The Mandate for Packet Flow Data

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

INTRODUCTION

This IT Impact Brief is the first in a set of four briefs leading up to the Interop conference in Las Vegas at the end of April. This brief will focus on the need for packet flow data and the three subsequent briefs will focus on instrumentation, processes, and products.

To both introduce and position this IT Impact Brief I want to briefly mention a trend that has been discussed in several previous briefs as well as a trend that I introduced for the first time in the last impact brief. The familiar trend is the growing requirement to ensure acceptable application performance. One of the primary reasons why I write so much about this trend is its growing importance to the NetScout community. Another reason that I write so much about this trend is because ensuring acceptable application performance is extremely difficult today, and as will be discussed in this brief will only become more challenging in the next few years.

"I think that the growing interest in network management is due in large part to the fact that the job of the network manager is changing."

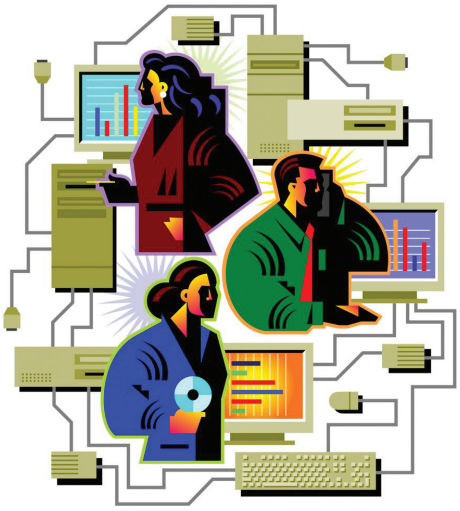

The second trend is the increased interest that I see in network management. As I mentioned in the last IT Impact Brief, over the last two years I have been a moderator at roughly a dozen of Network World's IT Roadmap conferences. In addition to the track that I moderate on network management, these conferences also have tracks on topics that have a lot of sizzle associated with them; i.e., Web 2.0, mobility, network and application acceleration. In every city, however, the track that gets the most attendees is always the network management track.

I believe that there is a strong interrelationship between these two trends. For example, I think that the growing interest in network management in due in large part to the fact that the job of the network manager is changing. In the not too distant past, the job used to focus entirely on the network. Today it is far more diverse and network managers are as likely to spend their time on topics such as security and managing application performance as they are to spend their time on purely networking issues. Because their jobs are changing so dramatically, network managers are aggressively trying to learn how to be successful in their new roles and as a result they are attending conferences such as the ones produced by Network World.

Given the importance of these two trends, the goal of the next few IT Impact Briefs is to drill down into how IT managers can successfully take on their new roles – roles that often include responsibility for managing application performance.

TODAY'S APPLICATION DELIVERY CHALLENGES

There are a number of factors that make ensuring acceptable application performance difficult today. Three of the primary factors are:

Ubiquitous WAN Access

Not that long ago the vast majority of employees worked in a headquarters facility and accessed applications over a high-speed, low-latency LAN. Now, in part because most workers are remote and in part because IT organizations have consolidated IT resources (e.g., storage, servers, applications) into a small number of central sites, the vast majority of employees access applications over a relatively low-speed, high-latency WAN.

Distributed Applications

In the late 1980s, client-server architectures began to be deployed. A client-server architecture is sometimes referred to as a two-tier architecture because the intelligence required to run the application resides on two separate computing devices – the client and the server. Within a few years of the initial deployment of a two-tier application architecture the industry began to deploy an n-tier application architecture. The phrase n-tier application architecture refers to an architecture in which the intelligence to run the application typically resides on three or four separate computing devices ("n" separate computing devices). For example, in a three-tier application architecture the application intelligence is split between a Web browser, a Web server and a database server.

WAN Vicious Protocols

Just over a year ago, I wrote an IT Impact Brief entitled "WAN Vicious Applications". (http://www.netscout.com/docs/itimpactbriefs/NetScout_iib_Metzler_1106_WAN_Applications.pdf). In that brief, I talked about chatty protocols. Chatty protocols are WAN vicious because they require tens or even hundreds of round trips just to complete a single transaction. XML is a WAN vicious protocol. The reason being that XML is very verbose and as a result the use of XML can consume five to ten times more WAN bandwidth than would be required if a traditional binary format was used.

The Emerging Challenges

I occasionally come across IT professionals who are in a state of denial relative to some of the challenges associated with ensuring acceptable application performance. For example, one argument that I often hear is that the cost of WAN bandwidth is going to become so low that we will soon be able to afford to throw WAN bandwidth at application performance problems. I have two concerns with that argument. My first concern is that even if it were true, adding WAN bandwidth will not make much of an improvement in the performance of any application that uses a chatty protocol. My second concern is that in most cases WAN traffic volumes are increasing faster than WAN tariffs are decreasing. As is discussed below, the adoption of new application architectures will dramatically increase the difficulty associated with managing application delivery in general, as well as increase the amount of traffic that transits the WAN.

"The adoption of new application architectures will dramatically increase the difficulty associated with managing application delivery in general, as well as increase the amount of traffic that transits the WAN."

