

Unique Local IPv6 Unicast Addresses

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

This document defines an IPv6 unicast address format that is globally unique and is intended for local communications, usually inside of a site. These addresses are not expected to be routable on the global Internet.

Table of Contents

1. Introduction	2
2. Acknowledgements	3
3. Local IPv6 Unicast Addresses	3
3.1. Format	3
3.1.1. Background	4
3.2. Global ID	4
3.2.1. Locally Assigned Global IDs	5
3.2.2. Sample Code for Pseudo-Random Global ID Algorithm ...	5
3.2.3. Analysis of the Uniqueness of Global IDs	6
3.3. Scope Definition	6
4. Operational Guidelines	7
4.1. Routing	7
4.2. Renumbering and Site Merging	7
4.3. Site Border Router and Firewall Packet Filtering	8
4.4. DNS Issues	8
4.5. Application and Higher Level Protocol Issues	9
4.6. Use of Local IPv6 Addresses for Local Communication	9
4.7. Use of Local IPv6 Addresses with VPNs	10

5. Global Routing Considerations	11
5.1. From the Standpoint of the Internet	11
5.2. From the Standpoint of a Site	11
6. Advantages and Disadvantages	12
6.1. Advantages	12
6.2. Disadvantages	13
7. Security Considerations	13
8. IANA Considerations	13
9. References	13
9.1. Normative References	13
9.2. Informative References	14

1. Introduction

This document defines an IPv6 unicast address format that is globally unique and is intended for local communications [IPV6]. These addresses are called Unique Local IPv6 Unicast Addresses and are abbreviated in this document as Local IPv6 addresses. They are not expected to be routable on the global Internet. They are routable inside of a more limited area such as a site. They may also be routed between a limited set of sites.

Local IPv6 unicast addresses have the following characteristics:

- Globally unique prefix (with high probability of uniqueness).
- Well-known prefix to allow for easy filtering at site boundaries.
- Allow sites to be combined or privately interconnected without creating any address conflicts or requiring renumbering of interfaces that use these prefixes.
- Internet Service Provider independent and can be used for communications inside of a site without having any permanent or intermittent Internet connectivity.
- If accidentally leaked outside of a site via routing or DNS, there is no conflict with any other addresses.
- In practice, applications may treat these addresses like global scoped addresses.

This document defines the format of Local IPv6 addresses, how to allocate them, and usage considerations including routing, site border routers, DNS, application support, VPN usage, and guidelines for how to use for local communication inside a site.

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 [RFC2119].


2. Acknowledgements

The underlying idea of creating Local IPv6 addresses described in this document has been proposed a number of times by a variety of people. The authors of this document do not claim exclusive credit. Credit goes to Brian Carpenter, Christian Huitema, Aidan Williams, Andrew White, Charlie Perkins, and many others. The authors would also like to thank Brian Carpenter, Charlie Perkins, Harald Alvestrand, Keith Moore, Margaret Wasserman, Shannon Behrens, Alan Beard, Hans Kruse, Geoff Huston, Pekka Savola, Christian Huitema, Tim Chown, Steve Bellovin, Alex Zinin, Tony Hain, Bill Fenner, Sam Hartman, and Elwyn Davies for their comments and suggestions on this document.

3. Local IPv6 Unicast Addresses

3.1. Format

The Local IPv6 addresses are created using a pseudo-randomly allocated global ID. They have the following format:



Where:

Prefix	FC00::/7 prefix to identify Local IPv6 unicast addresses.
L	Set to 1 if the prefix is locally assigned. Set to 0 may be defined in the future. See Section 3.2 for additional information.
Global ID	40-bit global identifier used to create a globally unique prefix. See Section 3.2 for additional information.
Subnet ID	16-bit Subnet ID is an identifier of a subnet within the site.
Interface ID	64-bit Interface ID as defined in [ADDARCH].

3.1.1. Background

There were a range of choices available when choosing the size of the prefix and Global ID field length. There is a direct tradeoff between having a Global ID field large enough to support foreseeable future growth and not using too much of the IPv6 address space needlessly. A reasonable way of evaluating a specific field length is to compare it to a projected 2050 world population of 9.3 billion [POPUL] and the number of resulting /48 prefixes per person. A range of prefix choices is shown in the following table:

Prefix	Global ID Length	Number of /48 Prefixes	Prefixes per Person	% of IPv6 Address Space
/11	37	137,438,953,472	15	0.049%
/10	38	274,877,906,944	30	0.098%
/9	39	549,755,813,888	59	0.195%
/8	40	1,099,511,627,776	118	0.391%
/7	41	2,199,023,255,552	236	0.781%
/6	42	4,398,046,511,104	473	1.563%

A very high utilization ratio of these allocations can be assumed because the Global ID field does not require internal structure, and there is no reason to be able to aggregate the prefixes.

