Commentary on Distribution Mechanisms for Unique Local IPv6 Unicast Addresses

Geoff Huston
August 2003

Abstract

This memorandum examines the characteristics of Unique Local IPv6 Unicast addresses, as well as the requirements for address distribution mechanisms for this class of addresses. It is intended as a commentary on an Internet Draft currently under consideration in the IPv6 Working Group of the IETF..

Introduction

Current work within the IETF IPv6 working includes the drafting of a proposal to define part of the IPv6 unicast address space for local use. This is currently IETF work in progress being considered by the IPv6 Working Group, documented in an Internet draft, "draft-hinden-ipv6-global-local-addr-02.txt" (attached). These addresses are intended for various forms of local communications and are not expected to be routable on the global Internet. The proposal refers to such addresses as "Unique Local IPv6 Unicast Addresses".

There are a number of characteristics of such addresses that have been proposed in order to ensure that they can fulfill the role of a local-use address, and there are also a number of considerations relating to the distribution mechanisms for these addresses that distinguish them from globally routable unicast addresses. This document explores these intended characteristics in further detail as well as the associated distribution mechanisms.

Characteristics of Local Use Addresses

The characteristics listed in the draft proposal for such addresses are:

```
1    Globally unique prefix.
2    Well known prefix to allow for easy filtering at site boundaries.
3    Allows sites to be combined or privately interconnected without
     creating any address conflicts or require renumbering of interfaces
     using these prefixes.
4    Internet Service Provider independent and can be used for
     communications inside of a site without having any permanent or
     intermittent Internet connectivity.
5    If accidentally leaked outside of a site via routing or DNS, there is
     no conflict with any other addresses.
6    In practice, applications may treat these address like global scoped
     addresses.
7    These addresses are also candidates for end-to-end use in some
     classes of multihoming solutions.
```

It could be argued that, strictly, the third and fifth characteristics are a consequence of the first, as they can be all grouped under the overall characteristic of "use of a common unique prefix". The second, forth and sixth characteristics commonly refer to unique use of a local address block drawn from the global unicast address pool. Also, strictly speaking, the seventh characteristic is not a characteristic per se, but flags a special case that involves further consideration in the context of multi-homing.

Restating this list of characteristics gives:

1. Exclusive use of a common prefix drawn from the global unicast address space for all local use addresses.

2. Unique assignment of local use address blocks from within the pool of addresses defined by this prefix.

Section 3.1 of the Internet Draft proposal further refines the set of characteristics, by describing the address as a four part object:

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

where:

`prefix`	prefix to identify Local IPv6 unicast addresses. FC00::/7
`global ID`	global identifier used to create a globally unique prefix.
`subnet ID`	16-bit subnet ID is an identifier of a subnet within the site.
`interface ID`	64-bit Interface ID.

The length of the prefix + global ID part is 48 bits in length, allowing 16 bits for local assignation of subnet IDs and 64 bits for the interface ID. This allows for 2,199,023,255,552 assignable local use address blocks.

There is a further characteristic of the address block defined in this section of the draft, namely:

3. There is no internal structure within the global ID, and these global IDs cannot be aggregated in a routing context.

The proposal splits this address pool into two halves: locally and centrally assigned prefixes.

Locally Assigned Local Use Prefixes

One half, using the common prefix FD00::/8, is described as being "locally assigned". The proposal indicates that such locally assigned global IDs must be generated with a pseudo-random algorithm. The proposal notes that there is a high probability that the prefix will not conflict with another locally generated prefix, but there is no absolute assurance of this outcome. Analysis of the probability involved here indicates that the probability of a collision in the space using a random draw function exceeds 0.5 after 1.2 million random draws.

Probability P of a collision after *d* draws from *n* possible values

$$P = 1 - (n! / ((n^{**}d) ((n-d)!)))$$

This is likely to be too small a value for any assured level of uniqueness, particularly if there is some consideration that such values can be useable as unique prefixes within end-to-end contexts. Some further consideration should be given to this part of the proposal.

