

Rapid Problem Identification: A Cornerstone of **Effective Application Delivery**

By Jim Metzler

SPONSORED BY



Rapid Problem Identification: A Cornerstone of Effective Application Delivery

Introduction: The Need for Rapid Problem Identification

As recently as a few years ago, application delivery was not an important topic for most IT organizations. Over the last few years, however, this has changed dramatically. What has also changed is the role of the network operations group. The traditional role of the network operations group had been to provide fault management in an environment in which the diversity of endpoint devices centered on the type of desktop or laptop. Now fault management is simply a part of the contemporary network operations group's charter. This group must also support effective application delivery in an environment in which there are myriad endpoint devices, including desktop computers, laptops, VoIP phones, PDAs, IP video cameras, network-attached storage..

As this white paper will demonstrate, the mean time to repair (MTTR) associated with application delivery is notably longer than the MTTR that is associated with fault management. This extended MTTR, which is caused in large part by the lack of effective management tools, impacts business in a number of ways, including reduced revenue, lowered productivity, and lowered perceived value of the IT organization.

One of the goals of this white paper is to demonstrate the need for a trusted solution that can quickly and effectively perform rapid problem identification (RPI). **Throughout this white paper the phrase, *rapid problem identification*, or RPI refers to the ability to both identify that the performance of one or more applications is degrading, and to identify the likely cause of the degradation.** With this goal in mind, this white paper compares the day-to-day challenges that network operations teams face when trying to resolve the traditional issues associated with fault management and the challenges of trying to resolve the complex set of issues associated with application delivery. Another goal of this white paper is to detail the key characteristics of a solution that enables network operations teams to overcome these challenges and to describe how a solution from Xangati has helped a number of IT organizations improve their approach to application management.

To quantify the changes that are impacting network operations groups, Kubernan¹ has performed extensive market research. Throughout this white paper, that market

¹ Kubernan is an analyst and consulting joint venture of Jim Metzler and Steven Taylor.

IT Innovation Report

Published By

Kubernan
www.Kubernan.com

Cofounders

Jim Metzler
jim@ashtonmetzler.com

Steven Taylor
taylor@webtutorials.com

Design/Layout Artist

Debi Vozikis

Copyright © 2007

Kubernan

For Editorial and Sponsorship Information

Contact Jim Metzler
or Steven Taylor

Kubernan is an analyst and consulting joint venture of Steven Taylor and Jim Metzler.

Professional Opinions Disclaimer

All information presented and opinions expressed in this IT Innovation Report represent the current opinions of the author(s) based on professional judgment and best available information at the time of the presentation. Consequently, the information is subject to change, and no liability for advice presented is assumed. Ultimate responsibility for choice of appropriate solutions remains with the reader.

research is referred to as *The Kubernan Market Research*. To provide additional context for this white paper, four IT professionals were interviewed. Table 1 contains a brief listing of the people who were interviewed, along with the phrase that are used in this white paper to refer to them.

Table 1: The Interviewees

| Job Title | Industry | Reference Phrase |
|-------------------------|----------------|-----------------------------|
| CIO | Non-Profit | The CIO |
| IS Manager | ILEC | The IS Manager |
| Senior Engineer | Medical | The Medical Engineer |
| Senior Network Engineer | Transportation | The Transportation Engineer |

The Changing IT Environment

The IT environment is undergoing fundamental change. One component of the changing IT environment is structural and relates to the type of endpoints that must be supported, as well as the applications that are accessed using these endpoints. For example, while desktop computers and laptops still function as endpoint devices, in many instances so do other devices including VoIP phones, PDAs, IP video cameras, network-attached storage, etc. Relative to the growth in these new types of endpoints, The CIO noted that his organization has seen "an exponential growth in the number of handheld devices." He added that he believes that "the huge growth in the number of handheld devices will continue for at least another year."

