Network Working Group
Internet-Draft
Expires: August 15, 2006

G. Van de Velde
C. Popoviciu
Cisco Systems
T. Chown
University of Southampton
February 11, 2006

IPv6 Unicast Address Assignment Considerations
<draft-vandevelde-v6ops-addcon-00.txt>

Status of this Memo

Copyright Notice

Abstract

One fundamental aspect of any IP communications infrastructure is its addressing plan. With its new address architecture and allocation policies, the introduction of IPv6 into a network means that network designers and operators need to reconsider their existing approaches to network addressing. Lack of guideliness on handling this aspect of network design could slow down the integration of IPv6. This

draft aims to provide the information and recommendations relevant to planning the addressing aspects of IPv6 deployments. The draft also provides IPv6 addressing case studies for both an enterprise and an ISP network. In this first version of the draft we aim to provoke discussion on this important topic; more detailed case study texts will follow.

Table of Contents

1. Introduction

The Internet Protocol Version 6 (IPv6) Addressing Architecture [10] defines three main types of addresses: unicast, anycast and multicast. This document focuses on unicast addresses, for which there are currently three principal allocated types: Global Unique Addresses [12] ('globals'), Unique Local IPv6 Addresses [22] (ULAs) and 6bone address space [3].

The document covers aspects that should be considered during IPv6 deployment for the design and planning of an addressing scheme for an IPv6 network. The network's IPv6 addressing plan may be for an IPv6-only network, or for a dual-stack infrastructure where some or all devices have addresses in both protocols. These considerations will help an IPv6 network designer to efficiently and prudently assign the IPv6 address space that has been allocated to its organization.

The address assignment considerations are analyzed separately for the two major components of the IPv6 unicast addresses, namely 'Network Level Addressing' (the allocation of subnets) and the 'Subnet Prefix' (address usage within a subnet). Thus the document includes a discussion of aspects of address assignment to nodes and interfaces in an IPv6 network. Finally the document will provide two examples of a successfully deployed address plan in a service provider (ISP) and an enterprise network.

Parts of this document highlight the differences that an experienced IPv4 network designer should consider when planning an IPv6 deployment, for example:

- IPv6 devices will more likely be multi-addressed in comparison with their IPv4 counterparts.
- The practically unlimited size of an IPv6 subnet (2^64 bits) reduces the requirement to size subnets to device counts for the purposes of (IPv4) address conservation.
- The implications of the reduced threat of address-based host scanning, as discussed in [25].

We do not discuss here how a site or ISP should proceed with acquiring its globally routable IPv6 address prefix. However, one should note that IPv6 networks receive their global unicast address allocation from their 'upstream' provider, which may be another ISP, a Local Internet Registry (LIR) or a Regional Internet Registry (RIR). In each case the prefix received is provider assigned (PA); there is currently no provider independent (PI) address space for IPv6. Thus an IPv6 network which changes provider will need to undergo a renumbering process, as described in [21]. A separate document [27] makes recommendations to ease the IPv6 renumbering

process.

2. Network Level Addressing Design Considerations

This section discusses the kind of IPv6 addresses used at the network level for the IPv6 infrastructure. The kind of addresses that can be considered are Global Unique Addresses, ULAs and 6bone address space.

2.1. Global Unique Addresses

The most commonly used unicast addresses will be Global Unique Addresses ('globals'). No significant considerations are neccesary if the organization has an address space allocation and a single prefix is deployed through a single upstream provider.

However, a multihomed site may deploy addresses from two or more Service Provider assigned IPv6 address ranges. Here, the network Administrator must have awareness on where and how these ranges are used on the multihomed infrastructure environment. The nature of the usage of multiple prefixes may depend on the reason for multihoming (e.g. resilience failover, load balancing, policy-based routing, or multihoming during an IPv6 renumbering event). IPv6 introduces improved support for multi-addressed hosts through the IPv6 default address selection methods described in RFC3484 [9]. A multihomed host may thus have two addresses, one per prefix (provider), and select source and destination addresses to use as described in that RFC.

