

Network Working Group
Request for Comments: 2492
Category: Standards Track

G. Armitage
Lucent Technologies
P. Schulter
BrightTiger Technologies
M. Jork
Digital Equipment GmbH
January 1999

IPv6 over ATM Networks

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1999). All Rights Reserved.

Abstract

This document is a companion to the ION working group's architecture document, "IPv6 over Non Broadcast Multiple Access (NBMA) networks". It provides specific details on how to apply the IPv6 over NBMA architecture to ATM networks. This architecture allows conventional host-side operation of the IPv6 Neighbor Discovery protocol, while also supporting the establishment of 'shortcut' ATM forwarding paths (when using SVCs). Operation over administratively configured Point to Point PVCs is also supported.

1. Introduction.

This document is an ATM-specific companion document to the ION working group's, "IPv6 over Non Broadcast Multiple Access (NBMA) networks" specification [1]. Terminology and architectural descriptions will not be repeated here.

The use of ATM to provide point to point PVC service, or flexible point to point and point to multipoint SVC service, is covered by this document.

A minimally conforming IPv6/ATM driver SHALL support the PVC mode of operation. An IPv6/ATM driver that supports the full SVC mode SHALL also support PVC mode of operation.

2. Specification Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [16].

3. PVC Environments

When the ATM network is used in PVC mode, each PVC will connect exactly two nodes and the use of Neighbor Discovery and other IPv6 features is limited. IPv6/ATM interfaces have only one neighbor on each Link. The MARS and NHRP protocols are NOT necessary, since multicast and broadcast operations collapse down to an ATM level unicast operation. Dynamically discovered shortcuts are not supported.

The actual details of encapsulations, MTU, and link token generation are provided in the following sections.

This use of PVC links does not mandate, nor does it prohibit the use of extensions to the Neighbor Discovery protocol which may be developed for either general use or for use in PVC connections (for example, Inverse Neighbor Discovery).

Since ATM PVC links do not use link-layer addresses, the link-layer address options SHOULD not be included in any ND message [11]. If a link-layer address option is present in an ND message, then the option SHOULD be ignored.

A minimally conforming IPv6/ATM driver SHALL support the PVC mode of operation. PVC only implementations are not required to support any SVC mode of operation.

3.1 Default Packet Encapsulation

Following the model in RFC 1483 [2], AAL5 SHALL be the default Adaptation Layer service, and (LLC/SNAP) encapsulation SHALL be default encapsulation used by unicast and multicast packets across pt-pt PVC links. As defined in [1], the default IPv6 packet encapsulation SHALL be:

```
[0xAA-AA-03][0x00-00-00][0x86-DD][IPv6 packet]
      (LLC)      (OUI)      (PID)
```

3.2 Optional null encapsulation

IPv6/ATM drivers MAY also support null encapsulation as a configurable option. When null encapsulation is enabled, the IPv6 packet is passed directly to the AAL5 layer. Both ends of the PVC MUST be configured to use null encapsulation. The PVC will not be available for use by protocols other than IPv6.

3.3 PPP encapsulation

The concatenation of IPv6 over PPP with PPP over AAL5 PVCs is not covered by this specification.

3.4 MTU For PVC Environments

The default IP MTU size for PVC links is 9180 bytes as specified in [7]. Other IP MTU values MAY be used.

3.5 Interface Token Formats in PVC Environments

When the ATM network is used in PVC mode interface tokens SHALL be generated using one of the methods described in section 5. Interface tokens need only be unique between the two nodes on the PVC link.

4 SVC environments

4.1 SVC Specific Code Points

4.1.1 ATM Adaptation Layer encapsulation for SVC environments

Following the model in RFC 1483 [2], AAL5 SHALL be the default Adaptation Layer service, and (LLC/SNAP) encapsulation SHALL be the default encapsulation used by unicast and multicast packets across SVC links.

