Make IT Service Monitoring Simple & Proactive with AIOps Powered Intelligent Thresholding & Alerting

An eG Innovations Technical White Paper

www.eginnovations.com

❖ Introduction

IT managers often complain about two main types of problems with monitoring and management tools:

- Firstly, after they install the software, they start to receive many "false" alerts. A false alert refers to a situation in which the monitoring tool indicates a problem, but the IT manager determines that there is no real problem in the network. Thousands of alerts can result in distracting IT administrators, preventing them from focusing on the genuine issues that can impact IT service quality. This is commonly referred to as 'noise'.

- Secondly, to avoid false alerts, IT managers must define threshold values for the different metrics collected by the monitoring tool. A threshold is a limit set in the monitoring tool for the metric, so that if a metric crosses this value, an alert is raised. In a large enterprise, a monitoring tool that provides visibility into the different network, server, and application tiers can collect millions of metrics. Having to set thresholds manually for every single metric is a very time-consuming, monotonous exercise. As a result, enterprises end up spending a lot of time and money having consultants tune thresholds manually or implementing costly-to-maintain bespoke automation scripting.

What administrators really need from a monitoring and management system is the ability to make thresholds simple to configure and accurate to enforce, so there are few false alerts (and a minimum amount of 'noise').

The eG Enterprise IT service monitoring solution from eG Innovations addresses these key requirements of IT managers using a combination of automated baselined dynamic thresholding and intelligent alerting driven by our powerful patented AIOps platform. By doing this without requiring a lot of manual intervention, eG Enterprise makes it simple to implement IT service monitoring in an enterprise, and yet deliver proactive alerts that are essential for ensuring that the service level expectations of users are met.

Reducing False Positives in IT Application & Infrastructure Monitoring

https://www.techtarget.com/searchnetworing/tip/Reducing-false-positives-in-network-monitoring

"The million-dollar question is: How can you reduce false positives for counters that tend to fluctuate a lot? I have seen some administrators try to reduce the sampling frequency in an effort to reduce false positives. Indeed, this technique may reduce false positives, but it still has the same result. Counters that fluctuate a lot will still produce false positive alerts."

—— **Brien Posey**
SearchNetworking

In this white paper we will cover the architectural qualities associated with eG Enterprise's thresholding methodologies to automate and optimize thresholding and alerting. This white paper will enable readers to evaluate and compare the thresholding capabilities of monitoring solutions for a range of use cases including APM (Application Performance Monitoring), BTM (Business Transaction Monitoring), Digital Workspace monitoring (Citrix / VMware) and Cloud Monitoring (Azure Monitor, Amazon CloudWatch). Capabilities discussed will cover:

- Static vs. Dynamic Thresholds
- Combining Static and Dynamic Thresholds to increase the stability and reliability of dynamic thresholds
- Multi-level thresholds – the ability to associate multiple thresholds and alert severities with a single metric or property
- Threshold sensitivity – leniency (associated with statistical quality control) and granular temporal aggregation sensitivity (ensuring, if appropriate, that a momentary spike in a metric does not trigger alerts or alarms)
- Associating thresholds with aggregated metrics and grouped components
- Automated deployment and configuration including within auto-scaling, Kubernetes and microservice architecture environments
- Alert correlation and filtering
- Automated handling of alerts triggered by thresholds and ingestion into ITSM service and help desk systems
- Help desk and administrator insight and visibility into thresholds and associated alerts
- Ensuring alert thresholds are suppressed during controlled and understood change scenarios such as maintenance windows and rolling out from pre-production to production

The information in this white paper should assist anyone looking to implement a proactive alerting and monitoring strategy incorporating dynamic or static alerting thresholds or by combining both the thresholding approaches.

❖ Defining Static Thresholds for Metrics

Thresholds are upper and lower bounds that determine whether a metric is performing to expectation or not. Every time the actual value of the metric falls outside the prescribed limits, the monitoring system detects an abnormality.

Depending on the metric being collected, upper bounds are appropriate for some metrics, while lower bounds are appropriate for others. For example, a lower bound or minimum threshold is applicable for the free disk space metric. If the value drops below the lower bound, an alert will be generated. In contrast, an upper bound or maximum threshold is applicable for the CPU usage metric of a server. If the value exceeds the upper bound, an alert will be generated. In some cases, both lower and upper bounds may be appropriate. For instance, if the number of users accessing a server is much higher or much lower than normal, it could be an indicator of a problem.

For many metrics, thresholds can be set statically. For instance, based on the service level expectations and agreements, IT managers can set thresholds for metrics such as network availability and latency. Application availability and response time can also be handled in the same manner. For example, availability should be 100% whenever the metric is measured. If not, a violation should be detected. Likewise, a network latency of several seconds is usually an indicator of a problem, no matter what time of day the measurement is made at.

Thresholds can also be set based on industry standard best practices. For example, a rule of thumb when tuning an Oracle database server is that the database dictionary cache hit ratio should be 90% or more. If the hit ratio falls below this value, it indicates a need to tune the database server. This is another example where a threshold is set statically, without considering the time of day when the measurement is being made. eG Enterprise includes pre-specified threshold values for many metrics based on industry standard best practices.

IT managers also want to set different threshold levels to map to different levels of severity of problems. For example, when the space usage of a disk drive is close to 90%, a minor alert is desirable. When the metric's value crosses 95%, the alert severity should change to major, and when the value crosses 99%, the IT manager would need to receive a critical alert. To support such a requirement, eG Enterprise allows administrators to set different thresholds levels for a metric (see Figure 1). Based on the value of the metric and the threshold level it crosses, an alert with the appropriate priority is generated. Multiple levels of threshold settings provides support for alert escalation - a minor alert is generated when a problem starts, and the severity is upgraded when the problem becomes worse.

Stop Monitoring Tools from Crying Wolf

"In my opinion, a poorly tuned monitoring server is as bad or worse than no server monitoring at all. At least with no monitoring you are less likely to become complacent. If you don't have a car alarm and live in a bad neighborhood, you'll probably be more careful to put away valuables and lock your doors. But if you have a car alarm that goes off every time another car drives by, you will naturally start to ignore it over time."

