

Network Working Group  
Request for Comments: 2008  
BCP: 7  
Category: Best Current Practice

Y. Rekhter  
T. Li  
Cisco Systems  
October 1996

## Implications of Various Address Allocation Policies for Internet Routing

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.

IESG Note:

The addressing constraints described in this document are largely the result of the interaction of existing router technology, address assignment, and architectural history. After extensive review and discussion, the authors of this document, the IETF working group that reviewed it, and the IESG have concluded that there are no other currently deployable technologies available to overcome these limitations. In the event that routing or router technology develops to the point that adequate routing aggregation can be achieved by other means or that routers can deal with larger routing and more dynamic tables, it may be appropriate to review these constraints.

1 Abstract

IP unicast address allocation and management are essential operational functions for the Public Internet. The exact policies for IP unicast address allocation and management continue to be the subject of many discussions. Such discussions cannot be pursued in a vacuum - the participants must understand the technical issues and implications associated with various address allocation and management policies.

The purpose of this document is to articulate certain relevant fundamental technical issues that must be considered in formulating unicast address allocation and management policies for the Public Internet, and to provide recommendations with respect to these policies.

The major focus of this document is on two possible policies, "address ownership" and "address lending," and the technical implications of these policies for the Public Internet. For the organizations that could provide reachability to a sufficiently large

fraction of the total destinations in the Internet, and could express such reachability through a single IP address prefix the document suggests to use the "address ownership" policy. However, applying the "address ownership" policy to every individual site or organization that connects to the Internet results in a non-scalable routing.

Consequently, this document also recommends that the "address lending" policy should be formally added to the set of address allocation policies in the Public Internet. The document also recommends that organizations that do not provide a sufficient degree of routing information aggregation, but wish to obtain access to the Internet routing services should be strongly encouraged to use this policy to gain access to the services.

## 2 On the intrinsic value of IP addresses

Syntactically, the set of IPv4 unicast addresses is the (finite) set of integers in the range 0x00000000 - 0xDFFFFFFF. IP addresses are used for Network Layer (IP) routing. An IP address is the sole piece of information about the node injected into the routing system.

The notable semantics of an IP unicast address is its ability to interact with the Public Internet routing service and thereby exchange data with the remainder of the Internet. In other words, for the Public Internet, it is the reachability of an IP address that gives it an intrinsic value. Observe, however, that IP addresses are used outside of the Public Internet. This document does not cover the value of addresses in other than the Public Internet context.

The above implies that in the Public Internet it is the service environment (the Internet) and its continued operation, including its routing system, which gives an IP address its intrinsic value, rather than the inverse. Consequently, if the Public Internet routing system ceases to be operational, the service disappears, and the addresses cease to have any functional value in the Internet. At this point, for the Public Internet, all address allocation and management policies, including existing policies, are rendered meaningless.

## 3 Hierarchical routing and its implication on address allocation

Hierarchical routing [Kleinrock 77] is a mechanism that improves the scaling properties of a routing system. It is the only proven mechanism for scaling routing to the current size of the Internet.

Hierarchical routing requires that addresses be assigned to reflect the actual network topology. Hierarchical routing works by taking the set of addresses covered by a portion of the topology, and generating a single routing advertisement (route) for the entire set. Further,

hierarchical routing allows this to be done recursively: multiple advertisements (routes) can be combined into a single advertisement (route). By exercising this recursion, the amount of information necessary to provide routing can be decreased substantially.

A common example of hierarchical routing is the phone network, where country codes, area codes, exchanges, and finally subscriber lines are different levels in the hierarchy. In the phone network, a switch need not keep detailed routing information about every possible subscriber in a distant area code. Instead, the switch usually knows one routing entry for the entire area code.

