



Networking in the Penumbra

Geoff Huston AM
APNIC

The Trusted Network

Networks enjoyed a privileged position of being able to observe:

- Who is communicating with whom
- What they are saying to each other

The Trusted Network

Users have an expectation of privacy in their communications

- This expectation was often reinforced through regulatory measures intended to constrain public network operators from disclosing knowledge gained through network operation

The Erosion of Trust

Trust has been eroded by intrusive middleware that collects aggregate (and sometimes specific) data on user behaviours

- The general adoption of advertising revenue as a means of funding for service platforms acts as a major incentive to assemble detailed profiles of individual users: age, gender, location, educational level, marital status, income, interest, purchase history,...
- The better the profile, the higher the value of the user to the advertiser

The Erosion of Trust

This network position of trust was further eroded by leakage of the activities of US state-based actors performing various forms of mass surveillance on network users

- The Snowden Papers was a watershed moment for the Internet
- But it was by no means the first time, nor was it the last
- Large scale state-sponsored surveillance continues

So how did we react?

RFC 7258

Pervasive Monitoring is an attack on privacy:

“The IETF community's technical assessment is that PM is an attack on the privacy of Internet users and organisations. The IETF community has expressed strong agreement that PM is an attack that needs to be mitigated where possible, via the design of protocols that make PM significantly more expensive or infeasible.”

RFC 7258 – May 2014

What did this IETF position
mean for the Internet?

Changes to the Applications


Hiding Web Traffic

- Shift to use TLS for all web transactions – HTTPS
 - TLS authenticates the identity of the server to the client
 - Is this service name authentic? Can the service operator demonstrate to the client that it has knowledge of the private part of the key pair that is associated with this DNS service name?
 - Service transactions are encrypted
 - TLS generates a session key used to encrypt all subsequent on-the-wire data

HTTPS Today

HTTP vs. HTTPS

Distribution of HTTP vs. HTTPS requests  




Changes to the Applications

Hiding the DNS

- Hide the query and response in DNS resolution transactions from the network
- The initial work has concentrated on hiding the DNS query names from the network by encrypting the DNS data exchanged
 - DNS over TLS
 - DNS over QUIC
 - DNS over HTTPS/2 and DNS over HTTPS/3

DoH, DOT Today

Cloudflare Open Recursive Resolver DNS Query Profile for World (XA)



DNS over HTTPS
DNS over TLS

<https://stats.labs.apnic.net/edns>

Can we go further?

- Can we hide the two ends from each other such that at no point in the network (and even at the server) are the two ends of the transaction visible at once?
- Can we also selectively obscure the content of the transaction such that the endpoints and the content of the transaction are not simultaneously discoverable

