



Understanding the true costs of application performance problems

by Mary Petrosky

Sponsored Exclusively By:

FLUKE
networks.
.....

This Special Advertising Section Produced By:

NetworkWorld
www.nwfusion.com



Table of Contents

- 3 An accelerating problem**
- 3 The impacts of poor performance**
- 3 Impacts on internal operations**
- 4 Doing the fire drill**
- 5 Impacts on customers**
- 5 Crimping business growth**
- 6 The high cost of misdiagnosis**
- 6 Benefits of addressing performance problems**
- 7 Tools and tips**
- 8 Service-level management is key**
- 8 Taming the elephant**

The first step is to recognize their affect on the bottom line. The next step is to do something about it.

Poor application performance is the elephant in the room that many companies don't want to acknowledge. That's understandable given how difficult it is to identify the root cause of performance problems. Like the proverbial blind men describing an elephant, IT staff may attribute sluggish applications to everything from badly written code and overburdened servers to slow storage devices and low-bandwidth WAN links, depending upon which part of the IT infrastructure they don't happen to manage.

Ignoring the elephant is costly. Poor application performance can result in decreased employee productivity and customer satisfaction, as well as unnecessary capital expenditures and lost business opportunities. Organizations that want to limit the impacts of application degradation need to acknowledge the real impacts of application performance problems and to understand what they can do to curtail them. Given tight IT budgets and competitive pressures, organizations need to address performance problems both to reduce their overhead expenses and to expand their business.

"Today every company has issues around application performance," states Rick Sturm, president of Enterprise Management Associates (EMA). Studies bear out Sturm's assertion. For example, in a survey of 190 *Network World* subscribers conducted by Packeteer and Network World, 61% of all respondents said they had experienced significant application performance degradation, with that figure jumping to 85% for corporations with revenue of more than \$1 billion. And the problem is getting worse for many; 43% of respondents noted that application slowdowns had increased over the past 12 months.

About the Author



Mary Petrosky is a freelance writer and independent technology analyst based in San Mateo, CA. She's tracked the computer and networking markets for twenty years as a journalist, editor, and industry analyst.

An accelerating problem

Several factors are contributing to the rise in performance problems. The complexity of IT systems is certainly one. Interactions between various computing platforms, data resources, applications and the network infrastructure can be hard to untangle. Mark Her, EMA's research director, notes that the trend toward virtualization, which began at the circuit level and now encompasses storage and other systems, further complicates the picture. "The more that things become virtualized, the greater the management challenge becomes," he says.

Another major factor is the nature of the business climate, says Mark Fabbi, vice president, enterprise communications, for Gartner. In what Fabbi refers to as "integrated globalization," companies are integrating business processes on a global basis up and down the supply and distribution chain. "Now you're managing your supply chain on a real-time basis to a manufacturing plant that may be across the globe," he says.

EMA's Ehr concurs: "The hyper acceleration of the business process makes IT systems more and more critical." Because those systems are on the direct path for transacting business, if they're down or performing so poorly that users avoid them, there isn't an alternative for getting the work done.

Fabbi recalls one client that spent nearly a million dollars to extend a mission-critical business process to a manufacturing plant in China to get better control over these operations. Performance was so bad that workers at the plant refused to use the application. They circumvented the problem by telephoning their counterparts in the U.S. and asking them to input data on their behalf.

The move away from client-server computing to Web-based applications also is contributing to performance problems, Fabbi says. "Web-based design allows enterprises to extend the reach of applications to business partners, road warriors and other users," he notes. However, Web applications make much less efficient use of the underlying network, spawning more TCP sessions than client-server based applications. Secure Sockets Layer (SSL) and other security mechanisms that developers must layer on only further compound protocol overhead. Not only are Web-based application protocols heavy, but "enterprises have little control over how users connect to these applications," Fabbi observes.

Peter Firstbrook, program manager for Meta Group, agrees that the move to Web-based applications has exacerbated application performance problems in recent years. Business and application development teams are often so focused on the business functionality of a new application that they overlook performance until an application is deployed. "They don't think about how fast the screen will refresh over a WAN from L.A. to Kuala Lumpur."