Notice that the effect on scaling is dramatic. If we look at the space complexity of the different schemes, the switch that knows about every subscriber in the world needs  $O(n)$  space for  $n$  worldwide subscribers. Now consider the case of hierarchical routing. We can break  $n$  down into the number of subscribers in the local area ( $l$ ), the other exchanges in the area code ( $e$ ), the other area codes in the local country code ( $a$ ) and other country codes ( $c$ ). Using this notation, hierarchical routing has space complexity  $O(l + e + a + c)$ . Notice that each of these factors is much, much less than  $n$ , and grows very slowly, if at all. This implies that a phone switch can be built today that has some hope of not running out of space when it is deployed.

The fundamental property of hierarchical routing that makes this scalability possible is the ability to form abstractions: here, the ability to group subscribers into exchanges, area codes and country codes. Further, such abstractions must provide useful information for the ability to do routing. Some abstractions, such as the group of users with green phones, are not useful when it comes time to route a call.

Since the information that the routing system really needs is the location of the address within the topology, for hierarchical routing, the useful abstraction must capture the topological location of an address within the network. In principle this could be accomplished in one of two ways. Either (a) constrain the topology (and allowed topology changes) to match address assignment. Or, (b) avoid constraints on the topology (and topology changes), but require that as the topology changes, an entity's address change as well. The process of changing an entity's address is known as "renumbering."

#### 4 Scaling the Internet routing system

The enormous growth of the Public Internet places a heavy load on the Internet routing system. Before the introduction of CIDR the growth rate had doubled the size of the routing table roughly every nine months. Capacity of computer technology doubles roughly every 24 months. Even if we could double the capacities of the routers in the Internet every 24 months, inevitably the size of the routing tables is going to exceed the limit of the routers. Therefore, to preserve uninterrupted continuous growth of the Public Internet, deploying mechanisms that contain the growth rate of the routing information is essential.

Lacking mechanisms to contain the growth rate of the routing information, the growth of the Internet would have to be either limited or frozen, or the Internet routing system would become overloaded. The result of overloading routing is that the routing subsystem will fail: either equipment (routers) could not maintain enough routes to insure global connectivity, or providers will simply exclude certain routes to insure that other routes provide connectivity to particular sites. This document assumes that neither of the outcomes mentioned in this paragraph is acceptable.

Classless Inter-Domain Routing (CIDR) [RFC1518, RFC1519] has been deployed since late 1992 in the Public Internet as the primary mechanism to contain the growth rate of the routing information - without CIDR the Internet routing system would have already collapsed. For example, in October 1995, within AlterNet (one of the major Internet Service Providers) there were 3194 routes. Thanks to aggregation, AlterNet advertised only 799 routes to the rest of the Internet - a saving of 2395 routes (75%) [Partan 95]. In October 1995 the Internet Routing Registry (IRR) contained 61,430 unique prefixes listed, not counting prefixes marked as withdrawn (or 65,191 prefixes with prefixes marked as withdrawn). That is roughly a lower bound since many prefixes are not registered in the IRR. CIDR aggregation resulted in less than 30,000 routes in the default-free part of the Internet routing system [Villamizar 95].

CIDR is an example of the application of hierarchical routing in the Public Internet, where subnets, subscribers, and finally providers are some possible levels in the hierarchy. For example, a router within a site need not keep detailed routing information about every possible host in that site. Instead, the router maintains routing information on a per subnet basis. Likewise, a router within a provider need not keep detailed routing information about individual subnets within its subscribers. Instead, the router could maintain routing information on a per subscriber basis. Moreover, a router within a provider need not keep detailed routing information about

stub (single home) subscribers of other providers by maintaining routing information on a per provider basis.

Because of pre-CIDR address allocation, many routes in the Internet are not suitable for hierarchical aggregation. Moreover, unconnected sites with pre-CIDR address allocations exist. If these sites connect to the Internet at some point in the future, the routes to these sites are unlikely to be suitable for hierarchical aggregation. Also, when a site uses addresses obtain from its provider, but then later switches to a different provider (while continuing to use the same addresses), the route to the site may no longer be suitable for hierarchical aggregation.

