Network Connections

In a LightStream 2020 (LS2020) enterprise ATM switch, all traffic that passes over the network is connection-oriented. This means that connections must be established before any traffic (ATM cells) can be transmitted or received. This chapter describes connections in an LS2020 network.

Call Admission Control

The call admission control mechanism determines whether the network can support a requested VCC. An LS2020 network establishes these requested VCCs through one of the following methods:

- Provisioning — This is the method used to establish virtual channel connections (VCCs) for frame relay, frame forwarding, circuit emulation, and ATM UNI devices. It creates ATM, frame relay, frame forwarding, and circuit emulation permanent virtual circuits (PVCs).
- Implicit set-up — This is the method used to establish VCCs for the Ethernet and FDDI services. It creates connections as needed within the network.

Provisioning

Provisioning is the explicit creation of a VCC in which a user specifies its endpoints and other attributes in a configuration database. Provisioned VCCs are called permanent virtual circuits (PVCs).

When you specify the endpoints of a connection, the LS2020 network automatically sets up a pair of VCCs to provide bidirectional communication between the two endpoints. You can configure the traffic management parameters, including bandwidth, separately for each direction of these VCCs. In this document, the term PVC refers to the provisioned VCCs used to provide bidirectional communication.

Implicitly Establishing VCCs

Implicit establishment of VCCs occurs dynamically when a module recognizes the need for a new connection. For example, when a LAN port on an Ethernet interface module does not recognize an incoming packet as belonging to an existing VCC, the system may create a new VCC and route the data across it. After a period of inactivity, the system tears down the VCCs so that they no longer use network resources.

Services Provided by an LS2020 Network

LS2020 provides several methods of connecting external devices to the LS2020 network and passing traffic through the network.

- ATM UNI
- Bridging
- Frame relay
- Frame forwarding
- Circuit emulation

ATM UNI

LS2020's ATM user network interface (UNI) service supplies an ATM interface that allows non-LS2020 ATM networks or other ATM-capable devices to use the LS2020 backbone network. LS2020's ATM UNI interface conforms to the structure and field encoding conventions defined by the American National Standards Institute (see Table 2-1).The ATM UNI is managed using structures in the standard management information base (MIB) and the LS2020 enterprise-specific MIB.

An ATM UNI PVC is defined by two endpoints (ATM ports) on the edges of the network and the local virtual channel identifiers (VCIs) associated with the particular PVC that runs between the source port and the destination port. Figure 3-1 shows two ATM UNI PVCs: PVC 1 and PVC 2.

Figure 3-1 LS2020 Network Containing Two ATM UNI PVCs

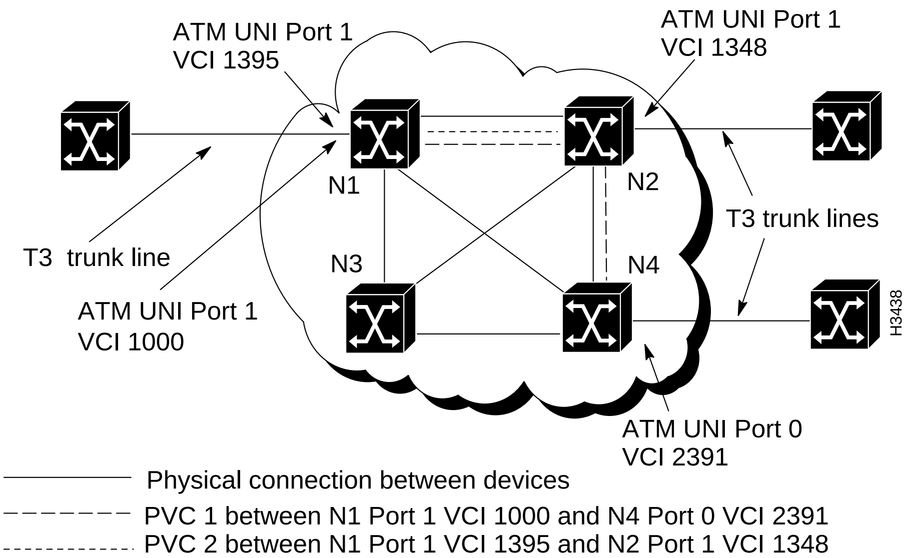
ATM UNI Port 1
VCI 1395
ATM UNI Port 1
VCI 1348
N1
N2
T3 trunk lines
T3 trunk line
N3
N4
ATM UNI Port 1
VCI 1000
H3438
ATM UNI Port 0
VCI 2391
Physical connection between devices
PVC 1 between N1 Port 1 VCI 1000 and N4 Port 0 VCI 2391
PVC 2 between N1 Port 1 VCI 1395 and N2 Port 1 VCI 1348

Since the incoming traffic is already in ATM cells, it is not necessary for the LS2020 switch to segment it into cells. The LS2020 switch at which the cells enter looks at the VCI and determines the PVC on which the traffic should be passed. Each cell is passed through the network on the selected PVC. When they reach the final LS2020 switch in the PVC, the cells are passed out of the LS2020 network on the correct destination port and VCI.

Figure 3-2 shows the ATM cells entering the network and Figure 3-3 shows the cells exiting the network.