2.2. Unique Local IPv6 Addresses

ULAs have replaced the originally conceived Site Local addresses in the IPv6 addressing architecture, for reasons described in [17]. ULAs improve on site locals by offering a high probability of the global uniqueness of the prefix used, which can be beneficial in the case of (deliberate or accidental) leakage, or where networks are merged. ULAs are akin to the private address space [1] assigned for IPv4 networks.

The ULA address range allows a network administrator to deploy IPv6 addresses on their network without asking for a globally unique registered IPv6 address range. A ULA prefix is 48 bits, i.e. a /48, the same as the currently recommended allocation for a site from the globally routable IPv6 address space [6].

ULAs provide the means to deploy a fixed addressing scheme that is not affected by a change in service provider and the corresponding PA global addresses. Internal operation of the network is thus

unaffected during renumbering events. Nevertheless, this type of address must be used with caution.

A site using ULAs may or may not also deploy globals. In an isolated network ULAs may be deployed on their own. In a connected network, that also deploys global addresses, both may be deployed, such that hosts become multiaddressed (one global and one ULA address) and the IPv6 default address selection algorithm will pick the appropriate source and destination addresses to use, e.g. ULAs will be selected where both the source and destination hosts have ULA addresses. Because a ULA and a global site prefix are both /48 length, an administrator can chgoose to use the same subnetting (and host addressing) plan for both prefixes.

As an example of the problems ULAs may cause, when using IPv6 multicast within the network, the IPv6 default address selection algorithm prefers the ULA address as the source address for the IPv6 multicast streams. This is NOT a valid option when sending an IPv6 multicast stream to the IPv6 Internet for two reasons. For one, these addresses are not globally routable so RPF checks for such traffic will fail outside the internal network. The other reason is that the traffic will likely not cross the network boundary due to multicast domain control and perimeter security policies.

In principal ULAs allow easier network mergers than RFC1918 addresses do for IPv4 because ULA prefixes have a high probability of uniqueness, if the prefix is chosen as described in the RFC.

The usage of ULAs should be carefully considered even when not attached to the IPv6 Internet due to the potential for added complexity when connecting to the Internet at some point in the future.

2.3. 6Bone Address Space

The 6Bone address space was used before the RIRs started to distribute 'production' IPv6 prefixes. The 6Bone prefixes have a common first 16 bits in the IPv6 Prefix of 3FFE::/16. This address range is deprecated as of 6th June 2006 [15] and should be avoided on any new IPv6 network deployments. Sites using 6bone address space should renumber to production address space using procedures as defined in [21].

2.4. Network Level Design Considerations

IPv6 provides network administrators with a significantly larger address space, enabling them to be very creative in how they can define logical and practical address plans. The subneting of

assigned prefixes can be done based on various logical schemes that involve factors such as:

- Geographical Boundaries - by assigning a common prefix to all subnets within a geographical area.
- Organizational Boundaries - by assigning a common prefix to an entire organization or group within a corporate infrastructure.
- Service Type - by reserving certain prefixes for predefined services such as: VoIP, Content Distribution, Internet Access, etc.

Such logical addressing plans have the potential to simplify network operations and service offerings, and to simplify network management and troubleshooting. A very large network would also have no need to consider using private address space for its infrastructure devices, simplifying network management.

The network designer must however keep in mind several factors when developing these new addressing schemes:

- Prefix Aggregation - The larger IPv6 addresses can lead to larger routing tables unless network designers are actively pursuing aggregation. While prefix aggregation will be enforced by the service provider, it is beneficial for the individual organizations to observe the same principles in their network design process.
- Network growth - The allocation mechanism for flexible growth of a network prefix, documented in RFC3531 [11] can be used to allow the network infrastructure to grow and be numbered in a way that is likely to preserve aggregation (the plan leaves 'holes' for growth).
- ULA usage in large networks - Networks which have a large number of 'sites' that each deploy a ULA prefix which will by default be a 'random' /48 under fc00::/7 will have no aggregation of those prefixes. Thus the end result may be cumbersome because the network will have large amounts of non-aggregated ULA prefixes.

