

Eliminating Roadblocks To Effectively Managing Application Performance

Jim Metzler

New technologies can help, but achieving the desired improvement involves factors that go beyond technology.

Over the last couple of years, assuring acceptable application performance has become a hot topic. As a result, a growing amount of attention has been paid recently to a wide variety of issues that affect the ability of the IT organization to manage application performance. The vast majority of this attention has focused on technologies that can be used to improve the performance of networks and applications. These technologies include compression, caching, protocol and application acceleration as well as server offloading.

Technology is clearly a critical component of a network and application management solution. However, the effectiveness of any network management solution has historically been affected significantly by issues that have no basis in the technology itself.

As will be demonstrated in this article, the primary impediments to effectively managing application performance have little to do with technology. In particular, the goal of this article is to examine some of the organizational dynamics that affect the ability of the IT organization to effectively manage application performance. We will also make recommendations for how IT organiza-

tions can overcome the roadblocks created by these dynamics.

Research Methodology

In the fall of 2006, a survey, sponsored by Network Physics, was given to the subscribers of Webtorials, and 215 responses were received. In order to gain additional insight into the topics covered by this report, five IT professionals were interviewed personally. As a general rule, IT professionals cannot be quoted by name or company in a report like this without having their input heavily filtered by their company.

With that in mind, Table 1 contains a brief listing of the people who were interviewed, along with the phrase that will be used in the report to refer to them.

The Importance Of Managing Application Performance

The conventional wisdom in the industry is that managing application performance is, in fact, an important task. In order to check on the validity of this wisdom, respondents were asked to indicate how the importance of managing application performance was viewed within their IT organizations (Figure 1).

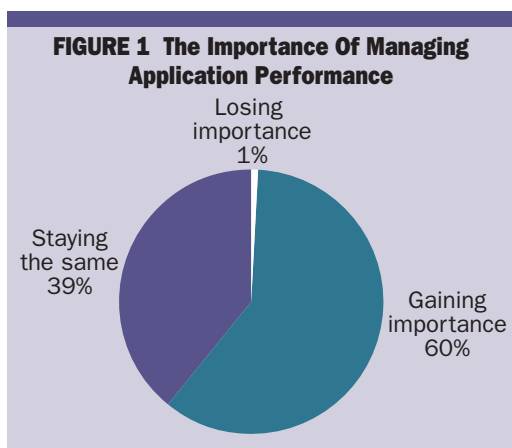
As the data in Figure 1 clearly indicates, the conventional wisdom is correct. Managing application performance is gaining in importance within the majority of IT organizations, and is losing importance in only a tiny percentage.

All five of the interviewees stated that within their IT organization, managing application per-

Dr. Jim Metzler has been an application developer, marketing manager, product manager, engineering manager, IT manager, research director and consultant. Along with Steven Taylor, Jim is a founding member of Kubernan, an analyst firm. He can be reached at Jim@Kubernan.com.

TABLE 1 List of Interviewees

Job Title	Industry	Reference Phrase
Manager of Network Services and Operations	Manufacturing	The Manufacturing Manager
Global Network Architect	Consulting	The Consulting Architect
LAN/WAN Integrator	Gaming	The Gaming Integrator
Enterprise Architect	Application Service Provider (ASP)	The ASP Architect
Global Infrastructure Engineering Manager	Automotive	The Infrastructure Engineering Manager



formance was gaining in importance. For example, The Gaming Integrator stated that managing application performance was gaining in importance in large part because the IT organization has recently deployed a number of new applications, and some of these applications did not initially perform the way that everyone had expected. As a result, “a lot of finger-pointing went on between the application group and the network group.”

According to The Consulting Architect, within the last couple of months, managing application performance has become the CIO’s number one priority. That CIO recently told our interviewee, “Managing application performance is the thing that I am getting the most flack on. We have to begin to work outside the silos and get away from the application and networking organizations pointing fingers at each other.”

Given the increasing importance of managing application performance, survey respondents were asked to indicate whether their company has a formalized set of processes for identifying and resolving application degradation (Table 2).

The data in Table 2 indicates that the vast majority of IT organizations either currently have formalized processes for identifying and resolving application degradation, or are working to develop these processes. This data is consistent with Figure 1. In particular, as the importance of managing application performance increases, a growing number of IT organizations have acknowledged the importance of having formalized processes for identifying and resolving application degradation.

