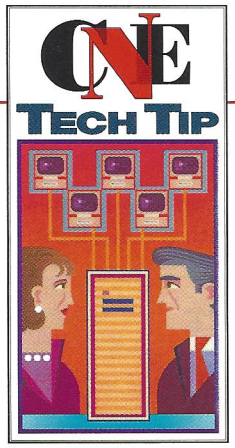

On the Wire Again . . .

By Laura Chappell

Your network seems to be running perfectly, but is it? Learn how to analyze network traffic and find out what's really happening on the network.

Editor's Note: Because this article is geared toward CNEs and Master CNEs, it assumes readers have a basic understanding of network communications and are familiar with network analyzers. For information about basic troubleshooting techniques, read "Troubleshooting: Identifying and Eliminating Problems on Ethernet Networks" (NetWare Connection, Nov./Dec. 1993), "Troubleshooting: Identifying and Eliminating Problems on Token Ring Networks" (NetWare Connection, Jan./Feb. 1994), and "Troubleshooting: Analyzing and Optimizing NetWare Communications" (NetWare Connection, Mar./Apr. 1994). You can download these articles from NetWare Connection's Networking Home Page (http://www.mcrlabs.com/net/documents.html).

Network communications problems can often be overlooked on a growing and constantly changing network. But don't be lulled into a false sense of security just because no one is complaining about poor performance or an inability to connect to a server. With few exceptions, the typical NetWare network has some hidden characteristic or configuration flaw that is affecting performance, albeit only slightly in some cases.

Over the past several years, my company has been hired to characterize the performance of both simple and sophisticated networks. We often find that overall the network is healthy, but some areas definitely warrant improvement.

To analyze network performance, we use a network analyzer to look directly at the communications occurring on the network. Although most network support providers use a network analyzer to troubleshoot problems that occur at the Data Link or Physical layer of the Open Systems Interconnection (OSI) model, we use an analyzer to troubleshoot problems that occur all the way up to the Application layer. This article explains the basic steps in characterizing your network traffic. In addition, it describes three examples of common network problems that we have found at our customer sites.

Network Characterization

When a client hires my company to evaluate a network, we request at least 24 hours of observation time on the network. A week would be better; a month would be ideal. Unfortunately, most clients want their network characterized and analyzed as of yesterday.

We use every analyzer tool we can find as well as every conceivable connector type. Our primary analysis tool is still LANalyzer for Windows, which a software-based analyzer available from Novell, and Network Probe (formerly called LANalyzer), which is a hardware/software-based network analyzer available from Network Communications Corporation.

Our LANalyzer for Windows system is a Toshiba portable computer with 8MB of RAM and a 250MB hard drive. Our Network Probe system is a COMPAQ portable (lunchbox) computer with a Token Ring Network Probe board, an Ethernet Network Probe board, 16MB of RAM, and a 500MB hard drive. And because we have found that our clients don't always supply the correct connector, we bring twisted pair and 9-pin connectors for Token Ring networks and twisted pair, BNC connectors, and thicknet connectors for Ethernet networks.

Pinpointing Basic Network Problems

When we begin our network analysis, we use a characterization checklist, as shown in Figure 1 on page 36 , and try to fill out as much information as possible during the first 24 hours. If the network is experiencing a basic problem, such as a bad network interface board or faulty cabling, the characterization checklist will help us spot it almost immediately. For example, one network we evaluated was receiving a high number of

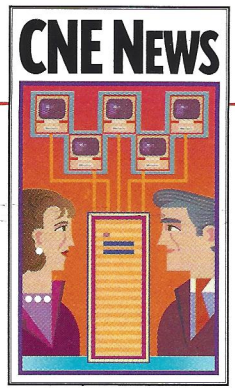

ECNEs certify as Master CNEs, the CNA program reaches 50,000, and Novell Education trains end-users.

ECNE Continuing Education Requirement

Novell Education has announced a continuing certification requirement for Enterprise CNEs (ECNEs) that will allow them to migrate to the Master CNE program. The Master CNE program, which was announced in April, is Novell's new high-level certification. It requires a CNE to demonstrate additional knowledge in a task-oriented area of specialty. These areas of specialty include Network Management, Infrastructure and Advanced Access, and GroupWare Integration.

