4
TCP/IP Over ATM

4.1 Introduction

The ATM network was designed to be used as both a LAN and a WAN. Consequently, it was once envisioned that the ATM network would become a ubiquitous network that would replace most of the current networks. However, because of the large installed base of legacy LANs (e.g., Ethernet) and associated TCP/IP infrastructure and the higher prices of ATM equipment, the current trend is to continue to use legacy LANs as much as possible and to use ATM as a high-speed backbone network to interconnect legacy LANs using TCP/IP. Therefore, TCP/IP and ATM are going to coexist and internetwork in the foreseeable future. This chapter describes various configurations and architectures that have been proposed for running TCP/IP over ATM networks.

4.2 ATM deployment in TCP/IP networks

Running TCP/IP over ATM networks involves, among other things, breaking large, variably-sized TCP/IP packets into small, fixed-size ATM cells.

There are two ways to deploy TCP/IP over ATM networks—ATM to the desktop and ATM in the backbone—depending on the point in the network at which the breaking of the IP packets takes place.

The *ATM to the desktop* approach is based on connecting the hosts directly to an ATM network using ATM network interface cards, as shown in Figure 4.1. The ATM network essentially replaces the legacy LAN, and ATM is configured to operate as an emulated LAN. The segmentation of the TCP/IP packets into ATM cells takes place at the hosts. In the *ATM in the backbone* approach, the hosts in an enterprise are connected using legacy LANs, and a gateway connects the enterprise to the ATM network using a multiprotocol router, as shown in Figure 4.2. In this case, the segmentation of the TCP/IP packets into ATM cells occurs at the gateway, and the hosts are not aware of the presence of ATM.

For IP hosts or routers interfacing directly with ATM, new mechanisms are necessary to support IP traffic over ATM. Section 4.3 explains two standard mechanisms for running IP over ATM.

4.3 Running IP over ATM

There are a number of differences between legacy LANs and ATM networks. For example, legacy LANs use connectionless shared media, and hence no connection setup is required before transmitting a packet. In contrast, ATM is based on setting up point-to-point (or point-to-

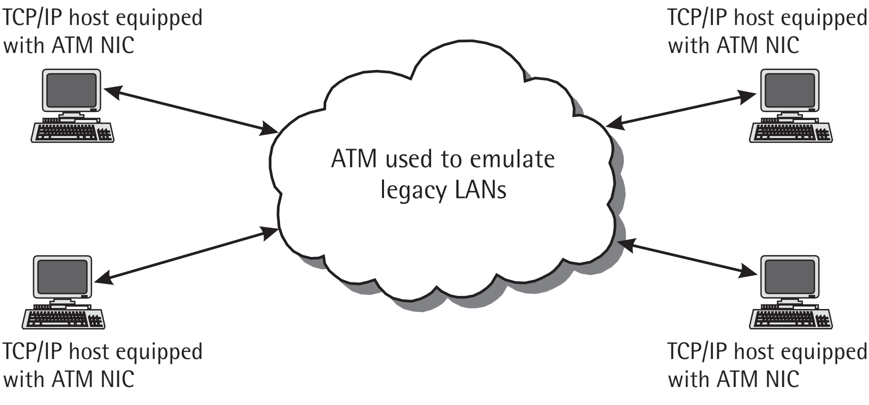

Figure 4.1 TCP/IP hosts connected to an ATM LAN using the ATM to the desktop approach.

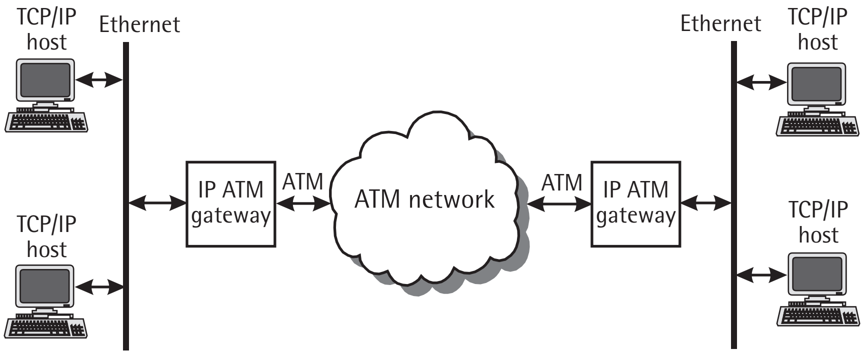

Figure 4.2 TCP/IP networking using the ATM in the backbone approach.

multipoint) connections before data can be sent over a connection. Therefore, mechanisms are necessary to run IP over ATM. The two standard approaches for running IP over ATM, LAN emulation (LANE) and classical IP over ATM, are described in Sections 4.3.1 and 4.3.2, respectively.

