RZ 3172 (#93218) 10/11/99
Engineering & Technology 25 pages

Research Report

IBM Mobile ATM Networking: IP Considerations

L. Fréléchoux, O. Stokes, S. Buchegger and R. Haas

IBM Research
Zurich Research Laboratory
8803 Rüschlikon
Switzerland

IBM Research
Almaden • Austin • Beijing • Haifa • T.J. Watson• Tokyo • Zurich

IBM Mobile ATM Networking:
IP Considerations

L. Fréléchoux, O. Stokes, S. Buchegger and R. Haas

IBM Research, Zurich Research Laboratory, 8803 Rüschlikon, Switzerland

Abstract

This document describes the recommended IP configuration to be used with IBM's standards-based implementation of support for mobile ATM networks. IBM's Mobile ATM Networking technology brings the following benefits in addition to the normal advantages of ATM:

- Seamless integration of an organization's mobile and fixed networking infrastructures.
- Optimal utilization of scarce resources such as the bandwidth of low-capacity satellite links.
- A robust connectivity service that eliminates the impact of mobility on new and existing connections as long as connectivity exists.

IBM's support for mobile ATM networks is described in detail in a companion White Paper, "IBM Mobile ATM Networking Technology Overview" [1]. The reader should be familiar with this document. In addition, IBM's suggested ATM architecture and addressing scheme are described in the White Paper entitled "IBM Mobile ATM Networking Network Design and Addressing Considerations" [2].

After providing a brief overview of the Mobile ATM network components and topology, this paper introduces IBM's solution for efficiently using IP over mobile ATM and discusses alternatives for providing quality of services at the IP layer in a mobile environment.

IBM's Mobile IP/ATM Networking solution is applicable in both commercial and military environments, and therefore the examples presented in this paper are intentionally generic.

1 Introduction

This document describes the recommended IP configuration to be used with IBM's standards-based implementation of support for mobile ATM networks. IBM's Mobile ATM Networking technology brings the following benefits in addition to the normal advantages of ATM:

- Seamless integration of an organization's mobile and fixed networking infrastructures.
- Optimal utilization of scarce resources such as the bandwidth of low-capacity satellite links.
- A robust connectivity service that eliminates the impact of mobility on new and existing connections as long as connectivity exists.

IBM's support for mobile ATM networks is described in detail in a companion White Paper, "IBM Mobile ATM Networking Technology Overview" [1]. The reader should be familiar with this document. In addition, IBM's suggested ATM architecture and addressing scheme are described in the White Paper entitled "IBM Mobile ATM Networking Network Design and Addressing Considerations" [2].

After providing a brief overview of the Mobile ATM network components and topology, this paper introduces IBM's solution for efficiently using IP over mobile ATM and discusses alternatives for providing quality of services at the IP layer in a mobile environment.

IBM's Mobile IP/ATM Networking solution is applicable in both commercial and military environments, and therefore the examples presented in this paper are intentionally generic.

2 ATM Network Topology

The recommended network ATM topology has three distinct components, as illustrated in Figure 1.

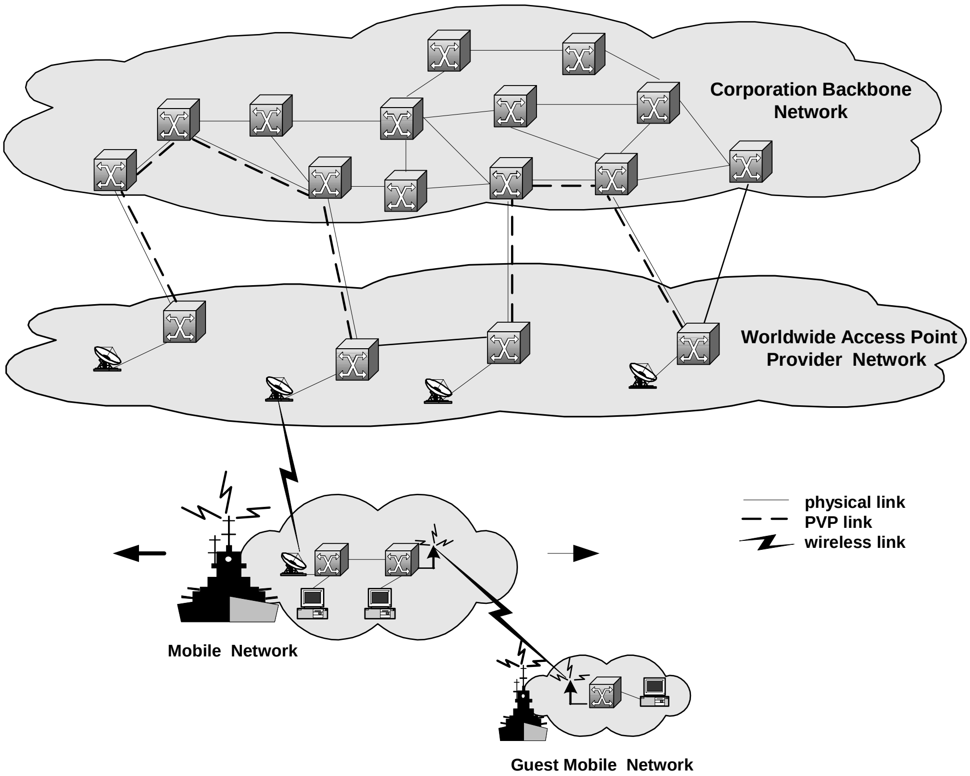

Figure 1: Mobile ATM Network Topology

- The **backbone network** represents the land-based networking infrastructure of an organization. This may include many campuses distributed throughout the world, interconnected by third parties, or an organization's MANs and WAN's, leased lines, etc. The backbone network is connected either ATM wise or through IP routers to the Access Point Provider Network.

  The backbone network is not aware of the location of individual mobile networks. This is essential because otherwise the network architecture does not scale. If a call is received that is addressed to a mobile network, that call is routed to the Access Point Provider Network.

- The **Access Point Provider Network** is a group of ATM switches located throughout the world that provide fixed-network access points to mobile networks. The Access Point Provider Network (APPN) must be able to route to the individual mobile networks and is therefore aware of the current location of each mobile network. Internally the switches of the APPN must be connected so that it is possible to reach all mobile networks from any point on the APPN. Therefore the switches in the APPN are connected either through physical links or through VP links that traverse the backbone network. The switches on the ground side of the wireless links are called the **access points**.

- Each large mobile platform, such as a ship or an aircraft, has its own onboard networking infrastructure, referred to as **mobile network**. When wireless links (e.g., satellite or line-of-sight links) are available, the mobile network has connectivity to the entire worldwide infrastructure via the Access Point Provider Network.

  Some mobile networks may not have direct contact to the APPN, but may have a wireless link to another mobile network. Such a network is illustrated as a "guest" mobile network in Figure 1. Support for guest nodes is an initial support to ***ad-hoc* networking**.

3 Overview of IP Overlay of Mobile ATM

As shown in Figure 2, the reference network architecture of a mobile environment consists of an ATM infrastructure, as described in section 2, populated with IP routers connected to the ATM switches. Mobile platforms contain one or several **mobile edge routers** located on board. **APPN edge routers** are located in the Access Point Provider Network and are distributed geographically. Each APPN edge router is in the proximity of one or a group of access points. The corporation backbone network is structured as usual, with its set of internal routers and with **gateway routers** (firewall) which connect the corporation backbone to the Internet. Some routers of the backbone are labeled as **home routers** of mobile edge routers. Ground servers providing services to a mobile platform, or to a group of mobile platforms, are connected to the home router(s) associated with these mobile platforms. A home router exchanges routing information with mobile edge routers, as if the latter were fixed.

IBM's approach to implementing mobile networking uses ATM as the transport layer for communication between mobile platforms through sites in the worldwide APPN to the corporation backbone. The APPN and the mobile networks are aware of and capable of adapting to mobility at the ATM layer. The IP layer takes full advantage of the PNNI mobility support provided by the ATM layer, which automatically reroutes point-to-point switched virtual connections (SVCs), including the router to router SVCs used to transport IP packets between two routers. The data transfer interruption caused by the rerouting of SVCs is minimal and only results in the loss of a few IP packets, which can be retransmitted by higher layer protocols, such as TCP. There is no indication of disruption to the routers, and the router to router sessions are maintained.

