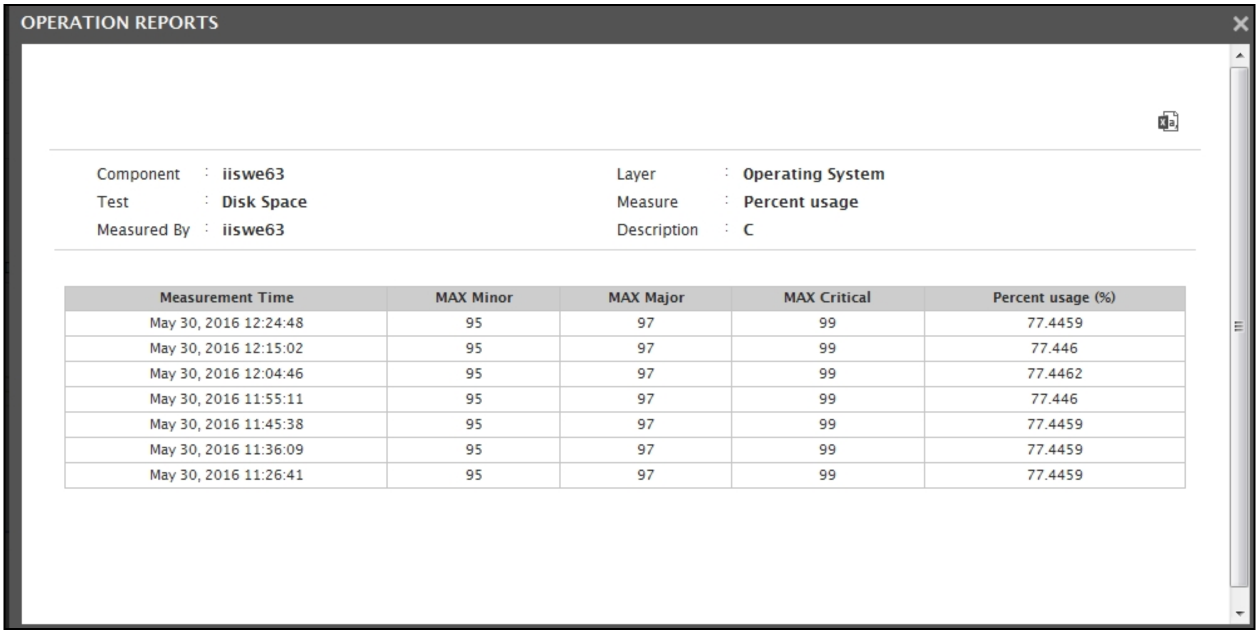
OPERATION REPORTS

Component : iiswe63
Test : Disk Space
Measured By : iiswe63
Layer : Operating System
Measure : Percent usage
Description : C

Measurement Time	MAX Minor	MAX Major	MAX Critical	Percent usage (%)
May 30, 2016 12:24:48	95	97	99	77.4459
May 30, 2016 12:15:02	95	97	99	77.446
May 30, 2016 12:04:46	95	97	99	77.4462
May 30, 2016 11:55:11	95	97	99	77.446
May 30, 2016 11:45:38	95	97	99	77.4459
May 30, 2016 11:36:09	95	97	99	77.4459
May 30, 2016 11:26:41	95	97	99	77.4459

Figure 12.2: Figure 12.2: Measurements depicting the variation of Percentage usage with time of day

Figure 12.3 depicts a summary graph that gives an overall picture of the percentage of time the measurements pertaining to a component were in a normal, critical, major, minor, or unknown state, over a period of time. The steps involved in plotting this graph are the same as that of a measure graph except for minor changes. Unlike the measurement graph, this graph pertains to a single measurement alone. The measurements reported by this test appear in the *Measurements* list box, one of which can be picked. The supermonitor can also opt to view the hourly, daily or monthly variations of measures in the case of summary graphs by selecting an appropriate option from the **Duration** list box. Like the measure graph, the summary graph can also be printed or saved.

In the graph, red indicates a critical state, orange indicates the existence of a major problem, yellow indicates the existence of a minor problem, green denotes normalcy, and blue implies that the state is unknown.

Figure 12.3: Figure 12.3: Summary graph showing the percentage of normal, critical, major, minor, and unknown measurements

Since Internet traffic is very bursty, using the measurement graphs over a long time window (greater than a week) to view trends in the measurements can be very difficult. To make it easier to analyze trends, eG Enterprise monitors and stores trend data on an hourly, daily, and monthly basis. The eG user interface allows the trend data to be plotted on a web browser. Figure 12.4 depicts a trend graph; by default, a trend graph takes the minimum and the maximum value of measurements over a period of time into consideration. Accordingly, the **Graph** list displays **Min/Max** by default. Alternatively, you can even generate a trend graph where the **average** values of a chosen measure are plotted over a period of time. To achieve this, simply select the **Average** option from the **Graph** list in Figure 12.5. For instance, you can plot a trend graph that depicts how many TCP connections on an average were established with a critical Terminal server, every day during a couple of weeks (see Figure 12.5); besides indicating the normal load on the Terminal server, such a graph also enables you to understand whether the Terminal server has been adequately tuned to handle higher loads, and thereby helps you make effective sizing recommendations for the future. Likewise, you can also choose **Sum** as the **Graph** type to view a trend graph that plots the sum of the values of a chosen measure for a specified timeline. For example, a **Sum** graph of TCP connections to a Terminal server (see Figure 12.6) serves as an accurate indicator of how busy the Terminal server was during a given period.

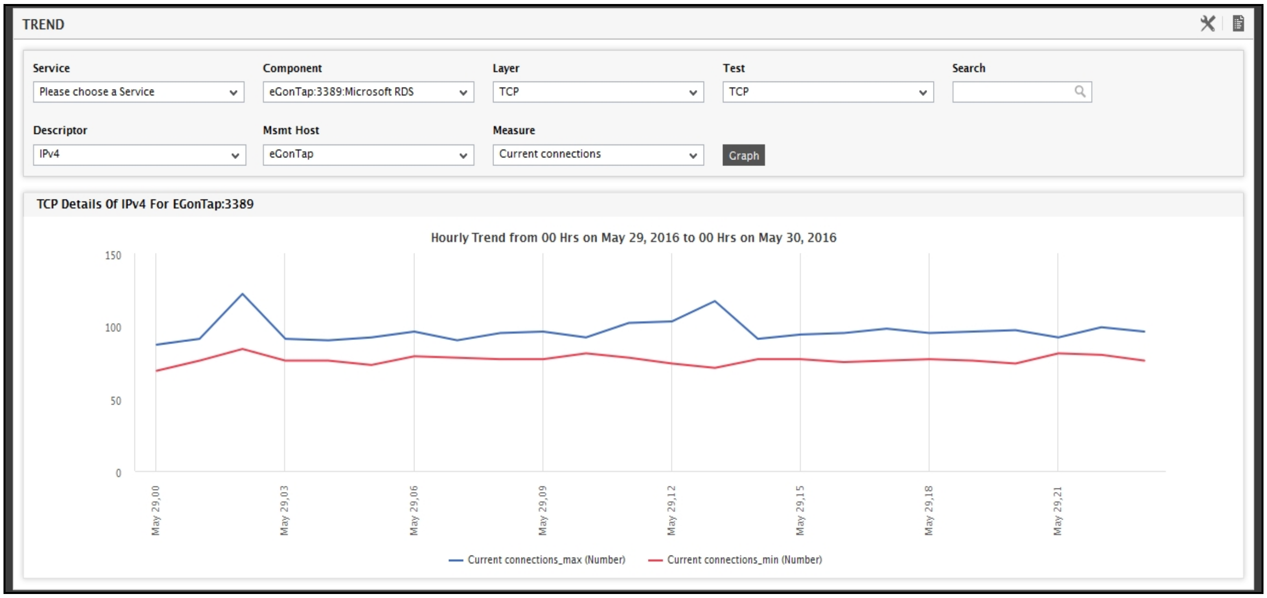

Figure 12.4: Figure 12.4: A Min/Max Trend Graph

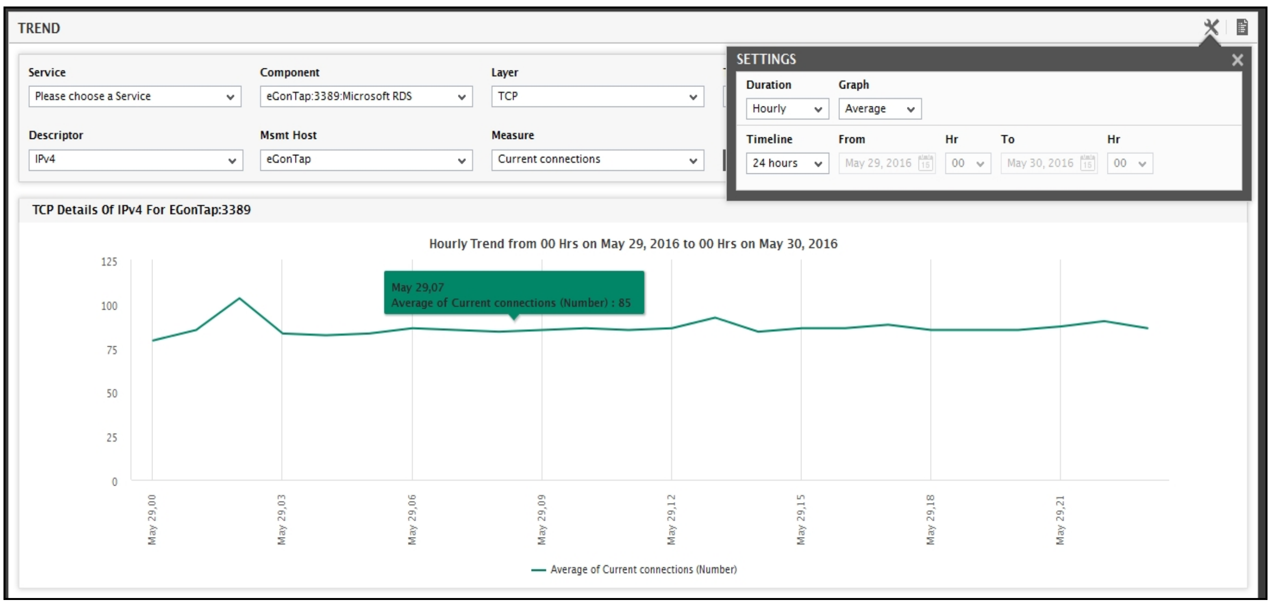

Figure 12.5: Figure 12.5: An Average Trend Graph

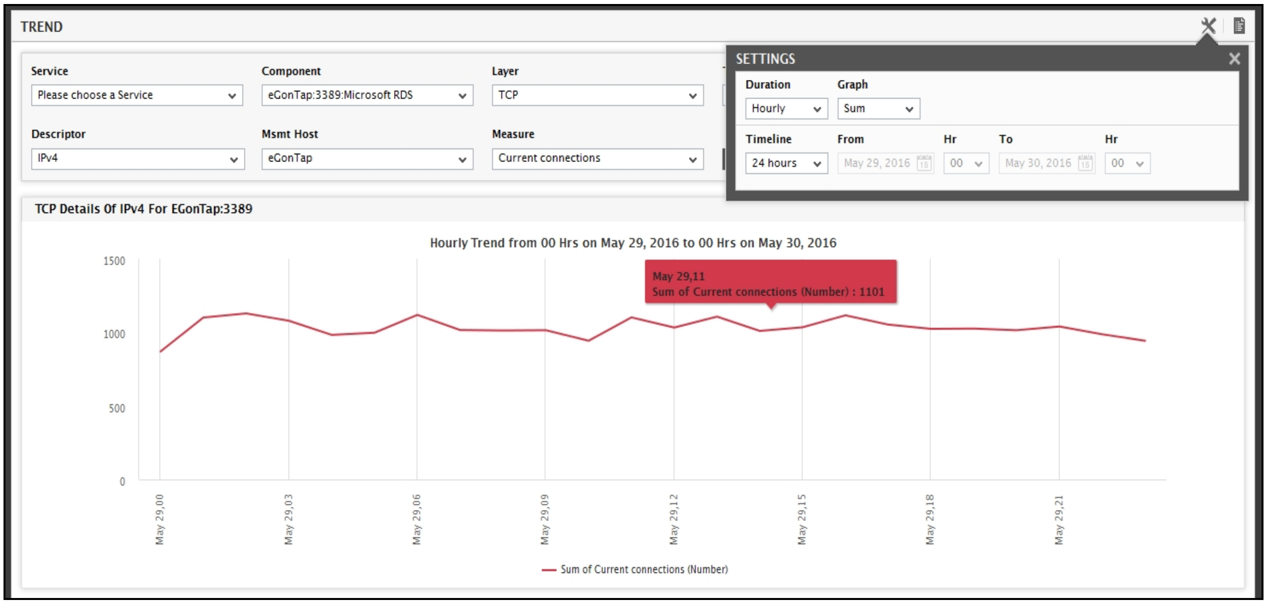

Figure 12.6: Figure 12.6: A Sum Trend Graph

Note:

The capability of the eG manager to compute the **Sum** and **Average** of metrics is governed by the **Compute average/sum of metrics while trending** flag in the **MANAGER SETTINGS** page in the eG administrative interface. By default, this flag is set to **Yes** indicating that eG Enterprise computes and stores the average and sum values of every performance metric in the database, by default. If, for some reason, you want to disable this capability, just set this flag to **No**, and **Update** the changes.

Note:

By default, the **Sum** trend graphs are generated for all measures. You can, if you so need, enable this capability for specific measures by following the steps given below:

- Edit the **eg_ui.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
- Set the **TrendSumForAllMeasures** flag in the **[MISC_ARGS]** section of the file to **false** (default is **true**).
- Save the **eg_ui.ini** file.
- Next, edit the **eg_tables.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory.
- In the **[MEASURE_TOTAL]** section of the file, provide an entry for the measure for which **Sum** trend graphs need to be generated, in the following format:

  <TestName>:<MeasureName>=<DisplayName>

  For example, to generate a **Sum** trend graph for the measure *Current_connections* reported by TcpTest, your specification would be:

  TcpTest:Current_connections=Current TCP Connections
- Finally, save the **eg_tables.ini** file.

Note:

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 in the **MONITOR SETTINGS** page of the eG administrative interface.