The impacts of poor performance

ChartOne, a San Jose medical records management firm, knows first-hand how poorly performing Web-based applications can affect employee productivity and customer satisfaction. Three years ago, to accommodate growth and to ease the support burden on its IT staff,

ChartOne migrated from a homegrown, client-server enterprise resource management application to a Web-based, off-the-shelf package. The productivity of the accounts receivable department took an immediate dive.

"It could take minutes to pull up an invoice," says Henry Svendblad, vice president of information technology and systems. "Employees were processing 20% to 40% fewer checks and invoices on a daily basis." And these were LAN-attached users.

Customers, who access online records via a variety of connection methods, found their performance impacted by slow or congested links. Despite this, users in medical records departments at the hospitals and other medical facilities that ChartOne serves generally found the application's performance acceptable. However, many doctors — who tend to access records via dial-up or DSL connections — were not as enthusiastic. "Doctors wield a lot of power," Svendblad notes. "If they're not happy, they're going to make their feelings known."

Poor application performance can have wide-ranging consequences, affecting user and IT staff productivity and morale, reducing customer satisfaction, and negatively impacting business growth opportunities. For example, among Packeteer/Network World survey respondents, 36% either delayed launching new applications or curtailed the roll out of a new application because of concerns about application performance.

EMA's Sturm puts it bluntly: "Anyone who's ignoring performance issues is inviting a) disaster and b) outsourcers into the executive suite."

Impacts on internal operations

Employees, both users and IT staff, are among the most severely affected by slow applications. For end users, "a poorly performing application hits on two levels," says Dave St. Clair, director and CIO of the International Assignment Services Technology

Faster troubleshooting

Quickly isolate application response time problems to the network, server or application

FLUKE
networks

Center of Deloitte Tax in Dallas. “That person is less productive and they now have a reticence to use the application because it’s too cumbersome.”

As **Figure 1** illustrates, 40% of Packeteer/NWW survey respondents said that application performance degradation had either a high or very high impact on individual employee productivity, while 31% indicated a moderate impact. Team productivity was also affected, with 32% of respondents reporting high or very high loss of team productivity and another 33% reporting moderate impact because of poor application performance.

Nearly 40% said the degradation was so severe as to affect employee willingness to use the application.

It can be difficult to quantify the productivity of office workers, which is why many companies let application problems that affect internal users languish, Firstbrook says. The impact on the bottom line depends on the application and who’s using it. It’s easiest to figure out losses for applications directly tied to revenue generation.

“You can start a ticker going and determine the problem is costing X,” Ehr says. “In general, if the users are internal, application degradation is going to impact the bottom line from the cost side. If performance is affecting external users of the system, it’s likely to impact the bottom line from the revenue side.”

Tracking the number of transactions that users can perform is a key metric. Like ChartOne’s Svendblad, Gartner’s Fabbi found that an insurance client saw a decrease in transactions per employee once a new Web-based application was rolled out. “You could tell where the remote office was based on response time, which impacted user productivity,” he says. Response time in the home office and nearby East Coast offices was good, but decreased the further west an office was. “By the time you got to Alaska, productivity had dropped off 50%,” Fabbi says.

User problems quickly become IT’s problem. According to Packeteer/NWW survey respondents, IT found out about problems through employee complaints, which topped the list (at 82%), followed by calls to the help desk (73%). Only 38% said they were clued into problems because they monitored application performance.

Doing the fire drill

Resolving performance issues is compounded because most organizations don’t track what’s on their networks. Seventy-five percent of the Packeteer/NWW respondents admitted they have only a partial understanding of what applications they’re actually running. Without monitoring tools, IT often doesn’t know there’s a problem until the phone rings, says Ehr, a former IT manager himself.

