Agilent Technologies

Ethernet/Fast Ethernet Problems and Solutions

Application Note 1336

Finding and Solving 10/100 Ethernet Problems with the Advisor

Agilent Technologies

Contents

Introduction

Today’s businesses rely heavily on data networks to distribute information and resources to users both within and without their organizations. Managers of multivendor, multiprotocol enterprise networks must meet their organization’s increasing demands for access and exchange of information. New nodes are added, users relocate, network resources are redistributed, the numbers of applications increase, and access to the Internet and the Web remains critical. In addition, more subnets and networks are interconnected causing networks to grow in complexity. It is not enough to simply troubleshoot problems. Network managers must also find ways to fine-tune the productivity and efficiency of their networks to ensure that users are never without valuable network resources.

Importance of Enterprise Network Testing

The critical and important job of troubleshooting and maintaining an enterprise network is highlighted in the following points:

- **Network Downtime Costs** - The reliability of your network has a large impact on the vitality and profitability of your organization. For example, airline reservation systems, financial institutions, and mail order businesses measure the cost of network downtime in thousands of dollars per minute. While the cost of downtime may be difficult to quantify for your organization, you know that prompt fault isolation and repair are key ingredients in restoring service as soon as possible. Communication test tools that can be connected to any part of the network, capture all the information, and perform comprehensive analysis are essential to this process.

- **Network Inefficiency Costs** – Similar to network downtime is network inefficiency. Long file transfer times, slow responses for server-based applications or the Web, and unreliable or slow email transmissions, are just a few examples of how inefficient networks can reduce the productivity of entire segments of an organization. Test equipment that help fine-tune network operations can prove invaluable.

- **Accelerating Pace of Technological Change** - This alone makes it difficult to stay in stride with the increasing variety of new and improved hardware and software products available for all phases of network growth. The emergence of Voice and Fax over IP technologies is only one example. Not only must you be equipped to meet the challenges associated with installing, configuring, and maintaining new network equipment and software, you must also be prepared to deal with problems when they arise. You must also use your testing tools in a proactive way to detect problems before they affect productivity.

Advisor – The Best Tool for the Job

Every organization strives to keep costs down. That is why investing in an intelligent LAN troubleshooting tool like the Agilent Technologies Advisor is one of the wisest decisions you can make. The Advisor can reduce the amount of time test personnel spend on the often repetitive and time-consuming tasks associated with troubleshooting, and allow test personnel to concentrate on the solution and its implementation. And when test equipment can be used by personnel with a wide range of skills (i.e. novice to expert), the resulting productivity gain more than justifies the cost of the tool.

For example, when a network fails, you must first gather information. While you can often gather an abundance of data on network performance, network configuration, and failure mode symptoms, only a small part of this data is actually relevant to the specific problem. Sifting through this large amount of data and turning it into useful information is made easier with intelligent communication test tools that help you eliminate the irrelevant and pinpoint the problem. The Advisor LAN, an intelligent tool designed to help you install, maintain, and troubleshoot data networks, can provide you with the information to formulate the best solution to the problem at hand.

How to Use This Application Note

This application note describes how the Agilent Technologies Advisor LAN can help you be more efficient at solving problem on your network. The Note adheres to a standard troubleshooting process, and demonstrates those features and capabilities of the Advisor that are useful in LAN testing situations. The Application Note is designed to be used along with the Advisor's own online Help system. That is, this Note presents high level test techniques for real-world situations, whereas the online Help system gives detailed procedural and user interface reference information. Used together, you will quickly learn how to apply this powerful test tool to many LAN troubleshooting and performance enhancement tasks.

Overview of LAN Troubleshooting

LAN troubleshooting can be simple or complicated, depending on the nature of the problem and the approach used to isolate it. While there is not a guaranteed recipe used to solve all problems in an Ethernet network, there are some general approaches that can make this process easier:

- **Gather Network Topology Information**
  It is essential to understand the topology of the network under test. Knowing the network components, connections, transmission media, and relationships between the different sections of the network will help you predict problems and interpret diagnostic information your Advisor LAN provides. For example, the traffic found in a LAN backbone is different than the traffic found in a printer segment, and thus, can be interpreted and evaluated in different ways.

- **Test the OSI Stack**
  One very key part of network troubleshooting is to test according to the OSI reference model. You have two choices: (1) start from the top layer of the OSI model (application layer) and test layer by layer until you get to the physical layer or layer one; or (2) start instead from the physical layer. It is crucial to follow a layer by layer approach because starting in the middle of the stack can give misleading results. Generally, a good recommendation is to follow the bottom-up strategy because most network problems occur at the physical level.

- **Proactive Testing or Baselining**
  Proactive network testing, sometimes referred to as baselining, is a very important task for efficient network managers. It is vital to have some idea of normal network operation documented so that you will have something with which to compare when trouble comes. The *Proactive Network Troubleshooting* section provides a network baselining strategy for a LAN environment.

- **Use the Advisor's Expert Analysis**
  Finally, the use of test tools that provide expert analysis can give you an advantage when troubleshooting a network. See the *Using the Advisor's Expert Analyzer* to learn more about this powerful analysis option.

Connecting, Starting, and Configuring the Agilent Advisor LAN

The first step in using a portable protocol analyzer such as the Advisor LAN is to physically connect it to the network you intend to test and to configure it correctly for that connection. The Advisor can be connected as a node on the network or it can be connected such that it monitors traffic between nodes. The situations in which one connection is preferred over the other are described in the following sections. Use the Advisor's online Help for additional information related to these subjects. **Note:** there is a distinction between configuring the Advisor itself to match the physical environment in which it is used, and configuring the individual measurements that you intend to run. This section is concerned with the former – please refer to the online Help or later sections of this Note for information on configuring the individual measurements.

Hubs vs. Switches

Before we discuss connecting the, it is a good idea to understand the differences between hub and switched environments. The connection and configuration you use is influenced not only by the kinds of tests you intend to run, but also the physical environment.

In general, the function of a hub is to transmit any packet that arrives at one port of the hub to all the other ports of the hub. This means that all nodes connected to a hub are able to listen to all other nodes. Switches, on the other hand, can "learn" the position of nodes in the network by mapping the physical addresses of the nodes localized in each segment of the network and then forwarding or filtering the packets depending on the destination address. When a packet reaches a switch, the switch compares the physical source and destination addresses of the conversation and isolates this conversation from the rest of the ports of the switch. Traditionally, hubs have been used in Ethernet networks because they appeared first in the market. However, switches have been taking the hub market share and this trend seems to continue.

Switch management and monitoring is sometimes an issue. One of the biggest challenges when testing in a switched network is the dynamic change in traffic patterns; that is, the switch will open and close ports depending on the traffic. When connecting the Advisor to the network, you must take this into account. For example, if you connect and configure the analyzer as a node in a switched environment, you will not see all the traffic you might otherwise hope to see. Since no other network element will 'know' about the existence of the analyzer, no traffic will be sent to the analyzer specifically and the switch will block the physical port to which the analyzer is connected. The only traffic that the analyzer will capture is broadcast traffic.

General Configuration Guideline:
Before connecting the Advisor, determine whether you are connecting in a switched environment or a hub environment. In most cases:

- When connecting in a switched environment, you will want to connect and configure the Advisor in a monitor-through mode so that all the traffic between a specific switch and a server/workstation is seen by the Advisor.

- When connecting in a hub environment, you will want to connect and configure the Advisor as a node so that the Advisor sees all the traffic destined for all the ports on the hub.

Configuring the Advisor to match the connection method is covered in the online Help and a later section of this Application Note.

Node Connections

The node connection shown in Figure 1 is used primarily in a hub environment. This connection, sometimes referred to as a point-to-point connection causes the Advisor to act as, and be seen as, a node or independent point on the network. The Advisor will see all the traffic that passes through the hub in the same way that any other Ethernet node would.

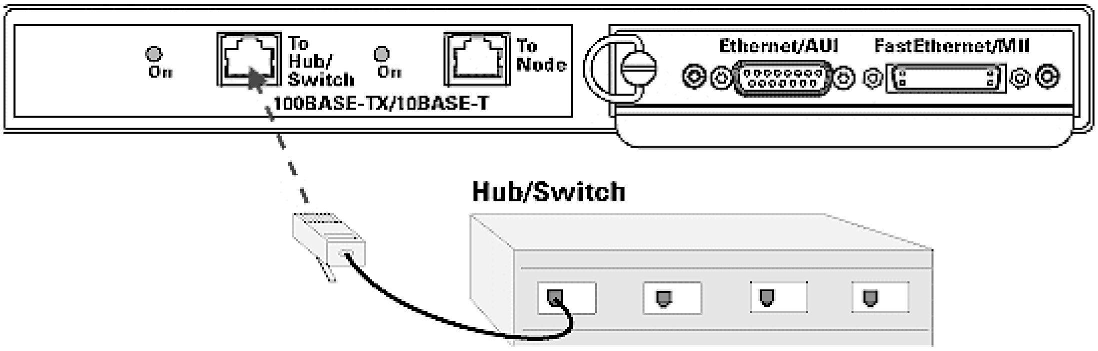

Figure 1: Node Connection

In the node or point-to-point connection, the Advisor is attached directly to an available hub port using a RJ-45 100Base-TX cable with a maximum length of 100 meters. The Advisor can monitor traffic from all the stations having the same collision domain as the hub port where the Advisor is connected. In this mode, the Advisor can also generate traffic onto the network.

Using this connection, the Advisor can operate in a half-duplex or full-duplex mode. Half-duplex is typically used in 10 Mbps Ethernet environments in which the node (the Advisor) transmits and receives at different times. Full-duplex is commonly used in Fast Ethernet environments in which both transmit and receive lines transmit simultaneously resulting in 200 Mbps throughput. You configure the Advisor for half or full duplex operation in the Interface/ Protocols folder discussed in the online Help and later in this Application Note.

Monitor (Through) Connections

The monitor (through) mode shown in Figure 2 is usually used in a switched environment in which the Advisor is placed between a switch port and a server or server segment to which the port normally connects.

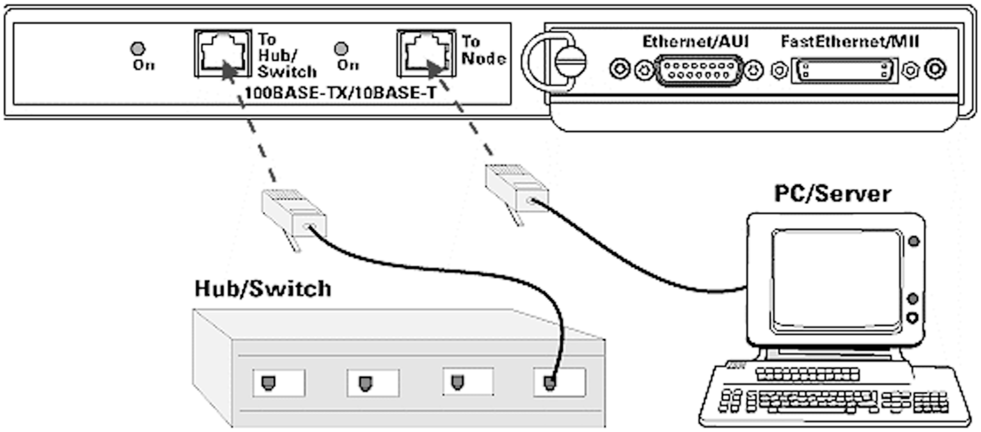

Figure 2: Monitor (Through) Connection

The switch is normally connected directly to the server through a crossover (or a straight) cable. In order to connect the Advisor between the server and the switch, disconnect the cable from the server and connect the cable to the node port of the analyzer. Connect a straight cable from the switch to the switch connector of the analyzer. This mode can be used for both 10 Mbps Ethernet and 100 Mbps Fast Ethernet.

Switched Fast Ethernet Monitor Mode
In order to configure the Advisor to test a switched network, you need to set it up with any of the monitor-through options set as the Line Mode parameter in the Interface/Protocols folder:

- The monitor-through mode does not allow traffic generation.

- The monitor-through mode allows for full-duplex monitoring in which switch-to-server and server-to-switch traffic will be captured. When viewing this traffic in the Decode view, the frames are color-coded to indicate which side (server or switch) the traffic is coming from. Please note that the configuration of the analyzer uses the terms ‘hub’ and ‘node’. This is equivalent to ‘switch’ and ‘server’ when you are working in a switched environment.

- Always use the existing cable that is connected between the server and the switch. The cable can be either a crossover or a straight-through cable depending on the specific equipment used. In addition to using the existing cable, the second cable should *always* be a straight-through cable so signals are not erroneously inverted.

- When using the monitor-through mode, the signal is not re-generated in the analyzer. Thus, the combined length of both cables should not exceed 100 meters.

- If the Advisor power is turned off, the connection between the switch and the server is maintained.

Switched Fiber Fast Ethernet Monitor Mode (Fiber)
In a switched fiber environment, the switch is normally connected directly to the server through a fiber cable. Figure 3 shows the Advisor connected between the server and the switch. To do this, disconnect the fiber cable from the server and connect the cable to the node port of the analyzer. You may encounter devices with ST type connectors instead of SC. The fiber module of the Advisor includes a fiber cable with a SC connector on one end and a ST connector on the other end to allow connection in such cases. The Interface/ Protocols folder should be configured for a fiber connection. Refer to the online Help for more information.

