

Internet Engineering Task Force
Internet-Draft
Updates: 1122, 1812, 3021 (if approved)
Intended status: Standards Track
Expires: 25 December 2025

S.D. Schoen
J. Gilmore
D. T辰ht
M. Karels
IPv4 Unicast Extensions Project
23 June 2025

Unicast Use of the Lowest Address in an IPv4 Subnet
draft-schoen-intarea-unicast-lowest-address-08

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 increase the number of unicast addresses in each existing subnet, by redefining the use of the lowest-numbered (zeroth) host address in each IPv4 subnet as an ordinary unicast host identifier, instead of as a duplicate segment-directed broadcast address.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 25 December 2025.

Copyright Notice

Copyright (c) 2025 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights

and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

| | |
|--|----|
| 1. Introduction | 2 |
| 1.1. Requirements Language | 3 |
| 2. Background and Current Standards | 3 |
| 2.1. Assumptions About the Lowest Host Address by Remote Systems | 5 |
| 2.2. Multicast Addresses; Point-to-Point Links | 6 |
| 2.3. Current Limitations on Subnet-Directed Broadcast Addresses | 7 |
| 2.4. Comparison to IPv6 Behavior | 7 |
| 3. Change to Interpretation of the Lowest Address | 8 |
| 3.1. Link-Layer Interaction | 8 |
| 3.2. Recommendations | 9 |
| 3.2.1. "Requirements for Internet Hosts -- Communication Layers" [RFC1122] | 9 |
| 3.2.2. "Requirements for IP Version 4 Routers" [RFC1812] | 9 |
| 3.3. Example | 11 |
| 3.4. Compatibility and Interoperability | 12 |
| 4. IANA Considerations | 12 |
| 5. Security Considerations | 13 |
| 6. Acknowledgements | 13 |
| 7. References | 13 |
| 7.1. Normative References | 13 |
| 7.2. Informative References | 14 |
| Appendix A. Implementation Status | 15 |
| Authors' Addresses | 16 |

1. Introduction

This document provides history and rationale to change the interpretation of the lowest address in each IPv4 subnet from an alternative broadcast address (or "network address") to an ordinary assignable host address, and updates requirements for hosts and routers accordingly. The decision taken in 1989 to reserve two forms instead of one for local IPv4 segment broadcasts is no longer necessary because of the obsolescence and disappearance of the software that motivated it. Unreserving the lowest address provides an optional extra IPv4 host address in every subnet, Internet-wide, alleviating some of the pressure of IPv4 address exhaustion.

1.1. Requirements Language

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 [RFC2119].

2. Background and Current Standards

IPv4 has long supported several mechanisms for broadcasting communications to every station on a network. One form of broadcast in IPv4 is "segment-directed broadcast" in which a broadcast is addressed to every station on a particular network (identified by its network number). This mechanism is obsolescent today, no longer supported at all in many contexts and having been supplanted for adjacent host discovery years ago by link-local multicasting (see Section 2.3).

Yet current standards still reserve a huge number of addresses for it: not just one, but two, addresses per subnet, Internet-wide. That represents a significant burden for small subnets such as a /30, where, according to current rules, two of the four addresses on the subnet are reserved for broadcasting.

The standard broadcast address on a subnet is the address whose "host part" consists of all ones in binary. For example, in a 24-bit subnet that starts at 192.168.3.0, the address 192.168.3.255 is the standard broadcast address. [RFC1122][RFC0894]

In addition to the standard broadcast address, RFC 1122 (October 1989) acknowledged a then-current implementation behavior in "4.2BSD Unix and its derivatives, but not 4.3BSD", whereby those operating systems "use non-standard broadcast address forms, substituting 0 for -1". (Note that there was no standard for IP broadcast when 4.2BSD was released in August 1983, more than a year prior to [RFC0919]. The -1 form was first proposed in [IEN212] in 1982 but was not yet a standard.) According to RFC 1122 and its successors, Internet hosts are expected to "recognize and accept [...] these non-standard broadcast addresses as the destination address of an incoming datagram", and not otherwise use them to identify Internet hosts. RFCs 1812 [RFC1812] (sections 4.2.3.1 and 5.3.5), and 3021 [RFC3021] (sections 2.2, 3.1, and 3.3) reiterate and further expand on this requirement.

