Unique Local IPv6 Unicast Addresses

<draft-ietf-ipv6-unique-local-addr-08.txt>

Status of this Memo

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 requiring 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, Christian Huitema, Tim Chown, Steve
Bellovin, Alex Zinin, Tony Hain, and Bill Fenner 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 pseudo-randomly
allocated global ID. They have the following format:

```
| 7 bits |1|  40 bits   |  16 bits  |          64 bits           |
+--------+-+-----------+-----------+----------------------------+
| prefix |L| global ID  | subnet ID |        interface ID        |
+--------+-+-----------+-----------+----------------------------+
```

   Where:

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

      L   Set to 1 if the prefix is locally assigned,
                          Set to 0 if it is centrally assigned. 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 40 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 not designed to
aggregate.

There are two ways to allocate Global IDs. These are centrally by a
allocation authority and locally by the site. The type of allocation
is distinguished by the L bit.

Two assignment methods are included because they have different
properties. The centrally assigned global IDs are uniquely assigned.
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.

This document only defines the allocation procedure for creating
global-IDs for locally assigned local IPv6 addresses (i.e., L=1).
The allocation procedure for centrally assigned local IPv6 addresses
(i.e., L=0) will be defined in a separate document.

3.2.1 Locally Assigned Global IDs

Global IDs can 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 3.2.2 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.

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
     creating 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 of that two or more of these IDs will collide can be

   approximated using 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 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 ULA
prefixes be contained either within one IGP domain or area. By
containing a ULA 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 using 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 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 be configured to 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 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. 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.

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.
The operational issues relating to this are beyond the scope of this
document.

For background on this recommendation, the concern 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 PTR record
might be registered by two different organizations. Due to this
concern, adding AAAA records is thought to be unwise because matching
PTR records can not be registered

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 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.

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 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.

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.

5.0 Advantages and Disadvantages

5.1 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.

5.2 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.

6.0 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.

7.0 IANA Considerations

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

8.0 References

8.1 Normative References

   [ADDARCH] Hinden, R., S. Deering, S., "IP Version 6 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., 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.

   [NTP]   Mills, David L., "Network Time Protocol (Version 3)
              Specification, Implementation and Analysis", RFC 1305,
              March 1992.

   [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.

   [SHA1]   D. Eastlake 3rd, P. Jones, "US Secure Hash Algorithm 1
              (SHA1)", RFC 3174, September 2001.

8.2 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.

   [POPUL]   Population Reference Bureau, "World Population Data Sheet
              of the Population Reference Bureau 2002", August 2002.

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

9.0 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
Johns Hopkins University
Applied Physics Lab
11100 Johns Hopkins Road
Laurel, MD 20723
USA

phone: +1 443 778 1319
email: brian@innovationslab.net

10.0 Change Log

   Draft <draft-ietf-ipv6-unique-local-addr-08.txt>
      o Moved sections 4-10, into one new section 4.0 titled
        "operational guidelines" to clarify the their scope.
      o Clarified routing requirements to make it clearer that these
        prefixes should not be routed on the global Internet.
      o Various editorial changes.

   Draft <draft-ietf-ipv6-unique-local-addr-07.txt>
      o Changed the format in section 3.1 in add a "L" (local/central)
        bit and reduced the size of the global-ID to 40 bits. This is
        equivalent to the previous separate prefixes and makes the
        document clearer.
      o Changed pseudo-random algorithm to use SHA-1 instead of MD5.
      o Moved [POPUL] to be an informative reference.
      o Added paragraph in Routing section to discuss the use of IGPs.
      o Various editorial changes.

   Draft <draft-ietf-ipv6-unique-local-addr-06.txt>
      o Clarify text to permit prefixes longer than /48 to be
        configured.
      o Changed text in section 7.0 to recommend that locally assigned
        ULA addresses are not installed in the global DNS and removed
        text about consequences of if they were installed in the global
        DNS.
      o Clarify the text in section 5.1 to state that there is a high
        probability that there will be no address conflict when
        renumbering.
      o Several minor editorial changes.

   Draft <draft-ietf-ipv6-unique-local-addr-05.txt>
      o Removed the definition and technical requirements for centrally
        assigned local address. The Centrally assigned local addresses
        will be defined in a separate document. This document defines
        the specific prefixes to be used for centrally and locally
        assigned IPv6 local addresses, but only the locally assigned
        local addresses are defined here.

   Draft <draft-ietf-ipv6-unique-local-addr-04.txt>

      o Clarified text in section 3.2.1 that central assigned prefixes
        should be assigned under the authority of a single allocation
        organization.
      o Added step to suggested pseudo-random algorithm that in the case
        of centrally assigned prefixes the computed global IDs should be
        verified against the escrow.
      o Added new text to section 3.2.2 that explains in more detail the

      need for pseudo-random global IDs (i.e., avoid duplicate
      allocations).
   o Rewrote section 7.0 to describe DNS AAAA and PTR records, and
      clarify when they might be installed in the global DNS.
   o Various editorial changes.

   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

11. Disclaimer of Validity

   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.

12. Copyright Statement