NetSense™

for ETHERPEEK and TOKENPEEK

User Guide

Copyright © 1996-1998 Net3 Group, Inc. All rights reserved.

Chapter 1

What is NetSense?

NetSense is a powerful software application that performs peer-to-peer packet transaction analysis of network protocols on Ethernet or Token Ring trace files. Designed specifically as a 32-bit Windows program, NetSense provides a wealth of information from packets contained in trace files captured from EtherPeek and TokenPeek. "Nuggets" of information from each packet are extracted by NetSense and displayed in several different formats.

For an overall view of your trace file, the NetSense main screen has four general information notebook tabs, including a general file summary, a packet size distribution chart, and a protocol distribution chart and table. Also available on the NetSense main screen are two expert function tabs: the Client/Server Expert and the DLC Error Expert. For more information on the NetSense main screen, see Chapter 2.

A key feature of NetSense is the ProAnalyst Toolbox. IP and IPX packets are automatically sorted by protocol and application transaction. The Packet Visualizer component provides unique insight into protocol operation displaying packets in a spatial fashion and allowing you to easily see packet sizes and the delays between them. From the ProAnalyst Toolbox, you can perform Latency Analysis, Throughput Analysis, bring up a Problem Analysis summary, and perform "What If" Analysis. Chapter 3 explains the ProAnalyst Toolbox in detail.

More specific information is available in the Station Screen, which contains the Station Peer Map and Station Statistics Table. The peer map displays all Data Link Control (DLC) stations' peers and which protocols are being used above the DLC layer (AppleTalk, DECnet, IP, IPX, NetBEUI, etc.). The Station Statistics Table contains ten columns of information for every DLC address in the trace file. Any of the columns can be sorted, allowing you to generate "top ten" lists on network utilization, average transmitted packet size, or any other statistic.

A detailed look at network layer communication is found in the Network Screen, which also has a peer map and a statistics table. The Network Peer Map shows which nodes are communicating with each other and over which protocols–IP (UDP, TCP, ICMP, etc.) or IPX (NCP, SPX, NetBIOS, RIP, etc.).

The Network Statistics table displays important IP and IPX information. Like the Station Statistics Table, any of the columns of information can be sorted to generate "top ten" lists.

The Station Screen and Network Screen are discussed in more detail in Chapter 4.

Supported Trace File Formats

NetSense for EtherPeek and TokenPeek supports trace file analysis from packets captured by either of these programs. When NetSense is used in conjunction with ProConvert™ a file conversion utility also developed by Net3 Group, packet traces saved from many popular packet analyzers, including Sniffer®, LanDecoder™, LANalyzer™ and others can be changed to EtherPeek or TokenPeek formats and read by NetSense.

System Requirements

The following are the minimum recommended system requirements:

Processor: 80486DX / Pentium 90 or better

Display: 640 x 480 / 800 x 600, 256 colors or better

Operating System: Windows 95/98/NT4.0.

System Memory: 16/32MB for Win 95/98; 24/48MB for NT.

Hard Disk Space: Approximately 2 MB for program files.

Installing NetSense

To install NetSense, you can locate and run the setup program (Setup.exe*)*, or follow these Windows-recommended installation steps:

1) Start Windows.

2) From the Start menu, highlight "Settings," select "Control Panel," and then double click the "Add/Remove Programs" icon.

3) In the "Add/Remove Programs Properties" window, go to the "Install/Uninstall" tab and click the "Install" button.

4) If you are installing from a floppy disk, insert it now. Click the "Next" button.

5) If necessary, locate the setup file executable (Setup.exe).

6) Click the "Finish" button after Windows locates the setup program command line.

7) Follow the on-screen instructions.

Uninstalling NetSense

NetSense comes with an automatic uninstaller. This uninstaller removes all NetSense directories and Start menu shortcuts. The uninstaller is located in the NetSense program folder. You may also follow these Windows recommended uninstallation steps:

1) Start Windows.

2) Make sure NetSense is not running before continuing.

3) From the Start menu, highlight "Settings," select "Control Panel," and then double click the "Add/Remove Programs" icon.

4) In the "Add/Remove Programs Properties" window, go to the "Install/Uninstall" tab and highlight NetSense from the list of applications. Click the "Add/Remove..." button.

5) Follow the on-screen instructions.

Basic Operations

Starting NetSense

With Windows 95/98 or NT 4.0, NetSense can be started from the NetSense program group via the Start menu or you may create a shortcut on your desktop.

The main menu bar allows you to perform all major NetSense functions. From the toolbar, you can open a trace file, print, open the ProAnalyst Toolbox, open the Station Screen, and open the Network Screen. The status bar at the bottom of the NetSense Main Screen displays a small summary of the currently opened trace file.

The first two functions on the tool bar, to open a file and print, are explained in this chapter. The remaining options are covered in Chapters 2 through 4.

Opening a Trace File

Once you have NetSense running on your computer, you can open a trace file. To do this, click the folder icon on the button bar or select "File" | "Open..." from the menu bar and the Open Trace File dialog box appears as shown in Figure 1-1.

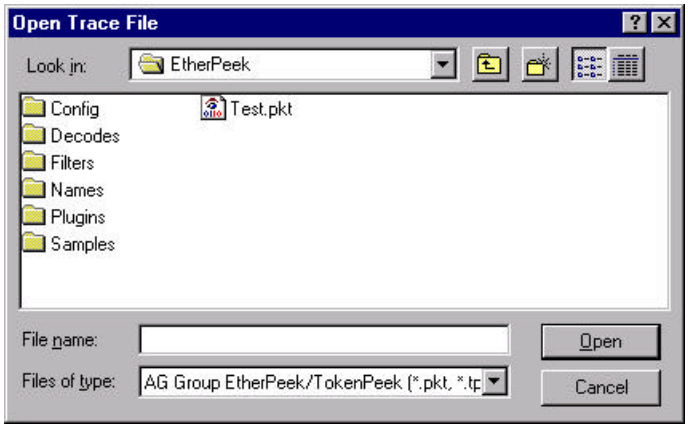


Figure 1-1: Open Trace File dialog

You can then select a file by clicking on it or entering the name into the "File name" box and pressing the "Open" button, or by simply double clicking on the file. NetSense also maintains a Most Recently Used (MRU) list of the last eight files you opened. This list is always available from the "File" | "Reopen" menu.

You can maintain and manage your trace files and directories (e.g. copy, delete, rename, etc.) from the Open Trace File dialog box. Simply right click on a file or directory to see a list of options.

IMPORTANT! Once a file is open, you should check the default bandwidth in the "Options" menu, so that a statistic such as network utilization is accurate. The default settings are 10 Mbps for Ethernet and 16 Mbps for Token Ring, so you'll need to change it if your trace is from a 100 Mbps Ethernet or a 4 Mbps Token Ring network.

Fly-over Hints

NetSense contains fly-over hints for most screen elements. By momentarily positioning the cursor over an icon or term, additional information will be provided. The hints are not just for icons, but for many technical terms as well.

Printing Screens

NetSense includes the capability to print screens by selecting the "File" | "Print..." item from the main menu, by clicking the printer button on the toolbar, or by using the keyboard shortcut combination Ctrl+P. Printing screens is a very useful feature, allowing you to keep records for baselining purposes, for example. Figure 1-2 shows the Print dialog box.

