A Study on the IPv6 Address Allocation and Distribution Methods

Disclaimer: Some of the thoughts and opinions invested in this study are elements of the research made by students towards their research studies at NAv6. The views here reflect the views of the authors at the time of writing this article. Even standards change over time, so may the views expressed here.

28 July 2009

Contents

Executive Summary ... **v**

1. Introduction ... **1**
1.1 Background ... 1
1.2 The Internet and Internet Addressing ... 2
1.3 IP Addresses ... 3
1.4 IPv6 Address architecture ... 4

2. Existing Internet Address Allocation Model ... 6
2.1 IPv4 Address Allocations ... 7
2.1.1 Existing IPv4 Address Allocation and Assignment Policies ... 7
2.2 IPv6 Address Allocations ... 8
2.2.1 Existing IPv6 Address Allocation and Assignment Policies ... 8

3. The Proposed Country based Internet Registries (CIR) Model ... **10**
3.1 IPv6 Address Allocation Size Policies ... 11
3.1.1 IPv6 address allocation size in the existing RIR model ... 11
3.1.2 IPv6 address allocation size in the proposed CIR model ... 13
3.2 Feasibility of the proposed CIR model ... 15
3.2.1 Technical feasibility ... 15
3.2.2 Economic feasibility ... 16
3.2.3 Social feasibility ... 16
3.2.4 Timing of the Implementation of the CIR model ... 17

4. Pros and Cons of the CIR model ... **18**
4.1 IPv6 address allocation ... 18
4.1.1 Technical requirements ... 18
4.1.2 Managerial requirements ... 20

5. Implementing the CIRs ... **27**

6. Summary ... **29**

Appendix A: IPv4 Address Architecture ... **31**

Appendix B: Allocated IPv4 Address Space ... **35**

Appendix C: IPv6 Address Space and Global Unicast Address Allocations ... **44**

Appendix D: Rationale for the size of address space to be allocated or assigned to end sites in the CIR model ... **48**

Appendix E: Dual-homing effects to routing table size and scalability ... **51**

Appendix F: BGP routing table size and scalability issues. ... **52**

List of Figures

Figure 1-1. General format of IPv6 Global Unicast address 5

Figure 2-1. Existing model for Internet address allocation... 6

Figure 2-2. Current IPv4 Address allocation hierarchy and policies on sizes 8

Figure 2-3. Current IPv6 Address allocation hierarchy and policies on sizes 9

Figure 3-1. Current IPv6 Global unicast address allocation to the RIRs by IANA........... 12

Figure 3-2. Proposed model of ITU allocating address blocks to countries 13

Figure 3-3. Proposed model of ITU IPv6 Address allocation hierarchy
and policies on sizes. 14

Figure A-1. General IP address format 32

Figure A-2. IPv4 Class based address format 33

Figure A-3. Subnet addressing structure 34

Figure D-1. Proposed IPv6 Global unicast address allocation to
national registries by ITU... 48

Figure F-2. Active BGP (FIB) entries for IPv6 56

Figure F-3. Allocation of IPv6 address blocks to RIRs... 56

Figure F-4. Comparing BGP routing table growth with Moore's Law. 58

List of Tables

Table 1-1. Address Type Identification 5

Table D-1. IPv6 Address delegation recommendations 48

Table F-1. BGP Table size for IPv4 53

Figure F-1. Active BGP (FIB) entries for IPv4 54

Table F-2. Summary of the IPv4 BGP network 54

Table F-3. BGP Table size for IPv6 55

Table F-4. Summary of the IPv6 BGP network 55

Executive Summary

This document has been created to propose and review the creation of an additional parallel structure to the RIRs for the allocation and distribution of IPv6 addresses to the global community as requested by some member nations to the ITU. This parallel structure would create the RIR equivalents within each country. The main aim of this document is to study the viability of this proposal from the technical and operational view point.

In this proposed structure, ITU obtains a pool of IPv6 addresses from IANA, similar to that of the existing RIRs. ITU then allocates IPv6 address blocks to requesting nations Internet Registry called CIR (Country Internet Registry). The CIRs then sub-allocate the received IPv6 addresses from the ITU to the ISPs and users within their country based on local policies.

Based on our finding after an in depth study, we can conclude that the proposed CIR model is absolutely viable both technically and operationally. The CIR model will not introduce any radical change in terms of IPv6 address allocations and assignment. The CIR model will serve as an addition to the existing RIR model and will co-exist such that the user has a choice from whom they wish to obtain the IPv6 addresses. The CIR model does not disturb the existing infrastructure nor introduces any new infrastructure. It follows the same routing architecture and routing algorithm for routing information in the Internet.

The overall number of prefixes added to the routing tables of the core routers in the Internet would remain the same whether RIR model alone exists or both the RIR and the CIR model exists. *As such the CIR model does not impact or threaten the global Internet stability and routability.*

Compared to the RIR model, the CIR model can follow a more fairly balanced aggregation and conservation goals through proper allocation of needed address space to end-sites. Also, a CIR being closer to the country's user would be able to cater better to the local needs of the end user in the country and also provide better check on the credentials of the applicants, thus enhancing an important Internet goal i.e., the conservation of IP addresses.

Existing Internet users are familiar with the RIR model. It will be an uphill task for ITU to challenge and change this mindset. However, with a proper implementation plan, and good and fair policies, the above perception can be defeated.

The ITU could be the best alternative to manage this additional parallel structure and to act as an intergovernmental, multilateral, multi-stakeholder international body to ensure that the Internet evolves in a direction that protects and advances the fair distribution of the global internet resources.

1. Introduction

1.1 Background

To seek for solutions to address the issues identified via TD14 Rev.6 (PLEN/2), many developing countries had requested that the Telecommunication Standardization Bureau (TSB) become an additional registry for Internet Protocol (IP) addresses so that countries could have the option of obtaining IP addresses directly from the International Telecommunications Union (ITU).[1] Based on this account, this document studies the possibility of having an alternative model to the existing Regional Internet Registries (RIR) to allocate IPv6 address block to countries. It is exemplified here how such an arrangement can be made for countries to have an option of obtaining IPv6 addresses from an alternative source via the ITU. This document studies the situation from the technical and managerial aspects of this alternate mode of IPv6 address allocation.

This document recognizes and appreciates the good work of the RIRs, who are regarded as service organizations for the contributions they have made in maintaining the stable operation and functioning of the Internet. With the expert advice through voluntary research organizations such as the Internet Engineering Task Force (IETF) and the Internet Architecture Board (IAB), the support of the Internet Assigned Number Authority (IANA) and Internet Service Providers (ISPs), and now with the coordination of the Number Resource Organization (NRO), the RIRs have developed standards and policies for the allocation and assignment of IP addresses. Along with the standards and policies, the RIRs have evolved with the Internet, matured with experiences and learned needs. The RIRs through the developed standards and address allocation policies have a strong influence on the structure and operation of the Internet.

Respecting the RIRs, this study is not to undermine their capabilities or existence, but to study the possibility of an IPv6 address allocation and distribution method that would serve in addition to the existing RIRs with emphasis given to the local needs of a country.

[1] see 6 of TD 30 Rev.1 (PLEN/3)

To this end, the ITU as an intergovernmental organization within the UN (United Nations) system that has a special partnership with governments and industry members ever since it was created in 1865 is the viable alternative to the RIRs. The ITU would obtain a large block of IPv6 address from IANA and delegate it to countries who request it through their Country-based Internet Registries (CIR) to be sub-allocated to users following their own local policies. The CIRs would retain the virtue of the RIRs and would work in close cooperation with them in the interest of the Internet, its services and the user community.

1.2 The Internet and Internet Addressing

Internet is an interconnected collection of networks. The Internet has evolved from a research based network in the 1970s to today as a critical public infrastructural capability for communication. The set of layers, protocols and standards defines the Internet architecture. There are different types of Internet architectures or models that include ISO/OSI (International Standards Organization/Open Systems Interconnection) model and the TCP/IP (Transmission Control Protocol/Internet Protocol) model. Of these, the TCP/IP mode is predominantly used over the Internet. The TCP/IP model defines a set of layered communication protocols used for the Internet and other similar networks popularly called as the Internet protocol suite, of which TCP and IP are the most important. To communicate using the Internet, a host must implement the Internet protocol suite. The Internet Protocol (IP) is a set of rules and procedure used to communicate between two hosts over the Internet.

The Internet Architecture requires a global addressing mechanism called IP address for a computer in a network to identify and communicate with computers on any other network. An IP address contains information to identify the receiving host, locate the host and identify the sequence of path called route to reach the destination.

1.3 IP Addresses

The Internet generally comprises of LANs and WANs as the elements where LANs comprise of hosts interconnected confined to a small geographical area. WANs encompass geographically dispersed hosts where LANs are interconnected by WANs. In the Internet Information is transported in terms of packets and these constituent elements provide the packet transport. The LANs and WANs are connected through routers or gateways for packet forwarding. The hosts, routers and gateways are uniquely identified by an IP address.

The global Internet address space offers hosts unique addresses within its defined space. For IPv4, there is 2^{32} = 4,294,967,296 possibilities and for IPv6 there are 2^{128} = 340,282,366,920,938,463,463,374,607,431,768,211,456 or $3.4*10^{38}$ possibilities of individual hosts.

In reality, blocks of addresses are allocated to organizations where these addresses allocated are globally reachable through ISPs. This reachability is ensured by routing protocols. Since each network is independent they may use different routing algorithm and since each network is independent of all the others, it is often referred to as an Autonomous System (AS).[2] Originally, an AS is controlled by a single entity, namely an ISP or a large organization. But according to RFC 1930, an AS is a collection of routing prefixes clearly defined by a single routing policy where the routing prefixes may be under the control of one or more organizations or network operators. As such, the Internet can be seen as an interconnected collection of subnetworks or ASs. Routing protocols such as OSPF (Open Shortest Path First) and BGP (Border Gateway Protocol) are used to advertise and route these collection of networks and addresses.

---

[2] Andrew S. Tanenbaum, Computer Networks, Fourth Edition, PHI Pvt Ltd., 2006. pp. 427.

1.4 IPv6 Address architecture[3]

IPv6 addresses are 128 bits in length. As such IPv6 addresses are very long and can be represented in the following format.

The preferred form of representing IPv6 addresses as text strings are `x:x:x:x:x:x:x:x`, where each `x` represents 16 bits or four hexadecimal digits. For e.g.

    ABCD:EF01:2345:6789:ABCD:EF01:2345:6789
    2001:DB8:0:0:8:800:200C:417A

In order to write addresses containing zero bits easier, the syntax "`::`" is used to compress the zeros. The use of "`::`" indicates one or more groups of 16 bits of zeros. The "`::`" can only appear once in an address and can also be used to compress leading or trailing zeros in an address. For e.g.

> A unicast address `2001:DB8:0:0:8:800:200C:417A` can be represented as `2001:DB8::8:800:200C:417A`
>
> A multicast address `FF01:0:0:0:0:0:0:101` can be represented as `FF01::101`

An IPv6 address prefix is written as, IPv6 address/prefix-length. Where, IPv6 address follows any of the notations explained above and the prefix-length is a decimal value specifying how many of the left most contiguous bits of the address comprises the prefix. For e.g. `2001:0DB2::CD3/60` denotes that the most significant 60 bits of this unicast address is the prefix.

The type of an IPv6 address is identified by the high-order bits of the address as given in Table 1-1.

[3] More information on IPv6 addressing architecture can be found at RFC4291.

Table 1-1. Address Type Identification

Address type	Binary Prefix	IPv6 notation
Unspecified	0…. 0 (128 bits)	::/128
Loopback	0…. . 1 (128 bits)	::1/128
Multicast	11111111	FF00::/8
Link-local unicast	1111111010	FE80::/10
Global unicast	(everything else)	

IPv6 unicast addresses are aggregatable with prefixes of arbitrary bit-length, similar to IPv4 addresses under Classless Inter-Domain Routing (CIDR)[4]. Currently, IPv6 Unicast addresses are generically structured as a two part address: a 64-bit Topology part, used by routers to forward a packet to its intended destination network, and a 64-bit Interface Identifier, that identifies a particular end point. The general format for IPv6 Global unicast address is as given in Fig. 1-1. The global routing prefix is a hierarchically structured value assigned to an organization and the subnet ID identifies a subnet within the organization. The Interface ID identifies the interface on a link. The interface IDs are 64 bits long constructed in modified EUI-64[5] format.

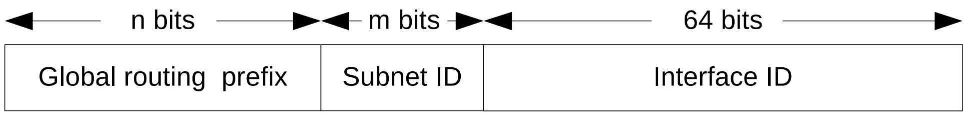

Figure 1-1. General format of IPv6 Global Unicast address

[4] Appendixes A details on IPv4 address architecture that includes class based addressing, subnets, and CIDR.
[5] IEEE, "Guidelines for 64-bit Global Identifier (EUI-64) Registration Authority", http://standards.ieee.org/regauth/oui/tutorials/EUI64.html, March 1997.

2. Existing Internet Address Allocation Model

IANA is responsible for global coordination of the IP addressing systems as well as Autonomous System (AS)[6] numbers used for routing Internet traffic.[7] In RFC 1881, the IETF recognizes IANA as the central authority on the management and allocation of IPv6 address space for the good of the Internet community with advice from the IAB (Internet Architecture Board) and IESG (Internet Engineering Steering Group). Both IPv4 and IPv6 addresses are generally assigned in a hierarchical manner on a provider based arrangement. IANA allocates IP address blocks to the RIRs for further allocation to their members. The hierarchical management structure of the IPv4 and IPv6 address space is that IANA allocates IP address space to RIRs, who in turn allocate address space to LIRs within their respective regions, who then assign them to end users. The existing model for IP address allocation is shown in Fig. 2-1.

