

<[draft-narten-iana-rir-ipv6-considerations-00.txt](#)>

June 16, 2005

Issues Related to the Management of IPv6 Address Space<[draft-narten-iana-rir-ipv6-considerations-00.txt](#)>

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress".

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/1id-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

This Internet-Draft expires on December 16, 2005.

Abstract

This document reviews the history of IPv6 address delegation and examines some of the technical implications of various allocation policies in terms of their impact on long-term address space utilization and aggregation. The purpose of this document is to stimulate discussion and foster improved understanding of the implications of various policies, so that appropriate policies are adopted that preserve the long-term scalability of IPv6.

Contents

Status of this Memo.....	1
1. Introduction.....	3
2. Definitions.....	3
2.1. Internet Registry (IR).....	3
2.2. Regional Internet Registry (RIR).....	4
2.3. National Internet Registry (NIR).....	4
2.4. Local Internet Registry (LIR).....	4
2.5. Allocate.....	4
2.6. Assign.....	4
2.7. Utilization.....	4
2.8. HD-Ratio.....	5
3. Background.....	5
3.1. IPv6 Address Structure.....	6
3.2. IPv6 Address Assignment History.....	7
3.3. Goals of IPv6 Address Space Management.....	7
3.4. Utilization.....	8
3.5. Aggregation.....	9
4. How Much Address Space Do We Really Have?.....	10
4.1. An example: Cable Modem/DSL Service in US.....	10
4.2. How Much Is Enough?.....	11
5. On RIR->LIR->End User Assignment and Allocation Policies.	12
5.1. On the HD Ratio Metric.....	12
5.2. On /48 Assignments to End Sites.....	12
5.3. On the /64 Boundary.....	15
5.4. Summing It All Up.....	15
6. On Reservations and Aggregation.....	17
6.1. On RIR Reservations.....	17
6.2. On IANA->RIR Assignments.....	18
7. Security Considerations.....	18
8. IANA Considerations.....	18
9. Acknowledgments.....	18
10. References.....	18
11. Appendix A : HD-Ratio.....	20
12. Author Çös Address.....	21

1. Introduction

The assignment and allocation of global IPv6 unicast address space to ISPs and end-users is implemented under policies maintained and implemented by the Regional Internet Registries (RIRs) and IANA. End users typically get IPv6 address space from their ISPs, who in turn obtain allocations from RIRs. RIRs in turn obtain address space from IANA.

This document reviews the history of IPv6 address delegation and examines some of the technical implications of various allocation policies in terms of their impact on long-term address space utilization and aggregation. The purpose of this document is stimulate discussion and foster improved understanding of the implications of various policies, so that appropriate policies are adopted that preserve the long-term scalability of IPv6.

Questions that this document attempts to shed light on include:

- How large should the chunks be that IANA allocates to RIRs?
- How much (if any) space should IANA hold in "reserve", to allow a RIR (or LIR) to "grow into" if/when it needs a further allocation?
- What algorithms should RIRs use to manage their space, and at what point are they justified in asking IANA for more address space?
- What are the long-term implications of these algorithms on overall IPv6 address space consumption, utilization and aggregability?

2. Definitions

[Note: the following definitions are taken (and in some cases slightly modified) from the "IPv6 Address Allocation and Assignment Policy" document that was produced and adopted by the APNIC, ARIN and RIPE communities in 2002.]

2.1. Internet Registry (IR)

An Internet Registry (IR) is an organization that is responsible for distributing IP address space to its members or customers and for registering those distributions. IRs are classified according to their primary function and territorial scope.

2.2. Regional Internet Registry (RIR)

Regional Internet Registries (RIRs) are established and authorized by respective regional communities, and recognized by the IANA to serve and represent large geographical regions. The primary role of RIRs is to manage and distribute public Internet address space within their respective regions. Currently, there are five RIRs: AFRINIC, APNIC, ARIN, LACNIC, and RIPE.

2.3. National Internet Registry (NIR)

A National Internet Registry (NIR) primarily allocates address space to its members or constituents, which are generally LIRs organized at a national level. NIRs exist mostly in the Asia Pacific region.

2.4. Local Internet Registry (LIR)

A Local Internet Registry (LIR) is an IR that primarily assigns address space to the users of the network services that it provides. LIRs are generally ISPs, whose customers are primarily end users or other ISPs.

