The Changing Internet

- Exhaustion of IPv4 address space
- Wireless and mobility
- Hypergiants and Centralisation
- Real-time and low-latency
- Overcoming ossification

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

The Changing Internet: Architectural Assumptions

- The original design of the Internet made certain assumptions:
  - Devices are generally located in a fixed location, and have a small number of network interfaces that are uniquely addressable
  - The network and services are decentralised
  - Best effort service provides sufficient quality
  - The network is trusted
  - Innovation can happen at the edges; the network is dumb
- Reasonable for the 1980s, when the IP architecture was defined, but the network has changed – how have the protocols changed?

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Changes in Addressing and Reachability

- Assumption: every network interface has unique IP address
  - Devices uniquely addressable with uniform connectivity
  - In principle, any device can connect to any other; policy enforced by firewalls, privacy protected by changing address assignments
- IPv4 has insufficient addresses to support this model
- IPv6 deployment is slow; NAT is widespread

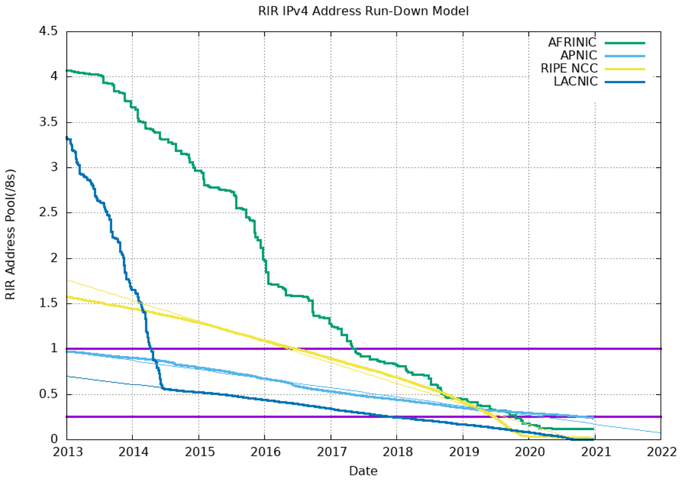

- Implication: connectivity becomes difficult
  - Dual-stack IPv4/IPv6 connection racing (“happy eyeballs”)
  - NAT traversal for peer-to-peer applications
  - Complicates software design and implementation
  - Forces reliance on cloud services and encourages centralisation

