C. Noronha
Winter 98-99

EE384A: Network Protocols and Standards
Homework #6 - Solutions
LAN Emulation and IP over ATM

1.) What are the reasons behind the ATM LAN Emulation and classical IP over ATM efforts?

Reasons behind LAN Emulation & Classical IP over ATM

- The interoperability with the existing network technology is the primary reason behind both LAN Emulation and IP over ATM. This would increase the acceptance of ATM technology in today's networks.
- The other reason is to avoid the need to change the applications' software and underlying networking software when bringing ATM technology into the picture. ATM will be hidden from existing software through the addition of emulation functionality. Furthermore, all the functionality that is normally assumed in existing networks (e.g., addressing, connectionless service, and in case of LAN emulation, multicasting and broadcasting) will be normally available.

Differences

The difference between the two is that they target different layers of the existing network architectures. LANE inserts the ATM structure underneath the MAC layer. The Emulation is transparent to the layers above MAC. IP over ATM attempts to replace the DLL (Data Link Layer) completely in the IP network by ATM so that the IP layer directly interfaces with ATM. However, classical IP over ATM as currently specified fails to provide the multicast functionality.

2.) What are the issues behind ATM LAN Emulation and IP Over ATM?

The two principal differences between ATM networks and LANs or IP networks are:

- Difference in addressing;
- Difference in service types (connectionless for LANs and IP, connection - oriented for ATM).

The issues then are the following:

ATM LAN Emulation:

- Address Resolution - Since LAN and ATM use different addressing schemes to identify the entities in their networks, LANE has to resolve the difference and make the two interchangeable with each other.
- Connectionless Services - LAN stations are able to send data without previously establishing connections. So LANE has to provide the appearance of such a connectionless service to the end systems.

- Multicast Services - Due to the connectionless nature of the MAC traffic, it is not difficult to support multicast services. So for emulated LAN, it should be able to provide this service despite of its connection-oriented nature.
- MAC Driver Interfaces in ATM Stations - Some of the commonly used interfaces to access MAC drivers are Microsoft's NDIS (Network Driver Interface Specification), Novell's ODI (Open Datalink Interface) and DLPI (Data Link Provider Interface) which is used in many Unix workstations, have to be supported.
- Emulated LANs - To satisfy the need to configure multiple ATM-attached devices to be a part of an "emulated LAN" which need not have any geographic limitations.
- Interconnection with existing LANs - Connectivity between LAN stations and ATM stations and also connectivity between two LAN stations across ATM have to be supported. This includes connectivity both from ATM stations to LAN stations as well as LAN stations to LAN stations across ATM.

IP over ATM

IP over ATM also faces issues like Addressing, Broadcast and Multicast, and connectionless nature of IP network.

- Addressing - Address resolution between IP addresses and ATM addresses.
- Broadcast and Multicast services - unfortunately, classical IP over ATM doesn't support those services, although extensions to it (RFC 2022) have been proposed to address this.
- Connectionless - IP is also connectionless. IP over ATM has to provide the appearance of such a connectionless service at the interface.

3.) Describe the architecture (the components introduced and their functionality) used in providing the solution to ATM LAN Emulation and IP Over ATM.