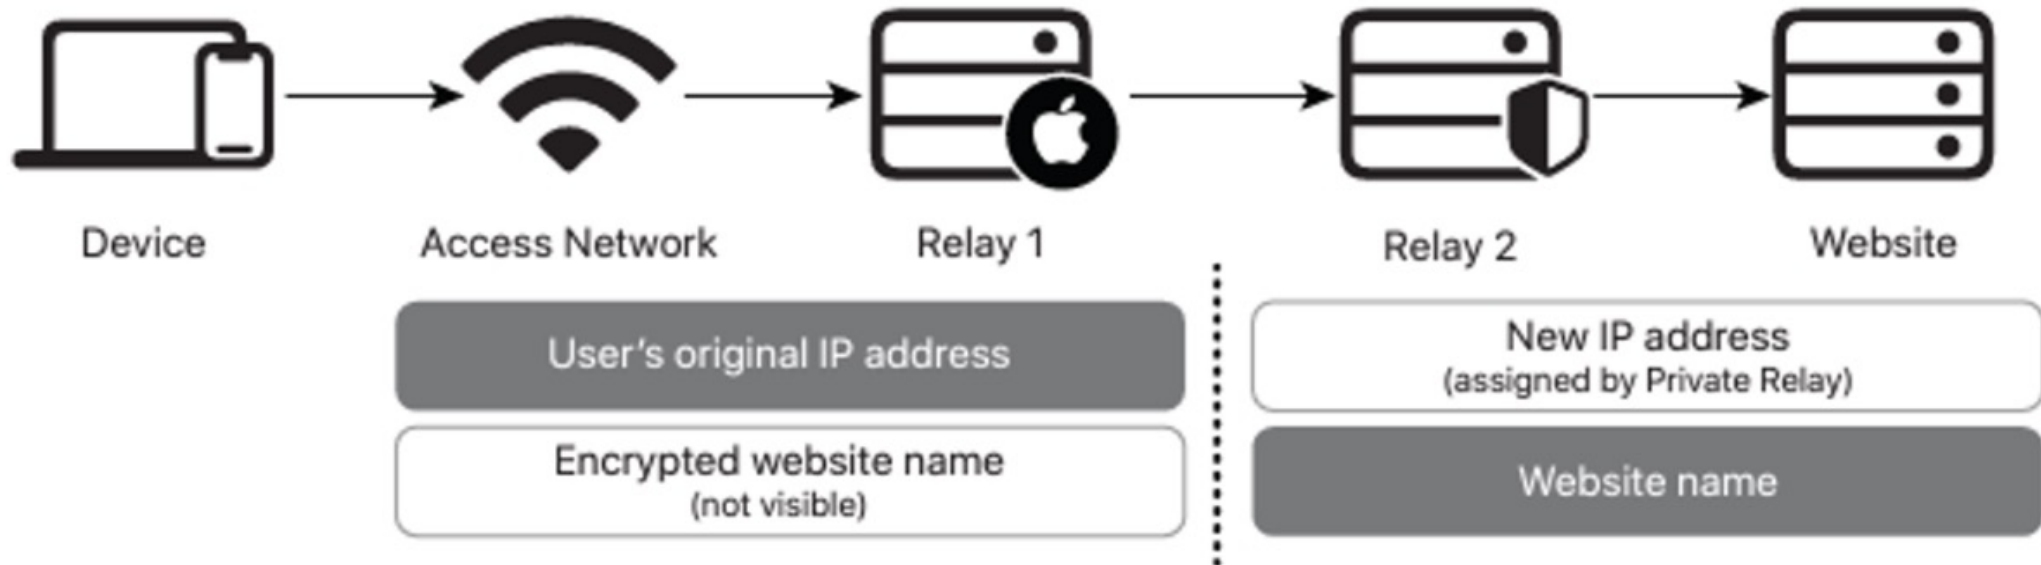
MASQUE and Relays

- With the use of 2-layer encryption and active relays then its possible to hide the endpoints from the network
- There is no single network observation point that can put together the combination of the service identity and the identification of the two endpoints of the service transaction
- Only the client endpoint knows its own identity and service, but does not know the identity used by the relay to present the service transaction to the server
- The server may use the application-level identity of the client, but does not know the client's network-level identity (IP address)
- This technique can be used in DNS resolution and HTTPS transactions

Apple Private Relay

When Private Relay is enabled, your requests are sent through two separate, secure internet relays.

- Your IP address is visible to your network provider and to the first relay, which is operated by Apple. Your DNS records are encrypted, so neither party can see the address of the website you're trying to visit.
- The second relay, which is operated by a third-party content provider, generates a temporary IP address, decrypts the name of the website you requested, and connects you to the site.



Sealing up the Peepholes

Attention has turned to the Server Name Indication (SNI) field in the TLS handshake

- This is the one last part of TLS that is still shown in the clear
- Efforts to encrypt this field in a robust manner are being studied
- The most effective way to securely communicate the public key that is used to encrypt the SNI (and the entire ClientHello message) appears to be a TLSA record in the DNS (DANE) using DoT or DoH, using a DNSSEC-signed record
- A shortcut hack is to use a trusted intermediary
(<https://blog.cloudflare.com/announcing-encrypted-client-hello/>)

Sealing up the Peepholes

- And there's the the Online Certificate Status Protocol, which can expose the IP address of the client and the name of the service that they are visiting to the CA
 - Which likely explains why Chrome browsers do not perform “live” certificate revocation checks, and rely instead on short validity periods for certificates*

* Which is probably just as bad, but in a different way!

Why are we doing this?

Who wants privacy?

Do users really care?

- Users cheerfully gave up email privacy in exchange for free email services
- Users happily tell Google Search way too much about themselves in exchange for instant answers
- In general, users will happily trade off privacy for access to services



If not users, then whom?