Figure 3-2 ATM Cells with Multiple Destinations Entering the LS2020 Network Through an ATM Port

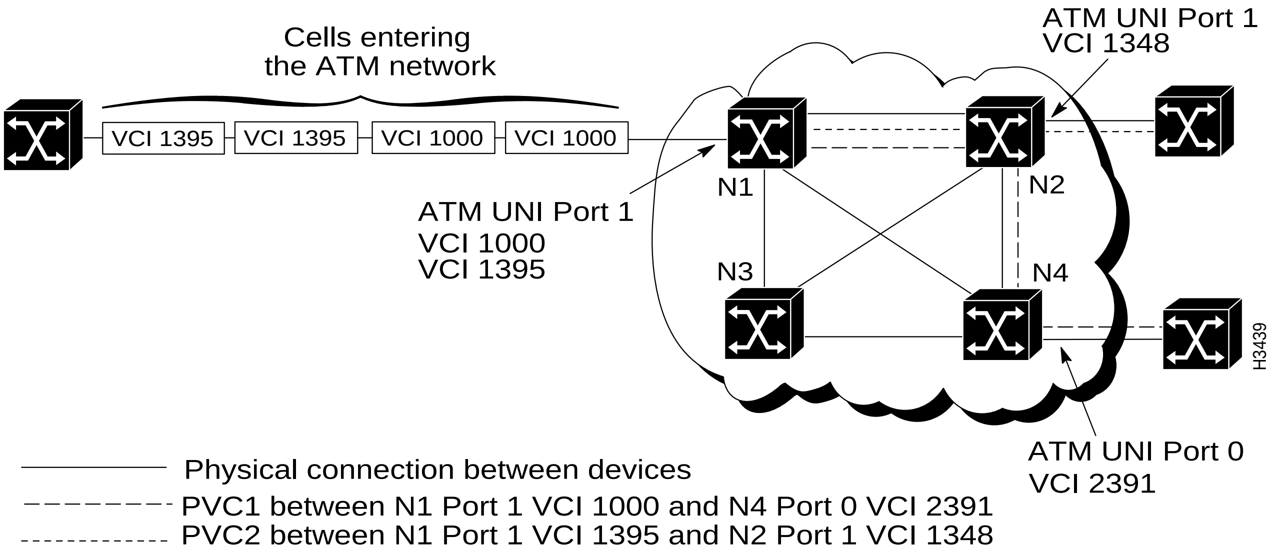

Figure 3-3 ATM Cells Exiting the LS2020 Network

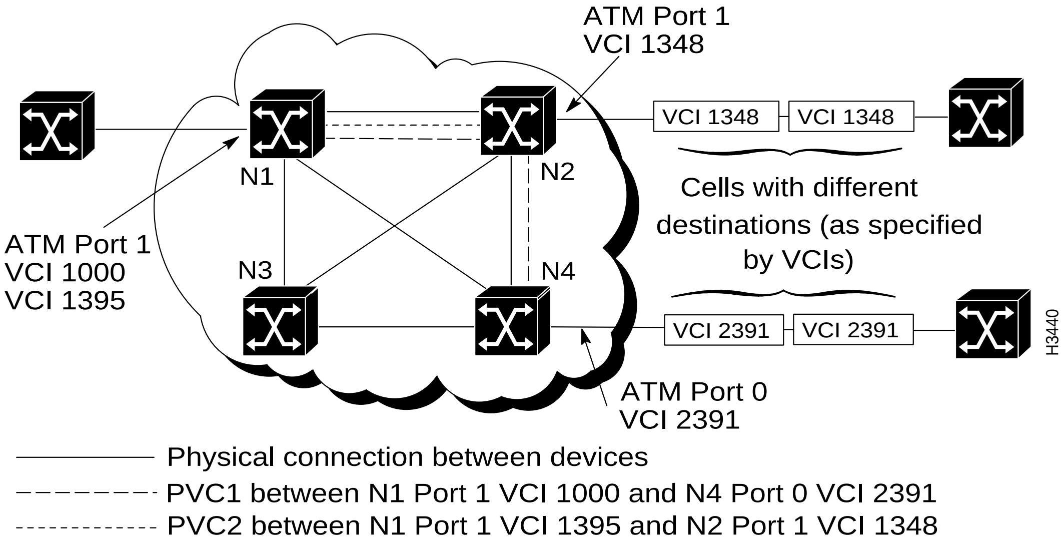
ATM Port 1
VCI 1348
VCI 1348
VCI 1348
N1
N2
Cells with different
destinations (as specified
by VCIs)
ATM Port 1
VCI 1000
VCI 1395
N3
N4
VCI 2391
VCI 2391
H3440
ATM Port 0
VCI 2391
Physical connection between devices
PVC1 between N1 Port 1 VCI 1000 and N4 Port 0 VCI 2391
PVC2 between N1 Port 1 VCI 1395 and N2 Port 1 VCI 1348

Bridging

An LS2020 network supports transparent and translation bridging. Specifically, it supports Ethernet-to-Ethernet, FDDI-to-FDDI, and Ethernet-to-FDDI bridging within a switch as well as across the ATM network.

LS2020 bridging is implemented as a cell internet with underlying ATM features, as described in the subsections that follow. Bridging service also includes the following features, each of which is described later in this section:

- Spanning Tree Protocol
- Custom filtering (layers 2 and 3)
- Static filtering
- Broadcast limiting
- IP packet fragmentation
- VirtualStream
  - AS/QoS
  - HPMS
  - Workgroups

From a user’s point of view, an LS2020 network is modeled as a collection of bridges connected through the ATM backbone with one bridge per LS2020 switch. Externally, all of the bridges in the network appear to be sharing a single broadcast medium on the inside of the ATM network. In the model, each bridge in an LS2020 network has one internal connection to an internal broadcast backbone.

Figure 3-4 shows the LS2020 bridging model.