This non-standard broadcast address is the address whose "host part" consists of all zeroes. (The quantity of zeroes depends, in present-day terminology, on the applicable subnet mask.) This address is the lowest expressible address within any particular numbered network. Following computer science tradition, it may also be referred to as

the "zeroth address" of its respective subnet. Many sources also refer to it as the "network address" of the subnet, because the subnet as a whole (pre-CIDR) was commonly referred to by its lowest address, while under CIDR the subnet as a whole is commonly referred to by its lowest address plus a prefix length.

The address triple syntax used in RFC 1122 looks unusual to modern eyes. These triples included the "{ network number, subnet number, host number }". The notation also used two's complement binary notation in referring to a host number "-1" as containing all binary 1-bits. After the widespread adoption of CIDR [RFC4632], network numbers no longer have an "address class" definition based on their high-order bits, and there is no distinction between a network number and a subnet number (except locally at the router where individual subnets are being routed). Instead, following RFC 4632, IPv4 network addresses are denoted with a dotted-decimal format containing one to four positive 8-bit integers, a slash, and a whole count of the bits in the network number portion, the so-called CIDR notation, reportedly devised by Phil Karn. So for example 192.1.2.0/28 has a 28-bit network number and a 4-bit host number (32 address bits total, minus those 28 bits). All of the bits of that particular host number are zero (because the whole fourth dotted number is 0), and thus the interpretation of this address would be affected by this document. We will use both notations as convenient. Where RFC 1122 and its successors use the terms "0 form" and "-1 form", we may refer respectively to "all-zeroes" and "all-ones" host numbers, since every bit in the binary representation of these two host numbers has the value 0 or 1, respectively.

4.2BSD, the operating system to whose behavior RFC 1122 required deference, was the first BSD operating system full release to include TCP/IP support; it was released in 1983. Its successor, 4.3BSD, was released in 1986, which is why RFC 1122 could already confirm that the broadcast behavior had been changed. See [BSDHIST]. RFC 1812 calls the old behavior "obsolete" in 1995, and RFC 3021 reiterates that it is "obsolete" in 2000, although both express the idea that the lowest address must continue to be reserved for historical reasons.

All subsequent operating systems used on the Internet implement the standard all-ones form of the broadcast address and use it by default. Continuing to reserve the non-standard all-zeroes form wastes one IPv4 address per subnet. No known operating system generates IP broadcasts in this format by default today, and documentation consistently encourages network administrators and software developers to use the standard form. The IPv4 protocol does not benefit from having two different broadcast addresses with the same functionality in every subnet, and the non-standard form has always been reserved primarily for backwards compatibility with systems that have not existed for decades on the Internet.

As IPv4 addresses were not perceived as particularly scarce through the 1980s, the prospect of wasting tens of millions of otherwise assignable addresses in order to achieve backwards compatibility with a particular operating system appeared reasonable. Today, those addresses are clearly valuable and could be put to good use as identifiers of Internet hosts in a time of IPv4 numbering resource exhaustion.

2.1. Assumptions About the Lowest Host Address by Remote Systems

In general, under CIDR [RFC4632], only hosts and routers on a network segment know that segment's netmask with certainty. Remote parts of the Internet are already expected not to make assumptions about whether or not a particular address is a broadcast address, since that determination is already only meaningful for devices connected to the segment containing that address. This document does not change any of these things. Thus, if the behavior of devices on a particular network segment has been updated in accordance with this memo, the lowest address on that segment can already be addressed by hosts elsewhere on the Internet without any changes to their own software.