2.5. Allocate

To allocate means to distribute address space to IRs for the purpose of subsequent distribution by them.

2.6. Assign

To assign means to delegate address space to an ISP or end-user, for specific use within the Internet infrastructure they operate. Assignments are made for specific documented purposes by specific organizations and are not sub-assigned to other parties.

2.7. Utilization

In IPv4, the utilization of a chunk of address space is defined as the ratio of the actual number of host assignments to the theoretical maximum number of host assignments. For example, a /24 in IPv4 can number 256 hosts, and if 200 have been assigned to hosts, the utilization would be 200/256 or 78%. In IPv4, a utilization ratio of 80% is the threshold value used to justify the granting of additional address space. When an entity has used up 80% of its space, it is

eligible to receive more.

Unlike IPv4, IPv6 is generally assigned to end sites in fixed amounts (e.g., in the /48 to /64 range per current address distribution policies). The actual "utilization" of an IPv6 address block in absolute terms (i.e., counting numbers of assigned hosts) will be exceedingly low. Thus, in IPv6, "utilization" is no longer the metric used for determining when additional address space can be obtained. Instead, an HD Metric is used, and it counts "links" or "subnets" rather than individual host assignments.

In the context of IPv6, the term utilization refers to the allocation of address blocks to end sites, and not the number of addresses assigned within individual /48s at end sites.

2.8. HD-Ratio

The HD-Ratio is a way of measuring the efficiency of address assignment [[RFC 3194](#)]. It is an adaptation of the H-Ratio originally defined in [[RFC1715](#)] and is expressed as follows:

$$\text{HD} = \frac{\text{Log (number of allocated objects)}}{\text{Log (maximum number of allocatable objects)}}$$

where the objects are IPv6 end site assignments (e.g., /48s) assigned from an IPv6 prefix of a given size.


The term "HD ratio" is somewhat misleading in the context of IPv6, since it is used to measure "site density" rather than "host density".

3. Background

The IETF has indicated that IANA may allocate IPv6 unicast address space out of the block 2000::/3, with the remaining 7/8ths of the address space to be held in reserve [[RFC3513](#)]. IANA, in turn, has been allocating small chunks of that address space to individual RIRs. The individual RIRs allocate addresses to ISPs (or more precisely, to LIRs), who then assign them to end sites. End sites are typically given assignments in the /48-/64 range, consistent with the recommendations given in [[RFC 3177](#)] and adopted RIR policy [RIR-IPV6].

3.1. IPv6 Address Structure

The general format of an IPv6 global unicast address is defined in [RFC3513](#) [[RFC3513](#)]:




IPv6 Address Structure

Figure 1

The 64-bit IPv6 interface ID is an architectural boundary defined by Stateless Address Autoconfiguration [[RFC2462](#)]. Thus, in terms of RIR/IANA address allocation policy, the left-most 64 bits of an IPv6 address are managed through the IANA/RIR processes. Under the current IPv6 allocation policies, the value for m in the figure above is commonly assumed to be 16, such that the global routing prefix is 48 bits in length and the per-customer subnet ID is 16 bits in length.

It is worth noting that the 16-bit subnet identifier is a policy choice, not an architectural one. Specifically, [RFC 3513bis] [Section 2.5](#) says:

A slightly sophisticated host (but still rather simple) may additionally be aware of subnet prefix(es) for the link(s) it is attached to, where different addresses may have different values for n :



Though a very simple router may have no knowledge of the internal structure of IPv6 unicast addresses, routers will more generally have knowledge of one or more of the hierarchical boundaries for the operation of routing protocols. The known boundaries will differ from router to router, depending on what positions the router holds in the routing hierarchy.

Except for the knowledge of the subnet boundary discussed in the previous paragraphs nodes should not make any assumptions about the structure of an IPv6 address.

The use of a 16-bit subnet ID is a policy rather than architectural boundary, and the arguments for choosing a subnet ID of 16 are documented in the IAB/IESG recommendations given in [[RFC3177](#)]. Although the RIR community adopted these recommendations, there have been on-going discussions within the RIR community as to whether giving a full 16 bits of subnet space to (essentially) all end sites is justified, necessary or prudent. The size of the subnet ID that is used by end sites has a significant impact on the overall utilization of IPv6 address space that can be achieved. For example, using a subnet ID of 8 bits (i.e., 256 values) would increase the effective size of the global routing prefix from 48 to 56 bits, an increase in size of two orders of magnitude.

