



The following paper was originally published in the
Proceedings of the Twelfth Systems Administration Conference (LISA '98)
Boston, Massachusetts, December 6-11, 1998

Wide Area Network Ecology

Jon T. Meek, Edwin S. Eichert, Kim Takayama
Cyanamid Agricultural Research Center/American Home Products Corporation

For more information about USENIX Association contact:

1. Phone: 510 528-8649
2. FAX: 510 548-5738
3. Email: office@usenix.org
4. WWW URL: <http://www.usenix.org>

Wide Area Network Ecology

Jon T. Meek, Edwin S. Eichert, Kim Takayama

– Cyanamid Agricultural Research Center/American Home Products Corporation

ABSTRACT

In an ideal world the need to provide data communications between facilities separated by a large ocean would be filled simply. One would estimate the bandwidth requirement, place an order with a global telecommunications company, then just hook up routers on each end and start using the link. Our experience was considerably more painful, primarily due to three factors: 1) The behavior of some of our applications, 2) problems with various WAN carrier networks, and 3) increasing Internet traffic. “Network Ecology” describes the management of these factors and others that affect network performance.

Introduction

American Home Products Corporation (AHP) is a global life sciences company with over 220 locations. This paper will examine the properties of Frame Relay Wide Area Network (WAN) connections between the Agricultural Research Center in Princeton New Jersey and two European facilities. Then the paper will look at the behavior of network applications on these links.

During the past year AHP started to switch its leased line based WAN to managed Frame Relay networks. Most of the previous WAN usage was for bulk file transfer, database synchronization, light interactive TTY sessions, and some http traffic.

Coincident with the start of Frame Relay implementation, several client-server applications went into testing at the two European sites. These are traditional client-server applications with client PCs in Europe interacting with Oracle databases in Princeton using Oracle’s SQL*Net protocol. At the same time, use of the Internet started to increase dramatically. Since the Internet access points for the Corporation are located in the US, this placed an additional load on the WAN.

The old lines did not have the bandwidth to gracefully handle the new demands, so complaints about the performance of the client-server applications were answered with “It should be better with Frame Relay.” As the European Agricultural Research sites came onto the Frame Relay network it became obvious that performance did not improve significantly.

We found that initial guesses about the cause of WAN performance problems were often incorrect. With work, they can usually be traced to one or more of the following factors.

- System and Network Administration Practices
- The WAN Carrier’s Network
- Commercial Hardware and Software Products
- In-house Application Programs
- Other Uses of the Network

This paper discusses what we learned about managing WAN links, what measurements and monitoring

have helped us, and how we worked with our Frame Relay carriers to improve performance.

Frame Relay Basics

A major advantage of Frame Relay is the ability to burst above the guaranteed bandwidth (committed information rate, or CIR) purchased from a carrier. In the case of the two connections discussed here, CIRs were 64kbps and 32kbps and the access lines varied from 128kbps to 512kbps. Bursting may be limited to multiples of the CIR such as 2x or 4x, or bursting to port speed may be possible. Our Carriers (OC) allow bursting to full port speed, depending on the availability of bandwidth in their core network and the customer’s recent usage history. Depending on the policies of the carrier, frames that exceed CIR may be sent with the Discard Eligible (DE) bit set. This allows the carrier to discard those frames if congestion is encountered while they flow through the network. Customers can build credits when usage runs below CIR which may allow bursting above CIR without frames being marked DE. Managing bandwidth use is clearly an important aspect of “Network Ecology.”

Network Parameters

In addition to bandwidth, other network performance variables include round-trip-time (RTT) or latency, dropped packet counts, and availability. According to OC packets are dropped only when traffic on a link bursts above the CIR (the DE is bit set and the frame encounters congestion). In our experience, availability is very high although regular monitoring is essential. Assuming that bandwidth utilization is under control this leaves RTT as the most important parameter to study.

Minimizing RTT is especially important for interactive TTY sessions and for applications that require a large number of acknowledgment packets. These acknowledgments, sometimes as many as one for each data packet, are due to both TCP and application flow control. In a session involving transfer of many packets, the “wait for acknowledgment” time

adds up quickly. We found that tuning systems and applications so that full-size packets were sent during bulk data transfer portions of a session resulted in the best performance.

On a LAN RTTs are typically <2 ms while trans-Atlantic link RTTs of 90-200 ms are typical. During times of over-utilization, or carrier network problems, RTTs may soar up to eight seconds.

Measuring Bandwidth Usage

It became apparent that we needed to do fairly high-resolution monitoring of network utilization and performance. OC does not normally provide access to the routers that they manage, even those located at customer sites. We were able to negotiate SNMP read-only access which provided several Frame Relay parameters for each PVC (permanant virtual circuit) served by a router.