4.3.1 LAN emulation (LANE)

LANE was proposed by the ATM Forum [1, 2] to allow hosts on legacy LANs to *transparently interoperate* with hosts directly connected to ATM switches or ATM networks. With LANE, the ATM network effectively becomes an extension of an existing legacy network where an ATM host can be treated as a legacy host and can be accessed using a legacy MAC address. Software that runs on legacy LAN hosts also runs on ATM hosts without any changes. Figure 4.3 illustrates the concept of transparent interoperability between an existing legacy LAN (Ethernet) and a newly deployed ATM LAN using LANE.

LANE allows ATM hosts to run existing legacy LAN software without making any changes. However, since legacy LAN software uses broadcast to communicate with other hosts on the LAN, a mechanism is necessary in LANE to support broadcast over VC-based ATM networks. Sections 4.3.1.1 and 4.3.1.2 describe the two possible approaches for communication between hosts connected through an ATM LANE.

4.3.1.1 Many point-to-multipoint connections

In this approach, each host opens a separate VC connection to every other host in the emulated LAN resulting in a fully connected mesh of VCs, as

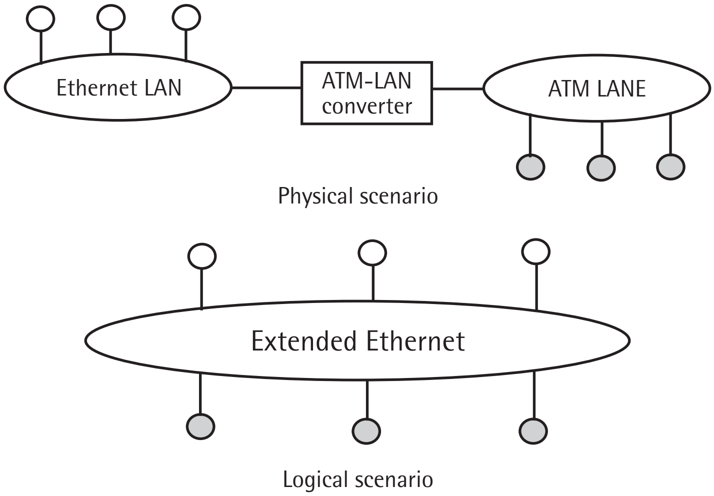

Figure 4.3 Interoperability between a legacy LAN and an ATM LAN using LANE.

shown in Figure 4.4. However, as the number of hosts in the emulated LAN increases, this method is not scalable since the number of connections that must be updated or opened every time a host joins or leaves the LAN becomes excessive.

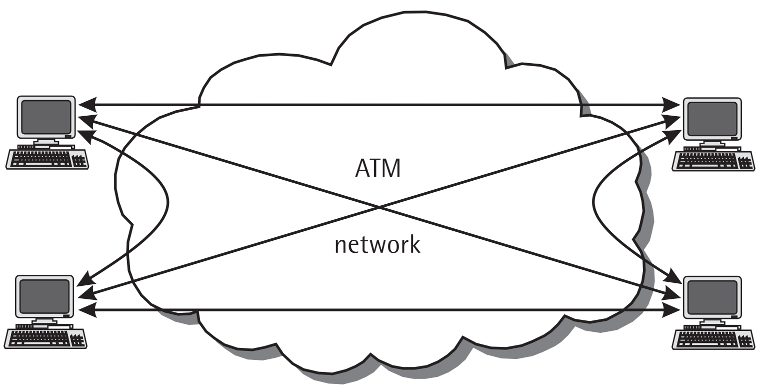

Figure 4.4 Many point-to-multipoint connections to implement an emulated LAN.

4.3.1.2 One point-to-multipoint connections

The scalability problem in the multipoint-to-multipoint approach gave rise to a second approach whereby all the hosts are connected to a multicast server as shown in Figure 4.5. A host aiming to broadcast data sends the data to the multicast server, which then sends it to all the hosts in the LAN. Here, the server has a point-to-multipoint connection, and the hosts only have point-to-point connections with the server.

The ATM Forum's LANE architecture, which consists of a number of hosts running the LANE client (LEC), uses the one point-to-multipoint approach. Each LEC in the emulated LAN is connected to three servers as shown in Figure 4.6. When a LEC joins an emulated LAN (ELAN), it sets up connections with all three servers. The functions of the three servers are described as follows.