[3.2.](#) IPv6 Address Assignment History

The earliest allocations from IANA to RIRs were made in July, 1999 and consisted of /23 allocations [[IAB-IANA98](#)]. Out of these allocations, the RIRs began allocating /35s to LIRs. In 2001, [RFC 3177] was developed, which called for assigning /48s to end sites in many cases. These recommendations were adopted by the RIRs in 2002 [[RIR-IPv6](#)]. Simultaneously, it was recognized that if LIRs were assigning /48s, but were assigning them out of a /35, the size of the RIR->LIR allocation also needed to be increased. Both changes are described in [[RIR-IPv6](#)], and that policy is still largely in place today.

From 1999 until December 2004, the size of the IANA->RIR allocation was a /23, an allocation size that had not been changed even though allocation policies had been significantly revised and updated since 1999. Moreover, in 2004, the RIRs began receiving requests from LIRs for (relatively) large amounts of address space that were justified under the in-place allocation policies, including requests for more than the default minimum size of /32. The combination of in-place allocation policies together with ISPs requesting IPv6 address allocations for numbering millions of end sites led to a need to revisit the default IANA->RIR allocation size, a topic under current discussion within the RIR community [[RIR-IANA-IPv6](#)].

A detailed listing of all the IPv6 assignments that have been made can be found at [[IPv6-STATS](#)].

[3.3.](#) Goals of IPv6 Address Space Management

Two important goals in the management of address space are to ensure reasonable levels of utilization (so that we don't exhaust the address space itself) as well as achieving good levels of

aggregation. In the current Internet routing architecture (i.e., CIDR) aggregation is necessary in order to keep the size of the routing table manageable. By "size", we refer not just to the size of forwarding table (in terms of numbers of individual entries), but also the cost and overhead of sending routing updates (in terms of bandwidth needed), the processing power needed for the routing algorithms to converge in a timely fashion, and on the overall understandability and manageability of the system from an operational perspective.

3.4. Utilization

In IPv4, there has been much emphasis on maintaining high utilization of the address space. Indeed, IPv4 address allocation policies use an "80% utilization" rule. Sites and LIRs are eligible for additional address space when they can demonstrate usage of 80% of the space they have been assigned. Likewise, LIRs obtain more space from RIRs when their utilization reaches 80%.

With IPv6, justification for obtaining address space is different in two important ways:

- 1) Rather than use a "absolute percentage utilization" metric to indicate thresholds, an HD-ratio metric is used. Under the policy rules adopted in 2002, a site or LIR is eligible for more address space once it shows that the amount of address space it has used exceeds an HD ratio of .8.
- 2) Rather than count "numbers of hosts assigned", the IPv6 HD ratio counts "assignments to end sites", typically /48 assignments; the actual number of devices assigned to those sites is not a factor.

Comparing just the HD-ratio metric to the 80% rule, an HD-ratio metric of .8 is significantly easier to satisfy than 80% rule, especially for LIRs (e.g., see [[huston-hd-metric](#)]). As the table in [Appendix A](#) shows, for a /32 assignment, an LIR need only assign 7132 /48s to reach the threshold, which corresponds to an absolute utilization of only 10.9%. Moreover, [Appendix A](#) also shows that the required level of utilization (in absolute percentage terms) decreases as the amount of address space the ratio is applied to increases. For example, one need only achieve a utilization of roughly 2-3% to justify allocations in the /20-/23 range, and this does not even consider that the IPv6 formula counts sites rather than host assignments.

The amount of address space available in IPv6 is immensely larger than that available in IPv4. The IETF has intentionally chosen to use

address space relatively inefficiently (compared with IPv4) in order to simplify other aspects of IPv6. For example, lower utilization levels have the benefit of simplifying addressing at end sites (e.g., through use of stateless address autoconfiguration [[RFC 2462](#)]); end sites get enough address space that they do not constantly have to renumber in order to maintain a high utilization, and get significantly more address space for a given sized-site than under the corresponding policies for IPv4 space.

Having said that, prudence calls for not being needlessly wasteful in terms of address space usage. It is hard to predict what the Internet will look like in 20-50 years, and poor decisions made now will almost certainly be very difficult to undo later.