The Infrastructure Engineering Manager said his IT organization does not currently have formalized processes for managing application performance,

but that they are working on it. He explained that they were motivated to develop these processes because application performance has become more of an issue recently, in large part because the IT organization is increasingly hosting applications in a single datacenter, and having users from all over the world access those applications. As a result, the parameters of the WAN that affect application performance (i.e., delay, jitter, packet loss) are more pronounced than they would be if there was less distance between the user and the application.

The Manufacturing Manager stated that as part of their processes for managing application performance, the IT organization offers an application service level agreement (SLA). He added that these SLAs focus primarily on the availability of the application, not on its performance. An exception is that they do have a performance SLA for their ERP system.

The Manufacturing Manager also stated that before the IT organization introduces a new application onto the network, they profile that application in a controlled environment in order to identify its requirements. They also perform a trend analysis on the performance of the network to see if they need to make any changes to the network in order to support the application.

The Organizational Quagmire

In order to better understand the organizational dynamics that affect the management of application performance, The Survey Respondents were asked to indicate which organizations have responsibility for the ongoing performance of applications in production. Their answers are contained in Table 3.

A number of conclusions can be drawn from


One CIO places a priority on getting beyond finger-pointing

TABLE 2 Existence of Formalized Processes

Response	Percentage of Respondents
Yes, and we have had these processes for a while	22.4%
Yes, and we have just recently developed these processes	13.3%
No, but we are in the process of developing these processes	31.0%
No	26.2%
Other	7.1%

TABLE 3 Group Responsible for Application Performance

Group	Percentage of Respondents
Network Group – including the NOC	64.6%
Application development group	48.5%
Server group	45.1%
Storage group	20.9%
Application performance management group	18.9%
Other	12.1%
No group	6.3%



Three of the four leading impediments are non-technical in nature

the data in Table 3. One obvious conclusion is that the organization that is most likely to have the responsibility for the ongoing management of application performance is the network group. Another conclusion is that typically, a number of organizations share that responsibility, and that the application development and server groups share a significant portion of that responsibility.

The Consulting Architect commented that application performance issues are usually found first by the end user, not the IT organization. He stated that once a problem has been identified, the task of determining the root cause of the problem bounces around within the IT organization and that, “It’s always assumed to be the network. Most of my job is defending the network.”

The Infrastructure Engineering Manager said that his organization has spent the last four years fighting against the assumption that if there is a problem with the performance of an application, the network must be at fault. During that time frame, his organization has added capacity to their global WAN and has reached the point that WAN bandwidth is not an issue.

He did point out, however, that explaining to either a user or an application developer the impact of latency on application performance “can be a very difficult conversation.”

The ASP Architect stated that within his company’s IT organization there is a group, referred to as the network performance team, that meets weekly. The purpose of this team is to deal with chronic performance issues. Multiple organizations compose the network performance team. The applications groups are not a member of the team, but are represented by the workstation group. The ASP Architect said that the workstation group is very knowledgeable about the company’s applications, and he further stated that the applications groups do get directly involved with the network performance team if there is an application performance issue that the workstation group cannot handle.

Impediments To Successful Application Management

One factor that limits the ability of IT organizations to successfully manage application performance is that many IT organizations regularly deploy applications with no thought to how those applications will perform over the WAN. This often results in the deployment of chatty applications—i.e., applications in which a given transaction requires tens or possibly hundreds of round trips, aka, application turns.

Survey respondents were asked the level of emphasis their IT organization places, during the development of an application, on how well that application will perform over the WAN. The question defined “moderate emphasis” to mean that application performance over the WAN gets as much attention as any other concern. Sixty-one

percent of respondents replied with answers that ranged from no emphasis to moderate emphasis.

The Manufacturing Manager said that his company develops some of their applications and acquires the rest from software vendors. In both cases, how well the application performs over the WAN “is not that important. What is important is how well the application meets the business need.”

The Gaming Integrator stated that his IT organization tends to acquire applications from a software vendor and that the organization does not tend to look at how well the application will run over the WAN prior to purchasing it. He also said his IT organization uses consultants to install the applications they acquire. As a result, when it comes to troubleshooting the performance of an application, few people inside the IT organization have a good understanding of the internals of the application.

The ASP Architect pointed out that an issue affecting application performance within his company is that the applications are developed on a high-speed LAN. As a result, processes such as an SQL query that worked well on the LAN do not necessarily work well over the WAN. He further said, “The applications group only needs to get burned that way a few times before they learn to write applications that are better suited to the WAN.”