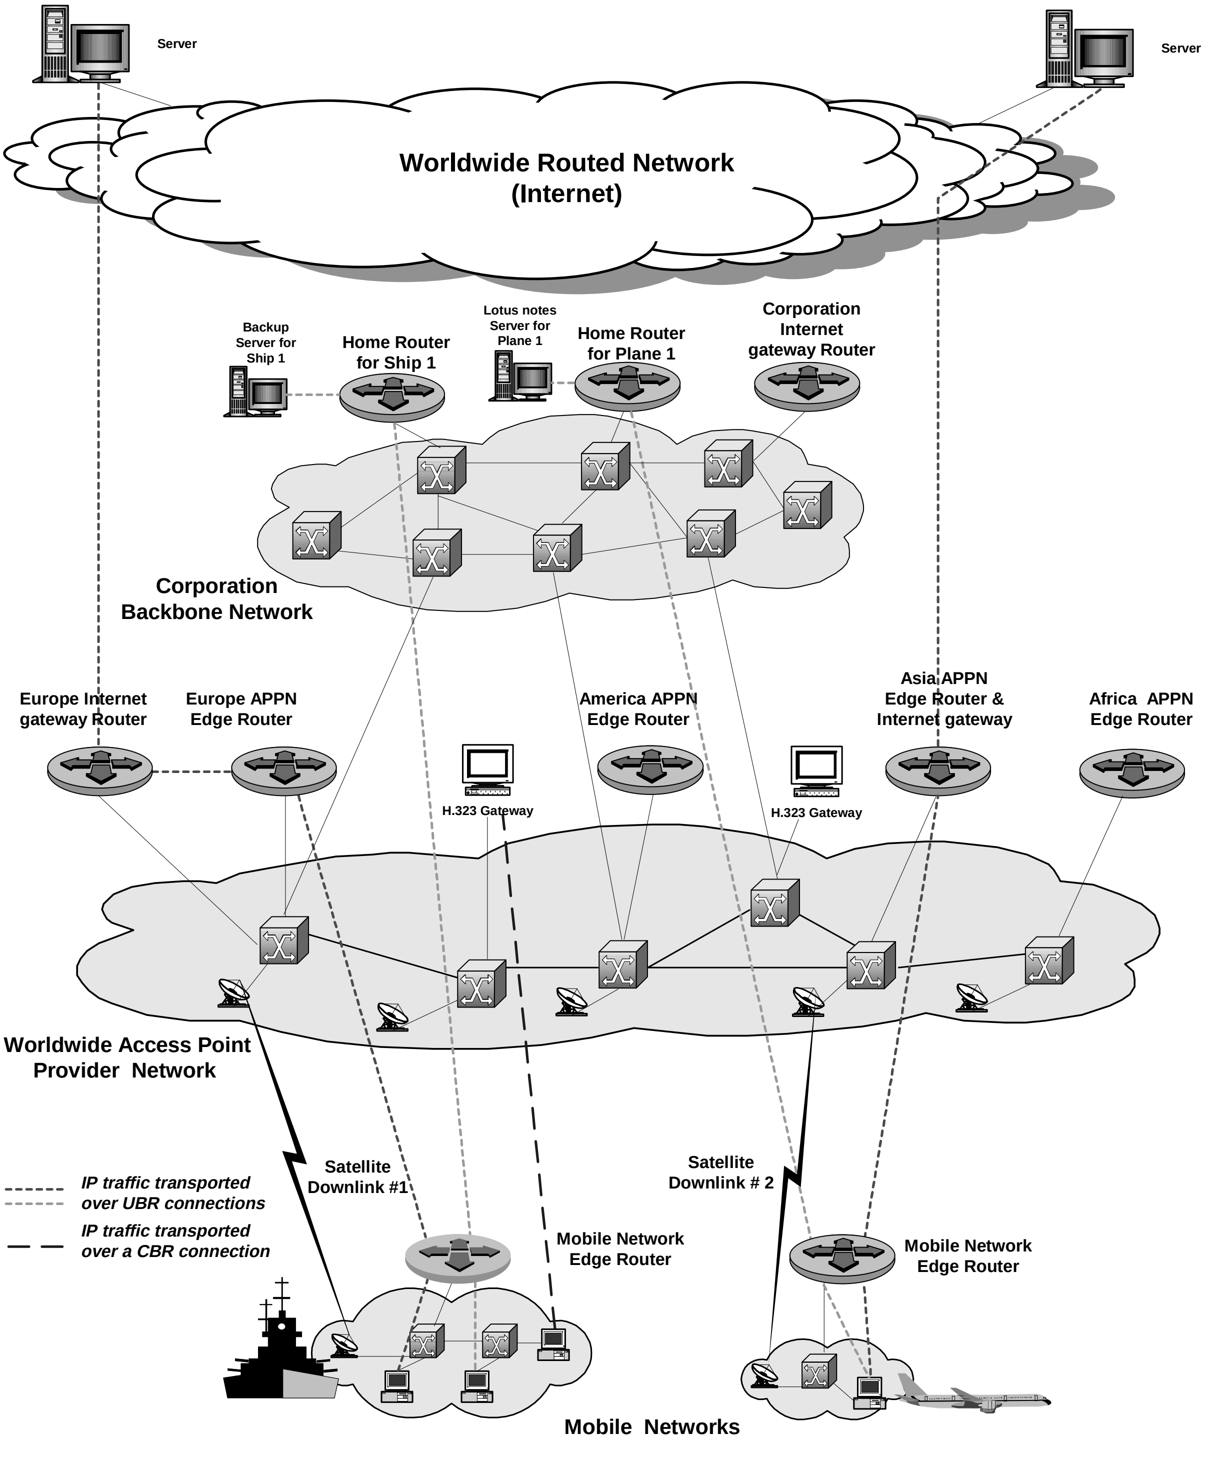

Figure 2: IP Overlay on the ATM APPN

Handover of connections is triggered automatically when a mobile platform is about to lose a satellite or a line-of-sight link. The connections are rerouted to a secondary link if one is available. If a mobile platform loses a direct link to the APPN, but is still connected to another mobile platform which has connectivity to the APPN, the IP traffic is redirected to use one of the remaining WAN links.

For ground services provided to a mobile platform or a group of mobile platforms, the mobile edge router can directly communicate with one of its home routers located in the backbone network. The home router then delivers the IP data directly to the servers. As the mobile platform moves and changes access points, the SVC between the mobile edge router and the home router is rerouted by the ATM layer, hiding the mobility from the IP layer. This configuration is referred to as the static IP configuration.

In the dynamic IP configuration, the IP traffic generated by a host in a mobile platform is sent to the mobile edge router which then forwards it to the geographically closest APPN edge router. As the mobile platform moves and roams between access points, the mobile edge router adapts its IP routing information to send the IP traffic to the geographically closest APPN edge router.

The APPN edge routers in turn forward the IP traffic to the backbone routers which deliver the packets to the servers. The IP traffic addressed to the Internet is directly forwarded by APPN edge router to the closest APPN Internet gateway, without going through the corporation network. IP traffic from the fixed network to the mobile networks takes the exact reverse path.

Implementation of Multiple Protocol Over ATM (MPOA) [3] or Next Hop Routing Protocol (NHRP) [4] can also be used to further eliminate interim routers by setting up dedicated ATM SVCs for IP flows corresponding to certain criteria.

Servers and end-devices requiring Quality of Services (QoS) can communicate either through native ATM connections, for example by establishing a direct Constant Bit Rate (CBR) SVC between a PC and an H323 gateway, or by using a form of IP QoS, such as Diffserv or RSVP. IBM believes that Diffserv will be an enabler for providing IP QoS in a mobile environment.

4 Logical configuration of IP over mobile ATM

4.1 IP network topology

Two logical IP network topologies are considered for the deployment of IP overlaying mobile ATM.

In the first model, the corporation network and the APPN are two distinct networks, managed as two autonomous systems. This model is the same as the ISP - private network model commonly used in the Internet. In this commercial model, a company provides the services of the APPN to a number of private customers, each with their own private network.

As shown in Figure 3, both the corporation backbone and the Access Point Provider Network run OSPF [5] to exchange IP routing information within their networking infrastructure. Each runs OSPF as its own autonomous system (AS). The corporation backbone and the APPN exchange IP routing information with an inter-AS routing protocol, for example BGP [6]. Both networks also use BGP sessions to exchange IP routing information with other autonomous systems of the Internet.

With this first model, the mobile networks, as described further on, can exchange IP routing information with both their corporation network and the APPN. A mobile edge router can have an interface on an OSPF area of the corporate network and an interface on an OSPF area of the APPN network. An internal BGP session bridges between the two autonomous systems.

In the second model, the corporate network and the APPN are integrated into the same network, and are under a common administration. This is reflected in the IP topology, by having both the corporate network and the APPN managed as one single autonomous system (see Figure 4). The integrated IP network runs OSPF as an intra-AS routing protocol, and exchanges, through BGP sessions, routing information with the other autonomous systems of the Internet.

With this second model, the mobile networks belong to the AS of the integrated corporate network - APPN and exchange routing information through OSPF with the routers on the ground infrastructure.

This white paper addresses both of these reference IP topologies. Although the design of an IP overlay for mobile ATM, as presented in the next sections, is very similar for integrated and separated topologies, the requirements for the mobile edge routers are different.

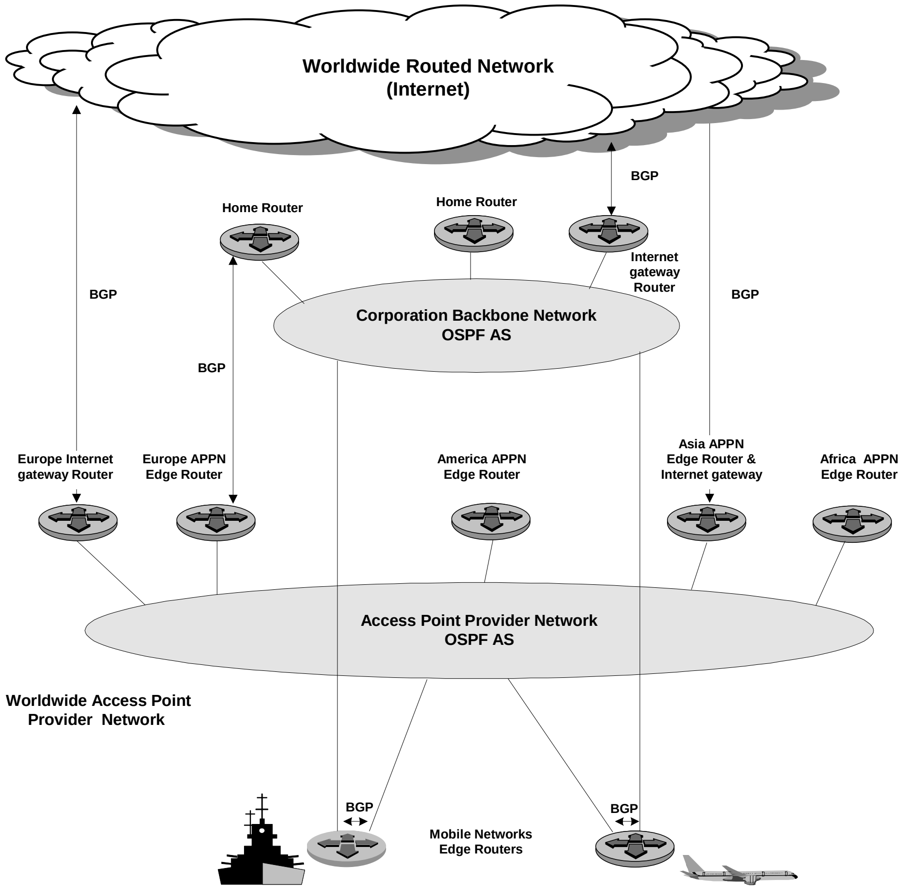

Figure 3: Network model with a separated corporation backbone and APPN

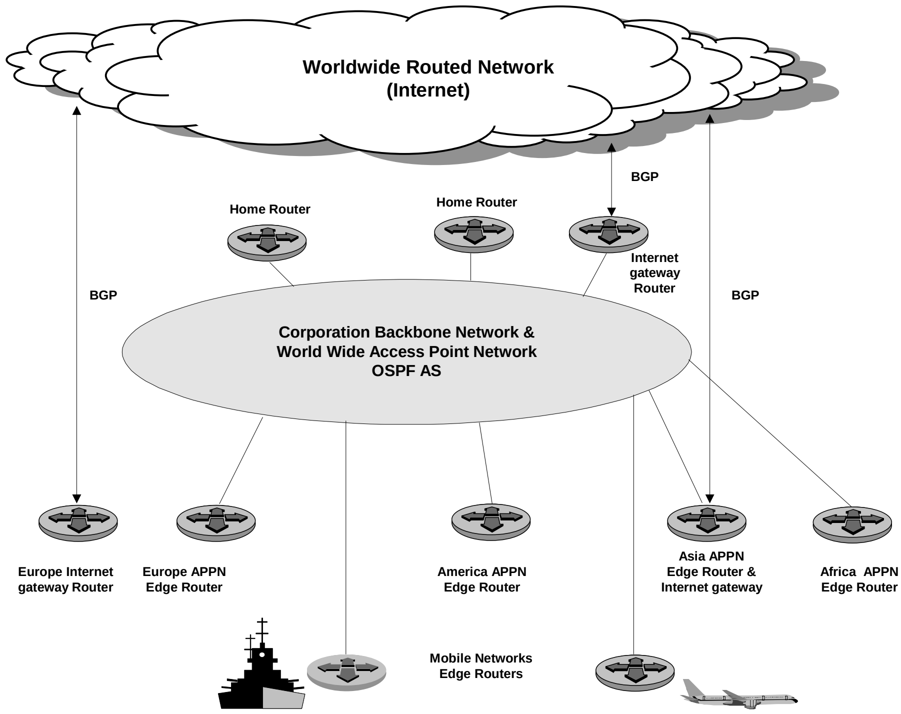

Figure 4: Network model with an integrated corporation backbone and APPN

4.2 Static IP configuration over ATM

The mobile ATM technology not only provides automatic rerouting of SVCs when a WAN link is lost, but the PNNI mobility extensions [7] also provide the same routing support to the IP layer as for static ATM networks. ATM calls are routed regardless of whether the destination address belongs to a mobile platform or is located in the backbone network.

This full support for mobility at the ATM layer can be used to completely hide the mobility to the IP layer. The logical IP network can therefore be configured as any static IP network. As shown in Figure 5, the static IP configuration uses the OSPF protocol to exchange routing information between the routers in the corporation backbone and the mobile edge routers. The mobile edge routers of one or several mobile platforms are grouped into the same OSPF area. In our example, the mobile edge routers in the mobile platforms L1 and S2 belong to the same OSPF area H.H.H.H. In the backbone network, the home routers of the mobile platform L1 and S2 have also IP interfaces that run OSPF and belong to the area H.H.H.H. The home routers of area H.H.H.H exchange routing information with the mobile edge routers belonging to the same area, as if the latter were fixed. As the mobile platform moves, the SVCs between the mobile edge router and the home routers are rerouted, without any involvement of the IP layer.