3.5. Aggregation

One important goal for IPv6 is to achieve good aggregation of addresses over significantly longer-term time frames than in IPv4. Specifically, thought should be given to ensuring aggregation over time frames on the order of 10-20 years (and longer), rather than just 2 or 3.

Routing in IPv6 is fundamentally equivalent to routing in IPv4, namely CIDR. Thus, the same pressures that have dominated IPv4 routing can be expected in IPv6. Namely, the desire for "provider independent" (PI) space, the desire to multi-home, etc. Although the IETF's multi6 and shim6 [[shim6](#)] efforts are developing a routing-infrastructure friendly (i.e., scalable) multi-homing solution, that work is still at a relatively early stage and is unlikely to result in widely-available, mature products for at least 5 years, even in optimistic scenarios. Moreover, it is unclear at the present time how broadly applicable the shim6 solution will ultimately be. Thus, it would be foolish at the present time to assume that routing in IPv6 will avoid any of the route scaling issues that continue to be a significant issue in IPv4.

One consequence to the 80% rule as it has been applied to IPv4 is that it has contributed to significant fragmentation of the IPv4 address space, which leads to an inability to aggregate routes as much as might be desired. To be fair, that is not a result of the 80% rule per se, rather, it results from a number of factors:

- RIRs do not hold in reserve large amounts of address space adjacent to an existing allocation to allow subsequent requests for additional space to be satisfied from an adjacent block. That is, policies governing when RIRs can request additional address space from IANA do not favor achieving good aggregation by

allowing large amounts of space to be held in reserve. Indeed, achieving good levels of aggregation have simply not been considered important in practice and are not part of the adopted policies. Consequently, there are no incentives in the system for encouraging aggregation over long periods.

- For various reasons (e.g., to get more optimal routing, traffic engineering, etc.), some LIRs deaggregate their allocations and advertise multiple longer prefixes. It should be noted that this type of fragmentation is reversible, since an LIR could fall back to advertising a single aggregate if necessary. That is, this type of fragmentation is an optimization, and isn't required in order to maintain basic connectivity for all end sites.
- Mergers and acquisitions results in administratively related sites being comprised of address space that is not aggregatable.

To support aggregation over time frames of 10-20 years (and beyond), and to avoid reliving the fragmentation experience of IPv4, RIR address policies must support the holding of adequately-sized "reservations" of address space adjacent to address space that has been delegated. Thus, the amount of space held in reserve by the RIRs, the manner in which it is held, and the time frame over which RIRs maintain that reserve are key policy questions with significant implications on long-term aggregation and fragmentation of the IPv6 address space.

4. How Much Address Space Do We Really Have?

4.1. An example: Cable Modem/DSL Service in US

In the hallway at a recent ARIN meeting, I was cornered by someone who had done a back-of-the envelope calculation that led him to believe the current policies needed adjustment. The argument went like:

If I assign 4M /48 CÖs of IPv6 (one to each cable modem on my network), according to the HD-ratio I am justified to obtain something around a /20 of IPv6 addresses. In other words, I am justified in getting 268M /48 CÖs even though I am only using 4M of them. That would be enough for me to assign at least two for every household in the US (not just the 19M on my network).

Now if all the cable providers (e.g., Comcast, Cox, Adelphia, Cablevision, Time-Warner, etc.) did the same for their networks; and each of the DSL companies made a similar move (SBC, Verizon,

Quest, etc.); perhaps we could easily see the broadband market in the US alone obtaining some 16 /20 ÇÖs of IPv6 or a total of /16. There are only 8192 of those available in the current global unicast space of 2001::/3.

Anyhow, you can see where this might lead...

4.2. How Much Is Enough?

At various times in the past, it has been asserted that we will "never run out of IPv6 addresses". For example, Wikipedia repeats the myth [[WIKIPEDIA](#)]:

If the earth were made entirely out of 1 cubic millimetre grains of sand, then you could give a unique address to each grain in 300 million planets the size of the earth.

These sorts of statements are made on the assumption that all of IPv6 ÇÖs 128-bit address space can be use efficiently for addressing devices. But as discussed earlier, only 64-bits are available for addressing links, since 64 bits are used for an interface identifier. Or, stated differently, the above figure is only true if one believes that one can have single links, on which 2^{64} hosts can be attached; in practice, the actual number of hosts connected to a single link is on the order of tens to thousands, with most links at the lower end of the range.