[RFC1812] noted in section 4.2.3.1 that whether a reserved address is treated specially at all depends on one's vantage point on the network:

```
| [a] router obviously cannot recognize addresses of the form {  
| <Network-prefix>, 0 } if the router has no interface to that  
| network prefix. In that case, the rules of the second bullet  
| [requiring a packet to be discarded] do not apply because, from  
| the point of view of the router, the packet is not an IP broadcast  
| packet.
```

It also noted in section 5.3.5.2 that

| in view of CIDR, such [packets addressed to broadcast addresses of
| distant networks] appear to be host addresses within the network
| prefix; we preclude inspection of the host part of such network
| prefixes.

To the extent that software continues to make assumptions about IP network classes today, it is out of compliance with RFC 4632. Ever since the adoption of CIDR in RFC 1519, it has been unknowable whether or to what extent the remote network would internally aggregate or deaggregate routes that were visible elsewhere on the Internet. Therefore, Internet hosts and routers MUST NOT assume that an IPv4 address on a remote network, other than 0.0.0.0, is invalid, unroutable, or inaccessible merely because it ends with a particular number of zeroes. In keeping with the Internet's end-to-end principle, decisions about possible invalidity of otherwise routable addresses belong as close to the endpoints as possible.

Although this memo does not change the broadcast status of the highest address on a segment, the same considerations discussed in this section correspondingly permit a network operator to unilaterally decide to use that address for unicast purposes. Distant networks are identically forbidden from making assumptions about the broadcast status of an address based on the number of ones it ends with. Unilaterally using both historic broadcast addresses as unicast addresses could be viewed as a generalization of the policy of [RFC3021], with similar benefits, albeit in a different context.

2.2. Multicast Addresses; Point-to-Point Links

Multicast addresses, as defined by RFC 1112 [RFC1112], do not have a network part and host part, nor do they have a netmask or CIDR prefix length. IPv4 multicast addresses, except 224.0.0.0 (which is "guaranteed not to be assigned to any group" by RFC 1112), could always end with any number of zeroes, and have never had any form of directed broadcast address.

[RFC3021], section 2.1, standardized that, in a point-to-point link using a 31-bit netmask, the all-zero and all-one forms of the host-part of the address MUST both be treated as unicast ("host") addresses.

The present document does not change the interpretation of multicast addresses or 31-bit subnet addresses in any way.

2.3. Current Limitations on Subnet-Directed Broadcast Addresses

Sending packets to a subnet-directed broadcast address is itself an obsolescent mechanism in today's Internet. It now works only for nodes attached directly to that subnet, and has also been replaced by link-local multicast for most application purposes. [RFC2644] discouraged routers from forwarding directed broadcast packets, to reduce their use in amplifying denial-of-service attacks, so they typically cannot be received when sent from distant hosts. Many hosts today completely ignore ICMP packets sent as broadcasts, so a directed broadcast ping is no longer a reliable means of enumerating all hosts attached to a network. As Informational [RFC6250] notes, "broadcast can only be relied on within a link"; moreover, some implementations no longer reply even to broadcasts within a link. Broadcasting for IPv4 autoconfiguration [RFC2131] uses the special addresses 0.0.0.0/32 and 255.255.255.255/32 for the destination.

Neighboring host discovery for other application-layer purposes is now most often handled by mechanisms such as Multicast DNS [RFC6762] and DNS-Based Service Discovery [RFC6763], which uses the hard-coded special multicast address 224.0.0.251 regardless of the local link's subnet configuration. The above-mentioned [RFC6250] says that, "[t]o a large extent, the addition of link-scoped multicast to the IP service model obsoleted the need for broadcast." In general, link-local multicast addresses can be used by applications that need to address an unknown group of other hosts on the same network. Hosts can indicate their interest in being contacted for a specific purpose by listening on its associated address within 224.0.0.0/24, like 224.0.0.251 for hosts that are prepared to reply to local mDNS queries.

