IPv6 Unicast Address Assignment Considerations

Status of this Memo

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 guidelines on handling this aspect of network design could slow down the deployment and integration of IPv6. This document aims to provide the information and recommendations

relevant to planning the addressing aspects of IPv6 deployments. The
document also provides IPv6 addressing case studies for both an
enterprise and an ISP network.

Table of Contents

---

1. Introduction

The Internet Protocol Version 6 (IPv6) Addressing Architecture [RFC4291] (Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture," February 2006.) defines three main types of addresses: unicast, anycast and multicast. This document focuses on unicast addresses, for which there are currently two principal allocated types: Globally Unique Addresses [RFC3587] (Hinden, R., Deering, S., and E. Nordmark, "IPv6 Global Unicast Address Format," August 2003.) ('globals') and Unique Local IPv6 Addresses [RFC4193] (Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses," October 2005.) (ULAs). In addition until recently there has been 'experimental' 6bone address space [RFC3701] (Fink, R. and R. Hinden, "6bone (IPv6 Testing Address Allocation) Phaseout," March 2004.), though its use has been deprecated since June 2006 [RFC3701] (Fink, R. and R. Hinden, "6bone (IPv6 Testing Address Allocation) Phaseout," March 2004.).

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 their 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 'interface-id' (the identification of the interface 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 provides two examples of deployed address plans 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 vastly increased subnet size on the
 threat of address-based host scanning and other scanning
 techniques, as discussed in [RFC5157] (Chown, T., "IPv6
 Implications for Network Scanning," March 2008.).

We do not discuss here how a site or ISP should proceed with acquiring
its globally routable IPv6 address prefix. In each case the prefix
received is either provider assigned (PA) or provider independent (PI).
We do not discuss PI policy here. The observations and recommendations
of this text are largely independent of the PA or PI nature of the
address block being used. At this time we assume that most commonly an
IPv6 network which changes provider will need to undergo a renumbering
process, as described in [RFC4192] (Baker, F., Lear, E., and R. Droms,
"Procedures for Renumbering an IPv6 Network without a Flag Day,"
September 2005.). A separate document [THINKABOUT] (Chown, T.,
Thompson, M., Ford, A., and S. Venaas, "Things to think about when
Renumbering an IPv6 network (draft-chown-v6ops-renumber-
thinkabout-05.txt)," March 2007.) makes recommendations to ease the
IPv6 renumbering process.
This document does not discuss implementation aspects related to the
transition between the ULA addresses and the now obsoleted site-local
addresses. Some implementations know about Site-local addresses even
though they are deprecated, and do not know about ULAs - even though
they represent current specification. As result transitioning between
these types of addresses may cause difficulties.

---

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 Globally Unique Addresses and ULAs. We also comment here
on the deprecated 6bone address space.

2.1. Globally Unique Addresses

The most commonly used unicast addresses will be Globally Unique
Addresses ('globals'). No significant considerations are necessary if
the organization has an address space assignment 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 [RFC3484] (Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)," February 2003.). A multihomed host may thus have two or more addresses, one per prefix (provider), and select source and destination addresses to use as described in that RFC. However multihoming also has some operational and administrative burdens besides chosing multiple addresses per interface [RFC4219] (Lear, E., "Things Multihoming in IPv6 (MULTI6) Developers Should Think About," October 2005.)[RFC4218] (Nordmark, E. and T. Li, "Threats Relating to IPv6 Multihoming Solutions," October 2005.).

2.2. Unique Local IPv6 Addresses

ULAs have replaced the originally conceived Site Local addresses in the IPv6 addressing architecture, for reasons described in [RFC3879] (Huitema, C. and B. Carpenter, "Deprecating Site Local Addresses," September 2004.). 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 [RFC1918] (Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets," February 1996.) assigned for IPv4 networks, except that in IPv6 networks we may expect to see ULAs used alongside global addresses, with ULAs used internally and globals used externally. Thus use of ULAs does not imply use of NAT for IPv6.
The ULA address range allows network administrators 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 [RFC3177] (IAB and IESG, "IAB/IESG Recommendations on IPv6 Address Allocations to Sites," September 2001.).
A site willing to use ULA address space can have either (a) multiple / 48 prefixes (e.g. a /44) and wishes to use ULAs, or (b) has one /48 and wishes to use ULAs or (c) a site has a less-than-/48 prefix (e.g. a /56 or /64) and wishes to use ULAs. In all above cases the ULA addresses can be randomly chosen according the principles specified in [RFC4193] (Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses," October 2005.). However, in case (a) the use of randomly chosen ULA addresses will provide suboptimal aggregation capabilities.
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 global addresses. 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 choose 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 Reverse Path Forwarding 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 principle 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.

2.3. 6Bone Address Space

The 6Bone address space was used before the Regional Internet
Registries (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
[RFC3701] (Fink, R. and R. Hinden, "6bone (IPv6 Testing Address
Allocation) Phaseout," March 2004.) and must not be used on any new
IPv6 network deployments. Sites using 6bone address space should
renumber to production address space using procedures as defined in
[RFC4192] (Baker, F., Lear, E., and R. Droms, "Procedures for
Renumbering an IPv6 Network without a Flag Day," September 2005.).

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 subnetting of assigned

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

*Using existing systems

    -translate the existing subnet number into IPv6 subnet id

    -translate the VLAN id into IPv6 subnet id

*Redesign

    -allocate according to your need

*Aggregation

    -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, wireless services, Internet Access, Security areas etc. This type of addressing may create dependencies on IP addresses that can make renumbering harder if the nodes or interfaces supporting those services on the network are sparse within the topology.

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 for networks with and without global connectivity:

*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 [RFC3531] (Blanchet, M., "A Flexible Method for Managing the Assignment of Bits of an IPv6 Address Block," April 2003.) 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.
 However, there is no rule to disallow large networks to use a
 single ULA prefix for all 'sites', as a ULA still provides 16
 bits for subnetting to be used internally

*It is possible that as registry policies evolve, a small site may
 experience an increase in prefix length when renumbering, e.g.
 from /48 to /56. For this reason, the best practice is number
 subnets compactly rather than sparsely, and to use low-order bits
 as much as possible when numbering subnets. In other words, even
 if a /48 is allocated, act as though only a /56 is available.
 Clearly, this advice does not apply to large sites and
 enterprises that have an intrinsic need for a /48 prefix.

*A small site may want to enable routing amongst interfaces
 connected to a gateway device. For example, a residential gateway
 which receives a /48, is situated in a home with multiple LANs of
 different media types (sensor network, wired, wifi, etc.), or has
 a need for traffic segmentation (home, work, kids, etc.) and
 could benefit greatly from multiple subnets and routing in IPv6.
 Ideally, residential networks would be given an address range of
 a /48 or /56 [reference2] (APNIC, ARIN, RIPE NCC, "www.ripe.net/
 ripe/docs/ipv6policy.html," July 2007.) such that multiple /64
 subnets could be used within the residence.

