Translation from RIETI Discussion Paper Series 01-J-006 (January 2002)

Is IPv6 Necessary?

IKEDA, Nobuo*
YAMADA, Hajime**

Synopsis

The "e-Japan Priority Policy Program" announced by the Japanese government has named, as one of its policy targets, the promotion of Internet Protocol version 6 (IPv6). The reason, it states, is because the IP addresses (IPv4) currently being used will soon run out. The authors believe this assumption to be very dubious. Actually, more than half of the addresses remain unused, and only three percent of all addresses have been used through connections with the Internet. It is likely that 15 years from now, some segments of addresses will still remain unused, and the effective use of unused addresses will enable almost unlimited use. A virtually infinite number of addresses can be used under the current IPv4, and it is unlikely that we will face a shortage. Also, in terms of functions, no important application has emerged to date that can be used only under IPv6, and it is doubtful that IPv6 is an essential technological innovation to deal with issues peculiar to the age of broadband communications. It is meaningless to rush to commercialize IPv6 under the current circumstances. Moreover, hasty attempts by the national government to promote it could hinder the healthy development of the Internet.

Acknowledgement

The authors are grateful for the valuable comments and criticisms of Takashi Arano, Masahiro Kusunoki, Hiroshi Fujiwara and Hiroyuki Morikawa, as well as participants in the Nikkei Digital Core mailing lists, and also to Tomoe Kiyosada for her cooperation in the surveys. It should be noted that the persons named above do not necessarily entirely support the views of the authors, and the contents and opinions stated in this paper are solely those of the authors, and do not necessarily reflect the official views of the Research Institute of Economy, Trade and Industry.

* Senior Researcher, Research Institute of Economy, Trade and Industry (an Independent Administrative Institute)
** Professor, Center for Global Communication, International University of Japan

Introduction

In a policy speech delivered during the extraordinary Diet session in November 2000, Prime Minister Yoshiro Mori surprised Japanese citizens by announcing that, "the Japanese government will actively pursue solutions for global issues on the Internet, and will work in the hope of making a significant global contribution to the development of the Internet, through IPv6." His speech was one of the reasons why Japanese people became greatly interested in Internet Protocol version 6 (IPv6), something that had been known among only a few people in the Internet community. In response, the "e-Japan Priority Policy Program" of the national government set out a policy that predicted, "By 2005, an Internet environment will have been built that will allow all Japanese people to obtain, process and transmit information they desire safely, promptly and easily, wherever they are."[1] Moreover, a Council to Promote and Advance IPv6 was formed with members from the public and private sectors. The IPv6 budget amounts to eight billion yen in the amended fiscal 2001 budget and 2.2 billion yen in the initial budget for fiscal 2002.

Proponents of IPv6 insist that Internet addresses are being consumed rapidly as always-on connections and mobile phones become popular, and in the age of Net home appliances that is to arrive in the near future, when all home appliances will be connected to the Internet, an enormous number of addresses will be needed. Thus, they argue, the IPv6 is the most powerful weapon for recovery for the Japanese IT industry in Japan. Mr. Mori's statement of beliefs shown on the home page of the Prime Minister's Office includes a note, which says, "Internet addresses have already run out in countries other than Japan, the United States, European nations and Australia, and it has become impossible to connect to the Internet in other countries."

But are IPv4 addresses really running out? Is IPv6 really a protocol that will deliver unprecedented functions, not yet available in the Internet today? And will the world adopt IPv6 and accept it as a new world standard? From our perspective, it seems that debate on these fundamental issues concerning IPv6 has been neglected in Japan, and instead the nationalistic argument that "the U.S. enjoyed an exclusive victory with IPv4, so Japan should strike back with IPv6" is being raised. It seems that the national government is virtually forcing people to adopt IPv6 by saying, "The national government will provide subsidies for IPv6-compatible technologies only." This paper considers the real situation concerning IP addresses, and looks at the necessity of IPv6 based on objective data.