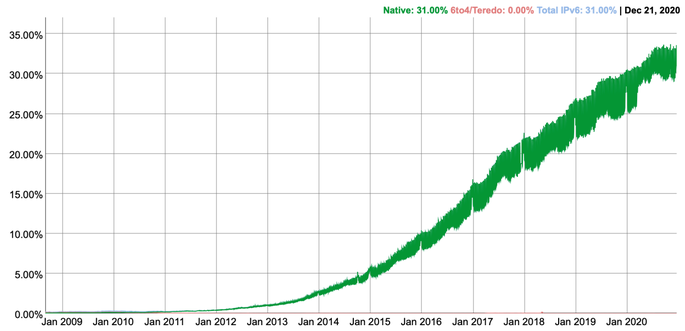

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

IPv4 Address Exhaustion

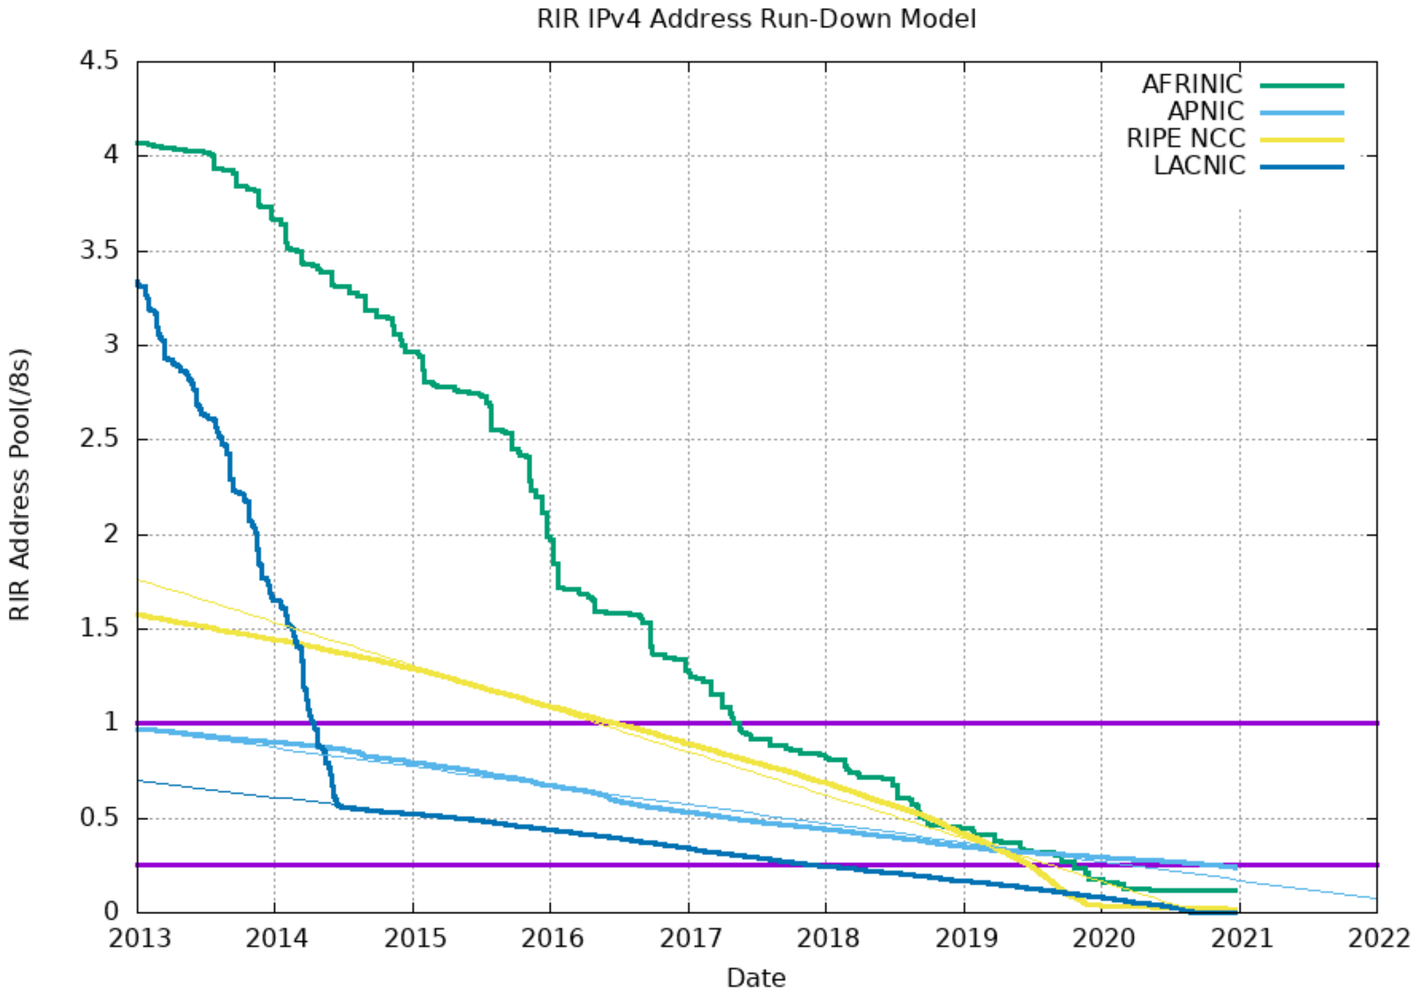

- IANA assigned last IPv4 block to regional Internet registries (RIRs) in 2011
- RIRs have exhausted the available IPv4 addresses, or soon will

