Accelerating a Diverse Mix of Traffic Across the WAN

By Jim Metzler, Cofounder, Webtorials Editorial/Analyst Division

Introduction

Over the last few years the topic of WAN acceleration has become a top of mind issue for virtually all IT organizations. A survey that was recently given to 337 IT professionals indicated that one of the reasons why this has become such an important topic is the sheer number of applications that transit the WAN. For example, as shown in Table 1, approximately thirty percent (30%) of the survey respondents indicated that they currently have more than fifty business applications running over their WAN and roughly sixteen percent (16%) of the survey respondents indicated that they currently have more than one hundred business applications running over their WAN.

When asked about how this situation will change over the next two years the responses are even more dramatic. Forty-four (44%) percent of the survey respondents expect that two years from now there will be more than fifty business applications running over their WAN and twenty-one percent (21%) of the survey respondents expect that there will be more than one hundred business application running over their WAN.

	More than 50 Business Applications	More than 100 Business Applications
Currently	30%	16%
Two years from now	44%	21%

Table 1: Number of Business Applications Transiting the WAN

According to the enterprises surveyed, the following applications were identified as the most important from an application delivery standpoint (in order of importance):

- VoIP
- Email
- Backup/Replication
- ERP
- Database applications
- Video
- Citrix[1]
- File related applications and protocols; i.e., CIFS, NFS, FTP

Each of these applications behave differently when confronted with the bandwidth, latency, and packet loss challenges that are endemic to enterprise WANs. As a result, different WAN acceleration techniques are required for each traffic type. To better explain this fact, this paper breaks enterprises' applications into common categories and discusses which optimization technologies are most effective at ensuring their successful delivery across a WAN.

[1] Throughout this brief, the term 'Citrix' will refer to Citrix presentation server traffic.

Key Traffic Categories

In the typical enterprise, it is common to see a variety of different types of applications. These can be broken down into the following categories:

Bulk TCP-Based Applications

A key characteristic of this category of application is that the data is typically sent in a large continuous stream utilizing the maximum TCP window. Typical applications include file services (i.e., CIFS, NFS, FTP) as well as storage replication, data base synchronization and email. Another application in this category that is growing in importance is backup/replication across the WAN. This traffic is often sent as a bulk TCP transfer between hosts[2].

Fifty-five (55%) percent of the survey respondents indicated that they perform remote backups over the WAN and fifty-three percent (53%) do data replication over the WAN. In addition, as noted above, two of the TCP applications whose performance over the WAN the survey respondents are most concerned about are email and backup & replication.

Real-Time Traffic

These applications are time-sensitive, and therefore show noticeable degradation in performance when insufficient bandwidth, latency or packet loss prevents timely packet delivery. Typical applications include Citrix and remote desktop, both of which run over TCP. Other real-time applications include Voice over IP (VoIP) and video conferencing, both of which run over the User Datagram Protocol (UDP). The big difference between TCP and UDP is that the latter does not throttle back bandwidth utilization in response to network congestion and does not have built-in mechanisms for retransmission. As such, UDP is better suited for latency sensitive traffic.

The survey respondents were presented with a number of real-time applications and were asked if their company ran these applications over their WAN. Their responses are contained in Table 2.

Real Time Application	Percentage of Companies Running this Applications over their WAN
Remote Desktop	82%
VoIP	65%
Citrix	58%
Video	56%

Table 2: Use of Real-Time Applications

In most cases, companies are running more than one of the applications shown in Table 2. For example, ninety-seven percent (97%) of the survey respondents run a thin client application over their WAN - Microsoft's remote desktop and/or Citrix. Seventy-nine percent (79%) of the survey respondents run VoIP and/or video over their WAN. As previously noted, the survey respondents have concerns about running VoIP, video and Citrix over their WAN. That concern is justified in part because real-time applications present a challenge for buffer-based optimization solutions that add latency.