The continuing education requirement for ECNEs consists of a single exam, *The Fundamentals of Internetwork and Management Design*. ECNEs who complete this requirement will be recognized as Master CNEs with the Infrastructure and Advanced Access specialty. To assist ECNEs in fulfilling this requirement, Novell will send each ECNE two complimentary coupons—one to take the *Fundamentals* test and one to receive the corresponding student kit, which includes the study materials necessary to prepare for the exam. The test coupon is redeemable at local Novell Authorized Education Centers (NAECs), and the student kit is redeemable at any Drake or Sylvan testing center.

ECNEs who have already taken *The Fundamentals of Internetwork and Management Design* can use the complimentary coupons they receive for any other test offered in the Master CNE electives and any other student kit. ECNEs have until January 15, 1996 to complete the continuing certification requirement.

50,000th CNA

The Certified Novell Administrator (CNA) program has certified 50,000 CNAs, the front-line support professionals who perform administrative tasks and troubleshoot basic network problems. Although the CNA program was implemented only 2 1/2 years ago, it is already Novell's fastest-growing program, adding 10,000 new CNAs every quarter.

The 50,000th CNA is Fiona Desruisseaux of Ottawa, Canada. Ms. Desruisseaux is putting her certification to work in her new job as a network administration consultant.

New End-User Courseware and Certification

Novell Education is now providing training and certification to end-users of Novell's applications products. These end-users can obtain a Novell Productivity Specialist certification by demonstrating their competency with Novell applications. Novell Productivity Specialists must be able to apply integrated, cross-product solutions to a variety of tasks that increase productivity.

To provide end-users with the skills they need to certify, Novell is offering training based on its new *End-User* courseware. Designed to meet the diverse needs of end-users, the courseware consists of modular blocks. This modular system includes one introduction module—the Perfect Office *Introduction for New Users*—and multiple add-on modules that encompass tasks performed in PerfectOffice, WordPerfect, Quattro Pro, Presentations, Paradox, Envoy, InfoCentral, GroupWise, InForms, and SoftSolutions.

As a result of the courseware's modular design, end-users can focus on those tasks that are important to their work, and instructors can tailor their classes to meet student needs. The courseware also provides flexibility by teaching tasks that span Novell's applications suite, rather than teaching single applications and features. This task-oriented focus enables instructors to build their classes around a specific audience, such as a word processing pool, a group of executive officers, or administrative assistants.

The first Novell *End-User* courseware modules began shipping in May 1995, and additional modules will be released through August 1995. Courseware and instructor-led training are available through NAECs, Novell Education Academic Partners (NAEPs), and Novell Application Training Providers (NATPs).

For more information about the Novell *End-User* courseware or the Novell Productivity Specialist certification, or to find a Novell training partner near you, call 1-800-233-EDUC. Outside the U.S. and Canada, call 1-801-429-5508. ■

Time Elapsed	1 hour	4 hours	24 hours	Other
1. Average Utilization %				
2. Average Error Rate per Second				
Type: _____				
Type: _____				
Type: _____				
Type: _____				
Type: _____				
Type: _____				
3. Protocol—Average Utilization %				
NetWare				
TCP/IP				
AppleTalk				
SNA				
_____				
_____				
_____				
_____				
4. Broadcast traffic %				
Protocols: _____				
5. Average Packets per Second				

6. Active Servers, Routers, and Stations (after 24 hours only)

Most active servers	1. _____	Utilization %: ____
	2. _____	Utilization %: ____
	3. _____	Utilization %: ____
	4. _____	Utilization %: ____
	5. _____	Utilization %: ____
Most active clients	1. _____	Utilization %: ____
	2. _____	Utilization %: ____
	3. _____	Utilization %: ____
	4. _____	Utilization %: ____
	5. _____	Utilization %: ____
Most active routers	1. _____	Utilization %: ____
	2. _____	Utilization %: ____
	3. _____	Utilization %: ____
	4. _____	Utilization %: ____
	5. _____	Utilization %: ____

Figure 1. *When you are analyzing a network, you should fill out as much information as possible on your characterization checklist during the first 24 hours.*

Cyclic Redundancy Check (CRC) errors attributed to a single network node. The user was aware that it was taking longer than normal to load applications but had not yet reported the problem to the IS group. The obvious solution was to replace the network interface board, and no more CRC errors occurred.