4.1.2 Unicast Packet Encapsulation

As defined in [1], the default IPv6 unicast packet encapsulation SHALL be:

```
[0xAA-AA-03][0x00-00-00][0x86-DD][IPv6 packet]
  (LLC)      (OUI)      (PID)
```

4.1.3 Multicast packet encapsulation

As defined in [1], the default IPv6 multicast packet encapsulation SHALL be:

```
[0xAA-AA-03][0x00-00-5E][0x00-01][pkt$cmi][0x86DD][IPv6
packet]
      (LLC)      (OUI)      (PID)      (mars encaps)
```

The IPv6/ATM driver's Cluster Member ID SHALL be copied into the 2 octet pkt\$cmi field prior to transmission.

4.1.4 Optional null encapsulation

IPv6/ATM drivers MAY also support null encapsulation as a configurable option. Null encapsulation SHALL only be used for passing IPv6 packets from one IPv6/ATM driver to another. Null encapsulation SHALL NOT be used on the pt-pt SVC between the IPv6/ATM driver and its local MARS.

If null encapsulation is enabled, the IPv6 packet is passed directly to the AAL5 layer. Both ends of the SVC MUST agree to use null encapsulation during the call SETUP phase. The SVC will not be available for use by protocols other than IPv6.

If null encapsulation is enabled on data SVCs between routers, inter-router NHRP traffic SHALL utilize a separate, parallel SVC.

Use of null encapsulation is not encouraged when IPv6/ATM is used with MARS/NHRP/ND as described in [1].

4.1.5 MARS control messages

The encapsulation of MARS control messages (between MARS and MARS Clients) remains the same as shown in RFC 2022 [3]:

```
[0xAA-AA-03][0x00-00-5E][0x00-03][MARS control message]
      (LLC)      (OUI)      (PID)
```

The key control field values are:

The mar\$afn field remains 0x0F (ATM addresses)

The mar\$pro field SHALL be 0x86DD (IPv6)

The mar\$op.version field remains 0x00 (MARS)

The mar\$spln and mar\$tpln fields (where relevant) are either 0 (for null or non-existent information) or 16 (for the full IPv6 protocol address)

The way in which ATM addresses are stored remains the same as shown in RFC 2022 [3]

4.1.6 NHRP control messages

The encapsulation of NHRP control messages remains the same as shown in RFC 2332 [4]:

```
[0xAA-AA-03][0x00-00-5E][0x00-03][NHRP control message]
      (LLC)      (OUI)      (PID)
```

The key control field values are:

- The ar\$afn field remains 0x0F (ATM addresses)
- The ar\$pro field SHALL be 0x86DD (IPv6)
- The ar\$op.version field remains 0x01 (NHRP)


The ar\$spln and ar\$tpln fields (where relevant) are either 0 (for null or non-existent information) or 16 (for the full IPv6 protocol address)

The way in which ATM addresses are stored remains the same as shown in RFC 2022 [3]

4.1.7 Neighbor Discovery control messages

Section 5.2 of [1] describes the ND Link-layer address option. For IPv6/ATM drivers, the subfields SHALL be encoded in the following manner:

[NTL] defines the type and length of the ATM number immediately following the [STL] field. The format is as follows:



The most significant bit is reserved and MUST be set to zero. The second most significant bit (x) is a flag indicating whether the ATM number is in:

ATM Forum AESA format (x = 0).
Native E.164 format (x = 1).

The bottom 6 bits represent an unsigned integer value indicating the length of the associated ATM address field in octets.

The [STL] format is the same as the [NTL] field. Defines the length of the subaddress field, if it exists. If it does not exist this entire octet field MUST be zero. If the subaddress exists it will be in AESA format, so flag x SHALL be zero.

[NBMA Number] is a variable length field containing the ATM address of the Link layer target. It is always present.

[NBMA Subaddress] is a variable length field containing the ATM subaddress of the Link layer target. It may or may not be present. When it is not, the option ends after the [NBMA Number] (or any additional padding for 8 byte alignment).