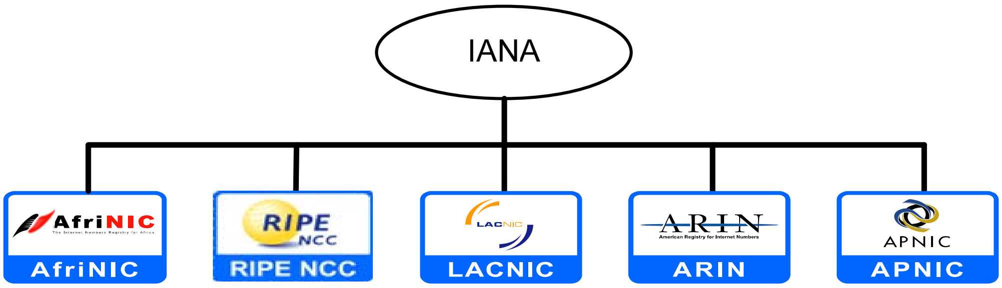

Figure 2-1. Existing model for Internet address allocation

IANA allocates IP addresses from the pool of unallocated addresses to the RIRs based on their needs as described by global policies and standard documents such as RFC's made by the IETF. Subsequent allocations are made by IANA to RIRs when they require more IP addresses for allocation or assignment within their region. IANA does not allocate IP addresses directly to ISPs or end users except in specific circumstances, such as allocations of multicast address or other protocol specific needs.

[6] See section 1.3 on IP Addresses in page no. 3 for definition on AS numbers.
[7] http://www.iana.org/numbers/

The four RIRs – APNIC, ARIN, LACNIC and RIPE NCC have formed the Number Resource Organization (NRO) through a MoU. The purpose of the NRO is to undertake joint activities of the RIRs, including joint technical projects, liaison activities and IP address allocation policy co-ordination.[8]

2.1 IPv4 Address Allocations

The allocation of IPv4 address space to various registries is listed in Appendix B. Also, RFC1466 documents most of these allocations.

2.1.1 Existing IPv4 Address Allocation and Assignment Policies

Originally, all the IPv4 addresses where directly assigned by IANA. Later parts of the address space were allocated to various registries to manage for particular purposes or regional areas of the world. Since the introduction of the CIDR system, IANA typically allocates address space in the size of '/8' prefix blocks for IPv4. The size of minimum IPv4 portable allocation varies from a /20 to /22 between RIRs, and the minimum size of a sub-allocation is /24. The Fig. 2-2 shows the current IPv4 address allocation hierarchy and policies on sizes.

More information on the IPv4 address allocation and assignment policies can be found at the respective RIRs website.

[8] Source: http://www.nro.net/about/index.html

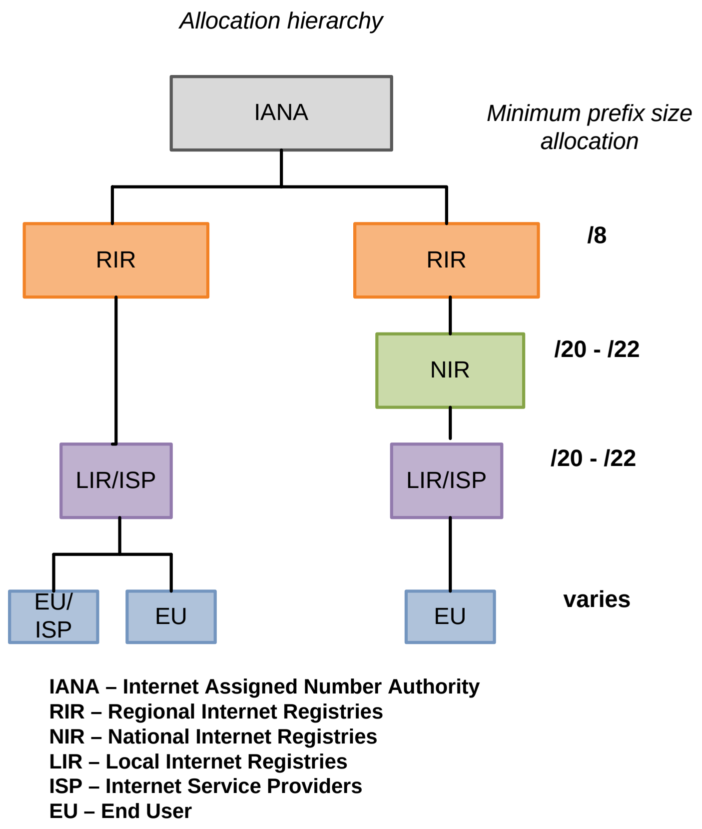

Figure 2-2. Current IPv4 Address allocation hierarchy and policies on sizes

2.2 IPv6 Address Allocations

The IPv6 address space and Global Unicast Allocations to various registries is given in Appendix C.

2.2.1 Existing IPv6 Address Allocation and Assignment Policies

The IPv6 Global Unicast space encompasses the entire IPv6 address range with the exception of ::/128, ::1/128, FF00::/8, and FE80::/10 [RFC4291]. Refer to section 1.4 for more information on these four address ranges. IANA Global Unicast address assignments are currently limited to the IPv6 unicast address range of 2000::/3.[9] These assignments are made to RIRs, and the address assignments from this block are registered in the IANA registry (http://www.iana.org/assignments/ipv6-address-space/). The

[9] RFC 4147, http://www.faqs.org/rfcs/rfc4147.html

fraction of the IPv6 address space occupied by the aggregatable global unicast addresses with the Format Prefix (001) is 1/8. The current set of unicast addresses that includes aggregatable global unicast address, link-local unicast address and site-local unicast address (deprecated), represent only 15 % of the entire IPv6 address space.[10]

Fig. 2-3 shows the current IPv6 address allocation hierarchy and policies on sizes. The current minimum IPv6 allocation from IANA to an RIR is /12, minimum IPv6 allocation from RIR to a NIR or LIR/ISP is /32 and the assignment of address space by LIR/ISP to end user could vary between /48 and /64. More information on the IPv6 address allocation and assignment policies can be found at the respective RIRs website.

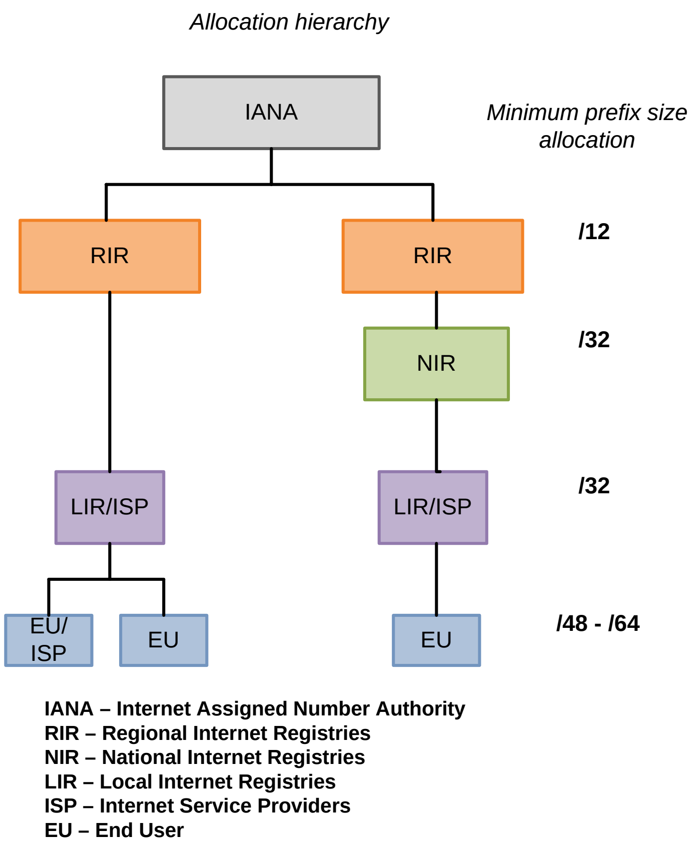

Figure 2-3. Current IPv6 Address allocation hierarchy and policies on sizes

[10] IPv6 Address Space, Updated Jan 21 2005, Microsoft TechNet. http://technet.microsoft.com/en-us/library/cc781652(WS.10).aspx

3. The Proposed Country based Internet Registries (CIR) Model

The structure of the proposed Country based Internet Registry (CIR) is similar to the existing RIRs. ITU obtains a pool of IPv6 addresses from IANA and delegates addresses from the obtained token space to the requesting nations Internet Registry called Country-based Internet Registry (CIR) following similar rules in terms of address aggregation so that the address allocation are scalable. The CIRs in turn would sub-allocate IPv6 addresses to the ISPs and users within their country based on the fundamental policies created by the RIRs. These policies can be further enhanced by local policies which can further benefit the nation's user community.

Address allocation system by ITU to CIRs

The allocated IPv6 address space to ITU would be delegated to the requesting member nations[11] using Growth based Address Partitioning Algorithm (GAP)[12] or similar address partitioning and allocation schemes. Equal or uniform partitioning of the total address space among member nations may not be efficient as each nation requires different address space owing to their size and have different growth rates. The allocation scheme would facilitate to allocate the IPv6 address space more efficiently and fairly among nations. The address allocation scheme can use the following criteria,

a) The size of population

b) Growth rate in terms of utilization of the IPv6 address space[13]

c) Business and organizational growth.

As such ITU reserves at a minimum, contiguous address blocks as distinct address sets from the IANA allocated address space looking at it as a common address pool such that each of the address sets or token space are distinguished by countries and subsequent

[11] As on date there are 192 member states recognized by UN. http://www.un.org/en/aboutun/index.shtml
[12] Mei Wang, A growth based address allocation scheme for IPv6, Networking, 2005.
[13] The measurement of growth rate in terms of utilization of the IPv6 address space is beyond the discussion of this study.

allocations to countries can be made contiguously to allow for aggregation as a single block.

National Internet Registries (NIRs)

Currently APNIC has a structure called NIRs[14]. However, the uptake of the NIR model has been very limited. Currently, the existing NIRs essentially process and approve IP address requests made by their countries ISPs and organizations. The address allocations however, are directly made by the respective RIRs and not by the NIRs themselves.

Alternative choices for the user: CIR and RIR

The above model of ITU allocating IPv6 address blocks to requesting nations through their RIRs must serve as an alternative registry to the existing RIRs. The users have a choice to choose among the two alternatives to source their IPv6 addresses and the existence of one will not diminish the roles or threaten the existence of the other. There will be definitely greater benefits to the country if both the CIR and RIR jointly work together, especially on baseline policies.

3.1 IPv6 Address Allocation Size Policies

3.1.1 IPv6 address allocation size in the existing RIR model

The unit of IPv6 allocation (and therefore the minimum IPv6 allocation) from IANA to an RIR is a /12.[15] The IPv6 Global Unicast Address Allocation, last updated 13-05-2008 shows that each of the five RIRs have been allocated a /12 each.[16] The current IPv6 address space allocation from IANA to the five RIRs is shown in Fig. 3-1 below.

[14] NIRs primarily operate in the Asia Pacific Region under the authority of APNIC. Under APNIC, there are NIRs in Japan, China, Taiwan, Korea, Indonesia and Vietnam. Source: http://en.wikipedia.org/wiki/National_Internet_registry and http://www.nro.net/about/rir-system.html

[15] Allocation of IPv6 Address Space by IANA Policy to RIRs. http://www.nro.net/policy/iana-rir-ipv6-allocation-proposal.html

[16] http://www.iana.org/assignments/ipv6-unicast-address-assignments/

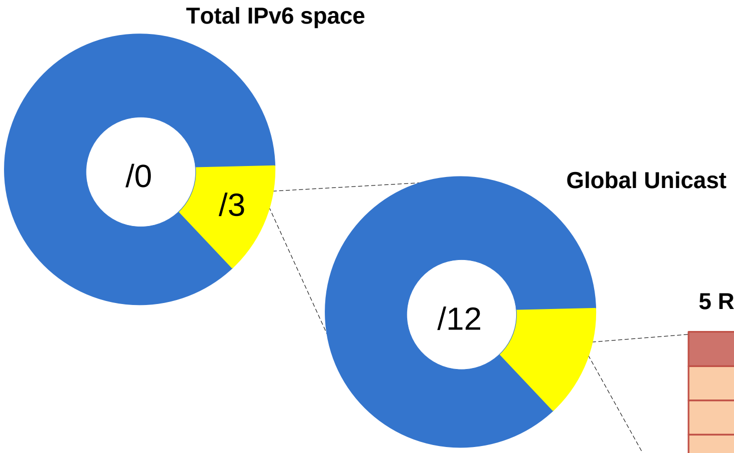

5 RIRs, each a /12 (Oct 2006)

RIR	IPv6 Address
AfriNIC	2C00:0000::/12
APNIC	2400:0000::/12
ARIN	2600:0000::/12
LACNIC	2800:0000::/12
RIPE NCC	2A00:0000::/12

Figure 3-1. Current IPv6 Global unicast address allocation to the RIRs by IANA

According to IPv6 Address Allocation and Assignment Policy defined by the RIRs, the minimum address allocation size for IPv6 address space is /32 to an LIR/ISP. This ranges to a block of 2^{96} addresses for the LIR/ISPs to allocate. This also means each LIR/ISP have up to 2^{32} = 4.3 billion subnets to allocate with 2^{64} possible addresses with each subnet.

ISP's are usually allocated with a /32 prefix which is the smallest prefix length that is globally routable without any problem. Based on the recommendations of IAB/IESG on IPv6 Address Allocations to end-sites [RFC 3177][17], the RIRs allocate a minimum of /48 to end-sites[18] in most cases. The exact choice of how much address space to assign end sites is a policy issue under the purview of the RIRs, subject to IPv6 architectural and operational considerations as substantiated in the work in progress document, "IPv6 Address Assignment to End sites by Narten et al.[19]" The said document also details the address space and rationale in allocating between a /48 to /64 to end sites.

[17] http://tools.ietf.org/html/rfc3177
[18] Information on end-sites is detailed in Appendix D
[19] http://tools.ietf.org/html/draft-narten-ipv6-3177bis-48boundary-04

3.1.2 IPv6 address allocation size in the proposed CIR model

Fig. 3-2 shows the proposed model of ITU allocating address blocks to countries and Fig. 3-3 shows the proposed model of ITU's IPv6 address allocation hierarchy and policies on sizes respectively. This proposed alternative would also facilitate in bridging the digital divide among developed and developing nations by more efficiently handling the management of IPv6 addresses.