A more realistic measure is to consider how many links, or end sites can be addressed. One way to frame the discussion is to ask the question: how many devices will a person or end site need to be able to assign addresses to?

World population projections for the year 2050 estimate a world population of 9.2B persons [[CENSUS](#)]. Under the current address allocation policies (specifically, /48 to end sites and use of an HD-ratio threshold of .8), the following observation can be made.

If we allow for one /48 per person, and assume there is only a single ISP for everyone, 9.2B /48s corresponds to a /7, or 1/128th of the entire available address space (and fully 1/16th of 2001::/3). Clearly, this does not indicate a need to panic or that IPv6 will run out of addresses. However, such numbers are also nowhere near "practically infinite". A key policy question is whether the community is comfortable with a projection that can result in fully 1% of the available address space used up in 50 years, given:

- the inherent uncertainties of predicting future address consumption rates (e.g., an appropriate number of devices per person, and indeed whether a per-person perspective is even the appropriate one).
- the inherent uncertainty in the length of time that IPv6 will be a key part of the underlying infrastructure (i.e., 50 vs. 100 vs. 200 years).
- a long history of infamous "predictions", e.g., "640K ought to be enough for anybody." - Bill Gates [[GATES-640K](#)]

5. On RIR->LIR->End User Assignment and Allocation Policies

5.1. On the HD Ratio Metric

The current use of the HD Metric is documented in [[RIR-IPV6](#)]. An HD ratio of .8 is the threshold for justifying additional allocations. Now that we have some operational experience with using the HD-Ratio, we are seeing some (on the surface surprisingly) large allocations being made. They are "surprising" in the sense that people who look at specific examples feel that the current policies grant far more address space than a service provider actually needs to (comfortably) number their infrastructure and customer networks. Using the example from [Section 4.1](#), a large cable/DSL provider could obtain a /20, even though that corresponds to a utilization of only 2.1% (In IPv4, 80% utilization would be required).

The topic of various HD Ratio values and their implications are discussed in more detail in [huston-hd-ratio]. That work suggests that changing the HD ratio threshold from .8 to .94 reduces the overall address space consumption by an equivalent of 3 bits of address space over the next 50 years, without imposing hardship on ISPs/LIRs in numbering their networks. In addition, the work also suggests that a small number of "big" allocations accounts for a disproportionately large amount of total projected address space consumption.

5.2. On /48 Assignments to End Sites

Per existing policy, the RIRs assume that end site assignments are /48s, and thus utilization measurements are based on an HD-ratio that counts numbers of /48 assignments. Granting a /48 to an end site, allows a site to have up to 65,536 subnets. While that number may make sense for larger sites, it is hard to imagine a typical home

user requiring so much space. Indeed, the default end site assignment today is in practice the same for home users and larger businesses.

Looking back at some of the original motivations behind the /48 recommendation [[RFC 3177](#)], one overriding concern was to ensure that end sites could easily obtain sufficient address space without having to "jump through hoops" to do so. For example, if someone felt they needed more space, just the act of asking would at some level be sufficient justification. As a comparison point, in IPv4, typical home users are given a single public IP address (even this is not always assured), but getting even a small number more is typically more expensive either in terms of the effort needed to obtain additional addresses, or in the actual monthly cost. (It should be noted that the "cost" of additional addresses cannot be justified by the actual cost of those addresses as charged by the RIRs, but the need for additional addresses is sometimes used to imply a different type or "higher grade" of service, for which some ISPs charge extra.) Thus, an important goal in IPv6 was to significantly change the default and minimal end site assignment, from "a single address" to "multiple networks".

Another concern was that if a site changes ISPs and subsequently renumbers, renumbering from a larger set of "subnet bits" into a smaller set is particularly painful, as it it can require making changes to the network itself (e.g., collapsing links). In contrast, renumbering a site into a prefix that has the same number (or more) of subnet bits is easier, because only the top-level bits of the address need to change. Thus, another goal of the [RFC 3177](#) recommendation is to ensure that upon renumbering, one does not have to deal with renumbering into a smaller subnet size.