- *Broadcast and unknown server (BUS):* A LEC that wants to broadcast a packet to all the hosts sends the packet to the BUS, which then forwards it to all the hosts. The BUS is also used when the sending LEC does not know the address of the receiving LEC.
- *LANE server (LES):* The LES is used to resolve the mapping between MAC addresses and ATM addresses. If a LEC does not have the ATM address of another LEC with which it wants to communicate, the LEC sends out a LANE address resolution protocol (LE_ARP) request to the LES, which either returns the ATM address from its

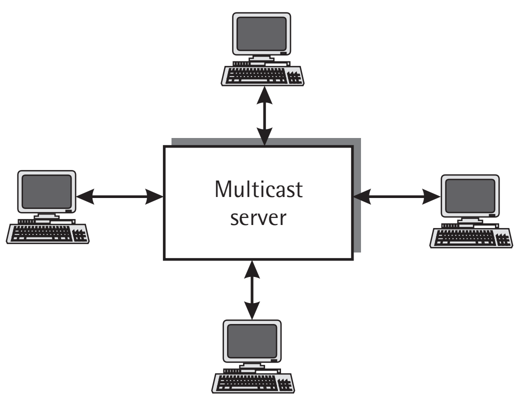

Figure 4.5 A number of hosts connected to the multicast server.

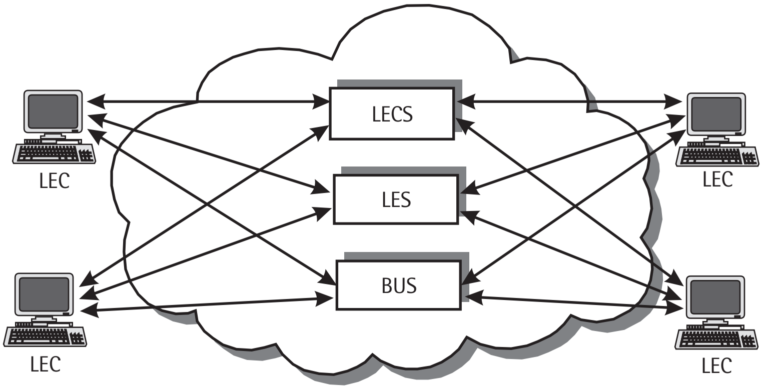

Figure 4.6 Servers required in the ATM Forum's LANE architecture.

cache or broadcasts the LE_ARP request to all LECs that have registered with the LES. There is one LES per ELAN.

- *LANE configuration server (LECS):* The LECS maintains a database of configuration information for the ELAN. For example, when a LEC joins the ELAN, the LECS informs the LEC of the ATM address of the LES with which it should register. To distribute the load, there could be a number of LECS in a distributed fashion within an ELAN. The LECS is configured by the systems administrator.

The primary advantage of LANE is that all existing protocols, such as IP, IPX, and Netbeui, can be supported readily without modification. The limitations of this approach, of course, are the lack of access to ATM signaling and the fact that there is no guarantee of QoS for the higher layers.

4.3.2 Classical IP over ATM

The classical IP over ATM method was standardized by the IETF [3] as a mechanism to run IP over ATM networks where IP treats ATM as a new technology (with no emulation of Ethernet or token ring) and has access to all ATM functions and features including the guaranteed QoS. In this architecture, the IP hosts that are connected to an ATM network form a logically independent subnet (LIS) with a single ATM address resolution protocol (ARP) server to resolve the translation between IP and ATM addresses as shown in Figure 4.7.

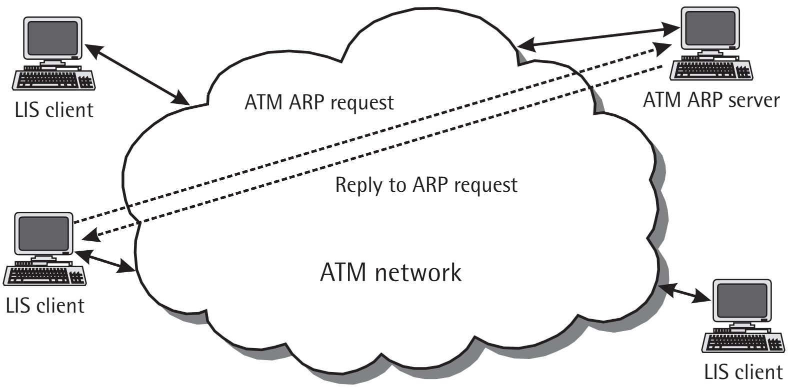

Figure 4.7 Connecting a number of hosts using classical IP over ATM.

