

Network Working Group
Request for Comments: 1335

Z. Wang
J. Crowcroft
University College London
May 1992

A Two-Tier Address Structure for the Internet:
A Solution to the Problem of Address Space Exhaustion

Status of this Memo

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

Abstract

This RFC presents a solution to problem of address space exhaustion in the Internet. It proposes a two-tier address structure for the Internet. This is an "idea" paper and discussion is strongly encouraged.

Introduction

Address space exhaustion is one of the most serious and immediate problems that the Internet faces today [1,2]. The current Internet address space is 32-bit. Each Internet address is divided into two parts: a network portion and a host portion. This division corresponds the three primary Internet address classes: Class A, Class B and Class C. Table 1 lists the network number statistics as of April 1992.

	Total	Allocated	Allocated (%)
Class A	126	48	54%
Class B	16383	7006	43%
Class C	2097151	40724	2%

Table 1: Network Number Statistics (April 1992)

If recent trends of exponential growth continue, the network numbers in Class B will soon run out [1,2]. There are over 2 million Class C network numbers and only 2% have been allocated. However, a Class C network number can only accommodate 254 host numbers which is too small for most networks. With the rapid expansion of the Internet and drastic increase in personal computers, the time when the 32-bit address space is exhausted altogether is also not too distant [1-3].

Recently several proposals have been put forward to deal with the

immediate problem [1-4]. The Supernetting and C-sharp schemes attempt to make the Class C numbers more usable by re-defining the way in which Class C network numbers are classified and assigned [3,4]. Both schemes require modifications to the exterior routing algorithms and global coordination across the Internet may be required for the deployment. The two schemes do not expand the total number of addresses available to the Internet and therefore can only be used as a short-term fix for next two or three years. Schemes have also been put forward in which the 32-bit address field is replaced with a field of the same size but with different meaning and the gateways on the boundary re-write the address when the packet crossed the boundary [1,2,5]. Such schemes, however, requires substantial changes to the gateways and the exterior routing algorithm.

In this paper, we present an alternative solution to the problem of address space exhaustion. The "Dual Network Addressing (DNA)" scheme proposed here is based on a two-tier address structure and sharing of addresses. It requires no modifications to the exterior routing algorithms and any networks can adopt the scheme individually at any time without affecting other networks.

The Scheme

The DNA scheme attempts to reduce the waste in using the Internet addresses. A useful analogy to our scheme is the extension system used in the telephone system. Many large organizations usually have extensive private telephone networks for internal use and at the mean time hire a limited number of external lines for communications with the outside world. In such a telephone system, important offices may have direct external lines and telephones in the public areas may be restricted to internal calls only. The majority of the telephones can usually make both internal calls and external calls. But they must share a limited number of external lines. When an external call is being made, a pre-defined digit has to be pressed so that an external line can be allocated from the pool of external lines.

In the DNA scheme, there are two types of Internet addresses: Internal addresses and External addresses. An internal address is an Internet address only used within one network and is unique only within that network. An interface with an internal address can only communicate with another interface with an internal address in the same network. An external address is unique in the entire Internet and an interface with an external address can communicate directly to another interface with an external address over the Internet. All current Internet addresses are external addresses.

In effect, the external addresses form one global Internet and the

internal addresses form many private Internets. Within one network, the external addresses are only used for inter-network communications and internal addresses for intra-network communications. An External Address Sharing Service (EASS) is needed to manage the sharing of external addresses. An EASS server reserves a number of external addresses. When a machine that only has an internal address wants to communicate a machine with an external address in other networks, it can send a request to an EASS server to obtain a temporary external address. After the use, the machine can return the external address to the EASS server.

We believe that, with the DNA scheme, a network can operate with a limited number of external addresses. The reasons are as follows:

- * In most networks, the majority of the traffic is confined to its local area networks. This is due the nature of networking applications and the bandwidth constraints on inter-network links.
- * The number of machines which act as Internet servers, i.e., running programs waiting to be called by machines in other networks, is often limited and certainly much smaller than the total number of machines. These machines include mail servers, domain name servers, ftp archive servers, directory servers, etc.
- * There are an increasingly large number of personal machines entering the Internet. The use of these machines is primarily limited to their local environment. They may also be used as "clients" such as ftp and telnet to access other machines.
- * For security reasons, many large organizations, such as banks, government departments, military institution and some companies, may only allow a very limited number of their machines to have access to the global Internet. The majority of their machines are purely for internal use.

In the DNA scheme, all machines in a network are assigned a permanent internal address and can communicate with any machines within the same network. The allocation of external addresses depends on the functions of the machines and as a result it creates three-level privileges:

- * machines which act as servers or used as central computing infrastructure are likely to have frequent communications with other networks therefore they may require external addresses all the time. These machines are allocated

permanent external addresses.

