1. Classical IP over ATM (CLIP)

Classical IP over ATM was defined by the IETF in January 1994. The specification RFC 1577 for transmitting IP datagrams and ATM Address Resolution Protocol (ATMARP) requests over ATM Adaptation Layer 5 (AAL5) [RFC 1577]. The term *Classical IP* is used for the network model, where the networks nodes are organised to the subnetworks, which share the same IP prefix and address mask. The Address Resolution Protocol (ARP) is used for mapping IP addresses to the appropriate MAC addresses and communication between the subnetworks is made through routers. The specification considers these problems in a case when the local LAN segments are replaced by ATM technology.

The model CLIP treats the ATM network as a number of separate logical IP subnets (LIS) connected through routers. A LIS has the following properties [RFC 1577] :

- all members have the same IP network/subnet number and address mask (independently of the physical location)
- all members within a LIS are directly connected to the same ATM network it communicates with each other through end-to-end ATM connections, either PVC or SVC
- all members outside of the LIS are accessed via a router

For transport of any network layer protocols over an ATM network (overlay model) it is needed to define :

- address resolution
- data encapsulation

1.1 Address Resolution

For operation of IP over ATM, a mechanism must be used to resolve IP addresses to the corresponding ATM addresses within an ATM logical IP subnet. CLIP uses ATMARP and InATMARP protocols for this purpose, that are based on the ARP and the InARP protocols (that are used in the standard IP networks). For this purpose an ATMARP server is defined within each LIS to resolve the requests of the LIS clients to map the IP address to the ATM address. All nodes within the LIS are configured with the unique ATM address of the ATMARP server. When a node comes up within the LIS, it establishes a connection to the ATMARP server using of the configured address. When the ATMARP sevrver accepts ATM calls/connections from a new LIS client, it transmits an Inverse ATMARP request to the new client and requests the node's IP and ATM addresses. It will add them into its ATMARP table. When LIS client (node A) wishes to send data to the node B in the same LIS, it sends an

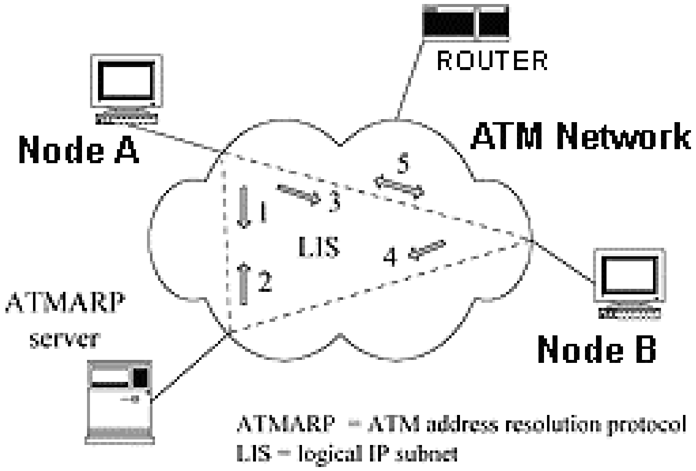

Figure 3.1.1 ATMARP protocol

ATMARP request to the server (1). Then it will generate the corresponding ATMARP response reply (if it has an entry in its ATMARP table (2)). The node A will use this address to establish the SVC connection to the node B (3). The node B (upon receiving the first IP packet from node A) the ATM address of the node A requests from ATMARP server(4). Then the node B resolves that it has connection already established with this node and it is not necessary to initiate the new one. The communication can begin (5).

The ATMARP server ages out its address table for robustness unless clients periodically refresh their entry with response to the server's inverse ATMARP queries. Server's ATMARP table entries are valid for a minimum time of 20 minutes.

The RFC 1577 defines using of UNI 3.0/3.1 signalization to establish SVC connection between nodes of a LIS.

1.2 Data Encapsulation

The RFC 1483 specification was defined by IETF in July 1993. It describes two encapsulation methods for carrying of network interconnect traffic over the ATM AAL5 [RFC 1483]. These two methods allows to transport network and link layer packets across an ATM connection (it can require QoS guarantees) and also multiplexing multiple packet types on the same connection (only UBR or ABR). It allows to conserve connection resource space (all data transfer between two nodes across the same connection) and to save in the connection latency after the first connection set-up [Minoli].

- *LLC/SNAP Encapsulation :* this encapsulation method is used for carrying multiple protocols over a single ATM VC connection. The packet type (IP, IPX, AppleTalk) is identified by a LLC header placed in front of the carried Protocol Data Unit (PDU).
- *VC Multiplexing :* in this method, each protocol is carried over a separate ATM VC, with the type of protocol identified at a connection set up. This approach is preferred in these environments, where dynamic creation of large ATM VCs is fast and economical.

The default packet size for IP members is 9180 octets plus 8 octets LLS/SNAP header. But the negotiation between members of a LIS is allowed and after the size can to be maximum of AAL 5 of 65K.

The CLIP has one great advantage - it is its simplicity. Any changes are required in conventional networks based on the standard routing architecture (end-to-end). The CLIP operation is straightforward, hence it has a number of limitation. One of these limitations is that communications between two nodes in two different LISs in the same ATM networks must go hop-by-hop through ATM routers (router with ATM interfaces) in the path from the source to the destination. These nodes can not established the direct connection across the ATM (routers can become bottlenecks). The second limitation the CLIP does not support broadcast. The IETF has been working on protocols that overcome these limitations and defined new specifications.

