

Network Working Group
Request for Comments: 1917
BCP: 4
Category: Best Current Practice

P. Nesser II
Nesser & Nesser Consulting
February 1996

An Appeal to the Internet Community to Return Unused IP Networks (Prefixes) to the IANA

Status of this Memo

This document specifies an Internet Best Current Practices for the Internet Community, and requests discussion and suggestions for improvements. Distribution of this memo is unlimited.

Abstract

This document is an appeal to the Internet community to return unused address space, i.e. any block of consecutive IP prefixes, to the Internet Assigned Numbers Authority (IANA) or any of the delegated registries, for reapportionment. Similarly an appeal is issued to providers to return unused prefixes which fall outside their customary address blocks to the IANA for reapportionment.

1. Background

The Internet of today is a dramatically different network than the original designers ever envisioned. It is the largest public data network in the world, and continues to grow at an exponential rate which doubles all major operational parameters every nine months. A common metaphor in engineering is that every time a problem increases in size by an order of magnitude, it becomes a new problem. This adage has been true over the lifetime of the Internet.

The Internet is currently faced with two major operational problems (among others). The first is the eventual exhaustion of the IPv4 address space and the second is the ability to route packets between the large number of individual networks that make up the Internet. The first problem is simply one of supply. There are only 2^{32} IPv4 addresses available. The lifetime of that space is proportional to the efficiency of its allocation and utilization. The second problem is mainly a capacity problem. If the number of routes exceeds the current capacity of the core Internet routers, some routes will be dropped and sections of the Internet will no longer be able to communicate with each other. The two problems are coupled and the dominant one has, and will, change over time.

The initial design of IP had all addresses the same, eight bits of network number and twenty four bits of host number. The expectation was of a few, large, global networks. During the first spurts of growth, especially with the invention of LAN technologies, it became obvious that this assumption was wrong and the separation of the address space into three classes (Class A for a few huge networks; Class B for more, smaller networks; and Class C for those really small LANs, with lots of network numbers) was implemented. Soon subnets were added so sites with many small LANs could appear as a single network to others, the first step at limiting routing table size. And finally, CIDR was introduced to the network, to add even more flexibility to the addressing, extending the split from three classes to potentially thirty different classes.

Subnets were introduced to provide a mechanism for sites to divide a single network number (Class A, B, or C) into pieces, allowing a higher utilization of address space, and thus promoting conservation of the IPv4 address space. Because of the built-in notion of classful addresses, subnetting automatically induced a reduction in the routing requirements on the Internet. Instead of using two (or more) class C networks, a site could subnet a single class B into two (or more) subnets. Both the allocation and the advertisement of a route to the second and succeeding class C's are saved.

Since 1993, the concept of classless (the "C" in CIDR) addresses have been introduced to the Internet community. Addresses are increasingly thought of as bitwise contiguous blocks of the entire address space, rather than a class A,B,C network. For example, the address block formerly known as a Class A network, would be referred to as a network with a /8 prefix, meaning the first 8 bits of the address define the network portion of the address. Sometimes the /8 will be expressed as a mask of 255.0.0.0 (in the same way a 16 bit subnet mask will be written as 255.255.0.0).

This scheme allows "supernetting" of addresses together into blocks which can be advertised as a single routing entry. The practical purpose of this effort is to allow service providers and address registries to delegate realistic address spaces to organizations and be unfettered by the traditional network classes, which were inappropriately sized for most organizations. For example the block of 2048 class C network numbers beginning with 192.24.0.0 and ending with 192.31.255.0 can be referenced as 192.24/19, or 192.24.0.0 with a mask of 255.248.0.0 (i.e. similar to a 19 bit subnet mask written in dotted decimal notation). The concept of "supernetting" allows the remaining Internet address space to be allocated in smaller blocks, thus allowing more networks and better efficiency. For a more detailed discussion refer to RFC 1518.