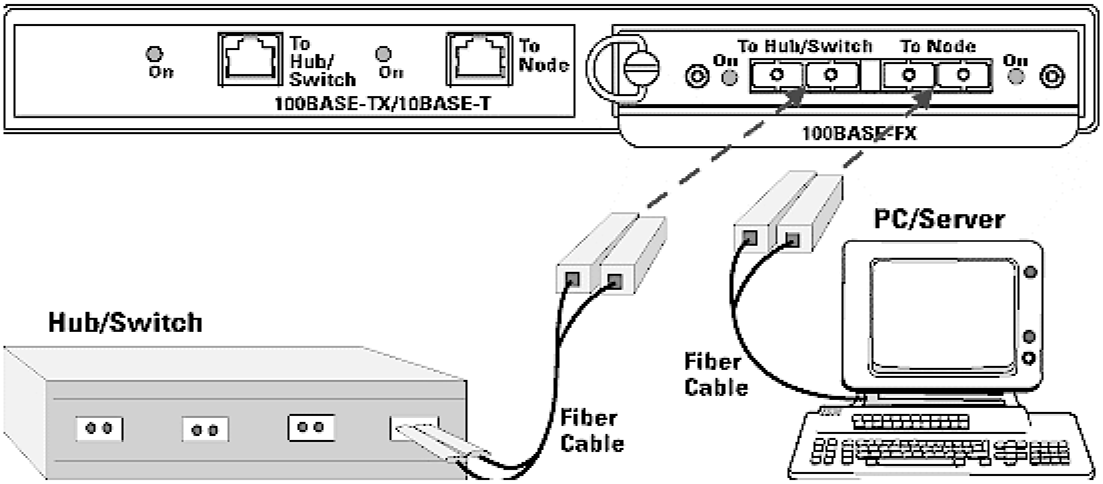

Figure 3: Fiber Connection in Switched Fast Ethernet Environment

- The monitor-through mode does not allow the Advisor to generate traffic onto the network.
- The monitor-through mode allows for full-duplex monitoring. The Advisor will monitor both switch-to-server and server-to-switch traffic. When viewing the traffic in the Decode view, the frames are color-coded to indicate which side (server or switch) the traffic is coming from.
- If the Advisor's power is turned off, the connection between the switch and the server is lost. Power must be maintained for the fiber interface to operate.

Starting the Analysis

Once the Advisor is connected to the network, you need to launch the LAN application and start the analysis. How to do this is covered in the *Advisor LAN Getting Started Guide*, but for convenience it's covered briefly here. You can refer back to this section as you proceed through this Application Note.

Launching the Application

To launch the Advisor LAN Ethernet or Fast Ethernet application, you use the start menu in the Windows desktop as shown in Figure 4. Depending on the type of Advisor mainframe you are using (or the type of undercradle you have attached), you can select from a number of physical layer options – for the purposes of this application note, you would select Ethernet or Fast Ethernet Undercradle. Once you click on either of these options, the Advisor LAN application will be launched.

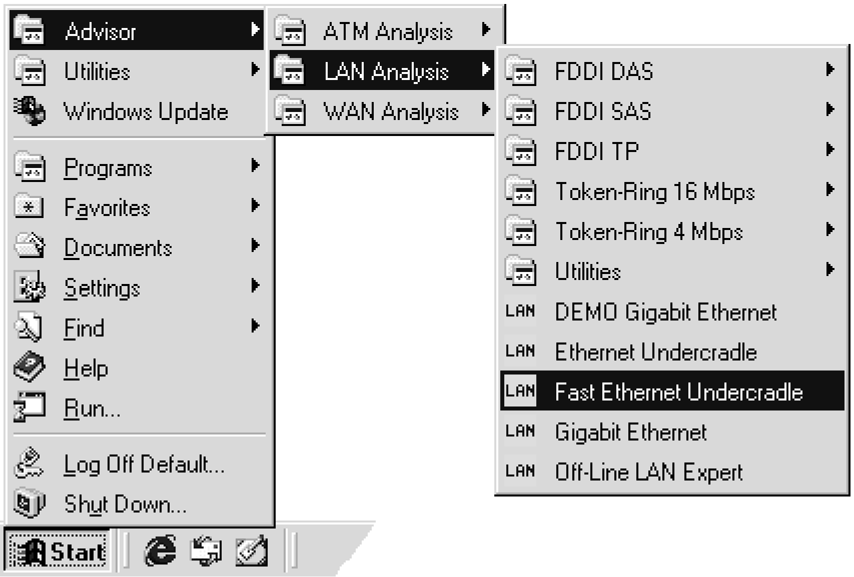

Figure 4: Start Menu

Configuring the Instrument and Measurements

Generally, there are two types of configuration you need to do when getting preparing to run Advisor LAN measurements. First, you need to configure the instrument itself. That is, you need to configure the physical interface to match the physical connection you have made. This is done in the Configuration view as shown in Figure 5. You can also configure the Advisor's capture filters and its logging operations.

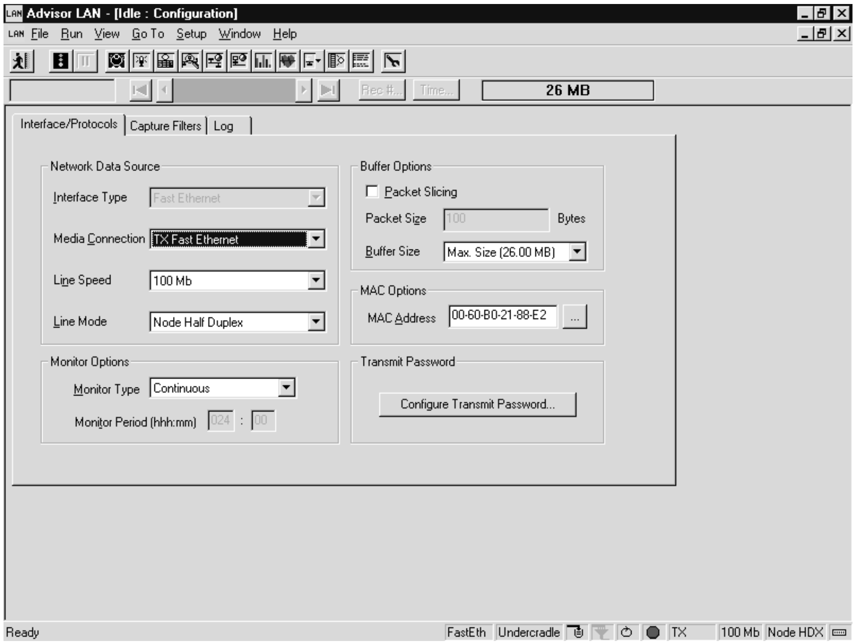

26 MB

Network Data Source
Interface Type: Fast Ethernet
Media Connection: TX Fast Ethernet
Line Speed: 100 Mb
Line Mode: Node Half Duplex

Buffer Options
☐ Packet Slicing
Packet Size: 100 Bytes
Buffer Size: Max. Size (26.00 MB)

MAC Options
MAC Address: 00-60-B0-21-88-E2

Monitor Options
Monitor Type: Continuous
Monitor Period (hhh:mm): 024 : 00

Transmit Password

Ready | FastEth | Undercradle | TX | 100 Mb | Node HDX

Figure 5: Configuration View

Next, and depending on the analysis you intend to perform, you will need to configure the measurement itself. You do this by clicking the Configuration tool bar button at the top of the measurement as shown in Figure 6. The configuration parameters differ depending on the measurement you intend to run.

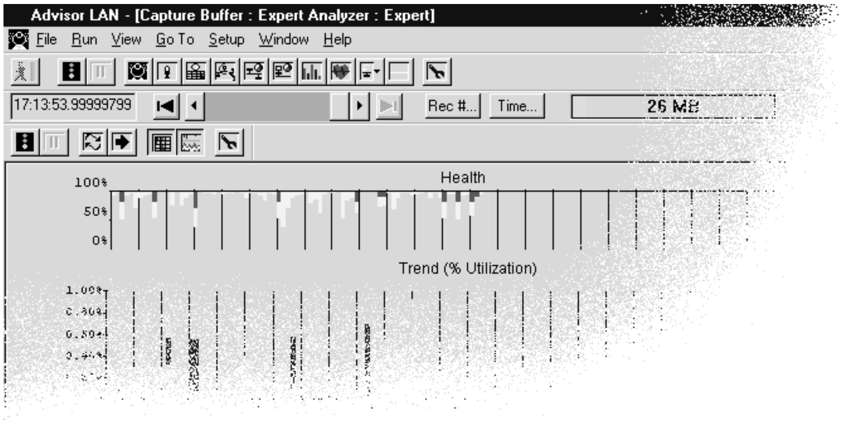

Figure 6: Tool Bar Buttons

Starting a Measurement Run

Similar to configuring the Advisor, there are two ways to start an analysis run. You can start all open measurements (once you launch the application, each measurement needs to be opened explicitly), or you can start a single, individual measurement. The Start tool bar button at the top of the main Advisor screen starts all open measurements, and the Start tool bar button in the measurement tool bar starts the individual measurement (see Figure 6).

Physical and Data Link Layer Analysis of the Network

A network's infrastructure consists of cables, connectors, and network interface cards. As important as these physical components is the protocol layer that provides the interface between the physical layer and upper layer applications - the Data Link layer. Most network problems (80 to 90%) are caused at the physical layer, at the interface between the physical and Data Link layers, or at the Data Link layer itself.

Many of the network inefficiencies caused by problems at layers 1 and 2 are masked by the large bandwidth available in most Ethernet networks (10 or 100 Mbps). These inefficiencies may not seriously affect the network's performance, but as the network grows, or new services are added, inefficiencies can begin to cause problems. It is best to recognize and address problems before they turn into network failure and downtime.

This section covers how to use the Advisor LAN to analyze the lower layers of the OSI reference model: the physical layer and the data link layer.

Physical Layer Analysis

Most troubleshooting scenarios begin with an attempt to isolate possible physical problems such as false contacts, loose cables, and defective NICs. Ensuring that the network's cabling and interface cards are working properly is essential to ensuring trouble-free network operation. Problems at the physical layer will most often manifest themselves as unreasonably high utilization, excessive collisions, bad frame check sequences (FCSs), runts, or jabbers. The best way to measure these key indicators is to use the Advisor's Line Vital Statistics measurement. The Line Vital Statistics measurement (shown in Figure 7) is a physical layer test that measures and displays statistics for utilization, frame counts, collisions, bad FCSs, runts, misaligns, broadcasts, multicasts, etc.

To test the physical and data link layers, connect, launch, and configure the Advisor as described in the previous section, in the *Getting Started Guide*, or in the online Help. Open the Line Vital Statistics measurement by clicking the Line Vital Statistics tool bar button, and start the measurement run as described earlier.

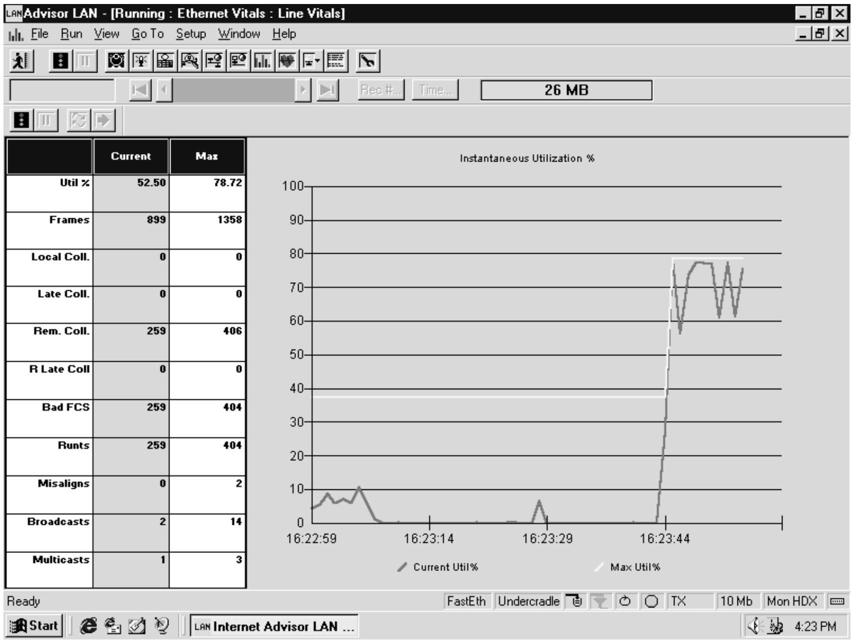

	Current	Max
Util %	52.50	78.72
Frames	899	1358
Local Coll.	0	0
Late Coll.	0	0
Rem. Coll.	259	406
R Late Coll	0	0
Bad FCS	259	404
Runts	259	404
Misaligns	0	2
Broadcasts	2	14
Multicasts	1	3

Figure 7: Line Vital Statistics Measurement

Once the Line Vital Statistics measurement is started, examine the conditions as described in the following paragraphs.

