*Excerpts taken from:*

# Network Troubleshooting
**By Othmar Kyas**

An Agilent Technologies Publication

## Section II

## Troubleshooting Local-Area-Networks

### *Chapter 7*

## 10/100/1,000 Mbit/s Ethernet

## 7.2 Troubleshooting in 10/100/1,000 Mbit/s Ethernet

### 7.2.1 Gathering Information on Symptoms and Recent Changes

The first step in any troubleshooting process is to gather information. The more information you have about the symptoms and characteristics of a problem—including *when* it first occurred—the better your chances of solving the problem quickly and efficiently. Typical questions you might ask at this stage include:

- Do the symptoms occur regularly or intermittently?
- Are the symptoms related to certain applications (running simultaneously with), or do they affect all network operations?
- Are similar applications malfunctioning?
- How many users are involved?
- Do the symptoms correlate to other activities in the network?
- When was the first occurrence of the symptom?
- Was there any change in any hardware or software network component?
- Has anyone connected or disconnected a PC (laptop or desktop) or any other component to or from the network?
- Has anyone installed an interface card in a computer?
- Has anyone stepped on a cable?
- Has any maintenance work been performed in the building recently (by a telephone company or building maintenance personnel, for example)?
- Has anyone (including cleaning personnel) moved any equipment or furniture?

In general, it is advised to restart a device immediately after a change has been made to it. Otherwise a problem resulting from a change may come into effect at a later time, when another user starts the device.

### 7.2.2 Starting the Troubleshooting Procedure

Troubleshooting in Ethernet LANs is primarily performed using protocol analyzers, network management software, and cable testers. If the network is still up and running, the first step in the troubleshooting procedure involves using a protocol analyzer to determine the main operating statistics of the network. These statistics include capacity use in bytes and as a percentage, packet throughput per second, the collision rate, the packet length distribution, FCS errors, the proportions of broadcast and multicast packets, the numbers of runts and jabbers and the number of transmitting stations.

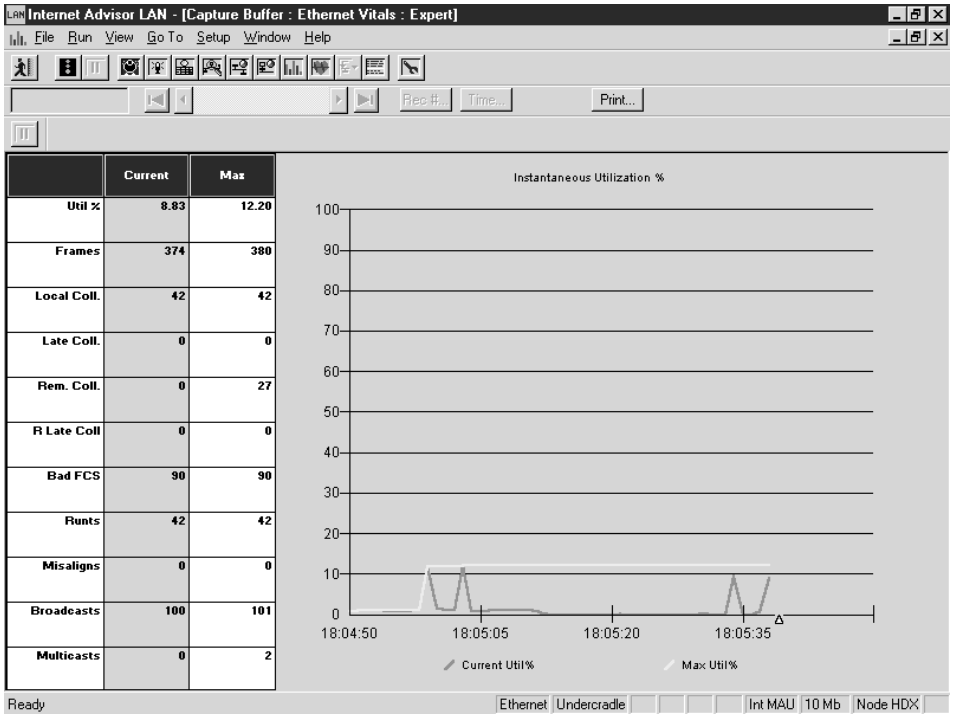


*Figure 7.14 Display of characteristic operating statistics using a protocol analyzer*

In many cases, the analysis of operating statistics leads directly to the source of the problem. However, because Ethernet networks use a bus topology, many problem symptoms that are detected in one segment can originate in an adjacent segment or in more remote parts of the network. If the problem cannot be localized, longer-term measurements are necessary. Characteristic operating statistics are recorded systematically over time and analyzed for correlated parameters. For example, if a correlation is observed between the number of active stations, the number of collisions and the network capacity use, it is highly likely that performance problems are related to the active components (network nodes, repeaters, etc.) in the network. If there is no such correlation, it is more likely that the source of the problem lies in the cabling infrastructure.

In many cases, such measurements must be taken in several network segments concurrently (using LAN probes, for example). When the results are displayed in a graph, an analysis of any temporal correlations can indicate whether events in a given segment caused symptoms in others or vice versa. This measurement and analysis procedure is repeated until the range of possible problem sources can be limited to a small area.

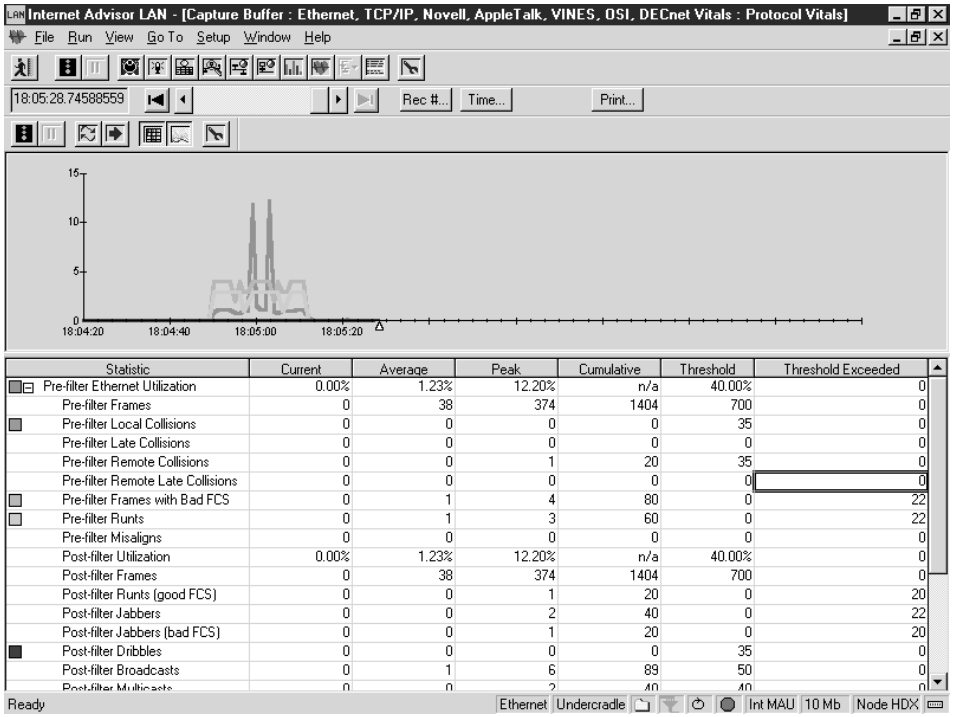


*Figure 7.15 Analysis of trend measurements and correlations*

What steps to take after a protocol analyzer performs the basic measurements depends on the nature of the symptoms. If the symptoms can be localized or occur periodically, or can at least be reproduced, then the troubleshooting process continues with the network component nearest the problem. If the problem source cannot be detected there, the range of analysis is successively expanded. For example, if the problems are found to be related to a single network node, the next step is to analyze the station's software and hardware components. If no fault is found, the examination progresses to the station's MAU and AUI cable, its power cable connector, the wall jack, the cable to the hub, the hub itself, the cable to the server, and so on. If the problem cannot be localized at all, or if problems that were thought to have been localized cannot be pinpointed, the only way to find the source of the problem is through systematic segmentation of the network. The first thing to do is to refer to the network documentation to identify how many hubs, switches or wire centers are involved. If a small number of hubs are involved, it makes sense to isolate one at a time to see if the problem disappears. If the segment comprises a large number of (cascaded) hubs, the network needs to be cut off physically into two (or more)

parts and monitored. The portion of the network found to produce the problem is then segmented again, until the segment containing the source of the trouble is found. Theoretically, this procedure can lead to the source of the problem very quickly—but only with some luck in the selection of segmentation points. In practice, this drastic method causes considerable disruption in network opera-