Like subnetting, CIDR also helps address the reduction of routing requirements, but it is not as automatic as the case of subnets. CIDR blocks are allocated in a way which promotes hierarchical routing. A provider is typically given a large block of addresses to redistribute to their customers. For example, if the provider P has been given the CIDR block 192.168/16, a block of 255 contiguous class C networks, they can provide one class C network to each of 255 customers (who may in turn subnet those class C networks into smaller pieces) yet still only advertise the single route 192.168/16. Thus CIDR only helps reduce the routing problem if blocks are assigned and maintained in a hierarchical manner.

RFC 1797 described a technical experiment designed to test the problems with allocating the currently reserved Class A network space. RFC 1879 described the results of this experiment. This effort shows that "supersubnetting" of a Class A network into numerous (even millions) of smaller networks is practical.

The dominating portion of the problem facing the Internet today is routing requirements. The following statements constitute a first order approximation based on current growth, a simple model of router resources, etc. Current routing technology can handle approximately twice the number of routes which are currently advertised on "core" Internet routers. Router capacity is doubling every 18 months, while routing tables are doubling every 9 months. If routes continue to be introduced at the current rate, the Internet will cease to function as a reliable infrastructure in approximately 2 to 3 years.

The good news is that CIDR is working. Address blocks are being allocated and assigned in a hierarchical manner, and the CIDR'ization of large portions of the address space which were assigned according to the guidelines of RFC 1466 resulted in a significant drop of advertised routes. However, recent growth trends show that the number of routes is once again growing at an exponential rate, and that the reduction with the introduction of CIDR was simply a sawtooth in the rate.

The growth in the number of routes can logically come from only two places, the extra routes generated with the breakup of CIDR blocks, and previously allocated and unannounced networks being connected. (Registries are still allocating a few addresses not within CIDR blocks, so a small thirdsource does exist.) With increasing popularity there is increasing competition between providers. If a site changes provider and retains the use of their CIDR block addresses, holes appear in the blocks and specific routes are added to the routing structure to accommodate these cases. Thus over time, CIDR will improve address utilization efficiency yet not help the routing requirements unless providers can keep their CIDR blocks

intact.

The second source for new route introduction is sites who had previously operated a private IP network, which had been registered and assigned a network number (or numerous networks), but have only recently connected to the global Internet. This RFC is a policy based attempt to help preserve the operation of the current Internet by addressing the issues of previously registered but unannounced IP networks.

An additional area of route introduction comes from non-aggregating router configurations. Aggregation is not automatic on most routers, and providers who may have intact CIDR blocks are, in many cases, advertising individual routes instead of an aggregate block without realizing.

In the context of this document, the phrase "Global Internet" refers to the mesh of interconnected public networks (Autonomous Systems) which has its origins in the U.S. National Science Foundation (NSF) backbone, other national networks, and commercial enterprises. Similarly, the phrase or any references to the "Core Routers" refer to the set of routers which carry the full set of route advertisements and act as interconnect points for the public networks making up the "Global Internet."

2. History

The IANA has historically managed the assignment of addresses to Internet sites. During the earliest days of the IANA, given a vast address space, the requirements for assignments of network address space were much less stringent than those required today. Organizations were essentially assigned networks based on their requests.

2.1 Class A Networks (/8 Prefixes)

The upper half of the Class A address space (64.0.0.0 - 126.0.0.0) (127.0.0.0 has traditionally been used by the Unix operating system as the "loopback" network, and is thus unavailable) has been reserved by the IANA for growth within the IPv4 address space. Of the lower half of the address space, 22 were assigned pre-1982, 6 were assigned between 1982 and 1987, 26 were assigned between 1988 and 1992, and 2 were assigned between 1993 and 1995. In May of 1995 four Class A networks previously assigned have been returned to the IANA. All remaining Class A addresses have also been reserved for growth within the IPv4 address space. The Class A address space is 50% of the total IPv4 address space.

2.2 Class B Networks (/16 prefixes)

From 1989 until 1993 approximately 80% of the currently assigned Class B IP networks were assigned or allocated. Allocations dropped dramatically in 1994 and 1995 due to the adoption of policies outlined in RFC 1466. 61.65% of the Class B address space is currently allocated. The class B address space is 25% of the total IPv4 address space.

