Keeping the “Inter” in “Internet”

Leslie Daigle
ldaigle@thinkingcat.com
May 2018

Presented at InteropITX, May 1, 2018.

Expected takeaways

- No network is immune to the challenges of the general Internet
  - Many of the same problems internally as the network scales to meet needs
  - Many of the Internet's challenges bleed into your reality
- Looking at the general Internet and understanding how it has grown and adjusted is instructive
  - More Inter, not less
- There will always be new wicked problems – it's important to know how to solve them, not just find the solution for an immediate one

What is the Internet

And how to talk about it

What is the Internet

- There's a lot of dogma around what is/isn't the Internet
  - End-to-end
  - Internet Protocol (IPv4 or IPv6!)
  - Open and public
  - No NATs
  - No DPI or middleboxes
  - "Network neutrality"
  - SOPA/PIPA
- Things that were (or are) controversial
  - Load balancers
    - Until anycasting became a thing
  - CDNs
    - Muck with the end-to-end notion
  - Cloud
    - "There is no cloud… there is only someone else's computer"
  - Split DNS

How helpful is that?

- Not very
- Many of the controversial things have grown up because pragmatic needs outstripped available "blessed" technology
- We need a better way to think about what needs to be unchanging in the Internet ("invariants")
  - To make sure they aren't messed with
  - And to see when some of the pragmatic shortcuts are really going to be dead ends

Invariant Properties of the Internet

- Global reach, integrity
- General Purpose
- Supports innovation without requiring permission (by anyone)
- Accessible
- Based on interoperability and mutual agreement
- Collaboration
- Technology – reusable building blocks
- There are no permanent favourites

https://www.internetsociety.org/internet-invariants-what-really-matters/

And the greatest of these is…

- And the greatest of these is: Collaboration
  - Looking outside your own corporate 4 walls, working with other companies – potentially even competitors, when you have a common goal
- Diversity – of networks and support systems

Collaboration Illustrations

Collaboration at multiple levels

Physical Internet

- Connecting the world
- Network resiliency

Logical Internet

- Build out of Internet to support actual usage
- Monoculture & single points of failure

