Infoblox
CONTROL YOUR NETWORK

WHITE PAPER

Unlocking Business Value with IPv6: A CIO's Guide

By Tom Coffeen, Chief IPv6 Evangelist at Infoblox

Introduction

North America formally exhausted its supply of routable IPv4 in the summer of 2015 (joining Europe, Asia Pacific and Latin America). Yet enterprise networks (and the Internet at large) didn't stop working. Given this, IT leaders could be forgiven for thinking IPv4 address exhaustion will not impact There is, however, an opportunity to build your company's business by understanding and leveraging the power of IPv6—the successor to IPv4. While many of the technical details of IPv6 can managed by networking teams, CIOs still need to understand the problems and how to take full advantage of the solution.

Let's start with the basics. It often goes unnoticed or unremarked that the Internet's foundational technology, the Internet Protocol (or IP for short), is the key component of all modern computer networks, public and private. This ubiquity of IP is a testament to the relative simplicity of its design and the engineering resilience and flexibility it delivers.

IP provides the unique addressing required for any device or endpoint to connect to the Internet and send and receive information. These addresses are formed with 32 binary bits. Their human-readable form is 4 decimal values (from 0-255) called octets, separated by dots: e.g., 192.168.10.1.

The need for IP infrastructure and more IP addresses is greater than ever within enterprise networks; the BYOD phenomenon and the adoption of cloud computing are just two more recent examples of technology trends requiring more IP addresses. The emerging Internet of Things promises more of the same, and at potentially enormous scale. But the traditional "silo-ization" of IT (and the mentality that accompanies it) frequently leaves enterprises unprepared to accommodate these trends. This legacy mindset leads to rigid solutions that are unable to maximize business agility and value while also ensuring the security and uptime of critical data and applications.

This is the critical context in which the issue of the exhaustion of legacy IP addresses must be considered. We are living and working through the period of the exhaustion of legacy IP—also known as IP version 4 (IPv4). Chances are you've recently heard or read about IPv4 run-out (also known as exhaustion or depletion) and either that it's a) already happened or b) is about to happen soon. Given such an apparent discrepancy, skepticism is understandable. So what are the bare facts? Is IPv4 exhausted already or will it run out sooner, later, or never? In fact, IPv4 ran out in North America in the summer of 2015. But what does that mean, given the earlier proclamations or predictions of exhaustion? Before answering, let's briefly review the history of IPv4 and how we arrived at this moment.

A (Very) Brief History of IPv4

The Internet Protocol's earliest incarnation appeared in 1973 as part of the U.S. government's DARPA (Defense Advanced Research Projects Agency) project to create a fault-tolerant network connecting computers across the country and ultimately the world. But it wasn't until 1978 that the address space size of 4.3 billion was chosen by IP's co-inventor, Dr. Vint Cerf. Given that the "Internet" of the time was still very much a research network and consisted of around 100 nodes, those 4.3 billion available addresses seemed more than sufficient. But the public and commercial expansion of the Internet in the 1990s made it obvious to the Internet engineering community of the time that this number would be insufficient. Their predictions suggested that IP exhaustion would happen within that decade. However, the invention and wide deployment of Network Address Translation (NAT), along with methods for more granular subnetting, greatly extended the life of IPv4. NAT makes a subset of "private" addresses available for use (and reuse) by enterprise networks, reducing the demand on routable public IPv4 addresses and, notably, breaking the original end-to-end model of the Internet. Several successor protocols were proposed, but IP version 6 (IPv6), first proposed in 1995, emerged as the one that would eventually replace IPv4.

Introducing IPv6

The IPv6 address space is really, really big. How big?

- The number of possible IPv6 addresses is five quadrillion times the number of estimated stars in the universe.
- A single LAN interface subnet in IPv6 contains the same number of total IPv4 Internet addresses *squared* (or roughly 18 billion *billion* addresses).
- If the total number of IPv6 addresses were the volume of the Earth, IPv4's equivalent would be the size of a deck of cards.