Communications off the mobile platform to the fixed network are achieved via a Classical IP (CIP) [8] Logical IP Subnet (LIS) between the mobile edge router and router(s) in the mobile platform’s home base. The mobile edge router has a pre-configured adjacent neighbor which is an interface on the corporation backbone home router. The ATM ARP server of the LIS is located on the fixed network, in the subnet of the mobile platform’s home router(s). CIP is preferred to LANE [9] in the mobile environment because LANE requires point to multipoint connections to be set up between the LAN Emulation Client (LEC) and both the LAN Emulation Server (LES) and the Broadcast and Unknown Server (BUS) in support of broadcasts that occur in a normal LAN environment. These point-to-multipoint connections cannot currently be rerouted with the mobile ATM technology when the mobile platform changes access points, requiring therefore the establishment of new router-to-router IP sessions when the mobile platform changes location.

Router adjacencies between the mobile edge router and its home router(s) can also be “hard coded” with both the IP and the ATM addresses to eliminate the need for an ATM ARP server.

If the static IP configuration is not combined with dynamic IP configuration (see below), the home routers based OSPF areas can be configured as stub areas, so that all of the routes known to the APPN edge routers are not advertised over the wireless link. The mobile edge router receives only default route information which indicates that all traffic leaving the mobile platform must be forwarded through its neighbor home router. This reduces the exchange of IP routing information, improves scalability and stability in the network, and saves bandwidth on the wireless link.

Although the static IP configuration is very straightforward to implement, this configuration can lead to suboptimal forwarding of IP packets, because the IP layer is completely unaware of the real location of a mobile network. It could very well be that a ship, currently located in Asia has its home router located in North America, requiring physically every IP packet sent from the mobile platform to travel at least half way around the world to reach its destination. From an IP perspective the mobile edge router is connected in one hop to its home router, but the ATM hides the physical length of this hop.

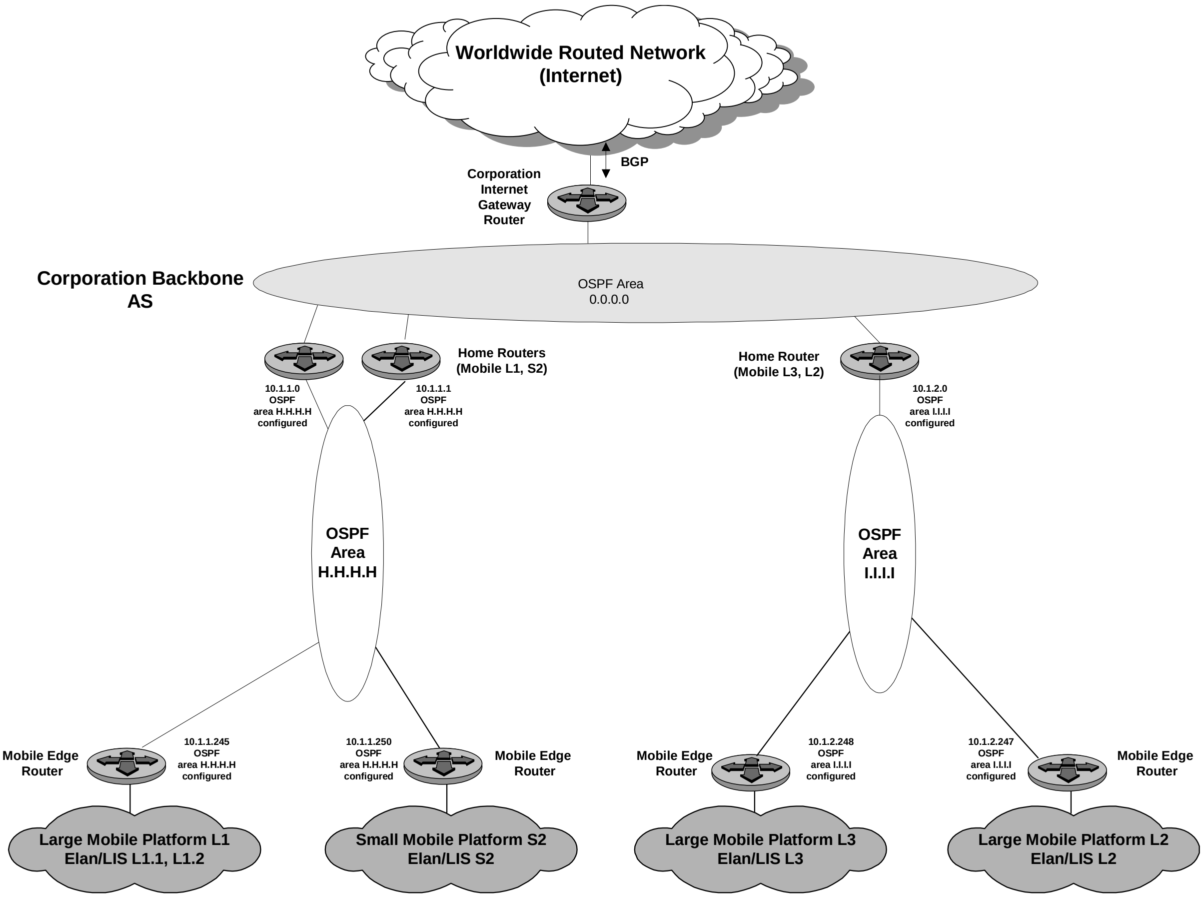

Figure 5: Static IP configuration

The static IP configuration is well suited when the mobile platforms have dedicated servers in the corporation network which provide centralized services that are not geographically distributed. These servers can be connected directly to the home routers of the mobile platforms.

4.3 Dynamic IP configuration over ATM

To alleviate the suboptimal routing problem described in the previous section, the IP layer has to be partially aware of mobility. The integration of mobile ATM technology with ATM Proxy PAR technology [10][11] enables the optimal routing of IP packets between the mobile platforms and the corporation backbone or the Internet.

The Proxy PAR mechanism allows routers to register upper layer protocol information (e.g. OSPF, BGP interfaces) and to get this information flooded throughout the ATM network by the PNNI protocol. A router can for example register an OSPF interface (IP address, ATM address, OSPF area) to its serving ATM switch, which then floods the information throughout the ATM network. Other routers attached to the same ATM network can retrieve this information by querying their serving ATM switches and initiate for example an OSPF neighbor relationship with the router that registered this interface. When registering information, Proxy PAR allows the specification of a scope which confines the flooding of the information to the portion of the ATM network within that scope. Proxy PAR can also be used to allow end systems to find ARP, DHCP, and DNS servers and any other service automatically.

In the mobile environment, each APPN edge router registers its IP routing interfaces (e.g. OSPF) with a scope that limits the flooding of the information to the mobile networks that are currently connected to the access point(s) collocated with this router. A mobile platform with connectivity to the APPN receives via Proxy PAR the information about the routers which are geographically close to its current access point. The mobile edge router(s) on board this platform can then initiate OSPF neighbor relationships with the routers that are local. Annex A gives further details about the Proxy PAR configuration to be implemented in a mobile environment.

The dynamic IP configuration, as shown in Figure 6, uses the OSPF protocol to exchange routing information between the APPN edge routers and the mobile edge routers. All the APPN edge routers (or group of APPN edge routers) have interfaces into the backbone OSPF area (0.0.0.0). In addition, each APPN edge router is configured with its own OSPF area and registers its OSPF interfaces via Proxy PAR. A mobile edge router, by querying its serving ATM switch, dynamically learns about the OSPF interfaces of APPN edge router(s) that are local to the current access point. The mobile edge router can then create an IP interface with a temporary IP address that belongs to the same subnet as the IP address of the interface registered by the APPN edge router. The IP interface is configured to be part of local OSPF area and the mobile edge router can build a router adjacency with the local APPN edge router.

The assignment of a temporary IP address is either implemented with a mechanism such as DHCP [12], or based on the configuration of a unique IP postfix in every mobile edge router. The IP postfix is unique within the APPN and the mobile edge router can build a unique temporary IP address by adding its postfix to the IP subnet prefix registered by the APPN edge router for the OSPF interface. IP v6 addresses [13] are structured according to this scheme, with the unique postfix contained in the Interface ID field.

Figure 6 shows that the mobile edge router of the large mobile platform L1 learns via Proxy PAR that the European APPN edge router has the IP interface 9.5.1.0 configured to run OSPF in the area A.A.A.A. The mobile edge router creates the IP interface 9.5.1.245 to run OSPF in the area A.A.A.A and builds a neighbor relationship with the European APPN edge router. The mobile edge router receives the IP topology of the European area, meaning of the mobile platforms currently connected

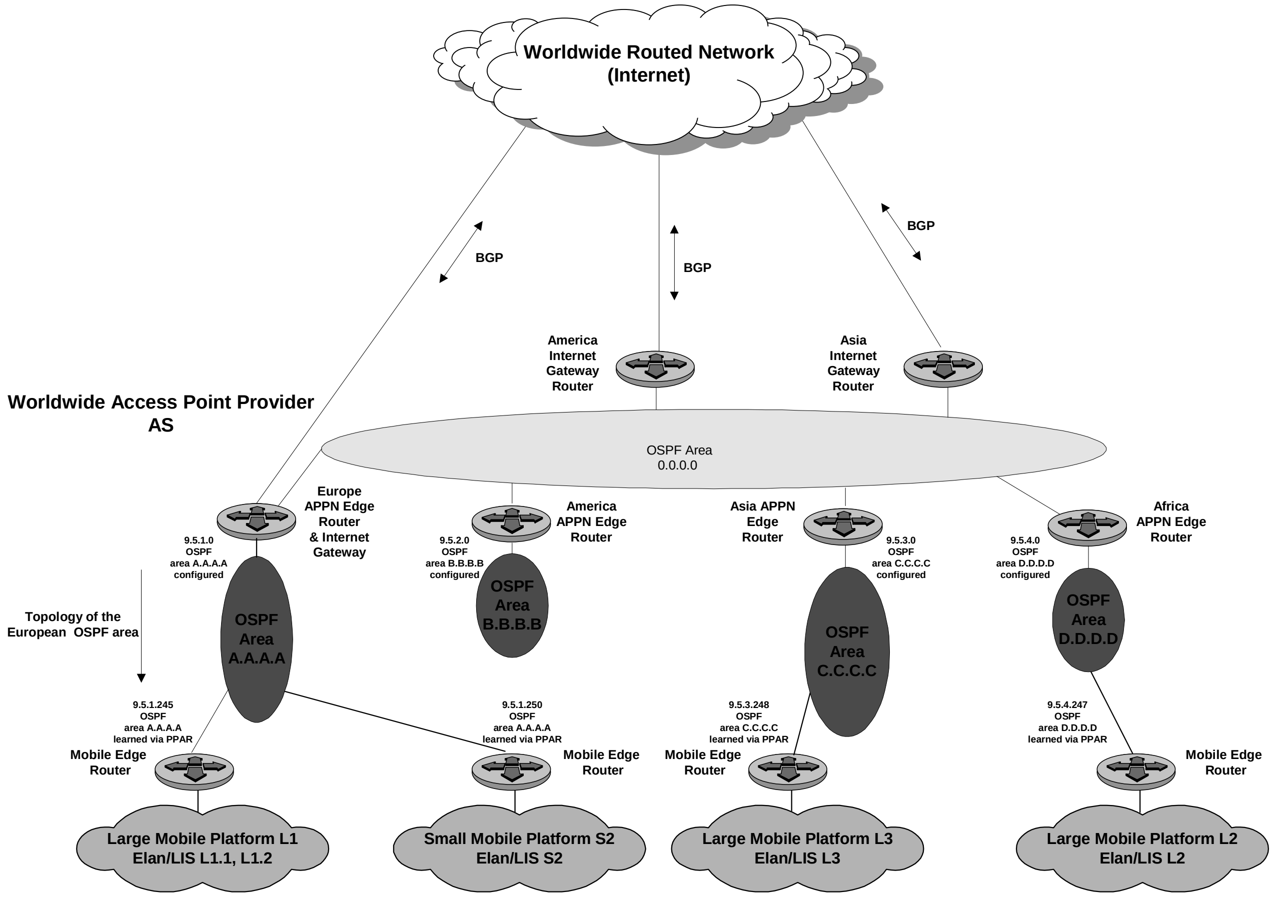

Figure 6: Dynamic IP configuration

to the European access point, and a summary of the IP reachabilities in the backbone area and in the Internet.