LAN Emulation

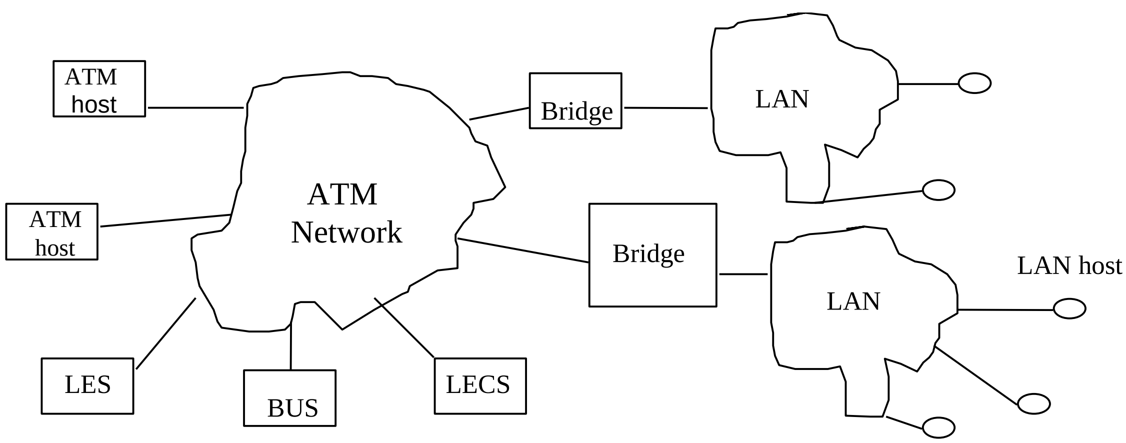

Figure 1. A network configuration consisting of an ATM network, ATM hosts, LANs, LAN host and bridges connecting LANs to the ATM network. Note the LES, BUS and LECs.

- *LEC* (LAN Emulation Client) - It is the entity in end systems (hosts or bridges acting as proxy to hosts on a LAN) which performs data forwarding, address resolution and other control functions.

- *LES* (LAN Emulation Server) - It does the address resolution from the MAC addresses or route descriptors to ATM addresses. LECs have to register themselves with the LES, and query it to resolve a MAC address into an ATM address.

- *LECS* (LAN Emulation Configuration Server) - It maps LECs to different Emulated LANs, based on their physical locations (ATM addresses) or some other criteria.

- *BUS* (Broadcast and Unknown Server) - All broadcast and multicast frames are sent to the BUS to be transmitted to all other stations using the existing PVCs or SVCs. The initial unicast packet for an address that is not cached in the LES is sent to the BUS to be forwarded to the destination.