Source: Geoff Huston, APNIC
December 2020 http://ipv4.potaroo.net/

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

IPv6 Deployment is Slow

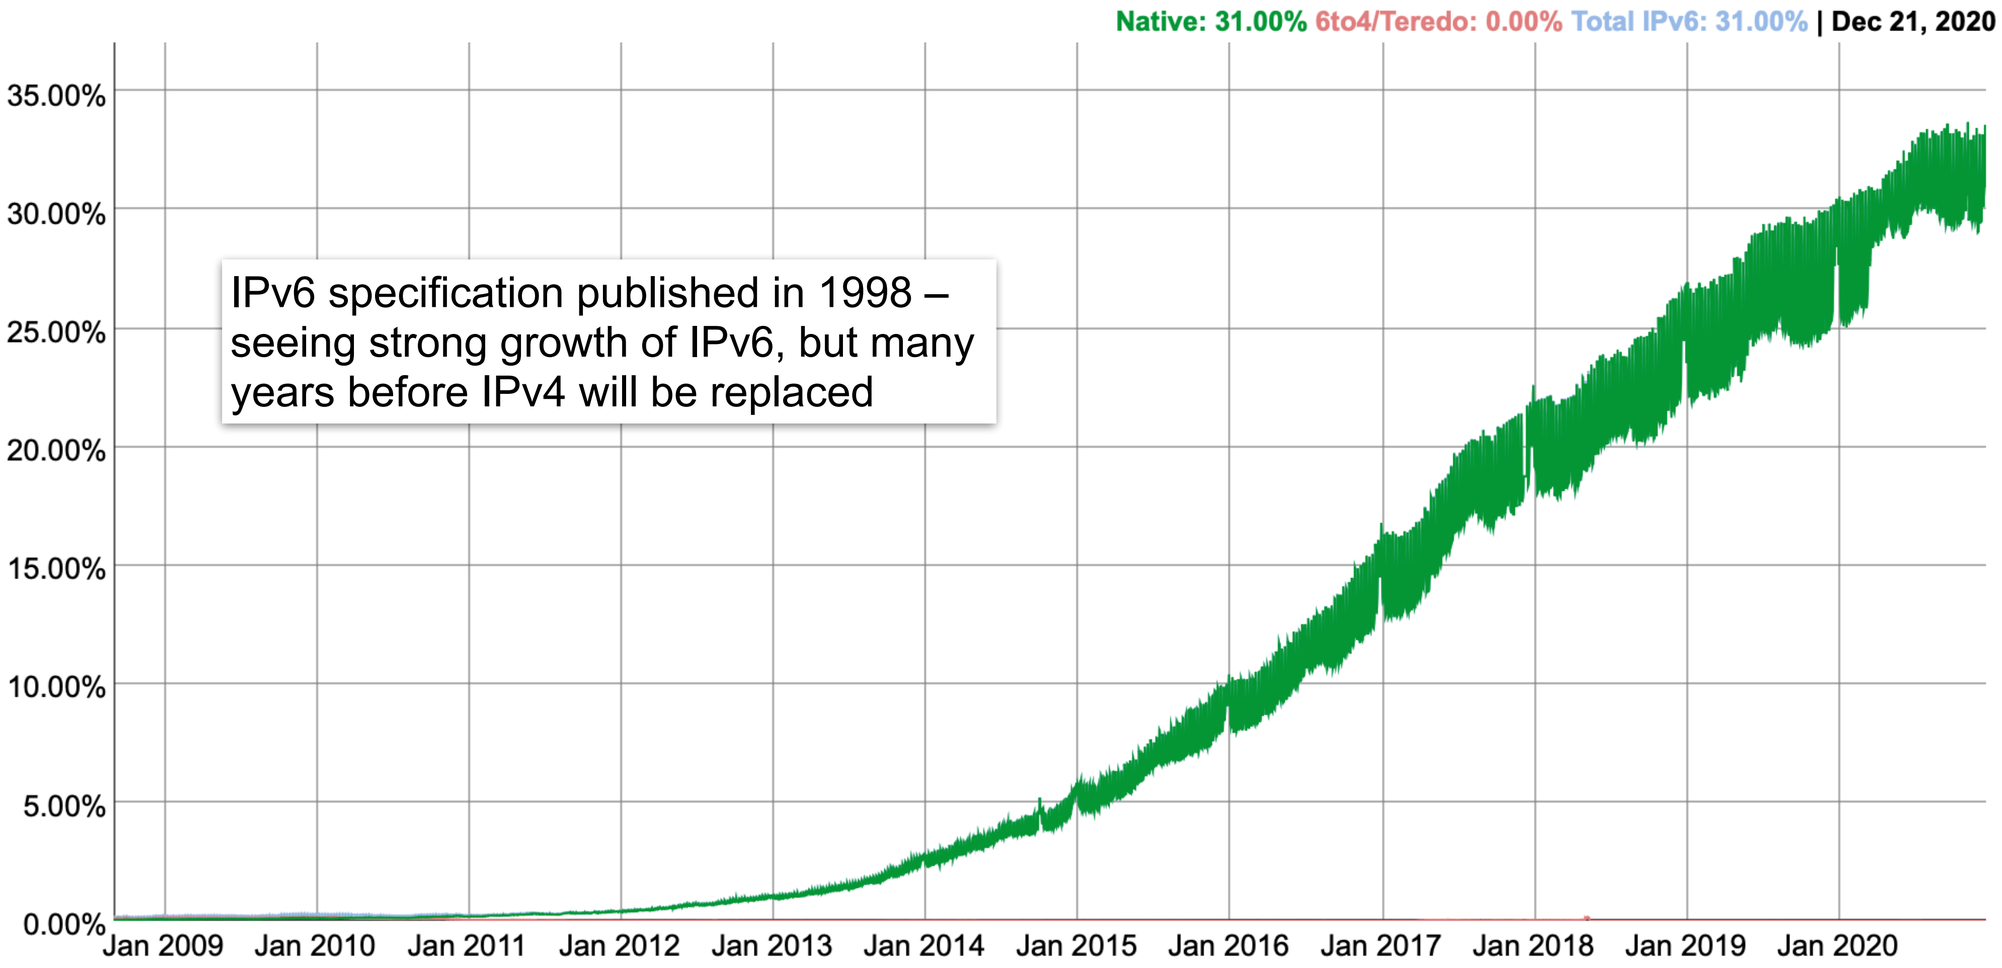

Source: Google, December 2020
https://www.google.com/intl/en/ipv6/statistics.html

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Establishing Connectivity in the Modern Internet

