

Network Working Group
Request for Comments: 3021
Category: Standards Track

A. Retana
R. White
Cisco Systems
V. Fuller
GTE Internetworking
D. McPherson
Amber Networks
December 2000

Using 31-Bit Prefixes on IPv4 Point-to-Point Links

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 (2000). All Rights Reserved.

Abstract

With ever-increasing pressure to conserve IP address space on the Internet, it makes sense to consider where relatively minor changes can be made to fielded practice to improve numbering efficiency. One such change, proposed by this document, is to halve the amount of address space assigned to point-to-point links (common throughout the Internet infrastructure) by allowing the use of 31-bit subnet masks in a very limited way.

1. Introduction and Motivation

The perceived problem of a lack of Internet addresses has driven a number of changes in address space usage and a number of different approaches to solving the problem:

- More stringent address space allocation guidelines, enforced by the IANA and the regional address assignment authorities [RFC2050].
- Use of Network Address Translators (NATs), where a small number of IANA-compliant addresses are shared by a larger pool of private, non-globally routed addresses topologically behind a NAT box [RFC1631].

- Deployment of a new Internet Protocol to increase the size of the address space. One such protocol, IPv6 [RFC2460], has been through the IETF process but has yet to see production deployment. Should it be, deployed, it will still face a many year transition period.

Prior to the availability of a larger address space, it seems prudent to consider opportunities for making more efficient use of the existing address space.

One such (small) opportunity is to change the way that point-to-point links are numbered. One option, which is used today on some parts of the Internet, is to simply not number point-to-point links between routers. While this practice may seem, at first, to handily resolve the problem, it causes a number of problems of its own, including the inability to consistently manage the unnumbered link or reach a router through it, difficulty in management and debugging of those links, and the lack of standardization [RFC1812].

In current practice, numbered Internet subnets do not use longer than a 30-bit subnet mask (in most cases), which requires four addresses per link - two host addresses, one all-zeros network, and one all-ones broadcast. This is unfortunate for point-to-point links, since they can only possibly have two identifying endpoints and don't support the notion of broadcast - any packet which is transmitted by one end of a link is always received by the other.

A third option is to use host addresses on both ends of a point-to-point link. This option provides the same address space savings as using a 31-bit subnet mask, but may only be used in links using PPP encapsulation [RFC1332]. The use of host addresses allows for the assignment of IP addresses belonging to different networks at each side of the link, causing link and network management not to be straight forward.

This document is based on the idea that conserving IP addresses on point-to-point links (using longer than a 30-bit subnet mask) while maintaining manageability and standard interaction is possible. Existing documentation [RFC950] has already hinted at the possible use of a 1-bit wide host-number field.

The savings in address space resulting from this change is easily seen--each point-to-point link in a large network would consume two addresses instead of four. In a network with 500 point-to-point links, for example, this practice would amount to a savings of 1000 addresses (the equivalent of four class C address spaces).

2. Considerations of 31-Bit Prefixes

This section discusses the possible effects, on Internet routing and operations, of using 31-bit prefixes on point-to-point links. The considerations made here are also reflected in Section 3.

For the length of this document, an IP address will be interpreted as:

<Network-number><Host-number>

where the <Host-number> represents the unmasked portion of the address and it SHOULD be at least 1 bit wide. The "-1" notation is used to mean that the field has all 1 bits. For purposes of this discussion, the routing system is considered capable of classless, or CIDR [RFC1519], routing.

2.1. Addressing

If a 31-bit subnet mask is assigned to a point-to-point link, it leaves the <Host-number> with only 1 bit. Consequently, only two possible addresses may result:

{<Network-number>, 0} and {<Network-number>, -1}

These addresses have historically been associated with network and broadcast addresses (see Section 2.2). In a point-to-point link with a 31-bit subnet mask, the two addresses above MUST be interpreted as host addresses.

2.2. Broadcast and Network Addresses

There are several historically recognized broadcast addresses [RFC1812] on IP segments:

(a) the directed broadcast

{<Network-number>, -1}

{<Network-number>, 0}

The network address itself {<Network-number>, 0} is an obsolete form of directed broadcast, but it may still be used by older hosts.

(b) the link local (or limited) broadcast

{-1, -1}

{0, 0}

The {0, 0} form of a limited broadcast is obsolete, but may still be present in a network.