12.1 Live Graph Display

Administrators of large, mission-critical environments are expected to be on high vigil 24 x 7, as even seemingly minor aberrations in performance could prove to be fatal for the infrastructure. Particularly, critical components of such infrastructures demand continuous attention. Therefore, it is essential for the monitoring solution in use to report even the smallest of variations in performance of such components, in real-time. Towards this end, the eG monitoring suite provides the **LIVE GRAPH DISPLAY** option which displays real-time graphs of key performance metrics relating to critical components in the infrastructure, allowing the administrator to keep a constant watch on the measure behavior, observe variations in the measures as they occur, and detect anomalies on-the-fly. Besides, eG Enterprise provides the option to plot historic data alongside the current data in the graph, so that an effective comparison of the past and the present performance can be performed, and appropriate performance decisions can be easily taken.

To view the live graphs, select the **Live** option from the **Graphs** tile in the eG monitor interface.

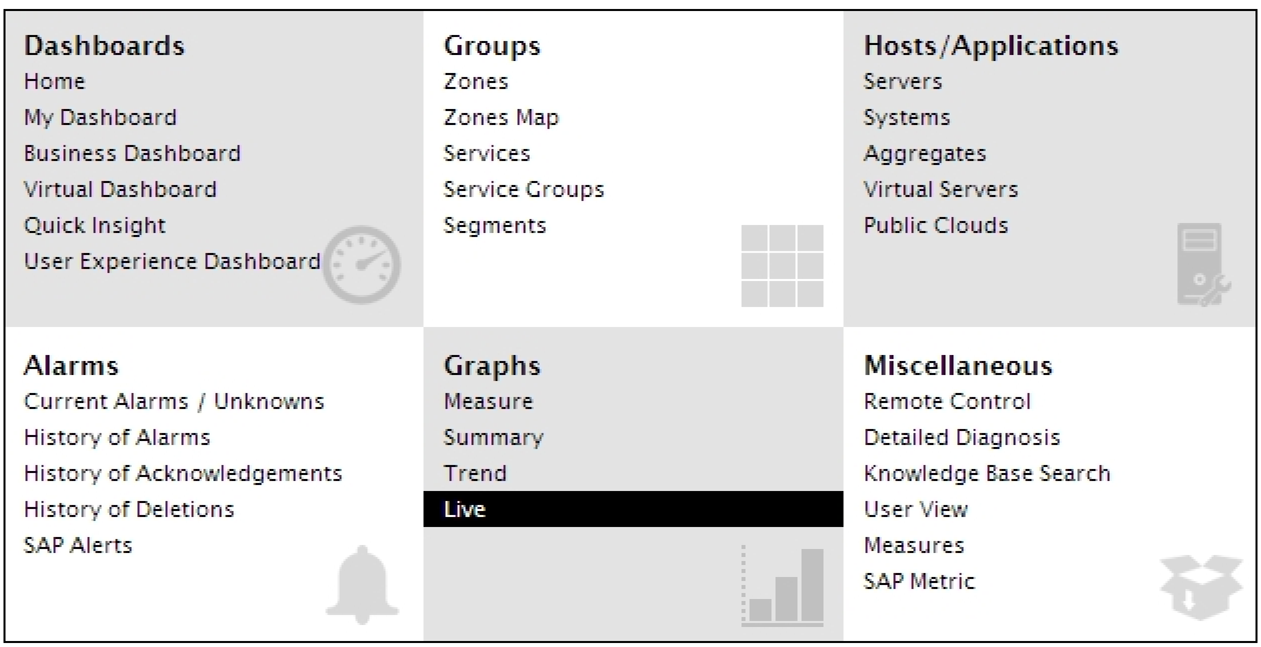

Figure 12.7: Selecting the LIVE option

If custom views for the live graphs pre-exist, then, when the **Live** option is clicked, the graphs configured for that custom view which has been set as the default view, will appear. If only one custom view has been defined so far, the eG Enterprise system will automatically set this view as the default view and display the graphs configured for it. If no custom views have been configured yet, then a message to that effect will appear.

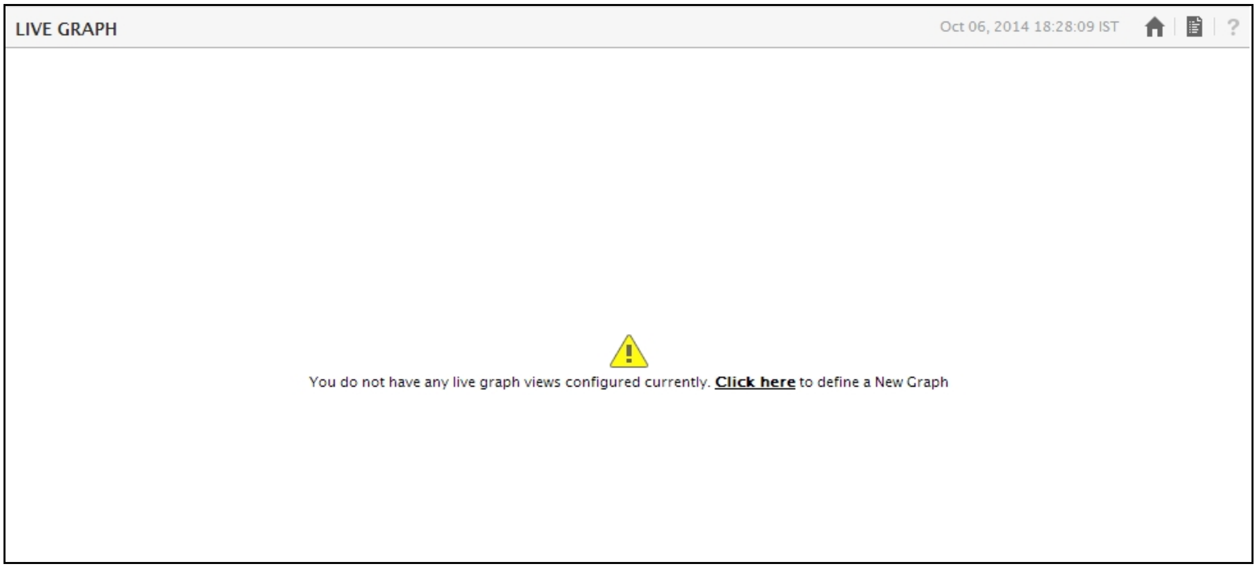

Figure 12.8: The message indicating that no custom views exist

To create a new view, click on the **Click here** hyperlink in 12.1. 12.1 will then appear allowing you to configure the new custom view.

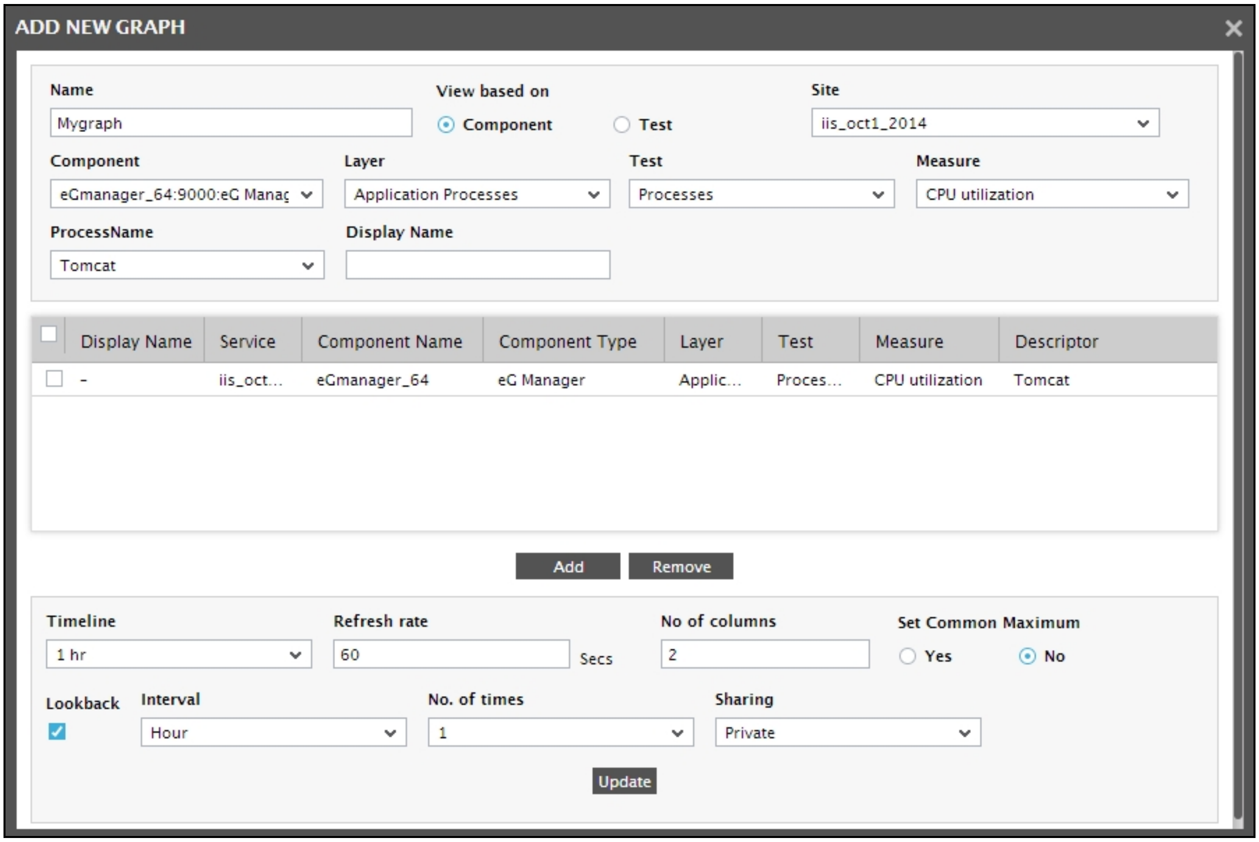

Figure 12.9: The CONFIGURE GRAPH page with the Component option chosen

To configure the graphs to be associated with the new view, do the following:

1. Provide a **Name** for the view.
2. Next, select the criterion based on which your view is to be designed. A view can be built on a **Component** or a **Test** basis. Select the **Component** option if you want to search for and pick those measures that are reported by the tests executing on a specific component(s). Then, proceed as follows:
   - If the measure to be plotted live pertains to a particular site, then select the site from the **Service** list box of Figure 12.10. This is an optional field, and will appear only if at least one service is configured in the environment.
   - If a **Service** is selected, the components associated with that service will appear in the **Component** list box. If no service is selected, then all managed components in the environment will appear in the **Component** list box. Choose the component for which a measure graph needs to be generated.
   - The layers of the chosen component will then be listed by the **Layer** drop-down list box. From this list box, select the layer with which the measure of interest is associated.
   - The **Test** list box will then be populated with the tests mapped to the chosen **Layer**. Select the test which collects the performance metric(s) to be plotted in the graph.
   - From the **Measure** list box that lists all the measures that are reported by the chosen **Test**, pick the measure for which a graph is to be generated.

- If the test that was selected previously takes descriptors, then select a descriptor from the next list box – the descriptor type will be the label of this list box. For instance, if the chosen test takes disk partitions as descriptors, then the label will be **DiskDrive**. To include all the descriptors in a single graph of the chosen measure, select the **All** option from this list box.
- An appropriate **Display name** for the measure can then be provided.
- Then, add the specification by clicking on the **Add** button in Figure 12.11.

3. Alternatively, you can use the **Test** option against **View based on** to monitor multiple components on which a particular test executes. Once the **Test** option is selected, proceed as discussed hereunder:

   - If the measure to be plotted live pertains to a particular site, then select the site from the **Service** list box of Figure 4.99. This is an optional field, and will appear only if at least one service is configured in the environment.
   - If a **Service** is selected, then all the tests associated with the components under the chosen service will appear in the **Test** list box. Choose the **Test** for which a measure graph needs to be generated.
   - From the **Measure** list box that lists all the measures that are reported by the chosen **Test**, pick the measure for which a graph is to be generated.
   - Next, choose a component from the **Component** list, and then pick the descriptors for which the graph is to be generated from the **Selection** list. To include all the descriptors in a single graph of the chosen measure, select the **All** option from this list box.
   - An appropriate **Display name** for the specification can then be provided.
   - Then, add the specification by clicking on the **Add** button in Figure 12.10.