— Kyle Rankin

Data Center Systems Management, Tech Target

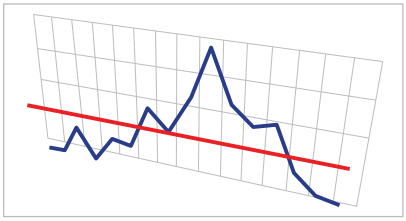

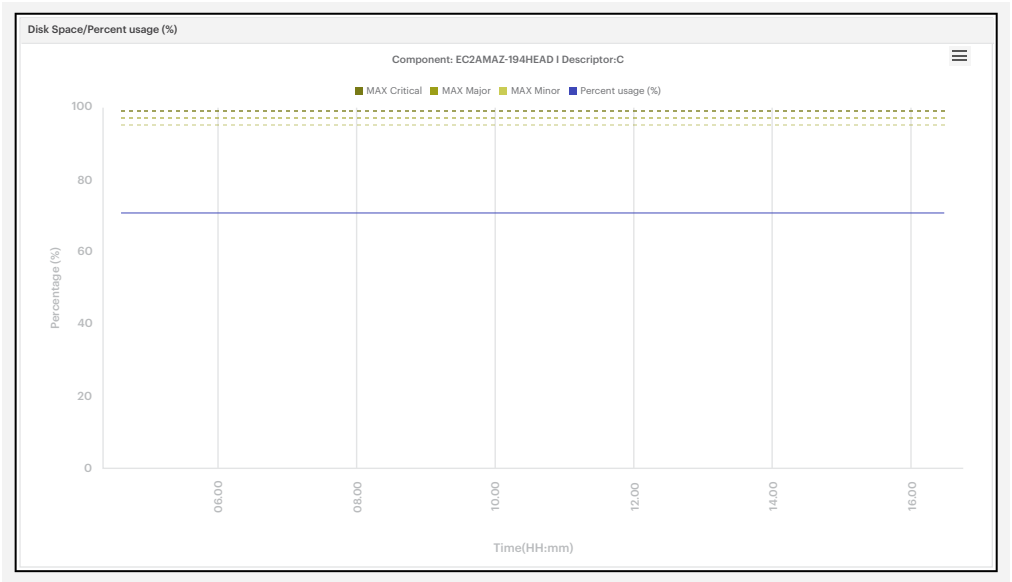

Figure 1: Multiple levels of static thresholds are set for a metric and the priority of an alert generated is based on the max/min threshold value that is crossed. The blue line is the actual metric, the yellow lines are different threshold settings.

❖ Defining Automatic, Self-Adjusting Dynamic Thresholds for Metrics

In infrastructures where a metric varies with time, a static absolute threshold value cannot serve as a reliable basis for judging performance. For example, consider a web server hosting a web site. The number of TCP connections to the web site could be high on one particular day and low on another. Similarly, it could be high during the working hours and low during the nights. In such situations where measurement values change with the time of the day, it is exceedingly difficult to set accurate maximum and minimum limits manually. In such cases, the threshold value for this metric ideally would be time variant.

Even when a metric is not time variant, its value may change from one server to another. For example, a high-end datacenter server may be able to handle hundreds of users, whereas a low-end standard server may be able to handle only a few tens of users.

Requiring IT teams to manually configure thresholds has many drawbacks:

- A monitoring tool that provides visibility into the different network, server, and application tiers can collect thousands of metrics even in a small/medium IT environment. Having to set thresholds manually for every single metric is a very time-consuming, monotonous exercise. This is why many enterprises end up spending a lot of time and money having consultants tune thresholds manually.
- If thresholds are to be configured based on server sizing, this requires experts to be involved.
- Modern IT environments are very dynamic. Servers and VMs can be powered on and off on-demand. Requiring manual intervention to configure thresholds is not ideal.

If the wrong thresholds are set, IT teams may either miss significant problems, or get too many issues reported when no problem exists (false positives).

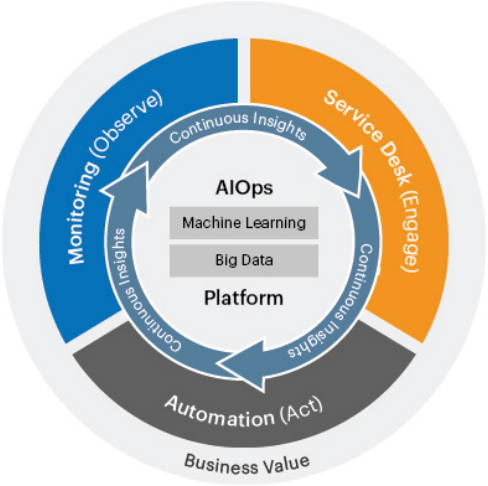

Most enterprise and cloud monitoring solutions acknowledge the limitations of static thresholds by implementing machine learning technology and including an AIOps (Artificial Intelligence for IT Operations) engine capable of learning about the normal behavior of systems over multiple timeframes. This means that normal is understood with the context of time-of-day, day-of-the-week, monthly and seasonal variations.

Once the real usage of a system has been established, this type of auto-baselining learns what is normal and dynamic thresholds can be applied, e.g., raise an alert if the bandwidth used by a server exceeds 200% of normal usage.

A dynamic system should learn that a high CPU load during a nightly backup is normal, but that 80% CPU utilization on the same server on a mid-week morning is abnormal. When such tuning is automatic, an IT monitoring strategy can include thousands of thresholds, even ones that change over time to follow business cycles or vary between similar components e.g., web servers serving different applications.

Additionally, AIOps driven dynamic thresholding will self-tune to learn the normal behaviors of individual components such as servers that may be supporting very differing workloads and business applications with vastly varying performance profiles and resource needs.

The benefits of dynamic thresholding leveraging AIOps (Artificial Intelligence for IT Operations) technologies is that millions of data points can be ingested and automatically adjusted to maintain thousands & thousands of metric thresholds providing coverage way beyond what a human IT team could maintain or manage.

❖ Problems with Dynamic Thresholds

Dynamic thresholds though do have some problems, TechTarget editor *Alistair Cooke covers these in an excellent article* in which he summarizes that - "Dynamic thresholds are not as intelligent as people". Some of the problems include:

- Dynamic thresholding can become confused when normal cyclic patterns are not adhered to. For e.g., a public holiday means only a small number of staff logon on a weekday.
- Dynamic monitoring tools deployed in a broken or poorly performing IT environment can learn that state as normal and even start to send alerts due to it getting better.
- Dynamic systems are also inclined to view things that get broken for a while as the new normal. If a storage array slowly gets overloaded and unresponsive, the dynamic threshold monitoring system will register the overloaded state as the new normal.
- Systems deployed in test or pre-production may have little realistic load, e.g., VMs are not accessed by real users using applications so dynamic thresholds may benchmark their normal usage as 3% CPU.

eG Enterprise is one of very few solutions that implements dynamic-static combination thresholds to overcome the limitations of dynamic thresholds and reduce time-wasting false positive alerts.

❖ Dynamic-Static Combination Thresholds

Dynamic-static (sometimes called auto-static) thresholds combine dynamic and static thresholds to override dynamic thresholds when they would not make sense.

To understand the limitations of just using dynamic thresholds, consider an application server in a staging environment. Typically, there is no load on the application server and the dynamic threshold is set accordingly. When someone logs in, the threshold will be breached, and an alert may be raised by the system. This is a false alert because one user logging in does not signify a situation of interest to an IT manager. This scenario shows that while dynamic thresholding reduces the effort involved in configuring the monitoring tool (because IT managers do not have to configure thresholds for every metric and server), it does not eliminate false alerts. This scenario is commonly encountered with tools like Azure Monitor that only provide IT admins a choice of static or dynamic thresholds.

To avoid this problem, eG Enterprise allows IT managers to use a combination of static and automatic/dynamic thresholds. A static threshold applied along with a dynamic threshold provides a realistic boundary that must be crossed before an alert is to be triggered. An IT manager can now configure an absolute maximum and a dynamic maximum threshold for a metric.

eG Enterprise compares the actual measurement value with the higher of the two maximum thresholds and generates an alert only when the higher threshold is violated. In the example of the staging application server, the IT manager can set a static limit of say 10 sessions. Once this is done, only if the actual load exceeds 10 current sessions will an alert be generated, even if the auto-computed threshold is less than 10. If the auto-computed threshold is greater than 10, this value is used as the actual threshold.

