Chapter 6

A Traffic Differentiation Method for IPv6 over ATM Networks

Jorge Sá Silva , Nuno Veiga , Sérgio Duarte and Fernando Boavida
{sasilva,nunov,sduarte,boavida}@dei.uc.pt
Centro de Informática e Sistemas da Universidade de Coimbra
Polo II - Pinhal de Marrocos
P-3030 Coimbra – PORTUGAL
Tel.: +351 39 790000 Fax.: +351 39 701266

Key words: Multimedia network management models and architectures

Abstract: At present, the number of Internet users continues to increase dramatically, accompanied by the large-scale introduction of new applications with multimedia capabilities. New users and applications are constantly demanding better performance from the Internet, if possible with some forms of quality of service (QoS) guarantees. Knowing that no single technology is able to provide such guarantees in the Internet environment – which, by definition is a multi-technology environment – the exploration of promising technologies capable of traffic differentiation that are applicable to parts of the Internet appears as a possible way forward. This paper addresses and explores the benefits of the Internet Protocol version 6 (IPv6) and Asynchronous Transfer Mode (ATM) integration in multimedia environments. A method for traffic differentiation in IPv6-over-ATM networks is proposed, based on the mapping of different IPv6 flows into the address resolution functions of ATM. The proposal is evaluated in terms of introduced overhead, by means of simulation studies, supported by an experimental platform and an IPv6-over-ATM simulator developed for this purpose.

1. INTRODUCTION

The current size of the Internet makes it very difficult to incorporate radical changes; yet, at least some changes are required in order to allow the support of multimedia and other QoS-requiring applications. Two of the technologies that, at the moment, have some potential to support multimedia applications are the ATM technology and the new version of the IP protocol, IPv6.

IPv6 was developed to overcome IPv4 limitations. The main improvements observed in IPv6 are concerned with address space – to overcome the current address exhaustion of the IPv4 – QoS support, self-configuration, security, efficient routing, multicast and anycast. On other hand, ATM is a proven technology, with good capabilities in terms of QoS support, widely deployed in Metropolitan Area Networks and Wide Area Networks. Nevertheless, when compared to other technologies, ATM is less deployed in local area network environments because of the complexity and the cost associated with it. On the other hand, ATM offers the flexibility to transport different services, respecting their needs and properties.

The combination of IPv6 and ATM places a set of new challenges, when compared to the IPv4 environment. As will be described in section 2, there are some key differences that require some extra functionality. In addition to identifying those differences, section 2 also summarises the IPv6-over-ATM "state of the art". Section 3 presents the 'Mars Extensions Developed for IPv6 over ATM' (MEDIA) project, in which scope the presented work was developed. In this section a new solution for the provision of quality of service in IPv6-over-ATM networks is also presented. Evaluation results concerning the proposed solution are presented and discussed in section 4. The conclusions and topics for further work are addressed in section 5.

2. IPV6-OVER-ATM NETWORKS

The main difference in IPv6 and ATM integration when compared to the IPv4 case is due to address resolution. In fact, in IPv4 environments using an overlay-model layer 3 to layer 2 address translation is made by auxiliary protocols like ARP or ATMARP. In contrast, IPv6 does this task by itself although, on other hand, it assumes a broadcast connectionless medium at layer 2 in order to implement Neighbour Discovery, Router Discovery and Address Configuration tasks [NAR96]. The apparent reason for this assumption is that IPv6 was developed assuming that the underlying network is an Ethernet.

In the Neighbour Discovery process, when the IPv6 layer needs to send an unicast message to some IPv6 destination address it looks for the corresponding layer 2 address in a translation table. If this layer 2 address is not found, the IPv6 layer sends a Neighbour Solicitation (NS) message to all in-link nodes. The node that has the corresponding IPv6 destination address receives the NS message and, if it accepts the connection, it establishes a point-to-point VCC and sends a Neighbour Advertisement (NA) to the calling node (Figure 1). When this is received, the source node can send the pending information, due to the bi-directionality of the established VCC.