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
[RFC3177] (IAB and IESG, "IAB/IESG Recommendations on IPv6 Address
Allocations to Sites," September 2001.) , however different RIR service
regions policies may suggest alternative default assignments or let the
ISPs to decide on what they believe is more appropriate for their
specific case [ARIN] (ARIN, "http://www.arin.net/policy/
nrpm.html#six54," .). The default provider allocation via the RIRs is
currently a /32 [reference2] (APNIC, ARIN, RIPE NCC, "www.ripe.net/
ripe/docs/ipv6policy.html," July 2007.). These allocations are
indicators for a first allocation for a network. Different sizes may be
obtained based on the anticipated address usage [reference2] (APNIC,
ARIN, RIPE NCC, "www.ripe.net/ripe/docs/ipv6policy.html," July 2007.).
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 subnetting, 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 Host-Density (HD) [RFC3194] (Durand, A. and C. Huitema, "The H-Density Ratio for Address Assignment Efficiency An Update on the H ratio," November 2001.) value for IPv6 is 0.94 compared to the current value of 0.96 for IPv4. Note that for IPv6 HD is calculated for sites (e.g. on a basis of /48), instead of based on addresses like with IPv4.

3. Subnet Prefix Considerations

An important part of an IPv4 addressing plan is deciding the length of each subnet prefix. Unlike in IPv4, the IPv6 addressing architecture [RFC4291] (Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture," February 2006.) specifies that all subnets using Globally Unique Addresses and ULAs always have the same prefix length of 64 bits. (This applies also to the deprecated 6Bone and Site Local addresses.)

The only exception to this rule are special addresses starting with the binary value 000, such as IPv4-Compatible IPv6 Addresses. These exceptions are largely beyond the scope of this document.

Using a subnet prefix length other than a /64 will break many features of IPv6, amongst other things Neighbor Discovery (ND), Secure Neighborship Discovery (SEND) [RFC3971] (Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)," March 2005.), privacy extensions [RFC4941] (Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6," September 2007.), parts of Mobile IPv6 [RFC4866] (Arkko, J., Vogt, C., and W. Haddad, "Enhanced Route Optimization for Mobile IPv6," May 2007.), PIM-SM with Embedded-RP [RFC3956] (Savola, P. and B. Haberman, "Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address," November 2004.), and SHIM6 [SHIM6] (IETF, "http://www.ietf.org/html.charters/shim6-charter.html," .). A number of other features currently in development, or being proposed, also rely on /64 subnet prefixes.

Nevertheless, many IPv6 implementations do not prevent the administrator from configuring a subnet prefix length shorter or longer

than 64 bits. Using subnet prefixes shorter than /64 would rarely be useful; see Appendix B.1 for discussion.

However, some network administrators have used prefixes longer than /64 for links connecting routers, usually just two routers on a point-to-point link. On links where all the addresses are assigned by manual configuration, and all nodes on the link are routers (not end hosts) that are known by the network administrators do not need any of the IPv6 features that rely on /64 subnet prefixes, this can work. Using subnet prefixes longer than /64 are not recommended for general use, and using them for links containing end hosts would be an especially bad idea, as it is difficult to predict what IPv6 features the hosts will use in the future.

Appendix B.2 describes some practical considerations that need to be taken into account when using prefixes longer than /64 in limited cases. In particular, a number of IPv6 features use interface identifiers that have a special form (such as a certain fixed value in some bit positions). When using prefixes longer than /64, it is prudent to avoid certain subnet prefix values so that nodes who assume that the prefix is /64 will not incorrectly identify the addresses in that subnet as having a special form. Appendix B.2 describes the subnet prefix values that are currently believed to be potentially problematic; however, the list is not exhaustive and can be expected to grow in the future.

Using /64 subnets is strongly recommended, also for links connecting only routers. A deployment compliant with the current IPv6 specifications cannot use other prefix lengths. However, the V6OPS WG believes that despite the drawbacks (and a potentially expensive network redesign, if IPv6 features relying on /64 subnets are needed in the future), that some networks administrators will use prefixes longer than /64.

3.1. Considerations for /64 Prefixes

Based on RFC3177 [RFC3177] (IAB and IESG, "IAB/IESG Recommendations on IPv6 Address Allocations to Sites," September 2001.), 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 [RFC3315] (Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)," July 2003.) [RFC3736] (Droms, R., "Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6," April 2004.) or by stateless autoconfiguration [RFC4862] (Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration," September 2007.).

Note that RFC3177 strongly prescribes 64 bit subnets for general usage, and that stateless autoconfiguration option is only defined for 64 bit subnets. While in theory it might be possible that some future autoconfiguration mechanisms would allow longer than 64 bit prefix lengths to be used, the use of such prefixes is not recommended at this time.

3.2. 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 [RFC4862] (Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration," September 2007.) based addresses. For the manual or dynamic options, the overlap with well known IPv6 addresses should be avoided.

3.2.1. Automatic EUI-64 Format Option

When using this method the network administrator has to allocate a valid 64 bit subnet prefix. The EUI-64 [RFC4862] (Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration," September 2007.) allocation procedure can from that moment onward 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.

3.2.2. Using Privacy Extensions

The main purpose of IIDs generated based on RFC4941 [RFC4941] (Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6," September 2007.) is to provide privacy to the entity using this address. While there are no particular constraints in the usage of these addresses as defined in [RFC4941] (Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6," September 2007.) there are some implications to be aware of when using privacy addresses as

documented in section 4 of RFC4941 [RFC4941] (Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6," September 2007.)

3.2.3. 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 (Appendix B.2.5.1.)

   *Reserved Subnet Anycast Address (Appendix B.2.5.2.)

   *Addresses used by Embedded-RP (Appendix B.2.6.)

   *ISATAP Addresses (Appendix B.2.7.)

When using an address assigned by human intervention it is recommended to choose IPv6 addresses which are not obvious 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.

3.3. IANA Considerations

There are no extra IANA consideration for this document.

3.4. Security Considerations

This document doesn't add any new security considerations that aren't already outlined in the security considerations of the references.
It must be noted that using subnet prefixes other than /64 breaks security mechanisms such as Cryptographically Generated Addresses

(CGAs) and Hash Based Addresses (HBAs), and thus makes it impossible to use protocols that depend on them.