The Limitations of Dynamic Thresholds – A Microsoft System Center Example

"In response to "Why do I get an alert when the terminal services active sessions metric exceeds 1.333333?":

There isn't a fix as such as it is behaving the way it should. It is just that self-tuning thresholds are a pain in the backside and should be avoided, especially when low values are returned as a small value change can represent a large percentage change. They are great in theory but just don't work very well in practice.

Any time a customer is not happy with the results of a self-tuning threshold monitor – they should simply create a static threshold monitor. This is very basic and provides the best solution."

—— Graham Davies
Microsoft System Center Forum

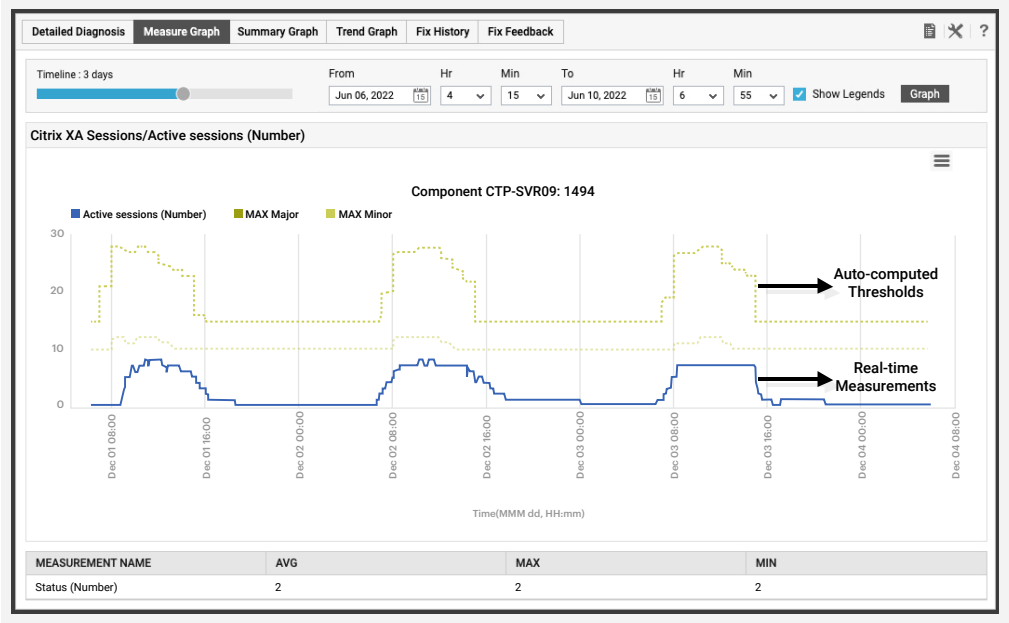

Figure 2: A performance graph showing the number of user sessions and the auto-computed, dynamic-static thresholds used for alerting.

Consider the example in Figure 2. An auto-static combination threshold is applied to this metric. In the morning hours, a static threshold is applied because the dynamic threshold is lower. The static value ensures that alerts are not generated as long as the number of sessions stays below 10. During the day (8am onwards), the automatic threshold takes over. The blue line in the figure denotes the metric's value over time. The yellow lines represent the upper threshold values. Notice that from 4pm to 8am, the threshold is static – with the minor value at 10 sessions and the major value at 15 sessions. Since the automatically computed value is less than both thresholds, the statically set threshold values apply in this case.

As in the case with the maximum thresholds, if a static minimum and an automatic minimum threshold are specified, then eG Enterprise will generate alarms only when the current value falls below the lower of the two threshold settings.

❖ Providing Leniency for Thresholds using a Sensitivity Slider

Even when dynamic thresholds are set automatically, an IT manager may want to choose a leniency factor for the thresholds. For example, an IT manager may want to allow for a 10% deviation from the norm. To accommodate such requests, eG Enterprise allows administrators to set a "sensitivity slider" for automatic thresholds.

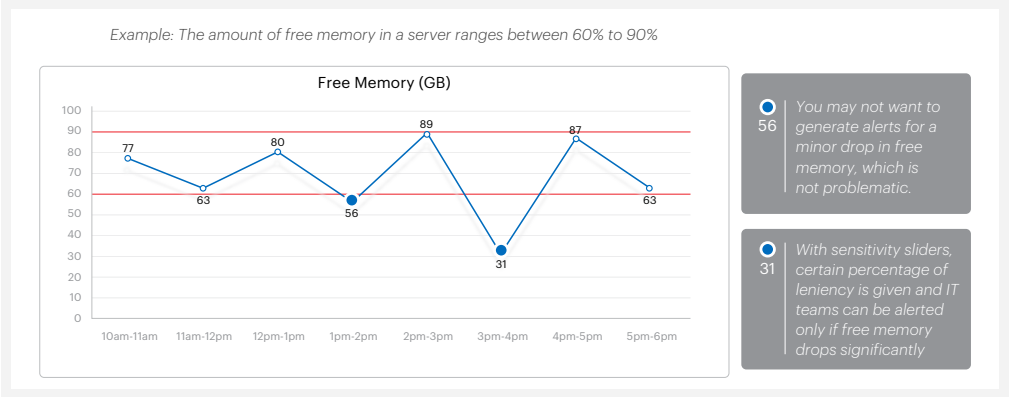

Figure 3: Sensitivity slider can be used to provide a degree of leniency when configuring thresholds for a metric

This slider should be specified as a multiple of the threshold value computed using statistical quality control (sqc). For example, consider the case of the "Free memory" measure, which is an indicator of the amount of free memory available on a server. Assume that on one of the managed servers, the free memory is known to decrease consistently and then grow back up (e.g., the operating system frees memory periodically). In such a scenario, the free memory threshold will be violated often (since the value decreases consistently), and this will result in a number of false alerts. In such a situation, the eG administrator can set the threshold to be a multiple of sqc - for example, if the minimum threshold is set to 0.7*sqc, it implies that the administrator has introduced a 30% leniency. That is, alerts are generated only if the free memory is 30% lower than what is the normal value. This capability allows administrators to fine-tune eG Enterprise's relative thresholding capability to suit their specific requirements.