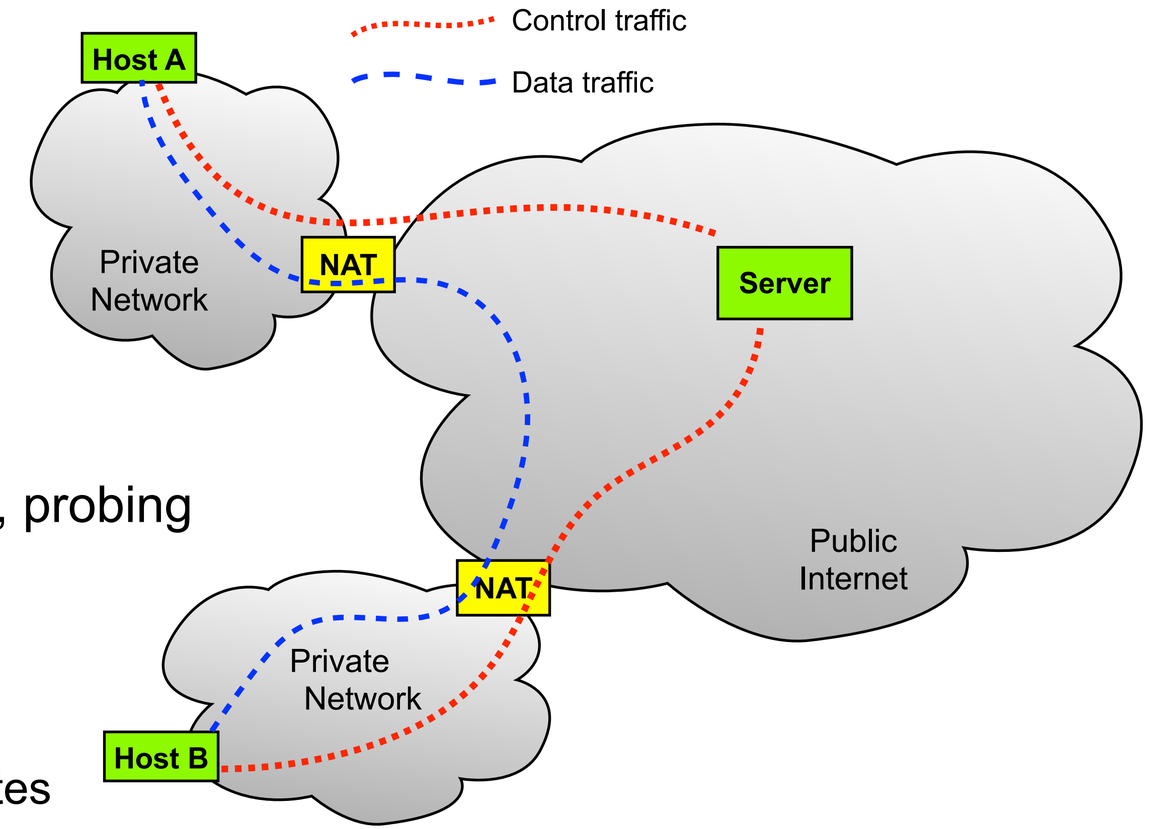

- NAT traversal → binding discovery, exchange, probing
  - STUN, TURN, and the ICE algorithm
  - Requires a server with public IP address for binding discovery; centralised point of control
  - Complex, slow, unreliable, wastes power, generates unnecessary traffic

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Establishing Connectivity in the Modern Internet

- Dual stack → connection racing
  - Some paths support IPv4, some IPv6; probing and connection timeouts slow for client-server applications
  - Race connections by probing in parallel: hard to implement, wastes resources

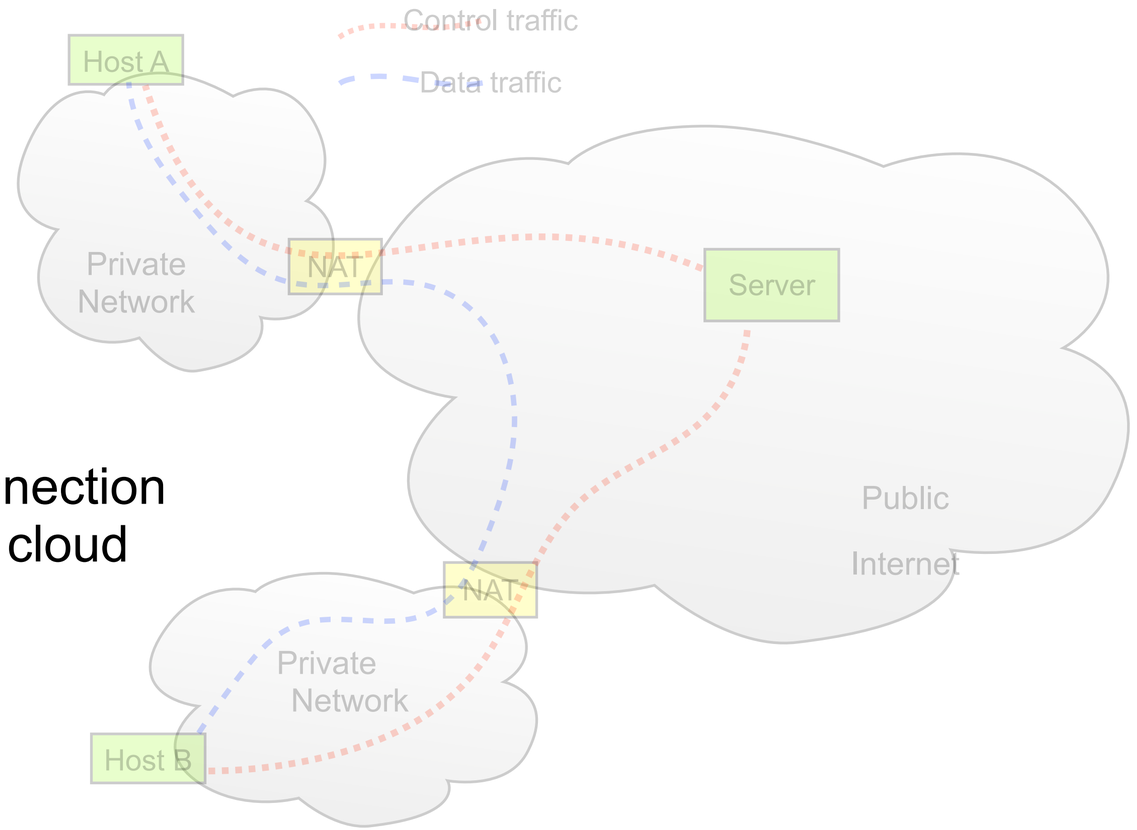

- NAT and dual stack operation complicates connection establishment; encourages centralisation onto cloud services

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Increasing Mobility

- IP addresses intended to encode location in the network
- Ubiquitous mobile devices – *smartphones* – complicate addressing and reachability since devices change their IP address each time they move:
  - TCP connections must be re-established after each move
  - UDP packets must be redirected after each move
  - Complex signalling and connection management if devices move frequently

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Hypergiants and Centralisation