The above concerns were met by the /48 recommendation, but could also be realized through a more conservative approach. For example, one can imagine "classes" of network types, with default sizes for each class. For example:

- a PDA-type device with a low bandwidth WAN connection and one additional PAN connection - a single network or /64 assignment. A /64 assignment allows for the addressing of a number of hosts, each connected to the same PAN (personal area network, e.g., Bluetooth) link as the device. This would be appropriate in deployments where the end device is not expected to provide connectivity for a larger site, but is intended to provide connectivity for the device and a small number of additional devices directly connected to the same PAN as the device.
- home user - expected to have a small number of subnets, e.g., less than 10 - a /56 (or /60) assignment

- small business/organization - one having a small number of networks, e.g., less than 100 - a /56 assignment
- large business/organization - an organization having more than 100 subnets - a /48.

A change in policy (such as above) would have a significant impact on address consumption projections and the expected longevity for IPv6. For example, changing the default assignment from a /48 to /56 (for the vast majority of end sites) would result in a savings of up to 8 bits, reducing the "total projected address consumption" by (up to) 8 bits or two orders of magnitude. (The exact amount of savings depends on the relative number of home users compared with the number of larger sites.)

One can, of course, imagine a policy supporting the entire range of assignments between /48 and /64, depending on the size or type of end site. However, there are a number of arguments in favor of having a small number of classes of sizes:

- It becomes easier for end users to identify an assignment size that is appropriate for them
- It increases the likelihood that there will be agreement among ISPs on the appropriate assignment sizes for the various customer classes.
- Having excess flexibility in selecting an appropriate assignment size for a given customer type can lead to different ISPs offering different assignment sizes to the same customer. This is undesirable because it may result in
 - an end site needing to renumber into a smaller subnet space when switching ISPs, or
 - it may lead to ISPs attempting to differentiate their service offerings by offering the most liberal address assignment policies (to attract customers), defeating the purpose of having multiple class assignment sizes.
- The operational management of the reverse DNS tree is simplified if delegations are on nibble boundaries (e.g., /48, /52, /56, and /64).

5.3. On the /64 Boundary

In theory, even more savings could be realized by reducing the size of the Interface identifier, i.e., making it smaller than 64 bits. While possible, this is the most difficult boundary to move in practice. Considerations include:

- Stateless address autoconfiguration [[RFC 2462](#)] assumes Interface identifiers are fixed at 64 bits. Changing this would require revising that specification and devising a transition plan for migrating existing implementations from 64-bit identifiers to something shorter.
- Stateless address autoconfiguration has been widely implemented, with additional implementations in the pipeline. Removing those implementations and replacing them with upgraded implementations will take many years.
- it is unclear how one would transition to Interface identifiers of shorter length in the short term while preserving stateless address autoconfiguration.
- one transition strategy might be to disable stateless address autoconfiguration completely (for generating addresses) and use DHCP to assign addresses. However, client implementation of DHCP for address configuration is not mandatory in IPv6, and it is believed that few IPv6 devices actually implement the client portion of address configuration via DHCP.
- some existing IPv6 multicast standards assume that an IPv6 routing prefix is no more than 64-bits in length, and include the 64-bit subnet prefix within an IPv6 multicast address [[RFC3306](#), [RFC3956](#)].

5.4. Summing It All Up

A fundamental question to ask is whether the Internet community is comfortable that the current allocation policies will result in a reasonable address consumption projections over the lifetime of IPv6, and if not where and how does it make sense to propose changes.

It should be noted that some have argued:

We are currently only allocating address space out of prefix 001 and are holding the remaining 7/8ths of the available space in reserve. If we later determine we have been too liberal in our address usage, we can just implement more conservative policies

for the remainder of the space.

There are a number of reasons to give pause to that line of thinking:

- The Internet has become a critical component of the world's infrastructure, and its overall stewardship is a matter of public concern. Indeed, IP addresses are a global public resource that must be managed prudently and fairly.
- It is far easier in practice to (later) liberalize a more conservative allocation policy than it is to convert a more liberal policy to one that is more conservative. Indeed, the IPv4 experience, in which early adopters were able to obtain large amounts of address space (e.g., class A and B blocks) while later adopters got significantly less, has led to significant and continuing unhappiness about unfairness in IPv4 address management.
- Relatively small policy adjustments now would appear to significantly reduce the projected consumption of IPv6 address space, without causing significant hardship.

In terms of ease of implementation, changing the HD ratio metric from .8 to something larger (e.g., .94) is the least painful. It has no impact on the assignments made to end sites and affects only the very largest of sites (at least in the short term). It would also be an easy value to change back downward (in the future), if needed. Making this type of change will reduce projected address space consumption by approximately one order of magnitude.