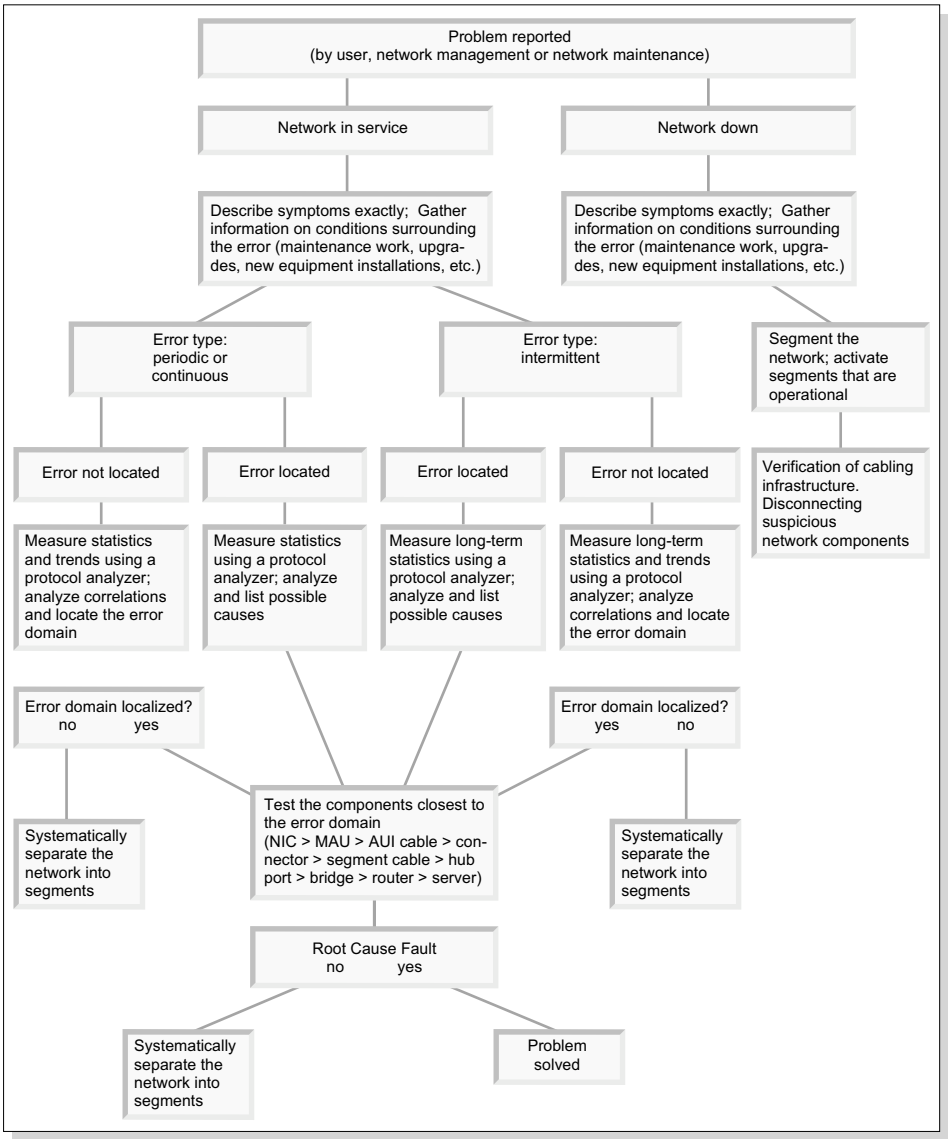


*Figure 7.16 Systematic troubleshooting in data networks*

tion, and is therefore only applied as a last resort when the problem itself severely impairs normal network operation.

If the symptoms occur intermittently, long-term measurements are necessary. These must be performed continuously until the basic network operating statistics have been measured during the occurrence of the fault. This information usually provides the first clue to the source of the error. You can also try defining filters and alarms for Ethernet frame errors: if you configure your protocol analyzer accordingly, you may be able to capture the Ethernet packets that are transmitted in the error situation. Furthermore, it is essential to log the exact time of each error event. Later this information can be used to find correlations with other events in the network or on a given node, such as backups, the start of specific applications, connections through routers, access to the Internet, working hours of neighboring departments (for example, an error might occur only when employees in the neighboring department boot their network nodes between 7:45 and 8:00 a.m.), etc. If this does not help to track down the error, you may have to resort to the segmentation method. Depending on which causes the least inconvenience to users, you can either systematically disable network functions and applications, or physically separate the segments. These methods usually leads to the error source.

Another simple, but sometimes effective troubleshooting method if nothing else is working, is the compare and analyze method. The principle idea is to compare the network element in question to similar or identical network elements that are working. However, in many cases this method has limited use because in today's networks routers and servers are rarely truly identical.

### 7.2.3 Error Symptoms in Ethernet

Many of the problem symptoms in Ethernet networks occur in normal network operation and are not necessarily due to errors. These conditions include high traffic loads, high collision rates due to large numbers of active stations per segment, or a high proportion of broadcasts due to the use of certain protocols, such as Windows NT resource sharing. However, the occurrence of certain types of defective packets can point directly to problems that have nothing to do with normal operating conditions. These defective packets are divided into the following categories: fragments from local, remote, or late collisions; undersized packets; jabbers (oversized packets); FCS errors; alignment errors; noise errors; and length errors.

#### Collisions

Collisions result when two stations on the same segment begin transmitting Ethernet frames at the same time. In Ethernet LANs operating in half-duplex

mode, the occurrence of collisions is perfectly normal. If the distance limitations are respected, it takes 25.51 µs in the worst case for the first bit of a 10 Mbit/s Ethernet packet to travel to the furthest node and collide with the frame that node is beginning to send. In the collision, the voltages of the two signals are added together and the resulting collision signal takes a maximum of 25.51 µs to return to the original sender. Thus the longest time it can take a station to recognize the occurrence of a collision and stop transmission is 51.2 µs or 512 bit times, which is equal to 64 bytes. In other words, normal collision fragments do not exceed a length of 64 bytes. Another characteristic of collision fragments is that the frame contents are completely destroyed so that no checksum is detectable.

#### *Remote and Local Collisions*

Collisions that occur in a local segment are called local collisions. These collisions happen within a cable section without a repeater unit in 10Base2 networks, or between hub and node in 10Base-T networks. These collision fragments are less than 64 bytes long, have a signal voltage higher than that of a normal signal, and carry an invalid FCS.

Remote collisions are collision fragments that are exported to other segments by repeaters or hubs. Their signal voltages are not higher than normal because they are corrected by the repeater hardware. In most cases, this type of collision ends with a sequence of jam bits. As soon as a collision is detected on one port, repeaters and hubs emit jam bits on all ports in order to prevent further collisions by stopping transmission in the connected segments.

#### *Late Collisions*

Late collision fragments are longer than 64 bytes. They can be recognized as collision fragments by the increased signal voltage (in local late collisions) or by the lack of a checksum and the presence of jam bits (in remote late collisions). Late collisions can be caused by defective interface hardware or by cable segments that exceed distance limitations. Because the transmitting station does not detect the collision (the packet had already been emitted in its entirety before the collision occurred), the consequences of this type of collision can be doubly harmful: retransmission is requested only by the higher-layer protocols, which can have adverse effects on application performance.

### Short Packets

A short packet is less than 64 bytes long, but carries a valid FCS. Faulty packets most often are the cause. In rare cases, active network components such as bridges may use such short packets to exchange data. Short packets are also referred to as runts with valid FCSs.

### Runts

'Runt' is not an officially defined term, but is a widely-used expression for a short, invalid packet. The term generally refers to any frame that is shorter than 64 bytes, and thus may indicate a local or remote collision fragment or a short frame with or without a valid FCS.

### Jabbers

According to IEEE 802.3, a jabber is any frame longer than 1,518 bytes, whether its FCS is valid or not. The most common causes of jabber frames include faulty interface cards, defective cabling, and grounding problems.

### Long Packets

Long packets are longer than 1,518 bytes, but carry a valid FCS. Long packets can be caused by faulty interface card drivers or by software problems in bridges or routers.

### FCS Errors

Frames that contain invalid frame check sequences usually indicate a faulty network interface card. In most cases the headers of these frames are readable, so that the station that transmitted them can be identified by capturing and decoding the packets with a protocol analyzer. FCS errors can also be caused by defective hub ports or by induction voltage in the cabling.

### Alignment Errors

Ethernet frames that do not end on an octet boundary (the number of bits in such a packet is not an integer multiple of 8) are called alignment error frames. This type of error is typically caused by defective software drivers or by collision fragments on the segment.

### Noise Errors

Noise errors are caused by voltage induced in the network cables by external sources, leading other nodes in the network to assume that the network medium is busy transmitting data. Repeaters and hubs sometimes even exacerbate the symptoms. Symptoms of noise errors (also known as "ghosts") are low capacity use combined with extremely low network performance. Typical causes of noise errors are ground loops, a MAU power supply leaking voltage into the cabling, or other wiring problems.

### Length Errors

A length error means that the length field of a given Ethernet frame does not match the actual length of its data field.

### 7.2.4 Cabling Problems

Problems with cabling are still very common in Ethernet networks. Typical causes include defective or low-quality cables, incorrect characteristic impedance, defective terminating resistors, wiring mistakes or electromagnetic interference (noise). These types of problems are discussed in detail in the section on cabling.

### 7.2.5 Problems with Ethernet Interface Cards

Two types of Network Interface Cards (NICs) are used in Ethernet networks: cards with integrated MAUs (transceivers), which are connected directly to the communication medium by BNC, RJ-45 or fiber-optic subscriber connectors (SC); and cards with external MAUs. In the latter case, a MAU is chosen to fit the medium (coaxial cable, twisted pair, or fiber) and connected to the network interface card by an AUI cable. The interface card is configured using the manufacturer's software or by means of hardware switches on the card. A number of symptoms, such as invalid FCSs, jabber packets or late collisions, are often caused by faulty interface cards.

The first step in localizing a faulty interface card is to identify suspicious nodes on the network. Begin by making a list of all network nodes that transmit defective packets. Most protocol analyzers provide this information using fully automatic test programs. If the source addresses of the defective packets are invalid and cannot be decoded, try the correlation method: begin by simultaneously charting the activity of the suspicious nodes and the number of defective packets on the segment. If you observe a correlation between the activity of a certain node and the number of defective packets, then you have probably found the faulty interface card. If this strategy also fails, the only alternative is to employ the segmentation method described previously, systematically disconnecting network segments (hubs) until the error ceases to occur.

Once suspicious nodes have been isolated, they are tested individually. Interface cards with AUI connectors are tested using an adapter installed between the AUI connector and the transceiver. The test adapter indicates the status of the AUI signal lines, the signal quality error (SQE) signal and the transceiver power voltage. If the packets sent do not activate the corresponding LED on the AUI test adapter, then the interface card is not configured correctly or is defective. The next step is to check all software settings and hardware switches. If the configuration is correct, the card must be replaced. If the card has an integrated MAU, test the interface card in mini-networks connected to the interface card, rather than use your actual network. In 10Base2 Ethernet, a mini-network consists of no more than a T-connector with two 50Ω terminating resistors. In

10/100/1,000Base-T, the minimum network is the computer with its NIC and a mini-hub. Transmit simple loopback packets (IP loopback to address 127.0.0.1, Ethernet Configuration Test Packets (CTPs), IEEE 802.2 LLC Test Packet, 802.2 LLC XID Test Packet) while monitoring the activity on the mini-network with a protocol analyzer. If the analyzer detects no defective packets, then the card is probably defective.