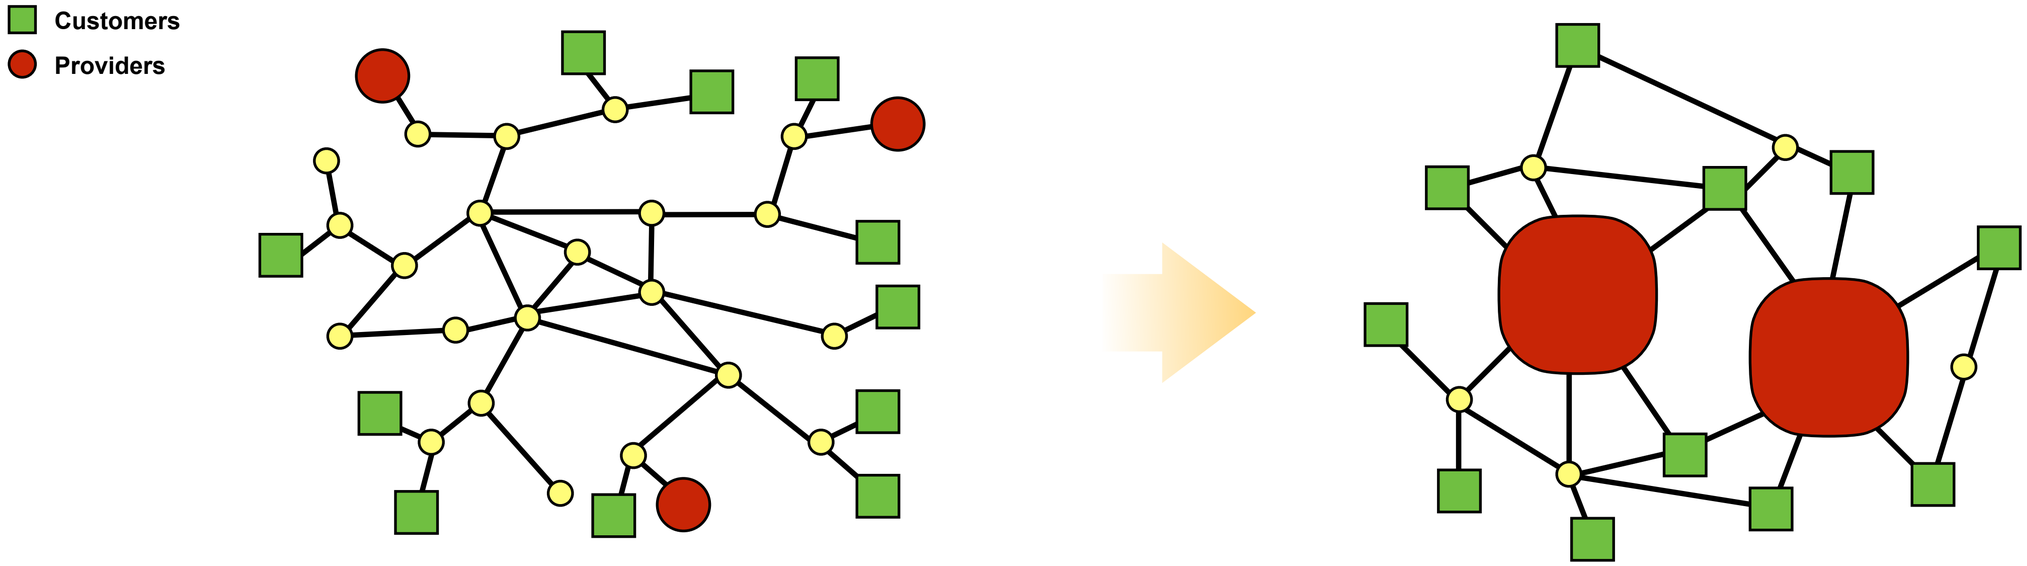

- Internet topology is flattening, becoming increasingly centralised
- Transit network ecosystem being replaced by direct connections from “eyeball” networks to content providers – Google, Facebook, Amazon, Akamai, …
- Implications for network neutrality, competition, innovation

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Supporting Real-time Traffic