It is concluded here that this 'random draw' is an inadequate response to item 2 of the required characteristics for Local Use addresses, particularly if such addresses are being contemplated to be used in the context of end-point identification. A probability of uniqueness is tangibly different to the property of assured uniqueness. If strong uniqueness is an essential characteristic of all elements of this address space, then it is necessary to drop the random self-selection mechanism from the draft proposal, and that all Local-Use addresses be distributed in such a manner that uniqueness is assured in every case.

Centrally Assigned Local Use Prefixes

The other half of the local use space is proposed in the draft to be "centrally assigned" using fixed size /48 blocks. This refines the second characteristic to read:

2. Unique assignment of fixed size local use address blocks from within the pool of addresses defined by this prefix, using a Global ID as the block prefix.

The proposal notes that these assignments can be escrowed to resolve any disputes regarding duplicate assignments. It is noted that escrow is a specific solution to a more general characteristic, and the desired characteristic being defined here is:

4. The assignment information must be recorded and stored in a reliable manner.

The assignment function is described in the proposal as one that treats sequential allocations in a random fashion, and explicitly notes that they should not be assigned accordingly to any particular structure, and therefore they cannot be aggregated in a routing environment.

5. Local Use Addresses are not intended to be passed within the global routing environment

The complete list of characteristics of this Centrally Assigned Local Use IPv6 Unicast address space is:

1. Exclusive use of a common prefix drawn from the global unicast address space for all local use addresses.

2. Unique assignment of fixed size local use address blocks from within the pool of addresses defined by this prefix, using a Global ID as the block prefix.

3. There is no internal structure within the global ID, and these global IDs cannot be aggregated in a routing context.

4. The assignment information must be recorded and stored in a reliable manner.

5. Local Use Addresses are not intended to be passed within the global routing environment

The potential for use of this address in end-to-end solutions relating to multi-homing is limited to the extent that this identity space is unstructured, so it cannot be used as a lookup key in any mapping system that maps identities into locators. If the intended use is through a sequence of mappings from domain name to identifier to current locator, then the last mapping (from identifier to locator) is not feasible in an unstructured identifier space. In this sense the role of such an address is limited to an assertion of a fixed, globally unique label that can be used in conjunction with dynamic change of location-based address to provide some form of transport session resiliency in a multi-homed environment.

Local Use Address Distribution Mechanisms

The proposal notes that:

```
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 unstated implication from the first requirement is that this is undertaken without consideration of the current or intended level of use of the address block, so that there are no qualifications regarding assignment of a Local Use Address block. The proposal also notes that such availability should include non-Internet access mechanisms as a desired additional mechanism.

The second and third aspects of this proposed distribution mechanism describe the use of a one-time fee for a one-time service transaction that has enduring consequences.

Allocation Fees

The first aspect here is the consideration of the allocation fee. The draft motivates this payment as a means of prevention of hoarding of blocks from within this pool by imposing a financial impost. While there are many forms of control over a distribution mechanism to prevent distortions such as hoarding, this pricing approach is seen as a lightweight and effective mechanism that has the potential to address the identified problem. However, there are some consequences of this aspect of the draft proposal that should be examined in further detail. The imposition of a charge without relation to service cost is seen in many regulatory regimes as an imposition that is likened to a monopoly rental or a form of taxation. Such forms of charges have no valid role, and should be avoided. It is more reasonable to allow the operator(s) of this distribution mechanism to be able to account for their costs in operating this service, and allow the operator to determine a service fee that is based on these costs.

The operator needs to consider that if this is to be a one-time fee for an unbounded service (so called 'cemetery plot' fees), the fee should cover both the processing component and the subsequent record maintenance component of the service.

Allocation Period

The proposal explicitly indicates that the allocation should be 'permanent'. This implies that there is no concept of return of a Local Use prefix once it has been allocated from the central registry, and that there is no concept of a registry-recorded transfer of an allocation. The implication of this service model is that there is no form of reuse of blocks from this address space. The implicit assumption here is that for the entire useful lifetime of the technology, under all conceivable allocation demand scenarios, that there will be adequate available address space to continue to meet demand from the Local Use address pool. Without any form of periodic renewal or similar opportunity to alter the terms of use of this address space then, if exhaustion of the space is considered to be a potential risk, the observations made in 1994 regarding the possible outcomes of the (then) IPv4 address allocation practices are once more relevant here:

