

Handwritten: Notebook

IEN - 212

Table of Contents

I Introduction..... 1

II IP to Local Network Address Mapping..... 2

III Broadcast Addressing..... 3

IV Appendix A - Revised Information Reply Message..... 4

V References..... 11

Robert Gurwitz
Robert Hinden

Bolt Beranek and Newman Inc.
10 Moulton St.
Cambridge, Massachusetts 02238

September 1982

prepared for
Defense Advanced Research Projects Agency
Information Processing Techniques Office
1400 Wilson Boulevard
Arlington, Virginia 22209

Table of Contents

1 Introduction..... 1
2 IP to Local Network Address Mapping..... 2
3 Broadcast Addressing..... 7
Appendix A - Revised Information Reply Message..... 9
References..... 11

Bob Barnek and Norman Inc.
10 Houlton St.
Cambridge, Massachusetts 02138

September 1982

prepared for
Defense Advanced Research Projects Agency
Information Processing Techniques Office
1400 Wilson Boulevard
Arlington, Virginia 22209

1 Introduction

This note discusses several issues in Internet Protocol (IP) addressing in local area networks, such as Ethernet [1]. Such networks typically have addressing characteristics and features that are different from earlier network implementations, such as the ARPANET. Specifically, we describe approaches for handling IP to local network address mapping and broadcast. The methods described below are meant to integrate these features of local networks into the DARPA internet within the spirit of IP, in a general manner that is not tied to a particular network implementation. The goal here is to minimize the need for changes to IP, by adopting conventions and methods that use existing mechanisms.

Of course, the approaches described here are not unique. Other methods have been described which are similar in spirit to that presented here. The purpose of this note is to stimulate discussion of these issues, with the goal of arriving at a consensus on how to resolve these and possibly other issues relating to the use of local area networks in the internet.

2 IP to Local Network Address Mapping

Ethernet differs from other networks in the internet in that it uses local addresses which are larger (48 bits) than IP addresses (32 bits). With other networks having smaller local address formats, the mapping from local network address to IP address is simple: embed the entire local address in the appropriate class of IP address. For example, in the ARPANET which has 24-bit addresses, the mapping to a class A IP address, which has 24-bits of local host address, is direct:

Host A on IMP B -> 10.A.B(hi).B(lo)

where 10 is the class A network number for the ARPANET, the host portion of the local address is an 8-bit value, and the IMP portion is a 16-bit value whose high order 8 bits are usually zero. Note that mapping a 24-bit ARPANET address to a non-class A address, would require some other approach, like truncation (say, dropping the high order 8 bits of the IMP part) or table lookup.

A similar, but more severe, problem occurs with 48-bit Ethernet addresses. In this case, there may be no simple direct mapping (like truncation) to the 32-bit IP address, since the

address of an individual Ethernet controller is a unique 48-bit value assigned by the manufacturer that could conceivably conflict with other addresses on the local network if a subset of the 48-bits is used to form the local host part of the IP address.

To solve this potential conflict, we choose a mapping that depends on table lookup. Each host on the local network is assigned an IP address with a unique local host part for that network. The IP module in that host "knows" its own IP address and uses it to send and receive IP datagrams. Other hosts on the local network (including gateways), maintain tables which map the IP addresses of the other hosts to their corresponding local addresses.

This table may be "hardwired" into the IP modules or initialized when the host comes up. A better implementation would be to maintain the table as a cache, whose entries are updated dynamically while the host is active. To do this, we can make use of a feature supported by many local area networks, "broadcast addressing". This feature allows messages to be received by all hosts on the local network, by sending them to a specially designated broadcast address (in Ethernet the broadcast

address is "all ones"). To map a given IP address to its corresponding local address, the IP address is used to look up the local address in the mapping table. If the mapping is found, the local address is used in the local network header. Otherwise, if no mapping is found in the table:

- 1) Send the datagram which was not in the mapping table to the local broadcast address (see figure 1). All hosts on the network including the correct one (the one specified by destination IP address in the datagram) will receive the datagram.




