Internet Engineering Task Force
Internet-Draft
Updates: 4291,5156 (if approved)
Intended status: Standards Track
Expires: January 31, 2013

M. Smith
IMOT

July 30, 2012

A Larger Loopback Prefix for IPv6
draft-smith-v6ops-larger-ipv6-loopback-prefix-00

Abstract

In IPv4, 127/8 is the loopback prefix, where as in IPv6 it is ::1/128. The significant difference between these two prefixes is the number of addresses they cover; 127/8 covers 2^24 or 16 777 216 addresses, where as ::1/128 covers just a single address.

IPv4's large number of loopback addresses has facilitated some novel uses of the loopback function that cannot be achieved with the single loopback address available in IPv6. This memo proposes a new larger loopback prefix for IPv6 so that these uses of the loopback function become available for IPv6.

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 http://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 January 31, 2013.

Table of Contents

1. Introduction

In IPv4, 127/8 is the internal host loopback prefix [RFC1122]. In IPv6, ::1/128 is the internal node loopback prefix [RFC4291]. Packets sent from addresses within these prefixes are not to leave the node, and packets destined to addresses from within these prefixes are to be returned internally within the originating node for local processing.

The significant difference between the IPv4 127/8 loopback prefix and the IPv6 ::1/128 loopback prefix is the number of addresses they each cover; 127/8 covers 2^24 or 16 777 216 addresses, where as ::1/128 covers just a single address.

The large amount of address space covered by 127/8 has facilitated some novel uses of the loopback function, which concurrently utilise multiple loopback addresses. These loopback function uses are not possible with the IPv6 ::1/128 loopback prefix.

The IPv4-Mapped IPv6 Address form of 127/8, ::ffff:127.0.0.0/104 [RFC4291], could be used for these loopback uses under IPv6. However, /104 is not a prefix length commonly used in native IPv6 addressing, and may create unacceptable constraints on these loopback uses when applied to IPv6. For example, 64 bit interface identifers [RFC4291] cannot be used within ::ffff:127.0.0.0/104.

This memo proposes a larger IPv6 loopback prefix to overcome the constraints of ::1/128.

The memo starts by describing some use cases of the loopback function that are currently not possible with ::1/128. Following these uses, the requirements a larger IPv6 loopback prefix should attempt to meet are provided. A proposed new larger IPv6 loopback prefix that meets the majority of these requirements is then specified. For this new larger IPv6 loopback prefix, the host and router processing rules for packets containing addresses within this larger IPv6 loopback prefix are described. Finally, relevant security considerations are discussed.

This memo, if published, updates [RFC4291] and [RFC5156].

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. Example Loopback Prefix Uses Not Possible With ::1/128

The following are examples of loopback function use under IPv4, facilitated by 127/8, that are not achievable under IPv6 with ::1/128.

2.1. Multiple Application Instances Listening On The Same Transport Layer Protocol Port

During network application development and testing, it can be useful to run multiple instances of the same application concurrently. Additionally, it can be useful to avoid these networked application instances being reachable via the host's external interfaces, for security and other reasons.

Network applications that use a well known transport layer protocol port will normally listen on that port for all addresses available on the host. Consequently, attempting to run another instance of the application will cause the second instance to fail, as the listening port is already in use.

This port reuse limitation can be overcome by either having each application instance listen on a different transport layer protocol port. Alternatively, each application instance could use the same well known transport layer port, bound to different and distinct addresses available on the host. The latter method can be more convenient, as networked applications tend to be written to more easily communicate with processes listening on different addresses, typically expressed as DNS names.

127/8 has provided many IPv4 addresses that can be used on a host to run multiple instances of applications listening on the same transport layer protocol port, while also preventing those application instances from being reachable through the host's external interfaces.

2.2. ntpd Reference Clock Device Drivers

ntpd [NTPD] supports local reference clocks, in addition to network located time sources. A local reference clock is typically a hardware device attached directly to the host, receiving time information from an external source, such as a satellite.

To simplify time source configuration, comparison and selection, local reference clocks are represented as though they were network attached time sources, by using addresses that fall within 127/8. These 127/8 addresses have the form 127.127.t.u, where 't' represents the reference clock driver type, and 'u' represents the reference

clock driver unit or instance. [NTPD-RCD]

2.3. Multiple Loopback Interfaces As Routing Next Hops

Some IPv4 router implementations allow multiple virtual link layer loopback interfaces to be created. These virtual interfaces can then be assigned different prefixes and addresses from within 127/8, with non-locally assigned addresses within these prefixes used as next hop addresses for static or dynamic IPv4 routes. The state of these routes will then depend on the administrative state of the corresponding virtual interface. Administratively enabling or disabling a virtual interface will add or remove the corresponding routes from the IPv4 route table. This can be useful for testing routing convergence performance or the simulation of networks with large numbers of routes.

