IT Performance Monitoring Tools: Reading Between the Lines

Key Capabilities for Effective Monitoring and Troubleshooting

An eG Innovations White Paper

www.eginnovations.com

Total Performance Visibility

Introduction

There are literally hundreds of monitoring products available in the marketplace today. All of them use similar terminologies to describe what they do – root cause analysis, topology views, proactive monitoring, etc. So, are they all the same?

Of course not. The capabilities of different monitoring tools can be widely varying, and you need to go beyond the terminologies used in marketing literature to decipher the exact capabilities of these tools. Here's a guide to some common performance monitoring terms and what you need to know about them in order to make the right choice for your needs.

Root Cause Analysis

Root cause analysis is an activity that identifies the root cause of an incident or problem. From a monitoring and event management perspective, the "root cause" is the event that, when corrected, will clear other events which occur as effects, rather than the actual cause of the event storm.

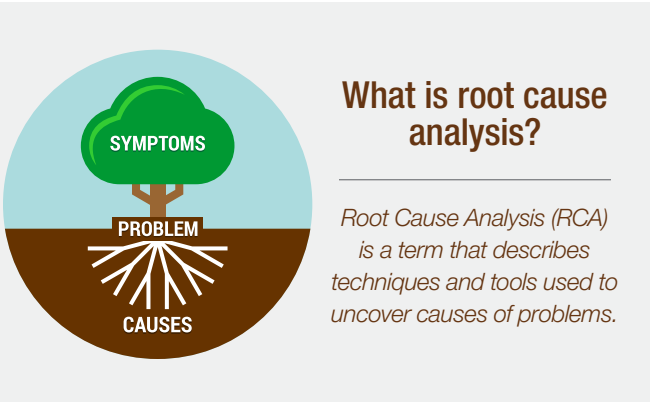

Figure 1: Understanding root cause analysis

What You Need to Know

Almost all monitoring tools can enable some form of root cause analysis. The key question is: **Who is doing the analysis?** Many tools provide all the events to the administrator and require him/her to analyze where the cause of a problem lies. To find the root cause of problems by analyzing all the events, administrators need to have a lot of domain expertise. Further, root cause analysis also requires a lot of time. Tools that do not provide automated root cause diagnosis cannot be effectively used by helpdesk/L1 support personnel.

> *Automated root cause analysis is where the real savings on a monitoring product is, and is why vendors use (and abuse) this term so much.*

Providing charts and graphs that an administrator has to analyze is not automated root cause analysis. When evaluating monitoring tools and hearing "root cause analysis", ask if it is automated and how it works? Also, what is required to get it to work, what level of accuracy does it enable and what work is needed to keep it working?

Topology

A topology shows an interconnection of devices, servers, applications etc. It provides a pictorial representation of how the infrastructure is set up and operating currently. Topologies often attractive images that present impressive-looking displays.

What You Need to Know

Today's IT infrastructures are heavily interdependent. A single web service may involve multiple web server front-ends that use middleware application servers, which in turn rely on database servers. Many of these interdependencies are logical (i.e., at the application-level), rather than physical interconnections. So, the questions to ask about the topology views provided by a monitoring tool are:

- Does the topology diagram present network-level or logical (inter-application) dependencies?
- What about other dependencies, such as a virtual machine to a physical machine, or an application to a virtual machine? How are these represented?
- Some of these dependencies are dynamic – e.g., a virtual machine can move between physical machines in a cluster. Does the monitoring tool take such dynamic dependencies into account?

Monitoring tools that represent physical interconnections alone are not sufficient to diagnose application performance issues. And, tools that do not consider dynamic interdependencies are not effective if your infrastructure is virtualized or hosted in the cloud (as with most current infrastructures).