2.4. Comparison to IPv6 Behavior

In IPv6 there are no reserved per-segment broadcast addresses (or, indeed, any broadcast addresses whatsoever). Instead, IPv6 hosts can address all hosts on a network segment by using the link-local multicast group address ff02::1 [RFC4291], which, for example, produces a multicast Ethernet frame when transmitted over an Ethernet-like link [RFC2464]. The lowest address on a subnet is, however, reserved in IPv6 by Section 2.6.1 of RFC 4291 [RFC4291] for the Subnet-Router address (a means of addressing "any router" on an indicated subnet).

3. Change to Interpretation of the Lowest Address

The purpose of this document is to designate the all-zeroes address in each subnet as a unicast address. All such addresses are now available for general non-broadcast use, treated identically to all host addresses in the subnet besides the "all-ones" broadcast address. This document therefore eliminates an element of the IPv4 protocol's historical adaptation to 4.2BSD's behavior. All hosts SHOULD continue to recognize and accept only the all-ones form of the IPv4 subnet broadcast address.

Host software that intends to transmit a segment-directed broadcast packet in an IPv4 network MUST use only the all-ones form as the destination address of the packet.

An IPv4 datagram containing a source or destination that is equal to the all-zeroes address in its subnet SHOULD be treated, by both hosts and routers, as a normal unicast datagram; it SHOULD NOT be treated as a local broadcast datagram.

Host software SHOULD allow a network interface to be configured with the lowest address on a subnet. A host with such an address configured MUST use this assigned address as a source address for datagrams just as it would with any other assigned interface address, and MUST recognize a datagram sent to that address as addressed to itself. Host software SHOULD be capable of generating unicast packets to the lowest address on a subnet when so requested by an application, and MUST encapsulate such packets into link-layer unicast frames when transmitted on a link layer that distinguishes unicast and broadcast.

Clients for autoconfiguration mechanisms such as DHCP [RFC2131] SHOULD accept a lease or assignment of the lowest address whenever the underlying operating system is capable of accepting it. Servers for these mechanisms SHOULD assign this address when so configured. The network operator of each subnet retains the discretion to number hosts on that subnet with, or without, the use of the lowest address, based on local conditions.

3.1. Link-Layer Interaction

The link layer always indicates to the IP layer whether or not a datagram was transmitted as a broadcast at the link layer. Hosts MUST continue to follow the RFC 1122 rule about link-layer broadcast indications:


```
| A host SHOULD silently discard a datagram that is received via a  
| link-layer broadcast [...] but does not specify an IP multicast or  
| broadcast destination address.
```

This rule is, among other things, intended to avoid broadcast storms. This document now defines the lowest address as a non-broadcast address. Therefore, a host SHOULD silently discard a datagram received via a link-layer broadcast whose destination address is the lowest IPv4 address in a subnet. This is true even if the interface on which the host received that datagram uses the lowest address as a unicast IPv4 address.

3.2. Recommendations

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

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

The new section numbered 3.2.1.3 (h) which was added by RFC 3021 is replaced with:

```
| (h) { <Network-number>, <Subnet-number>, 0 }  
  
| An ordinary unicast ("host") address in the subnet. May be used  
| as either a source or destination address. If a link-level  
| broadcast packet is received with this address (or any other  
| unicast address) as its destination, it MUST be silently  
| discarded. Such a packet may be sent by long-obsolete hosts on  
| the local network.  
  
| In applications using CIDR notation [RFC4632], this address, or  
| any other address in the subnet, may also be used together with a  
| prefix length to refer to the entire subnet.
```

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

The new section (numbered 4.2.2.11 (f)) added by RFC 3021 is replaced by:

```
| (f) { <Network-prefix>, 0 }  
|
```

An ordinary unicast ("host") address in the subnet. May be used as either a source or destination address. If a link-layer broadcast packet is received with this address (or any other unicast address) as its destination, it MUST be silently discarded. Such a packet may be sent by long-obsolete hosts on the local network.