3. Larger IPv6 Loopback Prefix Requirements

A new larger IPv6 loopback prefix should attempt to satisfy all of the following requirements.

3.1. Well Known Prefix

A new larger IPv6 loopback prefix should have a single value and be well known. This allows it to be automatically configured on hosts and routers upon system initialisation. For currently deployed hosts and routers that aren't aware of a new loopback prefix, a single well known prefix will simplify the configuration of an additional loopback prefix static route, additional loopback prefixes and addresses on a loopback virtual interfaces, and configuration of packet filters or firewall rules.

3.2. Within An Existing Special Purpose IPv6 Prefix

A new larger IPv6 loopback prefix should be located within an existing aggregate IPv6 prefix already in use for similar well known special use prefixes.

3.3. Easy For A Human To Use

As previously described, one use of a larger IPv6 loopback prefix will be during development and testing of new network applications. In this and similar cases, it is likely that a human will be entering the prefix or addresses that fall within it quite regularly. Therefore, the new larger IPv6 loopback prefix should be both easy to remember and easy to enter into a computer system.

A special use IPv6 prefix that is easy use has the following characteristics:

- o is numerically similar or significantly different to other special use IPv6 prefixes that serve similar purposes, as this assists with remembering the prefix

- o is as short possible, once leading zeros have been supressed and strings of zeros have been suppressed using "::", assisting with both remembering the prefix and the accurate entry of the prefix into a computing system

3.4. Covers the Existing IPv6 Loopback Prefix

The new larger IPv6 loopback prefix should cover the existing ::1/128 IPv6 loopback prefix, as a single loopback prefix suiting all loopback uses is less complex than two loopback prefixes, with one providing a single address, and the other providing many addresses.

3.5. Supports 64 bit Interface Identifiers

The network applications being developed and tested using the loopback prefix may perform IPv6 addressing related functions. To simulate native IPv6 addressing, the new larger IPv6 loopback prefix should accomodate 64 bit interface identifers.

3.6. Supports Multiple Subnets

New network applications may also perform IPv6 subnet related functions, or need to be tested with multiple IPv6 addresses from different subnets, including situations such as phasing in and phasing out of subnets via the preferred and valid lifetime mechanism. The new larger IPv6 loopback prefix should support multiple IPv6 subnets, typically /64s, with the number of supported subnets being large enough for most if not all conceivable uses.

4. Proposed Larger IPv6 Loopback Prefix

The proposed larger IPv6 loopback prefix is:

0001:0000:0000:0000:0000:0000:0000:0000/48

This prefix meets all but one of the previously described requirements; the exception being that it does not cover the existing ::1/128 loopback prefix. Specifically,

- it would be a well known prefix
- it would fall within an existing IPv6 prefix used for similar special purpose prefixes (0::/8)
- it would be easy for a human to use, as concisely it is 1::/48
- it would support 64 bit interface identifiers
- and would provide 2^16 or 65536 /64 subnets

It is not possible to meet the requirement that ::1/128 falls within the new larger loopback prefix, if the prefix length of the new prefix is /48. This would result in a new larger loopback prefix of ::/48, which would also cover the IPv4 mapped IPv6 address prefix, ::ffff:0.0.0.0/104. It would not be acceptable to locally loop traffic destined to these addresses, as it would prevent their use as described in [RFC4038]. A compromise of excluding the IPv4 mapped IPv6 address prefix from the loopback function applied to ::/48 is not feasible, as that would prevent a range of 64 bit IPv6 interface identifier values from being available within all loopback prefix subnets.

Some Internet Protocol implementations represent or perform the loopback function using a virtual link layer interface, commonly known as the "loopback" interface. Conceptually, or in actuallity, the node's operating system transmits packets out the virtual interface, and then the loopback interface device driver returns the packet to the host as though it had been received by the loopback interface. The packets are then processed by the local network layer protocol implementation.

Although all addresses within 127/8 are considered assigned to the host, it is common to have the individual address 127.0.0.1/8 automatically configured on the loopback interface during system initialisation, and to show this address when querying the loopback interface for assigned IPv4 addresses. This is for operational convenience rather than necessity. Similarly, ::1/128 is also typically automatically configured on the loopback interface.