Companies that have well-understood performance objectives for their business-critical applications can use these objectives to make decisions about the design of applications and of the network. Conversely, if companies do not have these objectives, application and network design decisions become highly arbitrary. More than half of respondents (55 percent) indicated that their company does not have these performance objectives.

The Infrastructure Engineering Manager stated that his organization does not have any targets for application performance. As a result, when it comes to prioritizing how they will respond to complaints of application degradation, “It comes down to who screams the loudest.”

As mentioned, The Manufacturing Manager indicated that his IT organization does offer an application SLA, but that it is primarily focused on the availability of the application, not the application’s performance. The Consulting Architect stated that they do not currently have well-understood performance objectives for their business-critical applications, and that this was part of what they were trying to accomplish. The CIO of that company highlighted the need for application performance objectives when he recently said, “What we have now is garbage. We do not have the right metrics.”

Respondents were given a list of possible impediments and were asked to indicate which two were the most significant barriers to effective application delivery. Table 4 indicates the four impediments that received the most responses.

TABLE 4 Impediments to Application Delivery

Impediment	Percentage
The processes that we have are inadequate	39.9%
The difficulty in explaining the causes of application degradation and getting any real buy-in	35.1%
The tools that we have are inadequate	32.7%
There is an adversarial relationship between the application development group and the rest of IT	23.6%

One way to look at the data in Table 4 is that three of the top four impediments to effective application delivery have little to do with technology. Another way to look at the data is that it balances the data in Table 2. In particular, the data in Table 2 indicates that the majority of IT organizations either already have formalized processes to identify and resolve application degradation, or are in the process of developing these. In contrast, the data in Table 4 indicates that in many cases, those processes are inadequate.

The Manufacturing Manager stated that his organization is behind where they would like to be relative to both the tools and the processes to manage application performance. He explained that his company is growing rapidly and the IT organization is never quite able to catch up to the business growth.

The Gaming Integrator indicated that the network organization had adequate processes for managing application performance, but that the rest of the IT organization does not. He added that the network organization has significant difficulty explaining the causes of application performance problems in part because they do not have the people within their organization who understand the details of the company's applications.

The ASP Architect stated that the infrastructure component of the IT organization has worked hard to improve their processes in general, and to improve their communications with the business units in particular. He pointed out that the infrastructure is now ISO certified and they are working on adopting an IT Infrastructure Library (ITIL) model for problem tracking. These improvements have greatly enhanced the reputation of the infrastructure organization, both within IT and between the infrastructure organization and the company's business units. It has reached the point that the applications development groups have seen the benefits and are working, with the help of the infrastructure organization, to also become ISO certified.

Given the importance of the concepts that comprise Table 4, survey respondents were asked a second, somewhat similar question. In particular, respondents were asked to indicate which of three choices presented the greatest difficulty relative to managing application performance. The choices were:

■ Diagnosing the situation to determine the source of the problem.

■ Identifying a solution once the problem had been diagnosed.

■ Getting support for the solution; i.e., funding and/or buy-in from other organizations.

The choice with the highest percentage of responses (42.4 percent) was choice #3: Getting support for the solution. This is further evidence that the primary impediments to effectively managing application performance have little to do with technology.

The ASP Architect provided insight into the challenges of determining the source of an application performance issue. He said, "We used to have a real problem with identifying performance problems. We would have to run around with sniffers and other less friendly tools to troubleshoot problems. The finger-pointing was often pretty bad."

He went on to say that in order to do a better job of identifying performance problems, the IT organization developed some of its own tools. These tools are currently used by the traditional IT infrastructure groups as well as by some of the application teams. He went on to say that the reports generated by these tools helped to develop credibility for the networking organization with the applications development organization.

The ASP Architect also provided insight into the difficulty of getting support for the solution once the problem has been diagnosed. Part of the problem his organization faces is that his company uses a lot of third-party applications. If it is determined that the source of the application performance problem is a badly written application from a third party, then his organization has to work with the software vendor to convince that vendor to modify the code.

Another situation in which it can be difficult to get support for a solution is when there is a problem with an application that does not directly affect the ASP's clients, but does affect multiple business units within the ASP. The problem in this case is the complexity of getting multiple business units to agree on a solution and then setting a high enough priority to ensure that the solution gets implemented quickly.

To exemplify the latter point, The ASP Architect described the applications group and said, "It is not like the applications group has a lot of free time on their hands. They are under a lot of pressure to deliver a product that helps the company make money."

