Workgroup: Internet Engineering Task Force
Internet-Draft:
draft-schoen-intarea-unicast-127-01
Updates: 1122, 1812, 2827, 3704 (if approved)
Published: 7 March 2022
Intended Status: Standards Track
Expires: 8 September 2022
Authors: S.D. Schoen
IPv4 Unicast Extensions Project
J. Gilmore
IPv4 Unicast Extensions Project
D. Täht
IPv4 Unicast Extensions Project

Unicast Use of the Formerly Reserved 127/8

Abstract

This document redefines the IPv4 local loopback network as consisting only of the 65,536 addresses 127.0.0.0 to 127.0.255.255 (127.0.0.0/16). It asks implementers to make addresses in the prior loopback range 127.1.0.0 to 127.255.255.255 fully usable for unicast use on the Internet.

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 8 September 2022.

Copyright Notice

Table of Contents

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 allow the unicast use of more than 16 million historically reserved addresses in the middle of the IPv4 address space.

This document provides history and rationale to reduce the size of the IPv4 local loopback network ("localnet") from /8 to /16, freeing up over 16 million IPv4 addresses for other possible uses.

When all of 127.0.0.0/8 was reserved for loopback addressing, IPv4 addresses were not yet recognized as scarce. Today, there is no justification for allocating 1/256 of all IPv4 addresses for this purpose, when only one of these addresses is commonly used and only a handful are regularly used at all. Unreserving the majority of these addresses provides a large number of additional IPv4 host addresses for possible use, 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

The IPv4 network 127/8 was first reserved by Jon Postel in 1981 [RFC0776]. Postel's policy was to reserve the first and last network of each class, and it does not appear that he had a specific plan for how to use 127/8. Apparently, the first operating systems to support a loopback interface as we understand it today were experimental Berkeley Unix releases by Bill Joy and Sam Leffler at the University of California at Berkeley. The choice of 127.0.0.1 as loopback address was made in 1983 by Joy and Leffler in the code base that was eventually released as 4.2BSD. Their earliest experimental code bases used 254.0.0.0 and 127.0.0.0 as loopback addresses. Three years later, Postel and Joyce Reynolds documented the loopback function in November 1986 [RFC0990], and it was codified as a requirement for all Internet hosts three years after that, in [RFC1122]. The substantive interpretation of these addresses has remained unchanged since RFC 990 indicated that the

> network number 127 is assigned the "loopback" function, that is, a datagram sent by a higher level protocol to a network 127 address should loop back inside the host. No datagram "sent" to a network 127 address should ever appear on any network anywhere.

Many decisions about IPv4 addressing contemporaneous with this one underscore the lack of concern about address scarcity. It was common in the early 1980s to allocate an entire /8 to an individual university, company, government agency, or even a research project.

By contrast, IPv6, despite its vastly larger pool of available address space, allocates only a single local loopback address (::1) [RFC4291]. This appears to be an architectural vote of confidence in the idea that Internet protocols ultimately do not require millions of distinct loopback addresses.

Most applications use only the single loopback address 127.0.0.1 ("localhost") for IPv4 loopback purposes, although there are exceptions. For example, the systemd-resolved service on Linux provides a stub DNS resolver at 127.0.0.53.

In theory, having multiple local loopback addresses might be useful for increasing the number of distinct IPv4 sockets that can be used for inter-process communication within a host. The local loopback /16 network retained by this document will still permit billions of distinct concurrent loopback TCP connections within a single host, even if both the IP address and port number of one endpoint of each connection are fixed.

3. Change in Status of Addresses Within 127/8

The purpose of this document is to reduce the size of the special-case reservation of 127/8, so that only 127.0/16 is reserved as the local loopback network.

Other IPv4 addresses whose first octet is 127 (that is, the addresses 127.1.0.0 to 127.255.255.255) are no longer reserved and are now available for general Internet unicast use, treated identically to other IPv4 addresses, and subject to potential future allocation.

All host and router software SHOULD treat 127.1.0.0 to 127.255.255.255 as a global unicast address range.

Clients for autoconfiguration mechanisms such as DHCP [RFC2131] SHOULD accept a lease or assignment of addresses within 127.1/16 to 127.255/16 whenever the underlying operating system is capable of accepting it. Servers for these mechanisms SHOULD assign this address when so configured.

4. Compatibility and Interoperability

Many deployed systems follow older Internet standards in rejecting externally-originating packets from addresses in 127/8, and in not generating packets addressed to them). RFC 3704 [RFC3704] (BCP 84) cites RFC 2827 [RFC2827] (BCP 38) to this effect:

> 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 127/8 address space is no longer reserved as a whole, an address within this space, other than those within 127/16, 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 127/8 (other than those in 127/16). Maintainers of lists of "Martian addresses" MUST NOT designate addresses from the 127/8 range (other than those within 127/16) as "Martian".

The filtering recommended by RFC 3704 is designed for border routers, not for hosts. To the extent that an ISP had validly allocated an address range from within 127/8 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 127/8 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 127/8 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 a portion of 127/8. It therefore requests IANA to update the IPv4 Special-Purpose Address registry [IANA4SP] by replacing the entry for 127.0.0.0/8 with 127.0.0.0/16, with authority of this document.

