APPENDIX A

Understanding Token Ring Switching

This appendix discusses several aspects of Token Ring switching and how they relate to the Catalyst 2600. This appendix provides information on the following:

- Switches versus Bridges and Routers
- Source-Route Switching
- Store-and-Forward
- Cut-Through
- Adaptive Cut-Through
- Dedicated Token Ring
- Spanning-Tree Protocol

Switches versus Bridges and Routers

Because the number of stations that can be connected to any single ring is limited, large Token Ring LANs are divided into smaller rings. Furthermore, because stations must contend for the token with other stations on the same ring, attaching fewer stations to a ring gives each one a greater number of opportunities to transmit and receive information. This results in a larger number of rings, or segments.

The traditional method of connecting multiple Token Ring segments is to use a source-routing bridge. For example, bridges are often used to link workgroup rings to the backbone ring. However, the introduction of the bridge can significantly reduce performance at the user's workstation. Further problems may be introduced by aggregate traffic loading on the backbone ring.

To maintain performance and avoid overloading the backbone ring, you can locate servers on the same ring as the workgroup that needs to access the server. However, this makes the servers more difficult to back up, administer, and secure than if they are located on the backbone ring, and limits the number of servers that particular stations can access.

Collapsed backbone routers may offer greater bandwidth than bridges, and can interconnect a larger number of rings without becoming overloaded. However, routers are best suited to WAN environments, and are not the best solution for connecting Token Ring LANs together.

Routers introduce a delay, or latency, between receiving the first bit of a frame and transmitting the first bit on another port. This is because routers operate in store-and-forward mode, buffering each packet into memory before determining the destination of the frame.

As a local collapsed backbone device, a Token Ring switch offers a lower per-port cost and incurs lower interstation latency than a router, and it supports the direct connection of workstations and network servers. Alternatively, a Token Ring switch can be used in conjunction with a router, providing a high-capacity interconnection between Token Ring segments while retaining the wide-area connectivity provided by the router.

Source-Route Switching

Similar to a transparent bridge, the Catalyst 2600 can forward broadcast, multicast, and unicast frames based on MAC address. If, however, you have source-route bridges in your network, the Catalyst 2600 can forward frames based on the routing information field. This dual frame-forwarding technology is called *source-route switching*.

In source-route switching, the switch learns and forwards frames based on source route descriptors for stations that are one or more source-route bridge hops away. A *route descriptor* is a portion of a routing information field (RIF) that indicates a single hop. It is defined as a ring number, bridge number, ring number tuple. When a source-routed frame enters the switch, the switch learns the route descriptor for the hop closest to the switch. Frames received from other ports with the same next hop route descriptor as their destination will be forwarded to that port.

Source-route switching provides the following benefits:

- The switch does not need to learn the MAC addresses of the devices on the other side of a source-route bridge. Therefore, the number of MAC addresses that the switch must learn and maintain is significantly reduced.
- The switch can support parallel source-routing paths.
- The switch can support duplicate MAC addresses if the stations reside on LAN segments with different LAN IDs (ring numbers).

A switch equipped with only source-route switching can be used to microsegment an existing Token Ring network that includes source-route bridges. It cannot, however, be used to replace existing source-route bridges without renumbering the rings.

Store-and-Forward

Store-and-forward is the traditional mode of operation for a bridge and is one of the modes supported by the Catalyst 2600. In store-and-forward, the port adapter reads the entire frame into memory, and then determines whether the frame should be forwarded. At this point, the frame is also examined for any errors (frames with errors are not forwarded). If it is determined that the frame should be forwarded and contains no errors, the frame is sent to the destination port for forwarding.

While store-and-forward reduces the amount of error traffic on the LAN, it also causes a delay in frame forwarding that is dependent upon the length of the frame.

Cut-Through

In cut-through mode, the Catalyst 2600 transfers nonbroadcast packets between ports without buffering the entire frame into memory. Instead, when a port on the Catalyst 2600 that is operating in cut-through mode receives the first few bytes of a frame, it analyzes the packet header to determine the destination of the frame, establishes a connection between the input and output ports, and, when the token becomes available, it transmits the frame onto the destination ring.

In accordance with specification ISO/IEC 10038, the Catalyst 2600 switch uses Access Priority 4 to gain priority access to the token on the output ring. This increases the proportion of packets that can be cut-through, and makes it possible for the Catalyst 2600 to reduce the average inter-station latency.

In certain circumstances, however, the cut-through technique cannot be applied and the Catalyst 2600 must buffer frames into memory.

For example, buffering must be performed when:

- The Catalyst 2600 has two packets to transmit onto the same ring.
- A packet is switched between 4- and 16- Mbps rings.
- The destination ring is beaconing.

Adaptive Cut-Through

With adaptive cut-through mode, the user can configure the switch to automatically use the best forwarding mode based on user-defined thresholds. In adaptive cut-through mode, the ports operate in cut-through mode unless the number of forwarded frames that contain errors exceeds a specified percentage. When this percentage is exceeded, the switch automatically changes the mode of the port to store-and-forward. Then, once the percentage of frames containing errors falls below a specified percentage, the operation mode of the ports is once again set to cut-through.

Dedicated Token Ring

Classic 4- and 16-Mbps Token Ring adapters must be connected to a port on a concentrator. These adapters are also limited to operating in half-duplex mode. In half-duplex mode, the adapter can only be sending or receiving a frame; it cannot do both simultaneously.

Dedicated Token Ring (DTR), developed by the IEEE, defines a method in which the switch port can emulate a concentrator port, thereby eliminating the need for an intermediate concentrator. In addition, DTR defines a new full-duplex data passing mode called Transmit Immediate, which eliminates the need for a token, allowing the adapter to transmit and receive simultaneously.

DTR is particularly useful for providing improved access to servers. A server can be attached directly to a switch. This allows the server to take advantage of the full 16 Mbps for sending and receiving, resulting in an aggregate bandwidth of 32 Mbps.

Spanning-Tree Protocol

If configured to do so, the Catalyst 2600 can use the IEEE Spanning-Tree Protocol to determine the best path for frames when there are multiple path routes in an internetwork. With the Spanning-Tree Protocol, the Catalyst 2600 can exchange bridge protocol data units (BPDU) messages with other Catalyst 2600s and bridges to detect loops and to then remove the loops by shutting down selected interfaces or ports.