Excerpts taken from:

Network Troubleshooting
By Othmar Kyas

An Agilent Technologies Publication

Section II

Troubleshooting Local-Area-Networks

Chapter 11

Switched LANs

Agilent Technologies

11 Switched LANs

"You can observe a lot by watching."

Yogi Berra

11.1 Switched LANs: Specification and Implementation

LAN switches connect network segments or nodes and use a direct forwarding technique for data transmission. The LAN switch forwards each incoming frame directly to the switch port that is connected to the destination device. Multiple ports on a given LAN switch can operate simultaneously (As long as the transmission paths involved do not overlap), which can help ease congestion in network traffic. The high-performance forwarding hardware that makes this multiple parallel forwarding possible is the main difference between a LAN switch and a router or multiport bridge. The way in which frames are forwarded by the LAN switch is configured for each switch port. Thus, any port on a given LAN switch can act as a repeater, a bridge, or a router (see Figure 11.1).

LAN switches can be deployed in all traditional LAN topologies, including 10/100/1,000 Mbit/s Ethernet, 4/16 Mbit/s Token Ring, and FDDI. LAN switches generally use one of two techniques for forwarding data packets: "cut-through" or "store-and-forward". When the cut-through technique is employed, the switch makes the forwarding decision for each incoming frame as soon as it has received the first 6 bytes (which contain the destination address). The benefit of this technique is that the switch can forward frames at a high rate; latency is reduced to around 40 µs. The drawback is that the switch begins forwarding the data packet before it can determine whether there are errors in the frame, which could result in the propagation of defective frames.

By contrast, with the store-and-forward technique, which is increasingly used in state-of-the-art switches, the switch does not make the forwarding decision until it has received the entire frame. This has the advantage of not propagating defective frames. The drawback, however, is greater latency that is dependent on the frame size. The latency of a 1,000-byte data packet in a 10Base-T network introduced by a store-and-forward switch, for example, is over 800 µs. If the

transmission path leads through multiple store-and-forward switch ports, the transmission delays could be significant and higher-layer timers might run out. The search for a solution to this problem has led to the development of "fragment-free" cut-through forwarding. In this method, the forwarding decision is not made until the first 64 bytes of a packet have been received. This eliminates the majority of defective packets because the most common errors (for example, including Ethernet runt packets) can be detected within the first 64 bytes. At the same time, the latency involved is not as long as in store-and-forward switching. Many of the LAN switches available today also offer an "adaptive forwarding" feature, which enables each switch port to change between the store-and-forward and cut-through techniques based primarily on the numbers of runts and defective packets received at that port.

