The transition to IPv6 DNS can be easier

Let's be honest: The prospect of having to move to IPv6 DNS may seem daunting. Like most enterprises, you're probably trying to remain IPv4 only as long as possible.

But if the corporate call comes to jump off the cliff and transition, BlueCat's platform is at the ready to help cushion your landing.

In this post, we'll provide a primer on IPv6 and IPv4, and look at why there's some resistance to the move. Then we'll delve into how BlueCat tools can help ease the transition.

The basics

Let's get some basic terminology out of the way first. IP, which stands for internet protocol, is the internet's principal form of communications. And IP addressing is a logical means of assigning addresses to devices on a network.

What is IPv4?

IPv4, or internet protocol version 4, has been in place for more than 35 years. IPv4 uses 32-bit addresses (for example, 192.0.2.246), to route most of today's internet traffic.

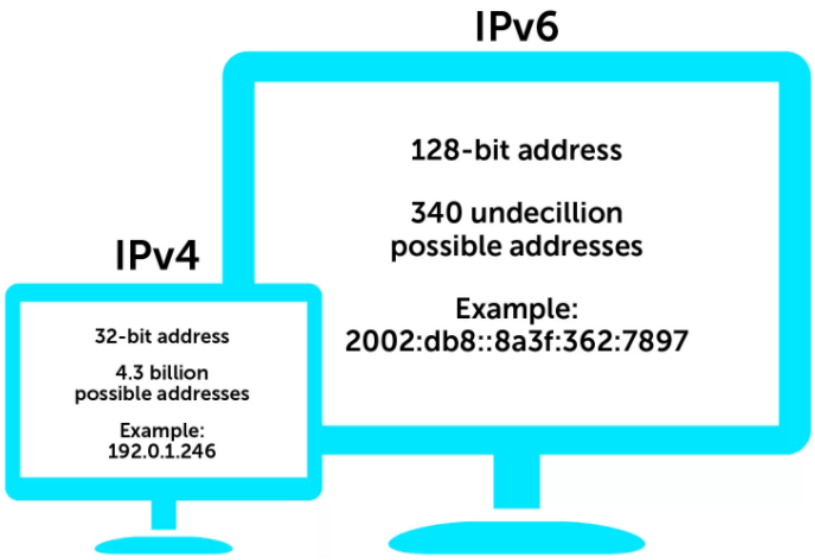

A 32-bit address space limits the number of unique hosts to 2^{32}, which is nearly 4.3 billion IPv4 addresses. But in today's ultra-connected world, 4.3 billion isn't nearly enough.

In 2011, the Internet Assigned Numbers Authority (IANA), the global coordinator of IP addressing, ran out of IPv4 addresses to allocate to regional registries. Since then, regional registries have exhausted those allocations.

In short, we've run out of address space.

Seeing that this would be a problem, IANA's Internet Engineering Task force (IETF) came up with a new way of providing Internet Protocol (IP) address assignments.

