In a new Internet, make new mistakes

Fred Baker

FG2030 Workshop 2018-12-18

Overview

- I have been asked to give you some benefit from having been involved in Internet development since the 1980's.
- This exhortation has two important parts:
  - The **admission** that mistakes were made in the development and deployment of the Internet as we know it
  - The **observation** that many of them were and are easy to make – and seem to be happening again.
- The fundamental recommendation:
  - **Learn from the past…**

Let me give you an example

- Content Centric Networking (CCN)
  - Information Centric Networking (ICN)
  - Named Data Networking (NDN)
- Van (whom I respect immensely) comments that CCN is immune from certain kinds of DDOS attacks because senders do not "**send**" data to anyone they please; receivers **invite** (express **interest** in) data they want to receive
  - Yes, but; even in today's IPv4/IPv6 network, receivers request data that is discovered after the fact to contain malware
  - Consider electronic mail, using IMAP/POP/Exchange/etc.
- So CCN/ICN/NDN may limit the effect of certain kinds of DDOS at the network layer, but it continues to permit attacks at the application layer in much the same way the current Internet does.

Example in pictures

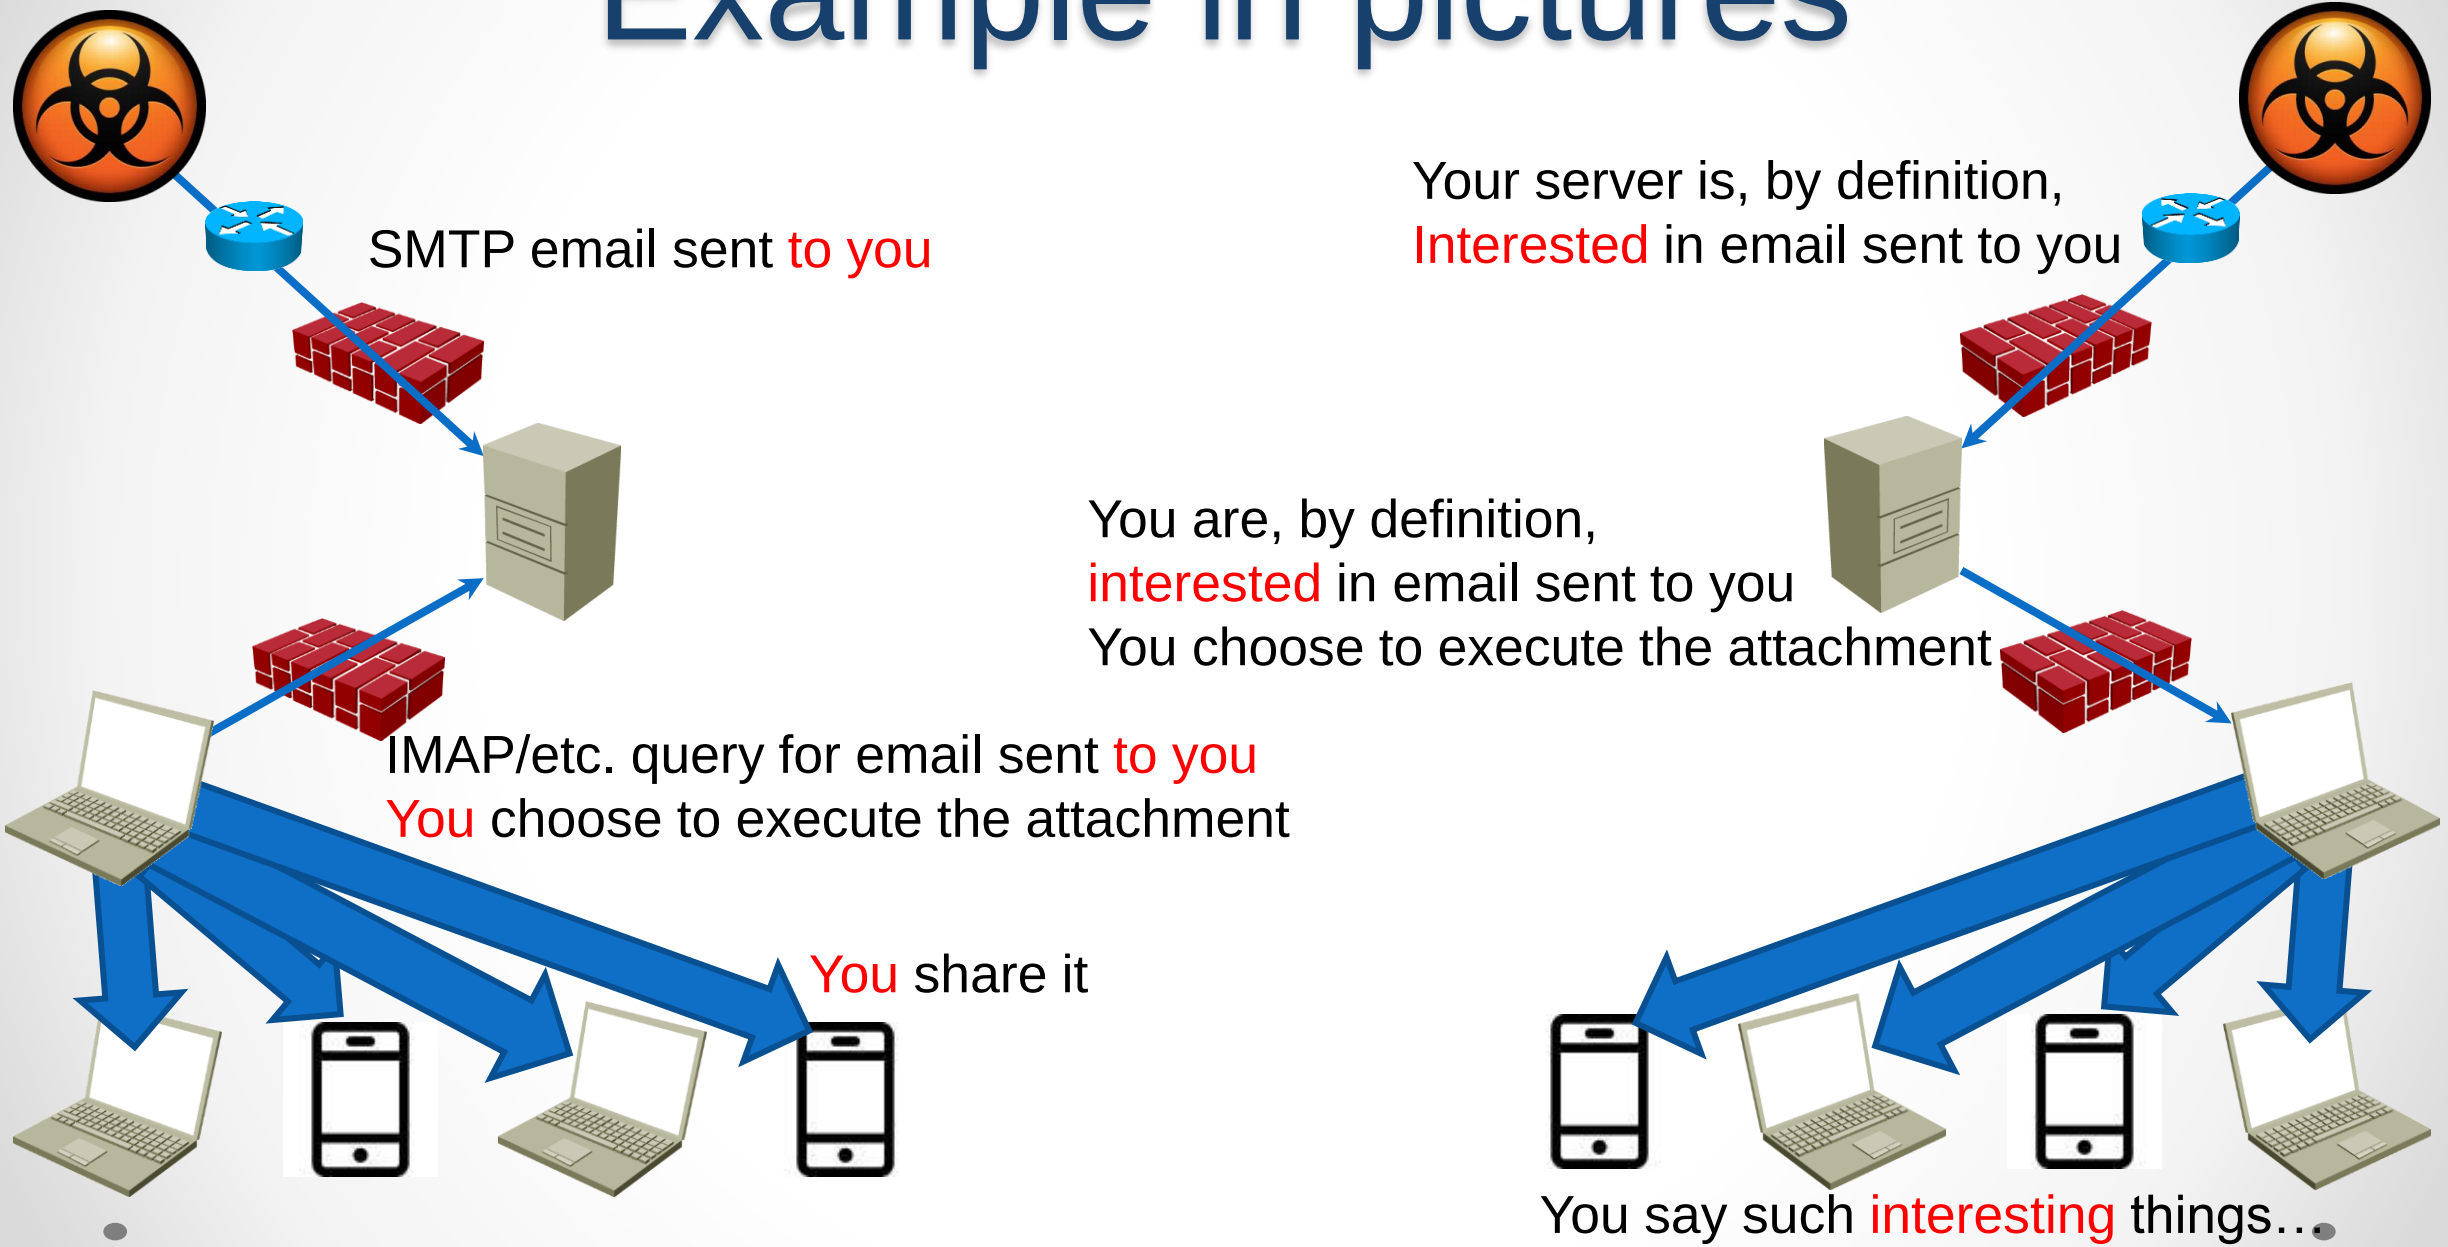

What was the mistake?