If not users, then whom?



The folk with the most to gain (or lose)!

2023 [\[edit \]](#)

This list is up to date as of 31 December 2023. Indicated changes in market value are relative to the previous quarter.

Rank	First quarter	Second quarter	Third quarter	Fourth quarter
1	 Apple ▲2,609,000 ^[30]	 Apple ▲3,050,000 ^[30]	 Apple ▼2,677,000 ^[30]	 Apple ▲2,994,000 ^[30]
2	 Microsoft ▲2,146,000 ^[31]	 Microsoft ▲2,532,000 ^[31]	 Microsoft ▼2,346,000 ^[31]	 Microsoft ▲2,795,000 ^[31]
3	 Alphabet ▲1,332,000 ^[32]	 Alphabet ▲1,530,000 ^[32]	 Alphabet ▲1,662,000 ^[32]	 Alphabet ▲1,764,000 ^[32]
4	 Amazon ▲1,058,000 ^[33]	 Amazon ▲1,337,000 ^[33]	 Amazon ▼1,312,000 ^[33]	 Amazon ▲1,570,000 ^[33]
5	 Nvidia ▲686,090 ^[34]	 Nvidia ▲1,044,000 ^[34]	 Nvidia ▲1,074,000 ^[34]	 Nvidia ▲1,223,000 ^[34]
6	 Berkshire Hathaway ▼677,770 ^[35]	 Tesla ▲829,670 ^[36]	 Tesla ▼794,200 ^[36]	 Meta ▲909,000 ^[37]
7	 Tesla ▲656,420 ^[36]	 Berkshire Hathaway ▲745,010 ^[35]	 Meta ▲772,490 ^[37]	 Tesla ▼789,930 ^[36]
8	 Meta ▲549,480 ^[38]	 Meta ▲735,450 ^[38]	 Berkshire Hathaway ▲769,260 ^[35]	 Berkshire Hathaway ▲783,550 ^[35]
9	 TSMC ▲482,410 ^[39]	 TSMC ▲523,410 ^[39]	 Eli Lilly ▲509,890 ^[40]	 Eli Lilly ▲553,370 ^[40]
10	 Visa ▲473,870 ^[41]	 Visa ▲497,370 ^[41]	 Visa ▼480,990 ^[41]	 TSMC ▲539,390 ^[39]

Surveillance Capitalism



Who cares about privacy?

- None of the entities who spend large sums to assemble detailed profiles of users want to leak that data to their competitors
- So, privacy is about protecting the core asset of gathering individual profiles of users from:
 - Other services
 - The common host platform
 - Common Infrastructure services
 - The network

Let me rephrase that:

We want to allow **the application** to operate in a mode that obscures its behaviour from:

- Other services
- The common host platform
- Common Infrastructure services
- The network

How do you do that?

By lifting out as much as you can from the lower levels of the protocol stack that are managed by common services and performing it within the application

So, how do you do that?

Transport Privacy

Which means we are looking at how to lift TCP out of the common parts of the host platform and and shift it across to the application

