Internet Engineering Task Force
Internet-Draft
Updates: 1122, 3704, 6890 (if approved)
Intended status: Standards Track
Expires: 25 December 2025

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

Unicast Use of the Formerly Reserved 240/4
draft-schoen-intarea-unicast-240-09

Abstract

This document redesignates 240/4, the region of the IPv4 address space historically known as "Experimental," "Future Use," or "Class E" address space, so that this space is no longer reserved. It asks implementers to make addresses in this range fully usable for unicast use on the Internet.

Table of Contents

1. Introduction 2
1.1. Requirements Language 2
2. Background 2
2.1. History of IPv4 Address Types 3
2.2. Reserved IPv4 Addresses in the RFC Series 4
2.3. Attempts to Use the "Future Use" Addresses 4
2.4. Recent Use as Ordinary Unicast Addresses 5
3. Change in Status of 240/4 6
3.1. Continued Special Treatment for 255.255.255.255/32 . . . 7
4. Compatibility and Interoperability 7
5. IANA Considerations 9
6. Security Considerations 10
6.1. Existing Unofficial Uses of 240/4 11
7. Acknowledgements 12
8. References 12
8.1. Normative References 12
8.2. Informative References 14
Appendix A. Implementation Status 17
A.1. Operating systems 17
A.2. Routers and Switches 18
A.3. DHCP implementations 19
A.4. Other implementations 19
A.5. Internet of Things 20
Authors' Addresses 20

1. Introduction

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 redefine the "Experimental" or "Future Use" 240/4 region (historically known as "Class E" addresses) as ordinary unicast addresses. These 268 million IPv4 addresses are already usable for unicast traffic in many popular implementations today. Standardization as unicast addresses will eventually allow them to be later deployed by Internet stewardship organizations to relieve address space scarcity.

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

2.1. History of IPv4 Address Types

When the Internet Protocol was being designed, it was unclear whether it would be a success, or which of its features might be the key features that led to success. The bulk of its address space was dedicated to ordinary "host addresses". Other blocks and corners of the address space were reserved, either for particular protocol functions such as loopback, LAN broadcasting, or host bootstrapping, or for future definition. A major allocation of 268 million addresses was later made for multicasting [RFC0988], while leaving another 268 million reserved for "future use". After the invention of broadcast and multicast, the original ordinary host addresses were later described as unicast addresses, which is now the usual terminology.

With decades of hindsight, we can now see that unicast has been the success story of the Internet. Trillions of unicast packets now move around the world daily. By contrast, the non-unicast addresses are seldom used. The use of routable broadcast packets in denial of service attacks has now limited broadcast packets to local-area networks [RFC2644], and to critical but highly-specialized protocol functions such as DHCP [RFC2131], routing updates [RFC1256], or neighbor discovery.

Wide-area multicast packets had a brief research heyday, but never reached critical mass. Today, the overwhelming majority of multiply-replicated media streams (such as popular songs and videos, television programs, conference calls, and video meetings) are carried in unicast packets mediated by application-level replication rather than IP-protocol-level multicasting or broadcasting.

The Internet became a rapid worldwide success. Partly due to the reduction in experimentation that accompanied that success, little effort has been paid to looking back at the historical allocations of reserved addresses. The success of unicast traffic has led to a huge demand for unicast addresses. By contrast, there is far more supply of reserved, ignored, loopback, and multicast addresses than any foreseeable IPv4 Internet will demand. Most of these historical accidents were not carried forward into the IPv6 protocol [RFC4291]. We propose simple, compatible changes to existing IPv4 implementations that will increase the supply of unicast addresses by redesignating addresses that today are almost completely unused on the Internet. The best and easiest "future use" of many of today's formerly reserved IPv4 addresses is as ordinary unicast addresses.

2.2. Reserved IPv4 Addresses in the RFC Series

The Assigned Numbers RFC series reserved various IP addresses or assigned them special meanings, starting in 1977 and continuing through the early 1990s. The detailed behavioral requirements for IPv4 implementations based on these designations are set out in October 1989's RFC 1122 [RFC1122]. As other special cases continued to be introduced on occasion, RFC 3232 [RFC3232] announced that IANA would track such information in an online database; the present-day version of this mechanism is the IPv4 Special-Purpose Address Registry [IANA4], as provided for by RFC 6890 [RFC6890]. A wide range of host and network software follows these designations by treating these Internet addresses specially.