Figure 3-4 LS2020 Bridging Model

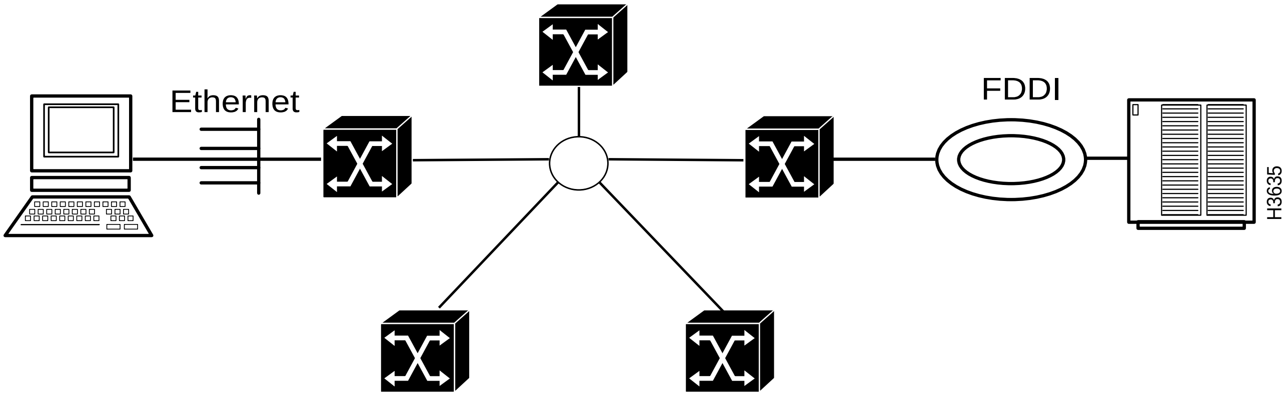

Underlying LS2020 ATM Services

To overlay a connectionless service such as bridging onto an ATM network infrastructure, the LS2020 network automatically manages bandwidth and VCCs for bridged traffic. These operations are invisible to the user of the network.

- All bridged traffic between a pair of LAN ports uses the same ATM VCC. The Maximum Rate for the ATM VCC is limited by the slowest link along the path used by the bridge traffic. To compensate for packet to cell segmentation overhead, bandwidth for bridged traffic is overallocated by a factor of 1.2.

  A stream of packets travelling from one LAN station to another is called a flow. Flows between the same two LAN ports have the same VCC.

  ATM VCCs are set up on demand (implicitly established VCCs). When a frame with a previously unseen destination address arrives, the network sets up a flow to that destination address. If a VCC to the required destination port is already available, the network automatically uses it.
- ATM VCCs are aged out. If a flow remains idle for longer than a specified interval, its ATM connection is torn down. Similarly, if a frame that contains a learned source MAC address has not been received on the port on which it was learned for longer than a specified interval, then the forwarding database removes the station location information and destroys all ATM flows for this MAC address.
- Flows and ATM VCCs are managed automatically. When changes occur to the LS2020 *internal* network topology (for example, a trunk fails) the affected ATM VCCs automatically disconnect. If another suitable path exists in the network, the connections re-establish on demand. When changes occur to the *external* network topology (for example, a station moves from one LAN to another), the affected flows are automatically destroyed. When the station resumes operation, the connections re-establish on demand.

Spanning Tree Protocol

The LS2020 bridging model is structured so that loops cannot occur in a network composed solely of LS2020 switches. However, spanning tree support is required for interoperability with bridges from other vendors. Ports that are configured for bridging implement the spanning tree algorithm defined in IEEE802.1d to eliminate loops that may be caused by external bridges or incorrect cabling attached to multiple LS2020 ports.

The logical network that the algorithm creates is always a spanning tree with these characteristics:

- There are no loops
- There is only one path between any two end stations
- All LANs are connected

If the spanning tree protocol detects a loop, one of the ports on the bridge goes into a blocking state to break the loop. While in the blocking state, the port discards all bridged traffic and stops the process of learning media access control (MAC) address information.

Custom Filtering

The LS2020 bridge supports custom filters on LAN interfaces. You create custom filters and assign them to ports using either the configurator or the CLI. Based on the filters applied to a port, the bridge drops or forwards incoming frames. You could, for instance, prohibit a particular protocol from passing between two ports by creating the appropriate filter for each port.

You create custom filters on a per-chassis basis. Creating a custom filter consists of defining the filter and then assigning the filter to a port or ports. You can assign multiple filters to one port, and you can assign the same filter to multiple ports. Custom filtering is applicable only at inbound ports.

Before you can create custom filters, you must configure a chassis and at least one FDDI, Ethernet, or Fiber Ethernet card with its associated ports.

The LS2020 bridge custom filtering capability for LAN flows supports:

- MAC layer header filtering
- Network layer header filtering (for IP and IPX traffic)
- The association of QoS and HPMS groups with LAN flows on both the data link and network link headers

For information about how to configure custom filters, see the *LightStream 2020 Configuration Guide.*

Static Filtering

The LS2020 bridging software supports static filtering (also called static bridge forwarding as defined in IEEE801.d). Through the configurator or the CLI, you can make static entries in the bridge’s filtering database. You may, for instance, want to make a static entry if you are directing a broadcast to specific ports in order to limit broadcast propagation. You would also make a static entry if you have an end station that only receives traffic, in which case the bridge cannot learn about the station.

For information about how to configure static filters, see the *LightStream 2020 Configuration Guide.*

Broadcast Limiting

To reduce the amount of broadcast traffic on the network, LS2020 bridging software provides the following capabilities:

- Per port broadcast rate limit — You can configure individual port parameters to limit the Maximum Rate at which the port forwards broadcast frames.
- Global addressing distribution — When a bridged port learns a new end station address, it notifies all of the other LS2020 nodes in the network of the location of the end station. This greatly reduces the need to flood unknown packets.
- ARP caching — The LS2020 bridge learns the association of MAC to IP addresses for stations directly attached to the LAN interfaces and propagates this data to the other LS2020 switches in the network. When the LS2020 bridge receives an IP address resolution protocol (ARP) request, it checks its local ARP cache. If the MAC to IP association is known, then an ARP reply is sent to the requesting host instead of flooding the ARP request into the network.
- IP Packet Fragmentation—When necessary, an LS2020 switch can fragment IP packets when bridging the packets between ports that have different maximum transfer unit (MTU) values, such as when going from FDDI to Ethernet.

VirtualStream

The following three sections describe VirtualStream, Cisco System's set of virtual LAN internetworking features: application-specific quality of service (AS/QoS), high performance multicast service (HPMS), and workgroups.

AS/QoS

AS/QoS allows you to assign traffic management attributes to LAN flows (a traffic profile). By associating a traffic profile with the use of a custom filter, you can determine which LAN flows should receive a specific type of service. These types of services are configurable.

