



All servers on the Internet, such as these, require a unique address from a limited pool and as of 2015 they have all been claimed.

On the Internet, four billion is not enough

BY JASON POTTERF

ALL PHOTOS: © JASON POTTERF

THE ESTA CONTROL PROTOCOLS WORKING GROUP stood on the shoulders of giants when it developed its first generation of Internet-Protocol-based lighting control standards. ESTA standards take advantage of decades of technological innovation in groups like the Internet Engineering Task Force (IETF), which develops the protocol standards that are the foundation of the Internet. Now, that foundation is in the process of a long overdue upgrade. Yes, that's right, the Internet is upgrading from Internet Protocol Version 4 (IPv4) to Internet Protocol Version 6 (IPv6). Please, try to contain your excitement. Why is this taking place? Because 4 billion wasn't enough.

Addressing the issue at hand

Just as sending a letter through the post office requires an address to ensure proper delivery, the Internet requires every connected computer to have a unique address. The addresses are, like everything in computing, binary numbers of finite length. Vint Cerf, the father of the Internet, settled an argument in the 1970s by setting

the address field length for IPv4 to 32-bits, which allowed for 4,294,967,296 unique addresses. In the 1970s it was unfathomable that there would even be four billion computers in existence, much less that they would all be connected to each other. Nonetheless, that's exactly what happened and it happened faster than anyone imagined. As of 2015, almost all of the four billion available addresses are in use. A tiny percentage in Africa is still up for grabs, but those are expected to be gone by 2019.

When the giant moves, so do we.

This problem, known as IPv4 address exhaustion, was foreseen by Cerf and his collaborators. Two approaches were identified early on to deal with this problem. The first was a form of address recycling, called Network Address Translation (NAT), which allowed many computers to share one IP address. NAT bought everyone some time. The permanent solution is to move from IPv4, which was originally designed as an experimental protocol, to the production-ready solution: IPv6. IPv6

allows 340 trillion trillion trillion addresses and brings with it a better Internet for all the people and things that it connects. But what does all this mean for ESTA?

Impact on architectural lighting

Entertainment lighting has been in transition to Ethernet networks for over a decade, but the architectural lighting industry is just now undergoing a major transition to Ethernet. Today you can buy an array of architectural lighting fixtures that get both their lighting control data and power via standard Category 5 Ethernet cables. Many vendors of these Power over Ethernet luminaires are choosing to run ahead and adopt an IPv6-only policy. Thus, when it came time a few years ago to choose control protocols for the next generation of IP-based architectural lights, only those protocols supporting IPv6 were considered. Various IoT protocols such as Constrained Application Protocol (CoAP) and Message Queue Telemetry Transport (MQTT) were favorites for this application. CoAP is showing the most promise, and several

PoE lighting vendors are choosing it as the path forward. This topic is deserving of its own discussion, but for the purpose of this article one must only appreciate that the driving force behind IPv6 adoption in architectural lighting is simple: assigning a unique IP address to every light in every building across the world is only possible using IPv6's massive pool of addresses.

typically utilized IPv4 address ranges that are explicitly forbidden from being carried across the Internet. Given this, why slow down the already excruciatingly slow standards process with a new technology when the old one works fine?

The problem with this mindset is related to our aforementioned standing on the shoulders of giants. When the giant moves,