This document is concerned with the largest special case in RFC 1122: the designation of an entire /4 block for Future Use. In retrospect, the flexibility offered by keeping these addresses unused was insightful for its time, but since they ended up never being needed for any special purposes, they have become the least productive portion of the Internet address space.

The largest block of original addresses reserved for future use in 1983 was called "Class D" in RFC 870 [RFC0870], and contained what would now be called 224/3. This contained about 536 million addresses, about 12.5% of the total available address space. By 1986, RFC 988 [RFC0988] split the former Class D in half, designating a multicast Class D block, now called 224/4, and a future-use Class E block, now called 240/4. Following the 1993 implementation of CIDR [RFC1519] and its 2006 clarification [RFC4632], we no longer speak of any IPv4 address as having an "address class," but the reservations of these specific addresses that were made by RFC 1122, were unaffected by the CIDR change in terminology and routing technology.

2.3. Attempts to Use the "Future Use" Addresses

Through the 1980s, there were many reasons to suppose that new forms of Internet addressing could emerge, so reserving a substantial number of addresses for them was prudent.

One likely candidate for some time was protocol translation methods between IP and other protocols using special surrogate IP addresses. This possibility was particularly significant during the time frame when IP coexisted widely on heterogeneous networks with other protocols. Special number ranges could have been used to facilitate interoperability, protocol translation, or encapsulation between IP and non-IP protocols.

This prospect received new salience with the adoption of IPv6, where some deployed or proposed transition mechanisms use special-purpose IPv4 addresses with a distinctive meaning in the context of IPv6 transition, such as NAT64 [RFC7050] and the deprecated 6to4 [RFC3068]. While IPv6 transition mechanisms could conceivably have used portions of 240/4, they ended up instead using very small amounts of special address space from the IETF Protocol Assignments block 192.0.0.0/24 or elsewhere within the unicast space.

Another form of addressing that was novel in 1989 is anycast addressing, in which the same address is used to identify servers at physically distinct locations and connected to the Internet at different points. It would have been possible to designate a new "class" of addresses for anycast operations. RFC 1546 [RFC1546], which first defined anycast, concluded that this would be a possible and even desirable approach:

> There appear to be a number of ways to support anycast addresses, some of which use small pieces of the existing address space, others of which require that a special class of IP addresses be assigned. [...] In the balance it seems wiser to use a separate class of addresses.

But anycast services turned out to work fine in most respects by using existing unicast routing protocols, existing unicast datagram delivery protocols, and ordinary unicast addresses. They are now widely used for specific applications [RFC7094] such as the Internet’s root nameservers.

2.4. Recent Use as Ordinary Unicast Addresses

Overall, 30 years of experience have demonstrated that no new addressing mechanism requires the use of 240/4; nor is any likely to require it in the future, particularly in light of the IPv6 transition. Other explicit reservations such as the IETF Protocol Assignments block at 192.0.0.0/24 have been sufficient. While it was reasonable to plan for an unknown future, the reserved block at 240/4 did not ultimately aid Internet innovation or functionality. The future has arrived, and it wants IPv4 unicast addresses far more than it wants permanently unusable IPv4 addresses.

The idea of making 240/4 addresses available for unicast addressing is not new. It was suggested by Lear on the influential TCP-IP mailing list in 1988 [Lear]. It was formally proposed to IETF more than a decade ago, both by Fuller, Lear, and Mayer [FLM], and by Wilson, Michaelson, and Huston [WMH]. While the idea of unicast use of 240/4 was merely being considered at IETF, the "running code" required was simple enough and compatible enough that this behavior change was implemented at that time in several operating systems. Then, when the protocol change was ultimately not standardized, those implementations remained, but were largely forgotten. (They are summarized in the "Implementation Status" section of this document.)

The unicast support created in about 2008 in those implementations is now running in millions of nodes on the Internet, and has not caused any problems over the past decade. As a result, the 240/4 space has been attracting "wildcat" use in private networks; see [VPC].

Although software support for unicast use of 240/4 is widespread, it is not yet universal. The present document moves this process further along by confirming the consensus that unicast is the preferred use for 240/4, documenting the exact behavior changes required for maximum interoperability, and calling on all vendors and implementers to adopt this behavior. Doing so will prepare for a future in which use of these addresses is anticipated and unsurprising, so that their allocation can be considered.