2.4.1. Sizing the Network Allocation

We do not discuss here how a network designer sizes their application for address space. By default a site will receive a /48 prefix [6]. The default provider allocation via the RIRs is currently a /32 [26]. These allocations are indicators for a first allocation for a network. Different sizes may be obtained based on the anticipated address usage [26]. There are examples of allocations as large as /19 having been made from RIRs to providers at the time of writing.

2.4.2. Address Space Conservation

Despite the large IPv6 address space which enables easier subneting, it still is important to ensure an efficient use of this resource.

Some addressing schemes, while facilitating aggregation and management, could lead to significant numbers of addresses being unused. Address conservation requirements are less stringent in IPv6 but they should still be observed.

The proposed HD [7] value for IPv6 is 0.94 compared to the current value of 0.96 for IPv4.

3. Subnet Prefix Considerations

This section analyzes the considerations applied to define the subnet prefix of the IPv6 addresses. The boundaries of the subnet prefix allocation are specified in RFC3513 [10]. In this document we analyze their practical implications. Based on RFC3513 [10] it is legal for any IPv6 unicast address starting with binary address '000' to have a subnet prefix larger than, smaller than or of equal to 64 bits. Each of these three options are discussed in this document.

3.1. Considerations for subnet prefixes shorter then /64

An allocation of a prefix shorter then 64 bits to a node or interface is bad practice. The shortest subnet prefix that could theoretically be assigned to an interface or node is limited by the size of the network prefix allocated to the organization.

A possible reason for choosing the subnet prefix for an interface shorter then /64 is that it would allow more nodes to be attached to that interface compared to a prescribed length of 64 bits. This however is unnecessary considering that 2^64 provides plenty of node addresses for a well designed IPv6 network. Layer two technologies are unlikely to support such large numbers of nodes within a single link (e.g. Ethernet limited to 48-bits of hosts)

The subnet prefix assignments can be made either by manual configuration, by a stateful Host Configuration Protocol [8] or by a stateful prefix delegation mechanism [14].

3.2. Considerations for /64 prefixes

Based on RFC3177 [6], 64 bits is the prescribed subnet prefix length to allocate to interfaces and nodes.

When using a /64 subnet length, the address assignment for these addresses can be made either by manual configuration, by a stateful Host Configuration Protocol [8] [16] or by stateless autoconfiguration [2].

Note that RFC3177 strongly prescribes 64 bit subnets for general usage, and that stateless autoconfiguration option is only defined for 64 bit subnets.

3.3. Considerations for subnet prefixes longer then /64

Address space conservation is the main motivation for using a subnet prefix length longer than 64 bits.

The address assignment can be made either by manual configuration or by a stateful Host Configuration Protocol [8].

When assigning a subnet prefix of more then 80 bits, according to RFC3513 [10] "u" and "g" bits (respectively the 81st and 82nd bit) need to be taken into consideration and should be set correctly. In currently implemented IPv6 protocol stacks, the relevance of the "u" (universal/local) bit and "g" (the individual/group) bit are marginal and typically will not show an issue when configured wrongly, however future implementations may turn out differently.

When using subnet lengths longer then 64 bits, it is important to avoid selecting addresses that may have a predefined use and could confuse IPv6 protocol stacks. The alternate usage may not be a simple unicast address in all cases. The following points should be considerated when selecting a subnet length longer then 64 bits subnet prefix length.

3.3.1. Subnet Router Anycast address

RFC3513 [10] stated that within each subnet, the highest 128 interface identifier values are reserved for assignment as subnet anycast addresses.