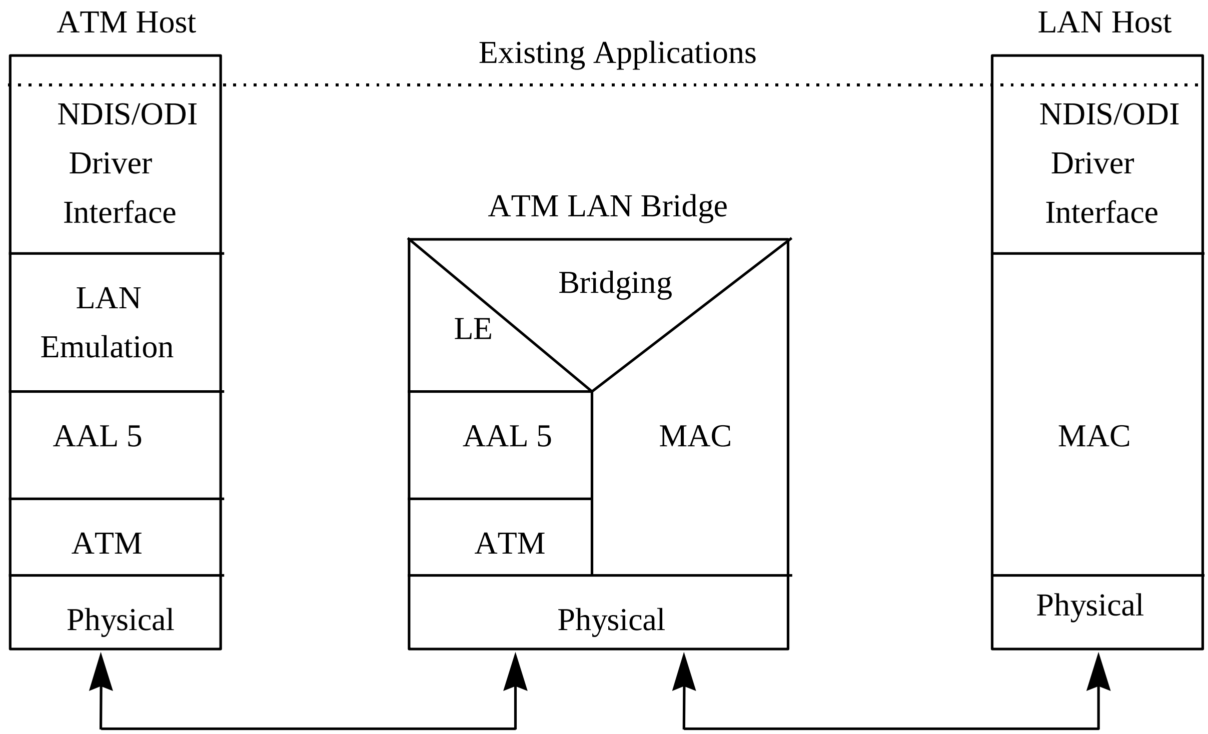

Figure 2. Places where LANE client software resides

IP over ATM

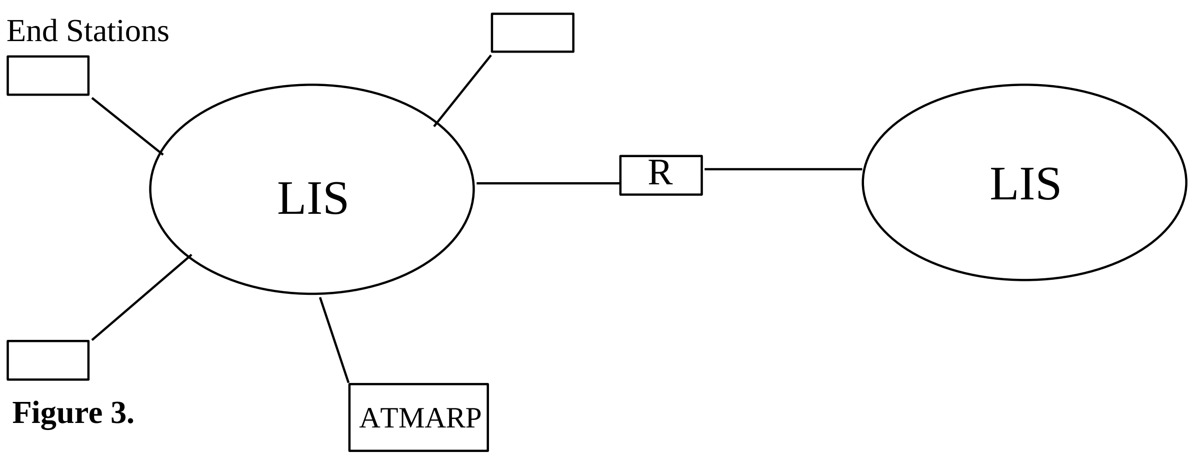

Figure 3.

- ATMARP server in every logical IP subnet (LIS): The server performs the address resolution between ATM addresses and IP addresses.
- End Stations: implement the ATMARP client, as well as InARP.

4.) What are the procedures followed in the operation of the ATM LAN Emulation and IP over ATM?

LAN Emulation

1. Initialization
The first requirement of a LEC is to determine the ATM address of the LECS and to establish a connection to it.

2. Configuration
Once a connection has been established to the LECS, the LEC transmits a message to the LECS with its own ATM address, MAC address and LAN type and frame sizes supported. The LECS will, upon validation of the request, return the address of the appropriate LES and the LAN type and frame size to use for this session.

3. Joining
Once the LEC has the ATM address of the LES, it attempts to create a connection to it (the Control Direct VCC). Once the connection is established, the LEC transmits a message a number of parameters, including its ATM address, LAN type and frame size information, and a MAC address to register.

4. Registration and BUS Initialization
Once the LEC has registered any of its other unicast MAC addresses, it requests the ATM address for the broadcast MAC address. The ATM address returned will be that of the Broadcast and Unknown Server. The LEC then initiates a Multicast Send VCC to this address and accepts the incoming Multicast Forward VCC from the BUS. The BUS knows the address to create the Multicast Forward VCC because all of the connection indications include both the calling and the called party addresses.