3.5. Acknowledgements

---

4. References

4.1. Normative References

4.2. Informative References

[RFC1918]	Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets," BCP 5, RFC 1918, February 1996 (TXT).
[RFC2526]	Johnson, D. and S. Deering, "Reserved IPv6 Subnet Anycast Addresses," RFC 2526, March 1999 (TXT).
[RFC3021]	Retana, A., White, R., Fuller, V., and D. McPherson, "Using 31-Bit Prefixes on IPv4 Point-to-Point Links," RFC 3021, December 2000 (TXT).
[RFC3053]	Durand, A., Fasano, P., Guardini, I., and D. Lento, "IPv6 Tunnel Broker," RFC 3053, January 2001 (TXT).
[RFC3056]	Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds," RFC 3056, February 2001 (TXT).
[RFC3177]	IAB and IESG, "IAB/IESG Recommendations on IPv6 Address Allocations to Sites," RFC 3177, September 2001 (TXT).
[RFC3180]	Meyer, D. and P. Lothberg, "GLOP Addressing in 233/8," BCP 53, RFC 3180, September 2001 (TXT).
[RFC3194]	Durand, A. and C. Huitema, "The H-Density Ratio for Address Assignment Efficiency An Update on the H ratio," RFC 3194, November 2001 (TXT).

[RFC3315]	Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)," RFC 3315, July 2003 (TXT).
[RFC3484]	Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)," RFC 3484, February 2003 (TXT).
[RFC3531]	Blanchet, M., "A Flexible Method for Managing the Assignment of Bits of an IPv6 Address Block," RFC 3531, April 2003 (TXT).
[RFC3587]	Hinden, R., Deering, S., and E. Nordmark, "IPv6 Global Unicast Address Format," RFC 3587, August 2003 (TXT).
[RFC3627]	Savola, P., "Use of /127 Prefix Length Between Routers Considered Harmful," RFC 3627, September 2003 (TXT).
[RFC3633]	Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6," RFC 3633, December 2003 (TXT).
[RFC3701]	Fink, R. and R. Hinden, "6bone (IPv6 Testing Address Allocation) Phaseout," RFC 3701, March 2004 (TXT).
[RFC3736]	Droms, R., "Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6," RFC 3736, April 2004 (TXT).
[RFC3879]	Huitema, C. and B. Carpenter, "Deprecating Site Local Addresses," RFC 3879, September 2004 (TXT).
[RFC3956]	Savola, P. and B. Haberman, "Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address," RFC 3956, November 2004 (TXT).
[RFC3971]	Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)," RFC 3971, March 2005 (TXT).
[RFC4192]	Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day," RFC 4192, September 2005 (TXT).
[RFC4193]	Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses," RFC 4193, October 2005 (TXT).
[RFC4218]	Nordmark, E. and T. Li, "Threats Relating to IPv6 Multihoming Solutions," RFC 4218, October 2005 (TXT).
[RFC4219]	Lear, E., "Things Multihoming in IPv6 (MULTI6) Developers Should Think About," RFC 4219, October 2005 (TXT).
[RFC4271]	Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," RFC 4271, January 2006 (TXT).
[RFC4291]	Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture," RFC 4291, February 2006 (TXT).
[RFC4477]	Chown, T., Venaas, S., and C. Strauf, "Dynamic Host Configuration Protocol (DHCP): IPv4 and IPv6 Dual-Stack Issues," RFC 4477, May 2006 (TXT).
[RFC4798]	

	De Clercq, J., Ooms, D., Prevost, S., and F. Le Faucheur, "Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)," RFC 4798, February 2007 (TXT).
[RFC4862]	Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration," RFC 4862, September 2007 (TXT).
[RFC4866]	Arkko, J., Vogt, C., and W. Haddad, "Enhanced Route Optimization for Mobile IPv6," RFC 4866, May 2007 (TXT).
[RFC4941]	Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6," RFC 4941, September 2007 (TXT).
[RFC5214]	Templin, F., Gleeson, T., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)," RFC 5214, March 2008 (TXT).
[RFC5157]	Chown, T., "IPv6 Implications for Network Scanning," RFC 5157, March 2008 (TXT).
[SHIM6]	IETF, "http://www.ietf.org/html.charters/shim6-charter.html."
[ARIN]	ARIN, "http://www.arin.net/policy/nrpm.html#six54."
[reference2]	APNIC, ARIN, RIPE NCC, "www.ripe.net/ripe/docs/ipv6policy.html," July 2007.
[reference3]	APNIC, ARIN, RIPE NCC, "http://www.ripe.net/ripe/docs/ripe-412.html," July 2007.
[reference4]	ARIN, "http://www.arin.net/policy/nrpm.html#ipv6," March 2008.
[reference5]	APNIC, "http://www.apnic.net/policy/ipv6-address-policy.html," March 2007.
[reference6]	LACNIC, "http://lacnic.net/en/politicas/ipv6.html."
[reference7]	AFRINIC, "http://www.afrinic.net/docs/policies/afpol-v6200407-000.htm," March 2004.
[THINKABOUT]	Chown, T., Thompson, M., Ford, A., and S. Venaas, "Things to think about when Renumbering an IPv6 network (draft-chown-v6ops-renumber-thinkabout-05.txt)," March 2007.

Appendix A. Case Studies

This appendix contains two case studies for IPv6 addressing schemas that have been based on the statements and considerations of this draft. These case studies illustrate how this draft has been used in two specific network scenarios. The case studies may serve as basic considerations for an administrator who designs the IPv6 addressing schema for an enterprise or ISP network, but are not intended to serve

as general design proposal for every kind of IPv6 network. All subnet sizes used in this appendix are for practical visualization and do not dictate RIR policy.

A.1. Enterprise Considerations

In this section one considers 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), focusing on a large department within that network. The deployment currently spans around 1,000 hosts and over 1,500 users.

A.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, via its local regional network LeNSE. JANET currently has a /32 allocation from RIPE NCC. 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. The regional network also uses its own allocation from the NREN provider.

No ULA addressing is used on site. The campus is not multihomed (JANET is the sole provider), nor does it 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 [RFC4192] (Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day," September 2005.) on two occasions, and this has proven adequate (with provisos documented in [THINKABOUT] (Chown, T., Thompson, M., Ford, A., and S. Venaas, "Things to think about when Renumbering an IPv6 network (draft-chown-v6ops-renumber-thinkabout-05.txt)," March 2007.). The campus 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. In some cases, use of private IP address space in IPv4 creates problems, so University of Southampton believe that the availability of ample global IPv6 address space for infrastructure may be a benefit for many sites.

No 6bone addressing is used on site any more. Since the 6bone phaseout of June 2006 [RFC3701] (Fink, R. and R. Hinden, "6bone (IPv6 Testing Address Allocation) Phaseout," March 2004.) most transit ISPs have begun filtering attempted use of such prefixes.

Southampton does participate in global and organization scope IPv6 multicast networks. Multicast address allocations are not discussed here as they are not in scope for the document. It is noted that IPv6 has advantages for multicast group address allocation. In IPv4 a site needs to use techniques like GLOP [RFC3180] (Meyer, D. and P. Lothberg, "GLOP Addressing in 233/8," September 2001.) to pick a globally unique multicast group to use. This is problematic if the site does not use Border Gateway Protocol (BGP) [RFC4271] (Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," January 2006.) and have an Autonomous System Number (ASN). In IPv6 unicast-prefix-based IPv6 multicast addresses empower a site to pick a globally unique group address based on its unicast own site or link prefix. Embedded RP is also in use, is seen as a potential advantage for IPv6 and multicast, and has been tested successfully across providers between sites (including paths to/from the US and UK).