2 This traffic can sometimes run over UDP or proprietary TCP protocols.

Non-TCP Applications

As mentioned above, many time-sensitive applications such as VoIP and video run over UDP instead of TCP. A number of data replication solutions also fall into this category; e.g., Aspera, Veritas Volume Replicator, Isilon and CLARiiON disk library. In addition, a number of applications such as Fibre Channel over IP (FCIP) are written to run directly over IP and a number of applications are encapsulated inside of IP; i.e., GRE, IP-IP, and IPv6 tunneled inside of IPv4.

The vast majority of the survey respondents stated that they have non-TCP based applications transiting their WAN. The most frequently mentioned applications were VoIP and Video.

Dynamic Data

Some applications scramble the way data is stored when a file is opened and saved. In many instances this can occur even if the file has not been modified. Examples of applications with dynamic data formats include newer versions of AutoCAD and MS Excel.

A recent article[3] discussed the challenges that dynamic data presents to the way that some optimization solutions implement deduplication. That article pointed out that in many instances AutoCAD 2007/2008 files are completely re-written when a user does a full save on the drawing. The result is that instead of only a few bytes changing with a simple edit, the whole byte structure of the file changes. How dynamic applications impact the effectiveness of deduplication techniques will be discussed later in this brief.

WAN Characteristics and Requirements

WAN Characteristics

The WAN has several characteristics that make it challenging for IT organizations to successfully run applications over it. These characteristics include bandwidth limitations, latency, and packet loss and re-ordering.

Bandwidth Limitations

Unlike the situation in the LAN, IT organizations pay for a WAN circuit on a monthly recurring basis and the cost of the circuit is proportional to the bandwidth of the circuit. As a result, few organizations can afford to pay for all of the WAN bandwidth it needs during peak periods. The lack of bandwidth results in the inability to transmit in a timely fashion the volumes of data that are associated with applications such as backup and data replication. In addition, a lack of bandwidth can lead to packet loss and retransmissions.

Latency

As shown in Figure 3, the delay associated with the LAN is significantly less than the delay associated with a terrestrial WAN link, which is less than the delay associated with a satellite link.

Network Type	Typical Latency
LAN	< 1 ms
WAN Link: Western US to Eastern US	60 to 100 ms
International WAN Link	100 to 400 ms
Satellite Link	500+ ms

Table 3: Latency Values

[3] Why a change to an AutoCAD file format is throwing some WAN accelerators for a loop, http://weblog.infoworld.com/tcdaily/archives/2008/04/why_a_change_to.html

Latency negatively impacts the behavior of all applications and is particularly detrimental to real-time applications. One of the curious aspects of latency is that even a small increase in WAN latency can have a significant impact on the overall performance of the application. For example, if the standard TCP window size (64KB) is used, then the maximum throughput on a single TCP stream is less than or equal to 64 KB/latency[4]. Table 4 demonstrates the relationship between latency and throughput.

Latency	Maximum Throughput on a Single TCP Stream
100 ms	5.12 Mbps
200 ms	2.56 Mbps
300 ms	1.71 Mbps
400 ms	1.28 Mbps
500 ms	1.02 Mbps
600 ms	853 Kbps

Table 4: Throughput as a Function of Latency

Packet loss/re-ordering

A recently published brief [5] pointed out that packet loss and packets that are received out of order will both cause a re-transmission of packets. As such, both packet loss and out of order packets have the same affect on goodput[6].

That brief also quantified the impact of packet loss and re-ordering, showing that with a 1% packet loss and a round trip time of 50 ms or greater, the maximum throughput on a single TCP stream is roughly 3 megabits per second no matter how large the WAN link is. This limited throughput would significantly reduce the IT organization's ability to successfully implement backup and replication applications.

WAN Acceleration Requirements

In order to overcome the WAN characteristics described above, a network and application optimization solution must support certain key functionality. This functionality includes network integrity features, network acceleration techniques and deduplication.

Network Integrity Features

One key network integrity feature is Quality of Service (QoS). QoS refers to the ability of the network to provide preferential treatment to certain classes of traffic, such as voice traffic. QoS is required in those situations where bandwidth is scarce and there are one or more delay sensitive, business critical applications.