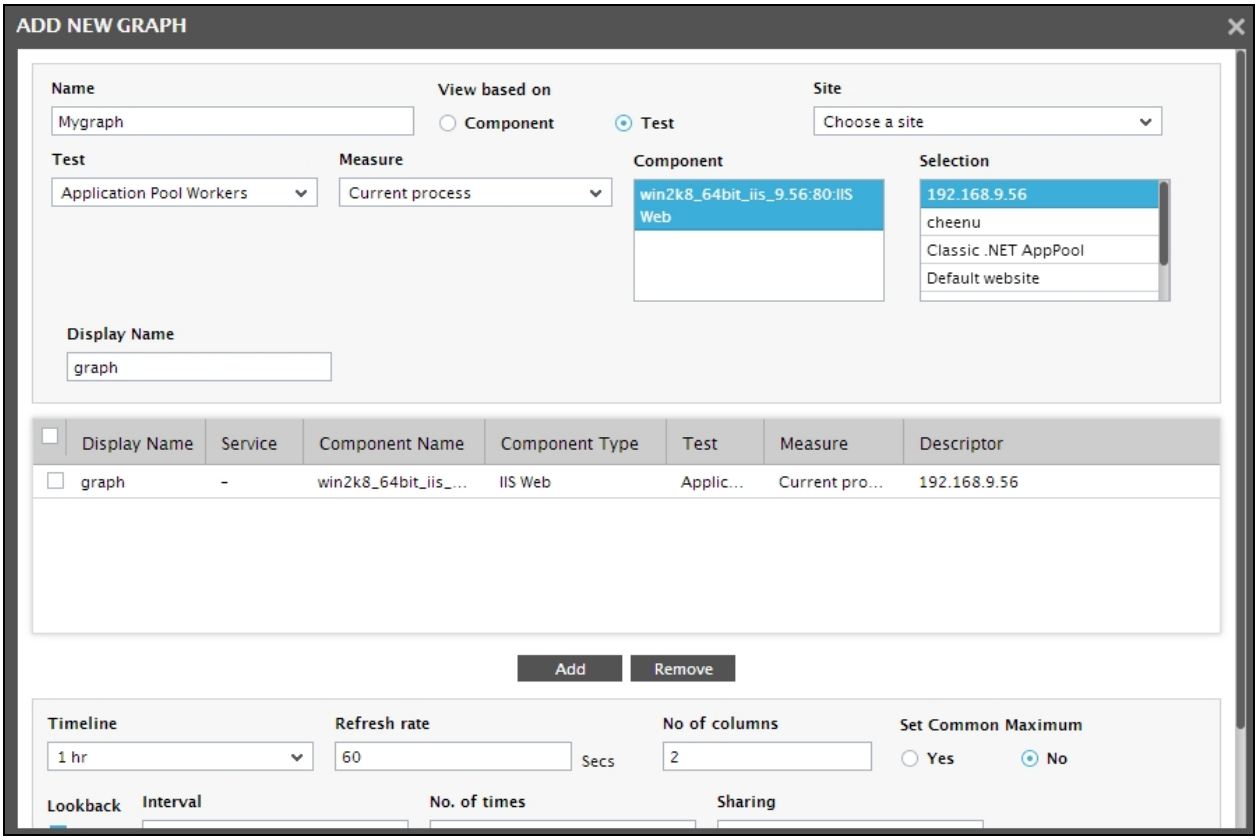

Figure 12.10: A Test-based view

4. Similarly multiple measures can be added (see Figure 12.10 and Figure 12.10).

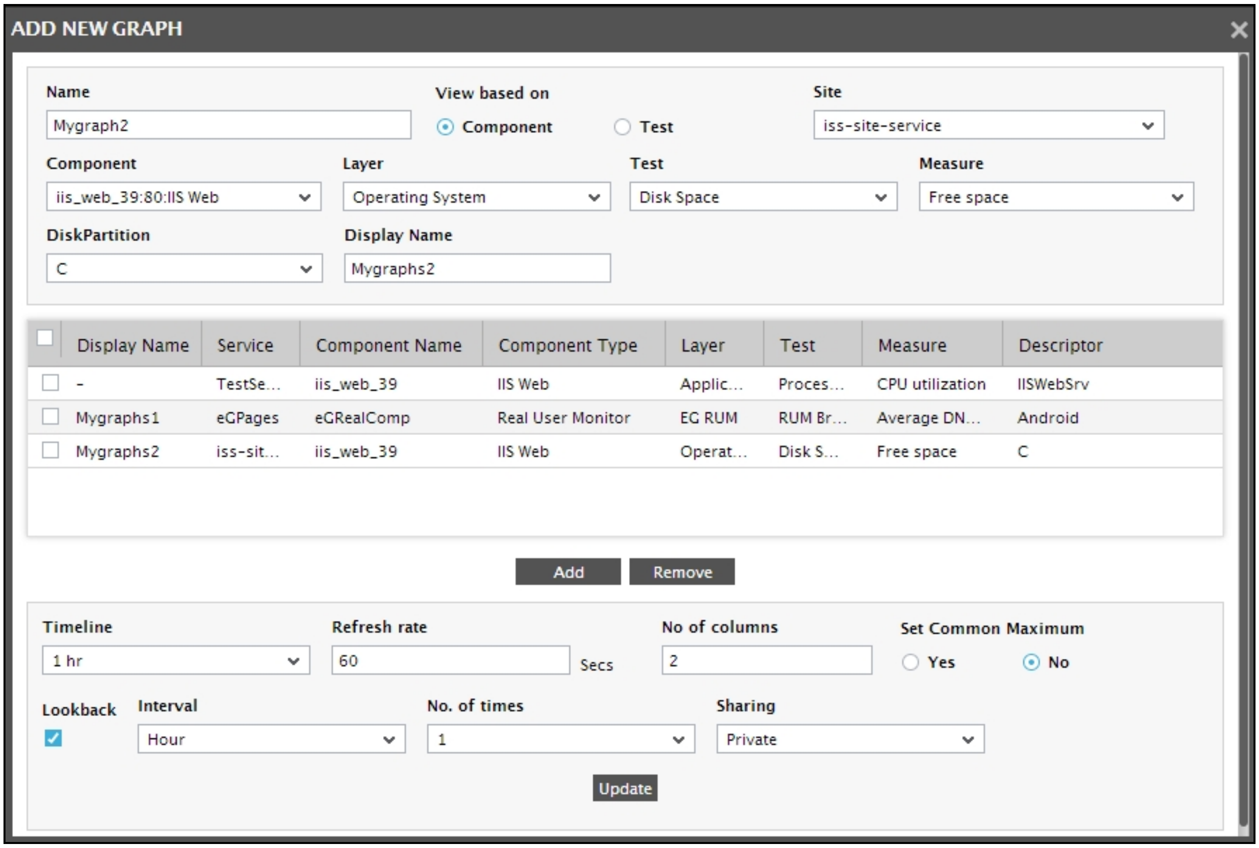

Figure 12.11: Multiple specifications in a Component-based view

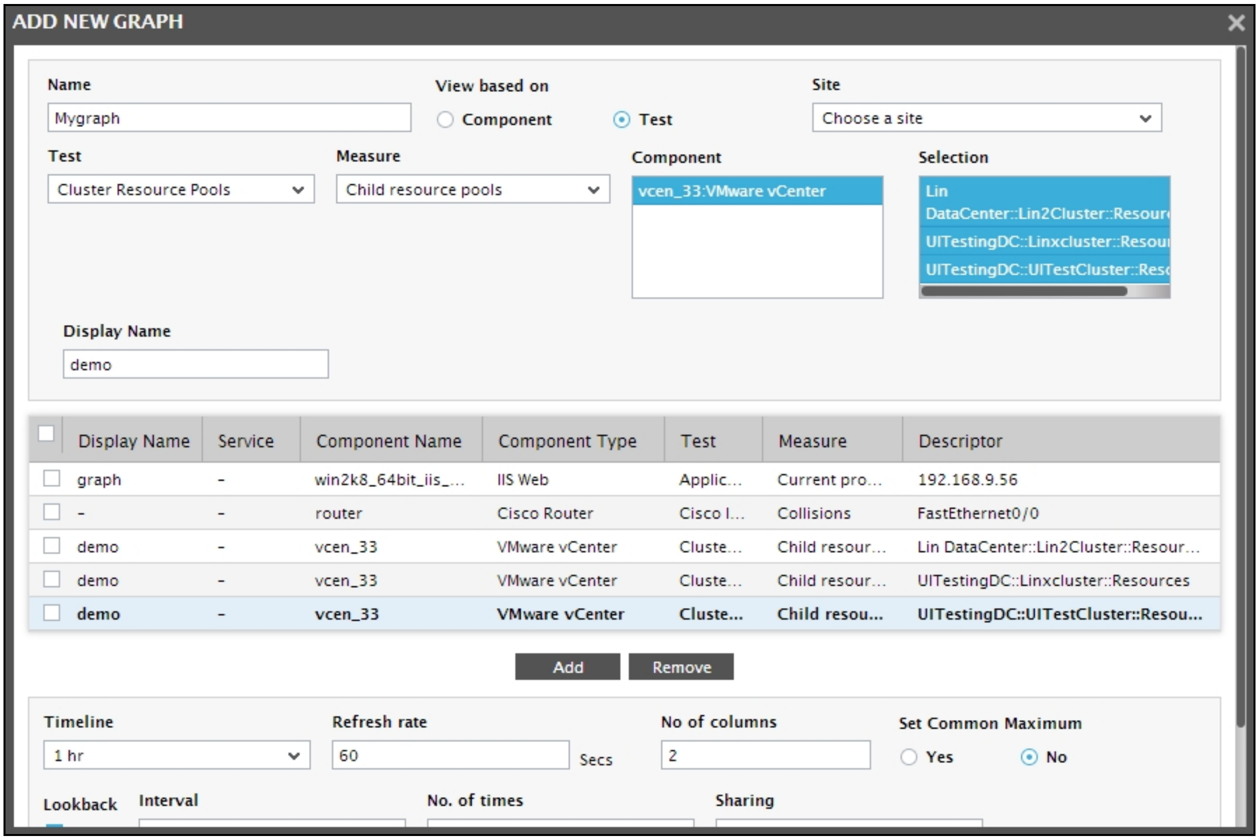

Figure 12.12: Multiple specifications in a Test-based view

5. To delete a measure, select the corresponding entry from the list box adjacent to the **Add** button and click the **Delete** button.

6. Next, select a **Timeline** for the graphs, and specify how frequently the measure values are to be refreshed in the **Refresh Rate** (see 12.1 or Figure 12.13) text box.

7. The **LIVE GRAPH DISPLAY** page will display a measure graph for every measure chosen in Figure 12.10 or Figure 12.9. These graphs will be displayed in a tabular format, characterized by rows and columns. Every graph will occupy a particular cell in the table. By default, the graph display will contain 2 rows and 2 columns.

8. Next, if you want the measure graph to plot the past values too, select the **Lookback** check box.

9. By default, the **Set Common Maximum** flag is set to **NO**. This indicates that by default, eG Enterprise does not set a common maximum value for all graphs in a view - in other words, the maximum value differs from graph to graph. This is because, the suite, by default, determines the maximum value for a graph by comparing all the values reported by its corresponding measure during the specified **Timeline**. If you want to set a common maximum value across all graphs in a view, set the **Set Common Maximum** flag to **YES**.

10. If the **Lookback** option is chosen, then proceed to select an **Interval** for the past values. For example, if the **Timeline** for the current measures is set to 1 hour, then the current measure graph will be plotted for the last one hour - say, for 12.00 PM to 1.00 PM of February 22, 2005. If the **Interval** is set to 1 day, then eG Enterprise will plot the values reported between the same hour (i.e., 12.00 PM and 1.00 PM), but for the previous day - i.e., February 21, 2005.

11. Optionally, you can pick the **Others** option from the **Interval** list and then choose a fixed timeline from the drop-down list that appears alongside (see Figure 12.15); for instance, you can select **Others** from the **Interval** list and then pick **3 hrs** from the drop-down list that appears alongside. This will result in a live graph in which the **past 3 hours of data** is plotted.

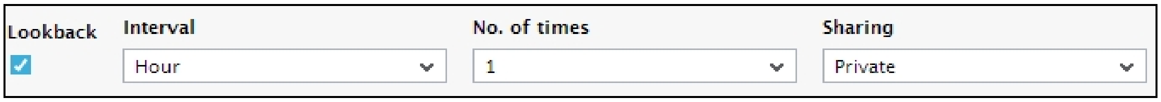

Figure 12.13: Choosing the 'Others' option from the 'Interval' list

To view meaningful data, it is recommended that the **Interval** set for the past values be greater than or equal to the **Timeline** chosen for the current values.

12. The number of past measurements to be plotted can be specified in the **No. of Times** text box. In the example above, if **2** is specified against the **No. of Times** text box, then the measure graph will plot 2 sets of past measurements. This includes:

    - The values for 12.00 PM to 1.00 PM on the February 21, 2005
    - The values for 12.00 PM to 1.00 PM on the February 20, 2005

    Note:

    If the **All** option is chosen from the list of descriptors, then the **Lookback** check box and its related options, such as the **Interval** and **No. of Times** fields, will become disabled.

13. To facilitate collaboration and knowledge sharing among administrators, eG Enterprise now allows users the flexibility to share useful live graph views they create with other users. For sharing a view, you need to first set a share type by selecting the desired option from the **Sharing** list box in 12.1. The options are as follows:

    - **Private**: If you set a view as **Private**, then all other users to the eG Enterprise system will be denied access to that view. Only the creator of the view will be able to access the view.
    - **Public**: If you set the view as **Public**, then only users with the following rights will have access to that custom view:
        - Users with access to all the managed components in the environment
        - Users with access to one/more components that are included in the view being shared

      Note:

        - If a view has been granted **Public** access, then users with the privilege to monitor only a few components in that view will be able to take a look at only those graphs in the view that pertain to the components in his/her monitoring scope.
        - Assume that a web site service, *www.abc.com*, has been configured, which is supported by an *IIS Web* server and an *Oracle Database* server. Say, a user *john* is created, who is allowed to monitor only the *Oracle Database* server that is involved in delivering www.abc.com. Assume that a live graph view is then designed by first selecting the **Site** www.abc.com, and then selecting key metrics from each of the components that are supporting the site. If this view is allowed **Public**

access, then user *john* will not be able to access the view at all. This is because, the live graph view that was shared was designed using the **Site** www.abc.com as the base. User *john* does not have access rights to the site www.abc.com. Therefore, eG Enterprise denies user *john* access to the site-based view.

On the other hand, if the live graph view was designed by directly selecting the *IIS Web* and *Oracle Database* components (involved in service delivery) from the **Component** list, then such a view can be accessed by user *john*. In this case however, note that user *john* will be able to view only the live graphs pertaining to the *Oracle Database* server assigned to him, and not the *IIS Web* server.

- If a view has been allowed **Public** access, then users whose monitoring scope does not include any of the components in that view will not be allowed access to the view.

- **Share:** On the other hand, if you choose the **Share** option from the **Sharing** list, then you can pick and choose the users who need to be granted access to that view. To map users to a view, do the following:

14. Select the users to be granted access from the **Available Users** list box.
15. Then, click on **Add selected (** **) icon** to grant the selected users access to the view. Doing so shifts the selection to the **Selected Users** list.
16. To revoke the access permission granted to a user, you can select the user from the **Selected Users** list, and click the **Remove from selected (** **)** icon.
17. Finally, click on **Update** to save the changes.

Note:

- A user without access to **Live Graphs** cannot view any of the views shared with him/her.
- A view can be modified / deleted only by the user who creates it. The shared users on the other hand, can only view the live graphs, and cannot modify/delete the view.
- If a user has been granted **Share** access to a view, but he/she has the right to monitor only a few components that are included in the view, then such a user will be able to take a look at only those graphs in the shared view that pertain to the components in his/her monitoring scope.
- Assume that a web site service, *www.abc.com*, has been configured, which is supported by an *IIS Web* server and an *Oracle Database* server. Say, a user *john* is created, who is allowed to monitor only the *Oracle Database* server that is involved in delivering www.abc.com. Assume that a live graph view is then designed by first selecting the **Site** www.abc.com, and then selecting key metrics from each of the components that are supporting the site. If this view is allowed **Share** access, then user *john* will not be able to access the view at all. This is because, the live graph view that was shared was designed using the **Site** www.abc.com as the base. User *john* does not have access rights to the site www.abc.com. Therefore, eG Enterprise denies user *john* access to the site-based view.

  On the other hand, if the live graph view was designed by directly selecting the *IIS Web* and *Oracle Database* components (involved in service delivery) from the **Component** list, then such a view can be accessed by user *john*. In this case however, note that user *john* will be able to view only the live graphs pertaining to the *Oracle Database* server assigned to him, and not the *IIS Web* server.

- If a user has been allowed **Share** access to a view, but he/she does not have the privilege to monitor any component included in the view, then such a view, despite being shared, will not be available to that user.

18. Finally, click on the **Update** button to register the changes.
19. Clicking on the **Update** button will open the **LIVE GRAPH DISPLAY** page that displays the configured measure graphs (see Figure 12.14).