You can configure the following traffic profile parameters:

- Rate

  This class of attributes specifies the amount and type of bandwidth. Attributes used in determining the amount and type of bandwidth are

  - Maximum rate
  - Maximum burst
  - Insured rate
  - Insured burst
  - Secondary scale
- Transmit priority
- Cell discard eligibility

Note Traffic profile variables are applied to forwarding filters.

For more information on setting a traffic profile, see the *LightStream 2020 Configuration Guide*.

HPMS

HPMS allows multicast and broadcast flows to be sent across an LS2020 network at wire speed. This feature supports multicast groups. Multicast groups (lists of destination ports) deliver LAN traffic using an ATM point-to-multipoint VCC. Members of multicast groups may be anywhere in the network and do not have to have the same media type. For example, you can have a mix of Ethernet and FDDI ports. You can have multiple multicast groups and a LAN port can belong to multiple groups.

Note Although it is possible to define a multicast group containing non-LAN ports, the multicast LAN traffic will be delivered only to LAN ports.

When you assign a custom filter to a port, it may optionally have associated with it a multicast group, assuming the action of the filter is to forward the matching LAN flow. At this time, a traffic profile must be assigned to it. You are provided with a set of default parameter values for the traffic profile associated with the multicast group, which may be used instead of explicitly configured traffic profile parameters.

When a LAN flow matching that custom filter is detected, then a point-to-multipoint VCC is created from that source port to each of the ports in the multicast group. If the source is also a member of the multicast group, it is not included as a destination of the point-to-multipoint VCC.

You cannot modify the definition of a multicast group while it is assigned to a filter. If you want to define a new multicast group (with a different ID), you need to change the assignment for the filter to the new ID. When you do this, the active flows terminate and rebuild with the new multicast group configuration.

For more information on assigning multicast groups, see the *LightStream 2020 Configuration Guide*.

Workgroups

A workgroup is a collection of LAN ports that are allowed to communicate with each other. By assigning groups of ports to different workgroups, you can provide privacy between groups or limit the impact of one group’s traffic on another. The workgroup membership for a port defines the membership of all the stations attached to it.

You can create workgroups through the configurator or CLI. By default, all ports in the network are assigned to a single workgroup. This makes the default behavior the same as that of an ordinary bridged network.

In an LS2020 network, ports can

- Belong to more than one workgroup
- Communicate only if their workgroup membership lists have at least one workgroup in common

Figure 3-5 shows a sample workgroup configuration.

Figure 3-5 Workgroup Concept

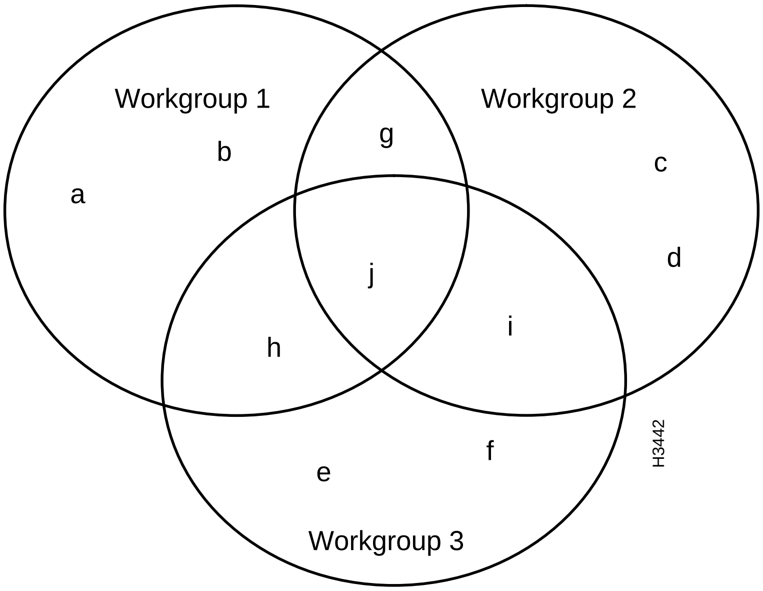

Ports a, b, c, d, e, and f belong to only one workgroup; therefore, they can only communicate with other ports in that workgroup.

Ports g, h, and i belong to two workgroups; therefore, they can communicate with ports in either of those workgroups.

Port j belongs to all three workgroups; therefore, it can communicate with all the other ports.

For more information about workgroup configuration, see the *LightStream 2020 Configuration Guide.*

Frame Relay

The LS2020 supports a frame relay DCE interface to which you can connect routers, packet switches, and other devices that have frame relay DTE interfaces. It also supports a frame relay NNI interface to which you can connect other frame relay switches or networks.

Using the frame relay service, the LS2020 network can accept traffic at a single port and send that traffic to multiple destinations. This is in contrast to the frame forwarding service, where all traffic received on a particular port is sent to one destination port.

A frame relay PVC is defined by two endpoints (frame relay ports) on the edges of the network and the local data link connection identifiers (DLCIs) associated with those endpoints. The LS2020 network uses the DLCI associated with each frame of the traffic to determine its PVC. LS2020 then segments each frame into cells and sends it to its destination.

Figure 3-6 shows three frame relay PVCs. As shown, there can be more than one frame relay PVC between the same LS2020 switches.