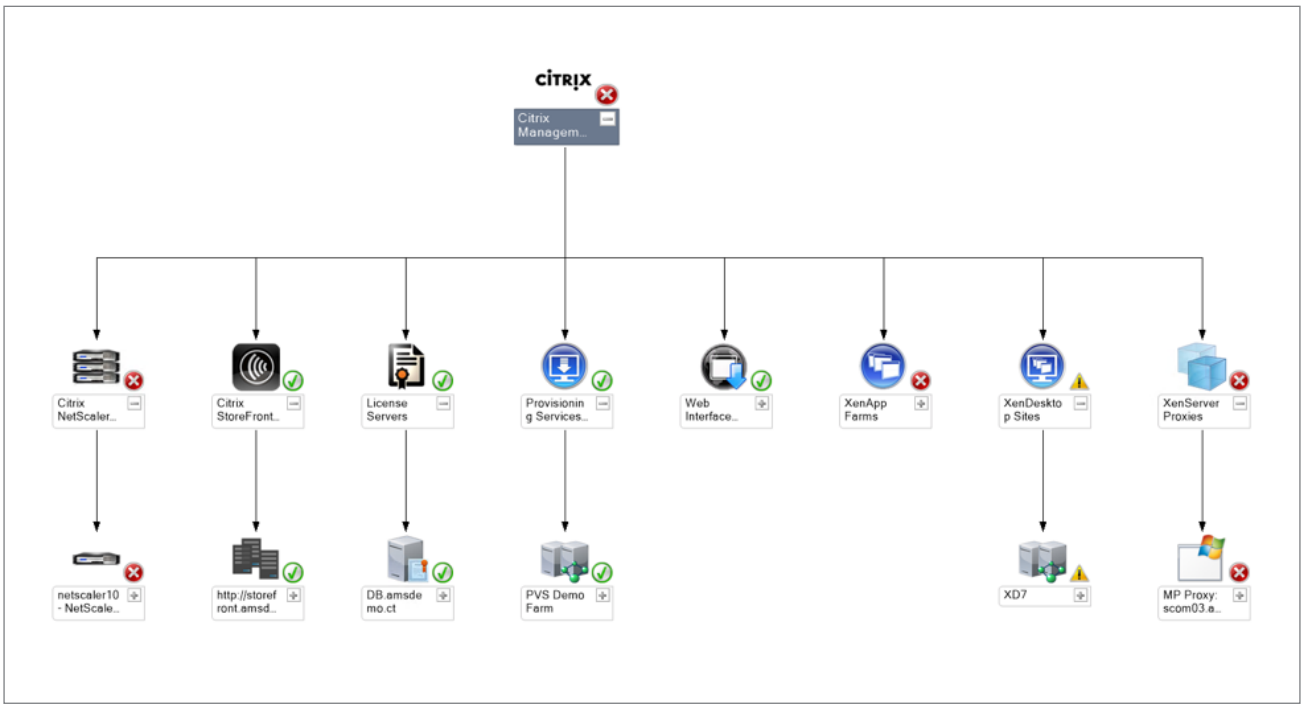

Figure 2: Example of a topology map without dynamic dependency mapping (SCOM)

> *Topologies are used in monitoring products, but many are just 'sizzle' and are not actively used as part of the monitoring and event management activities.*

Last but not the least, ask the question: **How is the topology used by the monitoring tool?** Is it just for displaying the state of the infrastructure, or does the monitoring tool use the interdependency information to help prioritize alarms?

To illustrate the point, consider how Microsoft SCOM shows topologies (see Figure 2). This topology is merely a grouping of components of different types. The color code for the group is a summary of the health of all components in that group. So, if one of the components has a severe problem, the group also reflects this state. Microsoft SCOM's topology diagram does not show any dependencies between components. In addition to this, there is no performance correlation between the different components in the topology.

Notice in the diagram above that there are multiple red alerts. So, where is the root cause? Is there one problem in the infrastructure triggering these alerts (which is the case in most environments), or are there really many mutually exclusive problems and alerts, as this topology indicates?