When the large mobile platform moves and connects to the American access point, the mobile edge router removes the IP interface 9.5.1.245 and creates the IP interface 9.5.2.245 to run OSPF in the area B.B.B.B.

4.4 Guest node connectivity

The Mobile ATM Technology supports line-of-sight communications between certain mobile platforms whether or not any mobile platform has connectivity back to the APPN. To do this, a hierarchical routing structure is implemented. Our structure conforms to the concept of a large mobile platform being higher in the structure than a small mobile platform. In this case, the large platform has multiple layers of PNNI hierarchy configured and the small platform normally has less levels. As two mobile platforms connect with a line-of-sight link, the small mobile platform integrates into the PNNI hierarchy of the large mobile platform.

If both a large and a small platform have a mobile edge router on board, the same mechanism as described in section 4.3 can be used to dynamically configure the neighbor relationship between these two routers. The mobile edge router of the large platform acts as the APPN edge router and registers via Proxy PAR its IP interface with a subnet and an OSPF area unique to that large platform. This OSPF area is configured as a stub area so that the mobile edge router of the large platform advertises a default route to the mobile edge routers of the small platforms. When a small platform establishes a line-of-sight link with the large platform, the mobile edge router of the small platform retrieves via Proxy PAR the IP interface registered by the mobile edge router of the large platform and dynamically creates an IP interface for OSPF within the advertised area. A BGP session, internal to the mobile edge router, ensures the exchange of IP reachabilities between the "guest node" OSPF area and the OSPF area of the APPN edge router. The "guest node" OSPF area is an autonomous system in itself, and therefore the routing of IP traffic within the "guest node" OSPF area also works when none of the mobile platforms has connectivity to the APPN.

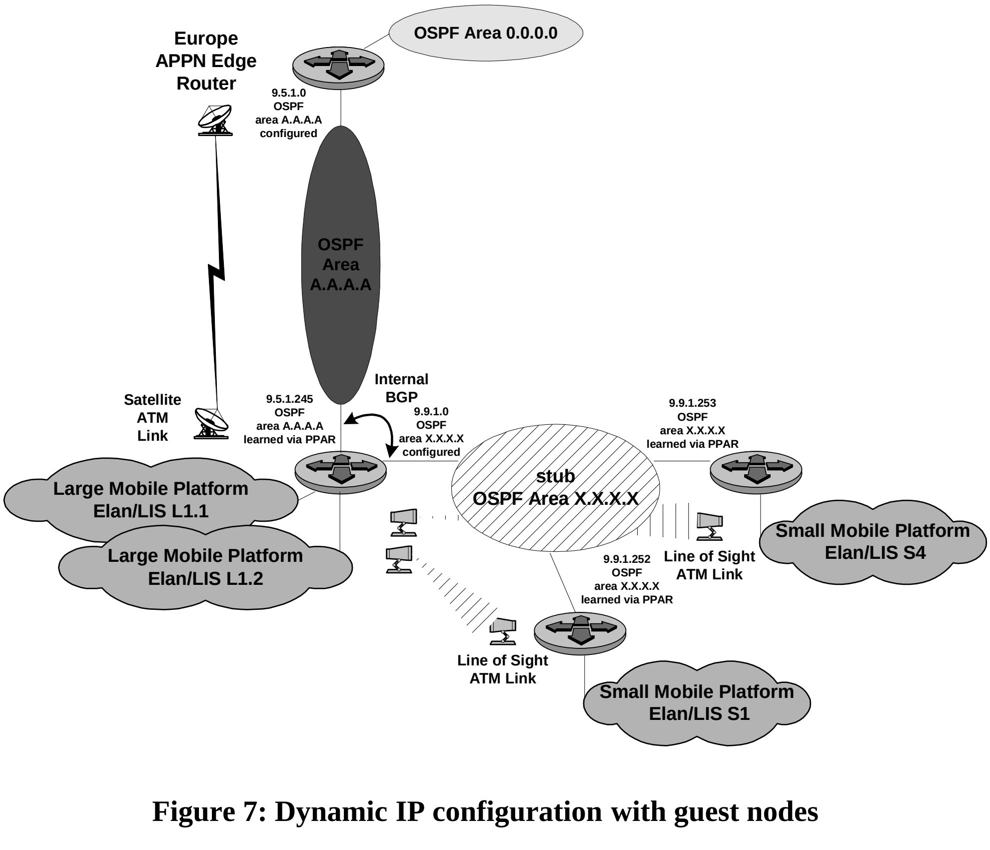

Figure 7: Dynamic IP configuration with guest nodes

As shown in Figure 7, the mobile edge router on board the large platform L1 has the IP interface 9.9.1.0 configured to run OSPF in the area of value X.X.X.X. The router registers this information with Proxy PAR. When the small platform S4 establishes a line-of-sight link to L1, the ATM switch on board the S4 receives the Proxy PAR information registered by the router of L1. The mobile edge router of S4, having retrieved this information, dynamically creates the IP interface 9.9.1.253 to run OSPF in the area X.X.X.X. When the second small mobile platform S1 establishes a line-of-sight link to L1, it creates the IP interface 9.9.1.254 and joins the OSPF area X.X.X.X through the same mechanism. Annex A describes in further detail the Proxy PAR configuration to support the logical IP configuration for guest nodes.

4.5 Static versus Dynamic IP configuration

The static and the dynamic logical IP configurations can be implemented together, to provide the best performance when forwarding IP traffic. In a mobile environment, the static IP configuration is very efficient when the hosts on board the mobile platform communicate with servers in the corporation backbone which provide centralized services that are not geographically distributed. The dynamic IP configuration, on the other hand, allows an optimal routing of the IP traffic addressed to servers in the Internet, provided that the Access Point Network Provider has several IP gateway routers located in the different regions of the earth. The communication between mobile platforms is also more efficient, preventing the IP traffic from being sent to the backbone network and forwarded back to the APPN.

The static configuration is thought to provide the default connectivity between the mobile platforms and the ground. With this configuration the mobile platform has access to all the ground services, even if no APPN edge router is located at the current access point or if the edge router is not available.

The examples given in section 4.2 and 4.3 are based on the IP network model where the corporation network and the APPN are two distinct autonomous systems. Figure 8 shows the IP configuration where the corporation network and the APPN are in the same autonomous system.

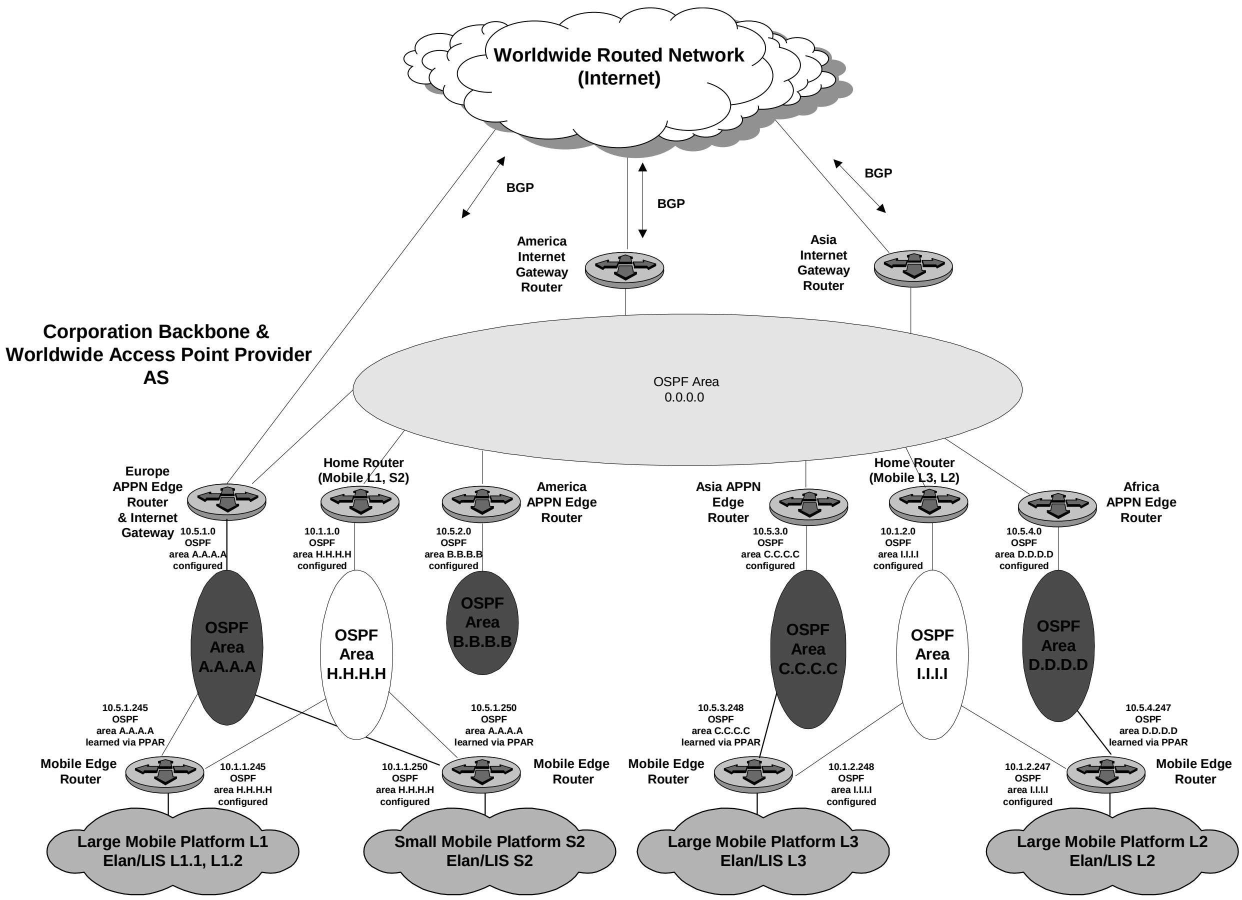

Figure 8: Static and Dynamic IP configuration in an integrated IP model

4.6 Inter-AS routing

To confine the impact of mobility and to build scaleable networks, it is important that the routing between the APPN and the other autonomous systems in the Internet is not affected by the motion of individual mobile networks. The address prefixes advertised by the APPN AS are therefore not a function of the location of the mobile networks, but rather indicate that a group of mobile hosts are reachable through the APPN AS, without specifying their current location.

As depicted in Figure 9, when considering the separated IP network model, the APPN AS advertises the network prefixes 9.5.*.* and 10.0.*.*. The prefix 9.5.*.* summarizes all the IP addresses reachable within the APPN while 10.0.*.* summarizes the IP addresses of all the mobile platforms of the corporation. The corporation backbone network AS advertises the address prefix 10.*.*.* which summarizes all the IP addresses of the hosts in the corporation, including the hosts on board the mobile platforms. An IP packet from the Internet and targeted to a mobile host is routed to the closest APPN gateway, which then uses the OSPF routing information to route it to the right APPN edge router and finally to the mobile edge router. The longest prefix match used in IP routing gives preference to the APPN gateways (advertising 10.0.*.*) over the corporation gateways (advertising 10.*.*.*) for IP traffic sent to the mobile networks, resulting in better paths. An IP packet generated by a server in the corporation backbone and targeted to a workstation on board a mobile platform L1, exits the corporation AS via the home router of L1.