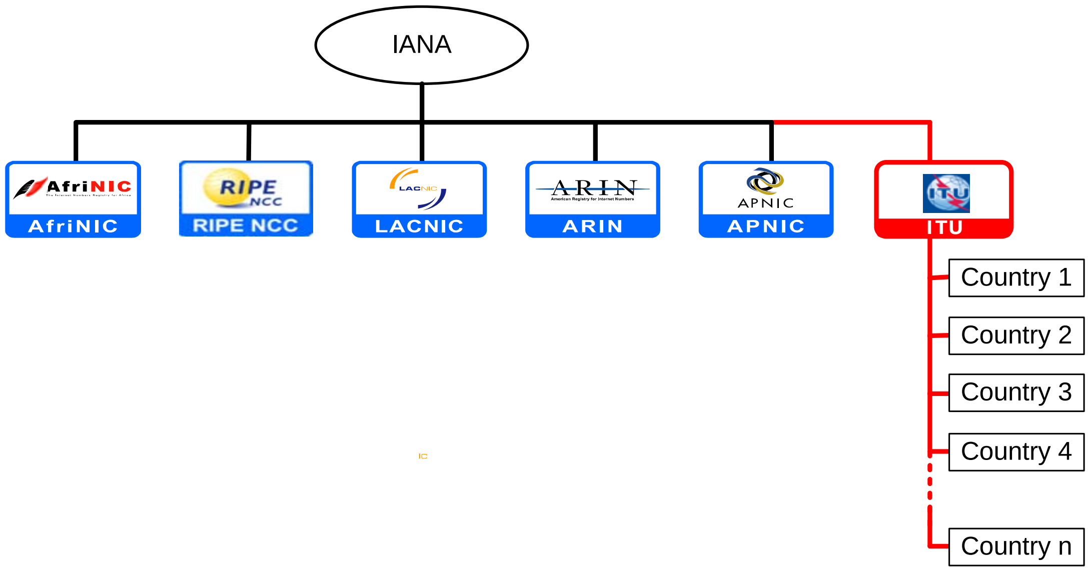

Figure 3-2. Proposed model of ITU allocating address blocks to countries

IANA similar to the 5 RIRs could allocate a minimum IPv6 allocation of /12 to ITU, which would then allocate contiguous, address blocks of /32 to each requesting nation based on their needs. ITU would delegate these IPv6 address blocks to the CIRs formed, which would then manage its national address block based on policies established preferably by ITU and also by the country.

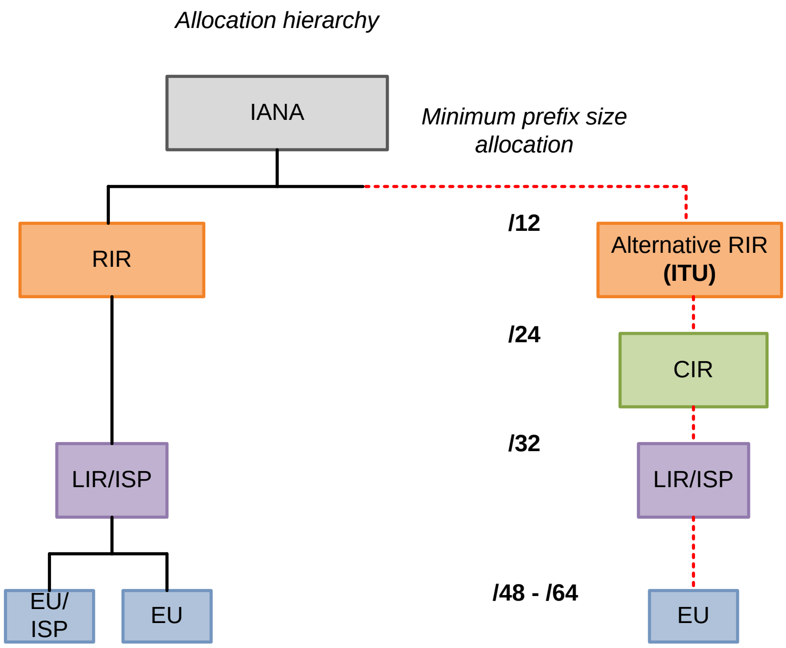

IANA – Internet Assigned Number Authority
RIR – Regional Internet Registries
ITU – International Telecommunication Union
NIR – National Internet Registries
CIR – Country based Internet Registries
LIR – Local Internet Registries
ISP – Internet Service Providers
EU – End User

Figure 3-3. Proposed model of ITU IPv6 Address allocation hierarchy and policies on sizes.

Considering a /12 prefix is initially allocated to ITU by IANA similar to the 5 RIRs, ITU allocates each nation with at least a /24 prefix. The address prefix size allocated to a country would vary based on their needs. There are 1,048,576 numbers of /32 prefix allocations possible in a /12. If a minimum of /48 is allocated to each end site, there are 65,536 number of /48 possible for a /32 prefix. A /48 prefix network would have 65,536 subnets to be used in routing with 65,536 unique host numbers within each network. This would be large enough for most enterprises or organizations allowing delegation of address blocks to aggregating entities.[20] An illustration of the above and the IPv6 address delegation size recommended for the CIRs and the ISPs within the CIRs to follow are given in Appendix D.

[20] Rationale for the size of address space to be allocated or assigned to end sites in the CIR model is given in Appendix D.

3.2 Feasibility of the proposed CIR model

3.2.1 Technical feasibility

The proposed CIR model is absolutely viable both technically and operationally. The introduction of the CIR model does not introduce any radical change in terms of IPv6 address allocation and assignment. The CIR model does not replaces the existing RIR model but serves as an alternative source for providing IPv6 addresses at the country level rather than the regional level. So both the RIRs and CIRs will co-exist such that the user has a choice from whom they wish to obtain the IPv6 addresses. The CIR model neither disturbs the current existing network infrastructure nor introduces any new network infrastructure, such as a country level gateway. It would use the existing network infrastructure established by the ISPs and will follow the current existing routing architecture and algorithms for routing information in the Internet.

Routing table size & scalability

As the CIR model,

- i) Does not introduce new network infrastructure,
- ii) Follows the current routing architecture,
- iii) Is assigned a large enough contiguous address block for the country,

It will not affect the routing table sizes on the Internet. As such the overall number of prefixes added to the routing tables of the core Internet would remain the same whether if we have a single supplier for IPv6 addresses namely the RIRs or multiple suppliers that includes the CIRs.

Provider Independent IP addresses

Today, the increased allocation of provider independent addresses and need for traffic engineering and multi-homing has led to the routing scalability problem in the core of the Internet. This existing scalability problem is said to be attributed to the growth in the demand for routing of individual address prefixes that cannot be aggregated into coarser routes and not by growth of the size of the Internet.[21] **It is important to note that this condition already occurs naturally with the current RIR model and is not affected by the introduction of the CIR model.**

More information on routing, BGP routing table size and routing table scalability issues in general can be found in Appendix E and Appendix F.

3.2.2 Economic feasibility

The CIR model would work with the existing network infrastructure provided by the ISPs and demands no additional hardware or software capabilities. Thus from an economic feasibility viewpoint, the introduction of the CIR model will not introduce any additional costs to the ISPs and those managing the Internet.

3.2.3 Social feasibility

Generally, people oppose change for different reasons. One reason among them is for no reasons they might oppose change.[22] The existing Internet users are familiar with the RIR structure. The perception that ITU is an intergovernmental organization might have an influence to resist change to the CIR model. This perception also includes a view that the CIR model would lead to dominant control by the requesting countries government over the allocation and management of IPv6 address space through their CIRs and force the existing ISPs to source IPv6 addresses from the CIRs.

The above perceptions can be defeated by conducting proper awareness programmes, and implementing well coordinated common standard address allocation and assignment policies to be followed by the CIRs. The RIRs baseline address allocation policies have worked well and should be followed by the CIRs. Additional policies must be created to benefit the Country's Internet community and must value add to policies already offered

---

[21] Routing & Addressing -- activities. source: http://www.ietf.org/mail-archive/web/ietf-announce/current/msg03255.html
[22] Fred Luthans, Organizational Behavior.

by the RIRs. These additional policies MUST NOT be restrictive, else the CIR model will face failure.

3.2.4 Timing of the Implementation of the CIR model

Is this is the right time to implement the CIR model?

The address space with IPv4 is nearing exhaustion so the CIR model would not feasible for the allocation and assignment of IPv4 addresses. However, with IPv6, which has only recently started to kick-off in terms of take-up of IPv6 addresses, the CIR model can be rightly implemented now. This would result in efficient management of IPv6 addresses and benefit the users who will now have an alternative to source there IP address leading to a healthy competition between the two models.

4. Pros and Cons of the CIR model

4.1 IPv6 address allocation

There are two general requirements to be satisfied in allocating IPv6 addresses, namely Technical requirements and Managerial requirements. The technical and managerial requirements are tied to each other as the managerial requirements are defined to the extent to satisfy technical requirements.

4.1.1 Technical requirements

i) Conservation:

IPv6 addresses are large but not infinite. As such, conservation should be practiced from an early stage. It is true that the practical number of addresses available is far less than the theoretical maximum. The reason being, of the 128 bits width of the address, the first 64 bits are used to identify the network and the lower 64 bits are used to identify the interface. The lower 64 bits usually represents the MAC address of the interface represented in EUI-64 bit format to support auto-configuration or plug and play. As such inefficient allocation of IPv6 addresses by a supplier to users would lead to early depletion of IPv6 addresses than what it would be normally possible.

Conservation with existing RIR model:

With the existing model a NIR, and an LIR/ISP can get a /32 address prefix at the minimum, and an end user site can get /48 at the minimum. A /32 and a /48 is sufficiently large that for most organizations it would be a colossal waste. For e.g. an end user site such as small organizations and home user never in this millennium would be using such a prefix size. Though IPv6 are abundant and as they say “There is no free lunch”, IPv6 address allocation could be done more prudently.

Conservation with proposed CIR model:

Looking at section 3.1.1 and 3.1.2 above it can be observed the proposed CIR model would allocate IPv6 addresses more efficiently than the existing RIR model. The CIR would be able to better understand the local needs (demonstrated needs) of the user for IPv6 addresses and allocate accordingly as they can easily check on the credentials of the applicant.

ii) Aggregation:

Routing aggregation is essential to ensure global routability of all IP addresses. IP address allocation should follow a hierarchy so that subsequent allocations can be made contiguously. This would allow a network to be identified by a single address prefix and help in reducing the routing table size.

Inappropriate IP address allocation algorithms would lead to creating new prefixes for subsequent allocations where a single user would be represented by multiple disjoint prefixes ending-up in address *fragmentation*.[23] Fragmentation of IP addresses should be avoided to the maximum as it will cause the routing system to fail resulting in discontinuation of services to many parts of the Internet.

Both the RIR and CIR model follow addressing hierarchy and strive for address aggregation. Irrespective of the RIR or CIR model, address aggregation purely depends on the address allocation algorithm and policies followed.

Aggregation in the existing RIR model:

The RIRs follow a sparse or binary algorithm to maximize aggregation of address blocks allocated. As such the chosen source address block from which allocations made are sufficiently large. This method treats all users equally in terms of size and evenly splits the allocated blocks to maximize the space between users. Being said so; the IPv6

[23] Address allocation is not the only factor that influences fragmentation; there are other issues such as multi-homing and traffic engineering which are common to both the RIR and CIR models.

address allocation and assignment policy defined by the RIRs say there can be no guarantee of contiguous allocation.[24]

Aggregation in the proposed CIR model:

In reality, different users would require different address space size and would have different growth rates. So, bisection algorithm for address space allocation may not be efficient for most users especially when they are small or they are fast growers. To overcome this inefficiency of address space utilization the CIR model can further study and simulate different address allocation algorithms.

The CIRs will take responsibility in allocating contiguous IP addresses to users when subsequent allocations are made from the distinct address space designated for a country. This may lead to better management of IPv6 addresses resulting in *better aggregation and conservation if done properly*. In conclusion to aggregation, the RIRs have allocated addresses fairly well till date. However, there is always room for improvement by both the RIRs and CIRs if newer and better algorithms are used.

4.1.2 Managerial requirements

The management of IPv6 address allocation and assignment should satisfy the technical requirement stated in 4.1.1 irrespective of whatever address allocation scheme is followed.

i) Uniqueness:

Every address space allocation and assignment of IPv6 address space must be globally unique so that every public host can be uniquely identified on the Internet.

[24] IPv6 Address Allocation and Assignment Policy, version: 006, dated 4 August 2008. Source: http://www.apnic.net

Uniqueness in the existing RIR model:

Uniqueness is ensured using registration of allocated and assigned address space in a registry database. This is a good method and has worked well.

Uniqueness in the proposed CIR model:

Uniqueness would be ensured following a similar registration procedure as followed by the RIRs.

ii) Registration:

The address space allocations and assignments must be documented in a registry database accessible to members of the Internet community. This is to ensure uniqueness of each Internet address and to provide reference information for trouble shooting at all levels of users.[25]

Registration in the existing RIR model:

The registration is available in terms of "Whois Database".

Registration in the proposed CIR model:

Similar to the existing RIR model a "Whois Database" would be made available.

iii) Fairness:

Allocation of IPv6 addresses should be done globally uniform irrespective of their location, nationality, size or any other factor.

Fairness in the existing RIR model:

The explanation follows the same as given for *Aggregation* in section 4.1.1

[25] IPv6 Address Allocation and Assignment Policy, source: http://www.apnic.net

Fairness in the proposed CIR model:

The explanation follows the same as given for *Aggregation* in section 4.1.1

iv) Minimized operational overhead:

To minimize operational overhead with obtaining and managing address space. This includes making requests frequently with the RIRs for additional address space and making a large number of small successive incremental expansions instead of making few, large expansions. This scenario existed with IPv4 as the address blocks allocated by the RIRs to a user was fairly small owing to the limited size of the address space available. So, fast growing users will have to make subsequent requests for IP addresses frequently which may be an overhead in terms of address management for both the user as well as the RIR. With IPv6, this scenario does not exist as the available IPv6 address space and the allocated address block to a user are very large.

Minimized operational overhead in the existing RIR model:

The RIRs allocate a minimum of /32 address prefix to NIR and LIR/ISPs, and a /48-/64 at the minimum to end-user sites. Though this would minimize address space allocation and management overhead, it would lead to a lot of wasteful addresses.

