NETANALYSIS.ORG, LLC

ONSITE ANALYSIS REPORT

ABC CORPORATION LIMITED (ABCCL)

ONSITE ANALYSIS REPORT

ABC CORPORATION LIMITED (ABCCL)

OVERVIEW

This network analysis report is respectfully submitted by Laura A. Chappell, Sr. Protocol Analyst for NetAnalysis.org, a Delaware Limited Liability Corporation located at 18724 Cox Avenue, Saratoga, California 95070. This report is based on the onsite analysis performed on the internetwork owned and operated by ABC Corporation Limited (hereinafter referred to as "ABCCL").

This report contains confidential information on the status of the ABCCL network and should be treated with care to ensure only trusted parties have access to the information contained herein. <This sample report does not use actual names/numbers.>

The onsite analysis was performed using an NCC LAN Probe. Offsite analysis was performed with Network Associates Sniffer Pro and Novell's LANalyzer for Windows 2.2.

The onsite visit and information gathered for this report was made possible through the efforts of Wally Dauber, Colin Dixon, and Drake Dougherty.

REPORT CONTENTS

This report contains the findings, suggestions and concerns covering:

- Token Ring Network Performance and Health
- Ethernet (Firewall Network) Performance and Health
- Broadcast/Multicast/Anycast/Unicast Levels and Concerns
- IPX/SPX Communications Performance and Health
- TCP/IP Communications Performance and Health
- Application and Upper-Layer Protocol Performance and Health
- Network Design and Data Flows
- Miscellaneous/Other Concerns

Dated: January 1, 2000

LAURA A. CHAPPELL, SR. PROTOCOL ANALYST, NETANALYSIS.ORG.

TABLE OF CONTENTS

SUMMARY OF FINDINGS

Overall, the ABCCL network consists of healthy rings that did not display any signs of congestion or ring faults. The token rotation times viewed were acceptable and the ring poll process completed within acceptable time limits. A failing server card was detected and backup card located to ensure fault tolerance of the ring and server communications when, and if, complete card failure occurs.

The firewall's Ethernet network is reporting excessive collisions (10%) which must be explored further to identify the cause of these collisions.

The IPX/SPX traffic displayed relatively small average packet sizes due to application configuration and programming. Minimal SAP and RIP broadcast traffic was observed; the ABCCL network does not require and would not benefit from NLSP routing at this time (it is understood that the IPX PING ability will be lost without NLSP functioning). Some rogue packets were found to come from devices attempting SPX II connections, but these packets were not utilizing high bandwidth percentages. NCP health was good with no visible Server Overload (0x9999) packets or Burst Mode System Packets. The primary area of concern on the IPX/SPX network is the general service connections that do not appear to be torn down at system shutdown time. Upgrading the client software should fix this problem.

☒ *This observation needs further study/analysis by the onsite analysts.*

There were some evident problems in the TCP/IP area that deal primarily with protocol stack problems and application faults. TCP/IP routing problems are most likely due to what appears to be a stack problem on the OS/2 clients and unnecessary broadcast ARP traffic was being generated by two IP devices contained in the training room. Abuse of the company's Internet connection caused excessive overhead of non-work related traffic – I have been advised that the guilty party is no longer employed by ABCCL.

The network design should be reviewed for the possibility of "single point of failure" occurring on either of the primary switches or on one of the three routers that connect the ABCCL network to other networks. The network should also be reconstructed to remove the WAN bridges and install WAN routers (allowing for DLSW encapsulation of non-routable protocols whenever possible).

The remainder of this report will detail these findings and present packet-level evidence of communication concerns whenever possible.

It is in the best interest of ABCCL to ensure the analysis students are provided with analysis devices to continue studying and optimizing this internetwork.

Laura A. Chappell
Sr. Protocol Analyst
NetAnalysis.org, LLC
lchappell@netanalysis.org
January 1, 2000

Onsite Token Ring analysis indicates very healthy rings. The following section details the health of the network and areas to be watchful of.

RANGE OF TOKEN ROTATION TIME (TRT)

The token rotation time (TRT) indicated a range of speed between 5 and 22 microseconds. The higher the TRT, the less often the ring devices have access to a token to transmit data. The higher TRT was noted on ring 301, as shown in Figure 1. Ring 3 connects to the primary A switch set and the outside trusted and untrusted world through routers and bridges.

***Note**: For your reference, an ABCCL network diagram is appended to the end of this report. <Not included in this sample report.>*

As the rings grow in number of interconnected devices, the TRT will increase. Consider other symptoms, such as receiver congestion, as a good indication of an overloaded ring.

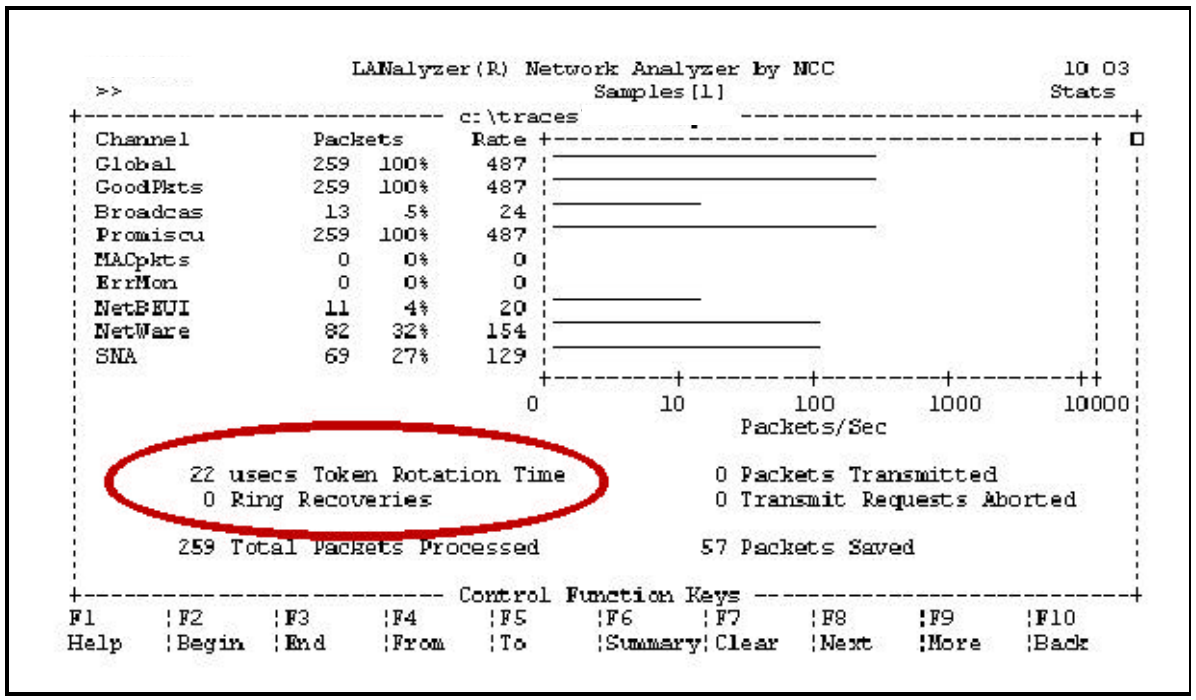

Figure 1: Token Rotation Time indicator.