#### Symptoms and Causes of Faulty Ethernet Interface Cards

Characteristic symptoms of faulty Ethernet interface cards include high collision rates, jabber packets, packets with invalid FCSs, and intermittent connection problems in individual network nodes. The most common problem sources are faulty MAUs, blown fuses on network interface cards, incorrectly configured card settings or card failures due to defective components.

##### *Carrier Sense Failures*

Jabber packets and high collision rates are most often caused by faulty MAUs. If the MAU is defective, then the card's carrier sense function does not work, and the station transmits packets regardless of traffic sent by other stations on the segment.

##### *Configuration Errors*

If the station cannot transmit or receive packets at all, the first thing to check is the card configuration. Common configuration errors are:

- The wrong port on the card is activated (the RJ-45 connector instead of the AUI connector, or vice versa).
- The interrupt level configured on the card is already in use by another device.

##### *MAU (Transceiver) Has No Power*

If a card is connected to the network by an external transceiver, use an AUI test connector to check for correct interaction between card and transceiver. If the card does not provide proper voltage to the transceiver, check whether the card's fuse (if it has one) has blown. If the fuse is intact, the card must be replaced.

##### *Illegal Coaxial Cable Connection*

In small 10Base2 networks, inexperienced operators sometimes attach an interconnection cable between the Ethernet bus T-connector and the station, rather than connecting the station's network card directly to the T-connector. This does not work. The interface card must be connected directly to the Ethernet bus cable. The maximum length of the connection between the T-connector and the interface card, according to the IEEE standard, is 4 cm.

## 7.2.6 Problems with Media Access Units (MAUs)

External MAUs allow more freedom of location when connecting nodes to the network, for example, when the Ethernet cable is several meters away from the station. The 15-pin AUI interface connector on the interface card is always the same, no matter what type of cabling is used (optical fiber, twisted pair, thin coaxial, thick coaxial). Multiport transceivers can be used to connect more than one node to the network at a single physical access point. Because MAUs (transceivers) and AUI cables provide the actual connection between the network node and the network, it is essential that they function correctly. Today's MAUs and AUI cables are suitable for both Ethernet v2.0 and 10/100 Mbit/s IEEE 802.3 networks. They have LEDs to indicate collisions, data transmission and signal quality error (SQE) signals. The SQE is sent by the MAU to the interface card when non-standard signals, such as collisions, are received. However, not all interface cards use the SQE in accordance with the standard. Older Ethernet 2.0 cards, for example, require an SQE signal (heartbeat) after every "read" or "write" operation. MAUs that are connected to such interface cards must be operated in the heartbeat mode. There is usually a switch located on the housing of the MAU to change its operating mode.

#### Connecting and Removing AUI Cables and MAUs

When a MAU is connected to a live network node, the computer's power supply may cause voltage spikes. This can cause loss of data on the hard disk or damage to the hardware itself. For this reason it is important to switch off network nodes before connecting or removing MAUs.

#### Equipment for Troubleshooting

The following equipment is required to test MAUs and AUI cabling:

- A mini-network of the appropriate topology (a terminated T-connector for 10Base2 and 10Base5; mini-hubs for 10/100/1,000Base-T)
- A fully functional MAU
- An AUI test connector with LEDs
- A functional AUI cable
- A multimeter

#### Symptoms and Causes of Faulty MAUs

A faulty MAU usually results in the complete breakdown of a network node's connection to the rest of the network, or in a significantly increased collision rate accompanied by jabber frames. The most likely causes are described here.

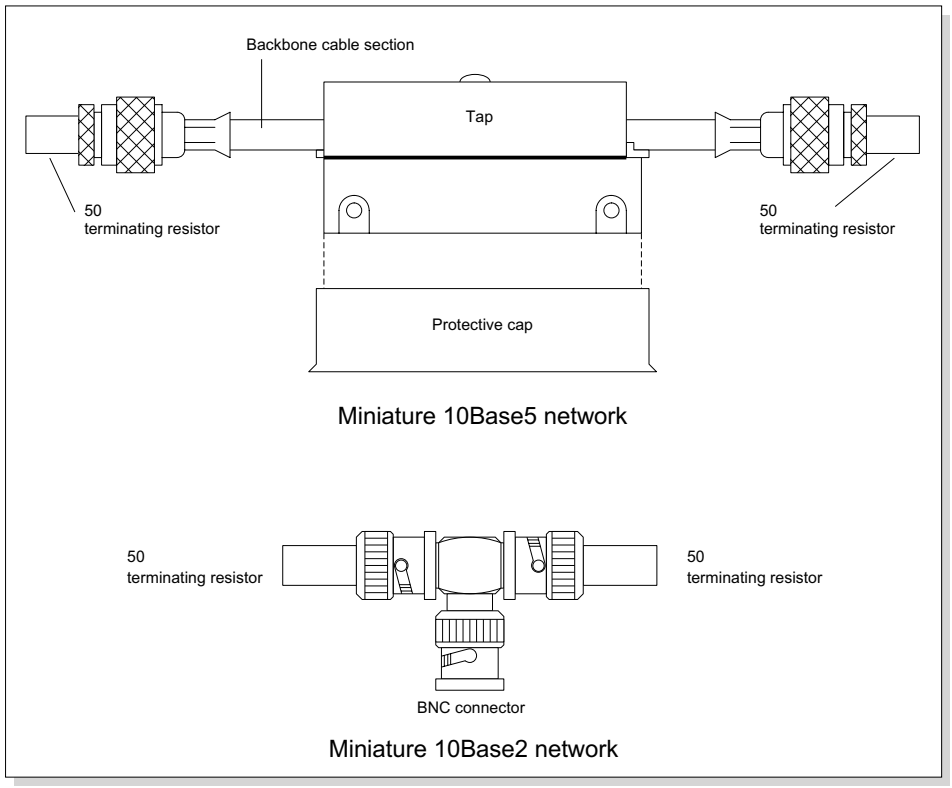


*Figure 7.17 Mini-networks for 10Base2 and 10Base5 networks*

### *IEEE 802.3 Instead of Ethernet v2.0 AUI Cabling*

Using IEEE 802.3-compliant AUI cabling together with Ethernet 2.0-compliant MAUs and interface cards can cause severe problems. According to the IEEE 802.3 standard, pin 4 of the AUI interface is grounded and pin 1 is assigned to the signal-in line. In Ethernet 2.0, these assignments are reversed: pin 4 is defined as signal-in and pin 1 is grounded (see Figure 7.18).

### *Signal Quality Error Problems*

As discussed earlier, MAUs can be operated in either heartbeat or SQE mode. If a MAU in heartbeat mode is connected to an IEEE 802.3-compliant interface card, it acknowledges every “read” and “write” operation with an SQE signal. The card interprets this as a collision, however, and attempts to retransmit data packets.

*Loose Connections*

Loose connectors often cause intermittent problems. Checking connectors, however, as simple as it may sound, can be difficult or even impossible in some cases. If the AUI cable is not accessible (if it is located under flooring or above ceiling panels, for instance) you may not be able to test the plug connection directly. In this case, proceed directly to loopback tests. In 10Base2 and 10Base5 networks, you should check the MAU's connection as well as the terminating resistor and its grounding. It is not altogether unusual to find that a MAU has been mistakenly connected to the end of the cable in place of a terminating resistor. In 10/100/1,000Base-T networks, you should also inspect the grounding of the cable shielding and of the wall jacks.

If the problem persists after you have checked all plug connections, you can systematically localize the error source by disconnecting the MAU from the live

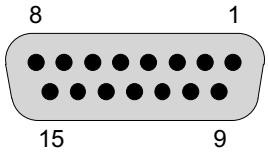


| Pin number | Signal abbreviation | Signal name |
|---|---|---|
| 3 | DO-A | Data Out A |
| 10 | DO-B | Data Out B |
| 11 | DO-S | Data Out Shield |
| 5 | DI-A | Data In A |
| 12 | DI-B | Data In B |
| 4 | DI-S | Data In Shield |
| 7 | CO-A | Control Out A |
| 15 | CO-B | Control Out B |
| 8 | CO-S | Control Out Shield |
| 2 | Ci-A | Control In A |
| 9 | CI-B | Control In B |
| 1 | Ci-S | Control In Shield |
| 6 | Vc | Voltage Common |
| 13 | VP | Voltage Plus |
| 14 | VS | Voltage Shield |
| Housing | PG | Protective Ground |

*Figure 7.18 Pin layout of media access units (MAUs)*

network and connecting it to a mini-network. If loopback packets can now be sent without problems, then the network node, the AUI cable and the MAU are not defective and the problem source must be in the live network. If the loopback test fails, replace the MAU, the AUI cable, and the interface card successively until the error is eliminated. To determine whether the MAU or the interface card is the defective device, you can measure the voltage the card supplies to the MAU between pins 13 and 6. It should be between 11.28 and 15.75 volts. Some cards have built-in fuses to prevent hardware damage due to voltage spikes. If this fuse blows, the power supply to the MAU fails and the node loses its connection to the network.

## 7.2.7 Problems with Repeaters and Hubs

Repeaters are used to connect Ethernet segments and to overcome distance limitations within segments. Repeaters that connect 10/100/1,000Base-T nodes in star topology are called hubs or multiport repeaters; their basic functions, however, are identical to those of repeaters as defined in IEEE 802.3. Today's repeaters are modular and capable of connecting segments of different transmission media (coaxial, twisted pair, and fiber optic). The maximum number of repeaters within the transmission path of a segment is limited : the transmission delay increases with each additional repeater and each additional cable segment, which increases the slot time. Longer signal transmission times also mean a higher probability of collisions and thus have a negative effect on network efficiency.

#### Repeater Functions

Repeaters that conform to the IEEE 802.3 standard fulfill the following functions:

- Regeneration of signal amplitude and removal of phase jitter in the pulse flanks
- Temporary storage of frames so that signals can be forwarded with correct timing
- Carrier sensing (checking for transmission media activity before transmitting signals)
- Replacement of missing bits in the preamble. Each Ethernet frame starts with a 62-bit preamble made up of a sequence of 0s and 1s (1010 1010 1010 1010), which is required by the receiving nodes to synchronize on the data stream. Because today's receiver chips need approximately 30 bit times to synchronize, and some bits get lost in the transmitter electronics on the NIC or on the medium, each repeater needs to replace missing bits in the

preamble. Otherwise the preamble would disappear by the time it passed through the second repeater.

