

Network Working Group  
Request for Comments: 2470  
Category: Standards Track

M. Crawford  
Fermilab  
T. Narten  
IBM  
S. Thomas  
TransNexus  
December 1998

## Transmission of IPv6 Packets over Token Ring 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 (1998). All Rights Reserved.

### 1. Introduction

This memo specifies the MTU and frame format for transmission of IPv6 packets on Token Ring networks. It also specifies the method of forming IPv6 link-local addresses on Token Ring networks and the content of the Source/Target Link-layer Address option used the Router Solicitation, Router Advertisement, Redirect, Neighbor Solicitation and Neighbor Advertisement messages when those messages are transmitted on a Token Ring network.

Implementors should be careful to note that Token Ring adaptors assume addresses are in non-canonical rather than canonical format, requiring that special care be taken to insure that addresses are processed correctly. See [CANON] for more details.

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 [KEYWORD].

### 2. Maximum Transmission Unit

IEEE 802.5 networks have a maximum frame size based on the maximum time a node may hold the token. This time depends on many factors including the data signaling rate and the number of nodes on the ring. Because the maximum frame size varies, implementations must

rely on manual configuration or router advertisements [DISC] to determine actual MTU sizes. Common default values include approximately 2000, 4000, and 8000 octets.

In the absence of any other information, an implementation should use a default MTU of 1500 octets. This size offers compatibility with all common 802.5 defaults, as well as with Ethernet LANs in an environment using transparent bridging.

In an environment using source route bridging, the process of discovering the MAC-level path to a neighbor can yield the MTU for the path to that neighbor. The information is contained in the largest frame (LF) subfield of the routing information field. This field limits the size of the information field of frames to that destination, and that information field includes both the LLC [LLC] header and the IPv6 datagram. Since, for IPv6, the LLC header is always 8 octets in length, the IPv6 MTU can be found by subtracting 8 from the maximum frame size defined by the LF subfield. If an implementation uses this information to determine MTU sizes, it must maintain separate MTU values for each neighbor.

A detailed list of the LF values and the resulting maximum frame size can be found in [BRIDGE]. To illustrate the calculation of IPv6 MTU, the following table lists several common values. Note that some of the 802.1D LF values would result in an IP MTU less than 1280 bytes. This size is less than the IPv6 minimum, and communication across paths with those MTUs is generally not possible using IPv6.

LF (base)	LF (extension)	MAC MTU	IP MTU
001	000	1470	1462
010	000	2052	2044
011	000	4399	4391
100	000	8130	8122
101	000	11407	11399
110	000	17749	17741
111	000	41600	41592

When presented with conflicting MTU values from several sources, an implementation should choose from those sources according to the following priorities:

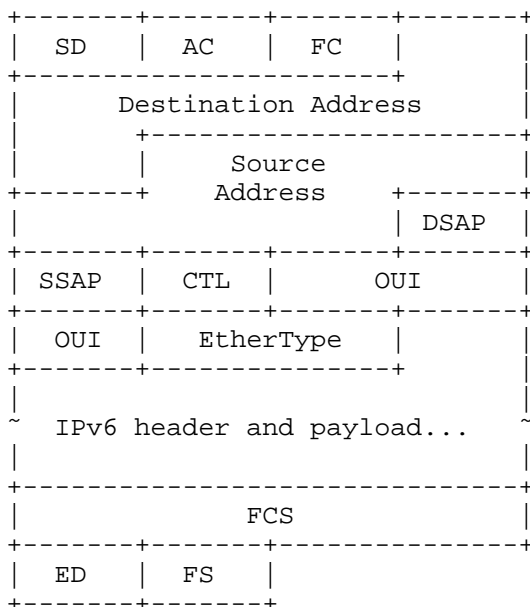
1. Largest Frame values from source route bridging (only for specific, unicast destinations), but only if not greater than value from any router advertisements
2. Router advertisements, but only if not greater than any manual configuration (including DHCP)

3. Manual configuration (including DHCP)

4. Default of 1500

### 3. Frame Format

IPv6 packets are transmitted in LLC/SNAP frames. The data field contains the IPv6 header and payload. The following figure shows a complete 802.5 frame containing an IPv6 datagram.



#### Token Ring Header Fields

SD: Starting Delimiter

AC: Access Control

FC: Frame Control

Destination Address: 48-bit IEEE address of destination station

Source Address: 48-bit IEEE address of source station

DSAP: Destination Service Access Point (for LLC/SNAP format, shall always contain the value 0xAA)

SSAP: Source Service Access Point (for LLC/SNAP format, shall always contain the value 0xAA)

CTL: Control Field (for Unnumbered Information, shall always contain the value 0x03)

OUI: Organizationally Unique Identifier (for EtherType encoding, shall always contain the value 0x000000)

EtherType: Protocol type of encapsulated payload (for IPv6, shall always contain the value 0x86DD)