NO SIGN OF AN OVERLOADED RING

The rings which were tapped into did not indicate an excess number of packets per second. Although an increased TRT indicates more stations active in the process of frame and token repeat, not a single station reported receiver congestion due to too many packets. The stations appear to be able to keep up with the current ring speed.

GOOD RING POLL PROCESS

The ring poll process on ring 301 indicated that the entire process (from initial AMP to final SMP) took 462.158 milliseconds. This process must complete, without fail, every 7 seconds.

LOW TOKEN RING ERROR COUNT

No ring errors were witnessed during the onsite analysis. I recommend that your analysis team periodically check ring health and look for the Token Ring errors defined in the onsite course (Section 3, "Token Ring Analysis"). <Analysis course was presented to the client before analysis report was submitted.>

CARD FAILURE IMMINENT

One of the IBM servers indicated an internal error count that was incrementing at a rate of approximately 5 failures per hour. Drake verified that a second identical card is in the server (unbound) and may be used when the primary card fails. All server card statistics should be reviewed to ensure internal errors are minimal or nonexistent.

☒ *This observation needs further study/analysis by the onsite analysis team.*

ETHERNET (FIREWALL NETWORK) PERFORMANCE AND HEALTH

According to the information I was provided during the onsite visit and the course, the Firewall is reporting an Ethernet collision count of 10%. As we discussed during the course, this counter must be questioned:

- Is this an indication of the number of times the firewall was directly involved in a collision (and therefore had to execute the backoff algorithm)?
- Is this an indication of the number of packets seen by the firewall that are considered fragments (less than 64 bytes with a bad CRC)?

In the first case, we must examine the traffic moving through the switch that the firewall could be witness to, such as broadcasts and multicasts. Are the other devices connected to the firewall's Ethernet switch causing a broadcast storm or sending a high number of packets to an unknown address (or to the firewall's address)?

In the second case, we must consider the effectiveness of the Ethernet switch that the firewall is connected to. The switch should be a 'fragment free' device that recognizes a fragment packet and does not forward it to all attached ports.

☒ *This observation needs further study/analysis by the onsite analysts.*

No live analysis was performed on the Ethernet network.

Concerns arise when too many broadcast and multicast packets are seen because the ABCCL network is a heavily switched (layer 2) network. Since switches typically forward broadcasts to all attached ports, a single broadcast storm could cause major outages on the network. Figure 2 shows a snapshot of the current broadcast traffic seen on Ring 301.

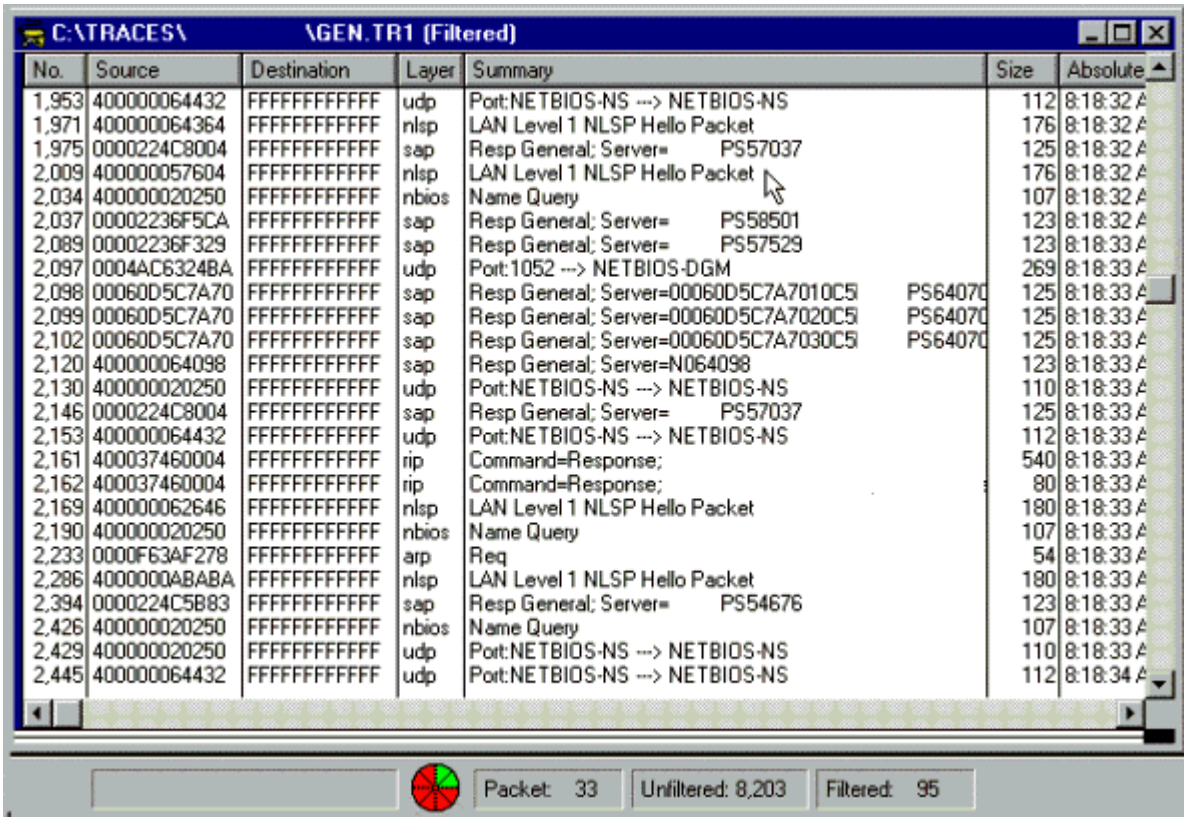

Figure 2: Network broadcasts support a variety of protocols.

BROADCASTS SHOULD BE WATCHED

As you can see from Figure 2, broadcasts are sent by a variety of protocols including NetBIOS, NLSP, SAP, IPX RIP and IP RIP. As the absolute time column indicates, however, the broadcasts are not very frequent (typically 18 broadcasts per second maximum).