Therefore, it is important to assess the topology diagrams provided by a monitoring tool and to see whether they present dependency information between components or not. To be useful, the dependency information should be used to assist with root-cause detection. For instance, consider the topology diagram shown in Figure 3. All of the software and hardware tiers supporting a business service are shown in the topology and interdependencies between tiers is shown as well. The states of the different tiers are represented in the topology diagram so administrators can simply follow the color cues to determine where the root-cause of a problem lies.

In the example in Figure 3, a critical alert in the VMware tier is affecting the performance of a Citrix XenApp server. The Citrix XenApp server's performance, in turn, is affecting the Citrix StoreFront. Using dependency information between the different tiers, the monitoring system has downgraded the alerts on the XenApp tier and the StoreFront tier, thereby highlighting the VMware server problem as the potential root cause. Automatic, dependency-based performance correlation greatly simplifies problem diagnosis in dynamic, multi-tier application infrastructures, and accelerates troubleshooting.

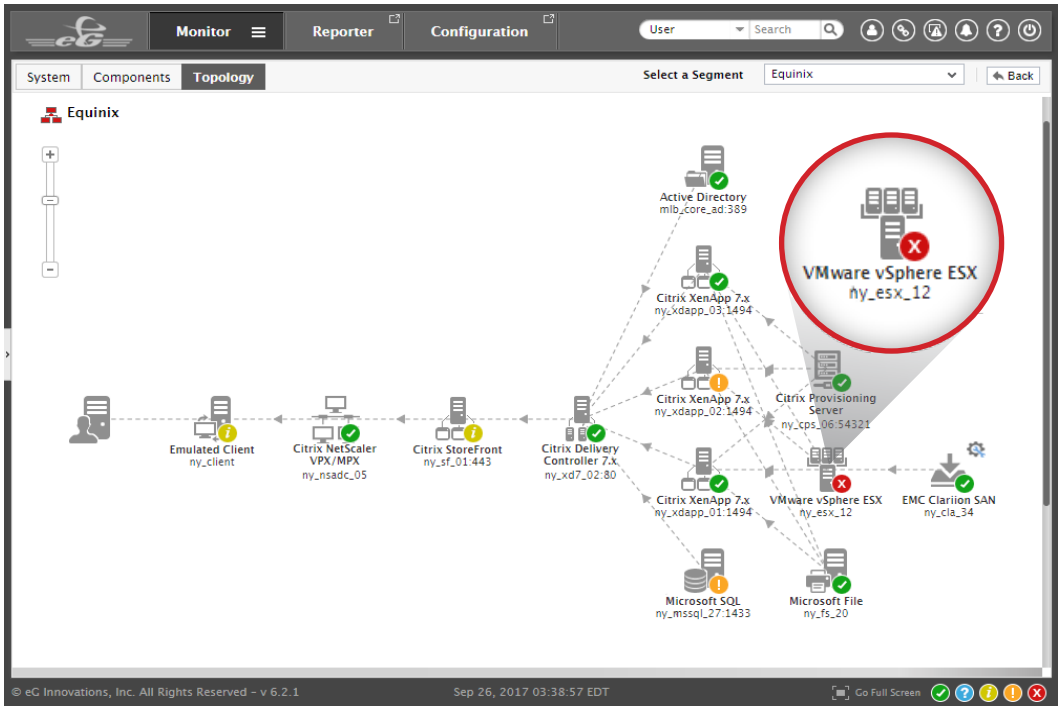

Figure 3: Example of a topology map with automated dependency mapping and root cause analysis

Dashboards

Dashboards are another feature that make monitoring products compelling for an organization. Dashboards are very useful for providing an "at-a-glance" view of key performance metrics, and again, most monitoring products will have some dashboard functionality.

What You Need to Know

To be useful, dashboards need to display real-time information, so as the state of the infrastructure changes, the dashboard reflects this. A key question relating to dashboards is: **How configurable are they?** What an IT executive wants to see is different from what an IT operations person wants to see. So, dashboards need to be configurable and tuned to specific personas, as well as facilitate their respective process activities, along with the rest of the monitoring solution. Dashboards that are static and cannot be configured are of limited use for IT operations. So, when evaluating monitoring tools, it is important to gauge the flexibility of the dashboards that a tool provides.