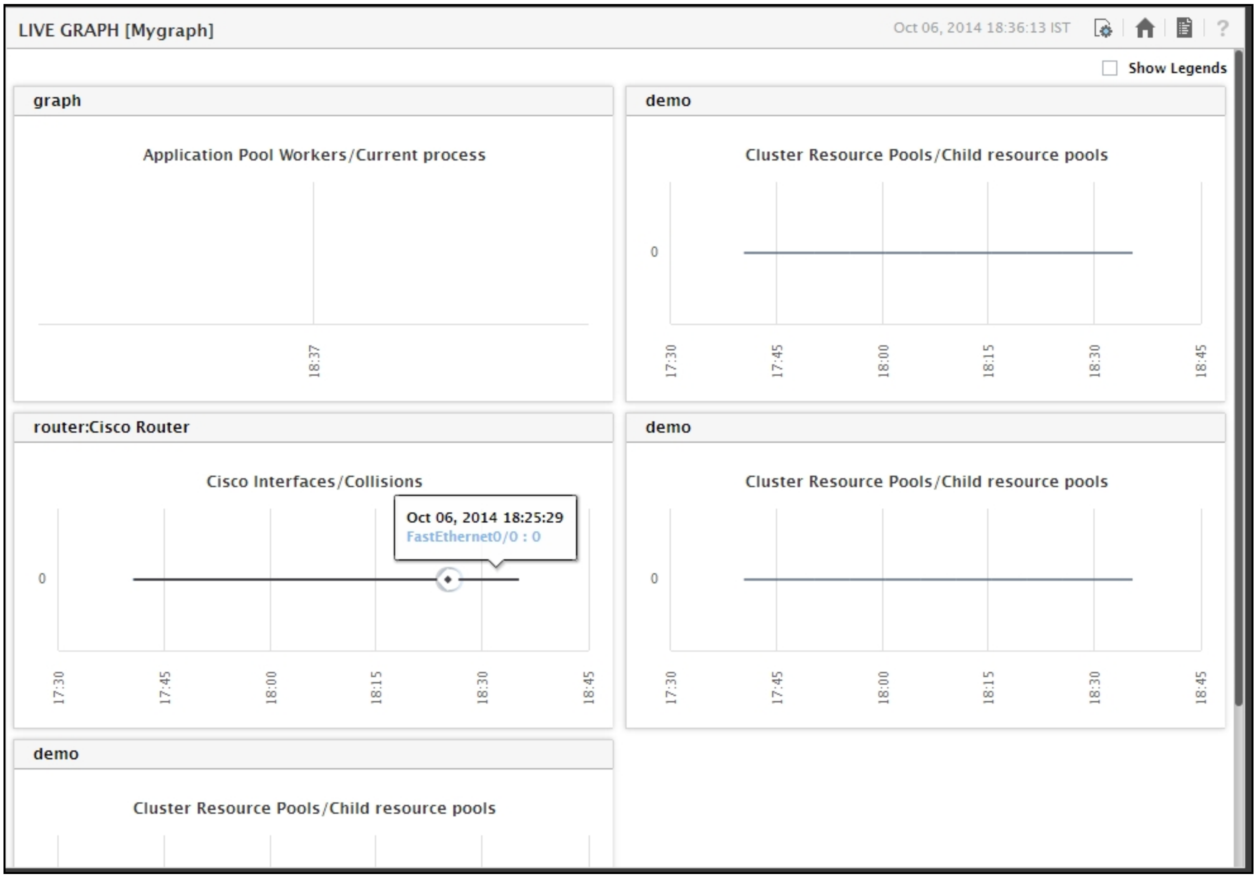

Figure 12.14: A Component-based Live Graph Display

20. If the values reported for a large number of descriptors are plotted in a single graph, the graph would naturally appear cluttered. To view and analyze such live graphs more clearly, just click on the corresponding graph title or on the graph itself. This will automatically enlarge the graph, thus enabling you to view it more clearly.
21. To make changes to the configuration, click on the **Views** button in the Live Graph menu at the right top corner of Figure 12.14.
22. Figure 12.14 will appear displaying the complete list of custom views (including the newly configured view).

Figure 12.15: The complete list of views

23. To add a new view, click the **Add New** button in Figure 12.15. To set any of the listed views as the default view, click the **Set As Default** ( ) corresponding to the view name in Figure 12.15. To delete a view, select the check box corresponding to the view name in Figure 12.15, and click the **Delete** ( )icon. Choose **Select All** and then click the **Delete** button, if all the listed views need to be deleted.

24. To make changes to the custom view configurations, click on the icon against a view in Figure 12.15. This will lead you back to Figure 12.15 where such configurations can be performed. However, in the modify mode, you cannot change the name of the custom view. Moreover, an additional **Delete** button will appear in Figure 12.15, which will allow you to delete the entire configuration.

25. To view the configured graphs in real-time, click the **View** option against the view name in Figure 12.15. This will take you to Figure 12.15 once again.

26. If other users have shared their views with you, then those views will be listed separately in the **Views Shared by other Users** section as depicted in the Figure 12.15.

27. Against every shared view, this section also displays the name of the view creator - i.e., the user who originally created the view and shared it with the current user. This user is represented by a user icon ( ).

28. To hide the unnecessary shared views select the views you want to hide by providing checkmark in the corresponding checkboxes.

29. Now click on **Hide Selected Views.** Doing so ensures that the chosen views are hidden as depicted in the Figure 12.15.

12.1.1 Customizing Live Graphs

The width, color, dimension, and background of the graphs displayed in the **LIVE GRAPH DISPLAY** page can be configured using the **eg_ui.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory. The **[LIVE_GRAPH]** section in the **eg_ui.ini** file lists the following parameters and their default values:

```
[LIVE_GRAPH]
```

```
LINEWIDTH=2
SAMPLECOLORS=#FF0000,#009900,#AAAAFF,#000000,#00ffff,#aa99aa,#ffee00,#DEC887,#ee9985,#ff5
599,
#6622aa,#960a90,#ee0aff,#3399ff,#88aaff,#0000FF,#0096af,#FA5A00,#7FFFD4
3DMODEON=FALSE
CHARTBACKGROUND=white
BACKGROUND=#d3d3d3
FOREGROUND=black
CHARTFOREGROUND=black
3DDEPTH=4
```

The **LINEWIDTH** parameter defines the width of the lines in a graph. By default, this is set to 2.

You can also configure the colors that each of the lines in a measure graph should take. A measure graph can plot a maximum of 8 sets of measurements, where 1 set corresponds to current measures, and the remaining 7 sets are historic data. Accordingly, a measure graph displayed in the **LIVE GRAPH DISPLAY** page can consist of a maximum of 8 line graphs. Using the **SAMPLECOLORS** parameter in the **eg_ui.ini** file, you can specify the color of each of the 8 line graphs displayed in the **LIVE GRAPH DISPLAY** page, as a comma-separated list. The first color specification in the comma-separated list will correspond to the graph of current measures, and this will be followed by the colors for each of the past measure graphs.

Setting the **3DMODEON** flag to **TRUE** will ensure that the graphs displayed are 3-dimensional. To view 2-dimensional graphs, set the **3DMODEON** flag to **FALSE**. If 3D graphs are enabled, then you can indicate the depth of these graphs using the **3DDEPTH** parameter. By default, this is set to 4.

To change the background of the graphs, set the **BACKGROUND** and **CHARTBACKGROUND** parameters to a color of your choice. Similarly, to change the foreground of the graphs, change the color specification against the **FOREGROUND** and the **CHARTFOREGROUND** parameters. By default, both will be set to *black*.

13

Role based Monitoring Views

13.1 Monitor Users

The previous sections explained how the supermonitor holds unrestricted access to monitor the entire environment. The interface presented by the system to the other monitor users is similar to the supermonitor's view. The key difference between the monitor and the supermonitor users is that the monitor users are restricted to the services, segments, zones, and independent components that have been associated with them. Moreover, using the virtual manager concept, the eG manager tracks alarms individually for each user, and can generate personalized alerts for each user. To demonstrate this difference, let's take the example of monitor users *john* and *elvis* who have been added via the eG administrative interface. Let us begin by exploring user *john's* view of the monitored infrastructure. Assume that the following infrastructure elements have been explicitly associated with user *john*:

Zone	east_coast
Service	buy.abc.com
Segment	None
Components	None

Figure 13.1: Table listing the infrastructure elements 'directly' assigned to user john

Figure 13.1 depicts the **Monitor Dashboard** for user john.

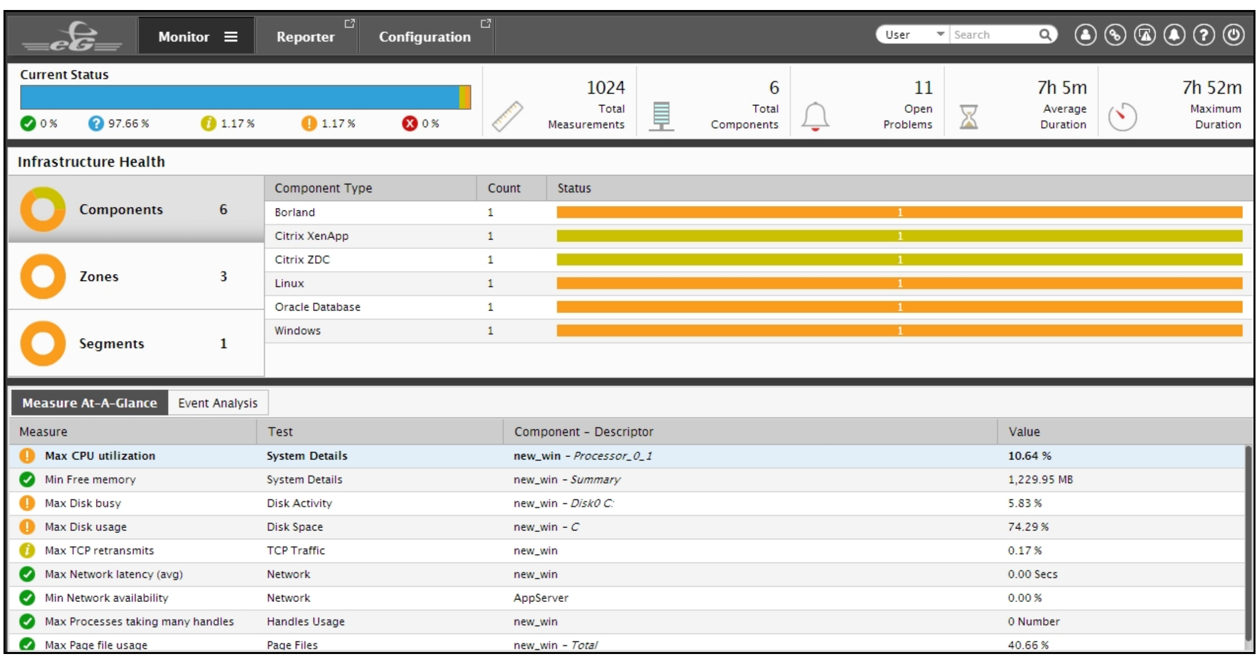

Figure 13.2: Monitor Dashboard of user *john*

The **Infrastructure Health** section of Figure 13.2 defines user *john's* monitoring scope. Let us first focus on the **Zones** bar in the **Infrastructure Health** section of Figure 13.2. According to the **Zones** bar, *john* has the right to monitor two zones, of which one is currently in a **Critical** state. According to the table in Figure 13.3 however, only one zone has been assigned to user *john* for monitoring. While it can be assumed that the *east_coast* zone is one of the zones, the key questions are what is the other zone, what is its current state, and why is there a mismatch between the zone assignments indicated by Figure 4.294 and the Monitor Dashboard of Figure 13.3. To find the answers, first, click on the **Zones** label in the **Infrastructure Health** section of Figure 13.2. Figure 13.3 then appears revealing the complete list of zones in *john's* monitoring purview.

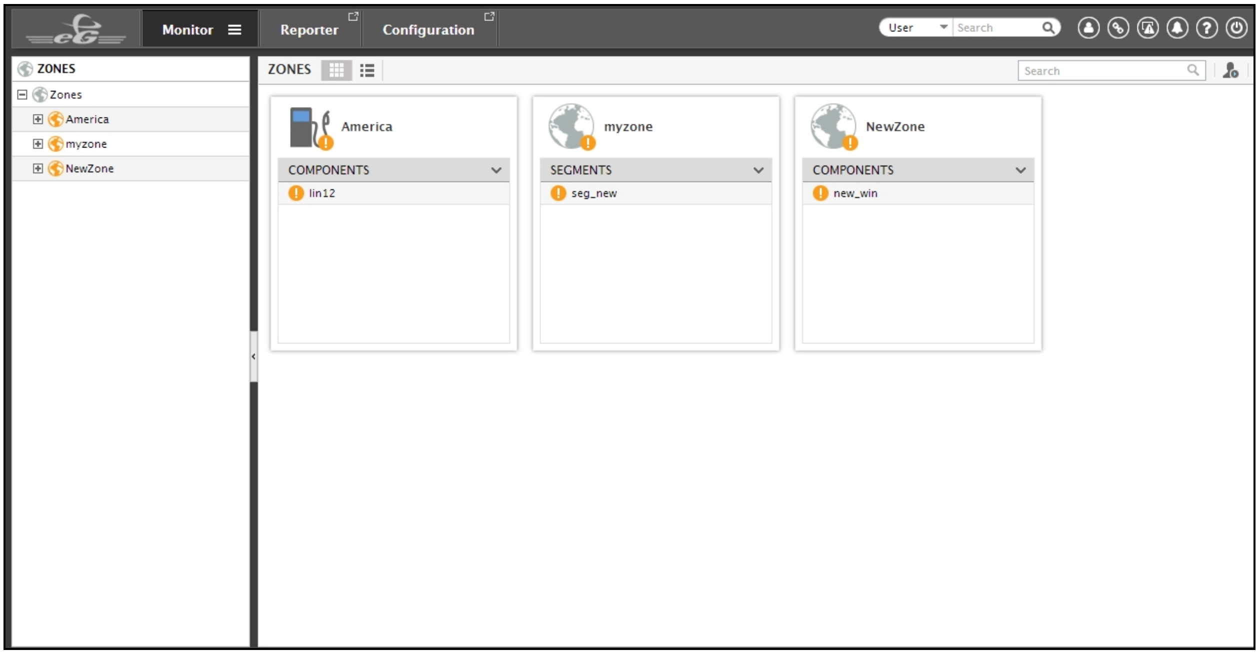

Figure 13.3: The zones monitored by user john

From Figure 13.3, it can be inferred that besides the *America* zone that is currently in a Major state, a *newzone* (which is in the Normal state) is also included in *john's* view. To know how the *newzone* got included in *john's* monitoring scope and to diagnose the root-cause of the Critical issues with the *east_coast* zone, click on *east_ coast* zone in Figure 13.3.