The broadcast rate should be consistently checked. Remember that all devices must process and buffer packets sent to the broadcast address (0xFF-FF-FF-FF-FF-FF). Broadcast traffic to watch includes NetBIOS and ARP. NetWare SAP, IPX RIP, and IP RIP are minimal.

SET A BROADCAST/SECOND ALARM THRESHOLD

RMON or analyzer alarms should be set to notify network management staff immediately if the rate exceeds 100 broadcasts/second.

CONSIDER LAYER 3 BROADCAST CONTROL

Consider adding a layer 3 switch or router at some time if broadcasts become a problem or the concern of a broadcast storm indicates the need for broadcast control.

STRANGE BROADCAST BLOCKING TO BE VERIFIED

We observed that not all broadcasts were witnessed on all sides of the 3Com A switch set. This is an indication that either broadcasts are not being forwarded by the switch due to buffering errors, or the switch is somehow filtering these broadcasts from attached rings.

☒ *This observation needs further study/analysis by the onsite analysts.*

OTHER 'CAST PACKETS

Multicast packets were minimal and required to support Token Ring MAC-layer operations.

Unicast packets were the majority of packets seen.

Anycast (used in IPv6) is not in use.

In general, the IPX/SPX network is in good health. There were some strange behavior witnessed that may require further study by the in-house analysis team, however.

ROUTING: RIP IS FINE

RIP's routing abilities will handle your network satisfactorily until you configure parallel paths or connect another 90 networks or so. Although RIP is a distance vector routing protocol and prone to selecting the least desirable path, your network contains no loops and would not benefit from that element at this time. The overhead of link state routing is unnecessary and NLSP should be turned off of the NetWare servers to remove the NLSP Hello traffic from the network. Use INETCFG and select "RIP and SAP only" as the routing type.

Although NetWare 4 and 5 RIP can be configured for different update intervals and packet sizes, your network does not suffer from too much RIP traffic; you would not feel any performance improvements from changing these settings.

SAP TRAFFIC IS ACCEPTABLE

Currently, your SAP rate is acceptable, although it contains extremely small packets. Split horizon technology is causing your SAP packet size to be minimal (with only internal services being advertised most of the time). Although NetWare 4 and 5 SAP can be configured for different update intervals and packet sizes, your network does not suffer from too much SAP traffic; you would not feel any performance improvements from changing these settings.

DUAL ATTACHED SERVER QUESTIONED

There appears to be a server, B4, that contains two network interface cards and is duplicating all broadcast traffic onto the same cabling system. Ensure that both channels into the server are being used for data and performance of this server is acceptable. If not, consider reconfiguring the dual-attached server.

☒ *This observation needs further study/analysis by the onsite analysis team.*

SPX II SOURCE SHOULD BE IDENTIFIED

There appears to be an SPX II application (possibly ManageWise) that is attempting to make an SPX II connection with a local device. The source of the SPX II connection setup packet is 0xF6-29-89-29:00-00-00-00-00-01. The destination is 0x00-00-0C-31:00-00-22-36-F5-9A. We can see in Figure 3 that this is a connection setup packet because the destination connection ID is 65536 (0xFFFF), just a padded field. Although this problem is minimal in overhead (throughput and bandwidth), it should be resolved easily.

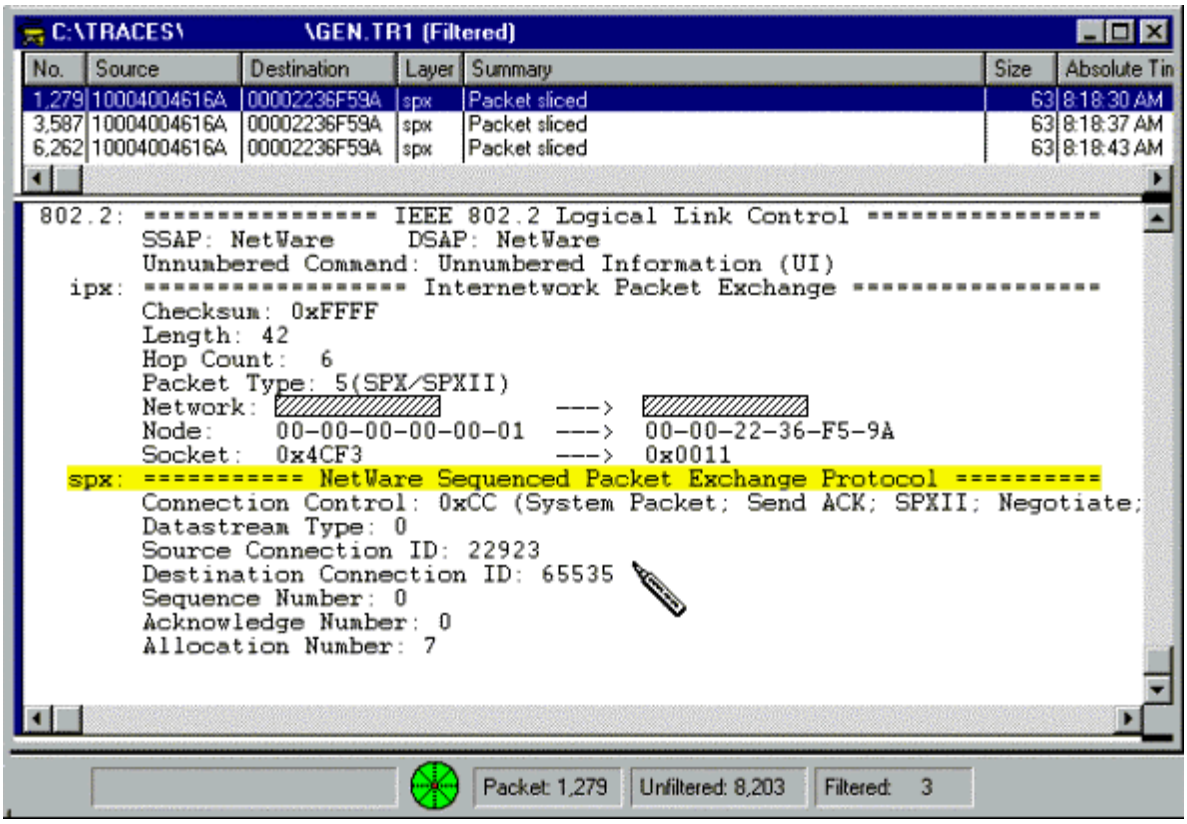

Figure 3: The SPX II connection setup packet never gets answered.