The octet ordering of the [NBMA Number] and [NBMA Subaddress] fields SHALL be the same as that used in MARS and NHRP control messages.

4.2 UNI 3.0/3.1 signaling issues (SVC mode).

When an IPv6 node places a call to another IPv6 node, it SHOULD follow the procedures in [6] and [7] for signalling UNI 3.0/3.1 SVCs [9] and negotiating MTU. The default IP MTU size on a LL is 9180 bytes as specified in [7].

Note that while the procedures in [7] still apply to IPv6 over ATM, IPv6 Path MTU Discovery [8] is used by nodes and routers rather than IPv4 MTU discovery. Additionally, while IPv6 nodes are not required to implement Path MTU Discovery, IPv6/ATM nodes SHOULD implement it. Also, since IPv6 nodes will negotiate an appropriate MTU for each VC, Path MTU should never be triggered since neither node should ever receive a Packet Too Big message to trigger Path MTU Discovery. When nodes are communicating via one or more routers Path MTU Discovery will be used just as it is for legacy networks.

5 Interface Tokens

For both PVC and SVC modes of operation, one of the following methods SHALL be used to generate Interface Tokens as required by section 5.1 of [1].

5.1 Interface Tokens Based on ESI values

When the underlying ATM interface is identified by an ATM End System Address (AESA, formerly known as an NSAPA), the interface token MAY be formed from the ESI and SEL values in the AESA as follows:

[0x00][ESI][SEL]

[0x00] is a one octet field which is always set to 0.

Note that the bit corresponding to the EUI-64 Global/Local bit [5] is always reset indicating that this address is not a globally unique IPv6 interface token.

[ESI] is a six octet field.

This field always contains the six octet ESI value for the AESA used to address the specific instance of the IPv6/ATM interface.

[SEL] is a one octet field.

This field always contains the SEL value from the AESA used to address the specific instance of the IPv6/ATM interface.

5.2 Interface Tokens Based on 48 Bit MAC Values

Where the underlying ATM NIC driver has access to a set of one or more 48 bit MAC values unique to the ATM NIC (e.g. MAC addresses configured into the NIC's ROM), the IPv6/ATM interface MAY use one of these values to create a unique interface token as described in [10].

5.3 Interface Tokens Based on EUI-64 Values

Where the underlying ATM NIC driver has access to a set of one or more 64 bit EUI-64 values unique to the ATM NIC (e.g. EUI-64 addresses configured into the NIC's ROM), the IPv6/ATM interface SHOULD use one of these values to create a unique interface token. after inverting the Global/Local identifier bit [10]. (Any relationship between these values and the ESI(s) registered with the local ATM switch by the ATM driver are outside the scope of this document.)

When EUI-64 values are used for IPv6 interface tokens the only modification allowed to the octet string read from the NIC is inversion of the Global/Local identifier bit.

5.4 Interface Tokens Based on Native E.164 Addresses

When an interface uses Native E.164 addresses then the E.164 values MAY be used to generate an interface token as follows:

[D14][D13D12][D11D10][D9D8][D9D6][D5D4][D3D2][D1D0]

[D14] A single octet containing the semi-octet representing the most significant E.164 digit shifted left four bits to the most significant four bits of the octet. The lower four bits MUST be set to 0. Note that the EUI-64 Global/Local indicator is set to 0 indicating that this is not a globally unique IPv6 interface token.

[D13D12] A single octet containing the semi-octet representing the second most significant E.164 digit [D13] shifted left four places to the most significant bits of the octet, and the third most significant semi-octet in the four least significant bits of the octet.

[D11D10] - [D1D0] Octets each containing two E.164 digits, one in the most significant four bits, and one in the least significant four bits as indicated.

5.5 Nodes Without Unique Identifiers

If no MAC, EUI-64, AESA, or E.164 value is available for generating an interface token, then the interface token SHALL be generated as described in Appendix A of [10].