- Extending collision fragments with jam bits to a length of 96 bits. This ensures clean collision detection in the connected segments.
- Collision detection and generation of jam bits. Once a repeater detects a collision it must transmit jam bits to all other repeater ports to prevent additional transmissions in the remaining segments. The jam signal must have a minimum length of 62 bits and consist of a 101010 sequence, of which the first bit must be a 1. The transmission of the jam signal must occur within 5 bit times after the end of the collision.
- Automatic partitioning. If high error rates occur in one segment, modern repeaters interrupt transmission and deactivate the corresponding port. The trigger for auto-partitioning of a port is 30 successive collisions or abnormally long collisions. After receiving 500 collision-free bits, the repeater reactivates the port.

The transmission delay introduced by 10 Mbit/s Ethernet repeater ports must not exceed the values listed in Figure 7.19.

| Ethernet type | Input + output delay | Total delay in bit times |
|---|---|---|
| 10Base5 | 6.5 + 3.5 | 14 |
| 10Base2 | 6.5 + 3.5 | 14 |
| FOIRL | 3.5 + 3.5 | 11 |
| 10Base-T | 8 + 5 | 17 |
| 10Base-FP | 3 + 4 | 11 |
| 10Base-FB | 2 + 2 | 4 |
| 10Base-FL | 5 + 5 | 11 |

*Figure 7.19 The maximum latency for valid signals on repeater ports*

Repeaters cannot detect MAC-layer errors such as invalid FCSs or illegal frame sizes.

### Diagnosing Repeater Problems

#### *Localizing Collision Domains*

Once a repeater detects a collision, it starts transmitting jam bits to all ports. A collision that occurs in one segment is visible as a jammed fragment or a remote collision in all other segments connected to the repeater. Remote collision packets are thus nothing but local collision packets exported by the repeater or hub of the segment.

If concurrent measurements in all segments show that the occurrence of remote collisions correlates with local collisions in another segment, then the source of the remote collisions is easily found. However, this strategy is only practical in 10Base2 and 10Base5 segments because in 10Base-T networks every node has a segment to itself. In 10/100/1,000Base-T networks, collision sources must be identified by correlating the network load due to the most active stations with collision rates. If the collision history seems to correlate with the packet output of a certain station, then that node could be the source of the collision problem.

*MAU Jabber Lockup Due to No Auto-Partitioning*

The MAU jabber lockup mechanism provides another important symptom of repeater problems. If no segment is connected to an open repeater port, the repeater transmits 5 ms of jam bits, alternating with 0.01 ms transmission gaps. The gaps are necessary to avoid activating the jabber lockup function of the interface cards. If the receiving electronics of an interface card detect jabbering that lasts longer than 5 ms, the card starts transmitting jam bits itself in order to stop the "jabber station". If jabber packets occur with a length of 6,250 bytes (equal to a 5 ms frame in 10 Mbit/s Ethernet) and a 10 µs gap (equal to 100 bits in 10 Mbit/s Ethernet), the cause is likely to be an open repeater port or a segment without a proper terminating resistor. In these cases, modern repeaters activate auto-partitioning and disable the port.

*Phantom Addresses*

Quite frequently the analysis of Ethernet frames using a protocol analyzer shows source and destination addresses with the following content:

*55555555* or

*AAAAAAAA*

Defective packets often contain these types of addresses. They appear at first glance to indicate a mysterious transmission error, but on closer examination of the binary representation of the hex format they are seen to be nothing but jam bits:

*5555 5555* is equivalent to

*1010 1010 1010 1010 1010 1010 1010 1010* in binary format and

*AAAA AAAA* is equal to

*0101 0101 0101 0101 0101 0101 0101 0101.*

These are the jam sequences sent by a repeater to all connected segments once a collision is detected on one of the ports. (Remember that a jam signal consists of alternating 0s and 1s.)

### Hub (Repeater) Problems: Symptoms and Causes

Typical symptoms of hub problems are low network performance, intermittent connection losses, high collision rates or the occurrence of jabber frames. The most likely causes are described here.

#### *Packets Missing Due to a Short Inter-Frame Gap*

If an interface card does not maintain the required minimum inter-frame spacing (9.6 µs in 10 Mbit/s, 0,96 Mbit/s Ethernet) between successive frames, slow repeaters cannot forward the frames fast enough and packets are lost. The data loss is detected in the higher-layer protocols, causing retransmission requests. This can result in significantly lower application performance for connections across one or more routers.

The inter-frame gap may be too short for any of several reasons. Some interfaces will transmit a data packet very soon after the occurrence of a collision, but otherwise perform in conformity with the 9.6 µs rule. Others stick very close to or drop below the 9.6 µs spacing. Packet fragments also can cause problems. Each hub has to extend fragments to the minimum length of 96 bits (including preamble). If a 64-bit fragment is received, the hub extends it to 96 bits by adding 32 jam bits, while at the same time reducing the gap between it and the following frame by 32 bit times or 3.2 µs. This can make the inter-frame gaps too short. However, modern hubs should be able to maintain constant inter-frame gaps by the use of buffers, whether a given packet needs to be extended or not. The symptoms of such hub problems are often only recognized at the application level. Typical cases are file transfers that take significantly longer between two stations that communicate through hubs than between stations located in the same segment. If all TCP packets are present with the correct sequence numbers before a repeater, for example, but sequence numbers are missing after the hub, then the hub is the cause of the problem. The file transfer throughput may then be only a few hundred Kbit/s instead of Mbit/s because the lost packets have to be retransmitted.

#### *10Base2 Repeater Grounding Problems*

Segments of 10Base2 or 10Base5 networks must be grounded only at one point. If a grounded repeater and a grounded terminating resistor form a current loop with the network cable, and if the two grounds have different electrical potentials, current will flow along the segment cable, resulting in severe interference, high collision rates and often a complete breakdown of network operation.

#### *Problems Due to an Excessive Number of Repeaters in the Transmission Path*

If a frame has to travel not only long distances but also across several repeaters, the maximum transmission delay (25.6 µs) may be exceeded because every

repeater adds latency. However, the higher the delay in an Ethernet network, the less efficient the CSMA/CD algorithm becomes. Collisions result, often in the form of late collisions, as well as reduced network performance and intermittent connection losses. As mentioned previously, the maximum number of repeaters per segment is limited to four.

#### *Installation and Configuration Errors*

Among the leading causes of problems with hubs are faulty installations or configurations of the equipment. Incorrectly configured ports (port not enabled; wrong operating mode, for example, 10 Mbit/s instead of 100 Mbit/s), loose connections (loose cables, connectors, or plug-in modules), and wiring defects on the back plane or in the wiring cabinet are the most common error sources.

#### *Hardware Problems*

To locate or rule out hardware problems, check the power supply and connectors and run the hub's self-test function.

### Diagnosing Problems with Optical Components in Ethernet Networks

Optical fiber is increasingly the medium of choice for long-distance connections in Ethernet networks. Fiber-optic links are set up using either active hub modules or simple passive star couplers.

#### *Passive Stars*

Passive star couplers have no active components and thus do not require a power supply. Their only purpose is to connect optical fiber in a star topology with minimum attenuation. To optimize the use of passive optical stars, the IEEE 802.3 10Base-FP standard defines a modified preamble for frames transported over 10Base-FP segments.

#### *Active Optical Stars*

Optical repeater components must be able to perform all functions defined for repeaters in the IEEE 802.3 specification. These include signal regeneration, fragment extension, collision detection, inter-frame gap verification and generation of jam bits.

#### *Optical Power Problems*

The optical power of some optical repeater modules can be adjusted. A power level that is too low can cause transmission problems, especially when a packet travels over a number of fiber lines. In this case, you need to measure the loss over the various fiber lines as well as the optical power that is emitted from the router port. Symptoms of low optical power or high fiber-line attenuation are generally low network performance between specific nodes and an increased incidence of short packets (runts).

*Problems with Non-Standardized Optical Repeater Components*

Simple electrical-to-optical converters are sometimes used instead of optical repeater ports that fulfill all the functions specified in the IEEE 802.3 standard. When these converters are deployed in high numbers, problems may occur. Typical symptoms include:

- Runts
- Insufficient inter-frame gaps
- Noise (especially when using converters without retiming boards because these neither regenerate amplitude nor remove phase jitter)

*Hardware Problems*

As with simple repeaters, if you suspect that there is a problem with the hardware, check the power supply and the connectors, and run the self-test function.

## 7.2.8 Problems with Bridges

Bridges are inter-network elements that connect network segments on OSI Layer 2 (the MAC layer). Independent of higher-layer protocols, bridges store and filter the Ethernet frames they receive from connected segments and transmit them to their destination segments. The main function of a MAC-layer bridge is to prevent the spreading of local traffic to neighboring segments and to overcome such limitations of the particular network topology as the maximum number of nodes per segment, the transmission delay or the distance limits. Packet filtering is based on an address table. This table is also called the "forwarding table", and can either be generated by the bridge in the learning mode (see the following) or manually defined by the operator. In the learning mode, the bridge acts as a repeater, except that the source address of each packet arriving at a bridge port is stored in the forwarding table together with the corresponding port number. This operating mode is also called "transparent bridging", and must be implemented in every bridge according to the IEEE 802.1 specification. After a short time the bridge knows which node is located in which segment and can begin operating in bridge mode. From that point on, frames are forwarded selectively. Each frame is forwarded only to the segment in which its destination node is located. To limit the size of the forwarding table and take changes in the network topology into account, all address entries are maintained according to an aging process. If no activity is detected for a given address entry for a certain period of time, the entry is automatically deleted.

Another important bridge function is checking Ethernet frames. Unlike repeaters, bridges discard defective packets, such as those with an illegal length, an invalid FCS or incorrect alignment.

