Understanding and Preparing for the IPv4 Address Exhaustion Issue

The IPv4 address exhaustion issue that has been a point of discussion for the past few years has now become a reality. Here we explain the issues surrounding IPv4 address exhaustion and the responses that will be required from ISPs and corporations.

3.1 Introduction

Communications over IP (Internet Protocol) consist of the exchange of data that is split into small fragments known as packets. A variety of data is delivered between network devices, with even large pieces of data segmented into small packets for transmission, and then reassembled into the original data on the recipient side. These packets include numbers that indicate the devices on both the sender and recipient sides. This information is like an address that shows the location of devices on a network, and is called an IP address.

There are a variety of topics to discuss regarding IP addresses, but here we examine the IPv4 address exhaustion issue that has been a popular topic regarding IP addresses from the perspectives of both ISPs and corporate users.

3.2 The Internet and IP Addresses

Devices communicating over the Internet are all assigned an IP address. Packets can be delivered over the Internet based on the recipient's IP address no matter how far away their device is. The technologies that make this possible are called route control and routing. On the Internet devices known as routers dynamically exchange information about which IP address is being used where, enabling packets to be delivered to the correct recipient at all times.

This means that IP addresses play a key role in delivering packets to the correct recipient. The delivery of packets to a recipient device over any distance using only an IP address depends on the fact that each of these IP addresses corresponds to a unique device on the vast network that is the Internet. If each user around the world set their own IP address it would cause overlap, making it no longer possible to deliver packets to the correct recipient. To prevent this and assure the uniqueness of IP addresses, some kind of management structure is required.

3.3 IP Address Management

3.3.1 IP Address Allocation

On the internet IP addresses are currently allocated through the Internet Registry, which has a hierarchical structure. The global free pool of available IPv4 addresses is managed by IANA (Internet Assigned Numbers Authority), which is operated by ICANN (Internet Corporation for Assigned Names and Numbers). IANA allocates IP addresses as necessary to RIR (Regional Internet Registry) such as APNIC (Asia-Pacific Network Information Centre), which are established in each region. LIR (Local Internet Registry) such as ISPs that actually allocate IP addresses to users receive IP address allocations via an RIR such as APNIC or an NIR (National Internet Registry) such as JPNIC (Japan Network Information Center).

ISPs use the IP addresses allocated to them in this way to construct networks and provide connection services to users. Until now ISPs allocated static IP addresses for always-on connection services such as dedicated lines and data centers, and dynamic IP addresses for each session of temporary connection services such as dial-up connections. Because the allocation of dynamic IP addresses can be carried out according to the level of demand, they had the benefit of enabling operation with a smaller number of IP addresses than actual subscribers. However, as the demand for broadband connections increases, IP addresses are not released even when subscribers are not using the Internet, making more IP addresses necessary even with the same number of subscribers.

At the time of writing, IP addresses are still being allocated based on existing allocation policies. For example, ISPs apply for the required volume based on prospective demand, and receive an allocation of IP addresses if this application is approved. However, a major issue is now taking place, as the number of available IPv4 addresses managed by IANA has been exhausted. When available IPv4 addresses are exhausted, the method of allocation according to demand that has been used up until now will no longer be possible. At present the APNIC has already adopted a policy of halting allocation of IPv4 addresses based on existing demand when the number of available addresses under their management drops below a certain level, and allocating a limited range of IPv4 addresses uniformly to their LIR members.

3.3.2 The IPv4 Address Exhaustion Issue

When IANA allocated IPv4 addresses to APNIC in early February, the number of IANA IPv4 addresses available dropped below the set level. In response, on February 3, 2011 the last remaining addresses were allocated evenly between five RIR based on the final allocation policy determined through previous discussions, exhausting the IPv4 addresses held by IANA for allocation. Considering past trends in demand and allocation, APNIC is expected run out of available addresses by May 2011 at the earliest, and by the end of 2012 at the latest. As JPNIC has almost no available addresses of their own, the exhaustion of APNIC addresses means that addresses will run out in Japan as well.

It is thought that the exhaustion of IPv4 addresses will cause a number of issues, such as ISPs looking to start operations not being able to secure the required IP addresses, and services for individual users not being able to connect new users. The timing of these events will differ depending on the number of available IP addresses held by each ISP, as well as usage growth. However, either way the IPv4 addresses available to ISPs acquiring new customers or expanding their network will inevitably be exhausted.