Business on the Internet

- Open Source
- Solving Wicked Problems – IPv6, for example

Physical Internet

Connecting the World

Diversity is good

Collaboration through multiply-connected resources promotes diversity

Routing failovers are essential

A submariner's perspective

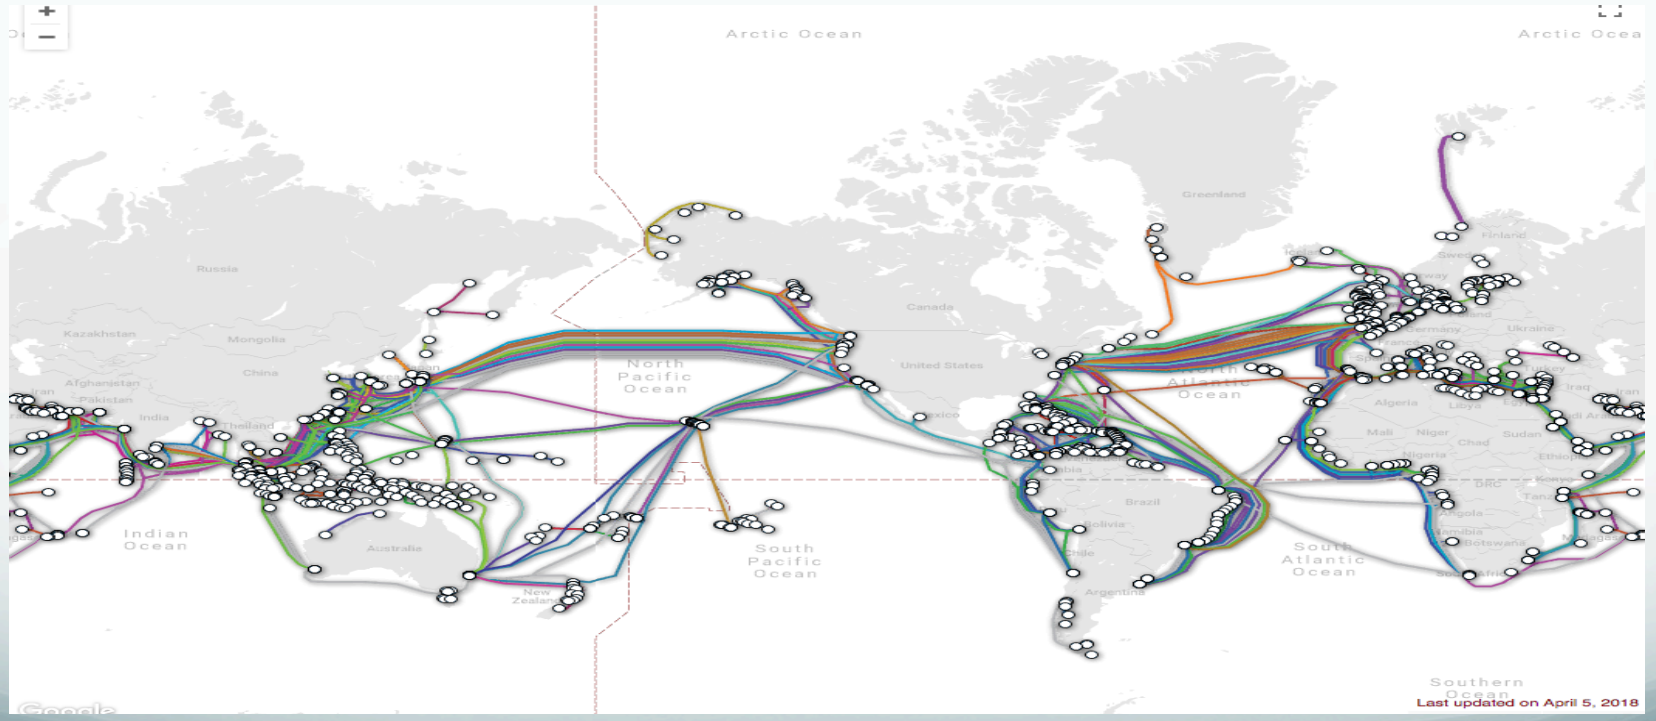

Telegeography: https://www.submarinecablemap.com/

The Atlantic links

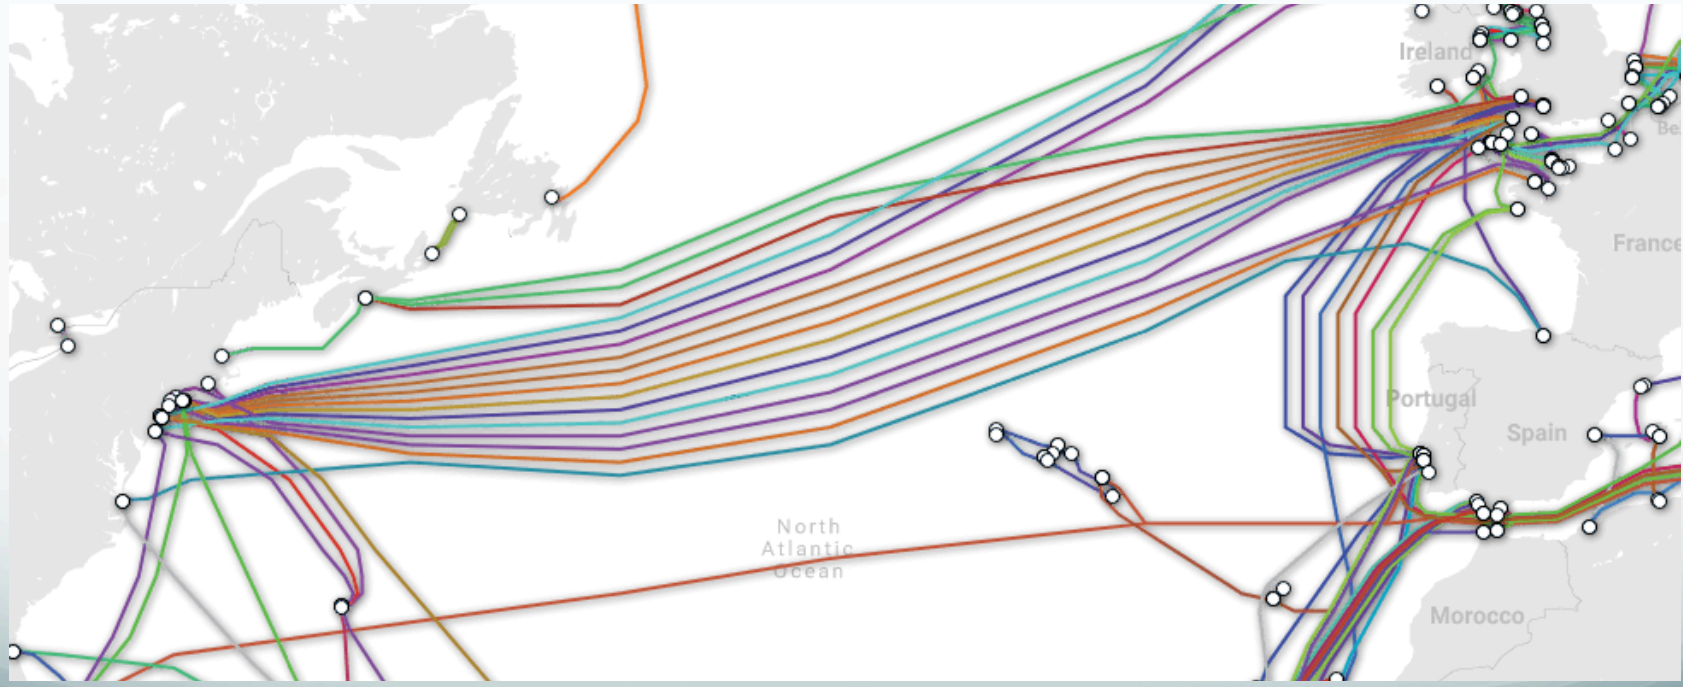

Telegeography: https://www.submarinecablemap.com/

A Pacific Perspective

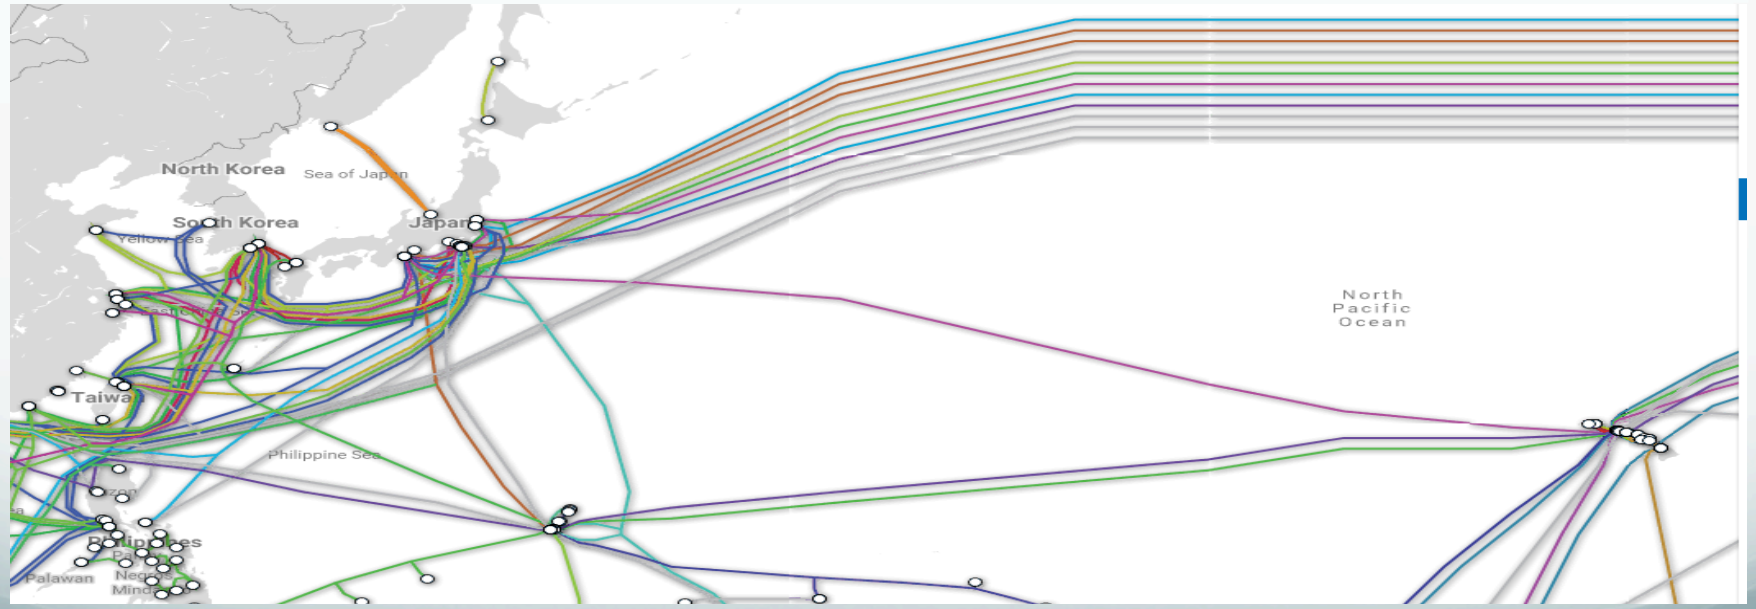

Telegeography: https://www.submarinecablemap.com

Resilience in the face of disaster

Diversity is good!

New York – 9/11

- New York was the landing site for many important trans-oceanic network links
- 9/11 – had a significant impact on key pieces of physical infrastructure
- Andy Ogielski and Jim Cowie of Renesys Corporation shared an analysis of the impact on network reachability and stability
  - March 5-6, 2002, "National Research Council, Workshop on The Internet Under Crisis Conditions"

Physical Impact of 9/11

- As noted by Ogielski & Cowie:
  - "World Trade Center 1 & 2: Below-ground-level fiber from the Telehouse to 60 Hudson St, and to transatlantic cables."
    - Cables were lost in the destruction of WTC buildings on 9/11
  - "60 Hudson St: A large carrier hotel, termination of multiple transatlantic cables, with many PoPs at 25 Broadway"