Every five minutes the following parameters are logged for each PVC: Frames Sent, Frames Received, Bytes Sent, Bytes Received, FECNs, and BECNs. Bytes sent and received are a direct measure of bandwidth usage. The last two parameters, Forward Explicit Congestion Notification and Backward Explicit Congestion Notification are indications of congestion on the network between the end points and may be useful to help detect problems on the carrier's network [Cava98]. The SNMP parameter log is run by cron to ensure that the periods are accurate five minute intervals. The log files are rotated monthly and old logs are retained indefinitely.

Measuring Round-Trip-Time

Since RTT is subject to variation depending on load and routing changes in the OC network, we measure it every five minutes. The RTT measurements double as a connectivity check and are implemented as a mon [Troc97] monitor.

The RTT check monitor sends five small (44 bytes including headers) UDP packets to the echo port of each end-point router. The minimum RTT is used as the reference, but we record the number of packets returned, minimum, mean, and maximum times. If the minimum RTT exceeds a set acceptable limit (currently two seconds), mon alarms are triggered.

If all five of the UDP packets are dropped, then a TCP connection to the echo port is attempted. If the TCP connection attempt times out, the link is considered down and a mon alarm is triggered. About three minutes is required for this process to fail, so we should alarm only on outages that last more than three minutes.

The use of these small probe packets, totaling less than 250 bytes per five minute period, has negligible impact on network capacity.

Communicating Measurement Results

The performance and utilization information collected every five minutes is made available to network managers through Web queries. This allows them to determine if too much bandwidth is being used or if there might be a problem in the carrier's network. Among the parameters supplied on the Web reports is percent of CIR used for both in-bound and out-bound directions. This calculation is based on the five minute average use and, while useful to network managers, is very different from the CIR computed by the Frame Relay switches. The switches use time periods on the order of seconds and compute CIR using algorithms that are not completely known by the carrier's customers.

Other WAN Quality Measurements

In addition to the regular RTT measurements discussed above, we found that measuring RTT vs. packet size is useful. These tests send 1000 random size UDP packets with between 0 and 1472 bytes of random data to the echo port of a router on the other end of a link. All of the results shown here were done at quiet times. The test packet rate was limited by the RTT since we wait (with a 15s timeout) for each packet to return before sending the next packet. The MD5 checksum of the data is computed before the packet is sent and after it is echoed back. This verifies the integrity of the link and eliminates any possible problems with packets that were assumed lost due to the timeout but eventually returned.




Figure 1a: Round-trip-time vs. UDP packet size. Good performance with 512 kbps and 192 kbps access lines.

By plotting measured RTTs on the y axis and packet sizes on the x axis it is possible to determine the fixed delay (y-intercept), serialization delays (slope), and consistency (scatter of points). The

serialization delay can be predicted quite accurately by just considering the speeds of the access lines on each end of the link (typically 192kbps to 512kbps). The best performing links will have a minimum y-intercept and most points lying close to a straight line. Figure 1 shows three examples of this test on different PVCs.

Table 1 below shows the result of fitting several sets of RTT vs. packet size data. The estimated value was computed using only the speed of the access lines at each end of the link. The measured value includes all serialization delays encountered in the path. The measured fixed delays vary here because the measurements were made over a three month period when the configuration of both our access lines and the core network were changing.

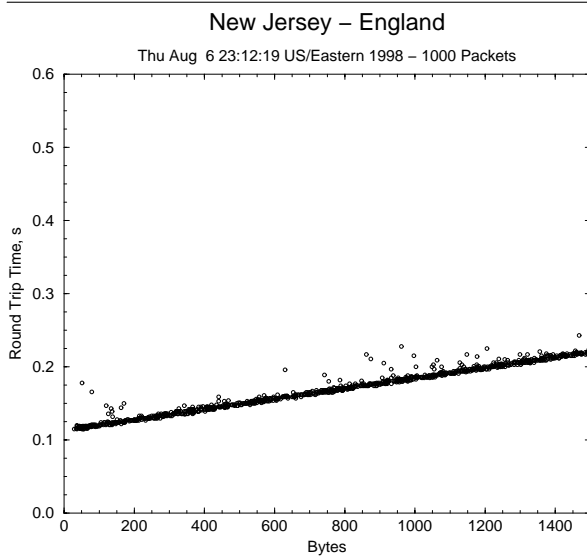


Figure 1b: Round-trip-time vs. UDP packet size. Good performance with 512 kbps access lines on each end.