- In part, incomplete analysis –
  - Ray Tomlinson, in 1974, had no idea that he was creating a security vulnerability
  - He thought he was enabling human communication in a manner comparable to postal mail
- Confirmation Bias
  - Lack of belief or understanding that you might be wrong
  - *Common among engineers and researchers*
- The indicated correction – test assumptions for correctness
  - “Red Team”
  - “An independent group that challenges an organization to improve its effectiveness by assuming an adversarial role or point of view.”
  - Common use of prototypes and test programs

Second mistake

- FG2030 asserts that it wants to create a new Internet, using new technology
- What does it mean to create a new “Internet”?

- The Internet is a *commercial service*.
  - If the new Internet is not, it will have fundamentally failed.
  - If you are not thinking in terms of operating a commercial service, you’re not going to (at least intentionally) create a replacement for the existing commercial service
  - You are therefore not recognizing that what you create will become a commercial service and has the requirements of a commercial service.

What does it mean to be a "commercial service"?

Commercial

- Concerned with or engaged in commerce: *a commercial agreement*.
- Making or intended to make a profit: *commercial products*.
- Having profit rather than artistic or other value as a primary aim: *their work is too commercial*.
- https://en.oxforddictionaries.com

Service

- A system supplying a public need such as transport, communications, or utilities such as electricity and water:
  - *a regular bus service*.
- In the Internet, a "service" is generally a repeatable offering from an ISP or other provider that one can purchase and have installed "right now", and they will be able to maintain.

A little history

- The Internet, when it was first created, was a research experiment exploring several theories.
  - Among others, Len Kleinrock's 1961 thesis asserted that a multipath packet network with dynamic routing among imperfect network elements could be a network its users would find reliable.
  - It was not intended to be *or become* a commercial service

- In the latter 1980's, Internet Service Providers decided to make the offering commercial.
  - Surprise! They had a plethora of new requirements that had not been directly envisioned or addressed before.
  - This is not because the Internet's designers knew what they were doing and did a bad job. This is because they were learning as they went along (the purpose of research) and made many mistakes along the way.

https://www.lk.cs.ucla.edu/data/files/Kleinrock/Information%20Flow%20in%20Large%20Communication%20Nets.pdf

Issues that came up commercializing the Internet

- Routing had to be dynamic and yet predictable, and we needed to be able to control it.
- MANRS: it needed to be possible to identify and reject an inappropriate route
- CIDR: we changed the structure of an address, and had to change all of the routing protocols
- DHCP: we needed to be able to centrally configure everything in a network
- We needed to be able to prove that communication was between the entities we thought it was between, or prove that it wasn't
- We needed to be able to manage, disregard, and avoid attack traffic and malware
- We needed to ensure that the network traffic received predictable and specifiable service
- We needed to be able to configure and manage changes to the network
- Congestion management was important in several forms
- Every application couldn't be built from scratch
- We needed to be able to innovate without permission or overhead from governments or IT managers
- We needed to be able to hide information when appropriate.
- We needed to be able to identify traffic that was inappropriate or hid information inappropriately.
- It became important to associate names with addresses and other information.
- Resources (such as address space, the ability to use names, etc.) needed to be scalable and sufficiently plentiful that they would not become exhausted.
- Governments had requirements of various kinds
- Lower layer services and entities couldn't be allowed to subvert the intention of high layer entities.

IPv6 will be important in 2030

I say this because people try hard to ignore the fact

History of IPv6

- 1990: IETF realized that IPv4 address space would run out
  - Took steps to alleviate that
- 1993, IETF requested proposals for “next generation” protocol
  - IPv6 proposed in 1994 (among other proposals)
  - Translation proposed in 1995; extended life of IPv4 Internet by ~15 years
  - IPv6 standardized in 1998
  - Research trial deployments…
  - Supporting work in DHCP, DNS, routing protocols, etc
  - Implementation in various operating systems; Windows late
- Uptake of prefixes started 2007
  - ICANN policy for prefix allocation 2006
  - Tokyo University report on reality of IPv4 exhaustion predictions
- IANA allocation of last IPv4 prefix in 2011

