

IPv4 Address Behaviour Today

Status of this Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

The main purpose of this note is to clarify the current interpretation of the 32-bit IP version 4 address space, whose significance has changed substantially since it was originally defined. A short section on IPv6 addresses mentions the main points of similarity with, and difference from, IPv4.

Table of Contents

1. Introduction.....	1
2. Terminology.....	2
3. Ideal properties.....	3
4. Overview of the current situation of IPv4 addresses.....	4
4.1. Addresses are no longer globally unique locators.....	4
4.2. Addresses are no longer all temporally unique.....	6
4.3. Multicast and Anycast.....	7
4.4. Summary.....	8
5. IPv6 Considerations.....	8
ANNEX: Current Practices for IPv4 Address Allocation & Routing..	9
Security Considerations.....	10
Acknowledgements.....	11
References.....	11
Authors' Addresses.....	13

1. Introduction

The main purpose of this note is to clarify the current interpretation of the 32-bit IP version 4 address space, whose significance has changed substantially since it was originally defined in 1981 [RFC 791].

This clarification is intended to assist protocol designers, product implementors, Internet service providers, and user sites. It aims to avoid misunderstandings about IP addresses that can result from the substantial changes that have taken place in the last few years, as a result of the Internet's exponential growth.

A short section on IPv6 addresses mentions the main points of similarity with, and difference from, IPv4.

2. Terminology

It is well understood that in computer networks, the concepts of directories, names, network addresses, and routes are separate and must be analysed separately [RFC 1498]. However, it is also necessary to sub-divide the concept of "network address" (abbreviated to "address" from here on) into at least two notions, namely "identifier" and "locator". This was perhaps less well understood when RFC 791 was written.

In this document, the term "host" refers to any system originating and/or terminating IPv4 packets, and "router" refers to any system forwarding IPv4 packets from one host or router to another.

For the purposes of this document, an "identifier" is a bit string which is used throughout the lifetime of a communication session between two hosts, to identify one of the hosts as far as the other is concerned. Such an identifier is used to verify the source of incoming packets as being truly the other end of the communication concerned, e.g. in the TCP pseudo-header [RFC 793] or in an IP Security association [RFC 1825]. Traditionally, the source IPv4 address in every packet is used for this.

Note that other definitions of "identifier" are sometimes used; this document does not claim to discuss the general issue of the semantics of end-point identifiers.

For the purposes of this document, a "locator" is a bit string which is used to identify where a particular packet must be delivered, i.e. it serves to locate the place in the Internet topology where the destination host is attached. Traditionally, the destination IPv4 address in every packet is used for this. IP routing protocols interpret IPv4 addresses as locators and construct routing tables based on which routers (which have their own locators) claim to know a route towards the locators of particular hosts.

Both identifiers and locators have requirements of uniqueness, but these requirements are different. Identifiers must be unique with respect to each set of inter-communicating hosts. Locators must be

unique with respect to each set of inter-communicating routers (which we will call a routing "realm"). While locators must be unique within a given routing realm, this uniqueness (but not routability) could extend to more than one realm. Thus we can further distinguish between a set of realms with unique locators versus a set of realms with non-unique (overlapping) locators.

Both identifiers and locators have requirements of lifetime, but these requirements are different. Identifiers must be valid for at least the maximum lifetime of a communication between two hosts. Locators must be valid only as long as the routing mechanisms so require (which could be shorter or longer than the lifetime of a communication).

It will be noted that it is a contingent fact of history that the same address space and the same fields in the IP header (source and destination addresses) are used by RFC 791 and RFC 793 for both identifiers and locators, and that in the traditional Internet a host's identifier is identical to its locator, as well as being spatially unique (unambiguous) and temporally unique (constant).

These uniqueness conditions had a number of consequences for design assumptions of routing (the infrastructure that IPv4 locators enable) and transport protocols (that which depends on the IP connectivity). Spatial uniqueness of an address meant that it served as both an interface identifier and a host identifier, as well as the key to the routing table. Temporal uniqueness of an address meant that there was no need for TCP implementations to maintain state regarding identity of the far end, other than the IP address. Thus IP addresses could be used both for end-to-end IP security and for binding upper layer sessions.

Generally speaking, the use of IPv4 addresses as locators has been considered more important than their use as identifiers, and whenever there has been a conflict between the two uses, the use as a locator has prevailed. That is, it has been considered more useful to deliver the packet, then worry about how to identify the end points, than to provide identity in a packet that cannot be delivered. In other words, there has been intensive work on routing protocols and little concrete work on other aspects of address usage.