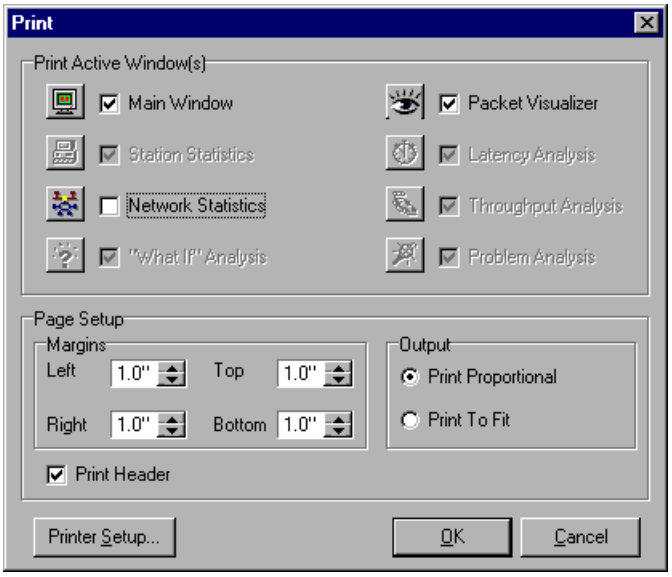


Figure 1-2: Print dialog

Make sure that you select a printer before attempting to print. You may specify the printer for output either through the "Printer Setup..." option in the "File" menu on the main screen or through the "Printer Setup..." button in the Print window.

In order to print a particular screen, that screen has to be open. However, it does not have to be the active window or even in view. To print one or more screens, open them up and then go to the Print window. From there, you can select the screens you would like to print by putting a check in the appropriate check box.

The Page Setup panel allows you to customize the output of your printing job. You can adjust the margins and proportions of the open screens. The "Print Proportional" radio button in the Output panel prints the screen as it appears on your monitor. If the Print to Fit option is selected, the screens will be enlarged or reduced to best fit the printer paper size. You can also select whether to have your printouts contain a trace file header.

Resizing Windows

Most of the windows within NetSense can be resized, minimized, and maximized like any standard Windows application. With peer maps, for instance, you may want to view them in a maximized window. Resolutions up to 1280 x 1024 are supported. When you exit NetSense, all the window sizes and positions on your display are saved

for the next time you run the program.

List Scrolling

Whenever a list is presented (such as when selecting a set of stations to view on the statistics table and peer map), clicking in that list (to make sure it has the “focus”) allows you to use the PgUp and PgDn keys in addition to the vertical scroll bar to scroll through the list.

Chapter 2

The NetSense Main Screen contains four general information tabs: File Summary, Packet Size Chart, Protocol Chart, and Protocol Table; and two expert function tabs: Client/Server Expert and DLC Error Expert.

File Summary Tab

The File Summary tab, Figure 2-1, opens a window that provides general information about the currently open trace file.

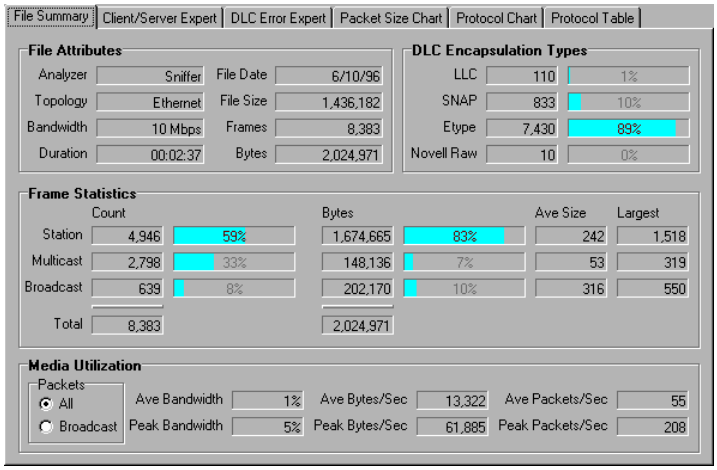


Figure 2-1: File Summary

This screen is self-explanatory except for the “Media Utilization” group box. Selecting the “All” option in the “Packets” group box calculates the “Ave Bandwidth,” “Ave Bytes/Sec,” etc., for all packets in the trace file. The “Broadcast” option means "perform these calculations for only the broadcast packets found in the trace file."

You can compare the "Broadcast" option statistics to the "All" option statistics to get an idea of the percentage of broadcast traffic on your network.

Packet Size Chart Tab

The Packet Size Chart tab displays a graph of the distribution of packet sizes or lengths in bytes (Figure 2-2).

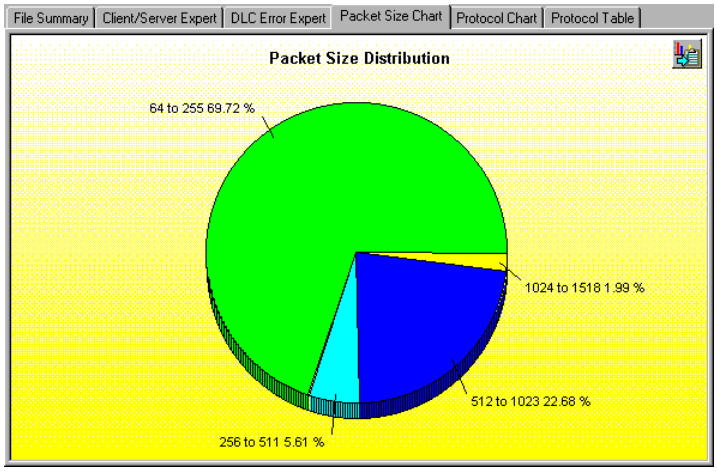


Figure 2-2: Packet Size Distribution Graph

Protocol Chart Tab

The Protocol Chart tab displays a graph of the distribution of protocols in your trace file (Figure 2-3).

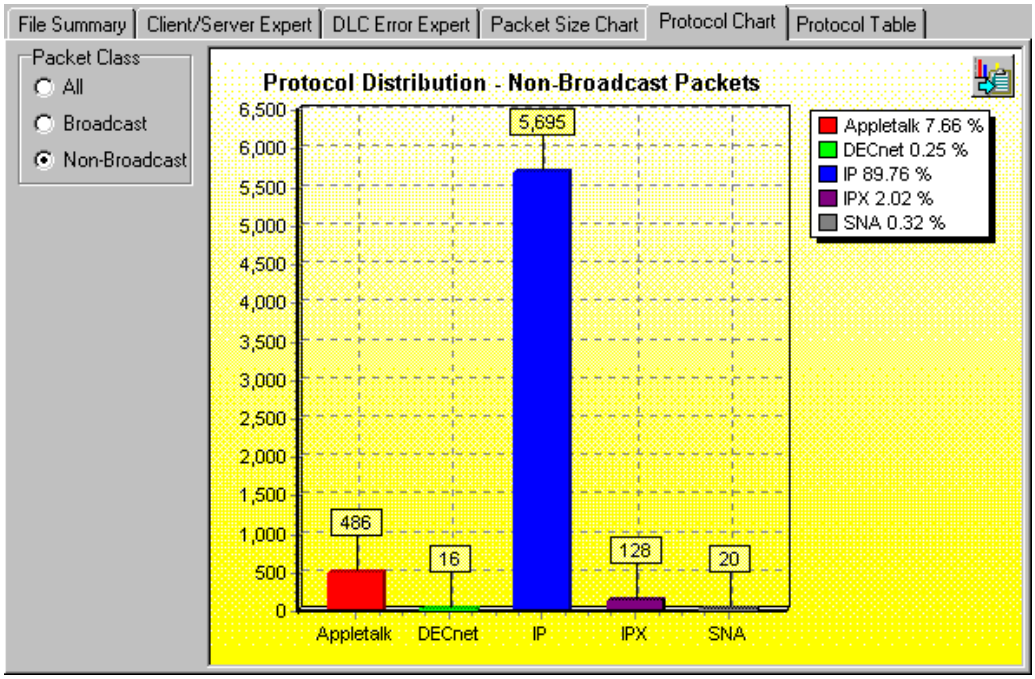


Figure 2-3: Protocol Distribution Chart

The “Packet Class” group box on the left-hand side of the window specifies what class of packets to graph.

