

Internet Engineering Task Force (IETF)
Request for Comments: 6319
Category: Informational
ISSN: 2070-1721

M. Azinger
Frontier Communications
Corporation
L. Vegoda
ICANN
July 2011

Issues Associated with
Designating Additional Private IPv4 Address Space

Abstract

When a private network or internetwork grows very large, it is sometimes not possible to address all interfaces using private IPv4 address space because there are not enough addresses. This document describes the problems faced by those networks, the available options, and the issues involved in assigning a new block of private IPv4 address space.

While this informational document does not make a recommendation for action, it documents the issues surrounding the various options that have been considered.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6319>.

Copyright Notice

Copyright (c) 2011 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 (<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	Large Networks	3
3.	Non-Unique Addresses	3
3.1.	Subscriber Use Network Address Translation	3
3.2.	Carrier-Grade Network Address Translation	4
4.	Available Options	4
4.1.	IPv6 Options	4
4.1.1.	Unique Globally Scoped IPv6 Unicast Addresses	4
4.1.2.	Unique Local IPv6 Unicast Addresses	5
4.2.	IPv4 Options	5
4.2.1.	Address Transfers or Leases from Organizations with Available Address Space	5
4.2.2.	Using Unannounced Address Space Allocated to Another Organization	5
4.2.3.	Unique IPv4 Space Registered by an RIR	6
5.	Options and Consequences for Defining New Private Use Space	6
5.1.	Redefining Existing Unicast Space as Private Address Space	6
5.2.	Unique IPv4 Space Shared by a Group of Operators	7
5.3.	Potential Consequences of Not Redefining Existing Unicast Space as Private Address Space	8
5.4.	Redefining Future Use Space as Unicast Address Space	8
6.	Security Considerations	8
7.	References	9
7.1.	Normative References	9
7.2.	Informative References	9
Appendix A.	Acknowledgments	12

1. Introduction

[RFC1918] sets aside three blocks of IPv4 address space for use in private networks: 192.168.0.0/16, 172.16.0.0/12 and 10.0.0.0/8. These blocks can be used simultaneously in multiple, separately managed networks without registration or coordination with IANA or any Internet registry. Very large networks can find that they need to number more device interfaces than there are available addresses in these three ranges. It has occasionally been suggested that additional private IPv4 address space should be reserved for use by these networks. Although such an action might address some of the needs for these very large network operators, it is not without consequences, particularly as we near the date when the IANA free pool will be fully allocated.

The overall conclusion is that allocating additional address space to be used as private address space has severe problems and would, for instance, impact any software or configuration that has built-in assumptions about private address space. However, it is also well understood that cascading Network Address Translation (NAT) deployments in the existing private address space will cause different types of severe problems when address spaces overlap. At this point, there is no clear agreement of the likelihood of various problems or the respective trade-offs.

2. Large Networks

The main categories of very large networks using private address space are: cable operators, wireless (cell phone) operators, private internets, and VPN service providers. In the case of the first two categories, the complete address space reserved in [RFC1918] tends to be used by a single organization. In the case of private internets and VPN service providers, there are multiple independently managed and operated networks and the difficulty is in avoiding address clashes.

3. Non-Unique Addresses

3.1. Subscriber Use Network Address Translation

The address space set aside in [RFC1918] is a finite resource that can be used to provide limited Internet access via NAT. A discussion of the advantages and disadvantages of NATs is outside the scope of this document, but an analysis of the advantages, disadvantages, and architectural implications can be found in [RFC2993]. Nonetheless, it must be acknowledged that NAT is adequate in some situations and not in others. For instance, it might technically be feasible to use NAT or even multiple layers of NAT within the networks operated by

residential users or corporations where only limited Internet access is required. A more detailed analysis can be found in [RFC3022]. Where true peer-to-peer communication is needed or where services or applications do not work properly behind NAT, globally unique address space is required. In other cases, NAT traversal techniques facilitate peer-to-peer like communication for devices behind NATs.

In many cases, it is possible to use multiple layers of NAT to re-use parts of the address space defined in [RFC1918]. It is not always possible to rely on Customer Premises Equipment (CPE) devices using any particular range, however. In some cases, this means that unorthodox workarounds including assigning CPE devices unallocated address space or address space allocated to other network operators are feasible. In other cases, organizations choose to operate multiple separate routing domains to allow them to re-use the same private address ranges in multiple contexts. One consequence of this is the added complexity involved in identifying which system is referred to when an IP address is identified in a log or management system.