But the abundance of IPv6 comes with a significant cost—it is not backwards compatible with IPv4. This means IPv6 deployment must rely on technologies that translate it to IPv4 (and vice versa), tunnel it over IPv4, or run it in parallel with IPv4. The relative complexity of IPv6 has led most enterprises to continue to rely on IPv4 and NAT and slowed the overall pace of IPv6 adoption. While the lack of any overt crisis has disproven the worst-case scenarios claimed by some in the Internet engineering and IPv6 adoption communities, it's also left many IT directors and CIOs with a false sense that IPv6 can be ignored indefinitely.

Since it's not possible for most organizations to migrate to IPv6 exclusively and turn IPv4 off, running IPv6 in parallel on an existing production IPv4 network is the configuration that is most often recommended and deployed. This configuration is commonly known as dual-stack.

Because IPv6 is just over 20 years old, most networking gear now supports this configuration. Running IPv4 and IPv6 on the same IT and network infrastructure offers a way to ensure the availability of critical applications that may not ever support IPv6. And since the depletion of IPv4 will impact different organizations at different times and in different ways, it should prove helpful to define "IPv4 exhaustion" more clearly.

Are We There Yet? The Many Phases of IPv4 Exhaustion

That there is much confusion surrounding IPv4 exhaustion is not surprising, in part because the definition varies depending on who's talking about it and why. The more formal and technical explanation is easy to understand. Let's explore it first before discussing the more subjective and controversial topic of the real-world definition and impact IPv4 exhaustion.

The Internet Assigned Numbers Authority is the global authority responsible for allocating and managing all, well, Internet assigned numbers. These days they don't have a lot to do where IPv4 is concerned, because in February of 2011 they handed out the last of their IPv4 supply to each of the five Regional Internet Registries (RIRs). For each of five global regions (see the below table), the RIRs administer the allocation of IP addresses (both IPv4 and IPv6) and facilitate the policies governing those allocations.

Four of these five RIRs officially exhausted their supply of IPv4 on the following dates:

RIR	**Exhaustion Date (</8 of IPv4)**
APNIC (Asia/Pacific)	April 2011
RIPE (Europe and the Middle East)	September 2012
ARIN (North America)	March 2014
LACNIC (Central and South America, including Mexico)	June 2014
AFRINIC (Africa)	2017 (estimated)

Looking at this you would logically conclude that IPv4 addresses are not available at all in most parts of the world. Depending on how many addresses you needed and of what type, that might be correct. But let's first clarify what the RIRs define as *IPv4 exhaustion*.

The dates above are when the RIRs dropped below their last 16.7 million available IPv4 addresses. But the policies for responding varies from RIR to RIR. For example APNIC, RIPE, and LACNIC can still provide a one-time-only allocation of 1,024 IPv4 addresses to a requesting organization. ARIN could choose to provide a larger allocation, though such requests, as you might imagine, were heavily scrutinized.

As a result, the date on which *total* IPv4 exhaustion happens for each RIR may arrive later for APNIC, RIPE, and LACNIC. Meanwhile, ARIN totally depleted its IPv4 supply in the summer of 2015.

Scraping the Bottom of the Barrel

So no more IPv4 for you, right? Well that depends.

The RIRs may also allow transfers of available IPv4 address space between regions. For instance, a block of IPv4 addresses belonging to ARIN can be transferred to a requesting organization in Asia through an inter-RIR transfer process.

It's worth noting all of this is much more critical for Internet services providers (ISPs) than enterprise networks. ISPs generally need routable addresses to connect mobile and broadband subscribers. By comparison, enterprises still heavily rely on private IPv4 and NAT for their internal networks. The small number of public, routable IPv4 addresses required to connect enterprise networks to the Internet is usually provided by the ISP.

"Psst! Hey Buddy. Over Here. Need Some IP Addresses?"