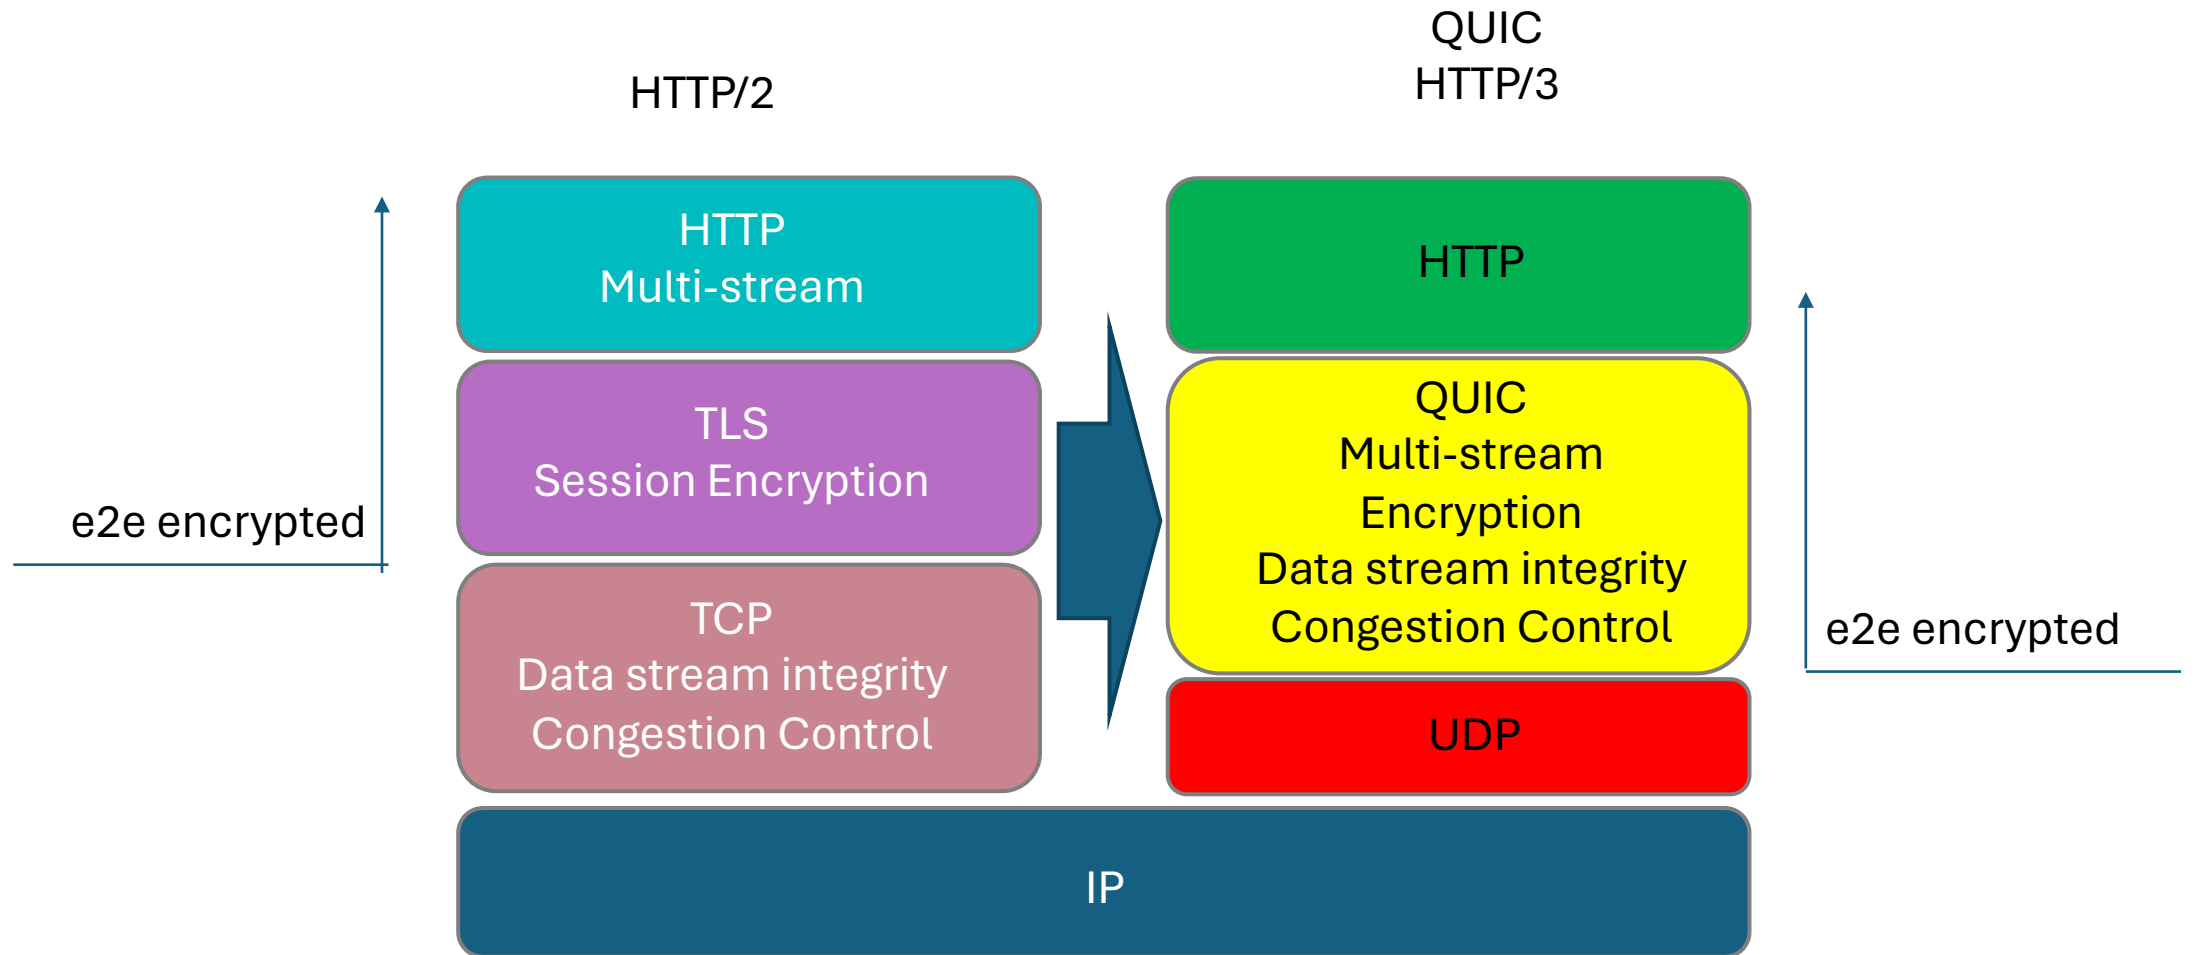
We need to change TCP!