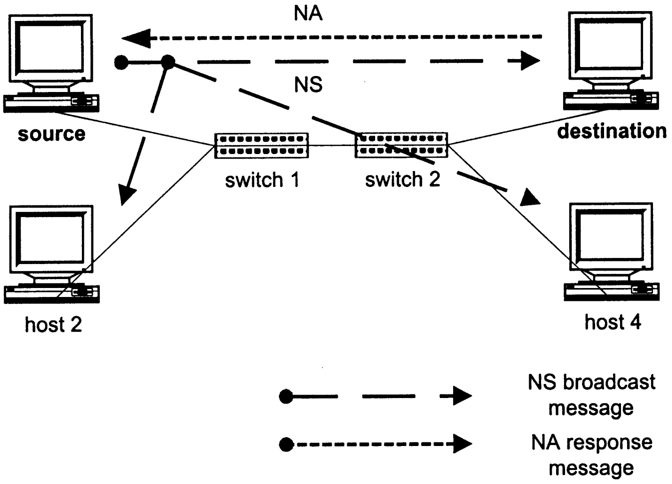

Figure 1. Neighbour Discovery process

The above mentioned process – originally based on the use of broadcast connectionless media – requires a new sublayer between the IPv6 layer and the ATM layer, hereafter called IPv6-over-ATM sublayer, that emulates the broadcast connectionless environment not natively provided by ATM networks.

Several approaches have been developed in recent years for the deployment of IPv6 over ATM. The Logical Link (LL) concept was introduced in [SCH96]. The LL is the ATM node set which can send and receive broadcast traffic. Each LL layer should emulate the layer 2 functionalities assumed by the IPv6 layer and, at the same time, should offer fault tolerance and redundant mechanisms. According to this proposal, each IPv6-over-ATM sublayer must capture all the broadcast traffic and send it to the Logical Link Server (LLS). This entity receives all broadcast traffic and sends it to all ATM nodes in the LL. There is no traffic differentiation and/or QoS mechanism present in this solution. In fact, this solution is totally based on the use of best-effort VCs. The LL concept can be seen as a "light-weight" solution, when compared to the MARS server approach that will be described next.

The Multicast Address Resolution Server (MARS) solution [ARM96] was originally proposed to translate IP multicast addresses to sets of ATM addresses, in multicast Classical IP environments [LAU94]. In [ARM99A] and [ARM99B], the use of the MARS approach by the IPv6-over-ATM sublayer is proposed. In a host, all the broadcast traffic sent by the IPv6 layer is captured by the IPv6-over-ATM sublayer and sent to the MARS server. After arrival at the server, the traffic is routed to all the cluster members.

The communication between off-link members uses, by default, *mrouters* as agglutination points. The use of shortcuts, using flow trigger or source trigger mechanisms, is also supported by Next Hop Resolution Protocol (NHRP) [LUC98]. In [TUR98], a study is presented where the use of a bit to differentiate shortcuts from the same source but with different QoS needs is proposed.

The choice of using permanent connections in the [ARM99A] and [ARM99B] model implies that the number of necessary static VCs is too high, as all the stations need to receive the broadcast traffic from the Neighbour Discovery process. On the other hand, the use of switched VCs implies a substantial transit of overhead messages whenever a station needs to transmit broadcast information. To overcome these problems [ARM96T] proposes the use of the control VCs or the ClusterControlVC to broadcast the Neighbour Discovery messages. It is a simple model, as it only needs few modifications to adapt the MARS server code.

[ATK96] proposes the introduction of an IPv6 layer version specific to the ATM technology. A new Neighbour Discovery process is presented, and it is also suggested the use of 6 tokens at the auto-configuration process. This proposal uses an NHRP extension to find the layer 2 address given the layer 3 address. Although the use of an ATM-specific IPv6 version could offer better performance results, it can be looked at as a major impediment for the IPv6 success, as it reduces its generality.