---

[1] http://www.kantei.go.jp/jp/it/network/dai3/3siryou42.html

1. What Are IP Addresses?

A Transfer Control Protocol/ Internet Protocol (TCP/IP) packet consists of a header and the main data portion, and routers of each site read the header and transfer it to the nearest site of the recipient. IPv4 addresses are designated using 32 bits, and are usually indicated using decimal numbers with separations at every eight-bit section, like 202.238.95.24. An address consists of a network address for identifying LAN and a host address that identifies a terminal, and addresses are managed in layers according to classes. In a Class A address 043.***.***.***, for example, the first eight bit section is used as the network address and the remaining 24 bit sections can be used entirely in the same network, so there are 2^{24}, or about 16.7 million addresses available. Similarly, Class B has 16-bit (about 65,000) and Class C has eight-bit (256) addresses. IP addresses are allocated by the Internet Assigned Number Authority (IANA), which controls worldwide allocation to three Regional Internet Registries (RIRs) in North and South America (ARIN), the Asia Pacific (APNIC) and Europe (RIPE). The addresses are then distributed to Internet Service Provider (ISP) and other users via Network Information Centers (NIC) in each country. The router of each site has a routing table (path database) to clarify which route to take to send data to the designated address, and sends data packets accordingly. Since it is impossible for all routers to have routing tables for all addresses worldwide, addresses are assigned to a site after collecting them by class as far as possible, and each site of a class takes care of its share of the routing. It is similar to making a phone call from Tokyo to Osaka, where an operator on the Tokyo side makes a connection to the Osaka area code 06, and an operator on the Osaka side makes the remaining connection within the city of Osaka. Currently, there are 2^{32}, or about 4.3 billion, IP addresses available. The figure was deemed sufficient when the Internet first came into existence in around 1970. The number of users has increased rapidly since then, and IPv6 was proposed to the Internet Engineering Task Force (IETF) in 1992 to expand the address space, and earned official approval in 1998 (as RFC1883). Since an IPv6 address comprises 128 bits, an almost infinite number of addresses, or 2^{128} (3.4×10^{38}), would be available. Recently, a few attempts to support IPv6 have been made, as exemplified by Windows XP, which supports IPv6 and offers free upgrades of router software (IOS) to IPv6 by Cisco Systems. However, both IPv4 and IPv6 need to be supported to ensure compatibility, because IPv4 sites can be viewed from IPv6 sites but not vice versa. Further, advanced routers are needed to activate IPv6, and site facilities need to be completely renovated. As such, only about 1,000 sites support IPv6 worldwide[2], and commercial services have yet to get underway.

2. Are We Really Running Out of IP Addresses?

According to the Japanese government, the major reason to promote IPv6 is because the consumption of IPv4 addresses has grown rapidly, and we will soon run out of them. This assumption is based on an estimate by ICANN, which predicts, "We will run out of IP addresses by 2008." (ICANN 2001) The estimate was derived by assuming that the number of remaining addresses as of 2000 was about 1.7 billion and demand for new IP addresses will be 75 million in 2000, and moreover that demand for IP addresses will increase in a geometric progression after 2001. Based on these assumptions, the addresses would be depleting by 2008 if demand grows by a factor of 1.3 each year, and by 2006 if it grows by a factor of 1.5. (Figure 1)

Figure 1: Estimates for IP Address Consumption (ICANN)

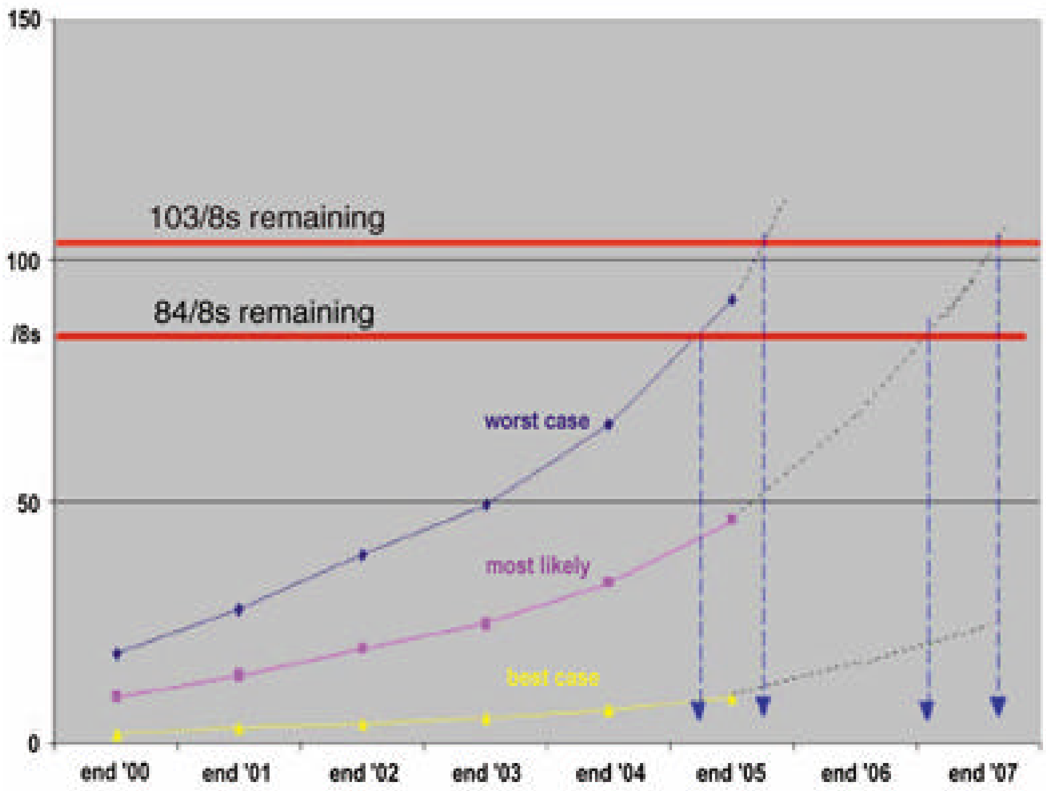

2 6bone Registry Status Report, http://whois.6bone.net

Figure 2: Long-term Tendencies in IP Address Allocation (ARIN)