From a medium- to long-term perspective, it is expected that the absolute number of IPv4 addresses will not be enough to cover the number of devices connected to the Internet. This means that the implementation of the IPv6 protocol, which increases the number of IP addresses available, will be necessary as a fundamental solution to IP address exhaustion. However, IPv6 is unfortunately not compatible with IPv4, so the addition of the new protocol will require work to implement. These constraints have meant that IPv6 implementation has not proceeded very swiftly. But the exhaustion of IPv4 addresses has become a reality, and it appears that IPv6 implementation is finally being looked at.

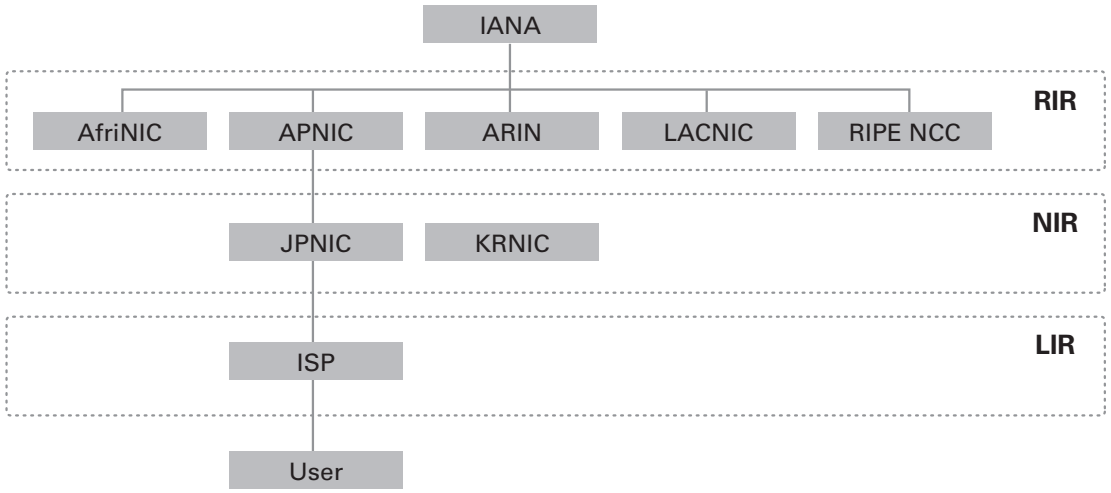

Figure 1: IP Address Allocation through a Hierarchical Structure

3.3.3 ISP Responses

Although the implementation of IPv6 is progressing, many sites are still operated using IPv4. With this being the case, it is necessary to provide IPv4 connectivity to users that want to connect to the Internet. For the time being IPv4 addresses are also necessary for servers that provide mail and Web services. In other words, IPv4 address operations must be extended by some means into the future. The main methods of approaching this are the effective use of IPv4 addresses, and the sharing of IPv4 addresses.

For example, one possible way of making effective use of IPv4 addresses is to recover IPv4 addresses not being used currently and use them in a different location. The mutual lending of IPv4 addresses between organizations is also being considered, and APNIC has already adopted a policy that makes it possible to transfer IPv4 addresses between its member organizations. Because it was judged that many of the connection services provided to users are for accessing external servers for Internet and mail, technology for sharing IPv4 addresses between users is also being examined.

3.4 Corporate Responses

We often receive queries from corporate information system administrators, such as whether they need to take action with regard to IPv4 address exhaustion issues, or whether their system needs to support IPv6. Should companies respond to IPv4 address exhaustion issues? Do companies need to support IPv6? Next we discuss viewpoints on policies for responding to the IPv4 address exhaustion issue, as well as specific examples of measures to take.

3.4.1 Analyzing the Impact on Your Company

First, it is necessary to properly understand the IPv4 address exhaustion issue. What exactly is the IPv4 address exhaustion issue? What is the status of the IPv4 address exhaustion issue, and how will it develop from now? Who will be impacted in what ways, and how will this affect the global climate? These are the kinds of things that must be understood correctly.

After assessing the current state of the IPv4 address exhaustion issue, look into how this will impact your company's system. The following two points should be considered.

- The number of IPv4 addresses available at your company
- The number of IPv4 addresses available globally

First check the number of available IPv4 addresses at your company and plans for future system enhancements to assess whether the consumption of IPv4 addresses will continue, as well as how long the currently available addresses will last.

However, even if the number of IPv4 addresses available at your company is sufficient, it does not mean that no response is required. When the IPv4 addresses available globally are exhausted and implementation of IPv6 proceeds, if your system only supports IPv4, users using IPv6 will not be able to connect. For this reason general companies as well as of course companies doing business via the Internet must implement IPv6 for public server groups and DMZ at a minimum.