You can select the “Broadcast” option to see what protocols are producing the majority of broadcast traffic and compare this to the “All” option or “Non-broadcast” option for broadcast to total packet count or broadcast to non-broadcast ratios.

Chart Properties

Both the Packet Size Chart and the Protocol Chart can be customized by right clicking on the chart and selecting "Chart Properties..." from the pop-up menu. You will then be presented with a window to modify chart properties such as: changing the display between a pie or bar chart, specifying the chart units, enabling and changing the background gradient, etc. Any changes made are reflected on both the Size Chart and the Protocol Chart

Charts may be copied to the clipboard by pressing the clipboard button in the upper right-hand corner of the chart, by selecting the "Edit" | "Copy Chart" item from the main menu, or by pressing the key combination Ctrl+C.

Pasting the charts into another application such as a word processor or spreadsheet can be useful in creating visually exciting network performance reports and baseline documents.

Protocol Table Tab

The Protocol Table tab provides an organized, statistical chart of the traffic in your trace file. For each protocol found in the trace file, statistics such as packet count, packet percent, average packet size, and DLC packet encapsulations used are generated and placed in the table.

Client/Server Expert Tab

The Client/Server Expert Tab (Figure 2-4) shows an entry for each client/server session established in the trace file.

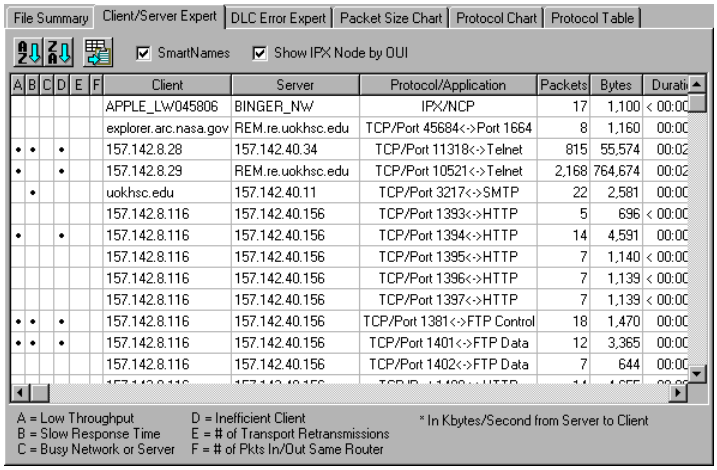


Figure 2-4: Client/Server Expert Session Screen

An example of a client/server session would be an ICMP Echo session. The Echo request originates from the client and the server sends back an Echo reply. Client/server transaction entries in the table may contain some expert diagnoses, such as “Low Throughput” or “Busy Network or Server” as indicated in columns A-F. The client/server problem key at the bottom of the screen explains columns A through F. Further details can be found later in this section.

General information such as the protocol and port, the number of bytes exchanged, the transaction duration, the response time, and throughput is calculated and reported.

Columns in the Client/Server Expert table can be sorted by double clicking or by

selecting a column and then pressing either the "sort ascending" or "descending" buttons in the upper left-hand corner of the screen. Selected cells within the table can be copied to the clipboard also.

Clicking on the ProAnalyst Toolbox button will take you right to the Packet Visualizer with the client/server from the highlighted row. Or, you can right click on a row.

The expert system bases its findings on thresholds and heuristics from human analysis of actual problems from hundreds of production networks. While some of your problems may fall outside the portion of the "bell curve" where most problems occur, the expert will alert you to possible deficiencies in many areas. Note that more details on any of the problems listed below can be found in the Packet Visualizer.

A = Low Throughput. Throughput is calculated from the server to the client. NetSense takes into account the Ethernet or Token Ring bandwidth of the trace and how much of that bandwidth is consumed while transferring data from the server back to the client If a very small percentage is being utilized (less than 0.01 percent), then the expert assumes that either the client is accessing the server infrequently or there is a more severe problem such as transport retransmissions that will be noted in another column. Thus, the throughput must be greater than 0.01 percent of the available local bandwidth but less than 10 percent to trigger low throughput. To see all the Client/Server connections with low throughput, perform a descending or an ascending sort on the throughput.

B = Slow Response Time. Also known as "round-trip response time" or "round-trip latency" from the client to server and back. The first thing to keep in mind is that response time is relative to where the trace was taken. The closer you are to the server, the lower the response time. Tracing from the server segment will provide the best average response time and a close measurement of the server efficiency, provided that there’s not a large amount of other traffic on that segment from other servers. Tracing from the client side will give you a good idea of the overall response time of the network+server. (For more details on this subject, please consult the section on Performing Latency Analysis in Chapter 2 of this manual.) The Expert requires at least 4 packet pairs with an average of more than 100 milliseconds to trigger slow response time.

C = Busy Network or Server. The Expert takes into account the minimum round-trip delay and the average round-trip delay to trigger a busy network or server. The more fluctuation there is in response time, the more likely there is contention for network or server resources. If this is flagged by the Expert, you should study the transactions in more detail in the Packet Visualizer or your protocol analyzer to see under what circumstances the variance (between minimum and round-trip delay) occurs, for example, during login authentication? File read/writes? You could also check other traces or SNMP monitoring tools for high utilization on segments between the client and server, such as a congested backbone or WAN link. Finally, you might check the

server itself for high CPU utilization and which processes are consuming the most processor time.

D = Inefficient Client. If the conditions for low throughput are met (see above) and the client is basically in a "ping pong" conversation with the server (i.e., the send-to-receive packet ratio or window is close to 1:1), then inefficient client is triggered. Look at the Packet Visualizer or your protocol analyzer for more details.

E = # of Transport Retransmissions. Transport retransmissions basically occur when the client does not receive a response to a request (like NCP) or receive an acknowledgment to sent data (TCP). Retransmissions can be caused by many factors including, but not limited to:

- packet(s) discarded on the way to the server or back due to CRC errors or buffering in a bridge/switch/router;
- packet(s) delayed on the way over or back due to bottlenecks along the way, especially with WAN circuits;
- packet(s) dropped by the server due to adapter congestion; or
- packet(s) held in server memory but delayed in processing.

A good technique for narrowing down this problem is to capture traces from different segments along the path to the server. To get a snapshot of the "top retransmitting clients," perform a descending sort on this column.

F = # of Pkts In/Out Same Router. In the case where you have multiple IP subnets defined on one router port or are using multiple IPX frame encapsulation types, you may encounter packets being "locally routed," i.e., a packet enters a router (or a NetWare server with internal IPX routing), then is sent back on the same segment. To see how serious this "packet doubling" is, take a "wide open" trace snapshot with your analyzer then analyze it with NetSense and perform a descending sort on this column. Technical solutions and trade-offs to this problem are beyond the scope of this manual, but consider using a flatter subnet for IP, go back to IP routing or layer 3 switching, or get down to one encapsulation type in the case of IPX.

Data Link Control (DLC) Error Expert Tab

NetSense can perform Data Link Control (DLC) expert analysis on both Ethernet and Token Ring traces.

Ethernet DLC Error Expert

For Ethernet, packets with Cyclic Redundancy Check (CRC) errors, Runt/Short packets, long packets, and fragments are totaled. For Ethernet promiscuous traces (all

packets), fragment analysis may also be performed. The Ethernet Error Expert window is shown in Figure 2-5.