Enter IPv6

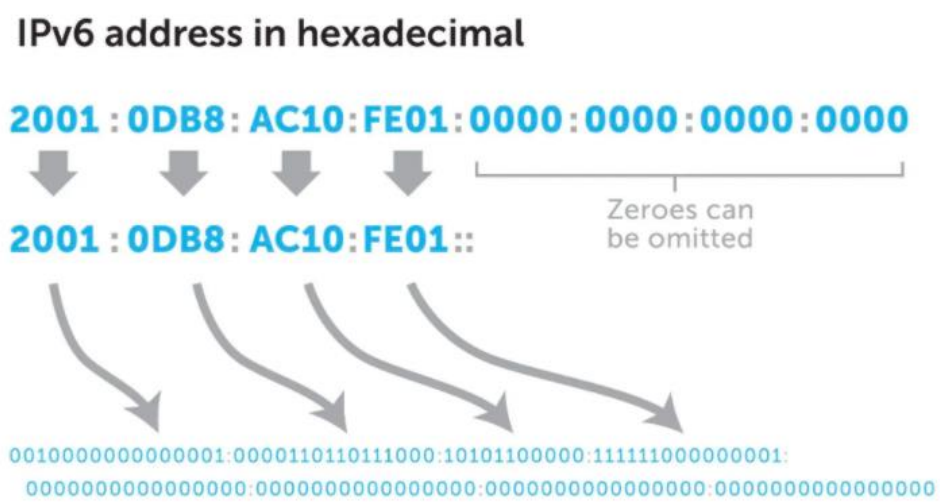

IPv6, the most recent version of the internet protocol, uses 128-bit address space. Unlike IPv4, both letters and numbers are used as identifiers (for example, 2002:db8::8a3f:362:7897). By implementing these changes, IANA created 2^{128} new IP addresses, which is about 340 undecillion or 340 billion billion billion billion. A whole lot.

With IPv6, a single network can have more addresses than the entire IPv4 address space. IPv6 exhaustion is basically impossible. (There is a hypothetical world IPv6 exhaustion counter out there. Nine million AD, anyone?).

Furthermore, routing tables are simpler. Admins can start from square one and be thoughtful and logical about deploying an addressing scheme. And there's plenty of room to add more.

Security was also at the forefront when the IPv6 address space was built, while IPv4 has modern-day security measures tacked on after the fact. However, that's not to say that you get a free pass to omit IPv6 space from your network security model. And the first IPv6 DDoS attack served as an important reminder.

Eliminating private networks

About 18 million IPv4 addresses were set aside for private addressing, drawn from a range known as RFC 1918. Most organizations use IPv4 private addresses on internal networks. However, devices with private addresses have no direct path to the public internet.

To access the public internet, these devices require a complex and resource-intensive workaround called network address translation (NAT).

IPv6 is NAT-free, enabling every device to communicate directly without intermediary steps.

The challenges of implementing IPv6 DNS

All of this change was born of necessity, but not everyone is on board. This is not just a configuration change. Think of it more like a challenging system migration.

Examples of IPv6 challenges

- **IPv4 and IPv6 are not directly interoperable.** Deployment of IPv6 is completely different from IPv4, requiring a steep learning curve to master. IPv6 is managed differently than IPv4, requiring a steep learning curve to master. IPv6 address formats are also longer, so they can't be easily memorized or transcribed.

- **It's a lot of work to test all of your applications end-to-end in an IPv6 environment.** And what may work well in a small test lab may fall apart when implemented at scale.

- **Every part of your network chain (including every IPv6 DNS server) has to be compliant.** Legacy network applications or devices hard-coded for IPv4 may lack IPv6 support.

- **Specifically, most IoT devices do not support IPv6.** If critical IoT devices on your network aren't IPv6-ready, then you can't transition your network at all. This a particularly tough conundrum for the healthcare industry.

- **Tertiary content addressable memory (TCAM) quickly gets depleted when adding IPv6 addresses.** TCAM stores access control lists on network routers. Routing vendors have allowed admins to tune how much TCAM to allocate to IPv4 and IPv6, with mixed results. Ultimately, enterprises end up having to buy more pricey TCAM.

Complex enterprise implementation

Enterprise implementation itself can be complex, with segmented steps and testing required at each point.

You might first start with your external-facing networks and services like web servers. Then go to your perimeter (or DMZ) networks and your data centers. And finally, your internal networks and devices. Just like your current network, you'll need an IPv6 nameserver, DNS server, and all the rest.

It's enough to say "thanks but no thanks" and stick with IPv4. Sure, the more workarounds that you add to your IPv4 network, the more you have to manage. But it works, you understand it, and you know the network won't break.

There is no driving event such as a government mandate forcing the transition *en masse*. As a result, institutional inertia will be strong enough in most organizations to keep the status quo in place. The work involved in a transition simply isn't worth it... yet.