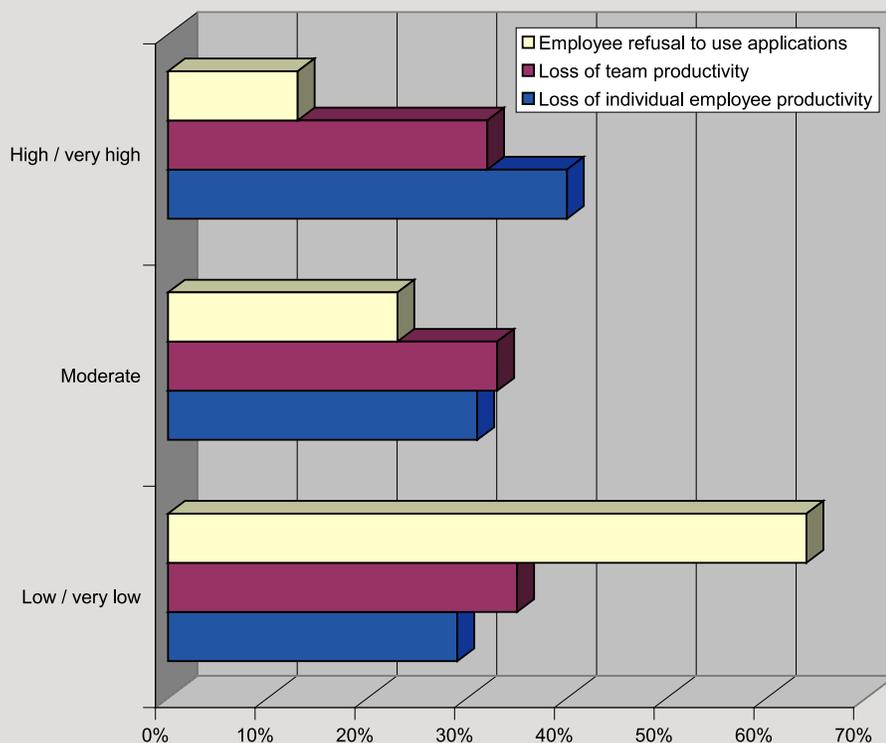
“And when the phone rings, you pull the red lever and do a fire drill, which is enormously wasteful and may involve eight to 10 people when all it requires is one to fix the problem, if you knew what the real problem was,” he says. This type of reactive behavior is demoralizing for IT to deal with day after day vs. “doing something fun, such as new projects that help the organization,” he says. It also gives IT a black eye.

In addition, the cost to fix an application performance problem is much higher if the problem is found in the deployment phase, rather than during design, development and testing, Deloitte’s St. Clair notes. Catching problems early means higher productivity for the development team as well as the IT infrastructure team, “and that equates to being able to do more — hopefully within normal hours, if there is such a thing,” he says.

Unfortunately, the compartmentalization of IT functions often hinders the resolution of application performance problems. “We see a lot of finger pointing,” Fabbi says. “IT groups that look at the problem on a silo by silo basis will never understand the problem. Unfortunately, it generally takes the failure of a pretty big project to get the different IT groups working together.”

Finger pointing definitely hindered the resolution of application performance problems for one financial services firm. The IT manag-

Figure 1: Impact of Performance Degradation on Employee and Team Productivity



er at this firm noted that a development team regularly creates applications related to new products and services. During a typical development cycle, an application goes through testing phases with various groups, including the applications development team, production support team and network group. Because each group was using its own test tools, pinpointing the source of a problem was difficult and led to finger pointing.

Several months ago this firm began using Fluke Networks' SuperAgent Application Performance Analyzer. The IT manager noted that the Fluke Networks' tool provides all the teams access to the same data, enabling them to quickly identify the root cause of a problem as well as to spot performance trends.

Without such tools, the blame game continues, Fabbi notes. IT's workload is often increased as a result. He notes that he often sees organizations roll out an application that works fine in the head office but performs poorly in remote offices. IT then ends up maintaining both the legacy system and the new one. "So you have a doubling up of operations costs as opposed to streamlining, which is what their initial intent was," Fabbi says.

Ensuring that users are happy is paramount, Ehr says. "IT staff really need to be thinking about how to defend against being out-sourced. One key way of doing that is to ensure that end users are happy — it's a fundamental tenet of a service-based philosophy. IT is a service to its consumers regardless of whether it's on-site, off-site or in India. If users aren't happy, something's going to change."