In applications using CIDR notation [RFC4632], this address, or any other address in the subnet, may also be used together with a prefix length to refer to the entire subnet.

The first paragraph on page 49 (which appears after section 4.2.2.11 (e) in the original RFC 1812, or after section 4.2.2.11 (f) in RFC 1812 as modified by RFC 3021) is changed from this original text

IP addresses are not permitted to have the value 0 or -1 for the <Host-number> or <Network-prefix> fields except in the special cases listed above. This implies that each of these fields will be at least two bits long.

DISCUSSION

Previous versions of this document also noted that subnet numbers must be neither 0 nor -1, and must be at least two bits in length. In a CIDR world, the subnet number is clearly an extension of the network prefix and cannot be interpreted without the remainder of the prefix. This restriction of subnet numbers is therefore meaningless in view of CIDR and may be safely ignored.

to this new text

Unicast IP addresses are permitted to have the value 0 for the <Host-number> field, and may have the value -1 in the special cases listed above. There is no requirement that the <Host-number> field be any particular length. In some cases using CIDR notation, a host may be designated with a /32 suffix (e.g. 192.0.2.34/32), indicating that the specific host rather than its subnet is being described.

DISCUSSION

Previous versions of this document also noted that subnet numbers must be neither 0 nor -1, and must be at least two bits in length. Other versions required that <Network-prefix> fields must be neither 0 nor -1, and must be at least two bits long.

Now that the Internet has fully transitioned to CIDR routing, there are no original classful <Network-number>s to be distinguished from <Subnet-numbers>. Each address only has a <Network-prefix> based on its network mask (or equivalently, the CIDR suffix specifying how many bits are in the <Network-prefix>). The former restrictions on subnet numbers and their sizes are meaningless in view of CIDR and are hereby repealed. For example, a route to 0.0.0.0/6 or even 0.0.0.0/1 is a viable CIDR route (for the aggregation of the blocks 0/8, 1/8, 2/8, and 3/8; or for the entire lower half of the IPv4 address space) and should not be considered invalid. 0.0.0.0/0 is standardized to mean "all unicast IPv4 addresses", e.g. in a default route, by section 5.1 of [RFC4632], which MUST also continue to work.

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

(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. 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.

A packet addressed to { <Network-prefix>, 0 } is an ordinary unicast packet, and MUST be treated as such.

(4) SHOULD NOT originate datagrams addressed to 0.0.0.0. SHOULD allow for the generation of datagrams addressed to {<Network-prefix>, 0 } since that is now defined as an ordinary unicast address.

3.3. Example

The only IPv4 broadcast address for 192.168.42.0/24 is 192.168.42.255 (the all-ones or "-1" host number). 192.168.42.0 (the all-zeroes or "0" host number) was formerly a second broadcast address on that subnet, but is now a unicast address.

The fact that the address identifier 192.168.42.0 can refer both to a network (as when it is regarded as a "network address" with or without an explicit prefix length) and to a specific host 192.168.42.0 is not unusual. Similarly, referring to a subnet as 192.168.42.0/24 and configuring a particular interface on that subnet as 192.168.42.0/24 is also not unusual. Computer scientists normally count all sorts of things starting at the zeroth (lowest) element in a sequence.[EWD831] For example, the initial element in an array is likely to be stored at a memory address equal to the memory address of the array itself.[ARRAY] Similarly, IPv4 hosts in a subnet MAY be enumerated starting with an address that matches the address of the subnet itself.

Similarly, the only IPv4 broadcast address for the subnet 192.168.42.96/28 is 192.168.42.111. The address 192.168.42.96 MAY be assigned to an individual host on this network.

3.4. Compatibility and Interoperability

Many deployed systems follow older Internet standards in not allowing the lowest address in a network to be assigned or used as a source or destination address. Assigning this address to a host may thus make it inaccessible by some devices on its local network segment. Network operators considering assigning this address to a host should investigate their own network environments to determine whether their interoperability requirements will be met. Interoperability with these addresses is likely to improve over time, following the publication of this document.