Implementations generally treat public and private addresses identically, with the differences occurring only in how routes, firewalls, and DNS servers are configured. The earlier draft [WMH] suggested designating the unreserved 240/4 range as [RFC1918]-style private address space. Like the [FLM] draft, this document does not attempt to decide or designate whether future allocations from this address range will be public or private addresses. Both options require that both hosts and routers be able to use these addresses, so the next section fully defines both host and router behavior.

3. Change in Status of 240/4

The purpose of this document is to make addresses in the range 240/4 available for active unicast use on the public Internet. This includes supporting them for numbering and addressing networks and hosts, like any other unicast address.

Host and router software SHOULD treat addresses in the 240/4 range in the same way that they would treat other unicast IPv4 addresses. Software SHOULD be capable of accepting datagrams from, and generating datagrams to, addresses within this range.

Clients for autoconfiguration mechanisms such as DHCP [RFC2131] SHOULD accept a lease or assignment of an address within 240/4 whenever the underlying operating system is capable of accepting it.

Other interoperability details related to address-based filtering are discussed in a separate section, below.

3.1. Continued Special Treatment for 255.255.255.255/32

The address 255.255.255.255/32 was given a special meaning as a local segment limited broadcast address by numerous prior Internet standards, starting with RFC 919 [RFC0919] and continuing consistently up to the present day. For example, 255.255.255.255 is used as a network-layer destination address in BOOTP [RFC0951] and DHCP [RFC2131] for address autoconfiguration broadcasts by hosts that don’t yet know anything about the networks to which they are connected. While some newer autoconfiguration or autodiscovery protocols use other addresses, the use of 255.255.255.255 remains widespread.

The special meaning of 255.255.255.255 was never restricted or affected by the reservation of 240/4. Accordingly, the existing distinctive meaning of 255.255.255.255 is unchanged by this document. This single address MUST NOT be assigned to an individual host, or interpreted as the address of an individual host, even if it would otherwise be part of an allocated or announced network block.

4. Compatibility and Interoperability

Merely implementing unicast treatment of addresses in 240/4 in routers and operating systems, as this document proposes, does not cause any compatibility nor interoperability issues. Hundreds of millions of IPv4 nodes currently contain this unicast treatment, and all are interoperating successfully with each other and with non-updated nodes.

Compatibility and interoperability issues only arise if and when an address from 240/4 is assigned to an interface on a node, and then IPv4 packets are exchanged which use such an address as a source or destination address. This document does not recommend doing these things, except for testing purposes.

Older Internet standards counseled implementations in varying ways to reject packets from, and not to generate packets to, addresses within 240/4.

RFC 1122 [RFC1122], section 3.2.1.3, states that a "host MUST silently discard an incoming datagram containing an IP source address that is invalid by the rules of this section." The same section states that Class E addresses are "reserved" (which might be taken, in context, to imply that they are "invalid"); the section further treats Class A, B, and C as the only possibly relevant address ranges for unicast addressing.

RFC1812 [RFC1812], section 5.3.7, states that a "router SHOULD NOT forward" a packet with such a destination address. (If section 4.2.2.11's reference to these addresses as "reserved" is taken to imply that they are "special," section 5.3.7 would also imply that a "router SHOULD NOT forward" a packet with such a source address.)

RFC 3704 [RFC3704] (BCP 84) cites RFC 2827 [RFC2827] (BCP 38) in asking providers to filter based on source address:

> RFC 2827 recommends that ISPs police their customers' traffic by dropping traffic entering their networks that is coming from a source address not legitimately in use by the customer network. The filtering includes but is in no way limited to the traffic whose source address is a so-called "Martian Address" - an address that is reserved, including any address within 0.0.0.0/8, 10.0.0.0/8, 127.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16, 224.0.0.0/4, or 240.0.0.0/4.

In this context, RFC 3704 specifies filtering of these addresses as source (not destination) addresses at a network ingress point as a countermeasure against forged source addresses, limiting forwarded packets' source addresses to only the set which have been actually assigned to the customer's network. The RFC's mention of these "Martian Addresses" is based on the assumption that they could never be legitimately in use by the customer network.

Because the 240/4 address space is no longer reserved as a whole, an address within this space is no longer inherently a "Martian" address. Both hosts and routers MUST NOT hard-code a policy of always rejecting such addresses. Hosts and routers SHOULD NOT be configured to apply Martian address filtering to any packet solely on the basis of its reference to a source (or destination) address in 240/4. Maintainers of lists of "Martian addresses" MUST NOT designate addresses from this range as "Martian". As noted elsewhere, the address 255.255.255.255 retains its special meaning, but is also not a "Martian" address.