While local bridges are used to connect neighboring segments, wide-area bridges consist of two separate components connected over a wide-area network. WAN bridges can link networks that are located at different sites. A unified standard for bridges, published in 1990, is part of IEEE 802.1. This document has been updated by the ISO standard IS 10038.

### Features and Functions of Bridges

*Bridge Networks with the Spanning Tree Algorithm*

If two subnetworks can be reached through more than one path, a mechanism is needed that selects the optimum route in each case. Furthermore, the creation of loops must be avoided. The spanning tree algorithm (STA) is designed to prevent loops in Ethernet networks. If a network that contains loops uses bridges that do not support STA, this can lead to "sputniks", or packets that circulate continuously in network loops and reproduce themselves. One copy of such a frame may reach its destination while one or more copies remain in the loop, unnecessarily increasing the network load.

The spanning tree algorithm recognizes one bridge as a root bridge for the network. On each remaining bridge, one port that leads to the root bridge is defined as the root port. The bridges are configured with the bridging protocol, which uses special frames (bridge protocol data units or BPDUs) to select the root bridge and the root ports. Initially, all bridges are defined by default as root bridges and begin broadcasting packets to all connected segments every two seconds. Each broadcast contains the sender's 8-byte identification together with notification of its status as a root bridge. The bridge with the lowest identification value is then selected as the root bridge, and each of the other bridges defines the root port it will use to connect to the root bridge. In each network that is not directly connected to the root bridge, the bridges inform each other of the identification of their root bridges and the corresponding transmission capacities. The transmission capacity is expressed as the reciprocal of the communication speed, or the sum of all transmission delays. We recommend using the expression 1,000/(data speed in Mbit/s). This yields transmission coefficients of 2,000 for a 512 Kbit/s line, 250 for 4 Mbit/s Token Ring, 100 for 10 Mbit/s Ethernet, 62 for 16 Mbit/s Token Ring, 10 for FDDI and 100 Mbit/s Ethernet, and 1 for Gigabit Ethernet. If there is more than one connection between a subnetwork and the root bridge, the path with the best transmission characteristics can be selected. The remaining ports for this path are then deactivated. Configuration requires both configuration BPDUs and topology BPDUs. Bridges periodically distribute the configuration BPDUs within the network. The starting age of packets sent out by the root bridge is defined as zero, and each bridge increases the age counter of forwarded packets by one. Packets age as the distance increases; the age value, however, must never

exceed 20. If it does, the information from the last configuration BPDU is discarded and a new reconfiguration of the root port is initiated. The bridge informs the remaining network of the change by transmitting a topology-change BPDU. The root bridge acknowledges the reconfiguration with a configuration BPDU in which a topology-change bit is set.

| | | | |
|---|---|---|---|
| Protocol ID = 0 | | | |
| Protocol version ID = 0 | | | BPDU type = 0 |
| TCA | Flags | TC | |
| Root ID | | | |
| Transport cost to root bridge | | | |
| Bridge ID | | | |
| Port ID | | | |
| Message age | | | |
| Maximum age | | | |
| Hello timer | | | |
| Forwarding delay | | | |

| | |
|---|---|
| TCA (Topology Change Acknowledgment): | The bridge acknowledges receipt of a topology change BPDU |
| TC (Topology Change): | The root bridge notifies the other bridges of a topology change |
| Bridge ID | Priority of the bridge (0 = highest; 65535 = lowest; default=32768) |
| Port ID | Throughput capacity to root bridge (1= highest; 65535 = lowest) |
| Message age | Delay between receipt of previous and current configuration BPDUs |
| Maximum age | Maximum message age (6 - 40 seconds; recommended value: 20 seconds) |
| Hello timer | Maximum interval between root bridge configuration BPDUs |
| Forwarding delay | Delay between learning and forwarding phases (4 - 30 seconds; recommended value: 15 seconds) |

**Format of a configuration BPDU**

| | |
|---|---|
| Protocol ID = 0 | |
| Protocol version ID = 0 | BPDU type = 0 |

**Format of a topology change BPDU**

*Figure 7.20 Configuration and topology change BPDUs*

Other important parameters for bridge operation are the maximum bridge latency, the BPDU transmission delay and the minimum time between configuration changes (hold time):

Bridge latency: 1 second recommended; 4 seconds maximum
BPDU transmission delay: 1 second recommended; 4 seconds maximum
Hold Time (minimum time between two BPDUs): 1 second

### *Bridge Filters*

For most bridges, manually-defined filters can also be created in addition to the forwarding table generated in the transparent bridging mode. These filters can be based on various criteria, including source or destination address, type field, or other bit masks. Logical operators (AND, OR, NOT) can also link the filters. Bridge filters can be a powerful tool for actively avoiding error situations and restricting network traffic. Their usefulness is limited, however, by the fact that they are static; higher-layer protocols that select port addresses dynamically, such as FTP or Telnet, cannot be filtered properly.

### *MAC Switching Bridges*

The latest generation of bridges combines bridge functions with MAC-layer switching. These bridges are able to make forwarding decisions and carry them out for several packets concurrently. Replacing hubs by MAC-layer switching bridges has the advantage that each frame is routed directly to its destination segment without blocking the other stations on the same network. This can significantly increase the communication bandwidth available in the network. The principles of segment switching are discussed in more detail in Chapter 11.

### *Remote Bridges*

Remote bridges usually consist of one LAN port and several serial ports, over which the LAN can be connected to remote locations using WAN links. The serial ports can be V series interfaces (such as V.24 or V.35), ISDN lines or PDH interfaces (E1/T1, E3/T3). The protocols used for transporting the bridged data over the wide-area links are usually Point-to-Point Protocol (PPP) or proprietary variants of the High-level Data Link Control Protocol (HDLC).

### *Load Balancing*

Network traffic can be distributed over several parallel transmission paths using the non-standard load-balancing protocol (distributed load sharing or DLS). When used in conjunction with WAN bridges, the load-balancing protocol calculates the transmission cost for each transmission path based on the available bandwidth. The traffic is then distributed to minimize the total cost. Because the data packets of a single communication process may be distributed over various lines, the packet sequence may change. Thus protocols or applications that rely on unchanged packet sequences should not be operated over links that use load balancing.

### *Connecting Different LAN Topologies*

If bridges are used to connect different LAN topologies (for example, Ethernet, Token Ring, FDDI, ATM), a number of special functions are required in order to

overcome the differences in data speed, packet format and media access technique. Because not only the packet lengths but also the entire frame formats are different for each of these topologies, each packet must be translated before transmission, which requires additional CPU performance. The difference in data speeds can also present difficulties. For example, if a long sequence of frames is being transmitted from a 16 Mbit/s Token Ring to a 10 Mbit/s Ethernet network, the bridge cannot forward the traffic at the same speed at which it arrives. Therefore, packets must be stored in the bridge buffer, which can overflow if the transmission sequence is too long. Timers in the higher-layer protocols present another significant problem, and one that is often overlooked. If a Layer 3 timeout occurs before a given packet has arrived at a destination in a slower LAN, the transmitting station may start retransmissions too soon. After several unsuccessful attempts, the sender gives up any further attempt on the assumption that the connection has been interrupted.

Another special bridge function is necessitated by the differences in packet lengths. While the maximum packet size in IEEE 802.3 Ethernet is 1,518 bytes, Token Ring packets can have up to 4,500 bytes (in 4 Mbit/s Token Ring) or 17,800 bytes (in 16 Mbit/s Token Ring); FDDI also allows 4,500 bytes. Because there is no OSI Layer 2 mechanism defined that would allow packets to be divided into several parts and then reassembled, packets that are too long for their destination LAN topology must be discarded by the bridge. Due to the problems involved in translating between two completely different transmission methods (translation bridging), encapsulation bridging was defined as an alternative in IEEE 801.2H. In encapsulation bridging, a backbone such as a FDDI ring transports Ethernet frames as a simple payload from one ring node to another, paying no attention to the frame headers. Analogously, RFC 1483 defines the encapsulation of LAN packets (Ethernet, Token Ring, FDDI) for transport in ATM cells without regard to network packet headers.

All in all, there are many difficulties and restrictions involved when bridges are used to couple different network topologies, which is one reason why most networks now use routers for this purpose.

### Diagnosing Bridge Problems

The challenge when analyzing bridge problems is to correlate the occurrence of symptoms in several different network segments. Using probe-based monitoring systems to provide concurrent measurements in several segments can be very helpful. Performance measurements of bridges by expensive, specialized multiport systems are less important: most modern bridges are capable of forwarding packets at line speed anyway, so performance measurements in most cases just confirm the manufacturer's technical data. It is more efficient to

request system specifications from the manufacturer based on standardized test methods, as specified in RFC 1242 and RFC 2544.

Most problems that affect bridges can best be located by a process of elimination that involves the correlation of specific measurements and an analysis of the network topology. Symptoms of bridge problems can include poor network performance in particular segments, intermittent or permanent loss of connection for particular stations, or the failure of certain protocols and services. The first step of the troubleshooting process is, as always, a review of all configuration changes that were made in the network before the error occurred and, of course, general information gathering. If the symptoms correlate to particular connections, you can begin by checking all bridges located along the corresponding transmission path. Otherwise, the next step is to prepare a list of all the stations, connections, protocols and services affected by the problems observed. To do this, measure the current parameters in the various network segments and compare the results with statistics gathered during normal operation. This involves recording and analyzing throughput and performance parameters of network nodes, protocols, and services as well as reviewing the log files that contain operating statistics on all bridges in the network. The log files provide bridge parameters such as CPU capacity use, port capacity use, buffer capacity use and error statistics. To measure the response times of connections across bridges, transmit loop-back packets—Ethernet configuration test packets (CTP), IEEE 802.2 LLC test packets, 802.2 LLC XID test packets, IP pings—from different network segments across the bridges. Automatic long-term response time measurements can be made using dedicated response time agents distributed throughout the network. This type of long-term measurement can be especially useful in the case of intermittent problems. Based on the results of such measurements, the range of potential sources of the error in question can usually be narrowed down to specific components.

#### Symptoms and Causes of Bridge Problems

This section summarizes typical symptoms of bridge problems and their most frequent causes.

