Is IPv4 Sufficient for Another 30 Years?

October 7, 2004

Abstract

TCP/IP was developed 30 years ago. It has been successful over the past 30 years, until recently when its limitation started emerging.

Recent years, the Internet grew more rapidly than any one might have imagined. Its rapid growth brings the IPv4's address shortage issue into picture. IPv4 uses only a 32-bit address space, and more than half of its addresses has already been assigned. People started to discuss: is IPv4 sufficient for another 30 years?

Our answer to the question is no – the current IP protocol is reaching its limit. It will run out of its address space in a decade or so.

In this paper, we argue that the Internet is now growing exponentially, and it will keep growing in that way for at least a decade. Because of that, IPv4 is running out of its addresses in a decade or so. We also argue that NAT is not a solution to this issue, instead, it is just a temporary workaround due to its shortcomings.

Finally, we briefly introduced the solution to this issue – IPv6.

I Introduction

IP is responsible for moving packets of data from host to host. It forwards each packet based on a destination address – the IP address.

The current version of the Internet Protocol, IPv4, uses 32-bit addresses to identify hosts connected to the Internet. 32-bit address space was supposed to be sufficient and it has been successful for the past 30 years, because theoretically, it can support four billion addresses.

But due to a history of inefficient address assignments, half of these addresses were already assigned. Only a fraction of the remaining addresses can be used because of shortcomings in the original IPv4 design.

On the other hand, the Internet is growing with an exponential speed. Comparing to this rapid growth, the IPv4 address space seems to be very limited, especially for countries other than the United States. Usually only a small portion of the address space is available for them to use.

Some ISPs are using Network Address Translation (NAT) technology to deal with the IP address problem. We argue that, however, this is only a short term workaround. Due to the shortcomings in its structure, many internet-based services are not available to users behind NATs. Hence NAT should not be considered as a solution to the IP address shortage problem.

II Brief History of TCP/IP and The Growth of The Internet

The history of the Internet goes back to the 60's.

In 1969, the first network, ARPANET, was constructed, linking only 4 nodes: University of California at Los Angeles, SRI (in Stanford), University of California at Santa Barbara, and University of Utah.

The development of TCP/IP started in 1973. From then on, people started to use IP address to identify hosts on the network.

In 1983, every host connected to the ARPANET was required to use TCP/IP. At that time, there were around 600 hosts. Since then on, TCP/IP started its domination of the Internet.

Over the next 14 years, the number grew by 600. But that pattern does not apply to today's Internet at all. In January 2004, there were 233,101,481 hosts[1]; while In July 2004, the number became 285,139,107. So for the year 2004, 600 is averagely just the increase of the number of hosts in 3 hours! [1].

Figure 1 shows the growth of the number of hosts connected to the Internet over the past 10 years [1]. The data table of Figure 1 is displayed in Table 1.

[1] The survey defines a "host" as a host that they could reach. This implies that hosts behind NAT are not counted in.

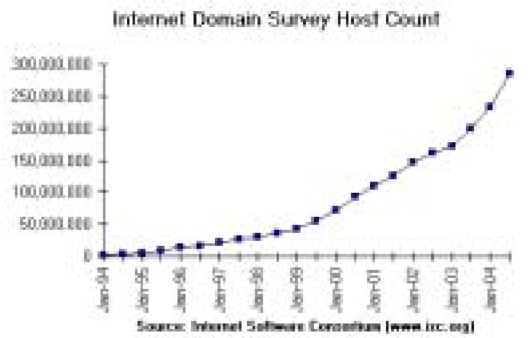

Figure 1 The growth of the Internet over the past 10 years

Year	# of Hosts	Year	# of Hosts
1993	1,313,000	1999	43,230,000
1994	2,217,000	2000	72,398,092
1995	5,846,000	2001	109,574,429
1996	14,352,000	2002	147,344,723
1997	21,819,000	2003	171,638,297
1998	29,670,000	2004	233,101,481

Table 1 data table of Figure 1

We can see from Figure 1 that the growth seems to be in an exponential manner. In fact it is, as we will see in section IV. This exponential growth pattern has been followed for at least the past 8 years.

