



January 5, 2010
(Circ. Online)

Q&A: Application Performance Management and the End User
Managing application performance is all about how the end user perceives performance, and current tech trends are having an impact on that perception.

By James Powell

Application performance management (APM) isn't just about making sure applications stay up and running. The focus has shifted to end users, since their opinion and perception of software performance is what matters. To learn more about this shift, what the term "end-user experience" means, and how trends in the data center and technology (such as virtualization and cloud computing) are affecting APM and where APM is headed, we spoke with Lori Wizdo, a vice president with Knoa Software, an end-user-experience management solutions provider.

Ms. Wizdo has worked in the enterprise software industry since 1979 at a diverse set of technology firms, from start-ups to such global corporations as Unisys, NCR, and BMC.

Enterprise Strategies: What is application performance management and what benefits does it offer IT? Lori Wizdo: Application Performance Management (APM) is a discipline that aims to measure, manage, and optimize the performance and service availability of software applications. This is a fundamental discipline because software applications are the interface point between the IT organization and the business, and improving performance of business applications is a key to delivering business value.

There are many IT stakeholders who benefit from APM metrics and application insight. Application managers use alerting and real-time dashboards to proactively identify quality of service issues. Other members of the IT operations team (e.g., network managers, server administrators, database administrators) leverage deeper analytics to both troubleshoot application issues and plan for ongoing usage.

The discipline of application performance management is still a young discipline. The management tools for APM are often not integrated and provide a fragmented view of application performance. It is still a challenge for most companies to measure "end-to-end" response times. Furthermore, industry studies frequently report that fewer than 50 percent of organizations frequently have implemented and frequently review their service-level-agreements.

There seems to be a marked shift in the APM market -- it has shifted to focus on the end-user experience. What is behind this trend?

It's almost stating the obvious to say that a user's opinion of the performance of a software application will be a function of the experience that application delivers to the user. Yet, the strategy of monitoring and managing end-user experience is still in the vanguard of application performance management practice.

The seminal approach to measuring application performance was to measure the resources used by the application, and the processing times, at each of the tiers of the back-end infrastructure. This approach assumed that if the execution of the application was not causing a resource constraint or delay at the database server or the network server, or the application server (e.g., IBM CICS, J2EE, or .NET), then the application must be performing well.

This first generation of APM (which is still the predominant method in use today) is far from optimal. Back-end monitoring often results in the condition where “all systems are green” on the back-end but the business constituency is complaining that the application is slow or non-responsive. Just as frequently, a volume of alerts on the back end are simply “false-positive noise” because the business users are not being impacted. Both of these problems highlight the reality that measuring application performance as experienced by real end users is the superior strategy.

The term "end-user experience" presumably includes availability and good response times. What else does it cover, and does IT have a good sense of what's on this list, or is the term just vague marketing rhetoric that can mean anything to anyone?

The term “application performance” traditionally relates to how quickly a software request or transaction delivers a response to the application user -- whether that response is additional information or a confirmation of completion. The definition of “end-user experience” is more dimensional, going beyond simple response times to capture the overall quality of the end user’s experience.

The measurement of end-user experience should cover four important dimensions. The first, of course, is availability and response time of the application. A robust user experience solution will also include application and system errors that have a significant impact on the ability of the user to complete a task. Since the user experience is often impacted by the performance of the user’s device, metrics about desktop/laptop performance are required for adequate root-cause analysis.

A final dimension of end-user experience is end-user workflow which provides important business context around incidents and can help evaluate the impact of end-user behavior on performance problems.

Since there is no de facto standard for end-user experience, it is important to understand each APM vendor’s definition of the term “end-user experience.”

How are trends in data center and infrastructure management such as virtualization and cloud computing affecting the APM market?

Virtualization and cloud computing represent significant disruption in the provisioning model for the infrastructure on which applications run. The benefits of increased agility, reduced costs, and sustainability can be potentially lost if the new provisioning paradigm introduces business risk -- as business-critical applications suffer stability, performance, or quality problems in the more dynamic, volatile operating environment. As a consequence, most enterprises are exercising prudent caution before rushing to “virtualize” their critical applications or migrate them to “the cloud.”

These disruptive technologies increase the velocity of the end-user experience monitoring imperative. Measuring end-to-end response time as experienced by the end-user is the only known way to ensure application performance is commensurate with business needs after the back-end infrastructure has been virtualized or migrated to the Internet. End-user monitoring provides unified visibility and “one version of the truth,” offering the risk mitigation that can accelerate the deployment, and benefits realization, of these alternate provisioning strategies.

How do current trends in application architecture affect the APM market?

Today’s application development paradigm is another factor pushing the end-user experience monitoring imperative. The idea of measuring application performance by measuring resources consumption at each tier in the architecture is rooted in mainframe systems management thinking -- when the mainframe was a closed, tightly integrated system and software was built with monolithic application architecture.