A.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 embracing the available address space and planning for future expansion.

In a dual-stack network, it was chosen 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. It is not expected to have IPv6-only subnets in production use for a while yet, outside the test beds and some early Mobile IPv6 trials. With congruent addressing, our firewall policies are also aligned for IPv4 and IPv6 traffic at the site border.

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. A /56 of IPv6 address space equates to 256 /64 size subnet allocations. Further /56 allocations were made for central IT infrastructure, for the network infrastructure and the server side systems.

A.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 ISATAP [RFC5214] (Templin, F., Gleeson, T., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)," March 2008.) for example.
For the Mobile IPv6 early trials there is one allocated prefix for Home Agent (HA) use. However there has been no detailed consideration yet how Mobile IPv6 usage may grow, and whether more or even every subnet will require HA support.
The university operates a tunnel broker [RFC3053] (Durand, A., Fasano, P., Guardini, I., and D. Lento, "IPv6 Tunnel Broker," January 2001.) service on behalf of UKERNA for JANET sites. This uses separate address space from JANET, not our university site allocation.

A.1.4. Node Configuration Considerations

Currently stateless autoconfiguration is used on most subnets for IPv6 hosts. There is no DHCPv6 service deployed yet, beyond tests of early code releases. It is planned to deploy DHCPv6 for address assignment when robust client and server code is available (at the time of writing the potential for this looks good, e.g. via the ISC implementation). University of Southampton is also investigating a common integrated DHCP/DNS management platform, even if the servers themselves are not co-located, including integrated DHCPv4 and DHCPv6 server configuration, as discussed in [RFC4477] (Chown, T., Venaas, S., and C. Strauf, "Dynamic Host Configuration Protocol (DHCP): IPv4 and IPv6 Dual-Stack Issues," May 2006.). Currently clients with statelessly autoconfigured addresses are added to the DNS manually, though dynamic DNS is an option. The network administrators would prefer the use of DHCP because they believe it gives them more management control. Regarding the implications of the larger IPv6 subnet address space on scanning attacks [RFC5157] (Chown, T., "IPv6 Implications for Network Scanning," March 2008.), it is noted that all the hosts are dual-stack, and thus are potentially exposed over both protocols anyway. All

addresses or published in DNS, and hence do not operate a two faced DNS.

There is internal usage of RFC4941 privacy addresses [RFC4941] (Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6," September 2007.) currently (certain platforms currently ship with it on by default), but may desire to administratively disable this (perhaps via DHCP) to ease management complexity. However, it is desired to determine the feasibility of this on all systems, e.g. for guests on wireless LAN or other user-maintained systems. Network management and monitoring should be simpler without RFC4941 in operation, in terms of identifying which physical hosts are using which addresses. Note that RFC4941 is only an issue for outbound connections, and that there is potential to assign privacy addresses via DHCPv6.

Manually configured server addresses are used to avoid address changes based upon 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).

A.2. Service Provider Considerations

In this section an IPv6 addressing schema is sketched that could serve as an example for an Internet Service Provider.

Sub-section A.2.1 starts with some thoughts regarding objective requirements of such an addressing schema and derives a few general rules of thumb that have to be kept in mind when designing an ISP IPv6 addressing plan.

Sub-section A.2.2 illustrates these findings of A.2.1 with an exemplary IPv6 addressing schema for an MPLS-based ISP offering Internet Services as well as Network Access services to several millions of customers.

A.2.1. Investigation of objective Requirements for an IPv6 addressing schema of a Service Provider

The first step of the IPv6 addressing plan design for a Service provider should identify all technical, operational, political and business requirements that have to be satisfied by the services supported by this addressing schema.

According to the different technical constraints and business models as well as the different weights of these requirements (from the point of view of the corresponding Service Provider) it is very likely that different addressing schemas will be developed and deployed by

different ISPs. Nevertheless the addressing schema of sub-section A.2.2 is one possible example.
For this document it is assumed that our exemplary ISP has to fulfill several roles for its customers as there are:

  *Local Internet Registry

  *Network Access Provider

  *Internet Service Provider

A.2.1.1. Recommendations for an IPv6 Addressing Schema from the LIR Perspective of the Service Provider

In their role as Local Internet Registry (LIR) the Service Providers have to care about the policy constraints of the RIRs and the standards of the IETF regarding IPv6 addressing. In this context, the following basic recommendations have to be considered and should be satisfied by the IPv6 address allocation plan of a Service Provider:

  *As recommended in RFC 3177 [RFC3177] (IAB and IESG, "IAB/IESG Recommendations on IPv6 Address Allocations to Sites," September 2001.) and in several RIR policies "Common" customers sites (normally private customers) should receive a /48 prefix from the aggregate of the Service Provider. (Note: The addressing plan must be flexible enough and take into account the possible change of the minimum allocation size for end users currently under definition by the RIRs.)

  *"Big customers" (like big enterprises, governmental agencies etc.) may receive shorter prefixes according to their needs when this need could be documented and justified to the RIR.

  *The IPv6 address allocation schema has to be able to meet the HD-ratio that is proposed for IPv6. This requirement corresponds to the demand for an efficient usage of the IPv6 address aggregate by the Service Provider. (Note: The currently valid IPv6 HD-ratio of 0.94 means an effective usage of about 31% of a /20 prefix of the Service Provider on the basis of /48 assignments.)

  *All assignments to customers have to be documented and stored into a database that can also be queried by the RIR.

  *The LIR has to make available means for supporting the reverse DNS mapping of the customer prefixes.

* IPv6 Address Allocation and Assignment Policies can be found at RIRs and are similar in many aspects: [reference2] (APNIC, ARIN, RIPE NCC, "www.ripe.net/ripe/docs/ipv6policy.html," July 2007.) [reference3] (APNIC, ARIN, RIPE NCC, "http://www.ripe.net/ripe/docs/ripe-412.html," July 2007.)[reference4] (ARIN, "http://www.arin.net/policy/nrpm.html#ipv6," March 2008.) [reference5] (APNIC, "http://www.apnic.net/policy/ipv6-address-policy.html," March 2007.)[reference6] (LACNIC, "http://lacnic.net/en/politicas/ipv6.html," .)

A.2.1.2. IPv6 Addressing Schema Recommendations from the ISP Perspective of the Service Provider

From ISP perspective the following basic requirements could be identified:

* The IPv6 address allocation schema must be able to realize a maximal aggregation of all IPv6 address delegations to customers into the address aggregate of the Service Provider. Only this provider aggregate will be routed and injected into the global routing table (DFZ). This strong aggregation keeps the routing tables of the DFZ small and eases filtering and access control very much.