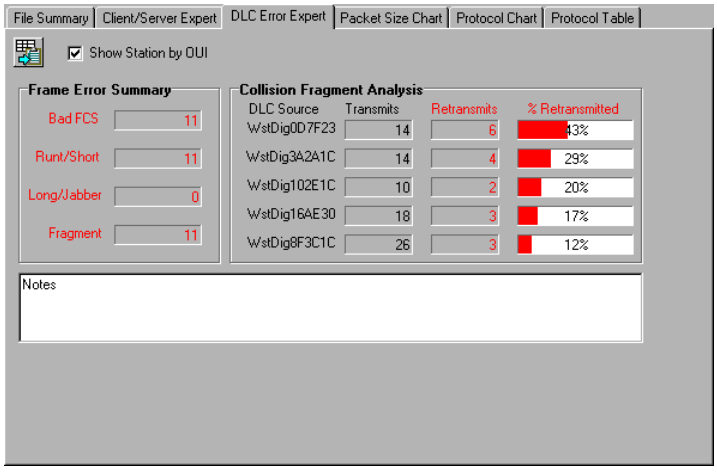


Figure 2-5: Ethernet Error Expert Window

Ethernet packets with FCS (Frame CheckSum) errors are those which your analyzer flagged as not passing the 32-bit FCS check when it received the packet. FCS errors may indicate a cabling, interference, transceiver, or repeater (such as a 10BaseT hub) problem.

Runt/Short packets are those that are shorter than the 64-byte minimum as required by both the Ethernet and the IEEE 802.3 standard. The standard includes the 32-bit or 4-byte CRC as part of the frame length, although some analyzers do not. Runt packets are often the result of collisions, and such packets will also contain a CRC error. A short packet without a CRC error is usually an illegally sized packet transmitted by a station.

A long packet, sometimes called a jabber, exceeds the standard maximum of 1518 bytes including the CRC. Long packets with CRC errors could be the result of a repeater or transmitter problem. If the packet does not have a CRC error, it may have been illegally sent by a station.

For trace files that contain Ethernet error packets, NetSense will also perform collision fragment analysis. NetSense counts a short frame with a bad CRC as a collision fragment.

In a "normal" Ethernet (one within the cabling and repeater specifications), most collisions will appear within the 64-bit preamble prior to sending the first field in the packet, the destination address. If your analyzer is collecting a lot of runt/short packets with CRC errors, your network is probably experiencing a lot of collisions,

and some of them are beyond the preamble, i.e., into the data portion of the packet. It is these collision fragments that your analyzer is able to capture.

A unique feature of NetSense is the ability to analyze these fragments to determine if they are caused by random transmissions over your Ethernet collision domain, or if a particular station is causing more collisions than others.

If a particular station has a high ratio of "bad" to "good" transmitted packets, then it is a suspect in not adhering properly to the Ethernet protocol. Perhaps its transceiver is bad or the Carrier Sense part of the CSMA (Carrier Sense, Multiple Access) mechanism is not functioning properly.

If a station is at the top of a "suspect list" as shown in Figure 2-4, then one should increase the traffic flow from that station to see if there is an increase in error counts on your analyzer (or cable tester).

Token Ring DLC Error Expert

For Token Ring, NetSense provides complete information about trace files containing soft error report Media Access Control (MAC) frames.

By capturing Token Ring MAC frames over a period of several minutes, you can perform a ring map diagnostic on your trace, as illustrated in Figure 2-6.

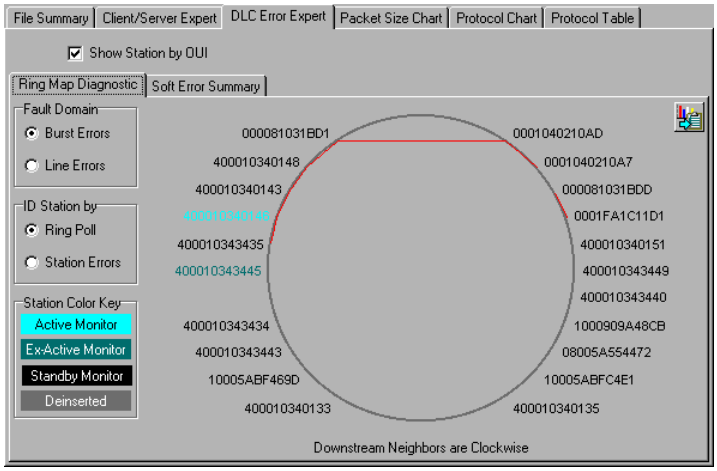


Figure 2-6: Token Ring Map Diagnostic

NetSense will draw the map based on ring poll frames, and also indicate any deviations in the station selected as the active monitor as well as stations that have uninserted or reinserted in the ring during the time the trace was taken.

Using lines between stations, NetSense will indicate the "hot spots" or fault domain where burst or line errors are occurring. Stations can be color coded by "Ring Poll" or by individual "Station Error" types.

NetSense groups the Token Ring soft error reports in a more logical fashion than the IBM "isolated" and "non-isolated" categories (a rule that many analyzers follow). As shown in Figure 2-6, NetSense breaks down each error report by the reporting station, that station's Nearest Active Upstream Neighbor (NAUN), also referred to as the Station's Upstream Address (SUA), and the total errors reported.

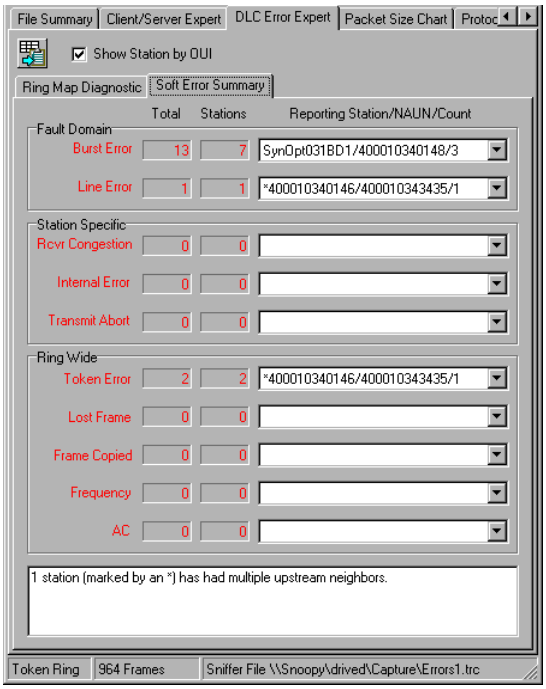


Figure 2-7: Error Report Summary

The "Fault Domain" error group shows errors that can be isolated using the station's NAUN. The error could be caused by the upstream station's transmitter, wiring, or hub port/; the receiving station's (the one reporting the error) hub port, wiring, or receiver; or any component in-between, such as a repeater or retiming module.

Specifically, burst errors are caused by a brief loss of signal such as when a station inserts or uninserts. A line error is the same as a CRC error, i.e., the packet fails the CRC as it is repeated by a station.

The "Station Specific" error group lists errors in which the NAUN is not applicable (even though it is included in the soft error report frame). These errors are specific to the reporting station.

Receiver congestion is when the station recognizes its address in the frame header (or a broadcast or functional address) but couldn't buffer the frame due to full memory on its adapter. Receiver congestion could be the result of a slow Token Ring adapter, slow adapter driver, slow computing platform, or a locked-up workstation that is still powered on (thus the adapter is still passing frames on the ring but the user can't do anything).

Internal errors suggest a marginal operating adapter or a conflict with another adapter in the system.

Transmit abort errors are when a station stops transmitting prior to the end of a frame.

The "Ring Wide" error group includes errors which could happen at any station on the ring. Token errors (lost or corrupted) and lost frame errors are only reported by the active monitor. A frame-copied error is when another station has already set the Address Recognized (AR) and Frame Copied (FC) bits, indicating a possible duplicate address on the ring. Frequency errors occur when the 4-Mbps or 16-Mbps clock derived from the received signal deviates from that station's internal clock. Finally, the AC error is when a station sees a standby monitor present frame but never saw the initial active monitor present frame at the beginning of a ring poll.

Chapter 3

The ProAnalyst Toolbox, the third button on the NetSense toolbar (illustrated in Figure 3-1) provides a number of powerful tools for analyzing latency/delay, application throughput, performing "What If" load and response time analysis, and diagnosing problems with packet transactions or exchanges.