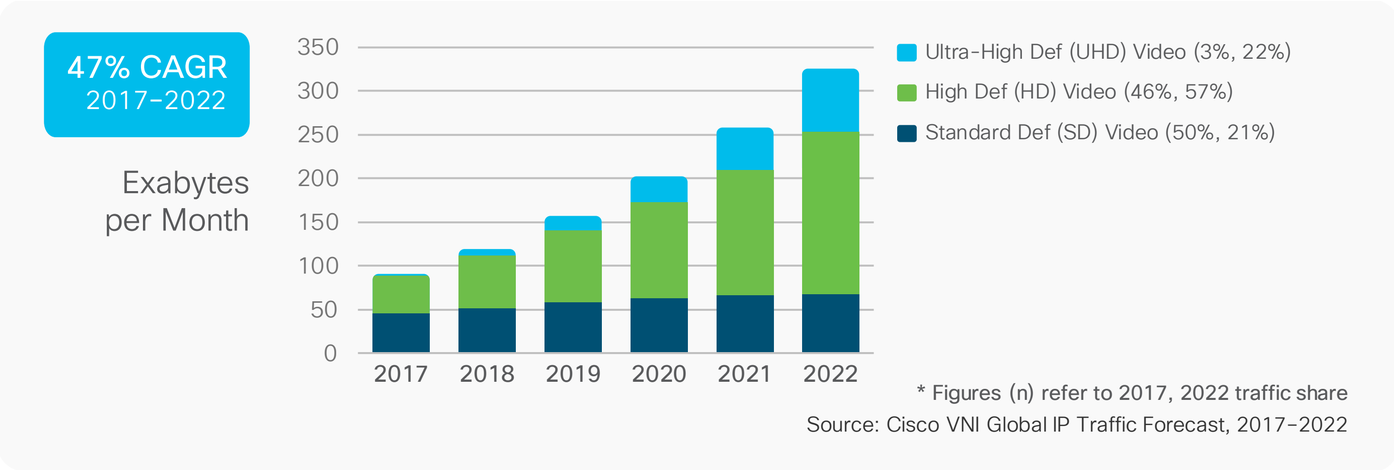

- Video streaming dominates Internet traffic, growing >40% per year
  - Packet loss → quality impairments or increased delay
  - Driving centralisation and direct CDN connections → easier to manage quality
  - Driving TCP congestion control and TCP replacements → BBR algorithm, QUIC; low-latency
- Interactive real-time media has stricter constraints
  - WebRTC, Zoom, VoIP, gaming, … → non-TCP transport to lower latency

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Supporting Real-time Traffic: Impact of COVID-19 (1/2)

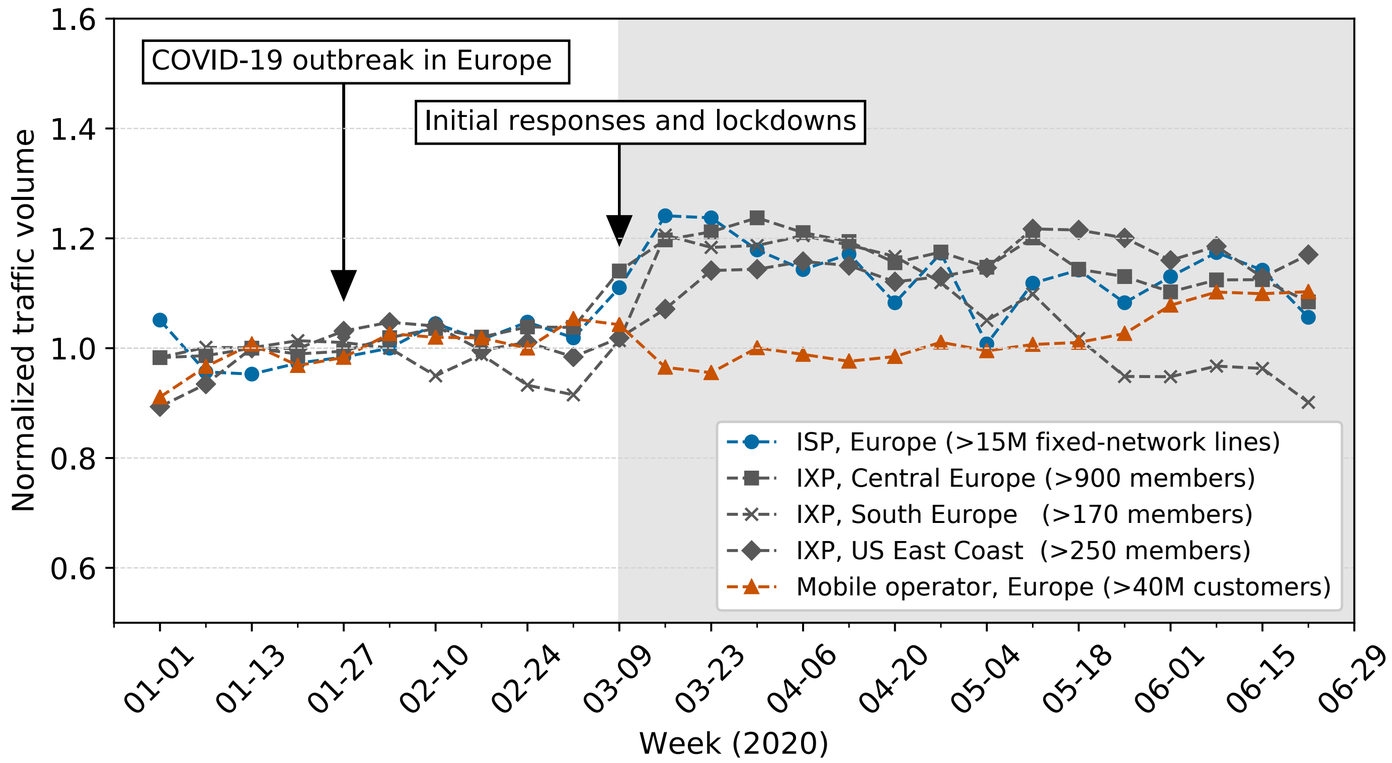

- Residential networks saw >20% increase in traffic overnight at the start of the COVID-19 lockdown
- Mobile networks saw ~10% drop in traffic

Anja Feldmann *et* al., "The Lockdown Effect: Implications of the COVID-19 Pandemic on Internet Traffic", ACM Internet Measurement Conference, October 2020, https://doi.org/10.1145/3419394.3423658

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Supporting Real-time Traffic: Impact of COVID-19 (2/2)

- Video conferencing providers saw massive growth in traffic
- The Internet was flexible enough to support this shift – can we maintain such flexibility while also improving quality?