These new endpoint devices are used to access a wide variety of applications. For example, The IS Manager stated that his organization has seen substantial growth in the deployment of PDAs and smartphones, and that today's primary use of these end devices is to access traditional applications, such as messaging and calendaring. The CIO noted that within his organization these new endpoint devices are being used to support a variety of new applications. As noted, one of those applications provides detectives (in the sheriff's department) the ability to access information in the National Crime Information Center before entering a crime scene. Another application allows the county to dynamically assign work to its remote maintenance workers and road crews.

Another component of the changing IT environment is the emerging role of the network operations group. Historically this group has focused their attention on the availability of the various components (i.e., switches, routers, WAN links) that comprise the network. While that is still an important task, **in the majority of organizations the network operations group has assumed a new and extremely challenging task—to ensure successful application delivery.** The Kubernan Market Research quantified this shifting role by surveying over three hundred IT professionals who indicated that 65% of the time the network operations group had ongoing responsibility for application performance.

The Increasing MTTR

In traditional fault management the MTTR problem is typically a couple of hours. Given the difficulty with both identifying whether an application is degrading, and if so, what the source of the degradation is, the MTTR that is associated with application management can be quite lengthy. In many cases, the MTTR associated with first identifying that an application performance exists, and then identifying the source of that problem, is measured in days or weeks.

The CIO stated that it is somewhat common in his organization to have the MTTR associated with managing application performance measured in weeks. The IS Manager stated that it was common to have the MTTR an application performance issue measured in weeks, and that on some occasions it can be measured in months.

Another reason that the MTTR that is associated with application management is so lengthy is that the current set of management tools does not provide effective insight into application performance, and a result, IT organizations are often left guessing if there is a performance issue. In addition, if it is determined that there is an application performance issue, IT organizations are left guessing as to the source of the issue. To put this in context, The CIO was asked what he thought of the ability of the current generation of tools to manage application performance. His response was "I feel that they are just in their infancy."

The IS Manager stated that, "The network is assumed to be the cause of application degradation until it is proven not to be." He added that because of this misconception, trouble tickets related to application degradation are usually assigned first to the network technician. As will be shown in a subsequent section of this white paper, the network is usually not the cause of application degradation. However, because the network group must spend time showing that it is not the network, before the real cause of the degradation can be identified, the MTTR that is associated with application delivery is lengthened.

Identifying Application Degradation

In traditional fault management, the network operations group is concerned with managing the availability of a small set of network components, typically switches and routers. One of the factors that make identifying a fault with a switch or router relatively easy is that a fault in a switch or router often leads to an outage, which is readily noticeable.

As noted *RPI* refers, in part, to the IT organization's ability to identify the performance of one or more applications is degrading. **Identifying that the performance of an application is beginning to degrade is significantly more difficult than is identifying that a switch or router is unavailable.** The extra difficulty associated with managing application performance is highlighted in The Kubernan Market Research. One of the most significant results uncovered is the dramatic lack of success that IT organizations have relative to identifying application degradation. In particular, Kubernan asked 345 IT professionals the following question: "If the performance of one of your company's key applications is beginning to degrade, who is the most likely to notice it first—the IT organization or the end user?" **Seventy-three percent of the survey respondents indicated that if the performance of an application is beginning to degrade, that it is first noticed by the end user and not by the IT organization.**

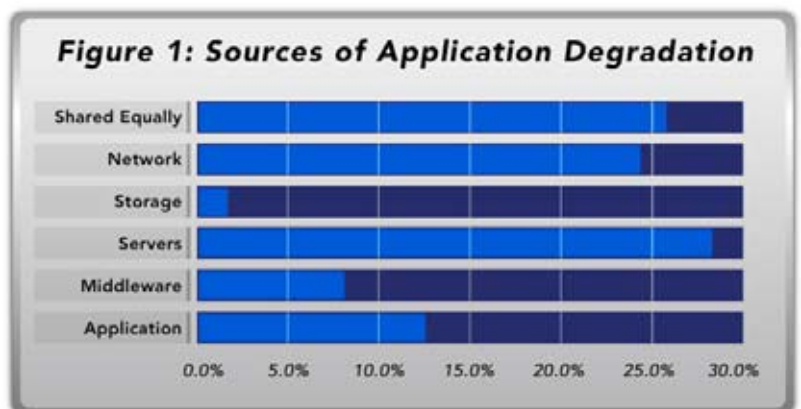
Relative to the question of who notices application degradation first, The IS Manager stated, "unfortunately it is the end user" and that "without an effec-

tive tool, even establishing that there is a problem is highly subjective." He added that since the IT organization does not use the majority of the corporate applications themselves, they have a tendency to discount the first few calls to the help desk as "the user is simply complaining."