  - "NYIIX at 25 Broadway Telehouse: A peering site serving about 40 Internet providers from NYC area, Europe, South America, and South Africa. Claims about 70% of Europe to US traffic."
    - 25 Broadway suffered power failure on 9/13 when backup generator failed; back on 9/15 and failed again.

Internet Impact of 9/11

- Reachability issues to fewer than 1,000 Internet destinations
  - On either side of the Atlantic
  - Due to physical implications
- No noticed routing stability issues
  - In fact, fewer transient route announcements – less normal maintenance on near, but unaffected networks?
- "There has been a sufficiently high peering redundancy, and so far single localized events such as the 9/11 attacks, or the 7/18 Baltimore tunnel fire, do not cause widespread global routing problems. Lack of redundancy is fatal: cf. South Africa"

Japan – Earthquake & Tsunami 2011

- "The Great East Japan Earthquake & Tsunami", March 11, 2011.
  - North and east of Japan
  - More than 15,000 dead, and 2,500 still missing (2016)
  - Fukushima reactor nuclear meltdown
  - 20 million tons of refuse & personal property were pulled out to sea

  ⇒Wide-scale, physical and infrastructural damage
- Kenjiro Cho, Cristel Pelsser, Randy Bush, Youngjoon Won, of Internet Initiative Japan (IIJ) wrote: "The Japan Earthquake: the impact on traffic and routing observed by a local ISP"

Internet Impact – in Japan

- Obvious impact on individual homes and businesses, local infrastructure, in impacted region
- IIJ reported its data center in Sendai (second largest city north of Tokyo; in hardest hit region) lost power, but ran on generator until power was restored (less than 24h later)
- Network links to Sendai went down, but connectivity was restored the day after the earthquake

Internet Impact -- Internationally

- Most intra-Asian links chiefly land on west side of Japan and were unaffected
- IIJ reported its US-connecting undersea cables and one Asian connector were initially fine, but started failing in the hours after the quake
  - Likely disrupted by undersea landslides in aftershocks
  - 3 of 8 trans-Pacific links were damaged
  - At all times, at least 3 were functional
    - And, routing updates adjusted traffic accordingly

3 trans-pacific links

Used with permission – Kenjiro Cho

The paper's conclusions

- "Despite many failures, the Internet was impressively resilient to the disaster, aside from the areas directly hit by the quake and ensuing tsunami."
- "We [...] observed how redundancy and over-provisioning worked for the Internet from traffic on trans-Pacific links. From the routing analysis, we have identified significant internal events within IIJ, but also observed that these events were masked to the rest of the world."
- "Our observations illustrate the resilience of the Internet to large scale disasters, and emphasize the importance of redundancy and over-provisioning in the network design."