Impacts on customers

For many organizations, degradation of customer-facing applications poses the greatest risk because competitors are only a click away. "The scary part about the impacts on external users is that it's a hard one to quantify. You may not know how much you're losing," Ehr says.

The worst situation, Ehr says, is where application performance degrades to the point of being slow and annoying but doesn't prompt complaints. "It's that brownout situation where performance may not be so bad that anyone calls the help desk, but it's impacting productivity, it's impacting your customers' impression of you. And if you're not aware that there's a problem, that could be very damaging, especially if it goes on for a while. That customer just quietly goes away."

Gartner's Fabbi agrees: "It's like closing your doors and turning the lights out. If you're making it difficult to do business with you, people go elsewhere."

Packeteer/NWW survey respondents also acknowledge that poor application performance often translates to reduced customer service, as **Figure 2** illustrates.

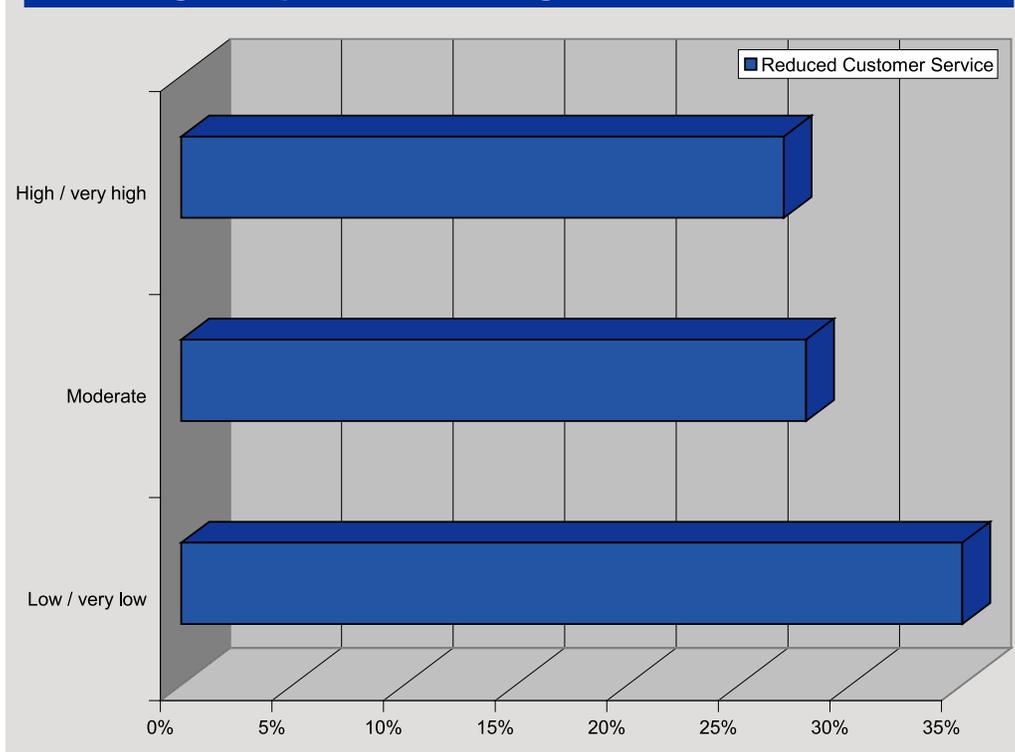
For service-based organizations such as ChartOne, Deloitte Tax and gedas USA, a positive customer experience is crucial to their business. For example, gedas USA, in Auburn Hills, Mich., provides IT solutions and services to the automotive and discrete manufacturing industries. Ethan West, team leader for gedas' data network, says unequivocally: "End user response time is everything."

Similarly, Deloitte Tax acts as an application service provider to 60,000 users in 70 countries, so it's imperative that application performance is good for all users. "It needs to be consistent within a reasonable degree of variance," St. Clair says. When a performance problem arises, "you have to view it from the client side," he notes. "The client doesn't care if it's a browser problem, a database problem, a network problem or an application problem. From their perspective, it's broke. So you have to ensure that you've taken all the precautions — common processes, common tools, whatever's necessary — to provide that high level of consistency."

