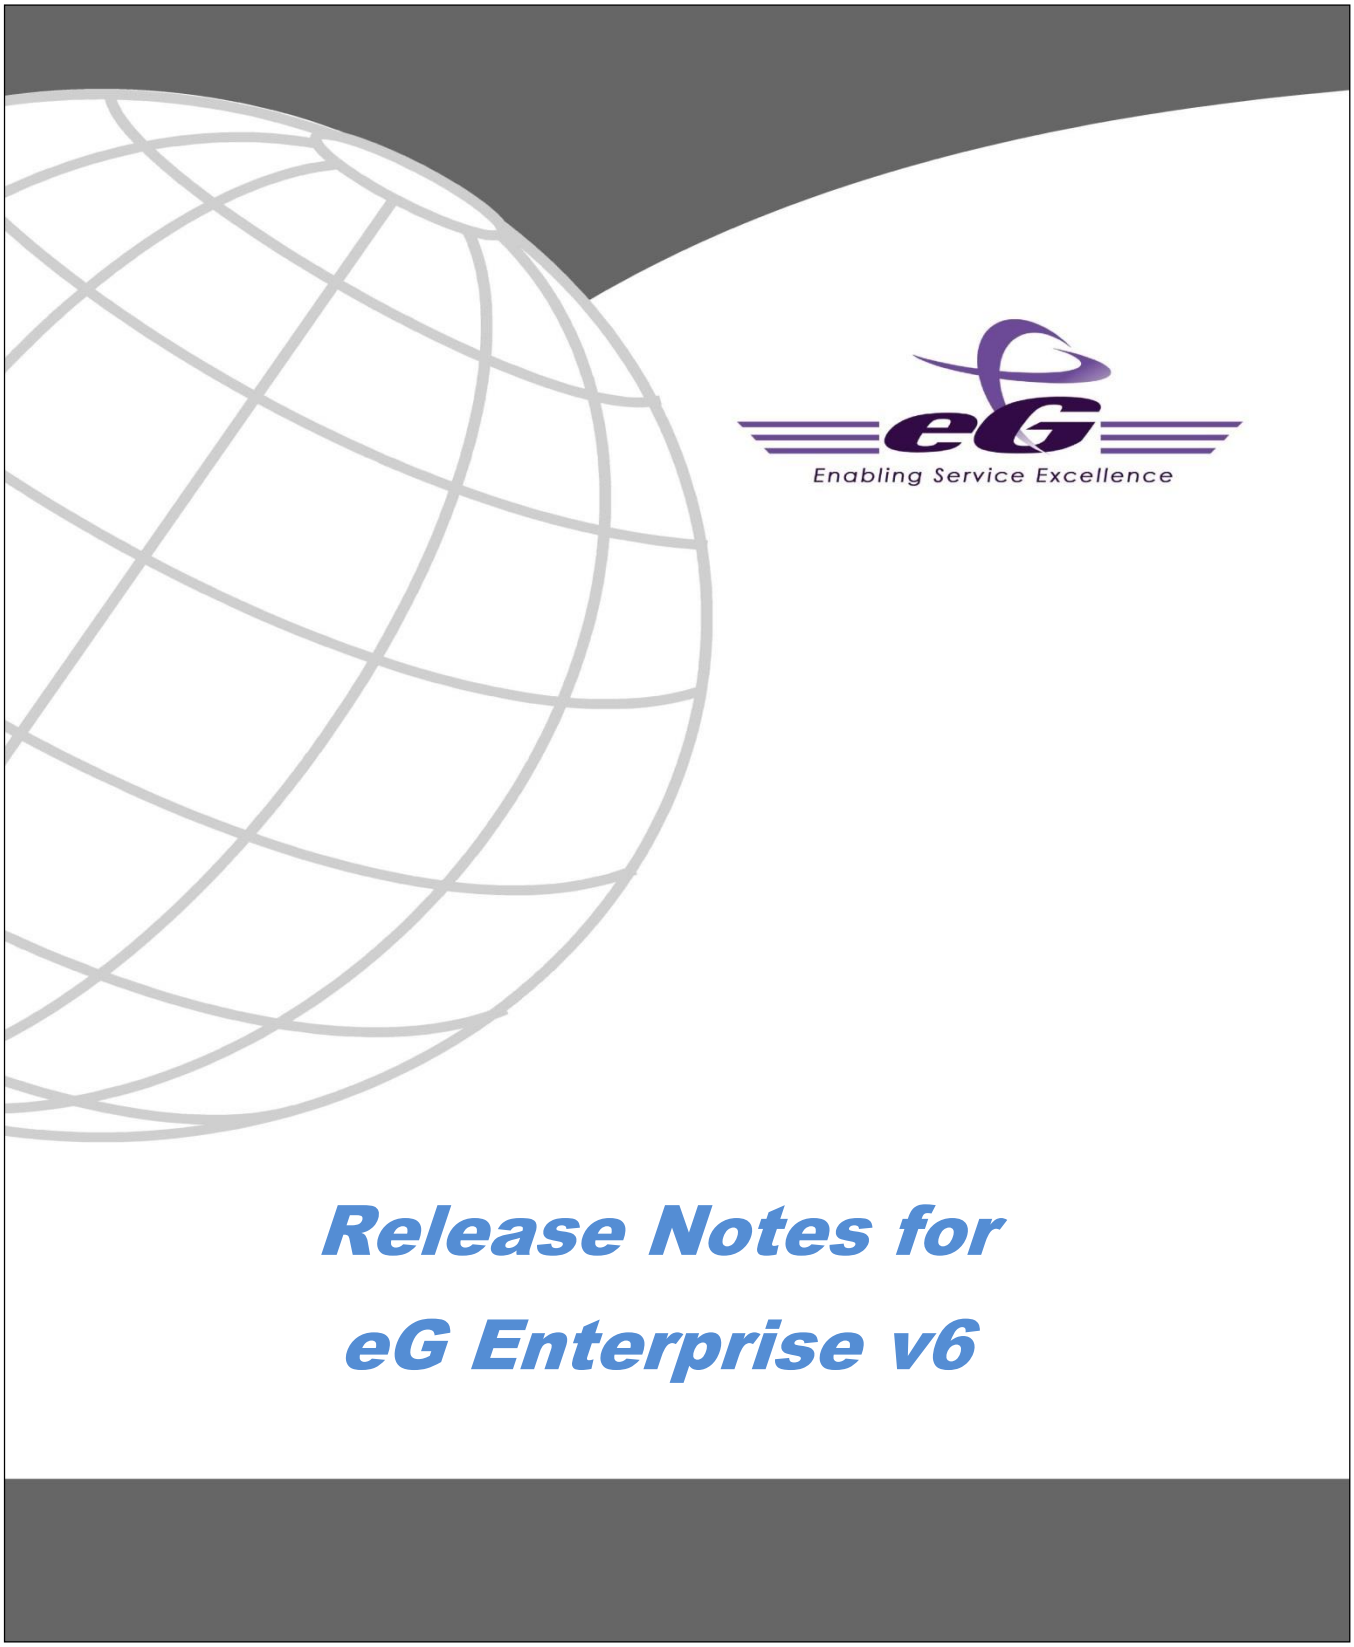

Release Notes for eG Enterprise v6

TABLE OF CONTENTS

Release Notes: eG Enterprise v6

Version 6 is a major release of the eG Enterprise performance monitoring, diagnosis and reporting solution for business-critical IT infrastructures. This document provides a comprehensive list of new features, enhancements and bug fixes that are part of this release.

1.1 eG Manager Changes/Enhancements

1.1.1 Admin Interface

- **Design changes:** In v6, the usability of the eG administrative interface has been enhanced, so that it is easier to navigate and operate the user interface. The key features of this new interface are as follows:
    - The **Admin** menu has been replaced in v6 with a sophisticated tile-view. Menu options are logically grouped, making browsing and selection a breeze. **Note that all modules of the user interface now use the tiled menu.**
    - Unlike the previous versions, where agent and manager discovery settings had to be defined using two different web pages, v6 of the eG administrative interface provides a single, central **DISCOVERY** tree that allows administrators to discover any component using any methodology they choose. This saves administrators the trouble of shuttling between multiple interfaces to perform discovery.
    - Version 6 includes a new and improved administrative interface for configuring specific thresholds for each application, server or device being managed. Administrators can now clearly understand what type of thresholds currently apply to which test – i.e., which tests are governed by global, group, default, and specific thresholds. Previously, administrators had to access multiple web pages to obtain this information.
    - Thresholds can be set for each measure separately using intuitive controls. These easy-to-use configurations allow little room for errors and reduce the training required to use the monitoring tool.
    - **HTTP** test and **Windows Shared Folders** test can now be configured in minimal time. The configuration page of both these tests has now been simplified, thus enabling administrators to quickly and easily configure multiple URLs/folders (as the case may be) for monitoring.
- **Logo per user:** Earlier versions of eG Enterprise allowed the eG administrator to configure a custom logo for the login screen and for every module of the eG manager. In v6, these logos can be personalized for each user to the eG Enterprise system. This way, different users can see different logos in the eG user interface. This capability allows service providers to customize the experience for different users.
- **Default role providing limited admin support:** With eG Enterprise v6, administrators can now create additional users with administrative privileges to configure the monitoring for the components that are assigned to them. These users can now configure tests, thresholds, alarm policies and maintenance policies for the components in their purview. The *MonitorwithLimitedAdmin* role included in eG Enterprise can be used to create such users. This capability allows delegated administration, which is a key requirement for many enterprises and service providers.
- **Instant installation of SSL certificate during AD integration:** Earlier, before integrating the

eG manager with an SSL-enabled Active Directory server, administrators had to manually install the SSL certificate on the eG manager. Version 6 saves administrators this trouble by allowing them to perform 'single-click' certificate installation from the eG admin interface itself. The **DOMAIN DETAILS** page of the eG administrative interface now provides an **Install SSL Certificate** button, which when clicked, enables the administrator to configure the name and location of the SSL certificate to be imported into the eG manager. The certificate so installed can be uninstalled at any later point in time using the **Uninstall SSL Certificate** button therein.

- **Support for AD groups:** Previously, if a user profile was created using the eG administrative interface, the monitoring preferences, scope, and privileges that were set as part of that profile applied to an individual user only. This means that if multiple users belonged to an Active Directory Group, then even if the rights and responsibilities of these users were the same, individual user profiles had to be created in eG for each of the users in this group. To avoid, v6 of the eG Enterprise Suite provides support for Active Directory Groups. This capability involves creating a profile in eG for the AD group as a whole, and not for every user in the group. This ensures that the monitoring preferences set for the group automatically apply to all the domain users in that group. This not only simplifies profile creation, but also significantly reduces the effort involved in modifying the profiles of users in an AD group or revoking the monitoring rights of the group users.
- **Multiple time zone support:** eG Enterprise is often deployed to manage servers in different geographies and time zones. For example, in a managed service provider environment, multiple customer infrastructures can be monitored from the same eG manager. In such situations, users (administrators in different geographies, customers of an MSP in different regions) prefer to see the performance metrics reported in their respective time zones. eG Enterprise v6 now allows time zones to be associated to each user's profile. By default, all users are associated with the local time zone of the location where the eG manager is hosted. However, users can change their time zone preferences to suit their requirements. When a user logs into the eG Enterprise console, all the metrics, alerts, and reports that the user accesses are displayed in the respective local time zone. This new capability ensures that eG Enterprise users receive a completely 'local' experience, regardless of which part of the world the eG manager is located in.
- **Automatic IP range discovery:** Previously, in the **DISCOVERY** page, administrators had to manually key in the IP range for discovery. In version 6 however, if eG Enterprise integrates with an Active Directory domain, then the eG manager automatically discovers the IP range for discovery and displays it in the **DISCOVERY** page. This greatly minimizes the time required to configure the discovery settings.
- **Separate page for auto-discovered segment topology:** The eG administrative interface now provides a separate page where auto-discovered segment topologies can be viewed. Earlier, no such page was available. To manage these discovered topologies, administrators can save them using this new page. This page also allows administrators to make changes to the discovered inter-relationships.
- **Auto-discovery enhancements:** Previously, when auto-discovery was enabled, the eG manager would detect all the possible components that could be mapped to a server and added all of these components to the eG manager's configuration. In eG Enterprise v6, priorities are pre-assigned to different applications for the purposes of auto-discovery. When discovering a server, the discovery process first looks for high priority applications on that host; high priority applications are discovered first and if any high priority application is detected, the discovery process for that host will stop. Other applications are discovered only if the higher priority discovery is not successful. This new capability ensures that eG Enterprise's discovery process only discovers the key components that administrators are likely to be interested in monitoring.
- **Automatic discovery of VMware vCenter:** Starting from v6, if a *VMware vCenter* server is added for monitoring using the eG administrative interface, that *VMware vCenter* server will be automatically available for configuring the discovery of *VMware vSphere/ESX* servers in the

environment. This was not the case earlier.