When a new IP host (LIS client) is connected to the network, it informs the ATM ARP server of its IP and ATM addresses, which are configured by the system administrator. When a LIS client wants to determine the ATM address of a second LIS client, it sends a request to the ATM ARP server, which returns a reply to the ARP request informing the LIS client of the ATM address. The LIS client can then communicate with the second LIS client using the ATM address obtained from the ATM ARP server. If a host is not part of the LIS, the message has to be sent to a router that will have to forward the packet to the appropriate subnet. Classical IP over ATM is an IP-only solution; similar but separate solutions are needed to run IPX or Netbeui over ATM.

4.4 Encapsulating IP packets into ATM cells

Section 4.3 discusses standard architectures to connect TCP/IP hosts using an ATM network. This section shows how the IP packets are encapsulated into ATM cells so that the underlying ATM network can carry them.

Two different methods of encapsulating connectionless data over an ATM network using AAL5 were suggested in RFC 1483 [4]. The first method, called LLC/subnetwork attachment point (SNAP) encapsulation, multiplexes a number of connectionless data streams (belonging to different protocols) over a single ATM VC using a logical link control (LLC) header. The second method uses individual ATM VCs to carry packets belonging to different protocols; hence, it is called VC-based multiplexing. Since both routers (or hosts) working at layer 3, and bridges (or LANE hosts) working at layer 2, can be connected to ATM networks, the two standard encapsulation methods use slightly different formats for carrying layer 3 (routed) and layer 2 (bridged) protocols. Sections 4.4.1 and 4.4.2 describe LLC/SNAP encapsulation and VC-based multiplexing for both routed and bridged protocols.

4.4.1 LLC/SNAP encapsulation

The LLC/SNAP multiplexing technique is used when several protocols are carried over the same ATM VC. Sections 4.4.1.1 and 4.4.1.2 discuss the AAL5 payload format for the routed and bridged protocols.

4.4.1.1 Routed protocols

To enable the destination to differentiate between the different protocols that are being carried over the same VC, the source appends a three-octet LLC header and a five-octet SNAP header; they are carried in the AAL5 payload of the transmitted data. Figure 4.8 shows the AAL5 payload for encapsulating an IP packet.

The SNAP header consists of a three-octet Organizationally Unique Identifier (OUI) and a two-octet Protocol Identifier (PID). The OUI identifies the organization that administers the meaning of the codes used to identify different protocols in the PID field.

4.4.1.2 Bridged protocols

The payload for bridged protocols differs slightly from the payload of the routed protocols described in Section 4.4.1.1. As an example, the

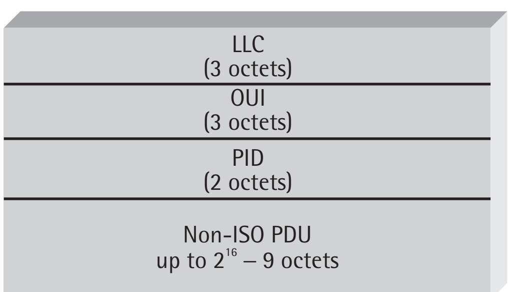

Figure 4.8 AAL5 payload for encapsulating an IP packet.

AAL5 payload for the bridged Ethernet is shown in Figure 4.9. A value of 0xAA-AA-03 in the LLC field indicates the presence of the SNAP header consisting of the OUI and PID fields. A value of 0x00-80-C2 in the OUI field represents the organizational code for the IEEE 802.1 working group. A PID value of 0x00-01 represents the presence of a LAN Frame Check Sequence (FCS) field and a bridged Ethernet protocol, whereas a PID value of 0x00-07 represents a bridged Ethernet protocol with no LAN FCS field. The LAN FCS field contains the FCS of the original PDU.

Since the PDUs of all the different protocols are carried over the same VC in the LLC/SNAP encapsulation method, the method is suitable when it is not convenient or possible to dynamically open a large number of VCs without incurring significant cost.

4.4.2 VC-based multiplexing

The VC-based multiplexing scheme is used when a large number of VCs can be opened without incurring significant cost. In the VC-based multiplexing scheme, a host opens a number of VCs to the destination, each VC being used to carry packets for a different protocol as shown in Figure 4.10. The destination host differentiates between the PDUs of the different protocols by the different VC numbers. There is no overhead (such as the LLC overhead used in the LLC multiplexing scheme) to differentiate packets from different protocols.