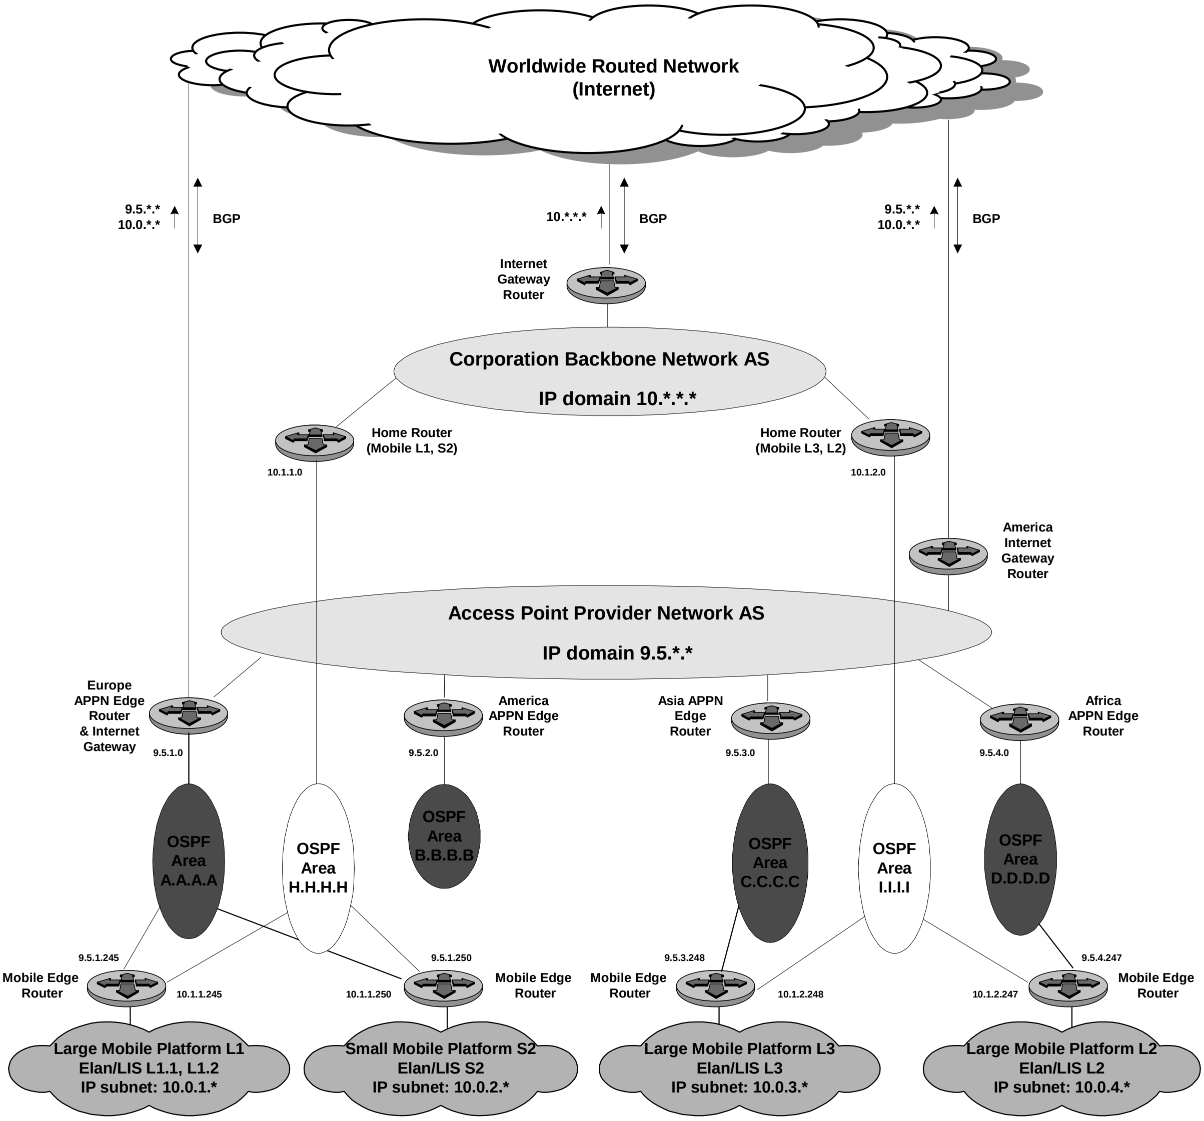

Figure 9: Inter-AS advertisement in the separated IP model

When considering an integrated IP model (see Figure 10), all the gateways advertise the same network prefix 10.*.*.* which summarizes the addresses of the fixed and mobile hosts in the corporation. An IP packet from the Internet, regardless of whether its destination is located on board a mobile platform or in the fixed network, is routed to the closest gateway. If the destination is a mobile host, this gateway then uses the OSPF routing information to route it to the right APPN edge router and finally to the mobile edge router.

In both the separated and the integrated IP models, the motion of a platform to a new access point does not impact the inter-AS routing. When a mobile platform moves, only the intra-AS routing information requires an update.

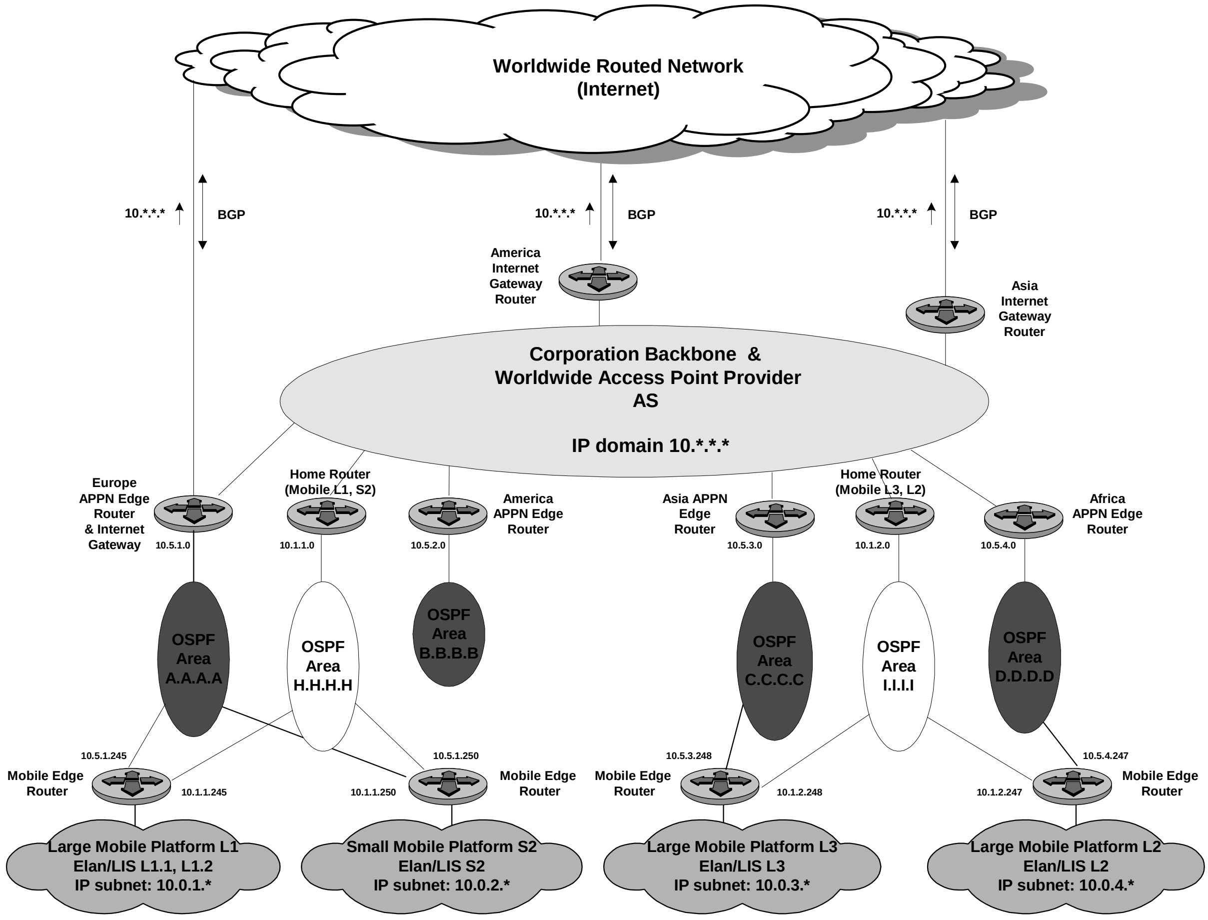

Figure 10: Inter-AS advertisement in the integrated IP model

4.7 Scenarios of mobility

The following examples show several mobility scenarios supported by IBM’s mobile networking solution. These scenarios are based on the integrated IP network model. The large ship (L1) with two sub networks on board is cruising on the Atlantic ocean (see Figure 11). The smaller ship (S4) is cruising in the same fleet and maintains a laser line-of-sight link with L1. The router on board L1 and the router on board S4 have instantiated IP interfaces and exchanged routing information using OSPF. IP traffic can be forwarded from hosts on board S4 to hosts on board L1 and vice versa. There is no connectivity to the ground networking infrastructure and the fleet of ships forms a stand-alone network.