- **Discovery using host name of components:** Previously, the eG manager, by default, auto-discovered components based on their IP addresses only. In DHCP-enabled environments however, the IP address may change frequently. Administrators of such environments therefore, preferred to manage components using their host names and not their IP addresses. To address this requirement, the **COMMON DISCOVERY SETTINGS** page of the eG admin interface v6 allows administrators to indicate whether/not the target environment is DHCP-enabled, so that discovery can be performed accordingly. **This setting also applies to agent-based component discovery.**
- **Automatic component management:** Prior to v6, every time an auto-discovered component had to be managed, administrators had to switch to the **MANAGE/UNMANAGE** page in the eG administrative interface to manually manage that component. In version 6 however, this procedure can be dispensed with! The eG manager can now be configured to automatically manage specific components upon discovery. A comma-separated list of component patterns to be included/excluded from monitoring can be provided. If the discovery process discovers components that match the 'included' component patterns, the eG manager will automatically manage those components. Likewise, if an 'excluded' component is discovered, it will not be auto-managed.
- **Import/export manager configuration:** Prior to v6, if multiple managers were installed in an environment and administrators wanted the configuration of these managers to be consistent, they had to manually synchronize the configuration of each manager with the other. To save the time and labour involved in this exercise, the eG administrative interface of v6 now provides a special web page using which administrators can export the key configurations of an eG manager and reimport these configurations to another manager. This way, the configurations of both managers can be synchronized with minimal effort and time.
- **Asset management:** In eG Enterprise v6, it is now possible to record asset information for every application, device, or server being managed. Administrators can use the administrative interface to record details such as the name of the asset, description, type, location and state. Additional details on manufacturer, serial number of the asset, maintenance information and license information can also be recorded. Ownership details including the person to be contacted in the event of an issue can also be recorded. Asset information can also be mass imported into eG Enterprise from CSV files. When an application, server, or device experiences performance degradation, through the Alarms window in the eG monitoring console, a help-desk person has single-click access to the asset information of each problematic application, server or device.
- **Address search in Google maps:** In older versions, when configuring the exact location of a zone using the Google Maps interface (provided by the eG administrative console), administrators had to manually browse Google Maps to identify the zone location and mark it. Version 6 simplifies this by providing an **Address** bar where the location to be searched for can be entered. Upon specifying the location, the Google Maps interface automatically zooms into that location, thus enabling administrators to quickly mark it.
- **Automatic test configuration:** In versions prior to v6, if a disabled test is enabled for a component type, then that test had to be manually configured for a managed component of that type. Version 6 saves the time and trouble involved in this, by automatically configuring a test enabled for a managed component.
- **Component filtering when applying test config to other components:** In earlier versions, if the **Apply to other components** button in the **SPECIFIC TEST CONFIGURATION** page was clicked, the test configuration was automatically applied to all components chosen from a list of managed components in the environment. In version 6 however, this component list can be filtered on the basis of segment, service, or zone name, so that component selection is easier.
- **Default threshold settings for descriptors:** Before v6, descriptor-based thresholds could be set only when component-specific threshold settings are defined. Because of this, thresholds for aggregate descriptors such as *Total* or *Summary* could only be set at the component-level and not

at the global component type-level. To cater to this requirement, version 6 now allows 'default' or 'component-type level' thresholds to be configured for descriptors.

- **Intersecting thresholds:** Typically, the state of a measure changes when the upper or lower bounds of performance of that measure are violated – i.e., when the value of the measure falls 'outside' the range specified by its threshold setting. For some measures however, administrators may want to be notified if measure values stay 'inside' the configured range – not when they stray 'outside'. To support such requirements, eG Enterprise v6 allows the configuration of **Intersecting thresholds**. By default, this capability is enabled on the eG manager. Accordingly, for a measure, administrators can configure a *Maximum threshold* that is lower than its *Minimum threshold*. In other words, for the *CPU usage* measure for instance, you can set *80%* as the *Maximum Threshold* and *90%* as the *Minimum Threshold*. This ensures that the measure switches to an abnormal state only if its value falls below *90%* or exceeds *80%* - i.e., if it remains within *80%* and *90%*. You can disable this capability by turning off the **IntersectingThresholds** flag in the **[AGENT_SETTINGS]** section of the **eg_tests.ini** file (in the **<EG-INSTALL_DIR>\manager\config** directory).
- **Making eG Enterprise's email alerting more reliable:** The eG manager must to be configured to use a mail server for routing email alerts to users. If this mail server fails for any reason, then important problem notifications may not reach administrators. In turn, this causes performance issues to remain undetected (and hence, unresolved!). eG Enterprise v6 allows administrators to configure more than one mail servers for routing email alerts to users. When an alert is generated, the eG manager will first attempt to send out an email alert using the primary mail server. If it is unable to do so, then the eG manager will automatically try and send the email alerts using each of the configured backup mail servers in sequence, until it succeeds. This ensures that no problem goes unnoticed by administrators, even if one mail server is unavailable. Moreover, the next time an email alert needs to be sent out, the eG manager intelligently picks the mail server that successfully sent out alerts during the last attempt and uses that server first to process the alert.
- **Message board:** Version 6 does away with the distracting alerts that pop-up time and again in the eG administrative interface informing administrators of license expiry, agent status, etc. Instead, a stylish message board is provided, where the eG manager posts important notifications related to eG license expiry, agent status, component management, and more.

1.1.2 Monitor Interface

- **Design changes:** In v6, eG Enterprise has a refreshing new monitoring interface that is designed based on Web 2.0 concepts. Some of the key features of its unique design are as follows:
    - eG Enterprise v6 allows users the flexibility to choose between a light and dark color theme for the monitoring interface. **This capability is available for all modules of the eG management console.**
    - eG Enterprise v6 is sensitive to the needs of the color-blind. Accordingly, the state of each component, test and metric is indicated both by colors and by distinctive icons, so that color-blind users can use eG Enterprise to detect problems in their environment.
    - eG Enterprise v6 clearly separates the visual representation from the data that is rendered. This way, it ensures that only changes in data values are sent over the network, leading to bandwidth optimization.
    - The eG management console v6 embeds an intelligent search capability. Regardless of which interface you are on (admin/monitor/reporter/configuration) or what you are doing, you can instantly check on the status of your mission-critical servers, services, segments, and zones using this intuitive search engine. All you need to do is specify the search condition, and within seconds, the element you are searching for and its current status will

be made available to you.

- Previously, if a layer is clicked in the **Layers** tab page of a component, the tests mapped to that layer appeared in a separate **Tests** panel. In v6 however, the tests associated with a layer are displayed under the layer name itself in the **Layers** tab page. The measures reported by a test appear to the right of the layer model in the **Layers** tab page.

- **New dashboards:** A wide variety of new dashboards are now available in v6. Some of the useful dashboards have been briefly discussed hereunder:

  - **MyDashboards:** The MyDashboard capability of eG Enterprise allows users to build completely customizable dashboards. In v6, MyDashboard has been enhanced to provide a completely modern new look and feel. A dashboard can now include a wide range of graphical elements from simple tables and line graphs to timeline charts, area charts, dial gauges, live measure displays, and tier health indicators. Each panel of the dashboard can be configured individually and even live aggregate metrics indicating the overall demand, quality, or consumption of each component, tier or IT service can be displayed. Dashboard panels can be resized, repositioned, added or removed at will - by stretching/shrinking each panel, or through an easy drag and drop interface. Users can even select from pre-specified color themes. Dashboard panels can also be configured to display the local weather or the latest news on a specific topic. Any of the created dashboards can also be published through Microsoft SharePoint to other users in the organization.

  - **User Experience Dashboard:** eG Enterprise v6 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 infrastructure. This way, end users can easily determine when they see a slowdown, is the problem being caused by connectivity to the Citrix 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. End-users do not have to login to the eG monitoring console to access the dashboard. By entering his/her domain user name, an end-user can get to see the performance of his / her Citrix or virtual desktop session. Historical performance can also be observed for all key metrics. Citrix/Virtual desktop administrators can also use the same dashboard to handle user complaints. This results in happier users, fewer complaints to helpdesk, and lower support costs.

  - **Business Dashboard:** The **Business Dashboard** of v6 provides IT executives with a high-level view of the performance of their critical business services 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.

- **eG as a mobile application:** eG Enterprise v6 is now available as a mobile application for Android and iPhone devices. Mobile device users can now get a fully-optimized mobile experience – they can connect to the eG monitoring console from their mobile device, see an overview of the state of their infrastructure from the Monitor Dashboard, view current alerts, drill-down to the layer model of the problem components, and even view detailed diagnostics and graphs, just the way they would on their desktops. This way, administrators will be able to stay in touch with the goings-on in their IT environment even when on the move.

- **Changing thresholds made easy:** Prior to v6, if administrators had to change the threshold settings for a measure displayed in the eG monitoring console, they had to switch to the **Admin** interface, navigate to the **THRESHOLD CONFIGURATION PAGE** and then make the change. Version 6

saves administrators this trouble by enabling them to get to the threshold configuration page from the monitor console itself with a single mouse click!

- **Zoomable graphs:** In v6, users can zoom into any portion of a graph by selecting a time range on the graph. This allows for more interactive analysis of performance metrics.

1.1.3 Reporter Interface

In v6, the eG Reporter interface has been packed with a wide array of features focused on minimizing the effort involved in identifying the report that one requires and quick report generation. Some of these key features have been discussed hereunder:

- **Logical grouping of reports:** Previously, in the eG Reporter interface, reports were grouped by function – for instance, you had Event Analysis reports, Virtualization Reports, Thin client reports, etc. In v6 however, in addition to these functional groups, you have component-wise report groups. For any component that you pick, the eG Reporter interface will present to you all the reports that you can generate for that component. These reports will also be logically grouped based on function, so that you can instantly identify the report that you are looking for. This saves the hours that you would otherwise spend in identifying the report that would best serve your purpose.
- **Reports in a tree-structure:** The report groups of v6 are not provided to you as menu options, as done previously. A tree-view is now available. For instance, in case of the functional report groups, the individual report categories form the nodes of this tree, which when expanded will display the report options. This tree-structure makes navigation a lot easier.
- **On-demand report settings:** When generating a component-specific report, the new eG Reporter interface dispenses with the need to provide detailed specifications prior to generating the report. Once you identify the report you need and click on it, the report will be instantly generated using default settings. You can alter these settings at any time using a **MORE OPTIONS** pull-down.
- **New reports:** Version 6 also offers an assortment of useful new reports. The most important of these introductions have been detailed below:
    - **VM Changes Report:**

      Large virtualized environments often consist of tens of VMs hosted on a number of virtual servers. In such environments, the dynamic migration of VMs poses a huge management challenge for administrators. This is because, owing to the constant in and out movement of VMs, administrators find it very difficult to determine the correct location of a VM at any given point in time. The deletion of unused VMs compounds the problem. This is why, when a user complains of a slowdown when working on a VM, administrators take hours to figure out whether that VM still exists or not, and if so on which virtual host that VM is currently running.

      Using the **VM Changes** report provided by v6, administrators can easily identify the virtual host from which a VM was migrated, the virtual host to which that VM was added, and the VMs that were deleted/removed from a virtual host. With the help this report, administrators can quickly determine where all the VMs in a virtualized farm are currently operating and which VMs no longer exist.

    - **Physical Server Resources Report for Hypervisors:**

      This report provides deep insights into the resource utilization of the physical servers on a day-to-day basis as well as on a monthly basis. Using this report, the administrators can easily figure out the following:

        - How well the CPU of the physical server has been utilized over a period of time? Is it

under-utilized or over-utilized?

- Has the maximum CPU utilization limit reached?
- How well the memory of the physical server has been utilized? Are adequate memory resources available?
- What is the network usage of the physical server? Is the network available to the physical server?
- How well the disk of the physical server is utilized? Is it adequately sized or if additional disk is required for the storage activity of the physical server?

- **The Virtual Capacity – VMs Report:**

  In large virtualized farms characterized by tens of VMs per virtual host, administrators often struggle to identify resource-intensive VMs across the farm. It is also difficult to understand the reason behind the abnormal/excessive resource usage of the VMs – is it due to the resource hungry processes of the VMs or due to the poor resource allocation of the VMs? .

  The **Virtual Capacity** report is ideal for such environments as it enables the administrators to:

  - Accurately identify, at a single glance, which VMs, across the entire farm, are utilizing resources excessively and which ones are utilizing them poorly;
  - Determine the amount of resources currently allocated to the top and in the farm, instantly compare it with the usage, and isolate those VMs that are over-sized and those that are under-sized;

  Using these inferences, administrators can evaluate the effectiveness of their current capacity plans. If the present capacity decisions are found to be inadequate, administrators can draw up new plans and resize the VMs to ensure better resource utilization.

- **User VMs Report:**

  This report is available to only those registered users of eG who have been explicitly assigned VMs for monitoring. For such a user, this report reveals how well each VM assigned to that user is using the resources allocated to it. Sporadic/consistent spikes in resource usage and root-cause for such abnormal usage patterns can thus be deduced.

- **VM Right Sizing Report:**

  If a VM is under-sized in terms of CPU resources, the performance of all the applications running on that VM will suffer. On the other hand, if a VM is over-sized, it may have too many resources allocated to it and could starve other VMs of key resources, thereby leading to performance degradation on applications running on those other VMs. Over-sizing of a VM also results in unnecessary wastage of resources, thereby resulting in lower return on investment.

  In a large virtualized environment, it is often a challenge to identify which VMs are over-sized and which ones are under-sized. To help administrators quickly and accurately isolate such VMs, eG Enterprise includes a **VM Right-sizing** report. This report highlights the VMs on a chosen host(s) that either have more CPU resources than required or less CPU resources than they need. From this report, administrators can also get valuable hints on how to right-size these VMs. Administrators can use this information to right-size their virtual infrastructure for maximum return on investment.

- **Citrix Application Launch Report:**

  Citrix administrators are often interested in auditing user activity on their XenApp farms.

They would like to understand who accessed each application in the farm and for how long. The new **Citrix Application Launch** in eG Enterprise v6 addresses this need. Using the information in this report, administrators can determine which applications are most accessed, by whom and for how long. This information can also be useful if Citrix administrators have to bill different organizations in an enterprise for usage of the different applications published in the Citrix XenApp farm.

- **KPI Health Reports:**

  Analyzing a large infrastructure with hundreds of servers can take a lot of time and effort, and it also requires a great deal of expertise. **New KPI health reports** in eG Enterprise allow administrators to analyze the performance of an IT infrastructure in a few clicks and highlight bottleneck areas in the infrastructure. Drilldowns provide more details of the bottlenecks.