5. Data Movement
Once the BUS connections have been established, the LEC may then begin forwarding frames. When a data frame is presented for transmission, an internal table is consulted to see if a connection already exists for that MAC address. If so, then the frame is transmitted on that VCC. Otherwise, the LES is queried for the ATM address that corresponds to the desired MAC address. While waiting for a response and possibly creating a Data Direct VCC, the frames destined for this MAC address may be forwarded to the BUS. In the case of a MAC address residing on the legacy side of a bridge and the bridge not knowing about that station, this will be the only way the legacy station will receive frames. Once the legacy station starts transmitting, the bridge will learn of the station's location and future requests to resolve the address will result in the bridge responding with its own ATM address. Connections are timed out after periods of inactivity in both directions.

After joining and registering itself and also connecting to the BUS, the LEC becomes operational.

Note: There are two timers used in LANE. One is for the cache entries in LEC, and the other is for the cache entries in LES. All the entries in the LEC and the LES have to be refreshed periodically. According to the LANE specification, the time-out period associated with the LEC is shorter than that of the LES. The entries that are not refreshed for a time-out period are dropped from the cache.

As an example of how an existing application might use LAN emulation, consider the case of a simple terminal session using telnet between two ATM hosts on the same IP subnet. Telnet is used to connect from one host to another with the following command:

telnet <address>

Where <address> is resolved into an IP address. In a conventional packet network, the request is routed to the relevant IP subnet, where an ARP may be necessary to locate the MAC address of the host that has that IP address. Once the MAC address has been determined, the packet is sent to that host, where the telnet connection request is processed.

In an emulated LAN, the telnet application and upper layer protocols (down through IP) are unchanged, with the differences showing up at the IP ARP.

Figure 4 shows how the IP ARP request is sent from the source LEC (LEC1) to the BUS to be broadcast to all LECs in the ELAN. Each LEC passes the ARP request up the protocol stack. The destination host responds with an IP ARP reply packet, which its LEC (LEC2) has to eventually send back to LEC1.

The procedure is:

1. LEC1 sends the IP ARP request to the BUS to be broadcast, just like the ARP request being broadcast in the traditional LAN.

2. BUS broadcasts the IP ARP request to the LECs in the emulated LAN.

3. LEC2 receives the IP ARP request and matches its IP to the IP address in the ARP request packet, so it replies by sending IP ARP Reply. There are two cases:
   - If LEC2 already has a direct connection set up to LEC1, the IP ARP reply is sent back over the same connection; or, if the ATM address of LEC1 is known, then the connection can be set up. (This case is trivial; it is not shown in the figure.)
   - If LEC2 does not have a connection or ATM address available, then either of two methods may be used:
     a) The destination (to LEC1) could be treated as unknown, and the IP ARP reply sent to the BUS for broadcast to all LECs (at some expense in broadcast traffic, but minimizing the response time).

     b) Alternatively, LEC2 could send an LE_ARP request to resolve LEC1 to the LES on the Control Direct connection. If the LES doesn't have an entry for LEC1, it would send a request packet to all LECs, LEC1 would reply back to LES, LES then sends the reply to LEC2. After getting the address of LEC1, LEC2 would thus set up a direct connection to LEC1, as shown in Figure 4. This approach minimizes the amount of broadcast traffic at the cost of some additional response time.

In either case, LEC1, the LEC that was the original source of the ARP request, receives its ARP reply packet and passes it up to the higher level ARP process. The ARP process can now pass the telnet packet to the LEC to send to the destination MAC address (that of LEC2). Since the direct connections between LECs are bi-directional, the connection should already exist and the packet will reach its LEC and host rapidly.