Crimping business growth

Analysts and IT managers indicate that beyond impacting employee and IT productivity and customer sat-

Figure 2: Impact of Performance Degradation on Customer Service



isfaction, poor application performance can stymie business growth by tying up IT resources that could be spent on new initiatives.

Fabbi notes that “the bigger problem associated with application performance problems is time to market or time to the next project.” Because IT is caught up in the cycle of fixing problems, they’re continually delayed from getting to the next thing on the priority list.

Ehr recalls a large energy company that was doing so many fire drills that its strategic projects were getting delayed, including a VoIP initiative. “So not only was the company having the impacts from slow IT response and application crashes, they were also missing out on millions in savings from the VoIP project.”

Sturm adds: “It’s about being able to pursue the strategic, not the urgent.”

For St. Clair, “it’s an opportunity cost. I’m spending a lot of time, energy and good will to throw people at problems, the majority of which should never have seen the light of day.” Because a key part of his business is developing applications for clients, his philosophy is to streamline each application to work in the most draconian environment. Otherwise, he says, “your growth opportunity is limited because the infrastructure required to run the application may be higher than your market can bear. It can be the greatest application since sliced bread, but no one will use it if you need a T-3 to each user to make it work.”

Nearly half of the respondents to the Packeteer/NWW survey indicated that application performance degradation led to a decline in revenue-producing activities. Nearly a third of respondents said they’d delayed rolling out CRM applications, for example, because of concerns about application performance.

The high cost of misdiagnosis

Perhaps the easiest aspect of application performance problems to quantify is the amount of money companies have spent — or mispent, in many cases — trying to fix a problem. If congested WAN links are seen as the source of performance troubles, IT might be tempted to proliferate data centers or buy more bandwidth. If a storage-area network is suspected of being a bottleneck, additional SAN gear is soon on order.

Server and network upgrades are among the most common “quick fixes” that organizations attempt. Gartner’s Fabbi notes that IT’s “silo” structure “tends to result in silo-driven expenditures.” For example, if the problem is perceived as server-related, server upgrades are seen as the solution.

ChartOne initially looked into upgrading its servers to boost performance of its newly deployed Web applications, Svendblad notes. “We were looking at about \$200,000 in servers,” he acknowledges. He did purchase one mid-range server from Sun, but saw only minimal performance improvement. This prompted him to investigate other solutions, including application front-end processors from Redline Networks. He invested \$50,000 in Redline gear two years ago and in that time the company has doubled the size of its accounts payable and receivable groups, which handle 40% more transactions without any upgrade to computing or network equipment.

Svendblad recently trialed a number of application performance acceleration appliances to see if he could improve performance for customers. ChartOne hosts compressed TIF files, which left Svendblad to question whether these devices would buy him much performance because many use compression, caching and pattern matching techniques to squeeze more data onto a network link. He was pleasantly surprised to find the Redline boxes reduced the bandwidth consumption of this traffic by roughly 35%.

“It’s given us more head room in our existing host environment and has improved the experience for end users, which translates to happier customers,” he notes. Not to mention a savings of \$150,000 in server hardware.

Ehr has seen several clients poised to spend money on the wrong solution. He recalled one CIO who was ready to do a multi-million-dollar server upgrade until a performance management tool properly identified the problem as a bottleneck in a data connection. “As the overall complexity of IT systems increases, the likelihood that a forklift upgrade or overprovisioning will solve the problem decreases proportionally, maybe more,” Ehr says.

Benefits of addressing performance problems

Being able to avoid or quickly resolve application performance problems yields a variety of benefits, from increased employee productivity and higher customer satisfaction to avoiding the need for bandwidth, server and other resource upgrades. For many organizations, tackling application issues more efficiently means freeing resources that then can be redirected to meet business requirements.