The construction of a reserved subnet anycast address depends on the type of IPv6 addresses used within the subnet, as indicated by the format prefix in the addresses.

The first type of Subnet Router Anycast addresses have been defined as follows for EUI-64 format:

The anycast address structure implies that it is important to avoid

creating a subnet prefix where the bits 65 to 121 are defined as "1111110111...111" (57 bits in total) so that confusion can be avoided.

For other IPv6 address types (that is, with format prefixes other than those listed above), the interface identifier is not in EUI-64 format and may be other than 64 bits in length; these reserved subnet anycast addresses for such address types are constructed as follows:

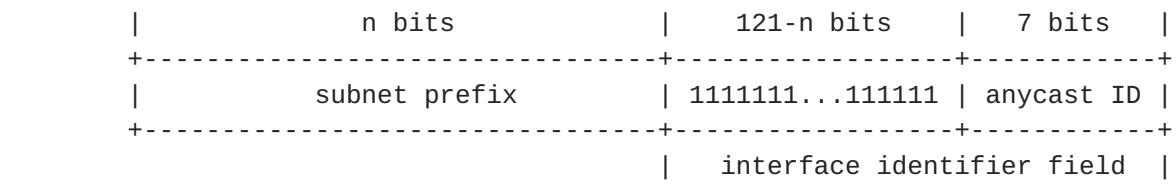

In the case discussed above there is no additional dependancy for the subnet prefix with the exception of the EUI-64 and an IID dependency. These will be discussed later in this document.

3.3.2. Addresses used by Embedded-RP (RFC3956)

Embedded-RP [18] reflects the concept of integrating the Rendezvous Point (RP) IPv6 address into the IPv6 multicast group address. Due to this embedding and the fact that the length of the IPv6 address AND the IPv6 multicast address are 128 bits, it is not possible to have the complete IPv6 address of the multicast RP embedded as such.

This limitation resulted in a restriction of 15 possible multicast addresses per subnet prefix. The space assigned for the embedded-RP is based on the 4 low order bits, while the remainder of the Interface ID is set to all '0'.

[IPv6-prefix (64 bits)][60 bits all '0'][RIID]

Where: [RIID] = 4 bit.

This leads to the constraint that when creating subnet lengths longer than 64 bits, the bits between bit 65 and the subnet boundary should not be set to be all "0".

3.3.3. ISATAP addresses

ISATAP [24] is an automatic tunneling protocol used to provide IPv6 connectivity over an IPv4 campus or enterprise environment. In order

to leverage the underlying IPv4 infrastructure, the IPv6 addresses are constructed in a special format.

An IPv6 ISATAP [24] address has the IPv4 address embedded, based on a predefined structure policy that identifies them as an ISATAP [24] address.

[IPv6 Prefix (64 bits)][0000:5EFE][IPv4 address]

When using subnet prefix length longer then 64 bits it is recommended that that the portion of the IPv6 prefix from bit 65 to the end of the subnet prefix does not match with the welknown ISATAP [0000:5EFE] address portion.

3.3.4. /126 addresses

The 126 bit subnet prefixes are typically used for point-to-point links similar to the RFC3021 [4] recommendations for IPv4. The usage of this subnet address length does not lead to any additional considerations other than the ones discussed earlier in this section, particularly those related to the "u" and "g" bits.

3.3.5. /127 addresses

The usage of the /127 addresses is not valid and should be strongly discouraged as documented in RFC3627 [13].

3.3.6. /128 addresses

The 128 bit address prefix may be used in those situations where we know that one, and only one address is sufficient. Example usage would be the offlink loopback address of a network device.