More elusive network problems, however, require more in-depth evaluation. To pinpoint these problems, you must carefully examine network traffic at the packet level.

Looking at Network Communications

You can often detect problems by taking a quick look at a network communications session. For example, using Novell's LANalyzer for Windows, you can fill up a trace buffer with general traffic (do not apply a packet filter). Then in the summary window (the top portion of the Packet Display window), you double-click on a line that looks like a standard NetWare Core Protocols (NCP)-type communication, using an NCP request or an NCP reply. (A summary window is shown in Figure 6 on p. 42.) The Display Filter window appears, with the desired communications session defined in the station box. (See Figure 2 on p. 38.)

A healthy communication consists of an NCP request followed by an NCP reply. If burst mode is being used and the client station is requesting a file, you may see a single request followed by multiple replies. If NCP requests are repeated without an NCP reply in between, however, this can signal a problem on the network. Let's take a look at what should happen when a client sends an NCP request to a server.

First, the client software (NetWare shell or NetWare DOS Requester) sends an NCP request to the server. The client's IPX retry counter is set to 0, and the IPX retry timer begins ticking. If no response is received by the time the retry timer reaches its timeout, the client increments the retry counter by 1 and reissues the request. (You can configure the IPX Retry Count in the NET.CFG file.) Then the client retry timer is reset and begins ticking again. If a reply is not received by the time the retry timer expires, the client follows the same steps: it increments the IPX retry counter, resets the retry timer, and reissues the request. This process is repeated until either the server responds or the retry counter limit is reached.

If the server responds, the client considers the NCP request successfully sent and answered, and begins the next task defined by the client software. If the server does not respond, however, this can indicate a problem with the cabling to the server, with a server component, or with a server function.

If a network is experiencing this kind of problem, you should first check to see if the server is dropping request packets. On the NetWare server, launch the MONITOR NetWare Loadable Module (NLM) and view the LAN driver statistics. Under the Generic Statistics heading, check the No ECB Available Count value. This statistic is incremented when packets are received, but the server discards the packets because it is out of available communication buffers to receive and process packets. If this number is