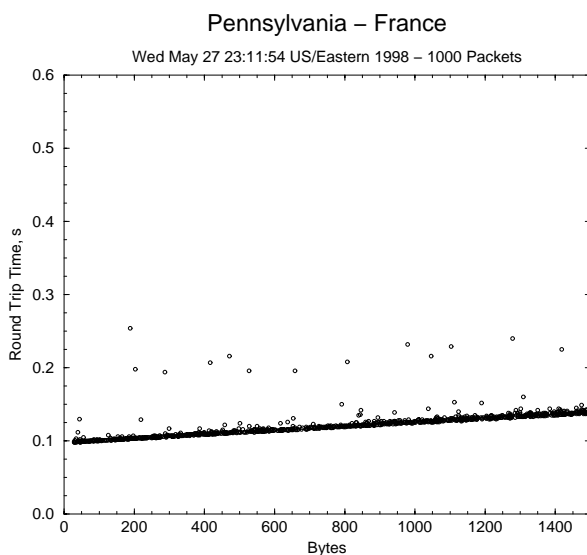


Figure 1c: Round-trip-time vs. UDP packet size. Good performance with T1 and E1 access lines.

Serialization delay improvements can be purchased (up to a point) by paying for faster access lines, while fixed RTT is usually specified only as a target value by WAN carriers and is limited by distance. The table includes measurements made before the New Jersey access line was upgraded from 128 kbps to 512 kbps. The last three table entries correspond to Figures 1a-1c. The difference between measured and estimated serialization delays will include a contribution due to serialization delays in OC's network where there are four additional serialization points per round trip with speeds between 2 and 16 Mbps.

End Points (Port Speed)	Estimated $\mu\text{s/bit}$	Measured $\mu\text{s/bit}$	Measured Fixed RTT, ms
New Jersey (128k) - Germany (192k)	25.4	27.7	150
New Jersey (128k) - England (512k)	19.0	23.1	151
New Jersey (512k) - Germany (192k)	14.0	17.8	155
New Jersey (512k) - England (512k)	7.6	8.9	114
Pennsylvania (T1) - France (E1)	2.3	3.4	99

Table 1: Estimated serialization delays based on access line speeds and measured serialization and fixed delays.

RTT vs. packet size plots can be useful as a measure of service uniformity. Figure 2 shows two plots of measurements taken while OC was experiencing some network instability. The New Jersey - Germany data might indicate route flapping between two, or more, different paths. It is possible that the results of Figure 2a could be due to congestion [Bolo93] either on the PVC, or on the carrier's network. Congestion on the PVC was unlikely in this case since the test packets were essentially the only traffic. Visual inspection of the plots in Figures 1 and 2 suggest that something has changed for the worse in Figure 2. Since the ultimate goal is a largely automatic monitoring system we investigated possible single number metrics that would indicate reduced quality-of-service. The RMS residual (square root of the sum of the squares of the difference between the fit line and the measurements) seems to be a good candidate for this metric. The RMS residuals are 0.306 ms, 0.155 ms, 1.017 ms, and 0.540 ms for Figures 1a, 1b, 2a, and 2b respectively. An OC engineer agreed that Figure 2a indicated a definite problem while Figure 2b was probably within normal operating limits. The best-fit line is shown on each plot.

WAN Quality Measurements - Dropped Packets

Another measure of network quality is the percentage of dropped packets when operating within CIR constraints (below CIR, or bursting with built-up credits). At one point we found that the size of successful ftp transfers from Germany to the US were limited to 25kB, but the reverse path allowed much larger files to be transferred with no problem. Using a custom Perl script that sent numbered UDP packets

we discovered that when packet size went above 966 bytes, every other packet was dropped. We were eventually able to demonstrate this problem using ping with the pre-load option that causes a specified number of packets to be sent as fast as possible.

Unfortunately, many versions of ping, including the Cisco version, do not have the pre-load option. This made it difficult to convince OC's first line support staff that there was a problem. Eventually OC discovered that a Frame Relay buffer size parameter was too small. After they increased the buffer size the problem was corrected.

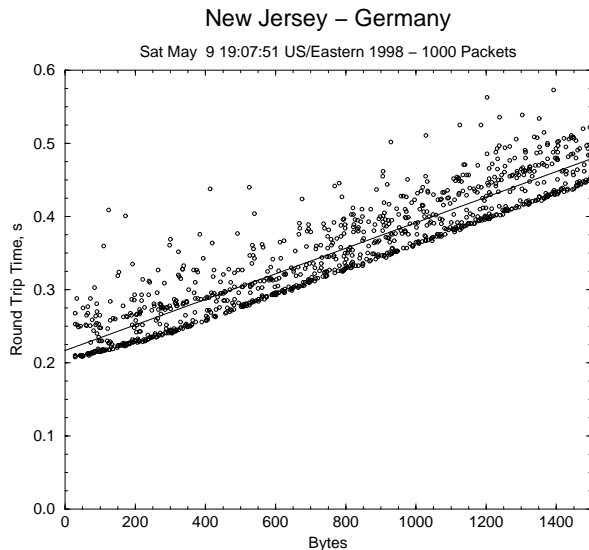


Figure 2a: Round-trip-time vs. UDP packet size, illustrating poor performance.

