WAN Vicious Applications

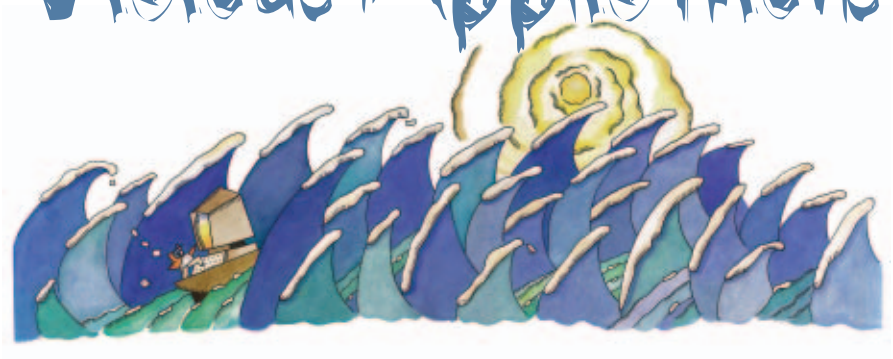

Jim Metzler
Ashton, Metzler & Associates
jim@ashtonmetzler.com

Introduction

Over the last few years, application delivery has become a hot topic. However, the media coverage of this topic has been focused primarily on a wide range of technologies that can improve the performance of networks, servers, and ultimately, of applications. This includes technologies such as caching, compression, protocol and application acceleration, and server offloads.

For many companies, implementing technologies that enhance the performance of their IT resources is a good thing. However, unto themselves these technologies don't do anything about the fact that today in the vast majority of cases if an application is degrading, it is noticed first by the user and not by the IT organization. That simple fact drives home a strong message: assuring acceptable application delivery is about more than just enhancing the performance of IT resources. Successful application delivery requires that IT organizations are able to identify the applications that are running on the network and are also able to ensure the acceptable performance of the applications that are relevant to the business, while controlling or eliminating applications that are not relevant.

The August 2006 IT Impact Brief looked at the issue of network misuse. In that brief I defined network misuse as those instances in which a significant portion of the company's network is being consumed by supporting traffic that is unauthorized and inappropriate. The brief described how often network misuse occurs and gave examples of some of the most unusual and outrageous examples of network misuse as reported by the NetScout community. The brief also described the variety of passive and active measures that IT organizations have implemented to control misuse and concluded that the first step in controlling network misuse is to have the visibility to understand what applications are running on the network.

I want to use this IT Impact Brief to discuss WAN vicious applications and the challenges that these applications present. Throughout this brief, the phrase WAN vicious will refer to an application that runs well over the LAN but that runs poorly over the WAN. Based on my definition of network misuse, it is a stretch to label WAN vicious applications as being an example of network abuse because in most cases they are very definitely authorized by senior management. However, a strong case can be made that WAN vicious applications make inappropriate use of the WAN.

Whether or not you think of WAN vicious applications as network misuse, this class of application is one more example of why organizations that have the mandate to assure acceptable application delivery need to have the ability to understand what applications are running on the network.

Research Methodology

In August we surveyed the NetScout community on a variety of topics, including some of the key components of their application development cycle. I also interviewed three members of the NetScout community in order to get their insight into these topics. One of the interviewees is the director of network infrastructure and support for a medical organization, another is a network engineer in the automotive industry, and the third is an IT management engineer in a company that provides network monitoring services. These interviews will be referred to respectively as The Infrastructure Director, The Network Engineer and The Management Engineer.

Why This Matters

One of the major reasons why we see so many WAN vicious applications is that in the typical application development environment the focus is on delivering the promised software functionality on time, on budget and with relatively few bugs. In the majority of cases, there is little if any focus during the design and development of an application on how well that application will run over a WAN. This is true whether the application is developed internally by the IT organization or acquired from a software vendor.

The Infrastructure Director commented that his organization tries to analyze any application that one of the departments is thinking about buying. They do this even though he acknowledged that with very few exceptions his organization couldn't tell any of the departments what they can and cannot run over the network. The approach that they take with these departments is to say, "If you deploy this application, we will have to double the bandwidth." That approach is effective because in most cases, the department absorbs the cost of the additional bandwidth.

There are many examples of WAN vicious applications. I came across one such application on a recent consulting engagement. That application sends a 2MB file from the Web server to the remote client just so that the client could open the file and extract a 10-digit code.

A class of applications that can be considered to be WAN vicious is chatty applications. The phrase chatty application refers to an application that requires tens or even hundreds of application turns just to complete a single transaction. There are two sources of the delay that are associated with each application turn. One source is the round-trip latency of the WAN - that is the amount of time that it takes a bit to travel from the origin to the destination and back again. In general, there is little that an IT organization can do to reduce the round-trip latency. The second source of delay is the transmission delay associated with sending whatever information is required.

To put this in context, assume that an application requires one hundred application turns to complete a particular transaction. If that application is run over a LAN that has ten microseconds of round-trip delay, then the total round-trip delay for that transaction is 1 millisecond - which is clearly not noticeable. However, if the same application is now run over a WAN that has 120 ms of round-trip delay, the total round-trip delay is 12 seconds - which clearly is very noticeable.

