IPv6 and content providers

Thomas Kernen
Cisco Systems

Earlier this year, the last five available Internet Protocol (IPv4) address blocks were allocated to the Regional Internet Registries. This signalled that the central IPv4 pool of address spaces is now exhausted and that it is only a matter of months before each of the regional pools will suffer the same fate.

A new Internet Protocol (IPv6) has been under development for a number of years, primarily to overcome IPv4 address-space exhaustion. However, as discussed here, IPv6 also allows for the integration of other components such as multicast, which may be of interest to content providers.

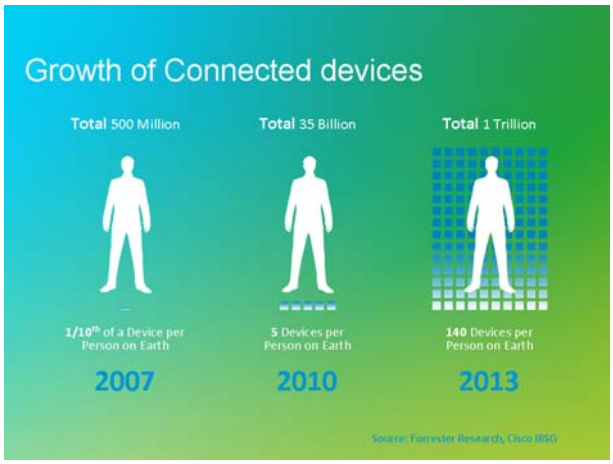

According to Cisco's Internet Business Solutions Group (IBSG), the growth in connected devices will reach one trillion (10^{12}) devices online by 2013. They will not only be limited to desktop and laptop computers, but also to smartphones, motor vehicles, sensor networks, power grids, water systems and home entertainment devices. Even light bulbs are getting their own IP addresses!

Back in the early 1990s, when the Internet started to attain mainstream status, researchers acknowledged the growth rate in the number of devices being connected to the Internet – with each one requiring its own unique IP address. This has meant that, over time, the current generation of the Internet Protocol (IPv4) address space, introduced in 1981, would suffer from address-space exhaustion.

Over the course of the following few years, the Internet Engineering Task Force (IETF) set up a number of efforts to specify the requirements for a new Internet Protocol effort. By 1996 the first series of Request for Comments (RFC) documents relating to IPv6 were published. Since then, these have been revised and further extensions to the IPv6 framework have been introduced.

The IPv4 address-space allocations are handed out by the Internet Assigned Numbers Authority (IANA) to the five Regional Internet Registries (RIRs), then onwards to their customers. On 3 February 2011, the last five available address blocks were allocated to the RIRs. This signalled that the central IPv4 pool is now exhausted and that it is a matter of months before each of the regional pools suffers the same fate. Indeed, this occurred on 15 April 2011 for the Asia-Pacific Region, governed by APNIC. The European / Middle Eastern region, governed by RIPE NCC, is expected to be next.

Whilst the primary intent was to overcome the address space exhaustion, the definition of the new protocol has allowed for the integration of other components, such as ***multicast***, that had been developed for IPv4 and may be of interest to content providers. Within IPv6, multicast is a native and mandatory requirement for many of the protocol interactions. Hence, the framework for end-to-end multicast services does exist, but remains dependent on service providers enabling those functionalities for their customers.

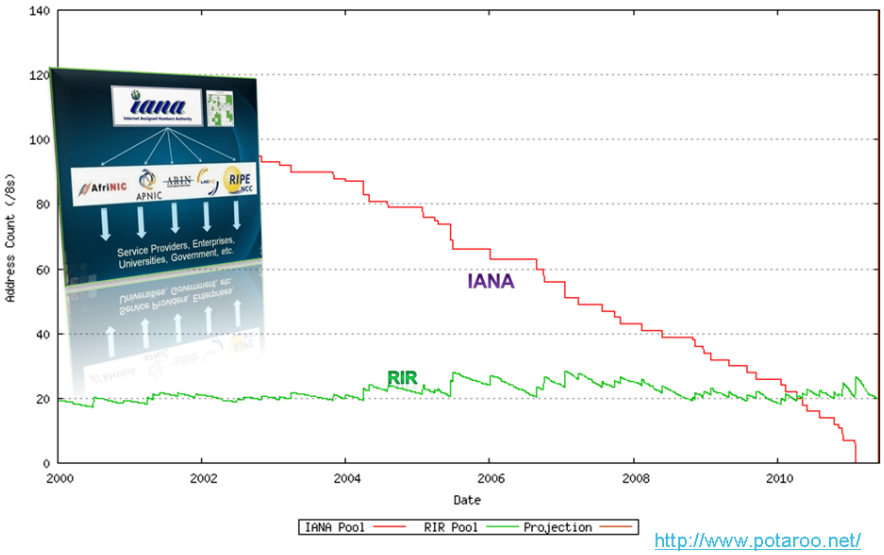

Impact on internet connectivity

The IPv4 address-space exhaustion will have a major impact on the growth of the Internet and on Internet Service Providers (ISPs). Any ISP that wishes to continue to grow its revenue by increasing its customer base will have to find a technique to add new Internet users without requiring additional globally unique IPv4 addresses. IPv6 was designed primarily to function over existing layer 2 technologies in the same way as does IPv4, and to have a larger address space, making it unlikely that the global Internet would ever suffer another such shortage. It is the networking industry's plan of action. **There is no alternative plan**.