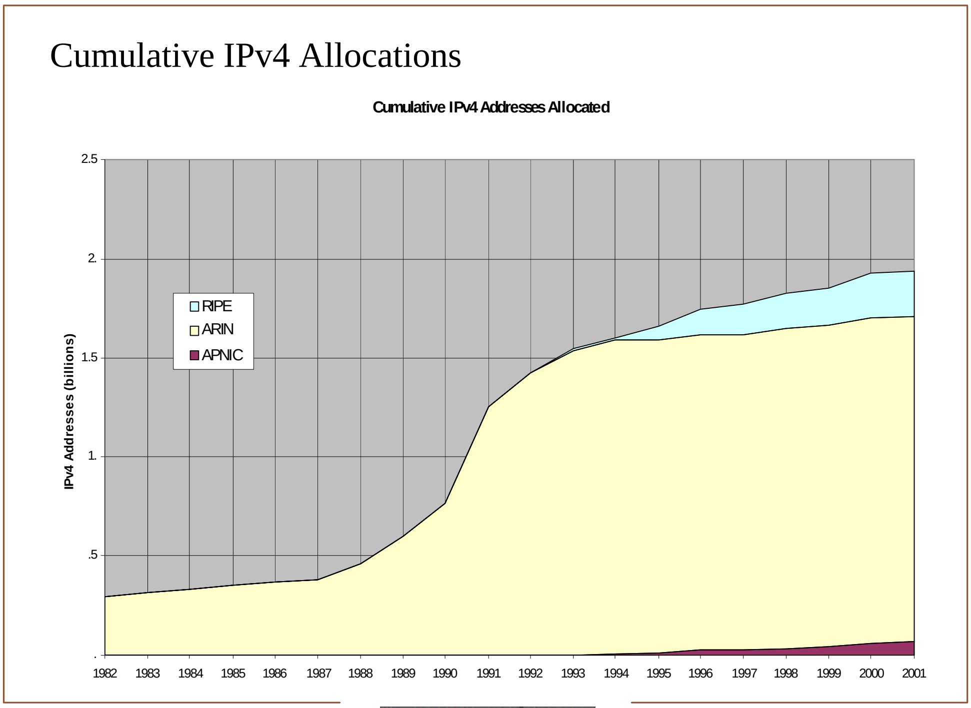

www.genuity.com

On the other hand, statistics taken by the American Registry of Internet Numbers (ARIN) (Marcus 2001) state that only about 1.9 billion IP addresses have been used, and we still have about 2.4 billion remaining. (Figure 2) Consumption has remained fairly stable at about 50 million each year. According to Scott Marcus, an ARIN Board Member who conducted the survey, the “shortage will become an issue only after 2010, at the earliest.”

Why do the figures differ so greatly? One questionable factor in the ICANN estimates is that they assume only 103 blocks (about 1.7 billion) of Class A addresses remain today. This was the number of addresses remaining in IANA as of 2000, but in reality there seems to be considerable “stock” of addresses allocated to each RIR and NIC, which have yet to be assigned to users. On the other hand, statistics taken by ARIN have counted the number of addresses actually assigned by RIR to NIC, so its estimate (for the number of remaining addresses, 2.4 billion) seems more reliable. If other stocks at NIC and ISP are included, the number of remaining addresses is more than that shown by ARIN’s data.

If we also take a look at estimates for future consumption, ICANN has estimated future consumption by assuming that it will grow in geometric progression, or by 1.3 times or 1.8 times each year. It is true that worldwide, IP address consumption rose rapidly between 1998 and 2000, as shown in Table 1. In 2000 in particular, the increase compared to the previous figure was a factor of 1.8. If IP address consumption is assumed to grow at the same rate of 1.8 times each year, we would run out of addresses (1.7 billion) in about three years.

But that two-year period discussed represented extraordinary years that marked a peak in the worldwide Internet bubble, and therefore should not be used as a long-term standard. In fact, consumption in 2001 fell below that of the previous year for RIPE, and growth has also slowed for APNIC, notwithstanding the rapid increase in the number of Internet users in China and Korea. In addition, the world growth rate has fallen to 11 percent. For JPNIC in Japan, too, consumption has remained stable in the most recent five-year period, at about two million each year.

Table 1: IP Address Allocation Worldwide (Total of RIRs; unit: 1 million)

	APNIC	RIPE	ARIN	WORD
1996	9	12	17	38
97	5	10	17	32
98	5	14	13	32
99	10	16	21	47
2000	21	28	34	83
01	29	25	38	92

It is impossible to forecast future trends based on the fluctuating figures shown in Table 1, and ICANN's estimates, which assumed growth in a geometric progression similar to that observed during the peak years, seem rather baseless. Even if the assumption were accepted, it will take at least nine years to consume the remaining 2.4 billion addresses. If the growth in 2001 (about 10 million) were to remain the same in the years to come, 15 years would be needed to use up the remaining addresses. If the growth rate decreases gradually and consumption is maintained at 92 million each year (the actual figure for 2001), then the number of remaining addresses would last for 26 years. (Figure 3) Therefore, IPv4 addresses should remain available for another 15 to 20 years, if the growth rate is assumed to remain constant at the average to date. Each RIR has the power to control the actual allocation, so it can restrict allocation if a shortage became a real problem. As such, it is unlikely that the rate of growth would remain the same each year and all the remaining addresses would be used up any time soon.