Utilization

One of the first things to examine is utilization levels. In particular, note the maximum utilization and the length of time for which the maximum utilization has been constant. A healthy network should have a *maximum* of 34% to 40% of constant utilization. If utilization levels are consistently higher than this, response time slows, the numbers of collisions increase, and overall performance is degraded.

Collisions

Another important physical layer parameter is collisions. Collisions are a normal part of Ethernet network operations (see Appendix A), but an excessive number of them severely reduces a network's throughput and suggests physical layer problems.

In a coaxial network, if the collision or error rate is very high, immediately suspect a cable problem, a bad or missing terminator, a loose barrel or T-connector, or a crushed cable. In order to isolate the causes of excessive collisions, you must determine if they are *normal* or *late* collisions. Late collisions occur after the normal collision window of 512 bits into a frame (8 bytes of preamble plus 56 frame bytes). Late collisions are caused by excessively long node-to-node propagation time (node-to-node propagation time is the sum of the total time for the signal to travel through all cable segments and repeaters/hubs). Network interface cards whose carrier sense circuitry is not functioning properly can also cause them. Since many Ethernet networks tend to grow in an evolutionary manner, an extra length of cable or new repeater can cause propagation time to exceed Ethernet design specifications. In most cases, this is remedied by reconfiguring the topology of the network.

Excessive “normal” collisions can be caused by any of the following:

- Impedance mismatches – bad terminators, loose barrel/T-connectors, too many connectors in a segment, cable kinks, and segments of other than 50-ohm coaxial cable (e.g. video cable).

- Noise – the most common cause is improper grounding (to prevent ground loops - the network should only be grounded in one location).

- Poor or intermittent connections.

Remote Late Collisions occur when a fragment is received which is likely to be the result of a late collision on another segment of the network. Bridges, repeaters, and 10BaseT hubs commonly filter out the signal levels that positively identify collisions. To be judged a remote late collision, the fragment must be longer than 64 bytes, have a bad FCS, and contain the jam pattern of alternating 1's and 0's.

On 10BaseT networks, virtually all late collisions appear as remote late collisions. On coaxial-based networks, these collisions will be common if the network involves heavy traffic and repeaters.

Remote late collisions should occur only rarely on a normal network. However, infrequent remote late collisions cause little damage. A regular occurrence of remote late collisions indicates a network that is too long, has too much delay from repeaters and bridges, or is very susceptible to noise interruptions.

If remote late collisions are appearing in surprising numbers on coaxial-based networks, try moving the Advisor to different segments to determine which segment has the LOCAL late collisions and investigate from there.

In the case where you have few collisions but a relatively large number of errors, excessive noise or perhaps a bad network interface card might be the cause. The next section discusses these errors and their causes in more detail.

Errors

Other indications of collisions or physical layer failure are errors such as runts (frames shorter than 64 bytes), jabbers (frames longer than 1518 bytes) and bad FCSs:

- A runt is most often a frame fragment resulting from two collided frames. This is a normal network condition, and in small coaxial-based networks runts will be almost nonexistent because the actual collision occurs within the frame preamble. You are more likely to see runts in a larger, or 10BaseT network. In larger networks the longer end-to-end propagation time causes collisions to occur within the transmitted frame. In a 10BaseT network, collisions occur "inside" the hub. The associated hub delays cause many frame fragments (i.e. runts) to be propagated throughout the network.

- Bad frame check sequences can come from many sources: collisions, cable noise, bad network interface cards, and poor connections. It is important to identify the source of these errors. If they are a result of collisions, and the collision rate is not too high, then it need not be a concern.

- Jabbers are not a normal network condition and generally are a serious problem. Jabbers (sometimes called 'giants') are frames that are longer than 1518 bytes, and are usually caused by a node generating frames outside Ethernet specifications or a faulty transceiver on the network. Other possible sources of jabbers include ground loops, or malfunctioning NICs or MAUs.

Once physical layer analysis has been done, it makes sense to move to data link layer analysis. This is covered next.

Data Link Layer Analysis

To perform basic data link layer analysis, close the Line Vitals Statistics measurement and open the Protocol Vital Statistics measurement (shown in Figure 8). Once you start this measurement, you will see protocol statistics graphed according to time across the top of the measurement window and tabulated in a spreadsheet along the bottom. You can graph statistical information by double clicking on the associated statistics option in the spreadsheet. Sometimes you will see information in red. This means that the corresponding data has exceeded a threshold. This threshold is set up in the configuration window of the measurement.

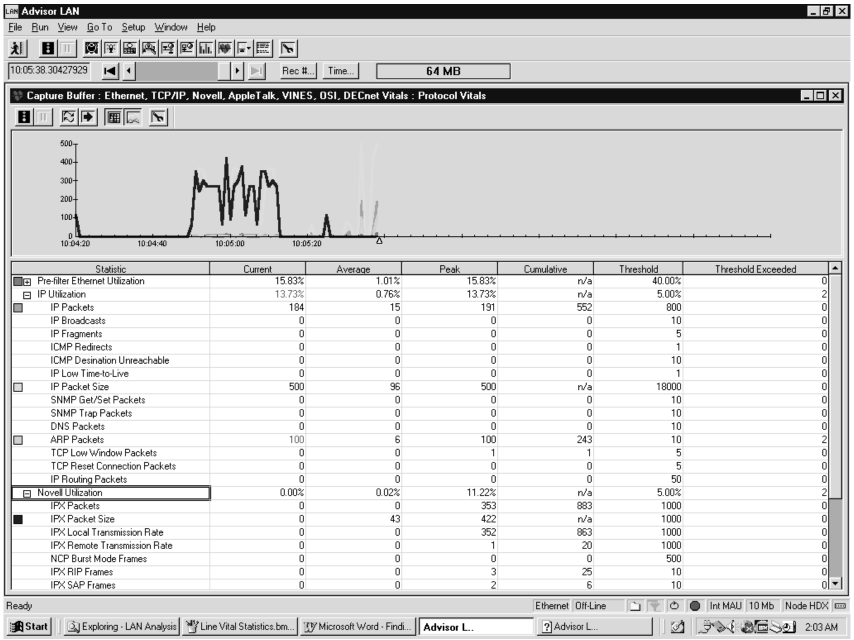
Advisor LAN
File Run View Go To Setup Window Help
10:05:38.30427929
Rec #...
Time...
64 MB
Capture Buffer : Ethernet, TCP/IP, Novell, AppleTalk, VINES, OSI, DECnet Vitals : Protocol Vitals
500
400
300
200
100
0
10:04:20
10:04:40
10:05:00
10:05:20

Statistic	Current	Average	Peak	Cumulative	Threshold	Threshold Exceeded
▨ ⊞ Pre-filter Ethernet Utilization	15.83%	1.01%	15.83%	n/a	40.00%	0
⊟ IP Utilization	13.73%	0.76%	13.73%	n/a	5.00%	2
▨ IP Packets	184	15	191	552	800	0
IP Broadcasts	0	0	0	0	10	0
IP Fragments	0	0	0	0	5	0
ICMP Redirects	0	0	0	0	1	0
ICMP Desination Unreachable	0	0	0	0	10	0
IP Low Time-to-Live	0	0	0	0	1	0
□ IP Packet Size	500	96	500	n/a	18000	0
SNMP Get/Set Packets	0	0	0	0	10	0
SNMP Trap Packets	0	0	0	0	10	0
DNS Packets	0	0	0	0	10	0
□ ARP Packets	100	6	100	243	10	2
TCP Low Window Packets	0	0	1	1	5	0
TCP Reset Connection Packets	0	0	0	0	5	0
IP Routing Packets	0	0	0	0	50	0
⊟ Novell Utilization	0.00%	0.02%	11.22%	n/a	5.00%	2
IPX Packets	0	0	353	883	1000	0
■ IPX Packet Size	0	43	422	n/a	1000	0
IPX Local Transmission Rate	0	0	352	863	1000	0
IPX Remote Transmission Rate	0	0	1	20	1000	0
NCP Burst Mode Frames	0	0	0	0	500	0
IPX RIP Frames	0	0	3	25	10	0
IPX SAP Frames	0	0	2	6	10	0

Figure 8: Protocol Vital Statistics Measurement

Evaluating data link layer integrity is more difficult and complex than the physical infrastructure because there are many more variables to consider. A few of the most important things to consider are the function of the segment in question (i.e. backbone or tributary), the types of nodes attached, and the location of spanning devices (repeaters, brides, routers or switches). These factors influence the utilization, broadcast traffic, and protocol mix. Although it is difficult to generalize, here are some guidelines:

- In a bridge backbone network you should expect a higher percentage of broadcast and multicast packets (broadcast and multicast packets are propagated throughout bridged networks because they are always forwarded). Additionally, average frame sizes will be smaller.

- Pay particular attention to the amount of broadcast traffic. Ideally it should be less than 20 packets per second. All nodes, regardless of the protocol stack being used, are required to process broadcast traffic. Broadcast traffic affects the performance of all nodes on the network.

- Link utilization is highly dependent upon the applications being used (e.g. file transfers vs. interactive processes).

Most of the problems in a network originate in the lower protocol layers. However, there can also be network problems in the upper protocol layers. To detect and troubleshoot these kinds of problems, you will need to use other features of the Advisor. The following example shows how one might examine upper protocol layers to see lower layer problems.

Decode View Exposes the Source of Physical Layer Problems

As an example of physical layer and data link layer analysis, consider a network that consists of three segments attached to a backbone via repeaters (see Figure 9). This network is laid out so that each of its segments could service a different functional area within an organization.

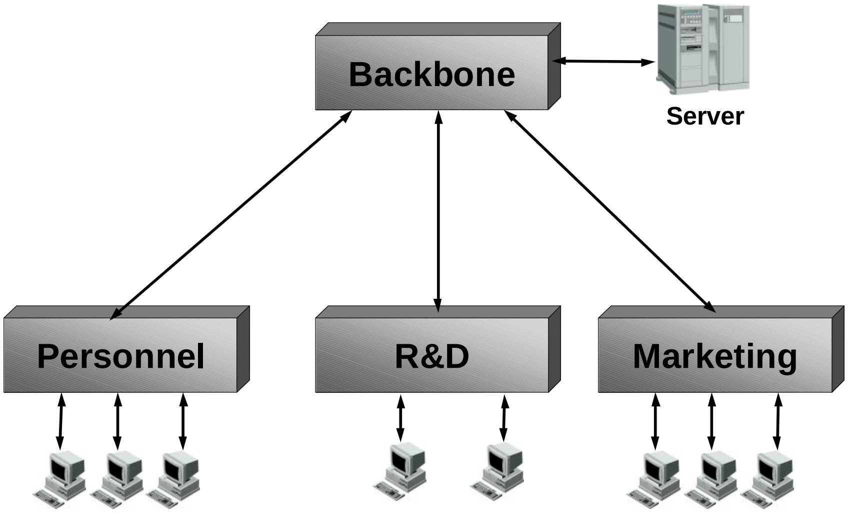

Figure 9: Network Topology Example

Suppose that users of the network begin to complain that their access to shared network devices has slowed. When a series of diagnostic tests such as the Line and Protocols Vitals Statistics measurements (described in the previous section) are run on the network, it is discovered that the collision rate has exceeded the predetermined threshold limit.

Often, a ‘divide and conquer’ approach is used to solve this kind of problem. This method isolates a problem on a network by progressively dividing the network in half physically until, by a process of elimination, the segment creating the problem is found. The process moves then to the segment itself in a similar fashion, continually physically dividing the segment in half until the station creating the problem has been found. Although effective, this process is slow and inconvenient. With the Advisor, this process is simplified considerably.

The threshold limit for local collisions in the Protocol Vital Statistics measurement is set by default to 35 collisions per second. In an Ethernet network, you normally allow a collision rate of 5% of the average network utilization measured in frames per second. Collision statistics that exceed the threshold limit are shown red in the spreadsheet. In additional to the Local Collisions indication of the measurement, there are other types of collisions: remote, late, and remote late collisions. Explanation of any of these collision categories can be found in the on-line Help system along with the probable cause of the collision. This can assist you in determining what course of action to take to fix the problem.

Once you know the kind of problem the network is having (i.e. excessive collisions), you then need to find the source of the problem. To do this, stop the measurement run and open the Decode view as shown in Figure 10. Examine captured traffic in a post-process mode by scrolling through the frames until you find a frame with an “E” on the left side of the decode entry - this indicates an error. An optional procedure is to use the Search button on top of the Decode view to look for Bad FCS Frames (see the online Help for details on searching decoded traffic for specific frames). As shown in the figure, frames involved in collisions have FF’s, 55’s, AA’s, A5’s or 5A’s in the address.