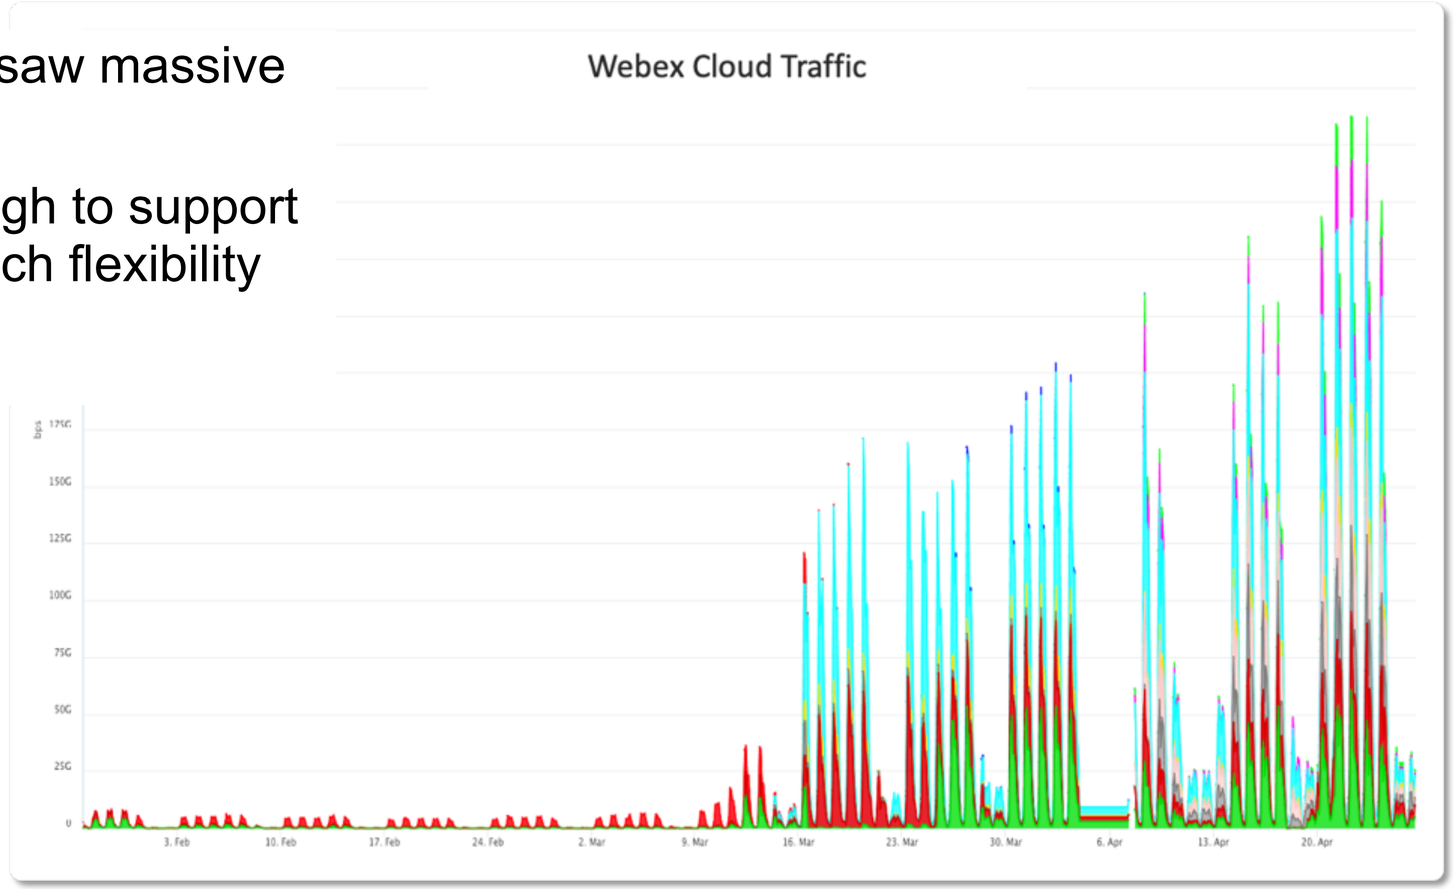

Source: Cullen Jennings and Panos Kozanian, “WebEx Scaling During Covid”, Proceedings of the Internet Architecture Board COVID-19 Network Impacts Workshop, 2020, https://www.iab.org/activities/workshops/covid-19-network-impacts-workshop-2020/

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Pervasive Monitoring and Trust in the Network



Edward Snowden

Internet Engineering Task Force (IETF) S. Farrell
Request for Comments: 7258 Trinity College Dublin
BCP: 188 H. Tschofenig
Category: Best Current Practice ARM Ltd.
ISSN: 2070-1721 May 2014

Pervasive Monitoring Is an Attack

Abstract

Pervasive monitoring is a technical attack that should be mitigated in the design of IETF protocols, where possible.

Status of This Memo

This memo documents an Internet Best Current Practice.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on BCPs is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc7258.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (http://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Farrell & Tschofenig Best Current Practice [Page 1]

S. Farrell and H. Tschofenig, "Pervasive Monitoring is an Attack", Internet Engineering Task Force, RFC 7258, May 2014. https://www.rfc-editor.org/rfc/rfc7258.html

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Innovation in the Network

- Can we still change the network protocols?
  - NATs, firewalls, middleboxes → TCP is *hard* to change now
  - Transport encryption, tunnel over UDP → QUIC? Pervasive encryption to defeat pervasive monitoring and surveillance?
  - Protocol ossification is a real concern
  - Is the model smart edges and a dumb network appropriate?
- Can new startups compete with the hypergiants?
  - Several forces pushing towards centralisation – can we evolve the network to counter these?
  - Can we enable a decentralised network?