Figure 3-6 LS2020 Network with Three Frame Relay PVCs

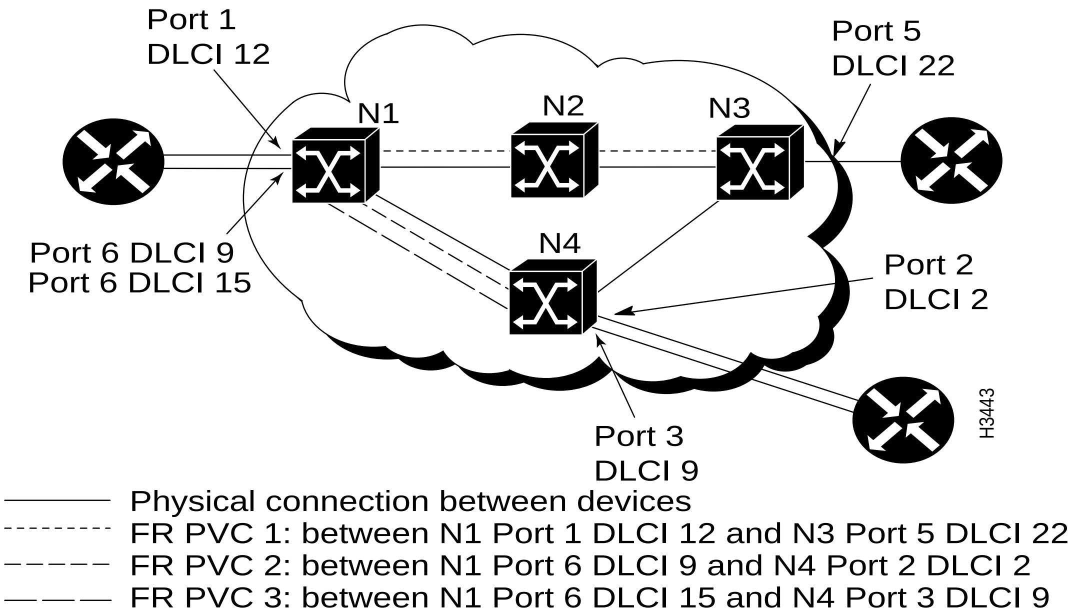
Port 1
DLCI 12
Port 5
DLCI 22
N1
N2
N3
N4
Port 6 DLCI 9
Port 6 DLCI 15
Port 2
DLCI 2
Port 3
DLCI 9
H3443
Physical connection between devices
FR PVC 1: between N1 Port 1 DLCI 12 and N3 Port 5 DLCI 22
FR PVC 2: between N1 Port 6 DLCI 9 and N4 Port 2 DLCI 2
FR PVC 3: between N1 Port 6 DLCI 15 and N4 Port 3 DLCI 9

Figure 3-7 and Figure 3-8 show how frames with multiple destinations are received on one port and passed through the LS2020 network to their correct destinations.

Figure 3-7 Frames with Multiple Destinations Passed Through the Network

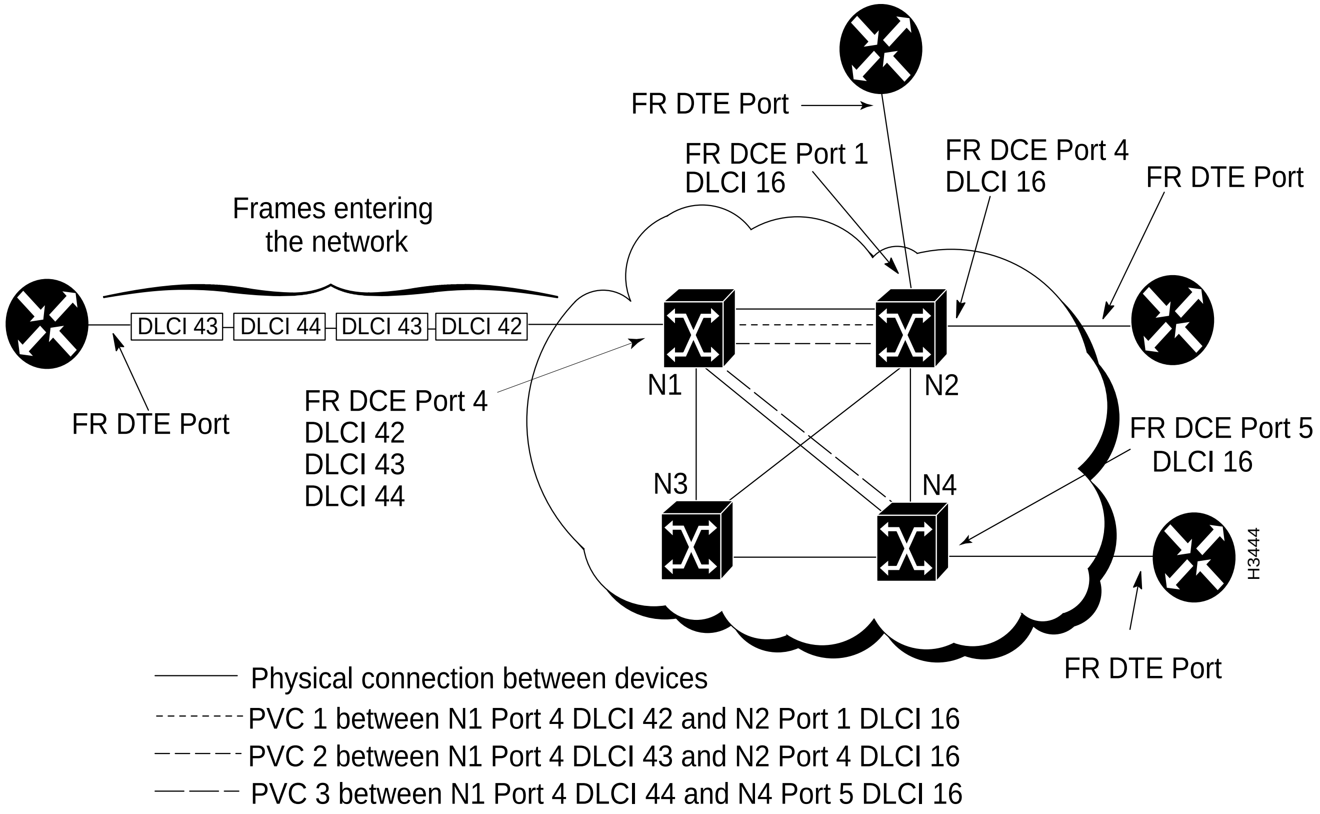

The LS2020 switch at which the traffic enters looks at each frame's DLCI and determines the PVC on which the traffic should be passed. The frame is then segmented into cells. Each cell is passed through the network on the selected PVC. When the cells reach the final LS2020 switch in the PVC, they are reassembled into a frame and passed out of the LS2020 network on the correct destination port and DLCI as shown in Figure 3-8.