Today’s business-critical applications have become modular and distributed. SOA has given us great flexibility and applications are essentially a set of loosely coupled services assembled at run-time into a coherent application. Rich Internet applications (RIA) are distributing processing to the desktop. With

applications that are assembled at the time of execution, you can make the point that the application doesn't actually exist as an entity to be monitored until the end-user hits the enter key. Hence, it makes sense to measure the performance of that application at the point of consumption by the end user.

APM solutions are frequently cited for their value in helping organizations achieve the elusive goal of "IT/business" alignment. Do you agree with this perspective, and if so, how does APM help this alignment?

Absolutely. The only use business has for IT (the technology and the organization) is when technology enables or facilitates better execution of a business process. Software applications are a reasonable proxy for a business process, and thus serve as the nexus between business and IT. When IT executives report improvements, investments, risks, and opportunities about an application, business customers will pay attention. When business executives discuss increasing the workforce, expanding geographically, or adding new business lines, IT must consider the impact to applications. Metrics about application availability and responsiveness are the lingua franca used by IT and business executives to discuss the quality of IT service delivery.

Increasingly APM vendors are talking about end-user productivity and performance. How is that relevant to an application performance management strategy?

It's a good trend and it does help give IT a richer vocabulary in discussing application performance with the business. A word of caution is that end-user productivity is not always seen as a harvestable source of value to the business. There are a number of incontrovertible sources of business value -- revenue growth, cost reduction, even something as seemingly soft as business agility. If employee productivity does not deliver one of those business outcomes, the business executives will discount IT's claims. Don't shy away from productivity improvement claims -- just be sure to link them to their ultimate business outcome.

What are some of the most common (and avoidable) mistakes you've seen people make when implementing APM? What strategies do you recommend to avoid them?

My advice here is controversial because it actually runs counter to advice that is positioned by traditional application management vendors as best practice.

The common mistake I see is that most APM strategies are built on the assumption that you can adequately manage application performance by managing a subset of the "most important" transactions.

The prevailing wisdom, promulgated by the APM vendor community, is that attempts to monitor every transaction are doomed to failure. The "best practice" advice is to engage with the business to ascertain which transactions are most critical based upon the impact if those transactions were unavailable or performing poorly. (That very conversation will not get you very far toward the goal of improving IT/business alignment.)

The reason that traditional APM vendors drive this approach is because of severe technology limitations in first generation APM solutions. Either because of where they collect their information in the stack or because of their technology platform design, these products require the customer to pre-define monitoring targets, then configure, script or instrument the solution to monitor these points.

There are a couple of serious problems with this approach. The first is that it is time-consuming and costly to deploy across the entire application landscape. When the application changes, which can happen weekly, these solutions have to be re-coded to stay current, so you have severe limitations on coverage and very high ongoing TCO. Second-generation APM products can identify meaningful transactions, across the entire application landscape, out of the box, and that's the way to avoid making this most common of APM mistakes.

The concept of a truly global APM strategy has also been constrained by the real difficulty in setting meaningful response time goals for the thousands of transactions that might be in a complex application. That is a real challenge that has been addressed in second-generation APM solutions. These solutions offer “dynamic baselining” which is a technique to compare real response times against historical averages. Dynamic baselining is an effective technique to provide meaningful insight into service anomalies without requiring the impossible task of setting absolute thresholds for every transaction.

Where is APM headed? What new features or functionality can we expect in the next few years?

The biggest changes in the APM market are going to be driven by factors external to the discipline and practice of APM. Change will be driven by some of the disruptive technologies that we have already discussed. With cloud computing, virtualization, composite applications, RIA, and other application methodologies and infrastructure options, it's become clear that the point of convergence, the point at which end-user performance needs to be measured, is the end user's desktop. You simply can't do it from anywhere else, period. Solutions that effectively measure cloud, composite, and virtualized applications from the end-user perspective will drive the innovation APM market from this point.

What role does Knoa play in APM?

Knoa provides end-user-experience monitoring solutions. Our end-user monitoring solutions enable management of many aspects of end-user experience, most significantly for APM, transaction response times, and system and application quality. A hallmark of the Knoa end-user monitoring strategy is the concept of global coverage – all transactions, all end users, all the time. If a company has an existing APM system, Knoa end-user metrics can be easily integrated to support a “single pane of glass” application management model. Knoa also offers a comprehensive portfolio of packaged end-user experience management solutions that offer role-based dashboards, flexible alerting, and comprehensive baseline, benchmark and trend reporting.

Knoa's vision for performance management reaches beyond infrastructure and application performance to monitor, measure and manage how end-users are utilizing the application to optimize business process execution. By monitoring the performance of people, process and technology, Knoa customers not only know if core enterprise applications are delivering an acceptable user experience, but also if application users are executing key processes effectively and efficiently - the key to achieving business value and ROI.

About the Author

James Powell is the editorial director of Enterprise Strategies (esj.com).