Also, dashboards need to be "active," meaning you must be able to click on a chart and drill down to other related information or more detail. This is important for performance monitoring – if the dashboard indicates corrective action is needed, you'll want to do so as quickly as possible, and within the context of the information provided by the dashboard interface.

> *Performance management dashboards need to display real-time information, be tailored to specific organizational and user needs, as well as enable performance management tasks easily.*

Another important question to ask about dashboards is: **Can they provide an aggregated view of the infrastructure?** Server-by-server views are important when diagnosing a problem. However, IT executives are more interested in farm-wide views – e.g. how many users are currently on your infrastructure (across all servers, rather than how many sessions are being supported server by server)? What is the average latency they are seeing? What are the utilization levels of your servers – what is the average CPU utilization of your servers currently? The monitoring tool and its dashboards must be able to provide such real-time aggregation capability. Taking this a step further, it is also important for an executive to know if the

total number of user sessions in progress now is normal or abnormal. To address this need, the monitoring tool must maintain history of aggregated metrics and compute health information for such metrics.

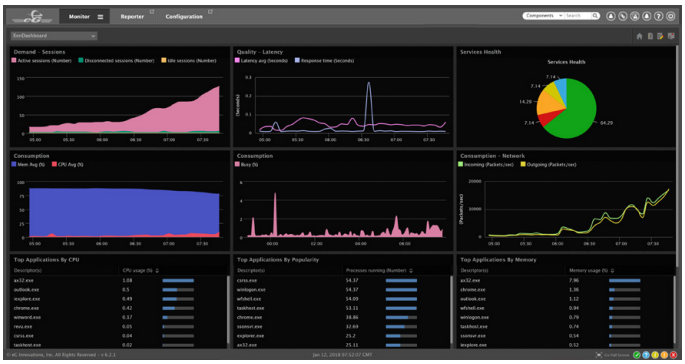

Figure 4: Example of monitoring dashboards with clickable, configurable, farm-wide views

Proactive

Root cause analysis is important, but getting a meaningful notification after a problem has happened is not as useful as learning about the problem before users are impacted. This is where proactive monitoring and management comes in.

Most monitoring tools can alert administrators when a metric exceeds a pre-specified threshold limit. This is not proactive monitoring. Specifying the threshold limit for every metric is very time consuming and requires a lot of expertise.

The questions to ask are:

- Can the monitoring tool detect changes in usage patterns and trends?
- Can the monitoring tool automatically determine what the normal values of a metric are?

Many tools present charts of metrics and rely on "eyes on glass," meaning that someone has to be reviewing the metrics at all times, looking for changes and trends. Ultimately, this is an impractical solution.

What You Need to Know

Monitoring tools that can analyze historical values of metrics and also learn the normal values of metrics can be truly proactive. These are capable of determining potential problem conditions in advance and alert administrators, so problems can be rectified before users notice them.

A key to proactive monitoring is understanding the norms of all measurement data, and knowing the deviations from normal behavior.

At the same time, note that the metrics collected by a monitoring tool themselves also determine whether it can be proactive or not. Metrics such as response time, availability and utilization levels may not necessarily be proactive indicators of issues. Look for metrics like errors, packet drops, memory leaks, queue lengths and so on that provide early warnings of impending problems. Such issues, if left unattended, can result in performance degradation and customer impact. A combination of auto-baselining and correlative intelligence will enable proactive performance monitoring.