Physical Internet: Takeaways

The Internet may not have been designed to withstand a nuclear attack, but the:

- highly interconnected,
- collaborative,
- resilient failover nature of the technology and provisioning

have in fact allowed it to substantially survive significant physical disasters

Logical Internet

The Internet is malleable – look about to see what makes more sense (or not)

Better ways to connect content

- Kurtis Lindqvist, APRICOT 2013, "A History of Peering and Learning for the Future"
  - Looking at traffic usage patterns, changes and better things for the future
- Sometimes you have to step out of your own perspective and look around to see how to make things work more rationally
- Things to note:
  - This is public data – unlike, say a phone network
    - And a public presentation
  - It's network/IXP operator perspective, but it shows the thought process

The Internet, ca. 1993

https://cdn0.vox-cdn.com/assets/4463763/world_large.gif

What changed in 20 years?

- Between 1993 and 2013: peering evolved
- As more local content and resources were developed, there were more local interconnections
  - And, as more local interconnections are built, more content can be kept locally

Logical next step for IXPs, CDNs

- As home broadband access grows, the capacity for consumption (and dissemination) grows
  - Uploading videos to YouTube – demonstrated async links were poor planning on ISPs' part!
- How to accommodate increased demand?
  - More, bigger pipes, or
  - More localization
    - IXPs for denser interconnects
    - CDNs put content closer
      - Google is attached to just about every major ISP, globally
- The conclusion (from Kurtis' talk): denser interconnects (collaboration)

Research in Motion: Blackberry

- Research In Motion built the market for mobile work platforms
  - 1999: the first truly usable e-mail over mobile networks
  - Conservative use of data made it responsive and affordable (data plan usage)
  - 8 years later, the iPhone doesn't touch it for performance or affordability

However…

- October 2011: RIM network failure halts BlackBerry services for all 70M customers, for three days
  - Finger pointed at a faulty router in their network
  - Because RIM built and relied on their own massive application network overlay, they had insufficient diversity or resiliency
  - From an enterprise perspective, that puts your threat surface outside your 4 walls
- Other mobile platforms, which use (chatty) Internet standard application protocols, leverage the resiliency of the Internet

Logical Internet: Takeaways

Globe-spanning actual network

- Exactly impossible when the Internet was built (why the Internet was built)
- Economies of scale – now it's possible
  - Doesn't mean it's necessarily a good thing

Improving Internet experience is largely tied to

- More diversity
- Collaboration

And, single points of failure… fail.

Business of/on the Internet

Collaboration – makes good business sense

Competition is not always the fastest path to your goal

Collaborative Open Source Software

- More than just “here’s some source code you can [review and then] compile and run”
  - Though that’s a good plan for trustable software
- More than just “here’s my source code, maybe I’ll accept your proposed changes”
- Commercial grade tools and environments are built through collaborative open source activities
  - Linux Foundation
  - Apache Software Foundation
  - Django, Drupal and Eclipse foundations…

Walmart and OSS

- "Why would a company spend resources on an open source program, and why is it really needed?"
  - "GitHub is where your talent pool lives
  - We decided upon node because not only was it a good technical fit, but we could bring in a world class team of engineers who were desperate to build very large scale node services." – Dion Almaer (WalmartLabs)
  - Benefit from bugs and feature development that happens through the community
- "By open sourcing Electrode, we're encouraging the OS community to help make it better — for us, and other developers who will use it." – Alex Grigoryan, the director of engineering of WalmartLabs

World IPv6 Events

World IPv6 Day (2011)
World IPv6 Launch (2012)

IPv6

- IPv6 – because there aren't enough IPv4 addresses, but it's a bear to get it deployed (not directly compatible with IPv4)
  - There's a story there about how the Internet was adequately academic in the mid-1990's when IPv6 was designed and standardized – everyone would naturally age out their equipment often enough that it was reasonable to assume deployment

World IPv6 Day

- June 8, 2011
  - For 24 hours, Google, Facebook, Yahoo! And Bing turned IPv6 on for their primary web pages
  - And then they turned it off
- More to the story
  - Sure – hundreds of other websites did, too
  - Some of them didn't turn it off
- But, this was not a "flag day", or even a "celebratory event"