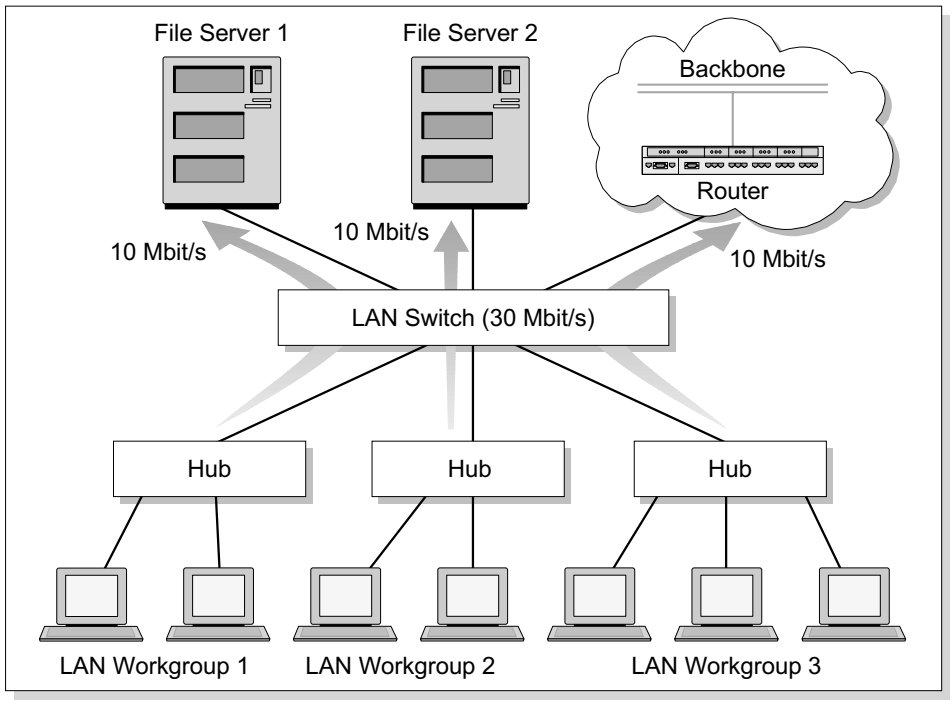

Figure 11.1 LAN switching increases throughput in inter-segment traffic

Another important difference between LAN switches is whether each switch port supports only a single MAC address or multiple MAC addresses. "Single-MAC" switches, also called desktop switches, are usually very economically priced and are primarily designed for direct connection to end devices. "Multi-MAC" or workgroup switches are used to connect multiple LAN segments. Workgroup switches support multiple hardware addresses for each port.

11.2 Design Guidelines for Switched LANs

To gain the full benefit of the high throughput capacity available with LAN switches, several basic guidelines must be followed when installing LAN switching components. It is essential to begin with an analysis of the network traffic, including the types of packets transmitted between the stations or segments to be linked by switches. A LAN switch cannot alleviate problems caused by broadcasts, for example, because switches distribute broadcasts over all ports just as bridges do, unless a Layer 3 switch capable of routing is used.

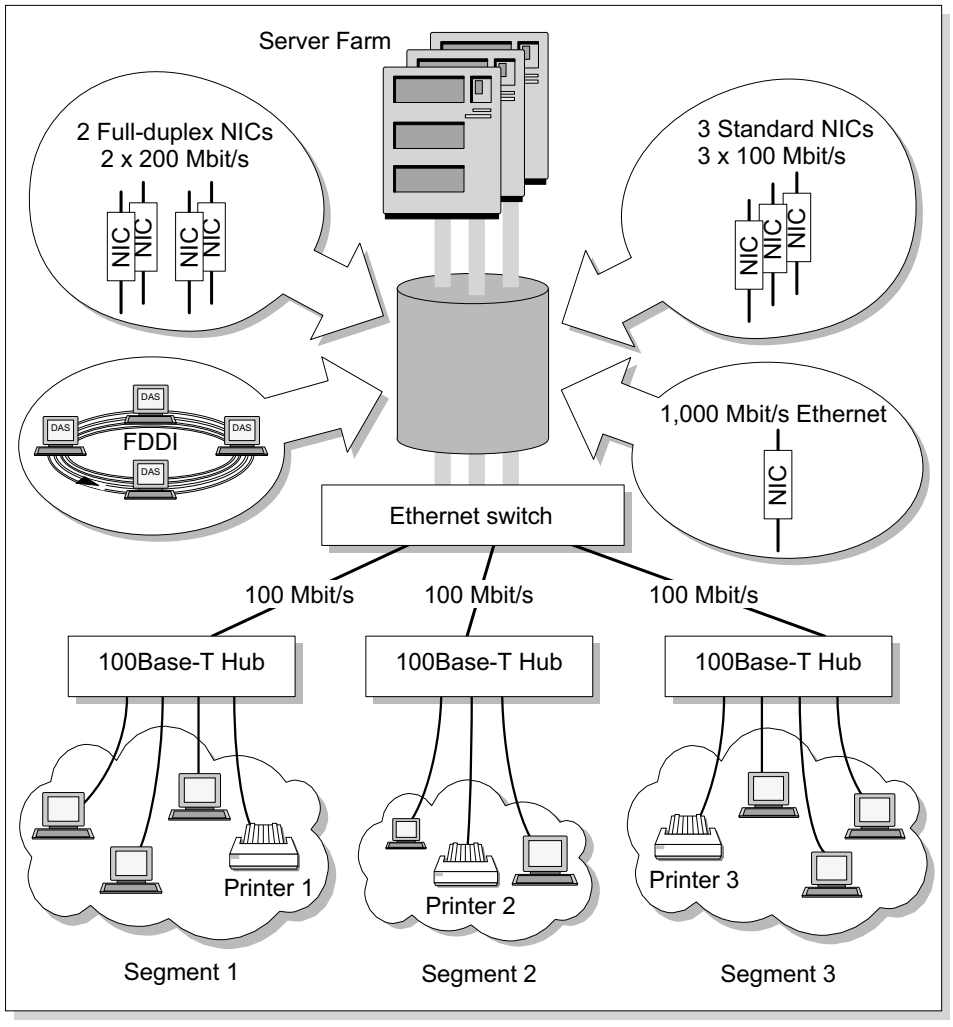

Figure 11.2 LAN switching in environments with asymmetrical load distribution

LAN Switching in Topologies with Symmetrical Load Distribution

If the network traffic is primarily symmetrical, you can distribute the nodes that share a given set of global resources (such as applications or servers) in one segment and connect this segment to the LAN switch. This allows the nodes in that segment, or workgroup, to access their main server while the nodes in another workgroup access another main server, without the two workgroups competing for bandwidth (see Figure 11.1).