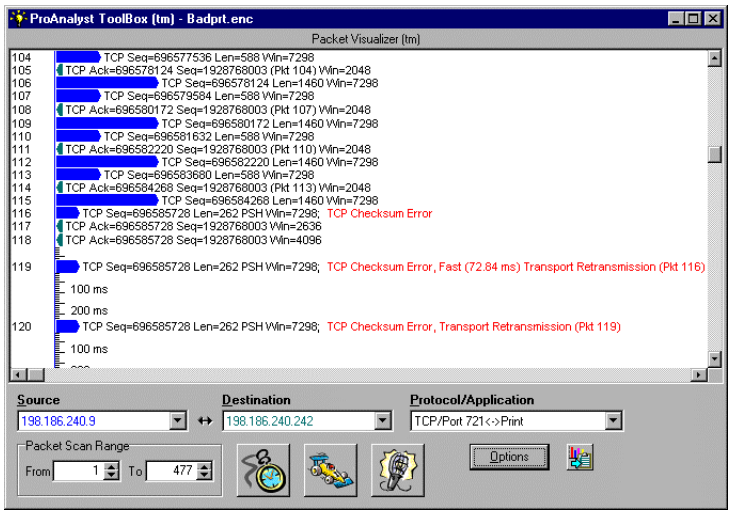


Figure 3-1: ProAnalyst Toolbox

The automatic filtering of protocols and applications (discussed below) makes it very easy to focus on studying performance and protocol problem areas between communicating nodes such as a client and a server. By tracking the latency and throughput figures over a period of time, baselines can be generated.

Automatic Conversation Filtering and Sorting

NetSense automatically breaks down a trace file by sorting peer-to-peer conversations (packet transactions) at the network layer and each protocol/application they are conversing with, down to the port level. For each source IP or IPX address, a choice of destination nodes communicating with that source is provided. The

Toolbox shows the protocols communicating directly over IP or IPX for each source/destination pair.

Two such examples for IP would be UDP and TCP. The applications using the protocol are identified by port number. An example is Telnet (Port 23) over TCP.

By selecting two IP addresses communicating with Telnet over IP for instance, you can easily perform latency (round-trip delay) analysis over frame relay.

Another useful feature is that multiple instances of an application are separated out in the application selection list. If a source has three separate FTP sessions, they will show up as FTP #1, FTP #2, and FTP #3.

Many of the common protocols over IPX including NetWare Core Protocol (NCP), SPX, and NetBIOS, are also separated for quick latency, throughput, and problem analysis.

Packet Visualizer Display

The Packet Visualizer portion of the ProAnalyst Toolbox screen provides a spatial view of a given set of protocol exchanges between two IP or IPX nodes. The relative packet sizes can be seen graphically for a quick visual glance of their payload. Packets are also spaced vertically depending on the delay between them. The time value for each tick mark is adjustable, with the ability to disable the packet spacing altogether. With packet spacing disabled, the packets are all shown back-to-back, similar to a conventional protocol analyzer.

The conversation shown in the Packet Visualizer is selectable via a pull-down list of source addresses. Once a source is selected, you may then choose the destination– NetSense will only list those addresses found in the trace file that are destinations communicating with the source, simplifying your selection. Finally, a protocol/application running over IP or IPX may be selected for the specific source/destination pair and values for the "Packet Scan Range" will be automatically set to the first and last packets of the current conversation found in the trace file. You can override this by editing the starting and ending packets in their respective boxes.

The Latency Analysis button (stop watch) performs instant latency (round-trip packet delay) analysis of the selected source/destination/application.

You can perform "What If" Analysis load and response time predictions, on the source and destination nodes from the Latency Analysis window.

The Throughput Analysis button (race car) calculates the minimum, maximum, and average number of bytes per second sent over a range of packets. The throughput contention factor is also computed.

The Problem Analysis button (trouble light) summarizes various errors along with

other protocol problems found inside the Packet Visualizer.

The "Options" button is used to set the delay tick mark scale between packets, set the relative packet lengths, turn the DLC or network layer summary on and off, and display the lower portion of an IPX address in OUI or hexadecimal format.

Performing Latency Analysis

Latency can be measured across all of the packets for a given source/destination pair and application. Often, it is desirable to limit the range to measure only a section of the packets (between user input delays, for example) for a more accurate calculation. Setting the packet scan range can be done in the ProAnalyst Toolbox window.

Clicking on the "Latency Analysis" button brings up a display similar to the one in Figure 3-2 below.

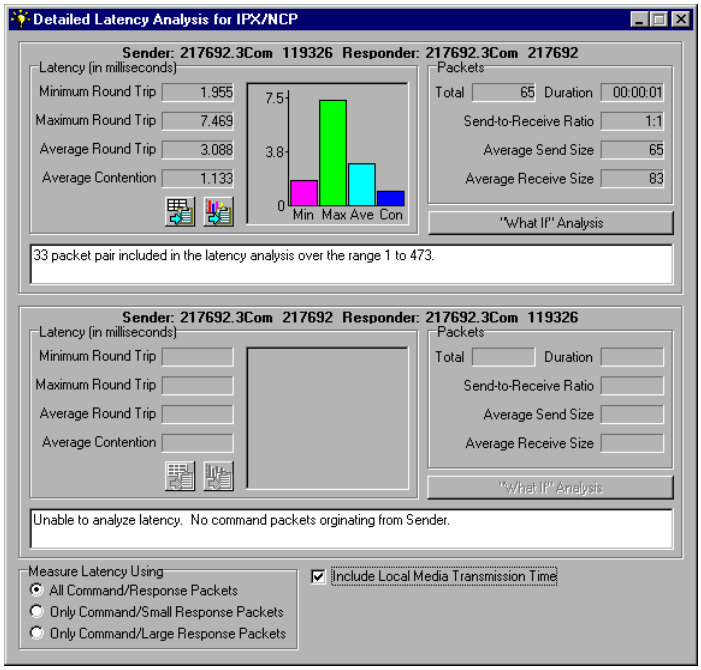


Figure 3-2: Latency Analysis Screen

Latency is measured by computing the delay between command and response packets. For protocols like TCP sliding window or NetWare packet burst, only the timing between the request and first reply packet is used. Received back-to-back packets are ignored. Capturing Telnet character echoes, Pings, or directory listings make excellent traces for latency analysis due to their "ping-ponging" of small

packets. This will give you a better understanding of the round-trip delay of a transmission path, especially if the trace has packets that traverse a WAN. Small packets will yield a close approximation of the best (minimum) latency of the circuit, while larger packets will increase the overall average because of their longer transmission time.

If your goal is to measure latency for client/server transactions over a WAN, you should capture packets on the source address of the client on the client's LAN segment. If your goal is to measure process time of a server for a set of packet transactions, capture packets locally on the server LAN segment with a filter to and from that server.

For precise measurements of packets confined to a local LAN segment, you may want to disable the "Include Local Media Factor in Calculation." On a 10 Mbps Ethernet for example, a maximum sized (1518 bytes including CRC) packet takes approximately 1.2 milliseconds to transmit. By removing this amount of time from a series of packets to and from a server, you get a more precise timing of the actual process time of the server (or workstation for that matter). Removing this factor for a measurement over a WAN will not have a significant impact on the timing.

The contention factor shown in the Latency Analysis window is the difference between the minimum or best case and maximum or worst case latency. When networks get busy, there is contention for bandwidth, router/bridge/switch resources, etc. By computing latency over a large sampling of packets, you get a pretty good idea of the best case response time based on the minimum latency value.

When there is little difference between the minimum and maximum latency values for an extended period of time, either your network load is very steady with virtually no instantaneous load fluctuations, or you have most of the network resources to yourself, such as during off-peak hours.