The business problem of IPv6

- Large website operators knew that they would lose some percentage of customers if they turned on IPv6, even while leaving IPv4 connectivity untouched
  - E.g., systems with badly configured IPv6, unresponsive proxies, etc, would simply time out the connection
- Large website operators' business decision makers knew this would mean customers would simply go to their competitor's services
  - "Oh, Google's down today. I'll go to my Yahoo! account"
- They didn't have a fix on how big a segment of their users that was
- They couldn't test it on their own

So, what WIPv6 Day actually was…

- Months of internal work at the major participant companies to get themselves ready
- Advance publicization of the event
  - With months of working with vendors to get known problems fixed and patches pushed out
    - "We're doing this on June 8; you don't want your software to look bad, do you?"
- Months of preparation of call centres at ISPs in case customers were confused
- A lot of measurement
  - So by the time June 8 rolled around and the event was live, there was a pretty clear expectation of success (small enough numbers to be "acceptable" for business purposes)

A year later…

- Major content providers knew they could do it
  - But – why would they? So few IPv6-enabled customers
- ISPs had to get in the game
  - And, why would they, absent content for people to reach over IPv6?
  - Heavy dependency on CPE vendors
  - Need to enable large chunks of their networks just to get a small fraction using IPv6
- Classic…

Chicken and Egg

- World IPv6 Launch asked (top Alexa) content providers to turn IPv6 on for good, and access providers to deliver a minimum of 1% of their traffic to those content providers over IPv6
  - 1% was a high bar
- June 6, 2012, IPv6 was launched
  - "IPv6 Is The New Normal"