LAN Switching in Topologies with Asymmetrical Load Distribution

If the load distribution in the LAN is strongly asymmetrical, as is the case between a server farm and several LAN workgroups, this must be taken into consideration when installing LAN switches. The nodes and segments that handle the majority of the intersegment load must be connected to LAN switches with high-bandwidth interfaces. This can be implemented by means of multiple interface cards in each node (n x 100 Mbit/s Ethernet, n x 16 Mbit/s Token Ring), full-duplex connections, or high-speed technologies such as ATM or Gigabit Ethernet. Most network element vendors support the usage of multiple ISLs (Inter Switch Links) to increase the inter-switch capacity. These bundles of ISLs are called trunks (see Figure 11.2).

11.3 Troubleshooting in Switched LANs

11.3.1 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:

- Was there any change in any hardware or software network component?
- Do the symptoms occur regularly or intermittently?
- When was the first occurrence of the symptom?
- Are the symptoms related to certain applications or limited to certain connections?
- Have any components or PCs been added or removed to or from the network?
- Have any new interface cards been added?
- 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?

11.3.2 Starting the Troubleshooting Procedure

The errors that occur in switched networks are the same as those that lead to problems in LAN topologies. The difficulties in troubleshooting switched environments lie not so much in finding individual error sources as in actually noticing the errors. Because a LAN switch forwards every data packet directly from the port it arrives at to the port that is connected to the destination node, you cannot monitor the data that flows through a switch by simply connecting a protocol analyzer to a switch port. In a switched network, a variety of measurement methods must be used to obtain the same information that is readily available in a shared-medium environment. These methods include evaluating switch management information, connecting a hub as an access point for a network analyzer, port tapping, circuit tapping and switch tapping.

The network management information in LAN switches is carried by Simple Network Management Protocol (SNMP) agents, which store information about switch operating states in the form of standard Monitoring Information Bases (MIBs), such as MIB I, MIB II, Remote Monitoring (RMON) MIB, and proprietary MIBs. The data about a given switch that can be obtained from the switch MIBs include the following:

- Traffic load on each port
- Error rates at each port
- Numbers of broadcasts and multicasts
- Number of discarded packets

In practice these monitoring functions are often implemented only in a very rudimentary fashion; and at peak network loads–just when the information is needed most–they either malfunction or are deactivated. Thus, in addition to evaluating the switch MIBs it is often necessary to use dedicated measurement systems, such as protocol analyzers or LAN probes. There are a number of methods for doing this. Some switches include a monitor port, which allows you to analyze network communications through port tapping, circuit tapping, or switch tapping. In port tapping, the LAN switch is configured to copy all traffic from a given port to the monitor port, where it can be monitored using a protocol analyzer. Theoretically, traffic from several ports can be copied to a monitor port. If the ports being monitored have the same bandwidth as the monitor port, however, there is a danger of losing data packets when traffic is

heavy. The circuit tapping technique is basically the same as port tapping, except that traffic is monitored between two ports rather than on a given port. With switch tapping, the traffic from all ports is copied to the monitoring port. If no monitor port is available on the LAN switch, a mini-hub can be inserted to connect the test instrument to the switch port you wish to monitor (see Figure 11.3).