Figure 3-8 Frames Sent from LS2020 Switch with New DLCI

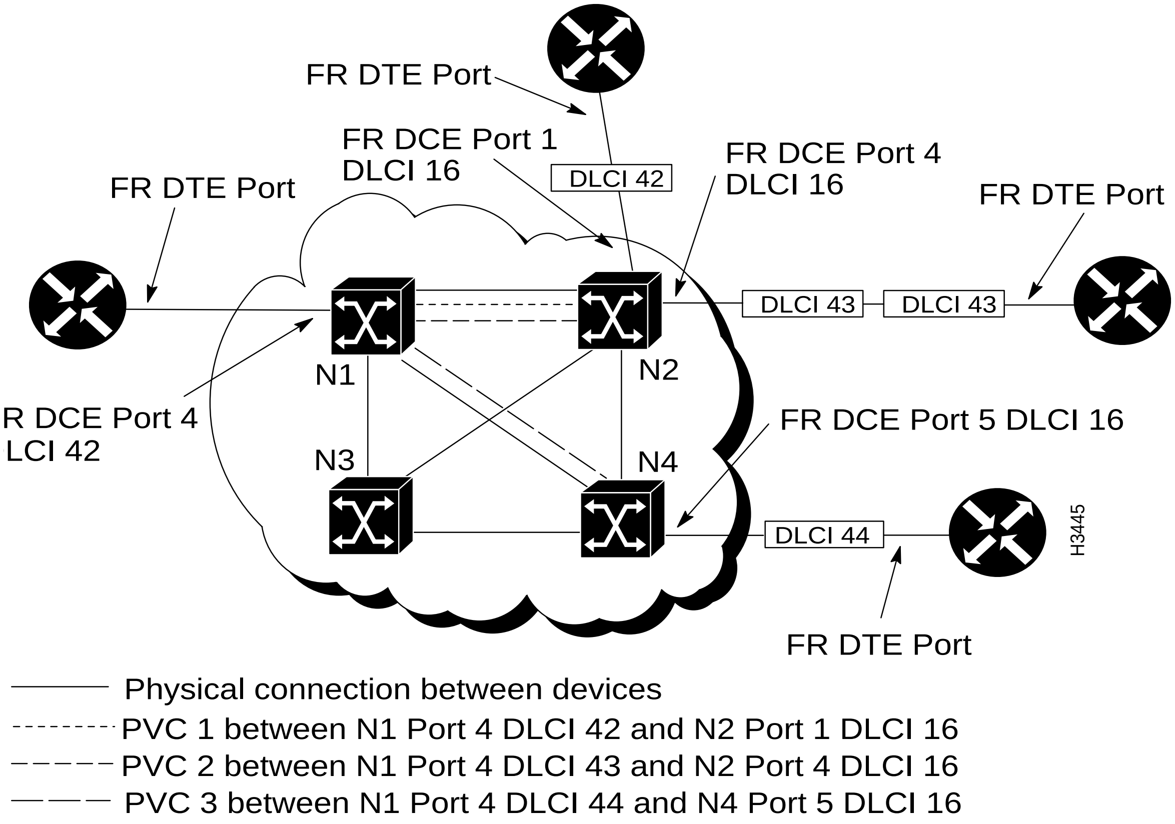
FR DTE Port
FR DCE Port 1
DLCI 16
DLCI 42
FR DCE Port 4
DLCI 16
FR DTE Port
FR DTE Port
DLCI 43
DLCI 43
N1
N2
R DCE Port 4
LCI 42
N3
N4
FR DCE Port 5 DLCI 16
DLCI 44
H3445
FR DTE Port
Physical connection between devices
PVC 1 between N1 Port 4 DLCI 42 and N2 Port 1 DLCI 16
PVC 2 between N1 Port 4 DLCI 43 and N2 Port 4 DLCI 16
PVC 3 between N1 Port 4 DLCI 44 and N4 Port 5 DLCI 16

Frame Forwarding

The frame forwarding service lets you replace direct connections between devices that support HDLC and SDLC with a connection through the LS2020 network; this allows you to connect older devices that do not support frame relay, ATM UNI, or LAN interfaces. For example, you can use the frame forwarding service to connect X.25 packet switching nodes or SNA devices through the LS2020 network.

Frame forwarding PVCs provide a “virtual wire” between two network ports on the edge of the LS2020 network. All traffic that enters the LS2020 network on a particular frame forwarding port is sent through the network to the port on the other end of the virtual wire. All traffic that enters the network on a particular frame forwarding port must have the same destination on the other side of the LS2020 network.

Unlike circuit switched connections which require permanent reservation of the bandwidth needed between the two ports, the frame forwarding function only uses internal network bandwidth when there is an actual frame to be sent, and does not use any bandwidth during the interframe gaps.

A frame forwarding PVC is defined by two endpoints (frame forwarding ports) on the edges of the network. Figure 3-9 shows two frame forwarding PVCs: PVC 1 and PVC 2. The endpoints of PVC 1 are port 1 on N1 and port 3 on N3. The endpoints of PVC 2 are port 2 on N1 and port 6 on N4. There may be any number of LS2020 switches between the endpoints. The LS2020 network selects the best route between the two endpoints and sends the ATM cells through that route.