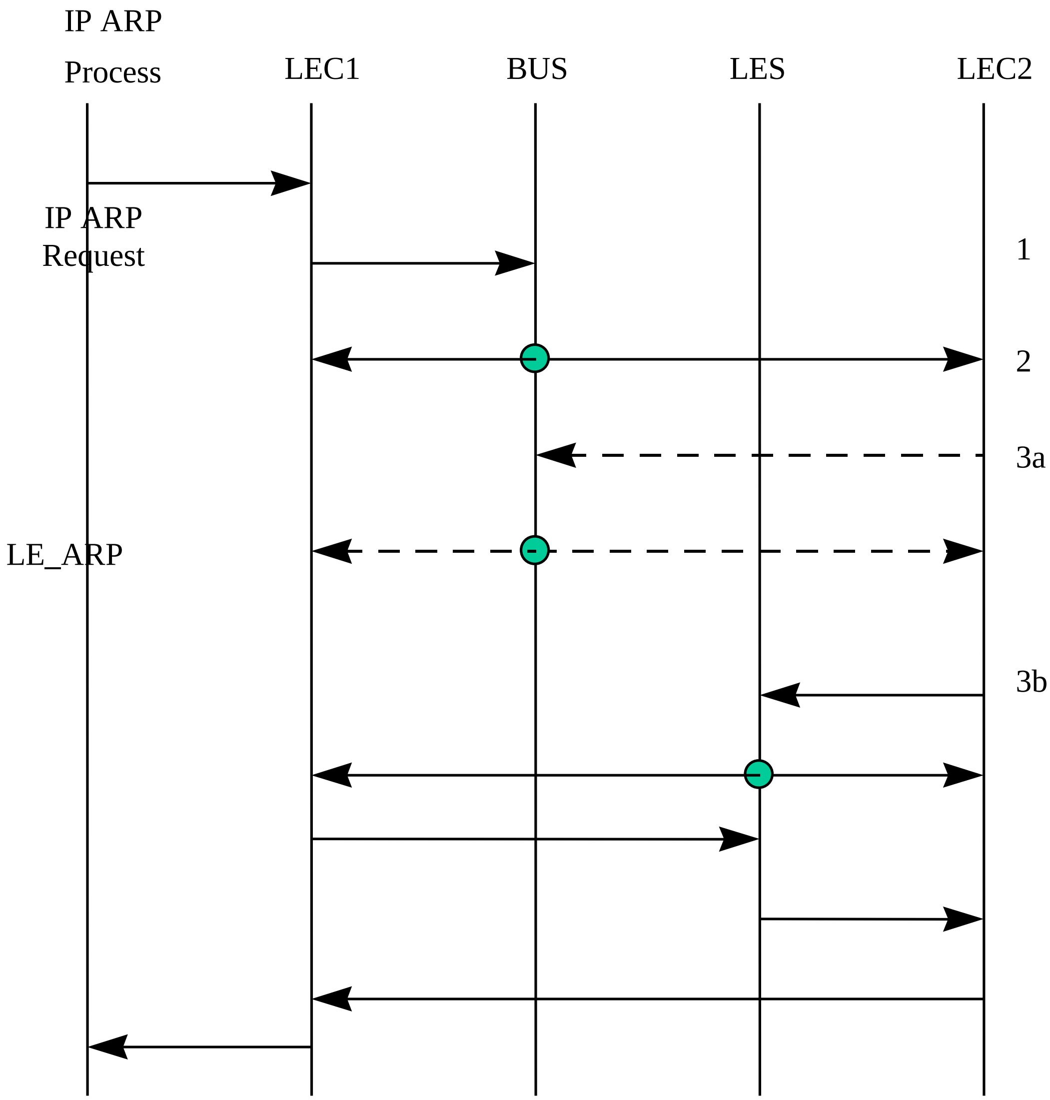

Figure 4. Telnet Across an ELAN

Of course, in an active ELAN, it is quite likely that connections between LECs will already exist, or that the MAC address to ATM address bindings would have been recently cached in the LES cache. Either scenario should significantly reduce the number of transactions necessary to transport a packet across an ELAN.

IP over ATM is intended to be used within a Logical IP Subnet (LIS). All communication within the LIS is performed using direct ATM connections between hosts. Hosts on the LIS communicate with external endpoints through an IP router, even though it may be possible to create a VC directly between the two endpoints.

Within each LIS there must be at least one ATMARP to perform the address resolution between ATM addresses and IP addresses.