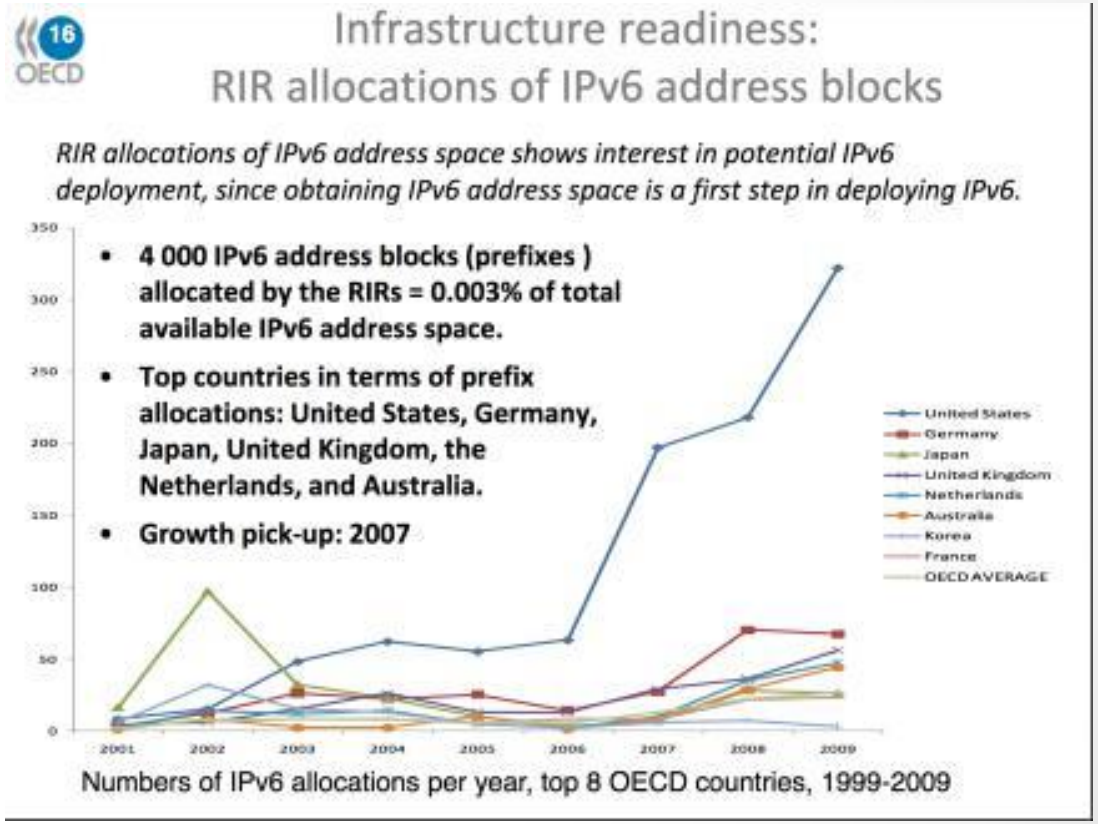

https://www.oecd.org/sti/ieconomy/44961688.pdf

Exhaustion timeframes

- All RIRs have now entered their respective end phases
  - Lots of IPv6 prefixes to allocate
  - IPv4 only for new entrants, and then in small quantities

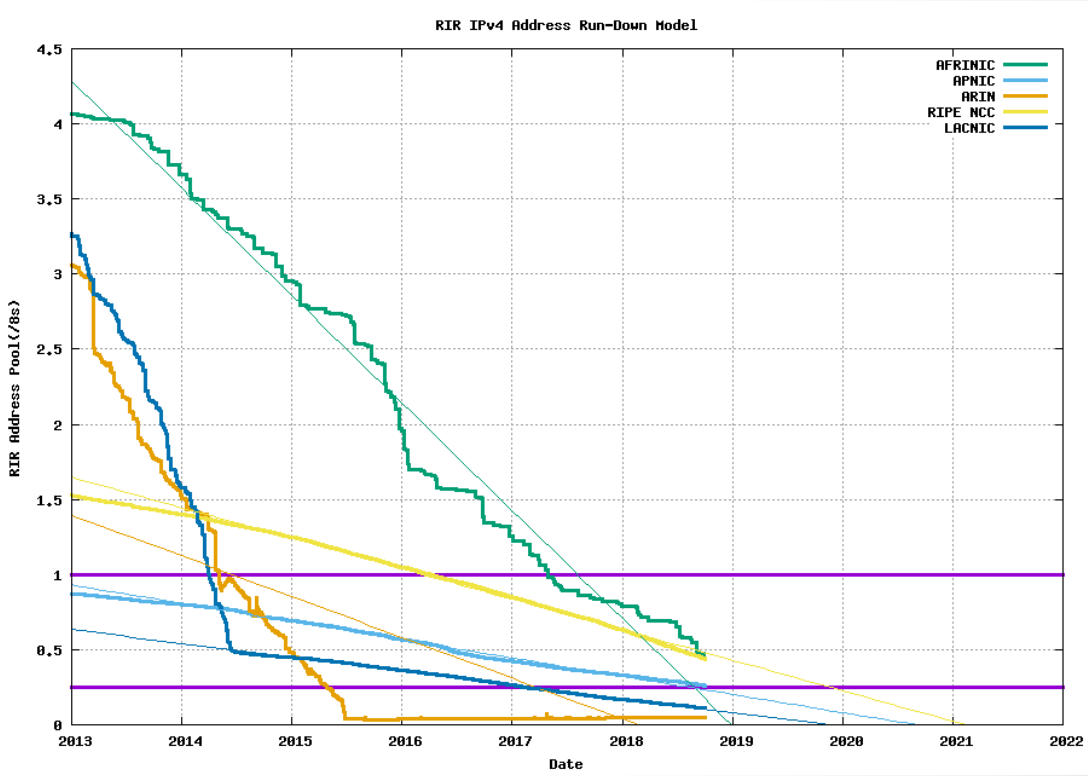

https://ipv4.potaroo.net/plotend.png

Why not make a market in IPv4 addresses?