2.3 Class C Networks (/24 Prefixes)

With the introduction of CIDR and RFC 1466 the allocation of Class C address space has skyrocketed since 1993. 27.82% of the Class C address space is currently allocated. The class C address space is 12.5% of the total IPv4 address space.

2.4 Class "D" and Beyond

Of the remaining 12.5% of the address space, the lower 6.25% is allocated for multicast applications (mbone, OSPF, etc.) and the upper half is reserved for future applications.

2.5 Totals

The weighted total shows that 40.99% of the total IPv4 address space is allocated and the remainder is reserved for future growth. It should be noted that careful extrapolations of the current trends suggest that the address space will be exhausted early in the next century.

3. Problem

Before the introduction of RFC 1466 and of CIDR, some 50,000 networks were assigned by the IANA, yet only a small percentage (30-40%) of the sites actually had connections to the global Internet and advertised those networks. As the popularity of the Internet is growing, a growing number of those sites are being connected, and increasing the size of the routing tables.

Current Internet sites have received their address assignments in various ways and steps. Some sites, through a little (or in some cases no) work, could donate unused IP nets back to the IANA.

Some organizations have made small requests at first and received a Class C assignment (or multiple Class C assignments), and after unexpected growth made subsequent requests and received Class B assignments.

Several Internet service providers were given blocks of the Class B address space to distribute to customers. This space was often provided to clients based upon a level of service purchased rather than actual need.

Many organizations have either merged or are associated with parent organizations which produce situations with large inefficiencies in address assignment.

Many organizations have requested addresses based on their need to run TCP/IP on internal machines which have no interest in connecting to the global Internet. Most vendors manuals have instructed (and provided copies of the application forms), sites to request IP address assignments.

Other organizations have large internal IP networks, and are connected to the Internet through application layer gateways or network address translators, and will never announce their internal networks.

4. Appeal

To the members of the Internet community who have IP network assignments which may be currently unused, the Internet community would like to encourage you to return those addresses to the IANA or your provider for reapportionment.

Specifically those sites who have networks which are unused are encouraged to return those addresses. Similarly to those sites who are using a small percentage of their address space and who could relatively easily remove network assignments from active use, the Internet community encourages such efforts.

To those sites who have networks which will never need to connect to the global Internet, or for security reasons will always be isolated, consider returning the address assignments to the IANA or your provider and utilizing prefixes recommended in RFC 1597.

In those cases where renumbering is required, sites are encouraged to put into place a plan to renumber machines, as is reasonably convenient, and work towards minimizing the number of routes advertised to their providers.

4.1 Suggestions to Providers

Many providers are currently advertising non-CIDR routes which encompass a large block of addresses, ie any Class A (0/1) or Class B (128/2) space. Some customers who are only using a percentage of

their address space (assuming they are subnetting using contiguous bits) may be willing to allow usage of the upper portion of their assigned address space by their providers other customers.

This scheme requires certain elements be installed or already in place to get the routing correct, but has the potential to gain the use of a large number of small networks without growth of the global routing tables. This would require additional measures of cooperation between providers and their customers but could prove to have both economic advantages, as well as good Internet citizen standing.

For example, large organization S has been assigned the class A block of addresses 10.0.0.0. and is currently using provider P for their connection to the global Internet. P is already advertising the route for 10.0.0.0 to the global Internet. S has been allocating its internal networks using a right to left bit incrementing model. P and S could agree that S will allow some /18 (for example) prefixes to be made available for P's other customers. This would impose no hardships whatsoever on S, presuming his router can speak BGP, and allow P to attach a huge number of small customers without the need to advertise more routes or request additional address blocks from the IANA or their upstream provider.

The "Net 39" experiment as outlined in RFC 1797 and summarized in RFC 1879 provided practical data on the implementation of the suggested schemes.

Additionally, providers are encouraged to release all unused networks which fall outside of their normal address blocks back to the IANA or the appropriate registry.

New customers, particularly those who may have recently changed providers, and who have small networks which are not part of CIDR'ized blocks, should be encouraged to renumber and release their previous addresses back to the provider or the IANA.