IANA is also requested to update the IPv4 Address Space Registry [IANA4] by changing the entry for 127/8 (IANA - Loopback) to read 127/16, and by adding a new entry 127.1/16-127.255/16 Unallocated [Date of this document] [blank] [blank] UNALLOCATED

Finally, IANA is requested to prepare for this address space to be addressed in 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 127.1.0.0 through 127.255.255.255. 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.

6. Security Considerations

The behavior change specified by this document could produce security concerns where two devices, or two different parts of the software on a host, or a software application and a human user, follow divergent interpretations of an address that was formerly a loopback address.

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. It could also lead to an application connecting to the local host when it expected to connect to a remote host, or vice versa.

One undesired case would arise where a local process on a host accepts connections on what it believes is a loopback address, in order to receive commands from other software on the same host, yet the bound address is actually reachable from outside that host. The traditional socket API present on most operating systems does not make this especially likely, since a listening process typically binds to either INADDR_ANY (which includes both loopback and nonloopback interfaces) or INADDR_LOOPBACK (which includes only the single address 127.0.0.1). The existence of an additional interface with a remotely addressable unicast address like 127.8.9.10 would not, in itself, change which hosts can communicate with either of these sockets. Nonetheless, an operating system or software library that provides some other interface with its own means of scoping the receipt of incoming connections must take care not to leave an ambiguity between host-only and non-host-only address scopes as a result of the change specified by this document.

The importance of the distinction just mentioned is underscored by practical examples of vulnerabilities when specific software relaxed the distinction between loopback and non-loopback addresses in a

different way. A 2017 vulnerability [CVE-2016-1551] related to the reference implementation of the Network Time Protocol v4 [RFC5905], and an analogous 2020 vulnerability [CVE-2020-8558] in the Kubernetes cluster management software, both involved the use of a Linux kernel option that removed the prohibition on sending or receiving packets over the wire with a 127/8 destination address. This, however, allowed other devices to reach and communicate with server processes that had deliberately listened on what they otherwise expected to be loopback addresses.

The change requested by this document does not have the same effect, because loopback addresses in the reduced 127.0/16 loopback range are still not permitted to appear on the wire, and should still be rejected by implementations. The ability to enforce the inaccessibility of loopback addresses by other hosts remains necessary for security. In particular, treating all of 127/8 as globally routable address space is not a safe behavior. Operating systems SHOULD continue to treat 127.0/16 as loopback-only and never route packets between 127.0/16 loopback addresses and any other interface. Addresses in 127.0/16 still SHOULD NOT appear on any network link and SHOULD NOT be accepted or generated over a network link. Applications MUST NOT use 127.1/16 to 127.255/16 for loopback purposes or assume that connections from these addresses necessarily originated from software on the local host.

Apart from that, firewall rules that assume that 127.1/16 through 127.255/16 are unroutable and/or local SHOULD be updated to take into account that they may be routable and/or non-local.

Software that assumes that all of 127/8, either as a source or a destination, refers to the local host SHOULD be updated to make that inference only for 127/16. Communications to or from 127.1/16 through 127.255/16 SHOULD NOT be treated as inherently more trusted than communications to or from the public Internet as a whole.

7. Acknowledgements

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

Members of the Internet History Mailing List helped us clarify the early history of 127/8.

8. References

8.1. Normative References

[IANA4] Internet Assigned Numbers Authority, "IANA IPv4 Address Space Registry", <https://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xhtml>.

[IANA4SP] Internet Assigned Numbers Authority, "IANA IPv4 Special-Purpose Address Registry", <https://www.iana.org/assignments/iana-ipv4-special-registry/iana-ipv4-special-registry.xhtml>.

[RFC0776] Postel, J., "Assigned numbers", RFC 776, DOI 10.17487/RFC0776, January 1981, <https://www.rfc-editor.org/info/rfc776>.

[RFC0990] Reynolds, J. and J. Postel, "Assigned numbers", RFC 990, DOI 10.17487/RFC0990, November 1986, <https://www.rfc-editor.org/info/rfc990>.

[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>.

[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>.

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

[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>.

[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>.

[RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010, <https://www.rfc-editor.org/info/rfc5905>.

8.2. Informative References

[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/>.

[CVE-2016-1551] NIST National Vulnerability Database, "CVE-2016-1551", January 2017, <https://nvd.nist.gov/vuln/detail/CVE-2016-1551>.

[CVE-2020-8558] NIST National Vulnerability Database, "CVE-2020-8558", July 2020, <https://nvd.nist.gov/vuln/detail/CVE-2020-8558>.

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

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

[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/>.

Appendix A. Implementation Status

To our knowledge, the behavior specified by this document is not currently the default in any TCP/IP implementation. We have prepared and tested small patches to the Linux and FreeBSD kernels, and achieved interoperability between patched versions of these systems when numbered with 127/8 addresses. The patched systems were otherwise usable normally.

The behavior of our patches contrasts with that of the existing route_localnet option in Linux. The route_localnet option is a Linux kernel feature which a user can enable in order to make all of 127/8 simultaneously addressable in both host and network address scopes, which, as described in the Security Considerations section, has had undesirable security consequences. Our patches instead retain

127.0/16 an exclusive loopback address range, continuing to forbid it from appearing on the wire at all.

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. Täht
IPv4 Unicast Extensions Project
Half Moon Bay, CA
United States of America

Email: dave@taht.net