Hierarchical routing requires that aggregation boundaries for the addressing information be formed along some hierarchy. As a result, many exceptions will be injected into the routing system in the future, besides those exceptions that currently exist. Each exception added to the routing system deters the scalability of the routing system. The exact number of exceptions that can be tolerated is dependent on the technology used to support routing. Unbridled growth in the number of such exceptions will cause the routing system to collapse.

## 5 Address allocation and management policies

IP address allocation and management policy is a complex, multifaceted issue. It covers a broad range of issues, such as who formulates the policies, who executes the policies, what is the role of various registries, what is the role of various organizations (e.g., ISOC, IAB, IESG, IETF, IEPG, various government bodies, etc.), the participation of end users in requesting addresses, and so on. Address allocation and management and the scalability of the routing system are interrelated - only certain address allocation and management policies yield scalable routing. The Internet routing system is subject to both technological and fundamental constraints. These constraints restrict the choices of address allocation policies that are practical.

### 5.1 The "address ownership" allocation policy and its implications on the Public Internet

"Address ownership" is one possible address allocation and management policy. The "address ownership" policy means that part of the address space, once allocated to an organization, remains allocated to the organization as long as that organization wants it. Further, that portion of the address space would not be allocated to any other organization. Often, such addresses are called "portable." It was assumed that if an organization acquires its addresses via the

"address ownership" policy, the organization would be able to use these addresses to gain access to the Internet routing services, regardless of where the organization connects to the Internet.

While it has never been explicitly stated that various Internet Registries use the "address ownership" allocation policy, it has always been assumed (and practiced).

To understand the implications of the "address ownership" policy ("portable" addresses) on the scalability of the Internet routing system, one must observe that:

(a) By definition, address ownership assumes that addresses, once assigned, fall under the control of the assignee. It is the assignee that decides when to relinquish the ownership (although the decision could be influenced by various factors). Specifically, the assignee is not required (but may be influenced) to relinquish the ownership as the connectivity of the assignee to the Internet changes.

(b) By definition, hierarchical routing assumes that addresses reflect the network topology as much as possible.

Therefore, the only presently known practical way to satisfy both scalable hierarchical routing and address ownership for everyone is to assume that the topology (or at least certain pieces of it) will be permanently fixed. Given the distributed, decentralized, largely unregulated, and global (international) nature of the Internet, constraining the Internet topology (or even certain parts of it) may have broad technical, social, economical, and political implications. To date, little is known of what these implications are; even less is known whether these implications would be acceptable (feasible) in practice. Therefore, at least for now, we have to support an Internet with an unconstrained topology (and unconstrained topological changes).

Since the Internet does not constrain its topology (or allowed topology changes), we can either have address ownership for everyone or a routable Internet, but not both, or we need to develop and deploy new mechanisms (e.g., by decoupling the address owned by the end users from those used by the Internet routing, and provide mechanisms to translate between the two). In the absence of new mechanisms, if we have address ownership ("portable" addresses) for everyone, then the routing overhead will lead to a breakdown of the routing system resulting in a fragmented (partitioned) Internet. Alternately, we can have a routable Internet, but without address ownership ("portable" addresses) for everyone.

## 5.2 The "address lending" allocation policy and its implications for the Public Internet

Recently, especially since the arrival of CIDR, some subscribers and providers have followed a model in which address space is not owned (not portable), but is bound to the topology. This model suggests an address allocation and management policy that differs from the "address ownership" policy. The following describes a policy, called "address lending," that provides a better match (as compared to the "address ownership" policy) to the model.

An "address lending" policy means that an organization gets its addresses on a "loan" basis. For the length of the loan, the lender cannot lend the addresses to any other borrower. Assignments and allocations based on the "address lending" policy should explicitly include the conditions of the loan. Such conditions must specify that allocations are returned if the borrower is no longer contractually bound to the lender, and the lender can no longer provide aggregation for the allocation. If a loan ends, the organization can no longer use the borrowed addresses, and therefore must get new addresses and renumber to use them. The "address lending" policy does not constrain how the new addresses could be acquired.

This document expects that the "address lending" policy would be used primarily by Internet Registries associated with providers; however, this document does not preclude the use of the "address lending" policy by an Internet Registry that is not associated with a provider.

This document expects that when the "address lending" policy is used by an Internet Registry associated with a provider, the provider is responsible for arranging aggregation of these addresses to a degree that is sufficient to achieve Internet-wide IP connectivity.

This document expects that when the "address lending" policy is used by an Internet Registry associated with a provider, the terms and conditions of the loan would be coupled to the service agreement between the provider and the subscribers. That is, if the subscriber moves to another provider, the loan would be canceled.

To reduce disruptions when a subscriber changes its providers, this document strongly recommends that the terms and conditions of the loan should include provision for a grace period. This provision would allow a subscriber that disconnects from its provider a certain grace period after the disconnection. During this grace period, the borrower (the subscriber) may continue to use the addresses obtained under the loan. This document recommends a grace period of at least 30 days. Further, to contain the routing information overhead, this document suggests that a grace period be no longer than six months.

To understand the scalability implications of the "address lending" policy, observe that if a subscriber borrows its addresses from its provider's block, then the provider can advertise a single address prefix. This reduces the routing information that needs to be carried by the Internet routing system (for more information, see Section 5.3.1 of RFC1518). As the subscriber changes its provider, the loan from the old provider would be returned, and the loan from the new provider would be established. As a result, the subscriber would renumber to the new addresses. Once the subscriber renumbers into the new provider's existing blocks, no new routes need to be introduced into the routing system.

Therefore, the "address lending" policy, if applied appropriately, is consistent with the constraints on address allocation policies imposed by hierarchical routing, and thus promotes a scalable routing system. Thus, the "address lending" policy, if applied appropriately, could play an important role in enabling the continuous uninterrupted growth of the Internet.

To be able to scale routing in other parts of the hierarchy, the "lending" policy may also be applied hierarchically, so that addresses may in turn be lent to other organizations. The implication here is that the end of a single loan may have effects on organizations that have recursively borrowed parts of the address space from the main allocation. In this case, the exact effects are difficult to determine a priori.

### 5.3 In the absence of an explicit "address lending" policy

Organizations connecting to the Internet should be aware that even if their current provider, and the provider they switch to in the future do not require renumbering, renumbering may still be needed to achieve Internet-wide IP connectivity. For example, an organization may now receive Internet service from some provider and allocate its addresses out of the CIDR block associated with the provider. Later the organization may switch to another provider. The previous provider may still be willing to allow the organization to retain part of the provider's CIDR block, and accept a more specific prefix



for that organization from the new provider. Likewise, the new provider may be willing to accept that organization without renumbering and advertise the more specific prefix (that covers destinations within the organization) to the rest of the Internet. However, if one or more other providers exist, that are unwilling or unable to accept the longer prefix advertised by the new provider, then the organization would not have IP connectivity to part of the Internet. Among the possible solutions open to the organization may be either to renumber, or for others to acquire connectivity to providers that are willing and able to accept the prefix.

The above shows that the absence of an explicit "address lending" policy from a current provider in no way ensures that renumbering will not be required in the future when changing providers. Organizations should be aware of this fact should they encounter a provider making claims to the contrary.

## 6 Recommendations

Observe that the goal of hierarchical routing in the Internet is not to reduce the total amount of routing information in the Internet to the theoretically possible minimum, but just to contain the volume of routing information within the limits of technology, price/performance, and human factors. Therefore, organizations that could provide reachability to a sufficiently large fraction of the total destinations in the Internet and could express such reachability through a single IP address prefix could expect that a route with this prefix will be maintained throughout the default-free part of the Internet routing system, regardless of where they connect to the Internet. Therefore, using the "address ownership" policy when allocating addresses to such organizations is a reasonable choice. Within such organizations this document suggests the use of the "address lending" policy.

For all other organizations that expect Internet-wide IP connectivity, the reachability information they inject into the Internet routing system should be subject to hierarchical aggregation. For such organizations, allocating addresses based on the "address ownership" policy makes hierarchical aggregation difficult, if not impossible. This, in turn, has a very detrimental effect on the Internet routing system. To prevent the collapse of the Internet routing system, for such organizations, this document recommends using the "address lending" policy. Consequently, when such an organization first connects to the Public Internet or changes its topological attachment to the Public Internet, the organization eventually needs to renumber. Renumbering allows the organization to withdraw any exceptional prefixes that the organization would otherwise inject into the Internet routing system. This applies to

the case where the organization takes its addresses out of its direct provider's block and the organization changes its direct provider. This may also apply to the case where the organization takes its addresses out of its indirect provider's block, and the organization changes its indirect provider, or the organization's direct provider changes its provider.

Carrying routing information has a cost associated with it. This cost, at some point, may be passed back in full to the organizations that inject the routing information. Aggregation of addressing information (via CIDR) could reduce the cost, as it allows an increase in the number of destinations covered by a single route. Organizations whose addresses are allocated based on the "address ownership" policy (and thus may not be suitable for aggregation) should be prepared to absorb the cost completely on their own.

Observe that neither the "address ownership," nor the "address lending" policy, by itself, is sufficient to guarantee Internet-wide IP connectivity. Therefore, we recommend that sites with addresses allocated based on either policy should consult their providers about the reachability scope that could be achieved with these addresses, and associated costs that result from using these addresses.

If an organization doesn't require Internet-wide IP connectivity, then address allocation for the organization could be done based on the "address ownership" policy. Here, the organization may still maintain limited IP connectivity (e.g., with all the subscribers of its direct provider) by limiting the distribution scope of its routing information to its direct provider. Connectivity to the rest of the Internet can be handled by mediating gateways (e.g., application layer gateways, Network Address Translators (NATs)). Note that use of mediating gateways eliminates the need for renumbering, and avoids burdening the Internet routing system with non-aggregatable addressing information; however they have other drawbacks which may prove awkward in certain situations.

Both renumbering (due to the "address lending" policy), and non-aggregated routing information (due to the "address ownership" policy), and the use of mediating gateways result in some costs. Therefore, an organization needs to analyze its own connectivity requirements carefully and compare the tradeoffs associated with addresses acquired via either policy vs. having connectivity via mediating gateways (possibly augmented by limited IP connectivity) using addresses acquired via "address ownership." To reduce the cost of renumbering, organizations should be strongly encouraged to deploy tools that simplify renumbering (e.g., Dynamic Host Configuration Protocol [RFC 1541]). Use of the DNS should be strongly encouraged.

## 7 Summary

Any address allocation and management policy for IP addresses used for Internet connectivity must take into account its impact on the scalability of the Public Internet routing system. Among all of the possible address allocation and management policies only the ones that yield a scalable routing system are feasible. All other policies are self-destructive in nature, as they lead to a collapse of the Internet routing system, and therefore to the fragmentation (partitioning) of the Public Internet.

Within the context of the current Public Internet, address allocation and management policies that assume unrestricted address ownership have an extremely negative impact on the scalability of the Internet routing system. Such policies are almost certain to exhaust the scalability of the Internet routing system well before we approach the exhaustion of the IPv4 address space and before we can make effective use of the IPv6 address space. Given the Internet's growth rate and current technology, the notion that everyone can own address space and receive Internet-wide routing services, despite where they connect to the Internet, is currently technically infeasible. Therefore, this document makes two recommendations. First, the "address lending" policy should be formally added to the set of address allocation policies in the Public Internet. Second, organizations that do not provide a sufficient degree of routing information aggregation to obtain access to the Internet routing services should be strongly encouraged to use this policy to gain access to the services.