For the 1::/48 loopback prefix, the address automatically configured on the loopback interface should be:

1::1/64

The implementation will still consider all addresses within 1::/48 to be locally assigned, such that removal of 1::1/64 from the loopback interface by a system administrator will not change the loopback

behaviour for 1::1/64, or any other address within 1::/48.

5. 1::/48 Processing Rules

The following processing rules apply to packets containing ::1/48 source and/or destination addresses.

5.1. Host Rules

The following rules apply to IPv6 hosts.

5.1.1. Packets Sent with 1::/48 Source and/or 1::/48 Destination Addresses

Packets with 1::/48 source and/or destination addresses MUST be returned to the host for processing by the local IPv6 protocol stack. They MUST NOT be sent over any external links attached to the host.

Processing of the locally returned packers is to occur as though they originated externally and had entered the host via a link layer interface. Standard incoming IPv6 packet processing occurs, which may include generating appropriate ICMPv6 error messages. For example, for an IPv6 packet with a 1::/48 source address, and a unicast destination address that is not assigned to the host, an ICMPv6 Destination Unreachable, Address Unreachable is likely to be generated. This ICMPv6 error message would be returned locally to the host for further processing, as it will have a 1::/48 destination address. (ICMPv6 error messages cannot be generated in response to received ICMPv6 error messages, preventing an endless loop of ICMPv6 error messages in this situation.)

In addition to unicast IPv6 addresses assigned to interfaces via other means, all destinations within 1::/48 MUST be considered assigned to the host.

5.1.2. Packets Received Externally With 1::/48 Source and/or Destination Addresses

Packets with 1::/48 source and/or destination addresses received over any of the external links attached to the host MUST be dropped. ICMPv6 error messages, such as Destination Unreachable messages, MUST NOT be generated for these dropped packets.

For these dropped packets, it may be useful to generate an appropriate system log message, indicating a packet with an invalid source or destination address (a "martian") was received over an external interface. By default, these messages should be suppressed.

If they are enabled, they should be appropriately rate limited, with the rate limit being able to be set by a system administrator. An appropriate rate limiting mechanism could be the one suggested for ICMPv6 messages, described in section 2.4, (f) of [RFC4443].

5.2. Router Rules

IPv4 loopback packet processing rules for routers, specified in [RFC1812], by default, prohibited forwarding of packets with 127/8 destinations, other than those originated locally by and returned back to the router itself. However, a software switch could be provided to disable this prohibition. This special case of allowing forwarding of packets towards 127/8 destinations has been taken advantage of by [RFC4379]. An equivalent function for IPv6 is provided by using the IPv4 mapped IPv6 prefix of ::ffff:127.0.0.0/104.

The existing loopback IPv6 packet processing rules for routers is the same as for IPv6 hosts; traffic towards ::1/128, not originated locally, and must not be forwarded by a router [RFC4291].

For the new 1::/48 loopback prefix, the IPv6 router processing rules are modified to match those of IPv4.

5.2.1. Packets Sent with 1::/48 Source and/or 1::/48 Destination Addresses

By default, an IPv6 router MUST follow the host processing rules, described previously, for packets sent with 1::/48 source and/or destination addresses. In summary, IPv6 packets with 1::/48 source and/or 1::/48 destination addresses are not to leave the router, and are to be looped back to the router for host oriented processing.

A software switch may be provided to permit packets with 1::/48 source and/or destination addresses to be sent via an external interface, to facilitate uses of 1::/48 similar to those described in [RFC4379]. If provided, this software switch MUST default to off.

5.2.2. Packets Received Externally With 1::/48 Source and/or Destination Addresses

By default, an IPv6 router must follow the host processing rules, described previously, for packets received externally with 1::/48 source and/or destination addresses. In summary, IPv6 packets with 1::/48 source and/or 1::/48 destination addresses are to be dropped, and ICMPv6 error messages are not to be generated in response. The ability to log reception of these types of packets could be provided, however, by default, they must not be logged.

A software switch may be provided to permit packets with 1::/48 source and/or destination addresses to be forwarded via an external interface, to facilitate uses of 1::/48 similar to those described in [RFC4379]. This software switch MUST default to off.

6. Acknowledgements

The following people provided useful comments on this memo:

7. IANA Considerations

IANA is requested to allocate 1::/48 from within 0::/8 of the Internet Protocol Version 6 Address Space, for use as a larger loopback prefix for IPv6 as described in this memo.