3. Ideal properties.

Whatever the constraints mentioned above, it is easy to see the ideal properties of identifiers and locators. Identifiers should be assigned at birth, never change, and never be re-used. Locators should describe the host's position in the network's topology, and should change whenever the topology changes.

Unfortunately neither of these ideals are met by IPv4 addresses. The remainder of this document is intended as a snapshot of the current real situation.

4. Overview of the current situation of IPv4 addresses.

It is a fact that IPv4 addresses are no longer all globally unique and no longer all have indefinite lifetimes.

4.1 Addresses are no longer globally unique locators

[RFC 1918] shows how corporate networks, a.k.a. Intranets, may if necessary legitimately re-use a subset of the IPv4 address space, forming multiple routing realms. At the boundary between two (or more) routing realms, we may find a spectrum of devices that enables communication between the realms.

At one end of the spectrum is a pure Application Layer Gateway (ALG). Such a device acts as a termination point for the application layer data stream, and is visible to an end-user. For example, when an end-user *Ua* in routing realm A wants to communicate with an end-user *Ub* in routing realm B, *Ua* has first to explicitly establish communication with the ALG that interconnects A and B, and only via that can *Ua* establish communication with *Ub*. We term such a gateway a "non-transparent" ALG.

Another form of ALG makes communication through the ALG transparent to an end user. Using the previous example, with a "transparent" ALG, *Ua* would not be required to establish explicit connectivity to the ALG first, before starting to communicate with *Ub*. Such connectivity will be established transparently to *Ua*, so that *Ua* would only see connectivity to *Ub*.

For completeness, note that it is not necessarily the case that communicating via an ALG involves changes to the network header. An ALG could be used only at the beginning of a session for the purpose of authentication, after which the ALG goes away and communication continues natively.

Both non-transparent and transparent ALGs are required (by definition) to understand the syntax and semantics of the application data stream. ALGs are very simple from the viewpoint of network layer architecture, since they appear as Internet hosts in each realm, i.e. they act as origination and termination points for communication.

At the other end of the spectrum is a Network Address Translator (NAT) [RFC 1631]. In the context of this document we define a NAT as a device that just modifies the network and the transport layer headers, but does not understand the syntax/semantics of the application layer data stream (using our terminology what is described in RFC1631 is a device that has both the NAT and ALG functionality).

In the standard case of a NAT placed between a corporate network using private addresses [RFC 1918] and the public Internet, that NAT changes the source IPv4 address in packets going towards the Internet, and changes the destination IPv4 address in packets coming from the Internet. When a NAT is used to interconnect routing realms with overlapping addresses, such as a direct connection between two intranets, the NAT may modify both addresses in the IP header. Since the NAT modifies address(es) in the IP header, the NAT also has to modify the transport (e.g., TCP, UDP) pseudo-header checksum. Upon some introspection one could observe that when interconnecting routing realms with overlapping addresses, the set of operations on the network and transport header performed by a NAT forms a (proper) subset of the set of operations on the network and transport layer performed by a transparent ALG.

By definition a NAT does not understand syntax and semantics of an application data stream. Therefore, a NAT cannot support applications that carry IP addresses at the application layer (e.g., FTP with PORT or PASV command [RFC 959]). On the other hand, a NAT can support any application, as long as such an application does not carry IP addresses at the application layer. This is in contrast with an ALG that can support only the applications coded into the ALG.

One can conclude that both NATs and ALGs have their own limitations, which could constrain their usefulness. Combining NAT and ALG functionality in a single device could be used to overcome some, but not all, of these limitations. Such a device would use the NAT functionality for the applications that do not carry IP addresses, and would resort to the ALG functionality when dealing with the applications that carry IP addresses. For example, such a device would use the NAT functionality to deal with the FTP data connection, but would use the ALG functionality to deal with the FTP control connection. However, such a device will fail completely handling an application that carries IP addresses, when the device does not support the application via the ALG functionality, but rather handles it via the NAT functionality.