- * machines which are not allowed to communicate with other networks have no external addresses and can only communicate with machines within their own network.
- * the rest of the machines share a number of external addresses. The external addresses are allocated by the EASS server on request. These machines can only be used as clients to call machines in other networks, i.e., they can not be called by machines in other networks.

A network can choose any network number other than its external network number as its internal network number. Different networks can use the same network number as their internal number. We propose to reserve one Class A network number as the well-known network number for internal use.

The Advantages

The DNA scheme attempts to tackle the problem from the bottom of the Internet, i.e., each individual network, while other schemes described in the first section deal with the problem from the top of the Internet, i.e., gateways and exterior routing algorithms. These schemes, however, do not need to be considered as mutually exclusive. The DNA scheme has several advantages:

- * The DNA scheme takes an evolutionary approach towards the changes. Different networks can individually choose to adopt the scheme at any time only when necessary. There is no need for global coordination between different networks for their deployment. The effects of the deployment are confined to the network in which the scheme is being implemented, and are invisible to exterior routing algorithms and external networks.
- * With the DNA scheme, it is possible for a medium size organization to use a Class C network number with 254 external addresses. The scheme allows the current Internet to expand to over 2 million networks and each network to have more than 16 million hosts. This will allow considerable time for a long-term solution to be developed and fully tested.
- * The DNA scheme requires modifications to the host software. However, the modifications are needed only in those networks which adopt the DNA scheme. Since all existing Class A and B networks usually have sufficient external addresses for all their machines, they do not need to adopt the DNA scheme, and therefore

need no modifications at all to their software. The networks which need to use the DNA scheme are those new networks which are set up after the Class A and B numbers run out and have to use a Class C number.

- * The DNA scheme makes it possible to develop to a new addressing scheme without expanding the 32-bit address length to 64-bit. With the two-tier address structure, the current 32-bit space can accommodate over 4 billion hosts in the global Internet and 100 million hosts in each individual network. When we move to a classless multi-hierarchic addressing scheme, the use of external addresses can be more efficient and less wasteful and the 32-bit space can be adequate for the external addresses.
- * When a new addressing scheme has been developed, all current Internet addresses have to be changed. The DNA scheme will make such a undertaking much easier and smoother, since only the EASS servers and those have permanent external addresses will be affected, and communications within the network will not be interrupted.

The Modifications

The major modifications to the host software is in the network interface code. The DNA scheme requires each machine to have at least two addresses. But most of the host software currently does not allow us to bind two addresses to one physical interface. This problem can be solved by using two network interfaces on each machine. But this option is too expensive. Note the two interfaces are actually connected to the same physical network. Therefore, if we modify the interface code to allow two logical interfaces to be mapped onto one single physical interface, the machine can then use both the external address and the internal address with one physical interface as if it has two physical interfaces. In effect, two logical IP networks operate over the same physical network.

The DNA scheme also has implications to the DNS service. Many machines will have two entries in the local name server. The DNS server must examine the source address of the request and decide which entry to use. If the source address matches the well-known internal network number, it passes the internal address of the domain name. Otherwise, the name server passes the external address.

An EASS server is required to manage the sharing of the external addresses, i.e., to allocate and de-allocate external addresses to the machines which do not have permanent external addresses. This service can be provided by using the "Dynamic Host Configuration Protocol (DHCP)" [6].

Many hosts do an inverse lookup of incoming connections. Therefore, it is desirable the entry in the DNS server be updated whenever a new external address is allocated. This will also allow an machine which currently has a temporary external address to be called by other machines. The updating of the entry in the DNS server can be done more easily if the EASS server and DNS server are co-located.

Acknowledgements

We would like to thank J. K. Reynolds for the network statistics, and V. Cerf, C. Topolcic, K. McCloghrie, R. Ullmann and K. Carlberg for their useful comments and discussion.

References

- [1] Chiappa, N., "The IP Addressing Issue", work in progress, October 1990.
- [2] Clark, D., Chapin, L., Cerf, V., Braden, R., and R. Hobby, "Towards the Future Architecture", RFC 1287, MIT, BBN, CNRI, ISI, UC Davis, December 1991.
- [3] Solensky, F., and F. Kastenholz, "A Revision to IP Address Classifications", work in progress, March 1992.
- [4] Fuller, V., Li, T., Yu, J., and K. Varadhan, "Supernetting: an Address Assignment and Aggregation Strategy", work in progress, March 1992.
- [5] Tsuchiya, P., "The IP Network Address Translator", work in progress, March 1991.
- [6] Droms, R., "Dynamic Host Configuration Protocol", work in progress, March 1992.

Security Considerations

Security issues are not discussed in this memo.

Authors' Addresses

Zheng Wang
Dept. of Computer Science
University College London
London WC1E 6BT, UK

EMail: z.wang@cs.ucl.ac.uk

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

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