Changing the default end site assignment sizes (e.g., moving from /48 to /56 for SOHO sites) has the potential for realizing even larger savings, i.e., on the order of two orders of magnitude. However, some issues apply:

- a (relatively) small number of assignments have already been made. What should be done with those? How should they be counted? Should an attempt be made to reclaim them?
- The existing minimum RIR->LIR allocation size of /32 is based on the assumption that end sites will be getting /48s. If the default size of end site allocation is changed, should the default minimum RIR->LIR allocation size also be adjusted downward?

Finally, some of the issues relating to changing the size of the Interface identifier have been discussed in [Section 5.3](#). In addition,

the above concerns related to changing the /48 boundary would also apply to changing the /64 boundary, and in all likelihood would be even more complex (e.g., to deal with transition issues).

In summary, in deciding which of the above factors to change, it is important to find an adequate balance between obtaining reasonable utilization for the long term balanced against the desire to maximize aggregatability within the routing table and to limit the amount of forwarding table route fragmentation. In addition, because IPv6 address space is a shared global resource, it must be managed in a prudent and fair manner.

6. On Reservations and Aggregation

6.1. On RIR Reservations

In order to allow subsequent allocations from an LIR to expand the size of an existing address block allocation (as opposed to requiring the allocation of a different, non-aggregatable block), RIRs need to hold some space in reserve adjacent to an allocated block to allow for future growth. The reservation can be an explicit reservation, or it can be implicit if sparse allocation is employed [[SPARSE](#)].

As an RIR allocates address blocks to LIRs, at some point it will need to go back to the IANA to obtain additional space. In order to allow for good aggregation over very long periods of time (e.g., 10 years or more), it is important that sufficient space be held in reserve to allow for potential growth (i.e., subsequent adjacent allocations) over long time frames. That is, to allow for future aggregation, it would be better for RIRs to obtain additional space from IANA than to have them potentially fragment the address space by allocating space that would better be held in reserve for future adjacent allocations .

An important policy question is how much space should be held for potential growth, the time period over which the space should be held in reserve, and how it is accounted for (e.g., when justifying the need for additional space from IANA). That is, what is the metric an RIR uses to justify going back to IANA for an additional address block. Does it count the reserved portion as being "allocated"? Does the RIR just use the same HD-ratio metric or a fixed percentage like IPv4?

At the present time, it is unclear how much IPv6 address space the RIRs are holding in reserve for subsequent allocations, but there is some evidence that it is not very much, e.g., on the order of a

single bit.

6.2. On IANA->RIR Assignments

A current proposal under discussion with the RIR community calls for having a /12 being the default IANA->RIR allocation unit [IANA-RIR-IPV6]. This size appears to be large enough to allow RIRs to implement sparse allocation or a related scheme and hold sufficient space in "reserve" for the time being.

[XXX need to fill in/explore this space in more detail.]

Note that it generally makes no sense for IANA to hold additional space in reserve when delegating space to an RIR, because an RIR will sub-allocate chunks to LIRs out its delegation. In order for a subsequent allocation for one of those LIRs to aggregate, free space must exist adjacent to the requesting LIR's existing allocation. Any space held in reserve by IANA can only be adjacent to the left most or right most sub-delegated block.

7. Security Considerations

8. IANA Considerations

This document makes no requests to IANA.

9. Acknowledgments

This work has been greatly facilitated and influenced by discussions held within the IAB ad-hoc IPv6 committee, hallway and public discussions at ARIN XV and RIPE 50 meetings, and by discussions with Geoff Huston in particular. It has also benefit from review comments from Andrew Dul, Geoff Huston, Lea Roberts and Leo Vegoda.

10. References

[GATES-640K]