<In some cases, network addresses have been crossed out to maintain confidentiality in this sample report.>

☒ *This observation needs further study/analysis by the onsite analysis team.*

PACKET SIZE DISTRIBUTION IS SMALL

Overall, the ABCCL network has relatively small packet sizes. There are several reasons why this network has such small packets.

- NetBIOS name registration and discovery packets
- SNA traffic
- SPX watchdog traffic
- NetWare 3.11 server backup
- Application coding (i.e., ABCTV)

Since small packet size does not provide the best usage of the network bandwidth, it would be best to try and use larger packet sizes whenever possible.

NCP HEALTH GOOD

At this time, the health of NCP-layer communications is healthy. There are no Server Delay packets (0x9999) or Burst System packets (a subcategory of 0x7777). These two NCP types indicate problems with the processing of client requests and may be due to insufficient RAM or CPU or file I/O problems. Analyzer systems should be set to alert the analyst when Server Overloads reach 5/minute.

The network analysts should be on the lookout for an excessive number of NCP Reply Failures (Completion Code 0x255) to spot NCP processing problems caused by drive mapping errors, misconfigurations or application faults.

CHECK SPX TIMERS

During the class, we noticed an SPX client timer setting that allowed the client to wait up to 127 seconds for an acknowledgment after transmitting data. This timer is excessive and should be reduced to a more reasonable timer of perhaps 5 seconds (setting=90).

SPX CONNECTIONS DENIED SOLUTION

Although I did not witness this problem during the onsite visit, it was brought to my attention that sometimes client stations receive an error indicating that SPX connections are being denied. Since SPX is the most-often patched area of the NetWare operating system, keep up-to-date on the current patches and fixes and refer to support.novell.com for the latest SPX problem listing. You could increase the SPX connections setting at the server by simply 1 and that should cause the server to clear out any old connections that are hanging around internally.

☒ *This observation needs further study/analysis by the onsite analysis team.*

GENERAL CONNECTIONS NOT BROUGHT DOWN

There was evidence of General Service Connections that were not torn down properly. This was most noticeable when we traced an OS/2 station connection process, as shown in Figure 4.

In order to determine the cause of this problem, it will be necessary to shut down the OS/2 station for at least 15 minutes (allowing Watchdog to terminate any residual connections). A new boot-up/connection trace should be taken and the number of General Service Connections granted should be documented. Follow this up with a login sequence and document any more connections that are set up. Finally, trace the connection teardown (turning off the OS/2 client completely) and see where the teardown process is incomplete. Back up and look at the processes that occurred directly after the persistent connection number was granted. This should indicate which process is at fault.

☒ *This observation needs further study/analysis by the onsite analysts.*

Based on the onsite analysis process, it appears that the OS/2 client was particularly susceptible to this connection hold problem. Additional studies should be done to verify that this problem is only linked to the OS/2 clients.

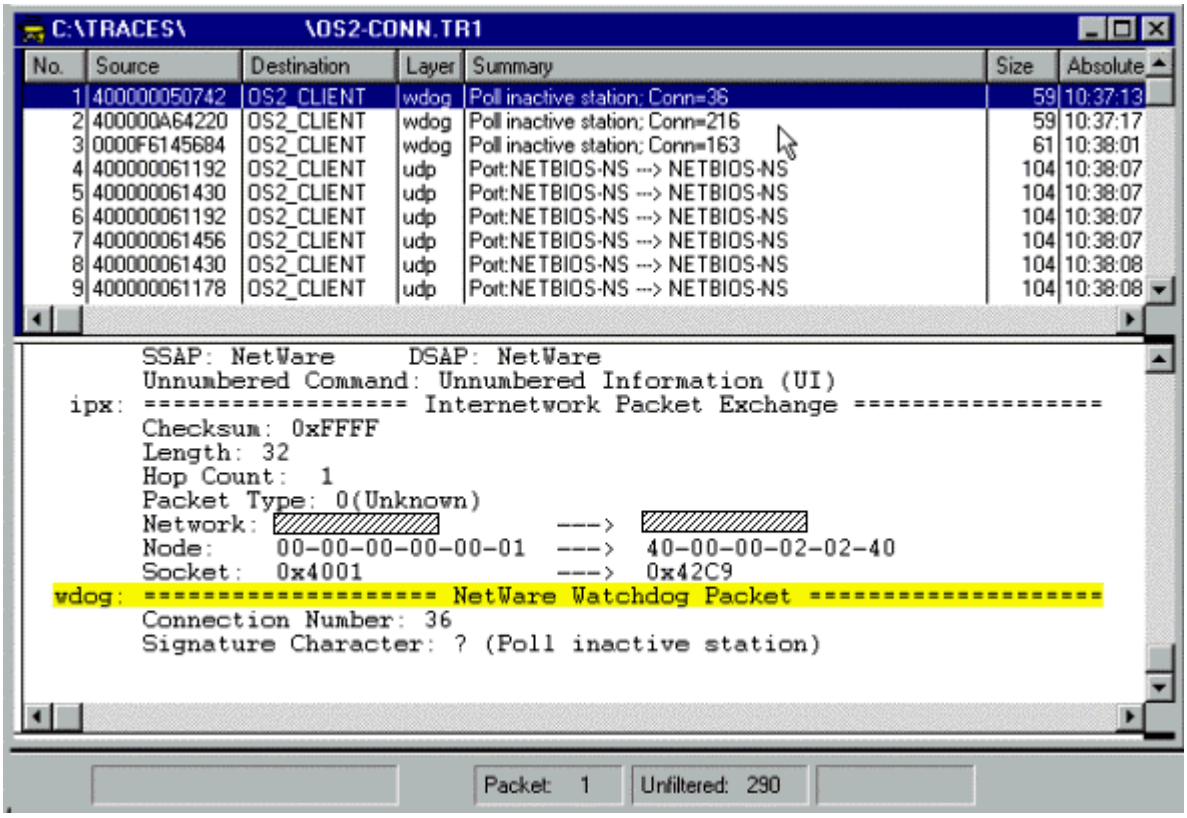

Figure 4: The Watchdog process indicates connections are still being held.

The ABCCL internetwork's TCP/IP communications appeared to have a number of problems that affect overall bandwidth and routing.

UNANSWERED, PERSISTENT ARPS; REMOTE ARP REQUESTS

There appeared to be some faulty stations/applications running during a class that was going on during the onsite analysis. Two stations in particular transmitted ARP packets consistently onto the network without ever receiving an answer. They send ARP requests from the null address (0.0.0.0), which indicates that they do not know their local address. This is not permitted according to the ARP specifications. In fact, several versions of Microsoft TCP/IP stack will not answer an ARP request that has come from the null address. Duplicate IP addresses can occur in this case. Figure 5 shows one of the station's strange behavior.