The MARS approach appears, then, as the most consistent approach for the provision of IPv6 over ATM networks. This was also recognised by IETF (Internet Engineering Task Force), which decided to adopt the MARS solution for IPv6-over-ATM environments. Nevertheless, the MARS model presents some limitations: it does not use some of the IPv6 benefits and there are some UNIv4 properties that are not sufficiently explored.

3. TRAFFIC DIFFERENTIATION

In the MEDIA project we are developing some extensions to the MARS approach, keeping full compatibility with the standard approach. The introduction of a new, original model would render difficult the success of the underlying ideas – no matter how interesting they would be – because of the wide acceptance of MARS and its support by IETF.

The MEDIA project intends to propose new extensions to CLIP (Classical IP) environments, that make use of the benefits inherent to IPv6 and UNIv4, not yet explored by other models. We are particularly interested in the use of additional QoS parameters in Neighbour Solicitation (NS) messages and in the IP multicast join requests – MARS_JOIN messages – for the support of traffic differentiation.

In unicast communications, we propose the use of the Flow Label Field or the Traffic Class Octet (TCO) IPv6 field in NS messages, using the Differentiated Services model [BER99], to transport the information concerning the requested QoS class. This value, when received by the destination node, is to be translated to the corresponding ATM QoS. The destination node then establishes a bi-directional point-to-point connection to the source node (figure 2).

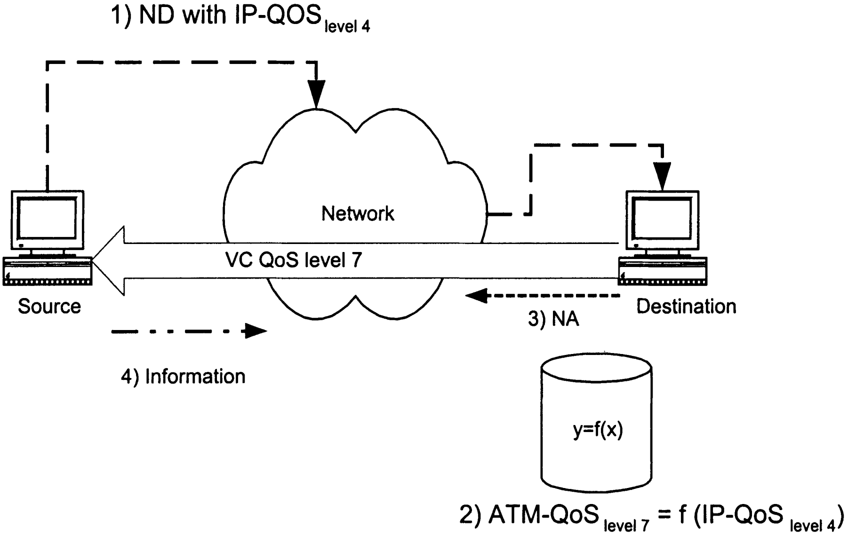

Figure 2. Proposed algorithm

As proposed in the Internet2 recommendations [INT97M], future services must support different QoS levels. When compared to the original process, the proposed solution respects QoS requirements.

As proposed by the Differentiated Services model, the QoS value at the IP layer is specified as an integer. The option between the use of the Flow

Label field and the TCO field depends on the number of bits necessary to distinguish the various QoS levels.

We propose the use of a single MARS server to all the ATM network. In our model the NS messages are directly sent to the destination nodes bypassing routers and the NHRP process, offering a more efficient solution when compared with [TUR98].

To not confuse the IDMR protocols and to not disbalance the multicast distribution tree, in case of multiple egress Mrouters, we use the single gateway principle, as proposed in EARTH [SMI97].