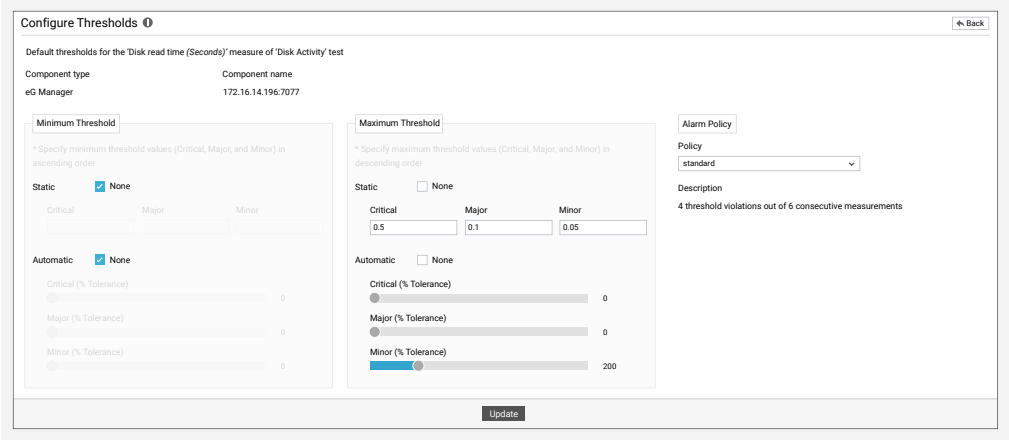

Figure 4: How eG Enterprise implements sensitivity sliders. Differing levels of tolerance can be set for different threshold levels

❖ Multiple Levels of Thresholds

As in the case of static threshold, multiple levels of dynamic-static thresholds can be set for any metric. Multiple levels of thresholds are also supported when setting automatic dynamic thresholds – for example, a minor alert can be generated when a 10% leniency limit is crossed, a major alert generated

when a 30% limit is crossed, and a critical alert generated when a 50% limit is crossed. Multiple levels of threshold settings allow proactive alarms to be generated when a metric is slightly out of conformance, and a severe alarm to be generated when the problem worsens.

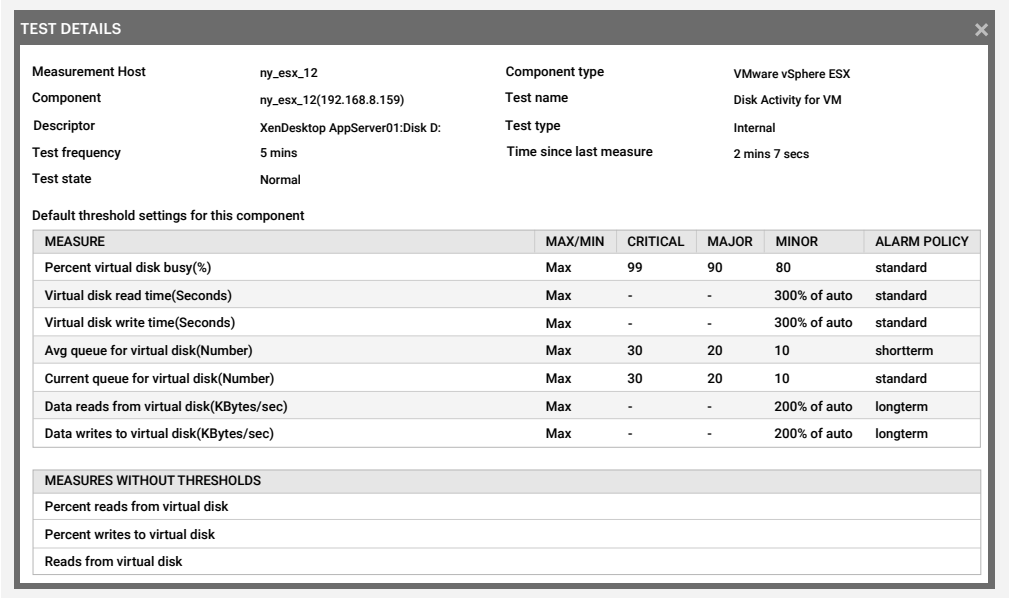

TEST DETAILS

Measurement Host	ny_esx_12	Component type	VMware vSphere ESX
Component	ny_esx_12(192.168.8.159)	Test name	Disk Activity for VM
Descriptor	XenDesktop AppServer01:Disk D:	Test type	Internal
Test frequency	5 mins	Time since last measure	2 mins 7 secs
Test state	Normal		

Default threshold settings for this component

MEASURE	MAX/MIN	CRITICAL	MAJOR	MINOR	ALARM POLICY
Percent virtual disk busy(%)	Max	99	90	80	standard
Virtual disk read time(Seconds)	Max	-	-	300% of auto	standard
Virtual disk write time(Seconds)	Max	-	-	300% of auto	standard
Avg queue for virtual disk(Number)	Max	30	20	10	shortterm
Current queue for virtual disk(Number)	Max	30	20	10	standard
Data reads from virtual disk(KBytes/sec)	Max	-	-	200% of auto	longterm
Data writes to virtual disk(KBytes/sec)	Max	-	-	200% of auto	longterm

MEASURES WITHOUT THRESHOLDS
Percent reads from virtual disk
Percent writes to virtual disk
Reads from virtual disk

Figure 5: Multi-level thresholds are set for many metrics associated with a VM on a VMware vSphere ESXi server which escalate alarms automatically – in this case "VM CPU Ready" will trigger a minor alert when the value exceeds 10%, which will be changed to a major alert at 20% and a critical alert at 40% .

❖ Flexible Alarm Policies: Granular Temporal Aggregation Sensitivity

Threshold configurations help determine when the state of a metric changes, but a threshold violation might not necessarily indicate a problem condition worthy of being reported to help desk. In other words, a single threshold violation might not always be reason enough for an alarm to be generated by the monitoring tool.

Set alerts for a length of time, not just usage

4 monitoring and alerting best practices for IT ops (techtarget.com)

A core best practice is to always include a duration -- not just a level -- of use for performance counters. Most applications spike resource utilization as they run, and alerts triggered by these momentary spikes will flood inboxes.

— Brian Kirsch

Milwaukee Area Technical College in TechTarget

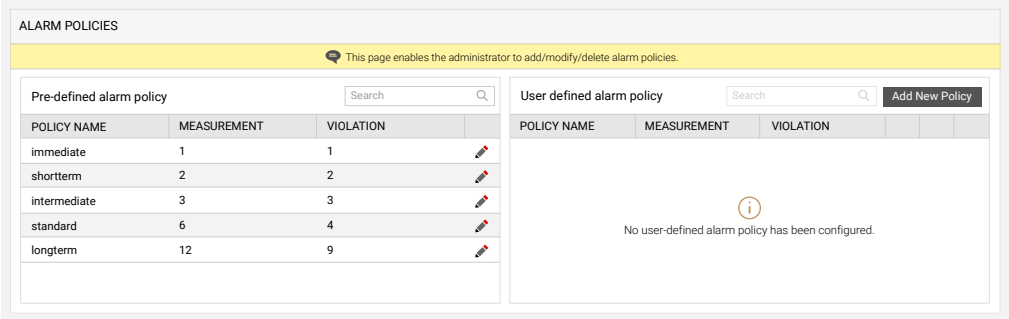

Figure 6: Out-of-the-box eG Enterprise sets numerous thresholds using the best choice of the pre-defined alarm policies. Administrators are free to add additional thresholds and modify those pre-configured; beyond this, administrators can also define custom alarm policies as needed.