3.2. Carrier-Grade Network Address Translation

Another option is to share one address across multiple interfaces and in some cases, subscribers. This model breaks the classical model used for logging address assignments and creates significant risks and additional burdens, as described in [CLAYTON] and more fully discussed in [FORD], and as documented in [DS-LITE].

4. Available Options

When a network operator has exhausted the private address space set aside in [RFC1918] but needs to continue operating a single routing domain, a number of options are available. These are described in the following sections.

4.1. IPv6 Options

4.1.1. Unique Globally Scoped IPv6 Unicast Addresses

Using unique, globally scoped IPv6 unicast addresses is the best permanent solution as it removes any concerns about address scarcity within the next few decades. Implementing IPv6 is a major endeavor for service providers with millions of consumers and is likely to take considerable effort and time. In some cases, implementing a new network protocol on a very large network takes more time than is available, based on network growth and the proportion of private space that has already been used. In these cases, there is a call

for additional private address space that can be shared by all network operators. [DAVIES] makes one such case.

4.1.2. Unique Local IPv6 Unicast Addresses

Using the unique, local IPv6 unicast addresses defined in [RFC4193] is another approach and does not require coordination with an Internet registry. Although the addresses defined in [RFC4193] are probabilistically unique, network operators on private internets and those providing VPN services might not want to use them because there is a very low probability of non-unique locally assigned global IDs being generated by the algorithm. Also, in the case of private internets, it can be very challenging to coordinate the introduction of a new network protocol to support the internet's continued growth.

4.2. IPv4 Options

4.2.1. Address Transfers or Leases from Organizations with Available Address Space

The Regional Internet Registry (RIR) communities have recently been developing policies to allow organizations with available address space to transfer such designated space to other organizations [RIR-POLICY]. In other cases, leases might be arranged. This approach is only viable for operators of very large networks if enough address space is made available for transfer or lease and if the very large networks are able to pay the costs of these transfers. It is not possible to know how much address space will become available in this way, when it will be available, and how much it will cost. However, it is unlikely to become available in large contiguous blocks, and this would add to the network management burden for the operator as a significant number of small prefixes would inflate the size of the operators routing table at a time when it is also adding an IPv6 routing table. These reasons will make address transfers a less attractive proposition to many large network operators. Leases might not be attractive to some organizations if both parties cannot agree to a suitable length of time. Also, the lessor might worry about its own unanticipated needs for additional IPv4 address space.

4.2.2. Using Unannounced Address Space Allocated to Another Organization

Some network operators have considered using IP address space that is allocated to another organization but is not publicly visible in BGP routing tables. This option is very strongly discouraged as the fact that an address block is not visible from one view does not mean that it is not visible from another. Furthermore, address usage tends to

leak beyond private network borders in e-mail headers, DNS queries, traceroute output and other ways. The ambiguity this causes is problematic for multiple organizations. This issue is discussed in [RFC3879], Section 2.3.

It is also possible that the registrant of the address block might want to increase its visibility to other networks in the future, causing problems for anyone using it unofficially. In some cases, there might also be legal risks involved in using address space officially allocated to another organization.

Where this has happened in the past, it has caused operational problems [FASTWEB].

4.2.3. Unique IPv4 Space Registered by an RIR

RIRs' policies allow network operators to receive unique IP addresses for use on internal networks. Further, network operators are not required to have already exhausted the private address space set aside in [RFC1918]. Nonetheless, network operators are naturally disinclined to request unique IPv4 addresses for the private areas of their networks, as using addresses in this way means they are not available for use by new Internet user connections.

It is likely to become more difficult for network operators to obtain large blocks of unique address space as we approach the point where all IPv4 unicast /8s have been allocated. Several RIRs already have policies about how to allocate from their last /8 [RIR-POLICY-FINAL-8], and there have been policy discussions that would reduce the maximum allocation size available to network operators [MAX-ALLOC] or would reduce the period of need for which the RIR can allocate [SHORTER-PERIODS].

5. Options and Consequences for Defining New Private Use Space

5.1. Redefining Existing Unicast Space as Private Address Space

It is possible to re-designate a portion of the current global unicast IPv4 address space as private unicast address space. Doing this could benefit a number of operators of large networks for the short period before they complete their IPv6 roll-out. However, this benefit incurs a cost by reducing the pool of global unicast addresses available to users in general.