A second key network integrity feature is the ability to correct in real-time for dropped packets. While it is not possible to eliminate packet loss, it is possible to implement a technique such as Forward Error Correction (FEC)[7] that can mitigate the impact of packet loss in real-time. The basic premise of FEC is that an additional error recovery packet is transmitted for every 'n' packets that are sent. The additional packet potentially enables the network equipment at the receiving end to reconstitute the lost packets and hence negates the need to re-transmit packets, which can hamper goodput. Additionally, techniques that re-order packets in real-time can also help preserve network integrity across the WAN.

---

4 http://en.wikipedia.org/wiki/TCP_Tuning

5 The Data Replication Bottleneck: Getting Narrowband Throughput on a Broadband WAN, http://webtorials.com/abstracts/KubernanBrief-1-4.htm

6 "Goodput" refers to the amount of data that is successfully transmitted. For example, if a thousand bit packet is transmitted ten times in a second before it is successfully received, the throughput is 10,000 bits/second and the goodput is 1,000 bits/second.

7 RFC 2354, Options for Repair of Streaming Media, http://www.rfc-archive.org/getrfc.php?rfc=2354

Network Acceleration Techniques

Many WAN protocols (e.g., TCP, CIFS, HTTP) have features that cause them to perform suboptimally when run over a WAN. For example the TCP windowing mechanism can make inefficient use of WAN bandwidth, especially over high latency paths. TCP can be accelerated by a variety of techniques that increase a session's ability to more fully utilize link bandwidth. Some of the available techniques are dynamic scaling of the window size, selective acknowledgement, and TCP Fast Start. Increasing the window size for large transfers allows more packets to be simultaneously in transit boosting bandwidth utilization. TCP selective acknowledgment (SACK) improves performance in the event that multiple packets are lost from one TCP window of data. With SACK, the receiver tells the sender which packets in the window were received, allowing the sender to retransmit only the missing data segments instead of all segments sent since the first lost packet. TCP slow start and congestion avoidance lower the data throughput drastically when loss is detected. TCP Fast Start remedies this by accelerating the growth of the TCP window size to quickly take advantage of link bandwidth.

WAN Deduplication

Conceptually deduplication is very simple – send less data over the WAN. The way WAN deduplication works is that the first time that a file is sent across the WAN, deduplication algorithms identify data patterns in the file and store these patterns in WAN optimization appliances on each end of the WAN link. On subsequent passes, the deduplication algorithms in the appliances identify these patterns and replace large pieces of data with notably smaller pieces of data (i.e., "references").

However, the way that WAN deduplication is implemented impacts how effective the solution will be. Some of the implementation options include:

- ***IP vs. TCP***

A WAN deduplication solution that works at the TCP layer[8] optimizes only TCP traffic. Any applications that use UDP, proprietary or encapsulated protocols (as discussed above) can only be optimized by a deduplication solution that works at the IP layer. In addition, TCP layer implementations often have difficulties handling large volumes of flows because they typically rely on the underlying operating system to perform TCP re-assembly and re-segmentation operations.

- ***Real-time vs. Buffered***

As previously noted, some applications (e.g., Citrix, Microsoft's RDP, voice, video) are notably sensitive to latency. Optimization solutions that rely on buffering add delay that can make real-time applications unusable.

- ***Disk vs. RAM***

Advanced compression solutions can be either disk or RAM-based. Disk-based systems typically can store as much as 1,000 times the volume of patterns in their dictionaries as compared with RAM-based systems, and those dictionaries can persist across power failures. The data, however, is slower to access than it would be with the typical RAM-based implementations, although the performance gains of a disk-based system are likely to more than compensate for this extra delay. While disks are more cost-effective than a RAM-based solution on a per byte basis, given the size of these systems they do add to the overall cost and introduce additional points of failure to a solution. Standard techniques such as RAID can mitigate the risk associated with these points of failure.