While a threshold policy determines how the thresholds for a metric are computed, an alarm policy determines when alarms are to be generated to inform IT managers about a problem. Depending on their criticality, different metrics may require different alarm policies. For instance, an instantaneous surge in the CPU usage of a system is a natural phenomenon in a production system. On the other hand, even a sporadic unavailability of a critical network router is a critical event that needs to be informed to the administrator. Alarm policies must also consider the frequency of threshold violations of a metric. E.g., while an instantaneous surge of the CPU usage is not a cause for concern, a prolonged set of surges of the same metric may indicate a problem situation that must be corrected.

Each of eG Enterprise's alarm policies defines a window size and number of crossings (see Figure 7). For example, an immediate policy has a window size of 1 and number of crossings value of 1. This means only one measurement value is considered in determining the state of a measurement. If the current value exceeds the upper threshold limit, the measurement is said to be in an abnormal state, since the number of crossings is 1. As its name indicates, this intelligent alarm policy is ideal for cases where the administrator needs to be alerted immediately when an anomaly occurs. Metrics such as network / application availability can be monitored using this policy. For some other metrics, an administrator may not wish to be bothered about a sporadic threshold violation and may prefer to be alerted if a problem remains for a period of time. The standard alarm policy could be ideal for this, as it has a window size of 6, with the number of crossings as 4.

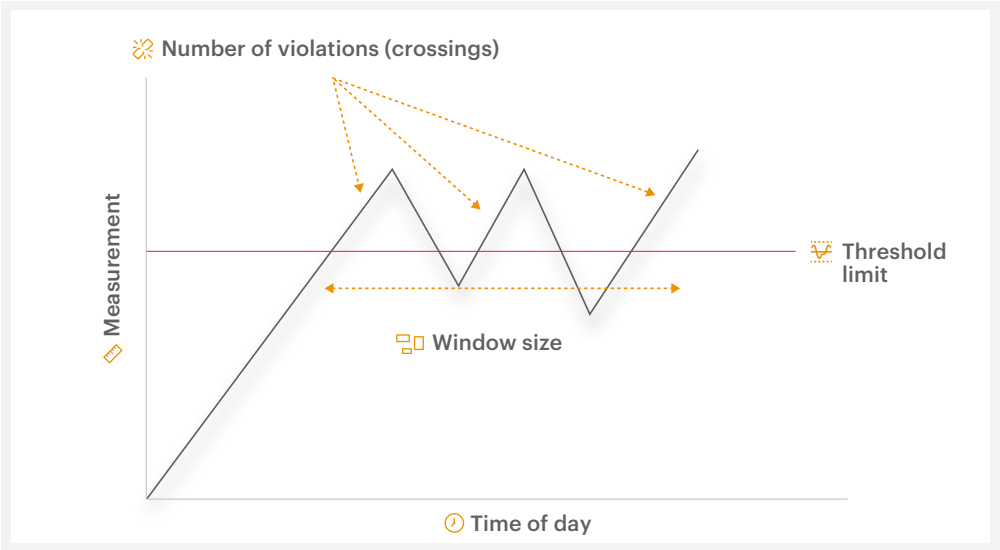

Figure 7: Defining Alarm Policies – This figure shows the concepts of "window size" and "number of crossings"

TEST DETAILS

Measurement Host	192.168.10.61	Component type	Oracle WebLogic
Component	Address-Validation-Service1:7001(10...	Test name	HTTP
Descriptor	HomePage	Test type	External
Test frequency	5 mins	Url Accessed	http://10.44.210.73.7001/
Time since last measure	2 mins 35 secs	Test state	Normal

Default threshold settings for this component

MEASURE	MAX/MIN	CRITICAL	MAJOR	MINOR	ALARM POLICY
Web availability(%)	Min	95	-	-	immediate
Total response time(Seconds)	Max	max(1, 125% of auto)	-	-	standard
TCP connection availability(%)	Min	95	-	-	immediate
TCP connect time(Seconds)	Max	max(0.5, 125% of auto)	-	-	standard
Server response time(Seconds)	Max	max(0.5, 125% of auto)	-	-	standard
Content validity(%)	Min	95	-	-	shortterm
DNS availability(%)	Min	95	-	-	immediate

MEASURES WITHOUT THRESHOLDS
Response code
Content length
Data transfer time

Figure 8: Alarm policies applied for different metrics collected for an Oracle WebLogic application server. Note that different metrics can use different policies.

❖ Choosing the Right Metrics is a Key to Being Proactive

Intelligent thresholding and flexible alarm policies are necessary, but not sufficient for a monitoring system to be effective. The metrics collected by the monitoring system are extremely important as well. If a monitoring system only collects availability and response time metrics, the metrics are not great early warning indicators of problems. This is because response time has an exponential distribution with load (see Figure 9) - i.e., as load increases, initially response time stays low, but as the load increases beyond the acceptable limit, response time shoots up dramatically even with a small variation in the load. This means that monitoring systems that use thresholds for response times are often not good at forecasting when problems are likely to occur.

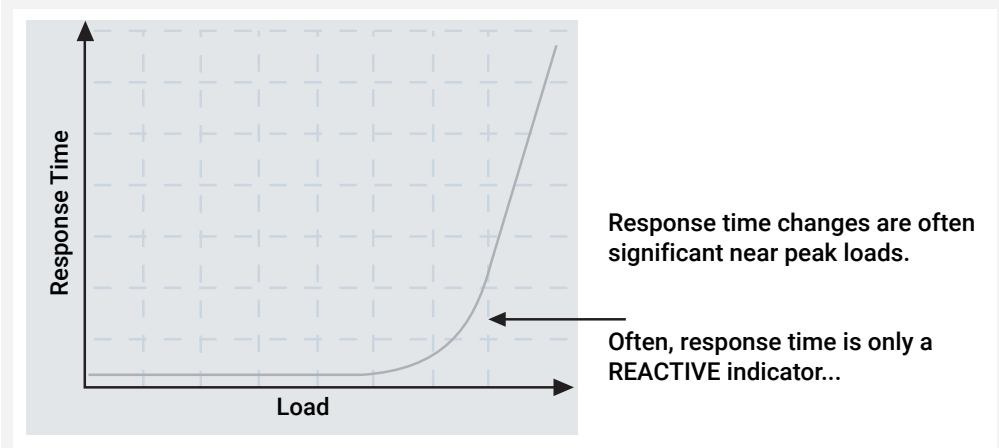

Figure 9: The variation of response time with load. The exponential relationship of response time with load means that response time is often not a good proactive indicator of problems.

A monitoring system that has domain and cross-domain knowledge built into it is often more effective than one that does not. Identifying which metrics are the most important to monitor and their inter-dependencies manually is untenable. For example, for a monitoring system monitoring VMware vSphere, it is essential that the monitoring system be able to look at "CPU ready times" of the virtual machines (VMs). By doing so, the monitoring system can determine times when the servers do not have sufficient CPU processing power. Early warning indicators can be provided to alert administrators if this issue starts to occur often. Likewise, tracking the number of requests waiting for I/O on a server can indicate a disk bottleneck, which if left unattended over time can result in catastrophic service outages.

Beyond Metrics

Monitoring thresholds determine IT performance alerts (techtarget.com)

A smart monitoring strategy uses more than just performance counters. Tools incorporate system logs to help identify issues and pair infrastructure monitoring with application monitoring.