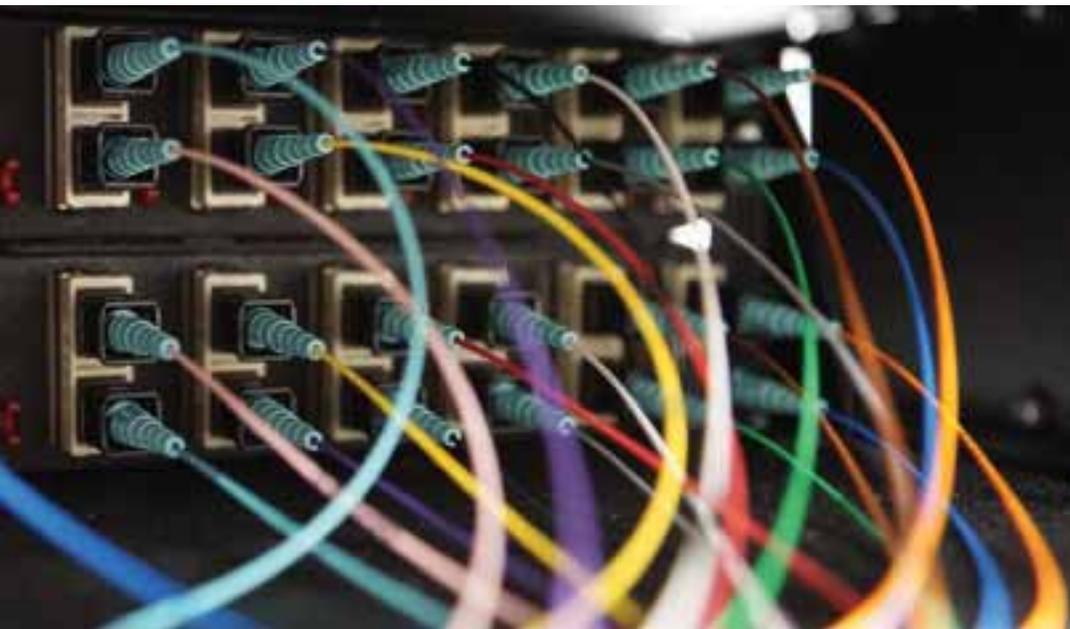
not without its early adopters of IPv6. MA Lighting saw the benefits of IPv6 early on, and has been using it for their inter-console communications for many years. Customers are amazed by the new plug-and-play features offered by IPv6. However, MA Lighting is still stuck with IPv4 for device control until *E1.31* is revised.

The good news is that adding support for IPv6 to ESTA protocols is fairly trivial, and a minor revision to the existing *ANSI E1.31 – 2016 – Lightweight Streaming Protocol for Transport of DMX512 Using ACN*, is currently out for public review. Depending on the comments received, *E1.31* may become the CPWG's first IPv6 enabled protocol standard as early as the end of 2017. *E1.33 RDMnet* has incorporated IPv6 support into its most recent draft and should support IPv6 upon becoming a standard. Some additional work is required to round out ESTA's IPv6 compliance strategy, but these two changes should enable the vast majority of entertainment lighting networks to support IPv6 operation. The ecosystem for developing IPv6-capable devices is well established. Further, with modern technology, this comes with little to no cost increase to ESTA members.

We've established that IPv6 is a good thing, but a reasonable question is: What about IPv6 has everyone moving in this direction? The answer is two-fold. First, it solves the addressing problem. Second, it adds new features that improve the user experience.

Can you count to 340 undecillion?

IPv6 once and for all solves the addressing problem. For IPv4 back in the 1970s, the math worked out as follows: Vint Cerf's team assumed that due to the enormous cost of computers and networking equipment at the time, only two networks per country were feasible. They guessed roughly 128 countries would have them, for



Fiber optic cables like these, made up of dozens of tiny strands of glass, connect the billions of computers around the world at the speed of light.

Traditionally, the cross-over between architectural and entertainment lighting has been minimal. The two markets in the past had negligible overlap, and the requirements of their control protocols were very different. Unfortunately, due to the architectural market's leap ahead with IPv6, this is likely to continue into the next generation of IP-based lighting control.

Impact on entertainment lighting

In the past, the ESTA Control Protocols Working Group has largely ignored IPv6 as irrelevant to the entertainment lighting industry. After all, most lighting networks are purposefully not connected to the Internet. In fact, lighting networks

so do we. In order to ensure the continued viability of the Internet, network equipment manufacturers, electronics vendors, and software companies are all shifting their attention to IPv6 as the primary path forward. These are the same companies that make the building blocks for our entertainment lighting systems. Most products today support what is known as "dual stack" IP, which support both IPv4 and IPv6 simultaneously. The customer then gets to choose IPv4, IPv6, or both. The bad news for the IPv6 naysayers is that customers are already demanding that IPv4 be turned off. This is happening in the architectural world, and by not supporting IPv6, ESTA protocols, and thus ESTA members' products, are shut out of those opportunities.

The entertainment lighting industry is

a total of 256 networks. Each network would have no more than 16 million computers per network, and 256 networks of 16 million computers would need 4 billion addresses.