Another interesting statistic provided is the send-to-receive packet ratio. This gives you an idea of whether or not the application is taking advantage of the TCP sliding window or IPX packet burst protocol. If not, or if the average packet size is small, the packet ratio will be 1:1. This is because an application must work with data units larger than the TCP Maximum Transmission Unit (MTU) or the IPX MTU for NetWare, to take advantage of the sliding window or burst protocol.

If you would like to include the graph of the computed latency values in a report, email, etc., simply click on the copy graph to clipboard button. Then go to your word processor or other Windows application and perform a paste operation. If you need the raw data, click on the copy data to clipboard button.

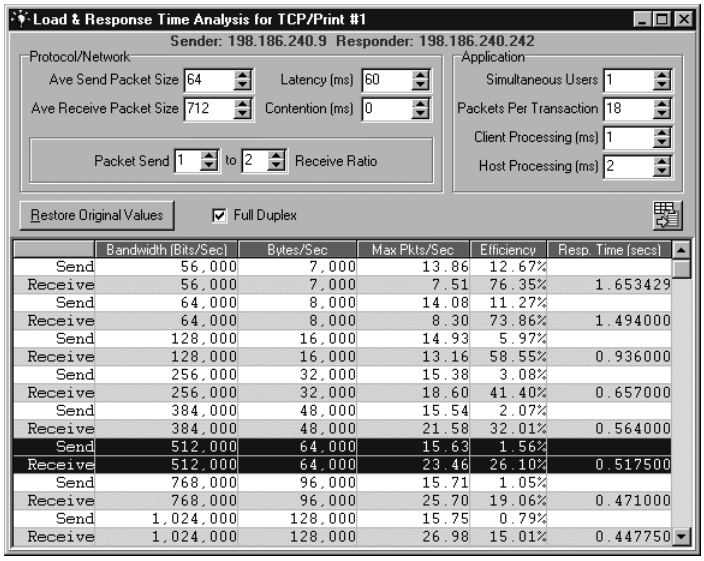
There can be many reasons for varying packet delay including high latency WAN networks (such as frame relay), contention for network bandwidth on any of the segments between the source and destination, or contention for resources in a host or server. NetSense computes these fluctuations in latency by showing the minimum,

maximum, and average latency, as well as the contention factor over a sampling of packet pairs.

Performing "What If" Analysis

Clicking on the "What If" button in the Latency Analysis window activates the "What If" window (Figure 3-3) where you can perform load and response time analysis under a variety of scenarios: "What If" my local Ethernet application ran over a WAN circuit? "What If" the application could be tuned to consolidate more data into larger packets? "What If" five users were running this application simultaneously? "What If" the host processing time is increased? "What If" the protocol can be tuned to improve the send-to-receive packet ratio? "What If" we use a dedicated leased-line instead of Frame Relay to reduce the circuit latency?

NetSense will take the actual Latency Analysis data as a starting point for predicting load and response time over a wide variety of WAN and LAN transmission speeds. As illustrated in Figure 3-3, there are a number of variables that can be changed to predict the response time of an application.



	Bandwidth (Bits/Sec)	Bytes/Sec	Max Pkts/Sec	Efficiency	Resp. Time (secs)
Send	56,000	7,000	13.86	12.67%	
Receive	56,000	7,000	7.51	76.35%	1.653429
Send	64,000	8,000	14.08	11.27%	
Receive	64,000	8,000	8.30	73.86%	1.494000
Send	128,000	16,000	14.93	5.97%	
Receive	128,000	16,000	13.16	58.55%	0.936000
Send	256,000	32,000	15.38	3.08%	
Receive	256,000	32,000	18.60	41.40%	0.657000
Send	384,000	48,000	15.54	2.07%	
Receive	384,000	48,000	21.58	32.01%	0.564000
Send	512,000	64,000	15.63	1.56%	
Receive	512,000	64,000	23.46	26.10%	0.517500
Send	768,000	96,000	15.71	1.05%	
Receive	768,000	96,000	25.70	19.06%	0.471000
Send	1,024,000	128,000	15.75	0.79%	
Receive	1,024,000	128,000	26.98	15.01%	0.447750

Figure 3-3: Response Time Variables for Latency Analysis

Figure 3-3 shows an application that ran locally on Ethernet, set the packets per transaction to 18, adjusted the client and host processing time, and adjusted the latency to approximate running the application over a frame relay connection (we set the contention to zero, as if no other traffic is contending for the link). Note that the

predicted response time for a 512 Kbps link is 517 milliseconds or about 1/2 second, with the application underutilizing the bandwidth. Doubling the bandwidth buys us very little in response time. Only by decreasing the latency or improving the send-to-receive packet ratio can we improve the response time. The same application running on a local half-duplex Ethernet segment with only 2 milliseconds of latency and contention, had a total response time of about 37 milliseconds (since most WAN circuits are full duplex, full duplex is the default setting; if your Ethernet is full duplex, you can use the default).

The table will automatically change when:

- Pressing the Enter key after typing in a new value
- Increasing or decreasing the value by 1
- Pressing the tab key to move to the next edit box

Performing Throughput Analysis

Throughput Analysis (Figure 3-4) computes the throughput in kilobytes per second, in both directions.

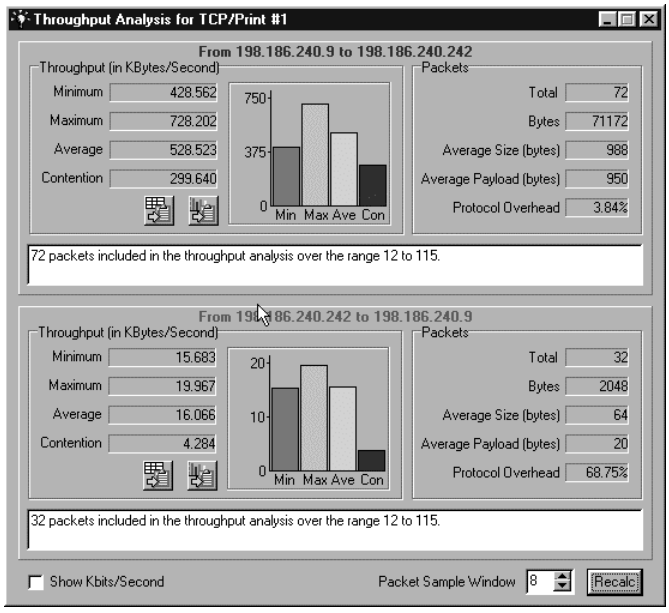


Figure 3-4: Throughput Analysis

The "Packet Sampling Window" edit box on the Throughput Analysis screen determines the minimum and maximum throughput calculation. If this value is set to eight, for example, then the throughput will be calculated in groups of eight packets in succession. NetSense actually performs thousands of calculations, since the first eight packets in the analysis range are used, then the second through ninth packets, third through tenth packets, etc. This gives the best case throughput of any succession of packets.

The larger the throughput sampling value, the closer the minimum and maximum throughput will be to the average over the entire range, thus smoothing out short-term fluctuations. Shorter sampling is useful to see what the maximum throughput could be in a tuned, or lightly loaded, network.

If there is not much contention for bandwidth, servers, routers, switches, and other resources, then the minimum, maximum, and average values will be close no matter what the sampling value is set to.

As with the latency analysis, the throughput graph may be copied to the clipboard by clicking on the "copy graphic to clipboard" button. Likewise, the data values may be copied to the clipboard by clicking on the "copy data to clipboard" button.

The protocol overhead compares the size of the header to the size of the packet payload of the network layer datagram (i.e., IP or IPX). The overhead is defined as the first byte in the packet through the end of the IP or IPX header.

Problem Analysis

You may see a number of packet errors and protocol problems displayed in the Packet Visualizer. Clicking the "Problem Analysis" button at the bottom of the ProAnalyst Toolbox provides a summary of these errors and problems. Additional information may be present, such as noting the number of packet retransmissions due to CRC errors.