When discussing re-designating a portion of the current global unicast IPv4 address space as private unicast address space, it is important to consider how much space would be used and for how long it would be sufficient. Not all of the large networks making full

use of the space defined in [RFC1918] would have their needs met with a single /8. In 2005, [HAIN] suggested reserving three /8s for this purpose, while in 2009 [DAVIES] suggested a single /10 would be sufficient. There does not seem to be a consensus for a particular prefix length nor an agreed basis for deciding what is sufficient. The problem is exacerbated by the continually changing needs of ever expanding networks.

A further consideration is which of the currently unallocated IPv4 unicast /8 blocks should be used for this purpose. Using address space that is known to be used unofficially is tempting. For instance, 1.0.0.0/8, which was unallocated until January 2010, was proposed in [HAIN] and is known to be used by a number of different users. These include networks making use of HIP LSIs [RFC4423], [WIANA], [anoNet], and others. There is anecdotal [VEGODA] and research [WESSELS] evidence to suggest that several other IPv4 /8s are used in this fashion. Also there have been discussions [NANOG] about some sections of these /8's being carved out and filtered, therefore unofficially enabling the use of these sections for private use.

Although new IPv4 /8s are allocated approximately once a month, they are not easy to bring into use because network operators are slow to change their filter configurations. This is despite long-running awareness campaigns [CYMRU] [LEWIS] and active work [ripe-351] to notify people whose filters are not changed in a timely fashion. Updating code that recognizes private address space in deployed software and infrastructure systems is likely to be far more difficult as many systems have these ranges hard-coded and cannot be quickly changed with a new configuration file.

Another consideration when redefining existing unicast space as private address space is that no single class of user can expect the space to stay unique to them. This means that an ISP using a new private address range cannot expect its customers not to already be using that address range within their own networks.

5.2. Unique IPv4 Space Shared by a Group of Operators

Where a group of networks find themselves in a position where they each need a large amount of IPv4 address space from an RIR in addition to that defined in [RFC1918], they might cooperatively agree to all use the same address space to number their networks. The clear benefit to this approach is that it significantly reduces the potential demand on the pool of unallocated IPv4 address space. However, the issues discussed in Sections 4.2.2 and 5.3 are of concern here, particularly the possibility that one operator might

decide to use the address space to number customer connections, rather than private infrastructure.

Nonetheless, this approach has the potential to create an unofficial new private address range without proper scrutiny.

5.3. Potential Consequences of Not Redefining Existing Unicast Space as Private Address Space

If additional private address space is not defined and the large network operators affected by this problem are not able to solve their problems with IPv6 address space or by segmenting their networks into multiple routing domains, those networks will need unique IPv4 addresses. It is possible and even likely that a single network could consume a whole IPv4 /8 in a year. At the time this document is being written, there are just 24 unallocated IPv4 /8s, so it would not take many such requests to make a major dent in the available IPv4 address space. [POTAROO] provides an analysis of IPv4 address consumption and projects the date on which the IANA and RIR pools will be fully allocated.

5.4. Redefining Future Use Space as Unicast Address Space

There have also been proposals to re-designate the former Class E space (240.0.0.0/4) as unicast address space. [WILSON] suggests that it should be privately scoped while [FULLER] does not propose a scope. Both proposals note that existing deployed equipment may not be able to use addresses from 240.0.0.0/4. Potential users would need to be sure of the status of the equipment on their network and the networks with which they intend to communicate.

It is not immediately clear how useful 240.0.0.0/4 could be in practice. While [FULLER] documents the status of several popular desktop and server operating systems, the status of the most widely deployed routers and switches is less clear, and it is possible that 240.0.0.0/4 might only be useful in very large, new green field deployments where full control of all deployed systems is available. However, in such cases it might well be easier to deploy an IPv6 network.

6. Security Considerations

This document has no security implications.

7. References

7.1. Normative References

- [RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
- [RFC2993] Hain, T., "Architectural Implications of NAT", RFC 2993, November 2000.
- [RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", RFC 3022, January 2001.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.