A **Zone Dashboard** then appears (see Figure 13.4) indicating the overall health of the *east_coast* zone. The **Infrastructure Health** section of the **Zone Dashboard** (see Figure 13.4) further reveals how many sub zones, segments, services, and components exist within the *america*, and the current state of these zone elements. From Figure 13.4, we can infer that the *America* zone consists of one subzone, one segment, and seven components.

You can even click on the respective bars in the **Infrastructure Health** section to know more about the problems (if any) affecting the performance of the corresponding zone elements.

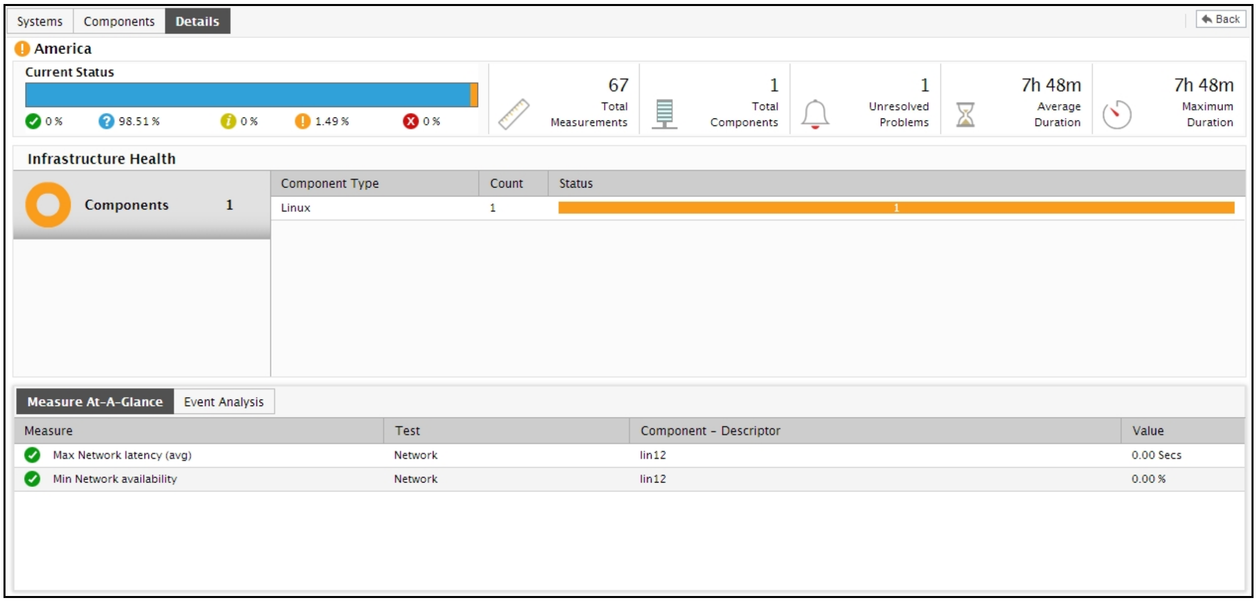

Figure 13.4: Dashboard view of the *America* zone assigned to user john

In the case of our example therefore, clicking on the **Sub Zones** bar will reveal that a *newzone* exists within the *America* zone and is currently in the **Normal** state (see Figure 13.4). This is the same zone that appeared in the **ZONE LIST** page of Figure 13.5, which displays the zones present in user *john's* view.

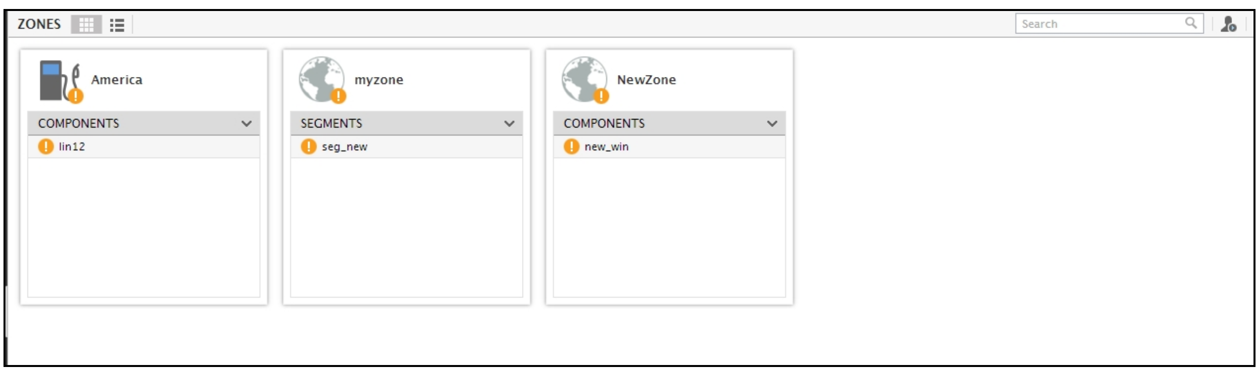

Figure 13.5: Health of the newyork_zone within the east_coast zone assigned to user john

We can thus conclude that the **Monitor Dashboard** of a user and the **ZONE LIST** page (of Figure 13.5) not only consider the zones directly assigned to that user, but also take into account the sub-zones within these zones. This explains why user *john's* **Monitor Dashboard** and **ZONE LIST** page included the sub-zone of *east_coast* zone, the *newyork_zone*.

Similarly, clicking on the **Segments** bar in Figure 13.6 will reveal that the *east_coast* zone consists of a segment named *seg-a* that is currently facing critical issues (see Figure 13.6).

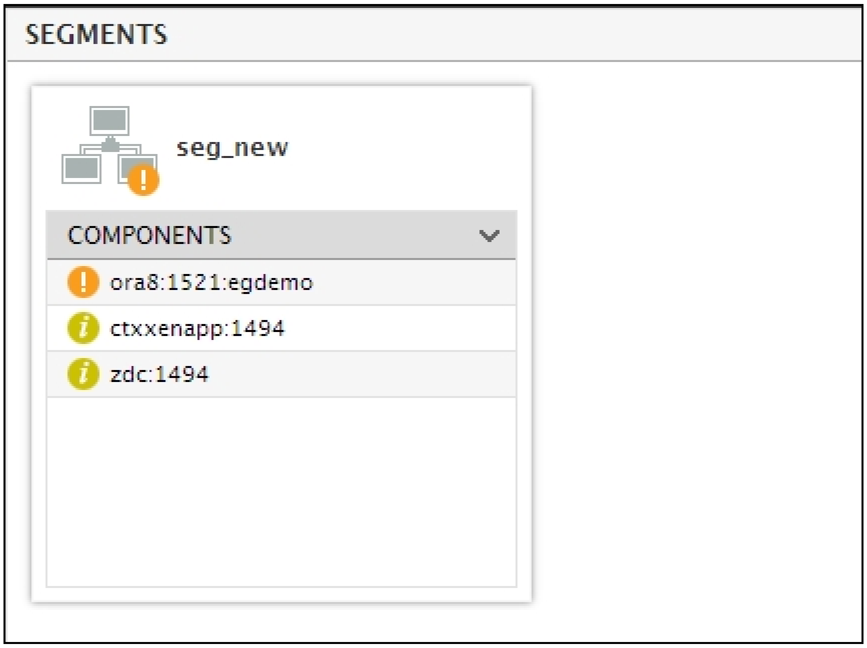

Figure 13.6: Health of seg-a within the east_coast zone assigned to user john

Let us now close the **Zone Dashboard** and return to the **Infrastructure Health** section of the **Monitor Dashboard** (see Figure 13.7) of user *john*. This time, let us focus on the **Services** bar. You can see that the service that has been directly assigned to user john is the cause for concern. According to Figure 13.6, a web site service named *buy.abc.com* is the problem service. Clicking on the **Services** bar invokes Figure 13.6, which also indicates that the *buy.abc.com* site is in a **Critical** state. You can further investigate the critical problem with the web site service, by clicking on *buy.abc.com* in Figure 13.6. This will reveal the **Transactions** tab page of the Service dashboard, where the names and state of the transactions mapped to the *buy.abc.com* web site will be listed. Clicking on a problem transaction there will lead you straight to the **Topology** tab page, using which you can accurately isolate the root-cause of the problem.

Figure 13.7: Services associated with user `john`

Once again, return to the **Monitor Dashboard** of Figure 13.7 This time, take a look at the **Segments** bar in the **Infrastructure Health** section of the dashboard. According to the **Segments** bar, user *john* is monitoring one segment, which is currently experiencing **Critical** issues. However, Figure 13.8 has a different story to tell - apparently, no segments were explicitly allocated to user *john* while his profile was created in the eG administrative interface. Similarly, while Figure 13.7 holds that no components were directly assigned to user *john*, the **Components** bar (in the **Infrastructure Health** section) and the **Components At-A-Glance** section (in Figure 13.9) indicate that 7 components are under john's monitoring purview. The question now is, where do these 7 components and the single segment come from? Let us proceed to figure this out.

In the eG Enterprise system, when a zone is directly associated with a user, all infrastructure elements included in that zone are automatically assigned to that user. This implies that while explicitly assigning the *east_coast* zone to *john*, all the infrastructure elements within the *east_coast* zone also got added to *john's* view automatically. We can thus infer that the **Segments** bar and the **Components** bar in the **Monitor Dashboard** of Figure 13.10 represent the state of the *seg-a* segment and the 7 components that are in the *east_coast* zone (see Figure 13.8 and Figure 13.9).

Let us now proceed to view the **CURRENT ALARMS** window of user john. Figure 13.9 depicts the alarm window corresponding to john. Only the alarms pertaining to the site buy.abc.com, which is being monitored by john, appear.

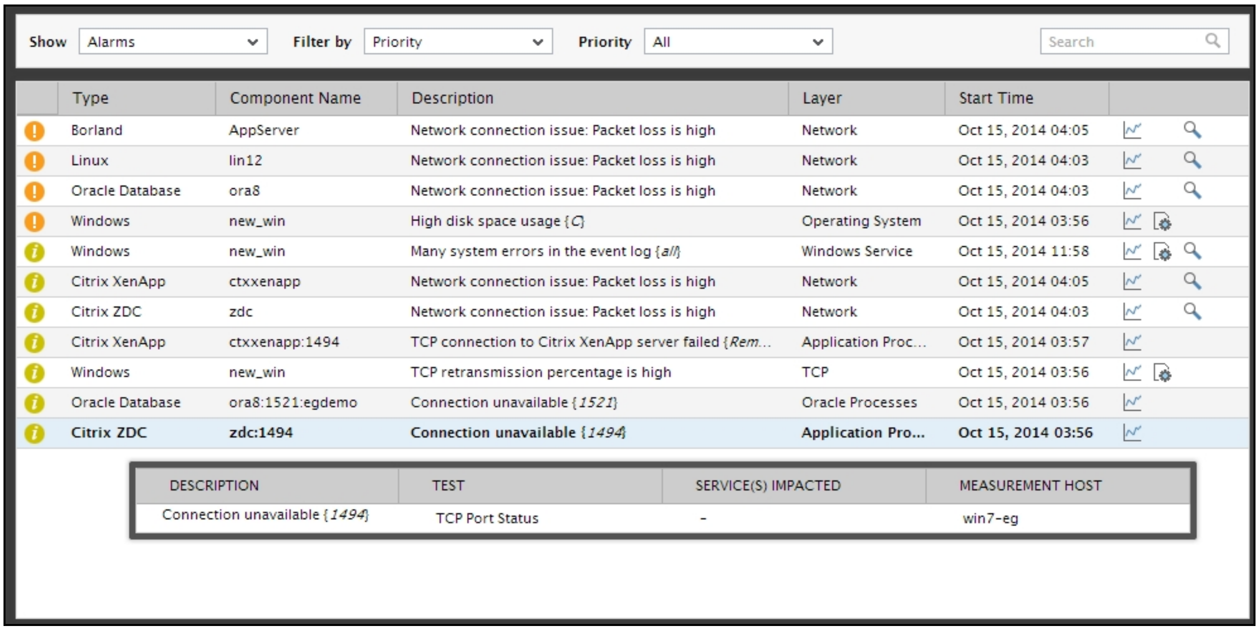

Type	Component Name	Description	Layer	Start Time
Borland	AppServer	Network connection issue: Packet loss is high	Network	Oct 15, 2014 04:05
Linux	lin12	Network connection issue: Packet loss is high	Network	Oct 15, 2014 04:03
Oracle Database	ora8	Network connection issue: Packet loss is high	Network	Oct 15, 2014 04:03
Windows	new_win	High disk space usage {*C*}	Operating System	Oct 15, 2014 03:56
Windows	new_win	Many system errors in the event log {*all*}	Windows Service	Oct 15, 2014 11:58
Citrix XenApp	ctxxenapp	Network connection issue: Packet loss is high	Network	Oct 15, 2014 04:05
Citrix ZDC	zdc	Network connection issue: Packet loss is high	Network	Oct 15, 2014 04:03
Citrix XenApp	ctxxenapp:1494	TCP connection to Citrix XenApp server failed {*Rem...*	Application Proc...	Oct 15, 2014 03:57
Windows	new_win	TCP retransmission percentage is high	TCP	Oct 15, 2014 03:56
Oracle Database	ora8:1521:egdemo	Connection unavailable {*1521*}	Oracle Processes	Oct 15, 2014 03:56
Citrix ZDC	**zdc:1494**	**Connection unavailable {*1494*}**	**Application Pro...**	**Oct 15, 2014 03:56**

DESCRIPTION	TEST	SERVICE(S) IMPACTED	MEASUREMENT HOST
Connection unavailable {*1494*}	TCP Port Status	-	win7-eg

Figure 13.8: Alarms specific to user `john`

If only 'critical' alarms pertaining to the buy.abc.com site assigned to user *john* had been configured for display in the monitor interface, then the **CURRENT ALARMS** window will appear as depicted by Figure 13.9.