III The Argument

IPv4 uses only 32 bits for IP address space, which allows only 4 billion hosts to be identified on the Internet. 4 billion sounds like a large number; on the other hand, however, it is even much less than the human population on the earth. Therefore, it is reasonable to believe that the IPv4 addresses will be used up eventually. The question is when.

Around 1992, the IETF became aware of a global shortage of IPv4 addresses. A large amount of discussion are going on in the Internet development community. Some people believe that IPv4 will be sufficient for at least another 30 years; while some others think that IPv4 is dying and are working on developing a new generation of IP protocol to replace IPv4.

Is IPv4 sufficient for another 30 years? Our answer to the question is NO. We believe that the Internet has been, is, and will be growing exponentially in terms of the number of hosts, thus IPv4 addresses will be used up in a decade or so. We know that there are some short term workarounds at work; however, we believe that those are not only just short term workarounds, but are also obstacles in the Internet development.

IV Counterclaims

Some people believe that IPv4 will be sufficient for another 30 years, because we have only used half of the IPv4 address space. They think the addresses being used so fast is partly because there were inefficient address assignment in the earlier years of the Internet history. By restricting the address assignment policy, the rest of the IP address space, which has about 2 billion unassigned addresses, should be able to last for another 30 years.

We do not agree with this argument. First of all, this argument ignores the fact that the Internet is growing in an exponential rate. We've used more than half of the address space for the past 30 years while the Internet was growing relatively slow. It is very likely that we use up all the addresses within a much shorter time because the Internet will be growing much faster than the past 30 years.

Secondly, there's always inefficiency in IP address assignment because the addresses are assigned in a "paging" fashion rather than a "segmentation" fashion. This is the same side effect as we all see in operating system's memory management.

Another argument agrees that IPv4 itself will not be able to survive under the exponentially increasing demand of IP addresses. But because we have alternative technologies that allow a large number of privately addressed hosts to be connected to the Internet using a small amount of public IP addresses, the pressure of lacking of public IP addresses could thus be relieved. NAT, being widely used under such circumstances, is a good example.

We do not agree with this argument either. Although NAT could be used as a temporary workaround, it should not be considered as a solution due to its born shortcomings. Servers behind NAT are not accessible from the outside network. Hosts behind different NAT cannot establish direct connections. NAT is not scalable. The internal network has to be small, or the NAT router will become a bottleneck in the network traffic.

These shortcomings especially the first two are actually obstructing the development of the Internet.

V IPv4 Address Space Issue

1. Internet is growing exponentially

As we have mentioned in earlier sections, the number of hosts in the Internet is growing exponentially. Here we want to convince you that its growth over the past few years was nothing but exponential.

To make this characteristic easier to see, we take logarithm on the number of hosts and the result is shown in Table 2 and in Figure 2 by pink line. By applying the Least Squares Fitting method on the data for the years from 1996 to 2004, we have

$$h = 10^{0.15567\times(y-1996)+7.18511},$$

where y stands for the year and h means the number of hosts connected to the Internet. This function is shown in Figure 2 by the dark blue line.

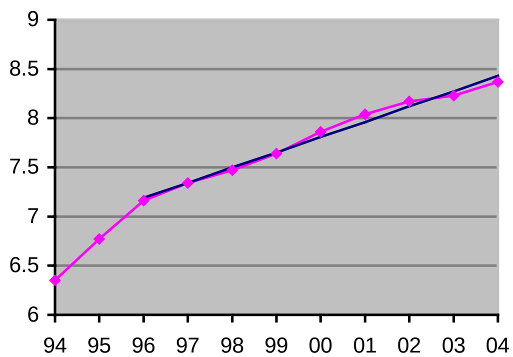

Figure 2 Logarithm of the number of hosts for the past 10 years

Year	# of Hosts	Log10	Year	# of Hosts	Log10
1993	1,313,000	6.12	1999	43,230,000	7.64
1994	2,217,000	6.35	2000	72,398,092	7.86
1995	5,846,000	6.77	2001	109,574,429	8.04
1996	14,352,000	7.16	2002	147,344,723	8.17
1997	21,819,000	7.34	2003	171,638,297	8.23
1998	29,670,000	7.47	2004	233,101,481	8.37

Table 2 data table of Figure 2

The actual data points are all very close to this fitted line. This tells us that this exponential function perfectly describes the characteristic of the growth of the Internet for the past 8 years.