Additionally, the implementation of IPv6 may have business advantages and disadvantages. For example, it is necessary to consider what business disadvantages you will face if a competitor company supports IPv6 and your company does not, and what advantages you will have if you support IPv6 while a competitor company does not. Cost is another important point to consider. The implementation of IPv6 can cost a great deal more than information system administrators envision. The decision of whether or not to implement IPv6 should be made after comparing business advantages and implementation costs.

Judge whether or not to respond to the IPv4 address exhaustion issue based on details such as these. If a response is called for, the timing and methods of response must also be looked into.

3.4.2 Measures Taken at Companies

A permanent measure for the IPv4 address exhaustion issue is the implementation of IPv6. However, other measures must be taken to extend the number of IPv4 addresses available.

One procedure for extending the number of IPv4 addresses available is the recovery of any IPv4 addresses that are assigned to servers or clients but are not actually being used. The optimization of any unnecessary network partitions also serves to extend the number of IPv4 addresses available. For example, let us say there are a total of eight IPv4 addresses in a network consisting of a subnet on /29. However, as two are used for the network address and broadcast address, and one to three for the default route, the number of addresses that can actually be used by the host is three. When there are many networks making up /29, by joining them together it is possible to use IPv4 addresses more economically. Network stocktaking through the recovery of unused addresses and the optimization of network partitions leads to more economical use of IPv4 addresses.

The most important point when implementing IPv6 on corporate information systems is the implementation plan. When systems were constructed using IPv4 up to now, it is likely that factors such as rules for assigning IP addresses, DNS naming conventions, and filter policies were designed according to previously-established rules.

When constructing a system using IPv6, it is necessary to decide how each of these factors that were taken for granted with IPv4 should be executed. Next, we examine a number of specific points to consider.

■ IPv6 Addressing

IPv6 addresses can be broadly categorized as either global addresses or link-local addresses. Additionally, methods for assigning IPv6 addresses to a device include the manual assignment of a fixed address, and the EUI-64 method of calculation from a MAC address. As this demonstrates, for each network device it is necessary to consider whether to assign a global address to the interface or just a link-local address, and whether to assign the IPv6 address manually or use the EUI-64 method. You must also evaluate whether or not to allow the use of the anonymous addresses unique to IPv6 (RFC3041) in server devices and client terminals.

■ DNS Naming Conventions

Normally, when implementing an IPv4 and IPv6 dual-stack system, the most standard name (e.g. www.example.co.jp) is specified for the dual-stack name. However, this kind of naming convention creates issues during operation because engineers operating the servers cannot tell whether communications are being carried out by IPv4 or IPv6. Methods for resolving this issue include naming rules that, for example, return www-v4.example.co.jp as the IPv4 name and www-v6.example.co.jp as the IPv6 name.