<http://www.brainyquote.com/quotes/quotes/b/billgates103747.html>

[IAB-IANA98] <http://www.iab.org/documents/docs/ipv6-address-space.html>

[IPV6-STATS] <http://www.ripe.net/rs/ipv6/stats/>

- [RFC 3194] The H-Density Ratio for Address Assignment Efficiency An Update
on the H ratio. A. Durand, C. Huitema. November 2001.
- [RFC 1715] The H Ratio for Address Assignment Efficiency. C. Huitema.
November 1994.
- [RFC 3513] Internet Protocol Version 6 (IPv6) Addressing
Architecture. R. Hinden, S. Deering. April 2003.
- [RFC 3513bis] [draft-ietf-ipv6-addr-arch-v4-04.txt](#) (Approved by IESG,
awaiting publication in RFC Editor Queue.)
- [RFC 2462] IPv6 Stateless Address Autoconfiguration. S. Thomson, T.
Narten. December 1998.
- [RFC3306] Unicast-Prefix-based IPv6 Multicast Addresses. B. Haberman,
D. Thaler. August 2002.
- [RFC 3177] IAB/IESG Recommendations on IPv6 Address Allocations to
Sites. IAB, IESG. September 2001.
- [RFC 3956] Embedding the Rendezvous Point (RP) Address in an IPv6
Multicast Address. P. Savola, B. Haberman. November 2004.
- [RIR-IPV6] ARIN: <http://www.arin.net/policy/index.html#six>; RIPE
Document ID: ripe-267, Date: 22 January 2003
<http://www.ripe.net/ripe/docs/ipv6policy.html>;
APNIC:
<http://www.apnic.net/docs/policy/ipv6-address-policy.html>
- [RIR-IANA-IPV6] http://www.arin.net/policy/proposals/2004_8.html
(similar discussion can be found at the other RIRs
as well)
- [shim6] shim6 is not yet a WG, but will be soon (hopefully). mailing
list: shim6@psg.com, subscribe:
majordomo:shim6@psg.com, archive:
<http://ftp.psg.com/pub/lists/shim6>.*
- [huston-hd-metric] [draft-huston-hd-metric-00.txt](#)
- [WIKIPEDIA] http://en.wikipedia.org/wiki/IP_address
- [CENSUS] <http://www.census.gov/ipc/www/worldpop.html>

[SPARSE] "IPv6 Address Space Management", Document ID: ripe-343,
 Date: 22 February 2005, Paul Wilson, Raymond Plzak,
 Axel Pawlik,
<http://www.ripe.net/ripe/docs/ipv6-sparse.html>

11. [Appendix A](#): HD-Ratio

The following table provides equivalent absolute and percentage address utilization figures for IPv6 prefixes, corresponding to an HD-Ratio of 0.8

P	48-P	Total /48s	Threshold	Util%
48	0	1	1	100.0%
47	1	2	2	87.1%
46	2	4	3	75.8%
45	3	8	5	66.0%
44	4	16	9	57.4%
43	5	32	16	50.0%
42	6	64	28	43.5%
41	7	128	49	37.9%
40	8	256	84	33.0%
39	9	512	147	28.7%
38	10	1024	256	25.0%
37	11	2048	446	21.8%
36	12	4096	776	18.9%
35	13	8192	1351	16.5%
34	14	16384	2353	14.4%
33	15	32768	4096	12.5%
32	16	65536	7132	10.9%
31	17	131072	12417	9.5%
30	18	262144	21619	8.2%
29	19	524288	37641	7.2%
28	20	1048576	65536	6.3%
27	21	2097152	114105	5.4%
26	22	4194304	198668	4.7%
25	23	8388608	345901	4.1%
24	24	16777216	602249	3.6%
23	25	33554432	1048576	3.1%
22	26	67108864	1825677	2.7%
21	27	134217728	3178688	2.4%
20	28	268435456	5534417	2.1%
19	29	536870912	9635980	1.8%
18	30	1073741824	16777216	1.6%
17	31	2147483648	29210830	1.4%
16	32	4294967296	50859008	1.2%
15	33	8589934592	88550677	1.0%

14	34	17179869184	154175683	0.9%
13	35	34359738368	268435456	0.8%
12	36	68719476736	467373275	0.7%
11	37	137438953472	813744135	0.6%
10	38	274877906944	1416810831	0.5%
9	39	549755813888	2466810934	0.4%
8	40	1099511627776	4294967296	0.4%
7	41	2199023255552	7477972398	0.3%
6	42	4398046511104	13019906166	0.3%
5	43	8796093022208	22668973294	0.3%
4	44	17592186044416	39468974941	0.2%

12. Author ÇÖs Address

Thomas Narten
IBM Corporation
3039 Cornwallis Ave.
PO Box 12195 - BRQA/502
Research Triangle Park, NC 27709-2195

Phone: 919-254-7798
EMail: narten@us.ibm.com

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Statement

Copyright (C) The Internet Society (2005).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