* The IPv6 addressing schema of the SP should contain optimal flexibility since the infrastructure of the SP will change over the time with new customers, transport technologies and business cases. The requirement of optimal flexibility is contrary to the recommendation of strong IPv6 address aggregation and efficient address usage, but at this point each SP has to decide which of these requirements to prioritize.

* Keeping the multilevel network hierarchy of an ISP in mind, due to addressing efficiency reasons not all hierarchy levels can and should be mapped into the IPv6 addressing schema of an ISP. Sometimes it is much better to implement a more "flat" addressing for the ISP network than to loose big chunks of the IPv6 address aggregate in addressing each level of network hierarchy. (Note: In special cases it is even recommendable for really "small" ISPs to design and implement a totally flat IPv6 addressing schema without any level of hierarchy.)

* Besides that a decoupling of provider network addressing and customer addressing is recommended. (Note: A strong aggregation e.g. on POP, aggregation router or Label Edge Router (LER) level limits the numbers of customer routes that are visible within the

ISP network but brings also down the efficiency of the IPv6
addressing schema. That's why each ISP has to decide how many
internal aggregation levels it wants to deploy.)

A.2.1.3. IPv6 Addressing Schema Recommendations from the Network Access provider Perspective of the Service Provider

As already done for the LIR and the ISP roles of the SP it is also
necessary to identify requirements that come from its Network Access
Provider role. Some of the basic requirements are:

*The IPv6 addressing schema of the SP must be chosen in a way that
 it can handle new requirements that are triggered from customer
 side. This can be for instance the growing needs of the customers
 regarding IPv6 addresses as well as customer driven modifications
 within the access network topology (e.g. when the customer moves
 from one point of network attachment (POP) to another). (See
 section A.2.3.4 "Changing Point of Network Attachment".)

*For each IPv6 address assignment to customers a "buffer zone"
 should be reserved that allows the customer to grow in its
 addressing range without renumbering or assignment of additional
 prefixes.

*The IPv6 addressing schema of the SP must deal with multiple-
 attachments of a single customer to the SP network infrastructure
 (i.e. multi-homed network access with the same SP).

These few requirements are only part of all the requirements a Service
Provider has to investigate and keep in mind during the definition
phase of its addressing architecture. Each SP will most likely add more
constraints to this list.

A.2.1.4. A Few Rules of Thumb for Designing an IPv6 ISP Addressing Architecture

As outcome of the above enumeration of requirements regarding an ISP
IPv6 addressing plan the following design "rules of thumb" have been
derived:

*No "One size fits all". Each ISP must develop its own IPv6
 address allocation schema depending on its concrete business
 needs. It is not practicable to design one addressing plan that

fits for all kinds of ISPs (Small / big, Routed / MPLS-based,
access / transit, LIR / No-LIR, etc.).

*The levels of IPv6 address aggregation within the ISP addressing
 schema should strongly correspond to the implemented network
 structure and their number should be minimized because of
 efficiency reasons. It is assumed that the SPs own infrastructure
 will be addressed in a fairly flat way whereas the part of the
 customer addressing architecture should contain several levels of
 aggregation.

*Keep the number of IPv6 customer routes inside your network as
 small as necessary. A totally flat customer IPv6 addressing
 architecture without any intermediate aggregation level will lead
 to lots of customer routes inside the SP network. A fair trade-
 off between address aggregation levels (and hence the size of the
 internal routing table of the SP) and address conservation of the
 addressing architecture has to be found.

*The ISP IPv6 addressing schema should provide maximal
 flexibility. This has to be realized for supporting different
 sizes of customer IPv6 address aggregates ("big" customers vs.
 "small" customers) as well as to allow future growing rates (e.g.
 of customer aggregates) and possible topological or
 infrastructural changes.

*A limited number of aggregation levels and sizes of customer
 aggregates will ease the management of the addressing schema.
 This has to be weighed against the previous "thumb rule" -
 flexibility.

A.2.2. Exemplary IPv6 Address Allocation Plan for a Service
Provider

In this example, the Service Provider is assumed to operate an MPLS
based backbone and implements 6PE [RFC4798] (De Clercq, J., Ooms, D.,
Prevost, S., and F. Le Faucheur, "Connecting IPv6 Islands over IPv4
MPLS Using IPv6 Provider Edge Routers (6PE)," February 2007.) to
provide IPv6 backbone transport between the different locations (POPs)
of a fully dual-stacked network access and aggregation area.
Besides that it is assumed that the Service Provider:

    *has received a /20 from its RIR

    *operates its own LIR

*has to address its own IPv6 infrastructure

*delegates prefixes from this aggregate to its customers

This addressing schema should illustrate how the /20 IPv6 prefix of the
SP can be used to address the SP-own infrastructure and to delegate
IPv6 prefixes to its customers following the above mentioned
requirements and rules of thumb as far as possible.
The below figure summarizes the device types in a SP network and the
typical network design of a MPLS-based service provider. The network
hierarchy of the SP has to be taken into account for the design of an
IPv6 addressing schema and defines its basic shape and the various
levels of aggregation.

```
    +-----------------------------------------------------------------+
    |            LSRs of the MPLS Backbone of the SP                  |
    +-----------------------------------------------------------------+
      |         |            |             |                |
      |         |            |             |                |
    +-----+ +-----+    +--------+    +--------+      +--------+
    | LER | | LER |    | LER-BB |    | LER-BB |      | LER-BB |
    +-----+ +-----+    +--------+    +--------+      +--------+
      | |   |   |        |    |     /      |           |     |
      | |   |   |        |    |    /       |           |     |
      | |   |   |   +------+  +------+  +------+        |     |
      | |   |   |   |BB-RAR|  |BB-RAR|  |  AG  |        |     |
      | |   |   |   +-----+   +------+  +------+        |     |
      | |   |   |    | |       | |       |  |           |     |
      | |   |   |    | |       | |       |  |           |     |
      | |   |   |    | |       | | +-----+ +-----+ +-----+ +-----+
      | |   |   |    | |       | | | RAR | | RAR | | RAR | | RAR |
      | |   |   |    | |       | | +-----+ +-----+ +-----+ +-----+
      | |   |   |    | |       | | | |     | |     | |     | |
      | |   |   |    | |       | | | |     | |     | |     | |
    +-----------------------------------------------------------------+
    |                    Customer networks                           |
    +-----------------------------------------------------------------+
```
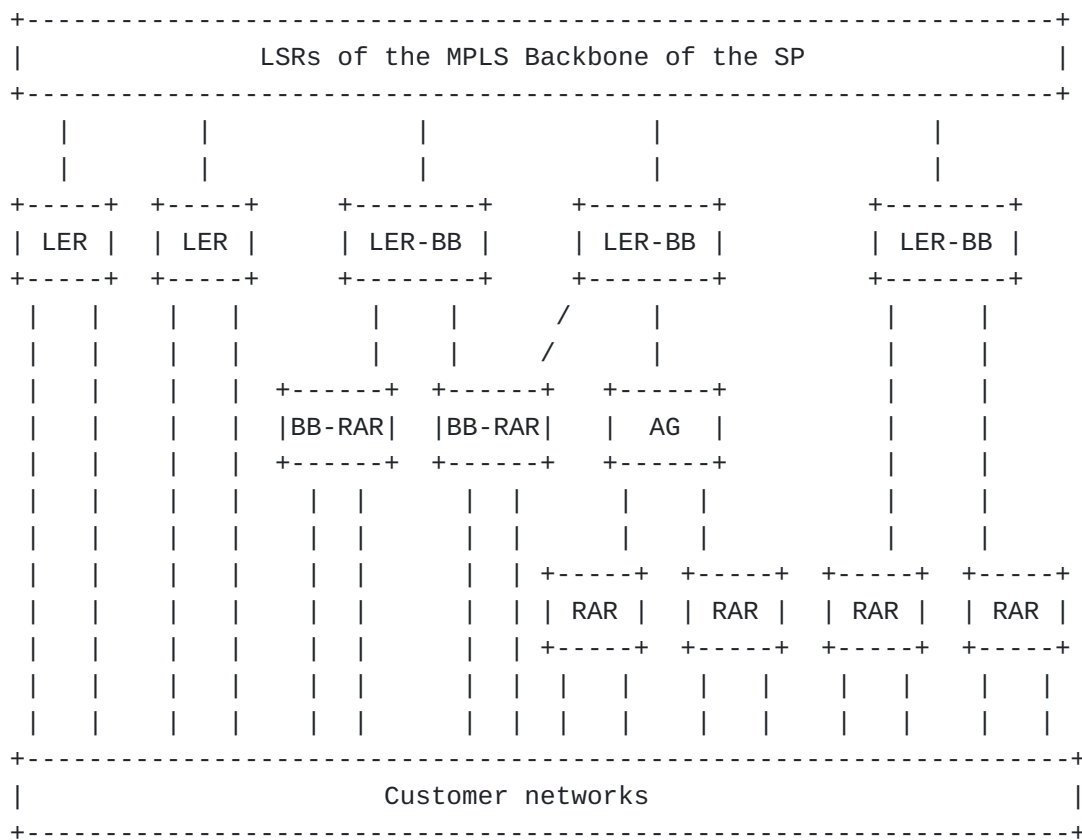Figure: Exemplary Service Provider Network

```
LSR     ... Label Switch Router
LER     ... Label Edge Router
LER-BB  ... Broadband Label Edge Router
RAR     ... Remote Access Router
BB-RAR  ... Broadband Remote Access Router
AG      ... Aggregation Router
```

Basic design decisions for the exemplary Service Provider IPv6 address
plan regarding customer prefixes take into consideration:

*The prefixes assigned to all customers behind the same LER (e.g.
 LER or LER-BB) are aggregated under one LER prefix. This ensures
 that the number of labels that have to be used for 6PE is limited
 and hence provides a strong MPLS label conservation.

*The /20 prefix of the SP is separated into 3 different pools that
 are used to allocate IPv6 prefixes to the customers of the SP:

-A pool (e.g. /24) for satisfying the addressing needs of really
 "big" customers (as defined in A.2.2.1 sub-section A.) that need
 IPv6 prefixes larger than /48 (e.g. /32). These customers are
 assumed to be connected to several POPs of the access network,
 so that this customer prefix will be visible in each of these
 POPs.

-A pool (e.g. /24) for the LERs with direct customer connections
 (e.g. dedicated line access) and without an additional
 aggregation area between the customer and the LER. (These LERs
 are mostly connected to a limited number of customers because of
 the limited number of interfaces/ports.)

-A larger pool (e.g. 14*/24) for LERs (e.g. LER-BB) that serve a
 high number of customers that are normally connected via some
 kind of aggregation network (e.g. DSL customers behind a BB-RAR
 or Dial-In customers behind a RAR).

-The IPv6 address delegation within each Pool (end customer
 delegation or also the aggregates that are dedicated to the LERs
 itself) should be chosen with an additional buffer zone of 100%
 - 300% for future growth. I.e. 1 or 2 additional prefix bits
 should be reserved according to the expected future growth rate
 of the corresponding customer / the corresponding network device
 aggregate.

**A.2.2.1. Defining an IPv6 Address Allocation Plan for
Customers of the Service Provider**

A.2.2.1.1. 'Big' Customers

SP's "big" customers receive their prefix from the /24 IPv6 address
aggregate that has been reserved for their "big" customers. A customer
is considered as "big" customer if it has a very complex network
infrastructure and/or huge IPv6 address needs (e.g. because of very
large customer numbers) and/or several uplinks to different POPs of the
SP network.
The assigned IPv6 address prefixes can have a prefix length in the
range 32-48 and for each assignment a 100 or 300% future growing zone
is marked as "reserved" for this customer. This means for instance that
with a delegation of a /34 to a customer the corresponding /32 prefix
(which contains this /34) is reserved for the customers future usage.
The prefixes for the "big" customers can be chosen from the
corresponding "big customer" pool by either using an equidistant
algorithm or using mechanisms similar to the Sparse Allocation
Algorithm (SAA) [reference2] (APNIC, ARIN, RIPE NCC, "www.ripe.net/
ripe/docs/ipv6policy.html," July 2007.).

A.2.2.1.2. 'Common' Customers

All customers that are not "big" customers are considered as "common"
customers. They represent the majority of customers hence they receive
a /48 out of the IPv6 customer address pool of the LER where they are
directly connected or aggregated.
Again a 100 - 300% future growing IPv6 address range is reserved for
each customer, so that a "common" customer receives a /48 allocation
but has a /47 or /46 reserved.
(Note: If it is obvious that the likelyhood of needing a /47 or /46 in
the future is very small for a "common" customer, than no growing
buffer should be reserved for it and only a /48 will be assigned
without any growing buffer.)
In the network access scenarios where the customer is directly
connected to the LER the customer prefix is directly taken out of the
customer IPv6 address aggregate (e.g. /38) of the corresponding LER.
In all other cases (e.g. the customer is attached to a RAR that is
themselves aggregated to an AG or to a LER-BB) at least 2 different
approaches are possible.
1) Mapping of Aggregation Network Hierarchy into Customer IPv6
Addressing Schema. The aggregation network hierarchy could be mapped
into the design of the customer prefix pools of each network level in
order to achieve a maximal aggregation at the LER level as well as at
the intermediate levels. (Example: Customer - /48, RAR - /38, AG - /32,
LER-BB - /30). At each network level an adequate growing zone should be
reserved. (Note: This approach requires of course some "fine tuning" of

the addressing schema based on a very good knowledge of the Service
Provider network topology including actual growing ranges and rates.)
When the IPv6 customer address pool of a LER (or another device of the
aggregation network - AG or RAR) is exhausted, the related LER (or AG
or RAR) prefix is shortened by 1 or 2 bits (e.g. from /38 to /37 or /
36) so that the originally reserved growing zone can be used for
further IPv6 address allocations to customers. In the case where this
growing zone is exhausted as well a new prefix range from the
corresponding pool of the next higher hierarchy level can be requested.
2) "Flat" Customer IPv6 Addressing Schema. The other option is to
allocate all the customer prefixes directly out of the customer IPv6
address pool of the LER where the customers are attached and aggregated
and to ignore the intermediate aggregation network infrastructure. This
approach leads of course to a higher amount of customer routes at LER
and aggregation network level but takes a great amount of complexity
out of the addressing schema. Nevertheless the aggregation of the
customer prefixes to one prefix at LER level is realized as required
above.
(Note: The handling of (e.g. technically triggered) changes within the
ISP access network is shortly discussed in section A.2.3.5.)
If the actual observed growing rates show that the reserved growing
zones are not needed than these growing areas can be freed and used for
assignments for prefix pools to other devices at the same level of the
network hierarchy.