```
"It is perhaps a sad reflection of the conflict of short term
objectives and longer term considerations that the evident short term
motivations of ready and equitable access to the IPv4 address (which
were the motivational factors in determining the current Internet
address allocation policies) run the consequent risk of monopoly-
based restrictive trade and barrier-based pricing as a longer term
outcome of unallocated address space exhaustion."

[RFC 1744 "Management of Internet Address Space"]
```

Of course if there is a high degree of confidence that exhaustion of the Local Use address pool is not a remotely possible eventuality, then such address prefixes can be considered in the same terms as a single-use disposable facility, and these considerations are not directly relevant.

Choice in Service Models

It is possible that clients of this allocation service want the choice between a single one-time permanent allocation (and a one-time service fee) and a defined period renewable service, where, at the end of the defined period the client has the choice of renewing the allocation or allowing it to lapse back to the pool. Given the central nature of the described distribution mechanism, allowing the client some choice in the form of service, rather than imposing a single service model is seen as a reasonable measure.

The model also proposes a single layer of distribution, where end clients interact with a proposed single central registry. Again this is an area where a different structure used for the distribution of many other forms of goods or resources, typically using some form of hierarchy in distribution with wholesale and retail roles. Such hierarchies often allow for a more efficient form of overall distribution than a single entity attempting to service a global consumer base. Current regulatory environments also look to competition as a means of ensuring that service regimes operate efficiently and that no single player can distort the price of the service through the imposition of monopoly rentals, artificial scarcity or selective servicing.

Recording Allocations

The proposal indicates that information relating to the 'ownership' of each individual allocation be private. This is not an easily achieved outcome, given that 'ownership' is a public claim to the unique ability to access and exploit the resource. Furthermore, this implies that the resource itself is a form of property, and that property can be traded, swapped or otherwise disposed of at the discretion of the owner, inferring that the address block, is in some form, an asset of the holder. It is unclear that this interpretation of the status of an address is the actual intent of the proponents of this approach, and that other forms of expression of unique and enduring interest in the address resource may be more appropriate for this resource. This observation is made in the context of the characterization of the larger protocol address space as a public good that is distinguished from concepts of ownership or the inferring of aspects of property and asset into this resource.

Reverse Mapping Local Use Addresses in ip6.arpa

It is unclear from the proposal whether Local Use Addresses could or should be entered into the ip6.arpa reverse mapping domain space. as a delegated domain.

Locally assigned prefixes cannot be entered into this domain space because of the lack of a condition of assured uniqueness.

The situation with respect to centrally assigned prefixes is not so clear. The considerations include:

- The potential size of the domain zone. Because of the lack of any structure beyond the 8th bit of the prefix, there is no ability to impose a hierarchy of zone files, and the reverse zone would need to list all assigned local use prefixes and their delegation points. There are obvious implications in terms of the potential size of this zone file. and some consideration as to tyhe efficiency of operation of a zone of such a potential size.

- The desired characteristic of Local Use prefixes where the "ownership" of the prefix is not public information. If the domain zone operator was distinct from the central registry operator, then the privacy of the address allocation information could preclude the domain operator from validating a delegation request for a Local Use address block.

- The potential use of these addresses in some classes of end-point identification may imply the need for an external entity, using the global DNS to be map from the local use identifier to a global use address, and one way to perform this mapping in the DNS is to use the reverse domain to map from the end point local use address to a global DNS name, and then map forward from this name to a global address. Precluding local use addresses from the global DNS would preclude this form of mapping.

For local use, a so called "two-faced" DNS can be configured to provide a local reverse mapping service for the local site.

Management Requirements for Local Use Addresses

In summary, the characteristics of the management of this space is where:

1. Every applicant may obtain an address block in this prefix space without providing any form of justification to the registry operator.

2. Every assigned Local-Use block is of the same size, namely a /48.

3. Each block is uniquely assigned to the applicant.

4. Each assignment is a randomly selected block from the entire remaining pool.

5. Each applicant may obtain an enduring assignment without further need to contact the registry or to pay further service fees (one-off service).

6. Any service fee, if used, should be high enough to make massive seizure financially undesirable, yet low enough to make it readily accessible to individuals as well as corporate entities on a global scale.

7. Any service fee, if used, should be clearly attributable to the costs associated with the provision of the service function for the lifetime of the provided service.

8. The service model is not restricted to a one-off assignment model, with the proviso that any other associated service models must have similar attributes of ease of accessibility.

9. The association of the assigned space and the identity of the applicant is not to be made public.

10. The assignment information is to be held in a way that is reliable and enduring.

Under the current arrangements, IANA is the IETF-selected registry for IPv4 and IPv4 global unicast space, and the RIRs undertake the associated distribution function, using policies that have been developed by an open process within each region.

Distribution Mechanisms

A complete consideration of the various regulatory and logistical considerations is considered to be well beyond the appropriate scope of the Internet Engineering Task Force to undertake within the defined scope and mission, and a more general statement of intent would be more fitting in this context.

An enumeration of the desired attributes of a distribution system is:

The adopted distribution mechanism should be:

- efficient,
- fair,
- generally accessible and imposing no barrier to access,
- undertaken in a manner that preserves the desired characteristics of the Local Use address space,
- one that uses a fee structure that fairly reflects the costs of efficient service delivery mechanisms,
- one that allows a choice of service models where feasible,
- one that prevents distortions of the distribution function through behaviours such as hoarding or selective reselling,
- one that does not place the operator(s) in contravention to various regulatory frameworks, and

- attuned to the long-term stable use of specific instances of this resource by consumers

Comment on IANA Considerations

The Local Use Address draft proposes that:

```
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.
```

It is noted that there are significant problems with this proposed approach to directions to IANA, particularly with the noted concept that this is a for-profit activity and IANA is, in effect, being directed to be in the position of selecting a global monopoly operator. Some of the lessons learned from DNS administration over the past decade would indicate that this is not a sensible directive to pass to IANA, as it is unlikely to be reasonable implemented in this precise form.

Relationship with Existing Address Distribution Mechanisms

The Local Use proposal's desire to operate the address space without any form of discernable structure by having all block assignments be drawn from a random selection from across the entire managed space precludes the reuse of the current distribution mechanism of an IANA allocation to each of the RIRs to service their particular region. In the context of assuming that the RIRs undertake this function, the proposed mechanism would see FD00::/8 allocated to the RIRs and managed via a single registry maintained by the RIRs working together. Each RIR would lodge a "draw request" for a block from this registry in response to individual customer requests, and the registry would respond with the selected block, using a random draw function.

The potential areas of difference between the current RIR practice and the requirements here are:

- the absence of any form of justification for the allocation,
- a fixed size of allocation,
- the potential to make extensive use of automated mechanisms in the registry allocation function
- public reporting of allocations from this space only in summary form (no detailed reports, such as currently published via Whois servers)
- consideration of adoption of a service model or models relating to the terms of the assignment.
- consideration of various forms of renewable allocations and the issue of whether permanent allocations are suitable for this intended role.
- determining a fee schedule where the registry service is operated in a manner that is cost neutral to the membership.
- adoption of a transaction-based fee-for service model (as distinct from a membership service model)
- specific consideration relating to long term reliable storage of individual allocation information

In this context, the areas of RIR liaison with the IETF would appear to be in understanding the role of coordinated RIR policies in this area, and the role of the IETF. As an example, the nomination of a fee schedule and a service model in the draft proposal would normally seen as prescribing matters that would normally be determined by the RIRs through the adoption of policy proposals rather than a matter for the IETF to determine, while the consideration of permanent allocations would be a matter that would entail some substantive consideration by the IETF.

On a purely pragmatic level there is no practical way that the IETF or the adopted distribution mechanism can totally prevent these address prefixes from leaking into the IPv6 global routing space. What is, or is not, carried in the routing space is largely a matter of convention from within the operator community. If the decision is taken not to publish the details of individual Local Use unique allocations, then this would be a factor in determining whether or not blocks drawn from this space may be carried in the global routing system, but it would not absolutely prevent such use.

The service model is again a relatively challenging concept. The original IPv4 address allocation system worked on a similar basis of enduring allocations, and this has proved to be problematic in terms of recovery of unused space in more recent times. While the draft proposal is explicit about attempting to prevent short term distortions such as hoarding, there is little doubt that any form of finite unmanaged resource will be placed under consumption pressures eventually. Attempting to set a global price that makes the resource generally accessible, while still attempting to make the price a deterrent to hoarding is not a completely reasonable exercise in global terms. What would be regarded as a trivially small fee