Figure 1 . Broadcast Datagram

- 2) The host on the network that recognizes its IP address in the datagram and notices that the datagram was received on the local network broadcast address sends an ICMP information reply message to the host that broadcast the request, with its local address filled in and its local network address in the data portion of the message (see figure 2). ** It also

←←←←←←←←←←←←

** Currently, there is no data area defined in the ICMP Information Reply message. We propose to add one for the purpose described in the document. Appendix A contains the format of ICMP Information Reply message with the data area.

accepts this broadcast datagram since it contains its IP address as the destination.




Figure 2 . Mapping Reply Format

- 3) When an ICMP information reply is received, an entry is made in the mapping table, using the local address found in the data portion of the message.

Finally, there must be some way of invalidating table entries, to handle cases like a host on the local network changing addresses:

Whenever a host on the local net comes up, it broadcasts an ICMP information reply with its source IP and local network addresses (figure 3). Other hosts on the network can use this information to update their tables.

This message contains the notion of an "IP broadcast address". This is simply a means of mapping an IP address to the local network broadcast address. It has the effect that any IP module on a particular network should accept datagrams with the broadcast address as the IP destination. The IP broadcast address




Figure 3 . Mapping Initialization Format

is discussed further below. Note that the use of this message also lowers the chances of hosts not having the mapping in their table, and reduces the number of broadcast mapping requests.

This approach has several advantages. 1) It is general enough to work on all local networks that support broadcast addressing. 2) It fits into IP without any major changes such as increasing the size of the addresses. 3) Hosts which do not implement this approach will still get all of their datagrams (at the cost of high overhead for all others on the net). 4) It allows hosts on the local network to assign their own local portion of their IP addresses and the mapping between their IP address and their local network address.


On the negative side, it can potentially cause adverse network and host performance when many hosts on the net are coming up, because of the use of broadcast packets. Also, performance can be degraded by hosts which do not implement the scheme fully, causing more broadcast packets to be sent. However, we expect the use of broadcast mapping requests to be a low frequency occurrence.

3 Broadcast Addressing

The scheme described above introduced the concept of IP broadcast address. We believe this is a generally useful notion, since it allows IP hosts to take advantage of an additional addressing capability supported by some local area networks. We propose to define the IP broadcast address to be the IP address in each class with all its local host part bits on. (E.G., A.255.255.255 for class A, A.B.255.255 for class B, and A.B.C.255 for class C.) In each case, the address would map to the local network broadcast address if broadcast addressing was supported. On messages coming from other networks, the mapping would be done in the gateway. If a network did not support broadcast addressing, an ICMP destination unreachable message would be

returned. The only "cost" of this mechanism is that it reserves one IP address from each class.

Appendix A - Revised Information Reply Message



IP Fields:

Addresses

The address of the source in a information request message will be the destination of the information reply message. To form a information reply message, the source and destination addresses are simply reversed, the type code changed to 16, and the checksum recomputed.

ICMP Fields:

Type

15 for information request message;
16 for information reply message.

Code

0 for normal information request/reply
1 for information reply with local network address in data portion.

Checksum

The checksum is the 16-bit ones's complement of the one's complement sum of the ICMP message starting with the ICMP Type. For computing the checksum, the checksum field should be zero. This checksum may be replaced in the future.

Identifier

If code = 0, an identifier to aid in matching request and replies, may be zero.

Sequence Number

If code = 0, a sequence number to aid in matching request and replies, may be zero.

Description

Code 0:

Code 0 may be received from a gateway or a host.

This message may be sent with the source network in the IP header source and destination address fields zero (which means "this" network). The replying IP module should send the reply with the addresses fully specified. This message is a way for a host to find out the number of the network it is on.

The identifier and sequence number may be used by the echo sender to aid in matching the replies with the requests. For example, the identifier might be used like a port in TCP or UDP to identify a session, and the sequence number might be incremented on each request sent. The destination returns these same values in the reply.

Code 1:

This message (Information Reply only) is used by a local area network host to inform another host what its local network address in relation to its IP address. The hosts local network address is sent in the data portion of the message.

References

- [1] "The Ethernet, A Local Area Network: Data Link Layer and Physical Layer Specifications", Version 1.0, Digital Equipment Corp., Intel, Xerox, Sept. 30, 1980.