Is it Still Possible to Extend TCP?

Michio Honda*, Yoshifumi Nishida*, Costin Raiciu†, Adam Greenhalgh‡,
Mark Handley‡, Hideyuki Tokuda*
Keio University*, Universitatea Politehnica Bucuresti†, University College London‡
{micchie,nishida}@sfc.wide.ad.jp, costin.raiciu@cs.pub.ro
{a.greenhalgh,m.handley}@cs.ucl.ac.uk, hxt@ht.sfc.keio.ac.jp

ABSTRACT

We've known for a while that the Internet has ossified as a result of the race to optimize existing applications or enhance security. NATs, performance-enhancing-proxies, firewalls and traffic normalizers are only a few of the middleboxes that are deployed in the network and look beyond the IP header to do their job. IP itself can't be extended because "IP options are not an option" [10]. Is the same true for TCP?

In this paper we develop a measurement methodology for evaluating middlebox behavior relating to TCP extensions and present the results of measurements conducted from multiple vantage points. The short answer is that we can still extend TCP, but extensions' design is very constrained as it needs to take into account prevalent middlebox behaviors. For instance, absolute sequence numbers cannot be embedded in options, as middleboxes can rewrite ISN and preserve undefined options. Sequence numbering also must be consistent for a TCP connection, because many middleboxes only allow through contiguous flows.

We used these findings to analyze three proposed extensions to TCP. We find that MPTCP is likely to work correctly in the Internet or fallback to regular TCP. TcpCrypt seems ready to be deployed, however it is fragile if resegmentation does happen—for instance with hardware offload. Finally, TCP extended options in its current form is not safe to deploy.

Categories and Subject Descriptors

C.2.2 [**Computer-communication Networks**]: Network Protocols; C.2.6 [**Computer-communication Networks**]: Internetworking

General Terms

Measurement, Design, Experimentation, Standardization

Keywords

Middleboxes, Measurements, TCP, Protocol design

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
IMC'11, November 2–4, 2011, Berlin, Germany.
Copyright 2011 ACM 978-1-4503-1013-0/11/11 ...$10.00.

1. INTRODUCTION

The Internet was designed to be extensible; routers only care about IP headers, not what the packets contain, and protocols such as IP and TCP were designed with options fields that could be used to add additional functionality. The great virtue of the Internet was always that it was *stupid*; it did no task especially well, but it was extremely flexible and general, allowing a proliferation of protocols and applications that the original designers could never have foreseen.

Unfortunately the Internet, as it is deployed, is no longer the Internet as it was designed. IP options have been unusable for twenty years[10] as they cause routers to process packets on their slow path. Above IP, the Internet has benefited (or suffered, depending on your viewpoint) from decades of optimizations and security enhancements. To improve performance [2, 7, 18, 3], reduce security exposure [15, 29], enhance control, and work around address space shortages [22], the Internet has experienced an invasion of middleboxes that *do* care about what the packets contain, and perform processing at layer 4 or higher *within* the network.

The problem now faced by designers of new protocols is that there is no longer a well defined or understood way to extend network functionality, short of implementing everything over HTTP[25]. Recently we have been working on adding both multipath support[11] and native encryption[5] to TCP. The obvious way to do this, in both cases, is to use TCP options. In the case of multipath, we would also like to stripe data across more than one path. At the end systems, the protocol design issues were mostly conventional. However, it became increasingly clear that no one, not the IETF, not the network operators, and not the OS vendors, knew what will and what will not pass through all the middleboxes as they are currently deployed and configured. Will TCP options pass unchanged? If the sequence space has holes, what happens? If a retransmission has different data than the original, which arrives? Are TCP segments coalesced or split? These and many more questions are crucial to answer if protocol designers are to extend TCP in a deployable way. Or have we already lost the ability to extend TCP, just like we did two decades ago for IP?

In this paper we present the results from a measurement study conducted from 142 networks in 24 countries, including cellular, WiFi and wired networks, public and private networks, residential, commercial and academic networks. We actively probe the network to elicit middlebox responses that violate the end-to-end transparency of the original Internet architecture. We focus on TCP, not only because it is by far the most widely used transport protocol, but also because while it is known that many middleboxes modify TCP behavior [6], it is not known how prevalent such middleboxes are, nor precisely what the *emergent behavior* is with TCP extensions that were unforeseen by the middlebox designers.

M. Honda, Y. Nishida, C. Raiciu, A. Greenhalgh, M. Handley, and H. Tokuda, "Is it still possible to extend TCP?" Proc. ACM Internet Measurement Conference, Berlin, Nov. 2011. https://dx.doi.org/10.1145/2068816.2068834

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Can the network evolve to address these challenges?

- How to establish connections in a fragmented network?
- How can encryption protect against pervasive monitoring and prevent transport ossification?
- How can we reduce latency and support real-time and interactive content?
- How can we adapt to the vagaries of wireless networks?
- How can we identify and distribute content? How can we manage the tussle for control of the DNS and naming?
- How can we manage interdomain routing to efficiently delivery content?
- How can we re-decentralise the network?

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow