Network Working Group
Request for Comments: 1918
Obsoletes: 1627, 1597
BCP: 5
Category: Best Current Practice

Y. Rekhter
Cisco Systems
B. Moskowitz
Chrysler Corp.
D. Karrenberg
RIPE NCC
G. J. de Groot
RIPE NCC
E. Lear
Silicon Graphics, Inc.
February 1996

Address Allocation for Private Internets

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.

1. Introduction

For the purposes of this document, an enterprise is an entity autonomously operating a network using TCP/IP and in particular determining the addressing plan and address assignments within that network.

This document describes address allocation for private internets. The allocation permits full network layer connectivity among all hosts inside an enterprise as well as among all public hosts of different enterprises. The cost of using private internet address space is the potentially costly effort to renumber hosts and networks between public and private.

2. Motivation

With the proliferation of TCP/IP technology worldwide, including outside the Internet itself, an increasing number of non-connected enterprises use this technology and its addressing capabilities for sole intra-enterprise communications, without any intention to ever directly connect to other enterprises or the Internet itself.

The Internet has grown beyond anyone's expectations. Sustained exponential growth continues to introduce new challenges. One challenge is a concern within the community that globally unique address space will be exhausted. A separate and far more pressing concern is that the amount of routing overhead will grow beyond the capabilities of Internet Service Providers. Efforts are in progress within the community to find long term solutions to both of these problems. Meanwhile it is necessary to revisit address allocation procedures, and their impact on the Internet routing system.

To contain growth of routing overhead, an Internet Provider obtains a block of address space from an address registry, and then assigns to its customers addresses from within that block based on each customer requirement. The result of this process is that routes to many customers will be aggregated together, and will appear to other providers as a single route [RFC1518], [RFC1519]. In order for route aggregation to be effective, Internet providers encourage customers joining their network to use the provider's block, and thus renumber their computers. Such encouragement may become a requirement in the future.

With the current size of the Internet and its growth rate it is no longer realistic to assume that by virtue of acquiring globally unique IP addresses out of an Internet registry an organization that acquires such addresses would have Internet-wide IP connectivity once the organization gets connected to the Internet. To the contrary, it is quite likely that when the organization would connect to the Internet to achieve Internet-wide IP connectivity the organization would need to change IP addresses (renumber) all of its public hosts (hosts that require Internet-wide IP connectivity), regardless of whether the addresses used by the organization initially were globally unique or not.

It has been typical to assign globally unique addresses to all hosts that use TCP/IP. In order to extend the life of the IPv4 address space, address registries are requiring more justification than ever before, making it harder for organizations to acquire additional address space [RFC1466].

Hosts within enterprises that use IP can be partitioned into three categories:

Category 1: hosts that do not require access to hosts in other enterprises or the Internet at large; hosts within this category may use IP addresses that are unambiguous within an enterprise, but may be ambiguous between enterprises.

Category 2: hosts that need access to a limited set of outside services (e.g., E-mail, FTP, netnews, remote login) which can be handled by mediating gateways (e.g., application layer gateways). For many hosts in this category an unrestricted external access (provided via IP connectivity) may be unnecessary and even undesirable for privacy/security reasons. Just like hosts within the first category, such hosts may use IP addresses that are unambiguous within an enterprise, but may be ambiguous between enterprises.

Category 3: hosts that need network layer access outside the enterprise (provided via IP connectivity); hosts in the last category require IP addresses that are globally unambiguous.

We will refer to the hosts in the first and second categories as "private". We will refer to the hosts in the third category as "public".

Many applications require connectivity only within one enterprise and do not need external (outside the enterprise) connectivity for the majority of internal hosts. In larger enterprises it is often easy to identify a substantial number of hosts using TCP/IP that do not need network layer connectivity outside the enterprise.

Some examples, where external connectivity might not be required, are:

- A large airport which has its arrival/departure displays individually addressable via TCP/IP. It is very unlikely that these displays need to be directly accessible from other networks.

- Large organizations like banks and retail chains are switching to TCP/IP for their internal communication. Large numbers of local workstations like cash registers, money machines, and equipment at clerical positions rarely need to have such connectivity.

- For security reasons, many enterprises use application layer gateways to connect their internal network to the Internet. The internal network usually does not have direct access to the Internet, thus only one or more gateways are visible from the Internet. In this case, the internal network can use non-unique IP network numbers.

- Interfaces of routers on an internal network usually do not need to be directly accessible from outside the enterprise.

3. Private Address Space

The Internet Assigned Numbers Authority (IANA) has reserved the following three blocks of the IP address space for private internets:

```
10.0.0.0        -   10.255.255.255  (10/8 prefix)
172.16.0.0      -   172.31.255.255  (172.16/12 prefix)
192.168.0.0     -   192.168.255.255 (192.168/16 prefix)
```

We will refer to the first block as "24-bit block", the second as "20-bit block", and to the third as "16-bit" block. Note that (in pre-CIDR notation) the first block is nothing but a single class A network number, while the second block is a set of 16 contiguous class B network numbers, and third block is a set of 256 contiguous class C network numbers.

An enterprise that decides to use IP addresses out of the address space defined in this document can do so without any coordination with IANA or an Internet registry. The address space can thus be used by many enterprises. Addresses within this private address space will only be unique within the enterprise, or the set of enterprises which choose to cooperate over this space so they may communicate with each other in their own private internet.

As before, any enterprise that needs globally unique address space is required to obtain such addresses from an Internet registry. An enterprise that requests IP addresses for its external connectivity will never be assigned addresses from the blocks defined above.

In order to use private address space, an enterprise needs to determine which hosts do not need to have network layer connectivity outside the enterprise in the foreseeable future and thus could be classified as private. Such hosts will use the private address space defined above. Private hosts can communicate with all other hosts inside the enterprise, both public and private. However, they cannot have IP connectivity to any host outside of the enterprise. While not having external (outside of the enterprise) IP connectivity private hosts can still have access to external services via mediating gateways (e.g., application layer gateways).

All other hosts will be public and will use globally unique address space assigned by an Internet Registry. Public hosts can communicate with other hosts inside the enterprise both public and private and can have IP connectivity to public hosts outside the enterprise. Public hosts do not have connectivity to private hosts of other enterprises.

Moving a host from private to public or vice versa involves a change of IP address, changes to the appropriate DNS entries, and changes to configuration files on other hosts that reference the host by IP address.

Because private addresses have no global meaning, routing information about private networks shall not be propagated on inter-enterprise links, and packets with private source or destination addresses should not be forwarded across such links. Routers in networks not using private address space, especially those of Internet service providers, are expected to be configured to reject (filter out) routing information about private networks. If such a router receives such information the rejection shall not be treated as a routing protocol error.

Indirect references to such addresses should be contained within the enterprise. Prominent examples of such references are DNS Resource Records and other information referring to internal private addresses. In particular, Internet service providers should take measures to prevent such leakage.

4. Advantages and Disadvantages of Using Private Address Space

The obvious advantage of using private address space for the Internet at large is to conserve the globally unique address space by not using it where global uniqueness is not required.

Enterprises themselves also enjoy a number of benefits from their usage of private address space: They gain a lot of flexibility in network design by having more address space at their disposal than they could obtain from the globally unique pool. This enables operationally and administratively convenient addressing schemes as well as easier growth paths.

For a variety of reasons the Internet has already encountered situations where an enterprise that has not been connected to the Internet had used IP address space for its hosts without getting this space assigned from the IANA. In some cases this address space had been already assigned to other enterprises. If such an enterprise would later connects to the Internet, this could potentially create very serious problems, as IP routing cannot provide correct operations in presence of ambiguous addressing. Although in principle Internet Service Providers should guard against such mistakes through the use of route filters, this does not always happen in practice. Using private address space provides a safe choice for such enterprises, avoiding clashes once outside connectivity is needed.

A major drawback to the use of private address space is that it may actually reduce an enterprise's flexibility to access the Internet. Once one commits to using a private address, one is committing to renumber part or all of an enterprise, should one decide to provide IP connectivity between that part (or all of the enterprise) and the Internet. Usually the cost of renumbering can be measured by counting the number of hosts that have to transition from private to public. As was discussed earlier, however, even if a network uses globally unique addresses, it may still have to renumber in order to acquire Internet-wide IP connectivity.

Another drawback to the use of private address space is that it may require renumbering when merging several private internets into a single private internet. If we review the examples we list in Section 2, we note that companies tend to merge. If such companies prior to the merge maintained their uncoordinated internets using private address space, then if after the merge these private internets would be combined into a single private internet, some addresses within the combined private internet may not be unique. As a result, hosts with these addresses would need to be renumbered.

The cost of renumbering may well be mitigated by development and deployment of tools that facilitate renumbering (e.g. Dynamic Host Configuration Protocol (DHCP)). When deciding whether to use private addresses, we recommend to inquire computer and software vendors about availability of such tools. A separate IETF effort (PIER Working Group) is pursuing full documentation of the requirements and procedures for renumbering.

5. Operational Considerations

One possible strategy is to design the private part of the network first and use private address space for all internal links. Then plan public subnets at the locations needed and design the external connectivity.

This design does not need to be fixed permanently. If a group of one or more hosts requires to change their status (from private to public or vice versa) later, this can be accomplished by renumbering only the hosts involved, and changing physical connectivity, if needed. In locations where such changes can be foreseen (machine rooms, etc.), it is advisable to configure separate physical media for public and private subnets to facilitate such changes. In order to avoid major network disruptions, it is advisable to group hosts with similar connectivity needs on their own subnets.

If a suitable subnetting scheme can be designed and is supported by the equipment concerned, it is advisable to use the 24-bit block (class A network) of private address space and make an addressing plan with a good growth path. If subnetting is a problem, the 16-bit block (class C networks), or the 20-bit block (class B networks) of private address space can be used.

One might be tempted to have both public and private addresses on the same physical medium. While this is possible, there are pitfalls to such a design (note that the pitfalls have nothing to do with the use of private addresses, but are due to the presence of multiple IP subnets on a common Data Link subnetwork). We advise caution when proceeding in this area.

It is strongly recommended that routers which connect enterprises to external networks are set up with appropriate packet and routing filters at both ends of the link in order to prevent packet and routing information leakage. An enterprise should also filter any private networks from inbound routing information in order to protect itself from ambiguous routing situations which can occur if routes to the private address space point outside the enterprise.

It is possible for two sites, who both coordinate their private address space, to communicate with each other over a public network. To do so they must use some method of encapsulation at their borders to a public network, thus keeping their private addresses private.

If two (or more) organizations follow the address allocation specified in this document and then later wish to establish IP connectivity with each other, then there is a risk that address uniqueness would be violated. To minimize the risk it is strongly recommended that an organization using private IP addresses choose randomly from the reserved pool of private addresses, when allocating sub-blocks for its internal allocation.

If an enterprise uses the private address space, or a mix of private and public address spaces, then DNS clients outside of the enterprise should not see addresses in the private address space used by the enterprise, since these addresses would be ambiguous. One way to ensure this is to run two authority servers for each DNS zone containing both publically and privately addressed hosts. One server would be visible from the public address space and would contain only the subset of the enterprise's addresses which were reachable using public addresses. The other server would be reachable only from the private network and would contain the full set of data, including the private addresses and whatever public addresses are reachable the private network. In order to ensure consistency, both servers should be configured from the same data of which the publically visible zone only contains a filtered version. There is certain degree of additional complexity associated with providing these capabilities.

6. Security Considerations

Security issues are not addressed in this memo.

7. Conclusion

With the described scheme many large enterprises will need only a relatively small block of addresses from the globally unique IP address space. The Internet at large benefits through conservation of globally unique address space which will effectively lengthen the lifetime of the IP address space. The enterprises benefit from the increased flexibility provided by a relatively large private address space. However, use of private addressing requires that an organization renumber part or all of its enterprise network, as its connectivity requirements change over time.

8. Acknowledgments

We would like to thank Tony Bates (MCI), Jordan Becker (ANS), Hans-Werner Braun (SDSC), Ross Callon (BayNetworks), John Curran (BBN Planet), Vince Fuller (BBN Planet), Tony Li (cisco Systems), Anne Lord (RIPE NCC), Milo Medin (NSI), Marten Terpstra (BayNetworks), Geza Turchanyi (RIPE NCC), Christophe Wolfhugel (Pasteur Institute), Andy Linton (connect.com.au), Brian Carpenter (CERN), Randy Bush (PSG), Erik Fair (Apple Computer), Dave Crocker (Brandenburg Consulting), Tom Kessler (SGI), Dave Piscitello (Core Competence), Matt Crawford (FNAL), Michael Patton (BBN), and Paul Vixie (Internet Software Consortium) for their review and constructive comments.

9. References

[RFC1466] Gerich, E., "Guidelines for Management of IP Address Space", RFC 1466, Merit Network, Inc., May 1993.

[RFC1518] Rekhter, Y., and T. Li, "An Architecture for IP Address Allocation with CIDR", RFC 1518, September 1993.

[RFC1519] Fuller, V., Li, T., Yu, J., and K. Varadhan, "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy", RFC 1519, September 1993.

10. Authors' Addresses

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

Robert G Moskowitz
Chrysler Corporation
CIMS: 424-73-00
25999 Lawrence Ave
Center Line, MI 48015
Phone: +1 810 758 8212
Fax: +1 810 758 8173
EMail: rgm3@is.chrysler.com

Daniel Karrenberg
RIPE Network Coordination Centre
Kruislaan 409
1098 SJ Amsterdam, the Netherlands
Phone: +31 20 592 5065
Fax: +31 20 592 5090
EMail: Daniel.Karrenberg@ripe.net

Geert Jan de Groot
RIPE Network Coordination Centre
Kruislaan 409
1098 SJ Amsterdam, the Netherlands
Phone: +31 20 592 5065
Fax: +31 20 592 5090
EMail: GeertJan.deGroot@ripe.net

Eliot Lear
Mail Stop 15-730
Silicon Graphics, Inc.
2011 N. Shoreline Blvd.
Mountain View, CA 94043-1389
Phone: +1 415 960 1980
Fax: +1 415 961 9584
EMail: lear@sgi.com