Using a 31-bit prefix length leaves only two numbering possibilities (see Section 2.1), eliminating the use of a directed broadcast to the link (see Section 2.2.1). The limited broadcast **MUST** be used for all broadcast traffic on a point-to-point link with a 31-bit subnet mask assigned to it.

The <Network-number> is assigned by the network administrator as unique to the local routing domain. The decision as to whether a destination IP address should be a directed broadcast or not is made by the router directly connected to the destination segment. Current forwarding schemes and algorithms are not affected in remote routers.

The intent of this document is to discuss the applicability and operation of 31-bit prefixes on point-to-point links. The effects (if any) on other types of interfaces are not considered.

2.2.1. Directed Broadcast

When a device wants to reach all the hosts on a given (remote, rather than directly connected) subnet, it may set the packet's destination address to the link's subnet broadcast address. This operation is not possible for point-to-point links with a 31-bit prefix.

As discussed in Section 6, the loss of functionality of a directed broadcast may actually be seen as a beneficial side effect, as it slightly enhances the network's resistance to a certain class of DoS Attacks [RFC2644, SMURF].

2.3. Impact on Current Routing Protocols

Networks with 31-bit prefixes have no impact on current routing protocols. Most of the currently deployed routing protocols have been designed to provide classless routing. Furthermore, the communication between peers is done using multicast, limited broadcast or unicast addresses (all on the local network), none of which are affected with the use of 31-bit subnet masks.

3. Recommendations

The considerations presented in Section 2 affect other published work. This section details the updates made to other documents.

3.1. "Requirements for Internet Hosts -- Communication Layers" [RFC1122]

Section 3.2.1.3 (e) is replaced with:

(e) { <Network-number>, <Subnet-number>, -1 }

Directed broadcast to the specified subnet. It MUST NOT be used as a source address, except when the originator is one of the endpoints of a point-to-point link with a 31-bit mask.

A new section (numbered 3.2.1.3 (h)) is added:

(h) { <Network-number>, <Subnet-number>, 0 }

Subnetwork number. SHOULD NOT be used as a source address, except when the originator is one of the endpoints of a point-to-point link with a 31-bit mask. For other types of links, a packet with such a destination SHOULD be silently discarded. If these packets are not silently discarded, they MUST be treated as IP broadcasts [RFC1812].

3.2. "Assigned Numbers" [RFC1700]

Sub-section (e) of the "Special Addresses" section in the "Introduction" is replaced with:

(e) { <Network-number>, <Subnet-number>, -1 }

Directed broadcast to specified subnet. Can only be used as a destination address. However, in the case where the originator is one of the endpoints of a point-to-point link with a 31-bit mask, it can also be used as a source address.

3.3. "Requirements for IP Version 4 Routers" [RFC1812]

Section 4.2.2.11 (d) is replaced with:

(d) { <Network-prefix>, -1 }

Directed Broadcast - a broadcast directed to the specified network prefix. It MUST NOT be used as a source address, except when the originator is one of the endpoints of a point-

to-point link with a 31-bit mask. A router MAY originate Network Directed Broadcast packets. A router MAY have a configuration option to allow it to receive directed broadcast packets, however this option MUST be disabled by default, and thus the router MUST NOT receive Network Directed Broadcast packets unless specifically configured by the end user.

The text above includes the update made by [RFC2644].

A new section (numbered 4.2.2.11 (f)) is added:

(f) { <Network-number>, <Subnet-number>, 0 }

Subnetwork number. SHOULD NOT be used as a source address, except when the originator is one of the endpoints of a point-to-point link with a 31-bit mask. For other types of links, a packet with such a destination SHOULD be silently discarded. If these packets are not silently discarded, they MUST be treated as IP broadcasts.

Sections 4.2.3.1 (1), (2) and (4) are replaced with:

(1) MUST treat as IP broadcasts packets addressed to 255.255.255.255 or { <Network-prefix>, -1 }.

In a point-to-point link with a 31-bit mask, a packet addressed to { <Network-prefix>, -1 } corresponds to one of the endpoints of such link, it MUST be treated as directed to the router on which the address is applied.

(2) SHOULD silently discard on receipt (i.e., do not even deliver to applications in the router) any packet addressed to 0.0.0.0 or { <Network-prefix>, 0 }. If these packets are not silently discarded, they MUST be treated as IP broadcasts (see Section [5.3.5]). There MAY be a configuration option to allow receipt of these packets. This option SHOULD default to discarding them.