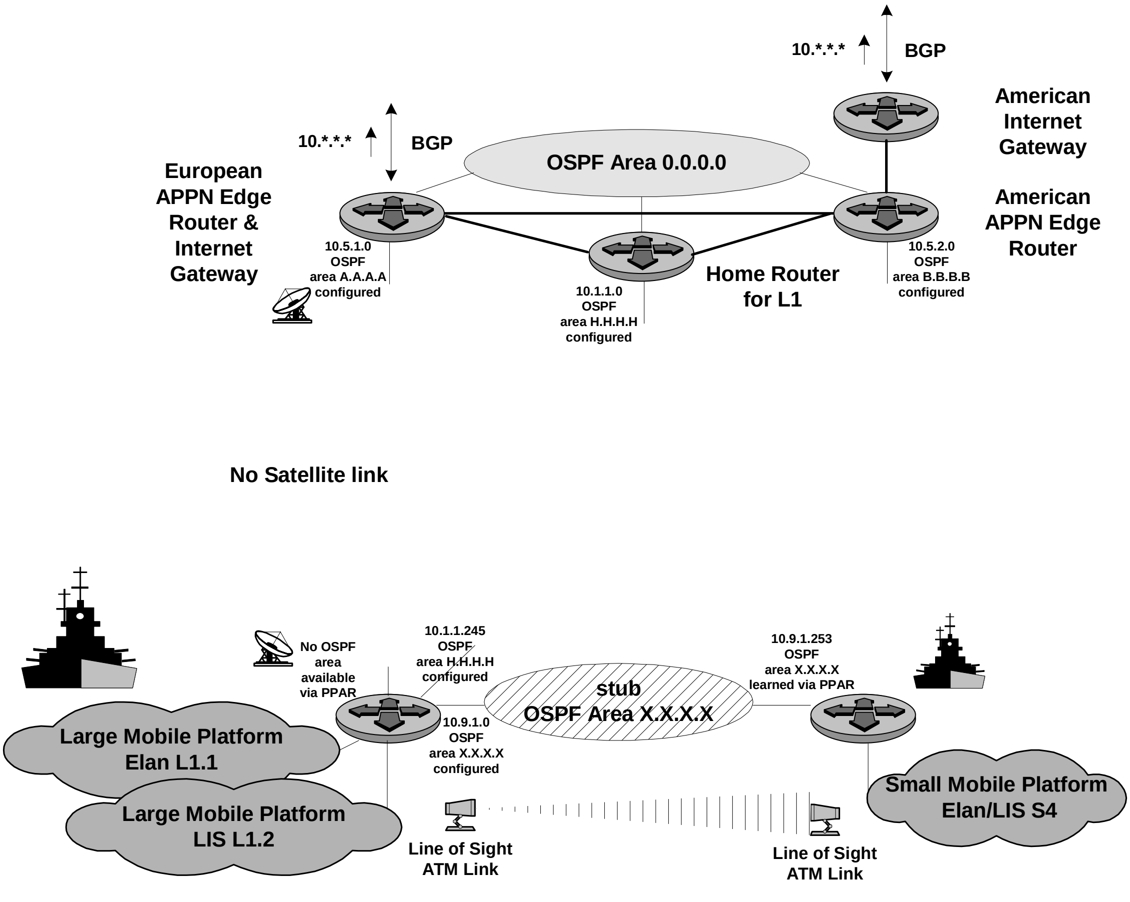

Figure 11: Mobile networks not connected to the backbone

As the fleet moves, the large ship enters the coverage area of a satellite (see Figure 12). Once the satellite link is up and the ATM connectivity is available, the mobile edge router of L1 activates the pre-configured IP interface running OSPF in the area H.H.H.H and establishes a neighbor relationship with its home router. At the same time, the mobile router learns via Proxy PAR of the existence on the European APPN edge router of an IP interface running OSPF in the area A.A.A.A. OSPF area X.X.X.X advertises the reachability to the subnet of S4 into the OSPF areas A.A.A.A and H.H.H.H through the internal BGP session of mobile router L1. Because area X.X.X.X is a stub area, S4 receives the default IP reachability configured at the mobile router L1. The IP packets can now be forwarded from hosts on S4 to the router on L1 which forwards them either directly to the home router of L1 or to the European APPN edge router. IP traffic targeted to servers in the home network is forwarded to the home router, while other IP packets are routed to the APPN edge router. IP traffic from the ground is forwarded either through the APPN edge router, or through the home router of L1 to the mobile platforms, depending on where it is generated. Traffic from S4 to L1 and vice versa is routed locally, without going over the satellite link.

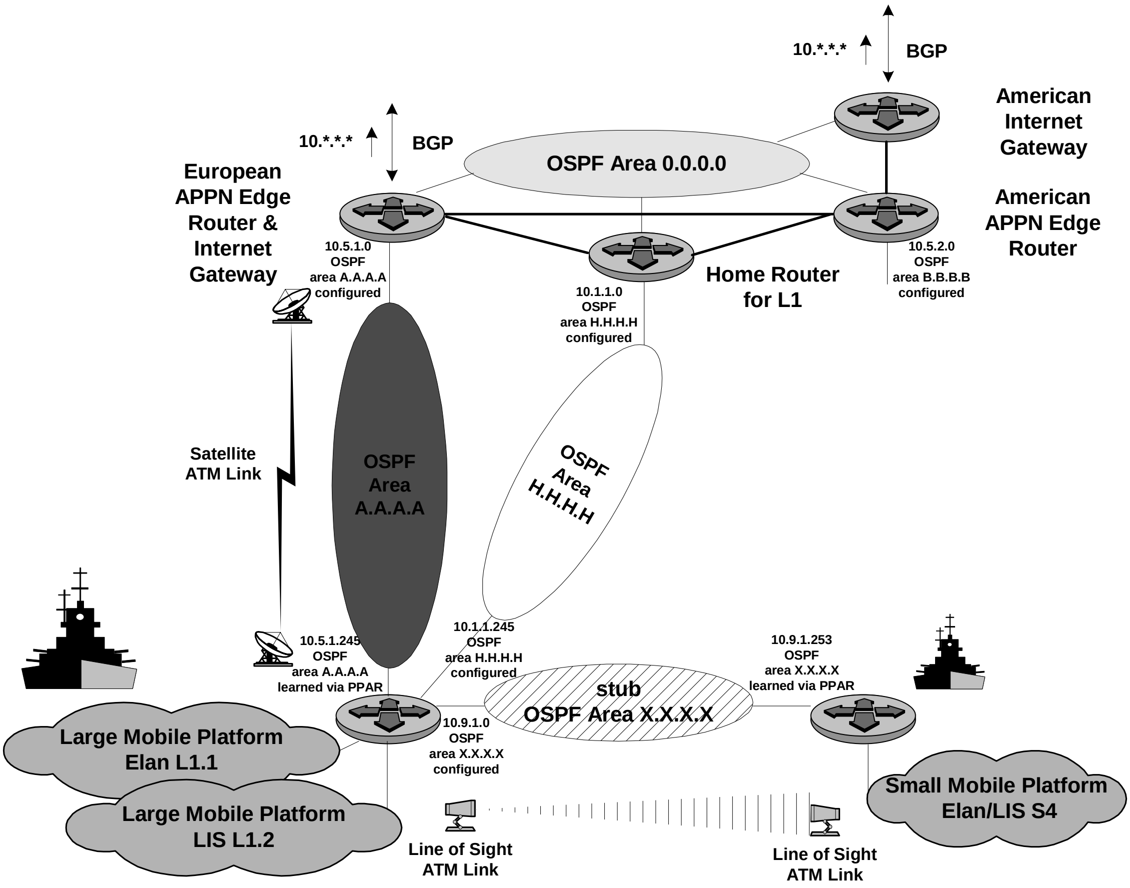

Figure 12: Large platform connects to the backbone

As the fleet continues its cruise, the satellite link fades and the handover to a new satellite becomes necessary (see Figure 13). Once the new satellite link is up and the ATM connectivity established, all the ATM connections on the fading satellite are automatically rerouted by the ATM layer. This includes the point-to-point SVCs between the mobile edge router of L1 and its home router and the SVCs between the mobile edge router of L1 and the European APPN edge router. When the mobile edge router on board L1 notices the change in the Proxy PAR information delivered by its serving ATM switch, the mobile router of L1 creates a new IP interface to run OSPF in the area B.B.B.B. The mobile edge router then removes the IP interface communicating with the European APPN edge router and from that point on the IP traffic targeted to the APPN is forwarded to the American APPN edge router.

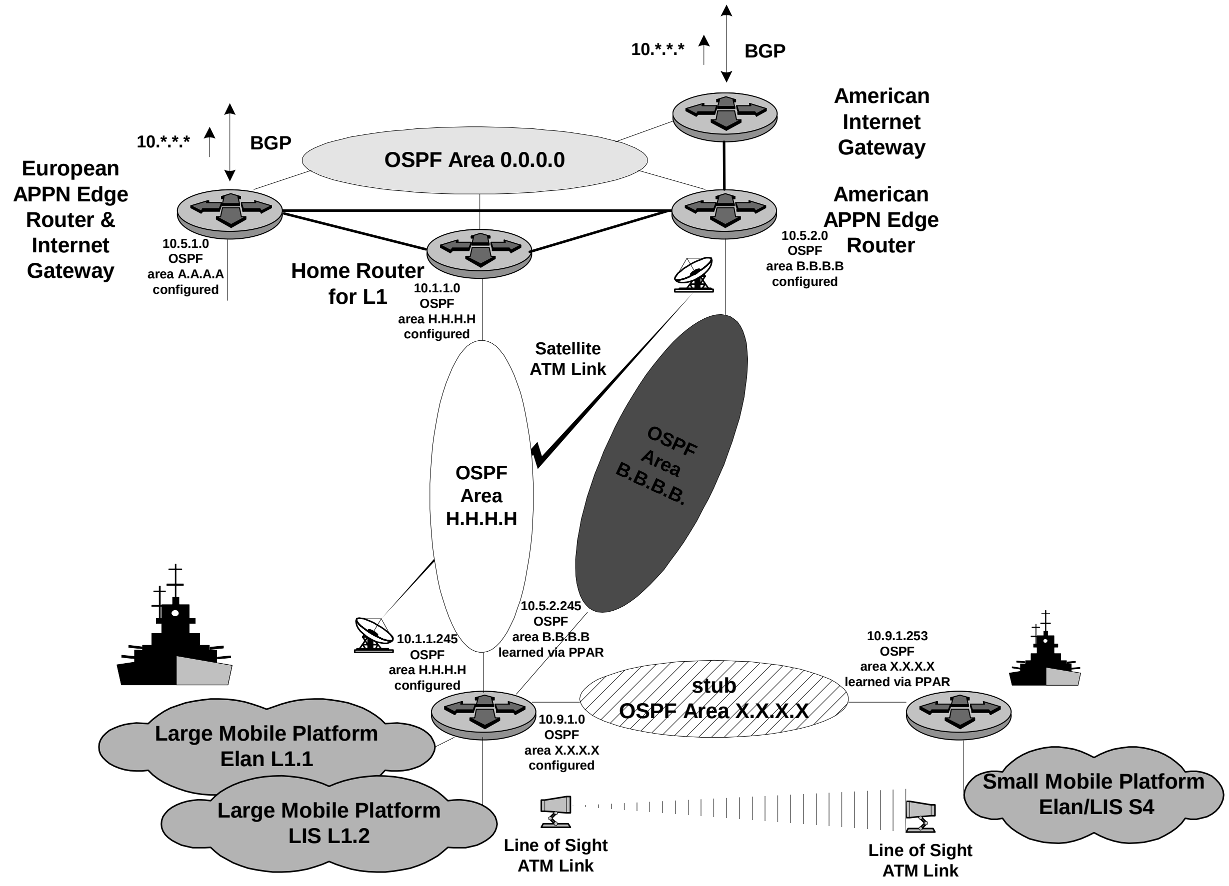

Figure 13: Large platform changes access points

Finally, the small ship S1 joins the fleet and establishes a line-of-sight link to the large ship L1 (see Figure 14). As with the mobile edge router on S4, the mobile edge router on S1 joins the OSPF area X.X.X.X. OSPF area X.X.X.X now advertises the reachability to the subnets of S4 and S1 into the OSPF areas A.A.A.A and H.H.H.H through the internal BGP session of mobile router L1.