— **Alastair Cooke**
TechTarget

Out-of-the-box eG Enterprise understands each key technology it supports and will configure thresholds to industry best-practices without the need for the administrator to manually apply thresholds and without the need to calculate or estimate the absolute threshold values.

❖ Automated Handling of Alerts and Ingestion into ITSM Service Desk Systems

If an enterprise adopts sophisticated thresholding technologies, they will minimize or even eliminate false positive alerts. Having reached this ideal state, it becomes practical and possible to automate the generation of service and help desk tickets as alerts can be assumed to be almost always genuine issues. eG Enterprise includes full API integrations with all major ITSM help and service desk tools to further eliminate the burden of manual tasks. Full API integrations ensure that alerts resolved and closed within eG Enterprise are automatically updated and closed in the ITSM systems ensuring the organization has a traceable single view of their operations.

"The ROI of eG Innovations exceeded our expectations. Not only did we get those chargebacks reduced, but we also benefitted IT as a whole because help desk tickets weren't created and time management was a lot better-- we weren't spending hours trying to pinpoint issues, we could reallocate those resources to other, more important activities."

— **Peter Dinh**
Senior, Virtualization Engineering Lead, eBay Inc.

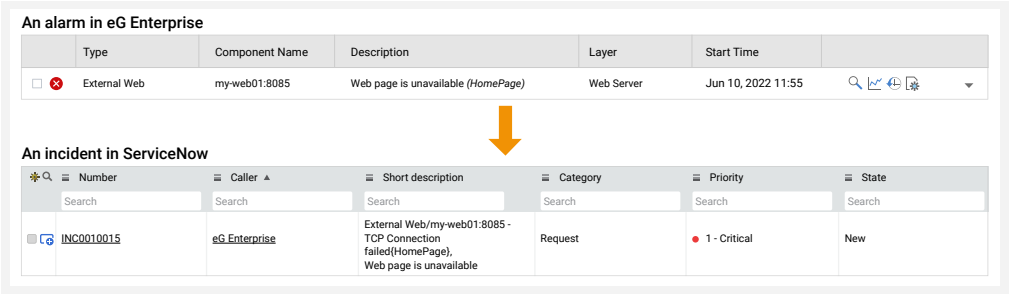

Figure 10: How alarms in eG Enterprise open incidents in ServiceNow ITSM without needing any human intervention

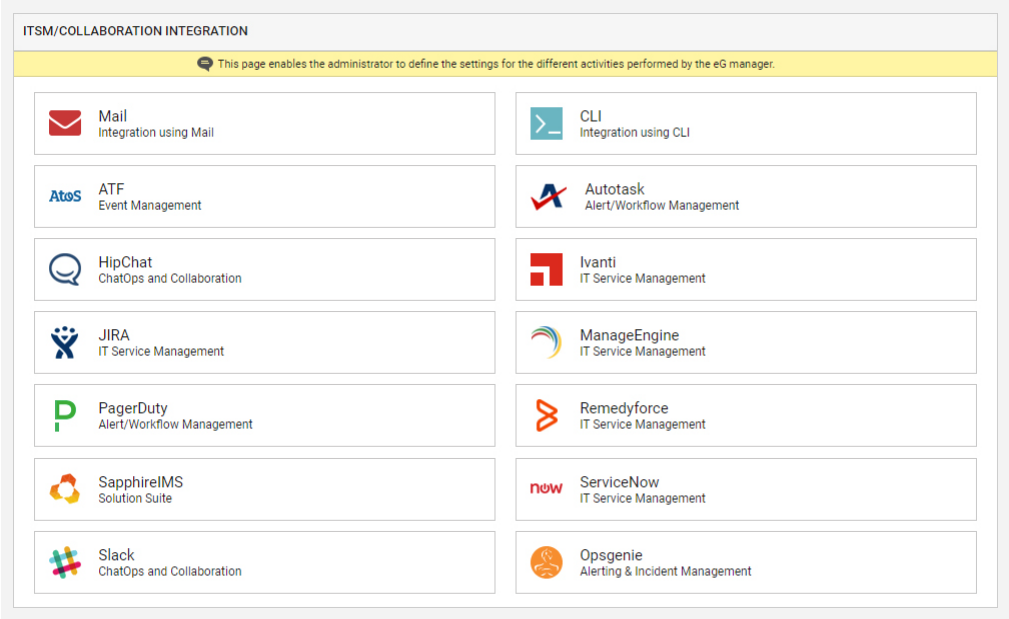

Figure 11: The different ITSM tool integrations supported by eG Enterprise

❖ Suppressing Alerts During Maintenance Windows

Enterprises that have adopted technologies that automate the definition and implementation of metric thresholds at scale must ensure that the associated alerting and alarming can be suppressed and temporarily halted with granular control for parts of their application and infrastructure landscapes. In larger organizations where teams may geographically remote and siloed, there must be mechanisms in place to ensure that taking a server down for routine maintenance does not trigger alerts and support tickets unnecessarily in other parts of the organization. Moreover, during controlled and planned changes such as rolling out pre-production to production systems it may be desirable to allow systems to stabilize and wait until all the components are in place before activating alerting.

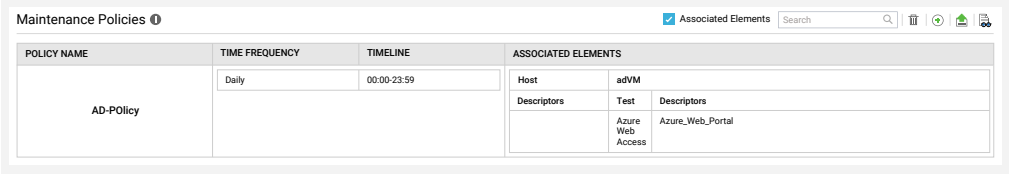

Figure 12: eG Enterprise allows administrators to schedule or spontaneously apply maintenance policies so that alerts are not generated to the help desk because "a server is down" when deliberate actions have been taken to take the server offline.

During some types of maintenance or change, administrators may want to leverage eG Enterprise to monitor and measure the impact of change but may be aware their actions may cause unintended consequences and impact their ITSM integrations. To avoid unnecessary tickets all our ITSM integrations include maintenance modes, which may also be used when refining thresholds.