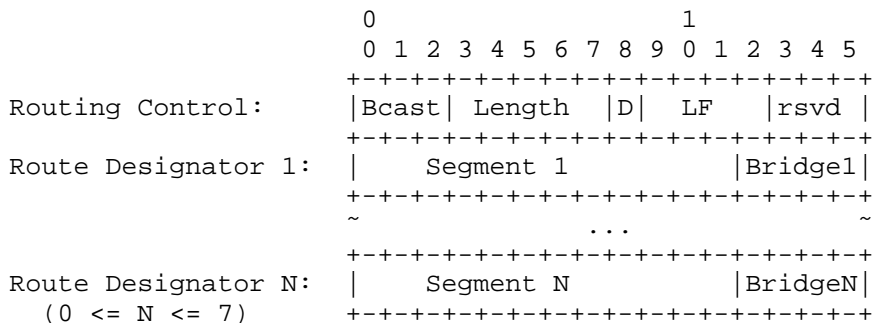
FCS: Frame Check Sequence

ED: Ending Delimiter

FS: Frame Status

In the presence of source route bridges, a routing information field (RIF) may appear immediately after the source address. A RIF is present in frames when the most significant bit of the source address is set to one. (This is the bit whose position corresponds to that of the Individual/Group bit in the Destination Address.)

The RIF is a variable-length field that (when present) contains a two-octet Routing Control (RC) header, followed by zero or more two-octet Route Designator fields:



Route Designator Fields:

- Bcast: Broadcast Indicator, Defined values:
- 10x: All Routes Explorer
  - 11x: Spanning Tree Explorer
  - 0xx: Specifically Routed Frame

Length: Total length of RIF field in octets

D: Direction of source route. A value of 0 means that the left-to-right sequence of Route Designators provides the path from the sender to recipient. A value of 1 indicates the sequence goes from recipient to sender.

LF: Largest Frame

rsvd: Reserved

On transmission, the Route Designator fields give the sequence of (bridge, LAN segment) numbers the packet is to traverse. It is the responsibility of the sender to provide this sequence for Specifically Routed Frames, i.e., unicast IP datagrams.

#### 4. Stateless Autoconfiguration

The Interface Identifier [AARCH] for a Token Ring interface is based on the EUI-64 identifier [EUI64] derived from the interface's built-in 48-bit IEEE 802 address. The OUI of the Token Ring address (the first three octets) becomes the company\_id of the EUI-64 (the first three octets). The fourth and fifth octets of the EUI are set to the fixed value FFFE hexadecimal. The last three octets of the Token Ring address become the last three octets of the EUI-64.

The Interface Identifier is then formed from the EUI-64 by complementing the "Universal/Local" (U/L) bit, which is the next-to-lowest order bit of the first octet of the EUI-64. Complementing this bit will generally change a 0 value to a 1, since an interface's built-in address is expected to be from a universally administered address space and hence have a globally unique value. A universally administered IEEE 802 address or an EUI-64 is signified by a 0 in the U/L bit position, while a globally unique IPv6 Interface Identifier is signified by a 1 in the corresponding position. For further discussion on this point, see [AARCH].

For example, the Interface Identifier for a Token Ring interface whose built-in address is, in hexadecimal and in canonical bit order,

34-56-78-9A-BC-DE

would be

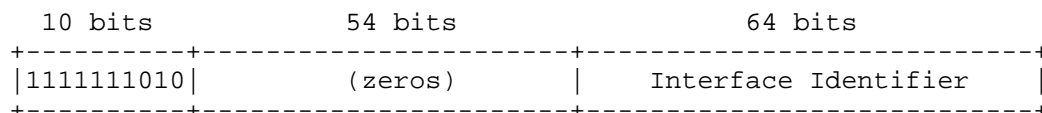
36-56-78-FF-FE-9A-BC-DE.

A different MAC address set manually or by software should not be used to derive the Interface Identifier. If such a MAC address must be used, its global uniqueness property should be reflected in the value of the U/L bit.

An IPv6 address prefix used for stateless autoconfiguration of a Token Ring interface must have a length of 64 bits.

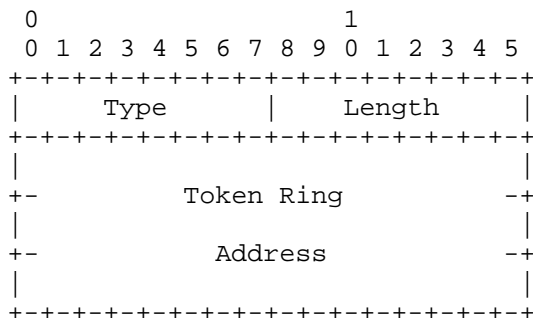
5. Link-Local Address

The IPv6 link-local address [AARCH] for a Token Ring interface is formed by appending the Interface Identifier, as defined above, to the prefix FE80::/64.



6. Address Mapping -- Unicast

The procedure for mapping unicast IPv6 addresses into Token Ring link-layer addresses is described in [DISC]. The Source/Target Link-layer Address option has the following form when the link layer is Token Ring.



Option fields:

Type:       1 for Source Link-layer address.  
          2 for Target Link-layer address.

Length:    1 (in units of 8 octets).