To initiate communication between two hosts within the LIS or between a host and an IP router for communication to an external endpoint, a host must first determine the ATM address of the destination to which it will send IP datagrams. There are two possible methods for doing this.

First, it is possible that the host already has a PVC established with the endpoint it is trying to communicate. To determine the IP address of the host on the other side of a PVC, the host sends an Inverse ATMARP message. The endpoint receives the InATMARP message and responds with its IP address.

Second, if a host does not have a PVC established with the desired endpoint, the host may query the ATMARP server for the LIS to determine the ATM address for the endpoint. The ATMARP server is similar to the LES used in LAN Emulation. All the hosts in the LIS register their IP address with this server. When the ATMARP accepts a connection from a station, it will generate an Inverse ATMARP request to determine the IP address for the host. The server will then add the (ATM address, IP address) pair to its map table. When a host wishes to determine the ATM address corresponding to a specific IP address, the host will send an ATMARP request to the ATMARP server over the VC that was initially established. The server will look up the IP address in its map table and return the corresponding ATM address. The host will then establish a separate SVC with the appropriate ATM station and begin transmitting the IP datagrams.

A step-by-step run-through illustrates how ATM connections are set up in a Classical IP over ATM network (see Figure 5). When Client 1 on LIS A wants to send data to Client 2 on the same LIS, the first IP packet sent by Client 1 triggers a request to the ATMARP server. The IP/ATM software module on Client 1 sends an ATMARP request to the ATMARP server, which looks up Client 2's IP address in its table and finds the corresponding ATM address. The server then sends back a response to Client 1 with the matching ATM address. Client 1 then uses the ATM address it just obtained to set up an SVC directly to Client 2. When Client 2 replies with an IP packet to Client 1, it also triggers a query to the ATMARP server.

- To establish a connection using Classical IP over ATM, a sending station (Client 1) sends an ATMARP request to the ATMARP server on the logical IP subnet (1).

- The server reconciles the IP address of the receiving station (Client 2) with its ATM address returns the ATM address to the sending client (2).

- Client 1 then uses the address to establish an ATM link (3).

- When Client 2 receives the first packet, it also sends an ATMARP request to the server to find the location of the sender (4). Client 2 may also use InARP on the link itself and bypass the server.

- Once the server returns the appropriate address, the two clients can communicate directly without further involvement by the server (5).

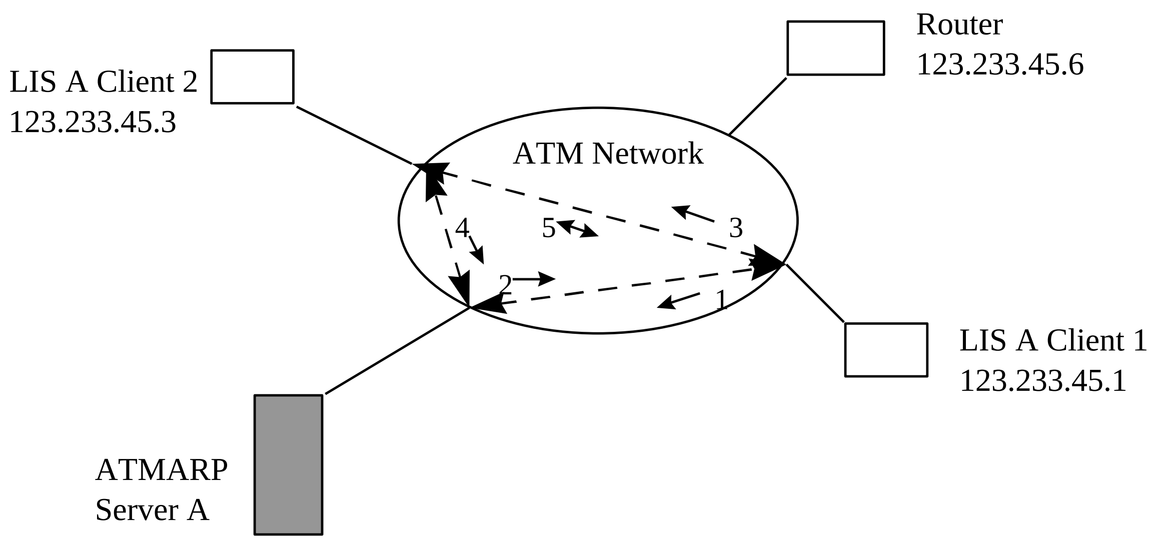

Figure 5. IP over ATM: Step by Step

When it receives the ATM address for Client 1, Client 2 generally will discover that it already has a call to Client 1's ATM address and will not set up another call. Once both clients know the connection, they communicate directly over the SVC, using standard protocols over IP.

A station can have more than one virtual circuit active at a time. A file server or e-mail server may have hundreds of connections within a short period of time, depending on how many client systems it serves. Connections that go unused for a specified amount of time (the default is 15 to 20 minutes) are automatically cleared to recover adapter and ATM network resources.

Classical IP over ATM is specified to work only within an ATM LIS. Work is now in progress to get around this limitation. In the meantime, there are two approaches for communications between LISs: through the router and bypassing the router with multiple LISs.

For two ATM stations on different LISs to be able to communicate, there must be a router on each subnet and an ATMARP server on each LIS. Multiple ATMARP servers may not be needed since one server often can be configured to handle more than one LIS independently.

When a client on one LIS tries to send an IP packet to a client on another LIS, the IP software on the sending client notes that the destination address is off-net and sends the packet to the default gateway (router) on the LIS. This triggers an ATMARP request to discover the router's ATM address, followed by a call setup from the sending client to the router. Once it receives IP packets from the sending client, the router issues an ATMARP request to find the ATM address of the receiving client. It then sets up the call to that client.

What this means is that when end-stations communicate through a router, there are two virtual circuits across the ATM network: one from the sending client to the ATM interface on the router, and another from the router to the receiving station. The sending station segments the IP packets into cells for transmission across the ATM network to be reassembled in the router. The router then makes a forwarding decision based on the IP packet header information, and then re-segments the packets for sending over another virtual circuit to the receiving station. Some routers, sometimes called one-armed routers, can be configured to route packets in and out of the same ATM interface.

Routers deployed in this way provide security filtering between different parts of the network, such as between a campus ATM network and a WAN. But if security isn't an issue, network managers may want to take advantage of the high speed and low latency of ATM to provide the highest performance for users.

One way to do this with Classical IP is to configure shared resource systems, such as file servers, to be members of more than one LIS, so that end-stations are in the same LIS with the resources they use most often. This network design bypasses the routers, eliminating unnecessary packet hops and providing higher performance. This very powerful technique also allows the network administrator to associate applications with separate subnets, each with its own quality of service, priority, and throughput.

Extra Credit Question: Comparison Between LANE V1.0 and LANE V2.0

The most important feature introduced in LANE V2 is the "selective multicast", which just means that it has functionality to have a multicast frame delivered only to nodes that are interested on this particular multicast group. In LANE V1, a multicast frame is sent by the LEC over the Multicast Send VCC, and the BUS repeats it to all LECs in the network either on their Multicast Send VCCs, or, more likely, over the Multicast Forward VCC (since this likely a point-to-multipoint VC, and then it is up to the ATM switch to replicate the cells). The LECs receive the multicast frame and send it up the stack, where it may be dropped if no application wishes to receive it. There are two drawbacks with this approach:

- Waste of bandwidth: unwanted multicast traffic is "eating" bandwidth in the inbound links to all the LECs. Depending on the traffic, this may be a problem – the link may run out of bandwidth.
- The LEC needs to deal with unwanted multicast frames. At a first glance, this situation may seem similar to what an Ethernet-attached device faces in the absence of GMRP: it needs to look at all multicast frames, and pick up only the ones it wants. However, all modern Ethernet chips implement some sort of multicast address filtering; the network interface hardware will weed out the unwanted multicast without CPU intervention. In LANE, however, there is no such filtering: the LEC is, almost always, a software process. Therefore, the LEC CPU utilization will grow with the volume of unwanted multicast traffic and may overwhelm a less-potent LEC.