One of the emerging application architectures that will dramatically increase the difficulty associated with managing application performance is the movement to Web services-based applications. To understand the difficulty of managing a Web services-based application, consider the three-tier application architecture that was previously discussed. In a three-tier application, the application server(s) and the database server(s) typically reside in the same data center. As a result, the impact of the WAN is constrained to a single traffic flow, that being the flow between the user's Web browser and the application server.

In a Web services-based application, the Web services that comprise the application typically run on servers that are housed within multiple data centers. As a result, the WAN impacts multiple traffic flows and hence has a greater overall impact on the performance of a Web services-based application that it does on the performance of an n-tier application. In addition, Web services-based applications are based on XML, which as previously noted consumes large volumes of WAN bandwidth.

However, it is not just Web services-based applications that will increase the difficulty of managing application delivery. Web 2.0 style applications, which came to life in order to support social networking sites such as MySpace, are beginning to find their way into the enterprise. A key component of Web 2.0 is the concept of an application that is itself the result of aggregating other applications together. This has become so common that a new term, mashup, has been coined to describe it. According to Wikipedia a mashup is a web application that combines data from more than one source into a single integrated tool. A typical example of a mashup is the use of cartographic data from Google Maps to add location information to real-estate data from Craigslist, thereby creating a new and distinct service that was not originally envisaged by either source. To enable Web 2.0, application platforms such as ASP.NET have been developed. The advantage of ASP.NET is that developers can use it to quickly develop applications. The disadvantage is that very often these applications do not perform well.

"In order to perform granular troubleshooting of complex IT environments, packet-level details are often necessary."

THE IMPORTANCE OF PACKET FLOW DATA

In a previous IT Impact Brief ("NetFlow – Gaining Application Awareness" http://www.netscout.com/docs/itimpactbriefs/NetScout_iib_Metzler_0106_NetFlow_Application_Awareness.pdf), I discussed some of the advantages of flow-based analysis. As mentioned in that brief, flow-based analysis enables IT organizations to perform tasks such as quantifying overall link utilization and identifying which network users or applications are consuming bandwidth.

Flow-based analysis is often sufficient to troubleshoot network problems. However, there are a number of factors that complicate network troubleshooting and which demand a deeper level of analysis. One of these factors is the complexity of contemporary networks. This complexity is partly a result of physical factors such as having a highly distributed infrastructure, as well as the fact that most IT organization continue to add new security functionality (e.g., firewalls, IDS, IPS) as well as functionality to optimize the performance of network and applications; e.g., compression, caching, protocol acceleration. This complexity is also a result of logical factors such as the movement to deploy both IPv6 as well as complex routing schemes such as asymmetric routing.

Other factors that complicate network troubleshooting include the distributed nature of applications, the emergence of new application architectures such as SOA and Web 2.0, as well the growing tendency of applications to not use their designated ports, but instead to use Port 80.
(see: The Port 80 Black Hole http://www.netscout.com/docs/itimpactbriefs/NetScout_iib_Metzler_0807_Port_80_Black_Hole.pdf)

In order to perform granular troubleshooting of complex IT environments, packet-level details are often necessary. In particular, in order to either eliminate a problem before it impacts end users or to reduce the MTTR once a problem has impacted end users, IT organizations need to have stored packet header and payload information. However, to avoid drowning in a sea of information, IT organizations also need automated expert analysis of the packet-level data that enables IT organizations to perform tasks such as:

- Application reconstruction and playback for transaction-level troubleshooting and compliance monitoring of IP voice and Web application sessions;
- Proactive scanning of recorded traffic using signature analysis, protocol semantic analysis and traffic pattern analysis to identify problems in network and application activity;
- Analyses and audits of traffic filtered by logical network definitions such as VLANs, DLCIs and PVCs.

This functionality allows IT organizations to reconstruct a session and understand what happened. It also enables IT organizations to identify dangerous and/or unapproved traffic.

SUMMARY

The NetScout community has been quite clear about the importance their organization places on ensuring acceptable application performance. There also cannot be much disagreement about the complexity of today's IT environment. Unfortunately that complexity will only increase in the near term as IT organizations add more functionality to the IT infrastructure and deploy new sophisticated application architectures, such as SOA and Web 2.0.

While there are strategies that IT organizations can use to be successful in this demanding environment, the bottom line is that IT organizations must have the ability to both collect packet-level data and to quickly perform an expert analysis of that data. As mentioned, the next three IT Impact Briefs will continue the discussion of how IT managers can successfully manage application performance. In particular, the next brief will examine the instrumentation choices that IT organizations have and discuss the pros and cons of each approach.

NetScout Systems, through its Sniffer and nGenius® solutions, offers large organizations cohesive views into application services delivered over today's complex, global networks, helping IT professionals optimize network and application performance and prevent misuse of critical enterprise resources. Based on granular, packet-flow performance information gathered across the enterprise, NetScout delivers key performance management functions, including application and network monitoring, capacity planning, troubleshooting, and user experience assurance, in a single integrated solution. For more information visit www.netscout.com.