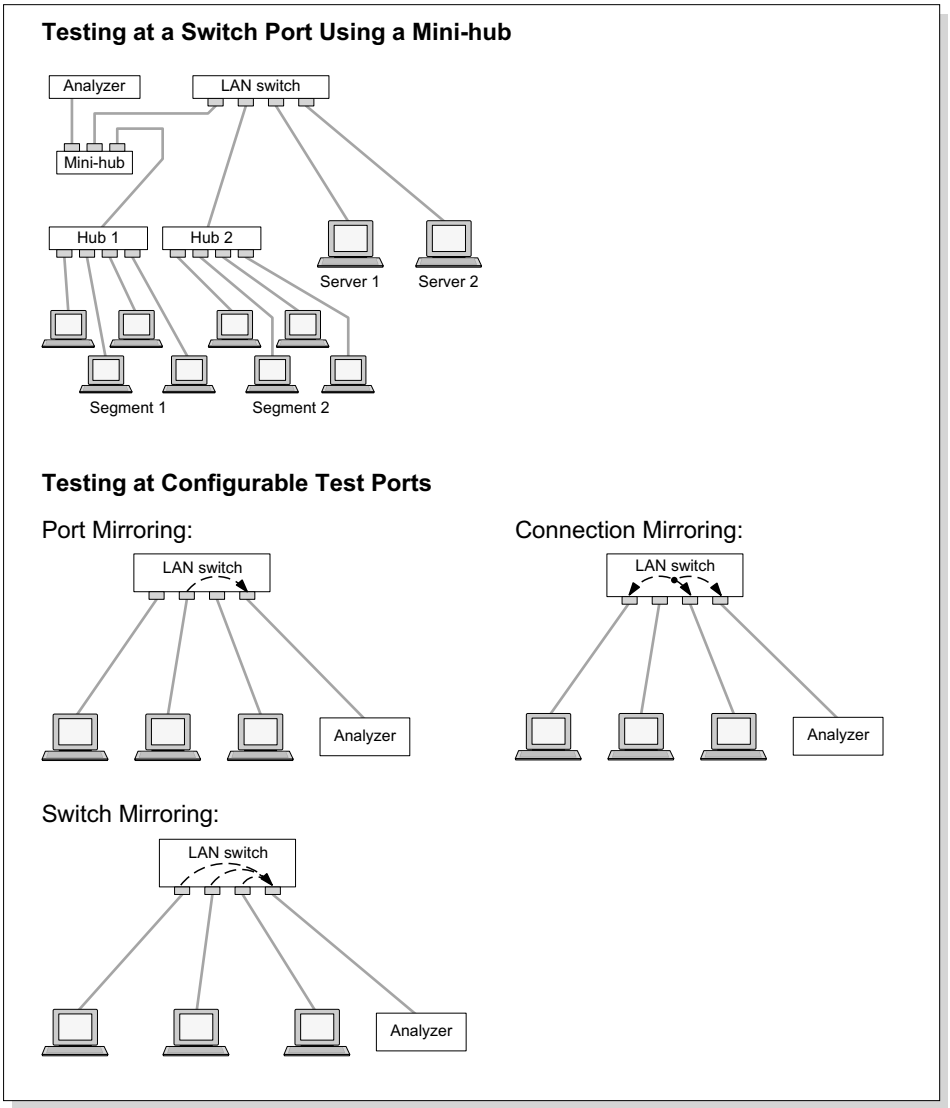

Figure 11.3 Testing methods in LAN switching environments

The first step in troubleshooting a switched LAN is to determine such parameters as network load and packet types (numbers of broadcasts and multicasts, packet lengths, error rates) at the switch ports, or call up this data from the integrated switch management software. The next step is to check the forwarding tables and port configurations in the switch. In many cases this information will narrow down the source of the problem to a particular segment or node. From this point, the troubleshooting procedure continues as described in one of the earlier chapters for the topology in question.

11.3.3 Error Symptoms in Switched Networks

Symptoms that indicate a fault involving a LAN switch include: interrupted connections between segments that are linked by LAN switches, broadcast storms, and throughput problems.

When connections over LAN switches break down, the most common causes are defective cabling, problems with the power supply, and faulty switch hardware. When Ethernet switches are used and extreme peaks occur in the number of broadcasts, bridge loops usually cause the problem. These loops result when the spanning tree protocol or algorithm is not supported or is deactivated. Broadcasts from one segment are then forwarded by a switch to a neighboring segment, and from there by another switch back to the original segment, and so on around the loop (see Figure 11.4).

