

Unique Local IPv6 Unicast Addresses

<[draft-hinden-ipv6-global-local-addr-02.txt](#)>

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#). Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

To view the list Internet-Draft Shadow Directories, see <http://www.ietf.org/shadow.html>.

This internet draft expires on January 5, 2004.

Abstract

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

INTERNET-DRAFT

Unique Local IPv6 Unicast Addresses

June 2003

[1.0](#) 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 given current routing technology. 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.
- Well known prefix to allow for easy filtering at site boundaries.
- Allows sites to be combined or privately interconnected without creating any address conflicts or require renumbering of interfaces using 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 address like global scoped addresses.
- These addresses are also candidates for end-to-end use in some classes of multihoming solutions.

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 [[RFC 2119](#)].

[2.0](#) Acknowledgments

The underlying idea of creating Local IPv6 addresses describe in this document been proposed a number of times by a variety of people. The authors of this draft do not claim exclusive credit. Credit goes to

	Length	Prefixes	per Person	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 no internal structure is required in the field nor is there any reason to be able to aggregate the prefixes.

The authors believes that a /7 prefix resulting in a 41 bit Global ID 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 was needed, then additional IPv6 address space could be allocated for this purpose.

For the rest of this document the FC00::/7 prefix and a 41-bit Global ID is used.

[3.2](#) Global ID

The allocation of global IDs should be pseudo-random [[RANDOM](#)]. They should not be assigned sequentially or with well known numbers. This 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 designed to not aggregate.

There are two ways to allocate Global IDs. These are centrally by a allocation authority and locally by the site. The Global ID is split into two parts for each type of allocation. The prefixes for each type are:

FC00::/8	Centrally assigned
FD00::/8	Locally assigned

Each results in a 40-bit space to allocate.

Two assignment methods are included because they have different properties. The centrally assigned global IDs have a much higher probability that they are unique and the assignments can be escrowed to resolve any disputes regarding duplicate assignments. The local assignments are free and do not need any central coordination or assignment, but have a lower (but still adequate) probability of being unique. It is expected that large managed sites will prefer central assignments and small or disconnected sites will prefer local assignments. Sites are free to choose either approach.

[3.2.1 Centrally Assigned Global IDs](#)

Centrally assigned global IDs MUST be generated with a pseudo-random algorithm consistent with [\[RANDOM\]](#). They should not be assigned sequentially or by locality. This 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 designed to not aggregate.

Global IDs should be assigned centrally by a single allocation authority because they are pseudo-random and without any structure. This is easiest to accomplish if there is a single source of the assignments.

The requirements for centrally assigned global ID allocations are:

- Available to anyone in an unbiased manner.
- Permanent with no periodic fees.
- One time non-refundable allocation fee in the order of 10 Euros per allocation.
- The ownership of each individual allocation should be private, but should be escrowed.

The allocation authority should permit allocations to be obtained without having any sort of internet connectivity. For example in addition to web based registration they should support some methods like telephone, postal mail, fax, telex, etc. They should also accept a variety of payment methods and currencies.

The reason for the one time 10 Euro charge for each prefix is to provide a barrier to any hoarding of the these allocations but at the same time keep the cost low enough to not create a barrier to anyone needing one. The charge is non-refundable in order to keep overhead low.

The ownership of the allocations is not needed to be public since the resulting addresses are intended to be used for local communication. It is escrowed to insure there are no duplicate allocations and in case it is needed in the future (e.g., to resolve duplicate allocation disputes).

An example of a allocation authority is a non-profit organization such as the Public Internet Registry (PIR) that the Internet Society has created to manage the .org domain. They already know how to collect small sums efficiently and there are safeguards in place for the appropriate use of any excess revenue generated.

Note, there are many possible ways of of creating an allocation authority. It is important to keep in mind when reviewing alternatives that the goal is to pick one that can do the job. It doesn't have to be perfect, only good enough to do the job at hand. The authors believe that PIR shows that this requirement can be satisfied, but this draft does not specifically recommend the PIR.

[3.2.2](#) Locally Assigned Global IDs

Global IDs can also be generated locally by an individual site. This makes it easy to get a prefix with out the need to contact an assignment authority or internet service provider. There is not as high a degree of assurance that the prefix will not conflict with another locally generated prefix, but the likelihood of conflict is small. Sites that are not comfortable with this degree of uncertainty should use a centrally assigned global ID.

Locally assigned global IDs MUST be generated with a pseudo-random algorithm consistent with [\[RANDOM\]](#). [Section 3.2.3](#) describes a suggested algorithm. It is important to insure a reasonable likelihood uniqueness that all sites generating a Global IDs use a functionally similar algorithm.

[3.2.3](#) Sample Code for Pseudo-Random Global ID Algorithm

The algorithm described below is intended to be used for centrally and 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 the birth date of the person running the algorithm (or one of his/her descendants or ancestors) in 64-bit NTP format.
- 3) Concatenate the time of day with the birth date resulting in a 128-bit value (i.e., TOD, Birthday).
- 4) Compute an MD5 digest on the 128-bit value as specified in [[MD5DIG](#)].
- 5) Use the least significant 40 bits as the Global ID.

This algorithm will result in a global ID that is reasonably unique and can be used as a Global ID.

[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 [ADDRARCH]. Rather, these prefixes are globally unique, and as such, their applicability exceeds the current 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, unlike site-locals, these prefixes can overlap

[4.0](#) 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](#)].

Any routing protocol that is used between sites is required to filter out any incoming or outgoing Local IPv6 unicast routes. The exception to this is if specific /48 IPv6 local unicast routes have been configured to allow for inter-site communication.