2. In the near future, Internet will keep growing exponentially

We believe that the Internet will keep expanding with the current growth pattern.

The industry of the Internet is still in the rapid expanding stage. It will keep in this stage for at lease another decade. More and more new technologies will still keep coming out each with a huge potential market value. This industry will only be growing faster and faster, rather than being stable or going down.

In countries that are young to the Internet, this industry is just started. The potential market is big enough to drive the Internet to keep growing for at least a decade.

While the first name server was created in the US in 1983, the first root name server was just installed in the Mainland China in October 2003.

This implies that China's Internet is still in a very early stage of development. Plus knowing that China's population is 7 times that of the US, we believe that China's Internet development has unpredictable potential.

3. We are running out of IPv4 addresses in a decade or so

With our assumption that this pattern will be followed in the near future, we could be able to figure out an upper bound of years after which we will be out of IPv4 address. We know that the theoretical address space of IPv4 is 2^{32}. However, the actual number of available addresses is smaller because of the way that the addresses are separated into classes, and because some addresses are set aside for multicasting, testing or other special uses. The actual number is somewhere between 3.2 and 3.3 billion.

By solving

$$10^{0.15567\times(y-1996)+7.18511} = 3.3\times10^{9},$$

we get this result: $y \cong 14.99$. This means if the assumption holds, we will be running out of IPv4 address in 15 years.

This falsifies the counterclaim and supports our claim by giving out an upper bound, 15 years, which is only half the time that is claimed in the argument, 30 years.

In fact, the above result is just a very loose upper bound. This is because it is also assuming that the IP addresses are perfectly efficiently assigned and any of the remaining addresses could be assigned to anywhere in this world, which are not true in reality.

3.1 In the US and North America area

The United States owns the source of the IP addresses. The Internet is rooted in the US – it was first built up in the US and its backbones and root servers were spread to the whole world from the US. The IANA (Internet Assigned Numbers Authority) is located in the US.

Because of these historical reasons, the US has much more available IP addresses than other countries. Once there are reserved addresses in the IANA registry, IP address requests from within the US could be satisfied.

Despite this "unlimited" [2] resource the US has, some experts believe that the US will run out of IP address in a decade.

We agree with this argument. It is accordant to our result, 15 years of upper bound. Further more, by taking into consideration the inefficiency in the IP addresses assignment, we believe this argument is true.

3.2 In other countries

Other countries are not as lucky as the US. For example, MIT has more IP addresses within its university system than the entire nation of China.

The reason of this is because the APNIC (Asia Pacific Network Information Center) only has limited [3] IP address resources available for its users.

"Recent years, China has experienced rapid IP address increase. China has a certain amount of IP addresses. However, the IP resource that China has still does not satisfy the needs of the development of the Internet service providers. Along with the increase of the Internet usage, this kind of mismatching will be more obvious and will become an obstacle in the Internet development in China." [2]

The IP address shortage has already started emerging in some countries such as China. End users in those country are already feeling the pinch.

ISPs have their workaround, NAT, to deal with this problem. However, NAT should not be considered as a solution.

VI Problems of NAT

Network Address Translation (NAT) is one of the technologies that allow computers in a private network to access the outside network, the Internet, without having registered IP addresses. It uses one set of IP addresses for internal traffic and another set of addresses for external traffic.

An NAT router maintains a mapping from every live TCP session to its corresponding internal host. When a privately addressed host sends a packet through an NAT router, the source address is converted to a registered IP address before the pocket is passes onto the Internet. When a packet sent from an external host reaches an NAT router, the NAT router forwards the packet to the

[2] comparatively to other countries

[3] the latest data shows APNIC current only has an address space of thirteen /8 addresses.

corresponding internal host according to the mapping.