The filtering recommended by RFC 3704 is designed for border routers, not for hosts. To the extent that an ISP had allocated an address range from within 240/4 to its customer, RFC 3704 would already not require packets with those source addresses to be filtered out by the ISP's border router.

Since deployed implementations' willingness to accept 240/4 addresses as valid unicast addresses varies, a host to which an address from this range has been assigned may also have a varying ability to communicate with other hosts.

Such a host might be inaccessible by some devices either on its local network segment or elsewhere on the Internet, due to a combination of host software limitations or reachability limitations in the network. IPv4 unicast interoperability with 240/4 can be expected to improve over time following the publication of this document. Before or after allocations are eventually made within this range, "debogonization" efforts for allocated ranges can improve reachability to the whole address block. Similar efforts have already been done by Cloudflare on 1.1.1.1 [Cloudflare], and by RIPE Labs on 1/8 [RIPElabs18], 2a10::/12 [RIPElabs2a1012], and 128.0/16 [RIPElabs128016]. The Internet community can use network probing with any of several measurement-oriented platforms to investigate how usable these addresses are at any particular point in time, as well as to localize medium-to-large-scale routing problems. (Examples are described in [Huston], [NLNOGRing], and [Atlas].) Any network operator to whom such addresses are made available by a future allocation will have to examine the situation in detail to determine how well its interoperability requirements will be met.

5. IANA Considerations

This memo unreserves the address block 240/4. It therefore requests IANA to update the IANA Special-Purpose Address Registry by removing the entry for 240/4, whose existing authority is RFC 1122, Section 4. Additionally, it requests IANA to update the IANA IPv4 Address Space Registry by changing the status of each /8 entry from 240/8 through 255/8 from "Future Use, 1981-09, RESERVED" to "Unallocated, [Date of this RFC], UNALLOCATED". Finally, IANA is requested to prepare for this address space to be addressed in the reverse DNS space in in-addr.arpa.

This memo does not effect a registration, transfer, allocation, or authorization for use of these addresses by any specific entity. This memo’s scope is to require IPv4 software implementations to support the ordinary unicast use of addresses in the newly unallocated range 240.0.0.0 through 255.255.255.254. During a significant transition period, it would only be prudent for the global Internet to use those addresses for experimental purposes such as debogonization and testing. After that transition period, a responsible entity such as IETF or IANA could later consider whether, how and when to allocate those addresses to entities or to other protocol functions such as private addresses.

6. Security Considerations

The change specified by this document could create a period of ambiguity about historical and future interpretations of the meaning of host and network addresses in 240/4. Some networks and hosts currently discard all IPv4 packets bearing these addresses, pursuant to statements in prior standards that packets containing these addresses have no agreed-upon meaning. Such implementations have protected themselves from possible incompatible future packet formats that might have eventually used these addresses.

Disparate filtering processes and rules, both at present, and in response to the adoption of this document, could make it easier for rogue network operators to hijack or spoof portions of this address space in order to send malicious traffic.

Live traffic, accepted and processed by other devices, may legitimately originate from these addresses in the future. Network operators, firewalls, and intrusion-detection systems may need to take account of this change in various regards, to avoid permitting either more or less traffic from such addresses than they expected.

Automated systems generating reports, and human beings reading those reports, SHOULD NOT assume that the use of a 240/4 source address indicates spoofing, an attack, or a new incompatible packet format. At the same time, they SHOULD NOT assume that the use of 240/4 is impossible or will be precluded by other systems' behavior.

An important concern about the [FLM] and [WMH] drafts was that discrepant behavior between systems could create security problems, as when a middlebox fails to detect or report an attack or policy violation because it believes that an address involved cannot be used or cannot be relevant. Similarly, a logging system could fail to log traffic related to 240/4 addresses because it incorporates an assumption that no such traffic can ever occur. Such discrepancies between multiple systems' views of communication semantics are a common security antipattern. (Compare [Sherr], exploiting discrepancies in telephony equipment's recognition and interpretation of DTMF signals.) Any change to the meaning or status of a group of addresses can introduce such a discrepancy.