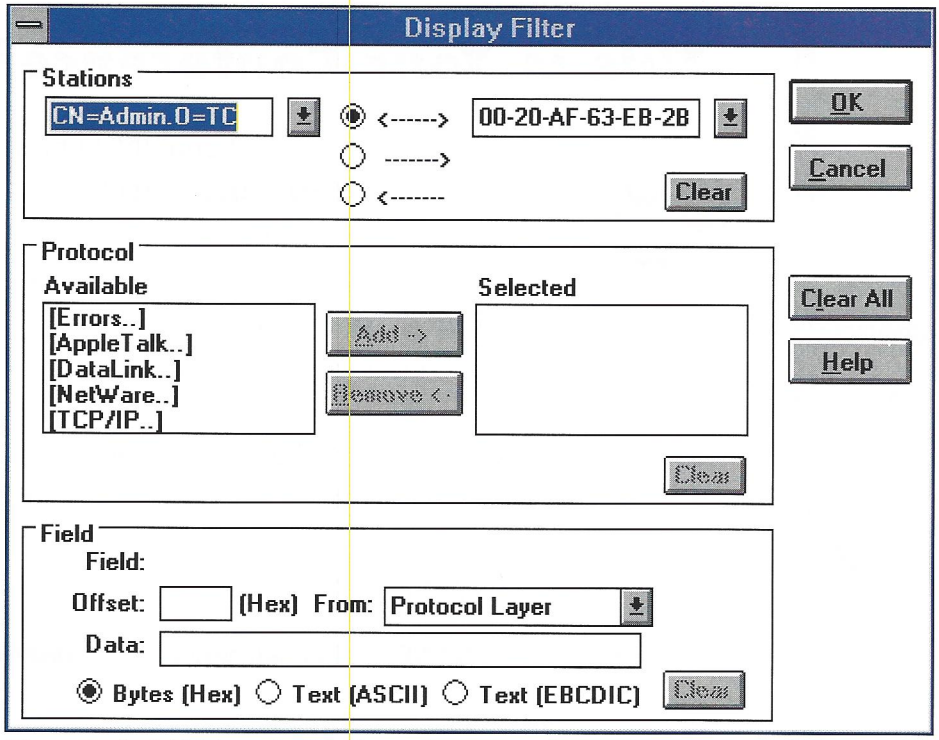

Figure 2. In LANalyzer for Windows, the station information is automatically defined when you double-click on a line in the Packet Display window.

increasing at a rapid rate, you should probably increase the number of maximum packet receive buffers. You may have to add server memory to allow for more buffers.

If the No ECB Available Count is not incremented when these unanswered requests are happening, you should check to see if the request is making it to the server's segment. To do this, move your analyzer to the affected network segment and check for traffic being sent from the client station to the server. Be sure you are looking in the IPX header for the source address because the local router replaces the source Media Access Control (MAC) header address with its own. (With LANalyzer for Windows, simply open up any packet, double-click on the source node address field in the IPX header, and input the transmitting station's node address.)

If no packets from the transmitting station appear on the local segment, the problem is not occurring at the NetWare server. You must then check the cabling and intermediary devices along the path between the client and the server. Follow the packet from the client station to the server to see where it is discarded. The packet may be discarded because a router is overloaded and is discarding packets or because a concentrator along the path has a faulty port and is corrupting packets.

Once you have examined a basic network communications session, you should look at some overall traffic. Run your network analyzer without any capture filters to get a sample of current traffic. The summary information should give you some idea if anything unusual is happening on your network. To illustrate how analyzing network traffic can pinpoint problems, the rest of this article outlines some of the more common problems we have seen on our clients' networks this year.

Duplicate Requests, Duplicate Replies

One common problem we have seen is wasted bandwidth on a single-segment network due to unnecessary duplication of requests and replies. For example, packet 52 in Figure 3 shows that Christy issued an NCP request to get the effective directory rights to a file named DEMO (the complete filename does not appear in the window). This request was transmitted to server CORP1.

Packet 53 shows the same NCP request, but it came from server CORP1, and it was addressed to server SALES1. In packet 54, SALES1 answered CORP1's request, and in packet 55, CORP1 replied to Christy's request (packet).

If you examine the packets more closely, you will notice that packets 52 and 53 (the request packets) are the same packet except for the MAC header and the hop count field in the IPX header. (See Figure 4 on p. 40.) The IPX header of packet 53 shows that the actual source of the packet is Christy and that the packet crossed one router (hence the hop field count of 1). Packets 54 and 55 (the reply packets) show similar differences: The MAC header is different, and the hop count field value is different, but all other field values are the same.

Why are there duplicate requests and duplicate replies? A simple frame type misconfiguration is the culprit. Christy and the server she is communicating with use different frame types. Because Christy's workstation uses the Ethernet 802.3 frame and SERV1 uses

Capture Buffer

No.	Source	Destination	Layer	Summary
52	CHRISTY	CORP1	ncp	Req Get Effective Dir Rights DE
53	CORP1	SALES1	ncp	Req Get Effective Dir Rights DE
54	SALES1	CORP1	ncp	Rply Get Effective Dir Rights
55	CORP1	CHRISTY	ncp	Req Get Effective Dir Rights
56	CHRISTY	CORP1	ncp	Req Set Dir Handle DEMO
57	CORP1	SALES1	sap	Req Set Dir Handle DEMO
58	SALES1	CORP1	ncp	Rply Set Dir Handle
59	CORP1	CHRISTY	ncp	Req Set Dir Handle
60	CHRISTY	CORP1	ncp	Req Search for File DEMO\SETUP.
61	CORP1	SALES1	ncp	Req Search for File DEMO\SETUP.
62	SALES1	CORP1	ncp	Rply Search for File SETUP.EXE
63	CORP1	CHRISTY	ncp	Req Search for File SETUP.EXE
64	CHRISTY	CORP1	ncp	Req Open File DEMO\SETUP.EXE
65	CORP1	SALES1	ncp	Req Open File DEMO\SETUP.EXE
66	SALES1	CORP1	ncp	Rply Open File SETUP.EXE
67	CORP1	CHRISTY	ncp	Req Open File SETUP.EXE
68	CHRISTY	CORP1	ncp	Req Close File SETUP.EXE
69	CORP1	SALES1	ncp	Req Close File SETUP.EXE
70	SALES1	CORP1	ncp	Rply Close File
71	CORP1	CHRISTY	ncp	Req Close File