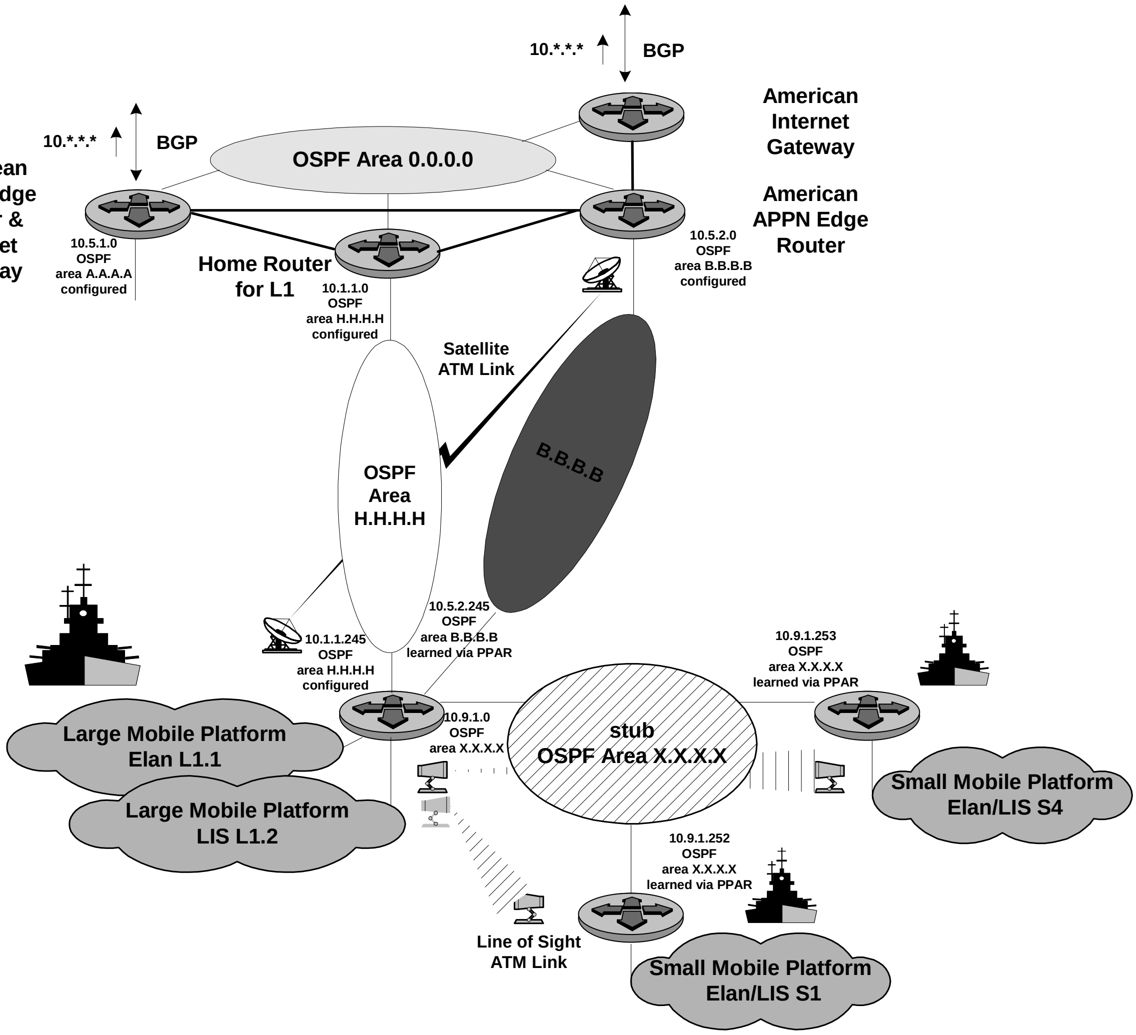

Figure 14: Small platform joins the mobile group

5 QoS Considerations

The need for Quality of Service for mobile networks is accentuated by the introduction of mobility and the utilization of wireless links. Mobility introduces frequent changes of network topology while wireless links provide lower and fluctuating bandwidth, increased delays, a higher error rate and less reliability than fixed links.

The mobile environment also adds challenges to the maintenance of Quality of Service over time. For example, consider what happens when a mobile platform changes links. The resources formerly reserved on the old link have to be reallocated. However, in some cases, the originally requested resources might not be available on the new link. It is therefore important to provide some QoS guarantees for applications used on mobile platforms in order to distribute the scarce resources.

IBM deals with QoS in mobile networks at both the IP and ATM layers as well as with the interaction between the IP and ATM layers.

5.1 IP QoS

There are currently two approaches to provide Quality of Service in the Internet: RSVP and Diffserv. The standardization process of RSVP is almost completed, while the standardization of Diffserv has just started. However, both approaches have already shown that QoS must be implemented together with policy mechanisms, service-level agreements, etc. to provide IP QoS.

Resource Reservation Protocol (RSVP)

RSVP as defined by the IETF [14] is part of the integrated services framework (int-serv). It is a protocol to reserve resources for unidirectional data streams, strictly distinguishing between senders and receivers. Multiple senders advertise themselves by sending messages downstream to create an RSVP session. Since RSVP is receiver-oriented, the receivers reserve the path upstream to the senders. When multiple receivers make reservations for that session along the path towards the senders, these reservations are merged as soon as they share a partial path. Breaking a fundamental principle of IP, RSVP stores state information for each flow in routers. However, the reservations use soft states, which are not permanent but time out after some period. Therefore, reservations must be refreshed periodically. Int-serv offers two types of QoS: Guaranteed Rate (a peek-rate reservation) and Controlled Load (a mean rate reservation, with possible extra bandwidth).

RSVP can cope with topology changes and link failures, which is important in the context of mobile networks. It is clear that bandwidth changes can have a significant impact on RSVP reservations. Methods to solve that issue require further research.

An RSVP extension allows the assignment of a preemption priority and a defending priority for each reservation. The former is used when deciding to accept or reject a request for a new reservation, and the latter is used once the reservation is in place [15]. This offers more flexibility for the attribution of bandwidth than the existing First Come First Served policy.

Differentiated Services (Diffserv)

The Diffserv working group of the IETF is using a new framework to provide QoS [16]. The general concept is to have traffic classification and conditioning at the edge of the network, i.e. the traffic is classified, policed and marked upon ingress to the network. The assigned mark determines the treatment to be received by the packet as it traverses a Diffserv domain. Diffserv reuses the TOS (Type of Service) field of the IP header for this purpose and renames it to DS (Differentiated Services). Within the network, local behaviors referred to as Per Hop Behaviors (PHBs) are defined. Along the path, each node has to comply with a certain PHB to forward an IP packet. As opposed to ATM or RSVP, Diffserv is a per-hop and per-packet QoS approach instead of a per-connection end-to-end QoS approach. This minimizes the per-flow state information stored in the nodes. Only the effects of concatenated PHBs can be observed at connection end points. The architecture provides service differentiation in each direction of traffic flow independently. Each packet using Diffserv indicates the QoS it requires, using the DS field in its header.

Rather than a "quantitative QoS" like RSVP, Diffserv offers a "qualitative-QoS". A quantitative service guarantee is defined by specific values for one or more performance metrics throughout the domain, whereas a qualitative guarantee ensures that the metrics of one service are better or worse than those of another. A fundamental difference between the two service types is that a qualitative service does not deny requests but allows for the service level to gracefully degrade as load increases or network capacity decreases.

This means that changes in bandwidth can be better accommodated with Diffserv than with RSVP. Diffserv defines two PHBs: EF (for Expedited Forwarding) and AF (for Assured Forwarding). EF offers the equivalent of a virtual leased line, while AF only assures the relative priority between packets set to one of four AF classes. EF requires absolute guarantees on the bandwidth and should therefore not be allocated a large portion of the link bandwidth. The rest should then be assigned to AF. This makes Diffserv a simple solution to offer relative QoS guarantees, which fits particularly well in the mobile environment.

Diffserv is expected to be widely accepted and deployed soon. IBM has previously reported implementation experience with Diffserv [17].

5.2 ATM QoS

When the bandwidth available on a wireless link decreases significantly, mobile ATM looks for an alternate path to accommodate existing and new connections. Among the candidate links, the path with the most available resources is selected to ensure the highest QoS possible. However, there may be insufficient alternative bandwidth available for providing all of the of previously used resources.

Several means are used to cope with a scenario of reduced resources in mobile wireless networks. One way to handle a reduction of bandwidth is to release some virtual connections. To distinguish between degrees of importance of connections and to treat them accordingly, a priority scheme allows selective handovers to hand over high priority connections. Meanwhile, low priority ones are dropped if the alternative links can not accommodate all previous connections.

Whenever appropriate, a graceful decrease of QoS for virtual connections complements the above scheme, making it possible to decrease the QoS without breaking the QoS contract. This requires a relative instead of an absolute approach to service. Therefore the service is not defined by absolute numbers of performance metrics but rather by comparison to others. High priority connections or flows are guaranteed a ‘better’ service than those having lower priority.

To decrease bandwidth used by UBR connections is straightforward. For other service categories, like CBR, the QoS has to be renegotiated. Today, only ABR provides feedback about the current state of the network using back pressure to decrease cell flow. The ATM forum is working on QoS renegotiation using the MODIFY signaling message to also enable other categories to provide feedback.

5.3 IP over ATM QoS

To provide IP QoS over a mobile ATM network requires solutions for two issues. First, IP QoS has to be mapped to ATM virtual connections. Second, the ATM layer has to interact with the IP layer to reflect the status of the network. For example, a status update can indicate available bandwidth or congestion. This feedback from the ATM switches to the IP routers has to be defined. These two issues are interdependent. The accuracy of mapping IP QoS to ATM QoS also determines how the feedback to the IP layer must be done.

Mapping from IP QoS to ATM

The need for QoS in IP networks, and subsequently an adequate mapping of IP QoS to ATM, is not due solely to the introduction of mobility. Nevertheless, the mapping (see Fig. 15) is an important factor in IP QoS support in mobile networks. Different multiplexing strategies for IP over ATM allow for a variety of mappings: there can be one virtual connection for each pair of routers, one per IP conver- sation (e.g. TCP connection, UDP flow), one per application type (e.g. ftp, telnet) or one per service class (e.g. as defined in the IP DS field). In addition, ATM traffic categories such as CBR, ABR, UBR and VBR with their respective QoS parameters offer the possibility of a very fine-grained distinction between service classes. Therefore the possibility of an accurate mapping exists.

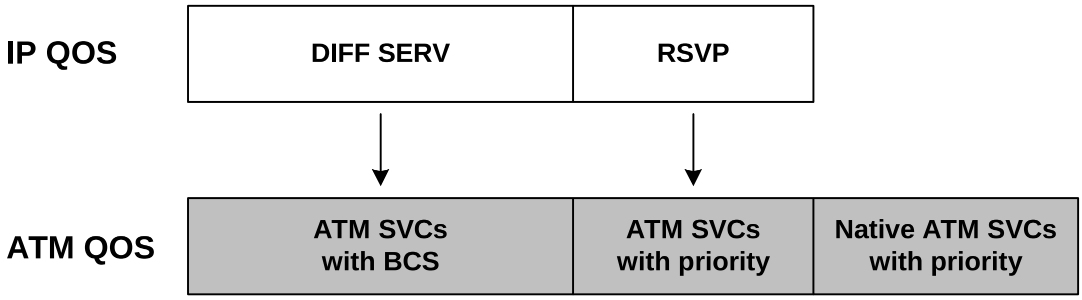

Figure 15: Mapping from IP QoS to ATM

The ATM Forum is working on support for Diffserv over ATM in the traffic management working group [18]. The approach is to map PHBs to ATM service categories. Part of this effort is to extend UNI, PNNI and AINI signaling by adding an optional behavior class selector (BCS) information element to SETUP and CONNECT messages [19]. This information element is used to convey through the network the information about the per-hop-behavior associated with the virtual connection.

Even though the current Diffserv architecture does not define end-to-end services, there are numerous services that have been discussed in literature and presented at the IETF. One example is the premium service, which is a low delay and low loss service. It is defined by a peak rate only. It is deployed using the EF PHB at the intermediate nodes of the network and by enforcing the peak rate at the network edges. Any traffic in excess of the defined peak rate is discarded at the network edge to meet the delay requirements of the service. Obviously, this premium service can be mapped to ATM CBR in order to meet its low loss and delay requirements. The peak rate of the premium service can be mapped in a straightforward fashion to the CBR PCR and the conformance definition of the CBR service is enforced in the usual manner. Other services can be mapped in an analogous way.

For the integration of RSVP and ATM signaling in support of the int-serv model, there are two issues to be solved. One is again the mapping of the int-serv QoS to ATM QoS. The other is VC management, which involves determining how many VCs are needed and which traffic flows are routed over which VC. Both the ATM forum and the IETF have proposed solutions to these issues [20]. The call priority scheme currently being discussed at the ATM forum [21] offers a possibility to map RSVP priorities to ATM call priorities for virtual connections.

Feedback from ATM to IP

Depending on how a bandwidth fluctuation is handled at the ATM layer, the QoS support at the IP layer, as well as the IP layer itself, will be impacted differently.

First of all, consider the straightforward case where the ATM technology is able to handle the bandwidth fluctuation without releasing any connections. This is transparent to the IP layer. If VCs are always rerouted and their bandwidth, as presented to the IP layer, remains the same, then nothing will change at the IP layer. QoS reservations will simply remain in place.