The only issue facing us is when and how a transition from IPv4 to IPv6 will occur.

IPv6 is not directly compatible with IPv4: an IPv4-only node cannot communicate with an IPv6-only node (and vice-versa). While IPv4 and IPv6 are not compatible, they can use the same network simultaneously, and various technologies aim at v4/v6 integration and coexistence.

Integration and coexistence strategies are briefly described below.

- **Dual-stack**

  All Internet users are given both a routable IPv4 and a routable IPv6 address (actually a network prefix). It is then up to the user's computer to select which address to use (most operating systems always prefer to use IPv6 when the corresponding node has both an IPv6 and an IPv4 address). The Dual-stack technique cannot be applied to all current and future Internet users, as the ISP must allocate one globally-routable IPv4 address to each of its customers, and very soon there will no longer be enough free IPv4 address blocks.

- **Shared-IPv4 address**

  The ISP shares a few globally-routable IPv4 addresses among several hundred or even thousand of its customers. Each customer is assigned an IPv4 address (e.g. RFC 1918) that is only used within the ISP network. If a customer wants to access the Internet, then the customer packets will go through a Network Address Translation (NAT) device that is implemented within the ISP network. The ISP can also provide IPv6 connectivity at the same time.

- **IPv6-only**

  The ISP does not give any IPv4 Internet access to the customer, who has access only to the IPv6 part of the Internet. This can be combined with some ISP-operated Address Family Trans-

lation (AFT) mechanism that allows an IPv6-only node to communicate with an IPv4-only node. Mobile operators in some countries intend to offer this connection on next-generation mobile handsets: Long Term Evolution (LTE).

Each ISP will adopt one or more different techniques, and each ISP will probably change techniques after a couple of years. It is expected that the Internet will slowly move to be IPv6-only, but not within at least five years.

Service providers are likely to implement any of these three approaches over the short to medium term. The risk of shared IPv4 addresses for the enterprise is that some applications may fail or may work poorly; in particular, applications that make use of many concurrent transport connections such as AJAX which is used for dynamic web pages. Content providers have no say in which service providers are selected by their customers and business partners, and hence have little say as to whether a customer has access to IPv6 or is using a shared IPv4 address space. The best course of action is therefore to take a conservative approach toward IPv4 (such as not using AJAX for IPv4 clients) and a more aggressive approach with IPv6 (such as using AJAX for IPv6 clients and for delivering a better service) as consumers become IPv6-enabled.

The coexistence of several techniques leads to classifying Internet users in the next 3 to 5 years as follows:

- **Public IPv4-only –** an Internet user who has had his/her public IPv4 address and is keeping it for the foreseeable future. He/she can only access IPv4 services.
- **Shared IPv4-only –** an Internet user whose connections to the Internet go through a NAT function operated by the ISP or the enterprise. He/she can only access IPv4 servers, and the use of NAT puts constraints on the applications he/she can use.
- **Public-IPv4 and IPv6 –** an Internet user who has public IPv4 and IPv6 addresses and can access both IPv4 and IPv6 services without any restrictions.
- **Shared-IPv4 and IPv6 –** an Internet user who has a public IPv6 address and a shared IPv4 address, can access all IPv6 services without any restrictions, and all IPv4 services through a NAT.
- **IPv6-only –** an Internet user who has only a public IPv6 address and can access only IPv6 services.

Impact on content providers

The existing Internet presence of content providers will not change for existing users that are currently connecting over IPv4. The question is: how will customers and business partners who are connected over IPv6 be able to access new or existing services that are being provided. Other questions to consider are:

- Do content providers wish to allow third-party service providers to handle the translation between IPv6 and IPv4 on their behalf, and therefore not be in control of delivering a native service to their users?
- Do local regulations enforce that content providers must ensure that all users have access to their content under some form of “universal access” to their services?
- Are there new or existing services that would not be compatible with a shared NAT for multiple-user services being provided by service providers?
- Is there any performance or resiliency benefit either to adding IPv6 or staying with IPv4?
- Could a unique identifier, such as an IP address, provide value to the services being delivered?
- Does their business require them to interconnect with partners at short notice and therefore all types of connectivity must be made available up-front to minimize any delay in setting up such new services?