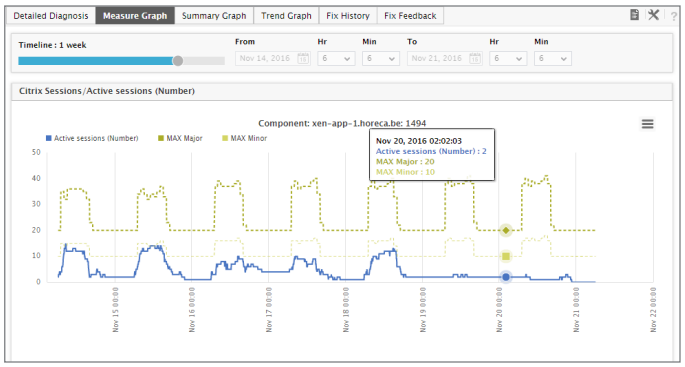

Figure 5: Example of auto-baselining of metrics

Shortage in resources on servers and VMs will affect performance of the applications running on them. Being able to analyze historical trends, extrapolate data and forecast future resource utilization patterns will help the administrator understand when a server/VM will run out of resources. A proactive monitoring strategy should also include forecasting capabilities to predict how long will existing resources last.

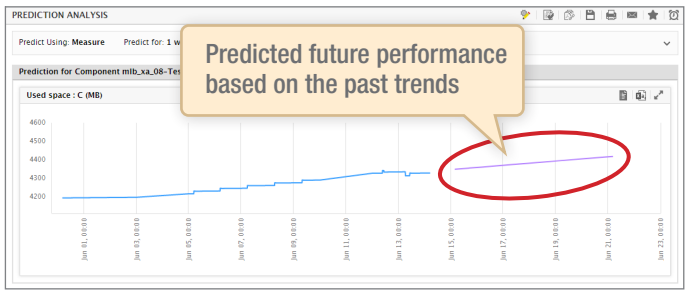

Figure 6: Example of forecasting capabilities in a monitoring tool

Navigating Your Way Around Monitoring Terminology

As we have seen, although different monitoring tools use the same terminologies, their actual features and capabilities can be very different. To draw a real-life analogy, you can drive from one place to another in multiple ways – you can use a map and trace your route, or you can use a GPS that not only gets you to your destination automatically, but it also chooses the best route you need to take based on current traffic conditions. The time taken for the travel and the effort you need to put in are very different in the two cases. In the same way, depending on the exact functionalities available in each monitoring tool, the results and ROI can be completely different.

Today's digital business services require performance monitoring solutions that are more like a GPS. The depth of metric collection, the breadth of visibility the tool has and the degree of automated analytics it embeds are what differentiate one tool from another.

Look for IT performance monitoring tools that simplify problem diagnosis and troubleshooting, allowing the IT administration teams to do their jobs better with increased efficiency.

Is your monitoring tool...

… like a map?

… or, like a GPS?

About eG Enterprise

eG Enterprise from eG Innovations is a robust IT performance monitoring tool for any size and any type of IT environment – on-premises, cloud, and hybrid. It automatically discovers all IT components, maps dependencies dynamically on intuitive topology maps and allows IT admins to get service-level topology views. eG Enterprise makes it simple for IT admins to detect and resolve performance problems by automatically correlating alerts across all IT tiers and pinpointing the root cause of issues in just minutes. The customizable dashboards in eG Enterprise simplify troubleshooting as IT admins can view real-time data on intuitive charts, track anomalies and get farm-wide views. eG Enterprise uses machine learning to baseline metrics based on customizable time-of-day, and day-of-week thresholds and proactively alert to performance deviations from the norm.

Learn more about eG Enterprise »

Next Steps

For more information, please visit https://www.eginnovations.com

Email us at info@eginnovations.com

LIVE DEMO

Request a personal walkthrough to learn first-hand how eG Enterprise can help improve performance and operations in your business environment.

FREE TRIAL

15-days of free monitoring and diagnosis, in your own infrastructure. Try it and learn exactly how eG Enterprise helps you ensure a great end-user experience and improve IT operations.

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 20th Century Fox, Allscripts, Anthem Blue Cross and Blue Shield, Aviva, AXA, Biogen, Cox Communications, Denver Health, eBay, JP Morgan Chase, PayPal, Southern California Edison, Samsung, and many more.

To learn more visit www.eginnovations.com.

Total Performance Visibility