A.2.2.2. Defining an IPv6 Address Allocation Plan for the
Service Provider Network Infrastructure

For the IPv6 addressing of SPs own network infrastructure a /32 (or /
40) from the "big" customers address pool can be chosen.
This SP infrastructure prefix is used to code the network
infrastructure of the SP by assigning a /48 to every POP/location and
using for instance a /56 for coding the corresponding router within
this POP. Each SP internal link behind a router interface could be
coded using a /64 prefix. (Note: While it is suggested to choose a /48
for addressing the POP/location of the SP network it is left to each SP
to decide what prefix length to assign to the routers and links within
this POP.)
The IIDs of the router interfaces may be generated by using EUI-64 or
through plain manual configuration e.g. for coding additional network
or operational information into the IID.
It is assumed that again 100 - 300% growing zones for each level of
network hierarchy and additional prefix bits may be assigned to POPs
and/or routers if needed.
Loopback interfaces of routers may be chosen from the first /64 of the
/56 router prefix (in the example above).

(Note: The /32 (or /40) prefix that has been chosen for addressing SPs
own IPv6 network infrastructure gives enough place to code additional
functionalities like security levels or private and test infrastructure
although such approaches haven't been considered in more detail for the
above described SP until now.)
Point-to-point links to customers (e.g. PPP links, dedicated line etc.)
may be addressed using /126 prefixes out of the first /64 of the access
routers that could be reserved for this reason.

A.2.3. Additional Remarks

A.2.3.1. ULA

From the actual view point of SP there is no compelling reason why ULAs
should be used from a SP. Look at section 2.2.
ULAs could be used inside the SP network in order to have an additional
"site-local scoped" IPv6 address for SPs own infrastructure for
instance for network management reasons and maybe also in order to have
an addressing schema that couldn't be reached from outside the SP
network.
In the case when ULAs are used it is possible to map the proposed
internal IPv6 addressing of SPs own network infrastructure as described
in A.2.2.2 above directly to the ULA addressing schema by substituting
the /48 POP prefix with a /48 ULA site prefix.

A.2.3.2. Multicast

IPv6 Multicast-related addressing issues are out of the scope of this
document.

A.2.3.3. POP Multi-homing

POP (or better LER) Multi-homing of customers with the same SP can be
realized within the proposed IPv6 addressing schema of the SP by
assigning multiple LER-dependent prefixes to this customer (i.e.
considering each customer location as a single-standing customer) or by
choosing a customer prefix out of the pool of "big" customers. The
second solution has the disadvantage that in every LER where the

customer is attached this prefix will appear inside the IGP routing
table requiring an explicit MPLS label.
(Note: The described negative POP/LER Multi-homing effects to the
addressing architecture in the SP access network are not tackled by
implementing the Shim6 Site Multi-homing approach since this approach
targets only on a mechanism for dealing with multiple prefixes in end
systems – the SP will nevertheless have unaggregated customer prefixes
in its internal routing tables.)

A.2.3.4. Changing Point of Network Attachement

In the possible case that a customer has to change its point of network
attachment to another POP/LER within the ISP access network two
different approaches can be applied assuming that the customer uses PA
addresses out of the SP aggregate:
1.) The customer has to renumber its network with an adequate customer
prefix out of the aggregate of the corresponding LER/RAR of its new
network attachement. To minimise the administrative burden for the
customer the prefix should be of the same size as the former. This
conserves the IPv6 address aggregation within the SP network (and the
MPLS label space) but adds additional burden to the customer. Hence
this approach will most likely only be chosen in the case of "small
customers" with temporary addressing needs and/or prefix delegation
with address auto-configuration.
2.) The customer does not need to renumber its network and keeps its
address aggregate.
This apporach leads to additional more-specific routing entries within
the IGP routing table of the LER and will hence consume additional MPLS
labels - but it is totally transparent to the customer. Because this
results in additional administrative effort and will stress the router
resources (label space, memory) of the ISP this solution will only be
offered to the most valuable customers of an ISP (like e.g. "big
customers" or "enterprise customers").
Nevertheless the ISP has again to find a fair trade-off between
customer renumbering and sub-optimal address aggregation (i.e. the
generation of additional more-specific routing entries within the IGP
and the waste of MPLS Label space).

A.2.3.5. Restructuring of SP (access) Network and Renumbering

A technically triggered restructuring of the SP (access) network (for
instance because of split of equipment or installation of new
equipment) should not lead to a customer network renumbering. This

challenge should be handled in advance by an intelligent network design
and IPv6 address planing.
In the worst case the customer network renumbering could be avoided
through the implementation of more specific customer routes. (Note:
Since this kind of network restructuring will mostly happen within the
access network (at the level) below the LER, the LER aggregation level
will not be harmed and the more-specific routes will not consume
additional MPLS label space.)

A.2.3.6. Extensions Needed for the Later IPv6 Migration Phases

The proposed IPv6 addressing schema for a SP needs some slight
enhancements / modifications for the later phases of IPv6 integration,
for instance in the case when the whole MPLS backbone infrastructure
(LDP, IGP etc.) is realized over IPv6 transport and an IPv6 addressing
of the LSRs is needed. Other changes may be necessary as well but
should not be explained at this point.

Appendix B. Considerations for Subnet Prefixes Different then /64

B.1. Considerations for Subnet Prefixes Shorter then /64

An allocation of a prefix shorter then 64 bits to a node or interface
is considered bad practice. One exception to this statement is when
using 6to4 technology where a /16 prefix is utilized for the pseudo-
interface [RFC3056] (Carpenter, B. and K. Moore, "Connection of IPv6
Domains via IPv4 Clouds," February 2001.). 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 for most networks considering that 2^64 provides plenty
of node addresses.
The subnet prefix assignments can be made either by manual
configuration, by a stateful Host Configuration Protocol [RFC3315]
(Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney,
"Dynamic Host Configuration Protocol for IPv6 (DHCPv6)," July 2003.),
by a stateful prefix delegation mechanism [RFC3633] (Troan, O. and R.

[Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6," December 2003.)](#) or implied by stateless autoconfiguration from prefix RAs.

B.2. Considerations for Subnet Prefixes Longer then /64

The following subsections describe subnet prefix values that should be avoided in deployments, because nodes who assume that the subnet prefix is /64 could treat them incorrectly.

B.2.1. /126 Addresses

126 bit subnet prefixes are typically used for point-to-point links similar to a the IPv4 address conservative /30 allocation for point-to-point links. 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.

B.2.2. /127 Addresses

The usage of the /127 addresses, the equivalent of IPv4's RFC3021 [RFC3021] (Retana, A., White, R., Fuller, V., and D. McPherson, "Using 31-Bit Prefixes on IPv4 Point-to-Point Links," December 2000.) is not valid and should be strongly discouraged as documented in RFC3627 [RFC3627] (Savola, P., "Use of /127 Prefix Length Between Routers Considered Harmful," September 2003.).

B.2.3. /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 off-link 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

B.2.4. EUI-64 'u' and 'g' bits

When using subnet prefix lengths other than /64, the interface identifier cannot be in Modified EUI-64 format as required by [RFC4291] (Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture," February 2006.). However, nodes not aware that a prefix length other than /64 is used might still think it's an EUI-64; therefore, it's prudent to take the next considerations to set the bits into account. Address space conservation is the main motivation for using a subnet prefix length longer than 64 bits, however this kind of address conservation is of little benefit compared with the additional considerations one must make when creating and maintain an IPv6 address plan.

The address assignment can be made either by manual configuration or by a stateful Host Configuration Protocol [RFC3315] (Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)," July 2003.).

When assigning a subnet prefix of more then 70 bits, according to RFC4291 [RFC4291] (Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture," February 2006.) 'u' and 'g' bits (respectively the 71st and 72nd bit) need to be taken into consideration and should be set correct.

The 'u' (universal/local) bit is the 71st bit of IPv6 address and is used to determine whether the address is universally or locally administered. If 0, the IEEE, through the designation of a unique company ID, has administered the address. If 1, the address is locally administered. The network administrator has overridden the manufactured address and specified a different address.

The 'g' (the individual/group) bit is the 72st bit and is used to determine whether the address is an individual address (unicast) or a group address (multicast). If '0', the address is a unicast address. If '1', the address is a multicast address.

In current IPv6 protocol stacks, the relevance of the 'u' and 'g' bit is marginal and typically will not show an issue when configured wrongly, however future implementations may turn out differently if they would be processing the 'u' and 'g' bit in IEEE like behavior.

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 considered when selecting a subnet length longer then 64 bits.

B.2.5. Anycast Addresses

B.2.5.1. Subnet Router Anycast Address

RFC4291 [RFC4291] (Hinden, R. and S. Deering, "IP Version 6 Addressing
Architecture," February 2006.) provides a definition for the required
Subnet Router Anycast Address as follows:

```
|                 n bits                | 128-n bits  |
+---------------------------------------+-------------+
|             subnet prefix             | 00000000000000 |
+---------------------------------------+-------------+
```

It is recommended to avoid allocating this IPv6 address to a device
which expects to have a normal unicast address. There is no additional
dependency for the subnet prefix with the exception of the 64-bit
extended unique identifier (EUI-64) and an Interface Identifier (IID)
dependency. These will be discussed later in this document.
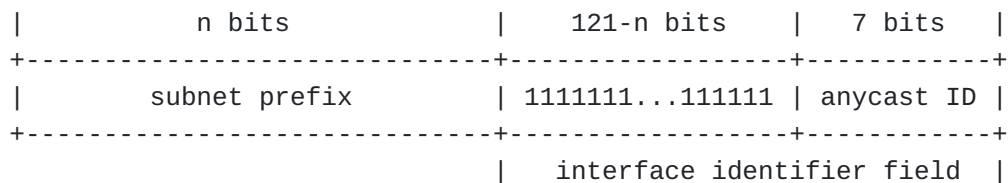

B.2.5.2. Reserved IPv6 Subnet Anycast Addresses

RFC2526 [RFC2526] (Johnson, D. and S. Deering, "Reserved IPv6 Subnet
Anycast Addresses," March 1999.) 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 Anycast addresses have been defined as follows
for EUI-64 format:

```
|           64 bits           |     57 bits     |  7 bits  |
+-----------------------------+-----------------+----------+
|        subnet prefix        | 1111110111...111 | anycast ID |
+-----------------------------+-----------------+----------+
```

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 64-bit extended unique identifier (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:

```
|            n bits            |   121-n bits     |  7 bits   |
+------------------------------+------------------+-----------+
|         subnet prefix        | 1111111...111111 | anycast ID |
+------------------------------+------------------+-----------+
                               |   interface identifier field  |
```

It is recommended to avoid allocating this IPv6 address to a device which expects to have a normal unicast address. There is no additional dependency for the subnet prefix with the exception of the EUI-64 and an Interface Identifier (IID) dependency. These will be discussed later in this document.

B.2.6. Addresses Used by Embedded-RP (RFC3956)

Embedded-RP [RFC3956] (Savola, P. and B. Haberman, "Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address," November 2004.) 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 resulted in a restriction of 15 possible RP-addresses per prefix that can be used with embedded-RP. The space assigned for the embedded-RP is based on the 4 low order bits, while the remainder of the Interface ID (RIID) is set to all '0'.

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

                 Where: (RIID) = 4 bit.
```

This format implies that when selecting subnet prefixes longer then 64, and the bits beyond the 64th one are non-zero, the subnet can not use embedded-RP.

In addition it is discouraged to assign a matching embedded-RP IPv6 address to a device that is not a real Multicast Rendezvous Point, even though it would not generate major problems.

B.2.7. ISATAP Addresses

ISATAP [RFC5214] (Templin, F., Gleeson, T., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)," March 2008.) is an experimental 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 address has the IPv4 address embedded, based on a predefined structure policy that identifies them as an ISATAP address.

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

When using subnet prefix length longer then 64 bits it is good engineering practice that the portion of the IPv6 prefix from bit 65 to the end of the host-id does not match with the well-known ISATAP [0000:5EFE] address when assigning an IPv6 address to a non-ISATAP interface.
Note that the definition of ISATAP does not support multicast.

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	
	Olaf Bonness	
	T-Systems Enterprise Services GmbH	
	Goslarer Ufer 35	
	Berlin, 10589	
	Germany	
Phone:	+49 30 3497 3124	
Email:	Olaf.Bonness@t-systems.com	
	Christian Hahn	
	T-Systems Enterprise Services GmbH	
	Goslarer Ufer 35	
	Berlin, 10589	
	Germany	
Phone:	+49 30 3497 3164	
Email:	HahnC@t-systems.com	