- **VDI Assessment Reports:**

  VDI assessment reports help administrators analyse the performance of virtual desktops. In an infrastructure with hundreds of virtual desktops, administrators need to quickly understand which desktops are consuming the highest amount of resources and which ones are configured with excessive resources. This is where the VDI assessment report helps. By identifying which desktops are taking resources, administrators can determine what action needs to be taken – e.g., prevent specific applications from running on the desktop, ensure that resource consuming desktops run on different physical hosts, etc.

1.1.4 Configuration Management

- **Configuration changes over email:** Previously, if eG Enterprise noticed configuration changes in a component at around the same time that a performance issue was observed on that component, it allowed users to quickly access the details of these changes by launching the Configuration Management portal directly from the **Alarms** window in the eG monitoring console. This enabled users to instantly diagnose whether/not the configuration change caused the performance bottleneck. For faster root-cause diagnosis, v6 includes the details of such configuration changes in the email/SMS alerts sent out for performance issues.
- **Comparison of configuration XML files:** In earlier versions, users to the **Configuration Management** console could use an intuitive interface to easily compare the configuration of two components and isolate discrepancies. In v6, this comparison capability has been extended to XML formatted configuration files. For instance, administrators looking to synchronize the configuration of two of their Tomcat servers can use this specialized interface to quickly compare the **server.xml** file of both the Tomcat servers and understand how they are different.

1.1.5 Integration

- **Trouble ticket integration:** Many trouble ticketing (TT) systems support web services APIs that monitoring tools can use to create, update and delete trouble tickets. eG Enterprise v6 can now be easily configured to route its alarms to a a TT system using a web services API. A generic framework supported in v6 allows any TT systems not supported out of the box to be easily integrated with eG Enterprise.
- **Enhancements to eG CLI:** The eG command line interface (CLI) can now be used to enable/disable tests and test descriptors. Multiple tests and descriptors can also be enabled/disabled in one shot using the CLI. Moreover, starting from v6, user profiles can also be created and maintained easily

using the eG CLI. This paves the way for touch-free administration of the eG Enterprise system.

- **Enhancements to eG SCOM Connector:** Starting from v6, eG Enterprise supports two-way integration with Microsoft SCOM. In earlier versions, the eG SCOM Connector collects state and alarm information pertaining to eG-managed components from the eG manager and transmits it to the SCOM manager. In version 6, administrators can optionally configure the connector to also check the SCOM server at configured intervals (default: 3 minutes) for eG alarms that may have been closed in the SCOM management console. If there are closed alarms in SCOM, the connector communicates their closure to the eG manager, thus enabling the manager to automatically close the same alarms at its end. This way, eG SCOM connector eliminates the need to manually synchronize the status of eG alarms between the eG and the SCOM managers.

1.2 eG Agent Changes/Enhancements

1.2.1 Monitoring Citrix/Microsoft RDS Environments

Version 6 adds more power to eG's Citrix XenApp and Microsoft RDS monitoring capabilities. The details are provided below:

- **Support for new components in the Citrix infrastructure:**
    - Out-of-the-box monitoring support is now available for the core components of the Citrix XenMobile infrastructure – i.e., Citrix XenMobile MDM, Citrix AppController, Citrix ShareFile, and On-premise Storage Zones.
    - Monitoring support is now available for Citrix CloudBridge.
    - Netscaler version 10.1 is now supported.
    - Monitoring support is now available for Citrix License servers on Unix (Linux, Solaris) as well.
- **Browser monitoring:** With eG Enterprise v6, the eG agent tracks all browser instances running on a Citrix XenApp server or a virtual desktop. While only IE (Internet Explorer) browser instances can be monitored on a XenApp server, IE, Firefox, and Chrome instances can be monitored on virtual desktops. The URLs accessed using the browser and the resource usage of each URL is revealed. This enables administrators identify which URL/web site is causing excessive resource usage by the browser.
- **Detailed User logon monitoring:** In version 6, eG Enterprise tracks user logons to the Citrix XenApp/Microsoft RDS server and provides logon time breakdowns for every logon. This enables administrators to quickly troubleshoot logon problems – is the slowdown due to Active Directory authentication or due to profile loading or due to group policies (if so, which policy), etc.
- **Monitoring the quality of HDX connections:** In v6, when monitoring the quality of user connections to applications installed on Citrix XenApp via ICA, eG Enterprise additionally reports the bandwidth used by every user, the input and output session line speeds for every user, the bandwidth utilized by each user for incoming and outgoing thinwire traffic, bandwidth used by each user when accessing multimedia content, and also the count of resource shared used by every user. With the help of these metrics, administrators can determine which user is using bandwidth excessively and what type of bandwidth intensive operations that user is engaged in.
- **Monitoring Citrix Receivers:** In version 6, eG Enterprise discovers and reports the count of the types of client devices that are connecting to the Citrix XenApp server via Citrix Receiver. Using the detailed diagnostics provided by eG, administrators can also determine which user is accessing using which device, and in the process, figure out if any device-related issues are contributing to a user's

unsatisfactory experience with Citrix.

- **GDI object monitoring:** An object is a data structure that represents a system resource, such as a file, thread, or graphic image. GDI objects support graphics. If an application creates a lot of these objects, without properly destroying references to the object (by closing the associated handle), then there will be multiple GDI objects occupying memory on the system for each object created. If this GDI leak is really bad, this can eventually bring a server to its knees, and cause all types of problems (slow logons, registry issues, system hangs, and so on). To help administrators avoid such eventualities, v6 of the eG Enterprise Suite periodically checks the GDI object handles created by each user to the Microsoft RDS server, reports the total number of handles created per user, and promptly notifies administrators if any user is creating more GDI handles than permitted. This way, probable GDI leaks can be proactively detected by administrators. In addition, detailed diagnostics are provided which reveal the process responsible for the GDI leak (if any).
- **Monitoring Citrix ICA/RDP listeners:** The listener component runs on the XenApp/Terminal server and is responsible for listening for and accepting new ICA/RDP client connections, thereby allowing users to establish new sessions on the XenApp/Terminal server. If this listener component is down, users may not be able to establish a connection with the XenApp server! This is why, eG Enterprise v6 tracks the status of the default listener ports and reports whether any of the ports is down. This check can be performed on Citrix XenApp, Microsoft RDS, and 2x Terminal Server components.
- **Monitoring multimedia event logs:** Using v6 of the eG Enterprise Suite, administrators can monitor the audio, RAVE, and flash event logs on Citrix XenApp server v7.5 (and above) and can capture critical problems impacting a user's multimedia experience with XenApp.
- **Monitoring Citrix broker agent:** A broker agent lies at the heart of any VDI deployment, and is the key component for assigning resources to end users. For a Citrix XenApp Server v7.5 (and above), version 6 of the eG Enterprise Suite additionally monitors the Citrix broker agent and reports the count of clients registered with the Citrix broker, the session load imposed by these clients on the Citrix server, and the nature of this load - – i.e., are they application sessions? or are they desktop sessions?
- **Details of disconnected sessions:** Detailed diagnostics are now available for the *Total disconnected sessions* measure of the **Citrix Disconnects** test, **Terminal Disconnects** test, and **Citrix XA Disconnects** test. The sessions that disconnected and the users who launched these sessions can be gleaned from these additional diagnostics.
- **User names pre-fixed with domain names:** From v6 of eG Enterprise, Citrix XenApp tests that take user names as descriptors or display user names in detailed diagnosis now prefix these user names with their domain name by default. This way, when analyzing user activity, Citrix administrators can quickly figure out which user belongs to which domain.
- **Ignoring down session IDs when monitoring user sessions:** Version 6 of the eG Enterprise Suite has introduced a new **IGNORE DOWN SESSION IDS** parameter to the **Citrix XA Sessions** test, **Citrix Sessions** test, and **Terminal Sessions** test. By default, this parameter is set to *65536,65537,65538* – these are nothing but the default ports at which the listener component listens. If any of these ports go down, then by default, this test will not count any of the sessions that failed when attempting to connect to that port as a **Down session**. You can override this default setting by adding more ports or by removing one/more existing ports.
- **Displaying application path in detailed diagnosis:** Previously, the detailed diagnosis of tests monitoring applications executing on and users logged into Citrix XenApp/Microsoft RDS servers, reported only the names of applications and their resource usage. In v6, the detailed diagnosis of these tests additionally reports the path from which the applications were accessed. If multiple versions of an application are published in different locations on the XenApp/RDS server and a user runs each of these versions, then this path information will indicate the exact application version

each process instance corresponds to – resource-hungry versions can thus be identified.

- **Monitoring Citrix Provisioning Services:** Monitoring of Citrix Provisioning Services has been enhanced in this version to additionally report the following:
    - The status of vDisks in a PVS store;
    - The target devices assigned to each vDisk;
    - The size and type of write cache of each vDisk;
    - The availability and responsiveness of the default store path and the default write cache path;
    - The streaming request load on the Citrix PVS and whether/not the server is configured right to handle the load;
    - Whether/not each vDisk in the PVS store is locked and if locked, how long it has been locked
    - Configuration metrics are now been collected from Citrix Provisioning servers.
- **Monitoring PVS write cache of Provisioned XenApp servers:** Where the Citrix Provisioning Server is used, all modifications, such as application installations, are written to a temporary file known as the Write Cache. When read requests for the newly written files come in, they are read from the write cache. In case of provisioned XenApp servers, administrators typically use the server's hard drive for storing the write cache. If the write-cache does not have enough disk space resources to grow, then many modifications to the vDisk will be lost. To avoid this, eG Enterprise v6 monitors the size and usage of the write cache and proactively alerts administrators when the write-cache runs out of space; this way, administrators are enabled to promptly initiate measures to prevent the data loss.
- **Monitoring using Citrix ODATA API:** Citrix XenApp and XenDesktop version 7 and higher support an Open Data (ODATA) API that third party applications can use to access the same metrics that administrators have access to from the Citrix Director tool. These metrics include data related to connection failures to virtual desktops, machines in a failure state, session usage, user logon duration with breakups of the login duration, and load balancing data. In v6, eG Enterprise leverages the ODATA API and reports the same metrics published in the Director console in the eG monitoring console. This way, Citrix administrators need not have to work with multiple consoles and all the key performance information about the Citrix infrastructure is available from the eG Enterprise console itself.
- **Citrix StoreFront Monitoring:** Additional statistics are now collected from the Citrix StoreFront server. These include the following:
    - In a Storefront server group, configuration information and details of users' application subscriptions are stored on and synchronized between all the servers in that group. With v6, administrators can promptly capture any data non-sync that may exist between a monitored server and the server group to which it belongs and can also determine if it is owing to latencies in synchronization.
    - StoreFront stores enumerate and aggregate desktops and applications from XenDesktop sites, XenApp farms, and AppController, making these resources available to users. With the help of eG Enterprise v6, administrators can periodically verify the availability and responsiveness of each store configured on StoreFront, and thus isolate the unavailable and unresponsive stores.

- **Monitoring Citrix Netscaler:** Citrix Netscaler monitoring has been improved in v6 to report the following:
    - The integrated cache is monitored, the type of requests served by the cache is tracked, and

the ratio of cache hits to misses is reported for each request type. This way, administrators can figure out how effectively the cache was utilized.

- Compression requests to the Netscaler are tracked, how well the Netscaler handles these requests is monitored, and the percentage of compression achieved is reported, so that issues in compression come to light.
- In this version, the **NS Global Memory Usage** test additionally reports the total memory and free memory of the NetScaler device, and thus reveals if the Netscaler is running short of memory resources.
- The **NS HTTP** test reports the count of HTTP requests and responses received during the last measurement period, enabling administrators to instantly capture a consistent increase in the HTTP workload.
- For each Netscaler interface, the **NS Interfaces** test additionally reports the input and output throughput of that interface during the last measurement period, and thus enables administrators to identify interfaces with low throughput.
- In this version, the **NS ICA Connections** test also reports the number of active user sessions to the Netscaler device via ICA. This serves as a good indicator of the ICA user load on the device.

1.2.2 Monitoring Virtualized Environments

- **Support for new virtualization platforms:** Monitoring support is now available for KVM server, Quality Virtual Desktop, Oracle VM Server, and Oracle VM Manager.
- **Monitoring of virtual host and "outside view" of VMs:** Version 6 of the eG Enterprise Suite reports additional metrics at the virtual host-level and as part of the "outside view" of performance of individual VMs. These new metrics have been detailed hereunder:
  - For a Citrix XenServer, the eG agent now reports the total space usage across all volume groups attached to the XenServer host, thus indicating the overall space usage by the host and accurately pointing to potential space contentions.
  - Similarly, the eG agent for Citrix XenServer also reports the network throughput (rate of data sent and received over the network) across all network interfaces supported by the xenserver host, thereby enabling administrators to proactively detect a probable slowdown in data transmissions/receptions over the network or high bandwidth usage.
  - For a Microsoft Hyper-V server, the eG agent v6 reports the aggregated space usage and network throughput across all the disk partitions and network interfaces (as the case may be) of the target Hyper-V server. In the process, administrators can determine how much disk space is available overall and how bandwidth-efficient the network traffic is.
  - When monitoring the datastores used by a VMware vSphere/ESX server or managed by a VMware vCenter server, the eG agent additionally reports the total capacity and usage across datastores, and in the process points to probable usage excesses, low throughput, and I/O processing bottlenecks.
  - For a Citrix XenServer, VMware vSphere/ESX server, and Microsoft Hyper-V server, the eG agent now reports the rate at which each VM handles read/write requests. This value is aggregated across all the virtual disks allocated to each VM and reveals the I/O load per VM and the ability of each VM to process the load. Likewise, the rate at which each VM processes network I/O is also revealed, so that the VM where processing is bottlenecked

can be identified.