When choosing a 128 bit prefix, it is recommended to take the "u" and "g" bits into consideration and to make sure that there is no overlap with either the following well known addresses:

- Subnet Router Anycast Address
- Reserved Subnet Anycast Address
- Addresses used by Embedded-RP
- ISATAP Addresses

4. Allocation of the IID of an IPv6 Address

In order to have a complete IPv6 address, an interface must be associated a prefix and an Interface Identifier (IID). Section 3 of

this document analyzed the prefix selection considerations. This section discusses the elements that should be considered when assigning the IID portion of the IPv6 address.

There are various ways to allocate an IPv6 address to a device or interface. The option with the least amount of caveats for the network administrator is that of EUI-64 [2] based addresses. For the manual or dynamic options, the overlap with well known IPv6 addresses should be avoided.

4.1. Automatic EUI-64 Format Option

When using this method the network administrator has to allocate a valid 64 bit subnet prefix. The EUI-64 [2] allocation procedure can from that moment onwards assign the remaining 64 IID bits in a stateless manner. All the considerations for selecting a valid IID have been incorporated in the EUI-64 methodology.

4.2. Using Privacy Extensions

The main purpose of IIDs generated based on RFC3041 [5] is to provide privacy to the entity using this address. While there is no particular restraints in the usage of these addresses as defined in [5] there are some implications to be aware of when using privacy addresses as documented in section 4 of RFC3041 [5]:

- The privacy extension algoritm may complicate flexibility in future transport protocols
- These addresses may add complexity to the operational management and troubleshooting of the infrastructure (i.e. which address belongs to which real host)
- A reverse DNS lookup check may be broken when using privacy extensions

4.3. Cryptographically Generated IPv6 Addresses

Cryptographically Generated Addresses (CGAs) are based upon RFC3972 [20] and provide a method for binding a public signature key to an IPv6 address in the Secure Neighbor Discovery (SEND) protocol [19].

The basic idea is to generate the interface identifier (i.e. the rightmost 64 bits) of the IPv6 address by computing a cryptographic hash of the public key. The resulting IPv6 address is called a cryptographically generated address (CGA). The corresponding private key can then be used to sign messages sent from that address.

Implications to be aware of when using CGA addresses are found in section 7 of RFC3972 [20]:

- When using CGA addresses the values of the "u" and "g" bits are ignored however it does not add any security or implementation implications
- There is no mechanism for proving that an address is not a CGA
- When it is discovered that a node has been compromised, a new signature key and a new CGA SHOULD be generated

Due to the fact that CGA generated addresses are indistinguishable from a privacy address the same considerations as with privacy addresses are also valid for CGA generated addresses.

4.4. Manual/Dynamic Assignment Option

This section discusses those IID allocations that are not implemented through stateless address configuration (Section 4.1). They are applicable regardless of the prefix length used on the link. It is out of scope for this section to discuss the various assignment methods (e.g. manual configuration, DHCPv6, etc).

In this situation the actual allocation is done by human intervention and consideration needs to be given to the complete IPv6 address so that it does not result in overlaps with any of the well known IPv6 addresses:

- Subnet Router Anycast Address
- Reserved Subnet Anycast Address
- Addresses used by Embedded-RP
- ISATAP Addresses

When using an address assigned by human intervention it is recommended to choose IPv6 addresses which are not abvious to guess and/or avoid any IPv6 addresses that embed IPv4 addresses used in the current infrastructure. Following these two recommendations will make it more difficult for malicious third parties to guess targets for attack, and thus reduce security threats to a certain extent.

5. Case Studies

tbc.

5.1. Enterprise Considerations

In this section we consider a case study of a campus network that is deploying IPv6 in parallel with existing IPv4 protocols in a dual-stack environment. The specific example is the University of Southampton (UK). The case study is a 'work in progress' as the deployment is an evolving one, currently covering around 1,500 hosts.

5.1.1. Obtaining general IPv6 network prefixes

In the case of a campus network, the site will typically take its connectivity from its National Research and Education Network (NREN). Southampton connects to JANET, the UK academic network. JANET currently has a /32 allocation from RIPE of 2001:630::/32. The current recommended practice is for sites to receive a /48 allocation, and on this basis Southampton has received such a prefix for its own use, specifically 2001:630:d0::/48.