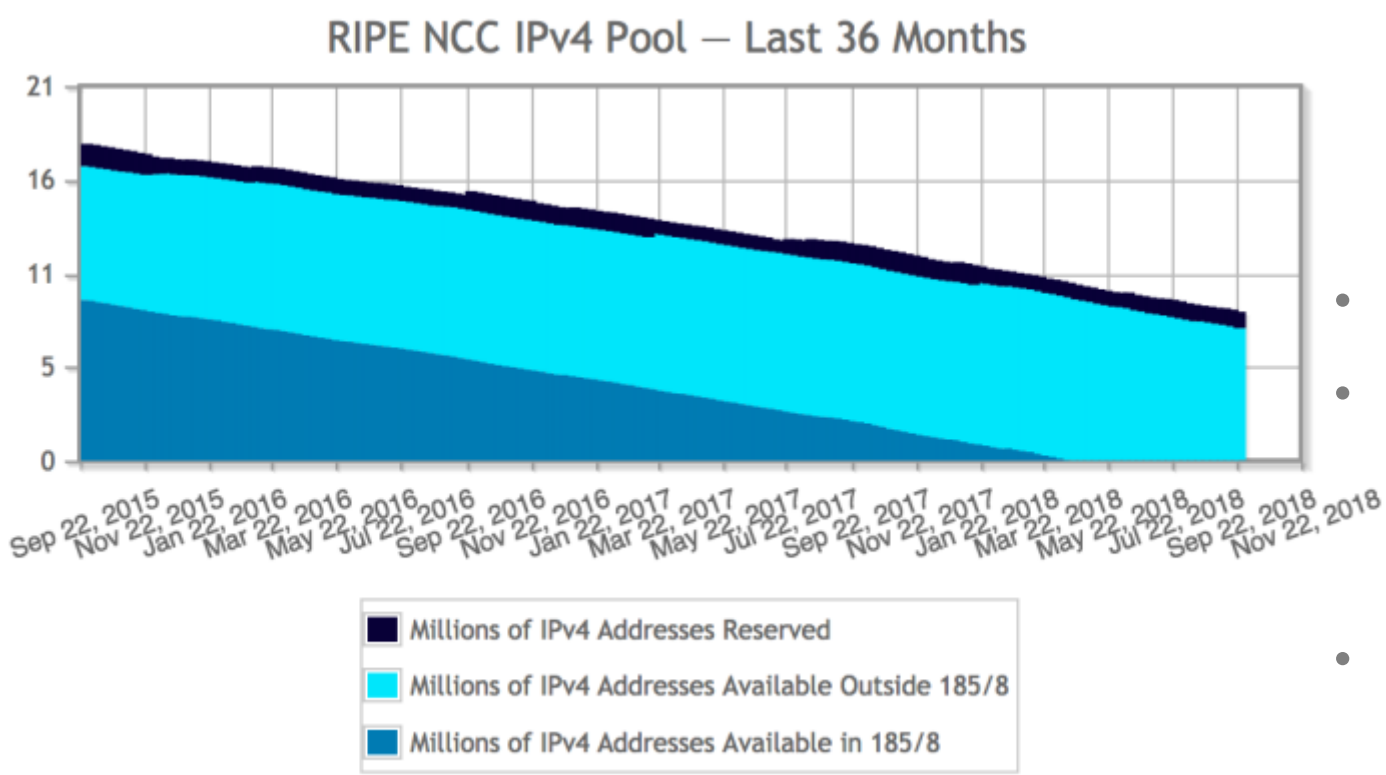

- People have.
- 90% of sales are to CDNs and large social media sites.

- The largest blocks have already been sold; increasingly, the blocks that remain are small ones, which can be difficult to manage.

https://www.ripe.net/publications/ipv6-info-centre/about-ipv6/ipv4-exhaustion/ipv4-available-pool

Sure, but it's time to sell, not to buy

http://ipv4marketgroup.com/ipv4-pricing/

- And then there's the price…
- ISC recently sold a /15 at $14/address.
  - $14*2^15 = $458,752
- AWS, recently purchased half of MIT's address space at (reportedly) $20/address
  - $20*2^24/2 = $16.8M
  - MIT using it to fund IPv6 deployment
  - Amazon using it to hold ground while deploying IPv6

https://www.networkworld.com/article/3191503/internet/mit-selling-8-million-coveted-ipv4-addresses-amazon-a-buyer.html

Why not just add layers of translation?