Since the first introduction of CIDR in April of 1994, many providers have aggressively pursued the concepts of aggregation. Some providers actively persuaded their customers to renumber, while others pursued peering arrangements with other providers, and others did both. Providers should continue to actively and routinely pursue both methods to streamline routing table growth. Cooperation between providers is absolutely essential to short (and long) term management of routing requirements.

Providers should regularly verify the routes they are advertising to their upstream provider(s) to validate their router configurations and confirm correct aggregation is occurring.

4.2 Suggestions to the IANA and Address Registries

In cases where addresses are returned to the IANA, or any other address registry, which fits into another registry or providers block, the addresses should be turned over to the appropriate authority. This will help maximize the availability of addresses and minimize routing table loads.

4.3 How to Return a Block of Address Space to the IANA

Send the following form to Hostmaster@internic.net & iana@isi.edu, changing the \$NET_PREFIX to the network being returned.

Please update the contact information on the following net as follows:

Netname: RESERVED

Netnumber: \$NET_PREFIX

Coordinator:

Reynolds, Joyce K. (JKR1) JKRey@ISI.EDU
(310) 822-1511

Alternate Contact:

Postel, Jon (JBP) POSTEL@ISI.EDU
(310) 822-1511

4.4 How to Return a Block of Address Space to another Address Registry

Each registry will have its own forms and addresses. Please contact the appropriate registry directly.

5. Conclusion

Rationalizing the global addressing hierarchy is a goal which should be supported by any organization which is currently connected or plans to connect to the Internet. If (and possibly when) the situation ever reaches a critical point, the core service providers whose routers are failing and losing routes will be forced to make one of two choices, both painful to the user community.

They could begin blocking routes to their customers who are advertising too many disjoint routes, where "too many" will be set at the level necessary to keep their routers functioning properly. This is a domino effect since the next level of providers will be forced to make the same effort, until individual organizations are forced to only advertise routes to portions of their networks.

The second option the core providers have is to charge for advertised routes. The price level will be set at a point which reduces the number of routes to a level which will keep their routers functioning properly. Once again a domino effect will take place until the price increases will effect individual organizations.

Some planning and efforts by organizations and providers now while there is a some time available can help delay or prevent either or the two scenarios from occurring.

This system has already produced very favorable results when applied on a small scale. As of this writing 4 Class A networks have been returned to the IANA. This may not seem significant but those 4 networks represent over 1.5% of the total IPv4 address capacity.

6. References

1. Gerich, E., "Guidelines for Management of the IP Address Space", RFC 1466, May 1993.
2. Topolcic, C., "Status of CIDR Deployment in the Internet", RFC 1467, August 1993.
3. Rekhter, Y., and T. Li, "An Architecture for IP Address Allocation with CIDR", RFC 1518, September 1993.
4. Fuller, V., Li, T., Yu, J., and K. Varadhan, "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy", RFC 1519, September 1993.
5. Rekhter, Y., Moskowitz, R., Karrenberg, D., and de Groot, G., "Address Allocation for Private Internets", RFC 1597, March 1994.
6. Lear, E., Fair, E., Crocker, D., and T. Kessler, "Network 10 Considered Harmful (Some Practices Shouldn't be Codified)", RFC 1627, July 1994.
7. Huitema, C., "The H Ratio for Address Assignment Efficiency", RFC 1715, November 1994.

8. IANA, Class A Subnet Experiment, RFC 1797, April 1995.

7. Security Considerations

Security issues are not discussed in this memo.

8. Acknowledgements

I would like to thank the members of the CIDRD mailing list and working groups for their suggestion and comments on this document. Specific thanks should go to Michael Patton, Tony Li, Noel Chiappa, and Dale Higgs for detailed comments and suggestions.

9. Author's Address

Philip J. Nesser II
Nesser & Nesser Consulting
16015 84th Avenue N.E.
Bothell, WA 98011-4451

Phone: (206) 488-6268
Fax: (206) 488-6268
EMail: pjnesser@martigny.ai.mit.edu