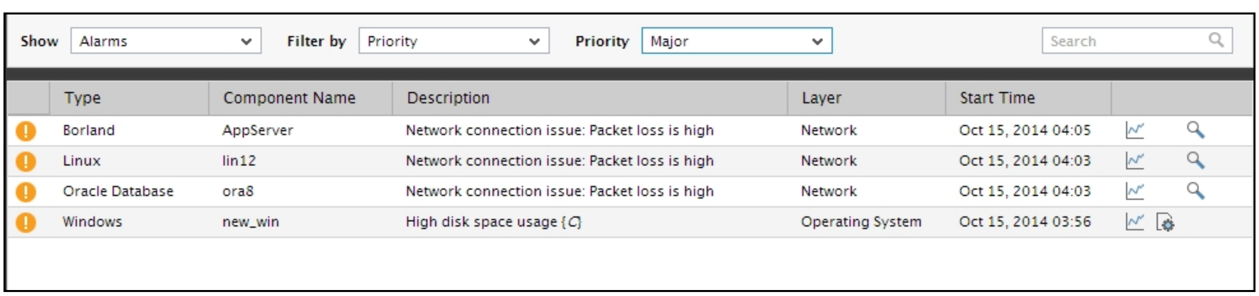

Type	Component Name	Description	Layer	Start Time
Borland	AppServer	Network connection issue: Packet loss is high	Network	Oct 15, 2014 04:05
Linux	lin12	Network connection issue: Packet loss is high	Network	Oct 15, 2014 04:03
Oracle Database	ora8	Network connection issue: Packet loss is high	Network	Oct 15, 2014 04:03
Windows	new_win	High disk space usage {*C*}	Operating System	Oct 15, 2014 03:56

Figure 13.9: The Alarms window displaying only the Major alarms of user john

Move your mouse over the first alarm and you will see the details of the alarm as depicted by Figure 13.10 below. The details include a brief description of the problem and the test that is reporting the problem. Note that the first alarm pertains to a NAS server, the layer that is problematic is the **Application Processes** layer, the test reporting the problem is **Processes** test, and the problem is with the **kjs** process, which is not running.

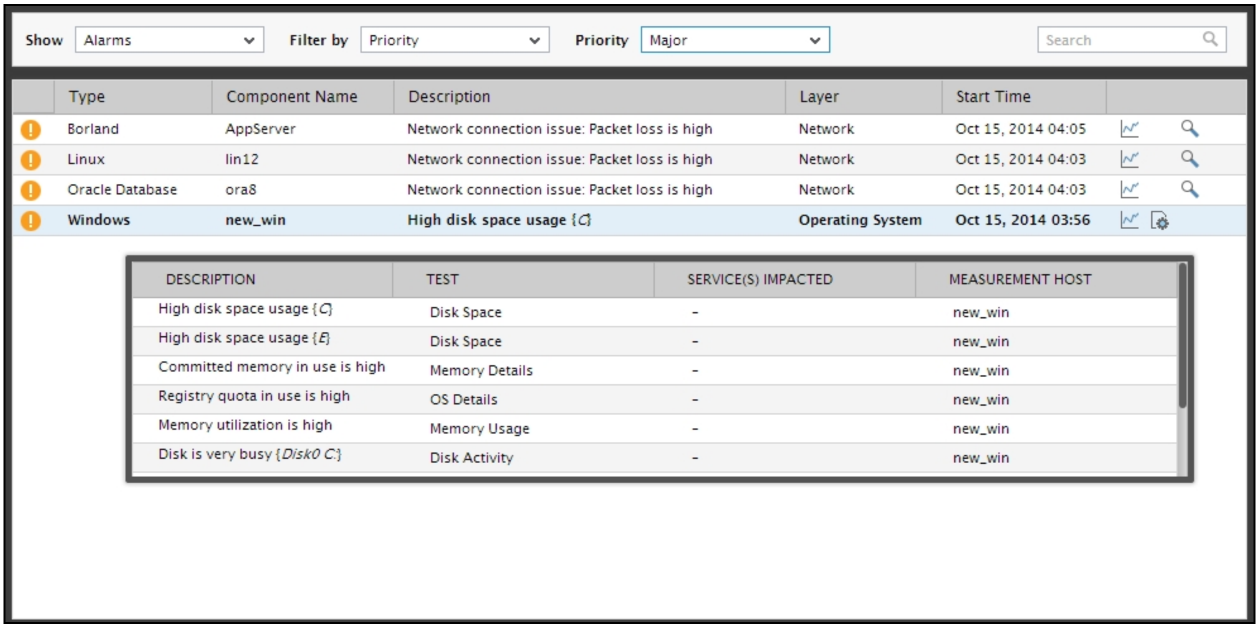

Figure 13.10: Description of the alarm

13.2 Role of an AlarmViewer

While the supermonitor users require an overall perspective of the health of a monitored environment, **AlarmViewers** are concerned with only the problems that are affecting the performance of the components associated with them. Therefore, when an AlarmViewer logs into the monitor interface, all he gets to view immediately is the **CURRENT ALARMS** window (see Figure 13.11), listing only those alarms that pertain to the components of interest to him/her.

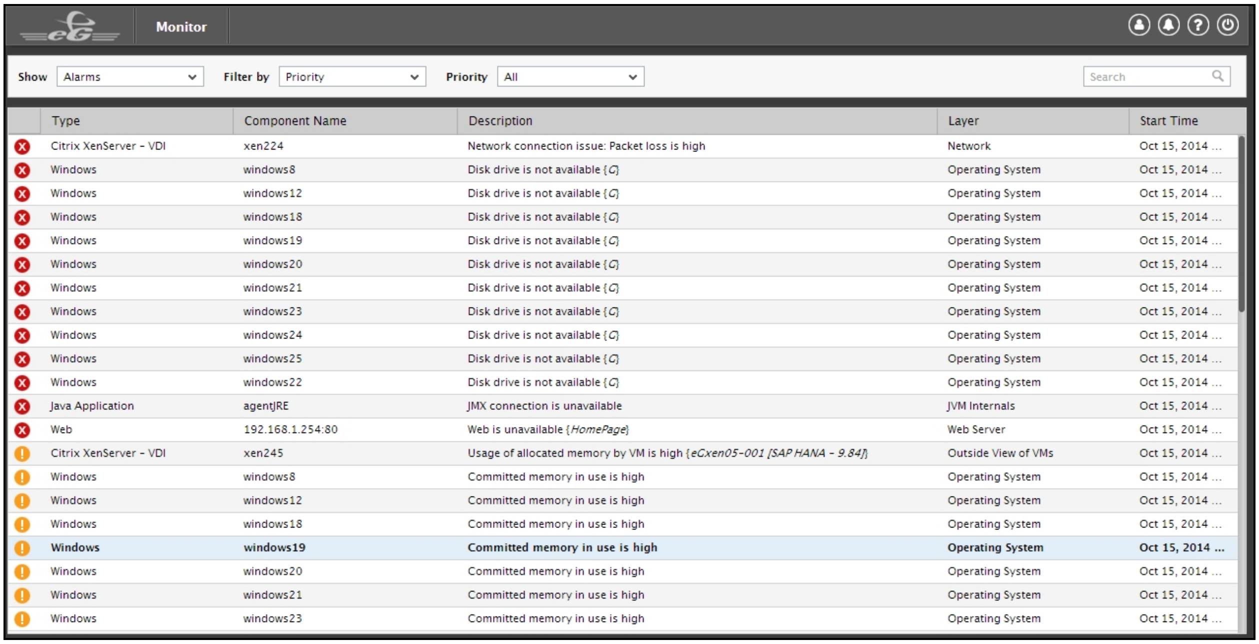

Figure 13.11: The CURRENT ALARMS window of an AlarmViewer

To know more about an issue, move your mouse pointer over the corresponding alarm in Figure 13.11. Figure 13.12 will appear displaying a brief description of the problem, the problem test, the problem service (if any), and the measurement host.

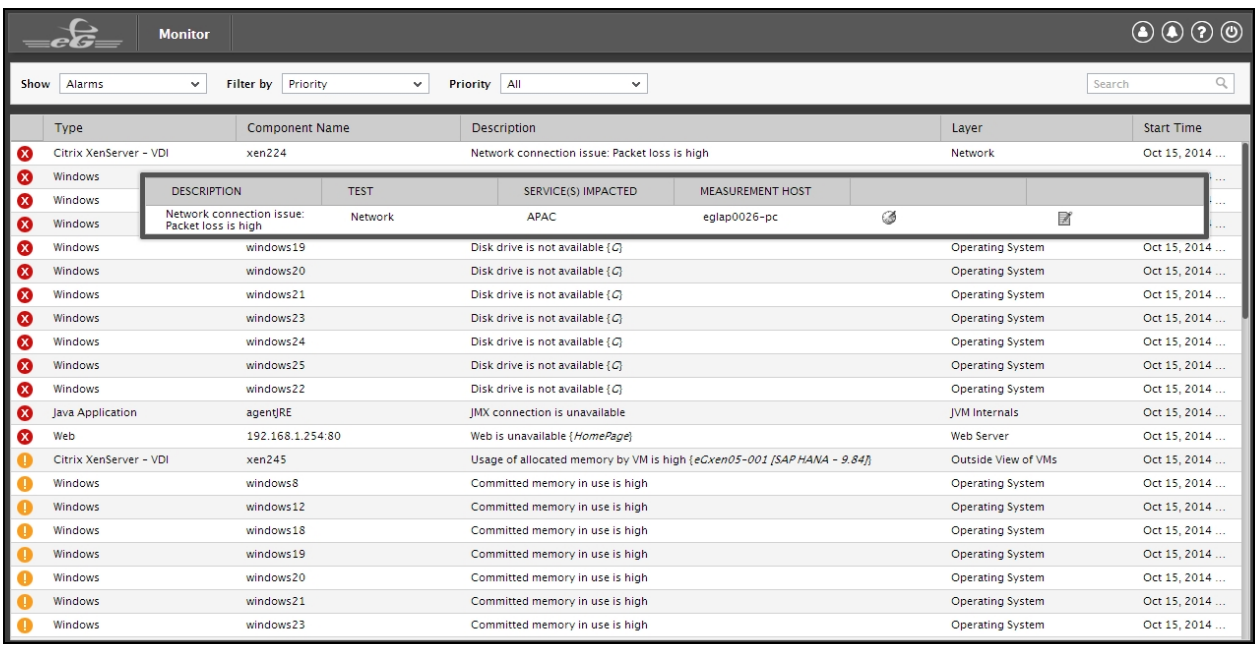

Figure 13.12: Alarm description

In addition to the above, Figure Figure 13.12 also reveals two icons. Clicking on the alarm and then on the icon will lead you to the **FEEDBACK** screen wherein information such as why the problem occurred, and when and how the problem was resolved can be recorded (see Figure 13.13, and a knowledge base of fixes can thus be maintained.

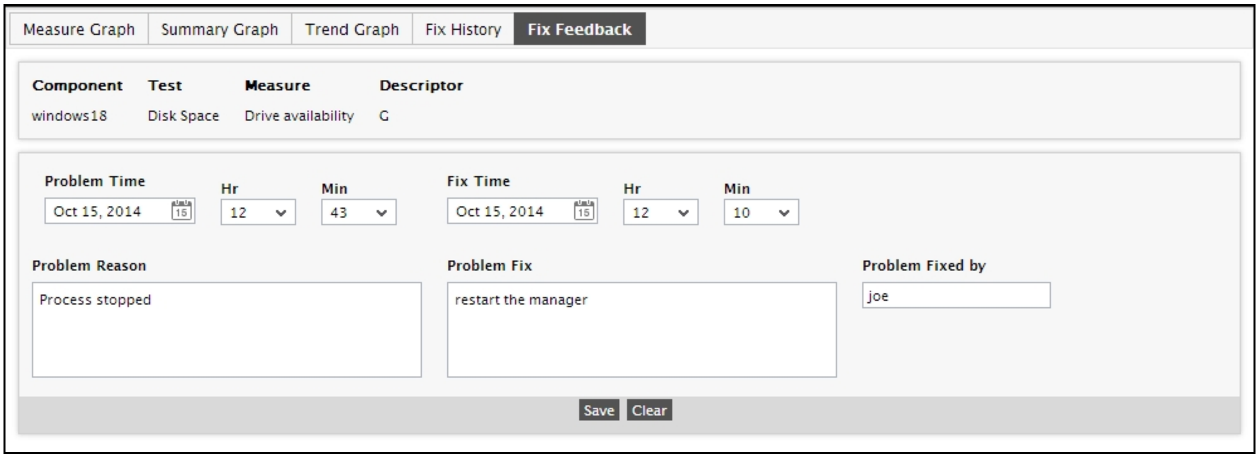

Figure 13.13: Recording how a problem was fixed

If the problem recurs, then clicking on the alarm and then on the  icon helps users to view a brief history of how the same problem was resolved in the past.

Besides receiving email and visual alerts of problems and recording feedback on fixes, the AlarmViewer also has the right to change his profile, if so required. To do this, the AlarmViewer will have to click on the **Profile** option at the right top corner of the monitor interface.

13.3 Role of a SuperAlarmViewer

While an *AlarmViewer* is alerted of problems pertaining to **Limited** components, the *SuperAlarmViewer* receives alerts of problems in the **Complete** monitored environment. Like an *AlarmViewer*, the *SuperAlarmViewer* too can record feedback on fixes, view the history of fixes, and change his/her profile.

14

Control Actions

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.

The benefits of eG's remote control actions are:

- Enable remote control in addition to routine monitoring, thereby offering a quick and easy way of initiating corrective actions
- Remote diagnosis and control of any component in the monitored infrastructure is enabled from anywhere, using just a web browser
- Remote control capability is selectively enabled for users based on their access rights

The control actions are enabled with no change in the eG architecture. The agents do not listen on any TCP ports. Hence, security risks in the target environment are minimum. Furthermore, since control actions can be initiated from a web browser, they can be triggered from anywhere, at any time.

To understand this concept better, consider the example of a web server shown in Figure 14.1.

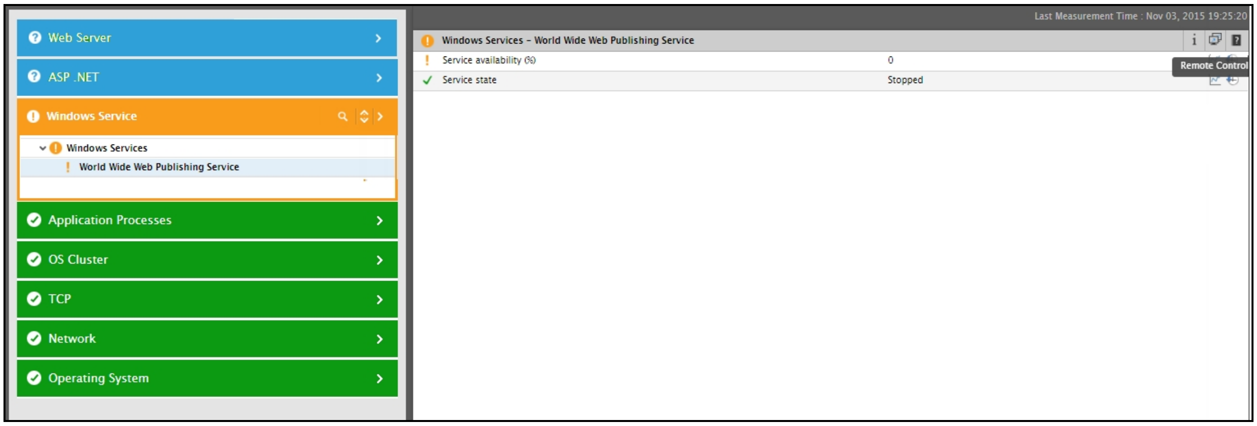

Figure 14.1: A web server indicating a problem in the Windows Service layer

Figure 14.1 indicates that the World Wide Web Publishing service of a web server has stopped running and hence, the web server is not available. Starting this service in the web server host will solve this problem. For this, you will not be required to manually access the host and start the WWW service in it.