In this case, because 240/4 is already commonly supported by several widely-used implementations, and is already used for private network communications, such discrepancies are already a reality. If routers follow this document's request to cease filtering this address range, they will increase the variety of contexts in which implementations may receive ordinary unicast packets containing these addresses. (Such packets are still unlikely to arrive from distant hosts until some of these addresses are eventually allocated for experimental or production use, and until the global routing table receives announcements for subnets in this range.)

The adoption of this document will converge on an explicitly shared understanding that implementations should prepare for this possibility. Since unofficial private use of 240/4 addresses is a reality today, while any public allocations from this range are still distant and contingent on further study, implementers are receiving considerable advance notice of this issue.

6.1. Existing Unofficial Uses of 240/4

Some organizations are using portions of 240/4 internally as RFC 1918-type private-use address space, for example for internal communications within datacenters. RIPE’s ATLAS project experimentally detected the use of this address space in several large institutions [RIPE240] in 2022. Amazon subsequently described [Dale] its internal use of 240/4 at the NANOG88 conference in 2023. As an Amazon engineer explained:

> [A]long the way, we actually ran out of RFC 1918, 3330, every piece of address space out there, so we did something that some people have noticed. [...] Some of those hops have class E addresses. Yeah, we did do that. We started numbering things out of Class E. Now that was very kind of deliberate in terms of doing that. We had no address space left.

Google has informally advised hosting customers [VPC] for several years that they may use this address space this way, and in September 2024 made that advice more explicit [GKE], saying that customers

> can use [240/4] for private use within Google Cloud VPCs, for both Compute Engine instances or Kubernetes pods/services in GKE. [...] [These] addresses offer a significantly larger pool of IP addresses compared to traditional RFC 1918 private addresses [...]

and providing technical details about how to do so, including some information about vendor support. The same document notes that Google cloud customer Snap Inc.’s use of 240/4 addresses

> mitigated IP exhaustion and provided Snap with multiple years of IP address scale needed to support future growth and reduce operational overhead. In addition, this approach also aligned with Snap窶冱 long-term strategy for migrating to IPv6.

A Snap engineering manager is quoted as saying that these address resources "doubled [Snap’s] IP capacity [...] and eliminated address exhaustion and subnet management issues".

These organizations’ use of 240/4 address space provides further demonstration that this address space can actually be used in practice by many devices and operating systems today, but it also represents a failure to coordinate this use with the global Internet community. That failure may be harmful in the future. For instance, Google predicts [GKE] that "adoption [of 240/4 addressing] is expected to grow as the need for more IP addresses increases", but now encourages customers to use them privately for internal data center communication without coordination, and confirms that large customers are already doing so as of 2024. Coordination could make it more likely that some or all of this address space can be used for interdomain or global communications in the future without renumbering.

Future allocations of 240/4 could result in use of this space on the public Internet in ways that overlap these unofficial private-use addresses, creating ambiguity about whether a particular host intended to use such an address to refer to a private or public network. Among other unintended outcomes, hosts or firewalls that have extended greater trust to other hosts based on their use of a certain unofficial network number (that was considered to imply presence on a LAN or within an organization) may eventually receive legitimate traffic from an external network to which this address space has been allocated. On the other hand, unofficial private use may reduce the ability of networks to successfully communicate with each other in the future if the global Internet community publicly allocates address blocks from 240/4 in the future.

Operators of networks that are making unofficial uses of portions of 240/4 may wish to plan to discontinue these uses and renumber their internal networks, or to request that IANA formally designate certain ranges as additional Private-Use areas.

7. Acknowledgements

This document directly builds on prior work by Dave Tテ、ht (1965-2025) and John Gilmore as part of the IPv4 Unicast Extensions Project.

We thank our late colleague Michael J. Karels (1956-2024) for his comments on this draft and related FreeBSD implementation efforts.

8. References

8.1. Normative References

[IANA4] Internet Assigned Numbers Authority, "IANA IPv4 Special-Purpose Address Registry", <https://www.iana.org/assignments/iana-ipv4-special-registry/iana-ipv4-special-registry.xhtml>.

[RFC0870] Reynolds, J. and J. Postel, "Assigned numbers", RFC 870, DOI 10.17487/RFC0870, October 1983, <https://www.rfc-editor.org/info/rfc870>.

[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>.

[RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. J., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <https://www.rfc-editor.org/info/rfc1918>.

[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>.

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827, May 2000, <https://www.rfc-editor.org/info/rfc2827>.

[RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, DOI 10.17487/RFC3704, March 2004, <https://www.rfc-editor.org/info/rfc3704>.

[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>.

[RFC6890] Cotton, M., Vegoda, L., Bonica, R., Ed., and B. Haberman, "Special-Purpose IP Address Registries", BCP 153, RFC 6890, DOI 10.17487/RFC6890, April 2013, <https://www.rfc-editor.org/info/rfc6890>.

[RFC7050] Savolainen, T., Korhonen, J., and D. Wing, "Discovery of the IPv6 Prefix Used for IPv6 Address Synthesis", RFC 7050, DOI 10.17487/RFC7050, November 2013, <https://www.rfc-editor.org/info/rfc7050>.

8.2. Informative References

[Arista-user-manual]
Arista Networks, "EOS 4.28.1F User Manual", 18 April 2022, <https://www.arista.com/en/um-eos/eos-ipv4#ipv4_routable_240.0.0.0_4>.

[Atlas] RIPE Network Coordination Centre, "RIPE Atlas", <https://atlas.ripe.net/>.

[Cloudflare]
Strong, M., "Fixing reachability to 1.1.1.1, GLOBALLY!", 4 April 2018, <https://blog.cloudflare.com/fixing-reachability-to-1-1-1-1-globally/>.

[Cox] Cartwright-Cox, B. and Z. Pospテュchal, "On "Reclaiming" 240.0.0.0/4", 23 May 2024, <https://ripe88.ripe.net/archives/video/1367/>.

[Dale] Dale, L. and F. Korsbテ、ck, "Dive Deep on AWS Networking Infrastructure (presentation)", NANOG88, June 2023, <https://www.youtube.com/watch?v=0tcR-iQce7s&t=1709s>.

[FLM] Fuller, V., Lear, E., and D. Meyer, "Reclassifying 240/4 as usable unicast address space", Work in Progress, Internet-Draft, draft-fuller-240space-02, 25 March 2008, <https://datatracker.ietf.org/doc/html/draft-fuller-240space-02>.

[GKE] Amarkhed, B. and G. Balakrishnan, "Leveraging Class E IPv4 Address space to mitigate IPv4 exhaustion issues in GKE", 26 September 2024, <https://cloud.google.com/blog/products/containers-kubernetes/how-class-e-addresses-solve-for-ip-address-exhaustion-in-gke/>.

[Huston] Huston, G., "Detecting IP Address Filters", 13 January 2012, <https://labs.ripe.net/author/gih/detecting-ip-address-filters/>.

[JUNOS-release-notes-9.6]
Juniper Networks, "Juniper Networks JUNOS 9.6 Software Release Notes", 1 June 2010,

<https://www.juniper.net/documentation/en_US/junos9.6/information-products/topic-collections/release-notes/9.6/junos-release-notes-9.6.pdf>.

[JUNOS-routing-properties]
Juniper Networks, "JUNOS OS: Protocol-Independent Routing Properties User Guide", 14 March 2022, <https://www.juniper.net/documentation/us/en/software/junos/static-routing/static-routing.pdf>.

[Lear] Lear, E., "Re: Running out of Internet addresses?", TCP-IP mailing list, 27 November 1988, <https://web.archive.org/web/20120514082839/http://www-mice.cs.ucl.ac.uk/multimedia/misc/tcp_ip/8813.mm.www/0146.html>.

[NLNOGRing]
NLNOG RING, "10 Years of NLNOG RING", <https://ring.nlnog.net/post/10-years-of-nlnog-ring/>.

[RFC0919] Mogul, J., "Broadcasting Internet Datagrams", STD 5, RFC 919, DOI 10.17487/RFC0919, October 1984, <https://www.rfc-editor.org/info/rfc919>.

[RFC0951] Croft, W. and J. Gilmore, "Bootstrap Protocol", RFC 951, DOI 10.17487/RFC0951, September 1985, <https://www.rfc-editor.org/info/rfc951>.

[RFC0988] Deering, S., "Host extensions for IP multicasting", RFC 988, DOI 10.17487/RFC0988, July 1986, <https://www.rfc-editor.org/info/rfc988>.

[RFC1256] Deering, S., Ed., "ICMP Router Discovery Messages", RFC 1256, DOI 10.17487/RFC1256, September 1991, <https://www.rfc-editor.org/info/rfc1256>.