Chapter 4

Station Screen

By selecting the fourth button on the NetSense toolbar, you open up the Station Screen window where specific DLC layer information is available. The Station Screen contains both the Station Peer Map, which graphically represents DLC communication, and Station Statistics Table, which calculates and reports information and statistics for each DLC source address.

Station Peer Map

The Station Peer Map (Figure 4-1) displays all DLC stations' peers and which protocols are being used above the DLC layer (AppleTalk, DECnet, IP, IPX, NetBEUI, etc.).

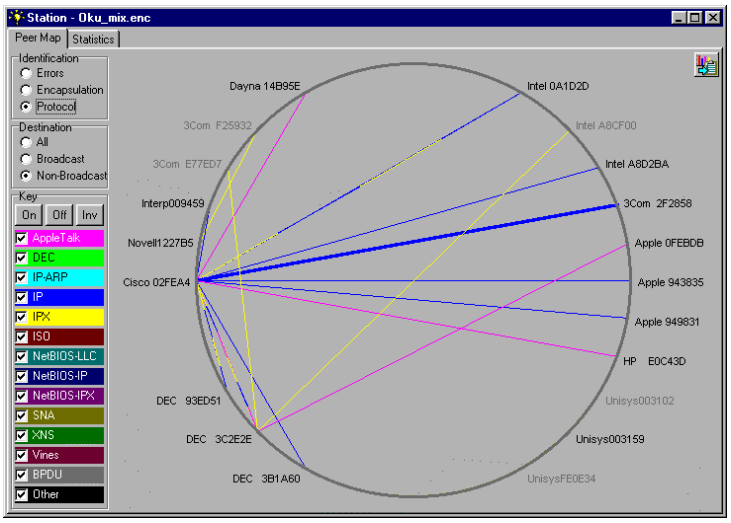


Figure 4-1: Station Peer Map

You may wish to display the peer map by “Errors” (Ethernet CRC errors only), by DLC “Encapsulation” type, or by “Protocol” by selecting the appropriate radio button in the “Identification” group box.

The “Destination” group box selects which type of destination addresses are displayed on the Station Peer Map: “All”, “Broadcast”, or “Non-broadcast”. The

“All” option displays both broadcast and non-broadcast destination addresses, the “Broadcast” option displays only broadcast traffic, and the “Non-broadcast” option displays only point-to-point traffic.

You may select which protocols are displayed on the peer map by checking or de-checking the corresponding check boxes in the “Key” group box. The “On”, “Off”, and “Inv” buttons turn on all protocols, turn off all protocols, and invert the selection state of all protocols, respectively.

Any DLC address appearing in a trace file as a destination address only appears on the map in grayed-out text. This makes it easy to spot broadcast addresses as well as flooded packets in a bridged/switched environment.

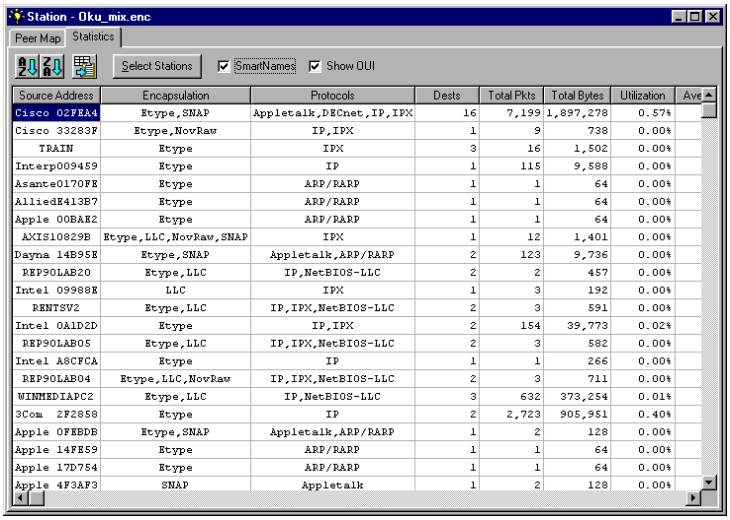
Right-clicking on a source station’s DLC address brings up a pop-up menu allowing you to easily switch to that station's entry in the Station Statistics Table.

The Station Peer Map displays lines of communication between stations in different thicknesses. The thickness of a line corresponds to the amount of communication between the two stations, i.e., thicker lines represent more traffic flowing between stations than would be indicated by thinner lines.

The Station Peer Map graphic can be copied to the clipboard by pressing the clipboard button in the upper right-hand corner of the screen.

Station Statistics Table

The Station Statistics Table in Figure 4-2 shows ten columns of data for all source DLC addresses in the trace file.



Source Address	Encapsulation	Protocols	Dests	Total Pkts	Total Bytes	Utilization	Ave
Cisco 02FEA4	Etype,SNAP	Appletalk,DECnet,IP,IPX	16	7,199	1,897,278	0.57%	
Cisco 33283F	Etype,NovRaw	IP,IPX	1	9	738	0.00%	
TRAIN	Etype	IPX	3	16	1,502	0.00%	
Interp009459	Etype	IP	1	115	9,588	0.00%	
Asante0170FE	Etype	ARP/RARP	1	1	64	0.00%	
AlliedE413B7	Etype	ARP/RARP	1	1	64	0.00%	
Apple 00BAE2	Etype	ARP/RARP	1	1	64	0.00%	
AXIS10829B	Etype,LLC,NovRaw,SNAP	IPX	1	12	1,401	0.00%	
Dayna 14B95E	Etype,SNAP	Appletalk,ARP/RARP	2	123	9,736	0.00%	
REP90LAB20	Etype,LLC	IP,NetBIOS-LLC	2	2	457	0.00%	
Intel 09988E	LLC	IPX	1	3	192	0.00%	
RENTSV2	Etype,LLC	IP,IPX,NetBIOS-LLC	2	3	591	0.00%	
Intel 0A1D2D	Etype	IP,IPX	2	154	39,773	0.02%	
REP90LAB05	Etype,LLC	IP,IPX,NetBIOS-LLC	2	3	582	0.00%	
Intel A8CFCA	Etype	IP	1	1	266	0.00%	
REP90LAB04	Etype,LLC,NovRaw	IP,IPX,NetBIOS-LLC	2	3	711	0.00%	
WINMEDIAPC2	Etype,LLC	IP,NetBIOS-LLC	3	632	373,254	0.01%	
3Com 2F2858	Etype	IP	2	2,723	905,951	0.40%	
Apple 0FEBDB	Etype,SNAP	Appletalk,ARP/RARP	1	2	128	0.00%	
Apple 14FE59	Etype	ARP/RARP	1	1	64	0.00%	
Apple 17D754	Etype	ARP/RARP	1	1	64	0.00%	
Apple 4F3AF3	SNAP	Appletalk	1	2	128	0.00%	

Figure 4-2: Station Statistics Table

Any column in the Station Statistics Table can be sorted either by double-clicking on the column or selecting the column and pressing the sort ascending or sort descending buttons in the upper left-hand corner of the screen.

Selected cells in the Station Statistics Table can be copied to the clipboard by pressing Ctrl+C or by pressing the clipboard button in the upper left-hand corner of the window.

If your trace file contains a lot of DLC addresses, you may want to limit the stations displayed in the Station Peer Map and Station Statistics Table, by going to the Station Statistics Table and selecting the "Select Stations" button. You will be presented with a list from which to select the desired stations by address.

Network Screen

Specific network layer information over IP and IPX is available in the Network Screen window by choosing the fifth button on the NetSense toolbar. The Network Screen contains both the Network Peer Map, which graphically represents network layer communication, and Network Statistics Table, which calculates and reports information and statistics for each network source address.