Transport Surgery

How do you change TCP?

- TCP is a kernel function that is defined at the platform level
- Applications have no intrinsic ability to alter the TCP characteristics for the application on a customized basis
- You could try to define a new transport protocol (such as SCTP)
- But the deployed infrastructure (NATs) tends to discard all packets that are not protocol 6 (TCP) or protocol 17 (UDP)
- If you want to bypass kernel handling of TCP and get through existing network filters and middleware then you're forced into using UDP
- So, you change TCP by using UDP!

QUIC is the new TCP



QUIC is:

- A logical evolutionary step for transport services, providing more flexibility, faster connection setup, and a larger set of transport services
- It's what we should expect from a capable modern transport protocol!


Cloudflare's Numbers

Cloudflare report on observed use levels



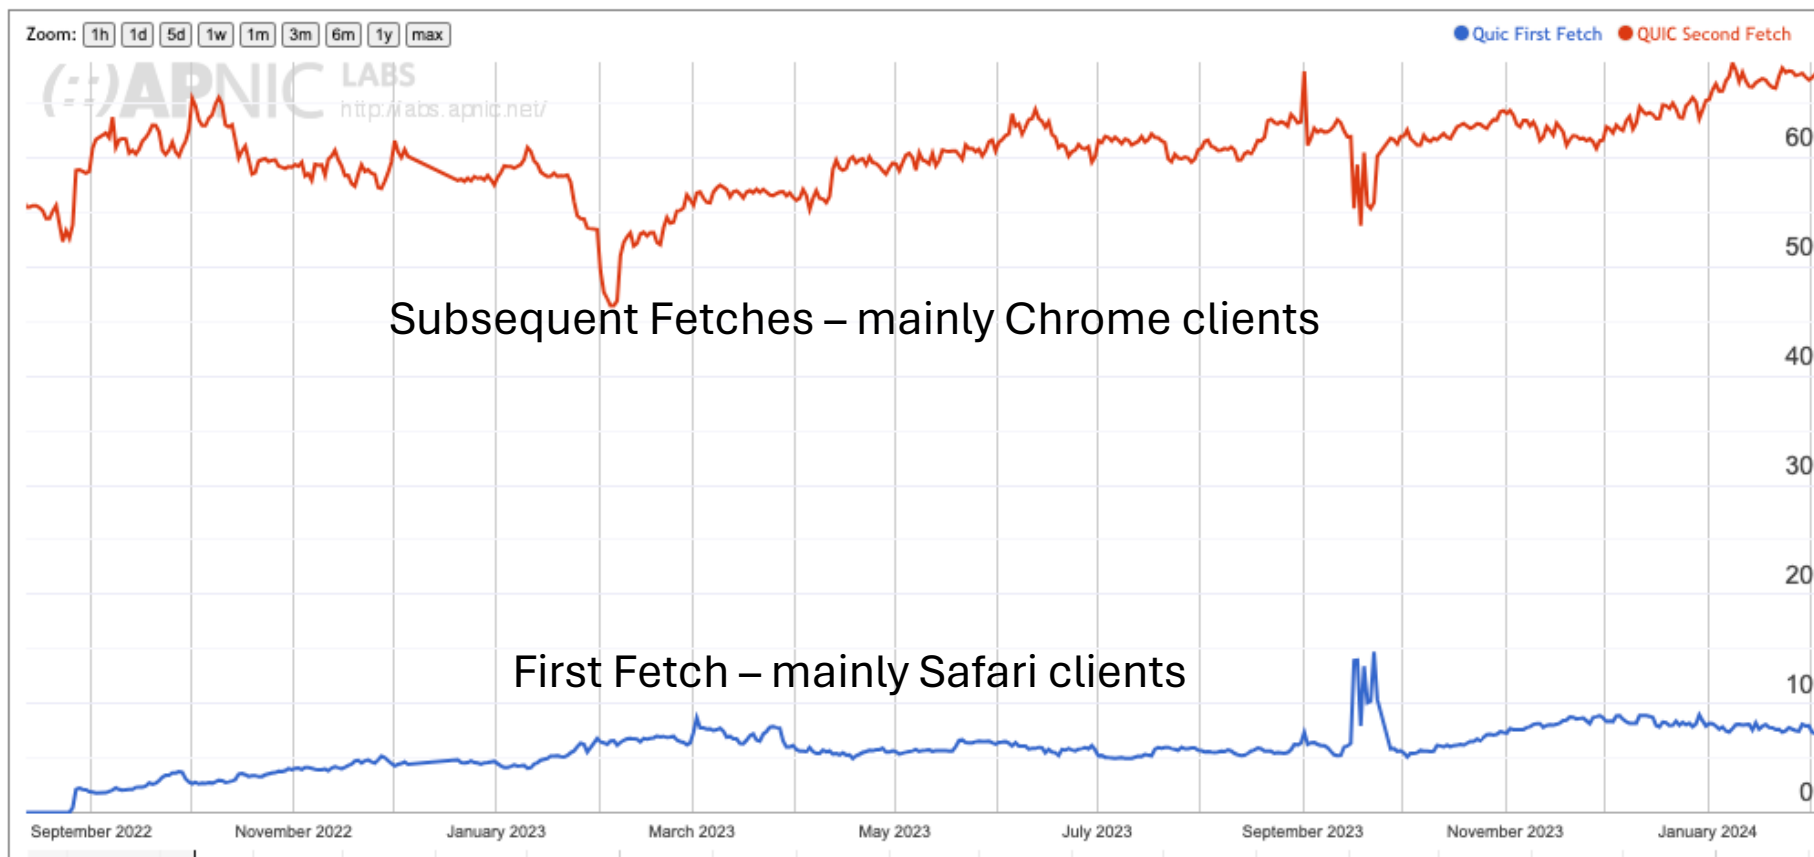
TLS 1.2 vs. TLS 1.3 vs. QUIC

Distribution of secure traffic by protocol [?](#) [🔗](#)