The CIO agreed with the statement that in the vast majority of instances the end users notice application degradation before the IT organization does. He stated that in his organization the fact that **the IT organization does not know when an application has begun to degrade has lead to the perception that IT is "a bunch of bumbling idiots."** He further revealed that this situation has also fostered an environment in which individual departments have both felt the need and been allowed to establish their own shadow IT organizations.

Identifying the Source of Application Degradation

In addition to referring to the ability of the IT organization to identify that the performance of one or more applications is degrading, *RPI* also refers to the ability to identify the likely cause of the application degradation. The Kubernan Market Research provides insight into the complexity associated with identifying the likely *cause* of the application degradation. As part of that research, Kubernan asked the survey respondents to indicate what component of IT was the biggest cause of application degradation. Figure 1 summarizes their answers. In Figure 1, the answer *shared equally* means that multiple components of IT are equally likely to cause application degradation.



One of the conclusions that can be drawn from the data in Figure 1 is that **virtually any component of IT could be the cause of application degradation**. Another conclusion that can be drawn from the data in Figure 1 is that while most organizations assume that the network is the source of any application degradation—that is only true in a minority of instances.

The IS Manager was previously quoted on the difficulty of identifying that an application was degrading. He further commented that once his organization is convinced that there is a performance problem, identifying the cause of that problem can be a lengthy process because "it seems like a million different components of IT could be at fault."

The data in Figure 1 also speaks to the technical and organizational complexity associated with managing application performance. In particular, unlike traditional fault management, the management of application performance requires that IT organizations transcend both technical and organizational boundaries. Regrettably, most IT organizations operate in a stove-pipe fashion, i.e., the organizations that comprise IT do not have common tools and solutions, processes, terminology, and goals.

The CIO stated that, like most IT organizations, his organization is stove-piped and that "communications across IT groups has been a struggle." He added that having a solution that all of the groups trust is a very effective way to encourage effective communications amongst the groups.

The Mandate for Change

The preceding sections highlighted that within the majority of organizations, the network operations group is responsible for the ongoing management of application performance. Those sections also highlighted out the fact that in the many instances the MTTR application performance issues is measured in days or weeks. As The CIO called out, the lengthy MTTR application performance issues often leads to a situation where the creditability of the IT organization is called into question.

Network operations groups will not be regarded as being successful as long as they continue with the current approach to application management. In order

to be successful, network operations groups must adopt a fundamentally new approach to application delivery that includes a strong emphasis on RPI. The goal of this new approach is to enable the network operations group to reach a point where the MTTR application performance issues are measured in hours.

A key component of this new approach to application management is the migration away from using the current generation of management tools and towards using the emerging generation of RPI solutions. One of the limiting characteristics of the current set of management tools is that they require significant manual intervention, which can be costly, error prone, and add to the MTTR. The Transportation Engineer used sniffers to illustrate some of the limitations of the current generation of management tools. He said that sniffers provide a real value but, in order to be useful, his organization must direct the sniffer to the location that experiences a problem. This takes time and often when they do get the sniffer running, the application degradation is not occurring, which means letting the sniffer continue to run until the problem reappears. All of this, he emphasized, adds to the MTTR.

Another one of the limiting characteristics of the current generation of management tools is that they are narrowly focused. In most cases, the narrow focus of the current generation of management tools mirrors the stove-piped IT organization and encourages continuing with a stove-piped approach to application management. As part of the current stove-piped approach to application management, many networking groups use a management tool to set thresholds for a metric, such as the utilization of their WAN links. IT organizations that use this approach to managing network and application performance implicitly make two assumptions:

- If the network is heavily utilized, the applications are performing poorly.
- If the network is lightly utilized, the applications are performing well.