Since the current IPv6 address allocation architecture is based on CIDR, recommendations presented in this document apply to IPv6 address allocation and management policies as well.

## 8 Security Considerations

Renumbering a site has several possible implications on the security policies of both the site itself and sites that regularly communicate with the renumbering sites.

Many sites currently use "firewall" systems to provide coarse-grained access control from external networks, such as The Internet, to their internal systems. Such firewalls might include access control decisions based on the claimed source address of packets arriving at such firewall systems. When the firewall policy relates to packets arriving on the firewall from inside the site, then that firewall will need to be reconfigured at the same time that the site itself renumbers. When the firewall policy relates to packets arriving at the firewall from outside the site, then such firewalls will need to

be reconfigured whenever an outside site that is granted any access inside the site through the firewall is renumbered.

It is highly inadvisable to rely upon unauthenticated source or destination IP addresses for security policy decisions. [Bellovin89] IP address spoofing is not difficult with widely available systems, such as personal computers. A better approach would probably involve the use of IP Security techniques, such as the IP Authentication Header [RFC-1826] or IP Encapsulating Security Payload [RFC-1827], at the firewall so that the firewall can rely on cryptographic techniques for identification when making its security policy decisions.

It is strongly desirable that authentication be present in any mechanism used to renumber IP nodes. A renumbering mechanism that lacks authentication could be used by an adversary to renumber systems that should not have been renumbered, for example.

There may be other security considerations that are not covered in this document.

## 9 Acknowledgments

This document borrows heavily from various postings on various mailing lists. Special thanks to Noel Chiappa, Dennis Ferguson, Eric Fleischman, Geoff Huston, and Jon Postel whose postings were used in this document.

Most of the Section 5.3 was contributed by Curtis Villamizar. The Security section was contributed by Ran Atkinson.

Many thanks to Scott Bradner, Randy Bush, Brian Carpenter, Noel Chiappa, David Conrad, John Curran, Sean Doran, Dorian Kim, Thomas Narten, Andrew Partan, Dave Piscitello, Simon Poole, Curtis Villamizar, and Nicolas Williams for their review, comments, and contributions to this document.

Finally, we like to thank all the members of the CIDR Working Group for their review and comments.

## 9 References

- [Bellovin89] Bellovin, S., "Security Problems in the TCP/IP Protocol Suite", ACM Computer Communications Review, Vol. 19, No. 2, March 1989.
- [Kleinrock 77] Kleinrock, L., and K. Farouk, K., "Hierarchical Routing for Large Networks," Computer Networks 1 (1977), North-Holland Publishing Company.
- [Partan 95] Partan, A., private communications, October 1995.
- [RFC 1541] Droms, R., "Dynamic Host Configuration Protocol", October 1993.
- [RFC 1519] Fuller, V., Li, T., Yu, J., and K. Varadhan, "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy", September 1993.
- [RFC 1518] Rekhter, Y., and T. Li, "An Architecture for IP Address Allocation with CIDR", September 1993.
- [RFC 1825] Atkinson, R., "IP Security Architecture", RFC 1825, August 1995.
- [RFC 1826] Atkinson, R., "IP Authentication Header (AH)", RFC 1826, August 1995.
- [RFC 1827] Atkinson, R., "IP Encapsulating Security Payload (ESP)", RFC 1827, August 1995.
- [Villamizar 95] Villamizar, C., private communications, October 1995.

## 10 Authors' Addresses

Yakov Rekhter  
cisco Systems, Inc.  
170 Tasman Dr.  
San Jose, CA 95134  
Phone: (914) 528-0090  
EMail: yakov@cisco.com

Tony Li  
cisco Systems, Inc.  
170 Tasman Dr.  
San Jose, CA 95134  
Phone: (408) 526-8186  
EMail: tli@cisco.com