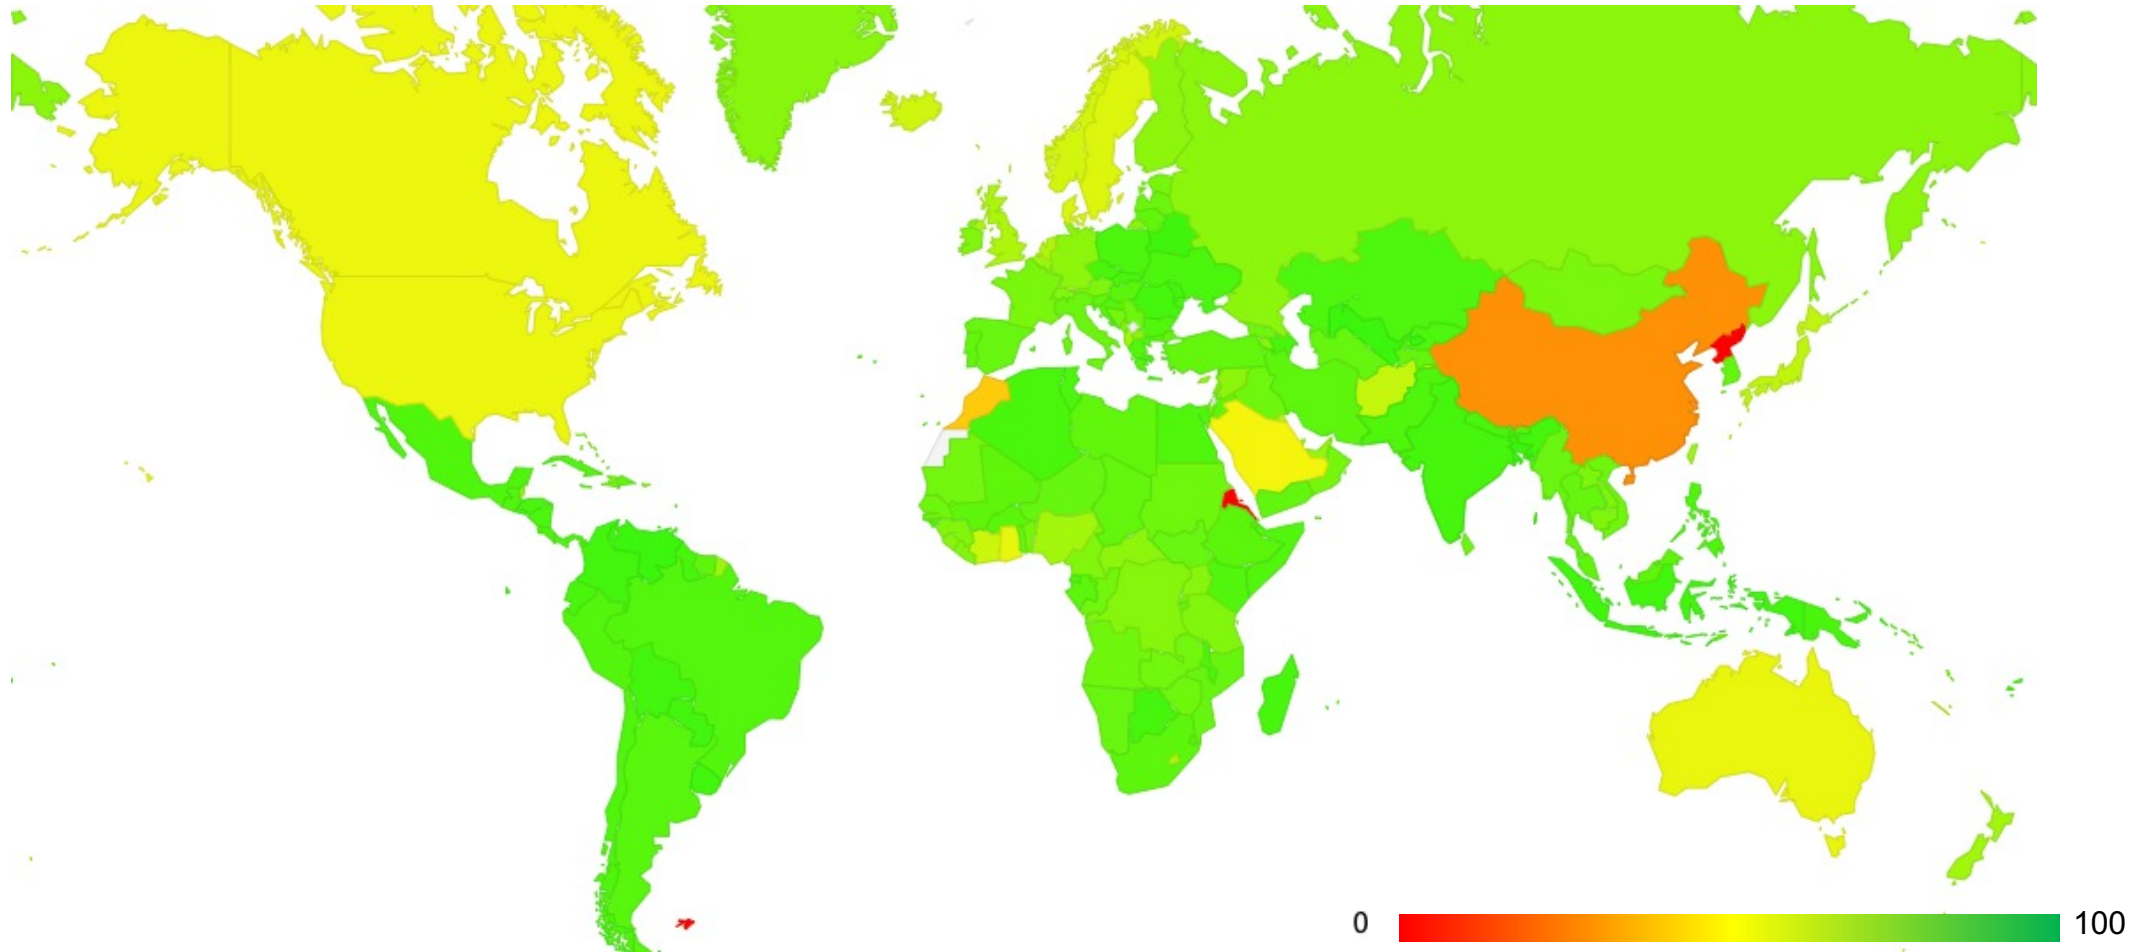


APNIC's Numbers


APNIC reports on *potential* use levels



APNIC's QUIC Numbers



Cisco's Numbers: Traffic Volume



*source EU Operator 2022

Why is QUIC important?

Because QUIC is **faster**

Because QUIC **encrypts everything**

- No visible transport control settings
- No visible Server Name Indication in the crypto-setup
- No visible traffic profile other than inter-packet timing

Because QUIC is an **application capability**

- QUIC can interact with the platform through the UDP API, so all of QUIC can be implemented within the application. This gives the application more control over its service outcomes and reduces external dependencies

What does this mean for TCP?


It's not looking all that good for TCP's prospects

- QUIC not only does faster start up, but it supports multi-channel in a frictionless manner
- QUIC resists network operator efforts to perform traffic shaping through direct manipulation of TCP control parameters
- QUIC allows the application service provider to control the congestion behaviour of its sessions

The new Networking Space

QUIC is pushing both network carriage and host platform into commodity roles in networking and allowing applications to effectively customize the way in which they want to deliver services and dominating the entire networked environment

QUIC is the application's view of what Transport should be!



What does this mean for the Internet?

- The relationship between applications, hosts and networks has soured into mutual distrust and suspicion
- The application now defends its integrity by wrapping up as much of the service transaction with encryption and indirection