- The **CPU – Esx** test mapped to a VMware vSphere/ESX server now reports the percentage of time a SMP virtual machine was ready to run, but incurred delay due to co-vCPU scheduling contention. Also, when reporting the outside view of performance of each VM on a VMware vSphere/ESX server, version 6 of the eG Enterprise Suite reports the percentage of time each VM was in the **COSTOP** state, the percentage of time the virtual machine was interrupted to perform system services on behalf of that virtual machine or other virtual machines, and the percentage of time each VM was consuming CPU resources. This helps administrators identify the VM that is hogging the CPU resources and the VM(s) that might be contending for limited CPU resources.

- **Enhancements to "inside view" monitoring:** Inside view monitoring has been enhanced in this version to additionally report the following:
  - For every application running on a virtual desktop, v6 additionally reports the rate of I/O reads, writes, read operations, write operations and total I/O operations performed by that application. This helps administrators isolate I/O-intensive applications running on virtual desktops.
  - Version 6 additionally measures page file usage of every Windows VM on a hypervisor, so that abnormal page file usage patterns inside VMs can be promptly captured. The time difference between the reference clock and each Windows VM is also reported, so that administrators can identify that Windows VM with a time setting that does not synchronize with the reference clock.
  - The **Memory Usage – VM** test for each VM on a virtual host now additionally reports the amount of memory, immediately available for allocation to a process or for system use, the amount of memory that is allocated to the modified page list, the amount of memory assigned to the standby list, and the amount of cached memory per VM.
  - The personal vDisk retains the single image management of pooled and streamed desktops while allowing people to install applications and change their desktop settings. If the personal vDisk runs out of space, users will no longer be able to hold on to their customizations, allowing them access to only the base VM and the applications installed therein. This outcome beats the entire purpose of having personal vDisks. To enable administrators avert such anomalies, eG Enterprise v6 tracks the status and space usage of the personal vDisk of each VM and promptly reports errors / abnormal space usage. This way, administrators can accurately identify personal vDisks with very limited space, which VM such personal vDisks are associated with, and what is consuming too much disk space – user profiles? Or user applications?
  - The **Disk Activity – VM** test for Linux and Solaris VMs can now be configured to run a *sudo* command to retrieve detailed diagnostics for the *Percent virtual disk busy* measure of the test.
  - The **System Details – VM** test for Linux VMs can now be configured to use the */usr/bin/top* command, instead of the default *usr/bin/ps* command for reporting the detailed diagnosis of the *Virtual CPU utilization* measure, if the latter does not function properly.
  - In v6, when monitoring the quality of user connections to desktops via ICA, eG Enterprise additionally reports the bandwidth used by every user, the input and output session line speeds for every user, the bandwidth utilized by each user for incoming and outgoing thinwire traffic, bandwidth used by each user when accessing multimedia content, and also the count of resource shared used by every user. With the help of these metrics, administrators can determine which user is using bandwidth excessively and what type of bandwidth intensive operations that user is engaged in.
  - In graphic intensive environments, the XenServer hypervisor is often configured with a

separate graphics processing unit (GPU) to enhance application performance. In such cases, graphics intensive portions of the application are handled by the GPU while the central processing units (CPUs) handle the compute intensive portions. In such cases, it is not sufficient to monitor CPU utilization alone – doing so will not reveal any bottlenecks in graphics processing. Therefore, in v6, eG Enterprise includes monitors for the GPU as well. GPU usage is tracked at the host level to determine bottlenecks. The processing requirements are also monitored at the virtual machine level, so administrators can understand which of the VMs has more graphic intensive applications running on them. Processing, power and memory usage levels of the GPU are reported in real-time, so any graphics processing bottlenecks are reported instantly to the administrator.

- For those customers who are only interested in the internal health and operations of each VM on a virtual host, a dedicated **ESX VMs Server** model has been newly introduced in v6. This model reports only the inside view of performance of the VMs on a virtual host.
- Inside view tests configured to run by default have been minimized in v6 to reduce the overhead of inside view monitoring.
- eG Enterprise now supports monitoring of virtual desktops even when the VMware BLAST protocol is used for virtual desktop access by users.

- **Microsoft Hyper-V Monitoring:** Monitoring of Microsoft Hyper-V servers now includes the following:
  - The heartbeat service installed on each VM is monitored and the health of the service and the VM it is operating on is reported. This way, administrators can determine whether a VM is accessible or not, and if so, what is causing it.
  - Hyper-V Replica allows virtual machines running at a primary site to be efficiently replicated to secondary location (Replica site) across a WAN link. In version 6, eG Enterprise monitors the replication activity performed by Hyper-V Replica for each VM on a Hyper-V host, instantly detects latencies or inconsistencies in the process, and proactively warns administrators of the same, so that the necessary corrective/control action can be taken. In addition, the status of the replication is also checked at frequent intervals and errors in the replication process captured and reported.
  - The amount of memory that the Hyper-V host has reserved for itself is periodically checked and reported in this version, thus enabling administrators to figure out whether/not the host has sufficient memory for its own operations.
  - The **Memory Usage** test of the Microsoft Hyper-V server has now been replaced with the **Hyper-V Memory Usage** test. This new test reports how much physical memory has been allocated to the VMs on the Hyper-V host. With the help of this test, administrators can proactively detect a potential memory contention on the Hyper-V host caused due to improper resource allocation to the VMs. The detailed diagnosis of this test also sheds light on that VM that could be over-sized in terms of memory.
  - Virtual machine snapshots are file-based snapshots of the state, disk data, and configuration of a virtual machine at a specific point in time. By reporting the count of large and aged snapshots per VM and revealing the names of these snapshots, eG Enterprise v6 helps administrators decide whether/not these snapshots can be removed to make more storage space available for a VM.
  - eG Enterprise v6 also performs in-depth monitoring of useful Hyper-V logs such as Hyper-V Config logs, Hyper-V High Availability logs, Hy per-V integration logs, Hyper-V VMMS logs, and Hyper-V Worker logs and helps administrators quickly capture critical errors/warning

events related to Hyper-V performance.

- In this version, the **Hyper-V Processor** test now reports detailed diagnostics revealing the CPU usage of each VM on the Hyper-V server. Using this information, administrators can quickly identify that VM that is hogging the CPU resources of the server.

- **Changes to VMware vCenter monitoring:** In v6, eG's VMware vCenter monitoring has been changed to include the following parameters:
  - In this version, tests mapped to the *VMware vCenter* component take an additional **SHOW VMS IN DD** flag. If this flag is set to **Yes**, then the detailed diagnosis of tests will include detals of VMs managed by vCenter. If set to **No**, then tests tests will not report VM details as part of detailed diagnostics.
  - In v6, the *vCenter Events* test mapped to a VMware vCenter component supports an additional **SHOW ALL ONLY** parameter. If this parameter is set too **Yes**, then all vCenter events will appear as descriptors of this test. This parameter applicable only when the **SHOWALLEVENTS** parameter is set to **Yes**.

1.2.3 Monitoring Connection Brokers

- **Monitoring desktop status:** When monitoring VMware Horizon View brokers, eG Enterprise v6 reports the count of desktops in various states of activity – in the READY state, in the error state, in the maintenance state, and the deleting state. The names of the VMs that are in each state are also provided as part of detailed metrics.
- **Reporting utilizing desktops:** In version 6, the Desktop Pools Usage **test** additionally reports the percentage of desktops actively used in each desktop pool managed by the VMware Horizon View broker. Detailed metrics reveals the names of the desktops that are being actively used.

1.2.4 Monitoring Databases

- **Support for new database servers:** Monitoring support is now available for PostGreSQL server v. 9.3.2.
- **Monitoring Microsoft SQL Servers:** eG's Microsoft SQL server monitoring capabilities have been enhanced in v6 to provide the following insights:
  - eG Enterprise now tracks the resource usage of the sessions to Microsoft SQL server. In the process, it turns the spotlight on resource-intensive sessions and the queries executed by such sessions that may require fine-tuning. Additionally, it also reports the average wait time of sessions, leads you to that session that has been waiting for the maximum time, and points you to the exact query that the session has been taking too long to execute. Inefficient queries are thus revealed, enabling you to quickly initiate query optimization measures.
  - eG Enterprise v6 monitors the uptime of the SQL server and in the process, captures sudden breaks in server availability and unscheduled server reboots.
  - In this version, when monitoring Microsoft SQL servers running Windows 2008 R2 (and above), the **SQL Database Space** test additionally reports whether/not auto-growth is enabled for a database. Where auto-growth is enabled, this test also reports the maximum size to which that database can grow and the percentage of free space in the database in

relation to the max size. This not only reveals rapid erosion of space in a database, but also prompts administrators to fine-tune the auto-growth settings so that more free space is always available to the database.

- When monitoring blocking processes on a Microsoft SQL database, eG Enterprise v6 additionally reports the count of processes that are being blocked by the root blockers. Detailed diagnostics reveal the names of the blocked processes. This way, administrators can identify those processes that are affected by the blocking processes.
- eG Enterprise v6 monitors the system and application processes running on the active node of a SQL cluster and reports the resource usage of and IOPS generated by the processes. Abnormal resource usage and I/O activity on the cluster can thus be captured. Detailed diagnostics reveal which processes are the leading resource consumers.

- **Monitoring Oracle database servers:**
  - The CPU and memory usage of client sessions to the Oracle database server are also monitored in this version. Sessions and queries consuming resources excessively are highlighted.
  - Oracle uses temporary tablespaces as work areas for tasks such as sort operations for users and sorting during index creation. To ensure that sufficient space is always available in the temporary tabelspaces for carrying out these critical operations, version 6 tracks the space usage of the temporary tablespace and alerts administrators if free space in the tablespace is alarmingly low.
  - The **Oracle Alerts** test in v6 now supports an **EXCLUDE PATTERN** parameter where administrators can provide a comma-separated list of alert log message patterns to be excluded from monitoring. The eG agent thus monitors only those alert messages that do not match the excluded pattern configuration.
  - Version 6 enables administrators to track queries to their Oracle, DB2, and DB2 DPF databases, and identify those queries that are resource hogs. By pinpointing the queries that generate the maximum disk I/O, that consume the maximum memory and CPU resources, and that which take the longest time to execute, this version brings to light those queries that are resource-intensive/time-consuming and are hence candidates for optimization.
  - In v6, eG Enterprise reports the count of nodes in an Oracle cluster and indicates the number and names of those nodes that are currently accessible. This way, the nodes that are inaccessible/unavailable can be identified. The percentage of available nodes is also reported, and administrators alerted if only very few nodes in the cluster are able to service the client requests. This signals a potential overload.
  - OraSqlNetTest DD query had been optimized recently to use sys.user$ system table. But this required that eG user must have permission to access this view. If there is no permission, we have to use the old query itself. Issue with the previous change that was checked-in is that we will keep executing new query every time in computeMeasures and if it fails we use original query. We have avoided this now, and if found that the user does not have sys.user$ table one time, we will go ahead with the previous query for subsequent executions.
  - A PGA is a memory region that contains data and control information for a server process. It is nonshared memory created by Oracle Database when a server process is started. In this version, eG Enterprise additionally monitors the memory consumption by the PGA and proactively alerts administrators if one/more server processes are draining memory from the PGA rapidly.
  - Version 6 also enables administrators to effectively analyse the workload of Oracle database

servers. For this purpose, eG Enterprise v6 reports the current CPU usage of the server and indicates its current load. In addition, it reveals the number and type of transactions the server processes every second, so that administrators can understand how well the server handles the load and can accurately identify where bottlenecks lie. Moreover, this version also reports how many key database operations – eg., data modifications, block changes, reads/writes, parses, rollbacks, etc. – are performed on the server per transaction. This way, the real workload of the server can be determined. These statistics also enable administrators to compare current CPU usage with the real workload, so that they can figure out whether/not the server needs to be resized to handle its load.

- eG's **Oracle SGA** monitoring capabilities have been enhanced in this version to report the amount of allocated space that the SGA is currently using, with pointers to the exact SGA component that could be over-utilizing the available space resources. Additionally, the percentage of buffer nowait requests, the percentage of hard and soft parses, the ratio of executes to parses, the percentage of CPU used for parsing and for other activities, are also reported, so that administrators can clearly gauge the health of the Oracle SGA.
- When monitoring system wait events on Oracle, version 6 additionally reports the total time for which waits on each wait event persisted, thus pointing administrators to those events that may have contributed to a server slowdown.
- eG's **Oracle Redo Log** test now additionally reports the percentage of redo entries for which there was space immediately available in the redo log. In the process, the test alerts administrators if a high percentage of redo entries are waiting for space to become available in the redo log.
- This version also tracks the growth in the size of Oracle alert logs and SQL error logs. Abnormal growth in the size of these logs can thus be promptly captured and controlled, before it increases the IOPS on the server and degrades its performance.
- In this version, the **Oracle Instance Status** test additionally reports the uptime of the Oracle instance and also indicates whether/not it was rebooted recently, thus shedding light on unplanned instance restarts.
- Previously, the **Oracle Root Blockers** test reported even that session that was blocking just one other session as a *root blocker*. However, since blocking is considered an issue in the real world only when the root blocker blocks more sessions, version 6 introduces two additional parameters to this effect to the **Oracle Root Blockers** test. While the **BLOCKED SESSION COUNT** parameter can be used to indicate the number of sessions that a blocker should block for it to be counted as a root blocker, the **MAX BLOCKING TIME** parameter can be used to indicate how long a blocker should block a session for it to be considered a root blocker. Any session that fulfils one of the two conditions above will be treated as a blocking session.

- **Monitoring SAP HANA database servers:** In this version, the **HANA Service** test and the **HANA Uptime** test mapped to a *SAP HANA Database* server additionally report whether/not the service/server (as the case may be) was rebooted in the last measurement period. This indicates whether the reboot was a scheduled – i.e., expected one – or unscheduled/unexpected.