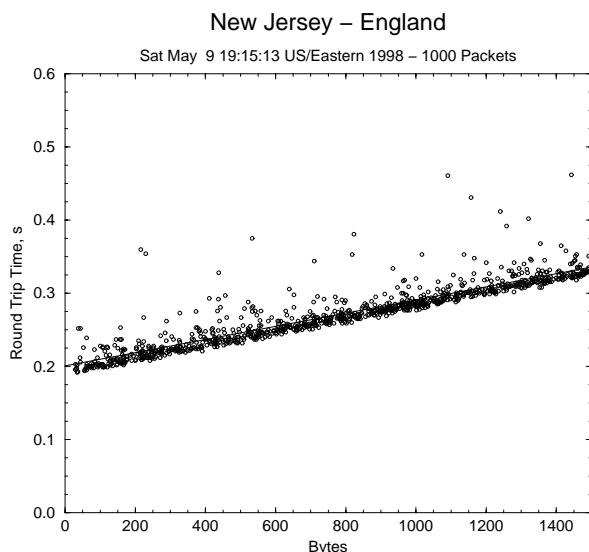


Figure 2b: Round-trip-time vs. UDP packet size, illustrating degraded performance.

Our monitoring program records the number of UDP packets successfully echoed during RTT tests. This provides one measure of the drop rate at a given

time. It would be better to count re-transmitted packets and the number of Frame Relay frames sent and received with the discard eligible bit set. These numbers are not available via SNMP from the routers we are using but could be obtained from another monitoring technique.

What Does the Carrier Monitor?

It took several months before we fully understood what network parameters were pro-actively monitored by OC. It turned out that they only watched for connectivity outages. If their network monitoring system could ping each end of a link, then all was considered well. Furthermore, transient outages were likely to be missed if someone was not watching the network management screen at the right time.

When OC is informed of a customer's negative feelings ("the network seems slow today"), they manually probe deeper to look for problems. Clearly this was not enough; we needed regular measurements of RTTs and bandwidth usage. These measurements are used to establish baselines, trigger alarms when some limit is exceeded, provide reports to assist network management, and build credibility with OC by reporting only real problems.

What's on the Wire? Who's Using the Wire?

After observing that the two European links often had a lot of traffic and that RTTs increased with load, we started to characterize the packets. Using a combination of tcpdump [McCa97], the libpcap Perl module [List97], Network Flight Recorder [Ranu97], firewall logs, and a Network General Sniffer, we were able to determine that some traffic could be eliminated.

There were a lot of routing broadcasts, http traffic to the Internet, and, on one link, Novell broadcast traffic. The Novell traffic was especially interesting since we do not use Novell on either side of that WAN link. It turned out that another Division was using this link to get to their European facilities.

Some of the actual problems discovered via monitoring were:

- For a period of several days, one of the links was fully saturated. Investigation revealed that the Division using our link for Novell traffic was sending large ping packets from three servers in St. Louis to two machines located in Ireland every second. They had been doing some troubleshooting, and forgot to stop the pings. After this had been going on for a week, we brought it to the attention of the responsible network managers and the pings were stopped.
- The printing of purchase orders and other documents from an ERP (Enterprise Resource Planning) system between European sites was very slow. The IT staff responsible for the application suggested that there was "a problem" with

OC's network. By capturing a print session with `tcpdump` and extracting the PostScript data, it was determined that more than half of the 1.1MB print job was due to trailing spaces that padded every line out to column 140. Another 350kB was due to multiple bitmapped company logos that could be replaced with a 6kB scalable PostScript object. The IT staff is working with the print software vendor to solve these problems.

- Quite a few other high usage problems have been due to "mis-configured" Web browsers, or users leaving their browsers pointing at automatically refreshing pages. The automated firewall reports described below were effective in solving many of these problems.

Where is the Wire?

In our quest to improve WAN performance it seemed that fixed delay was a good parameter to pursue. On the two European links discussed here, fixed delays varied from 150 ms to 225 ms on stable, quiet lines. In contrast, another AHP Division's Pennsylvania to France link had 100 ms fixed delay and the RTT to OC's public Web server located in England, over the Internet was only 90 ms. Therefore, improvement seemed possible.