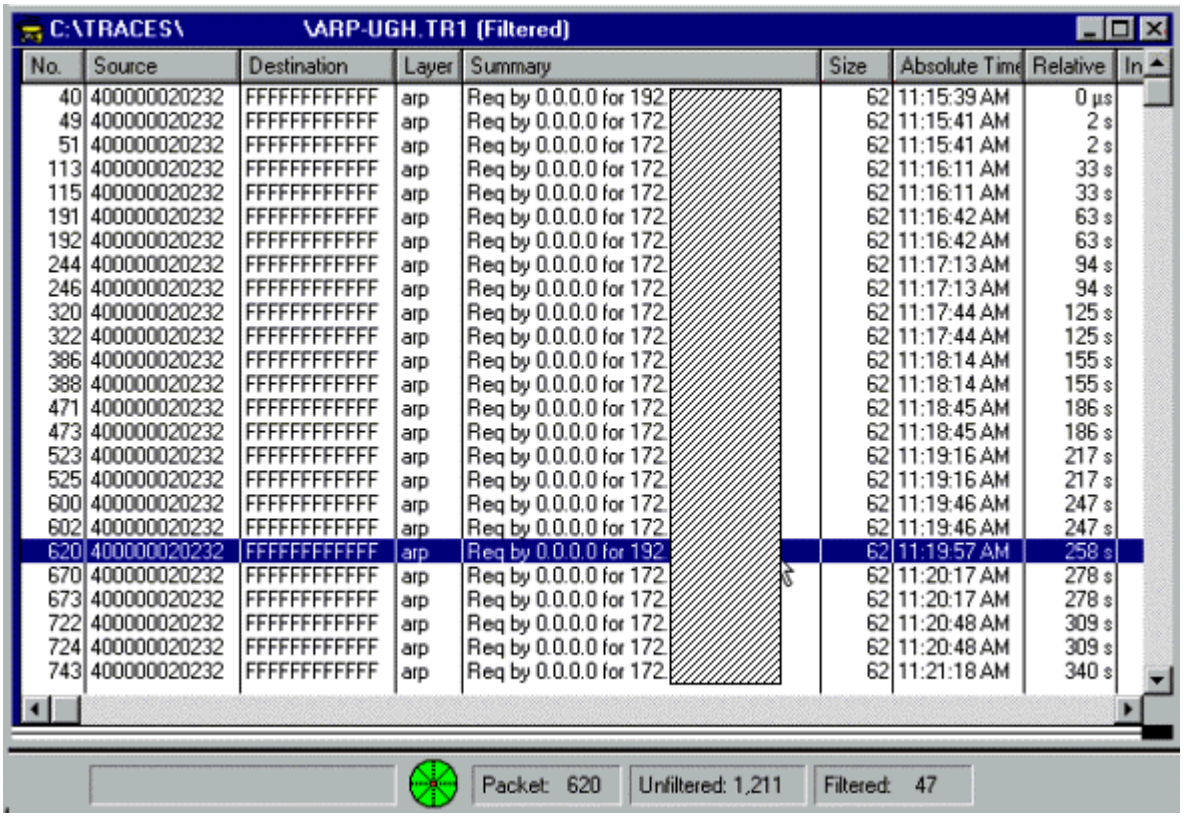

Figure 5: The station's ARP requests go unanswered.

Notice also that the stations transmit an ARP for 192.x.x.x, a device that is obviously on a separate network. This is an indication of an application failure, a routing table failure, or possibly a default gateway problem.

I recommend that you identify the application and protocol stack that was in training during that class and ensure that it is isolated and fully tested before rolling out onto the network.

☒ *This observation needs further study/analysis by the onsite analysis team.*

UNANSWERED ARPS FROM KNOWN IP ADDRESS

Another station (0x00-00-F6-3A-F2-78) transmits ARP requests for a variety of IP addresses, as shown in Figure 6 as if it is building a table or map of IP addresses-to-MAC addresses.

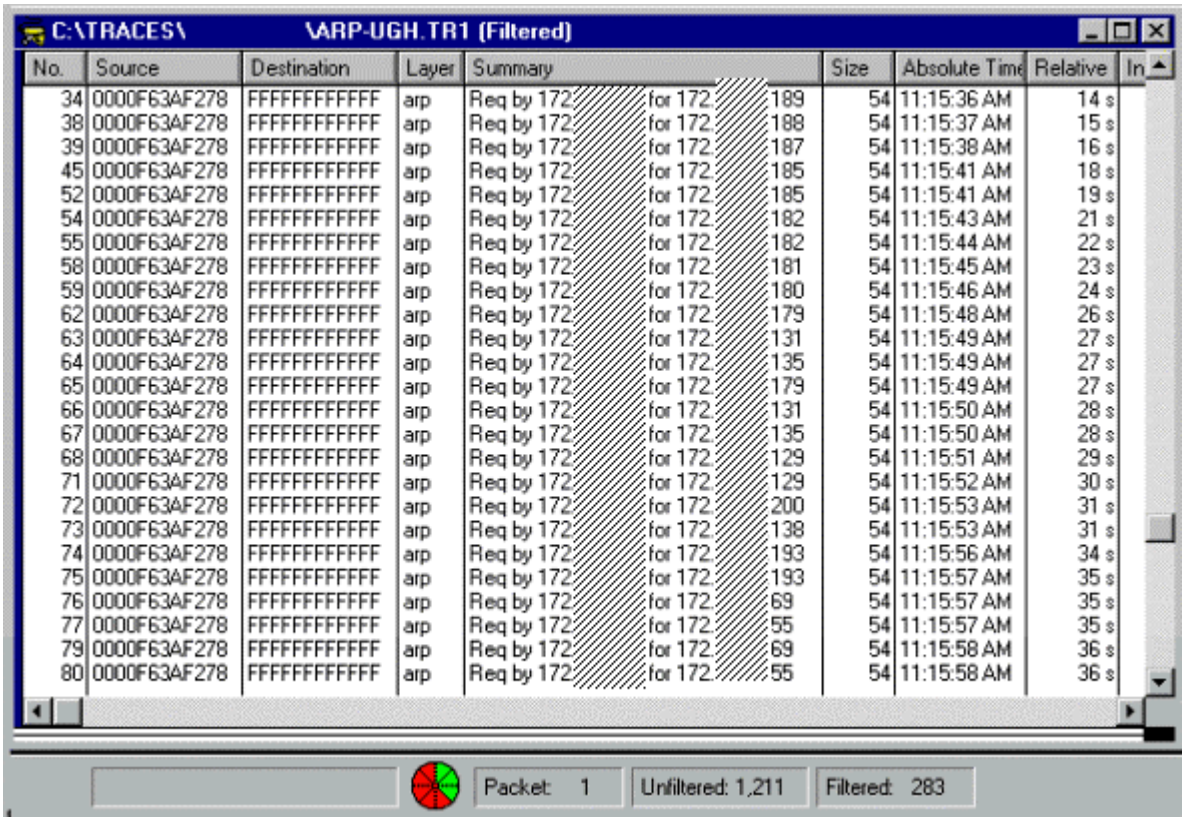

Figure 6: A station appears to be building a MAC-IP table using ARP.