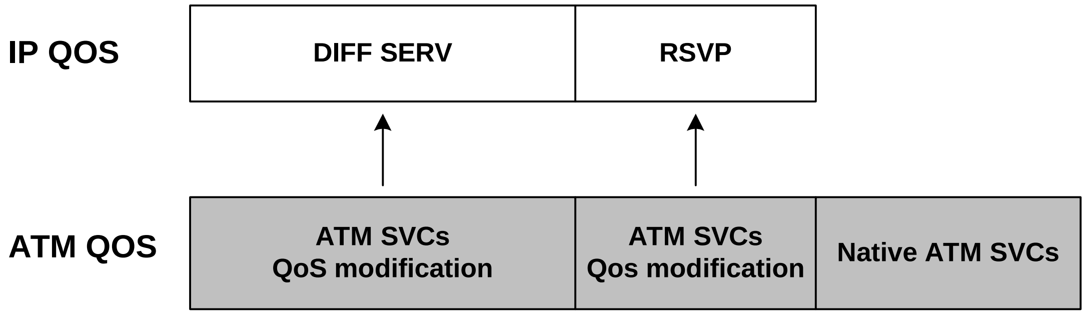

Figure 16: Feedback from ATM to IP

However, following a change in the mobile network, there may be a significant decrease in total available bandwidth. This will require additional actions from the IP layer compared to current QoS reservations. In RSVP this means that certain existing reservations must be dropped following the preemption priority assigned to each reservation. In Diffserv, packets belonging to the lowest service class will be discarded first. If service classes are mapped to different virtual connections, then ATM can release particular connections carrying the lower class IP traffic.

For IP to be able to react optimally to changes in the mobile network, information about available bandwidth is useful (see Fig. 16). This information can be provided directly by the ATM switch to which the IP router is attached.

In most cases, mobility at the ATM layer impacts the IP layer as well. While IP is designed to function with changing topologies, special attention has to be given when dealing with RSVP reservations (like preemption). Since qualitative guarantees better accommodate the change of bandwidth in mobile links, a rather conservative mix of RSVP and Diffserv traffic on these links will most likely be the solution of choice (smaller portion of RSVP traffic). Besides the actual QoS-guarantees mechanism, it is the policy function that will make the whole system controlable.

6 Conclusion

IBM is committed to developing an industry-leading technology for supporting mobile networks. The key features of this technology are:

- Seamless integration of mobile and fixed networking infrastructure.
- Optimal utilization of network resources.
- Robust connectivity despite mobility.

The IBM Mobile IP/ATM solution provides a seamless IP architecture that takes full advantage of the underlying ATM network. The architecture uses Mobile ATM and Proxy PAR protocols to ensure that IP traffic uses the optimal network path to its destination. When centralized IP-based services exist, mobility is hidden from the IP layer so that fast, efficient ATM switching is used to maintain connectivity to these “home” services. When distributed IP-based services are desired, Proxy PAR is used to locate available local access to these services. Therefore, no manual IP reconfiguration is required as a mobile network changes location. The IP addressing architecture ensures that knowledge about the location of a mobile network is confined to the access point network. This provides a scaleable solution for large networks and for connecting to the Internet. The architecture operates with either a separate access point network provided by a vendor or an access point network integrated with a corporation’s traditional backbone network. It even permits a “guest” network to access the services of a larger mobile platform in the absence of a connection to a fixed network access point. Thus, this combination of IP and ATM technologies provides a flexible, seamless design for incorporating mobile networking into a customer’s existing corporate network.

IBM will not restrict a QoS solution for IP over ATM to one type of effort. Several types will coexist and be used in the future. Nevertheless, initial emphasis will be given to support Diffserv in the mobile environment. Once the proposed solutions, Diffserv and RSVP, are ready for deployment over ATM, they will be integrated in IBM's support for Quality of Service in mobile wireless networks. This integration will include a mapping of IP QoS to ATM QoS, handling ATM QoS accordingly and providing feedback from the ATM layer to IP.

7 References

[1] "IBM Mobile ATM Networking, Technology Overview v2.0", IBM Research Laboratory, August 1998

[2] "IBM Mobile ATM Networking, Network Design and Addressing Considerations v2.0", IBM Research Laboratory, January 1999

[3] "Multi-Protocol Over ATM Specification v1.0", af-mpoa-0087.000, ATM Forum, July 1997

[4] J. Luciani, D. Katz, D. Piscitello, B. Cole, N. Doraswamy, "NBMA Next Hop Resolution Protocol (NHRP)", Internet RFC 2332, April 1998

[5] J. Moy, "OSPF v2.0", Internet standard, STD 0054, April 1998

[6] Y. Rekhter, T. Li, "A Border Gateway Protocol 4 (BGP-4)", Internet RFC 1771, March 1995

[7] "PNNI addendum for mobility extensions v1.0", af-ra-0123.000, ATM Forum, April 1999

[8] M. Laubach, J. Halpern, "Classical IP and ARP over ATM", Internet RFC 2225, April 1998

[9] "LAN Emulation over ATM 1.0", af-lane-0021.000, ATM Forum, January 1995

[10] "PNNI Augmented Routing (PAR) v1.0", af-ra-0104.000, ATM Forum, January 1999

[11] P. Droz, T. Przygienda, "Proxy PAR", Internet draft, draft-ietf-ion-proxypar-arch-01.txt, February 1999 (work in progress)

[12] S. Alexander, R. Droms, "DHCP Options and BOOTP Vendor Extensions", Internet RFC 2132, March 1997

[13] R. Hinden, M. O'Dell, S. Deering, "An IPv6 Aggregatable Global Unicast Address Format", Internet RFC 2374, July 1998

[14] J. Wroclawski, "The Use of RSVP with IETF Integrated Services", Internet RFC 2210, September 1997

[15] S. Herzog, "Signaled Preemption Priority Policy Element", Internet Draft, draft-ietf-rap-signaled-priority-03.txt, February 1999

[16] S. Blake et al. "An Architecture for Differentiated Services", Internet RFC 2475, December 1998

[17] S. Nadas, et al., "IBM Diffserv Implementation Experiments", Internet Draft, draft-nadas-diffserv-experience-00.txt, June 1999

[18] "Enhancements to Support IP Differentiated Services and IEEE 802.1D over ATM Draft", btd-tm-diff-01.00, ATM Forum, July 1999

[19] M. Wiles, "Signalling support for Diffserv and 802.1D", 99-0320, -1,-2, ATM Forum, July 1999

[20] L. Berger, et al., "A Framework for Integrated Services and RSVP over ATM", Internet RFC 2382, "RSVP over ATM Implementation Requirements", Internet RFC 2380, August 1998

[21] G.R. Ash, et al., "Motivation for Support of Priority Services in ATM Networks" and "Signaling Requirements for the Support of Priority Services in ATM Networks", 99-0348 and 99-0350, ATM Forum, July 1999

ANNEX A Proxy PAR configuration in a mobile environment

Proxy PNNI Augmented Routing (Proxy PAR) [10,11] is a protocol standardized by the ATM Forum and being standardized by the IETF to improve and simplify the operations of IP over ATM networks. Using Proxy PAR, any application or service running on an ATM network can register itself with the network. Any Proxy-PAR capable device can query the ATM network to find about particular applications or services registered by other devices. As an example, all routers in a specified part of an ATM network can automatically find their peers and bring up an IP network. Proxy PAR can be used to find OSPF, BGP interfaces and ARP, DHCP, or DNS servers or any other network service. In a mobile environment, Proxy PAR can be used to locate peer routers or peer servers when a mobile platform moves and connects to a different part of the network.

Routers can indicate a scope when registering a network service. The description of the service is only flooded within the ATM network up to the PNNI level matching the scope specified at the registration. This scoping mechanism is used to ensure that a mobile edge router only receives the services registered by the APPN edge router local to the current access point. As shown in Figure 17, APPN edge routers register their IP routing services with a scope equal to the PNNI level 64. Level 64 is the level at which large mobile platforms can integrate the PNNI hierarchy of the APPN. All the mobile platforms currently attached to this access point (i.e. which integrated the PNNI hierarchy at this location at level 64 or lower) receive the services registered. The scope, however, prevents the IP routing services of an APPN edge router from being advertised to other APPN edge routers (and by extension to mobile platforms currently attached to other access points), because the common PNNI level at which access points are logically connected is level 32.

Similarly, the IP routing service for the “guest node” OSPF area is registered by each mobile edge router of a large mobile platform with a scope equal to the PNNI level 72. Level 72 is the level at which small mobile platforms can integrate the hierarchy of a large mobile platform. Therefore only the small mobile platforms that have integrated into the hierarchy of the large mobile platform receives the advertisement for the IP routing services available on the large platform.

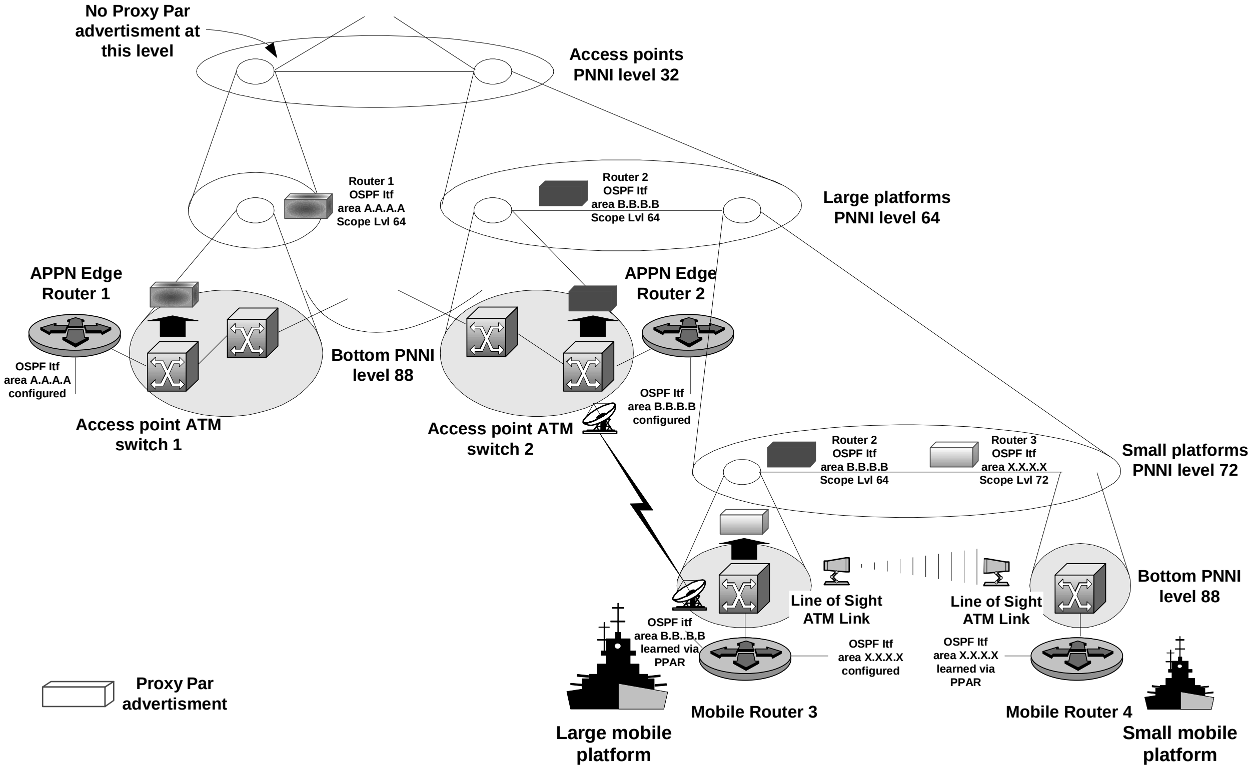

Figure 17: Use of the Proxy PAR scope to match the IP and ATM topology

One should note that the mobile edge router of a small platform receives the advertisements for both the services of the large platform’s router and those of the APPN edge router. The mobile edge router is however configured to consider only the services with the lowest scope.