IPv4 exchanges have sprung up to facilitate the permissible transfer of address space between organizations, when allowed by RIR policies. The IPv4 transfer market is still relatively new but is heating up, with costs ranging from $9 to $20 per address and an average price of around $10. While only service providers generally need publicly routable IPv4 addresses, the economics of trying to add millions of new users at $10 per address has led many of them to adopt IPv6. For instance, in the U.S., Verizon's level of IPv6 traffic is over 60 percent while T-Mobile recently surpassed the 50 percent mark. Other service providers are turning to Carrier-Grade NAT (CGN), a technology that, while extending the life of IPv4, comes with its own deficiencies and costs.

Why You Need to Adopt IPv6

Most enterprises have delayed IPv6 adoption based on a generally valid recognition that their need for routable IPv4 is and will remain minimal. The use of NAT and private IPv4 addresses will continue to be a viable configuration for the foreseeable future. But it will prove a costly mistake for enterprises to conclude that IPv6 adoption can be avoided or ignored entirely. Here are some of the reasons:

IPv6 Is Already Running on Your LAN…And You Can't Turn It Off

All modern routers and switches offer robust IPv6 support. This is great news for more recent adopters of IPv6. Similarly, all modern operating systems likely to be found in a corporate LAN environment—Windows, Mac OS X, Linux, IOS, and Android—have IPv6 enabled by default. In fact, if you're running Windows at all, you cannot disable IPv6 without being out-of-scope for support from Microsoft.

The BYOD phenomenon has resulted in a flood of employee smartphones and tablets connecting to the corporate network. Since most of these devices have IPv6 enabled by default, IT security and access policy needs to be updated to include IPv6. Network management tools need visibility into IPv6 to make sure compromised or unauthorized devices aren't using leveraging IPv6 to cause harm to the network or access sensitive data.

IPv6 Unlocks the Value of Tomorrow's Networks

The tens of billions of devices constituting the emerging Internet of Things hold an ocean of data that is currently offline or isolated. The analytics applied to this data, and the commercial intelligence that results, promises to drive business agility to new heights. None of that can happen without a set of standardized methods and interfaces for connecting IoT devices to existing corporate networks and the Internet.

Only IPv6 offers sufficient address space to facilitate not just the addressing of the devices themselves but also the development of a mature, robust, and scalable design and operational practice to keep these devices online and sharing their data.

A similar challenge exists in cloud-based IT. The full benefit of clouds—efficient, robust, and scalable designs and operational practices—will be difficult to achieve for hybrid and private cloud models using only the limited addressing resources of IPv4. Even the largest blocks of IPv4 private addressing are infinitesimal when compared with the smallest LAN IPv6 subnet. And that's disregarding the operational challenges of managing overlapping private IPv4 address space that can result from having to reuse the same private blocks. By comparison, IPv6 provides unlimited unique addressing.

Get Started with IPv6 Adoption Today

For most enterprises new to IPv6, the map to its adoption seems a lot scarier than the actual terrain turns out to be. And a surprising number of enterprises have already made their websites available over IPv6. Most organizations are able to take a phased approach to IPv6 adoption that makes the transition to a dual-stack network cost-effective, without impacting the existing IPv4 production network—opening up their organization to exciting new opportunities offered by the Internet of Things and cloud architecture.

In short, the end of IPv4 really is in sight, but the IPv6 future is bright.

Infoblox offers excellent resources to get you started with IPv6 adoption. Along with other white papers, our IPv6 Readiness center provides videos, solution notes, and a regularly updated blog with posts from industry-recognized IPv6 experts.

About Infoblox

Infoblox (NYSE:BLOX), headquartered in Santa Clara, California, delivers network control solutions, the fundamental technology that connects end users, devices, and networks. These solutions enable more than 7,000 enterprises and service providers around the world to transform, secure, and scale complex networks. Infoblox (www. infoblox.com) helps take the burden of complex network control out of human hands, reduce costs, and increase security, accuracy, and uptime.

CORPORATE HEADQUARTERS
+1.408.986.4000
+1.866.463.6256
(toll-free, U.S. and Canada)
info@infoblox.com
www.infoblox.com

EMEA HEADQUARTERS
+32.3.259.04.30
info-emea@infoblox.com

APAC HEADQUARTERS
+852.3793.3428
sales-apac@infoblox.com