No ULA addressing is used on site. The campus does not expect to change service provider, and thus does not plan to use ULAs for the (perceived) benefit of easing network renumbering. Indeed, the campus has renumbered following the aforementioned renumbering procedure [21] on two occassions, and this has proven adequate (with provisos documented in [27]. We also do not see any need to deploy ULAs for in or out of band network management; there are enough IPv6 prefixes available in the site allocation for the infrastructure.

No 6bone addressing is used on site. This was phased out some time ago. We note that as of 6th June 2006 transit ISPs will likely filter any attempted use of such prefixes.

Southampton does participate in global and organisation scope IPv6 multicast networks. Multicast address allocations are not discussed here as they are not in scope for the document. Embedded RP is in use, and has been tested successfully across providers between sites.

5.1.2. Forming an address (subnet) allocation plan

The campus has a /16 prefix for IPv4 use; in principle 256 subnets of 256 addresses. In reality the subnetting is muddier, because of concerns of IPv4 address conservation; subnets are sized to the hosts within them, e.g. a /26 IPv4 prefix is used if a subnet has 35 hosts in it. While this is efficient, it increases management burden when physical deployments change, and IPv4 subnets require resizing (up or down), even with DHCP in use.

The /48 IPv6 prefix is considerably larger than the IPv4 allocation already in place at the site. It is loosely equivalent to a 'Class A' IPv4 prefix in that it has 2^16 (over 65,000) subnets, but has an effectively unlimited subnet address size (2^64) compared to 256 in the IPv4 equivalent. The increased subnet size means that /64 IPv6 prefixes can be used on all subnets, without any requirement to resize them at a later date. The increased subnet volume allows subnets to be allocated more generously to schools and departments in the campus. While address conservation is still important, it is no longer an impediment on network management. Rather, address (subnet)

allocation is more about planning for future expansion.

In a dual-stack network, we chose to deploy our IP subnets congruently for IPv4 and IPv6. This is because the systems are still in the same administrative domains and the same geography. We do not expect to have IPv6-only subnets in production use for a while yet, outside testbeds and our early Mobile IPv6 trials. The firewall would ideally be a single dual-stack device with consistent policies (by host rather than IP version), however this is currently implemented as a firewall per IP protocol due to vendor limitations (Nokia/Checkpoint for IPv4, BSD pf tool for IPv6).

The subnet allocation plan required a division of the address space per school or department. Here a /56 was allocated to the school level of the university; there are around 30 schools currently. Further allocations were made for central IT infrastructure, for the network infrastructure and the server side systems.

5.1.3. Other considerations

The network uses a Demilitarized Zone (DMZ) topology for some level of protection of 'public' systems. Again, this topology is congruent with the IPv4 network.

There are no specific transition methods deployed internally to the campus; everything is using the conventional dual-stack approach. There is no use of tools such as ISATAP for example.

For the Mobile IPv6 early trails, we have allocated one prefix for Home Agent (HA) use. We have not yet considered how Mobile IPv6 usage may grow, and whether more or even every subnet will require HA support.

The university operates a tunnel broker service on behalf of UKERNA. This uses separate address space from JANET, not the main university allocation.

5.1.4. Node configuration considerations

We currently use stateless autoconfiguration on most subnets for IPv6 hosts. There is no DHCPv6 service deployed yet, beyond tests of early code releases. We do seek a common integrated DHCP/DNS management platform, even if the servers themselves are not colocated. Currently we add client statelessly autoconfigured addresses to the DNS manually. Our administrators would prefer the use of DHCP because they believe it gives them some management control.

Regarding the [25] implications, we note that all our hosts are dual-stack, and thus are potentially exposed over both protocols anyway. We publish all addresses in DNS, and do not operate a two faced DNS.