The fixed delay time for packet transit is due to switching delays and the distance that the signal must travel. We decided to concentrate on distance first.

Several of the US-Europe trans-Atlantic fiber optic cables leave from New Jersey, at least one leaves from Long Island, and some leave from Rhode Island. Our traffic is carried over a number of these cables although we don't know which ones. We were, however, able to learn more about the routes our packets traveled on their way to the trans-Atlantic cables.

Initially our Frame Relay access line was connected to a switch in Maryland, a seemingly round-about way to get to any of the trans-Atlantic cables. When we upgraded the access line speed from 128kbps to 512kbps (to handle capacity requirements and reduce serialization delay) an additional 50 ms was immediately added to the fixed delay. Detective work revealed that because there were no 512k ports in Maryland we were now connected to a switch in Georgia, adding at least 2500 miles to our packet's round-trip. This is especially sad considering that some of the trans-Atlantic cables are located only 35 miles from our Princeton facility.

The signal propagation speed in a fiber optic cable is about 0.66 times the speed of light. This results in a physics limited delay of about $8 \mu\text{s}/\text{mile}$. The extra 2500 miles thus represents about 20 ms of fixed RTT. Since the 2500 miles is based on straight

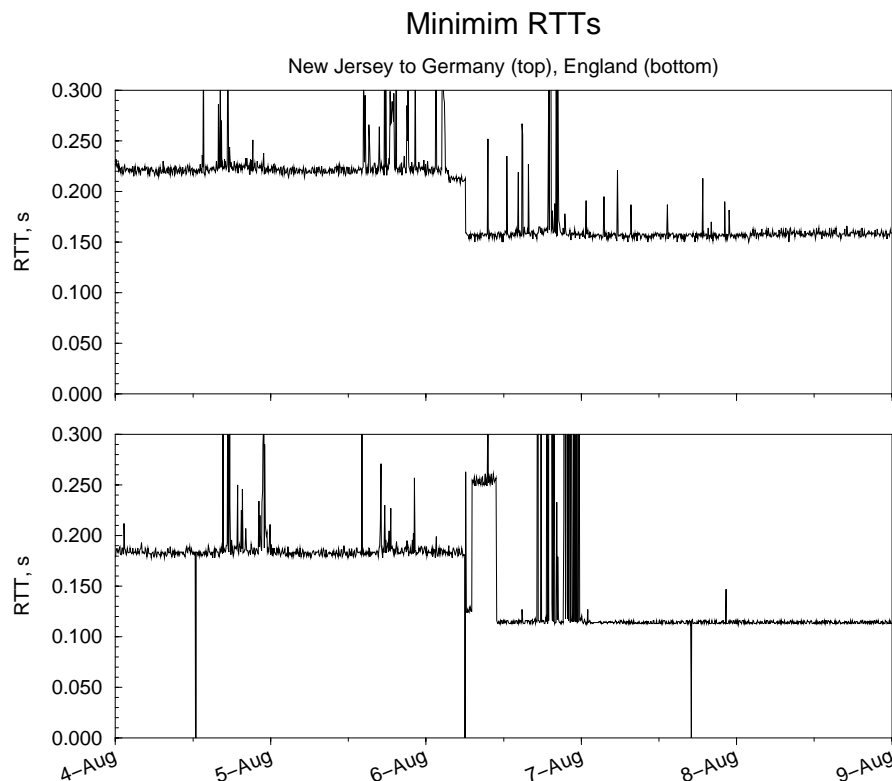


Figure 3: Minimum round-trip-times per five minute interval illustrating improvement due to geographical move of access line.

line distance, and since there must be additional switching delays on this long path, the 50 ms is a reasonable total RTT addition due to the access point change.

After waiting three months we were able to get the access line moved back to Maryland and a 60 ms RTT improvement was immediately realized (10 ms more than we “lost”) as shown in Figure 3 (on

6-Aug-1998). While we hoped to get more direct access to the transatlantic cables than passing through Maryland, we were told that the Maryland site is a required stop for Frame Relay packets, unless we wanted to visit Chicago on the way between New Jersey and Europe.