Communicating through either ALGs or NATs involves changes to the network header (and specifically source and destination addresses), and to the transport header. Since IP Security authentication headers assume that the addresses in the network header are preserved end-to-end, it is not clear how one could support IP Security-based authentication between a pair of hosts communicating through either an ALG or a NAT. Since IP Security, when used for confidentiality, encrypts the entire transport layer end-to-end, it is not clear how an ALG or NAT could modify encrypted packets as they require to. In other words, both ALGs and NATs are likely to force a boundary between two distinct IP Security domains, both for authentication and for confidentiality, unless specific enhancements to IP Security are designed for this purpose.

Interconnecting routing realms via either ALGs or NATs relies on the DNS [RFC 1035]. Specifically, for a given set of (interconnected) routing realms, even if network layer addresses are no longer unique across the set, fully qualified domain names would need to be unique across the set. However, a site that is running a NAT or ALG probably needs to run two DNS servers, one inside and one outside the NAT or ALG, giving different answers to identical queries. This is discussed further in [kre]. DNS security [RFC 2065] and some dynamic DNS updates [dns2] will presumably not be valid across a NAT/ALG boundary, so we must assume that the external DNS server acquires at least part of its tables by some other mechanism.

To summarize, since RFC 1918, we have not really changed the spatial uniqueness of an address, so much as recognized that there are multiple spaces. i.e. each space is still a routing realm such as an intranet, possibly connected to other intranets, or the Internet, by NATs or ALGs (see above discussion). The temporal uniqueness of an address is unchanged by RFC 1918.

4.2. Addresses are no longer all temporally unique

Note that as soon as address significance changes anywhere in the address space, it has in some sense changed everywhere. This has in fact already happened.

IPv4 address blocks were for many years assigned chronologically, i.e. effectively at random with respect to network topology. This led to constantly growing routing tables; this does not scale. Today, hierarchical routing (CIDR [RFC 1518], [RFC 1519]) is used as a mechanism to improve scaling of routing within a routing realm, and especially within the Internet (The Annex goes into more details on CIDR).

Scaling capabilities of CIDR are based on the assumption that address allocation reflects network topology as much as possible, and boundaries for aggregation of addressing information are not required to be fully contained within a single organization - they may span multiple organizations (e.g., provider with its subscribers). Thus if a subscriber changes its provider, then to avoid injecting additional overhead in the Internet routing system, the subscriber may need to renumber.

Changing providers is just one possible reason for renumbering. The informational document [RFC 1900] shows why renumbering is an increasingly frequent event. Both DHCP [RFC 1541] and PPP [RFC 1661] promote the use of dynamic address allocation.

To summarize, since the development and deployment of DHCP and PPP, and since it is expected that renumbering is likely to become a common event, IP address significance has indeed been changed. Spatial uniqueness should be the same, so addresses are still effective locators. Temporal uniqueness is no longer assured. It may be quite short, possibly shorter than a TCP connection time. In such cases an IP address is no longer a good identifier. This has some impact on end-to-end security, and breaks TCP in its current form.

4.3. Multicast and Anycast

Since we deployed multicast [RFC 1112], we must separate the debate over meaning of IP addresses into meaning of source and destination addresses. A destination multicast address (i.e. a locator for a topologically spread group of hosts) can traverse a NAT, and is not necessarily restricted to an intranet (or to the public Internet). Its lifetime can be short too.

The concept of an anycast address is of an address that semantically locates any of a group of systems performing equivalent functions. There is no way such an address can be anything but a locator; it can never serve as an identifier as defined in this document, since it does not uniquely identify host. In this case, the effective temporal uniqueness, or useful lifetime, of an IP address can be less than the time taken to establish a TCP connection.

Here we have used TCP simply to illustrate the idea of an association - many UDP based applications (or other systems layered on IP) allocate state after receiving or sending a first packet, based on the source and/or destination. All are affected by absence of temporal uniqueness whereas only the routing infrastructure is affected by spatial uniqueness changes.

4.4. Summary

Due to dynamic address allocation and increasingly frequent network renumbering, temporal uniqueness of IPv4 addresses is no longer globally guaranteed, which puts their use as identifiers into severe question. Due to the proliferation of Intranets, spatial uniqueness is also no longer guaranteed across routing realms; interconnecting routing realms could be accomplished via either ALGs or NATs. In principle such interconnection will have less functionality than if those Intranets were directly connected. In practice the difference in functionality may or may not matter, depending on individual circumstances.