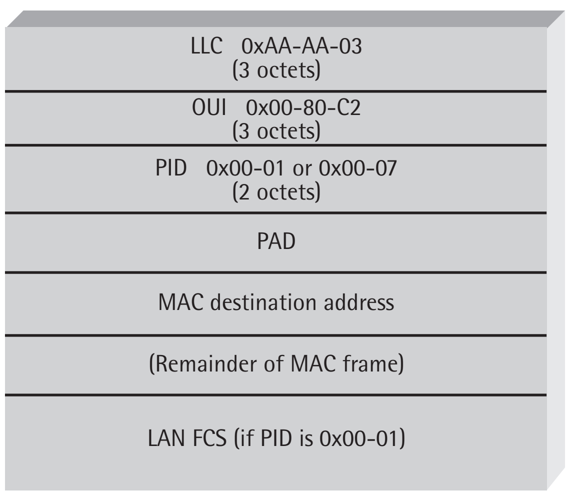

Figure 4.9 AAL5 payload for the bridged Ethernet.

The advantage of this method is that it requires minimal bandwidth to transmit data and only a small overhead to process headers. This scheme is better than the LLC multiplexing scheme when a large number of VCs can be dynamically opened very quickly without incurring much cost. It is anticipated that this scheme will prevail in the private ATM networks. Sections 4.4.2.1 and 4.4.2.2 describe techniques to implement the VC-based multiplexing scheme in the routed and bridged protocols.

4.4.2.1 Routed protocols

Since the destination does not need to differentiate between PDUs carried in a particular VC, the AAL5 payload field using VC-based multiplexing and routed PDUs will consist of only the IP PDU as shown in Figure 4.11 for IP over ATM.

4.4.2.2 Bridged protocols

When PDUs of bridged protocols are carried using the VC-based multiplexing scheme, the payload format is similar to when they are carried using the LLC encapsulation scheme except that the LLC, OUI, and PID fields are no longer required. Figure 4.12 shows the payload format for a bridged Ethernet carried using the VC-based multiplexing scheme.

4.5 Summary

This chapter discusses techniques to connect TCP/IP hosts using ATM. The first half of the chapter describes two ATM adoption approaches—ATM to

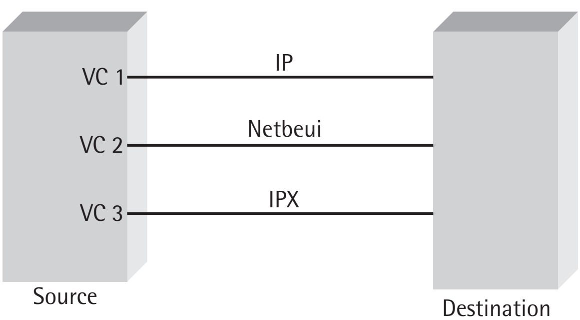

Figure 4.10 Two hosts exchanging multiprotocol data over ATM using VC-based multiplexing.

Figure 4.11 AAL5 payload format for encapsulating IP datagrams in the VC-based multiplexing scheme.

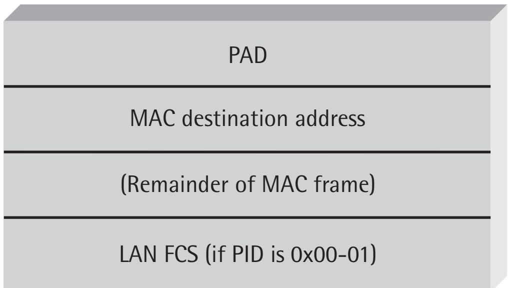

Figure 4.12 Payload for a bridged Ethernet in the VC-based multiplexing scheme.

the desktop and ATM at the backbone—and techniques to implement them. The second half of the chapter discusses techniques to break and encapsulate large variably-sized IP packets into small, fixed-size ATM cells.

References

[1] ATM Forum, "LAN Emulation Over ATM Specifications," Version 1, ATM Forum Specifications, February 1995.

[2] Sun, H., K. Huang, and L. Li, "Supporting IP on the ATM Networks: An Overview," *Computer Communications,* Vol. 21, No. 11, August 1998, pp. 1020–1029.

[3] Laubach, M., and J. Halpern, "Classical IP and ARP Over ATM," *Internet RFC 2225*, IETF, April 1998.

[4] Heinanen, J., "Multiprotocol Encapsulation Over ATM Adaptation Layer 5," *Internet RFC 1483,* IETF, July 1993.