Although NAT has been commonly employed at places that require more IP addresses than they are assigned to solve the IP address shortage problem, our argument is that NAT should only be considered as a temporary workaround to the IPv4 address shortage problem.

1. Servers behind NAT are not accessible from the outside

Because of the NAT mechanism, internal IP addresses are not exposed to the outside. Therefore it is usually impossible to run a server behind NAT, if the server is supposed to provide services clients outside the NAT.

There is a possible way to this. Use static NAT to associate a port on the NAT router to the appropriate port on the server. However, first, if there two servers both need to open the same port, one of them has to give up. Secondly, NAT is usually run by the ISP. For security and other reasons, ISPs usually does not open ports for end users.

This shortcoming is crucial to companies providing Internet-based services. For those companies, their servers must be publicly addressed in order to allow users to be able to access their services.

2. Hosts behind different NATs cannot establish direct connections

To establish a direct connection, there must be an initiator. When both of the two hosts are behind different NATs, none of them could know the other one's address. Thus none of them could be able to act the role of initiator. Namely direct connection is impossible to be established in this scenario.

Recent years, online chatting software, such as MSN Messenger, has become very popular. Friends use them to not only chat by exchanging plain text information, they could also exchange files through the software.

If both or one of the two ends are outside of any NAT, they can exchange both text data and binary data without a problem. However, if both of them are behind NAT and are not in the same LAN, the file data exchange functionality breaks.

This is because while text information could have the server, which is publicly accessible, as an intermediator, the binary data cannot due to the designer's performance concern.

A lot of users in countries lacking of IP addresses are experiencing this problem. File transfer is one of the problems due to the fact that direct connections cannot be established between two hosts that are both behind different NATs.

Due to this problem, a whole bunch of Internet-based applications will not be able to function as expected when running on hosts behind NAT. This includes online chatting software, online games, remote desktop systems, etc.

3. NAT scalability problem

There is a lot of work done in the NAT router while packets are going through it.

When an NAT receives an outgoing packet, it first needs to pick a free TCP port and create a TCP session between this local port and the target port on the destination host, if such a session is not present. Then it needs to create a new entry in the mapping table to associate the picked port on the NAT and the source port on the sender, an internal host. Finally, it converts the source address in the packet into the NATs public IP address and transfers it from the internal NIC (network interface card) to the external NIC in order to send the packet out onto the Internet.

When an NAT receives an incoming packet, it looks up in the mapping table to figure out which internal host is the destination and which port on that host is targeted. It overwrites the destination address to the host's internal address. Then it transfers the packet from the external NIC to the internal NIC to send the packet onto the internal network.

Although these work seem not to be time consuming on a regular computer, they make NAT difficult to scale. When having big number of internal hosts, these work will make NAT, as the only bridge connecting the inside and the outside, to become a bottleneck.

VII Solution – IPv6

The IPng (IP next generation) project was started more than 10 years ago. The goal was to work out the next generation of IP protocol which solves the IPv4 issues including the address shortage problem. After a large amount of discussion, around 1995, IPv6 was picked as the final IPng proposal.

Besides deploying more recent technologies, IPv6 also solves the address space shortage issue that IPv4 has by allowing 128 bit for IP address space, and cures the backbone routing table growth issue which has been a big headache to IPSs and backbone operators.

VIII Conclusion

The rapid growth of the Internet makes the address space shortage of IPv4 to happen earlier than its designers had expected. Some people argue that IPv4 will be sufficient for another 30 years. We believe that the address space shortage issue will kick IPv4 out of the stage in a decade or so. Although a workaround, NAT, seems to relieve this pressure on IPv4, we believe that the correct way of solving this problem is to bring the next generation of IP protocol, IPv6, into the picture.

References

[1] ISC (Internet Systems Consortium), "ISC Internet Domain Survey", http://www.isc.org/ops/ds/

[2] CNNIC (China Internet Network Information Center), China Internet Development and Usage Statistics, http://www.blogchina.com/idea/cnnic-14th-internet-report/001.htm