1.2.5 Monitoring Windows/Unix/Other Systems

- **Support for MAC systems:** Agentless monitoring support, based on SNMP, is now available for *Apple Macintosh OS*.
- **Java Runtime of the eG agent:** The eG agent for AIX, HPUX, Linux and Solaris is now bundled with JRE 1.6. Windows agents in earlier versions already had JRE 1.6.
- **Monitoring Windows systems/servers:** In eG Enterprise v6, the following

changes/enhancements have been introduced in eG's Windows server monitoring capabilities:

- The time difference between the reference clock and a Windows host is also reported in this version, so that administrators can be alerted if the time setting of the host does not synchronize with the reference clock.
- Besides application and system event logs, administrators can now monitor error/warning/general information events logged in custom event logs on a Windows host.
- If a process does not collect **Windows messages** from its queue - say, using the *GetMessage* function or *PeekMessage* function - it will be tagged as "not responding" - because it is not responding to **user interface events**. Using eG Enterprise v6, administrators can quickly figure out the number and names of the 'not responding' processes on a Windows host, so that reasons for lack of responsiveness can be investigated and treated.
- The **Windows Scheduled Tasks** test now takes an additional **EXCLUDE FOLDER** parameter. **This parameter applies only to Windows 2008 (and above) systems**. Using this parameter, administrators can exclude all scheduled tasks that reside in specific folders from the monitoring purview of this test. For this, a comma-separated list of folders to be excluded should be provided in the **EXCLUDE FOLDER** text box.
- If the **Fail over cluster service** is enabled on a Windows server, then the *Windows* monitoring model will automatically change to include an additional **OS Cluster** layer. This layer indicates which nodes are part of a Windows cluster, the current state of each node, and alerts administrators to the failure of the cluster service on any node. In addition, the layer auto-discovers services/applications that have been clustered, reports service/application failures, and pinpoints the probable reasons for the same. Cluster networks that are down and cluster resources that are offline can also be identified. Moreover, the layer tracks the current capacity and usage of cluster disks and cluster shared volumes and proactively detects potential space crunches.
- To enable the **Disk Space** test on Windows to discover and report the status of NFS drives mapped to the target Windows host, a new **DISCOVER NFS** flag has been introduced in the **Disk Space** test. Set this flag to **Yes** to monitor NFS drives. In this case therefore, you will have to additionally configure the test with domain user permissions to enable the discovery.
- The detailed diagnosis of the **Handles Usage** test now reports the number and types of open handles, thus leading administrators to what could be causing a memory leak on the Windows server/host being monitored.
- Detailed diagnosis of the **Windows Processes** test additionally reports the name of the 'User' running resource-intensive processes on the monitored host.

- **Monitoring Unix systems/servers:** In version 6, the following changes/enhancements have been effected in eG's Unix server monitoring capabilities:
  - The **Memory Usage** test (for both Windows and Unix) now additionally reports the amount of memory, immediately available for allocation to a process or for system use, the amount of memory that is allocated to the modified page list, the amount of memory assigned to the standby list, and the amount of cached memory. These metrics reveal how effectively the target system is using the memory resources.
  - Previously, the **System Details** test of a Linux server was reporting metrics for only the *Summary* descriptor. In this version, this test reports metrics per processor supported by the Linux server.
  - The **ZFS Pools** test mapped to a Solaris server now reports an additional *Scrub status*

measure, which indicates the status of ZFS scrubs that may have been performed on each ZFS pool during the last 8 days. ZFS Scrubs allows you to schedule and manage scrubs on a ZFS volume. Tracking the status of scrubs reveals the pools on which scrubs are in progress and the pools on which they have completed.

- Detailed diagnosis of the **Memory Usage** test (for both Windows and Unix) additionally reports the name of the 'User' running resource-intensive processes on the monitored host.
- The **Disk Activity** test for Linux and Solaris servers can now be configured to run a *sudo* command to retrieve detailed diagnostics for the *Disk busy* measure of the test.
- The **System Details** test for Linux servers can now be configured to use the */usr/bin/top* command, instead of the default *usr/bin/ps* command for reporting the detailed diagnosis of the *System CPU utilization* measure, as the latter may not function properly on some flavors of Linux.
- The **Processes** test for Linux servers can now be configured to */usr/bin/top* command, instead of the default *usr/bin/ps* command for reporting process health metrics, as the latter may not function properly on some flavors of Linux.
- Previously, Cron jobs could be monitored only on Solaris servers. In this version, cron job monitoring is supported in all Unix operating systems.
- The **Processes** test on Unix has now been configured to perform even case-insensitive searches for configured **PROCESSPATTERNS**.
- The **Disk Space** test for Unix servers now supports a **USE SUDO** flag. Turning this flag on enables the discovery of all drives supported by the Unix server. Without the *sudo* permission, some drives supported by the target Unix server may not be discovered owing to lack of access rights.

1.2.6 Monitoring Active Directory Servers

- **Monitoring FSMO roles:** FSMO stands for Flexible Single Master Operations, and FSMO roles (also known as operations master roles) help you prevent conflicts in your Active Directory. Version 6 of the eG Active Directory Monitor monitors domain controllers configured with the FSMO roles, and help you rapidly detect the unavailability of an FSMO domain controller over the network, isolate potential network connectivity issues and latencies, and spot real/probable delays in LDAP binding, so that such issues can be promptly remedied and conflicts prevented.
- **Monitoring AD domains:** For a configured list of fully-qualified domain names, the eG Active Directory Monitor v6 reports the number and names of all domain controllers that manage the servers and users in each of the configured domains. **This capability is available for AD servers on Windows 2008 only**.
- **Performing DNS health checks:** Active Directory uses DNS as its domain controller location mechanism and leverages the namespace design of DNS in the design of Active Directory domain names. In v6, the eG AD Monitor leverages the **DCDIAG** utility to run a set of DNS-centric health checks, and in the process, reports on the overall DNS health of the domain controllers. It also periodically emulates a client accessing DNS to issue a query, and in the process, reports the availability and responsiveness of the DNS server.
- **Monitoring the AD replication service:** The following key metrics on replication are additionally reported by eG's monitoring model for Active Directory servers in version 6:
    - **Replication** is the process by which the changes that are made on one domain controller are synchronized with all other domain controllers in the domain that store copies of the same information or replica. The eG AD Monitor v6 closely monitors the replication activities on the domain controller and promptly reports replication failures, so that administrators

can investigate such failures, discover the reasons for the same, fix them, and restore normalcy.

- In v6, the eG AD Monitor continuously tracks the replication queue length, so that AD administrators can be alerted as soon as the number of work items in the queue crosses an acceptable limit. Administrators can also use the detailed diagnostics reported by this Monitor to know what type of synchronization requests are in queue, so that they can figure out why the requests are taking too long to be processed.
- Version 6 of the eG AD Monitor also summarizes the replication state and relative health of an Active Directory forest by inventorying and contacting every domain controller in the forest, and collecting and reporting information such as replication deltas and replication failures. This enables administrators to accurately identify the domain controllers that are prone to frequent failures.

- **Monitoring lingering objects:** A lingering object is a deleted AD object that re-appears ("lingers") on the restored domain controller (DC) in its local copy of Active Directory. The eG AD Monitor v6 scans the event logs for replication events related to lingering objects, and promptly alerts administrators upon the occurrence of such events. Using the detailed diagnosis reported by this Monitor, administrators can easily determine the location of the lingering objects, so that they can immediately proceed to remove them. This way, the test ensures that the replication engine operates without a glitch.
- **Monitoring user account lockouts:** Account lockout is a feature of password security that disables a user account when a certain number of failed logons occur due to wrong passwords within a certain interval of time. In v6, the eG AD Monitor has been engineered to periodically report the count of locked user accounts and the names of the users who have been affected by this anomaly.
- **Monitoring global catalogs:** The global catalog is a distributed data repository that contains a searchable, partial representation of every object in every domain in a multidomain active directory domain services (AD DS) forest. Version 6 of the eG AD Monitor additionally monitors the global catalogs in the target domain controller and reports the number of catalogs that are currently available and unavailable. This way, the test enables administrators to determine whether/not adequate global catalogs are available in the domain controller to handle the request load.
- **Monitoring user account status:** In v6, the eG AD Monitor also reports the status of user accounts configured in the Active Directory server and thus, quickly points administrators to 'unused' accounts that can be deleted to make room for those that are actively used.
- **Monitoring the 'Audit account management events':** Time and again, it is important to verify whether the addition of new users/computers/groups to an Active Directory domain, changes to existing user/computer/group accounts, and deletion of accounts were performed only by authorized personnel and with no malicious intent. To track such operations, "Audit account management events" provides specific event IDs. In v6, the eG AD Monitor continuously tracks events with the event IDs grouped under *Audit account management events*, and proactively alerts administrators to the sudden addition/modification/deletion of users/groups/computers in the Active Directory. The Monitor additionally reports which user performed the addition/modification/deletion and when.
- **Monitoring computers in domain:** In v6, the eG AD Monitor takes stock of the total number of computers managed by the AD server and the status of these computers, so that administrators can determine from a single glance which computers are inactive/unused.

1.2.7 Monitoring SAP Environments

eG Enterprise now provides deep visibility into the performance of the SAP ABAP stack. Enhancements in this release are focused on many areas – tracking the health of all the communications from the SAP ABAP stack

to external applications, the health of the internal processes of the SAP ABAP system and understanding and analysing the user experience and activity on the SAP ABAP system. The details of these enhancements are as follows:

- **Monitoring of IDocs:** IDocs are structured ASCII files (or a virtual equivalent). They are the file format used by SAP ABAP to exchange data with foreign systems. eG Enterprise now monitors the inbound and outbound IDocs generated and reports the rate at which these IDocs were processed at various stages of transmission/reception, thus accurately pointing to processing slowdowns and where exactly the processing was bottlenecked. In addition, the Monitor also reports the number of IDocs that were found to be erroneous every second and the exact stage of transmission/reception at which the rate of errors peaked. This way, administartors can easily determine why and where errors have occurred in electronic data exchange.

- **Monitoring of RFCs:** Communication between applications within an SAP system and also with a remote system can basically be achieved using the Remote Function Call (RFC). tRFC or Transactional RFC is appropriate for communication between two independent SAP systems. In this scenario, data is transferred by tRFC, meaning that each function module sent to the target system is guaranteed to be executed one time only. qRFC performs a serialization of tRFC (Transactional RFC) using wait queues. While the actual sending process is done by the tRFC, inbound and outbound queues are added to the tRFC, thus resulting in a qRFC (queued Remote Function Call). qRFC can be used for communications between a SAP system and a non-SAP system. In v6, eG Enterprise tracks the tRFC and qRFC calls sent and received by the SAP ABAP system and reports the status of the calls, captures errors in the process, and highlights long running calls, thus pinpointing what is ailing communication between SAP systems and with non-SAP systems.

- **Monitoring of ICM:** The Internet Communication Manager (ICM) facilitates communication between SAP system(s) and the internet using the HTTP, HTTPS, and SMTP protocols. eG Enterprise periodically checks the availability, thread pool usage, and connection utilization of the ICM, and promptly reports the non-availability of the ICM, abnormal usage of worker threads by the ICM, and the over-utilization of ICM connections.

- **Monitoring background jobs:** In this version, background jobs are monitored, and their status, runtime and startup delays are reported. This way, failed jobs and those that have been running for long hours can be isolated.

- **Database consistency monitoring:** Inconsistencies in data between a SAP ABAP system and the database can cause critical SAP transactions to fail. To help administrators avoid such failures, eG Enterprise runs periodic consistency checks on primary indexes, secondary indexes, tables, and views, and proactively alerts administrators to inconsistencies.

- **Workload analysis:** In v6, the SAP workload is analyzed, overload conditions are isolated, and probable processing bottlenecks are detected. Moreover, the SAP task types and transactions contributing to the heavy load, excessive resource usage, and processing slowdown of the SAP ABAP system are highlighted.

- **Workflow linkages monitoring:** When defining a SAP workflow, you can create event linkages to specify the provider for the events that your application raises, to define the conditions based on which the events must be raised, and to map the event data type and event provider operation type. eG Enterprise scans workflows for event linkages, reports the total number of linkages, and brings linkage failures to the attention of administrators. Additional diagnostics also reveal the details of the failed linkages.

- **Work process monitoring:** Work processes of the SAP ABAP stack are components that are able to execute an application (that is, one dialog step each). When the SAP ABAP stack is initialized, it is configured with a pre-set maximum number of work processes. If the number of work processes configured is fully utilized at any point of time, new requests will not be processed and applications

using the SAP ABAP stack will fail. Hence, the utilization of work processes is a key metric of SAP ABAP health. In v6, eG Enterprise tracks the utilization of work processes and alerts administrators if the utilization of work processes is high. Using eG's detailed diagnosis capability, administrators can determine which applications are responsible for the high utilization of work processes - e.g., is a situation where the work processes are fully utilized caused by heavy load or is a specific application utilizing all of the work processes.