5. IPv6 Considerations

As far as temporal uniqueness (identifier-like behaviour) is concerned, the IPv6 model [RFC 1884] is very similar to the current state of the IPv4 model, only more so. IPv6 will provide mechanisms to autoconfigure IPv6 addresses on IPv6 hosts. Prefix changes, requiring the global IPv6 addresses of all hosts under a given prefix to change, are to be expected. Thus, IPv6 will amplify the existing problem of finding stable identifiers to be used for end-to-end security and for session bindings such as TCP state.

The IAB feels that this is unfortunate, and that the transition to IPv6 would be an ideal occasion to provide upper layer end-to-end protocols with temporally unique identifiers. The exact nature of these identifiers requires further study.

As far as spatial uniqueness (locator-like behaviour) is concerned, the IPv6 address space is so big that a shortage of addresses, requiring an RFC 1918-like approach and address translation, is hardly conceivable. Although there is no shortage of IPv6 addresses, there is also a well-defined mechanism for obtaining link-local and site-local addresses in IPv6 [RFC 1884, section 2.4.8]. These properties of IPv6 do not prevent separate routing realms for IPv6, if so desired (resulting in multiple security domains as well). While at the present moment we cannot identify a case in which multiple IPv6 routing realms would be required, it is also hard to give a definitive answer to whether there will be only one, or more than one IPv6 routing realms. If one hypothesises that there will be more than one IPv6 routing realm, then such realms could be interconnected together via ALGs and NATs. Considerations for such ALGs and NATs appear to be identical to those for IPv4.

ANNEX: Current Practices for IPv4 Address Allocation & Routing

Initially IP address structure and IP routing were designed around the notion of network number classes (Class A/B/C networks) [RFC 790]. In the earlier 90s growth of the Internet demanded significant improvements in both the scalability of the Internet routing system, as well as in the IP address space utilization. Classful structure of IP address space and associated with it classful routing turned out to be inadequate to meet the demands, so during 1992 - 1993 period the Internet adopted Classless Inter-Domain Routing (CIDR) [RFC 1380], [RFC 1518], [RFC 1519]. CIDR encompasses a new address allocation architecture, new routing protocols, and a new structure of IP addresses.

CIDR improves scalability of the Internet routing system by extending the notion of hierarchical routing beyond the level of individual subnets and networks, to allow routing information aggregation not only at the level of individual subnets and networks, but at the level of individual sites, as well as at the level of Internet Service Providers. Thus an organization (site) could act as an aggregator for all the destinations within the organization. Likewise, a provider could act as an aggregator for all the destinations within its subscribers (organizations directly connected to the provider).

Extending the notion of hierarchical routing to the level of individual sites and providers, and allowing sites and providers to act as aggregators of routing information, required changes both to the address allocation procedures, and to the routing protocols. While in pre-CIDR days address allocation to sites was done without taking into consideration the need to aggregate the addressing information above the level of an individual network numbers, CIDR-based allocation recommends that address allocation be done in such a way as to enable sites and providers to act as aggregators of addressing information - such allocation is called "aggregator based". To benefit from the "aggregator based" address allocation, CIDR introduces an inter-domain routing protocol (BGP-4) [RFC 1771, RFC 1772] that provides capabilities for routing information aggregation at the level of individual sites and providers.

CIDR improves address space utilization by eliminating the notion of network classes, and replacing it with the notion of contiguous variable size (power of 2) address blocks. This allows a better match between the amount of address space requested and the amount of address space allocated [RFC 1466]. It also facilitates "aggregator based" address allocation. Eliminating the notion of network classes requires new capabilities in the routing protocols (both intra and inter-domain), and IP forwarding. Specifically, the CIDR-capable

protocols are required to handle reachability (addressing) information expressed in terms of variable length address prefixes, and forwarding is required to implement the "longest match" algorithm. CIDR implications on routing protocols are described in [RFC 1817].

The scaling capabilities of CIDR are based on the assumption that address allocation reflects network topology as much as possible, especially at the level of sites, and their interconnection with providers, to enable sites and providers to act as aggregators. If a site changes its provider, then to avoid injecting additional overhead in the Internet routing system, the site may need to renumber. While CIDR does not require every site that changes its providers to renumber, it is important to stress that if none of the sites that change their providers will renumber, the Internet routing system might collapse due to the excessive amount of routing information it would need to handle.

Maintaining "aggregator based" address allocation (to promote scalable routing), and the need to support the ability of sites to change their providers (to promote competition) demands practical solutions for renumbering sites. The need to contain the overhead in a rapidly growing Internet routing system is likely to make renumbering more and more common [RFC 1900].