For Gensler, which provides architecture, planning and design services worldwide, finding a solution for their performance problem has opened a whole new way of doing business, according to Principal

More Effective Change Management

Test and measure the impact
of infrastructure changes on
end-user response time

FLUKE
networks

and CTO Bruce Bartolf. The San Francisco-based company has 26 offices globally and sharing data among remote sites was a huge challenge because of the size of the company files, which include CAD files, movies, large image and plot files and standard architectural detail objects. Users weren't working on centralized files in real time because it could take 30 to 50 minutes to open a file across a T-1 link.

Rather, the IT staff would copy files that users needed to share, sometimes physically to tapes which would be overnighted to the remote user, other times using the Robocopy utility at night. Once the modified file was sent back to its "home" site, it had to be manually reconciled with the original to avoid accidental overwrites and other consistency issues.

"Bandwidth wasn't our problem," Bartolf says. "Latency and protocol inefficiency was the problem." After installing Riverbed Technology's Steelhead transaction acceleration appliances, Bartolf says it now takes 3.5 to 4.5 minutes to transport a 35Meg CAD file from New York to San Francisco.

"We weren't able to do remote work before," he says. "And now we do it. We have a large number of highly skilled, highly vertical-market focused individuals all over the globe that we can bring to bear for client work. This allows us to staff projects with the best possible people and give our clients the best possible experience. It's a business changing model for us — that's huge." He notes that within five weeks of installing the Riverbed devices, Gensler was able to staff projects using individuals from multiple offices.

In addition, Bartolf now can centralize servers and "the IT staff, instead of babysitting file servers and doing tape changes, can now add value." He cites project quality, client satisfaction, project turnaround and team productivity as returns on his investment. But ultimately the rhetorical question he asks is: "What's an additional business capability worth?"

gedas has found addressing performance issues to be similarly business changing, says West. An international firm that hosts applications, such as customer relationship and project management, for its clients, gedas wanted flexibility in locating its computing resources. When the company decided to move its mainframe services from a LAN-attached set-up in the U.S. to Brazil, West knew that latency and bandwidth would affect performance. FTP (File Transfer Protocol) jobs that had taken 20 minutes were taking two to four hours, while an Internet Control Messaging Protocol ping took roughly 20 times as long. West knew that this level of performance was going to make it a challenge to meet service-level agreements (SLA) with customers.

After investigating WAN acceleration products, he selected devices from Peribit Networks. "With the Peribits in place, customers are interacting with the hosts as if those systems were local. The end user experience is nine-tenths the battle, and we're meeting our SLAs," West says. gedas is saving millions of dollars per year by having the flexibility to relocate its hosting services. gedas can more easily deploy VoIP to various sites now because the Peribit gear supports traffic prioritization, West notes.

"The Peribit boxes are allowing us to deploy computing resources where we need to without the expense of upgrading bandwidth, and to put more traffic on those links," he says. "They allow us flexibility in growing the business."

ChartOne's Svendblad also notes that the Redline performance acceleration technology he's deployed provides "more headroom in our existing host environment and has improved the experience for end users, which translates to happier customers."

Likewise, for a major financial services firm, Fluke Networks' application monitoring tools have streamlined application development, giving the company a competitive edge and freeing up IT resources to focus on strategic projects, such as VoIP deployment.

Tools and tips

While there's no silver bullet that will solve all application performance problems in all situations, there are a number of things organizations can do to limit their exposure to such problems. Understanding the types of performance enhancement and monitoring products on the market and when to deploy them is crucial.

For example, application acceleration devices from companies such as Redline can offload SSL and TCP processing from application servers, freeing CPU cycles for data crunching. Similarly, WAN optimization devices, such as those from Peribit and Riverbed, perform compression, data caching and other functions that allow more data to flow over a given WAN link.

"One of the reasons these devices are flying off the shelves is that they pay for themselves in less than six months in most cases," Meta's Firstbrook says. "It's a low-risk proposition." However, in the case of bandwidth optimization tools, Firstbrook cautions that "simply adding more bandwidth without controlling it may not solve your problem. You may double your bandwidth only to find that most of it's going to IM chatting or peer to peer Napster stuff."