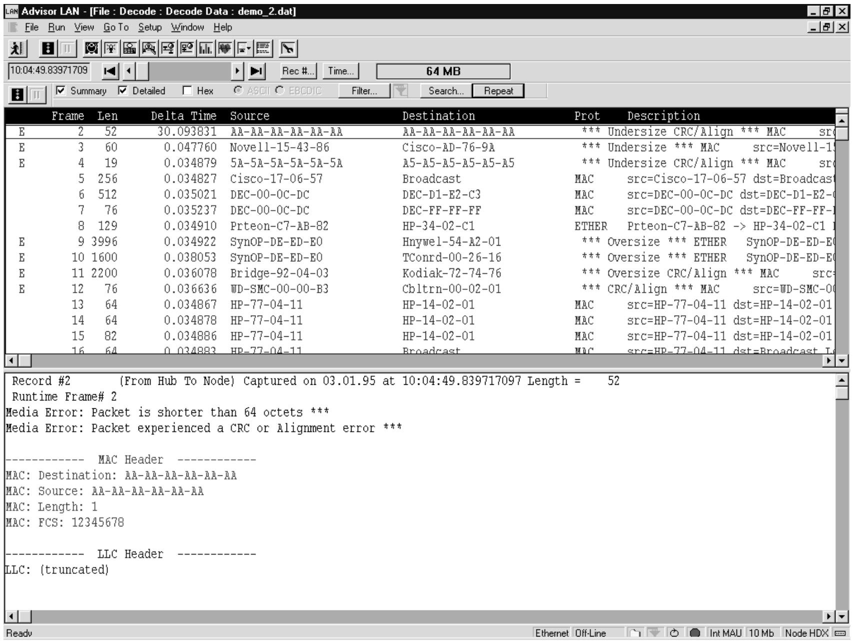
Advisor LAN - [File : Decode : Decode Data : demo_2.dat]
File Run View Go To Setup Window Help
10:04:49.83971709
Rec #...
Time...
64 MB
☑ Summary ☑ Detailed □ Hex ● ASCII ○ EBCDIC
Filter...
Search...
Repeat
Ready
Ethernet
Off-Line
Int MAU
10 Mb
Node HDX

```
     Frame  Len     Delta Time  Source              Destination          Prot    Description
 E       2   52      30.093831  AA-AA-AA-AA-AA-AA   AA-AA-AA-AA-AA-AA     *** Undersize CRC/Align *** MAC     sr
 E       3   60       0.047760  Novell-15-43-86     Cisco-AD-76-9A        *** Undersize *** MAC    src=Novell-1
 E       4   19       0.034879  5A-5A-5A-5A-5A-5A   A5-A5-A5-A5-A5-A5     *** Undersize CRC/Align *** MAC     sr
         5  256       0.034827  Cisco-17-06-57      Broadcast            MAC     src=Cisco-17-06-57 dst=Broadcas
         6  512       0.035021  DEC-00-0C-DC        DEC-D1-E2-C3         MAC     src=DEC-00-0C-DC dst=DEC-D1-E2-
         7   76       0.035237  DEC-00-0C-DC        DEC-FF-FF-FF         MAC     src=DEC-00-0C-DC dst=DEC-FF-FF-
         8  129       0.034910  Prteon-C7-AB-82     HP-34-02-C1          ETHER   Prteon-C7-AB-82 -> HP-34-02-C1
 E       9 3996       0.034922  SynOP-DE-ED-E0      Hnywel-54-A2-01       *** Oversize *** ETHER   SynOP-DE-ED-E
 E      10 1600       0.038053  SynOP-DE-ED-E0      TConrd-00-26-16       *** Oversize *** ETHER   SynOP-DE-ED-E
 E      11 2200       0.036078  Bridge-92-04-03     Kodiak-72-74-76       *** Oversize CRC/Align *** MAC     src
 E      12   76       0.036636  WD-SMC-00-00-B3     Cbltrn-00-02-01       *** CRC/Align *** MAC    src=WD-SMC-0
        13   64       0.034867  HP-77-04-11         HP-14-02-01          MAC     src=HP-77-04-11 dst=HP-14-02-01
        14   64       0.034878  HP-77-04-11         HP-14-02-01          MAC     src=HP-77-04-11 dst=HP-14-02-01
        15   82       0.034886  HP-77-04-11         HP-14-02-01          MAC     src=HP-77-04-11 dst=HP-14-02-01
        16   64       0.034883  HP-77-04-11         Broadcast            MAC     src=HP-77-04-11 dst=Broadcast
```

```
 Record #2       (From Hub To Node) Captured on 03.01.95 at 10:04:49.839717097 Length =    52
 Runtime Frame# 2
Media Error: Packet is shorter than 64 octets ***
Media Error: Packet experienced a CRC or Alignment error ***

------------  MAC Header  ------------
MAC: Destination: AA-AA-AA-AA-AA-AA
MAC: Source: AA-AA-AA-AA-AA-AA
MAC: Length: 1
MAC: FCS: 12345678

------------  LLC Header  ------------
LLC: (truncated)
```

Figure 10: Decode View

Note: because of the custom front-end processing hardware of the Advisor, it is possible to capture frames that are less than eight bytes long. Standard processors cannot capture frames less than 14 bytes long. Without this capability, collided frame fragments could not be captured along with their arrival times. It would also be impossible to discover which stations are causing collisions.

Now that you have identified the frames in the decoded traffic that have been involved in collisions, you need to identify where they are coming from. Typically, the frame immediately following the errored frame will be from one of the two stations causing the collision. To verify this, look at the timestamp. If the time difference between the errored frame and the next frame is less than about 150 milliseconds (the time is dependent upon the location of the station on the network and the length of the segment), it probably is one of the frames that collided. If traffic levels are high on the network, another station could have sent a frame prior to the offending station sending its frame.

You will also want to check the frames following several other errored frames. If the collision rate is high, there should be a large number. One of the stations causing the collisions should keep reappearing after the errored frame. Having discovered the offending station using the Advisor's Decode view, you can take corrective action. You can then repeat the network tests to check that the collision rate is once more in an acceptable range.

Using the Advisor's Expert Analyzer

Many network problems are not obvious and sometimes do not present any apparent symptoms. You might suspect a problem, however, due to constant retransmissions, high numbers of collisions, aborted connections, and so on. The question is how to isolate a problem that could originate from any number of places. You could begin troubleshooting each of the layers of the protocol stack, or you could use the Advisor LAN's Expert Analyzer. The Expert Analyzer can save you troubleshooting time because it collects information from the network, analyzes the information, and alerts you to error conditions. More importantly, it recommends possible solutions to the problem.

Performing a Checkup on the Network

In this example, we demonstrate how to use the Advisor's Expert Analyzer (shown in Figure 11) to evaluate the overall health of the network.

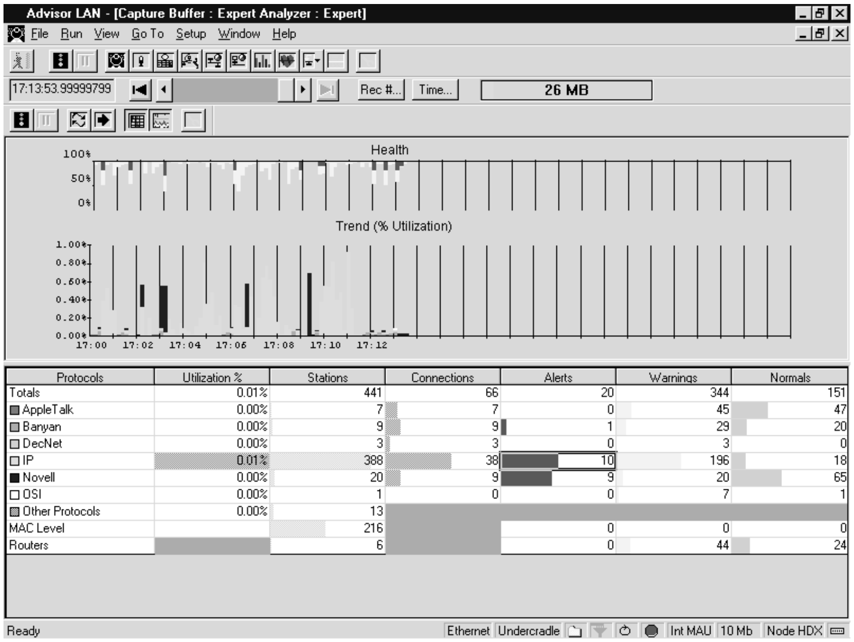

Protocols	Utilization %	Stations	Connections	Alerts	Warnings	Normals
Totals	0.01%	441	66	20	344	151
■ AppleTalk	0.00%	7	7	0	45	47
■ Banyan	0.00%	9	9	1	29	20
□ DecNet	0.00%	3	3	0	3	0
□ IP	0.01%	388	38	10	196	18
■ Novell	0.00%	20	9	9	20	65
□ OSI	0.00%	1	0	0	7	1
▨ Other Protocols	0.00%	13				
MAC Level		216		0	0	0
Routers		6		0	44	24

Figure 11: Expert Analyzer

Begin by connecting, launching, and configuring the Advisor as described in the *Connecting, Starting, and Configuring the Advisor* section of this Application Note and in the Advisor's *Getting Started Guide* and online Help.

Open the Expert Analyzer by clicking the Expert Analyzer tool bar button and then start the measurement. Figure 11 shows that the network utilization is low - about 1%. However, red and yellow lines in the graph at the top of the window suggests that there are traffic events affecting the overall health of the network. To learn more about these events, you can click on the colored lines in the graph to see more exact textual information.

As shown in Figure 11, there are currently 441 stations and 66 connections. You can also see that IP is the most predominant protocol by looking at the biggest horizontal color bars representing utilization, stations, and connections.

To learn more about the TCP/IP traffic on this network - that is, to find out more about the protocols and traffic carried by the IP packets - double-click the IP row of the protocol column. The result, shown in Figure 12, indicates that the HTTP protocol (shown as WWW in the pie chart) is the most used, followed by TELNET and FTP. This might be normal in a segment that contains WEB servers. On the other hand, if it is a production segment of a sales office, this might indicate that users are surfing the Web.

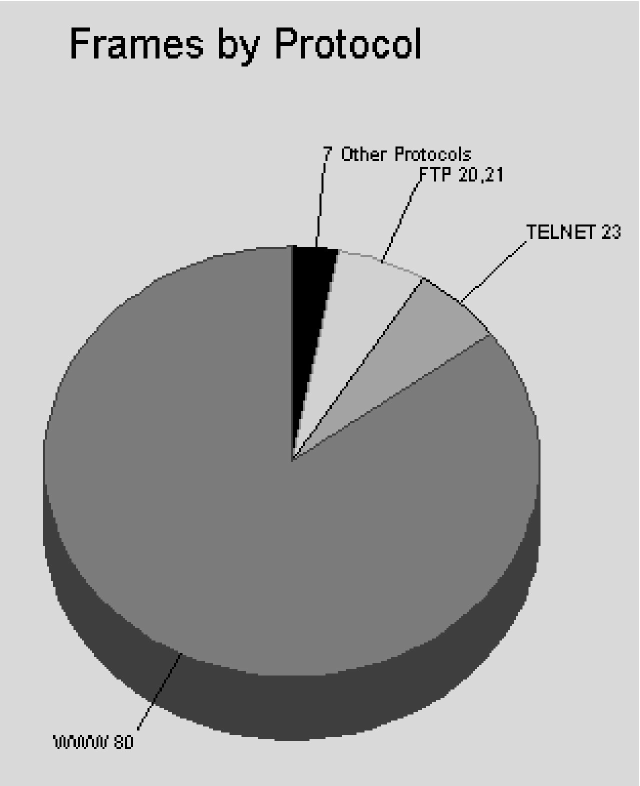

Figure 12: Protocol Distribution

To learn more about the utilization and to determine which user(s) are causing the excess HTTP traffic, return to the main Expert Analyzer view by clicking on the ‘x’ button in the upper right corner of the currently open window. Double-click the cell where the utilization column and IP row in the spreadsheet intersect to show which nodes are using HTTP or the WWW 80 TCP port. You can sort the information using the options in the View menu (shown in Figure 13).

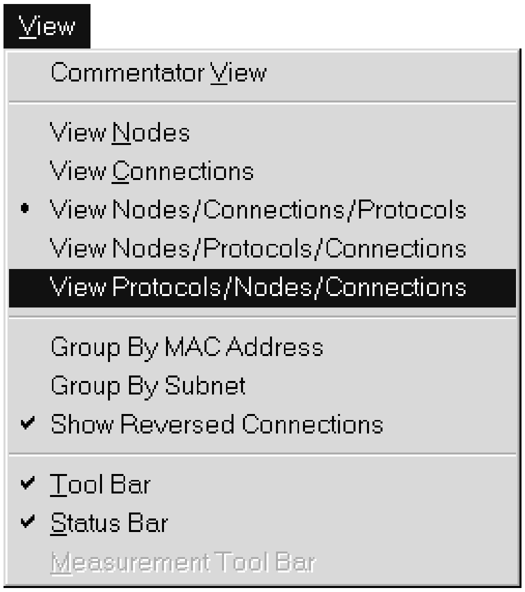
View
Commentator View
View Nodes
View Connections
• View Nodes/Connections/Protocols
View Nodes/Protocols/Connections
View Protocols/Nodes/Connections
Group By MAC Address
Group By Subnet
✔ Show Reversed Connections
✔ Tool Bar
✔ Status Bar
Measurement Tool Bar

Figure 13: View Menu

Note: the Advisor uses a database that allows you to get to any information from any point within the user interface. For example, you can access protocol utilization or network errors from node information.