The Network Engineer said that his organization has a problem with chatty applications because they host their exchange servers in centralized data centers. As a result, they run CIFS (Common Internet File System) over the WAN. CIFS is well known to be a very chatty protocol. The Infrastructure Director stated that they used to have problems with chatty applications but that they adopted Citrix, which allowed them to reduce the impact of these applications. He did point out that they recently migrated the user interface of a couple of their applications from a thick client to a web browser. According to The Infrastructure Director this migration drove the need for a lot more bandwidth.

A WAN vicious application like this can be very difficult to troubleshoot. In most instances if a user of this application complained about poor performance, the implicit assumption would be that the network was somehow at fault. However, tests of the network could well show that the network was exhibiting very low levels of delay, jitter and packet loss. In too many cases, this would then result in the issue bouncing around the IT organization as each group struggled to show that it was not their problem. The only solution to problems like this is for the IT organization to have visibility into the primary applications that run over the network and at least some understanding of how these key applications perform.

It is reasonable to ask, "Why would anybody develop an application that takes 100 turns for a single transaction?" One answer is that the developer may not have known how chatty the application was becoming. In particular, some of the more popular content management systems are inherently chatty. Hence, if the application developer uses one of these systems, the application will be chatty (a.k.a., WAN vicious) and there is little that the developer can do to prevent that other than to use a content management system that is less chatty.

Survey Data

The respondents to the August survey were asked of their IT organization set performance requirements on the performance of new applications (Table 1). This is an important question because 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. They can also use these objectives to drive more effective monitoring of network and application performance. Conversely, if companies do not have these objectives, application and network design decisions become highly arbitrary and monitoring is done based on subjective criteria.

Table 1: Setting Performance Requirements

	Yes	No	Don't Know
Are bandwidth requirements set	52%	35%	13%
Are response time requirements set	55%	34%	12%

The data in Table 1 is encouraging. In particular, over half of the survey respondents indicated that their company sets bandwidth and response time requirements for new applications.

The Management Engineer stated that while his company does not set specific bandwidth requirements for new applications, they do have a guideline that states that applications will use as little bandwidth as possible. The Network Engineer echoed this thought. He said that the vast majority of his company's WAN was comprised of links that run at 512 Kbps or less. He said that the application developers are aware of that and try to develop applications that do not consume large amounts of bandwidth.

The survey also asked the respondents if their IT organization set any other requirements on network and application performance. A number of the survey respondents indicated that their organization sets latency requirements for the network. Apparently

influenced by the types of issues caused by chatty applications, one survey respondent indicated that his organization sets a limit on the number of TCP turn-arounds per transaction.

The survey respondents were asked whether or not their organization tested applications prior to deploying them and were also asked if they piloted an application prior to putting it into full-scale production. Their answers are contained in Table 2.

Table 2: Testing and Piloting of Applications

	Yes	No	Don't Know
Are applications tested in the lab before being put on the production network?	81%	12%	8%
Are network engineers included in the testing process?	60%	30%	10%
Are applications piloted before being put into production use?	81%	10%	10%
Are network engineers included in the pilot process	59%	30%	12%

It is very encouraging to see how often IT organizations test an application in a lab, or conduct a pilot prior to putting the application into full-scale production. Having said that, it is a little puzzling that 81% of companies test an application in a lab before putting it on the production network and yet only 55% of companies have a performance objective for that application. It is also somewhat disappointing that roughly 30% of the time that an IT organization tests or pilots an application, that network engineers are not included.

The Infrastructure Director stated that he does not have the resources to do a lot of testing or piloting of applications. As a result, when his company is in the process of acquiring an application he asks the vendors for a lot of references and he follows up with the references to find out how the application runs over the WAN. Relative to testing applications prior to putting them into production, The Network Engineer said, "We are required to go through a lot of hoops." He went on to say that sometimes the testing was helpful, but that if the application development organization was under a lot of management pressure to get the application into production, that the application development organization often took the approach of deploying the application and then dealing with the performance problems later.

Conclusion

In our industry we often use the phrase "crossing the chasm" to refer to a technology that is just beginning to become a mainstream technology. When I look at assuring acceptable application delivery, I see that it both has and has not crossed the chasm. I believe that it has crossed the chasm in terms of expectations and responsibilities. By that I mean that within the vast majority of companies there is a growing awareness of the importance of assuring acceptable delivery and that in the majority of instances the network organization has assumed responsibility for this task.

I also believe that assuring acceptable delivery is just at the point of crossing the chasm in terms of how well the majority of IT organizations perform this task. In particular, it is rare for an IT organization to tell me how well they perform this task. However, I hear frequently from IT organizations that tell me that while they have processes designed to better manage application performance, in many cases these processes are inadequate. I also hear fairly often about the animosity that exists between the application group and the network group.

So while we all struggle with how we can do a better job of assuring acceptable application delivery, I want to repeat a point that I made at the beginning of this report. The point is that successful application delivery requires IT organizations to be able not only to identify the applications that are running on the network but also to ensure the acceptable performance of the applications that are relevant to the business while controlling or eliminating applications that are not relevant.

I want to close this IT Impact brief by mentioning that in the December IT Impact Brief I will discuss the ten most important network related events of 2006. I invite your input. Kindly send me an email (jim@ashtonmetzler) and tell me which network-related events or issues? of 2006 you think were the most important and why.