The zero RTT spikes in the New Jersey to England plot indicate short outages. The jump in RTT on

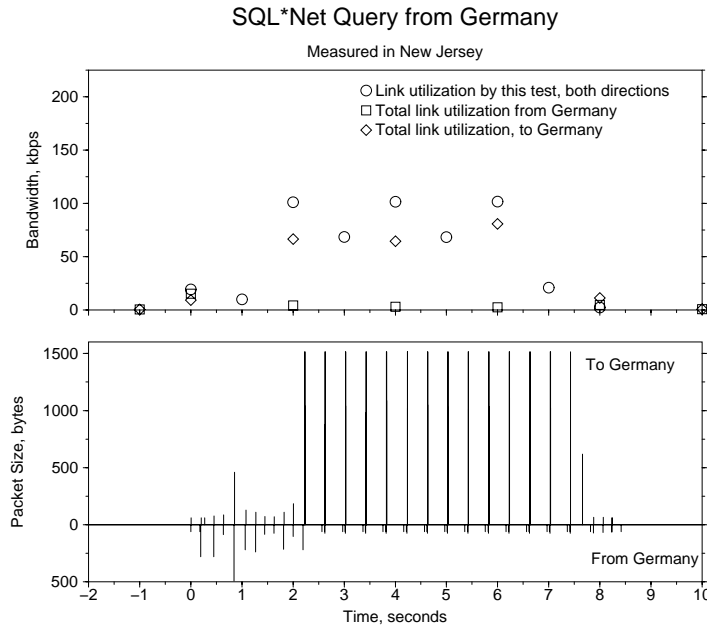


Figure 4a: Bandwidth usage and packet flow for remote database access. Each bulk data transfer cycle consists of about three large back-to-back packets followed by an equal number of acknowledgments.

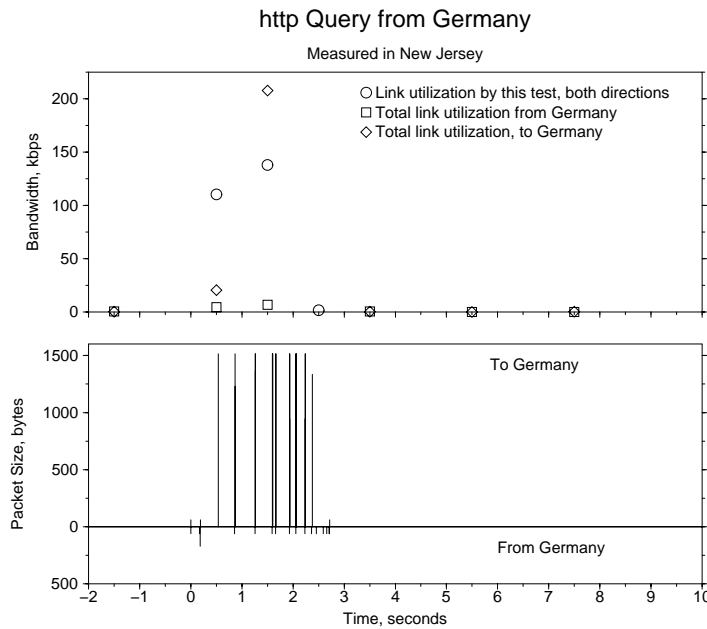


Figure 4b: Bandwidth usage and packet flow for http access of same data as above. The bulk data transfer cycle is similar to Figure 4b except that each set of data packets is followed by a single acknowledgment. Back-to-back packets overlap in both figures.

6-Aug-1998 soon after the access line move was due to an outage in England that caused the probe packets to be routed through Germany. Figure 3 shows quite a few spikes above 300 ms RTT that illustrate how the minimum RTT can increase significantly during times of heavy traffic, usually during working hours.

Analyzing Application Network Usage

In response to complaints about the performance of client-server applications, we captured and analyzed packets for sample sessions. Tpdump was used for packet capture and our own program for analysis. The test client was a SGI workstation located in Germany. On the client side, tests of SQL*Net were done using a Perl/DBD::Oracle [Bun98] script and the http tests used GNU wget [Nik97]. The servers were Sun SPARC Solaris 2.6 systems running Oracle 7.3.4 and Apache 1.2.6.

The first striking result was that bulk data transfer portions of Oracle/SQL*Net sessions sometimes consisted of many small packets with an acknowledgment for every data packet. On a fast LAN, with sub-millisecond RTTs, this is hardly noticeable; but on a WAN with 100-200 ms RTTs response time quickly adds up to multiple seconds. The Oracle/SQL*Net application was significantly improved by increasing the size of the Oracle row cache on the client side. Figure 4a shows the packet flow in the improved application. Packets in the bulk data transfer portion were mostly full-size, with several sent back-to-back. The client side still sent an acknowledgment for every data packet but several acknowledgment packets were now transmitted back-to-back.