Minimized operational overhead in the proposed CIR model:

The CIR's will allocate IPv6 addresses prudently to the LIR/ISP and end-user sites based on their needs and requirements. The address prefix allocation and assignment range is as given in section 3.1.2. The overhead in terms of verifying the IPv6 address requirements or needs of a user may be higher for a CIR. But owing to the lower volume handled by a CIR as compared to an RIR this would not be a significant problem.

Conflicting Goals

Some of the requirements of IPv6 address allocation that includes aggregation and conservation have conflicting goals. For instance, more importance to aggregation of routing information would waste address space and address conservation could not be achieved. On the other hand, an approach of maximizing conservation would threaten aggregation, increase administrative overheads and reduce fairness in actual address policies.[26] Thus address policies can be developed balancing the above goals in the CIR model. Of these conflicting goals route aggregation is considered to be most important.

Further benefits of the CIR Model

The additional benefits of the CIR model in contrast to the existing RIR model is identified under the entities listed below.[27]

i) Region:

The CIR model could facilitate in the formation of regional agencies or Internet Registries at a smaller regional level grouping together CIRs. For example, North African Countries can group together and form a North Africa Internet Registry (NAIR). These regional level registries would be able to bring together CIRs with similar agenda and interest within their region, and coordinate closely in addressing concerns with respect to the Internet address resources. Strong neighbors could help the needy neighbor in terms of capacity building in facilities, resources and infrastructure on IPv6 through the formation of a coalition of the CIRs at the smaller regional level. The ITU could assist in coordination and bringing in these alliances for the larger benefit of the public.

ii) Country:

As the Internet and IP based services are becoming widely deployed and accepted as public infrastructure and commercial entity of national importance, overseeing the address allocation of this infrastructure at the national level should be given to a national authority for policy control that reflects the sovereignty rights of the user's within those economies. This would be similar in functionality to the liberalized radio spectrum distribution and privatized telecommunications system introducing non monopolistic governing structure of the infrastructure facility.

---

[26] Establishment of global IPv6 address policies, Takashi Arano.
http://www.isoc.org/briefings/012/briefing12.pdf

[27] Most of the thoughts here are excerpted from the article, ITU and Governance written by Houlin Zhao. Source: WG-WSIS-7/6 Rev 1. 15 Dec 2004.

The CIR of the requesting country gets to manage its own IP addresses and sovereignty connected to the registration of addresses would be safeguarded. The ITU could help coordinate to establish Internet Exchange (IX or IXP)[28] points so that the countries could consolidate internet infrastructure and routes.

With ubiquitous Internet communication and countries with varied cultural and linguistic needs, domain names in Roman script would be disadvantageous. The CIR model could facilitate multilingual domain name registration with the local language support. The CIR can provide better customer support in multilingual catering to the local needs.

It would also be more efficient for the address allocation body (CIR or RIR) to be closer to the user. In this manner, both voice, video and physical customer support can be made available to the end users by the CIR.

iii) Users:

The users of CIR model would include ISPs and organizations. The Internet being a public infrastructure has become a source of revenue and also a system for transferring funds. Economic models developed by ITU would help in bringing fair revenues to all parties involved in utilizing the Internet infrastructure.

---

[28] An Internet exchange point (IX or IXP) is a physical infrastructure that allows different Internet service providers (ISPs) to exchange Internet traffic between their networks (autonomous systems) by means of mutual peering agreements, which allow traffic to be exchanged without cost.

IXPs reduce the portion of an ISP's traffic which must be delivered via their upstream transit providers, thereby reducing the Average Per-Bit Delivery Cost of their service. Furthermore, the increased number of paths learned through the IXP improves routing efficiency and fault-tolerance.

The primary purpose of an IXP is to allow networks to interconnect directly, via the exchange, rather than through one or more 3rd party networks. The advantages of the direct interconnection are numerous, but the primary reasons are cost, latency, and bandwidth. Traffic passing through an exchange is typically not billed by any party, whereas traffic to an ISP's upstream provider is. The direct interconnection, often located in the same city as both networks, avoids the need for data to travel to other cities (potentially on other continents) to get from one network to another, thus reducing latency. Definition: Wikipedia.

Note: A national level gateway is different from an IX or IXP.

Assigning address to countries enables users to choose their preferred source of IP addresses, either from the corresponding RIR or the ITU. This would lead to a healthy competition between the CIR and the RIR by which the users would be the beneficiaries.

iv) ITU:

ITU as a member-oriented organization is always driven by inputs from its members. As such it works in a bottom-up manner. With the successful implementation of this CIR model, it would have heeded to and satisfied the needs of some of its member states that have been voicing concerns on IP address allocation in relevance to TD14 Rev.6 (PLEN/2).

Similar to the role played in coordinating public telecommunications infrastructure and services, ITU could play an important role with Internet and IP based services. The implementation of the CIR model would facilitate in this. This would also advance the non-monopolistic control of the Internet.

Summary

Together with the RIR model, the CIR model can potentially value add to the creation of a more fairly balanced IPv6 address allocation model.

> *Historically, it has been desirable to de-centralize the resource distribution function to some extent for several reasons:*
>
> - *to improve scalability*
> - *to bring the function closer to the resource users*
> - *to ease the establishment of appropriate funding structures, and*
> - *to obtain greater support from the local community.*
>
> *These important goals have motivated the establishment of the five regional registries that we have today. Decentralizations has been shown to be necessary.*

[A Fine Balance: Internet Number Resource Distribution and De-Centralization, Internet Society][29]

The ITU as intergovernmental, multilateral, multi-stakeholder international body should be able to ensure close co-ordination between the CIRs as well as the CIRs and the RIRs. As the de-centralization is at the country level, the local needs of a country can be better understood by the CIRs. As such, the CIRs would be able to further enhance the four reasons needed above to de-centralize the resource distribution function.

[29] Available at, http://www.itu.int/dms_pub/itu-t/oth/3B/01/T3B010000010001PDFE.pdf

5. Implementing the CIRs

In moving forward, ITU needs to consider the real possibility of implementing the CIRs. The challenges faced in implementing the CIRs, would include the following:

Policies: This is very important as the CIR structure proposed within this document stress on the importance of an open market concept of allocating IPv6 addresses. This means the CIRs will be open to users for address allocation request, and there will be direct competition with the RIRs. Policies formulated must reflect certain criteria which are:

- *Baseline policies consistent with the current RIR policies*
- *Provides better and more incentives than the current RIR policies*
- *Encouraging policies which promote the use of IPv6*
- *Additional, that will further benefit the users within the country*

ITU should take the lead in formulating a baseline policy plan for the CIRs. These policies must be end-user friendly and must not involve any form of censorship at the level of the CIRs. Censorship is a country right and is currently enforced at the level of the ISPs. This must be kept the same as it is now and the CIRs must not be dragged into such policies. The CIR should only focus on IPv6 allocation policies, which is their main agenda.

Services: The services offered by the CIR should exceed the services currently offered by the RIRs. The services should include

- *Web registration which is developed in the local language/languages of the CIR.*
- *Phone and Internet support, with technical staff capable of assisting the end users in their own national language.*
- *Proper and detailed evaluation process, so that the address space is not wasted or depleted.*

Clustering: It may be difficult to establish country based national registries for some of the member economies. Such member economies can group together and address allocation to these countries can be made by the group facilitated by ITU, until if and when they are ready to form their own CIRs. For example, North African Countries can get together and form a North Africa Internet Registry (NAIR). This can help the smaller and developing countries by optimizing the resources needed to run the CIRs. Most small CIRs can be a 3-6 person operation, once a fully automated web system is in place for managing the address resources.

Cost Structures: The cost structures of the way the CIRs charge their LIRs/ISPs for the IPv6 addresses must be competitive. It must, in general, be equivalent or less than the current price structure that is provided by the RIRs. Only then will users find this to be a viable alternative, as most users are price sensitive.

Expectations: As IPv6 is a new growth area in most countries, the expectation of IPv6 address allocation may not be too high in the beginning. Some small economies may be able to only allocate 5-10 blocks per annum. Thus, indicators of success must be measured based on the Internet penetration and IPv6 awareness respectively. However the rate will grow, as shown by the IPv4 uptake and the new trends in the uptake of IPv6. It is only a matter of time, and sustenance in the case of the new CIRs.

Awareness: An awareness program for the promotion of IPv6 needs to be put in place. This is in direct relations to the expectations mentioned earlier. Such awareness programs will create the need for IPv6 as well as introduce cost saving measures, that includes ensuring all new ICT equipment and software purchases meet the IPv6 requirements.

A detailed CIR Implementation Plan must be created and reviewed first, and agreed upon by all involved parties before the actual implementation takes place.

6. Summary

***Routing Issues: Myth or Fact*:** It has been a long standing myth that if parallel systems like CIRs are proposed and implemented, it would lead to address fragmentation, and routing table size would grow beyond manageability, causing the Internet to become unstable. After a closer and detailed technical study we realized that this myth is really a myth. The routing table will grow in size but, this growth is normal growth and would be the same as if only the RIRs were assigning the addresses. There will be no additional fragmentation beyond the normal expected. An observation of the BGP routes indicates that the cause for more specific routes that increases the routing table size are mainly due to site multi-homing and traffic engineering.[30] Neither of these will increase with the introduction of the CIRs.

Additional routes will be created irrespective of whether one party or two parties are allocating IPv6 addresses in parallel. There is only one case which may help further optimize the routing table size. This is when the additional blocks requested are:

i) By the same LIR/ISP

ii) To the same RIR

iii) And contiguous address allocation is provided for by the RIR to the requesting LIR/ISP.

All the above 3 conditions must be met for a contiguous address allocation, which will save ONE extra route on the routing table. But chances for all the above 3 conditions being met are slim and such minor increases in the routing table size can easily be handled by the core Internet routers of today.

***Is it too late to start the CIRs*:** This is another question often asked? Is it now too late for another entity or entities to enter the picture of IP address allocation? In the case of IPv4, it may be the case but we have not studied the IPv4 situation in detail. However, for IPv6, it is definitely not too late and in fact, we are still at infancy in the IPv6 world.

[30] More information on BGP routing table size and scalability issues can be found in Appendix F.

IPv6 is not a technology with a normal linear growth. The growth process was extremely slow in the beginning, as was seen from the year 2000 till 2006 ($y=x/n$ type growth). Then there was a year or two of linear increase ($y=x$ type growth). Now, in 2009 the growth has been exponential. The exponential growth ($y=x^n$) will keep increasing as the demand for IPv6 addresses will soon rise quickly. This is a very positive indication that IPv6 has been accepted globally and the real demand has started.

Diminishing the RIR Monopoly**:** The CIRs should not take over the exiting RIR structure in allocating the IPv6 addresses to requesting users. Rather, work together with the RIRs and offer an alternative for the 2nd tier IPv6 service providers (LIR/ISP) to source their blocks of IPv6 addresses. This will break the current monopolistic manner of obtaining IPv6 address blocks from only the RIRs. When such monopolies are broken, the following will automatically happen:

i) A choice for better service

ii) An Alternative for users and a possible appeal process

iii) Healthy competition

iv) Bring down the charges currently incurred by the applicants.

Implementing the CIRs*:* In summary, the need for implementing the CIR structure is a real need. The implementation details need to be planned well. The interesting fact is the implementation of the CIRs can be done systematically and gradually. The CIRs should be viewed as an addition to the existing IPv6 address allocation infrastructure and not replacing or creating a radical change in any form or policy to the current methods of IPv6 address allocations. A detailed CIR Implementation Plan must be created and reviewed first, and agreed upon by all involved parties before the actual implementation takes place.

Appendix A: IPv4 Address Architecture

IP Addresses

The Internet generally comprises of LANs and WANs as the elements where LANs comprise of hosts interconnected confined to a small geographical area. WANs encompass geographically dispersed hosts where LANs are interconnected by WANs. In the Internet Information is transported in terms of packets and these constituent elements provide the packet transport. The LANs and WANs are connected through routers or gateways for packet forwarding. The hosts, routers and gateways are uniquely identified by an IP address.

The global Internet address space offers hosts unique addresses within its defined space. For IPv4, there is 2^{32} = 4,294,967,296 possibilities and for IPv6 there are 2^{128} = 340,282,366,920,938,463,463,374,607,431,768,211,456 or 3.4*1038 possibilities of individual hosts.

In reality, blocks of addresses are allocated to organizations where these addresses allocated are globally reachable through ISPs. This reachability is ensured by routing protocols. Since each network is independent they may use different routing algorithm and since each network is independent of all the others, it is often referred to as an Autonomous System (AS).[31] Originally, an AS is controlled by a single entity, namely an ISP or a large organization. But according to RFC 1930, an AS is a collection of routing prefixes clearly defined by a single routing policy where the routing prefixes may be under the control of one or more organizations or network operators. As such, the Internet can be seen as an interconnected collection of subnetworks or ASs. To route within a site or a domain or an AS, Interior Gateway Protocol for e.g. OSPF (Open Shortest Path First) is used. Between ASs Exterior Gateway Protocol (EGP) now called as Border Gateway Protocol (BGP) for e.g. BGP-4 is used as the routing protocol. BGP is used as an abstraction of the global routing system.

[31] Andrew S. Tanenbaum, Computer Networks, Fourth Edition, PHI Pvt Ltd., 2006. pp. 427.

IPv4 Address architecture

IP addresses have generally two parts namely the network or prefix part that identifies the network and the host part that identifies the interface connecting the host as shown in Fig. A-1.

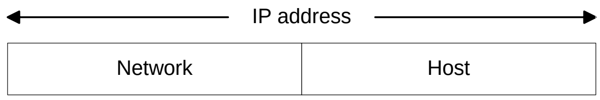

Figure A-1. General IP address format

IPv4 (Internet Protocol version 4) is the current version of the IP address that is predominantly used over the Internet now, uses a 32 bit binary address. Each address was organized as four 8-bit numbers separated by dots and called an octet. Each octet in the IP address is represented by a decimal number from `0` to `255` separated by dots. As such the entire range of IPv4 address goes from `0.0.0.0` to `255.255.255.255`.