[RFC1519] Fuller, V., Li, T., Yu, J., and K. Varadhan, "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy", RFC 1519, DOI 10.17487/RFC1519, September 1993, <https://www.rfc-editor.org/info/rfc1519>.

[RFC1546] Partridge, C., Mendez, T., and W. Milliken, "Host Anycasting Service", RFC 1546, DOI 10.17487/RFC1546, November 1993, <https://www.rfc-editor.org/info/rfc1546>.

[RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, DOI 10.17487/RFC2131, March 1997, <https://www.rfc-editor.org/info/rfc2131>.

[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>.

[RFC3068] Huitema, C., "An Anycast Prefix for 6to4 Relay Routers", RFC 3068, DOI 10.17487/RFC3068, June 2001, <https://www.rfc-editor.org/info/rfc3068>.

[RFC3232] Reynolds, J., Ed., "Assigned Numbers: RFC 1700 is Replaced by an On-line Database", RFC 3232, DOI 10.17487/RFC3232, January 2002, <https://www.rfc-editor.org/info/rfc3232>.

[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>.

[RFC7094] McPherson, D., Oran, D., Thaler, D., and E. Osterweil, "Architectural Considerations of IP Anycast", RFC 7094, DOI 10.17487/RFC7094, January 2014, <https://www.rfc-editor.org/info/rfc7094>.

[RIPE240] Lone, Q., "240/4 As Seen by RIPE Atlas", 23 August 2022, <https://labs.ripe.net/author/qasim-lone/2404-as-seen-by-ripe-atlas/>.

[RIPElabs128016]
Aben, E., "The Curious Case of 128.0/16", 6 December 2011, <https://labs.ripe.net/author/emileaben/the-curious-case-of-128016/>.

[RIPElabs18]
Schwarzinger, F., "Pollution in 1/8", 3 February 2010, <https://labs.ripe.net/author/franz/pollution-in-18/>.

[RIPElabs2a1012]
Aben, E., "The Debogonisation of 2a10::/12", 17 January 2020, <https://labs.ripe.net/author/emileaben/the-debogonisation-of-2a1012/>.

[Sherr] Sherr, M., Cronin, E., Clark, S., and M. Blaze, "Signaling vulnerabilities in wiretapping systems", IEEE Security & Privacy November-December 2005, <https://www.mattblaze.org/papers/wiretap.pdf>.

[VPC] Google Inc., "Virtual Private Cloud: Subnets overview: Valid IPv4 ranges", <https://cloud.google.com/vpc/docs/subnets>.

[WMH] Wilson, P., Michaelson, G., and G. Huston, "Redesignation of 240/4 from "Future Use" to "Private Use"", Work in Progress, Internet-Draft, draft-wilson-class-e-02, 29 September 2008, <https://datatracker.ietf.org/doc/html/draft-wilson-class-e-02>.

Appendix A. Implementation Status

The IPv4 protocol update proposed by this document has already been implemented in a variety of widely-used software platforms. In many cases, implementers were persuaded of the value of the suggestions contained in [FLM] and [WMH].

All known TCP/IP implementations either interoperate properly with packets with sources or destinations in the 240/4 range, or ignore these packets entirely, except FreeBSD and NetBSD, which have support for 240/4 for some purposes while blocking it for others.

A.1. Operating systems

240/4 has been supported for transmitting and receiving ordinary unicast packets in Linux kernels since linux-2.6.25 was released in January 2008. Creating interfaces in the 240/4 range also worked fine using the iproute2 api (as used by the "ip" command) in that release. A kernel patch that allows properly configuring interfaces in the 240/4 range using the busybox ifconfig command was released in linux-4.20 and linux-5.0 in December 2018.

240/4 as unicast was released in Fedora 9 in May 2008, and in Ubuntu 8.10 in October 2008.

240/4 has been supported as ordinary unicast in the Android mobile operating system since Android 1.5 Cupcake (April 2009, using linux-2.6.27).

240/4 has been supported as ordinary unicast in the OpenWRT router OS since OpenWRT 8.09 (September 2008, using linux-2.6.26). A December 2018 kernel patch that allows properly configuring interfaces in the 240/4 range using the ifconfig command was merged into OpenWRT 19.01, along with two other patches to netifd and bcp38 that improve support for 240/4.

240/4 has been supported as ordinary unicast in Apple’s macOS (formerly OS X) operating system and iOS mobile operating system since about 2008.