The authors believe that a /7 prefix resulting in a 41-bit Global ID space (including the L bit) is a good choice. It provides for a large number of assignments (i.e., 2.2 trillion) and at the same time uses less than .8% of the total IPv6 address space. It is unlikely that this space will be exhausted. If more than this were to be needed, then additional IPv6 address space could be allocated for this purpose.

3.2. Global ID

The allocation of Global IDs is pseudo-random [RANDOM]. They MUST NOT be assigned sequentially or with well-known numbers. This is to ensure that there is not any relationship between allocations and to help clarify that these prefixes are not intended to be routed globally. Specifically, these prefixes are not designed to aggregate.

This document defines a specific local method to allocate Global IDs, indicated by setting the L bit to 1. Another method, indicated by clearing the L bit, may be defined later. Apart from the allocation method, all Local IPv6 addresses behave and are treated identically.

The local assignments are self-generated and do not need any central coordination or assignment, but have an extremely high probability of being unique.

3.2.1. Locally Assigned Global IDs

Locally assigned Global IDs MUST be generated with a pseudo-random algorithm consistent with [RANDOM]. Section 3.2.2 describes a suggested algorithm. It is important that all sites generating Global IDs use a functionally similar algorithm to ensure there is a high probability of uniqueness.

The use of a pseudo-random algorithm to generate Global IDs in the locally assigned prefix gives an assurance that any network numbered using such a prefix is highly unlikely to have that address space clash with any other network that has another locally assigned prefix allocated to it. This is a particularly useful property when considering a number of scenarios including networks that merge, overlapping VPN address space, or hosts mobile between such networks.

3.2.2. Sample Code for Pseudo-Random Global ID Algorithm

The algorithm described below is intended to be used for locally assigned Global IDs. In each case the resulting global ID will be used in the appropriate prefix as defined in Section 3.2.

- 1) Obtain the current time of day in 64-bit NTP format [NTP].
- 2) Obtain an EUI-64 identifier from the system running this algorithm. If an EUI-64 does not exist, one can be created from a 48-bit MAC address as specified in [ADDARCH]. If an EUI-64 cannot be obtained or created, a suitably unique identifier, local to the node, should be used (e.g., system serial number).
- 3) Concatenate the time of day with the system-specific identifier in order to create a key.
- 4) Compute an SHA-1 digest on the key as specified in [FIPS, SHA1]; the resulting value is 160 bits.
- 5) Use the least significant 40 bits as the Global ID.
- 6) Concatenate FC00::/7, the L bit set to 1, and the 40-bit Global ID to create a Local IPv6 address prefix.

This algorithm will result in a Global ID that is reasonably unique and can be used to create a locally assigned Local IPv6 address prefix.

3.2.3. Analysis of the Uniqueness of Global IDs

The selection of a pseudo random Global ID is similar to the selection of an SSRC identifier in RTP/RTCP defined in Section 8.1 of [RTP]. This analysis is adapted from that document.

Since Global IDs are chosen randomly (and independently), it is possible that separate networks have chosen the same Global ID. For any given network, with one or more random Global IDs, that has inter-connections to other such networks, having a total of N such IDs, the probability that two or more of these IDs will collide can be approximated using the formula:

$$P = 1 - \exp(-N^{**2} / 2^{**}(L+1))$$

where P is the probability of collision, N is the number of interconnected Global IDs, and L is the length of the Global ID.

The following table shows the probability of a collision for a range of connections using a 40-bit Global ID field.

Connections	Probability of Collision
2	1.81*10 ⁻¹²
10	4.54*10 ⁻¹¹
100	4.54*10 ⁻⁰⁹
1000	4.54*10 ⁻⁰⁷
10000	4.54*10 ⁻⁰⁵

Based on this analysis, the uniqueness of locally generated Global IDs is adequate for sites planning a small to moderate amount of inter-site communication using locally generated Global IDs.