##### *Throughput Capacity Problems*

The throughput capacity of bridges is expressed in frames per second. The nominal value usually refers to the shortest possible frame length for the network topology in question. However, when throughput is given in bytes per second, the figures usually refer to an average or even the maximum possible frame length. Throughput capacity and the actual throughput rate should thus be checked, especially when dealing with older bridge models. State-of-the-art bridges have throughput capacities well in excess of the frame rates that occur

in practice, so that bottlenecks caused by limited throughput capacity are becoming increasingly rare.

### *Loss of Frames*

In addition to the loss of packets due to capacity problems, a correctly functioning bridge discards damaged frames (runts, jabbers, collisions, collision fragments) as well as frames that have exceeded their lifespan. Although a bridge does not examine the time stamp on a frame (this is left to higher-layer protocols), it can place a limit on the time a packet is buffered in the bridge. If this latency limit (maximum 4 seconds) is exceeded, the packet is discarded. A latency limit that is too short could cause problems in a network with heavy traffic.

### *Problems with Bridge Filters*

The following are potential sources of problems when filters are used:

- Overly complex filter structures, whose effects cannot be foreseen in every operational situation, may result in unwanted filtering under certain operating conditions.
- When multiport bridges are used in redundant areas of the network, keep in mind that data that normally has another transmission path may also be transmitted over the backup port in certain situations. Filters must therefore be checked to ensure that they do not block packet streams in backup mode.
- Depending on their architecture, some bridges show a significant performance reduction once a given number of filters are active. Check with the manufacturer for information on how active filters may affect the operating performance of the bridge.

### *Buffer Overflow*

To handle short-term peak traffic, bridges store incoming packets in a buffer until the CPU is ready to forward them to their final destination. If the buffer is full, the bridge has to start discarding packets.

### *Excessive Data Packet Length*

Often, when bridges translate between different network topologies (between Ethernet and FDDI, for example), the packets to be forwarded exceed the maximum packet length allowed in the destination network. These packets must be discarded because OSI Layer 2 does not define a fragmentation mechanism for bridges. The maximum packet length to be transmitted by the bridge can also be set manually. However, setting the value too low can lead to performance

problems, especially when using time-critical protocols such as Local Area Transport Protocol (LAT).

*Changes in Packet Sequence*

Bridges usually forward packets in the same order in which they are received. However, if the bridge is using a distributed load-sharing (DLS) mechanism, packets are sent over a number of redundant links or bridge ports and do not arrive in exactly the same order as they are sent (non-FIFO load sharing). In this case, only those protocols that can tolerate changes in the packet sequence (TCP/IP, XNS, IPX, etc.) will be able to communicate successfully across the bridge. The User Datagram Protocol (UDP), for example, requires packets to be in the correct sequence. Any irregularity in the transfer sequence would have to be corrected by the application that uses UDP as a transport protocol, but this is often not done. For this reason, UDP applications sometimes function properly within a local segment, but have major problems when working across bridges or routers.

*Problems Related to the Address Table (Forwarding Table)*

As discussed previously, each entry in the address table of a bridge is coupled with an "aging" algorithm. Entries that reach a certain age without being used are deleted from the address table. The maximum age of an entry can be configured manually, and usually has a value between 10 and 1,000,000 seconds. In addition to such dynamic address tables, most bridges also have static address tables. These must be set up manually and often contain special addresses, such as broadcast or group addresses. When the operation of a bridge is based entirely on the use of static address tables, the bridge is said to be operating in "protected mode". In this mode, addresses of new stations (stations just beginning network activity) are not automatically added to the address table, and are therefore unable to communicate across the bridge. Manually creating or editing such static tables often leads to errors. Together with incorrectly configured bridge filters, address tables account for the majority of bridge problems. The first step in your troubleshooting process should be to check the address and filter tables.

*Wrong Operating Mode*

Bridge ports operating in Ethernet v2.0 mode instead of IEEE 802.3 mode sometimes cause problems with various older network interface cards. However, state-of-the-art NICs are able to adjust to either operating mode automatically. Another common configuration error is setting a port for 10 Mbit/s Ethernet instead of 100 Mbit/s Ethernet or vice versa. Again, most of today's network components identify and adjust to the required Ethernet frame type automatically.

*Remote Bridge Problems*

The wide-area lines used to connect remote bridge ports are often too slow or of poor quality. In these cases, timeouts and lost packets result in retransmissions at peak loads. This problem may be solved by using data compression or by upgrading the wide-area line to a higher bandwidth. Frequent checks of the line's capacity use and quality by means of a protocol analyzer are recommended in order to avoid the sudden occurrence of serious communication problems.

*Problems with Bridges Connecting Different Network Topologies*

Bridges that connect different network topologies are faced with a number of basic problems due to differences in line speed and access mechanisms:

Token Ring to Ethernet Connection: The A and C bits in the frame status byte of IEEE 802.5 frames inform the transmitting station whether the destination node was able to receive the packet and copy it into its receive buffer. The bridge can be configured to set these bits to "packet received" and "frame copied" by default, but this will obviously create problems in the event that the receiving station is not ready to receive: the transmitting station will assume its packet was delivered successfully even if the destination station was not active and never received the frames in question.

FDDI to Ethernet Connection: As with Token Ring, the A and C bits in the frame status bytes inform the sender whether its packet was delivered successfully. Again, a default setting for these bits in the bridge is only useful as long as the packet really is delivered; otherwise the transmitting FDDI station may receive the message that its packet arrived even though the Ethernet station was not active.

*Installation and Configuration Errors*

Among the leading causes of bridge problems are errors due to faulty installation or configuration of the equipment. Incorrectly configured ports (port not active; wrong operating mode, for example, 10 Mbit/s instead of 100 Mbit/s); faulty connections (loose cables, connectors, or plug-in modules); and wiring errors on the back plane or in the wiring cabinet are the most common error sources.

*Hardware Problems*

To locate or rule out hardware problems, check the power supply and connectors and run the bridge's self-test function.

### 7.2.9 Problems with Routers

Routers are inter-networking components that connect network segments on OSI Layer 3. Because they operate on this level, routers can link networks regardless of their topologies. The router extracts and analyzes the network layer contents (IP packets for example) from the OSI Layer 2 frames it receives (Ethernet, Token Ring, FDDI, ATM, X.25, Frame Relay, ISDN), and forwards them encapsulated in the Layer 2 packet format of the destination network. Analysis of the Layer 2 packets is, therefore, of no use in diagnosing routers because no data transport takes place at this layer. Only the Layer 3 protocols involved (IP, IPX, NS, IDP, OSI, and others) are of interest here. Most of today's routers are multiprotocol routers, which are able to process several protocol stacks in parallel. Single-protocol routers specialize in forwarding just one Layer 3 protocol. Because some Layer 3 protocols are poorly—or not at all—suited for routing, modern systems support both routing and bridging. The system then decides whether to route or bridge each packet. Such systems are sometimes referred to as brouter (bridge + router) systems. Special routing protocols have been defined to ensure efficient route selection for packets forwarded through routers.

#### Router Protocols

The most widely used routing protocols are:

- RIP (routing information protocol), OSPF (open shortest path first) for TCP/IP
- RIP, NLSP (NetWare link services protocol) for NetWare/IPX
- DRP (digital routing protocol) for DECNet
- IS–IS (intermediate system–intermediate system) for OSI/CLNS
- RTMP (routing-table maintenance protocol) for AppleTalk
- IGRP (inter-gateway routing protocol)
- BGP (border gateway protocol)

Basically, there are two different methods for selecting routes within a network: static routing and dynamic routing. Static routing is used for static connections with fixed transmission characteristics (bandwidth, availability, bit-error rate). Static routes are rarely used unless called for by a very stringent security policy or if the connections between the routers are established over wide-area, packet-switched networks, such as ATM or X.25. Apart from these cases, dynamic routing methods are used, especially in local-area networks. Today's dynamic routing protocols are based either on the vector distance (Bellman-Ford) routing algorithm or the "shortest path first" (link-state) algorithm (SPF).

*Vector Distance (Bellman-Ford) Routing*

In vector distance routing the distance between the routers is given in hops, where each section of the overall link between two routers is one hop. Each router maintains a table that contains all possible routes to the various destinations and the distance in hops along each route. This table is sent periodically to all adjacent routers, which also broadcast their tables. The disadvantages of the vector distance routing algorithm are that route changes are distributed relatively slowly over the network, and that the size of the routing tables is proportional to the size of the network. The algorithm considers only the number of hops for its routing decisions, not the transport capacity of the links. The updates are time-based and are not based on available routes.

*Shortest Path First (SPF) Routing*

In SPF routing, each router has two tasks: the first is to test the status of the links to adjacent routers, and the second is to broadcast this routing table to all routers in the network. In this algorithm, the size of the routing update packets is independent of the size of the network because only the status of the adjacent neighbors is transmitted. However, each router in the network has the latest routing table, valid for the entire network, available at all times. The local router calculates the most efficient route and is not dependent on routing calculations made by other routers. The routing table updates are link-state-change based and not time-based, as is the case in vector distance routing.

*Routing Internet Protocol (RIP)*

In IP networks, the two most popular routing protocols are RIP and OSPF. Due to its simplicity, RIP (RFC 1058) is used in the majority of IP networks. It uses routing tables that are generated by the individual routers and contain the subnet addresses (not the host addresses) for all subnetworks. The size of the routing table is proportional to the number of segments, not to the number of hosts. Based on the routing tables it receives, each router can calculate the distance to the destination network and select the best route. This technique allows for no more than 14 routers between the transmitter and the receiver. An address that is 15 or more hops away is considered unreachable.

Alternatively, the routing table can be made smaller by including only the directly accessible subnets and designating a default route for all other subnets. All packets that cannot be forwarded directly to their destination subnet are then transmitted to the default gateway. In general RIP is primarily used for smaller networks due to its hop count limitation and due to its relatively bandwidth-intensive update mechanisms.