2. Next-Hop Resolution Protocol

With the CLIP, (as mentioned before), inter-LIS communications have to go through routers. This is not an optimal solution when both nodes communicate with each other are attached to the same ATM network. A mechanism is needed for an end system to resolve the IP address of another end system in a foreign LIS into the corresponding ATM address. The NBMA (Non Broadcast, Multi Access) NHRP protocol, developed by the IETF, overcomes this limitation and provides this mechanism. The term NBMA means an network technology such as ATM, Frame Relay or the X.25, which does not easily permit the use of broadcast mechanism and which allows the nodes to establish direct communication with each other.

A physical NBMA network may be partitioned into several logical NBMA subnetworks. There are one or more entities within the NBMA subnetwork, that implement the NHRP protocol : Next Hop Servers (NHSs) and NHRP Clients (NHCs). NHS is capable of answering to NHRP Resolution Request and to maintain «next hop resolution » cache which contains IP to ATM address mappings of all those nodes associated with particular NHS and table of IP address prefixes reachable through nodes served by the NHS. The cache is constructed from information obtained from registration packets (NHC are configured with the IP and ATM addresses its NHS(s)) and learning from NHRP resolution request/reply packets. The NHCs also maintains a cache of IP-to-ATM address mappings.

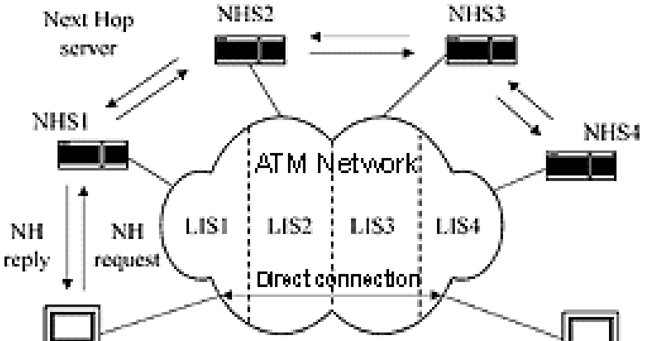

Figure 3.1.1 NHRP operation

The protocol works as follows (Figure 3.2.1). When a node has a packet to transmit across the NBMA network, and hence it needs to resolve a particular ATM address of destination, it sends NHRP request packet to its NHS. If requested destination is served by this NHS, it returns the address in a NHRP reply message to the requester. So far, the NHS behaves as an ATMARP server, and in LISs (where NHC and ATMARP clients coexist) NHSs are coupled with the function of ATMARP server. Limitation of an ATMARP server is that it can resolve an IP address that belongs only to the LIS, not to another LIS and NHS can.

If the NHS does not know an answer, it looks at its routing table (NHS works as a router, too) to determine the NHS next on the path to the destination address and forwards a request. At this next NHS, the same algorithm is used until a NHS is reached, that serves the destination and it will reply with corresponding ATM address to the requester. The reply travels back through the NHSs and an intermediate NHS may cache the IP to ATM mapping. Once the sender know the ATM address of the receiver, the direct connection can be established.

When resolution process is triggered, the source may (while awaiting reply), choose from : drop the packet, retain the packet until reply (NHRP resolution Reply) arrives and a more optimal data path is available, or it can forward the packet along the default path toward the receiver [NHRP].

NHRP also allows some optional features as route recording, detection of loops within NBMA network, address aggregation (with the address a network mask is provided that is associated with this address cached and used for protection for nonauthoritative request).

3. IP Multicast

One of the limitations of CLIP lies on its only support of only IP unicast over ATM. Current work on supporting of IP multicast over ATM is based on the Multicast Address Resolution Servers (MARS). With MARS, end systems directly connected to an ATM network are partitioned into "clusters". All end systems within the cluster are configured with the ATM address of the MARS. Distribution of multicast traffic within a cluster is done either Multicast Servers, or via a mesh of point-to-multipoint VCs. A MARS answers the queries for multicast addresses from end systems in the same way as ATMARP server for unicast addresses. The list of ATM addresses corresponding to the members for each multicast group is created by a host registration process. An end system join or leaves from a particular multicast group are made by sending Internet Group Multicast Protocol (IGMP) packets (includes MARS_JOIN or MARS_LEAVE message) to the MARS. When end system needs to transmit to a particular multicast group, it opens the connection to the MARS and issues a message for particular group. If the MARS has registered one or more other nodes for that multicast address, the two case can occur. The requested multicast address is configured to be served by a multicast server or by multicast mesh.

In the Multicast Server (MCS) case, the MARS returns message a that contains a server map of the one or more multicast servers serving the group. The MCS(s) will establish a point-to-multipoint connection or multiple point-to-point connection to the group members for forwarding the packets received from the end system to all members of the group. The requesting node then sets up a connection to the set of multicast servers and transmits its multicast packets.

In a case of multicast mesh, the MARS returns a message that contains the addresses of other nodes already registered as members of that group. After receiving the response, the requesting node will establish a point to multipoint connection to that set of nodes and begins to transmit packets in that connection.