During the course, this transmitting device was identified as the HP OpenView station. Apparently, patches have been applied to this software and certain addresses have been specifically filtered out from processing.

☒ *This observation needs further study/analysis by the onsite analysis team.*

Check the OpenView station to ensure proper configuration and consider taking the trace file to HP for further investigation.

ARPING FROM 0.0.0.0

The OS/2 workstation also appeared to have a problem of sending ARP packets from 0.0.0.0. Check with IBM to see if an IP stack patch exists. The source IP address of 0.0.0.0 is not permitted. Even if a station does not have an address yet and must seek out a DHCP host, it will not ARP for the DHCP host. It will transmit a broadcast DHCP discovery packet immediately (without a preceding ARP).

ICMP REDIRECTS INDICATE PROBLEM

An interesting ICMP redirect packet indicates a definite routing problem on the network. Notice in Figure 7, that the router (0x00-00-40-D8-2F-3F is sending an ICMP reply indicating that the best route to 172.x.x.x is through 172.x.x.x. This does not make sense at all. The device advertising its IP address as 172.x.x.x should be checked to verify all routing configuration and stack functionality. This is most likely part of the problem that has plagued the OS/2 devices and their routing processes.

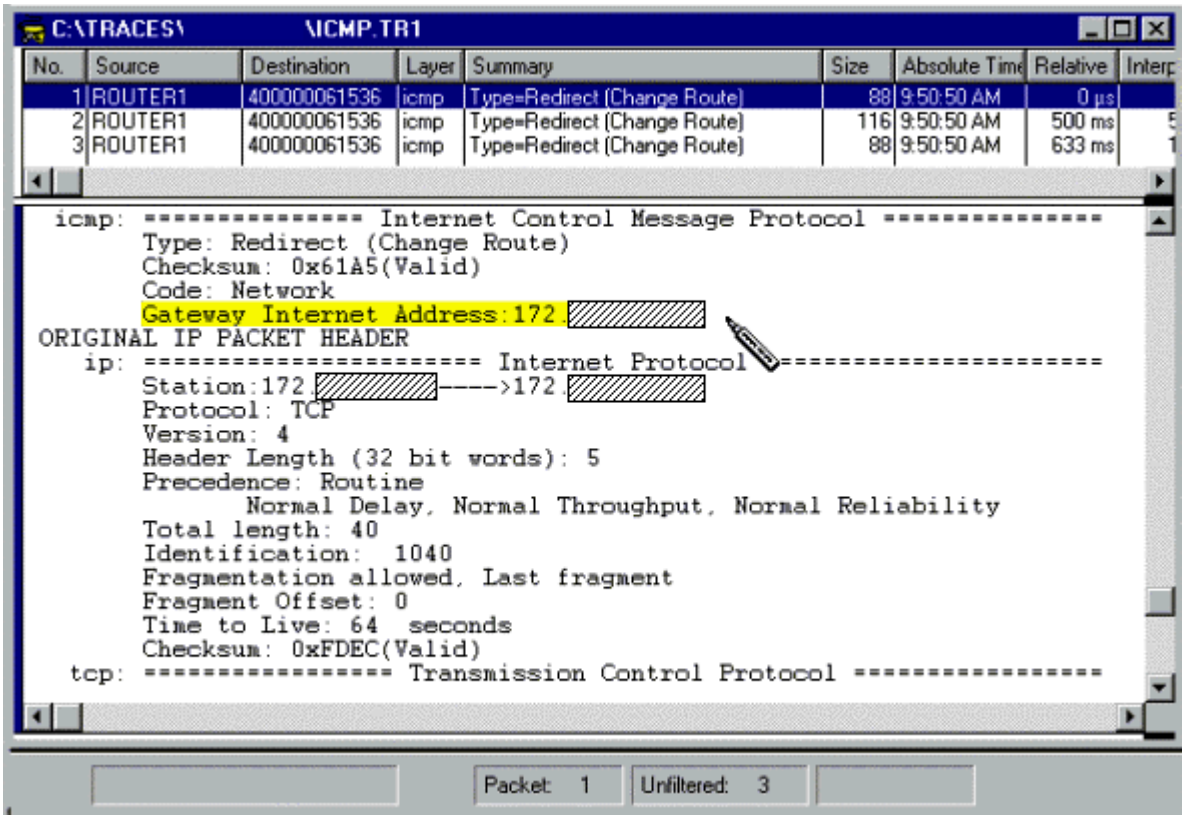

Figure 7: The redirect packet points to the sender as the proper gateway.

STATION HAS PROBLEMS WITH ROUTING UPDATES

RIP 1 updates occur every 30 seconds by default. Non-routing stations should silently discard these broadcast updates. Routing stations absorb the RIP information and compare the incoming data with the existing table entries. One of the stations (0x40-00-00-06-28-78) sends an ICMP reply indicating that the desired route information port (UDP port 0x0208) does not exist on the receiving device, as shown in Figure 8. Why does it answer the broadcast update? This station should be checked for any configuration faults or applications that require routing information.

