

A Caution On The Canonical Ordering Of Link-Layer Addresses

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1998). All Rights Reserved.

Abstract

Protocols such as ARP and Neighbor Discovery have data fields that contain link-layer addresses. In order to interoperate properly, a sender setting such a field must insure that the receiver extracts those bits and interprets them correctly. In most cases, such fields must be in "canonical form". Unfortunately, not all LAN adaptors are consistent in their use of canonical form, and implementations may need to explicitly bit swap individual bytes in order to obtain the correct format. This document provides information to implementors to help them avoid the pitfall of using non-canonical forms when canonical forms are required.

Table of Contents

1. Introduction.....	2
2. Canonical Form.....	2
3. Implementors Beware: Potential Trouble Spots.....	3
3.1. Neighbor Discovery in IPv6.....	3
3.2. IPv4 and ARP.....	3
4. Security Considerations.....	3
5. References.....	4
6. Authors' Addresses.....	4
7. Full Copyright Statement.....	5

1. Introduction

Protocols such as ARP [ARP] and ND [DISCOVERY] have data fields that contain link-layer addresses. In order to interoperate properly, a sender setting such a field must insure that the receiver extracts those bits and interprets them correctly. In most cases, such fields must be in "canonical form". Unfortunately, not all LAN adapters are consistent in their use of canonical form, and implementations may need to explicitly bit swap individual bytes in order to obtain the correct format.

2. Canonical Form

Canonical form (also known as "LSB format" and "Ethernet format") is the name given to the format of a LAN adapter address as it should be presented to the user according to the 802 LAN standard. It is best defined as how the bit order of an adapter address on the LAN media maps to the bit order of an adapter address in memory: The first bit of each byte that appears on the LAN maps to the least significant (i.e., right-most) bit of each byte in memory (the figure below illustrates this). This puts the group address indicator (i.e., the bit that defines whether an address is unicast or multicast) in the least significant bit of the first byte. Ethernet and 802.3 hardware behave consistently with this definition.

Unfortunately, Token Ring (and some FDDI) hardware does not behave consistently with this definition; it maps the first bit of each byte of the adapter address to the most significant (i.e., left-most) bit of each byte in memory, which puts the group address indicator in the most significant bit of the first byte. This mapping is variously called "MSB format", "IBM format", "Token-Ring format", and "non-canonical form". The figure below illustrates the difference between canonical and non-canonical form using the canonical form address 12-34-56-78-9A-BC as an example:

	12	34	56	78	9A	BC
In memory,						
canonical:	00010010	00110100	01010110	01111000	10011010	10111100
	1st bit appearing on LAN (group address indicator)					
On LAN:	01001000	00101100	01101010	00011110	01011001	00111101
In memory,						
MSB format:	01001000	00101100	01101010	00011110	01011001	00111101
	48	2C	6A	1E	59	3D

The implication of this inconsistency is that addresses extracted from adaptors, assigned to adaptors, or extracted from link-layer packet headers obtained from adaptors may need to be bit-swapped to put them into canonical form. Likewise, addresses in canonical form that are handed to adaptors (e.g., to set an address, to specify a destination address in a link-layer header, etc.) may need to be bit-swapped in order for the adaptor to process the request as expected.

3. Implementors Beware: Potential Trouble Spots

3.1. Neighbor Discovery in IPv6

All of the IPv6 over specific link layers documents specify that link-layer addresses must be transmitted in canonical order [IPv6-ETHER, IPv6-FDDI, IPv6-TOKEN]. As far as the authors can tell, all Ethernet LAN adaptors use canonical order and no special processing by implementations is needed. In contrast, some FDDI and all Token Ring adaptors appear to use non-canonical format. Implementors must insure that any addresses that appear in link-layer address options of Neighbor Discovery [DISCOVERY] messages are sent in canonical order and that any link-layer addresses extracted from ND packets are interpreted correctly on the local machine and its adaptors.

3.2. IPv4 and ARP

Ethernet addresses that appear in ARP packets are in canonical order. In contrast, when running ARP over Token Ring, the de facto practice is to transmit addresses in non-canonical order. Because all Token Ring adaptors assume non-canonical ordering, no interoperability problems result between communicating nodes attached to the same Token Ring.

In some environments, however, Token Rings and Ethernets are connected via a bridge. When a node on the Token Ring attempts to communicate with a node on the Ethernet, communication would normally fail, since the Ethernet will misinterpret the Token Ring address (and vice versa). To get around this problem, bridges that forward packets between dissimilar network types perform bit swaps of the addresses in the address fields of ARP packets that are forwarded from a network of one type to one of the other.

4. Security Considerations

There are no known security issues raised by this document.

5. References

- [ARP] Plummer, D., "An Ethernet Address Resolution Protocol", STD 37, RFC 826, November 1982.
- [DISCOVERY] Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.
- [IPv6-ETHER] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, December 1998.
- [IPv6-FDDI] Crawford, M., "Transmission of IPv6 Packets over FDDI Networks", RFC 2467, December 1998.
- [IPv6-TOKEN] Crawford, M., Narten, T. and S. Thomas, "Transmission of IPv6 Packets over Token Ring Networks", RFC 2470, December 1998.

6. Authors' Addresses

Thomas Narten
IBM Corporation
3039 Cornwallis Ave.
PO Box 12195
Research Triangle Park, NC 27709-2195

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

Charles F. Burton, III
IBM Corporation
3039 Cornwallis Ave.
PO Box 12195
Research Triangle Park, NC 27709-2195

Phone: 919-254-4355
EMail: burton@rtp.vnet.ibm.com

7. 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.