Prior standards required hosts and routers to ignore, and to refrain from generating, non-broadcast datagrams from or to the lowest address. So when a single network contains a device that has been assigned the lowest address as specified by this document, along with one or more devices that follow the traditional behavior, the traditional devices will not be able to communicate with the lowest-address device at all. Other sorts of malfunctions are unlikely, because the former standards (RFC 1122) required traditional hosts to drop any unicast packet addressed to the secondary broadcast address that they implemented at the lowest address.

4. IANA Considerations

This memo includes no request to IANA.

5. Security Considerations

The behavior change specified by this document could produce security concerns where two devices, or two different pieces of software on a single host, or a software application and a human user, follow divergent interpretations of the lowest address on a network. For example, this could lead to errors in the specification or enforcement of rules about Internet hosts' connectivity to one another, or their right to access resources.

Firewall rules that assume that the lowest address on a subnet cannot be addressed SHOULD be updated to take into account that it can be addressed, so as to avoid either unintentionally allowing or unintentionally forbidding connections involving it. Other security, monitoring, or logging systems that treat the lowest address as an inaccessible bogon address SHOULD likewise be updated.

Host software SHOULD make the distinction between lowest-address (considered individually) and subnet (considered as a group) clear to users, where this distinction is relevant and could be a subject of confusion.

6. Acknowledgements

This document directly builds on prior work by Dave T辰ht and John Gilmore as part of the IPv4 Unicast Extensions Project.

We thank our late coauthors Michael J. Karels (1956-2024) and Dave T辰ht (1965-2025) for their work on this draft and related implementation efforts.

We thank Eric Siegerman for helpful drafting fixes and Chris Hall for suggestions and information about Windows support.

7. References

7.1. Normative References

- [RFC0919] Mogul, J., "Broadcasting Internet Datagrams", STD 5, RFC 919, DOI 10.17487/RFC0919, October 1984, <<https://www.rfc-editor.org/info/rfc919>>.
- [RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, DOI 10.17487/RFC1122, October 1989, <<https://www.rfc-editor.org/info/rfc1122>>.

- [RFC1812] Baker, F., Ed., "Requirements for IP Version 4 Routers", RFC 1812, DOI 10.17487/RFC1812, June 1995, <<https://www.rfc-editor.org/info/rfc1812>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC4632] Fuller, V. and T. Li, "Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan", BCP 122, RFC 4632, DOI 10.17487/RFC4632, August 2006, <<https://www.rfc-editor.org/info/rfc4632>>.

7.2. Informative References

- [ARRAY] Wikipedia, "C Programming Language: Array-pointer interchangeability", <[https://en.wikipedia.org/wiki/C_\(programming_language\)#Array%E2%80%93pointer_interchangeability](https://en.wikipedia.org/wiki/C_(programming_language)#Array%E2%80%93pointer_interchangeability)>.
- [BSDHIST] Wikipedia, "History of the Berkeley Software Distribution", <https://en.wikipedia.org/wiki/History_of_the_Berkeley_Software_Distribution>.
- [EWD831] Dijkstra, E.W., "Why Numbering Should Start at Zero", August 1982, <<https://www.cs.utexas.edu/users/EWD/transcriptions/EWD08xx/EWD831.html>>.
- [IEN212] Gurwitz, R. and R. Hinden, "IP - Local Area Network Addressing Issues", IEN 212, September 1982, <<https://www.postel.org/ien/pdf/ien212.pdf>>.
- [RFC0894] Hornig, C., "A Standard for the Transmission of IP Datagrams over Ethernet Networks", STD 41, RFC 894, DOI 10.17487/RFC0894, April 1984, <<https://www.rfc-editor.org/info/rfc894>>.
- [RFC1112] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, DOI 10.17487/RFC1112, August 1989, <<https://www.rfc-editor.org/info/rfc1112>>.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, DOI 10.17487/RFC2131, March 1997, <<https://www.rfc-editor.org/info/rfc2131>>.