Figure 3: When Would We Run Out of IP Addresses?

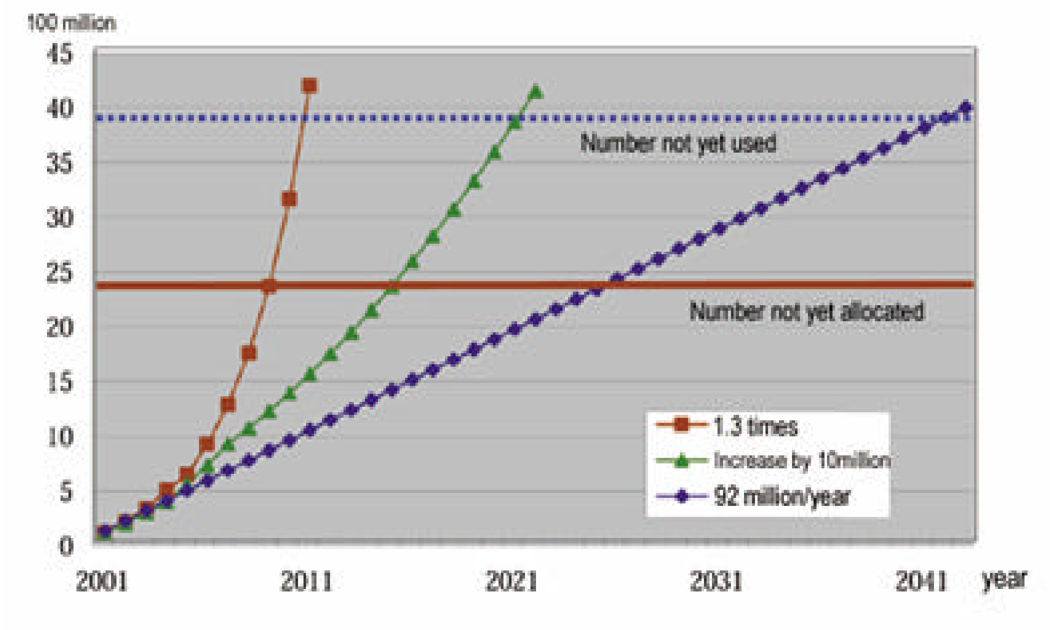

3. Inequality of Allocation

The real problem, however, is not the number of remaining addresses. To date, 1.9 billion IP addresses have been allocated, but the number of computers (hosts) connected to the Internet is estimated to be 130 million worldwide[3]. In other words, only three percent of the 4.3 billion addresses have actually been used. Naturally, it is impossible to use all of the remaining addresses, because of class structures. However, it is clear that we do not need to worry about an IP address shortage. The problem is, why has an enormous number of 1.9 billion IP addresses been allocated only to 130 million computer units?

As can be seen from Figure 2, about 1.5 billion addresses were allocated before 1995, and about 400 million after that time and up to 2001. On the other hand, the world's Internet population, which was 30 million in 1995, grew rapidly to over 400 million in 2001. In other words, while recent users have been given less than one address, as many as 50 addresses had been allocated to each person during the early years of the Internet.

This inequality is the reason for the apparent address shortage.

---

[3] "Host" refers to computers and devices that have IP addresses. (The Internet does not differentiate between servers and clients.) The statistics are obtained by counting host names, so there are other devices without names. However, the figure is likely to be close to the actual number, because many computers have several names.

The Internet is said to be an autonomous and dispersed network. Its address allocation system is similar to that of phone numbers, and addresses are centrally controlled and assigned by IANA to each RIR, and then to users. Addresses have been structured into classes, making partial changes difficult. As such, many addresses remain unused and are wasted. In the early days, a site which initially required a Class B number of addresses (about 65,000) was suddenly given a Class A number of addresses (about 16.7 million) as soon as the required number of addresses exceeded that of the Class B. Larger blocks of addresses were assigned as necessary, without confirming whether such a huge number of addresses were really necessary. (Hori, et al. 2001: p. 74) As a result, the U.S. Department of Defense alone owns six blocks of Class A addresses (about 100 million), and Bolt Beranek & Newman (BBN), a data service corporation which designed ARPANET (the predecessor of the Internet) owns three blocks of Class A addresses (about 50 million). The Massachusetts Institute of Technology (MIT) owns one block of Class A addresses, which is more than the total number of addresses allocated to China. (See the Appendix.) Currently, 74 percent of all the IP addresses in the world are being allocated to users in the Americas, while the number in Asia represents nine percent.

Under these circumstances, a report came out in 1992, asserting that, "IP addresses will be used up in two years if no measure is taken," and standardization procedures for IPv6 got underway. But while the standardization of IPv6 has fallen behind schedule, it has now become possible to allocate IP addresses for each subnet, in units of bits, in accordance with the volume consumed[4]. Inter-class routing by Classless Inter-Domain Routing (CIDR) is now also possible. Technologies like Network Address Translation (NAT) and Dynamic Host Configuration Protocol (DHCP) have improved the efficiency of IP address use, and when IPv6 was finally completed in 1998, many people doubted whether it would really prove useful.