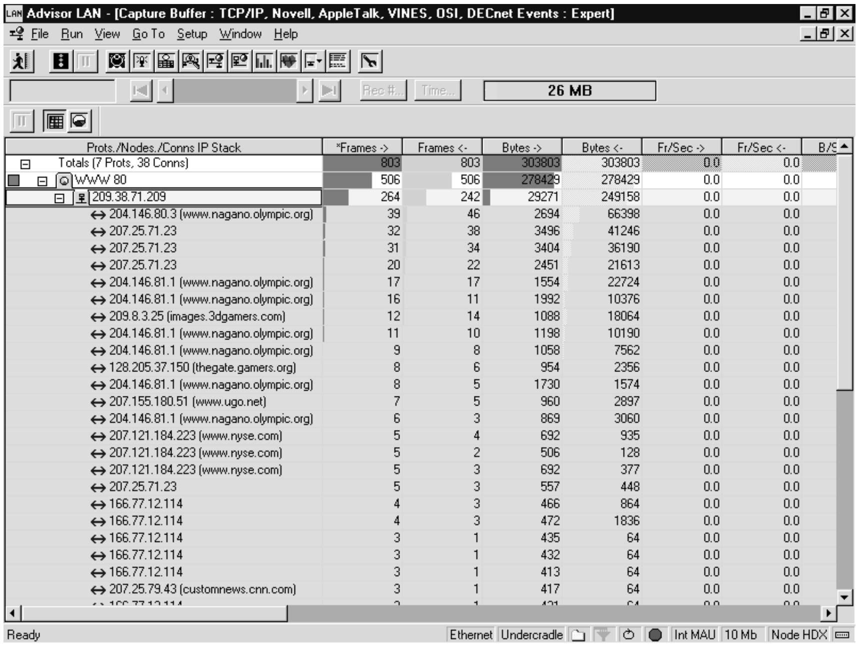

Prots./Nodes./Conns IP Stack	*Frames ->	Frames <-	Bytes ->	Bytes <-	Fr/Sec ->	Fr/Sec <-	B/S
⊟ Totals (7 Prots, 38 Conns)	803	803	303803	303803	0.0	0.0	
⊟ WWW 80	506	506	278429	278429	0.0	0.0	
⊟ 209.38.71.209	264	242	29271	249158	0.0	0.0	
↔ 204.146.80.3 (www.nagano.olympic.org)	39	46	2694	66398	0.0	0.0	
↔ 207.25.71.23	32	38	3496	41246	0.0	0.0	
↔ 207.25.71.23	31	34	3404	36190	0.0	0.0	
↔ 207.25.71.23	20	22	2451	21613	0.0	0.0	
↔ 204.146.81.1 (www.nagano.olympic.org)	17	17	1554	22724	0.0	0.0	
↔ 204.146.81.1 (www.nagano.olympic.org)	16	11	1992	10376	0.0	0.0	
↔ 209.8.3.25 (images.3dgamers.com)	12	14	1088	18064	0.0	0.0	
↔ 204.146.81.1 (www.nagano.olympic.org)	11	10	1198	10190	0.0	0.0	
↔ 204.146.81.1 (www.nagano.olympic.org)	9	8	1058	7562	0.0	0.0	
↔ 128.205.37.150 (thegate.gamers.org)	8	6	954	2356	0.0	0.0	
↔ 204.146.81.1 (www.nagano.olympic.org)	8	5	1730	1574	0.0	0.0	
↔ 207.155.180.51 (www.ugo.net)	7	5	960	2897	0.0	0.0	
↔ 204.146.81.1 (www.nagano.olympic.org)	6	3	869	3060	0.0	0.0	
↔ 207.121.184.223 (www.nyse.com)	5	4	692	935	0.0	0.0	
↔ 207.121.184.223 (www.nyse.com)	5	2	506	128	0.0	0.0	
↔ 207.121.184.223 (www.nyse.com)	5	3	692	377	0.0	0.0	
↔ 207.25.71.23	5	3	557	448	0.0	0.0	
↔ 166.77.12.114	4	3	466	864	0.0	0.0	
↔ 166.77.12.114	4	3	472	1836	0.0	0.0	
↔ 166.77.12.114	3	1	435	64	0.0	0.0	
↔ 166.77.12.114	3	1	432	64	0.0	0.0	
↔ 166.77.12.114	3	1	413	64	0.0	0.0	
↔ 207.25.79.43 (customnews.cnn.com)	3	1	417	64	0.0	0.0	

Figure 14: IP Node Identified

As Figure 14 shows, the “surfer” in the network is identified by the IP node 209.38.71.209 and the visited Web sites include those concerned with games and the olympics. You might want to try different views to get additional information. When you finish, close this window.

As a final analysis step, you can investigate the abnormal events in the network that cause the network health to decrease. These events are reported in the Alarm, Warnings, and Normal columns of the Expert Analyzer. The Warning and Alarm events are of the most interest. A warning event means that there is a possible error in the network. Alarm events are those that are truly errors and need attention. To investigate these events, double-click on any row or column in which you are interested. As shown in Figure 15, you can select from Node Events, Connection Events, and All Events. Expand the Display All Events option. You can enable or disable the type of events by clicking on the A, W, or N icons. Figure 15 shows an event called IP Broadcast Storms. The criteria (broadcast/seconds) and IP addresses the analyzer uses to detect broadcast storms are user-configurable and are described in the online Help. One cause of IP broadcast frames may be the BOOTP program used by some diskless workstations to obtain temporary IP addresses. If the BOOTP timeout value is misconfigured, an IP broadcast storm may result. An IP broadcast storm may also indicate misconfigured multicast addresses on an Ethernet network, or it may indicate the need to implement multicast addressing on the network.

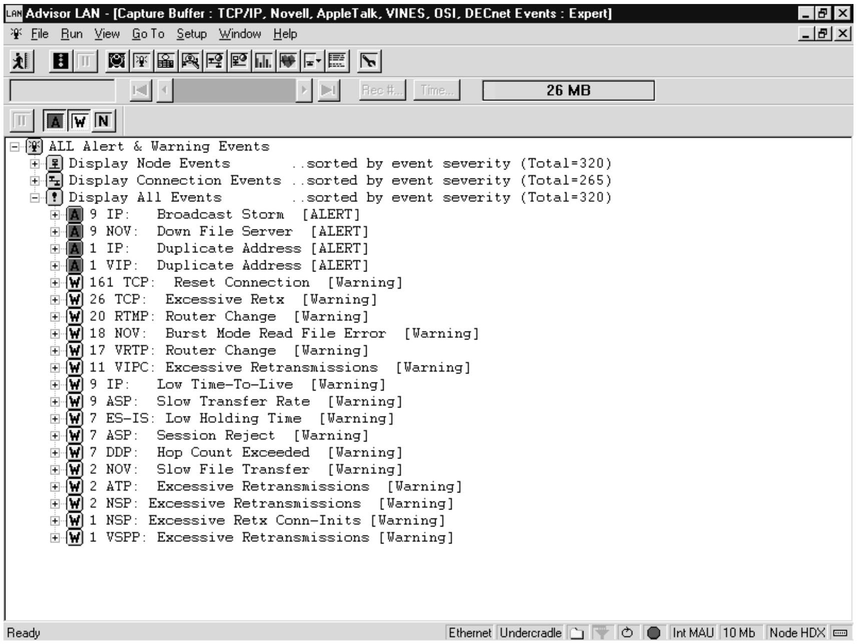
Advisor LAN - [Capture Buffer : TCP/IP, Novell, AppleTalk, VINES, OSI, DECnet Events : Expert]
File Run View Go To Setup Window Help
Rec #... Time...
26 MB
A W N
Ready
Ethernet Undercradle Int MAU 10 Mb Node HDX

```
⊟ ALL Alert & Warning Events
  ⊞ Display Node Events       ..sorted by event severity (Total=320)
  ⊞ Display Connection Events ..sorted by event severity (Total=265)
  ⊟ Display All Events        ..sorted by event severity (Total=320)
    ⊞ A 9 IP:   Broadcast Storm  [ALERT]
    ⊞ A 9 NOV:  Down File Server  [ALERT]
    ⊞ A 1 IP:   Duplicate Address [ALERT]
    ⊞ A 1 VIP:  Duplicate Address [ALERT]
    ⊞ W 161 TCP:  Reset Connection  [Warning]
    ⊞ W 26 TCP:  Excessive Retx  [Warning]
    ⊞ W 20 RTMP: Router Change  [Warning]
    ⊞ W 18 NOV:  Burst Mode Read File Error  [Warning]
    ⊞ W 17 VRTP: Router Change  [Warning]
    ⊞ W 11 VIPC: Excessive Retransmissions  [Warning]
    ⊞ W 9 IP:   Low Time-To-Live  [Warning]
    ⊞ W 9 ASP:  Slow Transfer Rate  [Warning]
    ⊞ W 7 ES-IS: Low Holding Time  [Warning]
    ⊞ W 7 ASP:  Session Reject  [Warning]
    ⊞ W 7 DDP:  Hop Count Exceeded  [Warning]
    ⊞ W 2 NOV:  Slow File Transfer  [Warning]
    ⊞ W 2 ATP:  Excessive Retransmissions  [Warning]
    ⊞ W 2 NSP: Excessive Retransmissions  [Warning]
    ⊞ W 1 NSP: Excessive Retx Conn-Inits [Warning]
    ⊞ W 1 VSPP: Excessive Retransmissions [Warning]
```

Figure 15: Alert and Warning Events

Another example of events in this analysis is the Novell Burst Mode Read File Error also shown in Figure 15. This is a ‘warning’ event that indicates that an error occurred on a file read process in burst mode. This can occur when a user does not have file read privileges for a specific file. Another reason can be that the server is out of disk space.

Configuration Tip: You can specify the number of broadcasts per second that should be considered a storm in the IP Broadcast Storm (#/sec) field of the Configuration dialog box. You can also configure the IP broadcast address in the Configuration dialog box. The Broadcast address is also setup in the configuration. The default is all routes broadcast of 255.255.255.255. This may be edited to examine specific subnet broadcasts. When set to 255.255.255.255, the older style broadcast of 0.0.0.0 is also considered a broadcast address in storm determination.

Troubleshooting a Slow Connection

Because the potential causes are many, isolating the cause of slow network response can be difficult. There could be excessive traffic from other nodes that affects the performance of the communication. Or there could be physical problems like defective network interface cards (NICs) or loose cables. Often the problem is in the node itself due to protocol and application configuration problems. The Expert Analyzer can provide the means to quickly determine the cause.

To isolate a slow connection problem, connect, launch, and configure the Advisor as discussed in the *Connecting, Starting, and Configuring the Advisor* section, the Advisor's *Getting Started Guide*, or the online Help. Open the Expert Analyzer view (as described in the last section) and start the measurement to get an overview of the network's health and to isolate the problem.

Drilling down to the particulars of Warning events as described in the previous section, you will obtain a screen similar to that shown in Figure 16. As shown in the figure, a specific node (address 15.6.72.20) has 24 warnings consisting of 22 TCP Reset Connections and 2 TCP Low Window events. The TCP (Transmission Control Protocol) Reset Connection message is considered a 'warning' event since it is the result of an abnormal termination of the TCP connection. This might occur when a server drops a connection because of misconfigured resources or because of router problems. Another possibility is when a user aborts a Web connection due to slow response from a server. Many older Web browsers use Reset Connection to shut down their HTTP connections. However, newer browsers are better at complying with the TCP disconnection procedure. Because many older browsers are still in use, TCP Reset Connection events are very common in many of today's LANs.

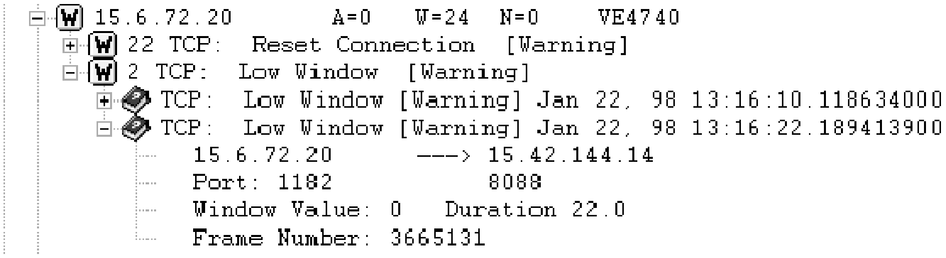

```
⊟ W 15.6.72.20       A=0   W=24  N=0    VE4740
   ⊞ W 22 TCP:  Reset Connection  [Warning]
   ⊟ W 2 TCP:  Low Window  [Warning]
      ⊞ TCP:  Low Window [Warning] Jan 22, 98 13:16:10.118634000
      ⊟ TCP:  Low Window [Warning] Jan 22, 98 13:16:22.189413900
            15.6.72.20      ---> 15.42.144.14
            Port: 1182           8088
            Window Value: 0   Duration 22.0
            Frame Number: 3665131
```

Figure 16: Node-specific Information

Now, click to expand the row that indicates the Low Window Warning. Double-click the book icon to open the on-line Commentator as shown in Figure 17. With the information described in the Commentator, you can begin to isolate and solve the problem.