Pattern 1: Tunnel Connection Method

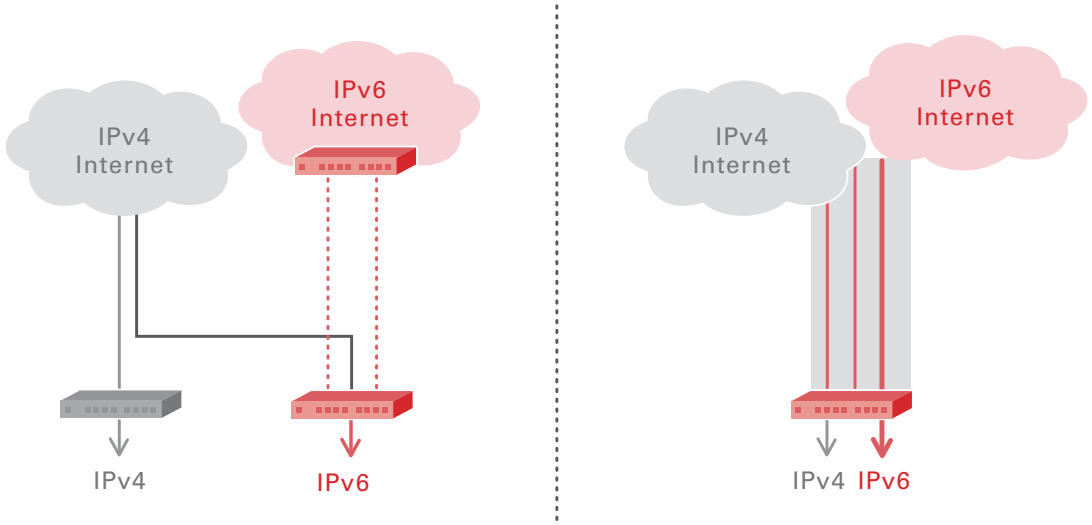

Figure 2: IPv6 Network Implementation Methods

■ Improving IPv6 Skills

When implementing IPv6 on an information system, it is necessary for all those involved with the system to acquire skills relating to IPv6. However, at present there is a shortage of engineers that can handle IPv6 at many companies, so improving the IPv6 skills of in-house engineers will be a key point.

■ Implementing IPv6 on a Network

Methods for implementing IPv6 on a network can be divided into two main categories. These are the tunnel method and the dual-stack method. The tunnel method involves installing an IPv6 tunnel router on an existing IPv4 network, and encapsulating IPv6 packets into IPv4 packets for transfer. Meanwhile, the dual-stack method involves making lines and routers capable of transferring a mix of both IPv4 and IPv6.

■ Implementing IPv6 on a Server

Methods for implementing IPv6 on a server can be divided into two main categories. These are the translator method and the dual-stack method. The translator method involves installing a translator device at the server, and not changing anything about the existing server itself. IPv6 packets are converted into IPv4 packets by the translator device, and transferred to an existing server. The dual-stack method involves making the server itself capable of handling both IPv4 and IPv6.

3.5 Conclusion

As discussed above, the IPv4 address exhaustion issue is now upon us. The exhaustion of available addresses is expected to cause changes that we have not experienced up until now, such as revisions to IPv4 address allocation policies, and user environments being altered through IPv4 address sharing and the implementation of IPv6. IIJ will respond to the exhaustion issue through its industry activities and implement IPv6 at an early stage to continue to make the Internet a user-friendly environment. We will also continue to provide support, giving customers the information they require and responding to the needs of customer organizations.

Pattern 1: Translator Method

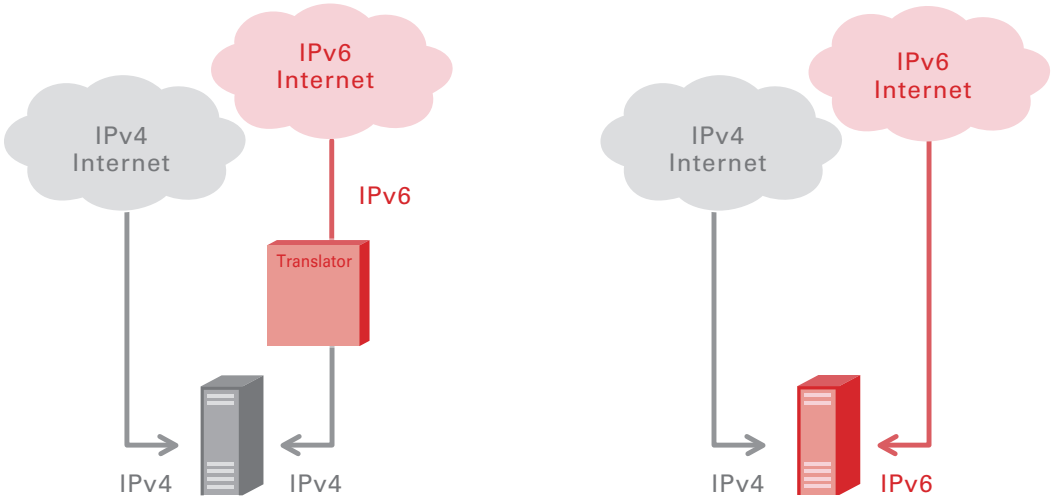

Figure 3: IPv6 Server Implementation Methods

Authors:
Yoshinobu Matsuzaki
Mr. Matsuzaki is a Senior Engineer in the Technology Promotion Section of the Network Service Division in the IIJ Network Service Department. Mr. Matsuzaki is always finding things that pique his interest while striving at his work. He is an IIJ-SECT member, co-chair of The Asia Pacific OperatorS Forum, chair of APNIC IPv6 SIG, and an expert advisor for JPCERT/CC.

Yasuyuki Kaga
Mr. Kaga is a Consultant at the Service Management Section of the Service Integration Division in the IIJ Service Department. He is engaged in the construction of corporate networks and consulting. He currently provides consulting regarding the implementation of IPv6.