So what should have been done before calling for IPv6 is to survey the situation of IP address use, and have users return their unused addresses. At the same time, the size of block to be allocated should be altered to suit the situation of use. Assuming maximum losses due to class structure and that 90 percent of the remaining addresses (about 3.9 billion) are usable, there would be a sufficient number of addresses to last 11 years even under the unrealistic assumptions made by ICANN, which simply extended the needs during the bubble years. If a more common-sense assumption is adopted,

---

[4] Subnet is a division of a class in units of bits. For example, if a unit of 20 bits or less is used as a host address, it is represented as /20. Recent allocations by the JPNIC and others has been in units of /20.

there are sufficient addresses to last for more than 20 years. (Figure 3) In the world of the Internet, 20 years is a veritable eon. In short, the lives of currently available addresses could be extended almost infinitely if proper allocation policies were to be adopted.

In fact, Class A addresses (043) possessed by the WIDE project in Japan was entrusted to the IPv6 Council, and Stanford University has returned its Class A addresses (036) to IANA. (See the Appendix.) When addresses are altered and segmented, the volume of routing tables grows, lowering the search efficiency. But this is merely an issue concerning the capacity of the hard disks of routers and CPU processing capability. The issue may have been more serious during the early years of the Internet, when computer capacities were much more limited. In recent years, however, hard disks with terabyte capacities have become commonplace. It is now senseless to waste addresses to reduce the size of routing tables.

Users do not own IP addresses; the RIR manages them. Instead of hoping that users voluntarily return unused addresses, a rule should be established to mandate the return.

Recently, allocation to ARIN is said to have been restricted. If allocation to the United States is stopped, the reuse of unused addresses is likely to be encouraged. Another idea is to charge a “commission” for address use, because it is difficult to confirm future plans even though the number of hosts currently used is known. If a dollar were charged for an address, for example, the total for Class A addresses would exceed two billion yen, so it is likely that the owner would return unused addresses. The return can be offset if the returned portion is purchased at a price commensurate with the commission.

The major reason why Japanese users await IPv6 is because the allocation by JPNIC has been done in a bureaucratic manner. The organization has over-governed users by checking every detail of their future business plans, and many users have felt they cannot use addresses as they wish. Currently, “commissions” are charged irregularly, but users should be charged according to the volume of use, including users who already have addresses. The address shortage has been an issue because IP addresses have become a rare resource, so allocation should be made according to a pricing mechanism. Of course, using this method will accrue costs, but it can be ignored when compared with the enormous cost necessary to alter all addresses in the world to make them IPv6-compliant. Only professionals who wish to communicate end-to-end (E2E) should possess fixed addresses, and normal users can communicate using servers. So 4.3 billion addresses should be sufficient. For more than 40 years, predictions have been made that petroleum reserves would be depleted "in 40 years." But the apocalyptic prophecies have never proven true. The reason is simply that nobody has taken the functions of the pricing mechanism into consideration. The "shortage" of IP addresses is not an issue of technology, but primarily one of economics that concerns allocation efficiency.

4. Do Address Spaces Need To Be Enlarged?

Aren't 4.3 billion global addresses sufficient? The ICANN report has cited, as reasons for the rapid increase in address consumption, the expected increase in the number of mobile phone users who connect to the Internet and also of always-on connection users who use DSL and cable Internet. In the report, demand is estimated "without considering savings effects by NAT and DHCP." Rather, it is a rather rough estimate. Even though DSL provides a always-on connection, not all users are connected to the Internet all the time. Normally, private addresses are allocated automatically by the DHCP when users log in, and are disconnected after a certain time lapses, thus saving address consumption.

Private addresses are similar to "extensions," which each user may use freely within LAN, among Class B addresses that correspond to 192.168.***.***. Some 2^{16}, or about 65,000 such addresses are available, sufficient even for very large networks. Theoretically, all global address can be extended into private addresses, so we already have an address space with a size of 4.3 billion times 65,000, or about 280 trillion. Moreover, these addresses can also be used after nesting, so the address space of IPv4 is virtually infinite in size.

However, private addresses can be used only within LAN, so they need to be translated into global addresses by NAT when accessing the Internet.

NAT is similar to PBX, which translates private addresses into global addresses when making external connections. Even when there is large number of phone units for extension, the necessary number of external lines is only the number sufficient for transferring simultaneous calls. Similarly, global addresses are necessary only in a number sufficient for simultaneous accesses, if NAT were used[5].

In short, the necessary number of addresses is determined not by the number of devices, but by the number of accesses. Therefore, the number of global addresses required would be only 10 to 20 percent of the number of users, even when DSL and

---

[5] If technologies called Network Address Port Translation (MAPT) or IP Masquerade were used in addition, one address may be supplemented with a port number and then divided into sections, to enable use by multiple devices.

cable Internet were used. NICs worldwide have recently allocated addresses on the assumption that private addresses will be used, so even the LAN of large companies have been given only several blocks of Class C addresses (involving several hundred addresses). This is why address consumption remained constant during the late 1990s, even as the number of Internet users increased rapidly.

So how many addresses will ultimately be needed? There are now about 130 million hosts connected to the Internet and about 400 million users worldwide. The number of users is likely to top one billion in the near future. However, if private addresses were used, the number of simultaneous accesses could never exceed the total number of users, so the number of global addresses required would equal the number of accesses, or one per person at most. That figure corresponds to the current number of actual allocations. In other words, 4.3 billion addresses would be sufficient if addresses were allocated efficiently, unless as many as 70 percent of the 6.3 billion people on the entire Earth tried to connect to the Internet simultaneously.

5. What Can IPv6 Do?

Even if we have enough IP addresses, it would be meaningful to shift to IPv6 if it delivered functions that were not conventionally available. The e-Japan Priority Policy Program asserts that, "The shift to an Internet equipped with IPv6 should be promoted for improved protection of privacy and security." It has also named other new functions enabled by IPv6, such as multicasting, real-time control and plug-and-play. But all of these functions are available with IPv4.

The security function of IPv6 is called IPsec, but IPv4 also has enabled this[6]. Information security also concerns measures to prevent altering of website content and concentrated accesses to a server, and means to track such invaders. However, these are not included in the scope of IPv6. For multicasting, a test service called MBone has already started with IPv4. For label switching that supports voice communication, streaming and other real-time communications, in which packet priorities are designated using labels, Multi-Protocol Label Switching (MPLS) has already become an international standard.

Plug-and-play, which sets addresses automatically, has also been possible with DHCP. The e-Japan Priority Policy Program assumes that for "net home appliances," where

[6] IPsec utilizes two methods: verification and encryption. Verification is used to confirm whether the sender has really sent the received data, or whether the data have been modified during transmission. For encryption, the sender and recipient need to decide in advance which encryption method to use, and data are sent after encryption based on the set rules.

refrigerators, microwave ovens and other home appliances are all connected to the Internet, each appliance needs to be allocated an IP address. The “IPv6 boom” took offer after Sony President Nobuyuki Idei announced that Sony strongly supports IPv6 and “all Sony products hereafter will be made IPv6-compliant.” However, even if each net appliance did require an address, DHCP can make infinite number of private addresses available even now. In short, there is no decisive killer application that is enabled only with IPv6, and it is significant only for the reason that IPv6 does not use NAT.

NAT has indeed caused many troubles in network operation. All communications are made through NAT, causing the loads to become concentrated. When a conversion table fails, all communications are disconnected. If digital signatures and other encryption methods are used for security reasons, NAT destroys encryptions when it converts the address. Since private addresses cannot be directly accessed from outside, it would be difficult to identify the other party when IP telephone and other bi-directional means are used. However, these problems only concern network managers, and can be solved with IPv4.

In Realm Specific Internet Protocol (RSIP), for which standardization procedures are underway, servers “rent” global addresses for a certain period to clients. If this technology were used, addresses can be allocated dynamically while preserving the E2E architecture. For IP telephones, a protocol called Media Gateway Control Protocol (MGCP) is already in use, which records private addresses used by users on NAT, and connects to the address when a call is made from outside. In the Internet, element technologies have been turned into modules that can be improved separately, so most of the functions can be realized by combining these modules.

An oft-discussed dream to be realized by IPv6 is the “ubiquitous network,” in which every computer, home appliance and other device worldwide is assigned an IP address and operated peer-to-peer (P2P) in a highly dispersed manner, without a server. But a worldwide network that connects many appliances directly is highly unstable in terms of security. When global addresses have actually been used, frequent invasions have been observed because each appliance can be targeted directly from any place in the world. Someone may break into your house by opening the automatic lock through the Internet, and refrigerators that are constantly supplied with electricity may be used as a “stepping stone” for network crimes.

Problems such as spamming have barely been controlled on the server side, but much more serious damage may result if each computer terminal were to be directly connected worldwide. Besides, it may be hard to take measures against these problems.

IPv6 can provide excellent security and privacy protection only when it is used skillfully, and if users are not thoroughly adept in computer troubleshooting, not only users but also other people in society may suffer harm. In an Internet age in which refrigerators and other appliances are networked, many users might specify their names or birthdays as passwords. The Internet appliances should be protected by a firewall, and this kind of server-client formation is possible even now.

The only significance of IPv6 worthy of consideration is the “philosophical” argument that a structure like NAT, in which ISP controls the network, is against the E2E design concept of the Internet, in which end users control the network. (Lessig, 2001: p. 171) If NAT were used in 1990, the World Wide Web (WWW) may not have been born. The greatest value of the Internet is that it enables users to fully leverage their creativity, without any license, as communication service providers. If technological innovations in the future were headed toward highly dispersed applications like P2P, the current IP architecture that needs to pass through servers each time may hinder the Internet from developing into highly dispersed and serial computers. Such arguments are worthy of note, but they would not appeal to general users unless a powerful P2P application that takes full advantage of the features of IPv6.

The actual Internet has already begun to deviate from E2E. For example, the use of firewalls for security runs counter to E2E, but only a limited number of users are able to establish a server by themselves for security. E2E is a design concept used for the Internet when it was operated by professionals only, in its early years. Today, even e-mail is neither sent nor received E2E. Many users have created their own websites in ISP servers to send out information around the world, even though the network is NAT based. The majority of general users would find it difficult to express themselves without the support of Web servers. TCP/IP has fully played its role as the divide in the fundamental transition of communication from telephone networks to the Internet. Even if it were corrected, we could not return to conventional telephone networks.