These are some of the questions that a content provider may need to take into consideration with regard to planning accessibility to their services or in turn their accessibility to services provided by third parties. Each case will be different, and will most likely evolve over time. Some specific services may not suffer from end users using NAT between IPv4 and IPv6, whilst others may be totally incompatible with such capabilities.

Some of these issues are not limited to the Internet presence but may also arise in IP-based contribution networks whereby more and more content is exchanged with third-party networks. Even if these are generally closed entities that are not directly connected to the public Internet, some partners may be using IPv6 for their own infrastructure as a forward-looking solution. Therefore the need to set up occasional-use fixed services, live or file-based, may dictate the need for IPv6 connectivity.

Enabling the IPv6 Internet presence

Whilst content providers may direct their focus towards their media-rich web services and media file exchange services, one must keep in mind that enabling an IPv6 Internet presence should also take into account two other important services: e-mail and Domain Name System (DNS). In order for these two services to be fully functional, some back-office and management systems will also need to become IPv6-aware.

In many cases the method for building and rolling out these IPv6 services will differ for each content provider, and some hosting partners may or may not be able to deliver IPv6 for those services at this point in time. Therefore a careful review of the services will most likely be required.

Here are some of the key steps to enable an IPv6 presence, directly managed by the content provider or via hosting partners.

Getting public IPv6 address space

As IPv6 is fundamentally not very different from IPv4, in order to build an Internet presence one must first acquire a block of addresses that can be routed on the Internet. IPv6 address blocks are distributed just as IPv4 blocks are. Service providers will likely include provider-allocated address space as part of the service. Unless the service itself is multi-homed, Provider Assigned (PA) address space is sufficient. Otherwise, organizations must procure provider-independent (PI) address space from their Regional Internet Registry. Note that different registries may have different policies and cost structures relating to PI address space.

Reviewing application needs

Procedures and requirements will vary widely across enterprises based on what services are publicly exposed. We discuss some of the more common ones below.

Abbreviations

AFT	(IPv6) Address Family Translation	**FQDN**	Fully-Qualified Domain Name
DNS	Domain Name System	**MTA**	Mail Transfer Agent
IETF	Internet Engineering Task Force http://www.ietf.org/	**NAT**	Network Address Translation
		RFC	Request For Comments (IETF standard)
IANA	Internet Assigned Numbers Authority http://www.iana.org/	**RIR**	Regional Internet Registry
		SMTP	Simple Mail Transfer Protocol
ISOC	Internet Society	**TCP**	Transmission Control Protocol
ISP	Internet Service Provider	**UTC**	Universal Coordinated Time

Adding IPv6 to web servers

In order to get an IPv6 web presence, it is usually enough to implement IPv6 on the front end of all web servers; there is no immediate need to upgrade any back-end database or back-end server, as those servers are never directly accessed from the Internet.

There are multiple ways of adding IPv6 connectivity to a web server farm:

- **Adding native IPv6 to existing web servers**

  Configure IPv6 on the web server itself (Apache, IIS and most other modern web servers have supported IPv6 for several years) as well as on the load balancers. This is the clean and efficient way to do it, but some applications or scripts running on the web servers may need some code changes (notably if they use, manipulate or store the remote IP address of their clients).

- **Adding a set of standalone native IPv6 web servers**

  Configure standalone web servers separately from your IPv4 infrastructure. This has the benefit of reducing dependencies on other components, perhaps even allowing selection of different hosting providers for IPv4 and IPv6. Of course, back-end processing must still be taken into account. Whether that happens using IPv4 or IPv6 is a separate decision.

- **Using Address Family Translation (AFT) in load balancers**

  Some modern load balancers are able to have clients connecting over IPv6 while the servers still run IPv4; those load balancers translate back and forth between the two address families (IPv4 and IPv6). This is probably the easiest way to add IPv6 to the web servers. Without a specific configuration, some information is lost in the web servers' logs because all IPv6 clients will appear as a single IPv4 address.

- **Using AFT in reverse web proxies**

  If reverse proxies are used (for example to enforce some security policies), then they can be used similarly to perform address family translation (with the same caveat as for load balancers).

- **Using AFT in network devices**

  In 2010, the IETF finalized the specification of AFT (either stateless or stateful) carried out in network devices when the connection is initiated from an IPv6-only host to an IPv4-only server. This is often named XLATE – previously known as NAT64 (address translation from IPv6 client to IPv4 server). It is expected that vendors will add this function to their routers in 2011 or 2012. This is another easy way to get an IPv6 Internet presence without touching the actual web front-end servers.

Adding IPv6 to e-mail

The sending and receiving of e-mail over the Internet occurs through Simple Mail Transfer Protocol (SMTP) over the TCP protocol. Most popular Mail Transfer Agents (MTAs) are fully capable of using IPv6. However, some of the support functions now common in these servers are not yet present for IPv6. This includes blacklisting and reputation services notably used for antispam. When more traffic, and hence more spam, moves to IPv6, these tools can be expected to become available.

