

Number Crunch

by Jonathan Zittrain, Leslie Daigle
January 10, 2011

How the Internet is running out of room, and what we must do about it

"CDT Fellows Focus" is a series from CDT that presents the views of other notable experts on tech policy issues. This week, CDT Fellow [Jonathan Zittrain \[1\]](#) and Leslie Daigle write about the end of IPv4 address space. Guest posts featured in "CDT Fellows' Focus" don't necessarily reflect the views of CDT; the goal of the series is to present diverse, well-informed views on significant tech policy issues.

The Internet's framers famously designed it without predicting much about how, or how much, it would be used. For example, the network's capacity was conceived less in a count of precisely how many could participate at once – the way traditional phone circuits worked – and more in flexibly divisible bandwidth. As that bandwidth got saturated, it would degrade gracefully: data might move slower for everyone, but no one would get an "all circuits are busy" message. In ways large and small, what animates Internet protocol design is a [procrastination principle \[2\]](#): if something can work well, it doesn't have to be perfect, and not every problem or limit must be anticipated and preempted. Potential but still speculative flaws can be fixed later – possibly somewhere other than inside the network.

Unfortunately "later" is arriving now for a crucial piece of the Internet: its ability to tell one attached device from another. Internet architects designed a simple way to identify participating computers and route data among them: assign each a unique number: an Internet Protocol (IP) address. No IP address, no delivery. The routers in between you and your friend use your friend's number the way a postal service would – the number says something about where she is. That's made possible because IP addresses are clustered together, just like street addresses grouped in a ZIP or other postal code.

The system has an Achilles' heel: there are a limited number of numbers. It might seem that you could add 1 to whatever the last number is and keep going, but there's a hard cap in venerable Internet Protocol Version 4 (IPv4): 4 billion IP addresses, which the Internet is outgrowing in much the same way that applications outstripped the original 64K of memory expected for a PC running Microsoft DOS. There is now general agreement among Internet technologists that the end days are upon us: the last block of fresh IPv4 addresses will likely be allocated to the Internet's North American address warehouse [in early 2011 \[3\]](#), to be passed out to Internet Service Providers here by mid- to late-2011.

Worse, because of the clustering of addresses, we can't squeeze the last bit of digital toothpaste out of the tube as fresh numbers become scarce. There's a gray market for chunks of already-allocated numbers despite restrictions against selling them – and some telecommunications providers are rumored to have been purchased only for their numbers! – but such used numbers carry their own risks. Anyone who has inherited the former phone number of a pizza shop will appreciate that some numbers are less desirable than others. Moreover, some IP addresses have at one time been the source of cyberattacks, or hosted politically sensitive content, and the resulting blocking of traffic originating from them by various ISPs is rarely revisited by those ISPs. (Who wants an old Wikileaks IP address?)

Running out of fresh numbers will not stop the Internet from working. But, unchecked, it will greatly complicate growth. As new computers and devices come online, something has to give – making more use of existing addresses, or finding a new way to address things.

In the first category – making do – the procrastination principle has bought us some time. Enterprising engineers developed an ingenious baling-wire-and-twine workaround to the one-number-per-computer rule. Known as network address translation, or NAT, it allows the holder of

a single IP address to share it among a group of computers. This happens nearly every time you hook up a wireless router and access point at home: your ISP only gives you one number, and you use your inexpensive router to share it with everyone who connects to your network. Cable and DSL ISPs are considering the same thing to put larger networks of multiple customers behind a single address, at least as an [interim measure](#) [4]. Unfortunately, like most such workarounds, it doesn't really work as well as having one number per machine: the fancy footwork required to share a number around can limit the kinds of applications you can run, and greatly increase the complexity of some software, such as Skype Internet telephony, if it's to work at all. NAT has bought us some time – much of Qatar has been known to share one IP address – but it's spackle covering a rapidly-rusting architecture stretched far beyond its creators' wildest ambitions.

Which brings us to a more comprehensive solution. Internet technologists did not sit idly by when it became clear IPv4 could not last. Over a decade ago, they specified its successor, IPv6 (don't ask what became of IPv5), with a few hundred trillion trillion trillion addresses. Such huge swaths of address space promise something even better than a well-functioning market for valuable but limited assets: abundance so great that no market is required, only careful administration. Unfortunately, for IPv6 to work, nearly every piece of networking software and hardware from one end of a data transmission to the other needs to be upgraded. If just one link in the chain hasn't been upgraded to understand the new numbers, IPv4 will still have to be used.

The idea for transition was that systems would work with both protocols for awhile, and gradually IPv4 would end not with a bang, but with a whimper – fading away like, say, the telegraph or telex addresses that used to share letterhead with telephone and fax numbers. However, even though many operating system and hardware vendors have been anticipating IPv6 for years (current Mac and Windows systems now support it out of the box), there are still gaps in available products and little business dependency on it, and there has been remarkably little deployment. This is consistent with the procrastination principle: the only networks that have deployed IPv6 are those that have found a business model for which it is a requirement. And, because the benefits are, generally, global rather than local to one network, the procrastination principle becomes a Prisoner's Dilemma: we're all better off if we all move to IPv6, but the worst case is if you pay to move while others don't. So why not wait – forever, if others act similarly – for everyone else to do it before making the investment?