Class based Addressing

IPv4 address architecture has gone through a number of iterations. The early Internet used the Class based addressing as illustrated in Fig. A-2, where the total address space was divided into five address classes which are no longer used now. Class A, B and C addresses supports unicast communication, Class D supports multicasting and Class E is reserved for experimentation. Class A, B and C had an 8/24, 16/16 and a 24/8 split respectively, where the first part identifies the network and second part indentifies the host or better a network interface. For e.g. with the legacy class C network address of 192.168.9.0, can be defined as 192.168.9.0/24 where the most significant 24 bits are ones and the least significant 8 bits are zeroes.

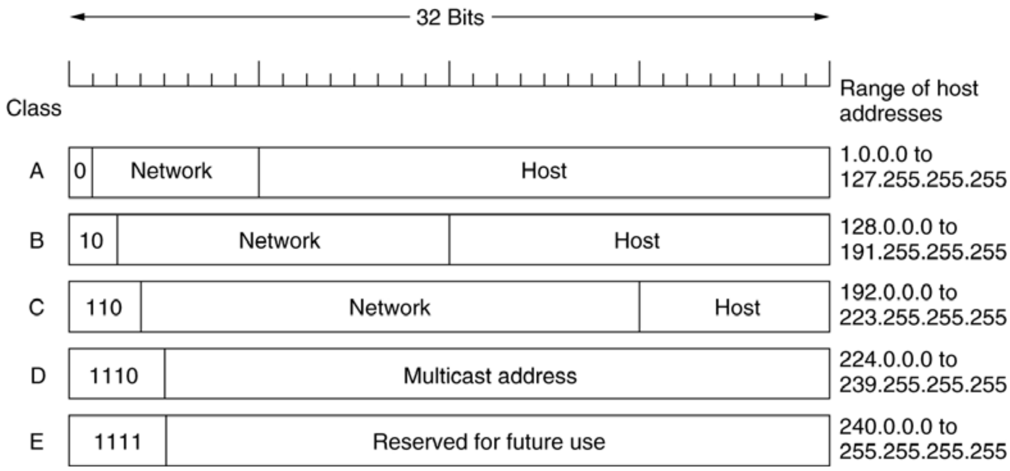

Figure A-2. IPv4 Class based address format[32]

Subnets

The class based addressing was extended by the concept of subnets[33] which allowed a network to be split into a number of subnetworks with unicast addresses for internal use but can be seen as a single network by the outside world. Subnets provided a multilevel hierarchical routing structure for the Internet allowing IPv4 addresses to be used efficiently reducing the number of network prefixes and routing complexity. The subnet and host addresses are derived from the original IP host address portion where the subnets are identified by a subnet mask. The subnet mask is a 32 bit number that identifies the network part and subnet part of the IP address while the remaining bits identify the host or the interface within the subnet as illustrated in Fig. A-3. In reality, the need for subnets is not homogeneous. It is not realistic to expect an organization or a network to be divided into uniform sized smaller organization or subnetworks. As such, using a fixed size mask would result in wasteful assignment of IP addresses within a subnet. The solution to this is to have a variable sized mask instead. Using a classless prefix with explicit prefix lengths would allow flexible matching of address space blocks according to actual need.

[32] Source: http://authors.phptr.com/tanenbaumcn4/

[33] RFC950

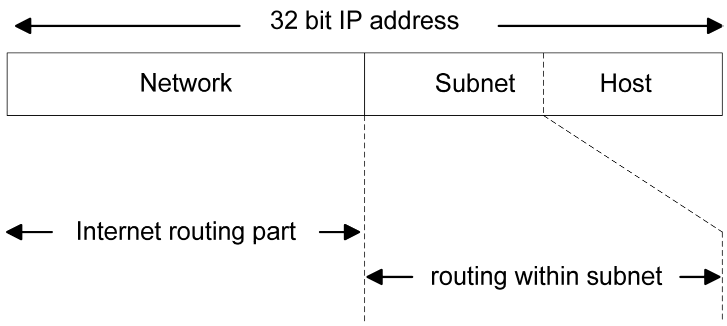

Figure A-3. Subnet addressing structure

Classless Inter-Domain Routing (CIDR)

In the early 1990s, owing to the exhaustion of Class-B addresses, routing scalability and eventual exhaustion of the 32-bit IPv4 address space led to the introduction of Classless Inter-Domain Routing[34] which used classless addresses. With classless addresses, the network part of an IP address can be of any size without being restricted by class boundaries. A classless IP address is represented in the form of `w.x.y.z / n`, where `n` denotes the number of bits in the network part of the address. Subnetting can be still done on the host part of the IP address. The scaling of routing is influenced by the way in which addresses were assigned and CIDR was deployed to improve routing scalability by improving route aggregation. CIDR is characterized by topologically significant address assignment, enables route aggregation, and performs packet forwarding based on longest prefix matching. To collapse routing information, the Internet is divided into addressing domains and as such all the Internet need not use network prefixes uniformly. Detailed information within a constituent network is available within a domain, but outside the domain only the common network prefix is advertised.

[34] More information on CIDR is available at RFC4632

Appendix B: Allocated IPv4 Address Space

IANA IPv4 Address Space Registry

Last Updated 2009-04-28

Description

The allocation of Internet Protocol version 4 (IPv4) address space to various registries is listed here. Originally, all the IPv4 address spaces were managed directly by the IANA. Later parts of the address space were allocated to various other registries to manage for particular purposes or regional areas of the world. RFC 1466 [RFC1466] documents most of these allocations.

Prefix	Designation	Date	Whois Status [1]	Note
000/8	IANA - Local Identification	1981-09		RESERVED [2]
001/8	IANA			UNALLOCATED
002/8	IANA			UNALLOCATED
003/8	General Electric Company	1994-05		LEGACY
004/8	Level 3 Communications, Inc.	1992-12		LEGACY
005/8	IANA			UNALLOCATED
006/8	Army Info Systems Center	1994-02		LEGACY
007/8	Administered by ARIN	1995-04	whois.arin.net	LEGACY
008/8	Level 3 Communications, Inc.	1992-12		LEGACY
009/8	IBM 1992-08			LEGACY
010/8	IANA - Private Use	1995-06		RESERVED [3]
011/8	DoD Intel Information Systems	1993-05		LEGACY
012/8	AT&T Bell Laboratories	1995-06		LEGACY
013/8	Xerox Corporation	1991-09		LEGACY
014/8	IANA			UNALLOCATED [4]
015/8	Hewlett-Packard Company	1994-07		LEGACY
016/8	Digital Equipment Corporation	1994-11		LEGACY
017/8	Apple Computer Inc.	1992-07		LEGACY
018/8	MIT	1994-01		LEGACY
019/8	Ford Motor Company	1995-05		LEGACY
020/8	Computer Sciences Corporation	1994-10		LEGACY

021/8	DDN-RVN	1991-07		LEGACY
022/8	Defense Information Systems Agency	1993-05		LEGACY
023/8	IANA			UNALLOCATED
024/8	ARIN	2001-05	whois.arin.net	ALLOCATED
025/8	UK Ministry of Defence 1995-01		whois.ripe.net	LEGACY
026/8	Defense Information Systems Agency	1995-05		LEGACY
027/8	IANA			UNALLOCATED
028/8	DSI-North	1992-07		LEGACY
029/8	Defense Information Systems Agency	1991-07		LEGACY
030/8	Defense Information Systems Agency	1991-07		LEGACY
031/8	IANA			UNALLOCATED
032/8	AT&T Global Network Services	1994-06		LEGACY
033/8	DLA Systems Automation Center	1991-01		LEGACY
034/8	Halliburton Company	1993-03		LEGACY
035/8	MERIT Computer Network	1994-04		LEGACY
036/8	IANA			UNALLOCATED
037/8	IANA			UNALLOCATED
038/8	Performance Systems International	1994-09		LEGACY
039/8	IANA			UNALLOCATED
040/8	Eli Lily & Company	1994-06		LEGACY
041/8	AfriNIC	2005-04	whois.afrinic.net	ALLOCATED
042/8	IANA			UNALLOCATED
043/8	Administered by APNIC	1991-01		LEGACY
044/8	Amateur Radio Digital Communications	1992-07		LEGACY
045/8	Interop Show Network	1995-01		LEGACY
046/8	IANA			UNALLOCATED
047/8	Bell-Northern Research	1991-01		LEGACY
048/8	Prudential Securities Inc.	1995-05		LEGACY
049/8	IANA			UNALLOCATED
050/8	IANA			UNALLOCATED
051/8	UK Government Department for Work and Pensions	1994-08	whois.ripe.net	LEGACY
052/8	E.I. duPont de Nemours and Co., Inc.	1991-12		LEGACY
053/8	Cap Debis CCS	1993-10		LEGACY
054/8	Merck and Co., Inc.	1992-03		LEGACY
055/8	DoD Network Information Center	1995-04		LEGACY
056/8	US Postal Service	1994-06		LEGACY

057/8	SITA	1995-05		LEGACY
058/8	APNIC	2004-04	whois.apnic.net	ALLOCATED
059/8	APNIC	2004-04	whois.apnic.net	ALLOCATED
060/8	APNIC	2003-04	whois.apnic.net	ALLOCATED
061/8	APNIC	1997-04	whois.apnic.net	ALLOCATED
062/8	RIPE NCC	1997-04	whois.ripe.net	ALLOCATED
063/8	ARIN	1997-04	whois.arin.net	ALLOCATED
064/8	ARIN	1999-07	whois.arin.net	ALLOCATED
065/8	ARIN	2000-07	whois.arin.net	ALLOCATED
066/8	ARIN	2000-07	whois.arin.net	ALLOCATED
067/8	ARIN	2001-05	whois.arin.net	ALLOCATED
068/8	ARIN	2001-06	whois.arin.net	ALLOCATED
069/8	ARIN	2002-08	whois.arin.net	ALLOCATED
070/8	ARIN	2004-01	whois.arin.net	ALLOCATED
071/8	ARIN	2004-08	whois.arin.net	ALLOCATED
072/8	ARIN	2004-08	whois.arin.net	ALLOCATED
073/8	ARIN	2005-03	whois.arin.net	ALLOCATED
074/8	ARIN	2005-06	whois.arin.net	ALLOCATED
075/8	ARIN	2005-06	whois.arin.net	ALLOCATED
076/8	ARIN	2005-06	whois.arin.net	ALLOCATED
077/8	RIPE NCC	2006-08	whois.ripe.net	ALLOCATED
078/8	RIPE NCC	2006-08	whois.ripe.net	ALLOCATED
079/8	RIPE NCC	2006-08	whois.ripe.net	ALLOCATED
080/8	RIPE NCC	2001-04	whois.ripe.net	ALLOCATED
081/8	RIPE NCC	2001-04	whois.ripe.net	ALLOCATED
082/8	RIPE NCC	2002-11	whois.ripe.net	ALLOCATED
083/8	RIPE NCC	2003-11	whois.ripe.net	ALLOCATED
084/8	RIPE NCC	2003-11	whois.ripe.net	ALLOCATED
085/8	RIPE NCC	2004-04	whois.ripe.net	ALLOCATED
086/8	RIPE NCC	2004-04	whois.ripe.net	ALLOCATED
087/8	RIPE NCC	2004-04	whois.ripe.net	ALLOCATED
088/8	RIPE NCC	2004-04	whois.ripe.net	ALLOCATED
089/8	RIPE NCC	2005-06	whois.ripe.net	ALLOCATED
090/8	RIPE NCC	2005-06	whois.ripe.net	ALLOCATED
091/8	RIPE NCC	2005-06	whois.ripe.net	ALLOCATED
092/8	RIPE NCC	2007-03	whois.ripe.net	ALLOCATED
093/8	RIPE NCC	2007-03	whois.ripe.net	ALLOCATED
094/8	RIPE NCC	2007-07	whois.ripe.net	ALLOCATED

095/8	RIPE NCC	2007-07	whois.ripe.net	ALLOCATED	
096/8	ARIN	2006-10	whois.arin.net	ALLOCATED	
097/8	ARIN	2006-10	whois.arin.net	ALLOCATED	
098/8	ARIN	2006-10	whois.arin.net	ALLOCATED	
099/8	ARIN	2006-10	whois.arin.net	ALLOCATED	
100/8	IANA		UNALLOCATED		
101/8	IANA		UNALLOCATED		
102/8	IANA		UNALLOCATED		
103/8	IANA		UNALLOCATED		
104/8	IANA		UNALLOCATED		
105/8	IANA		UNALLOCATED		
106/8	IANA		UNALLOCATED		
107/8	IANA		UNALLOCATED		
108/8	ARIN	2008-12	whois.arin.net	ALLOCATED	
109/8	RIPE NCC	2009-01	whois.ripe.net	ALLOCATED	
110/8	APNIC	2008-11	whois.apnic.net	ALLOCATED	
111/8	APNIC	2008-11	whois.apnic.net	ALLOCATED	
112/8	APNIC	2008-05	whois.apnic.net	ALLOCATED	
113/8	APNIC	2008-05	whois.apnic.net	ALLOCATED	
114/8	APNIC	2007-10	whois.apnic.net	ALLOCATED	
115/8	APNIC	2007-10	whois.apnic.net	ALLOCATED	
116/8	APNIC	2007-01	whois.apnic.net	ALLOCATED	
117/8	APNIC	2007-01	whois.apnic.net	ALLOCATED	
118/8	APNIC	2007-01	whois.apnic.net	ALLOCATED	
119/8	APNIC	2007-01	whois.apnic.net	ALLOCATED	
120/8	APNIC	2007-01	whois.apnic.net	ALLOCATED	
121/8	APNIC	2006-01	whois.apnic.net	ALLOCATED	
122/8	APNIC	2006-01	whois.apnic.net	ALLOCATED	
123/8	APNIC	2006-01	whois.apnic.net	ALLOCATED	
124/8	APNIC	2005-01	whois.apnic.net	ALLOCATED	
125/8	APNIC	2005-01	whois.apnic.net	ALLOCATED	
126/8	APNIC	2005-01	whois.apnic.net	ALLOCATED	
127/8	IANA - Loopback	1981-09		RESERVED	[5]
128/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY	
129/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY	
130/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY	
131/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY	
132/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY	