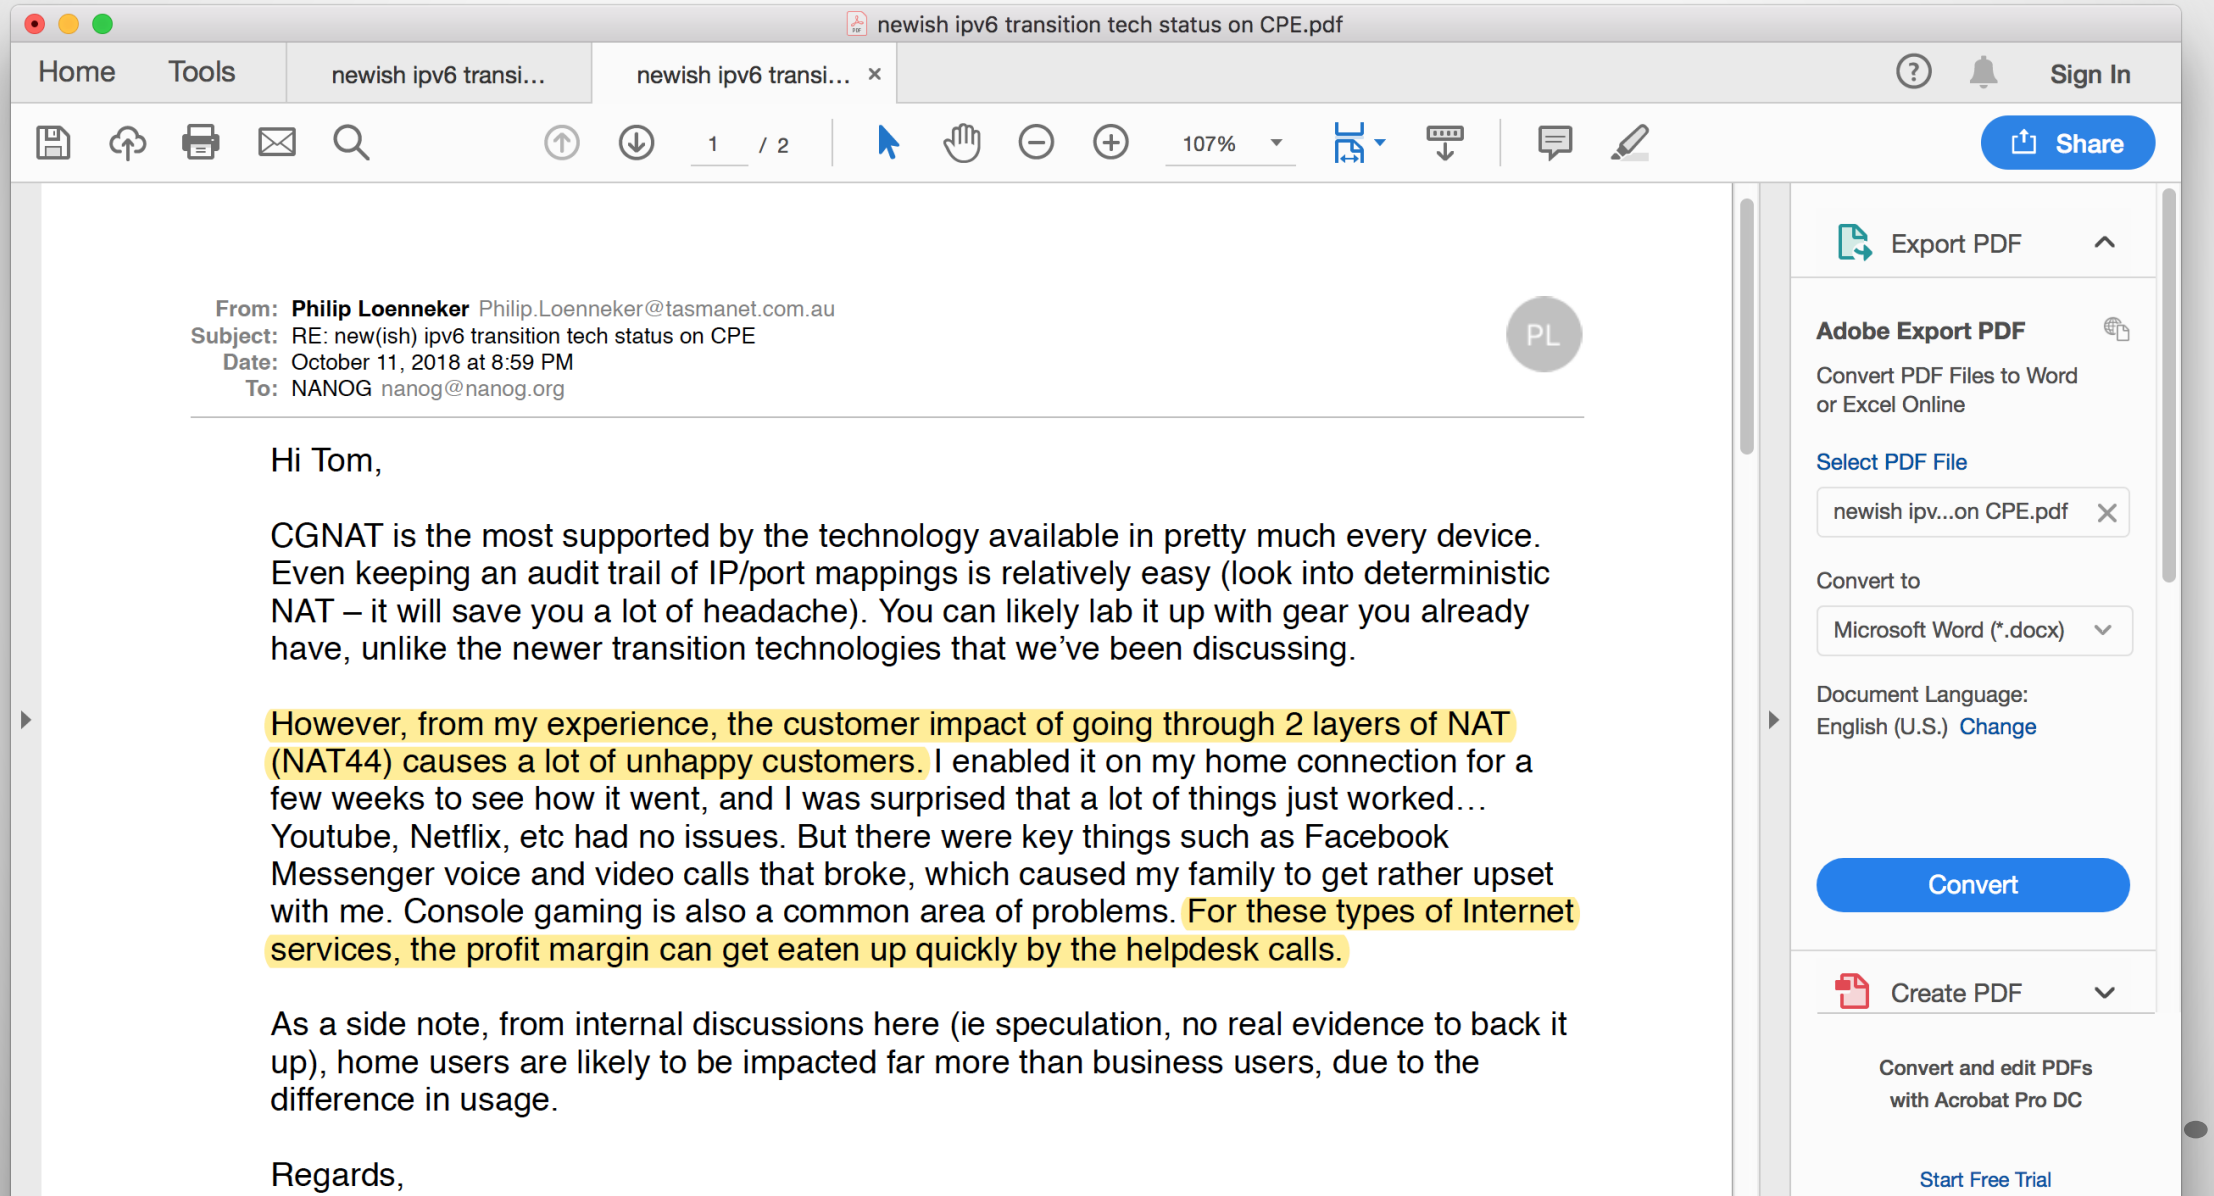

Greenfield Network business case

- IPv4
  - 5000 /24 prefixes = 1,280,000 addresse
  - $45,000 Annually, **but LACNIC doesn't have them to provide**
  - Open Market:
    - At $14/address, $17,920,000
    - At $20/address, $25,600,000
- IPv6
  - $2100 Annually