We've spent a decade with few networks taking the plunge to deploy IPv6.

This holding pattern is not likely to persist. With the larder dry, in the absence of fundamental innovation in Internet Protocol, we'd see an unfortunate ramp up in the use of NAT and its complications, coupled with parties' tussles over existing 'pure' IP addresses like rats fighting over crumbs. Demonstrated shortcomings of the type of IPv4 address sharing include degraded performance of network-intensive web services: web pages where different pieces show up slowly, rather than seamlessly. Customers will not see a poor network connection – they will perceive poor service from the product or company.

More directly, IPv6 is gaining ground among new entrants (who have little choice), so the days of an all-IPv4 Internet are numbered. In developing its broadband strategy, India went for IPv6. New industries looking at wide scale networking are also looking to IPv6 in order to have access to adequate address space, and to be able to build novel network architectures, unencumbered by the structural assumptions needed [to support address sharing](#) [5].

The best future for the Internet is for all networks to deploy IPv6, and pay the price of working in a dual IPv6 / IPv4 world for a period of transition. If companies wait until the business impacts of degraded IPv4 network experience or the identification of opportunities to work with new (IPv6) networks are upon them, the need to make a transition more quickly than a multi-year equipment refresh cycle will likely be more costly and difficult. So how to encourage enough entities to take the plunge?

One way out of a classic problem demanding collective action is through regulation. A government can incent or compel everyone to contribute. However, this would require coordinated regulation across boundaries not recognized by network traffic – the intricacies are daunting, and for the

Internet without precedent. And if successful, governments might gain an appetite for controlling the direction of an Internet which previously managed growth and innovation through elective uptake. Few are enthusiastic about mandated transitions.

Another way out is through leadership by big players. For example, governments aren't just regulators of information technology, they're purchasers of it. By insisting that government- and military-run subnetworks are IPv6, they'll stimulate demand for the newer technologies and encourage intertwined private parties to follow suit. The US government's Office of Management and Budget followed just such a route in 2005, requiring all government services to be [IPv6 capable by 2008](#) [6]. In September, Vivek Kundra crystallized requirements for government websites to be [IPv6 capable](#) [7].

China has been leading IPv6 adoption for years, in part because it may otherwise feel the IPv4 number crunch most acutely, and perhaps because the government has determined that it's in the country's [best commercial interests](#) [8]. Some large companies have placed bets on an upgrade. Google has been public about its activities to deploy IPv6, and a business rationale to not be last to market with [IPv6 support](#) [9].

A cold calculus on such investments for many Net-connected enterprises may indeed suggest holding off. But what has made the Internet better than the more proprietary networks that it eclipsed is that its participants have had a sense of stewardship of the space, justifying the absence of government planners and sheriffs, or a single corporate umbrella. Engineers from the public and private sectors labor on Internet protocols with loyalty to a network functioning as a commons, not simply to their employers' particular business models. An investment in IPv6 from enough corners is sensible if each corner decides to factor in the benefit to the overall ecosystem – not just itself.

If such capacious thinking comes through, the Internet won't run out of space – and we can go back to procrastinating on its future.

Jonathan Zittrain is Professor of Law at Harvard Law School, where he co-founded its Berkman Center for Internet & Society, and Professor of Computer Science at the Harvard School of Engineering and Applied Sciences. He is a member of the Board of Trustees of the Internet Society. Leslie Daigle is Chief Internet Technology Officer for the Internet Society.

-
- [IPv6](#)
- [IPv4](#)
- [CDT Fellows Focus](#)

Copyright © 2013 by the Center for Democracy & Technology. This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. CDT can be freely copied and used as long as you make no substantive changes and clearly give us credit. [Details.](#)

Source URL: <https://cdt.org/blogs/cdt/cdt-fellows-focus-number-crunch>

Links:

- [1] <http://cdt.org/personnel/jonathan-zittrain>
- [2] <http://yupnet.org/zittrain/archives/10#49>
- [3] <http://www.potaroo.net/tools/ipv4/index.html>
- [4] <http://www.networkworld.com/news/2008/072108-comcast-ipv6.html?page=2>
- [5] <http://www.ipso-alliance.org>
- [6] <http://www.networkworld.com/news/2005/080105-ipv6.html>
- [7] <http://gcn.com/articles/2010/09/28/kundra-sets-new-ipv6-deadlines.aspx>
- [8] http://news.cnet.com/China-launches-largest-IPv6-network/2100-1025_3-5506914.html
- [9] <http://www.networkworld.com/news/2009/032509-google-ipv6-easy.html?hpg1=bn>



Number Crunch

Published on Center for Democracy & Technology (<https://cdt.org>)