When, after receiving the NS message, the destination node cannot establish the requested VC with the required set of QoS parameters, it must ignore the NS message. The source node must, then, implement a timer that is activated whenever an NS message is sent. If the corresponding NA message does not arrive within a specified time interval, the source node must retransmit the NS message.

Although the majority of the existing approaches for QoS provision use the Integrated Services paradigm, Differentiated Services systems are now gaining importance, as this approach is far more compatible with the general Internet paradigm. The model presented in this paper is applicable when the ATM sub-network is part of a larger network respecting the Differentiated Services model [BER99]. The different QoS levels – that are at the basis of traffic differentiation – must be pre-configured, negotiated with the system administrator, and must be known by all the systems. The use of a set of pre-configured quality of service levels is a simplification that requires configuration by systems management functions. The use of a dynamical protocol to offer non-static QoS levels would complicate the proposed model and is outside the scope of the present proposal.

In the MEDIA project we are implementing an IPv6-over-ATM testbed using an experimental ATM network. The testbed is composed of two switches – a Fore200ASX switch and a 3COM 7000HD switch – and six hosts equipped with 155 Mbps ATM NICs (Figure 3).

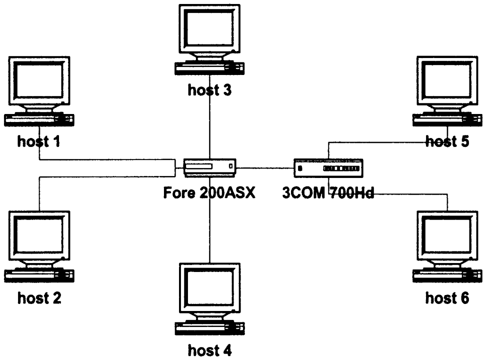

Figure 3. The MEDIA project testbed

We also developed the IPv6-over-ATM sublayer module and the MARS modules – the client module, the server module and the Multicast Server module. Each host has the modules presented in Figure 4.

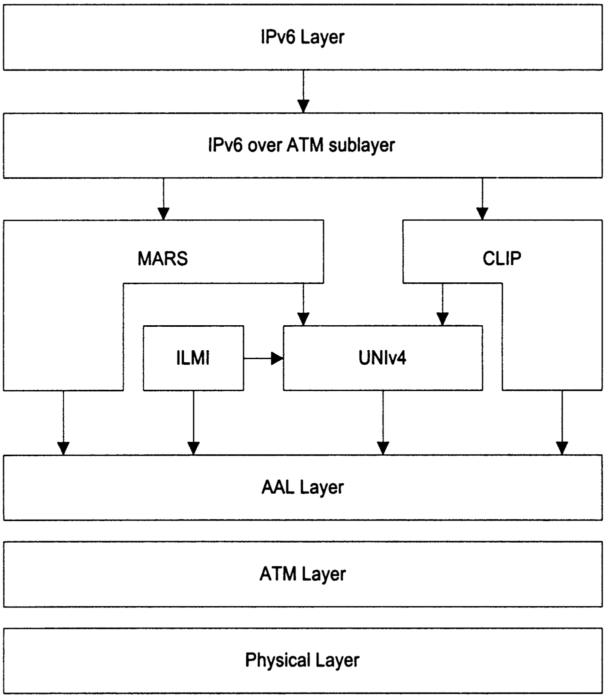

Figure 4. Protocol stack of IPv6 over ATM hosts in the MEDIA testbed

The developed CLIP version does not respect the original version. As IPv6 performs the address translation process, ATMARP is no longer present.

To complement the results obtained in the testbed, the MEDIA project team developed Sianet - an IPv6-over-ATM network simulator. Sianet is a C++ discrete event-driven simulator that offers a user-friendly graphical interface (Figure 6).

The simulator is composed of the engine module – that implements all of the previously described IPv6 and ATM functionalities – the user interface module, and the PNNI module. The user builds the network to be simulated using the graphical interface to specify hosts, switches and links, and selects a number of applications that run in each host (Figure 5). For each network element, different parameter sets can be specified.

Up to the present, we have developed some CBR and VBR applications. We plan to build some IPv6 native applications, as they have some different statistical properties when compared to IPv4 applications.]

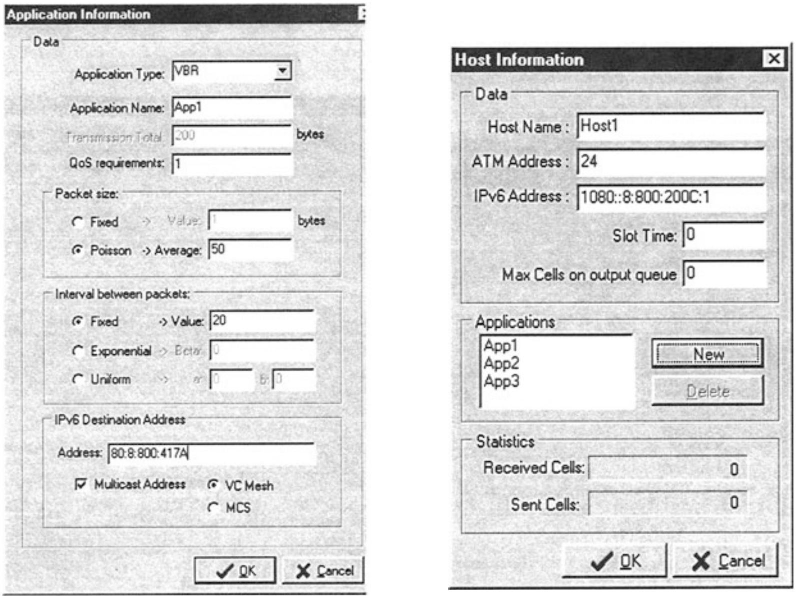

Application Type: VBR
Application Name: App1
Transmission Total: 200 bytes
QoS requirements: 1
○ Fixed -> Value: 1 bytes
● Poisson -> Average: 50
● Fixed -> Value: 20
○ Exponential -> Beta: 0
○ Uniform -> a: 0 b: 0
Address: 80:8:800:417A
☑ Multicast Address ● VC Mesh
○ MCS

Host Name : Host1
ATM Address : 24
IPv6 Address : 1080::8:800:200C:1
Slot Time: 0
Max Cells on output queue 0
App1
App2
App3
Received Cells: 0
Sent Cells: 0

Figure 5. Examples of host and application parameter specifications

The validation of Sianet was done using the developed testbed. Also, some size and time overheads were obtained from the real platform.

In the context of the MEDIA project, Sianet is being used for the study of the proposed algorithm behaviour in scenarios other than the experimental testbed scenario.

In this project we are also studying and proposing traffic differentiation algorithms for multicast environments. We are particularly interested in the use of a new set of TLVs in MARS_REQUEST and MARS_JOIN messages to convey QoS requirements information in multicast communications. We are also developing some studies to evaluate the use of ATM multicast addresses in IPv6-over-ATM environments. ATM multicast addresses were introduced with UNIv4 but only for anycast environments. Although some PNNI modifications are necessary, the use of one-to-one mapping of IP multicast addresses and ATM multicast addresses is very powerful and can substantially reduce the number of MARS_MULTI messages.

4. OVERHEAD ANLYSIS

We made some simulation studies in order to quantify the worst possible overhead of the proposed solution in relation to the original ND process. When compared to the original process, the proposed algorithm should need some extra time to translate the IPv6 QoS requirements into QoS characteristics (more precisely, non best-effort VCs that correspond to the required QoS level) at the ATM layer. We assumed the use of hashing algorithms in the translation tables.

The following elementary mean time-overhead values were used for the simulation. These values were measured in the MEDIA testbed.