Applications groups don't have a lot of free time—they have to help the company make money

Datacenter consolidation may require network optimization

Success Rate Of Key Application Management Tasks

The overall process that IT organizations use to manage application degradation is composed of many individual sub-processes or tasks. As such, the success that IT organizations achieve with their overall process for managing application performance will not be any higher than the lowest success rate of any of the tasks that comprise the process.

Survey respondents were given a number of these tasks and asked to indicate if their organization:

- Performs that task today
- Performs that task well today

For each task, the success rate was quantified by dividing the number of IT organizations that currently perform the task well by the number of companies that currently perform the task. Table 5 contains the answers to this question.

To understand how to interpret Table 5, consider the task of discovery. As shown, 55 percent of IT organizations currently perform this task, but only 41.5 percent of these are successful with it. Taken together, this indicates that only 23 percent of IT organizations perform discovery and perform it well.

The Infrastructure Engineering Manager said that his organization pays some attention to discovery, from both a proactive and reactive stance. He did point out, however, that his organization does nothing relative to the proactive alerting of network and application performance issues.

The Gaming Integrator said that discovery used to be a weakness for his organization, but that they recently acquired a new tool that greatly increases their ability in this area. In contrast, The Consulting Architect said, “On a scale of 1 to 5, I would give discovery a 1 or a 2.” He went on to state that his organization only does discovery in a reactive fashion, as part of troubleshooting a problem.

Plans To Enhance Application Management

The same set of tasks that were used to create the success rate metric were shown to the survey respondents, and they were asked to indicate if their organization intends to either implement that task or do it better sometime in the next year (Table 6).

The way to interpret the data in Table 6 is that 35.1 percent of The Survey Respondents indicated that over the next year their organization would either begin to implement discovery processes (i.e., who is on the network and what are they doing?) or would attempt to get better at this process.

The Infrastructure Engineering Manager said that his organization has implemented very stringent measures to block non-approved and inappropriate applications such as Internet radio. He did point out, however, that sometimes this traffic still gets through.

The ASP Architect said that quantifying the impact of deploying optimization solutions is not something they currently do, but is something they need to do.

The Gaming Integrator said that his company is interested in consolidating servers out of branch offices and into a centralized datacenter. As a result, his IT organization is in the middle of a trial to measure the impact of implementing a network optimization solution that is intended to overcome the issues related to server consolidation.

The ASP Architect said that the tools his organization developed to do a better job of identifying performance problems are also helpful in terms of measuring the performance of an application before and after a major change. He indicated that his organization intends to develop additional tools, but that the process of developing these tools is laborious. To add to the difficulty, the development process requires involvement of the business units, as they are the ones who understand which components of a complex application

TABLE 5 Success Rate

Task	Performs The Task	Success Rate
Discovery—who is on the network and what are they doing	55.0%	41.5%
Capacity Planning	51.6%	56.3%
Measuring the performance of an application before and after a major change	41.7%	52.3%
Isolate the problem source—network, servers, application, etc.	56.0%	61.1%
Drill down into the problem source once the source of the problem has been isolated	49.7%	69.9%
Quantify the impact of network parameters (loss, delay, jitter) on the performance of an application	38.0%	71.9%
Quantify the impact of optimization (caching, compression, protocol acceleration) on the application	31.7%	65.9%
Proactive alerting of network and application performance issues	47.3%	56.4%
Traffic Management/QOS	44.3%	56.4%
Baselining the performance of the network	42.9%	58.3%
Identifying non-approved and inappropriate applications	41.1%	61.2%

are important and what is an acceptable level of performance for those components.

The Infrastructure Engineering Manager said that his IT organization is trying to take some proactive steps to reduce the number of times that an application degrades. For example, they are in the process of developing a set of best practices around a wide range of common IT tasks, such as the best way to access a database. The intention is that these best practices will drive an approach to performing key IT tasks that is common across the disparate applications groups and the infrastructure organization, and which will also reduce the causes of application degradation.

A second step is that they are trying to do a better job of testing an application before deploying it. He said that this testing could result in minor changes to the application, or could result in the deployment of some sort of network optimization techniques; i.e., caching or compression.

A third step is that they are working more closely with the application development teams—around both the selection of tools such as content management systems, as well as the actual development of applications. Relative to tools, his organization is getting involved early in the selection cycle. Their goal is to identify how well each tool runs over a WAN and to discourage the adoption of any tool that performs badly over the WAN.