If BGP is being used at the site border with an ISP, by default filters **MUST** be installed in the BGP configuration to keep any Local IPv6 address prefixes from being advertised outside of the site or for these prefixes to be learned from another site. The exception to this is if there are specific /48 routes created for one or more Local IPv6 prefixes.

[5.0](#) Renumbering and Site Merging

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

When merging multiple sites none of the addresses created with these prefixes need to be renumbered because all of the addresses are unique. Routes for each specific prefix would have to be configured to allow routing to work correctly between the formerly separate sites.

[6.0](#) Site Border Router and Firewall 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 they be filtered to keep any packets with Local IPv6 destination addresses from leaking outside of the site and to keep any site prefixes from being advertised outside of their site.

Site border routers **SHOULD** install a black hole route for the Local IPv6 prefix FC00::/7. This will insure that packets with Local IPv6 destination addresses will not be forwarded outside of the site via a default route.

Site border routers and firewalls **SHOULD 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 other Local IPv6 prefixes. The default behavior of these devices SHOULD be to filter them.

[7.0](#) DNS Issues

Local IPv6 addresses SHOULD NOT be installed in the global DNS. They may be installed in a naming system local to the site or kept separate from the global DNS using techniques such as "two-faced" DNS.

If Local IPv6 address are configured in the global DNS, no harm is done because they are unique and will not create any confusion. They may not be reachable, but this is a property that is common to all types of global IPv6 unicast addresses.

For future study names with Local IPv6 addresses may be resolved inside of the site using dynamic naming systems such as Multicast DNS.

[8.0](#) 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 addresses may not be reachable, but that is no different from other types of IPv6 global unicast addresses. Applications need to be able to handle multiple addresses that may or may not be reachable any point in time. In most cases this complexity should be hidden in APIs.

From a host's perspective this difference shows up as different reachability than global unicast and could be handled by default 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.

9.0 Use of Local IPv6 Addresses for Local Communications

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) and configured in the DNS. They are not created automatically 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. Likewise, 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 the DNS. For these reasons, it is straight forward to control their usage in a site.

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.

[10.0](#) 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.

[11.0](#) Advantages and Disadvantages

[11.1](#) 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 using 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.

[11.2](#) 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.

[12.0](#) 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.

[13.0](#) IANA Considerations

The IANA is instructed to allocate the FC00::/7 prefix for Unique Local IPv6 unicast addresses.

The IANA is instructed to delegate, within a reasonable time, the prefix FC00::/8 to an allocation authority for Unique Local IPv6 Unicast prefixes of length /48. This allocation authority shall comply with the requirements described in [section 3.2](#) of this document, including in particular the charging of a modest one-time fee, with any profit being used for the public good in connection with the Internet.

[14.0](#) Change Log

Draft <[draft-hinden-ipv6-global-local-addr-02.txt](#)>

- o Changed title and name of addresses defined in this document to "Unique Local IPv6 Unicast Addresses" with abbreviation of

- "Local IPv6 addresses".
- o Several editorial changes.

Draft <[draft-hinden-ipv6-global-local-addr-01.txt](#)>

- o Added section on scope definition and updated application requirement section.
- o Clarified that, by default, the scope of these addresses is global.
- o Renumbered sections and general text improvements
- o Removed reserved global ID values
- o Added pseudo code for local allocation submitted by Brian Haberman and added him as an author.
- o Split Global ID values into centrally assigned and local assignments and added text to describe local assignments

[draft-hinden-ipv6-global-local-addr-02.txt](#)

[Page 11]

INTERNET-DRAFT Unique Local IPv6 Unicast Addresses

June 2003

Draft <[draft-hinden-ipv6-global-local-addr-00.txt](#)>

- o Initial Draft

REFERENCES

Normative

- [ADDARCH] Hinden, R., S. Deering, S., "IP Version 6 Addressing Architecture", [RFC3515](#), April 2003.
- [GLOBAL] Hinden, R., S. Deering, E. Nordmark, "IPv6 Global Unicast Address Format", Internet Draft, <[draft-ietf-ipv6-unicast-aggr-v2-03.txt](#)>, June 2003.
- [IPV6] Deering, S., R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC2460](#), December 1998.
- [MD5DIG] Rivest, R., "The MD5 Message-Digest Algorithm", [RFC1321](#), April 1992.
- [NTP] Mills, David L., "Network Time Protocol (Version 3) Specification, Implementation and Analysis", [RFC1305](#), March 1992.

- [POPUL] Population Reference Bureau, "World Population Data Sheet of the Population Reference Bureau 2002", August 2002.
- [RANDOM] Eastlake, D. 3rd, S. Crocker, J. Schiller, "Randomness Recommendations for Security", [RFC1750](#), December 1994.
- [RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", [RFC2026](#), [BCP00009](#), October 1996.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [RFC2119](#), [BCP14](#), March 1997.

Non-Normative

- [ADDAUTO] Thomson, S., T. Narten, "IPv6 Stateless Address Autoconfiguration", [RFC2462](#), December 1998.
- [DHCP6] Droms, R., et. al., "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", Internet Draft, <[draft-ietf-dhc-dhcpv6-28.txt](#)>, November 2002.

[draft-hinden-ipv6-global-local-addr-02.txt](#)

[Page 12]

INTERNET-DRAFT

Unique Local IPv6 Unicast Addresses

June 2003

- [ADDSEL] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", [RFC3484](#), February 2003.

AUTHOR'S ADDRESSES

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

phone: +1 650 625-2004
email: bob.hinden@nokia.com

Brian Haberman
Caspian Networks

1 Park Drive, Suite 300
Research Triangle Park, NC 27709
US

phone: +1-929-949-4828
email: brian@innovationslab.net