Table 1. Mean time overhead values

Overhead	Time (μs)
IPv6 layer: downwards direction	30
IPv6 layer: upwards direction	4
IPv6 layer: address table look up	0.6
IPv6 over ATM sublayer: unicast downwards direction	0.5
IPv6 over ATM sublayer: multicast downwards direction	384
IPv6 over ATM sublayer: upwards direction	0.5
CLIP layer: downwards direction	3
CLIP layer: upwards direction	2.2
MARS Server: processing of MARS_REQUEST	189
MARS Client: processing of MARS_MULTI	154
MARS Server: broadcasting of NS messages	120
MCS: broadcast messages	120
ATM layer: downwards direction	50
ATM layer: upwards direction	50

The measurement of each time value involved the identification of the initial and final points in the code of the operating system and driver. In these points the microtime() function was used to obtain the time values.

Table 2 presents some simulation results, in terms of the time needed to send 200 bytes of information from a host A to a host B, using a congested network composed of a varying number of intermediate switches (figure 6).

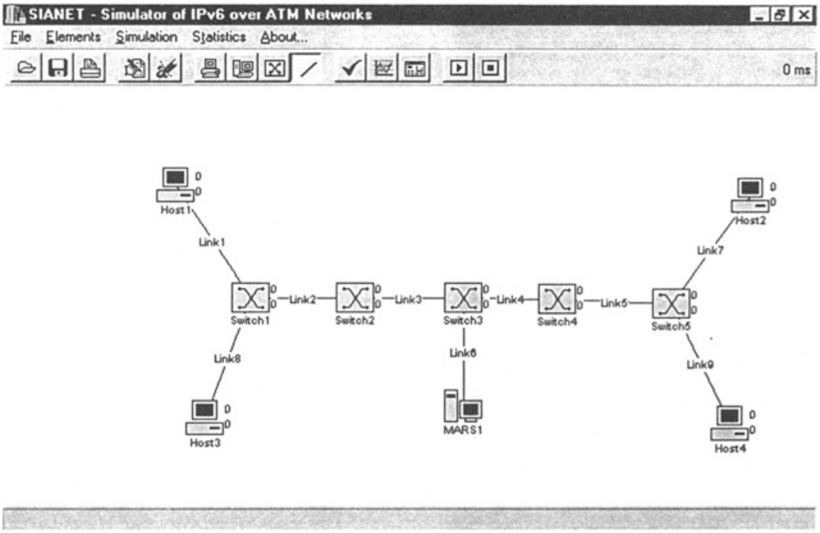

Figure 6. Sample configuration of a simulated network (5 switches)

The values obtained for each study were the result of the average of the values obtained in 10 simulation runs. The relative accuracy (for this number of simulations and for a confidence interval of 95 percent) varied between 1 and 4 percent. That is, it can be affirmed with 95 percent certainty that the real time average differs from the measured average in any given simulation run by a maximum of 1-4 percent.

Table 2. Time to send 200 bytes of unicast information

Number of switches	Time (μs)		
	original process	proposed process	overhead
1	20867	21447	580
2	41096	41688	592
3	61325	61929	604
4	81553	82170	617
5	101782	102412	630
10	202926	203618	692
20	405213	406030	817

As expected, the proposed solution does increase the time spent in the transmission of the information, but this increase is more or less constant (as the QoS mapping is done only once, prior to the establishment of the virtual channel) and negligible.

This same observation was confirmed by a second study. In this case, a network composed of 3 switches was used and the message lengths varied from 1 byte to 5000 bytes. Figure 7 presents the simulation results.