Many sites also run scripts that parse mail server logs. IPv6 will change the format of those logs. As with other services, some care should be taken to ensure that such service management functions are IPv6-capable.

Today, IPv4-based mail systems will reject incoming mail from servers in domains that are not properly configured. The same can be expected for IPv6-based mail systems. That is, properly configured DNS will be just as much a prerequisite for IPv6-based systems as it is today for IPv4.

Adding IPv6 to DNS

DNS is of course a critical piece of any Internet presence, as it is used to announce the IP addresses of the web and e-mail servers. There are two steps to fully support IPv6 on a DNS server:

- **IPv6 information in the DNS zones**

  Adding the IPv6 addresses of all public servers in the DNS database. This is simply done by adding specific Resource Records (RRs) with the IPv6 address (those records are called AAAA). In order to facilitate debugging and operation, it is also advised to add the reverse mapping of IPv6 addresses to Fully Qualified Domain Names (FQDNs). For dual-stack servers, there are two RRs per FQDN: one IPv4 address (type A) and one IPv6 address (type AAAA).

- **IPv6 transport of DNS information**

  The DNS server accepts DNS requests over IPv6 and replies over IPv6. It is more common to have a dual-stack DNS server accepting requests and replies over IPv4 and IPv6.

It should be noted that these two steps are independent; one can be done without the other. In order to have an Internet IPv6 presence, only the first step must be done; that is, the enterprise must publish the IPv6 addresses of all its Internet servers in its DNS zone information. All major DNS server implementations (including ISC BIND, Cisco Network Registrar, Microsoft DNS Server) have supported IPv6 for several years.

Whichever methods are selected to enable the IPv6 presence, the management and security processes must be updated and adapted to the new capabilities.

World IPv6 Day

On 12 January 2011, the Internet Society (ISOC) announced with Google, Facebook, Yahoo!, Akamai and Limelight Networks that World IPv6 Day would take place on 8 June 2011. Since the initial announcement, over 200 companies around the world announced that they would be taking part in the event.

The goal was to motivate organizations across the industry – Internet service providers, content providers, hardware manufacturers, operating system vendors, etc. – to prepare their services for IPv6 by running a 24-hour IPv6 “test flight” operation.

The event acted as a focal point to bring existing efforts together. As a coordinated effort, it ensured that specific channels are set up to alert, discuss and focus on potential IPv6 issues that may arise, such as brokenness in end-user networks and incomplete IPv6 interconnection from end to end.

During the event, IPv4 connectivity was untouched, so IPv4-only users could continue to operate as normal. Only users that have some form of IPv6 connectivity were able to participate in this event.

Thomas Kernen is a Consulting Engineer working for Cisco’s European Innovation Consulting Engineering Team. His main area of focus is in defining video architectures and transmission solutions for content providers, broadcasters, telecom operators and IPTV service providers.

Prior to joining Cisco, Mr Kernen spent ten years with different telecoms operators, including three years with an FTTH triple play operator, for whom he developed their IPTV architecture. He is a member of the IEEE, SMPTE and is active in the TM-AVC group within the DVB Project.

Whilst this was just a 24-hour event running on 8 June 2011, 00:00 to 23:59 UTC, many of the sites that enabled IPv6 connectivity on that day are expected to keep their IPv6 services operating henceforth.

It is to be noted that in the content provider industry, Google (including their YouTube service), Facebook and Netflix have been operating services over IPv6 for some time. From a European perspective, RTBF in Belgium, Heise.de in Germany, A-pressen Digitale Medier and VG Multimedia in Norway already provide their content over IPv6.

Conclusions

It is clear that the time has come for content providers around the world to review their online strategy and the impact of IPv6 on their business. This impacts the ability of their users to reach their services independently of the means provided to access the Internet by their current service provider.

Whilst initially confined to research environments and core networks, we are now beyond the stage of experimentation and have moved into the world of real deployments in order to ensure service and business continuity.

Preparing for these changes and implementing them in the manner that makes the most sense for their business will be key to the success of content providers in the years to come.

Further information

[1] IPv6 Act Now: **http://www.ipv6actnow.org/**

[2] World IPv6 Day: **http://isoc.org/wp/worldipv6day/**

[3] Test your IPv6 Connectivity: **http://test-ipv6.com/**

[4] Cisco IPv6: **http://cisco.com/go/ipv6/**

This version: 23 June 2011

Published by the European Broadcasting Union, Geneva, Switzerland

ISSN: 1609-1469

Editeur Responsable: Lieven Vermaele

Editor: Mike Meyer

E-mail: **tech@ebu.ch**

EBU•UER

The responsibility for views expressed in this article rests solely with the author