3.3. Scope Definition

By default, the scope of these addresses is global. That is, they are not limited by ambiguity like the site-local addresses defined in [ADDARCH]. Rather, these prefixes are globally unique, and as such, their applicability is greater than site-local addresses. Their limitation is in the routability of the prefixes, which is limited to a site and any explicit routing agreements with other sites to propagate them (also see Section 4.1). Also, unlike site-locals, a site may have more than one of these prefixes and use them at the same time.

4. Operational Guidelines

The guidelines in this section do not require any change to the normal routing and forwarding functionality in an IPv6 host or router. These are configuration and operational usage guidelines.

4.1. Routing

Local IPv6 addresses are designed to be routed inside of a site in the same manner as other types of unicast addresses. They can be carried in any IPv6 routing protocol without any change.

It is expected that they would share the same Subnet IDs with provider-based global unicast addresses, if they were being used concurrently [GLOBAL].

The default behavior of exterior routing protocol sessions between administrative routing regions must be to ignore receipt of and not advertise prefixes in the FC00::/7 block. A network operator may specifically configure prefixes longer than FC00::/7 for inter-site communication.

If BGP is being used at the site border with an ISP, the default BGP configuration must filter out any Local IPv6 address prefixes, both incoming and outgoing. It must be set both to keep any Local IPv6 address prefixes from being advertised outside of the site as well as to keep these prefixes from being learned from another site. The exception to this is if there are specific /48 or longer routes created for one or more Local IPv6 prefixes.

For link-state IGPs, it is suggested that a site utilizing IPv6 local address prefixes be contained within one IGP domain or area. By containing an IPv6 local address prefix to a single link-state area or domain, the distribution of prefixes can be controlled.

4.2. Renumbering and Site Merging

The use of Local IPv6 addresses in a site results in making communication that uses these addresses independent of renumbering a site's provider-based global addresses.

When merging multiple sites, the addresses created with these prefixes are unlikely to need to be renumbered because all of the addresses have a high probability of being unique. Routes for each specific prefix would have to be configured to allow routing to work correctly between the formerly separate sites.

4.3. Site Border Router and Firewall Packet Filtering

While no serious harm will be done if packets with these addresses are sent outside of a site via a default route, it is recommended that routers be configured by default to keep any packets with Local IPv6 addresses from leaking outside of the site and to keep any site prefixes from being advertised outside of their site.

Site border routers and firewalls should be configured to not forward any packets with Local IPv6 source or destination addresses outside of the site, unless they have been explicitly configured with routing information about specific /48 or longer Local IPv6 prefixes. This will ensure that packets with Local IPv6 destination addresses will not be forwarded outside of the site via a default route. The default behavior of these devices should be to install a "reject" route for these prefixes. Site border routers should respond with the appropriate ICMPv6 Destination Unreachable message to inform the source that the packet was not forwarded. [ICMPV6]. This feedback is important to avoid transport protocol timeouts.

Routers that maintain peering arrangements between Autonomous Systems throughout the Internet should obey the recommendations for site border routers, unless configured otherwise.

4.4. DNS Issues

At the present time, AAAA and PTR records for locally assigned local IPv6 addresses are not recommended to be installed in the global DNS.

For background on this recommendation, one of the concerns about adding AAAA and PTR records to the global DNS for locally assigned Local IPv6 addresses stems from the lack of complete assurance that the prefixes are unique. There is a small possibility that the same locally assigned IPv6 Local addresses will be used by two different organizations both claiming to be authoritative with different contents. In this scenario, it is likely there will be a connection attempt to the closest host with the corresponding locally assigned IPv6 Local address. This may result in connection timeouts, connection failures indicated by ICMP Destination Unreachable messages, or successful connections to the wrong host. Due to this concern, adding AAAA records for these addresses to the global DNS is thought to be unwise.

Reverse (address-to-name) queries for locally assigned IPv6 Local addresses MUST NOT be sent to name servers for the global DNS, due to the load that such queries would create for the authoritative name servers for the ip6.arpa zone. This form of query load is not specific to locally assigned Local IPv6 addresses; any current form

of local addressing creates additional load of this kind, due to reverse queries leaking out of the site. However, since allowing such queries to escape from the site serves no useful purpose, there is no good reason to make the existing load problems worse.

The recommended way to avoid sending such queries to nameservers for the global DNS is for recursive name server implementations to act as if they were authoritative for an empty `d.f.ip6.arpa` zone and return RCODE 3 for any such query. Implementations that choose this strategy should allow it to be overridden, but returning an RCODE 3 response for such queries should be the default, both because this will reduce the query load problem and also because, if the site administrator has not set up the reverse tree corresponding to the locally assigned IPv6 Local addresses in use, returning RCODE 3 is in fact the correct answer.

4.5. Application and Higher Level Protocol Issues

Application and other higher level protocols can treat Local IPv6 addresses in the same manner as other types of global unicast addresses. No special handling is required. This type of address may not be reachable, but that is no different from other types of IPv6 global unicast address. Applications need to be able to handle multiple addresses that may or may not be reachable at any point in time. In most cases, this complexity should be hidden in APIs.

From a host's perspective, the difference between Local IPv6 and other types of global unicast addresses shows up as different reachability and could be handled by default in that way. In some cases, it is better for nodes and applications to treat them differently from global unicast addresses. A starting point might be to give them preference over global unicast, but fall back to global unicast if a particular destination is found to be unreachable. Much of this behavior can be controlled by how they are allocated to nodes and put into the DNS. However, it is useful if a host can have both types of addresses and use them appropriately.

Note that the address selection mechanisms of [ADDSEL], and in particular the policy override mechanism replacing default address selection, are expected to be used on a site where Local IPv6 addresses are configured.

4.6. Use of Local IPv6 Addresses for Local Communication

Local IPv6 addresses, like global scope unicast addresses, are only assigned to nodes if their use has been enabled (via IPv6 address autoconfiguration [ADDAUTO], DHCPv6 [DHCP6], or manually). They are

not created automatically in the way that IPv6 link-local addresses are and will not appear or be used unless they are purposely configured.

In order for hosts to autoconfigure Local IPv6 addresses, routers have to be configured to advertise Local IPv6 /64 prefixes in router advertisements, or a DHCPv6 server must have been configured to assign them. In order for a node to learn the Local IPv6 address of another node, the Local IPv6 address must have been installed in a naming system (e.g., DNS, proprietary naming system, etc.) For these reasons, controlling their usage in a site is straightforward.

To limit the use of Local IPv6 addresses the following guidelines apply:

- Nodes that are to only be reachable inside of a site: The local DNS should be configured to only include the Local IPv6 addresses of these nodes. Nodes with only Local IPv6 addresses must not be installed in the global DNS.
- Nodes that are to be limited to only communicate with other nodes in the site: These nodes should be set to only autoconfigure Local IPv6 addresses via [ADDAUTO] or to only receive Local IPv6 addresses via [DHCP6]. Note: For the case where both global and Local IPv6 prefixes are being advertised on a subnet, this will require a switch in the devices to only autoconfigure Local IPv6 addresses.
- Nodes that are to be reachable from inside of the site and from outside of the site: The DNS should be configured to include the global addresses of these nodes. The local DNS may be configured to also include the Local IPv6 addresses of these nodes.
- Nodes that can communicate with other nodes inside of the site and outside of the site: These nodes should autoconfigure global addresses via [ADDAUTO] or receive global address via [DHCP6]. They may also obtain Local IPv6 addresses via the same mechanisms.

4.7. Use of Local IPv6 Addresses with VPNs

Local IPv6 addresses can be used for inter-site Virtual Private Networks (VPN) if appropriate routes are set up. Because the addresses are unique, these VPNs will work reliably and without the need for translation. They have the additional property that they will continue to work if the individual sites are renumbered or merged.

5. Global Routing Considerations

Section 4.1 provides operational guidelines that forbid default routing of local addresses between sites. Concerns were raised to the IPv6 working group and to the IETF as a whole that sites may attempt to use local addresses as globally routed provider-independent addresses. This section describes why using local addresses as globally-routed provider-independent addresses is unadvisable.

5.1. From the Standpoint of the Internet

There is a mismatch between the structure of IPv6 local addresses and the normal IPv6 wide area routing model. The /48 prefix of an IPv6 local addresses fits nowhere in the normal hierarchy of IPv6 unicast addresses. Normal IPv6 unicast addresses can be routed hierarchically down to physical subnet (link) level and only have to be flat-routed on the physical subnet. IPv6 local addresses would have to be flat-routed even over the wide area Internet.

Thus, packets whose destination address is an IPv6 local address could be routed over the wide area only if the corresponding /48 prefix were carried by the wide area routing protocol in use, such as BGP. This contravenes the operational assumption that long prefixes will be aggregated into many fewer short prefixes, to limit the table size and convergence time of the routing protocol. If a network uses both normal IPv6 addresses [ADDARCH] and IPv6 local addresses, these types of addresses will certainly not aggregate with each other, since they differ from the most significant bit onwards. Neither will IPv6 local addresses aggregate with each other, due to their random bit patterns. This means that there would be a very significant operational penalty for attempting to use IPv6 local address prefixes generically with currently known wide area routing technology.