- **User activity monitoring:** SAP is a shared infrastructure where multiple user simultaneously login and perform activities and transactions. If a specific transaction is very resource intensive, it can slow down the performance seen by other users. Background jobs can consume server resources. eG Enterprise v6 provides deep visibility into user activity and user experience in a SAP ABAP environment. The number of users logged in and their session details (time logged in, time logged out, etc.) are tracked, so administrators can determine which users access the SAP servers for the longest period of time. Login failures are monitored and the reasons for failure are detected and reported. The resource usage of each user, the number of transactions executed and detailed response time for each step of a transaction are all monitored. This way, eG Enterprise can help SAP administrators answer some of the key questions about the SAP infrastructure including:
    - Which user is executing resource-intensive transactions on the ABAP system?
    - Which user is overloading the system?
    - Which user is experiencing slowness when running transactions on the system? Where did this delay occur? – in the dispatcher queue? when loading/generating objects? when rolling-in/rolling out user contexts? in the database? when performing enqueue operations? or when waiting for RFC calls to complete?
    - What scheduled jobs are running in the SAP ABAP stack?
- **Monitoring TemSe:** TemSe is the SAP store location for temporary sequential data. Temporary sequential data are objects that are not normally permanently held in the system and are stored in the TemSe. The SAP spool system uses the TemSe to store output data temporarily. Version 5.6 of the eG Enterprise Suite monitors each TemSe – i.e., each location that stores temporary sequential data – measures the rate at which this data is created and is growing in every location, and points to the modules that create this data – eg., spool, background jobs, HR module, etc. In the process, abnormal growth in TemSe size and the module responsible for this come to light.
- **Monitoring the SAP's CTS:** The Change and Transport System (CTS) helps administrators organize development projects in ABAP Workbench and in customizing, and then transporting the changes between the SAP systems in the system landscape. Version 6 monitors SAP's Change and Transport system and reports the status of changes exported to other systems by CTS.
- **Monitoring spool requests:** Version 6 tracks spool requests to the SAP ABAP system, reports the count of such requests, the rate at which the SAP ABAP serviced these requests, and the number of requests that failed. Bottlenecks in spool request processing can be proactively detected and failed requests can be captured.
- **Monitoring Syslogs:** In this version, syslogs are monitored and the errors/warning events captured by the logs are reported, so that issues affecting SAP ABAP performance can be captured early.
- **Monitoring SAP locks:** For each type of lock active in the SAP ABAP system, eG Enterprise v6 reports the count of locks of each type and the duration for which the locks were held. Detailed metrics provide the complete details of locks that have been held for long time periods.
- **Monitoring batch input sessions:** BDC stands for ***Batch Data Communication*** and is also referred to as Batch Input. It is a technique for mass input of data by simulating user inputs in screens of existing transactions. Version 6 of the eG Enterprise Suite monitors the BDC sessions and

reports the number of sessions in various operational states – say, sessions with errors, sessions that are being created, background sessions, and sessions in progress. In the process, sessions with errors can be identified and their details viewed as part of detailed metrics.

1.2.7.1 Monitoring Exchange Mail Servers

- **Support for new versions:** eG Enterprise v6 provides out-of-the-box monitoring support to Exchange 2013. A dedicated monitoring model is available in version 6 for this.
- **Monitoring Exchange 2010:** For Exchange 2010, version 6 reports a wide variety of new metrics. These have been discussed below:
    - Exchange ActiveSync connectivity is checked at every stage (a.k.a scenario) of synchronization with a user device – eg., the Logon stage, the FolderSync stage, the Options stage, etc. - and latencies (if any) in connectivity are reported, thus leading administrators to the exact stage at which the failure/slowdown occurred.
    - The email traffic to and from the Exchange server is continuously tracked, so that potential overload conditions can be detected and the source of the overload can be isolated – is it because of mails sent/received by users in the intranet? Or mail traffic over the internet?
    - All aspects of replication and replay are monitored in this version and reports on the health of each aspect – for instance, the replication pipeline, the availability of the Active manager, and the health and status of the underlying cluster service, quorum, and network components.
    - eG Enterprise v6 attempts to connect to the Outlook Web App URL from the intranet, and for every stage (a.k.a scenario) of the connection process, it reports whether/not that stage completed successfully and the time taken for completion. This not only points to a potential OWA connectivity failure/slowdown, it also points administrators to the exact stage at which the failure/slowdown may have occurred. This brings connectivity issues in the internal network and their probable causes to light.
    - The Mailbox Replication Service (MRS), which resides on all Microsoft Exchange Server 2010 Client Access servers, is the service responsible for moving mailboxes, importing and exporting .pst files, and restoring disabled and soft-deleted mailboxes. eG Enterprise v6 periodically checks whether/not MRS is running and how quickly it responds to an RPC ping check.
- **Monitoring Exchange ActiveSync:** Exchange ActiveSync monitoring has been enhanced in v6. The details are as follows:
    - The number of hits/accesses made by each device synchronizing with the Exchange server and the amount of data transmitted and received by that device, so as to pinpoint:
        - Devices that are currently connected to the Exchange server; unauthorized devices can thus be quickly captured;
        - Devices that are accessing the Exchange server mailboxes frequently and those that seldom use the mailboxes; sizing and policy decisions can be taken based on this observation
        - Devices that are consuming excessive bandwidth resources and could hence be contributing to the sluggish quality of the network;
        - The number and percentage of accesses that returned each HTTP status code, thus shedding light on ActiveSync errors that occurred frequently.
        - The number of ActiveSync accesses and the number and size of items transmitted and received by each device type synchronizing with Exchange; this pinpoints

those device types that are utilizing the available network and server resources excessively, thus degrading the experience of some or all device users.

        - Devices that are not OK – i.e., devices that are either unable to connect to ActiveSync or cannot synchronize with Exchange Server via ActiveSync
    - The error/warning conditions captured recently by ActiveSync logs;
    - Count of requests that fail to comply with or are only partially compliant with ActiveSync policies, so that the probable reason for synchronization failures can be ascertained.

- **Monitoring Exchange mailbox database space:** Mailbox Database can be considered a container that stores and maintains all mailboxes of the users. Lack of space in an exchange database can cause the mailboxes in that database to stop receiving emails! To avoid such an outcome, eG Enterprise v6 continuously tracks the space usage in each of the Exchange mailbox databases configured on the Exchange server and rapidly isolates the mailbox database that is running out of space. The test also helps administrators differentiate between actual free space and white space.
- **Test configuration made easy:** In eG v6, tests mapped to managed Exchange 2007 and Exchange 2010 servers can be configured quickly and easily. This is because, eG auto-discovers and auto-configures tests with the location of the Exchange management shell, using which the tests collect metrics.

1.2.8 Monitoring Storage Devices

- **Support for new storage devices:** Out-of-the-box monitoring support is available for the following new storage devices in version 6:
    - HP P 2000 SAN
    - IBM Storewize v 7000
    - Atlantis ILIO
    - QNAP NAS
    - Clustered NetApp
    - Data Domain
- **Monitoring block storage pools on EMC VNX:** Using eG Enterprise v6, EMC VNX administrators can monitor block level storage levels, track the capacity and space usage of each pool, and bring pools running out of storage space to light.

1.2.9 Monitoring vCloud Director

eG's vCloud monitoring capabilities have been enhanced in v6. The details are as follows:

- **Monitoring the message bus:** The message bus in the vCloud Director cell is used for inter-cell communication, based on ActiveMQ (the message broker) and JMS (which defines the message format). The eG agent now reports the health of the message bus connection and helps the administrator to figure out the number of vCloud Director cells that are currently reachable and not reachable in a vCloud Director infrastructure.
- **Monitoring the console proxy:** The Console Proxy along with the VMRC i.e., the VMware Remote Console client in the vCloud Director cell enables a user to view the screen of a specified vApp (VM). In v6, the eG agent additionally reports the total number of Console Proxy connections and the number of connections that are currently active, and thus holds a mirror to the current workload of the vCloud Director Cell
- **Monitoring the vCloud Director database:** vCloud Director cells use a database to store shared

information. Long running transactions to the database can keep resources locked unnecessarily and can cause unprecedented delays in processing subsequent transactions. Version 6 of the eG agent proactively alerts administrators to transactions and sessions that have been active for a long time, so that they can quickly deep dive into the database server to know which sessions and transactions have been running for a long time and why.

- **Monitoring web requests:** In this version, the eG agent also monitors the monitors the web requests to the vCloud Director UI and the REST API and reports the number of active requests, so that the workload on the vCloud Director can be measured. In the event of an overload, this will help administrators determine which type of requests is high - UI requests? or REST API requests?
- **Monitoring queries via the query service:** Using the vCloud API query service, administrators can build typed and packaged queries that can be executed on the vCloud Director database for retrieving information about objects in the cloud. The eG agent v6 enables administrators to track the status of every query built and executed via the query service, so that query failures can be instantly detected.
- **Monitoring user sessions:** In this version, the eG agent tracks the user sessions to the vCloud Director cell, and in the process, alerts administrators to login failures and potential overload conditions.
- **Monitoring validators:** The validators configured on the vCloud Director cell are auto-discovered and the count of vCenter (VC) object executions that are available for validation per validator are reported in this version, so that the load per validator can be ascertained. Also, the length of the validation queue is measured continuously, so that probable processing bottlenecks in the validator can be captured early.
- **Monitoring user connections:** The eG agent v6 monitors cell user connections to each vCenter server instance and promptly brings disconnected instances to light. In addition, it also reports reconnects and the number of times the VC listener had to be started owing to the disconnects.
- **Monitoring inventory events:** For every type of inventory events that were handled by the VCD cell, version 6 of the eG agent reports the number of inventory events that were dispatched. This enables administrators to identify the type of inventory events that occurred the maximum on the VCD cell.
- **Monitoring the VIM inventory event manager:** The eG agent now measures the processing ability of the VIM Inventory Event Manager by indicating the number of object updates that could not be processed by the event manager. In addition, the eG agent v6 also monitors the PropertyCollector - the managed object type that clients use to retrieve a set of properties from one or more managed objects. In the process, it reports how frequently objects were updated and were ready for processing by the collector.
- **Monitoring vCloud Director tasks:** vCloud Director tasks represent long-running operations on a vCloud Director cell. In version 6, the eG agent periodically checks the vcloud Director cell for long running tasks, and for each task it auto-discovers, it promptly reports the count of failed instances of that task. This way, administrators can identify tasks that fail frequently. In addition, the eG agent also reveals the total time for which each task has been running, thereby pinpointing the tasks that have been running for the longest time. The Virtual Center Task Manager statistics are also reported in this version. From these metrics, the failure of tasks can be promptly detected, the count of tasks waiting for inventory can be determined, and the time taken for task completion can be measured. This information helps isolate bottlenecks in task execution.
- **Monitoring organizations:** An Organization is the fundamental vCloud Director grouping that contains users, the vApps that they create, and the resources the vApps use. The eG agent v6 reports the number of user sessions that are currently active on each organization configured on the vCloud Director cell and thus indicates the user load on each organization.
- **Monitoring access to database objects:** Version 6 of the eG agent also monitors accesses to

database objects and reports whether the last database object access was slow or not. In the process, the test points you to road-blocks (if any) to database access.

1.2.10 Monitoring Network Devices

- **Support for new network devices:** eG Enterprise v6 provides out-of-the-box monitoring support to Sonic and Watchguard firewalls, the Coyote load balancer, Infoblox, Juniper EX switch, and the Delta UPS.
- **Monitoring Juniper Netscreen firewall:** eG's Netscreen firewall monitoring capabilities have been improved in this version to report on the health of the fan and power units, to monitor temperature fluctuations, and track the battery status. This enables administrators to promptly detect fan failures, abnormal spikes in temperature of the firewall, and poor battery status.
- **Monitoring the F5 Traffic Manager:** For an F5 traffic manager, eG Enterprise v6 additionally tracks the connections to each F5 pool, reports the percentage of connections that are currently active, and alerts if connection usage is abnormally high. The virtual servers are also auto-discovered and the network traffic to and from each virtual server is reported, so that the busy servers can be isolated. Connections to the virtual servers are also tracked and the connection usage is reported, so that administrators can capture abnormal usage trends.

1.2.11 Monitoring Java Web/Application/Messaging servers

- **Support for new J2EE-enabled application servers:** eG Enterprise v6 provides useful insights into the performance of OC4J, WebSphere JCA, WebSphere JEUS, and JBoss 7.x.
- **SOAP Monitoring:** Using SOAP, eG Enterprise v6 monitors the availability, responsiveness, and the quality of content returned by web services that are supported by WebLogic, WebSphere, Oracle 9i, JBoss, GlassFish, JRun, SunONE, and Tomcat application servers.
- **Monitoring SIB-enabled WebSphere servers:** A service integration bus (SIB) is a group of one or more application servers or server clusters in a WebSphere® Application Server cell that co-operate to provide asynchronous messaging services. A destination is a point of addressability within the bus. Messages are sent to and received from destinations. A destination can be a queue or topic. Version 6 of the eG agent monitors the queues and topics on SIB-enabled WebSphere application servers, tracks the message queue length, measures how long messages have been in queue, checks how quickly messages are getting published on topics and consumed by subscribers to the topics, and thus points to bottlenecks in the messaging function and where the bottleneck lies.
- **Monitoring message queues on JBoss:** For each message queue on JBoss, v6 additionally reports the queue length, count of messages delivered per queue, the time that elapsed since the last message was sent from each queue, the number of messages added per queue, and the number of subscribers to messages in each queue. Administrators can thus capture consistent spikes in queue length, which would typically indicate a message delivery bottleneck. They will also be able to detect delays in message delivery by observing variations in the time elapsed since the last message was sent in each queue. Detailed metrics reveal who are the subscribers to the messages in each queue.
- **Monitoring the channels of WebSphere MQ servers:** The **WebSphere MQ Channel** test of version 6 additionally reports the channel status, thereby pointing administrators to those channels in an abnormal state – this could be the 'Retrying' or 'Stopped' states.
- **Monitoring Java Transactions:** The following enhancements have been included in the Java transaction monitoring capabilities of v6:
    - Previously, whenever a transaction encountered any error/failure, the eG JTM automatically included that transaction in its *Error transactions* count. In some cases, administrators may want to disregard certain errors. In such cases, administrators may want to make sure that

the transactions in which such errors occur are not counted as *Error transactions*. To enable this, eG JTM v6 allows administrators to specify errors that can be excluded from monitoring.