The first assumption is often, but not always true. For example, if the company is primarily supporting email or bulk file transfer applications, heavy network utilization is unlikely to cause unacceptable application performance.

The second assumption is often false. It is quite possible to have the network operating at relatively low utilization levels and still have the application perform poorly. An example of this is any application that uses a chatty protocol² over the WAN. In this case, the application can perform badly due to a large number of application turns, even though the network is exhibiting low levels of delay, jitter, and packet loss.

Rapid Problem Identification Solutions

As noted, network operations groups will be unable to successfully manage application performance using the current set of management tools. To successfully manage application performance, network operations groups need to deploy a solution that is focused directly on reducing the MTTR for performance issues through an emphasis on rapid problem identification. Hereafter, these solutions will be referred to as RPI solutions.

The sub-topics that follow describe the primary characteristics that distinguish an RPI solution from the current set of management tools.

Automation

Automation refers to a wide range of functionality. For example, an effective RPI solution should be able to automatically discover the endpoints, the topology, and the applications—both those applications that are sanctioned and those that are not.

Automation also enables an IT organization to implement an approach to application management in which the RPI solution is continually aware of the real-time environment. In this approach, the RPI solution locates the source of application performance issues and IT professionals are focused not on staring at computer screens, but on quickly resolving the identified application performance issues.

Focus on Applications and Endpoints

To identify application degradation before the end user does, network operations groups must deploy

² A chatty protocol uses tens or hundreds of application turns for a given transaction.

a solution that can measure the behavior of every endpoint and every application. A solution with this functionality can provide the network operations group with a holistic view of the inter-relationships between the endpoints and the applications. This is in contrast to the current set of management tools that typically focuses on the individual components of the IT infrastructure; i.e., switches, routers, application servers, and WAN links.

The information gathered by the current set of management tools is extremely important. As was previously described, however, when network operations groups focus exclusively on metrics, such as the utilization or latency associated with a WAN link or a particular application, they are incapable of understanding the source of complex application performance problems. The IS Manager stated that the tools that they had been using provided information only at the component level. He added that this "does not tell us about the experience of the user" and that this approach "makes it difficult for us to make any intelligent conclusions about what was happening in the network."

The CIO said that they are attempting to migrate away from a stove-piped approach to management and move towards an end-to-end approach based on the experience of the end user. He stated that his organization has recently deployed the Xangati solution and that his goal is to use that solution to quickly reach the point where the IT organization notices half of the instances of application degradation before the end user does. He believes that if he can achieve this goal he can turn around the perception of the IT organization as being "bumbling idiots."

Contextual Awareness of the Application Delivery Infrastructure

The preceding section pointed out that when network operations groups focus exclusively on metrics, such as the utilization or latency that is associated with a WAN link or a particular application, they are not in a position to understand the source of complex application performance problems. Analogously, focusing just on the behavior of every endpoint and application does not fully position the network operations group to understand the cause of application degradation. **What is needed is for the solu-**

tion to understand the behavior of every endpoint and application, and simultaneously be thoroughly aware of the behavior of the application infrastructure. For example, an effective RPI solution must be able to note changes in the application delivery infrastructure in case these changes have a resulting effect on the applications and the endpoints.

The Transportation Engineer highlighted the need for this functionality when he stated that what he needs is not additional management tools, but a solution that, "tells me in real time that there is a problem and that points me in the direction of the source of the problem."

Profiling

The preceding section mentioned the current practice of setting thresholds for a metric, such as the utilization of an organization's WAN links. To not be inundated with alarms, the network operations group typically sets thresholds at a high value. As a result, the network operations group misses the vast majority of events, because the severity of these events is less than the conservative thresholds that are typically established.

To avoid this limitation, **an RPI solution must be able to understand the normal behavior of endpoints, applications, and the application delivery infrastructure. In addition, an RPI solution must be able to compare normal behavior to real-time behavior in order to identify abnormal behavior.** This capability is often referred to as *profiling*. According to The IS Manager, "detecting variances from the norm is important." Because of the importance that he places on this capability, his organization has recently deployed the Xangati solution. He said that the profiling capabilities of this tool identifies variances and helps to notify him if something out of the ordinary happens.