In a point-to-point link with a 31-bit mask, a packet addressed to { <Network-prefix>, 0 } corresponds to one of the endpoints of such link, it MUST be treated as directed to the router on which the address is applied.

(4) SHOULD NOT originate datagrams addressed to 0.0.0.0 or { <Network-prefix>, 0 }. There MAY be a configuration option to allow generation of these packets (instead of using the relevant 1s format broadcast). This option SHOULD default to not generating them.

In a point-to-point link with a 31-bit mask, the configuration of such a mask SHOULD allow for the generation of datagrams addressed to { <Network-prefix>, 0 }.

The following text is added to section 4.3.3.9:

The 255.255.255.255 IP broadcast address MUST be used for broadcast Address Mask Replies in point-to-point links with 31-bit subnet masks

4. Operational Experience

The recommendations presented in this document have been implemented by several router vendors in beta code. The implementation has been tested by at least three ISPs with positive results (i.e., no problems have been found). Among the routing protocols tested successfully are OSPF, IS-IS, BGP and EIGRP.

It is expected that the implementation will be officially released within the next few months and that other vendors will adopt it.

5. Deployment Considerations

The intent of this document is to discuss the applicability and operation of 31-bit prefixes on point-to-point links. The effects (if any) on other types of interfaces are not considered. Note that a point-to-point link in which only one end supports the use of 31-bit prefixes may not operate correctly.

6. Security Considerations

In the light of various denial of service (DoS) attacks on various networks within the Internet, security has become a major concern. The use of 31-bit subnet masks within the core of the Internet will reduce the number of physical links against which a DoS attack relying on packet replication through the use of directed broadcasts can be launched [RFC2644, SMURF].

Overall, implementation of this document recommendation will improve the Internet's resilience to these types of DoS attacks.

7. Acknowledgements

The authors of this document do not make any claims on the originality of the ideas described. Among other people, we would like to acknowledge Alex Zinin for his comments, and the many people who have tested 31 bit subnet masks in their labs and networks.

8. References

- [RFC950] Mogul, J. and J. Postel, "Internet Standard Subnetting Procedure", STD 5, RFC 950, August 1985.
- [RFC1122] Braden, R., "Requirements for Internet Hosts -- Communication Layers", STD 3, RFC 1122, October 1989.
- [RFC1332] McGregor, G., "The PPP Internet Protocol Control Protocol (IPCP)", RFC 1332, May 1992.
- [RFC1519] Fuller, V., Li, T., Yu, J. and K. Varadhan, "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy", RFC 1519, September 1993.
- [RFC1631] Egevang, K. and P. Francis, "The IP Network Address Translator (NAT)", RFC 1631, May 1994.
- [RFC1700] Reynolds, J. and J. Postel, "Assigned Numbers", STD 2, RFC 1700, October 1994.
- [RFC1812] Baker, F., "Requirements for IP Version 4 Routers", RFC 1812, June 1995.
- [RFC2050] Hubbard, K., Kusters, M., Conrad, D., Karrenberg, D. and J. Postel, "Internet Registry IP Allocation Guidelines", BCP 12, RFC 2050, November 1996.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [RFC2644] Senie, D., "Changing the Default for Directed Broadcasts in Routers", BCP 34, RFC 2644, August 1999.
- [SMURF] Huegen, C., "The Latest in Denial of Service Attacks: 'Smurfing': Description and Information to Minimize Effects", URL:
<http://users.quadrunner.com/chuegen/smurf.cgi>

9. Authors' Addresses

Alvaro Retana
Cisco Systems, Inc.
7025 Kit Creek Rd.
Research Triangle Park, NC 27709

EMail: aretana@cisco.com

Russ White
Cisco Systems, Inc.
7025 Kit Creek Rd.
Research Triangle Park, NC 27709

EMail: riw@cisco.com

Vince Fuller
GTE Internetworking
3801 E. Bayshore Rd.
Palo Alto, CA, 94303

EMail: vaf@valinor.barrnet.net

Danny McPherson
Amber Networks
2465 Augustine Drive
Santa Clara, CA 95054

EMail: danny@ambernetworks.com

Full Copyright Statement

Copyright (C) The Internet Society (2000). 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.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