Figure 3-9 LS2020 Network with Two Frame Forwarding PVCs

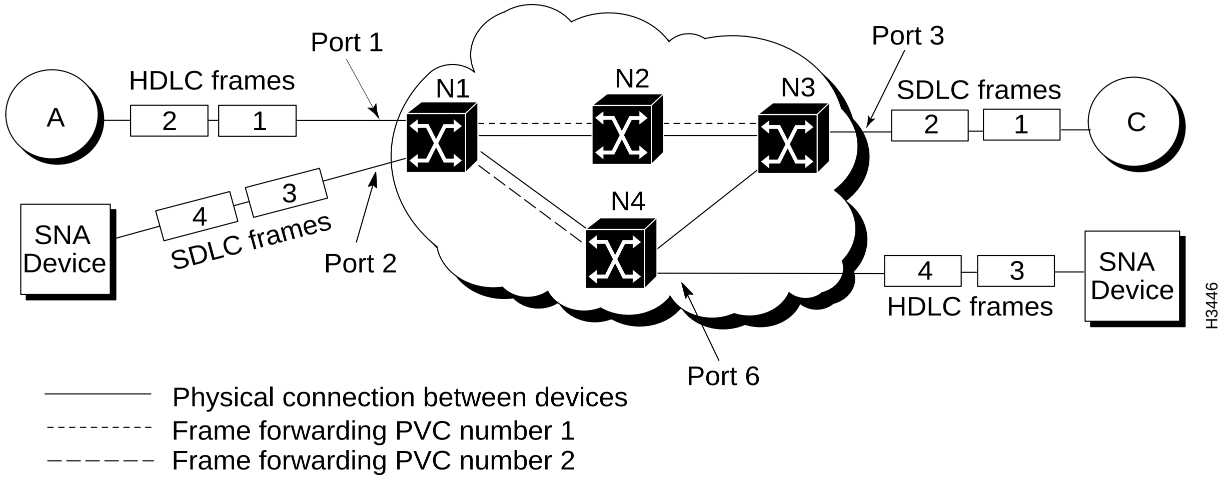

Circuit Emulation

Circuit emulation service lets you interconnect existing T1/E1 interfaces and other kinds of constant bit rate (CBR) equipment. Some CBR services include such features as PBX interconnect, consolidated voice and data traffic, and video conferencing.

With circuit emulation, data received from an external device at the edge of the LS2020 network is converted to ATM cells, sent through the network, reassembled into a bit stream, and passed out of the LS2020 network as shown in Figure 3-10. T1/E1 circuit emulation does not interpret the contents of the data stream. All the bits flowing into the input edge port of the ATM network are reproduced at one corresponding output edge port.

An emulated circuit is carried across the LS2020 network on a PVC, which is configured through the network management system.

Figure 3-10 LS2020 Network with One CBR PVC

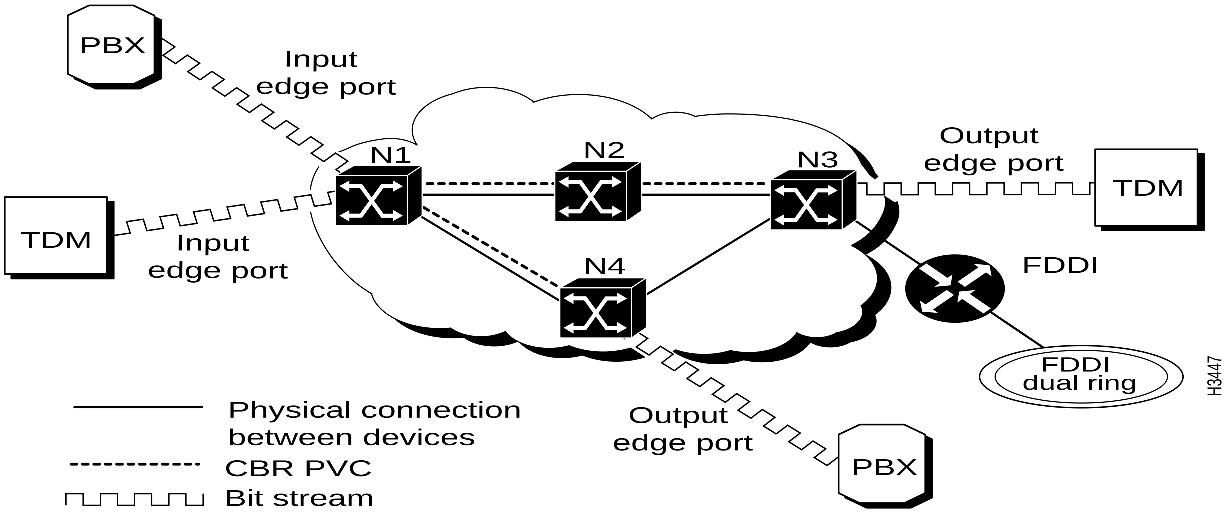

Limited IP Routing for Network Management Traffic

The LS2020 network offers limited IP routing capability to allow for the flow of SNMP, Telnet, and FTP traffic between LS2020 switches and an external network management system (NMS). An NMS station can attach directly to an NP Ethernet port, or can attach through an Ethernet or FDDI edge interface. Every NP has an internal IP address, and the network routing database contains enough information to route incoming IP packets between any NP in the network and any FDDI or Ethernet port, including the Ethernet ports on the NPs.

Note These IP routing services are provided only for monitoring and network management activities. They are not available for carrying user traffic.

Types of Service for VCCs

The LS2020 has a comprehensive traffic management subsystem which supplies various services with configurable alternatives for carrying user traffic. This section presents the configurable aspects of that subsystem using the terminology and concepts of the user interface. For any given VCC, the traffic service is determined by a composite of the independently configurable attributes described below.

Internal Mechanisms

Many of the internal mechanisms that govern the service supplied to individual VCCs are affected by the settings of configurable parameters. These mechanisms are summarized here.

Note For a more detailed explanation of these mechanisms, see the “Traffic Management” chapter.

Transmit Priority

There are five priority levels for servicing cell queues wherever they exist within the network. All cells waiting to be forwarded at a given level are serviced before any at a lower priority. The highest priority is reserved for CBR traffic. The next highest is for internal control traffic. The remaining three are available for user traffic.

Bandwidth Allocation

To manage its bandwidth resources, the LS2020 network tracks two kinds of available bandwidth—allocated and best effort. The allocated bandwidth is increased (or decreased) as a result of call establishment (or teardown). Its magnitude is determined by the requirements of VCCs for a specific traffic capacity under any circumstance. The best effort bandwidth is a dynamic quantity whose magnitude is the sum of the unallocated bandwidth (the difference between the allocated bandwidth and total capacity) and the currently unused allocated bandwidth. It represents statistically sharable capacity for carrying bursty traffic.