7.2. Informative References

- [CLAYTON] Clayton, R., "Practical mobile Internet access traceability", January 2010, <<http://www.lightbluetouchpaper.org/2010/01/13/practical-mobile-internet-access-traceability/>>.
- [CYMRU] Greene, B., "The Bogon Reference", <<http://www.team-cymru.org/Services/Bogons/>>.
- [DAVIES] Davies, G. and C. Liljenstolpe, "Transitional non-conflicting reusable IPv4 address block", Work in Progress, November 2009.
- [DS-LITE] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", Work in Progress, August 2010.
- [FASTWEB] Aina, A., "41/8 announcement", May 2006, <<http://www.afnog.org/archives/2006-May/002117.html>>.
- [FORD] Ford, M., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", Work in Progress, March 2010.

- [FULLER] Fuller, V., Lear, E., and D. Meyer, "Reclassifying 240/4 as usable unicast address space", Work in Progress, March 2008.
- [HAIN] Hain, T., "Expanded Address Allocation for Private Internets", Work in Progress, January 2005.
- [LEWIS] Lewis, J., "This system has been setup for testing purposes for 69/8 address space", March 2003, <<http://69box.atlantic.net/>>.
- [MAX-ALLOC] Spenceley, J. and J. Martin, "prop-070: Maximum IPv4 allocation size", January 2009, <<http://www.apnic.net/policy/proposals/prop-070>>.
- [NANOG] Dickson, B., "1/8 and 27/8 allocated to APNIC", January 2010, <<http://mailman.nanog.org/pipermail/nanog/2010-January/017451.html>>.
- [POTAROO] Huston, G., "IPv4 Address Report", <<http://www.potaroo.net/tools/ipv4/index.html>>.
- [RFC3879] Huitema, C. and B. Carpenter, "Deprecating Site Local Addresses", RFC 3879, September 2004.
- [RFC4423] Moskowitz, R. and P. Nikander, "Host Identity Protocol (HIP) Architecture", RFC 4423, May 2006.
- [RIR-POLICY] Number Resource Organization, "RIR Comparative Policy Overview, October 2009, Section 1.3.2 Transfer of Custodianship", <<http://www.nro.net/rir-comparative-policy-overview/rir-comparative-policy-overview-2009-03#1-3-2>>.
- [RIR-POLICY-FINAL-8] Number Resource Organization, "RIR Comparative Policy Overview, October 2009, 2.6. Use of Final Unallocated IPv4 Address Space", October 2009, <<http://www.nro.net/rir-comparative-policy-overview/rir-comparative-policy-overview-2009-03>>.
- [SHORTER-PERIODS] Karrenberg, D., O'Reilly, N., Titley, N., and R. Bush, "RIPE Policy Proposal 2009-03", April 2009, <<http://www.ripe.net/ripe/policies/proposals/2009-03>>.

- [VEGODA] Vegoda, L., "Awkward /8 Assignments", September 2007, <http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_10-3/103_awkward.html>.
- [WESSELS] Wessels, D., "Searching for Evidence of Unallocated Address Space Usage in DITL 2008 Data", June 2008, <<https://www.dns-oarc.net/files/dnsops-2008/Wessels-Unused-space.pdf>>.
- [WIANA] WIANA, "The Wireless Internet Assigned Numbers Authority", <<http://www.wiana.org/>>.
- [WILSON] Wilson, P., Michaelson, G., and G. Huston, "Redesignation of 240/4 from "Future Use" to "Private Use"", Work in Progress, September 2008.
- [anoNet] anoNet, "anoNet: Cooperative Chaos".
- [ripe-351] Karrenberg, D., "De-Bogonising New Address Blocks", October 2005, <<http://www.ripe.net/ripe/docs/ripe-351>>.

Appendix A. Acknowledgments

The authors would like to thank Ron Bonica, Michelle Cotton, Lee Howard, and Barbara Roseman for their assistance in early discussions of this document and to Maria Blackmore, Alex Bligh, Mat Ford, Thomas Narten, and Ricardo Patara for suggested improvements.

Authors' Addresses

Marla Azinger
Frontier Communications Corporation
Vancouver, WA
United States of America

EEmail: marla.azinger@ftr.com
URI: <http://www.frontiercorp.com/>

Leo Vegoda
Internet Corporation for Assigned Names and Numbers
4676 Admiralty Way, Suite 330
Marina del Rey, CA 90292
United States of America

Phone: +1-310-823-9358
EEmail: leo.vegoda@icann.org
URI: <http://www.iana.org/>