

Power of the End-User View

Net Forecasts – Peter J. Sevcik
BCR Volume 36, Number 11
November 2006

In my July column, I described the terrible state of affairs in IT management and made a case for managing applications end-to-end rather than according to each IT silo. I suggested that a key to better management is to identify a single person who is in charge of application performance. In this column I describe the importance of using true end-user response time measurements as the foundation of good application performance management.

Performance is a broad topic with as many definitions as the number of people you ask. Some vendors will say that they measure end-user response time, but when you dig deeper you find that many of their tools don't measure response time itself, but instead derive a proxy measurement by watching packets or TCP protocol performance. They claim that their measurement of a part of the user's traffic, somewhere within the enterprise, is "good enough" because "a shift in this number is indicative of a shift in end-user performance."

Other vendors will claim that they provide the end-user perspective through the use of agents that operate simulated transactions "just like a real user would."

The problem with all of these partial approaches to measuring end-user application response time is that they are inherently prone to error and misinformation. A shift in packet statistics may have little or no impact on application performance. A synthetic agent can't be programmed to do the unusual things real users do. Furthermore, all of these approaches lack any knowledge of what is occurring on a user's desktop when performance changes.

There is simply no substitute for the real thing. Actual end-user performance is the solid truth about what is going on. It is invaluable in many ways.

A Tool For The CIO

Measuring actual end-user performance is the only way to resolve conflicting messages coming from various parts of the enterprise--for example,

business managers complaining about productivity, help desks getting calls from users about slow performance, or silo managers claiming that performance is fine. Real end-user measurements can sort out whether a few users are just loud complainers, or if there is a true performance problem. It distinguishes real application performance from user perception.

Real data about real users also clears the air of extraneous data that does not matter to senior management or the bottom line. Without real data, managers for each silo will simply report what they care about and what might support their quest for new technology. But if, for example, the server group can't show good correlation between end-user response time and server memory utilization, then why comply when this group requests that you buy more memory?

You can't manage what you don't measure, so by extension you can't manage well what you don't measure well. This kind of in-depth understanding of performance across many applications and over time provides the foundation for decisions about investments, cost savings and staff realignment.

Real user data enables the CIO to align IT resources to business requirements. It provides means to understand performance of the applications that are essential to revenue generation and/or customer support. Real user data also gives the CIO the means to manage the business and make sure that the people in his/her organization are not playing the blame game.

A Key To Diagnosing Problems

A user-centric view encompasses the entire infrastructure, casting light on problems within all components--whereas other types of management tools only see a few components. If there really is a problem in the desktop, LAN, WAN, datacenter, server, application, or back-end system, then it will show up in the end-user measurements.

A tool that instruments the desktop can also report on the condition of the desktop and other

applications. For example, it can inform the network manager that slow performance a user is experiencing is due to the fact that their desktop backup program is operating in the background without the knowledge of the user.

Furthermore, a desktop view documents the exact nature of the application behavior. This is a good way, and in some cases the only way, to understand how the application works. The performance problem may be that certain transactions with the application require too many application turns, or force the desktop to stop and wait for data from a service that no one knew was in the application. For example, the application may have a slow feature because the software is checking with its developer to see if it is authorized to execute the feature.

Many people think that the desktop view is interesting but limited in value, believing that if there is a problem downstream from the desktop, then measurements from this endpoint will be of little diagnostic value. The saying goes, “You can’t fix what you can’t see,” and some people would argue you can’t see too far into the infrastructure from the desktop. However, this is a mistaken perspective.

The value from desktop measurements comes from the aggregate knowledge that is found in gathering data from all the desktops. Let us look at an example of a company that operates 10 applications on 5 servers to a user population of 1,000 desktops in 20 locations connected by a WAN. If all of the users accessed all of the applications over a work day, then you have at least 10,000 individual elements of the complete performance picture.

By contrast, the server performance management tool is supplying 5 views of application performance. If each server supports all 10 applications, then there are 50 views from this perspective.

The IT manager who must diagnose a performance problem should leverage the 10,000 end user data points. This is the best starting point for problem triage.

If the measurement tool has a good user interface into its data, then the IT manager can quickly zero in on a likely area of detailed investigation: Is the

performance problem limited to a single application, geographic region, one location or server? If not, then what are the common aspects of the multiple desktops that see the problem? Is the problem related to a specific period of the day? An enterprise can develop a performance-based problem diagnosis decision tree that speeds this type of analysis.

Of course, once the triage phase is complete, the proper group within IT will be dispatched with their specialized tools to further diagnose and fix the problem. In fact, the most important result from such a process is that the right group be assigned to the problem. The organization avoids finger-pointing among groups or the wasteful “all hands on deck” exercise where every group looks into the problem.

End-user performance measurement should not be viewed as something that is turned on to solve a problem and then turned off. User-centric response time performance must be comprehensive and continuously gathered in order to provide the foundation for all other diagnostics and confirmation of problem resolution.

Good Tools Are Available

Getting this level of data is not difficult. Real user task-level application response time can be accurately measured by tools from several vendors including Compuware (ClientVantage), Citrix (EdgeSight), Quest Software (Foglight), Coradiant, Auditec, Symphoniq and Centrisoft. CA and IBM have purchased similar real end-user performance measurement vendors and integrated them into their enterprise management platforms.

Since there are so many choices, and in some cases from well established vendors, it is surprising that more enterprises are not using these tools. It is probably the lack of a focus on true application performance and the lack of a single person who’s in charge of the function.

There is also a general lack of understanding that good application performance management requires good process, not just a tool. None of these tools makes an organization better, more responsive or more productive just by running and reporting, if the person to whom it’s reporting can’t properly act on the information supplied.

Summary

Real end-user application performance measurement is the foundation for aligning IT with the business. It is also the first (triage) and last (confirmation) authority for properly managing application performance.

With so much benefit from an annual cost of about \$40 per seat, all enterprises should take a serious look at gathering and leveraging end-user performance data. Vendors of these tools could also help their customers by promoting the understanding of processes required to get maximum benefits from the data they supply.

Companies Mentioned

Auditec (www.auditec-newtest.com)
CA (www.ca.com)
Centrisoft (www.centrisoft.com)
Citrix (www.citrix.com)
Compuware (www.compuware.com)
Coradiant (www.coradiant.com)
IBM (www.ibm.com)
Quest Software (www.quest.com)
Symphoniq (www.symphoniq.com)

Peter Sevcik is president of NetForecast and is a leading authority on Internet traffic, performance and technology. Peter has contributed to the design of more than 100 networks, including the Internet, and holds the patent on application response-time prediction. He can be reached at peter@netforecast.com.

NetForecast is an internationally recognized engineering consulting company that assess and improves voice, video and data application performance. We help enterprises align application performance with business needs using a process based on the Apdex standard. NetForecast also advises vendors about customer requirements, technology issues, and the business value of their products. Visit our library of educational reports and articles about performance engineering.

www.netforecast.com

