the Availability Digest

www.availabilitydigest.com

IPv6 Is Here – Like It or Not

April 2011

The IPv4 Doomsday is here. Are you prepared?

In our article almost two years ago,[1] we repeated the alarm raised by many that the decrepit IPv4 address space of 32 bits was about to be depleted. IPv6, with its 128-bit address space, was set to become king. At the time of our article, the assignment of the last available IPv4 address was expected to be in mid-2011. We are now approaching mid-2011, and the IPv4 addresses have, in fact, been depleted. Will the Internet collapse?

Of course not, but the transition to the new protocol specification is not trivial. If you elect not to convert, you will survive for a while; but your Internet experience will become more and more painful.

Help is here. In a recent article published in the Communications of the ACM[2] (the Association for Computing Machinery), Vinton Cerf, who is recognized as one of the fathers of the Internet along with Robert Kahn, collaborates with fellow Google associate Thomas Limoncelli to relate the experiences of Google, Comcast, and Nokia as they successfully navigated the conversion from IPv4 to IPv6. The result is one simple message. As the old KISS saying goes, “Keep it simple, stupid.”

In this article, we review the IPv4/IPv6 conundrum and the successful approaches suggested by Cerf and Limoncelli.

The History of IPv4

The roots of IPv4 go back to 1973 when Vinton Cerf and Robert Kahn started to explore packet-switching networks for the United States’ Department of Defense Advanced Research Projects Agency (DARPA). The result was a dual-protocol solution that comprised the Internet Protocol (IP) that carried datagrams end-to-end through packet networks and the Transmission Control Protocol (TCP), which managed and sequenced packets to deliver complete messages to hosts.

The first commercial version of the Internet was introduced in 1981 with version 4 of IP. It became known as IPv4. The address space for IPv4 was set at 32 bits, which provided for a nearly inexhaustible (or so the developers thought) four billion addresses, about one for every other man, woman, and child on earth. Except for Internet Service Providers (ISPs) and large corporations, these addresses are normally split up into a 24-bit network address and an 8-bit field for host addresses on the network.

[1] The IPv4 Doomsday, *Availability Digest*; August 2009.
[2] Successful Strategies for IPv6 Rollouts. Really, *Communications of the ACM*; April 2011. http://cacm.acm.org/magazines/2011/4/106582-successful-strategies-for-ipv6-rollouts-really/fulltext

However, the Internet grew with a fury that was never anticipated. As of the date of Cerf's and Limoncelli's article in April, 2011, the Internet Corporation for Assigned Names and Numbers (ICANN), which assigns IP address space to end users, announced that it had just allocated the last five of the 24-bit network address subsets.to its five worldwide Regional Internet Registries (RIRs), which in turn will distribute them to their customers.

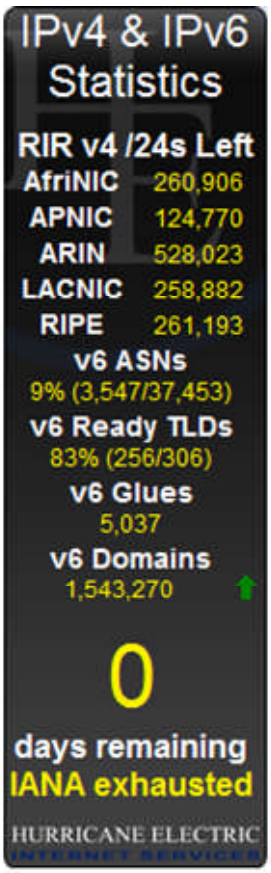

In our previous article, we showed the results of an application for iPhones and other devices available from Hurricane Electric, an IPv6 network provider. The application showed the estimated number of days left before IPv4 addresses would be exhausted. On July 22, 2009, it was estimated that this would occur in 699 days. On April 4, 2011, about 21 months or 630 days later, that time has virtually arrived. All of the 24-bit network addresses have been allocated to the regional registries, and the Registries have only about 1.4 million or so addresses left for customer assignment.

Enter IPv6

The Internet Engineering Task Force (IETF) recognized in the early 1990s that the rapid growth of the Internet would exhaust the supply of IP addresses. In response, it specified a new IP address format that quadrupled the address length to 128 bits. Being version 6 of the IP specification, IPv6 provides an immense address space – enough to give every known star in the universe a quadrillion IP addresses.

Unfortunately, the IPv6 format is not backwards compatible with IPv4. Systems that must communicate with both must use dual protocol stacks. ISPs in particular are going to have to support dual-mode operation for some years to come.