Instead, simply click on the (**Remote Control**) button indicated by Figure 14.1. Upon clicking, a page depicted by Figure 14.2 will appear.

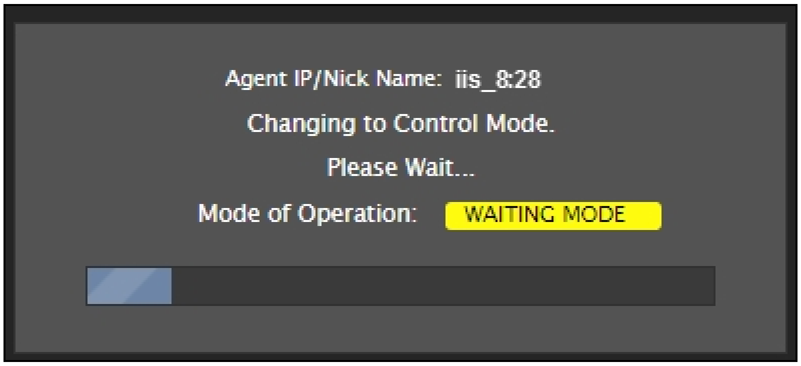

Figure 14.2: A page displaying the IP/nickname of the host to be controlled and current mode of the agent

This page will display the **Agent IP/Nickname** (192.168.10.173) to be controlled remotely and the current **Mode of Operation** of the agent. Figure 14.2 indicates that the **Mode Of Operation** of the agent on the host is **WAITING**. By default, an eG agent is in the **NORMAL** mode. No remote action can be performed on a host when the agent is in this mode. To start remote control activities, the agent has to be moved into a **CONTROL** mode. The waiting mode is a transition mode between the **NORMAL** and **CONTROL** modes. An agent remains in this state for a very short time, and finally changes to the **CONTROL** mode. Once the agent switches to the **CONTROL** mode, Figure 14.3 will appear, wherein you need to specify the command to be executed remotely on the agent host.

Figure 14.3 displays the **Agent IP/Nickname** of the host and its current **Mode Of Operation**. To execute a command, specify the following:

1. From the **Command** list box, select the command to be executed remotely. By default, the **Command** list contains a set of pre-configured, ready-to-use commands. If the command to be executed is already available in the **Command** list, select it, and specify the **Arguments** (if any), that may be supported by the chosen command (see Figure 14.3). In our case however, the command for starting the WWW service does not pre-exist. We are hence faced with the following options:

   a. First, add the command for starting the WWW service to the **Command** list (using the Agents -> Settings -> Remote Control menu sequence in the eG administrative interface), and then select that command from the **Command** list, OR,

   b. Directly issue the command for starting the WWW service using the **REMOTE CONTROL** page.

2. If you opt for (a), then follow the procedure discussed in this manual.

3. If you pick option (b), you need to first check whether you are authorized to execute a remote command that is outside of the **Command** list. By default, only the *admin* user has the right to directly execute a command on the agent host. To grant such a right to any other user(s), follow the steps given below:

- Edit the **eg_controls.ini** file in the **<EG_INSTALL_DIR>\manager\config** directory
- In the **[CONTROL_DEFAULTS]** section of the file, you will find an *AllowOtherCommands* parameter that is set to *admin*, by default. This indicates that, by default, only the *admin* user is privileged to execute commands other than the ones listed in the **Commands** list.
- To extend that right to other users, specify a comma-separated list of user names against the *AllowOtherCommands* parameter:
  AllowOtherCommands=admin,john,elvis,brian
- Finally, save the **eg_controls.ini** file.

4. Once the right is granted to you, you can proceed to directly execute an 'unregistered' remote command by selecting the **Other** option from the **Command** list, and providing the **Command description** (i.e., the syntax of the command) (see Figure 14.4).

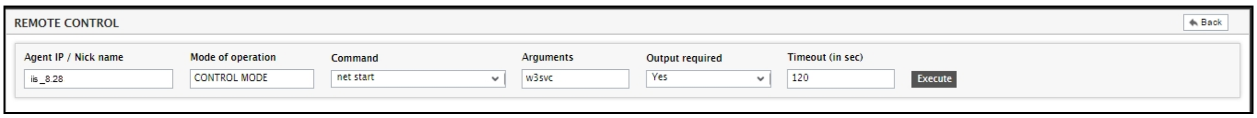

Figure 14.3: Directly issuing the command to be executed

5. In our example, the command which will start the WWW service is **net start W3SVC**, where **W3SVC** is the service name of the World Wide Web Publishing service. Therefore, against **Command description**, specify this command.
6. Then, set the value of **Output required** to **Yes** to view the output of the specified command. Otherwise, select **No**.
7. Then, enter the **Timeout** period. This is the duration for which the eG manager will wait for an output for the specified command. Upon the expiry of the period, the command execution will be automatically terminated.
8. Finally, execute the command by clicking the **Execute** button in Figure 14.4.
9. If **Output required** is set to **Yes**, then you will view the output of the executed command (see Figure 14.4).

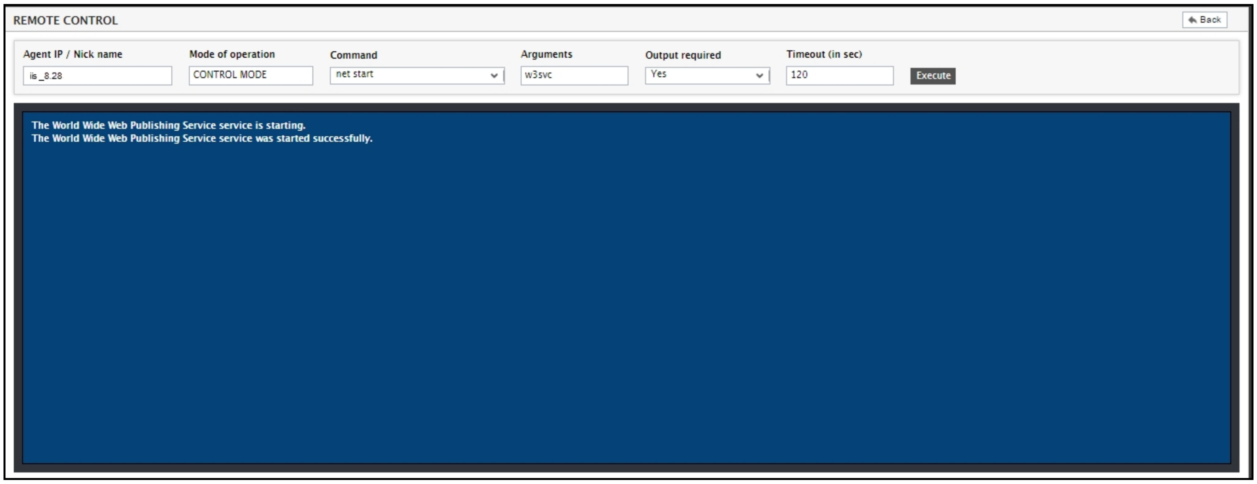

Figure 14.4: Output of the specified command

10. Once the command executes successfully, the WWW service will be started, automatically resolving the problem reported by the **Windows Service** test (see Figure 14.5).

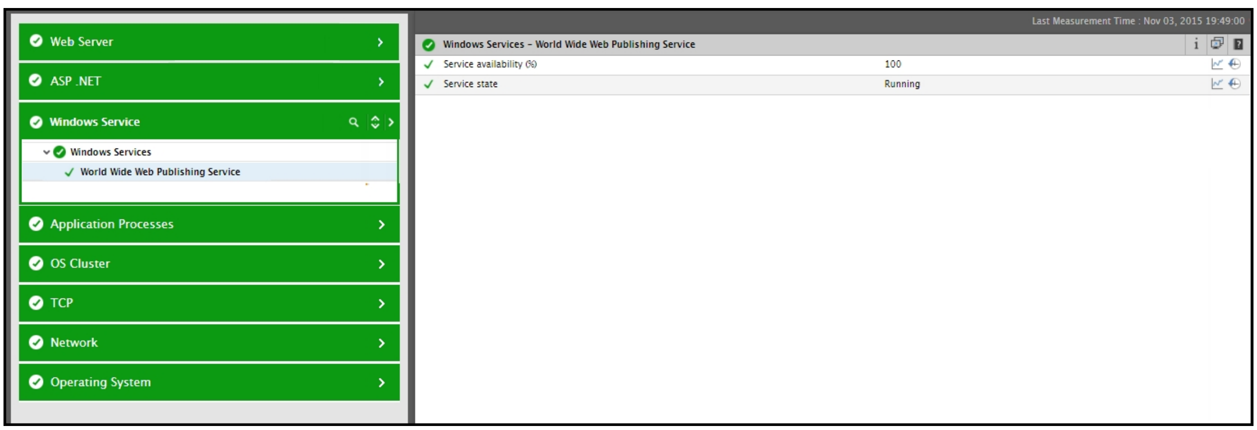

Figure 14.5: The Windows Service test problem is resolved

Note:

An agent will remain in the Control mode for 10 minutes, after which, it will return to the Normal mode.

This simple example highlights the utility of eG Enterprise's control actions capability.

The eG monitor interface also allows users to view the different eG agents and their modes of operation. To do so, select the **Remote Control** option from the **Options** menu of the eG monitor interface.

The resulting page (see Figure 14.6) lists the agents in the **Control** mode.

Figure 14.6: eG agents in the Control mode

Note that the agent on 192.168.10.173 in our example appears here. Clicking on the displayed agent will lead users to the page depicted by Figure 14.7, using which specific actions can be performed on the host. To search for a particular agent host, specify the whole/part of the IP/nickname of the host in the **Agent IP / Nickname** text box, and then, click the right arrow button (see Figure 14.6) next to it. All agents with IPs/nick names that embed the specified search string will then appear in this page.

Now, if you select **Normal** from the **Agents mode** text box in Figure 14.6, then all the agents in the Normal mode will be displayed (see Figure 14.6).

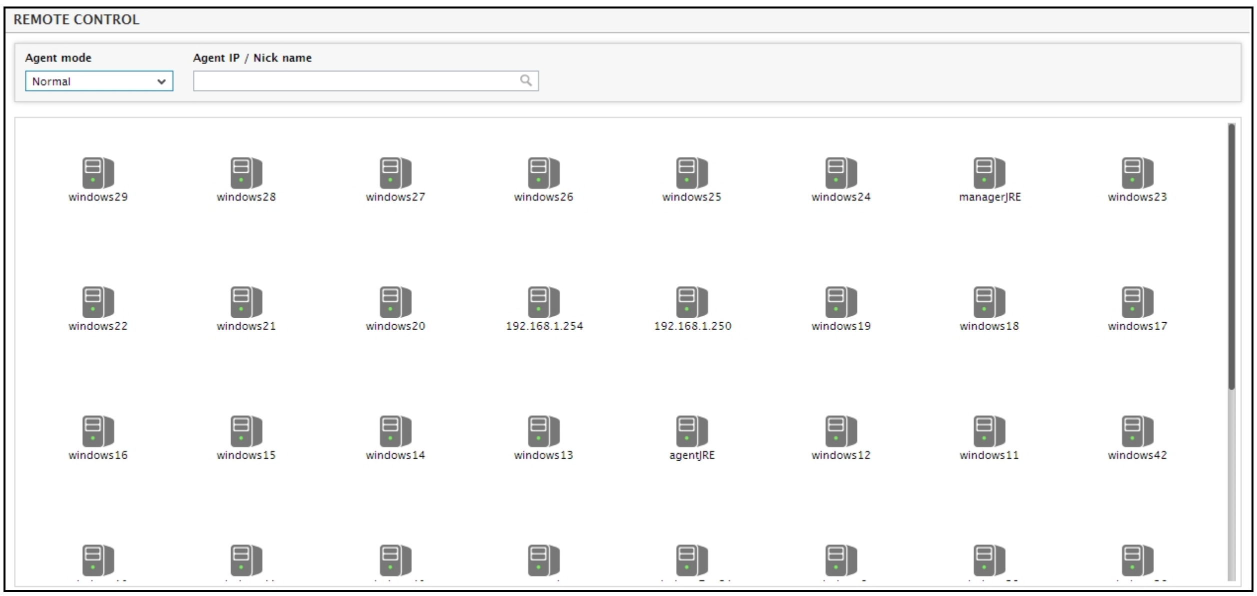

Figure 14.7: Agents in the Normal mode

Clicking on an agent displayed in this page (see Figure 14.7) will first lead users to the **WAITING** mode (see Figure 14.2); once the transition is successful, then the **CONTROL** mode page (see Figure 14.7) appears. Similarly, if you select **Waiting** from the **Agents mode** list box, then all the agents in the **WAITING** mode will appear. Clicking on the agent displayed here will lead users to Figure 14.7 directly.

14

Miscellaneous

14.1 User View

While managing a large number of users to the eG Enterprise system, administrators often find it very difficult to ascertain the overall health of a user's infrastructure, and whether there are any critical unresolved issues in the user environment. eG Enterprise provides a **USER VIEW** page that displays the number of segments/services/service groups/zones/components associated with a chosen user, and the alarms associated with these infrastructure elements. To access this page, follow the menu sequence: *Miscellaneous -> User View*.

