DNS "Openness"

Geoff Huston AM
APNIC Labs

Openness?

When speaking about *openness*, we do not mean *open* in a competitive sense, but rather:

> "Can users access and distribute information and content, use and provide applications and services, and use terminal equipment of their choice, irrespective of the user's or provider's location or the location, origin or destination of the information, content, application or service?"

But that was the entire POINT of the DNS!

The DNS was engineered to deliver the same answer to the same query, irrespective of the querier

- The answer did not depend on who was asking, where they were asking from, what platform they were using to generate the query, the resolvers they used to handle the query
- The answer did not depend on the origin of the information used to form the response, the platform used to serve this information, nor the location of information servers

Openness?

When speaking about openness we do not mean in a competitive sense, but rather:

> “Can users access and distribute information and content, use and provide applications and services, and use terminal equipment of their choice, irrespective of the user’s or provider’s location or the location, origin or destination of the information, content, application or service?”

The answer is ”Yes!”

Really?

- Is that really true? Do we all see the same DNS?
- Really?

Really?

- Is that really true? Do we all see the same DNS?
- Really?

Really? Not!

Why not:

- Government regulatory requirements to block the “correct” resolution of certain DNS names
  - It happens in China, UK, America, Australia, India, Russia, Syria, Iran, Vietnam, France, Turkey,…..
  - Its VERY widespread, for all kinds of national motives
- Occasional ISP desires to monetise the DNS
  - NXDOMAIN substitution to direct traffic from named destinations or services that do not exist to other destinations or services of their choosing
- Threat mitigation where the DNS names associated with malware are blocked
  - Such as Quad9 threat intelligence informed DNS resolution

Is this a "problem"?

Generally not:

- Nation states have a sovereign right to make such rules and bind their citizens to such rules
- Threat mitigation is typically regarded as a Good Thing rather than an incursion against the utility of a single DNS

- It's a problem when it gets used as a lever in a different fight
  - Such as the Australian rule to force Australian ISPs to block the DNS resolution of "thepiratebay.org"
    - It only pushes determined users to alternate name resolution strategies and ultimately is a comprehensive waste of everyone's time!
  - But even this is a relative sideshow to the larger DNS

Another interpretation of "Openness"?

When speaking about *openness* we might not mean *open* in a sense of consistency across providers and across client transactions, but rather:

Is the DNS an “open” system?

This is an echo of the 1980’s move by the industry embrace “open” technology, as in *Open* Systems Interconnect (OSI) for networking, *open* computer architectures, *open* operating systems. “*Open*” is being used here in the sense that it is intentionally positioned as the opposite of a vendor-proprietary “closed” technology.

Is the DNS "Open"?

Yes!

- The DNS name resolution protocol is openly specified without any IPR encumbrance
- Fully functional implementations of the DNS protocol are available as open source
- DNS name servers are configured as open “promiscuous” responders and will provide the same response to a query irrespective of the identity of the querier
- DNS information is openly available
  - There is some subtle qualification here in that the collection of a zone file may not be openly available, but the individual records in a zone can be queried
- DNS queries and responses are “open”

Is the DNS "Open"?

Yes!

- The DNS name resolution protocol is openly specified without any IPR encumbrance
- Fully functional implementations of the DNS protocol are available as open source
- DNS name servers are configured as open “promiscuous” responders and will provide the same response to a query irrespective of the identity of the querier
- DNS information is openly available
  - There is some subtle qualification here in that the collection of a zone file may not be openly available, but the individual records in a zone can be queried
- DNS queries and responses are “open”
  - **Which is a MASSIVE problem**!

When "openness" is a weakness

The DNS is used by **everyone** and **everything**

- Because pretty much everything you do on the net starts with a call to the DNS
- If we could see your stream of DNS queries in real time we could easily assemble a detailed profile of you and your interests and activities - as it happens!
- If we could edit your DNS responses we could make services disappear from your Internet!

https://xkcd.com/1361/

Lets look into this further

The DNS is mapping system which takes human-use labels that name services and maps these labels to IP addresses

Lets look into this further

The DNS is mapping system which takes human-use labels that name services and maps these labels to IP addresses