The Transportation Engineer commented that they have only had the Xangati solution deployed for a couple of weeks. He stated that in that timeframe the solution has helped them to better manage their WAN bandwidth by identifying applications, such as database backups that are meant to occur in off hours that were occurring during peak times.

Broad Scope

As noted, the current set of management tools is typically used for a narrowly defined function, such as monitoring the utilization of an organization's WAN links. However, as previously described, a wide variety of factors can cause the performance of an application to degrade. Another limitation of the current set of management tools is that the information produced by a given tool is typically not available to other tools.

To determine the source of application degradation, **an effective RPI solution needs the ability to profile multiple aspects of the applications, the endpoints, and the application delivery infrastructure.** At a minimum, the application solution should be able to establish profiles around performance characteristics, such as:

- Bit rate
- Packet rate
- Burstiness

To identify the likely cause of the degradation, however, an effective RPI solution must generate profiles that are much broader in scope than just performance. These profiles must also include factors, such as the number of endpoints normally communicated with, which applications are most commonly used, and at what performance levels are they used during relevant time periods; i.e., time of day and day of week.

The Medical Engineer stated that the primary value of the Xangati solution is its ability to profile applications and the supporting infrastructure. He further stated that once you have an accurate profile of your environment, getting alerted when there is behavior that is outside of the norm is "insanely useful."

Correlating Symptoms to a Core Problem

As noted, a major part of the challenge facing network operations groups is that identifying the source of application degradation is difficult to do as long as IT organizations employ a stove-piped approach to managing application performance. **A stove-piped approach to application management typically leads an IT organization to confuse a symptom of the**

problem with the cause of the problem. For example, an IT organization might determine that a sluggish e-mail server was the result of some problem with the server itself, when in fact an intermittently failing DNS server was causing the e-mail server to appear less than responsive.

An effective RPI solution must be able to guide the network operations group to differentiate between a symptom and a core problem. Therefore, the solution should have the capability to see the relationships between the various symptoms across dimensions, such as time, location, and the applications used. A solution with this functionality will allow the network operations group to quickly filter to the source of the problem and demonstrate how the problem source lead to a number of symptomatic effects at the endpoint and/or the application level.

As previously noted, The CIO stated that one factor that drives up the MTTR application performance issues is that with the current set of tools the problem is often misdiagnosed and sent to the wrong group for action. As a result of deploying the Xangati solution, he intends to reduce the number of times that a trouble ticket is misrouted, and to significantly reduce the amount of time it takes to resolve a problem. One of his goals is to reach a point where problems, such as minor configuration errors, are quickly resolved and his organization can focus on the truly challenging problems.

Summary and Call to Action

As recently as a few years ago, application delivery was not an important topic for most IT organizations. Over the last few years, however, this has changed dramatically. What has also changed is the role of the network operations group. In particular, in the majority of organizations, the network operations group is now responsible for the ongoing performance of applications.

Some of the key characteristics of the current approach to application management are that:

- The vast majority of times that an application is degrading, the degradation is noticed first by the end user and not by the IT organization

- Once it has been determined that an application is degrading, the IT organization is often not aware of the source of the degradation
- The MTTR poor application performance is often measured in days or weeks

Network operations groups will not be regarded as being successful as long as they continue with the current approach to application management. To be successful, network operations groups must adopt a fundamentally new approach to application delivery that includes a strong emphasis on RPI. The goal of this new approach is to enable the network operations groups to get to a situation where the MTTR application performance issues are measured in hours.

The current set of network management tools is an impediment to successful application delivery. In order to be successful, network operations groups must adopt a new generation of RPI solutions. The key characteristics of these solutions are:

- Automation
- Focus on Applications and Endpoints
- Contextual Awareness of the Application Delivery Infrastructure
- Profiling
- Broad Scope
- Correlating Symptoms to a Core Problem

Each of the characteristics is necessary to ensure RPI, which as outlined in this white paper, is a cornerstone of effective application delivery.