Figure 3. Duplicate requests and duplicate replies indicate a possible communications problem on the network.

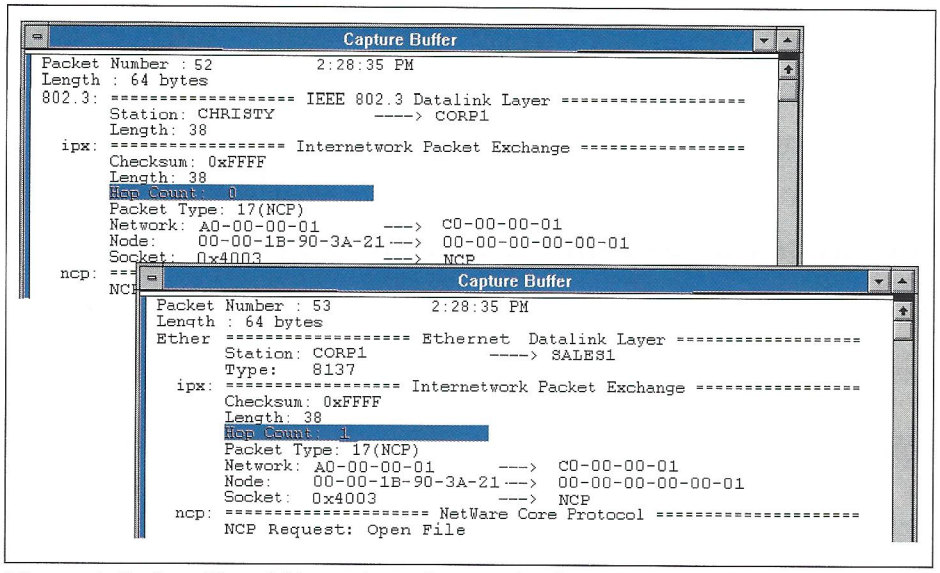

Figure 4. *Packets 52 and 53 are identical except for the MAC header and the IPX header hop count field value.*

the Ethernet 802.2 frame, CORP1 is acting as a router between these two different virtual networks.

The solution is simple if Christy's station is the only client using CORP1 as a router between frame types: You would simply edit her NET.CFG file to use the Ethernet 802.2 frame type.

If all stations on the local network are using CORP1 as a router to get to SERV1, however, it might be easier to configure SERV1 to use the Ethernet 802.3 frame type so the server supports both frame types. If all stations are using CORP1 as a router, you should see utilization drop by almost 50 percent once the problem is fixed.

Ideally, you should configure all the clients and servers to use the more flexible frame type, Ethernet 802.2. All of Novell's current products now support the 802.2 frame. Because the Ethernet 802.3 frame only supports NetWare's IPX/SPX protocol, this frame will be limiting in the long run because it does not enable you to use multiple protocols on your network.

SPX Everywhere!

As we analyze networks, we also frequently see communications sessions that are riddled with SPX packets that don't contain any data. In such instances, an SPX application may be creating problems. Ideally, SPX applications establish a connection just before they need to send data and close the connection immediately after the data has been transmitted successfully. Some applications, however, use SPX Keep Alive (or SPX Watchdog) packets to keep the SPX session active 24 hours a day, 365 days a year. These applications can cause problems if they operate across an on-demand WAN link. If SPX sessions are running properly, they should follow these steps:

Step 1: Establish the Connection

The source SPX device first transmits a connection request packet to the destination device. This packet is actually an ACK (acknowledgment) packet that is requesting an acknowledgment in response. The source device assigns itself a connection ID number for use in these SPX communications. As shown in Figure 5, Fred's station has assigned a source connection ID number of 36926.