5.2. From the Standpoint of a Site

There are a number of design factors in IPv6 local addresses that reduce the likelihood that IPv6 local addresses will be used as arbitrary global unicast addresses. These include:

- The default rules to filter packets and routes make it very difficult to use IPv6 local addresses for arbitrary use across the Internet. For a site to use them as general purpose unicast addresses, it would have to make sure that the default rules were not being used by all other sites and intermediate ISPs used for their current and future communication.

- They are not mathematically guaranteed to be unique and are not registered in public databases. Collisions, while highly unlikely, are possible and a collision can compromise the integrity of the communications. The lack of public registration creates operational problems.
- The addresses are allocated randomly. If a site had multiple prefixes that it wanted to be used globally, the cost of advertising them would be very high because they could not be aggregated.
- They have a long prefix (i.e., /48) so a single local address prefix doesn't provide enough address space to be used exclusively by the largest organizations.

6. Advantages and Disadvantages

6.1. Advantages

This approach has the following advantages:

- Provides Local IPv6 prefixes that can be used independently of any provider-based IPv6 unicast address allocations. This is useful for sites not always connected to the Internet or sites that wish to have a distinct prefix that can be used to localize traffic inside of the site.
- Applications can treat these addresses in an identical manner as any other type of global IPv6 unicast addresses.
- Sites can be merged without any renumbering of the Local IPv6 addresses.
- Sites can change their provider-based IPv6 unicast address without disrupting any communication that uses Local IPv6 addresses.
- Well-known prefix that allows for easy filtering at site boundary.
- Can be used for inter-site VPNs.
- If accidentally leaked outside of a site via routing or DNS, there is no conflict with any other addresses.

6.2. Disadvantages

This approach has the following disadvantages:

- Not possible to route Local IPv6 prefixes on the global Internet with current routing technology. Consequentially, it is necessary to have the default behavior of site border routers to filter these addresses.
- There is a very low probability of non-unique locally assigned Global IDs being generated by the algorithm in Section 3.2.3. This risk can be ignored for all practical purposes, but it leads to a theoretical risk of clashing address prefixes.

7. Security Considerations

Local IPv6 addresses do not provide any inherent security to the nodes that use them. They may be used with filters at site boundaries to keep Local IPv6 traffic inside of the site, but this is no more or less secure than filtering any other type of global IPv6 unicast addresses.

Local IPv6 addresses do allow for address-based security mechanisms, including IPsec, across end to end VPN connections.

8. IANA Considerations

The IANA has assigned the FC00::/7 prefix to "Unique Local Unicast".

9. References

9.1. Normative References

- [ADDARCH] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", RFC 3513, April 2003.
- [FIPS] "Federal Information Processing Standards Publication", (FIPS PUB) 180-1, Secure Hash Standard, 17 April 1995.
- [GLOBAL] Hinden, R., Deering, S., and E. Nordmark, "IPv6 Global Unicast Address Format", RFC 3587, August 2003.
- [ICMPV6] Conta, A. and S. Deering, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 2463, December 1998.

- [IPV6] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [NTP] Mills, D., "Network Time Protocol (Version 3) Specification, Implementation and Analysis", RFC 1305, March 1992.
- [RANDOM] Eastlake, D., 3rd, Schiller, J., and S. Crocker, "Randomness Requirements for Security", BCP 106, RFC 4086, June 2005.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [SHA1] Eastlake 3rd, D. and P. Jones, "US Secure Hash Algorithm 1 (SHA1)", RFC 3174, September 2001.

9.2. Informative References

- [ADDAUTO] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.
- [ADDSEL] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, February 2003.
- [DHCP6] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [POPUL] Population Reference Bureau, "World Population Data Sheet of the Population Reference Bureau 2002", August 2002.
- [RTP] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, July 2003.

Authors' Addresses

Robert M. Hinden
Nokia
313 Fairchild Drive
Mountain View, CA 94043
USA

Phone: +1 650 625-2004
EMail: bob.hinden@nokia.com

Brian Haberman
Johns Hopkins University
Applied Physics Lab
11100 Johns Hopkins Road
Laurel, MD 20723
USA

Phone: +1 443 778 1319
EMail: brian@innovationslab.net

Full Copyright Statement

Copyright (C) The Internet Society (2005).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.