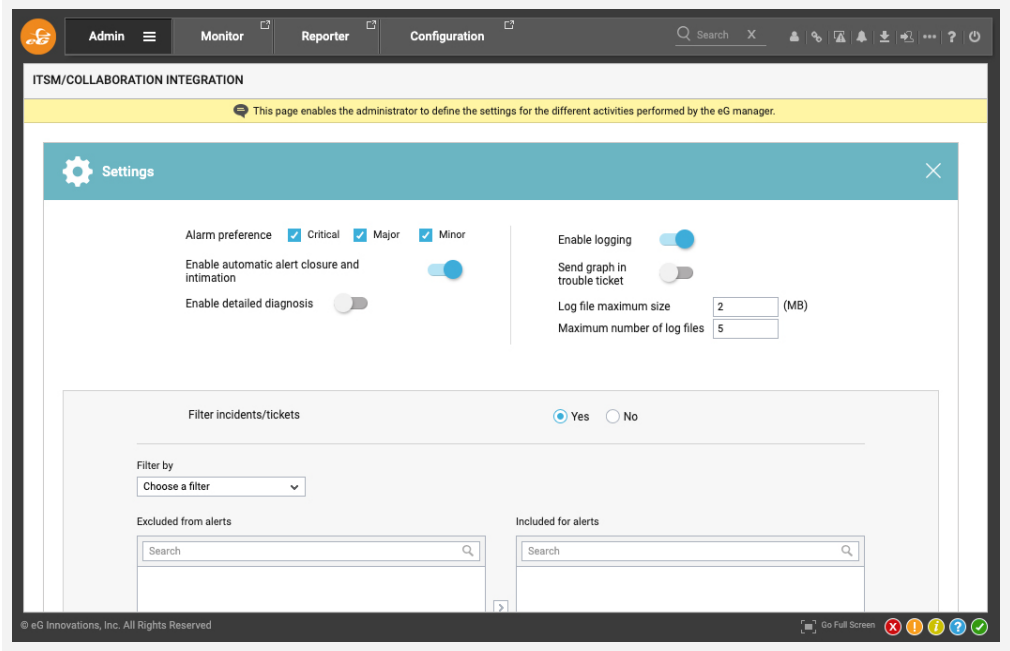

Figure 13: Different choices available when integrating eG Enterprise with ITSM tools for incident management

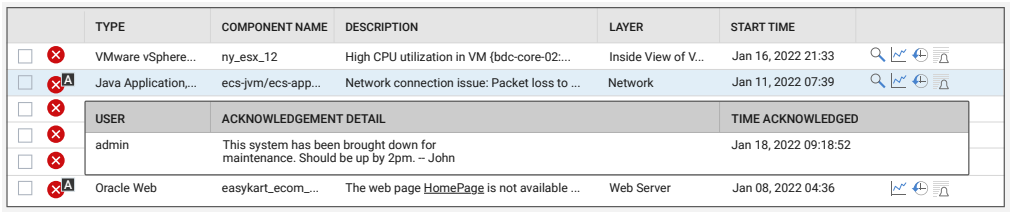

Figure 14: IT admins can acknowledge alerts in eG Enterprise. This serves as a simple way to communicate operational status between IT personnel

❖ Enabling Help Desk Staff Respond to Alerts

Once thresholds have triggered alarms and alerts, it is critical that the help desk operator or system administrator has instant access to understand which threshold has triggered the alert and the history of that metric alongside sufficient information for a non-domain expert to evaluate its significance.

eG Enterprise's context sensitive help feature embedded in the tool provides help desk staff with immediate details about all metrics – what each metric means, how to interpret the values of the metrics and recommended actions (if any) in case there are any abnormalities.

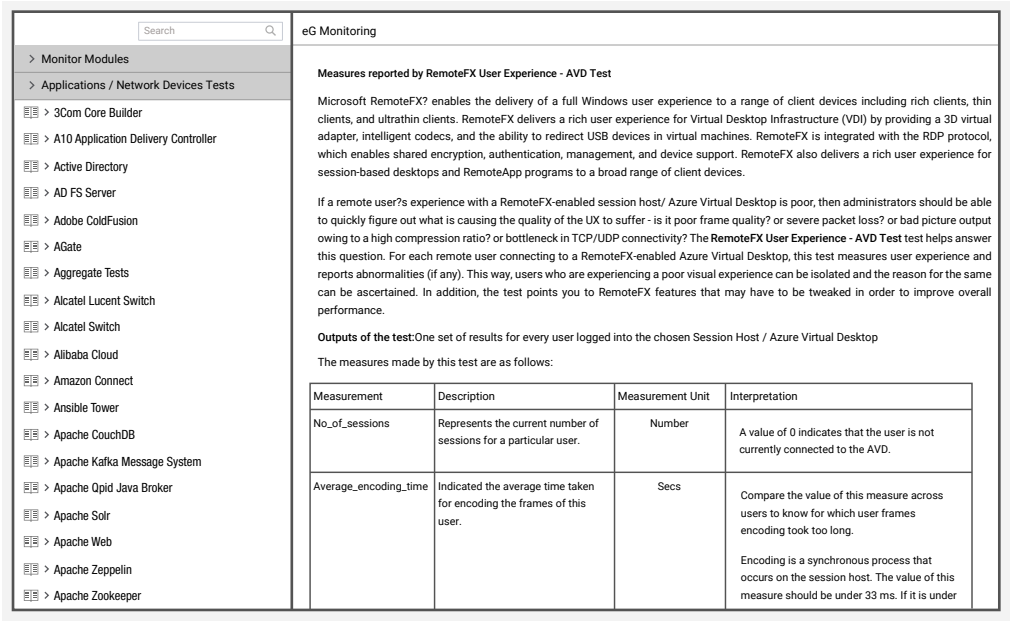

Fig 15: Context sensitive help provides immediate access to detailed information on metrics

Moreover, the built-in eG Enterprise "Knowledge Base" provides the frontline helpdesk operator with instant access to a site specific database of problems and known fixes. By searching this database, helpdesk operators can quickly identify what they need to do to resolve a problem. This knowledge base feature, if used as it is intended to, ensures that problem resolution is not an art known only to a few experts in the organization.

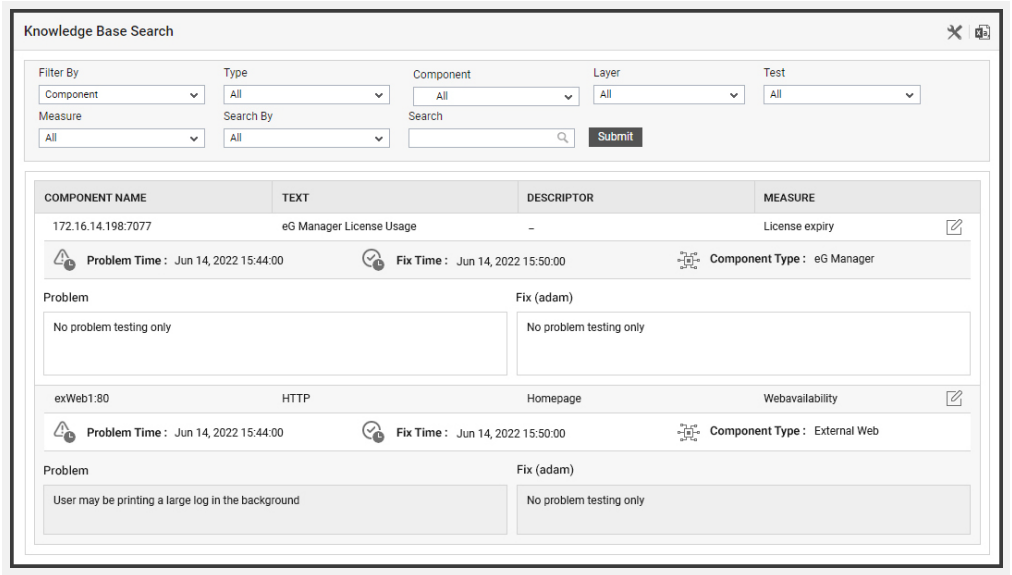

Figure 16: Knowledge Base search allows causes of problems and details of fixes to be logged and made searchable. This way, operational intelligence is not just with a few admins and IT teams can adopt processes to ensure knowledge sharing between their staff members

❖ Other AIOps Features

When looking to adopt an AIOps technology that can offer dynamic thresholding, you will also want to review what other AIOps features have been implemented. Products such as eG Enterprise leverage AIOps technologies further to reduce false alarms by filtering and correlating alerts to pinpoint the root-cause rather than secondary issues and symptoms. For example, if a VMware vSphere server host fails within eG Enterprise, this will be highlighted as the primary issue rather than the numerous in-VM user experience issues happening because of the host failure those sessions reside upon.