The source station fills the destination connection ID field in the SPX header with 65535 (in binary code, one, two-byte field filled with 1s) because it does not know what connection ID number the destination will assign itself.

Upon receipt, the destination station responds with an ACK packet. In the acknowledgment packet, the station includes its own source connection ID number (the station is the source of this packet), and it includes the connection ID of the intended recipient in the destination connection ID field. When Fred's station receives this ACK packet, it considers the connection established and can now send or request data.

Step 2: Send Data

Using an SPX session, a station transmits data and includes an ACK request in each packet it sends. The receiving station transmits an ACK packet back for each data packet successfully received.

Step 3: Close the Connection

When the stations have transmitted all their data to each other, the station that started the SPX connection makes a request to terminate the session. The End Connection request is answered with an End Connection acknowledgment packet.

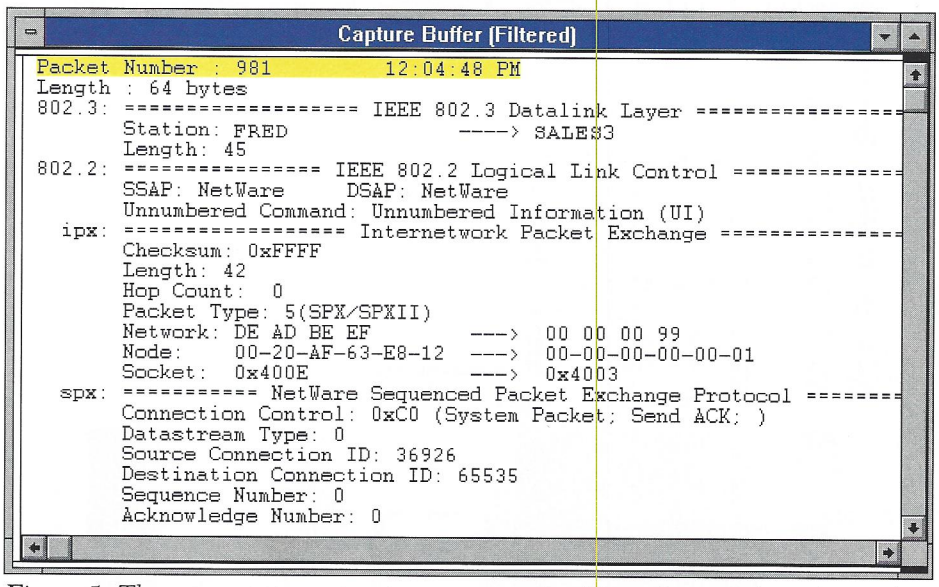

Figure 5. *The source station assigns its own connection ID number and places the number in the SPX header's source connection ID field.*