In contrast, performing the same Oracle query using a Web based approach where the SQL*Net traffic stays on the LAN, and only the http traffic passes over the WAN, resulted in improved performance. The http packets were full size without any need for tuning, and up to six packets were transferred before a single acknowledgment was transmitted. The Web based approach was about three times faster than using SQL*Net over the WAN. The Web method transferred about half as much data (due to considerable padding of SQL*Net data). It would, however, not be possible to convert all of the client-server applications to Web technology in the near future. It should also be noted that fancy formatting of the data, such as in a HTML table, would likely result in about the same number of bytes being transferred by both techniques. The SQL*Net vs. http tests are compared in Figure 4. During these tests we monitored the total out-bound bandwidth used on the link (diamonds) and the bandwidth used by the applications under test (circles). Http caused a burst well above the 64k CIR, but finished quickly.

During these tests the time between a burst of data and the associated acknowledgment was usually

between 170 and 350 ms, while the same tests on the LAN gave times between 1 and 8 ms.

References [Stev94] [Stev96] discuss some of the more subtle effects of RTT on network performance such as its effect on TCP window size, timeout, and retransmission, but our simple packet trace analysis made it apparent that RTT was a critical network performance parameter for our client-server applications. We also saw that a significant improvement would result if something could be done on the Oracle/SQL*Net side to enable transmission of more full-size packets.

Setting the Oracle SQL*Net server parameter SDU (Session Data Unit) to 1461 had a much smaller effect than increasing the client's row cache size but resulted in the direct one-to-one mapping of SQL*Net packets to TCP/IP packets. RTT still remains an important parameter that directly impacts performance.

The Role of Internet Traffic

We have found that Internet traffic often consumes a very large portion of the available WAN bandwidth. While there is controversy over the use of Internet usage logs due to privacy and related issues, we have found them to be a very useful tool for managing bandwidth.

At the end of each day we automatically produce a summary of Internet use from firewall logs. The summary includes "Number of Connections and Total Bytes by Network Segment," the "Top 100 Clients" by Number of Connections, Bytes Sent, Bytes Received, and a number of other parameters that do not identify the client's subnet.

The summaries are immediately available via Web pages, and custom reports are e-mailed to network managers with only the information that pertains to the subnets they manage. After being informed of possible problems (by client IP address) through the automated reports, network managers at remote sites have been very successful at reducing unnecessary Internet traffic.

WAN Implementation Suggestions

The following points may be helpful while negotiating with prospective WAN carriers:

- Understand the carrier's network. Get network maps and lists of possible access points.
- Determine what the carrier actually monitors, especially if you are considering a "managed solution." Consider monitoring all important parameters yourself. This will enable your organization to know that they are getting what they pay for, and if they are over utilizing the resource.
- Be sure that you will be able to have read-only SNMP and login access to the carrier's router located on your premises, even if a "managed

solution” is being considered.

- In addition to the usual service level agreement items, such as up-time, repair response, etc., find out what the carrier can specify for RTTs, both minimum and average.
- Make sure that you understand the carrier’s problem resolution procedures and how to escalate a problem to a higher level.
- Find out how the carrier notifies customers of system-wide problems. Is there a Web site with network status information?

Future Plans

We expect to develop the ideas presented here further before going into an automatic-only monitoring mode. In particular we want to investigate the following:

- Installation of Web proxy cache servers at many remote sites.
- Implement priority queuing at the routers to lower the priority of packets with destination addresses outside the corporate network (Internet traffic).
- Lowering the priority of packets based on protocol, such as smtp, ftp, and lpd to give interactive traffic the highest priority in router queues.
- Test setting the DE bit for all Internet traffic so that these frames will not count against CIR.
- Comparing the performance of VPNs over the Internet to the private Frame Relay service. We have already made measurements of minimum RTTs to England over the Internet that beat the Frame Relay minimum RTTs by 10 to 40 ms.
- Consider diverting bulk Internet-bound http traffic to Internet access points provided by the WAN carrier.
- Implementing statistical process controls to provide reasonable alarm triggers when a quality-of-service parameter changes significantly. We have already applied this technique to other types of alarms, such as the number of messages waiting in mail queues with good results.

Conclusion

We have discussed a number of techniques, both technical and administrative, that were employed to improve the performance of two trans-Atlantic WAN links. We also described the analysis of application behavior over these relatively low speed network connections, and the impact of several problems that were uncovered by this study.

Among the goals of this work was to keep the two links running smoothly, to develop methods that could be applied to other WAN links in our company, and to determine the ultimate best-case performance of a given link [Bell92]. Knowing the best-case performance, primarily the minimum RTTs, will help choose technology for future client-server applications (i.e., SQL*Net with PC client, other database

protocols with PC client, remote displays on PCs, Web based, or replicated database servers). By tracking the average and worst-case performance we can estimate how often application performance might be unacceptable. Our efforts have already paid off by eliminating the need to install replicate database servers with their high administration costs at the two European locations.