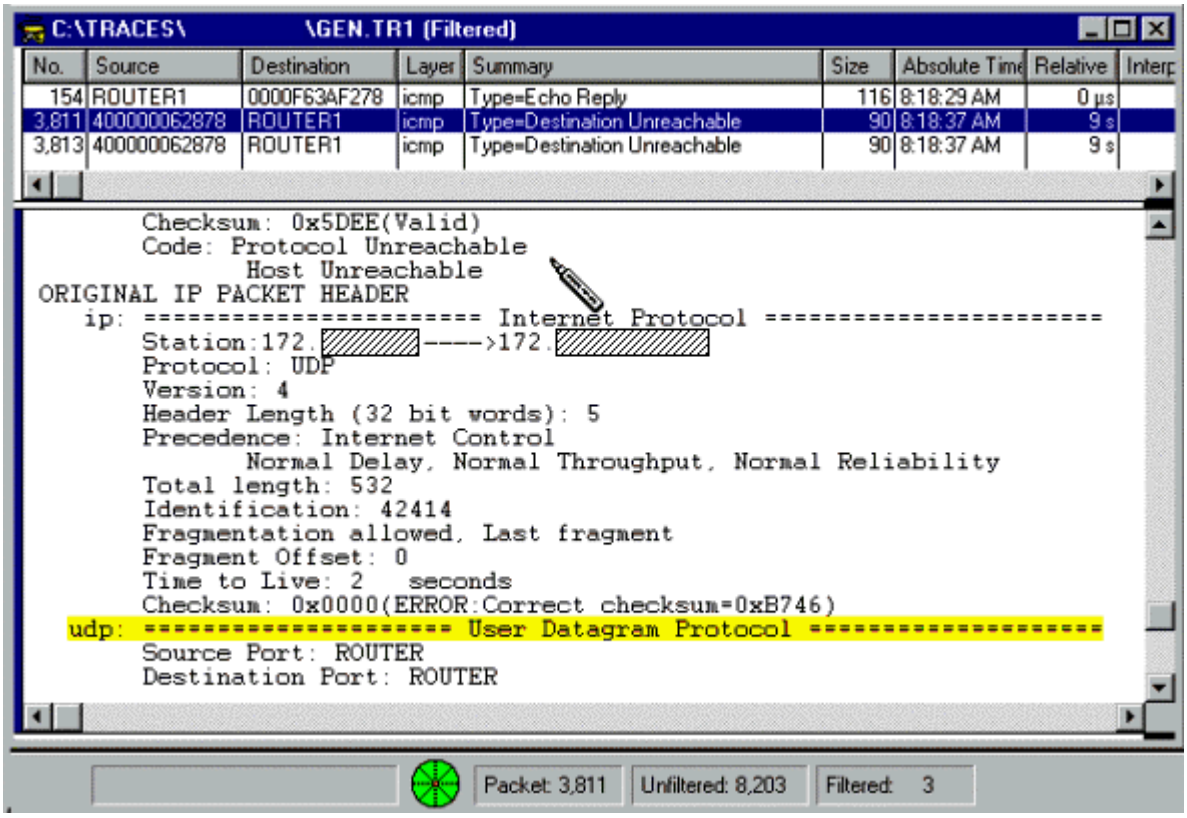

Figure 8: The station does not appreciate the RIP broadcasts.

INTERNET DOWNLOADS: PROMISCUOUS MODE

Consider the bandwidth of the network your information artery. The bandwidth should be protected from unnecessary traffic whenever possible. On the morning of the onsite analysis, a user was downloading an X-rated video from the Internet. The site URL, banner, and advertising information downloads in ASCII text, easily visible on an analyzer screen, as shown in Figure 9.

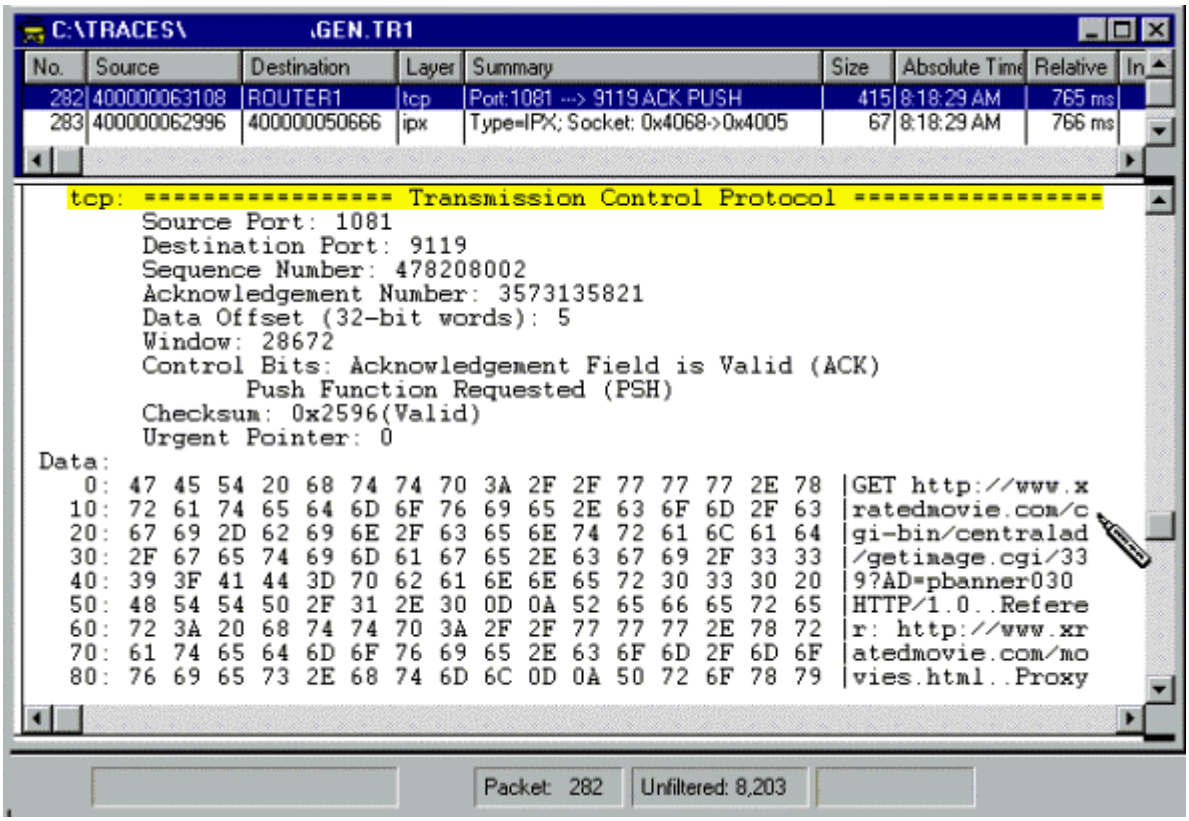

Figure 9: The site URL is easily visible on an analyzer.

RIP1 OK IF NO SUBNETTING NEEDED

One of the natural questions on the TCP/IP network is: OSPF or RIP? Currently, the ABCCL internetwork supports RIP1 only (distance vector-based routing without subnet masking ability). Since the network has very few network addresses being advertised locally, there is little advantage to supporting a link state routing protocol such as OSPF. Consider the same argument for IPX RIP v. NLSP.