How AIOps monitoring eases modern IT challenges (techtarget.com)

"IT ops teams must maintain visibility into these modern and dynamic architectures, and it's not enough to monitor key metrics in isolation. Instead, they need a holistic view of infrastructure resources, complete with dependency mapping, automated root cause analysis and predictive insights. These goals can be achieved with an AIOps monitoring strategy"

Kristin Knapp
Editorial Director - TechTarget

An overview of AIOps features is covered in our *articles* and a one-stop eBook: *AIOps Solutions and Strategies for IT Management.*

❖ Evaluating the Effectiveness of Alerting

There are a number of ways of evaluating the effectiveness of a monitoring tool. The trend of alerts over time is one measure. Alarm duration is another metric. With proactive monitoring in place, the number of alerts over time should reduce in a stable infrastructure. Alarm duration should also be within the acceptable service level. Also focus on the number of problems in the minor and major category that have not escalated to become critical. Ideally, if your alerting strategy is working, you will see more major and minor alerts and fewer critical alerts.

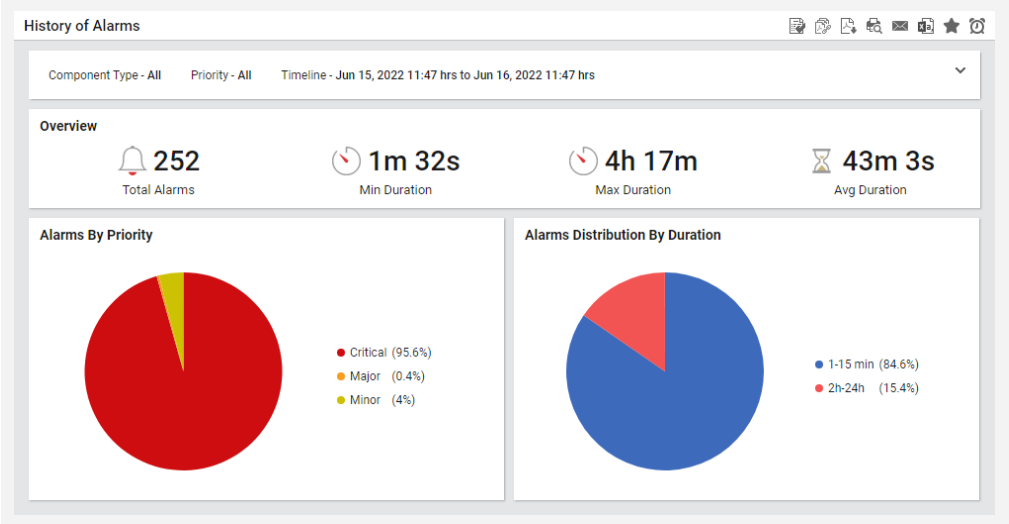

Figure 17: Historical reports in eG Enterprise allow IT admins to evaluate the effectiveness of alerts and to fine tune them

eG Enterprise also provides ways for admins to determine where threshold tuning may need to focus on. Review the alerts being generated by tier, by specific server, and by metric to determine if a specific tier, server or metric is responsible for a large number of alerts.

❖ Summary

In summary, to be effective, a monitoring system should include capabilities for thresholding intelligently, without requiring manual intervention, provide flexible policies for alerting administrators and be able to track key metrics in the underlying infrastructure. Having all three of these capabilities is key to enabling enterprises to manage their IT services effectively.

Early warning indicators provided by such systems can direct administrators to potential problems which if not fixed can have catastrophic consequences. The benefits of such a system are manifold. By being proactive, the monitoring system ensures it provides administrators with the indicators they need to fix problems without impacting the business services they are responsible for supporting. Intelligent thresholding makes the configuration and implementation simple, thereby ensuring that the monitoring system can be up and running quickly in a cost-effective manner.

On-going usage and maintenance of the monitoring system is also simplified by the intelligence built into the AIOps based monitoring system. Since they receive fewer false alerts, administrators can focus their attention on the key problems in the infrastructure, rather than being distracted by a large number of meaningless alerts and long-term alert-fatigue is avoided.

For more information on eG Enterprise, please visit *www.eginnovations.com*

❖ Learn More

To learn about how to understand, evaluate and leverage AIOps features, see
- AIOps Solutions and Strategies for IT Management | eG Innovations

- eG Enterprise can be deployed on premises, in a cloud of your choice or via ready-to-go SaaS (Software as a Service). For those looking to explore the benefits of static and dynamic thresholds and combining them a free trial on our SaaS option may be the best choice:
How to Deploy eG Enterprise – Choices and Models | eG Innovations

- TechTarget has several articles on the benefits of static vs dynamic alerts, including:
Monitoring thresholds determine IT performance alerts (techtarget.com) and
How AIOps monitoring eases modern IT challenges (techtarget.com)

- Read about integrating alerts with help and service desk tools such as Slack, ServiceNow, Autotask, JIRA, and others: *Service and Help Desk Automation Strategies*

- We have *a series of short (2-3 min) videos* covering the eG Enterprise interface including many aspects of alerting and thresholding, including: *How to Review and Interpret Alarms, Understanding and Modifying Thresholds, Understanding and Modifying Alarm Policies,* and *Creating Group Thresholds*

❖ Next Steps

✉ |	To contact eG Innovations sales team	: sales@eginnovations.com
🌐 |	Get a free trial of eG Enterprise	: www.eginnovations.com/FreeTrial
✉ |	For support queries and feature requests	: support@eginnovations.com

❖ About eG Innovations

eG Innovations provides the world's leading enterprise-class performance management solution that enables organizations to reliably deliver mission-critical business services across complex cloud, virtual, and physical IT environments. Where traditional monitoring tools often fail to provide insight into the performance drivers of business services and user experience, eG Innovations provides total performance visibility across every layer and every tier of the IT infrastructure that supports the business service chain. From desktops to applications, from servers to network and storage, eG Innovations helps companies proactively discover, instantly diagnose, and rapidly resolve even the most challenging performance and user experience issues.

eG Innovations' award-winning solutions are trusted by the world's most demanding companies to ensure end user productivity, deliver return on transformational IT investments, and keep business services up and running. Customers include Anthem, Humana, Staples, T-Mobile, Cox Communications, eBay, Denver Health, AXA, Aviva, Southern California Edison, Samsung, and many more.
To learn more visit www.eginnovations.com.