The need to scale the Internet routing system, and the use of CIDR as the primary mechanism for scaling, results in the evolution of address allocation and management policies for the Internet. This evolution results in adding the "address lending" policy as an alternative to the "address ownership" policy [RFC 2008].

IP addressing and routing have been in constant evolution since IP was first specified [RFC 791]. Some of the addressing and routing principles have been deprecated, some of the principles have been preserved, while new principles have been introduced. Current Internet routing and addresses (based on CIDR) is an evolutionary step that extends the use of hierarchy to maintain a routable global Internet.

Security Considerations

The impact of the IP addressing model on security is discussed in sections 4.1 and 5 of this document.

Acknowledgements

This document was developed in the IAB. Constructive comments were received from Ran Atkinson, Jim Bound, Matt Crawford, Tony Li, Michael A. Patton, Jeff Schiller. Earlier private communications from Noel Chiappa helped to clarify the concepts of locators and identifiers.

References

- [RFC 791] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.
- [RFC 790] Postel, J., "Assigned Numbers", September 1981.
- [RFC 959] Postel, J., and J. Reynolds, "File Transfer Protocol", STD 9, RFC 959, October 1985.
- [RFC 1035] Mockapetris, P., "Domain Names - Implementation and Specification", STD 13, RFC 1035, November 1987.
- [RFC 1112] Deering, S., "Host Extensions for IP Multicasting", STD 5, RFC 1112, September 1989.
- [RFC 1380] Gross, P., and P. Almquist, "IESG Deliberations on Routing and Addressing", RFC 1380, November 1992.
- [RFC 1466] Gerich, E., "Guidelines for Management of IP Address Space", RFC 1466, May 1993.
- [RFC 1498] Saltzer, J., "On the Naming and Binding of Network Destinations", RFC 1498, August 1993 (originally published 1982).
- [RFC 1518] Rekhter, Y., and T. Li, "An Architecture for IP Address Allocation with CIDR", RFC 1518, September 1993.
- [RFC 1519] Fuller, V., Li, T., Yu, J., and K. Varadhan, "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy", RFC 1519, September 1993.
- [RFC 1541] Droms, R., "Dynamic Host Configuration Protocol", RFC 1541, October 1993.
- [RFC 1661] Simpson, W., "The Point-to-Point Protocol (PPP)", STD 51, RFC 1661, July 1994.
- [RFC 1771] Rekhter, Y., and T. Li, "A Border Gateway Protocol 4 (BGP-4)", RFC 1771, March 1995.

[RFC 1772] Rekhter, Y., and P. Gross, "Application of the Border Gateway Protocol in the Internet", RFC 1772, March 1995.

[RFC 1817] Rekhter, Y., "CIDR and Classful Routing", RFC 1817, September 1995.

[RFC 1825] Atkinson, R., "Security Architecture for the Internet Protocol", RFC 1825, September 1995.

[RFC 1900] Carpenter, B., and Y. Rekhter, "Renumbering Needs Work", RFC 1900, February 1996.

[RFC 1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. J., and E. Lear, "Address Allocation for Private Internets", RFC 1918, February 1996.

[RFC 1933] Gilligan, R., and E. Nordmark, "Transition Mechanisms for IPv6 Hosts and Routers", RFC 1933, April 1996.

[RFC 2008] Rekhter, Y., and T. Li, "Implications of Various Address Allocation Policies for Internet Routing", RFC 2008, October 1996.

[kre] Elz, R., et. al., "Selection and Operation of Secondary DNS Servers", Work in Progress.

[RFC 2065] Eastlake, E., and C. Kaufman, "Domain Name System Security Extensions", RFC 2065, January 1997.

[dns2] Vixie, P., et. al., "Dynamic Updates in the Domain Name System (DNS UPDATE)", Work in Progress.

Authors' Addresses

Brian E. Carpenter
Computing and Networks Division
CERN
European Laboratory for Particle Physics
1211 Geneva 23, Switzerland

EMail: brian@dxcoms.cern.ch

Jon Crowcroft
Dept. of Computer Science
University College London
London WC1E 6BT, UK

EMail: j.crowcroft@cs.ucl.ac.uk

Yakov Rekhter
Cisco systems
170 West Tasman Drive
San Jose, CA, USA

Phone: +1 914 528 0090
Fax: +1 408 526-4952
EMail: yakov@cisco.com