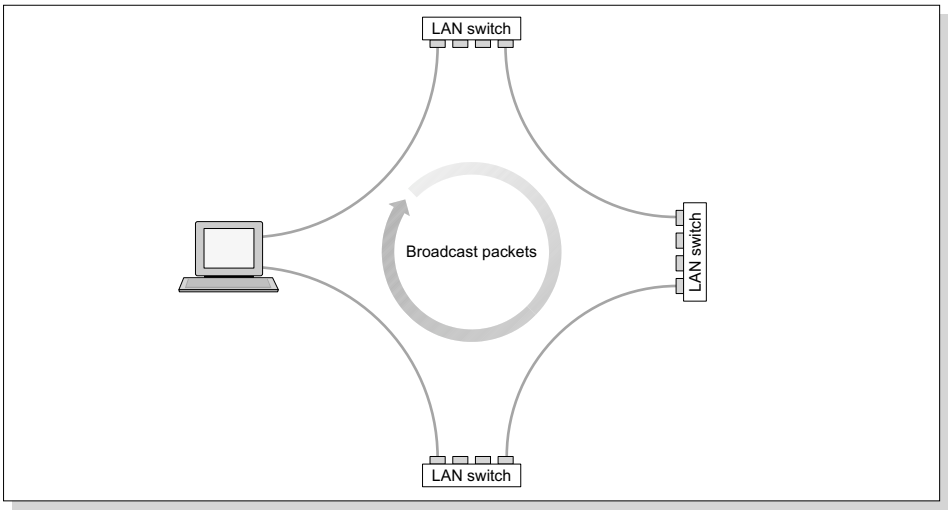

Figure 11.4 Broadcast storms due to bridge loops in switched networks

Bottlenecks at switch ports are usually caused either by design errors in networks with asymmetrical data loads, or by configuration errors at switch ports. For example, in Ethernet switches, ports mistakenly configured as full-duplex lead to frame check sequence (FCS) and alignment errors, while ports mistakenly configured as half-duplex lead to late collisions. In Token-Ring switches, errors commonly result when source routing is not activated or the ring speed configured at the port is incorrect.

11.3.4 Symptoms and Causes: LAN Switching

Symptom: No Connection Between Segments Linked by a LAN Switch

Cause (1): Defective cabling.
Cause (2): Switch power supply failure.
Cause (3): Faulty switch hardware.
Cause (4): Incorrectly configured switch: for example, wrong ring speed (Token Ring), 10 Mbit/s rather than 100 Mbit/s Ethernet, half-duplex rather than full-duplex Ethernet, etc.
Cause (5): Incorrect IP address, subnet mask, or default gateway setting in the switch.
Cause (6): Incorrect VLAN configuration; one of the nodes that cannot communicate is located in a different VLAN.
Cause (7): Source routing deactivated (Token Ring).
Cause (8): Duplicate FDDI address configured for an FDDI switch port.
Cause (9): Duplicate Token-Ring address configured for a Token-Ring switch port.
Cause (10): Defective interface card in the node that cannot communicate.

Symptom: Broadcast Storms

Cause (1): Transmission paths form a loop because the spanning tree algorithm is not activated or not supported.

Symptom: Low Throughput

Cause (1): Poor network design; asymmetrical loads on symmetrical bandwidth at the switch ports.
Cause (2): Incorrectly configured switch ports (10 Mbit/s rather than 100 Mbit/s Ethernet; half-duplex rather than full-duplex Ethernet, etc.).
Cause (3): High number of defective frames generated by a defective switch port.
Cause (4): Cable length exceeds specifications.

The following list summarizes the most frequent sources of problems in LAN-switched networks (in alphabetical order):

- Broadcast storms due to loops through several switches; spanning tree algorithm not activated
- Defective switch hardware
- Duplicate FDDI address configured for an FDDI switch port
- Duplicate Token-Ring address configured for a Token-Ring switch port
- Faulty cable infrastructure; excessive cable lengths (see Chapter 6)
- Faulty network design, asymmetrical traffic over switch ports with symmetrical bandwidth
- Faulty physical switch installation (loose cable, connectors or plug-ins; faulty wiring on the backplane)
- Incorrect Layer 3 switch settings: IP address, subnet mask, or default gateway
- Incorrect router or bridge settings for switch ports operating in router/bridge mode
- Incorrect VLAN configuration; nodes that cannot communicate are located in different VLANs
- Source routing deactivated (Token Ring)
- Switch overloaded
- Switch power supply failure
- Switch settings incorrectly configured: port not activated; wrong ring speed (Token Ring); wrong Ethernet speed; half-duplex instead of full-duplex Ethernet or vice versa

Figure 11.5 The most common causes of problems in LAN switched 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