within some economies would be seen as a prohibitively expensive price in other economies. More worryingly, the concept of an enduring assignment is that there is no opportunity to make any form of correction in later times to the extant assignments, and, as in IPv4, there is the distinct risk of giving early adopters a long term advantage that may not be enjoyed by later players who may be working under more restrictive allocation polices. A shorter term lease arrangement (such as 2 - 5 years) allows for regular renewal of the relationship with the registrar, allowing for assignment information to be updated to reflect the current state of the assignee, but would entail greater levels of registry activity. As this entire operation is intended to be sufficiently low in cost that it is generally accessible, and that the value here is not in routeable address space, but in the attribute of assured uniqueness for the address space, the consideration of the level of registry activity is a critical one. It may be that the distribution mechanism adopts both service models, allowing an enduring application to be undertaken at any time at one fee level, and a shorter identity-validated application and renewal to be undertaken on a biannual basis at a lower fee, This is obviously a matter for further consideration.

Overall, it is observed that this appears to be an area of activity that is entirely appropriate for the RIRs to undertake as part of their general role of global distribution of Internet number resources. The base infrastructure to support this particular function is already installed, and the scaling to accept higher volumes of transactions is balanced by the almost complete level of automation that may be applied to the processing of each transaction.

The further steps that are suggested here for the RIRs are:

- interaction with the authors concerning the details raised in this note
- interaction within the IPv6 working group concerning the acceptability of the pseudo-uniquenenss in the proposed random self-selection mechanism
- interaction with the IESG over the IANA considerations section if/when this draft is passed from the Working Group to the IESG for publication
- consideration of any issues of address management policy that should be passed to the open policy forums of the RIRs for consideration
- development of process and tools to implement this registry

Attachment: draft-hinden-ipv6-global-local-addr-02.txt

INTERNET-DRAFT R. Hinden, Nokia
June 29, 2003 Brian Haberman, Caspian

Unique Local IPv6 Unicast Addresses

<draft-hinden-ipv6-global-local-addr-02.txt>

Status of this Memo

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.

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 Brian Carpenter, Christian Huitema, Aidan Williams, Andrew White, Michel Py, Charlie Perkins, and many others. The authors would also like to thank Brian Carpenter, Charlie Perkins, Harald Alvestrand, Keith Moore, Margaret Wasserman, Michel Py, and Shannon Behrens for their comments and suggestions on this draft.

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:

```
| n      |
| bits   |  m bits    |  16 bits  |          64 bits            |
+--------+------------+-----------+-----------------------------+
| prefix | global ID  | subnet ID |        interface ID         |
+--------+------------+-----------+-----------------------------+
```

Where:

prefix   prefix to identify Local IPv6 unicast addresses.

global ID   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 IID as defined in [ADDARCH].

There are 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] to compare the number of resulting /48 prefixes per person. A range of prefix choices is shown in the following table:

Prefix	Global ID Length	Number /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 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 choice 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. The 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 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 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 <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.

[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