- Earlier, when monitoring application transactions, the eG JTM monitored every method in every class file to locate the exact method that caused application performance to degrade. Some administrators however may want certain packages or classes to be excluded from monitoring – for instance, administrators may want the eG JTM to monitor only the application-related packages and not the server-related ones. This is why, the eG JTM v6, by default, ignores all application server-specific packages (eg., class files related to Tomcat, WebLogic, WebSphere, JBoss, etc.) from its monitoring purview. Optionally, you can include/exclude more packages/classes for monitoring by editing the exclude.props **file** file in the directory that contains the eg_jtm.jar **file. Once done, t**his will significantly minimize monitoring overheads.
- Additional monitoring power has been built into eG JTM v6. The details are as follows:
  - The outbound HTTP requests per URL can now be tracked and the average responsiveness of the requests can be measured, so that administrators can accurately isolate the URL for which HTTP calls are slow.
  - The requests made by a Java transaction to web services and the response time of these web service requests are now monitored to enable administrators to precisely pinpoint the transaction for which web service interactions are slow.
- Typically, the eG JTM automatically discovers all transactions to a Java/J2EE application by default and reports metrics for each transaction. Most environments however may not be interested in monitoring all their Java transactions and may prefer to focus on only a few. To cater to these requirements, eG JTM v6 allows administrators to choose between automatic transaction (or URL) discovery and manual transaction (or URL) configuration. If automatic transaction discovery is chosen, then administrators should indicate how the discovered transactions need to be grouped – i.e., at which URL segment-level – so that the eG JTM monitors the groups alone (and not the individual transactions); this helps in minimizing monitoring overheads. On the other hand, if manual transaction configuration is chosen, then administrators can configure a comma-separated list of URL patterns that they want the eG JTM to monitor.

1.2.12 Monitoring Web Servers

eG's web server monitoring capabilities have been enhanced in this version. The details are as follows:

- **Monitoring SSL-enabled web sites:** SSL-enabled web sites are typically secured by a private key, public key, or a public-private key pair. Previously however, the HTTP test mapped to a web server did not provide support for SSL-enabled web sites that are additionally secured using a private key. Version 6 however emulates HTTP/S connections to such sites also using the HTTP test and reports metrics. For this purpose, two new parameters have been introduced for the **HTTP** test in version 6 – namely, **Private key file path** and **Password**. If the web page configured for this test is SSL-enabled – i.e., if an HTTPS URL is specified against **URL** parameter of the **HTTP** test – and the contents of this web page can only be accessed using a private key, then the full path to the private key file will have to be provided against **Private key file path** and the password of the private key file should be specified against **Password**. If no such private key protects the contents of the configured **URL**, then set the **Private key file path** and its **Password** to *none*.
- **Checking DNS health:** In version 6, the **HTTP** test additionally reports whether/not the DNS server was able to respond to requests successfully and the time taken for data transfer between the drive

and the host system.

- **Reporting state of threads in IIS application pools:** In version 6, the **Application Pool Workers** test of an IIS web server additionally reports the status of threads in each application pool. The count of threads in the Ready, Running, Standby, Wait, Transition, and Unknown states are reported and the percentage of running threads is revealed. Using this information, administrators can quickly identify application pools that do not have adequate worker threads in the Running state; such pools may not be able to handle request load efficiently, until the pools are resized.
- **Extended Apache web server status**

1.2.13 Enhanced Self-Monitoring Capabilities of eG Enterprise

eG's self-monitoring capabilities have been enhanced in this version to additionally report on the following:

- **Managed infrastructure:** The number and names of components managed, unmanaged, and deleted;
- **Trigger information:** The number of enabled, disabled, and missing triggers in the eG database;
- **Monitoring redundancy:** If eG manager being monitored is part of a redundant setup, then eG's Self Monitor will additionally indicate whether the monitored manager is the primary or secondary manager, reports which manager is currently running and which is not, and reveals the amount of data and number of files stored by the monitored manager for the other manager in the cluster.
- **Checking status of self-monitoring processes:** The self-monitoring and recovery capabilities of the eG manager and agent are governed by two processes – the **egmon** and the **egagentmon** respectively. If these processes fail to run, then the eG manager and agent will not be auto-detect the failure of any of their core components and will hence not be able to recover from the failures. Version 6 therefore checks whether these processes are up and running and promptly alerts administrators to sudden failures/restarts of these processes.
- **Monitoring critical eG manager functions:** eG Enterprise can be optionally configured to run the critical eG manager functions such as email alert management, threshold computation, trending, and database cleanup activities as separate Java processes. If this is done, then each of these processes should run for the corresponding activity to be performed. If the process does not run, then critical eG manager functions will not be performed – eg., email alerts will not be sent out, thresholds will not be computed, etc. Using version 6, administrators can monitor the status of each of these processes continuously and can be alerted if even a single instance of a process is not running.

1.2.14 Monitoring Microsoft RADIUS Servers

Network Policy Server (NPS) allows administrators to create and enforce organization-wide network access policies for client health, connection request authentication, and connection request authorization. Administrators use NPS as a Remote Authentication Dial-In User Service (RADIUS) proxy to forward connection requests to a server running NPS or other RADIUS servers that you configure in remote RADIUS server groups. eG Enterprise v6 targets environments where NPS is used and reports a wealth of performance information pertaining to connections to NPS and its usage.

- Where NPS is used as a proxy, eG Enterprise v6 monitors accounting and authentication requests proxied to the RADIUS server by NPS, and reports how well NPS handled these requests. The rate at which requests were sent, challenged, and responded to by the proxy, requests retransmitted, pending requests, packet drops, and packets with invalid information are captured and reported, so that issues in proxying can be proactively detected and promptly addressed.
- Interactions between the Microsoft RADIUS server and NPS are observed. In the process, the time taken for the NPS to respond to requests from the RADIUS server and the number of requests

pending at any given point in time are reported, so that bottlenecks in connections/request processing can be isolated.

- Where NPS is used as the authentication server, eG Enterprise v6 tracks the authentication and accounting requests received by NPS and measures how well NPS services these requests – were all requests responded to? how quickly were responses sent? Were any requests retransmitted? Were any data packets to NPS dropped? If so, how frequently? Are too many authentication/accounting requests still pending on NPS? eG Enterprise v6 provides accurate answers to these questions and enables help desk plan further action.
- The health validators on NPS are monitored, and for every health validator, eG Enterprise v6 reports the rate at which client communication and component failures occurred, the rate at which non-compliances were noticed per validator, and the time taken by each validator to respond to requests. The failure-prone validators and the relatively slow validators can be isolated.
- Where NPS is used as a proxy, eG Enterprise v6 monitors accounting and authentication requests proxied to the RADIUS server by NPS, and reports how well NPS handled these requests. The rate at which requests were sent, challenged, and responses received via the proxy, requests retransmitted, pending requests, packet drops, and packets with invalid information are captured and reported, so that issues in proxying can proactively detected and promptly addressed.

1.2.15 Monitoring Tuxedo Domain Servers

eG's Tuxedo Monitor reports the following additional metrics in version 6:

- The services running on the Tuxedo server are auto-discovered and the availability and request handling rate of each service is reported, so that breaks in service availability and bottlenecks in request processing can be promptly detected.
- eG Enterprise v6 monitors each server on which a Tuxedo application runs and alerts administrators if any server becomes unavailable, is slow in processing requests, and has hence completed servicing only a small fraction of its request load.
- Requests in queue are monitored and any increase in queue length is brought to the immediate attention of administrators, as this indicates a request processing bottleneck on the Tuxedo server.
- For every client process on the Tuxedo server, eG Enterprise v6 reports the number of clients connecting to the server using that process. Detailed metrics reveal the name of the user connecting using that client, when the user connection was established, and the current status of the connection. This way, active and idle user connections can be identified.

1.2.16 Other Changes/Enhancements

- **Other new component models added in v6:** Specialized monitors have been made available in eG Enterprise v6 monitoring the following:

    - App-V Client and Server
    - Microsoft Lync
    - Microsoft DFS
    - NTP server
    - NGINX web server
    - Veeam Backup
    - Siebel application

- Microsoft CRM
- Microsoft Project
- Sharepoint 2013
- Open VPN Access Server
- IBM Integration Bus
- Postfix mail server

- **Topology discovery optimized:** Earlier, the eG agent was taking too long to perform topology discovery. In v6, the eG agent has been optimized to discover the topology within 15-20 minutes the very first time it runs and take over 6 hours thereafter for rediscovering the topology.
- **Use of signed powershell scripts:** Earlier, the eG agent could not run signed Powershell scripts on target components to pull out metrics. In version 6 however, the eG agent can support both signed and unsigned powershell scripts, but at any given point in time, it can execute only of the aforesaid script types. For this, an **UsedSignedPSScripts** flag has been introduced in the **eg_services.ini** file (in the **<EG_INSTALL_DIR>\manager\config** directory). Set this flag to **yes** to run signed scripts and **no** to run only unsigned scripts.
- **Email alerts from third-party management tools:** In v6, the eG agent can be configured to report the count and contents of email alerts sent out by third-party management tools such as Microsoft SCOM, IBM Tivoli, etc., so that problem conditions identified by such tools can be viewed, analyzed, and interpreted from the single, central eG monitoring console. To read these mail alerts, the eG agent has to be configured with the name and access credentials of the POP3/IMAP account where the mails are stored.
- **Hardware monitoring for Windows Servers:** Earlier versions of eG Enterprise could monitor the hardware on which Windows operating systems were hosted. To do so, eG Enterprise required administrators to have a hardware agent (SNMP-based) from the hardware vendor - Dell, HP, or IBM - running on the server. In v6, eG agents can monitor the hardware status of physical servers running Microsoft Windows OS without requiring any third party agent to be present. Additionally, eG agents also integrate with integrated lights out (ILO) hardware management solutions from HP and IBM for hardware monitoring.
- **Custom selection of network adapter when monitoring Client Desktop component:** Previously, the *Client TCP* test and the *Client Service* test mapped to a Client Desktop component, could be configured to automatically discover all network adapters supported by the monitored Client Desktop. Of the auto-discovered adapters, the eG Enterprise system picked one adapter at random and used this for packet capture. In this version however, the auto-discovered adapters are made available for selection in the test configuration page of the *Client Service* and *Client TCP* tests, so that the administrator can pick and choose the adapter that the tests should use for packet capture.

1.3 Other Changes

- IE 10 is supported as a browser from v6.
- Starting from v6, the eG database can be configured on Oracle 12c, Microsoft SQL server 2014, and a Microsoft SQL Server on Azure.
- **Support for NTLM v2:** eG Enterprise v6 can use SQL database configured with NTLM authentication as backend.
- **Alarm retention on restart:** Previously, if the eG manager was restarted, the state of all managed infrastructure elements was by default re-initialized to **Normal** after the restart. Changes occurred in the state only later, when agents started reporting metrics to the manager and problems were

noticed. In this process, the eG manager ended up generating alarms once again for all those problems that had been open at the time of the restart. As a result, users were alerted twice to the same problem – once before the restart and once after it. To avoid this confusion, the eG manager has now been configured to retain the state of managed elements even after a restart. This way alarms and state are preserved even after a manager restart.

- **Heartbeat option:** If an agent stops running or reporting metrics, the eG manager detects the status change and intimates administrators of it. In earlier versions, the eG manager would detect the status change of an agent only when all of the tests that the agent was supposed to execute did not return any metrics. Also, if the eG manager was busy processing requests (e.g., in an environment with several hundred agents), it would take the manager a longer period of time to determine that an agent was not running. To ensure that administrators are notified of such anomalies promptly, the eG manager v6 can be optionally configured with an explicit heartbeat capability. If this capability is enabled, at configured frequencies (default: 5 minutes) the eG manager checks whether each agent is sending heartbeats in a timely manner or not. If an agent does not send a heartbeat message at the configured frequency, the manager marks the agent as not running and alerts administrators. With this capability, the eG manager can now detect and report on agent failures immediately.
- **Faster installation of the eG manager:** In earlier versions, configuration information was maintained in the eG manager both in its database and in configuration files. When the manager was installed and during the normal operation of the manager, both of these configurations had to be updated and kept synchronized. I v6, the duplication of eG manager configurations in the database and configuration files has been removed. This allows the eG manager installation to be faster.
- **Licensing changes:** Previous versions of eG Enterprise offered administrators two licensing models. The default licensing model is per operating system monitored. For Citrix, virtual desktops and Microsoft Terminal/RDS environments, eG Enterprise also includes a named user licensing option. In this model, the number of unique users accessing Citrix/VDI/Microsoft RDS is tracked and servers hosting these applications can be monitored using named user licenses (instead of operating system based licenses). The named user licensing model is ideal for virtualized environments, where the number of users per server is low. In some virtualized environments however, a large number of users could be accessing the servers every day, but only a small subset of this user population may access the servers concurrently – i.e., at the same time. In such a scenario, with the named user licensing model here, one would have to obtain licenses for all the users who log into their desktops each day; this could be expensive. To provide a cost-effective licensing option for such virtualized environments, eG Enterprise v6 now supports a **concurrent user licensing** model. In this model, the maximum number of concurrent users logged into Citrix/Microsoft RDS/virtual desktops is tracked and compared with the licensed limit. Customers will need to ensure that they have sufficient licenses procured to handle their current workload. Customers can also use a mix of server-based and one of the user-based licenses (either named user license or concurrent user license).
- **Quick updation of threshold changes:** In v6, unlike the previous versions, administrators will no longer have to wait for 24 hours for threshold changes to take effect – even if these changes apply to a group of servers or to the default configurations across the entire managed infrastructure. Whenever threshold settings are modified, the changes will take effect within the next hour itself.
- **SNMP monitoring over TCP is now supported:** In some environments, to ensure reliability, SNMP is configured over the reliable TCP protocol. eG Enterprise v6 now supports monitoring of network devices that are configured to respond to SNMP over TCP. All the out of the box SNMP monitoring and alerting capabilities can be configured this way. Custom SNMP monitors too can be configured to use TCP for SNMP communication instead of the UDP protocol.
- **User session timeout:** Previously, user sessions to the eG management console timed out after

30 minutes of inactivity, by default. This timeout setting is now configurable.