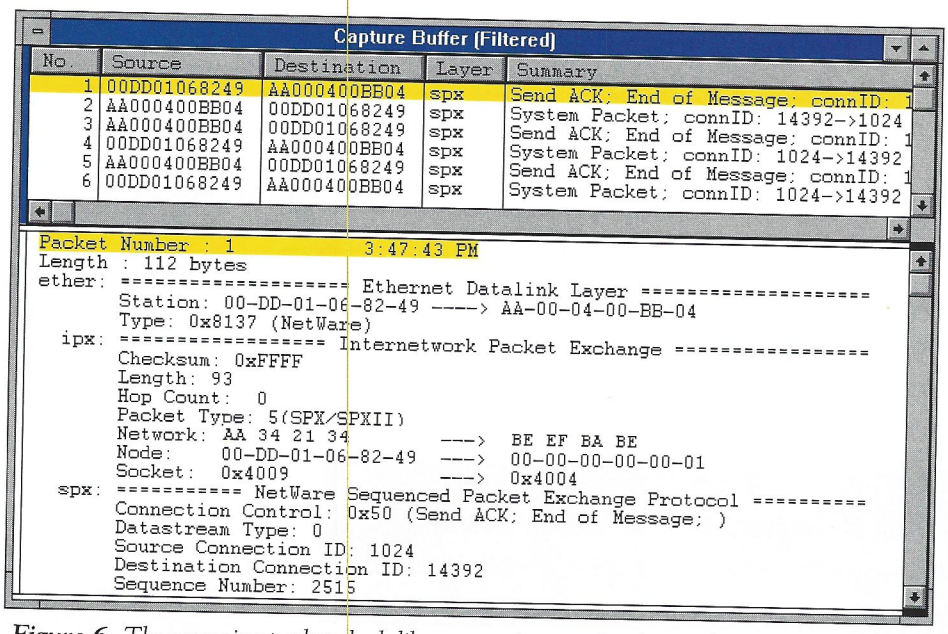

Figure 6. *The recurring packets look like connection termination packets.*

This SPX session was open only as long as necessary and was closed immediately after the last data was transmitted. Unfortunately, not all SPX applications perform in this manner. For example, Figure 6 shows a large number of minimum-sized (64 bytes for Ethernet) SPX packets on the wire. However, these packets were not used to exchange any data.

If you examine the packets more closely, they look like SPX connection termination packets. First, you see an End-of-Message packet that requests an acknowledgment. (See Figure 6.) The response is a simple ACK packet. These packets are transmitted back and forth on the network every two seconds.

The result is greater bandwidth utilization, an on-demand link that is kept open, or extra costs on an X.25 link that charges on a per packet basis. Nearly one-half of the networks we analyze currently have this problem. What causes it? An SPX application that keeps the SPX connection open full-time instead of breaking down the connection and reestablishing it as necessary.

One product that uses this type of communication is the Hewlett-Packard (HP) JetDirect card. In Queue Server mode (in which the JetDirect card emulates a Novell print server), the JetDirect card sends an SPX Keep Alive packet to the server. This packet is sent every two seconds, 24 hours a day, 365 days a year, allowing the JetDirect card to continually poll the queues for print jobs.

At this time, the only way to reduce the SPX traffic caused by JetDirect cards is to use HP's JetAdmin program (configuration > Advanced Settings) to change the polling time to ten seconds (the maximum setting). Although this does not solve the problem entirely, it reduces the SPX Keep Alive traffic significantly.

You should always keep an eye on SPX communications. We have seen several poorly written SPX applications appear over the last year. (For more information about SPX communications, read *Novell's Guide to NetWare LAN Analysis*, which is published by Novell Press.)

Beware of Routers in Parallel

The third problem we frequently see is clients communicating across network segments unnecessarily. We recently evaluated a network that was experiencing high utilization on a segment that supported few clients and was not the quickest path to any servers.

A closer look showed each client that communicated with the server on the local network sent its packets through a router and through another network. The network was configured with routers in parallel. (See Figure 7.) All communications between devices on network B and the server, SERV1, passed through the Cisco router onto network C. Yet the devices were on the same network segment as SERV1. The result was an unnecessary load on the Cisco router and network C.

Why were the stations communicating via the Cisco router and network C? The first step to determine the cause of the problem was to find out what happened when the clients on network B made their initial attachment to SERV1. To do this, you capture traffic from one of the stations and view the summary screen, as shown in Figure 8.

As we can see in packet 1, Fred sends out a Get Nearest Server (GNS) request to network B. SERV1 responds with a Get Nearest Server Reply. In that reply packet, SERV1 includes the internal IPX number for itself—in this case, 99. Next, in packet 3, Fred sends a Routing Information Protocol (RIP) packet looking for the route to use to get to network 99. This is the point at which things go awry.