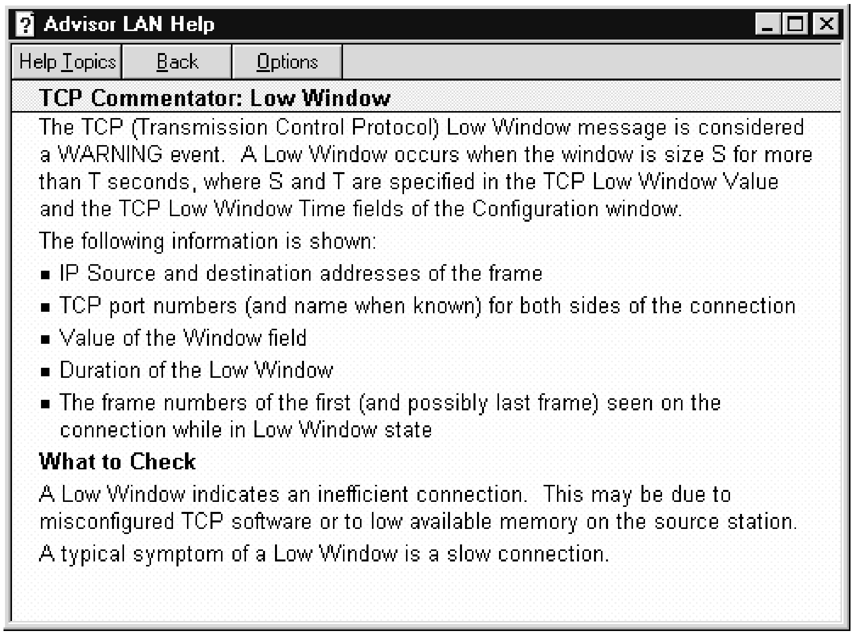

TCP Commentator: Low Window

The TCP (Transmission Control Protocol) Low Window message is considered a WARNING event. A Low Window occurs when the window is size S for more than T seconds, where S and T are specified in the TCP Low Window Value and the TCP Low Window Time fields of the Configuration window.

The following information is shown:

- IP Source and destination addresses of the frame
- TCP port numbers (and name when known) for both sides of the connection
- Value of the Window field
- Duration of the Low Window
- The frame numbers of the first (and possibly last frame) seen on the connection while in Low Window state

What to Check

A Low Window indicates an inefficient connection. This may be due to misconfigured TCP software or to low available memory on the source station.

A typical symptom of a Low Window is a slow connection.

Figure 17: On-line Commentator

Troubleshooting a Duplicate IP Address

An IP (Internet Protocol) duplicate address is a conflict between logical addresses that occurs when two physical MAC addresses erroneously share the same logical IP address. Sometimes when changes are made in a network, or when the network administrator does not have strict control of assigned IP addresses, errors in configuration can occur resulting in this type of duplication.

Symptoms of duplicated IP addresses can include users complaining of erratic TCP connection losses or ICMP warning messages from routers in the network. The easiest way to corroborate a duplicate IP address is to use ARP to see if more than one MAC station on your network is using the suspect IP address. Typically, this can be done by pinging the suspected IP address using ping commands on a network workstation or by using the Advisor's ping utility. Figure 18 shows a duplicate IP address displayed in the Advisor's Expert Analyzer. The information that is shown is the IP Network Address, the conflicting MAC Addresses and the frame numbers (as assigned by the Advisor) of the two captured frames. When the Advisor identifies routers, these routers are excluded from the duplicate IP reporting. However, an erroneous duplicate IP address may be reported for a short time span if a router has not yet been identified. Router MAC addresses are ignored as sources of duplicates. The IP Duplicate Address Delay parameter in the TCP/IP folder of the Commentator configuration allows the system to learn about IP routers (via RIP, IGRP, OSPF) before any IP duplicate event is generated. Setting this value too low may result in false duplicates due to router packet load balancing algorithms.

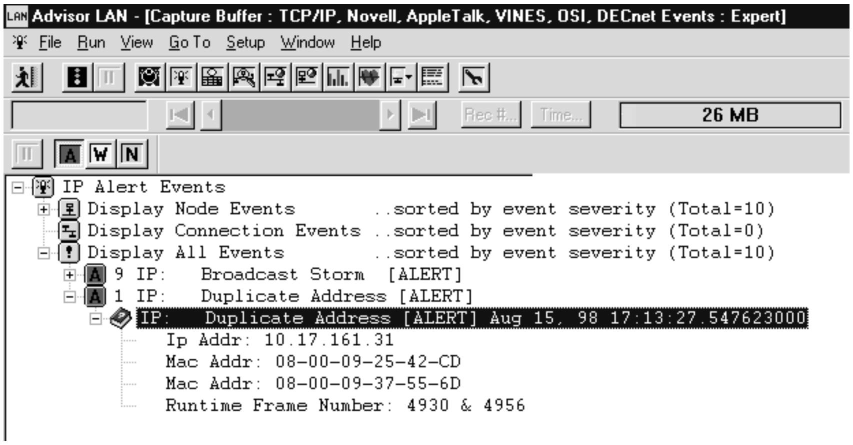
Advisor LAN - [Capture Buffer : TCP/IP, Novell, AppleTalk, VINES, OSI, DECnet Events : Expert]
File
Run
View
Go To
Setup
Window
Help
Rec #...
Time...
26 MB
A
W
N

```
IP Alert Events
  Display Node Events        ..sorted by event severity (Total=10)
  Display Connection Events  ..sorted by event severity (Total=0)
  Display All Events         ..sorted by event severity (Total=10)
    9 IP:   Broadcast Storm  [ALERT]
    1 IP:   Duplicate Address [ALERT]
      IP:   Duplicate Address [ALERT] Aug 15, 98 17:13:27.547623000
          Ip Addr: 10.17.161.31
          Mac Addr: 08-00-09-25-42-CD
          Mac Addr: 08-00-09-37-55-6D
          Runtime Frame Number: 4930 & 4956
```

Figure 18: Duplicate IP Address

Proactive Network Analysis

Because of the combination of distributed systems and network based applications, and the use of 10BaseT to replace the coaxial infrastructures of the mid-eighties, network administrators are less worried about catastrophic network failures than they were in the past. A bigger concern now is the optimization of network performance and efficiency that can result in real increases in user productivity. In addition, network administrators cannot spend their time 'fighting fires'. This is a stressful activity that consumes administrator's time without providing any value-add service to their customers.

Proactive network analysis, also known as 'baselining', is when you use the protocol analyzer to measure networks under normal operating conditions. Proactive analysis, if done regularly, can expose trends and potential problems before they become critical – that is, it provides a method to both increase the network's performance and to avoid problems that contribute to crisis management. Proactive network analysis allows you to understand the current usage of the network in order to eliminate inefficiencies, optimize the network performance, plan network growth, and reduce network downtimes. The next sub-section is an example of how proactive network analysis can make troubleshooting more successful, and the following one describes a general process that can be used to actively fine-tune your network and avoid problems.

Network Errors May Be Only Part of the Problem

Suppose that an end user of a network is complaining of poor response time from the file server, and when the Advisor is used to determine the source of the problem, it detects 200 collisions per second in the Expert Analyzer and Network Statistics views. In some cases, if you were to address only the collisions issue, the problem may still remain. For example, suppose the poor response time is not only caused by the 200 collisions per second but also by a new user who is consuming 30% of the network's bandwidth. By simply looking at error statistics, this new user would not have been seen and the problem would have remained.

Consider, on the other hand, if the troubleshooting scenario above also included an examination of previous network baselines. In this case, a trend comparison between the old baseline and recently measured utilization statistics would have shown that the level of activity in the network was increased, thereby slowing down communication. Notice, that an old baseline of the network would not have existed had a network administrator not proactively measured the network during times of normal operation. This is an example of how proactive network analysis in the form of baselining can increase the likelihood of solving the problem the first time.

A Proactive Testing Process

Generally, there are three steps to baseline a network: collect information, create a report, and interpret the results. These steps, as well as a tool that can assist in baseline analysis and reporting, are discussed next.

Agilent Advisor Reporter

Baseline analysis is enhanced by a data correlation and reporting tool called the Agilent Technologies Advisor Reporter that will be referred to throughout the remainder of this section. While the Advisor itself provides a way to save and print measurement data, it is often better to use a dedicated tool. The Advisor Reporter is a software application that augments Microsoft Excel®, allowing you to take Advisor measurement data and statistics and create simple graphics showing network conditions. You can also create documents using Microsoft Word®. Please contact your Agilent Technologies representative to obtain the Agilent Advisor Reporter application note for more information.

Collect Information

To collect information, you connect the Advisor to the network under test and start gathering traffic statistics and decodes to be saved and examined later. A comprehensive baseline should contain information for all segments of the network. This might also include WAN and ATM links if the network has a wide area component. Mission critical server or communications segments are of particular importance. Various types of measurements ranging from physical and data-link layer analysis to statistical and decode analysis should be performed to get a complete view of how the network behaves.

In addition to the location and scope of data gathering, a consistent schedule should be used when performing network characterizations in order to expose short and long term trends and time dependent network behavior. Critical segments should be measured more often. When a modification in the network is scheduled, it is recommended to create a baseline of the network before and after the change.

To make configuring the Advisor LAN to gather baseline information easier, start by closing all measurement windows. Then open only the measurements you want logged to disk. **Note:** The Advisor Reporter does not support Expert Advisor statistics, so if you want all gathered data to be formatted and presented by the Advisor Reporter, you will have to use the Advisor's other statistics views.

Now, go to the Log folder in the Configuration view and select the measurements you want to log, along with a filename, comments, logging interval, and logging period. It is very important to select adequate values for the interval and period samples. For troubleshooting; you normally use short interval samples, like 1 or 10 seconds. But for long term trends and baselining, you typically use long interval times ranging from 10 to 30 minutes or longer. Figure 19 shows the difference between the logging period and the logging interval.

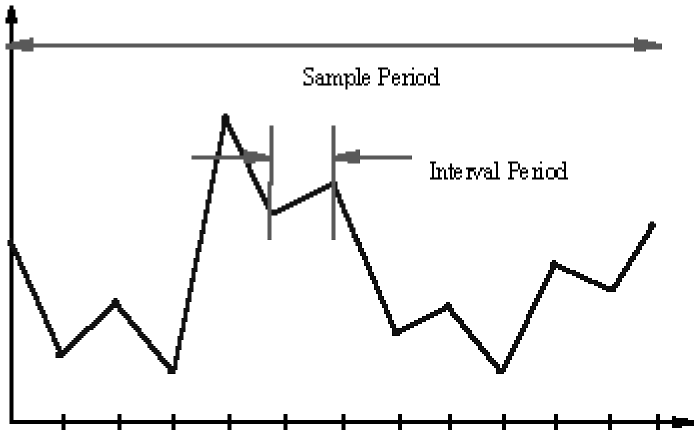

Figure19: Logging Period vs. Interval

Consider the following as you decide how to configure the Advisor for logging:

- You may want to limit what the Advisor captures during the analysis run. You can limit capture to specific protocols or network elements. Use the Capture Filters folder in the Configuration view to do this.
- It is not necessary to run the Line Vitals measurement if you plan to use the Protocol Vitals measurement because Line Vitals is a subset of Protocol Vitals.
- Except for utilization statistics, all values in the tables that Advisor Reporter generates are presented as total counts per sample interval.
- If you are not planning on reporting on particular protocol stacks or specific statistics within a stack, be sure to disable them in the Protocol Vitals measurement Configuration dialog. Only enabled protocols and statistics appear in Advisor Reporter tables.

- All forms of the Protocol Stats measurement are supported by Advisor Reporter, not just the default measurement on the Advisor command bar. You can load other specialized forms of statistics measurements by selecting Open Measurement from the File menu and selecting a .msx file. For example, opening the measurement file “Protocol Stats IP.msx” will load a version of the Protocol Stats measurement that focuses mostly on IP related statistics.

- You can have multiple forms of Protocol Stats running and saved to a file. This makes it possible to process them selectively with the Advisor Reporter. Once you have the the Advisor configured, start the analyzer by clicking the top most Start tool bar button (this starts all open measurements simultaneously). Let the analyzer run for the sample period. If you need to stop the analysis before the end of the sample period, just click on the Stop tool bar button. Captured data up to that point will be preserved.

Create a Report

While the collection of information itself is important, if it is not clearly and understandably presented, it will be much more difficult to efficiently analyze and present to colleagues and management. Graphics that show two-variable information simultaneously are very important because they provide correlated data that is easier to interpret. It is a good idea to save copies of the network baseline, because you can then detect trends over long periods of time. Once you have captured statistics and decodes into a log file with the Advisor, you are ready to process this file to create a report. You can use the Advisor Reporter to easily generate professional graphics and reports as shown in Figure 20.