WORLD
IPv6
LAUNCH

Google accesses over IPv6

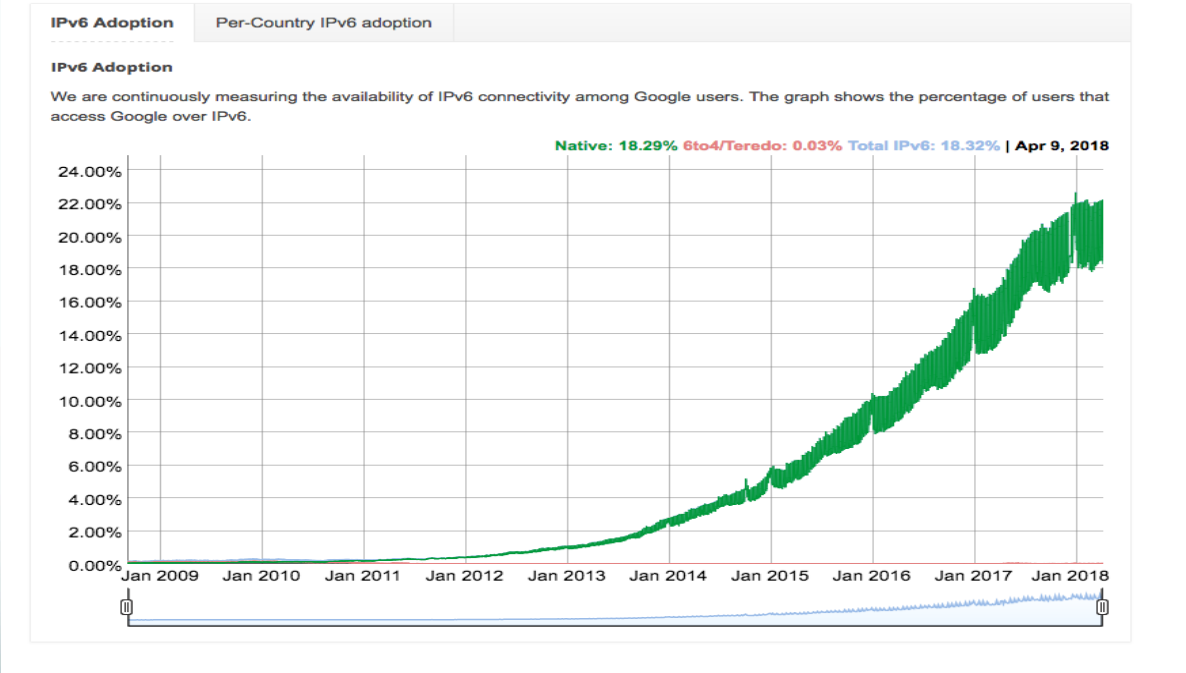

Source: https://www.google.com/intl/en/ipv6/statistics.html

IPv6 in 2018

Rank	Participating Network	ASN(s)	deployment
1	Comcast	7015, 7016, 7725, 7922, 11025, 13367, 13385, 20214, 21508, 22258, 22909, 33287, 33489, 33490, 33491, 33650, 33651, 33652, 33653, 33654, 33655, 33656, 33657, 33659, 33660, 33661, 33662, 33664, 33665, 33666, 33667, 33668, 36732, 36733	61.77%
2	KDDI	2516	43.60%
3	RELIANCE JIO INFOCOMM LTD	55836, 64049	87.91%
4	SoftBank	17676	31.99%
5	ATT	6389, 7018, 7132	65.33%
6	Charter Communications	7843, 10796, 11351, 11426, 11427, 12271, 20001, 20115, 33363	29.82%
7	Verizon Wireless	6167, 22394	82.64%
8	T-Mobile USA	21928	91.43%
9	Deutsche Telekom AG	3320	53.12%
10	SKTelecom	9644	34.98%

http://www.worldipv6launch.org/measurements/ March 14, 2018

Business of/on the Internet: Takeaways

At a business level, there are times when it makes most sense to collaborate – even with competitors

- This isn't just "old school Internet"
- It's not about ideals.

What does this mean for you

Wicked problems of the Internet – need your help!

We know how to solve wicked problems

- We have seen the problem… and the answer is us
- The Internet faces some rather major challenges, and we can solve them by going back to basics: more collaboration and interconnection, not less

Inter-networking means global challenges

- The Internet is not a globally managed network with universally defined service level assurances
  - This diversity is what gives it all the resiliency we've been discussing so far
- When there are updates and challenges, like IPv6 deployment, the best and only way to ensure success in overcoming them is through collaboration

Security of routing packets

- When you send packets from your network into the wilds of the Internet, you can worry whether they will:
  - arrive at the intended destination;
  - be copied or read; and/or
  - be tampered with en route.
- When you take in packets from the wilds of the Internet, you may wonder whether
  - you actually want them; and/or
  - they are who they say they are.

InSecurity of routing packets

- When you send packets from your network into the wilds of the Internet, you can worry whether they will:
  - arrive at the intended destination;
    - Route hijacking
  - be copied or read; and/or
    - Crypto can help; though not for metadata
  - be tampered with en route.
    - Crypto can help; though not for metadata
- When you take in packets from the wilds of the Internet, you may wonder whether
  - you actually want them; and/or
    - Unwanted traffic – DOS attack?
  - they are who they say they are.
    - Spoofed source addresses

Where the issues are

- Generally – business agreements between networks mean it's not completely "the wild West" and these are not normal, expected conditions
  - But, bad actors can still cause problems
- The problems
  - Route hijacking
    - Or, fat-fingering
  - Address spoofing
    - Facilitates DDoS