His group is also trying to get involved early in the application development cycle so that they can exert greater influence over how applications get developed, with the goal of eliminating most of the factors that cause an application to run badly over the WAN. He said the primary factor limiting his success in these endeavors is having enough people in his organization that have a deep understanding of software and the factors that affect application performance.

Summary And Recommendations

Many factors are making management of application performance important to virtually all IT

organizations, with this issue gaining in importance in more than half the IT organizations. One of these factors is that additional applications are continually being deployed on the network. This includes voice, Internet commerce and business critical applications such as ERP. Another factor is that IT organizations are increasingly hosting applications in a single datacenter, and users are accessing those applications from all over the world. As a result, the impact of the WAN on application performance is more pronounced.

The approach to managing application performance that has the best likelihood of success is the approach that was pointed out in interviews with The Consulting Architect. The CIO in his company has made managing application performance the IT organization's number one priority, and has clearly said, "We have to begin to work outside the silos and get away from the application and networking organizations pointing fingers at each other."

Regrettably, it is relatively rare to have a CIO set that clear an edict relative to managing application performance and eliminating organizational stove pipes. As a result, in most cases the dual tasks of improving the management of application performance and minimizing organizational stress is often more of a bottom-up than a top-down initiative. In addition, since it is common to have the network deemed to be the root of the problem until proven otherwise, it is typically up to the network organization to lead this effort.

The research presented in this report gave clear insight into the factors that were the most serious impediments to effective application delivery. Those factors were:

- Inadequate processes.
- The difficulty of getting buy-in to a proposed solution.
- Inadequate tools.
- The adversarial relationship that sometimes exists between the application development group and the rest of IT.

Network groups are trying to get involved earlier in the application development process

Task	Percentage
Discovery—who is on the network and what are they doing	35.1%
Capacity Planning	29.6%
Measuring the performance of an application before and after a major change	41.7%
Isolate the problem source – network, servers, application, etc.	18.7%
Drill down into the problem source once the source of the problem has been isolated	23.0%
Quantify the impact of network parameters (loss, delay, jitter) on the performance of an application	40.0%
Quantify the impact of optimization (caching, compression, protocol acceleration) on the application	48.2%
Proactive alerting of network and application performance issues	38.2%
Traffic Management/QoS	41.5%
Baselining the performance of the network	35.7%
Identifying non-approved and inappropriate applications	39.9%



Creating effective processes will give IT greater credibility

The ASP Architect provided insight into how his organization transitioned from finger-pointing between the application development group and the network group, to an environment characterized by cooperation. One of the steps his organization took was to implement tools that provided detailed, accurate insight into the source of application performance issues. They also improved their internal processes.

Over time, both of these steps resulted in an increase in the credibility of his organization with the application groups and also with the company's business units. These steps also resulted in the elimination of the adversarial relationship that had existed between the application development groups and the network group, and has made it somewhat easier to get buy-in for a proposed solution to an application performance problem.

Other network organizations looking to make a similar transition should choose tools that allow the organization to implement the type of effective processes that will enable them to successfully manage application performance. Many of these processes were discussed in this article, including:

- **Discovery**—The identification of parameters such as who is using the network, the applications that they are using, as well as how much utilization they are generating.

- **Trending**—The identification of how the network is currently being utilized and how that utilization has changed over time.

- **Quantifying WAN impact**—The quantification of WAN latency and packet loss, as well as the impact that these factors have on application performance.

- **Quantifying the change in application performance**—Measurement of the performance of key applications before and after a major change.

- **Isolating the problem source**—Identification of the likely source of application degradation: i.e., network, server, application or database.

- **Troubleshooting the problem**—The capability, once the source of the problem has been identified, to drill down into the actual cause of the problem.

- **Proactive alarming**—The setting of alarms that indicate that an IT resource has reached a threshold.

- **Application characterization**—The characterization of an application in order to identify which type of optimization technique is likely to improve the performance of the application. It also refers to the ability to quantify the impact of implementing an optimization technique.

There are frameworks, such as IT Infrastructure Library (ITIL, see *BCR*, December 2006, pp. 49–52), that network organizations can use to help redesign their key processes. However, the choice of whether to use a framework is less important than the decision to gain credibility by implementing effective processes based on tools that provide detailed, accurate insight into the source of application performance issues. As noted, this credibility eases the burden of getting buy-in from other organizations, minimizes the amount of finger pointing, reduces the organizational stress, and in some instances can lead to the deployment of applications that are designed to run well over the WAN □