133/8	Administered by APNIC	1997-03	whois.apnic.net	LEGACY
134/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
135/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
136/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
137/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
138/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
139/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
140/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
141/8	Administered by RIPE NCC	1993-05	whois.ripe.net	LEGACY
142/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
143/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
144/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
145/8	Administered by RIPE NCC	1993-05	whois.ripe.net	LEGACY
146/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
147/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
148/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
149/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
150/8	Administered by APNIC	1993-05	whois.apnic.net	LEGACY
151/8	Administered by RIPE NCC	1993-05	whois.ripe.net	LEGACY
152/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
153/8	Administered by APNIC	1993-05	whois.apnic.net	LEGACY
154/8	Administered by AfriNIC	1993-05	whois.afrinic.net	LEGACY
155/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
156/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
157/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
158/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
159/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
160/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
161/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
163/8	Administered by APNIC	1993-05	whois.apnic.net	LEGACY
164/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
165/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
166/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
167/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
168/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
169/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY [6]
170/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY
171/8	Administered by APNIC	1993-05	whois.apnic.net	LEGACY

172/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY [7]
173/8	ARIN	2008-02	whois.arin.net	ALLOCATED
174/8	ARIN	2008-02	whois.arin.net	ALLOCATED
175/8	IANA		UNALLOCATED	
176/8	IANA		UNALLOCATED	
177/8	IANA		UNALLOCATED	
178/8	RIPE NCC	2009-01	whois.ripe.net	ALLOCATED
179/8	IANA		UNALLOCATED	
180/8	APNIC	2009-04	whois.apnic.net	ALLOCATED
181/8	IANA		UNALLOCATED	
182/8	IANA		UNALLOCATED	
183/8	APNIC	2009-04	whois.apnic.net	ALLOCATED
184/8	ARIN	2008-12	whois.arin.net	ALLOCATED
185/8	IANA		UNALLOCATED	
186/8	LACNIC	2007-09	whois.lacnic.net	ALLOCATED
187/8	LACNIC	2007-09	whois.lacnic.net	ALLOCATED
188/8	Administered by RIPE NCC	1993-05	whois.ripe.net	LEGACY
189/8	LACNIC	1995-06	whois.lacnic.net	ALLOCATED
190/8	LACNIC	1995-06	whois.lacnic.net	ALLOCATED
191/8	Administered by LACNIC	1993-05	whois.lacnic.net	LEGACY
192/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY [8]
193/8	RIPE NCC	1993-05	whois.ripe.net	ALLOCATED
194/8	RIPE NCC	1993-05	whois.ripe.net	ALLOCATED
195/8	RIPE NCC	1993-05	whois.ripe.net	ALLOCATED
196/8	Administered by AfriNIC	1993-05	whois.afrinic.net	LEGACY
197/8	AfriNIC	2008-10	whois.afrinic.net	ALLOCATED
198/8	Administered by ARIN	1993-05	whois.arin.net	LEGACY [9]
199/8	ARIN	1993-05	whois.arin.net	ALLOCATED
200/8	LACNIC	2002-11	whois.lacnic.net	ALLOCATED
201/8	LACNIC	2003-04	whois.lacnic.net	ALLOCATED
202/8	APNIC	1993-05	whois.apnic.net	ALLOCATED
203/8	APNIC	1993-05	whois.apnic.net	ALLOCATED
204/8	ARIN	1994-03	whois.arin.net	ALLOCATED
205/8	ARIN	1994-03	whois.arin.net	ALLOCATED
206/8	ARIN	1995-04	whois.arin.net	ALLOCATED
207/8	ARIN	1995-11	whois.arin.net	ALLOCATED
208/8	ARIN	1996-04	whois.arin.net	ALLOCATED
209/8	ARIN	1996-06	whois.arin.net	ALLOCATED

210/8	APNIC	1996-06	whois.apnic.net	ALLOCATED
211/8	APNIC	1996-06	whois.apnic.net	ALLOCATED
212/8	RIPE NCC	1997-10	whois.ripe.net	ALLOCATED
213/8	RIPE NCC	1993-10	whois.ripe.net	ALLOCATED
214/8	US-DOD	1998-03	LEGACY	
215/8	US-DOD	1998-03	LEGACY	
216/8	ARIN	1998-04	whois.arin.net	ALLOCATED
217/8	RIPE NCC	2000-06	whois.ripe.net	ALLOCATED
218/8	APNIC	2000-12	whois.apnic.net	ALLOCATED
219/8	APNIC	2001-09	whois.apnic.net	ALLOCATED
220/8	APNIC	2001-12	whois.apnic.net	ALLOCATED
221/8	APNIC	2002-07	whois.apnic.net	ALLOCATED
222/8	APNIC	2003-02	whois.apnic.net	ALLOCATED
223/8	IANA		UNALLOCATED	
224/8	Multicast	1981-09	RESERVED [10]	
225/8	Multicast	1981-09	RESERVED [10]	
226/8	Multicast	1981-09	RESERVED [10]	
227/8	Multicast	1981-09	RESERVED [10]	
228/8	Multicast	1981-09	RESERVED [10]	
229/8	Multicast	1981-09	RESERVED [10]	
230/8	Multicast	1981-09	RESERVED [10]	
231/8	Multicast	1981-09	RESERVED [10]	
232/8	Multicast	1981-09	RESERVED [10]	
233/8	Multicast	1981-09	RESERVED [10]	
234/8	Multicast	1981-09	RESERVED [10]	
235/8	Multicast	1981-09	RESERVED [10]	
236/8	Multicast	1981-09	RESERVED [10]	
237/8	Multicast	1981-09	RESERVED [10]	
238/8	Multicast	1981-09	RESERVED [10]	
239/8	Multicast	1981-09	RESERVED [10]	
240/8	Future use	1981-09	RESERVED [11]	
241/8	Future use	1981-09	RESERVED [11]	
242/8	Future use	1981-09	RESERVED [11]	
243/8	Future use	1981-09	RESERVED [11]	
244/8	Future use	1981-09	RESERVED [11]	
245/8	Future use	1981-09	RESERVED [11]	
246/8	Future use	1981-09	RESERVED [11]	
247/8	Future use	1981-09	RESERVED [11]	

248/8	Future use	1981-09	RESERVED	[11]
249/8	Future use	1981-09	RESERVED	[11]
250/8	Future use	1981-09	RESERVED	[11]
251/8	Future use	1981-09	RESERVED	[11]
252/8	Future use	1981-09	RESERVED	[11]
253/8	Future use	1981-09	RESERVED	[11]
254/8	Future use	1981-09	RESERVED	[11]
255/8	Future use	1981-09	RESERVED	[11]

[1] Indicates the status of address blocks as follows:
RESERVED: designated by the IETF for specific non-unicast purposes as noted.
LEGACY: allocated by the central Internet Registry (IR) prior to the Regional Internet Registries (RIRs). This address space is now administered by individual RIRs as noted, including maintenance of WHOIS Directory and reverse DNS records. Assignments from these blocks are distributed globally on a regional basis.
ALLOCATED: delegated entirely to specific RIR as indicated.
UNALLOCATED: not yet allocated or reserved.

[2] 0.0.0.0/8 reserved for self-identification [RFC3330]

[3] Reserved for Private-Use Networks [RFC1918]

[4] This was reserved for Public Data Networks [RFC1356]
See: http://www.iana.org/assignments/public-data-network-numbers
It was recovered in February 2008.

[5] 127.0.0.0/8 is reserved for Loopback [RFC3330]

[6] 169.254.0.0/16 reserved for Link Local [RFC3330]

[7] 172.16.0.0/12 reserved for Private-Use Networks [RFC1918]

[8] 192.0.2.0/24 reserved for Test-Net [RFC3330]
192.88.99.0/24 reserved for 6to4 Relay Anycast [RFC3068]
192.168.0.0/16 reserved for Private-Use Networks [RFC1918]

[9] 198.18.0.0/15 reserved for Network Interconnect Device

Benchmark Testing [RFC3330]

[10] Multicast (formerly "Class D") [RFC1700]
See: http://www.iana.org/assignments/multicast-addresses

[11] Reserved for future use (formerly "Class E") [RFC1700

Appendix C: IPv6 Address Space and Global Unicast Address Allocations

Internet Protocol Version 6 Address Space

(last updated 2008-05-13)

IPv6 Prefix	Allocation	Reference	Note
0000::/8	Reserved by IETF	[RFC4291]	[1] [5]
0100::/8	Reserved by IETF	[RFC4291]	
0200::/7	Reserved by IETF	[RFC4048]	[2]
0400::/6	Reserved by IETF	[RFC4291]	
0800::/5	Reserved by IETF	[RFC4291]	
1000::/4	Reserved by IETF	[RFC4291]	
2000::/3	Global Unicast	[RFC4291]	[3]
4000::/3	Reserved by IETF	[RFC4291]	
6000::/3	Reserved by IETF	[RFC4291]	
8000::/3	Reserved by IETF	[RFC4291]	
A000::/3	Reserved by IETF	[RFC4291]	
C000::/3	Reserved by IETF	[RFC4291]	
E000::/4	Reserved by IETF	[RFC4291]	
F000::/5	Reserved by IETF	[RFC4291]	
F800::/6	Reserved by IETF	[RFC4291]	
FC00::/7	Unique Local Unicast	[RFC4193]	
FE00::/9	Reserved by IETF	[RFC4291]	
FE80::/10	Link Local Unicast	[RFC4291]	
FEC0::/10	Reserved by IETF	[RFC3879]	[4]
FF00::/8	Multicast	[RFC4291]	

Notes:

[0] The IPv6 address management function was formally delegated to IANA in December 1995 [RFC1881].

[1] The "unspecified address", the "loopback address", and the IPv6 Addresses with Embedded IPv4 Addresses are assigned out of the 0000::/8 address block.

[2] 0200::/7 was previously defined as an OSI NSAP-mapped prefix set [RFC4548]. This definition has been deprecated as of December 2004 [RFC4048].

[3] The IPv6 Unicast space encompasses the entire IPv6 address range with the exception of FF00::/8. [RFC4291] IANA unicast address assignments are currently limited to the IPv6 unicast address range of 2000::/3. IANA assignments from this block are registered in the IANA registry: iana-ipv6-unicast-address-assignments.

[4] FEC0::/10 was previously defined as a Site-Local scoped address prefix. This definition has been deprecated as of September 2004 [RFC3879].

[5] 0000::/96 was previously defined as the "IPv4-compatible IPv6 address" prefix. This definition has been deprecated by [RFC4291].

References

[RFC1881]   The IAB and IESG, "IPv6 Address Allocation Management", RFC 1881, December 1995.

[RFC1888]   J. Bound et al, "OSI NSAPs and IPv6", RFC 1888, August 1996.

[RFC3879]   C. Huitema and B. Carpenter, "Deprecating Site Local Addresses", RFC 3879, September 2004.

[RFC4048]   B. Carpenter, "RFC 1888 is obsolete", RFC 4048, April 2005.

[RFC4147]   G. Huston, "Proposed changes to the format of the IANA IPv6 Registry", RFC 4147, August 2005.

[RFC4193]   R. Hinden and B. Haberman, "Unique Local IPv6 Unicast Addresses" RFC 4193, October 2005.

[RFC4291]   R. Hinden, Nokia, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.

[RFC4548]   E. Gray, J. Rutemiller and G. Swallow, "Internet Code Point Assignments for NSAP Addresses", RFC 4548, May 2006.

IPv6 Global Unicast Address Assignments

[last updated 2008-05-13]

Global Unicast Prefix	Assignment	Date	Note
2001:0000::/23	IANA	01 Jul 99	[1]
2001:0200::/23	APNIC	01 Jul 99	
2001:0400::/23	ARIN	01 Jul 99	
2001:0600::/23	RIPE NCC	01 Jul 99	
2001:0800::/23	RIPE NCC	01 May 02	
2001:0A00::/23	RIPE NCC	02 Nov 02	
2001:0C00::/23	APNIC	01 May 02	[2]
2001:0E00::/23	APNIC	01 Jan 03	
2001:1200::/23	LACNIC	01 Nov 02	
2001:1400::/23	RIPE NCC	01 Feb 03	
2001:1600::/23	RIPE NCC	01 Jul 03	
2001:1800::/23	ARIN	01 Apr 03	
2001:1A00::/23	RIPE NCC	01 Jan 04	
2001:1C00::/22	RIPE NCC	01 May 04	
2001:2000::/20	RIPE NCC	01 May 04	
2001:3000::/21	RIPE NCC	01 May 04	
2001:3800::/22	RIPE NCC	01 May 04	
2001:3C00::/22	RESERVED	11 Jun 04	[3]
2001:4000::/23	RIPE NCC	11 Jun 04	
2001:4200::/23	AfriNIC	01 Jun 04	
2001:4400::/23	APNIC	11 Jun 04	
2001:4600::/23	RIPE NCC	17 Aug 04	
2001:4800::/23	ARIN	24 Aug 04	
2001:4A00::/23	RIPE NCC	15 Oct 04	
2001:4C00::/23	RIPE NCC	17 Dec 04	
2001:5000::/20	RIPE NCC	10 Sep 04	
2001:8000::/19	APNIC	30 Nov 04	
2001:A000::/20	APNIC	30 Nov 04	
2001:B000::/20	APNIC	08 Mar 06	
2002:0000::/16	6to4	01 Feb 01	[4]
2003:0000::/18	RIPE NCC	12 Jan 05	
2400:0000::/12	APNIC	03 Oct 06	[7]
2600:0000::/12	ARIN	03 Oct 06	[8]
2610:0000::/23	ARIN	17 Nov 05	
2620:0000::/23	ARIN	12 Sep 06	
2800:0000::/12	LACNIC	03 Oct 06	[6]
2A00:0000::/12	RIPE NCC	03 Oct 06	[5]
2C00:0000::/12	AfriNIC	03 Oct 06	

Notes:

[0] The assignable Global Unicast Address space is defined in [RFC4291] as being the address block defined by the prefix 2000::/3. All address space in this block not listed in the table above is reserved by IANA for future allocation.

[1] IANA Special Purpose Address Block [RFC4773].
See: http://www.iana.org/assignments/iana-ipv6-special-registry

[2] 2001:0DB8::/32 has been assigned as a NON-ROUTABLE

range to be used for documentation purpose [RFC3849].