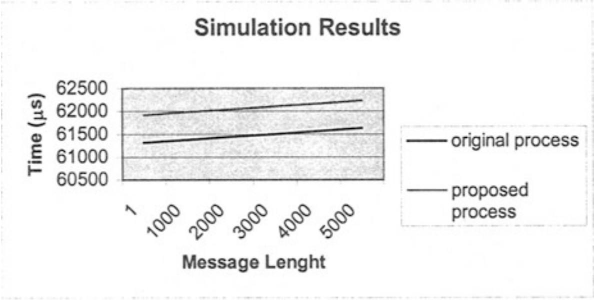

Figure 7. Transmission time as a function of message length, in IPv6-over-ATM unicast communication

As the message length increases, the relative overhead (that is, the ration overhead / message length) decreases, due to the fact that the introduced overhead affects only the establishment phase and is practically constant.

5. CONCLUSIONS

As multimedia and other QoS demanding applications become increasingly used in IP networks, traffic differentiation mechanisms take on a leading role in the efficient use of network resources. Several technologies have a good potential for the support of traffic with varying QoS requirements. Two of these – IPv6 and ATM – have been experiencing a growing acceptance over the last few years.

The existing proposals for the use of IPv6 in ATM networks were developed for best-effort environments and, thus, do not explore the traffic differentiation capabilities that are inherent to each of those technologies. The MEDIA project, briefly presented in the present paper, explores these capabilities, by proposing an approach to the mapping of IPv6 QoS requirements into VCs with appropriate quality of service characteristics. This mapping is done by an intermediate sublayer – the IPv6-over-ATM sublayer – using the standard, IETF-adopted address resolution mechanism.

Although the MEDIA project intends to analyse and propose new methods for unicast and multicast systems, this paper focused on the former case. After the presentation of the proposed QoS mapping solution, an analysis of the incurred overhead is performed, by means of simulation studies. The simulations showed that the proposed method introduces a small and almost constant overhead, which points to the fact that the proposed solution is feasible and that its inherent benefits have a negligible cost. Future developments of the proposed solution will address other IPv6 over ATM scenarios, like LAN emulation and MPOA.

ACKNOWLEDGEMENT

The work presented in this paper is partially financed by the Portuguese Foundation for Science and Technology, FCT, in the scope of the PRAXIS XXI program.

REFERENCES

[ARM96] G. Armitage "Support for Multicast over UNI 3.1 Networks based ATM Networks" RFC 2022 - 1996
[ARM96T] G. Armitage "Transient Neighbors for IPv6 over ATM" draft-armitage-ion-tn-01.txt - November 1996
[ARM99A] G. Armitage, P. Schulter, M. Jork and G. Harter "IPv6 over Non-Broadcast Multiple Access Networks" RFC 2491 - January 1999
[ARM99B] G. Armitage, P. Schulter and M. Jork "IPv6 over ATM Networks" RFC 2492 - January 1999
[ATK96] R. Atkinson "IPv6 over NBMA Networks" draft-ietf-ion-ipv6-nbma-00.txt - June 1996
[BER99] Y. Bernet et all "A Framework for Differentiated Services" draft-ietf-diffserv-framework-02.txt - February, 1999
[INT97M] Internet 2 "Internet Protocol Multicast Problem Statement", draft-bradner-multicast-problem-00.txt - September 1997
[LAU94] M. Laubach "Classical IP and ARP over ATM" RFC 1577 - January 1994
[LUC98] J. Luciani, D. Katz, Piscitello, B. Cole and N. Noraswamy "NBMA Next Hop Resolution Protocol (NHRP) " RFC 2461 - April 1998
[NAR96] T. Narten, E. Nordmark and W. Simpson "Neighbor Discovery for IP Version 6 (IPv6) " RFC 1970 - August 1996
[SCH96] P. Schulter "A Framework for IPv6 over ATM" draft-ietf-ion-ipv6atm-framework-00.txt
[SMI97] M. Smirnov GMD FOKUS "Efficient and Easy IP Multicast Routing Through ATM Networks", Proceedings of 7th IFIP Conference High Performance Networks 1997
[TUR98] S. Turner, and B. Staegemeir "Differentiated Services over Symetric NHRP Shortcuts" draft-turner-diff-nhrp-00.txt - March 1998