If, however, the company intends to install one or more routers inside the 172.x.x.x network while maintaining one single IP network address for the entire internetwork, then subnet masking will be required. In this case, RIP 2 (supporting subnet masking information inside the updates) or OSPF will be required.

There were some applications that used minimal packet sizes and didn't support windowing. Maximum packet size usage ensures the most efficient use of bandwidth, card usage (frame repeat mode) and switch CPU time. Larger packet payload (data portion) and fewer packet headers need to be created or processed by the network stations and interconnecting devices.

ABCTV is a perfect example of an application that chews up throughput by reading 84 bytes of data at a time (perhaps caused by reading a single line of data at a time). Figure 10 shows a typical ABCTV communication sequence.

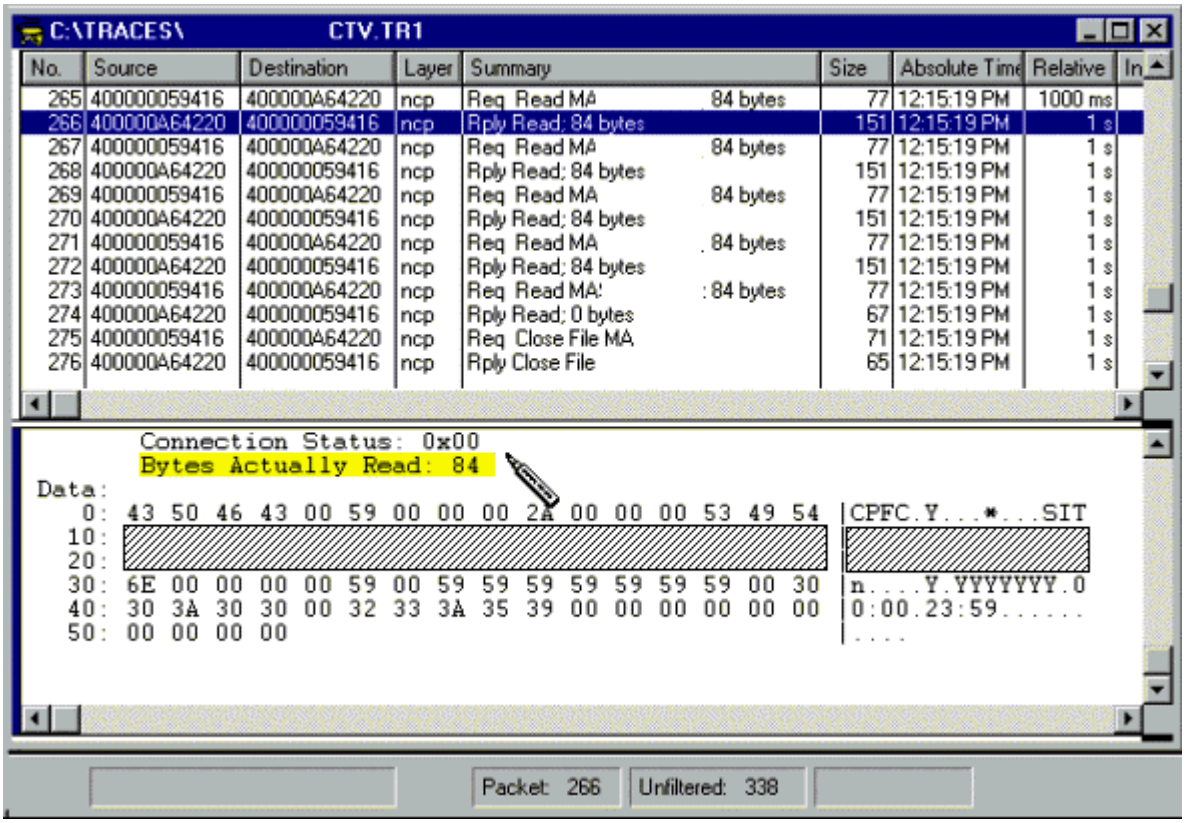

Figure 10: ABCTV uses minimal packet sizes during a read operation.

The onsite analysis team should check out any applications that are going to be deployed company-wide to check the packet size and windowing capabilities. In many cases, applications run over TCP/IP provide better use of the bandwidth through TCP's data streaming capabilities. Some applications have configuration settings that can affect performance. Run an analyzer on the cabling system after configuration settings have been altered; document the affect.

Another example of an application that has a negative affect on throughput is the IPX-based backup that occurs in the middle of the day. This backup does not take advantage of burst mode communications because it is connecting to a NetWare 3.11 server which does not support burst mode by default. It is highly recommended that you upgrade the NetWare 3.11 server.

The final page in this report contains the network diagram presented during the onsite visit.

Predominantly switched networks are susceptible to broadcast storms. It will be important to track broadcasts and multicasts to ensure they do not cause communication bottlenecks. Of course, if broadcast filtering is enabled on the switches, this could alleviate this problem.

Because there are two WAN connections setup with bridges, a broadcast storm will also affect devices on the other side of the WAN bridges. Once the WAN bridges are replaced with routers, the broadcast storm concern will be lessened. Broadcasts on each side of the WAN link will be answered locally or silently discarded.

The ABCCL internetwork supports a high number of NetBIOS packets. The NetBIOS protocol is a Type 20 forwarded protocol. If a loop occurs on the network, these packets can circulate around the network for up to 8 hops before being discarded. They will be propagated across WAN links as well. If you configure any network loops (for redundancy perhaps), consider the NetBIOS traffic patterns and take actions to ensure only a single path for data is resolved. An easy way to test this is through a 5-minute analysis session. If the same packet appears in the trace buffer (except for an increment in the hop count field), you may have circulating NetBIOS traffic.

MISCELLANEOUS/OTHER CONCERNS

Currently, there is an outstanding issue that is still being researched. There appears to be an 802.2 poll consistently checking stations. This poll is active at the moment that a workstation boots up, so it appears that an external process is ongoing and not affected by the connection establishment or teardown process of NetWare.

Results of this research should be available within the next several weeks. A report addendum will follow.

SUMMARY

In summary, the ABCCL network was surprisingly clean at the Data Link layer (Token Ring), with minor concerns of an 802.2 polling process. The network broadcast rate should be checked often to ensure the rate does not become excessive. The IPX/SPX protocol stack appears to have some faults which are credited to application-layer issues. The TCP/IP stack, on the other hand, indicates some configuration, IP stack and possibly application-layer problems.

Areas that should be inspected and updated by the onsite analysts are marked herein.

ABCCL NETWORK DIAGRAM