[3] 2001:3C00::/22 is reserved for possible future allocation to the RIPE NCC.

[4] 2002::/16 is reserved for use in 6to4 deployments [RFC3056].

[5] 2A00:0000::/21 was originally allocated on 19 Apr 05. 2A01:0000::/23 was allocated on 14 Jul 05. 2A01:0000::/16 (incorporating the 2A01:0000::/23) was allocated 15 Dec 2005. The more recent allocation (03 Oct 2006) incorporates these previous allocations.

[6] 2800:0000::/23 was allocated on 17 Nov 05. The more recent allocation (03 Oct 06) incorporates the previous allocation.

[7] 2400:0000::/19 was allocated on 20 May 05. 2400:2000::/19 was Allocated on 08 Jul 05. 2400:4000::/21 was allocated on 08 Aug 05. 2404:0000::/23 was allocated on 19 Jan 06. The more recent allocation (03 October 06)incorporates all these previous allocations.

[8] 2600:0000::/22, 2604:0000::/22, 2608:0000::/22 and 260C:0000::/22 Were allocated on 19 Apr 05. The more recent allocation (03 Oct 06) incorporates all these previous allocations.

References

[RFC2471] Hinden, R., R. Fink, J. Postel, "IPv6 Testing Address Allocation", RFC 2471, December 1998.

[RFC2928] Hinden, R., Deering, S., Fink, R., Hain, T., , "Initial IPv6 Sub-TLA ID Assignments", RFC 2928, September 2000.

[RFC3056] Carpenter, B., K. Moore, "Connection of IPv6 Domains via IPv4 Clouds without Explicit Tunnels", RFC 3056, February 2001.

[RFC4291] Hinden, R., "IP Version 6 Addressing Architecture", RFC 4291, February 2006.

[RFC3849] Huston, G., A. Lord, P. Smith, "IPv6 Address Prefix Reserved for Documentation", RFC 3849, July 2004.

[RFC4147] Huston, G., "Proposed changes to the format of the IANA IPv6 Registry", RFC 4147, August 2005.

[RFC4380] C. Huitema, "Teredo: Tunneling IPv6 over UDP through NATs", RFC 4380, February 2006.

[RFC4773] G. Huston, "Administration of the IANA Special Purpose Address Block", RFC 4773, December 2006.

Appendix D: Rationale for the size of address space to be allocated or assigned to end sites in the CIR model

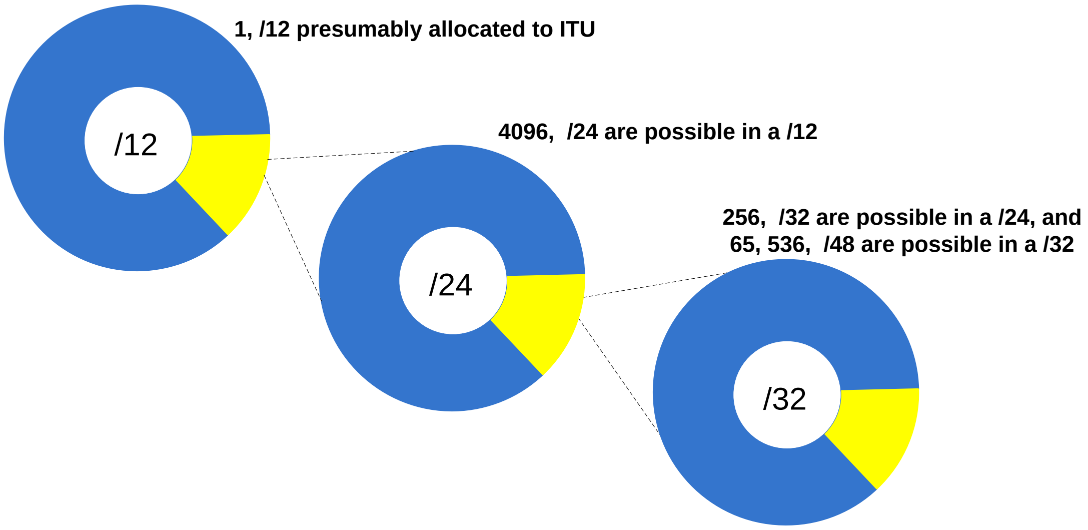

Figure D-1. Proposed IPv6 Global unicast address allocation to national registries by ITU

Table D-1. IPv6 Address delegation recommendations

Prefix length	Number of allocations within a /32	Number of Subnets in an allocation	Recommendations
/48	65,536	65,536	For very large organizations. > 1000 subnets
/52	1,048,576	4,096	For medium size organizations. < 1000 subnets
/56	16,777,216	256	For small organizations. < 100 subnets
/60	268,435,456	16	Home networks/users
/64	4,294,967,296	1	When one and only one subnet is needed
/128	Huge!	0	When one and only device is connecting

All end sites may not get recursively subnetted such as a home PC or user and a telephone in a cellular network. The future possibilities are that home networks, mobile cellular networks and organizations of any size would eventually have multiple always-on hosts with at least one subnet and possibilities for additional subnets therein need for internal routing capabilities. Owing to this the subscriber allocation units are viewed as a site and not always as a host.

By judiciously applying GAP algorithms, ITU can allocate the single /12 address space allocated to it to each of the requesting member nation in the world with many /32 prefix allocations. If we equally partition a single /12 address space with a /32 prefix among presumably 256 possibilities[35], each country would get about 4096 numbers of /32 prefix allocations. As such in total, a nation can get 268,435,456 numbers of /48 prefixes (65,536 x 4096). The total number of /32 that can be allocated by ITU from its current single /12 address space would represent about 2% of the whole allocatable IPv6 address space currently used by IANA to the RIRs, and 0.024% of the whole IPv6 address space.

Taking into consideration the population of a country, the huge number of devices that may get an IP address due to the ubiquitous nature of the Internet and IP explosion, 268,435,456 number of /48 prefixes would be sufficient even for the largest member nation of ITU. This observation is made based on the proposed IPv6 address delegation recommendations given in Table D-1. Also, not to be forgotten that this proposed model is going to be an alternative source for the requesting countries to get IPv6 addresses apart from the existing RIR model. So this number would be sufficient.

In case it so happens the 268,435,456 number of /48 prefixes allocated by ITU to a requesting nation is not sufficient, additional IPv6 address space could be sought by the CIRs from ITU based on the additional allocation policies defined by ITU to CIRs. If need be, ITU can request with IANA for additional IPv6 address space based on the additional allocation policies defined by IANA to the RIRs.

[35] Though as on date there are only 192 member states recognized by UN, taking bit boundary $2^8 = 256$ possibilities.

IANA follows sparse algorithm in allocating IPv6 address space to the RIRs so that subsequent allocations are contiguous, avoids fragmentation and maximizes aggregation of address space and the NIR or the LIR/ISP can retain a single prefix as they grow. IANA would continue delegating IPv6 address space to the existing 5 RIRs, and ITU where a user can have a choice of taking the IPv6 address from the relevant RIR or ITU through its CIR. This necessitates the establishment of CIRs in countries that wishes to follow this alternate model of obtaining IPv6 address from the ITU. In that case to conserve IPv6 address space and route aggregation, the CIR allocation would have to follow certain IP address allocation and assignment policies which would include those defined jointly by ARIN, RIPE and APNIC.

Appendix E: Dual-homing effects to routing table size and scalability

If a user has sourced the IPv6 address from both the RIR and ITU, then that user would have at least two root prefixes instead of one which would add-up to the internet core routing table. A similar situation could happen where a user can end up with two or more prefixes even when it sources its IPv6 address from the same RIR alone as the supplier. The RIRs do not guarantee contiguous allocation of IP addresses even though they advice addressing policies should seek to avoid fragmentation of address ranges.[36]

Thus the number of routing prefixes that could add-up to the routing table due to the above situation is trivial. Moreover, the capability of the present day routers[37] and Moore's law proves that the above increase in routing table size in the core routers is not a concern at the present. *As such the above issue would not impact or threaten the global Internet stability and routability.*

The increase in more specific prefixes or prefix de-aggregation due to traffic engineering and user site multi-homing will exist with both the existing RIR and proposed CIR, IPv6 address allocation models. Solutions to this end are defined in the direction of location identifier split methods that includes, Mike O Dells GSE proposal, 8+8 addressing architecture, LISP (Locator/Identifier Separation Protocol, ILNP (Identifier, locator network protocol), HIP (Host Identity Protocol), FARA (Forwarding directive, Association and Rendezvous Architecture, and ISLAY – A new routing & addressing architecture. All these models tend towards defining a new addressing architecture with changes to the existing Internet architecture itself.

---

[36] Section 3.4 Aggregation, IPv6 Address Allocation and Assignment Policy. Source: http://www.apnic.net

[37] Cisco CRS-1: Throughput, 1.2 Tbps. 4GB DRAM for system process and routing tables. Can handle millions of routes.

Juniper, T1600: Throughput, 1.6Tbps, TX Matrix with 4 * T640 – Throughput, 2.5Tbps and TX Matrix Plus with 16 T1600 – Throughput, 25.6 Tbps. 2 – 4GB DRAM for system process and routing tables. Can handle millions of routes.

Appendix F: BGP routing table size and scalability issues.

Routing

To transfer messages from a source to a destination, apart from the address or location of the destination, information on how to get there should be known. This knowledge is known as routing. The main purpose of addressing is routing as the address is important for determining the route to get to the destination. Hierarchical addressing is preferred as routing can be made hierarchical. This would facilitate to have optimal routing knowledge within each router to reach the destination.

Core Router BGP Routing table size

The current size and growth rate of the core router BGP routing table size has importance to both network operators and researchers. It also highlights the impact of customer needs such as multi-homing, Traffic Engineering and mobility on the limited resources available inside the core routers. The ISPs feed the routers with prefixes and maintain them. As such the router table size is of importance due to its increased growth and for an ISP, the routers memory capacity, processing speed, routing latency is of concern. They would like to estimate how big a router they need, what would be the expected size of their Forwarding Information Base (FIB) and Routing Information Base (RIB), and the memory processing capability such as Ternary Content Addressable Memory (TCAM).

Interestingly, our recent discussions with a CISCO Internet routing expert indicated that all the above is no longer a concern as their current Internet class core routers can handle large address tables at very fast look-up speeds. Thus the technical constraints for Internet routing table growth has already been overcome before it has even happened.

In Oct 2006, the workshop by the Internet Architecture Board studied on routing and addressing [RFC4984]. The workshop participants observed that *routing scalability* is the most important problem facing the Internet today.

IPv4 BGP Table

The Table F-1 shows the BGP table size for IPv4 and Fig. F-1 shows the active BGP entries for IPv4 since 1989. All tables and figures under this section are sourced form http://www.potaroo.net

Table F-1. BGP Table size for IPv4

Date	Prefixes	CIDR Aggregated
05-06-09	293617	183706
06-06-09	294620	183567
07-06-09	294066	182882
08-06-09	294012	183175
09-06-09	294091	183252
10-06-09	293965	183136
11-06-09	292531	183769
12-06-09	294158	183566

The Fig. F-1 below shows that there is stable increase in the active BGP (FIB) entries. From this figure it can be observed that the routing table growth rate is greater than the rate at which IP addresses are allocated. It can also be observed that the use of CIDR helps to aggregate larger prefix into smaller and this demands the addressing structure to be hierarchical. It also can be seen that along with the prefix count the AS count also is increasing. As the number of AS is increasing, the interconnection degree of AS would also increase. The growth is not at the edge of the network and the network is not growing any larger in terms of average AS path change. The growth is actually happening by increasing the density of the network by attaching new networks.

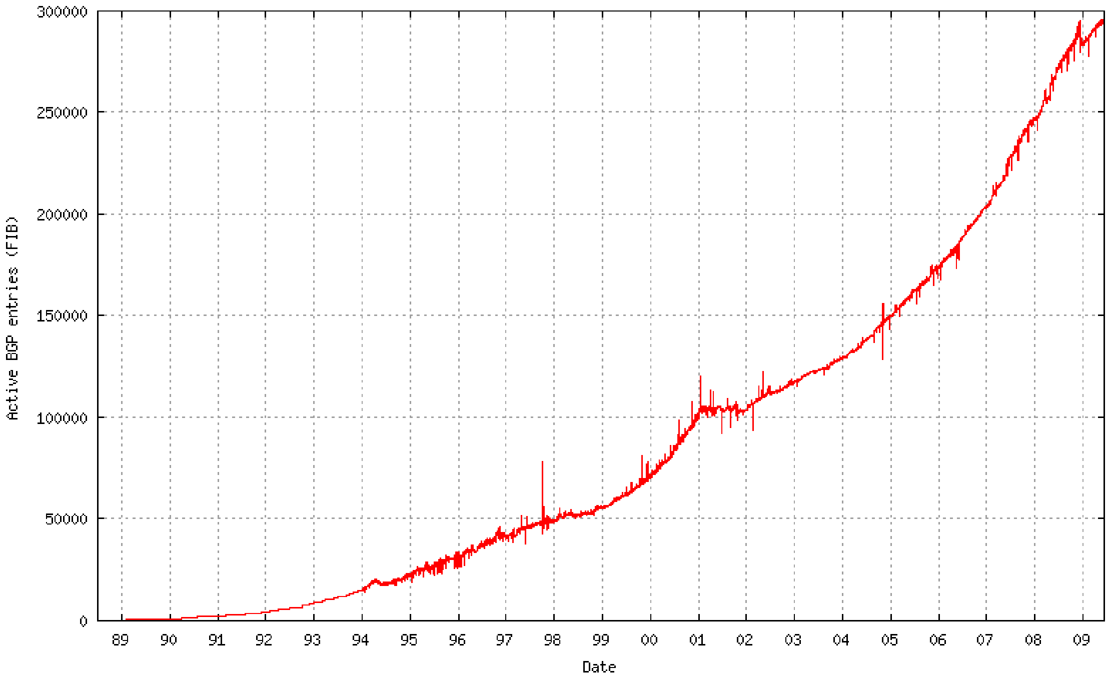

Figure F-1. Active BGP (FIB) entries for IPv4

Table F-2. Summary of the IPv4 BGP network[38]

	Dec 2008	Jun 2009
Prefix Count	286000	294180
Root Prefixes	133000	142846
More Specifics	152000	151334
Address Spans (/8)	118.44	122
AS Count	30200	3155
Transit AS Count	4100	105
Stub AS Count	26200	3050

IPv6 BGP Table

The Table F-3 and Fig. F-2 show the BGP table size and active BGP (FIB) entries for IPv6, while Fig. F-3 shows allocations of IPv6 address blocks for the different registries. From Fig. F-2 and Fig. F-3 it can be observed that previous to 2007 the growth has been mostly linear, while the current growth is exponential. IPv6 uptake is very recent and

[38] BGP data from AS65000, last updated 12 Jun 2009. Source: http://bgp.potaroo.net/as2.0/bgp-active.html

with the depletion of IPv4 nearing and as IPv6 opens up new opportunities the growth would be highly exponential.