- [RFC2464] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, DOI 10.17487/RFC2464, December 1998, <<https://www.rfc-editor.org/info/rfc2464>>.
- [RFC2644] Senie, D., "Changing the Default for Directed Broadcasts in Routers", BCP 34, RFC 2644, DOI 10.17487/RFC2644, August 1999, <<https://www.rfc-editor.org/info/rfc2644>>.
- [RFC3021] Retana, A., White, R., Fuller, V., and D. McPherson, "Using 31-Bit Prefixes on IPv4 Point-to-Point Links", RFC 3021, DOI 10.17487/RFC3021, December 2000, <<https://www.rfc-editor.org/info/rfc3021>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.
- [RFC6250] Thaler, D., "Evolution of the IP Model", RFC 6250, DOI 10.17487/RFC6250, May 2011, <<https://www.rfc-editor.org/info/rfc6250>>.
- [RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, DOI 10.17487/RFC6762, February 2013, <<https://www.rfc-editor.org/info/rfc6762>>.
- [RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, DOI 10.17487/RFC6763, February 2013, <<https://www.rfc-editor.org/info/rfc6763>>.

Appendix A. Implementation Status

The behavior specified in this document has been implemented by OpenBSD since version 4.7, released in May 2010.

The behavior specified in this document has been implemented by the Linux kernel since version 5.14, released in August 2021.

This behavior has also been implemented in the OS releases Fedora 35 (released in November 2021) and Ubuntu 22.04 (released in April 2022). In addition, nodes that run Fedora 34 and install the standard online updates also implement this feature. This behavior has also been included in OpenWrt as of OpenWrt 23.05 (released in October 2023) and in Yocto Linux as of Yocto Linux 4.0 "Kirkstone" (released in May 2022).

This behavior was an option in NetBSD since 1999 (via the `net.inet.ip.hostzerobroadcast sysctl`), and became the default behavior in 2016.

This behavior has also been implemented in FreeBSD since release 13.1 of May 2022. The FreeBSD, OpenBSD, NetBSD, and Linux implementations interoperate successfully.

Popular embedded Internet-of-Things environments such as RIOT and FreeRTOS have implemented this behavior for many years.

This behavior is optionally available in Android 13 (released in August 2022) when using the android13-5.15 kernel version, and is implemented by default in Android 14 beta.

Microsoft Windows 11 supports non-broadcast use of the lowest address on a subnet: it allows assigning the lowest address on a subnet to a Windows network device, as well as communicating with a host or router on the local network using such an address. Windows 11 was released in October 2021, but additional tests or documentation are needed to confirm whether this behavior was already present as of its original release date.

Huawei VRP 5.160 (released 2014) does not allow the lowest address on a subnet to be manually applied to a router interface, and does not successfully communicate with such an address when it is applied to another device on a locally-attached subnet.

Compatibility with lowest-address assignment may be typical of many DHCP implementations (because the lowest address special case has often been handled at the kernel level). If the underlying operating system supports lowest-address assignment, the final official ISC DHCP release (4.4.3) supports lowest-address allocation as both client and server, as do Busybox DHCP udhcp/udhcpd (release 1.1.15), and ISC Kea (which currently includes only a DHCP server implementation).

Authors' Addresses

Seth David Schoen
IPv4 Unicast Extensions Project
San Francisco, CA
United States of America
Email: schoen@loyalty.org

John Gilmore
IPv4 Unicast Extensions Project
PO Box 170640-rfc
San Francisco, CA 94117-0640
United States of America
Email: gnu@rfc.toad.com

David M. T辰ht
IPv4 Unicast Extensions Project
Email: dave@none.org

Michael J. Karels
IPv4 Unicast Extensions Project
Email: mkarels@none.org