[IPv6] adds new features that improve the user experience.

There were only a few problems with this math. First, there are more than 128 countries, and on today's Internet, every company and household operates its own network. Also, in 2017, Apple sells more than 16 million iPhones each month, and each requires its own unique IPv4 address. To be fair, this could not have been foreseen back then, but we're stuck with the problem none the less.

In the 1990s when IPv6 was designed by the IETF, they had learned from the mistakes of IPv4. In doing so, they chose an address size of 128 bits, double the number of bits generally accepted as sufficient for the foreseeable future. In fact, 64 bits would have allowed 18.4 quintillion (that's 18.4×10^{18}) addresses, which is enough to give a unique address to all eight billion people on earth today and every hair on their body while only using 0.2% of the available IPv6 address space. The final IPv6 address space of 128 bits allows 340 undecillion addresses, or $340 \times 1 \text{ trillion} \times 1 \text{ trillion} \times 1 \text{ trillion}$. Numbers this size defy comprehension, but suffice to say, we're unlikely to run out of addresses with IPv6, so we will only have to change Internet protocols once.

IPv4 headaches of the past

Before discussing the new features that might motivate you to adopt IPv6, let's take a moment to remember what we suffer with now in IPv4. As you may know, in order to get a device on the network, you first have to

set the IPv4 address. This is not always easy.

If you're lucky, there is a Dynamic Host Configuration Protocol (DHCP) server running on the network that does this dirty work for you. However, the address assigned today may not be the address assigned tomorrow, and if the DHCP server is out of IP addresses because the precious few were already consumed by other devices, you are out of luck.

The alternative is static IP address assignment, which requires the user to input all the parameters by hand. Anyone who has tried to set an IPv4 address on an LCD display with just four buttons will agree this process is ripe for improvement. You have to specify a unique IP address, and if you get this wrong, you could knock another device offline by accident. You also needed a subnet mask, which consists of a seemingly inexplicable number of 255s followed abruptly by a zero that seemingly has nothing to do with the IP address. Then you need to know your gateway IP address, DNS server IP addresses, and your domain name. Then, and only then, could you finally connect to the Internet and successfully use its vast resources to download cat videos.

IPv6 advantages for the future

The user experience is much improved in IPv6. If you don't require Internet access, and most lighting networks do not, you simply plug all your devices into the network, and you're done. Each IPv6 device upon connection to a network assigns itself a guaranteed-unique link-local IPv6 address. No router needed, no DHCP server required. This auto-generated link-local IPv6 address is all that's necessary for each device on a link to talk to its neighbors. That said, you can't download a cat video with a link-local IPv6 address, though. For that you need Internet access, and for that you need to be assigned a Global Unicast IPv6 address.

One key feature of IPv6, and a major reason for the huge number of bits assigned, is the network prefix. Much like the postal

system uses a zip code to get a piece of mail to the right post office for final sorting, the Internet uses the Network prefix in an IPv4 or IPv6 address to get a packet to the correct router. IPv4 achieved this in a very convoluted way due to the small number of address bits available to the protocol. This scheme required that a subnet mask must be entered alongside every IP address.

But since IPv6 has bits to spare, it can afford to waste a few and dispense with all of this subnet mask tomfoolery. The network prefix and a user customizable subnet ID are always in the same places in the first 64 bits of a Global Unicast IPv6 address. This is then followed by the interface ID in the last 64 bits. Because of this, you will never have to enter a subnet mask in an IPv6 network.

Through hard work, and a little luck, the ESTA CPWG will have its first IPv6 compliant protocol standardized this year.

As an aside, IPv6 addresses look very different when compared to IPv4 addresses. An IPv4 address is noted in dotted decimal format. A common one that you may recognize is 192.168.0.1 and is used on most home routers. In IPv6, there are four times as many numbers, and you can thank the architects of the Internet for not forcing us to type 32.1.13.184.0.0.0.0.0.0.0.254.237.250.206. They defined a shorthand, which allows you to write this address as 2001:DB8::FEED:FACE and is only marginally longer than an IPv4 address. Congratulations! You know how to



A view inside a high-end router's supervisor engine. Routers like these can handle both IPv4 and IPv6 traffic simultaneously and at high speed.

recognize and write an IPv6 address. Now go ahead and forget it, because IPv6 was designed so you will never need to type one in.