The Cisco routers are faster than the NetWare servers and respond to the RIP request first, as shown in packet 4. As a result, the NetWare client station

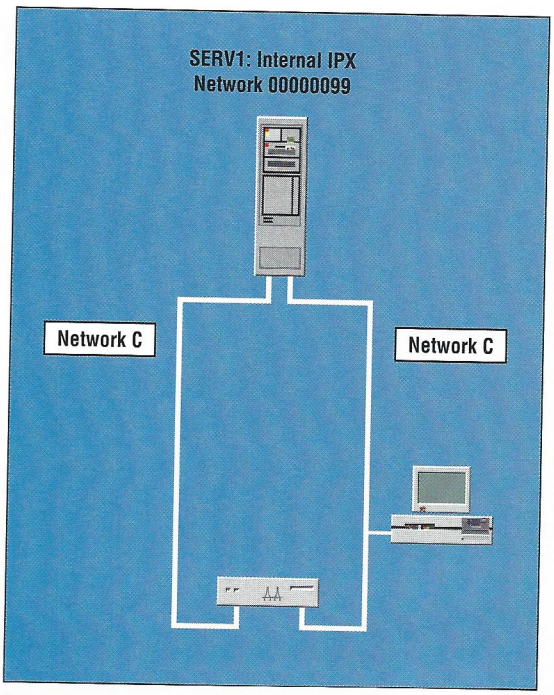

Figure 7. *The NetWare server SERV1 is configured in parallel with the Cisco router.*

FROM	TO	PACKET TYPE
Fred	Broadcast	SAP Request - GNS
SERV1	Fred	SAP Reply - Network 99
Fred	Broadcast	RIP Request - Network 99
CISCO RTR	Fred	RIP Reply
Fred	CISCO RTR	.... all remaining communications

Figure 8. *View the summary of the attachment process to determine how stations are communicating with each other.*

believes the Cisco router is the best way to get to network 99.

This problem occurs because of the way Cisco routers handle multiple, identical entries in their router tables. At our client site, the Cisco router received information about network 99 from two ports (the port connected to network C and the port connected to network B). If a router hears about a route from two or more ports and the routes have equal hops and ticks, the router assumes it is in parallel with another router or routers on the network. This is clearly the case on the network shown in Figure 8.

According to Novell specifications, if a router finds it is in parallel with another router, the router should not advertise the route. As you can see in Figure 8, the Cisco router is clearly in parallel with SERV1, which is acting as both a server and a router, but the Cisco router is still advertising its route. In our experience, we have seen network segments that have up to 10 percent unnecessary traffic due to this situation. We have discussed this situation with Cisco and found that there is a very simple solution.

Cisco routers include a configuration option called "Novell maximum-paths," which sets the maximum number of paths that the router will remember. By default, the router remembers a single entry if there are duplicate, *identical* paths. As a result, the Cisco router answers RIP requests even though it is in parallel with another router.

To fix this problem, you must change the path's setting to 2, and the router will no longer respond when it knows it is in parallel with another router. It will store both identical entries but will not "talk" about either, as outlined by the IPX router specification. After this setting is changed, the attachment process should be as follows:

- Fred's station broadcasts a GNS request.
- SERV1 responds with its internal IPX network number.
- Fred's station broadcasts a RIP for the internal IPX network number.
- SERV1 responds with a RIP reply, indicating it is the preferred route.

Now all traffic from Fred's station to SERV1 travels only on network B.

Conclusion

Analyzing a network takes time and research; you must carefully study the behavior of applications and hardware and software products. Always start by completing a characterization checklist to ensure you don't miss basic network communications problems such as utilization and excessive errors. The next step is to take a look at a network communications session to see how things are working. Finally, you must examine any communications that look suspicious, such as duplicate request and reply packets, excessive SPX traffic carrying no data, or a high number of packets between routers that are in parallel.

Above all, keep a record of your findings! Keeping a record should speed up the resolution of similar problems in the future and provide you with a valuable reference book for more advanced network troubleshooting and optimization tips.

Laura Chappell is a Product Manager for ImagiTech, Inc., a corporation that designs and develops interactive multimedia educational products. Laura is also the codeveloper of VirtuaLab, an online interactive multimedia network learning resource center for network technicians and administrators. Laura can be reached at CIS: 72000,3333 or via the internet at lchappell@imagitech.com. ■