Category	IPv4 Prefix	IP Addresses (up to and including)	Initial Assignment Fee (USD)	Annual Renewal Fee (USD)	Payment before due date (discount)	Payment more than 30 days overdue (surcharge)
*Nano	Smaller than a /22	1,023	600	600	570	630
*Micro	Smaller than a /20	4,095	1,000	1,000	950	1,050
Small	Smaller than a /18	16,383	IPv4 Depletion Phases: This number of IP addresses may not be requested	2,100	1,995	2,205
Medium	Smaller than a /16	65,535	IPv4 Depletion Phases: This number of IP addresses may not be requested	5,700	5,415	5,985
Large	Smaller than a /14	262,143	IPv4 Depletion Phases: This number of IP addresses may not be requested	14,000	13,300	14,700
X Large	Smaller than a /12	1,048,575	IPv4 Depletion Phases: This number of IP addresses may not be requested	28,000	26,600	29,400
2X Large	Smaller than a /11	2,097,151	IPv4 Depletion Phases: This number of IP addresses may not be requested	45,000	42,750	47,250

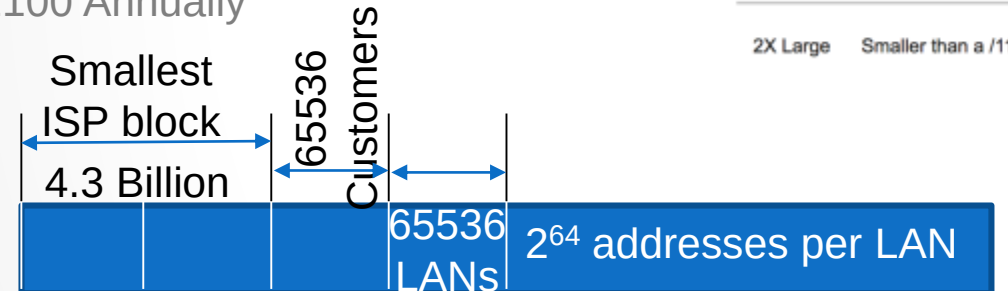

http://www.lacnic.net/2399/2/lacnic/membership-categories-and-fees

Category	IPv6 Prefix	Initial Assignment Fee (USD)	Annual Renewal Fee (USD)	Payment before due date (discount)	Payment more than 30 days overdue (surcharge)
Small	Smaller than or equal to a /32	2,100	2,100	1,995	2,205
Medium	Smaller than a /30	5,700	5,700	5,415	5,985
Large	Smaller than a /28	14,000	14,000	13,300	14,700

IPv6 Traffic Studies

- 12/14/2018, Google reported that more than 5% of its traffic from each of 50 countries used IPv6
  - Google, Akamai, and APNIC each report several countries in which they see 50% or more of traffic using IPv6
  - Several networks, including T-Mobile USA and DT Terastream, have **no IPv4 configuration, and are therefore IPv6-only**

https://www.vyncke.org/ipv6status/compare.php?metric=p&countries=be,de,gr,us,in,my,uy,ch,lu,jp,br,vn,ax,fr,ee,mx,fi,tw,gb,ca,hu,ec,tt,pt,ie,th,nz,pe,sa,au,lk,nl,pr,sx,ro,no,cz,si,pl,mo,gt,ar,sg,zw,at,fo,bo,ga,se,ba