- QUIC (and MASQUE) is an intrinsic part of this process of wrapping up traffic in encryption and redirection
- For the network operator there is little left to see or do. It's just packet shovelling!
- And I suspect that there is no coming back from here!

What can a Network Operator Do?

- When **all** customer traffic is completely obscured and encrypted?
 - Traffic Shaping?
 - Regulatory Requirements for traffic interception?
 - Load Balancing / ECMP

The new Internet Space

“What you can’t dominate, you commoditise*”

* A related quote is Peter Thiel’s “Competition is for Losers!”

- Vertically integrated service providers have faded away into history - the deregulated competitive service industry continues to specialize rather than generalize at every level
- Carriage is no longer an inescapable monopoly - massively replicated content can be used as a substitute for many carriage service elements
- Control over the platform is no longer control over the user. Operating systems have been pushed back into a basic task scheduling role, while functions are being absorbed into the application space

The new Internet Space

- Each service has an ability to define its own operational behaviours that are intrinsic to that service
 - Which is the antithesis of “interoperability”
- We have managed to minimize and commoditize the common parts of the Internet and push the valued functionality and service delivery up into each application
- Which means:
 - “Standards for Interoperability” is dead!
 - “Open” is a quaint historic notion!

The New Internet Space

The hyperscalers have won **everything!**

Thanks!



2024
APRICOT
APNIC 57

BANGKOK, THAILAND

21 February – 1 March 2024

[#apricot2024](#)