240/4 has been supported as ordinary unicast in Sun’s Solaris operating system since about 2008.

240/4 traffic has been partly supported in OpenBSD for many years and is substantially fully supported since OpenBSD 7.2 (released October 20, 2022).

240/4 traffic is partly supported for local interface assignment in the FreeBSD operating system. However, ICMP and packet forwarding are not supported by default. Full support for 240/4 addresses is disabled by default, but can be enabled by "sysctl net.inet.ip.allow_net240=1". This option is included in FreeBSD 14.0, released in November 2023, and was available in development versions since July 13, 2022.

We have prepared a patch which enables 240/4 support this on NetBSD. It has not been merged as of December 2023.

240/4 traffic is blocked by default in all versions of the Microsoft Windows operating system. Windows will not assign an interface address in this range, if one is offered by DHCP.

A.2. Routers and Switches

Unless otherwise noted, support in this section reflects interface assignment and packet-forwarding support, not BGP support, which may involve separate bogon logic.

240/4 has been tested to interoperate as ordinary unicast in 2019 in a Cisco router using IOS release-XR 6.5.2.28I, which was also released in 2019. Older and newer releases are also likely to work. Cisco also has two other router OS variants which have not been tested.

240/4 traffic is blocked by default in Juniper’s router operating system, but can be enabled with a simple configuration switch, starting from the JUNOS 9.6 release in June 2010. See page 50 of [JUNOS-release-notes-9.6]. It notes, "The JUNOS Software now allows Class E addresses to be configured on interfaces. To allow Class E addresses to be configured on interfaces, remove the Class E prefix from the list of martian addresses by including the [edit routing-options martians 240/4 orlonger allow] configuration statement." See also chapter 5, "Martian Addresses" on page 129 through 136 of the 2022 documentation [JUNOS-routing-properties]. It includes a completely worked example on "Removing the Class E Prefix on Martian Addresses".

Arista switches running EOS 4.25.2F (from February 2021), and later releases, include the command "ipv4 routable 240.0.0.0/4" which enables the use of 240/4 addresses on interfaces and in packet routing. The default is to disable this ability. [Arista-user-manual]

The Belkin AX3200 router (with firmware 1.0.01 build 101415 Oct 14, 2020) cannot use addresses from 240/4 locally, but is happy to route packets to such addresses elsewhere in the Internet.

A 2024 presentation by Ben Cartwright-Cox [Cox] reported on his experiments on support for 240/4 in several router environments (including BGP). According to Cartwright-Cox, RouterOS 7.7 and IOS XR fully support 240/4. Arista(v)EOS 4.29.0.2F allows opting in to 240/4 support and is believed to work when this is enabled, although it could not be fully tested in Cartwright-Cox's virtualized environment. JunOS 22.X similarly allows opting in to 240/4 support, and can properly use and route addresses in this range statically and dynamically, although its built-in DHCP server is not willing to assign them to other devices on a locally attached network. Finally, EXOS 31.1.1.6-1, IOS XE, Nokia SR-OS, and Huawei VRP generally did not support using or routing 240/4, while the latter three exhibited an apparent bug in which they would ostensibly accept dynamic routes within 240/4 but not actually apply these into the RIB or FIB.

Huawei VRP 5.160 (released 2014) does not allow addresses from the 240/4 range to be manually applied to a router interface.

A.3. DHCP implementations

Support for 240/4 addressing may be typical of many DHCP implementations (because the 240/4 address assignment special case has often been handled at the kernel level). If the underlying operating system supports 240/4 assignment to an interface, the final official ISC DHCP release (4.4.3) supports 240/4 allocation as both client and server, as do Busybox DHCP udhcpc/udhcpd (release 1.1.15), and ISC Kea (which currently includes only a DHCP server implementation).

A.4. Other implementations

Routing of subnets in the 240/4 range is fully supported by the Babel routing protocol and by its main implementation, as of 2020 (or earlier).

Routing of subnets in the 240/4 range is supported by the Gobgp routing daemon, as of release 3.0.0 in 2022-03 (or earlier).

Routing of subnets in the 240/4 range is supported by the BIRD routing daemon, as of release 2.0.10 in 2022-06.

A.5. Internet of Things

Popular embedded Internet-of-Things environments such as RIOT and FreeRTOS already support 240/4 as unicast.

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テ、ht
IPv4 Unicast Extensions Project
Email: dave@none.org