Table F-3. BGP Table size for IPv6

Date	Prefixes	CIDR Aggregated
05-06-09	1915	1748
06-06-09	1917	1757
07-06-09	1919	1757
08-06-09	1922	1753
09-06-09	1924	1755
10-06-09	1921	1762
11-06-09	1883	1761
12-06-09	1928	1765

Table F-4. Summary of the IPv6 BGP network[39]

	Dec 2008	Jun 2009
Prefix Count	1600	1916
Root Prefixes	1300	1554
More Specifics	300	362
Address Spans	16.65	27.34
AS Count	1230	1456
Transit AS Count	310	51
Stub AS Count	920	1405

[39] BGP data from AS6447, last updated 12 Jun 2009. Source: http://bgp.potaroo.net/v6/as6447/

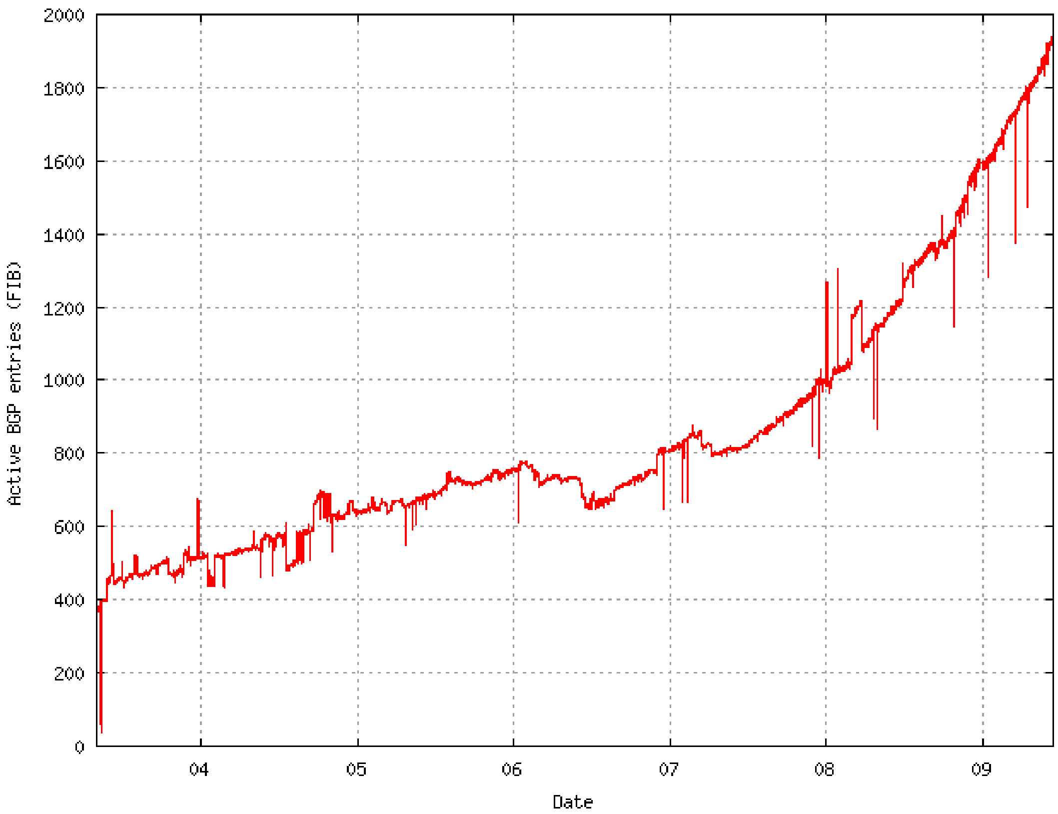

Figure F-2. Active BGP (FIB) entries for IPv6

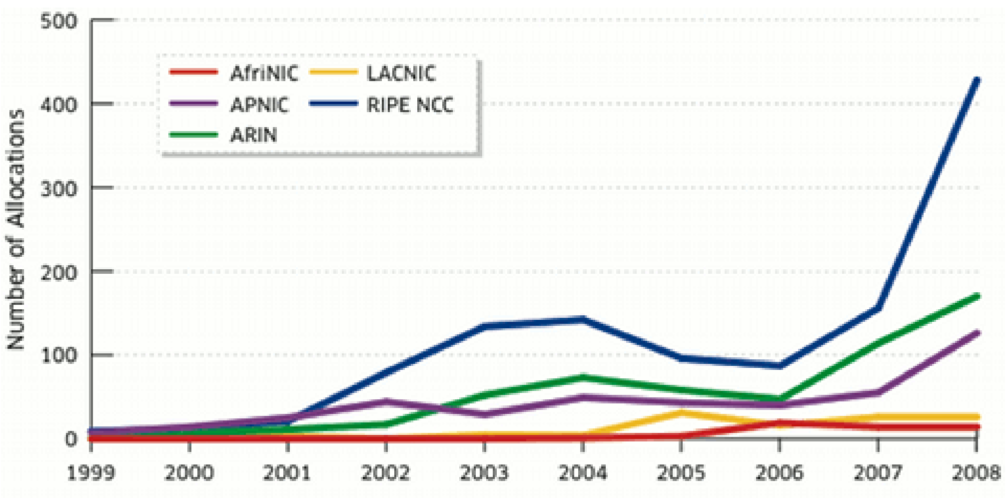

Figure F-3. Allocation of IPv6 address blocks to RIRs

General observation on the BGP routing table

In routing table entry, each prefix allocated to an ISP gets recorded into the routing table. End user assignments do not find an entry into the routing table unless they are made specifics by the ISP. As an ideal case, at the public topology there should be only 2^{32} = 4,294,967,296 prefixes for a single /32 prefix allocation. But in reality it would not be the case and the ISP would add more specific prefixes to which it wishes to have reachability globally. If we look at Table F-2 comprising a summary of the IPv4 BGP network, it can be observed that the more specific prefixes are nearly half of the total prefix count. The root prefixes are the one's actually allocated by the ISPs to organizations from its address space allocated by the RIRs. Similar is the situation with IPv6 but, observing Table F-4 summarizing the IPv6 BGP network it can be noted that the more specific prefixes are less than 1/4th of the actual root prefixes. This shows that route aggregation is more effective leading to less fragmentation with IPv6 than IPv4 as the size of the IPv6 address is large, so a much larger block can be allocated to a user. In our observation, as a comparison, the size of the IPv6 BGP routing table growth rate in terms of fragmentation would be less when compared with IPv4. As the size of the IPv6 address (128 bits) is bigger than IPv4 (32 bits), the minimum size of the prefix allocated to the ISP would be large. Eventually these route prefixes added to the more specific prefixes would result in lengthier routing tables.

The stable growth of the Default Free Zone (DFZ) or core router table size persists for long. With the deployment of IPv6 because of its large address space, it is expected that it would make this situation worst. More IPv6 prefix allocated would eventually end-up with more routing table entries in the DFZ. To this end, specific prefixes would be added by the ISPs to meet the user requirements or theirs in terms of site multi-homing, traffic engineering and mobility.

As each coin has two sides, the above scenario of the increasing DFZ router table size can be seen in a different perspective. *If it is assumed that the unit cost of routing* *continues to decline due to technological development on routers or at least the cost remains constant then this situation is not much of concern*.

Another assumption could be that applying Moore's Law. If the routing table growth parameters is within the limitation of Moore's Law then the unit cost of routing and switching hardware is no longer an issue as Moore's law indicates routing power will far exceed the routing table growth.

The Fig. F-4 below compares the quadratic projection model of the size of the DFZ routers with an exponential model defined by Moore's Law. It can be seen from the above figure that there is no real cause for alarm at this stage and the BGP table size appears to fit well within the limits of these parameters.[40]

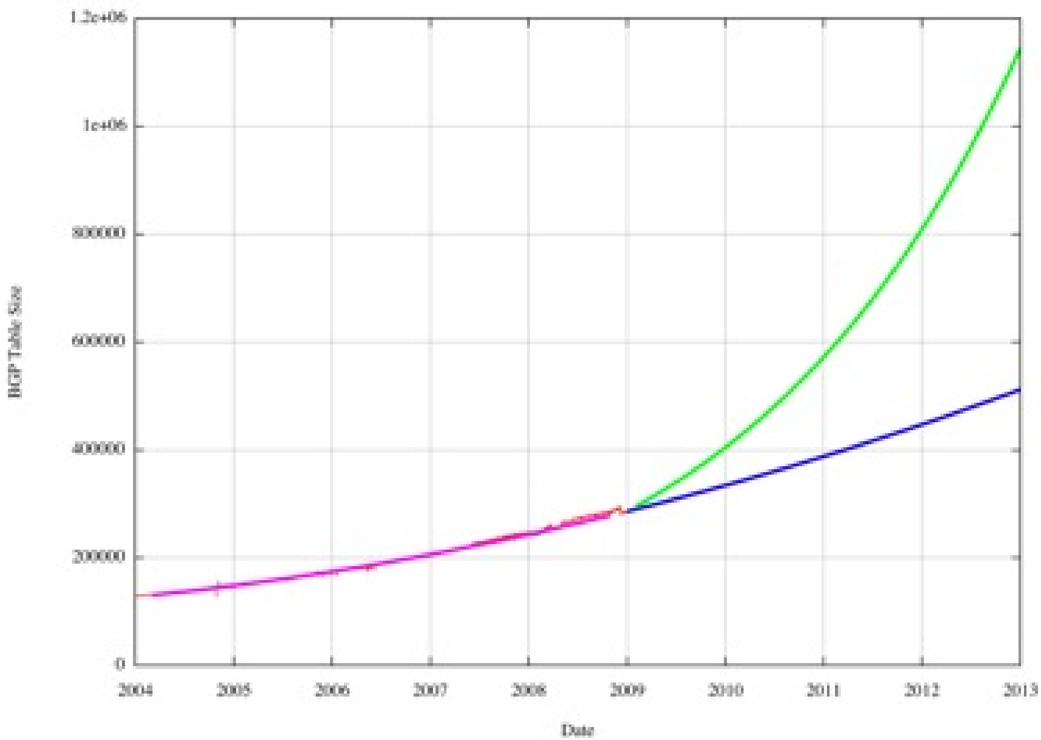

Figure F-4. Comparing BGP routing table growth with Moore's Law.

[40] G. Huston, BGP in 2008. The ISP Column, ISOC, March 2009.

Causes of Routing table growth

i) There are two mechanisms that shape the routing tables currently, namely
   a. Allocation practices, which define allocation of prefixes by the registries
   b. Routing practices define advertisement of those prefixes more specifically in BGP tables by ISPs and their customers.

ii) IP address allocation by registries directly, failure to aggregate properly, assignment by ISPs, Multi-homing and load balancing

iii) Allocation practice determines how and to whom address blocks are assigned. Routing practice determines which of these address blocks appears in the routing table and in what form (as a block or sub-block).

iv) RIRs try to ensure but do not guarantee that an allocation to an LIR is aggregatable with prior allocations.

v) ISPs might be using a first-fit or best fit algorithm to make address assignments to customers.

vi) The ideal case by CIDR is that each allocation to an ISP has exactly one entry in the routing table and no assignments appear in the routing table. Normally a customer advertises its assignment as a route in BGP or uses a static route or shares an IGP with the ISP. The ISP aggregates these routes and advertises to its peers or upstream provider its address prefix of its allocated space. In reality there are deviations to this ideal case caused by variety of routing practices that includes Multi-homing, customers connected to multiple upstream providers, and ISPs buying transit from multiple providers.

vii) Splitting allocations by ISPs and advertising the split prefixes separately in BGP would increase the routing table size. Splitting would cause a collection of sub-prefixes to appear in the routing table.

viii) LIRs are allocated IP addresses based on needs, and as these allocations are not guaranteed to be contiguous and the LIR may end up with several distinct prefix assigned to it.

ix) Every prefix assigned by the ISP is recorded in the routing table. As such, the size of the routing table gets increased as the number of allocations increases. That is one of the reasons to go for aggregation such that subsequent allocations can be aggregated to be represented as a single prefix by the ISP.

Analysis

Analysis of the routing tables in the DFZ shows that there is a stable growth in the table size drawing attention of the researchers. It is observed from the data collected of the various BGP routers that the DFZ routing table size growth rate is much higher than the allocated routing prefixes. As the available IPv4 address space becomes scarcer and with the expected wide scale deployment of IPv6, the DFZ routing table size is expected to face rapid growth.

With the present Internet routing and addressing architecture, both users and service providers share the same address and routing space. But, users and service providers do not share the same goals and challenges. With requirements focused on customer needs,
The requirements of the user and the service providers in terms of address space management are conflicting. From the service provider's perspective, the routing system needs topologically aggregateable addresses which would reduce fragmentation, more meaningful entries in the routing table and optimize the table size. To facilitate this RIRs follow a set of IP address policies, and the IP address allocation is Provider Based. But, customers expect Provider Independent (PI) addresses to ease site multihoming and avoid renumbering. Service providers and users follow different approaches on address allocation and routing announcements, each side having their own technical and economic reasons to support with. This has increased the number of fragmented prefixes, affecting aggregation resulting in growing table size at the DFZ routers.

The growing concern is not the growing size of the routing table itself as it is understood that as the IP addresses scale, obviously the routing table size is bound to scale. In an aspect the scaling routing table size reflects a healthy growth of the Network. The concern is what effect the scaling routing table size would have on the parameters such as memory, processing scalabilities, routing convergence delay and cost. Moore's law and the current Internet core routers capabilities have proven that the above is no longer a thing of concern.