To enable access to the Internet via IPv6, you need an IPv6 capable router. For most users, this router is provided by your service provider or IT department, and a growing number of service providers already support IPv6. Being the router in charge of the local link, it is able to answer questions for attached devices. It knows what the network prefix should be for the local link. It knows what the gateway, a.k.a. default router, address should be for all attached devices. It can also advertise the DNS server addresses and domain names for the local

link. All of this is communicated to the end device shortly after it is connected, giving it the full gateway and DNS server addresses, the domain name, as well as the first half of its IP address. Thus the end device has everything required to connect to the Internet and download feline films.

But wait—you said it only received the first half of its IP address! We just learned this contains only the network prefix and subnet ID. What about the other half, that all important unique interface ID? This is the beauty of the 64-bit length of the host ID. Every Ethernet-based device come from the factory with a guaranteed unique 48-bit identifier. IPv6 defines an algorithm for inflating this 48-bit address to fit the 64-bit

```
Frame 34: 70 bytes on wire (560 bits), 70 bytes captured (560 bits) on interface 1
Ethernet II, Src: Mistroni_10:e2:15 (54:ee:75:10:e2:15), Dst: IPv6mcast_02 (33:33:00:00:00:02)
  Destination: IPv6mcast_02 (33:33:00:00:00:02)
  Source: Mistroni_10:e2:15 (54:ee:75:10:e2:15)
  Type: IPv6 (8x86dd)
Internet Protocol Version 6, Src: Fe80::2c48:3c92:5c87:25ac, Dst: ff02::1
  0118 ... = Version: 0
  ... 0000 0000 ... = Traffic class: 0x00 (DSCP: CS0, ECN: Not-ECT)
  ... 0000 0000 0000 0000 0000 0000 = Flow label: 0x000000
  Payload length: 16
  Next header: ICMPv6 (58)
  Hop limit: 255
  Source: Fe80::2c48:3c92:5c87:25ac
  Destination: ff02::1
  [Source GeoIP: Unknown]
  [Destination GeoIP: Unknown]
Internet Control Message Protocol v6
  Type: Router Solicitation (133)
  Code: 0
  Checksum: 0xa50b [correct]
  [Checksum Status: Good]
  Reserved: 00000000
  ICMPv6 Option (Source link-layer address : 54:ee:75:10:e2:15)
```

An IPv6 router solicitation packet as decoded by Wireshark®, a popular network packet analyzer. This packet is used to find the router in charge of the local link.

Interface ID field in a consistent, predictable manner. Not only is this address guaranteed to be unique, but it won't change during the course of the network's operation. Thus as long as a properly configured router is present on the network segment, just a few seconds after plugging your device in, you have access to the Internet. All this without ever typing in a single IPv6 address and without dealing with the hassle of setting up a DHCP server.

Moving the herd

New features alone are not enough to incentivize the whole world to change over from IPv4 to IPv6. Like any new technology, you need early adopters to incentivize the change. In 2005, the US Government required all federal networks to support IPv6 by 2008, forcing all of its suppliers to support IPv6 in their equipment. China, perhaps most impacted by the perfect storm of population growth and IPv4 address exhaustion, has mandated IPv6 support as part of its China Next Generation Internet plan. In 2016, Apple announced that all iOS App Store apps would require IPv6 support or face rejection. Google reports that in 2017, nearly 17% of search traffic is coming via IPv6, up from only 6% in 2015. And through hard work, and a little luck, the ESTA Control Protocols Working Group will have its first IPv6 compliant protocol standardized this year. The transition to IPv6 is inevitable. There is no time like the present to get on board and avoid being left behind. ■



Jason Potterf is a veteran of both the networking and entertainment lighting industries. He currently is a Technical Leader at Cisco where he leads hardware development for the Cisco Digital Building family of Power over Ethernet

Switches. He is a long-time member of ESTA's Control Protocols Working Group where he is the Co-Chair of the E1.31 Streaming ACN IPv6 Update Task Group, and the chair of the IPv6 Task Group. He can be reached at jason.potterf@cisco.com.