Moratel Leak, 2012

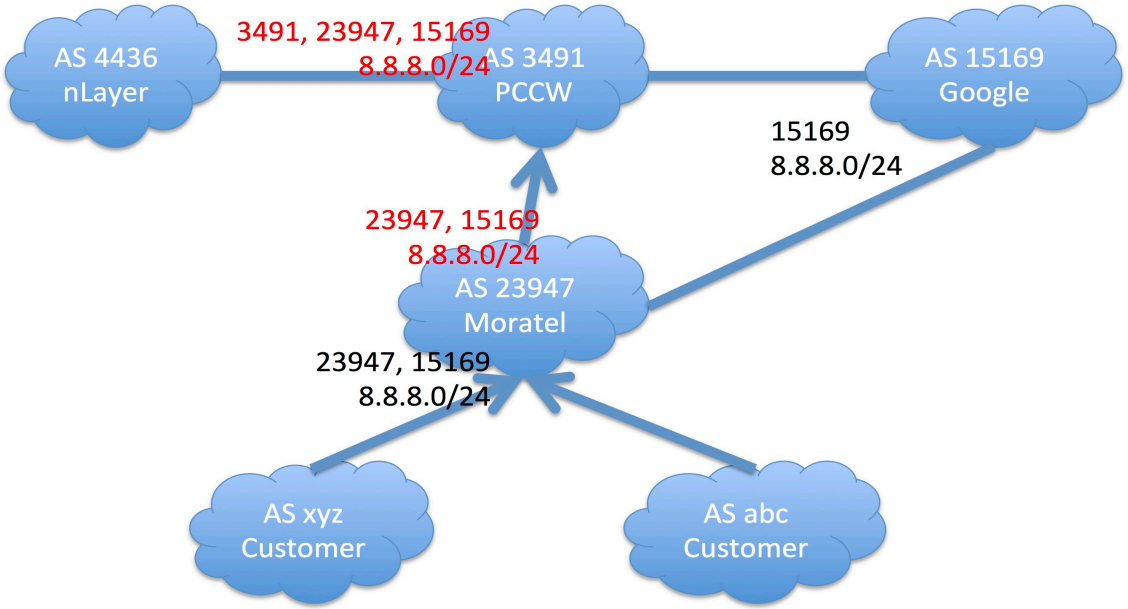

What is wrong in that picture…

- Clearly: Moratel (AS 23947) should not have advertised a path to Google (AS 15169)
- But, also, PCCW (AS 3491) should not have shared it
  - It did, because it was a customer link ($$)
  - And, how to detect it was inappropriate?
    - Not an advertisement for Moratel's address space
    - Not an advertisement for Moratel's customers

    ⇒PCCW should conclude Moratel should not be advertising it, and therefore drop it

But wait…

- How does PCCW know who is/isn't a Moratel customer?
  - Industry resistance to publicizing all neighbour links
    - Competitive intelligence
    - Some relationships may be ephemeral
      - Backup links
      - DDoS mitigation links
  - Lack of agreement on reliable sources
    - Which is *your* favourite IRR?
    - How to tell if data is up to date?

Current Work

- Crypto for routing
  - Using RPKI to verify ownership of resources, and to encase announcements in crypto
  - Resistance in deployment
    - Computationally expensive
    - Encasing -> reduced flexibility
    - Requires trust in external entities (RIRs)
    - Difficult to express actual networking relationships well
- Collaborative agreements for best practices: MANRS
  - "Mutually Agreed Norms for Routing Security"
    - https://www.manrs.org/
    - ~60 networks signed up (April 2018)
- URSA – Unwedging Routing Security Activity
  - Shameless plug :^) http://www.techark.org/ursa
  - Collaborative effort of operators to define practical steps forward

Wicked problems of the Internet need your help – and getting involved will help you

What have we learned?

- Surviving catastrophes through diversity, collaboration
  - 9/11
  - Japan Earthquake/Tsunami
- Creating catastrophe through lack of diversity and monopoly
  - Blackberry
- Collaboration solves business problems and is business-smart
  - Open source software
  - Looking at the Internet beyond your own network

And…

- Wicked problems can be solved through collaboration
  - IPv6 deployment
  - Routing security ← get involved!

So: get involved! Let's keep the "Inter-" in "Internet".

Additional Material

Resources (1)

- Internet Invariants
  - https://www.internetsociety.org/internet-invariants-what-really-matters/
  - https://www.cigionline.org/publications/nature-of-internet
- 9/11
  - March 5-6, 2002, "National Research Council, Workshop on The Internet Under Crisis Conditions", Andy Ogielski and Jim Cowie of Renesys Corporation
    - https://web.stanford.edu/class/msande237/viewgraphs/911.pdf
  - https://www.nap.edu/read/10569/chapter/4#29
- Japan earthquake/tsunami
  - https://www.computerworld.com/article/2506876/internet/japan-s-internet-largely-intact-after-earthquake--tsunami.html
  - Kenjiro Cho, Cristel Pelsser, Randy Bush, Youngjoon Won, of Internet Initiative Japan (IIJ) wrote: "The Japan Earthquake: the impact on traffic and routing observed by a local ISP"
    - http://www.iijlab.net/~kjc/papers/1569500743.pdf

Resources (2)

- Connecting content
  - Kurtis Lindqvist, APRICOT 2013
    - https://www.apricot.net/apricot2013/assets/history-of-peering-and-learning-for-future.key_1361802442.pdf
- Blackberry outage
  - https://www.theguardian.com/technology/2011/oct/14/blackberry-outage-faulty-router-suspected
- Walmart & Open Source Software
  - https://techcrunch.com/2016/10/03/walmartlabs-open-sources-the-application-platform-that-powers-walmart-com/
  - http://todogroup.org/blog/why-we-run-an-open-source-program-walmart-labs/
- Moratel Leak
  - https://arstechnica.com/information-technology/2012/11/how-an-indonesian-isp-took-down-the-mighty-google-for-30-minutes/