INTERNET-DRAFT   R. Hinden, Nokia
February 12, 2004   B. Haberman, Caspian

                 Unique Local IPv6 Unicast Addresses

              <draft-ietf-ipv6-unique-local-addr-03.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 August 17, 2004.

Abstract

   This document defines an IPv6 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.

Table of Contents

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. 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 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 [RFC 2119].

2.0 Acknowledgments

The underlying idea of creating Local IPv6 addresses described 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 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, and Christian Huitema for their comments
and suggestions on this document.

3.0 Local IPv6 Unicast Addresses

3.1 Format

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

```
| 7 bits |  41 bits   |  16 bits  |          64 bits              |
+--------+------------+-----------+-------------------------------+
| prefix | global ID  | subnet ID |         interface ID          |
+--------+------------+-----------+-------------------------------+
```

Where:

   prefix   FC00::/7 prefix to identify Local IPv6 unicast
                   addresses.

   global ID   41-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
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 should be pseudo-random [RANDOM]. They
should 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 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 are uniquely assigned
   and the assignments can be escrowed to resolve any disputes regarding
   duplicate assignments. The local assignments are self generated 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. It is recommended
   that sites planning to use Local IPv6 addresses for extensive inter-
   site communication, initially or as a future possibility, use a
   centrally assigned prefix as there is no possibility of assignment
   conflicts. 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 is to ensure that there is no
   relationship between allocations and to help clarify that these
   prefixes are not intended to be routed globally by eliminating the
   possibility of aggregation. 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 for the
   assignments.

   The requirements for centrally assigned global ID allocations are:

      - Available to anyone in an unbiased manner.
      - Permanent with no periodic fees.
      - Allocation on a permanent basis, without any need for renewal
        and without any procedure for de-allocation.
      - Provide mechanisms that prevent hoarding of these allocations.
      - 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, etc.

The allocation service should include sufficient provisions to avoid hoarding of numbers. This can be accomplished by various ways, for example, requiring an exchange of documents, a verbal contact, or a proof that the request is on behalf of a human rather than a machine. The service may charge a small fee in order to cover its costs, but the fee should be low enough to not create a barrier to anyone needing one. The precise mechanisms should be decided by the registration authority.

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 ensure there are no duplicate allocations and in case it is needed in the future (e.g., to resolve duplicate allocation disputes, or to support a change of the central allocation authority).

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.

This document directs the IANA, in section 13.0, to delegate the FC00::/8 prefix to an allocation authority to allocate centrally assigned /48 prefixes consistent with the requirements defined in this section.

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 without 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 ensure 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 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
     creating a key.
  4) Compute an MD5 digest on the key 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.

3.2.4 **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 the global ID is chosen randomly, it is possible that two or
more networks that have an inter-network connection using globally-
unique local addresses will chose the same global ID. The
probability of collision can be approximated based on the number of
connections between networks using globally-unique local addresses
and the length of the ID (40 bits). The formula

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

approximates the probability of collision (where N is the number
connections 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. Sites
planning more extensive inter-site communication should consider
using the centrally assigned 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
[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, unlike site-locals, a site may have more than
one of these prefixes and use them at the same time.

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 router that is used between sites must be configured 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, the default BGP
configuration must be set to 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 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 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 "reject" route for the Local
   IPv6 prefix FC00::/7. This will ensure that packets with Local IPv6
   destination addresses will not be forwarded outside of the site via a
   default route. 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.

   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 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.

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

7.0 **DNS Issues**

   AAAA records for 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.

   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.

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). 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, 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 the
   DNS or another naming system. 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 accidently 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 allocation on a permanent basis and
   with sufficient provisions to avoid hoarding of numbers. If deemed
   appropriate, the authority may also consist of multiple organizations
   performing the authority duties.

14.0 **References**

14.1 **Normative References**

[ADDARCH] Hinden, R., S. Deering, S., "IP Version 6 Addressing
          Architecture", RFC 3515, April 2003.

[GLOBAL] Hinden, R., S. Deering, E. Nordmark, "IPv6 Global Unicast
          Address Format", RFC 3587, August 2003.

[ICMPV6] Conta, A., S. Deering, "Internet Control Message Protocol
          (ICMPv6) for the Internet Protocol Version 6 (IPv6)
          Specification", RFC2463, December 1998.

[IPV6]   Deering, S., R. Hinden, "Internet Protocol, Version 6
          (IPv6) Specification", RFC 2460, December 1998.

[MD5DIG] Rivest, R., "The MD5 Message-Digest Algorithm", RFC 1321,
          April 1992.

[NTP]   Mills, David L., "Network Time Protocol (Version 3)
          Specification, Implementation and Analysis", RFC 1305,
          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", RFC 1750, December 1994.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
          Requirement Levels", RFC 2119, BCP14, March 1997.

14.2 **Informative References**

[ADDAUTO] Thomson, S., 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., et. al., "Dynamic Host Configuration Protocol
          for IPv6 (DHCPv6)", RFC3315, July 2003.

[RTP]   Schulzrinne, H., S. Casner, R. Frederick, V. Jacobson,
          "RTP: A Transport Protocol for Real-Time Applications"
          RFC3550, July 2003.

15.0 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
Caspian Networks
1 Park Drive, Suite 300
Research Triangle Park, NC 27709
USA

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

16.0 Change Log

   Draft <draft-ietf-ipv6-unique-local-addr-03.txt>

      o Removed requirement of a fee per central allocation and updated
        IANA considerations to reflect this. Rewrote text to focus on
        the requirement to avoid hoarding of allocations.
      o Changed "filtering" and "black hole routes" to "reject" routes.
      o Changed uppers case requirements (i.e., MUST, SHOULD, etc.) to
        lower case in sections giving operational advice.
      o Removed paragraph mentioning "Multicast DNS".
      o Various editorial changes.

   Draft <draft-ietf-ipv6-unique-local-addr-02.txt>

      o Removed mention of 10 euro charge and changed text in section
        3.2.1 and IANA considerations to restate the requirement for low
        cost allocations and added specific requirement to prevent
        hording.
      o Added need to send ICMPv6 destination unreachable messages if
        packets are filtered or dropped at site boundaries.
      o Changed format section to list prefix sizes and definition, and
        changed discussion of prefix sizes to new background section.
      o Various editorial changes.

   Draft <draft-ietf-ipv6-unique-local-addr-01.txt>

      o Removed example of PIR as an example of an allocation authority
        and clarified the text that the IANA should delegate the
        centrally assigned prefix to an allocation authority.
      o Changed sample code for generating pseudo random Global IDs to
        not require any human input. Changes from using birthday to
        unique token (e.g., EUI-64, serial number, etc.) available on
        machine running the algorithm.
      o Added a new section analyzing the uniqueness properties of the
        pseudo random number algorithm.
      o Added text to recommend that centrally assigned local addresses
        be used for site planning extensive inter-site communication.
      o Clarified that domain border routers should follow site border
        router recommendations.
      o Clarified that AAAA DNS records should not be installed in the
        global DNS.
      o Several editorial changes.

   Draft <draft-ietf-ipv6-unique-local-addr-00.txt>

      o Changed file name to become an IPv6 w.g. group document.
      o Clarified language in Routing and Firewall sections.

    o Several editorial changes.

   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 <draft-hinden-ipv6-global-local-addr-00.txt>

    o Initial Draft