*Open Shortest Path First (OSPF)*

The second most popular routing protocol in IP networks is OSPF (RFC 1247). OSPF differs from RIP primarily in that it has a hierarchical structure, can bridge more than 14 hops and uses the link-state algorithm SPF. The main strength of OSPF is its ability to pool several networks in so-called "areas". Each router receives a topology database (link-state database) for each area, which is identical for all routers in the area. Each router broadcasts the state of its own links (bandwidth, throughput, transmission delay) to all other routers within its area. Based on this topology database, each router calculates a tree structure locally, with itself as the trunk and the shortest routes to the various destinations as branches. OSPF provides fast network convergence because the routing table updates are not time-based, however, in large networks with frequent link-state changes it can consume bandwidth. Bandwidths of communication links are being taken into consideration for routing decisions.

*Inter-Gateway Routing Protocol (IGRP)*

IGRP is based on the distance vector algorithm with the addition of various enhancements that enable it to provide faster route selection while generating less overhead. Each router transmits routing table updates to all neighboring routers every 90 seconds. The neighboring routers interpret the updates but do not distribute them further. To evaluate the performance of the various links, the parameters describing response time, bandwidth, link length and availability are combined in the following formula:

$$Link\ cost = \{(K1 \cdot B)+(K2 \cdot B)/(256-L) + (K3 \cdot Dc)\} \cdot \{K5/(r+K4)\}$$

where

K1 = bandwidth weighting

K2 = capacity use weighting

K3 = delay weighting

K4, K5 = reliability weighting

B = bandwidth

L = capacity use (0 ... 255; 255 = 100% capacity use)

Dc = composite delay in units of 10 µs

r = number of frames sent / number of frames successfully transmitted (reliability)

The IGRP default values are K1=1, K2=0, K3=1, K4=0 and K5=0.

### Diagnosing Router Problems

Unless the errors are caused by hardware or installation problems, troubleshooting router problems begins by focusing on the protocols being used and routed. The first step is to gather protocol performance statistics using a protocol analyzer in order to get an idea of the current operating status of the network. In an IP network, for example, this would include statistics such as

- IP broadcasts
- ICMP redirects
- Low TTL messages
- Routing packets
- The proportion of IP traffic in the total network load (in percentages and kilobytes)
- Fragmented IP packets
- ICMP unreachable messages
- Other ICMP messages

In addition, the log files that contain the operating statistics of all active routers should be analyzed. There are a variety of router commands for retrieving data, such as CPU capacity use, memory capacity use, port capacity use, number of packets transmitted and received per protocol, timeouts, fragmentations, and numbers of connections and broadcasts (for Cisco routers, for example, these commands would include the following: show interfaces, show controllers, show buffers, show memory, show processes, show stacks).

If the symptoms can be related to a particular connection, you can begin by checking all routers located along its transmission path. Be sure to check the following items:

- Address tables (Are the address entries for the affected nodes correct?)
- Mapping tables (Is the mapping between network address and host name correct?)
- Routing tables (Is a route to the destination network available?)
- Filter entries
- Protocols (Are all protocols active?)
- Default gateways (Does the configuration define a default gateway or default route? This is represented by the destination network 0.0.0.0 and is used when a packet has a destination that is not contained in the routing table)
- Timer values (Are correct values configured for the active timers, such as the Hello timer or Dead timer in OSPF?)
- Static routes (Are static routes active? If so, are the links working?)
- WAN ports (Are all WAN links up and running?)

A check of the router's event log can provide additional clues to the cause of the problem.

### Symptoms and Causes of Router Problems

The following summarizes the typical symptoms of router problems and their most frequent causes.

#### *Throughput Capacity Problems*

Due to the trend towards high-speed network technologies, bottlenecks often occur during peak traffic even in state-of-the-art routers based on modern RISC processor technology. If the router manufacturer has provided detailed operating characteristics based on standardized test methods (RFC 1242, RFC 2544), you can compare this data with the actual peak loads that occur in your network. Based on the results of this comparison, you can make a rough estimate as to whether the traffic through the router might cause throughput problems or not. In addition, reviewing the operating statistics log may help you determine whether there are performance problems in your router.

#### *Address Table Problems*

Many router problems are caused by address tables that have not been updated or are otherwise incorrect. This type of error is often the result of configuration changes in the network that were not implemented on the routers. The problem symptoms are not seen until the next time the affected service is used, often hours or even days after the configuration changes. This can make it difficult or impossible to detect a correlation between the altered network configuration and the failure of a particular service, such as database or Internet access. Localizing such problems can be very time consuming, especially if the network configuration changes have not been documented, as is often the case.

#### *Faulty Subnet Masks*

Faulty subnet masks are another common cause of router problems. Unique host addresses, for example, can become subnet broadcasts due to a faulty subnet mask on the router. The only way to solve this type of problem is by systematic checking and documentation of all subnet masks.

#### *No Default Gateway*

A typical cause of partial connection losses—that is, when connections are possible between some nodes or subnets but not others—is the lack of a default gateway configuration on the router. In this case, connections to subnets that are directly linked to the router function properly, but connections through more distant routers do not because there is no default gateway available for the router to forward packets to.

*Faulty Timer Configuration*

Great care must be taken when setting protocol timer values on routers. Incorrect timer settings can lead to delayed distribution of routing information within the network (caused by incorrect setting of the OSPF Hello or Dead timer or the IGRP Active timer) as well as to early timeouts of particular connections. When combining routers from different manufacturers it is especially important to check the timer settings and adjust them if necessary.

*WAN Problems*

The second most common sources of router errors, after faulty address tables, are problems with WAN links. These may be problems in the public telecommunications network (link failures, high bit-error rates, long delays) or basic problems such as insufficient bandwidth in WAN links, or incorrect protocol settings (such as timers or window size).

*Routing Protocol Problems*

Another class of error sources involves problems in router-to-router communication. Analyzing the activity and content of the routing protocols with the help of a protocol analyzer may help you track down these faults.

*Installation and Configuration Errors*

As with hubs and bridges, the leading causes of router problems are errors due to faulty installation or configuration of the equipment. Incorrectly configured ports (port not active; protocol not active; wrong operating mode, for example, 10 Mbit/s instead of 100 Mbit/s), faulty connections (loose cables, connectors, or plug-in modules), and wiring errors on the back plane or in the wiring cabinet are the most common error sources.

*Hardware Problems*

To locate or rule out hardware problems, check the power supply and connectors and run the router's self-test function.

## 7.2.10 Symptoms and Causes: 10/100/1,000 MBit/s Ethernet

**Symptom: Diminished Network Performance in Conjunction with FCS Errors**

Invalid checksums (FCS errors) are a side effect of collisions, which in limited numbers are a normal consequence of the CSMA/CD algorithm. If FCS errors occur together with collisions, and if

their number is within reasonable limits, there is no reason to worry. Use a protocol analyzer to measure the number of collisions and the number of FCS errors over a period of time and compare the resulting curves. If there is no correlation between the collision and the FCS error curves, you might have one of the following problems:

Cause (1): Noise and interference on the network.
Noise results when the network is not grounded or if the grounding is faulty. Use a cable scanner or multimeter to check the noise level on your network. A 10Base2/10Base5 network segment must have no more than one ground connection. If there is a second ground connection, due to a faulty network interface card or a bad cable connection, for example, a voltage difference between the two grounds may cause current leak in the network cable.

Cause (2): Electromagnetic interference along the cable path.
Electromagnetic interference from devices such as photocopiers, mobile telephones, elevators or pagers can also cause FCS errors. Use a multimeter to check for interference and a cable tester to check for noise. If you detect interference, check whether the cable routes lead along elevator shafts, electric machinery, transformers, lighting bays, computer systems with high clock rates or X-ray equipment.

Cause (3): Faulty network interface card.
To determine whether a faulty NIC is the source of FCS errors, generate statistics of all defective packets sorted by network node (this is a standard report generated automatically by most protocol analyzers). If you find a suspicious station, measure its activity (for example, in packets/second) and the number of FCS errors occurring on the segment. If the two numbers seem to correlate, there is a good chance you have found the cause of your problems. Keep in mind that many faults on network interface cards occur only intermittently, for example, only after the card has reached a certain temperature. For this reason it may be necessary to take measurements over longer periods of time before you can obtain exact and repeatable results.

Cause (4): Defective or loose connectors (on NICs, wall jacks, MAUs, repeaters, hubs).
Check all connections in the network path.

**Symptom: Diminished Network Performance in Conjunction with Late Collisions**

An increase in the number of collisions is often caused by cable problems (cable segments too long), defective network interface cards, excessive repeater cascading, or defective or missing terminating resistors. Determining whether the collisions are "late" or normal collisions can help to narrow down the possible causes. Possible causes for late collisions include:

Cause (1): Cables longer than the specified maximum segment length for the given topology.
Measure length using a cable scanner.

Cause (2): Too many cascaded repeaters in network.
Replace one of the repeaters with a bridge, or change the network configuration.

Cause (3): Defective network interface card or MAU.
Use a protocol analyzer to collect statistics on the stations that send the most defective packets. Also gather statistics on numbers of collisions and active nodes, and look for correlations. If this does not help to localize the problem, the network segmentation method must be used.

**Symptom: Diminished Network Performance in Conjunction with Early Collisions**

Cause (1): Terminating resistor defective or not installed.
10Base2 and 10Base5 networks must be terminated by 50Ω resistors. Make sure all required terminating resistors are installed and use a multimeter to check the resistance (48 Ω < R < 52 Ω).

Cause (2): Loose or defective T-connector.
Check all connections in the network path.

Cause (3): Too many nodes in one segment.
Check the number of MAUs per segment; the number must not exceed 100 in a 10Base5 segment or 30 in a 10Base2 segment.

Cause (4): Kink in a cable.
Use a cable scanner to try to locate the damage and replace the affected cable.

Cause (5): Cable does not conform to IEEE 802.3.
IEEE 802.3 10Base5 cables are marked with a color code every 2.5 meters. In order to minimize the interference due to reflections at connection points, connectors should be inserted only at these markings. In addition, keep in mind that not all cables with BNC connectors are 50Ω cables. Although Ethernet works even on

75Ω cables over tens of meters, increasing network length will lead to problems sooner or later. Always check the specification of the cables you are using.