- **IPv6 support:** From v6, eG Enterprise supports both IPV4 and IPV6 addresses. IPV6 addresses can now be used during eG manager configuration, when adding components using the eG administrative interface, and during component discovery.
- **Alarm escalation-related changes:** Previously, the start time of the alarm was considered for escalation. Now, eG Enterprise v6 considers the time when an alarm was sent for each user when performing escalation.
- **Alarm acknowledgement and escalation:** In v6, the eG Enterprise system can be configured to not escalate alarms that have been acknowledged by a user. This option is provided because the acknowledgement of an alarm in itself is a declaration of a user's intention to resolve the issue raised, and does not warrant an escalation.

1.4 Bug Fixes/Optimizations

- In non-English OS, the eG manager and agent did not start sometimes. This issue has now been resolved.
- Previously, blocked threads inside the eG manager were causing frequent restarts/slowdowns of the manager. These threads have been eliminated in this version.
- On Windows 2012 managers, critical manager functions such as database cleanup, trending, email alerting, and threshold computation earlier failed to stop when the eG manager stopped. This issue has now been resolved.
- Sometimes, the threshold manager took a long time to run. This was because, the threshold manager engaged in unnecessary cleanup operations of last measure values from the eG database. In v6, this job has been transferred to the database cleanup process. Consequently, the threshold manager runs quickly in v6.
- Before v6, if the eG database was unavailable/shutdown, critical eG manager functions such as database cleanup, trend computation, threshold computation, and email alerting took 24 hours to complete. This was because of invalid connections not being removed from the free connections list. This issue has been resolved in v6 by tracking invalid connections and removing such connections from the free connections list.
- In earlier versions, the eG manager discovery process could not discover the Printer. This is no longer the case.
- Previously, if **Agent Discovery** was turned on, the eG agent on a Windows server auto-discovered all the listening ports on that server and accessed each port to discover the types of applications running on that server. In the process, meaningless errors were captured by the Windows event logs. To avoid these unnecessary errors, the eG agent v6 has been re-engineered to perform component discovery using only the default server ports and not the listening ports.
- Prior to v6, topology discovery did not work if the IIS web server spawned worker processes and those processes used the eG backend. This issue has now been resolved.
- Earlier, if a test reported metrics for descriptors that were suffixed by '++', such descriptors and their metrics were inserted in the eG database, but the state of such descriptors was not considered during state computation of the corresponding test/layer/component. This issue has now been resolved.
- Before v6, if the eG agent was configured to communicate with the eG manager via a proxy, then all tests run by that eG agent, which required the eG agent to interact with a web services API (eg., HTTP test, VMware-related tests, etc.), did not report metrics. This was because, these tests used

the proxy to communicate with the web services API. This issue no longer occurs.

- Prior to v6, if empty lines were found in the detailed diagnosis data files, the agent did not send these files to the eG manager. As a result, these files continued to grow to very large sizes on the agent host and caused the agent CPU usage to increase considerably. This issue has now been resolved.
- In older versions, if a user with *Limited admin rights* logged into the eG admin interface, he/she was able to view and modify all threshold groups configured and not just the ones that pertain to components explicitly assigned to him/her. This issue has now been resolved.
- In versions prior to v6, if a user who belonged to an Active Directory group logged into the eG monitoring console and recorded fix feedback for a problem, then, when viewing the history of fixes for that problem, the AD group name was displayed against 'Problem fixed by' – not the user name. This is no longer the case.
- In the older versions, if an **Admin** user used the **Search** option in the **MANAGE/UNMANAGE** page to pick a component, and then unmanaged it using the same page, the configuration of all the managed components of that type also got lost. This increased the number of unconfigured tests in the **Admin Home** page**.** This issue has now been resolved.
- Previously, if the managed infrastructure contained segments with many component groups, the **LIST OF SEGMENTS** page took a long time to load.
- Before v6, when creating segments from a zone that contains sub-zones, the components/groups inside sub-zones were not available for segment creation. This issue has now been fixed.
- Prior to v6, if detailed diagnosis was configured to be sent with email alerts to users, and the eG backend was Microsoft SQL server, then delays were noticed in the delivery of email alerts. This was owing to delays in retrieving detailed diagnostics from the eG backend. To reduce the time taken for retrieving the detailed diagnosis from the SQL database, the query executed for this purpose has now been optimized.
- In versions prior to v6, if email alerts were delayed, the eG manager sent old detailed diagnosis with the email alerts. This has been avoided in version 6. Now, if the time stamp of the problem measure and the detailed diagnosis information do not match, then the email alert will only carry a message stating DD has been delayed. On the other hand, if the time stamps match, then the actual detailed diagnosis information is sent with the email alert.
- Previously, if a VM used a cluster IP address, then inside view tests on that VM failed. In v6 however, if the eG agent discovers multiple IP addresses on a VM and one of which is the cluster IP address, the eG agent automatically disregards that address and runs the inside view tests using one of the other IP addresses on that VM.
- Earlier, the remote agent monitoring virtual hosts and VMs and the eG VM Agent deployed on Windows VMs consumed a lot of CPU resources on the Windows VMs when running inside view tests. The remote agent and the eG VM Agent have now been optimized to consume lesser CPU resources.
- In older versions, VMware view monitoring consumed CPU resources excessively. In v6, tests executed on the *VMware Horizon View* component have been optimized to reduce CPU consumption.
- Prior to v6, the eG remote agent could not run the **User Profile Management – VM** test on 64-bit Windows VMs. This was because, the remote agent ran 32-bit executables on the 64-bit VM to collect metrics. This issue has been resolved in v6. In this version, the eG remote agent embeds the intelligence to discover the VM architecture (whether 32-bit or 64-bit) and execute executables on that VM accordingly.
- Previously, even if a VM was explicitly excluded from monitoring using the **EXCLUDE VMS** parameter, the eG agent periodically accessed that VM to discover its IP address. Starting from v6 however, the eG agent will not perform even IP address checks for VMs that have been excluded

from monitoring.

- Prior to v6, if a large number of VMs operated on a monitored *RHEV Hypervisor*, the 'outside view' and 'inside view' tests of those VMs did not run at configured intervals – i.e., significant delays were also noticed in the execution of these tests. This issue has now been resolved.
- Sometimes, the **Hyper-V VMs** test of a *Microsoft Hyper-V* server wrongly reported that the *Virtual CPU usage* of a VM was over 100%. This issue has now been fixed.
- Previously, if physical NICs were not assigned to the vSwitch of a VMware vSphere/ESX server, the **Network – ESX** test of that server failed. This issue has now been fixed.
- In earlier versions, if a Citrix Provisioning server farm was not at all configured with a License server, the **PVS Farm** test wrongly reported *License server availability* as 0. This does not happen any more.
- Earlier, the **Citrix Users** test (mapped to a XenApp server and a XenDesktop Apps server) was not reporting metrics correctly on a Windows server with a non-English locale. This bug has now been fixed.
- Prior to v6, the **User Profile Management** test of a Citrix XenApp server reported all latency metrics in milliseconds and not seconds. Now, this test reports latencies in seconds only.
- In older versions, the **Ticket Request Status** test of Citrix XenApp was repeatedly reporting ticket request status as failed. In v6 however, the correct ticket request status is reported.
- Earlier, when computing the size of a Citrix user's profile, the **User Profile** test did not consider the size of the *AppData* folder. This test has now been changed to include the *AppData* folder's size in its profile computations.
- Sometimes, the **Microsoft SQL Net** test reported abnormally high response time values for a database. This happened if the query executed by the test took too long to report measures. To avoid this, this test has now been hard-coded to wait for 15 seconds for a response from the database. Beyond this time frame, the test will automatically time out and response time will not be reported by that test.
- Before v6, the **Oracle Processes** test did not allow users to configure process patterns for monitoring. This test now supports user-configured process patterns.
- Previously, the **Oracle SQL Network** test and the **Oracle Service** test mapped to Oracle database servers were taking too long to execute and report performance results. In this version, these tests have been optimized to run quickly and report metrics soon.
- In earlier versions, SSL-enabled Microsoft SQL servers could not be monitored properly by the eG agent. Starting from v6 however, such servers can be monitored by the eG agent without a glitch.
- In older versions, the **Oracle Processes** test, the **Oracle Client Connections** test, and the **Processes** test mapped to an Oracle database server did not report metrics properly. This issue has now been resolved.
- Before v6, many query-based tests did not report metrics for Microsoft SQL Server 2000. These tests now work.
- In versions prior to v6, the **SQL Datafile Activity** test incorrectly reported the value *0* for the *Size on disk* measure of each datafile. Valid values are now reported for this measure.
- Earlier, the **SQL Error Log** test reported the value 0 for all its metrics if no deadlocks were detected in the error log. This issue has now been fixed.
- In earlier versions, the **SQL Long Running Queries** test and the **SQL Blockers** test sometimes consumed too much CPU resources on the agent host. These tests have now been optimized to reduce CPU consumption.
- Before v6, the **Oracle RAC Temp Tablespaces** test was expressing the *Free percentage* measure as a ratio of *Free space* over *Total space.* While this computation took into account the space allocated

to the temp tablespace, it disregarded the maximum extent upto which the space in the temp tablespace can grow – i.e., the *Max space* of the temp tablespace. This is why, in v6, this test has been changed to express the *Free percentage* as a ratio of *Free space* over *Max space*.

- Previously, the **SQL Engine** test reported *Cpu usage* of over 100% on multi-processor systems. This issue has now been resolved.
- Earlier, the **Uptime** test did not work properly for Windows servers that were streamed from a gold image using a Citrix Provisioning server. This issue has now been fixed.
- Prior to v6, if an auto-correction script was associated with the **Windows Services** test, and this test was configured to monitor Citrix services with names containing white spaces, auto-correction did not work. This is not the case anymore.
- Earlier, the **TCP** test stopped reporting metrics and switched to the *Unknown* state on Windows 2008/2012 servers with a Danish locale. This issue has been fixed in v6.
- In older versions, the **Disk Activity** test did not report metrics for the *Data read rate from disk* and *Data write rate to disk* measures. This is not the case any longer.
- In versions prior to v6, the **AX Portal Statistics** test of a Microsoft Dynamics AX server did not work on Windows 2012. This is not the case any longer.
- In earlier versions, the **Cisco Interfaces** test of a *Cisco router* was by default configured to rediscover the network interfaces supported by the router once every hour. Because of this default configuration, new network interfaces and changes in the names of existing interfaces were reflected in the eG monitoring console only after an hour. To ensure that the eG Enterprise system is updated with such additions/changes quickly, the **Cisco Interfaces** test has now been hard-coded to perform interface discovery once every 5 minutes.
- Before v6, even if one measure of the **Host Processes** test, **Host Storage** test, **Host System** test, **Disk Status – NetSnmp** test, and **Memory Status** – **NetSnmp** test failed to report metrics, the entire test failed. This is no longer the case.
- Previously, the descriptors of the **Pix Hardware** test – i.e., the hardware units of a Cisco Pix firewall device - was hard-coded in the eG Enterprise system. If these standard (i.e., hard-coded) hardware units were not supported in the target environment, this test failed to report metrics. In version 6 therefore, this test has been changed to dynamically discover the hardware units the target firewall device supports and report the status of each such unit.
- In older versions the **WebSphere MQ Queue Details** test consumed CPU and memory resources on the eG agent host excessively, if one/more message queues that were monitored contained a large number of messages. To avoid this, this test has now been changed to include an additional **MESSAGE LIMIT** parameter. If a valid numeric value is provided against this parameter, then this test will monitor only those queues with a message count that is equal to or less than the configured value.
- Earlier, if a JVM thread was blocked for a very short duration, then the *Blocked time* measure displayed as part of the detailed diagnostics of the **JVM Threads** test (mapped to Java-based applications/J2EE-enabled application servers) returned values in exponential format. This issue has been now been resolved.
- Prior to v6, when monitoring an IIS web server on Windows 2008 R2, the **IIS Application Pools** test did not work. This issue has now been resolved.
- In versions prior to v6, the **PVM Host Disk Details** test mapped to a Marathon EverRun component was not reporting metrics properly for PVMs with multiple disks. In v6 however, this test can monitor PVMs with single and multiple disks well.
- In previous versions, when multiple NetApp USD storage devices were managed using the same remote agent, the same set of metrics were reported for all the NetApp devices. This issue has now

been fixed.

1.5 Known Issues

- If the **Windows Services** test or the **Processes** test is configured with an auto-corrective script, then this script will run and automatically correct issues only if the state of the associated measure changes. For instance, if the script is associated with the *Processing running* measure of the **Processes** test, then, the test will run only if the state of the test changes from Normal to Critical/Major/Minor or vice-versa.
- In this version, alarm acknowledgements will not be sent with email alerts, even if configured.
- Escalation mails will continue to be sent out for acknowledged alarms, even if the eG manager has been configured to not escalate such alarms.
- Trouble ticket integration, if configured via email, will not work for an eG manager in a redundant cluster.
- Service groups cannot be configured if the browser used for accessing the eG administrative interface is the Internet Explorer.
- Graphs will not get loaded properly if the browser used for accessing the eG monitoring/reporting console is the Internet Explorer.
- User-specific time zones will not be supported in some cases.
- The Remote Agent Controller (RAC) executable that is bundled with the Windows eG manager may not work properly. It is hence recommended that you download the RAC installable that has been uploaded to our site for remotely installing and controlling eG agents.
- If the configuration test, **Drives**, is executed on a Linux system using a remote agent on Windows, many shell scripts are unnecessarily spawned on the target Linux system.

1.6 Conclusion

This document provides a detailed list of the new features, enhancements, and bug fixes that have gone into v6 of the eG Enterprise Suite.

Feel free to contact eG Support in your region for technical issues or info@eginnovations.com for general enquiries related to the product.