6. Would Transition Be Possible?

Assuming that transition to IPv6 is meaningful, is it possible to achieve worldwide adoption of IPv6? In the case of zip codes, the national government changed the structure but IP addresses cannot be changed by an order from the national government, and users the world over need to modify their routers, servers and terminals or purchase new ones, at their own cost. Although they did for the Internet in its early days, current users centered on ISP do not provide implementation services voluntarily. The advantages of transition therefore need to outweigh the costs of modification. Considering that the current number of sites the support IPv6 account for only 1/30000 of the total 33 million sites worldwide, there is no business incentive to make the transition. Users in the United States, where many addresses remain unused, are especially reluctant. Enterprises that form hubs for the Internet (such as IBM, Hewlett-Packard and Apple) have a great number of addresses, so it is unlikely that these enterprises will make the transition. Ironically, IPv6 cannot be promoted unless the enormous stocks of IP addresses retained in the United States are discharged.

The greatest advantages of the Internet are its simplicity and universality, so the co-existence of IPv4 and IPv6 packets is undesirable as it may cause networks to become unstable. In reality, IPv4 addresses are allocated along with those for IPv6 to form "dual stacks," because users cannot use the addresses unless they are viewable from IPv4 sites. As such, it has not been useful at all for saving addresses. Therefore, the issue is whether it would be possible to entirely replace IPv4 with IPv6. Looking at the current speed of transition, it will take several centuries to make all sites worldwide IPv6-compliant.

The full transition of standards adopted by several hundred thousand people around the world into non-interchangeable specifications involves many difficult tasks, and only rarely results in success. In particular, in the market for communications protocols with strong "network externality," specifications that have become the standards bring about strong dependency on the historical path. As well known from examples like VHS and MS-DOS, functional excellence is neither required nor sufficient condition for making specifications a standard. Historically, there has been virtually no precedent of a successful replacement of a low-cost standard shared by a great many people with a costly and highly functional "sustaining technology," such as IPv6, achieved by improving the conventional technology. (Christensen, 1997) On the other hand, many examples of failure can be found, such as the four-channel stereo and L cassette in the past, and more recently the Advanced Photo System (APS). "Great progress" that causes international standards to switch over has been achieved only when an entirely new function became possible and the benefits of the new application are so enormous that users are willing to pay the cost for the switch-over. Compact discs and digital cameras are examples. On the other hand, the much-ballyhooed advantages of IPv6 only concern service providers, such as the elimination of restrictions from NAT and the availability of sufficient addresses. It is not clear how end users benefit. It may even end up in total futility, the fate of Open Systems Interconnection (OSI).

In reality, the only strategy for IPv6 to survive is to look for "niche markets," where interoperability is not much of a problem, instead of trying to totally replace the current system of IP addresses. Just as Macintosh has become a standard for education and art and Betamax boasts an overwhelming share in industrial applications, IPv6 may also flourish if it focused on clear targets to provide "high-grade Internet" services. However, technologies for the third-generation mobile phones and Bluetooth, which are expected to support IPv6, have yet to bear fruit. There is no plan for these technologies to be actually implemented. And is pointless to load net appliances with dual stacks, given that their computation resources are limited and maximum cost reduction is required.

After all, IPv6 seems usable only for P2P serial computations performed in an exclusive network of specific users, and for hobby applications like network games. However, we should recall that the WWW started as a hobby project. If an attractive application were to be created based on the concepts for IPv6, just like NCSA Mosaic in the past, IPv6 would instantly become popular without any aid from the national government. On the contrary, the risks are high that IPv6 would result in a failure if it were forced to be adopted widely as a "fixed protocol," because its applications are yet scarce and negative repercussions are foreseeable from its outward tendencies.

7. Limitations of IP

We shall now look at the matter from a very long perspective of time. Perhaps by 2020, we may run out of IP addresses, but IP itself is likely to have become obsolete by that time.

It is amazing that a technology that was in place in around 1970 still survives today, and it is unlikely the IP will remain a mainstream technology for another decade. Already in the U.S., optical switching (or optical routing) now forms core networks, and users in metropolitan areas have begun to turn to the "metro LAN" of the gigabit Ethernet. Given that broadband technologies require guaranteed bandwidth, it is better if routing is made for at low layers where possible. However, there is a danger that networks dependent on physical infrastructure may become closed, just as telephone networks are. So even though we need logical addresses like IP, it would be used as a local protocol that performs addressing at edge (terminal) routers. For these purposes, IPv4 is sufficient.

If radio waves are used effectively, a radio Internet in the 100 Mbps class may replace Fiber To The Home (FTTH), the mainstream technology for access systems. However, IP is not suited to mobile terminals. Since routing tables have been fixed at each site, network addresses are changed each time a terminal accesses a different base station (router). The current "mobile IP" carries out complex processes, in which data are transferred after conveying the address of the destination to the base station of origin. As such, a handover cannot take place at high speed. It is better if a network is structured in such a way that a terminal is given an absolute address that is not dependent on the router, which makes tracking possible at any time.