Through the concept of “Network Ecology,” which brings together the efforts of system and network administrators, applications programmers, and WAN carriers, we were able to improve the performance of our trans-Atlantic links. An important component of this effort was the development of methods to monitor network characteristics. We intend to continue this work by further automating network and application monitoring tools to keep a close watch over WAN performance with only a small demand on System and Network Administrator time.

Availability

The program for performing connectivity checks and routine RTT measurements (up_rtt.monitor) is part of the mon [Troc97] distribution. The programs to measure RTT as a function of packet size (net_validate) and to read tcpdump output (tcpd_read) may be made available in the future. Readers are directed to MRTG [Oet98] for a system that produces Web based reports on router traffic and other parameters.

Acknowledgments

The authors would like to acknowledge Jim Trocki for many valuable discussions and various pieces of software and Eric Anderson for his detailed review of this paper.

Author Information

Jon Meek is Senior Group Leader of Systems, Networks, and Telecommunications at the American Cyanamid Agricultural Products Research Division of the American Home Products Corp. He received BS and MS Degrees in Physics, and a PhD in Chemical Physics all from Indiana University and has worked in Nuclear and Chemical Physics, Analytical Chemistry, and Information Technology. His research interests include scientific applications of Web technology, systems and network management, data integrity, and laboratory data acquisition. He can be reached at <meekj@pt.cyanamid.com> or <meekj@ieee.org>.

Edwin Eichert is Associate Director of Computer Technologies at the American Cyanamid Agricultural Products Research Division of the American Home Products Corp. Ed received a BS in Electrical Engineering in 1970 and a Masters Degree in Management and Technology in 1991 both from the University of Pennsylvania. His early work, as an Engineer at Westinghouse, was in the design of computer systems to control electric power plants. After Westinghouse he

spent several years doing U.S. Navy sponsored research in holography and electro-sensing in fish. In 1976 returned to the computer industry at Fischer & Porter and FMC. His professional interests include scientific programming and managing technical specialists. He can be reached at <eicherte@pt.cyanamid.com>.

Kim Takayama is Network Manager at the American Cyanamid Agricultural Products Research Division of the American Home Products Corp. He received a BS degree in Microbiology from the University of Maine at Orono and has worked as a Genetic Toxicologist for Exxon Biomedical Sciences, followed by seven years of applications development. He is currently in his seventh year of managing networks and systems for Cyanamid. He can be reached at <takayamak@pt.cyanamid.com>.

References

- [Bell92] Steven M. Bellovin, "A Best-Case Network Performance Model," February 1992. <http://www.research.att.com/~smb/papers/index.html>.
- [Bolo93] Jean-Chrysostome Bolot, "Characterizing End-to-End Packet Delay and Loss in the Internet," *Journal of High Speed Networks*, Volume 2, Number 3, pp 305-323, 1993.
- [Bun98] Tim Bunce, "DBD::Oracle – an Oracle 7 and Oracle 8 interface for Perl 5," available from CPAN mirrors, see <http://www.perl.com>.
- [Cava98] James P. Cavanagh, "Frame Relay Applications: Business and Technology Case Studies," Morgan Kaufmann, 1998.
- [List97] P. Lister, "Net-Pcap-0.01," 1997.
- [Nik97] Hrvoje Niksic, "GNU wget" available from the master GNU archive site prep.ai.mit.edu, and its mirrors.
- [McCa97] Steve McCanne, Craig Leres, Van Jacobson, "TCPDUMP 3.4," Lawrence Berkeley National Laboratory Network Research Group, 1997.
- [Oet98] Tobias Oetiker, "MRTG, Multi Router Traffic Grapher," *12th Systems Administration Conference (LISA)*, 1998.
- [Ranu97] Marcus J. Ranum, Kent Landfield, Mike Stolarchuk, Mark Sienkiewicz, Andrew Lambeth, and Eric Wall. "Implementing a Generalized Tool for Network Monitoring," *11th Systems Administration Conference (LISA)*, 1997.
- [Stev94] R. Stevens, *TCP/IP Illustrated, Volume 1: The Protocols*, Addison-Wesley, 1994.
- [Stev96] R. Stevens, *TCP/IP Illustrated, Volume 3: TCP for Transactions, HTTP, NNTP, and the UNIX Domain Protocols*, Addison-Wesley, 1996.
- [Troc97] Jim Trocki, "mon, a general-purpose resource monitoring system," <http://www.kernel.org/software/mon/>.