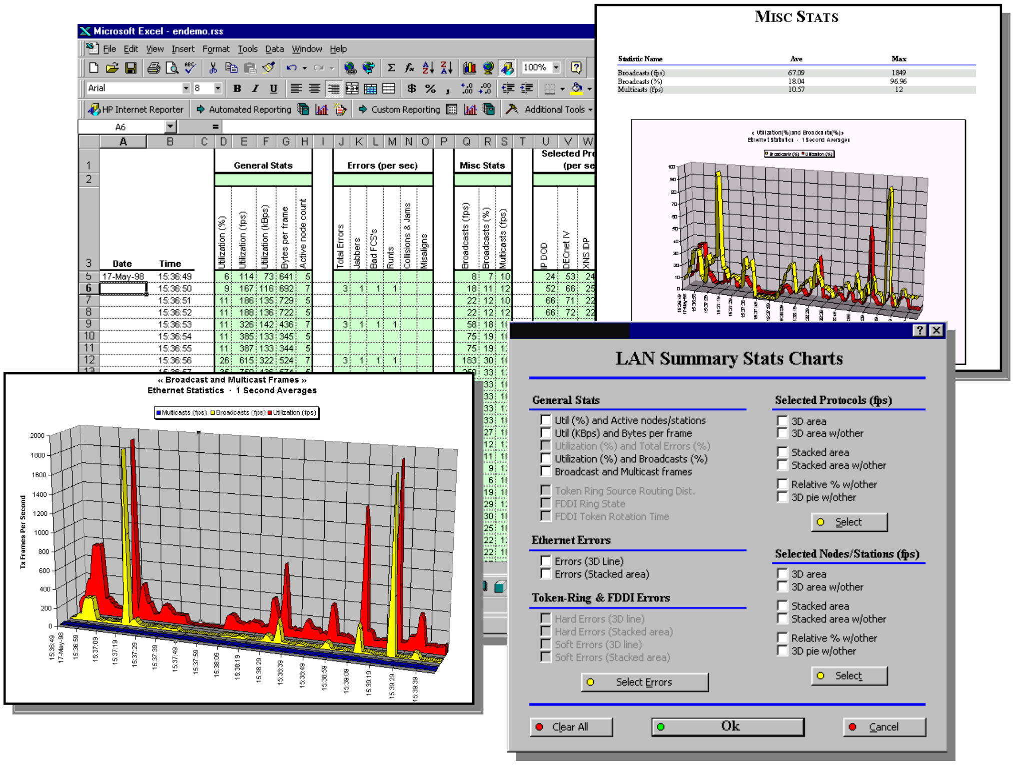
Microsoft Excel - endemo.rss
File Edit View Insert Format Tools Data Window Help
Arial 8
100%
HP Internet Reporter
Automated Reporting
Custom Reporting
Additional Tools
A6
« Broadcast and Multicast Frames »
Ethernet Statistics · 1 Second Averages
■ Multicasts (fps) ■ Broadcasts (fps) ■ Utilization (fps)
Tx Frames Per Second
2000
1800
1600
1400
1200
1000
800
600
400
200
0
15:36:49 17-May-98
15:36:59
15:37:09
15:37:19
15:37:29
15:37:39
15:37:49
15:37:59
15:38:09
15:38:19
15:38:29
15:38:39
15:38:49
15:38:59
15:39:09
15:39:19
15:39:29
15:39:39
« Utilization(%) and Broadcasts(%) »
Ethernet Statistics · 1 Second Averages
■ Broadcasts (%) ■ Utilization (%)
100
90
80
70
60
50
40
30
20
10
0

		General Stats					Errors (per sec)						Misc Stats			Selected Pr (per se		
Date	Time	Utilization (%)	Utilization (fps)	Utilization (kBps)	Bytes per frame	Active node count	Total Errors	Jabbers	Bad FCS's	Runts	Collisions & Jams	Misaligns	Broadcasts (fps)	Broadcasts (%)	Multicasts (fps)	IP DOD	DECnet IV	XNS IDP
17-May-98	15:36:49	6	114	73	641	5							8	7	10	24	53	24
	15:36:50	9	167	116	692	7	3	1	1	1			18	11	12	52	66	25
	15:36:51	11	186	135	729	5							22	12	10	66	71	22
	15:36:52	11	188	136	722	5							22	12	12	66	72	22
	15:36:53	11	326	142	436	7	3	1	1	1			58	18	[illegible]			
	15:36:54	11	385	133	345	5							75	19	[illegible]			
	15:36:55	11	387	133	344	5							75	19	[illegible]			
	15:36:56	26	615	322	524	7	3	1	1	1			183	30	[illegible]			

MISC STATS

Statistic Name	Ave	Max
Broadcasts (fps)	67.09	1849
Broadcasts (%)	18.04	96.96
Multicasts (fps)	10.57	12

LAN Summary Stats Charts

General Stats
☐ Util (%) and Active nodes/stations
☐ Util (KBps) and Bytes per frame
☐ Utilization (%) and Total Errors (%)
☐ Utilization (%) and Broadcasts (%)
☐ Broadcast and Multicast frames
☐ Token Ring Source Routing Dist.
☐ FDDI Ring State
☐ FDDI Token Rotation Time

Ethernet Errors
☐ Errors (3D Line)
☐ Errors (Stacked area)

Token-Ring & FDDI Errors
☐ Hard Errors (3D line)
☐ Hard Errors (Stacked area)
☐ Soft Errors (3D line)
☐ Soft Errors (Stacked area)

○ Select Errors

Selected Protocols (fps)
☐ 3D area
☐ 3D area w/other
☐ Stacked area
☐ Stacked area w/other
☐ Relative % w/other
☐ 3D pie w/other

○ Select

Selected Nodes/Stations (fps)
☐ 3D area
☐ 3D area w/other
☐ Stacked area
☐ Stacked area w/other
☐ Relative % w/other
☐ 3D pie w/other

○ Select

● Clear All ○ Ok ● Cancel

Figure 20: Internet Reporter Output

Interpret Results

Once you have collected baseline information and arranged it in easy to read reports, it is time to interpret the findings. The following are some things to look for when interpreting Ethernet network baseline information:

- Unusual traffic patterns like FTPs with mean packet size of 100 bytes. When a node sends a file using the File Transfer Protocol (FTP), it is more efficient to do it in large packets, instead of small ones. The largest Ethernet packet is 1518 bytes, much larger than 100 bytes.

- Sustained utilization greater than 35% indicates a congested network. While you can expect utilization to occasionally burst to 80% or 90%, this should not be a sustained condition. When the utilization is constantly high, the probability of collisions in the network increase resulting in undesirable retransmissions.

- Collision rates should not be greater than 5% of packet rate. For example, if the network is transmitting 100 packets per second, collisions should not occur more than 5 times per second. Collisions are part of the CSMA/CD access method (see Appendix A). They are necessary, but a large number indicates problems.

- Similar to the collision rate, the error rate should not be greater than 5% of the packet rate.

- Individual Ethernet segments should not contain more than 200 active stations because network demand will probably be high and the physical limits of the network can be easily reached.

- Broadcast and multicast traffic levels greater than about 20 pps (packets per second) are also a problem because a broadcast creates an interruption in the node CPU. This slows down the performance of the desktops in the network. The value 20 pps may be dependent upon the application.

- Top talkers should normally be routers and servers. Any other device in the network with high levels of utilization deserves a close look.

- Response time should vary less than 10% over time. If the response time varies more than this, there could be new users or applications affecting the network. Or there could be some physical problems or error in protocols.

- The protocol distribution should be in accordance with programmed protocols. If you have a Novell network which primarily uses IPX, and IP appears to be the most used protocol, encapsulation problems or improper configurations could be the problem.

Using the Advisor LAN, you can gather valuable information such as how your network looked before performance began to degrade or other failure modes began to appear. Proactive use of Advisor can save you time and money by helping to avoid problems before they start.

Virtual LAN Measurement

A virtual LAN (VLAN) is a logical broadcast domain spanning multiple physical segments. VLANs minimize backbone broadcast traffic because each Virtual LAN sends broadcast traffic to only its own domain. This helps to minimize CPU interrupts due to broadcasts. In addition, using VLANs can simplify modifications in the network. For example, a node in network A can easily be reassigned to network B with minimal effort and infrastructure change. Virtual LANs are port, MAC address, protocol, or rule based.

To make VLAN analysis easier, the Advisor LAN decodes Cisco's proprietary ISL (Inter-Switch Link Protocol), one of the most commonly used VLAN protocols. ISL identifies the VLAN association of each frame passed between switches by adding 30 bytes to the LAN frame consisting of a new header and CRC. Figure 21 shows VLAN statistics in the Protocol Vital Statistics view.

Note: the 802.1Q standard is a strict subset of Cisco's ISL capabilities. Both can coexist in the same network but cannot appear on the same port.

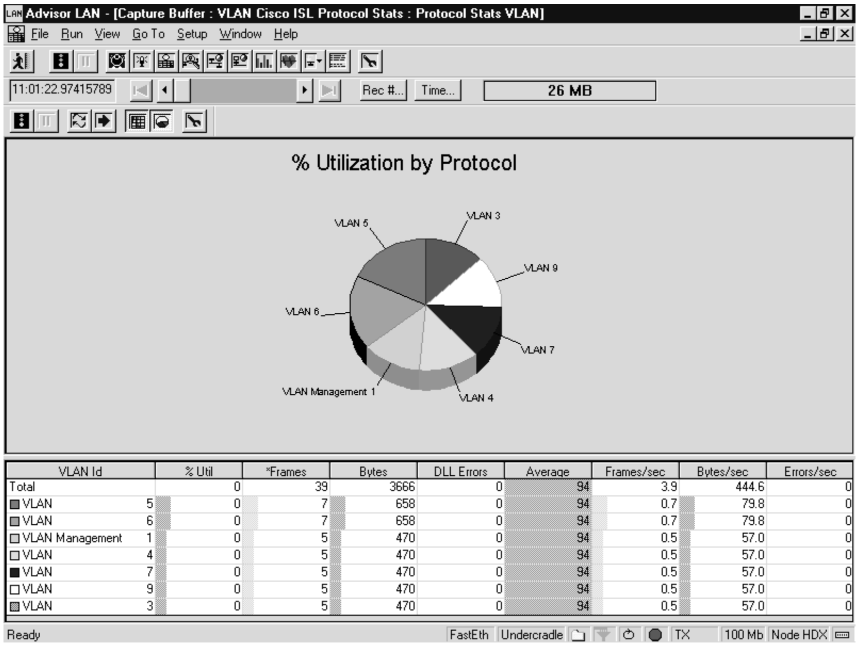
Advisor LAN - [Capture Buffer : VLAN Cisco ISL Protocol Stats : Protocol Stats VLAN]
File Run View Go To Setup Window Help
11:01:22.97415789
Rec #...
Time...
26 MB
% Utilization by Protocol
VLAN 5
VLAN 3
VLAN 9
VLAN 6
VLAN 7
VLAN Management 1
VLAN 4

VLAN Id		% Util	*Frames	Bytes	DLL Errors	Average	Frames/sec	Bytes/sec	Errors/sec
Total		0	39	3666	0	94	3.9	444.6	0
■ VLAN	5	0	7	658	0	94	0.7	79.8	0
■ VLAN	6	0	7	658	0	94	0.7	79.8	0
□ VLAN Management	1	0	5	470	0	94	0.5	57.0	0
□ VLAN	4	0	5	470	0	94	0.5	57.0	0
■ VLAN	7	0	5	470	0	94	0.5	57.0	0
□ VLAN	9	0	5	470	0	94	0.5	57.0	0
▨ VLAN	3	0	5	470	0	94	0.5	57.0	0

Ready
FastEth
Undercradle
TX
100 Mb
Node HDX

Figure 21: VLAN Statistics

Normally the best place to connect the Advisor for VLAN analysis is between two switches, because this is the link where the VLAN traffic makes sense and is also where most of the traffic is transmitted. Use the *Connections* section of this application note or the Advisor's online Help for specific connection information. Once you have connected the Advisor, open the Protocol Stats – VLAN measurement and start the analysis run. VLAN traffic statistics will be gathered and displayed, and VLAN frames will be automatically decoded and shown in the Advisor's Decode view. An example of decoded VLAN traffic is shown in Figure 22.

```
------- ISL Header -------
ISL: Destination Address = 01:00:0C:00:00
ISL: Encapsulated Frame Type = Ethernet (0)
ISL: User (Type Extension) = Normal Priority (0)
ISL: Source Address = 00:60:3E:97:14:01
ISL: Data Length (in bytes) = 76
ISL: AAAA03 Field = 0xAAAA03
ISL: High Order Bytes Of Source Address (HSA) = 0x01000C
ISL: Virtual LAN ID (VLAN) = 1
ISL: BPDU and CDP Indicator (BPDU) = 1
ISL: Switch Port Index (INDX) = 1
ISL: Reserved (RES) = 0x0
ISL: Encapsulated Packet Length = 60
ISL: CRC = 0x58CDCB38
ISL: Encapsulated Frame Follows :

------- MAC Header -------
MAC: Destination: 01-80-C2-00-00-00
MAC: Source: 00-60-3E-97-17-FD
MAC: Length: 38

------- LLC Header -------
LLC: Dsap: 0x42 (66)
LLC: Ssap: 0x42 (66) Command
LLC: Unnumbered frame: UI   (3)

------- BPDU Header -------
BPDU: Protocol Identifier = 0
BPDU: Protocol Version = 0
BPDU: BPDU Type = Configuration (0x00)
BPDU: Flags = 0x01
BPDU:      0... .... Not Topology Change Acknowledgment
BPDU:      .... ...1 Topology change flag
BPDU: Root Identifier = 0x800000603E971400
BPDU:      Priority = 32768
BPDU:      MAC Address = 0:60:3e:97:14:0
BPDU: Root Path Cost = 0
BPDU: Bridge Identifier = 0x800000603E971400
BPDU:      Priority = 32768
BPDU:      MAC Address = 0:60:3e:97:14:0
BPDU: Port Identifier = 0x8002
BPDU: Message Age = 0 (Seconds)
BPDU: Max Age = 20 (Seconds)
BPDU: Hello Time = 2 (Seconds)
BPDU: Forward Delay = 15 (Seconds)
01 00 0c 00 00 00 00 60    3e 97 14 01 00 4c aa aa
03 01 00 0c 00 03 00 01    00 00 01 80 c2 00 00 00
00 60 3e 97 17 fd 00 26    42 42 03 00 00 00 00 01
80 00 00 60 3e 97 14 00    00 00 00 00 80 00 00 60
3e 97 14 00 80 02 00 00    14 00 02 00 0f 00 00 00
00 00 00 00 00 00 ad 9e    33 05 58 cd cb 38
```

Figure 22: Decoded VLAN Traffic

Voice over IP Tests

The Agilent Advisor provides a comprehensive Voice Over IP (VoIP) decode suite to enable you to monitor and troubleshoot these new IP telephony protocols. One important set of protocols used for VoIP is that defined by the ITU-T H.323 recommendation. This section presents a VoIP troubleshooting example in which the Advisor's decodes play an important role.