Moving from IPv4 to IPv6

Transitioning from IPv4 to IPv6 is not trivial. Not only are the protocol stacks different, but the packet headers are also different. IPv6 requires the Internet Protocol Security (IPSec) protocol, which is optional with IPv4. IPv6 does not support fragmentation – the host must query the route to determine the maximum transmission unit (MTU) and must fragment its messages accordingly. Four-gigabyte jumbo packets for video transmission are supported by IPv6 if the link allows.

Fortunately, IPv4 addresses can be encapsulated in IPv6 addresses as the low-order 32 bits, with the higher-order bits set to a specified value.

Thomas Limoncelli suggests several approaches to the protocol conversion, some feasible and some fraught with problems.

Do Nothing (Problems)

One option is to do nothing. Rely on your ISP to do the conversion for you. After all, they will be running dual IP stacks to support their customers, won't they?

Though this approach may work for the short term, there are several problems.

- The translation between the protocols is expected to be slow. This will slow down the responsiveness of your web site, perhaps seriously impacting your users' experience.

- If your web site is geotargeted – that is, its functions depend upon knowing the location of the devices logged onto your web site (for instance, marketing sites) – the location of IPv4-enabled devices will now be that of your ISP. Geographical location will be lost.

- At some time in the future, your ISP will probably drop IPv4 support. At that point, you had better be converted.

Do Everything (Problems)

Another option is to propose to your management that you make a complete conversion to IPv6. This will require converting the network equipment – routers, switches, hubs, etc., the Domain Name System (DNS), the Dynamic Host Configuration Protocol (DHCP) system, applications, clients, desktops, and servers. It is a huge project that will involve every piece of equipment and software in the network.

Management may see little value in this Big Bang approach. It is a project with a high likelihood of failure and one which will probably aggravate your network users during a very long transition. Good luck on getting it approved.

Do One Thing (Feasible)

Pick a single application – your "one thing" – and focus on moving it to IPv6.

At the 2008 Google IPv6 Symposium, Comcast described a successful project that did just this – its set-top box management software. Every set-top box needs an IP address so that the central management facility can reach it for provisioning, testing, monitoring, billing, and other functions. The number of set-top boxes that Comcast needed to deploy was more than the number of IPv4 addresses that it could reasonably get. It therefore went to IPv6. All new Comcast set-top boxes now communicate over IPv6.

At the same conference, Nokia presented a case study in which it went to IPv6 for its new cell phones for, of all reasons, power consumption. It turns out that IPv4 requires pings to keep the Network Address Translation (NAT) session alive; IPv6 does not. Nokia's phones can now turn off their antennae until they have data to send. This saves power and extends the cell phone's life until its next recharge.

Work From the Outside In (Feasible)

Web farms are usually hidden from the public Internet by some hardware device such as a load balancer or a firewall. When external Web browsers connect to your web site, it is through one or more of these devices. They act as a "man in the middle" to the server farm and perform many functions. A new function that is available in many of these devices is IPv4/IPv6 protocol conversion.

The strategy is then simple. Upgrade every component (if any) to IPv6 in the path from your ISP to, for instance, your load balancer. Then use a load balancer that provides IP protocol conversion to convert incoming IPv6 traffic to IPv4, and outgoing IPv4 traffic to IPv6.

This is a manageable project that corrects the problems of the "do nothing" approach. Conversion is fast and has little if any effect on the user experience. Geotargeting is preserved.

The strategy gives you time to proceed with a full conversion at a measured pace. As each server becomes IPv6-ready with its IPv4/IPv6 dual stack, the load balancer no longer has to make the conversions for that server. When all servers have been converted, conversion is no longer

needed in the load balancer. The same holds true for routers, DNS servers, and other components, which can be upgraded in a controlled manner. The Web farm itself now handles both IPv4 and IPv6 traffic, and a full conversion has been made. No Big Bang – just a piece-meal approach over a comfortable period of time.

Summary

The transition from IPv4 to IPv6 is difficult; but by taking it in little steps, it is doable without a great deal of pain. In many cases, the difficulties imposed by not converting may be a great deal greater than the difficulties imposed by taking the transition steps.

It is no longer safe to ignore IPv6 as a future requirement. The IPv4 space is now depleted. Repeating the summary from our previous article:

> "But there is no problem with doomsdays of the future. As long as it won't happen on our watch – as long as we will have been retired for a while when it comes – don't take it out of our budget. Leave it for the next guy to figure out. That's how we handled the Y2K doomsday, and that's how we're handling the IPv4 doomsday. After all, that's the IT way."

The above strategy is no longer valid.

Acknowledgement

Thanks to our subscriber, Paul Holenstein, for pointing us to this article.