LANE V2 solves both problems by doing "The Right Thing": if a node is not interested in a certain multicast group, frames addressed to that group should not be sent to it in the first place. This is implemented by the following additional mechanisms:

- Besides registering unicast addresses with the LES, the LEC is now allowed to register multicast addresses as well. This indicates to the network that the LEC is interested in receiving frames addressed to that multicast group.
- A LEC is required to send a LE_ARP request on the group MAC address before sending frames to the group. The LES will respond with an ATM address (normally, the address of the BUS), and the LEC then opens a Multicast Send VCC to that address. In other words, the LEC in LANE V2 has multiple Multicast Send VCCs; one of them is the default, and the additional one(s) are group-specific.
- The BUS is allowed to open multiple Multicast Forward VCCs. One will be the default, and the others are group-specific.

When a LANE V2 LEC sends a multicast packet, it will do so on a group-specific Multicast Send VCC. The BUS will pick up the packet and relay it on a group-specific Multicast Forward VCC, which is composed only of the LECs that have registered interest on that group.

Other changes introduced in LANE V2 are:

Support for LLC-based multiplexing: LANE V1 only supported VC-based multiplexing. In other words, the VCCs created by LANE can only be used by the LANE services; they cannot be shared. LANE V2 introduced the possibility of using LLC multiplexing on the Data Direct VCCs (the Control VCCs and the Multicast VCCs cannot be shared). Using this technique, some pre-opened VCC can be used to carry both LANE and non-LANE traffic. The frame format generally conforms to RFC1483: it starts with the LLC SNAP header (AA-AA-03), followed by the ATM Forum's OUI (00-A0-3E) rather than the 802.1 organization OUI (00-80-C2). This allows complete interoperation. The objective of this feature is to allow sharing of VCCs in situations where these are scarce resources.

Support for multiple Data Direct VCCs at different QOS Levels: In LANE V1, a pair of LECs is joined by a single Data Direct VCC. If multiple VCCs are opened (for example, as a result of both sides creating the VCC more or less simultaneously), both sides switch to the VCC opened by the node with the lowest-numbered ATM address and let the other VCC(s) die out. In LANE V2, a pair of LECs may be joined by multiple data direct VCCs, each at a different negotiated QOS level. A LEC will register with the LES the QOS levels it is prepared to accept, and other LECs can retrieve that information as a part of the LE_ARP process. However, usage of this feature requires applications (on top of the LEC) that are QOS-aware, and that make use of some extended driver API features to indicate which QOS level should be used for the packets to be transmitted. There is a general trend to start supporting QOS from the "ground up" in the networking industry, and this feature is intended for these kinds of applications.

The Verify Protocol: The support for "selective multicast" described above requires the LEC to accept new Multicast Forward VCCs that can be created at any time. This optional protocol allows a LEC to check (verify) whether or not some ATM address is a valid address for the BUS. The LEC must accept a Multicast Forward VCC, but what if some intruder in the network is the source of that VCC? After accepting the connection (and receiving the other side's ATM address, which is provided by the switch), the LEC can ask the LES to verify the other side's address as being a valid BUS. If it is not, the LEC is allowed to drop the connection.

LE_ARP Extensions: LANE V2 extended LE_ARP to support generic TLVs (a TLV is generic variable, encoded as a **T**ype field, followed by a **L**ength field that specifies the number of bytes in the **V**alue field).

Targetless LE_ARP Request: This is a mechanism for a LEC to advertise bindings in its ARP cache. It is really not a request, because it does not have a target (the LEC is not asking anything). If well used, this may reduce control traffic.

No-source LE_NARP Request: This is a mechanism for a LEC to advertise the fact that it ***no longer*** represents a certain MAC address. It is a way to remove a binding. It is intended to reduce unwanted frames at the LEC sending the request.

Support for UNI 4.0 Signalling: LANE V2 tracks the new UNI 4.0 signalling, and the new traffic parameters available there.