Google

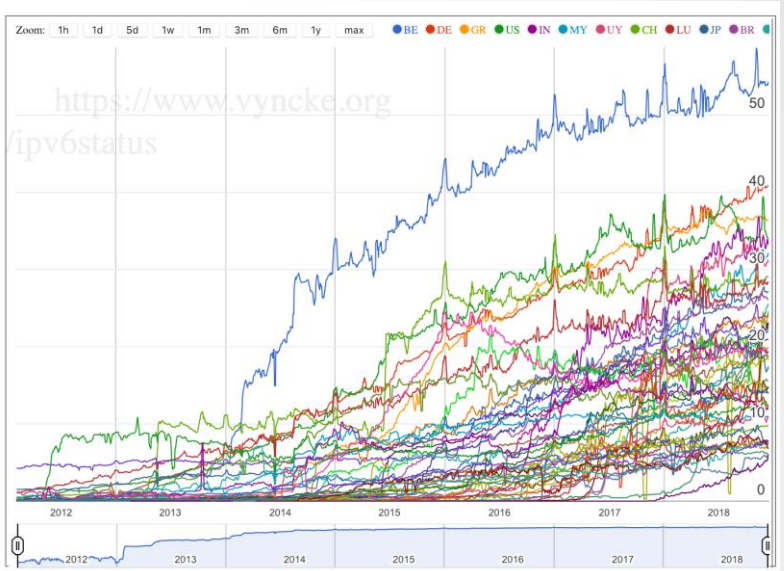

APNIC

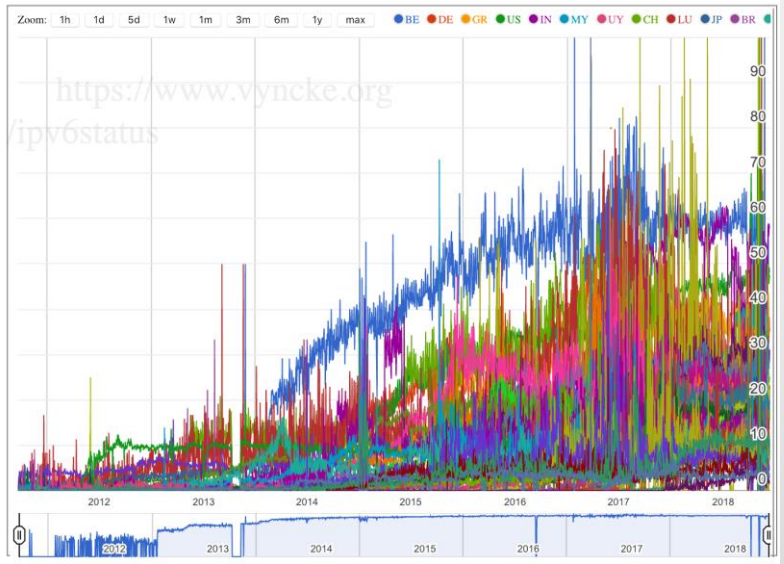

Akamai

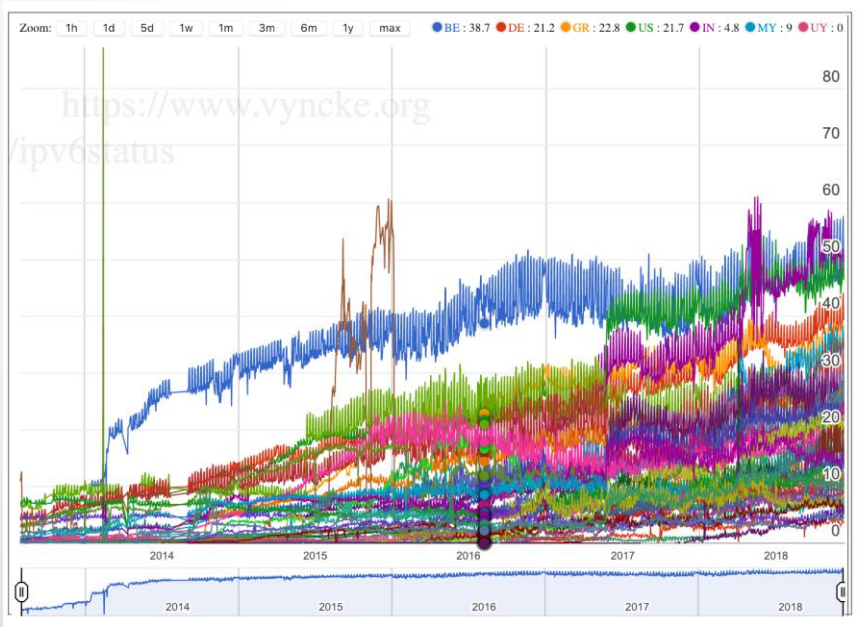

The IPv6 laggard: Enterprise

- I like Mythic Beasts' approach to web application hosting
  - IPv4 addresses and network complexity cost actual money
    - Pass cost along to customer
  - IPv6 addresses don't cost much, and from them are free
- I have ISPs worldwide telling me
  - They don't want to purchase IPv4 address space
  - Deploying 464xlat, lw4o6, etc. over IPv6-only networks
- **Fair to expect IPv6 traffic levels to rise as a result, as seen by published statistics**

Important things to think about up front – what to learn

• • •

"…the first deploying ISP should gain a competitive advantage through the ability to sell a service that is desirable even for the initial customers."

SCION

https://www.pqr.com/sites/default/files/bestanden/Downloads/6lessons_best-and-worst-of-the-internet-of-things.pdf

Scale

“If you’re not afraid, you don’t understand”

“The only real problem is scale”

Mike O’Dell

Chief Scientist, UUNET

Change Management

- From time to time, you will need to install or change something in your network
  - That might be a parameter, a download, a configuration…
- How will you do that in a manner any user can effectively use? Securely? On demand?

- **In order to fix the Mirai Botnet, we needed for all IOT devices to download a new software load**
  - Their creators might already be out of business.
  - Many had no download capability
  - Many could not be triggered to do so
  - Many could as easily download a corrupt software load from an improper source or be immediately reinfected

Align business needs with technology

- “But I’m a researcher, I’m not thinking about business”
  - Yes – but the user of your technology will.

- You need to provide the tools s/he will need, and leave the rest out
  - That includes flexibility to add new tools when the need becomes clear.
  - “But I like it” doesn’t make it useful
  - Unnecessary complexity/coupling can sink the technology
  - Too simple can be just as bad

Security and privacy

- Definitions:
  - Security: controlling access to information
    - Authentication vs Authorization
    - Needs to be mutual
  - Privacy: protecting identity
    - Private information can be data held securely
- Applies at each level of the architecture, not just the network layer or the application layer

Coexistence

- When you deploy your new system, much of what you need to access will be in the old system. The old system will likely not be *forward compatible* with the new.
  - When you update your system, the updated entities, virtual or physical, will be the “new” system and everything else the “old” system.

- **Enable the new system to interact cleanly with the old.**

- This has been a critical issue in IPv6 deployment:
  - Translation was eventually specified, but the IPv6 community has tried very hard to prevent or cripple it.
  - Only recently has it been recognized that the ability to communicate with existing IPv4 services is critical to convincing operators to deploy the new technology

In a new Internet, make new mistakes

Fred Baker

FG2030 Workshop 2018-12-18