Happier End Users

Continually track and ensure application performance over time using SLA reports

FLUKE
networks

He recommends companies look for bandwidth optimization products that provide traffic management, allowing IT to have clear visibility into what's on the network and to determine what gets priority handling. "You need traffic management tools with granular control," Firstbrook notes. "It's not enough to accelerate HTTP traffic because that means you've accelerated all your Web traffic as well as your mission-critical enterprise applications. You've got to get more detailed than that, down to the URL for a particular group of users, for example."

Monitoring tools, such as Fluke Networks' SuperAgent, are also crucial, providing insight into application, system and network behavior. Fifty-six percent of Packeteer/NWW survey respondents indicated that monitoring network performance clued them in that there was an application slowdown on their network, for example, while 38% noted that monitoring application performance tipped them off.

Deloitte uses a wide variety of tools to ensure good application performance, from automated quality assurance tools for application staging to tools that look at database, server and network utilization. "All of that's valuable," St. Clair says. "You need to make sure all those components are functioning." However, he doesn't stop there. "Once the application is deployed, you need to monitor it from the user's perspective," he says. "It's what's on the desktop that your client sees. So you need to be monitoring from the Ethernet port on the PC or whatever Internet device they have their browser on all the way through to where you know what their experience is."

Service-level management is key

To ensure that application performance meets user needs, IT must quantify "good" performance and then be able to validate that the specified level of performance is being delivered. These concepts are at the heart of service-level management. Ideally, IT should define a separate SLA for each mission-critical application, taking into account the variety of access methods users will employ. The baseline for acceptable performance should be based on user requirements, with an eye toward adjusting it over time to achieve ongoing operational improvement.

It's important to distinguish between application availability, or "uptime," and actual performance, EMAs Sturm says. "An application can be 'up' but performing so poorly as to be unusable," he notes. Any measurement of service levels "must encompass both availability and performance, implemented with end users' experience in mind. The only thing that matters is the end user experience. Nothing else matters."

Among the metrics that IT should track are the amount of time it takes traffic to travel between a client and server and back, and how much time it takes servers to process a user request. IT should have performance management tools that can break down a transaction into its component parts; for example, tracking the number of milliseconds a user request spent in transit across the WAN vs. the SAN, how much time the database or Web server needed to process the request.

Tools need to be detailed enough to indicate response by application, and should clearly depict any variance from the response standard. In addition, monitoring tools should provide

multi-tier reporting so that high-level summaries are available for business managers and non-technical audiences, while IT can drill down to the level of technical detail necessary to pinpoint the source of a problem.

For organizations that outsource application development or hosting, CIOs and IT managers should hold their application providers and systems integrators accountable, Fabbi says. "I'd have performance SLAs as part of the acceptance criteria of my contract," he advises. For companies developing applications in-house, he recommends instrumenting your infrastructure with monitoring and other management tools and rigorously testing an application before rolling it out. "Make sure you test in realistic conditions," he adds.

Taming the elephant

The first step in addressing application performance problems is to recognize how common they are, and what their true impacts are. A problem that at first blush might seem a temporary inconvenience for internal users is likely to have a ripple effect on IT and customers as well, escalating to a significant problem.

Following a systematic process that includes thorough testing of new applications before deployment and monitoring of key portions of the enterprise infrastructure can help IT rapidly pinpoint performance problems, if not head them off entirely. CIOs and IT managers need to work to open lines of communication and modify processes so that information is more readily shared and finger pointing minimized. Ultimately, monitoring actual end user experience is the only sure way to know whether application performance falls within acceptable bounds.

The downsides of application degradation are many. Corporations risk reduced employee productivity as well as a loss of business, both in terms of current customers and growth opportunities. The upside of addressing performance issues ranges from more productive employees and satisfied customers to new revenue options and ways of doing business.

"It's about lower cost, higher revenue and higher quality service," Sturm says. "Ignoring performance problems is inviting the elephant to sit on you."

© 2004 Network World, Inc. All rights reserved.

[To request reprints of this special report contact networkworld@reprintbuyer.com](mailto:networkworld@reprintbuyer.com)