- ***Instruction-Based vs. Token-Based***

Some WAN optimization solutions implement a token-based approach to deduplication while other solutions implement an instruction-based approach. A token-based approach relies on the use of a token to represent "chunks" of data. The tokens are created by a hash function[9] that tries to avoid having different chunks of data represented by the same token. In contrast, an instruction-based approach uses specific start-stop instructions to

---

[8] TCP runs at Layer 4 of the OSI protocol stack. This is often referred to as the transport layer.

[9] A hash function is any well-defined procedure or mathematical function for turning data into a relatively small integer, that may serve as an index into an array.

indicate where duplicate data can be found and retrieved. In an instruction based approach, no hashing or tokens are used.

These differences become most apparent when applications have dynamic data. As described above, applications like AutoCAD or MS Excel scramble the way that data is stored when a file is opened and saved. Even relatively minor changes can result in small changes being made throughout the file.

Deduplication solutions that rely on tokens to represent chunks of data may have a difficult time matching patterns when the data is scrambled. That is because the data represented by the tokens will not match up when one file is compared to another. This scrambling effect is less severe when using an instruction based schemes for deduplication. That is because this method provides better byte level granularity, enabling repetitive patterns to be detected even if they are not in the exact same location between files.

Summary

Virtually all organizations find that WAN acceleration is both important and complex. WAN acceleration is complex in part because of WAN characteristics such as limited bandwidth, latency and the combination of packet loss and out of order packets. Another factor that makes WAN acceleration complex is the fact that not all traffic has the same requirements and hence different traffic categories require different acceleration techniques. For example, real time applications such as VoIP and video conferencing require QoS. Bulk TCP-based applications typically require protocol acceleration as well as techniques such as FEC and packet re-ordering.

Further complicating the task of implementing WAN acceleration is the fact that how a technique such as WAN deduplication is implemented can have a significant impact on how well it performs. As mentioned, if WAN deduplication is implemented at the TCP layer, it only optimizes TCP traffic. If WAN deduplication is implemented using a token-based approach, it may not be very effective accelerating the performance of applications that have dynamic data.

In order to be successful with WAN acceleration, IT organizations must identify the primary categories of traffic that transit their WAN and must also identify the WAN acceleration requirements of those traffic categories. In addition, when comparing WAN acceleration solutions, IT organizations must understand how the various acceleration techniques are implemented and determine the impact of that implementation on the performance of their applications.

A Word from the Sponsor – Silver Peak

Silver Peak improves backup, replication and recovery between data centers and facilitates branch office server and storage centralization by improving application performance across the WAN. This is achieved using a variety of WAN optimization techniques, including disk based data reduction, compression, latency/loss mitigation and Quality of Service (QoS).

Silver Peak's award winning NX appliances bring unprecedented security and scalability to WAN acceleration, making them uniquely suited for the rigorous demands of high capacity WAN environments. Common IT initiatives facilitated by the Silver Peak solution include asynchronous data replication, network backup, disaster recovery, SQL transactions, file transfers, email, VoIP and video streaming.

Silver Peak is headquartered in Santa Clara, California. For more information, visit http://www.silver-peak.com.

About the Webtorials® Editorial/Analyst Division

The Webtorials® Editorial/Analyst Division, a joint venture of industry veterans Steven Taylor and Jim Metzler, is devoted to performing in-depth analysis and research in focused areas such as Metro Ethernet and MPLS, as well as in areas that cross the traditional functional boundaries of IT, such as Unified Communications and Application Delivery. The Editorial/Analyst Division's focus is on providing actionable insight through custom research with a forward looking viewpoint. Through reports that examine industry dynamics from both a demand and a supply perspective, the firm educates the marketplace both on emerging trends and the role that IT products, services and processes play in responding to those trends.

For more information and for additional Webtorials® Editorial/Analyst Division products, please contact Jim Metzler at jim@webtorials.com or Steven Taylor at taylor@webtorials.com.

Webtorials Briefs
Vol 2, Number 7

Published by Webtorials Editorial/Analyst Division
www.Webtorials.com

Division Cofounders:
Jim Metzler
jim@webtorials.com
Steven Taylor
taylor@webtorials.com

Professional Opinions Disclaimer
All information presented and opinions expressed in this publication 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.