Network Peer Map

The Network Peer Map (Figure 4-3) shows which network addresses are communicating with each other and which protocols over IP (UDP, TCP, ICMP, etc.) or IPX (NCP, SPX, NetBIOS, RIP, etc.) are in use.

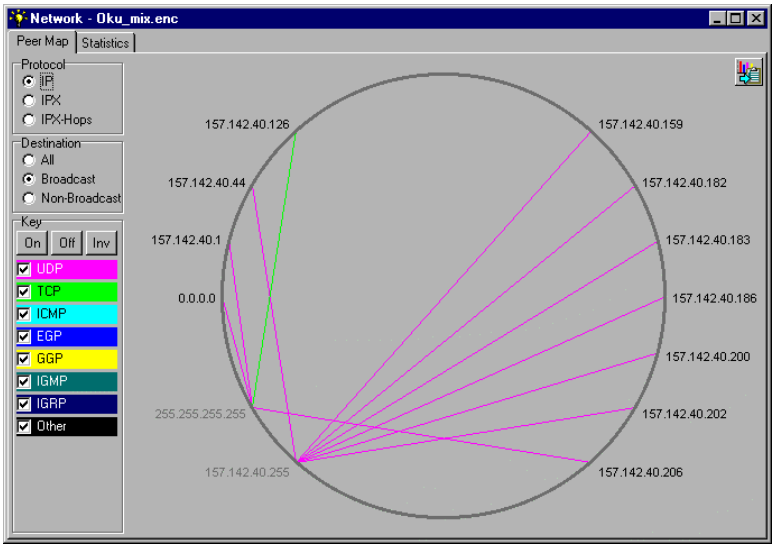


Figure 4-3: Network Peer Map

The “Protocol” group box specifies what type of graph you would like to see in the peer map. The “IP” option shows stations communicating to each other using the TCP/IP protocol stack. Both the “IPX” and “IPX-hops” options graph nodes running the NetWare protocol stack. The “IPX” option shows what types of protocols are being used on top of IPX while the “IPX-hops” option illustrates the distance between nodes by hop count.

The “Destination” group box selects which type of destination addresses are displayed on the Station Peer Map: “All”, “Broadcast”, or “Non-broadcast”. The “All” option displays both broadcast and non-broadcast traffic, the “Broadcast” option displays only broadcast traffic, and the “Non-broadcast” option displays only point-to-point traffic.

You may select which protocols are displayed on the peer map by checking or de-checking the corresponding check boxes in the “Key” group box. The “On”, “Off”, and “Inv” buttons turn on all protocols, turn off all protocols, and invert the selection state of all protocols, respectively.

IMPORTANT! Any network address appearing in a trace file as a destination address only appears grayed out on the peer map. This makes it easy to spot broadcast addresses as well as flooded packets in a bridge/switched environment.

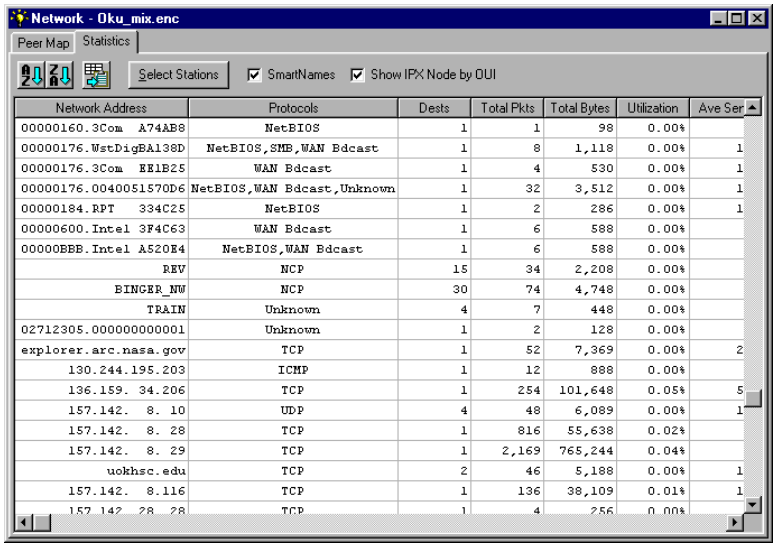
By right clicking on an address in the Network Peer Map, you can see its DLC

address as well as choose to go directly to the Network Statistics Table or the ProAnalyst Toolbox with that IP or IPX address.

The Network Peer Map can be copied to the clipboard in the same manner as the Station Peer Map–by pressing the clipboard button in the upper right-hand corner of the window.

Network Statistics Table

The Network Statistics Table (Figure 4-4) displays information from nine categories for IP and IPX.



Network Address	Protocols	Dests	Total Pkts	Total Bytes	Utilization	Ave Ser
00000160.3Com A74AB8	NetBIOS	1	1	98	0.00%	
00000176.WstDigBA138D	NetBIOS,SMB,WAN Bdcast	1	8	1,118	0.00%	1
00000176.3Com EE1B25	WAN Bdcast	1	4	530	0.00%	1
00000176.0040051570D6	NetBIOS,WAN Bdcast,Unknown	1	32	3,512	0.00%	1
00000184.RPT 334C25	NetBIOS	1	2	286	0.00%	1
00000600.Intel 3F4C63	WAN Bdcast	1	6	588	0.00%	
00000BBB.Intel A520E4	NetBIOS,WAN Bdcast	1	6	588	0.00%	
REV	NCP	15	34	2,208	0.00%	
BINGER_NW	NCP	30	74	4,748	0.00%	
TRAIN	Unknown	4	7	448	0.00%	
02712305.000000000001	Unknown	1	2	128	0.00%	
explorer.arc.nasa.gov	TCP	1	52	7,369	0.00%	2
130.244.195.203	ICMP	1	12	888	0.00%	
136.159. 34.206	TCP	1	254	101,648	0.05%	5
157.142. 8. 10	UDP	4	48	6,089	0.00%	1
157.142. 8. 28	TCP	1	816	55,638	0.02%	
157.142. 8. 29	TCP	1	2,169	765,244	0.04%	
uokhsc.edu	TCP	2	46	5,188	0.00%	1
157.142. 8.116	TCP	1	136	38,109	0.01%	1
157.142. 28. 28	TCP	1	4	256	0.00%	

Figure 4-4: Network Statistics Table

If your trace file contains many network addresses, you may want to limit the stations displayed in the statistics table as well as the Network Peer Map by pressing the "Select Stations" button. You will be presented with a list where you can choose to show or hide the desired addresses.

Similar to the Station Statistics Table, columns in the Network Statistics Table can be sorted and selected cells can be copied to the clipboard.

NetSense User Guide

Copyright © 1996-1998 Net3 Group, Inc. All rights reserved.

Copyrights & Trademarks

Net3 Group, NetSense, NetSense ProAnalyst, ProAnalyst Toolbox, and Packet Visualizer are trademarks of Net3 Group, Inc. The Net3 Group symbol, the NetSense light bulb, and the Packet Visualizer screen layout are Copyright © 1996-1998, Net3 Group, Inc. EtherPeek and TokenPeek are trademarks of AG Group, Inc. Network Associates and Sniffer are registered trademarks of Network Associates Corporation and/or its wholly owned subsidiaries. Windows is a registered trademark of Microsoft Corporation. All other company and product names used in this manual are trademarks or registered trademarks of their respective holders.

Technical Support

If you have any comments, questions, or problems to report, or if you encounter any technical difficulties during installation or while using NetSense, you may contact our technical support team. Please have your NetSense License Number ready.

Business hours: Monday - Friday, 9 am - 5 pm (PST)

Internet: http://www.aggroup.com
Email: support@aggroup.com
Phone: 925.937.7900 x27
Fax: 925.937.2479

AG Group, Inc.
2540 Camino Diablo, Suite 200
Walnut Creek, CA 94596