In a storage network like Content Delivery Network (CDN), data are transferred in a bundle to a cache server in advance, and only the requested address of a user is sent to the cache server via the Internet. Since the current Internet lacks a function to transfer data in bulk at high speed, CDNs have used communication satellites or dedicated lines. This may also cause closed "toll roads" to be formed on the Internet. To avoid this, the current "data gram" system that sends addresses and data together should be replaced with a new open protocol that separates the two and forms interconnections.

MPLS, mobile IP and CDN all carry out some form of network control in a physical layer. As such, none is an E2E network. However, today's Internet cannot operate without these systems. The reason why IPv6 has not become popular even after ten years, even though these "undesirable" technologies have spread rapidly, is because its objectives deviate from actual needs. As a result, the Internet, which started out as a simple E2E structure, has become a hopelessly complex network with many inconsistent protocols. To respond to these serious problems, a new open architecture needs to be developed, allowing different protocols to freely coexist, instead of sticking to IPv6.

8. Conclusion

Based on the above discussions, we have concluded that the commercialization of IPv6 need not be promoted in haste, and the national government's attempts to disseminate and promote it may have detrimental effects on the sound development of IPv6.

The IP addresses currently available are unlikely to run out in the foreseeable future. If we were actually running out of IP addresses, attempts to deal with the situation should have started worldwide. But only a small number of countries other than Japan seems interested in IPv6, and there are only several examples of a national government stepping in to promote it. It is not true that "connections to the Internet cannot be made in countries other than Japan, the United States, European nations, and Australia," and a needed IP address is allocated to any applicant in the world. Attempts to sell IPv6 based on the dubious theory that "we will run out of IP addresses soon" may be counterproductive because the significance of IPv6 itself may be questioned once the theory is shown to be false. Also, attempts to arouse nationalism among Japanese people by hinting that "IPv6 will enable the home appliance industry of Japan to strike back" will only intensify the impression that IPv6 is a protocol that serves only the national interests of Japan, thus preventing it achieving a worldwide footing. Even if the Japanese government had succeeded in pushing through IPv6 promotion in Japan, no one would find it worthwhile if IP addresses could be used only in Japan.

The limitations of IP have become apparent in many aspects, and IPv6 is capable only of expanding address spaces. It is not an essential technological innovation that can lead the world to switch all IP addresses. Of the technologies that the Japanese government supported through subsidies after the 1980s, such as high-definition TV, CAPTAIN System, OSI, fifth-generation computers, the Sigma Plan, TRON, ISDN, digital broadcasting and third-generation mobile phones, none has been truly successful. That is because the Japanese government is not aware of the fact that true technological revolution appears as a low-cost and functionally inferior "disruptive technology," one that is completely different from existing technologies.

The greatest issue that today's Internet faces is that the limitations of the existing architecture are becoming clear, as broadband technologies intensify demand for guaranteed bandwidth and mobile technologies necessitate dynamic routing. The E2E architecture, which performs logic routing without depending on any physical layer, has guaranteed the universality of the Internet and has driven freely the technological innovations. However, the freedom is costly as it makes network quality management difficult. Until now, such problems were disregarded with the promise "best efforts" will be made. But as the Internet itself becomes a social infrastructure, it cannot be entirely free of its various responsibilities to society. In short, the Internet needs to grow up.

The authors do not oppose IPv6 research, tests and development. Indeed, the development of new protocols is urgently needed, to replace IP with an alternative more suited for the broadband age. But focusing only on IP address spaces is too narrow as a vision. We are in a period of transition from 30 years of packet exchanges (or data grams) to a new architecture, and the route we should take remains unclear. Issues that the Internet faces now may be only solved with an entirely new architecture, and not by incremental improvements such as IPv6. We are at a stage where wide-ranging research needs to be conducted, to reveal a picture of the next-generation network. Attempts by the national government to promote a specific element technology may mislead. The Internet has evolved into its current state because users have managed it autonomously, unaffected by the interests of any nation or enterprise. Therefore, the national government should withdraw from the promotion of IPv6, and leave Internet issues to the private sector. We need to reconsider David Clark's famous message now: "We reject kings, presidents and voting. We believe in rough consensus and running code."

References

Christensen, C. M.: The Innovator's Dilemma, Harvard Business School Press, 1997 (Translated by Yumi Izuhara, Innovation-no-Dilemma, Shoeisha)
ICANN: "McFadden/ Holmes Report of the Ad Hoc Group on Numbering and Addressing", http://www.icann.org/committees/adhoc/mcfadden-holmes-report-08mar01.htm
Lessig,L.: The Future of Ideas, Random House, 2001
Marcus, J. S.: "Global Traffic Exchange among Internet Service Providers (ISPs)," 2001, http://wwwl.oecd.org/dsti/it/cm/act/Berlin/marcus.pdf/
Hori, Ikenaga, Mombayashi & Goto: Interconnection of Networks, Iwanami Shoten, 2001.

Appendix: Allocation of IPv4

* /8 means Class A; the numbers 058 and onwards are allocated by each RIR.