An area of potential difficulty in VoIP implementation is multi-vendor interoperability. For example, during call setup and disconnect, a series of signaling messages involving both Q.931 and H.245 message types are transmitted between gateways. If these gateways are supplied by different equipment manufactures, it is possible that their H.323 implementations differ. Suppose, for example, that when Gateway A closes a call connected to Gateway B, everything works as expected, but when Gateway B closes a call connected to Gateway A, Gateway A malfunctions and requires a power cycle reboot.

To troubleshoot this problem, you connect the Advisor to the gateway link (as a node) such that you can monitor and decode the signaling packets sent between the two gateways. You will be using the Advisor's Decode view. You can customize the Decode view to look at a summary, detailed, or hexadecimal version of decoded packets, and if you know the IP addresses of the gateways, you can set up hardware filters to limit captured traffic to that transmitted between the two gateways.

Start monitoring traffic with the Advisor and then initiate several call and disconnect conditions in the network itself. Figure 23 shows an example of the kind of traffic you will capture.

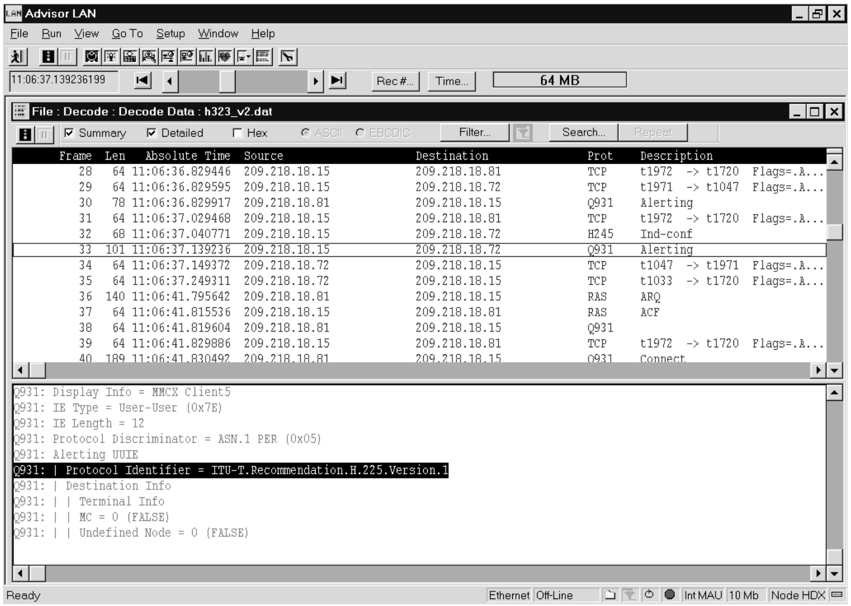

```
   Frame  Len  Absolute Time   Source          Destination     Prot  Description
      28   64 11:06:36.829446  209.218.18.15   209.218.18.81   TCP   t1972  -> t1720  Flags=.A...
      29   64 11:06:36.829595  209.218.18.15   209.218.18.72   TCP   t1971  -> t1047  Flags=.A...
      30   78 11:06:36.829917  209.218.18.81   209.218.18.15   Q931  Alerting
      31   64 11:06:37.029468  209.218.18.15   209.218.18.81   TCP   t1972  -> t1720  Flags=.A...
      32   68 11:06:37.040771  209.218.18.15   209.218.18.72   H245  Ind-conf
      33  101 11:06:37.139236  209.218.18.15   209.218.18.72   Q931  Alerting
      34   64 11:06:37.149372  209.218.18.72   209.218.18.15   TCP   t1047  -> t1971  Flags=.A...
      35   64 11:06:37.249311  209.218.18.72   209.218.18.15   TCP   t1033  -> t1720  Flags=.A...
      36  140 11:06:41.795642  209.218.18.81   209.218.18.15   RAS   ARQ
      37   64 11:06:41.815536  209.218.18.15   209.218.18.81   RAS   ACF
      38   64 11:06:41.819604  209.218.18.81   209.218.18.15   Q931
      39   64 11:06:41.829886  209.218.18.15   209.218.18.81   TCP   t1972  -> t1720  Flags=.A...
      40  189 11:06:41.830492  209.218.18.81   209.218.18.15   Q931  Connect
```

```
Q931: Display Info = MMCX Client5
Q931: IE Type = User-User (0x7E)
Q931: IE Length = 12
Q931: Protocol Discriminator = ASN.1 PER (0x05)
Q931: Alerting UUIE
Q931: | Protocol Identifier = ITU-T.Recommendation.H.225.Version.1
Q931: | Destination Info
Q931: | | Terminal Info
Q931: | | MC = 0 (FALSE)
Q931: | | Undefined Node = 0 (FALSE)
```

Figure 23: Decoded H.323 Traffic

While the analyzer is running, you can decode the information in real-time or you can stop the analyzer and perform post-process analysis on traffic held in the capture buffer. The Advisor's online Help provides detailed information on how to perform both real-time and post-process decoding of network traffic.

By following the signaling conversation between the two gateways, you find that gateway A closes its H.245 conversation before sending Q.931 disconnect instructions. However, gateway B never closes its H.245 conversation, leaving the logical channel open. Because gateway A expects something else, its operations are disrupted. By decoding and filtering both directions of the conversation with the Advisor, the difference between the two methods of handling the H.323 protocol suite can be exposed providing the data necessary for it to be resolved.

For more in depth information about Voice over IP tests with the Advisor, please refer to the *Advisor Troubleshooting VoIP Signaling* Application Note and other related VoIP documentation. Your Agilent Technologies representative can give you more information.

Appendix A: Ethernet Collision Specification

Collisions are a normal part of Ethernet transmission and occur when two stations attempt to transmit at the same time, causing their signals to collide with each other on the physical media and become garbled. When this happens and the transmitting station detects the collision, it stops transmitting data and sends out a jamming signal, four bytes in length. The jamming signal ensures that all other stations on the network detect the collision. All stations that have been transmitting cease transmission, wait a period of time, and, if the carrier is free, attempt retransmission of the frame.

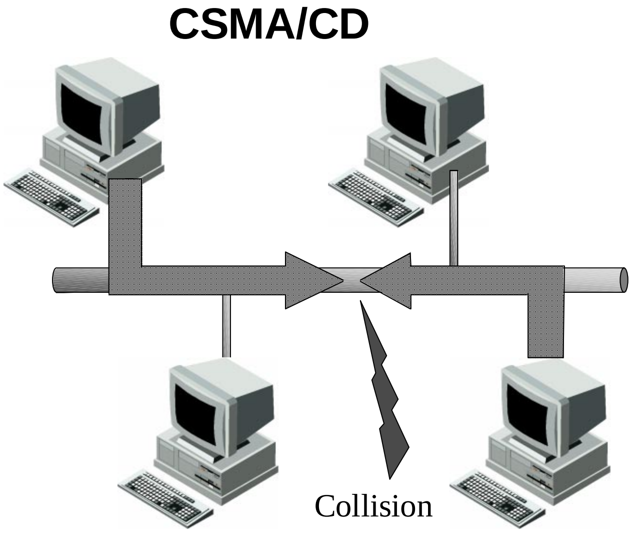

Figure 24: Ethernet Collisions

If all stations waited the same length of time before checking the carrier and starting transmission again, another collision would occur. To avoid this, each station generates a random number that determines the length of time it must wait before testing the carrier. This time period is known as the station's backoff delay. Backoff delay is calculated in terms of multiples of slot time - that is, the time it takes a signal to travel from one end of the network to the other plus some 'backoff' value. If the segment length is within specification, this time should be less than or equal to 51.2 milliseconds. The number of time slots each station must wait is determined by the random number (r) generated at each station. The smaller the range of values from which the random number is selected, the greater the likelihood that two stations will select the same number and have another collision. On the other hand, if the number is large, all the stations may wait for several time slots before any station transmits, causing transmission time to be wasted and reducing overall network response time.

To achieve a balance between these two considerations, the CSMA/CD standard uses an approach known as binary exponential backoff. The range of numbers is defined as $0 \leq r \leq 2^n$, where n reflects the number of retransmission attempts that the station has made. For the first ten attempts, n ranges from 1-10. For subsequent attempts, n continues to have a value of 10. This means that for the first attempt at retransmission, the range is 0-1; for the second attempt, 0-3; for the third, 0-7; and so on. If repeated collisions occur, the range continues to expand until the station successfully transmits without collision. If a station is unsuccessful in transmitting after 16 attempts, it reports an error condition. Binary exponential backoff results in minimum delays before retransmission when traffic on the network is light. When traffic is high, repeated collisions will cause the range of number to increase, thus lessening the chance of further collisions. Of course, when traffic is extremely high, repeated collisions could begin to generate error conditions.

Appendix B: The Internet Control Message Protocol (ICMP) system

To allow machines on the Internet to report errors or provide information about unexpected circumstances, protocol designers added a special-purpose message mechanism to the Internet Protocol (IP). The mechanism, known as the *Internet Control Message Protocol* (ICMP), is considered a required part of IP and must be included in every IP implementation. ICMP messages travel across the Internet in the data portion of the IP datagrams like all other traffic. The ultimate destination of an ICMP message is not a user process on the destination machine, but the Internet software on that machine. That is, when an ICMP error message arrives, the IP software module handles the problem itself - it does not pass the ICMP message to the application program whose datagram caused the problem. Initially designed to allow gateways to report the cause of delivery errors to hosts, ICMP is not restricted to gateways. An arbitrary machine can send an ICMP message to any other machine.

Although each ICMP message has its own format, they all begin with three fields:

- An 8-bit integer *message type* field.
- An 8-bit *code* field that provides further information about the message type.
- A 16-bit *checksum* field.

In addition, ICMP messages that report errors always include the Internet header and the first 64 data bits of the datagram causing the problem. The reason for returning more than the datagram header alone is to allow the receiver to determine more precisely which protocol(s) were used and which application was responsible for the datagram.

When a gateway or network station cannot deliver an IP datagram, it sends a destination unreachable message back to the original source. Destinations may be unreachable because:

- The hardware is temporarily out of service.
- The sender specified a nonexistent destination address.
- The gateway does not have a route to the destination address (in rare circumstances).

www.agilent.com

Note:
Subnet address is an extension of the Internet addressing scheme that allows a site to use a single Internet address for multiple physical networks. Outside of the sire using the subnet address, routing continues as usual by dividing the destination address into an Internet portion and local portion. Gateways and hosts inside a site using subnet addressing interpret the local portion of the address by dividing it into a physical network and a host portion.

Internet address is the 32-bit address assigned to hosts that want to participate in the Internet using TCP/IP. Internet addresses are the abstraction of physical hardware addresses, just as the Internet is an abstraction of physical networks. Actually assigned to the interconnection of a host to a physical network, and Internet address consists of a network portion and a host portion. This partitioning makes routing efficient.

There are numerous sources available for detailed information about IP and its addressing methods.

By internet, phone or fax, get assistance with all your Test and Measurement needs.

Online assistance:
http://www.agilent.com/find/assist

United States:
(Tel) 1 800 452 4844

Canada:
(Tel) 1 877 894 4414
(Fax) (905) 282 6495

China:
(Tel) 800-810-0189
(Fax) 1-0800-650-0121

Europe:
(Tel) (31 20) 547 2323
(Fax) (31 20) 547 2390

Japan:
(Tel) (81) 426 56 7832
(Fax) (81) 426 56 7840

Korea:
(Tel) (82-2) 2004-5004
(Fax) (82-2) 2004-5115

Latin America:
(Tel) (305) 269 7500
(Fax) (305) 269 7599

Taiwan:
(Tel) 080-004-7866
(Fax) (886-2) 2545-6723

Other Asia Pacific Countries:
(Tel) (65) 375-8100
(Fax) (65) 836-0252

Product specifications and descriptions in this document subject to change without notice.

Printed in U.S.A. October 8, 2001

5968-8428E

Use this link to go directly to our network troubleshooting solutions:
http://www.agilent.com/comms/onenetworks

Agilent Technologies