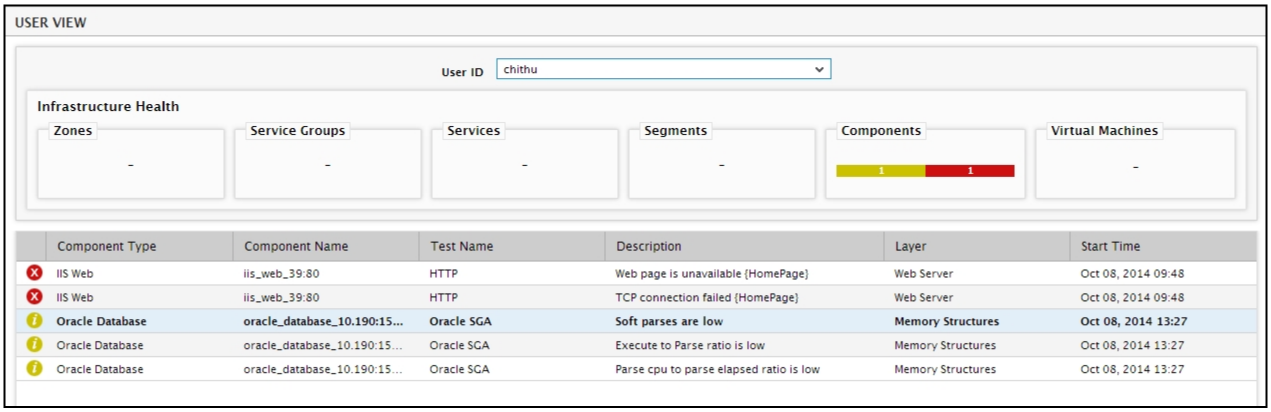

Figure 14.8: A selected user's view of the monitored infrastructure

Select a user from the **User Id** list and click the **Submit** button to view the associations of the selected user. Upon clicking the **Submit** button, an **Infrastructure health** section will appear that provides an overview of the selected user's environment.

This section displays bar graphs that indicate the number of zones/segments/services/service groups/components associated with the chosen user, and the current state of these elements. Clicking on a division in any of the bars listed in this section will lead the administrator to the corresponding list page. For instance, if you click on the *Critical* (i.e., RED) division in the **Zones** bar, then the **ZONE LIST** page will appear listing only those zones associated with the chosen user that are currently in a **CRITICAL** state.

Below the **Infrastructure health** section you will find the **Events** section that lists the current alarms pertaining to the components associated with the new user. This provides the administrator with an idea of the number and type of open issues in the user environment. Clicking on an alarm leads you to the layer model, tests, and measurements of the problem component.

Note:

The **User View** option is available only for users with the Admin or Supermonitor rights.

14.2 Detailed Diagnosis

In order to view the detailed measures for a particular component, the supermonitor will have to select the **Detailed Diagnosis** option from the **Options** menu of the eG monitor interface.

Note:

Users will be able to view the detailed measures only if the eG license enables the detailed diagnosis capability of the eG Enterprise system.

In the page that follows, specify the following information (see 14.2):

- **Service:** Select the service for which detailed measures are required. This is an optional field.
- **Component:** Select the component for which detailed measures need to be viewed. If there are too many components in the list to choose from, you can narrow your search further by using the **Search** text box. Specify the whole/part of the component name/type to search for in this text box, and click the right-arrow button next to it. The **Component** list will then be populated with all component names and/or component types that embed the specified search string. Select the component of your choice from this list.
- **Layer:** Choose the specific layer to which the measure requiring a detailed diagnosis is associated.
- **Test:** From this drop-down list box, pick the test that generates the detailed measures.
- **Measure:** From this drop-down list box, select the measurement for which a detailed diagnosis is required. The measures listed in this list box are the ones for which the corresponding test provides a detailed diagnosis.
- **Timeline:** The period for which the detailed diagnosis has to be analyzed can be specified using the **Timeline** list. You can either provide a fixed time line such as 1 hour, 2 days, etc., or select the **Any** option from the list to provide a **From** and **To** date/time for graph generation.

Finally, click the **Submit** button. Upon clicking, the detailed diagnosis of the specified measure for the given period, will appear (see 14.2):

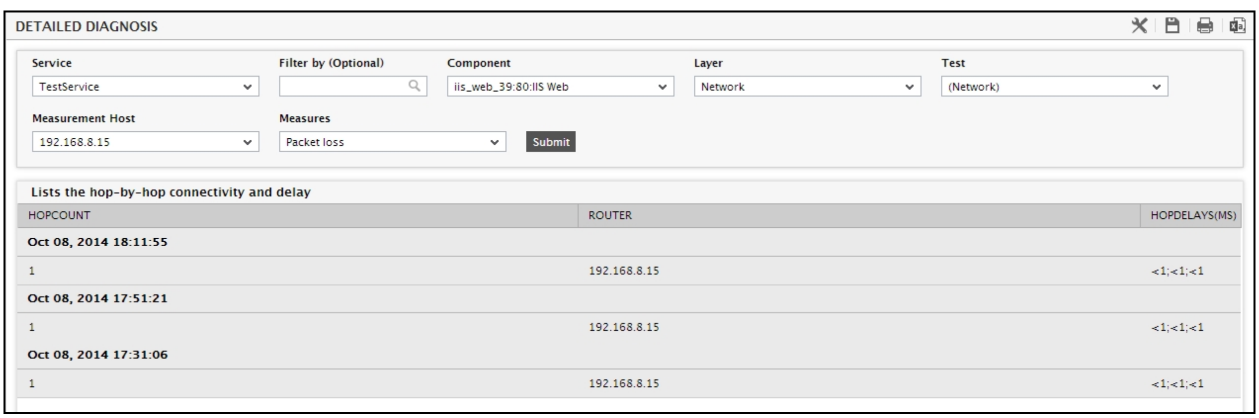

Figure 14.9: Detailed diagnosis of the *Cpu_util* measure for a given period

You can also save the **Detailed Diagnosis** information as a CSV file by clicking on the button at the top, right corner of 14.2, and even print the detailed diagnosis by clicking on the button.

14.3 Knowledge Base Search

The **Feedback History** page allows you to view the details of fixes that are available for a chosen component, test, measure, and descriptor only. Administrators however, may require single-click access to the fixes related to all the managed components in the environment. They may also demand the flexibility to query the knowledge base for viewing only that problem/fix information that interests them. Therefore, to enable administrators to perform quick infrastructure-wide searches on the knowledge ase and view the fix feedback of interest to them, the eG monitoring console offers the **KNOWLEDGE BASE SEARCH** page.

To access this page, follow the *Options -> Knowledge Base Search* menu sequence. 14.3 will then appear:

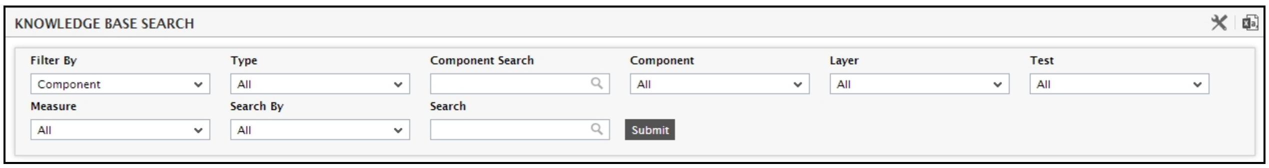

Figure 14.10: The Knowledge Base Search page

To search the knowledge base using this page, do the following:

1. First, select a search criterion from the **Search By** list. The options are as follows:
   - **Component**: Pick this option to view the details of fixes related to a single component or all managed components in the environment.
   - **Segment**: Pick this option to view the fixes related to one/all components that are part of a chosen segment. Once this option is chosen, a **Segment** list appears from which you will have to select the segment to be searched. Upon choosing a segment, the **Component** list will be populated with all the components that are part of the chosen segment.
   - **Service**: Select this option to view the fixes related to one/all components that are engaged in delivering a particular service. Choosing this option invokes a **Service** list from which you need to select the service of interest to you. Once a service is chosen, the **Component** list will be populated with all the components that support the selected service.
   - **Zone**: Select this option to view the fixes related to one/all components that are part of a specific zone. If this option is selected, then a **Zone** list will appear, from which you will have to pick the zone of interest to you. Doing so will populate the **Component** list with those components that are included in the chosen zone and its subzones (by default). If you want to exclude the components in the subzones from the purview of your search, then, set the **Include Subzones** flag to **No**. This will ensure that the **Component** list contains only those components that are associated with the parent zone.
2. From the **Component** list, select the component for which the fix details are to be viewed. Select **All** to view the fixes related to all components across the environment, or those that belong to a chosen service/segment/zone (as the case may be).
3. Select the **Layer** to which the tests of interest to you are mapped. Pick **All** to view the details of fixes, regardless of layer.

4. The **Test** list will then be populated according to the **Component** and **Layer** selections. To view fixes associated with all tests, select the **All** option.
5. Pick the **Measure** for which you want to view the fix history. Select **All** if the fix information related to all measures is to be retrieved.
6. Using this page, you can even search for those problems and/or fixes that embed a specified text string. For that, specify the string to search for in the **Search** text box, and indicate where the string should be searched - in the problem reason? the fix description? or in both. For that, select the **Problem**, **Fix**, or **All** options, respectively.
7. Then, provide a **Timeline** for your search.
8. Finally, click the **Submit** button.
9. The details of fixes that fulfill the specified filter criteria will then appear, as shown in Figure 14.11.

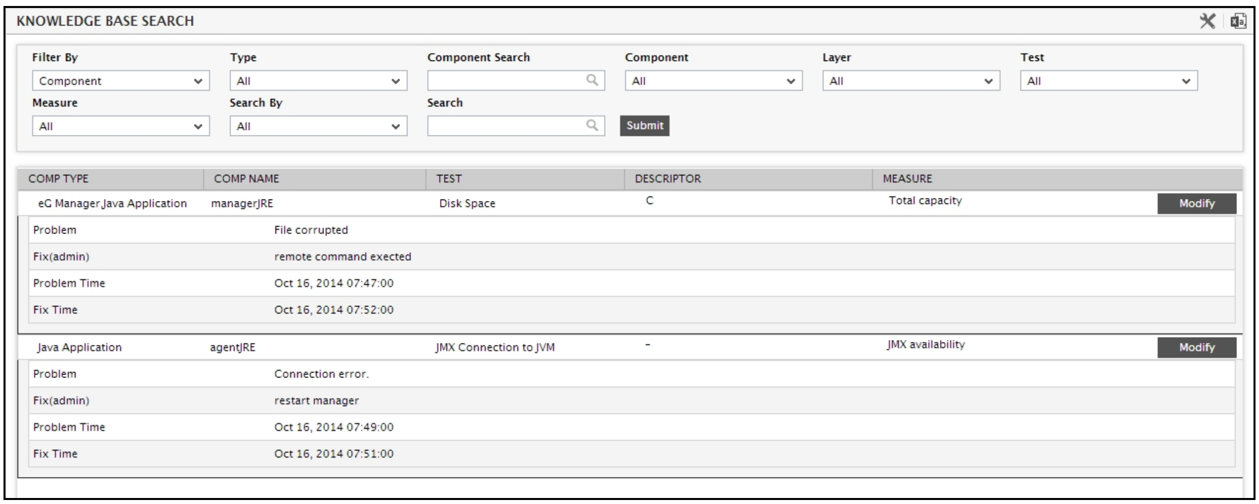

Figure 14.11: The details of fixes that match the specified filter criteria

10. You can even modify the problem reason and fix by clicking on the **Modify** button. Figure 14.11 will then appear using which you can make the necessary changes.

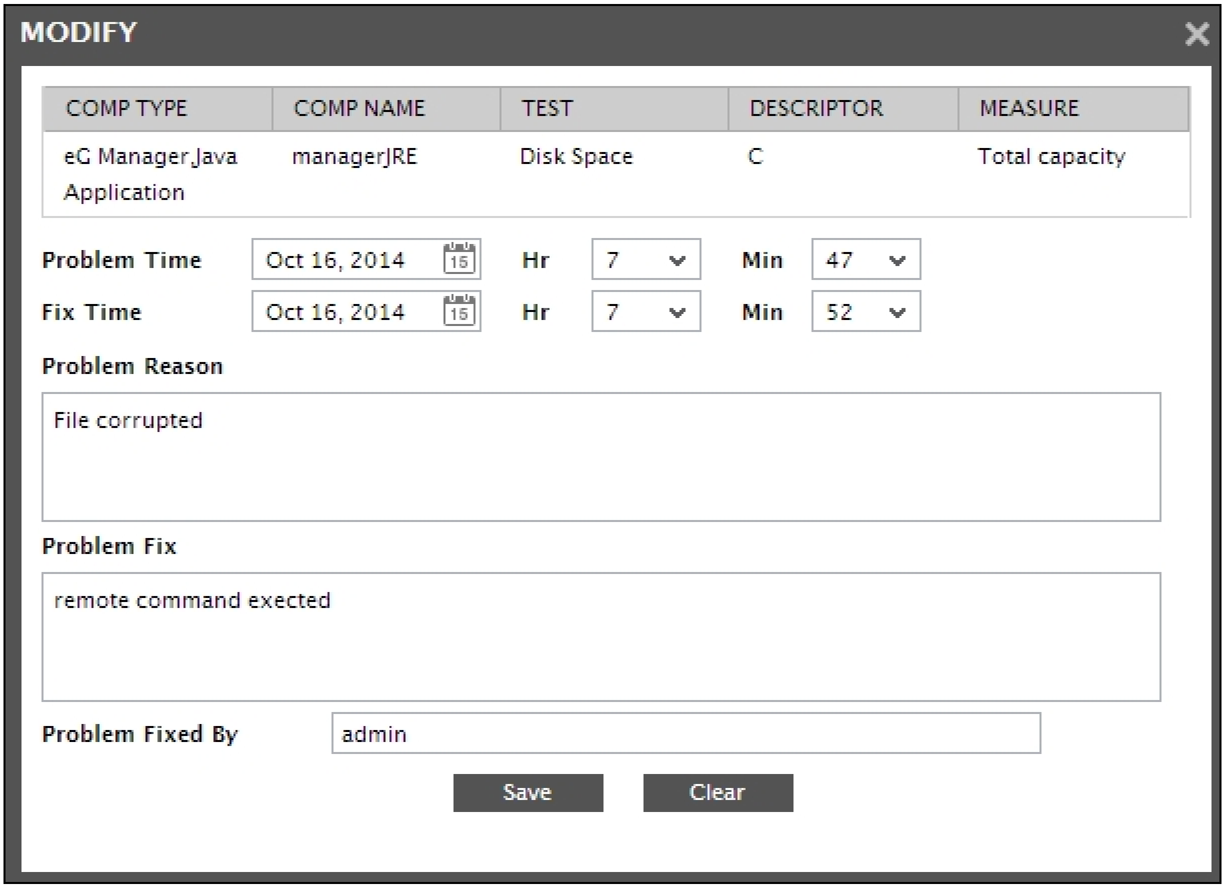

Figure 14.12: Modifying the problem reason and fix

15

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:

- A Virtual, Private Monitoring Solution for Multi-Domain IT Infrastructures
- The eG Installation Guide
- The eG Measurements Manual
- The eG Quick Reference Guide
- The eG Customization Manual

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.