We have internal usage of RFC3041 privacy addresses currently, but may wish to administratibely disable this (perhaps via DHCP), but we need to determine the feasibility of this on all systems, e.g. for WLAN guests or other user-maintained systems. Network management should be simpler without RFC3041 in opeation. Note RFC3041 is only an issue for outbound connections.

We manually configure server addresses to avoid address changes on a change of network adaptor. With IPv6 you can choose to pick ::53 for a DNS server, or can pick 'random' addresses for obfuscation, though that's not an issue for publicly advertised addresses (dns, mx, web, etc).

5.1.5. Observations

The site is not (yet) using prefix delegation tools for IPv6.

5.2. Service Provider Considerations

tbc.

6. Security Considerations

This IPv6 addressing documents does not have any direct impact on Internet infrastructure security.

7. References

7.1. Normative References

7.2. Informative References

[1] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[2] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.

[3] Hinden, R., Fink, R., and J. Postel, "IPv6 Testing Address Allocation", RFC 2471, December 1998.

[4] Retana, A., White, R., Fuller, V., and D. McPherson, "Using 31-Bit Prefixes on IPv4 Point-to-Point Links", RFC 3021, December 2000.

[5] Narten, T. and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 3041, January 2001.

[6] IAB and IESG, "IAB/IESG Recommendations on IPv6 Address Allocations to Sites", RFC 3177, September 2001.

[7] Durand, A. and C. Huitema, "The H-Density Ratio for Address Assignment Efficiency An Update on the H ratio", RFC 3194, November 2001.

[8] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.

[9] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, February 2003.

[10] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", RFC 3513, April 2003.

[11] Blanchet, M., "A Flexible Method for Managing the Assignment of Bits of an IPv6 Address Block", RFC 3531, April 2003.

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

[13] Savola, P., "Use of /127 Prefix Length Between Routers Considered Harmful", RFC 3627, September 2003.

[14] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.

[15] Fink, R. and R. Hinden, "6bone (IPv6 Testing Address Allocation) Phaseout", RFC 3701, March 2004.

[16] Droms, R., "Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6", RFC 3736, April 2004.

[17] Huitema, C. and B. Carpenter, "Deprecating Site Local Addresses", RFC 3879, September 2004.

[18] Savola, P. and B. Haberman, "Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address", RFC 3956,

November 2004.

[19] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.

[20] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, March 2005.

[21] Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day", RFC 4192, September 2005.

[22] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.

[23] Templin, F., Gleeson, T., Talwar, M., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", draft-ietf-ngtrans-isatap-24 (work in progress), January 2005.

[24] Templin, F., Gleeson, T., Talwar, M., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (draft-ietf-ngtrans-isatap-24.txt)", July 2005.

[25] Chown, T., "IPv6 Implications for TCP/UDP Port Scanning (chown-v6ops- port-scanning-implications-02.txt)", October 2005.

[26] APNIC, ARIN, RIPE NCC, "IPv6 Address Allocation and Assignment Policy (www.ripe.net/ripe/docs/ipv6policy.html)", January 2003.

[27] Chown, T., Thompson, M., Ford, A., and S. Venaas, "Things to think about when Renumbering an IPv6 network (draft-chown-v6ops-renumber-thinkabout-03.txt)", July 2005.

Authors' Addresses

Gunter Van de Velde
Cisco Systems
De Kleetlaan 6a
Diegem 1831
Belgium

Phone: +32 2704 5473
Email: gunter@cisco.com

Ciprian Popoviciu
Cisco Systems
7025-6 Kit Creek Road
Research Triangle Park, North Carolina PO Box 14987
USA

Phone: +1 919 392-3723
Email: cpopovic@cisco.com

Tim Chown
University of Southampton
Highfield
Southampton, SO17 1BJ
United Kingdom

Phone: +44 23 8059 3257
Email: tjc@ecs.soton.ac.uk

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

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