8. Security Considerations

Today, 1::/48 is an unallocated prefix. Traffic with source and/or destination addresses that fall within 1::/48 will be processed by hosts and routers using conventional unicast packet functions, rather than the processing rules specified in this memo. These types of hosts and routers will be described as "legacy" in this section. The result of this conventional unicast processing is that packets that are intended and expected to be looped back locally within the origin node may leave the legacy node ("leak") via an externally attached interface, and subsequently may be forwarded through the local routing domain, towards the global public Internet. This may disclose information to unauthorised parties, and therefore may have unacceptable security consequences, depending on local security policy.

A legacy node that does not have any external network attachments, while not looping packets for local processing, will inherently keep packets local to the node either by locally processing them or dropping them, eliminating the security implications of packets leaving the node.

For legacy nodes with externally attached interfaces, the following classes of packets will be forwarded by conventional unicast processing, contrary to the ruls specified in this memo:

1. non-1::/48 source, 1::/48 destination

2. 1::/48 source, non-1::/48 destination

3. 1::/48 source, 1::/48 destination

Packet filters (also commonly known as Access Control Lists, or ACLs), filtering on source and/or destination 1::/48 addresses, should be used to prevent these classes of packets being forwarded.

The ideal location to place packet filters for these classes of packets is as cl ose to the source of these packets as possible, which is on the origin hosts themselves. If the hosts support a packet filtering or more advanced firewalling capability, the filters would be applied to all externally attached interfaces and therefore to the packets traversing them. Preferably this should not be to distinct external interfaces, but rather to a class that contains active external interfaces, allowing the packet filter to be applied to dynamically created interfaces, such as those that may appear and disappear over time on mobile hosts. This is the best option to mitigate 1::/48 packet leaking if the hosts support this capability.

As a defence-in-depth measure, 1::/48 packet filters should also be applied to packets egressing and ingressing the local network, with the boundary of the local network likely to be where the local network has one or more attachments to the Internet or other external parties. These packet filters should be the first to be deployed, even if host based filtering is being used, to both cover the period during which filters are deployed to individual hosts, and to continue to act as a backup defence mechanism should the host filters fail or not be deployed on new hosts attached to the network.

It may also be useful to deploy packet filters at key, if not all routers within the local network. The chosen routers would likely correspond with security domain boundaries where it is important to drop packets with 1::/48 source and/or destination addresses. For example, if application developers are using 1::/48 addresses on their hosts, the router(s) where 1::/48 packet filters would be deployed as at the boundary between the application developer's sub-network and the rest of the local network.

Service Providers, in addition to deploying packet filters as above for their own 1::/48 use, should also apply 1::/48 packet filters to traffic received from their downstream customers' networks and their peer and upstream suppliers' networks. Service Providers are likely to have deployed ingress source address filtering to prevent denial of service attacks via source addres spoofing [BCP38], which will act as a filter for packets with 1::/48 source addresses. Service Provider routers that do not contain a default route (i.e. have complete knowledge of all Internet destinations) will drop packets with 1::/48 destinations, which may be an acceptable mitigation for 1::/48 packet leaks, if the customer, peer and upstream supplier

networks are attached to this class of routers. However, routing policy may change over time, so an explicit packet filter that drops packets with 1::/48 source and/or destination addresses, applied to all incoming packets would be wise.

In additon to packet filters, Service Providers must not accept route announcements for 1::/48. They must also ensure they do not announce route announcements for 1::/48 to their customers, peers and upstream providers.

9. Change Log [RFC Editor please remove]

draft-smith-larger-ipv6-loopback-prefix-00, initial version, 2012-07-24

10. References

10.1. Normative References

[RFC1122] Braden, R., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, October 1989.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

10.2. Informative References

[NTPD] "NTP: The Network Time Protocol", <http://ntp.org>.

[NTPD-RCD] "How to Write a Reference Clock Driver", <http://doc.ntp.org/4.2.6p5/howto.html>.

[RFC1812] Baker, F., "Requirements for IP Version 4 Routers", RFC 1812, June 1995.

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.

[RFC4038] Shin, M-K., Hong, Y-G., Hagino, J., Savola, P., and E. Castro, "Application Aspects of IPv6 Transition", RFC 4038, March 2005.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.

[RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, February 2006.

[RFC4443] Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, March 2006.

[RFC5156] Blanchet, M., "Special-Use IPv6 Addresses", RFC 5156, April 2008.

Author's Address

Mark Smith
In My Own Time
PO BOX 521
HEIDELBERG, VIC 3084
AU

Email: markzzzsmith@yahoo.com.au