

Network Working Group
Request for Comments: 1597
Category: Informational

T.J. Watson Research Center, IBM Corp.
Y. Rekhter
B. Moskowitz
Chrysler Corp.
D. Karrenberg
RIPE NCC
G. de Groot
RIPE NCC
March 1994

Address Allocation for Private Internets

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.

1. Introduction

This RFC describes methods to preserve IP address space by not allocating globally unique IP addresses to hosts private to an enterprise while still permitting full network layer connectivity between all hosts inside an enterprise as well as between all public hosts of different enterprises. The authors hope, that using these methods, significant savings can be made on allocating IP address space.

For the purposes of this memo, 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.

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 current practice is to assign globally unique addresses to all hosts that use TCP/IP. There is a growing concern that the finite IP address space might become exhausted. Therefore, the guidelines for assigning IP address space have been tightened in recent years [1]. These rules are often more conservative than enterprises would like, in order to implement and operate their networks.

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

- hosts that do not require access to hosts in other enterprises or the Internet at large;
- hosts that need access to a limited set of outside services (e.g., E-mail, FTP, netnews, remote login) which can be handled by application layer gateways;
- hosts that need network layer access outside the enterprise (provided via IP connectivity);
- hosts within the first category may use IP addresses that are unambiguous within an enterprise, but may be ambiguous between enterprises.

For many hosts in the second 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.

Only hosts in the last category require IP addresses that are globally unambiguous.

Many applications require connectivity only within one enterprise and do not even need external 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 organisations 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 (e.g., firewalls) 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 firewall hosts are visible from the Internet. In this case, the internal network can use non-unique IP numbers.
- If two enterprises communicate over their own private link, usually only a very limited set of hosts is mutually reachable from the other enterprise over this link. Only those hosts need globally unique IP 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 networks:

| | | |
|-------------|---|-----------------|
| 10.0.0.0 | - | 10.255.255.255 |
| 172.16.0.0 | - | 172.31.255.255 |
| 192.168.0.0 | - | 192.168.255.255 |

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 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 255 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.

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. Such hosts will be called private hosts, and 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 external host. While not having external network

layer connectivity private hosts can still have access to external services via application layer relays.

All other hosts will be called 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 external public hosts. 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.

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. When such an enterprise later connects to the Internet, it could potentially create very

serious problems, as IP routing cannot provide correct operations in presence of ambiguous addressing. Using private address space provides a safe choice for such enterprises, avoiding clashes once outside connectivity is needed.

One could argue that the potential need for renumbering represents a significant drawback of using the addresses out of the block allocated for private internets. However, we need to observe that the need is only "potential", since many hosts may never move into the third category, and an enterprise may never decide to interconnect (at IP level) with another enterprise.

But even if renumbering has to happen, we have to observe that with Classless Inter-Domain Routing (CIDR) an enterprise that is connected to the Internet may be encouraged to renumber its public hosts, as it changes its Network Service Providers. Thus renumbering is likely to happen more often in the future, regardless of whether an enterprise does or does not use the addresses out of the block allocated for private networks. Tools to facilitate renumbering (e.g., DHCP) would certainly make it less of a concern.

Also observe that the clear division of public and private hosts and the resulting need to renumber makes uncontrolled outside connectivity more difficult, so to some extent the need to renumber could be viewed as an advantage.

5. Operational Considerations

A recommended 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 is not fixed permanently. If a number of hosts require to change status later this can be accomplished by renumbering only the hosts involved and installing another physical subnet if required.

If a suitable subnetting scheme can be designed and is supported by the equipment concerned, it is advisable to use the 24-bit block of private address space and make an addressing plan with a good growth path. If subnetting is a problem, the 16-bit class C block, which consists of 255 contiguous class C network numbers, can be used.

Using multiple IP (sub)nets on the same physical medium has many pitfalls. We recommend to avoid it unless the operational problems are well understood and it is proven that all equipment supports this properly.

Moving a single host between private and public status will involve a change of address and in most cases physical connectivity. In locations where such changes can be foreseen (machine rooms etc.) it may be advisable to configure separate physical media for public and private subnets to facilitate such changes.

Changing the status of all hosts on a whole (sub)network can be done easily and without disruption for the enterprise network as a whole. Consequently it is advisable to group hosts whose connectivity needs might undergo similar changes in the future on their own subnets.

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.

Groups of organisations which foresee a big need for mutual communication can consider forming an enterprise by designing a common addressing plan supported by the necessary organisational arrangements like a registry.

If two sites of the same enterprise need to be connected using an external service provider, they can consider using an IP tunnel to prevent packet leaks from the private network.

A possible approach to avoid leaking of DNS RRs is to run two nameservers, one external server authoritative for all globally unique IP addresses of the enterprise and one internal nameserver authoritative for all IP addresses of the enterprise, both public and private. In order to ensure consistency both these servers should be configured from the same data of which the external nameserver only receives a filtered version.

The resolvers on all internal hosts, both public and private, query only the internal nameserver. The external server resolves queries from resolvers outside the enterprise and is linked into the global DNS. The internal server forwards all queries for information outside the enterprise to the external nameserver, so all internal hosts can access the global DNS. This ensures that information about private hosts does not reach resolvers and nameservers outside the enterprise.

6. References

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

7. Security Considerations

While using private address space can improve security, it is not a substitute for dedicated security measures.

8. 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.

9. Acknowledgments

We would like to thank Tony Bates (RIPE NCC), Jordan Becker (ANS), Hans-Werner Braun (SDSC), Ross Callon (Wellfleet), John Curran (NEARNET), Vince Fuller (Barrnet), Tony Li (cisco Systems), Anne Lord (RIPE NCC), Milo Medin (NSI), Marten Terpstra (RIPE NCC), and Geza Turchanyi (RIPE NCC) for their review and constructive comments.

10. Authors' Addresses

Yakov Rekhter
T.J. Watson Research Center, IBM Corp.
P.O. Box 218
Yorktown Heights, NY, 10598

Phone: +1 914 945 3896
Fax: +1 914 945 2141
EMail: yakov@watson.ibm.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: 3858921@mcimail.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