5.6 Multiple Logical Links on a Single Interface

A logical ATM interface might be associated with a different SEL field of a common AESA prefix, or a set of entirely separate ESIs might have been registered with the local ATM switch to create a range of unique AESAs.

The minimum information required to uniquely identify each logical ATM interface is (within the context of the local switch port) their ESI+SEL combination.

For the vhost case described in section 5.1.2 of [1], vhost SHALL select a different interface token from the range of 64 bit values available to the ATM NIC (as described in 4.1). Each vhost SHALL implement IPv6/ATM interfaces in such a way that no two or more vhosts end up advertising the same interface token onto the same LL. (Conformance with this requirement may be achieved by choosing different SEL values, ESI values, or both.)

6. Conclusion and Open Issues.

This document is an ATM-specific companion document to the ION working group's, "IPv6 over Non Broadcast Multiple Access (NBMA) networks" specification [1]. It specifies codepoints for the administratively configured PVC, and dynamically established SVC, modes of operation.

There are no major open issues. Comments to the ION mailing list are solicited (ion@nexen.com).

7. Security Considerations

While this proposal does not introduce any new security mechanisms all current IPv6 security mechanisms will work without modification for ATM. This includes both authentication and encryption for both Neighbor Discovery protocols as well as the exchange of IPv6 data packets.

Acknowledgments

The original IPv6/ATM work by G. Armitage occurred while employed at Bellcore. Elements of section 4 were borrowed from Matt Crawford's memo on IPv6 over Ethernet.

The authors would like to thank Kazuhiko Yamamoto, Kenjiro Cho, Yoshinobu Inoue, Hiroshi Esaki, Yoshifumi Atarashi, and Atsushi Hagiwara for their contributions based on actual PVC implementations.

Authors' Addresses

Grenville Armitage
Bell Laboratories, Lucent Technologies
101 Crawfords Corner Road
Holmdel, NJ 07733
USA

E-Mail: gja@lucent.com

Peter Schulter
BrightTiger Technologies
125 Nagog Park
Acton, MA 01720

E-Mail: paschulter@acm.org

Markus Jork
European Applied Research Center
Digital Equipment GmbH
CEC Karlsruhe
Vincenz-Priessnitz-Str. 1
D-76131 Karlsruhe
Germany

E-Mail: jork@kar.dec.com

References

- [1] Armitage, G., Schulter, P., Jork, M. and G. Harter, "IPv6 over Non-Broadcast Multiple Access (NBMA) networks", RFC 2491, January 1999.
- [2] Heinanen, J., "Multiprotocol Encapsulation over ATM Adaption Layer 5", RFC 1483, July 1993.
- [3] Armitage, G., "Support for Multicast over UNI 3.1 based ATM Networks", RFC 2022, November 1996.
- [4] Luciani, J., Katz, D., Piscitello, D., Cole, B. and N. Doraswamy, "NBMA Next Hop Resolution Protocol (NHRP)", RFC 2332, April 1998.
- [5] "64-Bit Global Identifier Format Tutorial",
<http://standards.ieee.org/db/oui/tutorials/EUI64.html>.
- [6] Perez, M., Liaw, F., Mankin, A., Hoffman, E., Grossman, D. and A. Malis, "ATM Signalling Support for IP over ATM", RFC 1755, February 1995.
- [7] Atkinson, R., "Default IP MTU for use over ATM AAL5", RFC 1626, May 1994.
- [8] McCann, J., Deering, S. and J. Mogul, "Path MTU Discovery for IP version 6", RFC 1981, August 1996.
- [9] ATM Forum, "ATM User Network Interface (UNI) Specification Version 3.1", ISBN 0-13-393828-X, Prentice Hall, Englewood Cliffs, NJ, June 1995.
- [10] Hinden, B. and S. Deering, "IP Version 6 Addressing Architecture", RFC 2373, July 1998.
- [11] Narten, T., Nordmark, E. and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.

Full Copyright Statement

Copyright (C) The Internet Society (1999). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.