Call Admission Control

For the network to support a requested VCC, it must be able to allocate bandwidth along the path, and impose a limit on the amount of traffic that the VCC will be allowed to carry. The allocation of bandwidth is required to meet service goals, and the limitation of traffic is necessary to protect the network from unruly traffic sources.

Traffic Policing

The policing function in an LS2020 network is done at the edges of the network for both frame- and cell-based traffic. It determines whether the traffic is allowed to proceed into the network, and if so, whether it should use allocated or best effort bandwidth within the network. For a given VCC, the policer operates with four static parameters: Insured Rate and burst, and Maximum Rate and burst. These represent, correspondingly, the largest average rate and instantaneous buffering associated with the insured traffic (the type which uses allocated bandwidth), and the largest average rate and instantaneous buffering associated with all traffic. In addition, the VCC policer also uses a dynamic parameter (controlled by the rate-based congestion avoidance mechanism) called Total Rate, which is never lower than the Insured Rate or higher than the Maximum Rate. Traffic which exceeds the Insured Rate and burst parameters, but is within the Total Rate and Maximum Burst parameters, is called excess and uses best effort bandwidth. Traffic which exceeds the Total Rate and Maximum Burst parameters is dropped.

Selective Cell Discard

Although traffic policing is the prevalent mechanism for discarding traffic that the network cannot handle, there can be occasional congestion within the network due to statistical fluctuations which cause local overload. When this happens, cells are discarded according to their cell drop eligibility, for example, cells with higher drop eligibility are dropped before cells with lower eligibility. This cell drop eligibility can be one of three levels, ranging from most to least eligible: Best Effort, Best Effort Plus, and Insured.

Rate-Based Congestion Avoidance

The rate-based congestion avoidance mechanism continuously monitors the best effort bandwidth availability within the network and adjusts the total rate parameter of each VCC policer. Its dual objectives are to maximize use of bandwidth resources (such as trunk lines) while preventing too much traffic from entering the network and causing congestion.

Configurable Attributes

The following attributes affect the operation of one or more of the internal mechanisms previously described for VCCs carrying user traffic. These attributes are explicitly configurable for frame relay, frame forwarding, CBR, and ATM UNI PVC VCCs. In addition, there is a predefined set of attribute values assigned to implicitly established VCCs carrying internal control traffic and bridged Ethernet/FDDI traffic. You can also set these attribute values for LAN traffic using traffic profiles.

Rate Parameters

There are five attributes which control traffic rate aspects of VCC service. These are Insured Rate, Insured Burst, Maximum Rate, Maximum Burst, and Secondary Scale.

The first four of these attributes establish the corresponding traffic policing parameters. The allocated bandwidth, used by the bandwidth allocation and call admission control mechanisms, is the sum of the insured rate plus a fraction (specified by the Secondary Scale) of the difference between the maximum and insured rates.

Principal Service Type

There are two principal service types, Guaranteed and Insured. They share control of the cell drop eligibility mechanism with the rate parameters.

If the rate is within the Insured Rate value, then the traffic is given lowest drop eligibility (Insured), whether the VCC is designated as having the Guaranteed or Insured principal type of service. In fact, the likelihood of any cell dropping of insured traffic is negligible, since all of its bandwidth has been allocated.

For best effort traffic, Insured principal service provides Best Effort (highest) drop eligibility and Guaranteed principal service provides Best Effort Plus (medium) drop eligibility.

Transmit Priority for User Traffic

The transmit priority attribute mainly controls the delay characteristics of traffic on a user VCC and has only two values: zero and one. Zero indicates the lowest of the five priorities maintained by the transmit priority mechanism, and one indicates the second of these. The highest priority is used for CBR traffic and the next is reserved for control traffic VCCs.The middle priority is currently unused.Traffic that is significantly delay-sensitive should use transmit priority one, while traffic that is less delay-sensitive or relatively delay-insensitive should use priority zero.

The transmit priority has a secondary effect on the selective discard mechanism, in that for a given cell drop eligibility, those cells that are assigned a higher transmit priority will be less likely to be dropped than those assigned a lower transmit priority.

Traffic Profiles

A traffic profile is a specific set of values of the configurable attributes. This mechanism allows the AS/QoS feature to allow user-configurable traffic parameter attribute values to be associated with bridged traffic flows.

Behind the Scenes

The LS2020 network performs two important services automatically: neighborhood discovery and global information distribution (GID). These services simplify network configuration and maintain a consistent, network-wide database of routing and address information.

Neighborhood Discovery

A neighborhood discovery process runs on every NP in an LS2020 network. It performs the following tasks:

- Continuously gathers information about the local topology of the network.
- Keeps track of the interface modules that are added or removed from service, in either planned or unplanned ways. Whenever you add a new interface module, the neighborhood discovery process automatically starts up the appropriate software on the network processor. Whenever you remove an interface module, the neighborhood discovery process terminates the associated NP process.
- Determines which NP controls each interface module.

Whenever you add or remove a local resource, the neighborhood discovery process informs the global information distribution (GID) system, which floods the information from NP module to NP module through the network. The neighborhood discovery process also keeps the local GID process informed about who its neighbors are so it can flood information properly.

Neighborhood discovery simplifies the network configuration process and eliminates the need to manually configure some of the interface module attributes in each switch and all the connections to other switches in the network.

Global Information Distribution

The GID system is a service that maintains a consistent network-wide database. All of the switches contribute to the database and all of the switches extract information from it. The GID system ensures that every switch has an up-to-date copy of all the information in the database.

NPs use a flooding algorithm to distribute the global information between neighboring NPs. The flooding algorithm is similar to the one used by the Open Shortest Path First (OSPF) routing system, but the updates are much more frequent. Flooding can occur only between NPs that have established a neighbor relationship, and therefore a communication path, between them. These relationships and communication paths are established, maintained, and removed by the neighborhood discovery process.

The GID system is represented by a process on every NP in the network. Each GID process serves several clients that produce and consume information. A GID process issues an update whenever a client contributes new information. The GID also has mechanisms for quickly initializing a GID database when a new switch enters the network.