**Symptom: Slow Network, High Response Time (No Excessive Collisions or FCS Errors)**

Cause (1): Buffer overflow in a bridge or router in the transmission path.
Check router and bridge statistics (CPU capacity use, port capacity use). Use a protocol analyzer to try to determine which nodes create the most traffic across the bridge or router. Do timeouts occur? Use pings to perform systematic measurements of response times across the bridge/router to check whether the interconnection devices are part of the problem. If this is the case, reconfigure the network (by moving a server or client to another segment, for example) to reduce the traffic over heavily-loaded interconnection devices.

Cause (2): Transmission problems over optical-fiber connections.
Fiber links bridging great distances can sometimes lead to performance problems without showing FCS errors if the line attenuation is too high or the light power emitted is too low. Use pings to check the response times of connections over the fiber connection in question. Check the settings of the fiber-optic couplers and the line attenuation.

Cause (3): Local segment routing.
Local routing is a common cause of slow networks. Local routing typically occurs for connections between two nodes with different subnet addresses connected to the same LAN switch, which is connected to a router (also called one-armed routing). In order to reach its destination a packet has to be switched to the router, then routed within the router before being transferred through the same switch again to the destination node.

**Symptom: Intermittent Problems with Connections and Network Performance, Alignment Errors**

Cause (1): Network interface card transmits a few extra bits after each FCS.
Use a protocol analyzer to capture the frames that have extra bits following the FCS (known as dribble frames or alignment-error frames). The source address of the captured packets identifies the faulty network interface card.

Cause (2): Maximum length of the transmission path exceeds that defined in the Ethernet specification.
Whether the signal arrives at its destination depends on the transmitting and receiving stations. Stations that are closer together (within the specified distance limits) can communicate without problems, while stations that have to communicate over a longer distance, but are still located in the same segment, have connection problems. Try to find a pattern in the connection problems to determine whether only certain nodes are affected. Use a cable tester to check the length and quality of the transmission path. Insert a bridge or router in the transmission path if necessary (see the section entitled "Network Design Guidelines" for details).

Cause (3): Too many bridges or routers are cascaded, resulting in long signal transmission delays and protocol timeouts (such as TCP timeouts).
Use pings or response time agents to check response times. Review the network design with regard to the maximum allowable cascading of bridges and routers (see the section entitled "Network Design Guidelines" for details).

**Symptom: Intermittent Connection Problems in Conjunction with Short Packets**

Cause: Faulty network interface card.
Use a protocol analyzer to try to capture the short packets and identify the emitting node by the source address. If the source address is corrupt, try to track down the defective card by evaluating correlation measurements (see the section titled "Troubleshooting in 10/100/1,000 Mbit/s Ethernet" for details).

**Symptom: Intermittent Connection Problems in Conjunction with Jabber Packets**

Cause (1): Double grounding in 10Base2 and 10Base5 networks, resulting in DC currents in the network cable.
Check the network grounding; use a cable tester to check for DC current.

Cause (2): Defective network interface card.
Defective interface cards sometimes generate jabber frames (excessively long frames), which lead to connection problems in the affected segment. Capture the jabber frames using a protocol ana-

lyzer and identify the faulty network interface card by analyzing the source addresses.

**Symptom: Intermittent Connection Problems in Conjunction with Short Inter-Frame Spacing**

Cause: Packet loss due to insufficient inter-frame spacing.
If a station does not maintain the required minimum inter-frame spacing gap (9.6 µs in 10 Mbit/s, 0.96 µs in 100 Mbit/s networks) some hubs will be unable to repeat the frames correctly. In such cases, packets sometimes mutate into jabber packets. Use a protocol analyzer to check the inter-frame gaps (calculated from the packet time stamps in the analyzer trace). The faulty network interface can then be identified by analyzing the source addresses.

**Symptom: Intermittent Connection Problems in Paths Across Bridges**

Cause: Change in packet sequence due to load-sharing mechanisms in the bridge.
Check the bridge configuration and deactivate load sharing if necessary.

**Symptom: Intermittent Connection Problems in Routes Across Routers**

Cause: Router connected to overloaded or low-quality WAN lines. Use a protocol analyzer to check capacity use, FCS rate and bit-error rates in the WAN link; analyze router port logs.

**Symptom: Loss of a Single Node's Connection**

Cause (1): Loose or faulty connection from the MAU to the network cable or from the NIC to the network.
A sudden complete failure of a single network node is often caused by one of the following:

- MAU plug not firmly connected
- Break, short circuit or noise in connecting cable
- Faulty network interface card

Check the cable and connector and the network interface card; replace if necessary. Replace the faulty node with a system known to be functioning correctly (such as a notebook). If the replacement node functions, the problem is inside the disconnected node; if not, the problem is on the network side.

Cause (2): Incorrectly configured network interface card: wrong connector activated (for example, AUI instead of twisted pair), or the selected interrupt is already assigned.

Send loopback pings (ping 127.0.0.0) to check whether the card is working and whether packets are being transmitted and received. Has anyone installed any hardware or software on the node recently? As described for Cause (1), replace the faulty node with a system known to be functioning correctly (such as a notebook) to determine whether the problem is inside the node or on the network side.

Cause (3): Defective network card, blown fuse.
Check whether the power supply to the MAU is intact (when using an external MAU). Send loopback pings (ping 127.0.0.0) to check whether the card is working and whether packets are being transmitted and received.

Cause (4): The MAU sends heartbeat signals (when working with an external MAU) but the interface card, in conformance with the standard, reads them as signal quality errors and aborts transmission.
Monitor the LEDs on the MAU. If the SQE LED lights up every time transmission is attempted, deactivate the heartbeat mode in the MAU (change it from Ethernet 2.0 mode to IEEE 802.3 mode).

Cause (5): Learning mode of a bridge not active because the bridge is operating in protected mode, and its aging function has deleted the address entry of the problem node.
Check the bridge address tables and the operating mode. (Is the learning mode on?)

Cause (6): Incorrectly configured bridge or router filters.
Check the filter settings and compare them with the address of the problem station. In particular, check the packet streams that occur when the bridge activates a backup path or load sharing.

Cause (7): MAC-IP address mapping problem, caused by change of static IP address or simultanous configuration with static IP address and Dynamic Host Configuration Protocol (DHCP).

**Symptom: An Entire Segment Has No Bridge Connection to the Rest of the Network**

Cause (1): Incorrectly configured bridge port (port not active; wrong operating mode, for example, 10 Mbit/s instead of 100 Mbit/s); faulty connections (loose cables, connectors, or plug-in modules); wiring errors on the back plane.
Check installation and configuration of the bridge.

Cause (2): Learning mode of a bridge is not active (that is, the bridge is operating in protected mode), and its aging function has deleted the address entry of the problem node.

Check the bridge address tables and the operating mode. (Is the learning mode on?)

Cause (3): Incorrectly configured bridge or router filters.
Check the filter settings; check wild card entries in particular.

**Symptom: An Entire Segment Has No Router Connection to the Rest of the Network**

Cause (1): Incorrectly configured router port (port not active; wrong operating mode, for example, 10 Mbit/s instead of 100 Mbit/s); protocol not active; faulty connections (loose cables, connectors, or plug-in modules); wiring errors on the back plane.
Check installation and configuration of the router.

Cause (2): Incorrectly configured address tables, mapping tables, routing tables.
Check router configuration.

Cause (3): Incorrectly configured router filters.
Check the filter settings. In particular, check wild card settings and filters that might block backup or load-sharing routes.

Cause (4): Failure of the wide-area connection on the router's WAN port.
Check whether the WAN line is up and running.

Cause (5): No default gateway setting.
Check whether a default gateway is set in the router configuration.

Cause (6): Incorrectly configured subnet mask.
Systematically check all subnet masks in the network.

Cause (7): Incorrectly configured timer settings.
Check the set timer values for the various protocols. Compare the default values, especially when using routers from different manufacturers.

**Symptom: Intermittent Connection Problems Between Client and Network**

Client connects, but loses connection periodically. Pings are returned, however, packet losses occur.

Cause (1): NIC or switch/router port misconfigured.
Both sides are not configured for the same operation mode. Check NIC, port settings.

Cause (2): NIC or switch/router port misconfigured (one side set to manual, one side set to auto negotiation).
Check NIC, port settings. Avoid using the auto-negotiate feature.

Cause (3): Host busy or overloaded, server experiencing problems.
Analyse server operating statistics and server response time.

### Common Errors

The most frequent sources of problems in Ethernet networks are listed in figure 7.21 in alphabetical order:

- AUI cable defective
- Bridge address list incorrectly configured; bridge in protected mode
- Bridge filter incorrectly configured
- Bridge overloaded
- Bridge's aging function deletes address entry
- Bridges or repeaters: too many are cascaded, resulting in timeouts and long response times
- Cable length exceeds specification
- Connectors, loose or defective: interface cards, wall jacks, MAUs, hubs, bridges, or routers
- Electromagnetic interference
- External MAU defective
- Faulty installation of physical router, bridge or hub (cable, connectors, plug-ins are loose; cable connections on the backplane are wrong)
- Grounding problems
- Inter-frame spacing gap too short
- Network grounded in more than one location
- NIC incorrectly configured
- Packets out of sequence due to bridge's load sharing function
- Signal power problems in optical components (optical hub ports)
- Router filters incorrectly configured
- Router incorrectly configured (port not active, protocol not active, wrong operating mode)
- Router overloaded
- Router protocol entries incorrectly configured (address tables, mapping tables, subnet masks, default gateways, routing tables, timer)
- Routing protocol problems (OSPF Hello timer, Dead timer, IGRP Active timer setting wrong)
- Terminating resistor defective or missing (10Base2, 10Base5)
- WAN connections down, overloaded, or of poor quality (high Bit-Error Rate (BER))

*Figure 7.21 The most frequent sources of problems in Ethernet networks*

**For additional excerpts from this chapter and other Network Troubleshooting book sections, be sure to regularly visit our web site at:**

## www.FreeTroubleshootingBook.com

New chapters will be posted every 2 to 3 weeks.
Be sure to visit our web site and vote for the chapters you would like to see posted!


Agilent Technologies