Token Ring Address: The 48 bit Token Ring IEEE 802 address, in canonical bit order. This is the address the interface currently responds to, and may be different from the built-in address used to derive the Interface Identifier.

When source routing bridges are used, the source route for the path to a destination can be extracted from the RIF field of received Neighbor Advertisement messages. Note that the RIF field of received packets can be reversed into a source route suitable for transmitting return traffic by toggling the value of the 'D' bit and insuring that the Bcast field is set to indicate a Specifically Routed Frame.

## 7. Address Mapping -- Multicast

All IPv6 packets with multicast destination addresses are transmitted to Token Ring functional addresses. The following table shows the specific mapping between the IPv6 addresses and Token Ring functional addresses (in canonical form). Note that protocols other than IPv6 may use these same functional addresses, so all Token Ring frames destined to these functional addresses are not guaranteed to be IPv6 datagrams.

MAC Addr (canonical)	IPv6 Multicast Addresses
03-00-80-00-00-00	All-Nodes (FF01::1 and FF02::1) and solicited node (FF02:0:0:0:0:1:FFXX:XXXX) addresses
03-00-40-00-00-00	All-Routers addresses (FF0X::2)
03-00-00-80-00-00	any other multicast address with three least significant bits = 000
03-00-00-40-00-00	any other multicast address with three least significant bits = 001
03-00-00-20-00-00	any other multicast address with three least significant bits = 010
03-00-00-10-00-00	any other multicast address with three least significant bits = 011
03-00-00-08-00-00	any other multicast address with three least significant bits = 100

- 03-00-00-04-00-00 any other multicast address with three  
least significant bits = 101
- 03-00-00-02-00-00 any other multicast address with three  
least significant bits = 110
- 03-00-00-01-00-00 any other multicast address with three  
least significant bits = 111

In a bridged token ring network, all multicast packets SHOULD be sent with a RIF header specifying the use of the Spanning Tree Explorer.

Note: it is believed that some (very) old bridge implementations do not properly support the Spanning Tree Explorer mechanism. In such environments, multicast traffic sent through bridges must use a RIF with the All Routes Explorer. Consequently, an implementation MAY wish to allow the sending of IP multicast traffic using an All Routes Explorer. However, such an ability must be configurable by a system administrator and the default setting of the switch MUST be to use the Spanning Tree Explorer.

## 8. Security Considerations

Token Ring, like most broadcast LAN technologies, has inherent security vulnerabilities. For example, any sender can claim the identity of another and forge traffic. It is the responsibility of higher layers to take appropriate steps in those environments where such vulnerabilities are unacceptable.

## 9. Acknowledgments

Several members of the IEEE 802.5 Working Group contributed their knowledge and experience to the drafting of this specification, including Jim, Andrew Draper, George Lin, John Messenger, Kirk Preiss, and Trevor Warwick. The author would also like to thank many members of the IPng working group for their advice and suggestions, including Ran Atkinson, Scott Bradner, Steve Deering, Francis Dupont, Robert Elz, and Matt Thomas. A special thanks is due Steve Wise, who gave the most relevant advice of all by actually trying to implement this specification while it was in progress.



## 10. References

- [802.5] 8802-5 : 1995 (ISO/IEC) [ANSI/IEEE 802.5, 1995 Edition] Information technology--Telecommunications and information exchange between systems--Local and metropolitan area networks--Specific requirements-- Part 5: Token ring access method and physical layer specification.
- [AARCH] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 2373, July 1998.
- [ACONF] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.
- [BRIDGE] 10038: 1993 (ISO/IEC) [ANSI/IEEE Std 802.1D, 1993 Edition] Information technology--Telecommunications and information exchange between systems--Local area networks--Media access control (MAC) bridges.
- [CANON] Narten, T. and C. Burton, "A Caution on Canonical Bit Order Of Link-Layer Addresses", RFC 2469, December 1998.
- [CONF] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 1971, August 1996.
- [DISC] Narten, T., Nordmark, E. and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.
- [EUI64] "64-Bit Global Identifier Format Tutorial", <http://standards.ieee.org/db/oui/tutorials/EUI64.html>.
- [IPV6] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [KEYWORD] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," BCP 14, RFC 2119, March 1997.
- [LLC] 8802-2 : 1994 (ISO/IEC) [ANSI/IEEE 802.2, 1994 Edition] Information technology--Telecommunications and information exchange between systems--Local and Metropolitan area networks--Specific requirements-- Part 2: Logical link control.

## 11. Authors' Addresses

Matt Crawford  
Fermilab MS 368  
PO Box 500  
Batavia, IL 60510 USA

Phone: +1 630 840 3461  
EMail: crawdad@fnal.gov

Thomas Narten  
IBM Corporation  
P.O. Box 12195  
Research Triangle Park, NC 27709-2195 USA

Phone: +1 919 254 7798  
EMail: narten@raleigh.ibm.com

Stephen Thomas  
TransNexus  
430 Tenth Street NW Suite N204  
Atlanta, GA 30318 USA

Phone: +1 404 872 4745  
EMail: stephen.thomas@transnexus.com

## Full Copyright Statement

Copyright (C) The Internet Society (1998). 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.