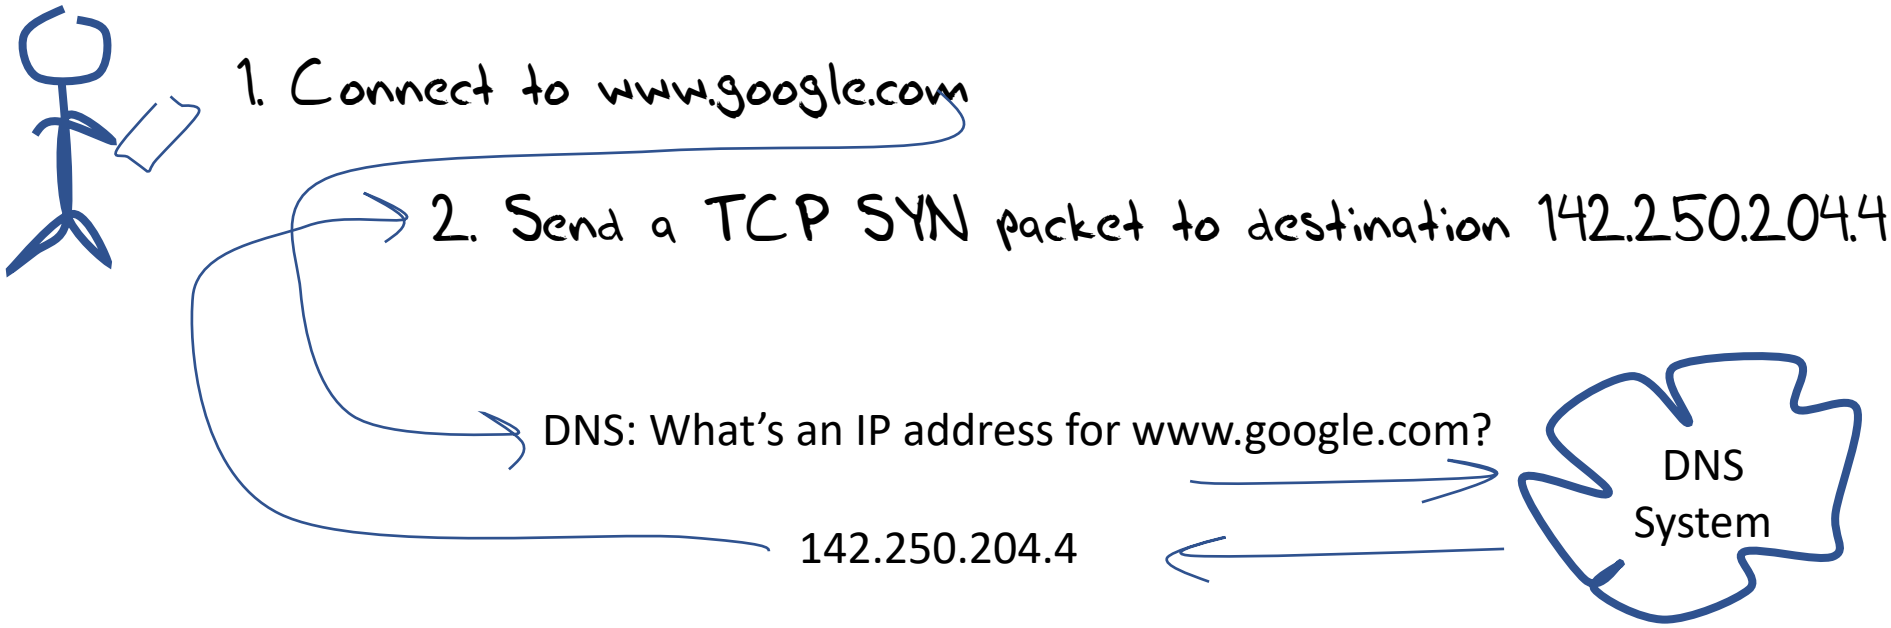

What's in that DNS "cloud"?

How the idealised model of the DNS works

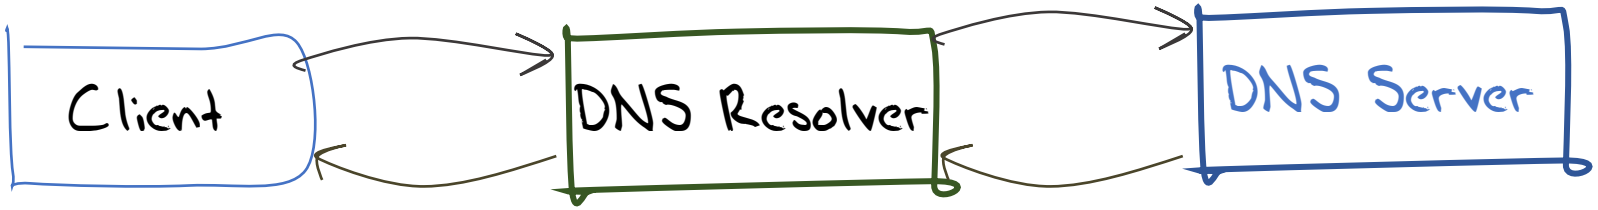

Clients, Resolvers and Servers

- Clients send their query to a resolver
  - The resolver's addresses was provided to the client by their ISP, or the user configured it directly into their device
- When it receives a query, the resolver first must work out whom to ask (discover the *authoritative server* for this domain name) and then it will direct a query to this server
- The resolver will use this response to answer the original query from the client
- And the resolver will also cache the answer to allow it to reuse this information if it is asked the same query in the future.

Clients, Resolvers and Servers

- Clients send their query to a resolver
  - The resolver's addresses was provided to the client by [illegible] user configured it directly into their device
- When it receives a query, the resolver will work out whom to ask (discover the *authoritative* servers for this domain name) and then it will direct a query to them
- The resolver will send a response to answer the original query from the client
- A resolver will also cache the answer to allow it to reuse this information if it is asked the same query in the future.

But you knew all this already!

Scaling DNS infrastructure

- Large scale DNS resolvers are generally implemented as a DNS server “farm”, where incoming queries are farmed across multiple resolver engines
- Very large DNS resolvers use a common front end service IP address farms, and use anycast to perform query load balancing over multiple distributed resolver farms
- Authoritative servers have also taken to anycast, both as a scaling option and a DOS defence

What's REALLY in that DNS "cloud"?

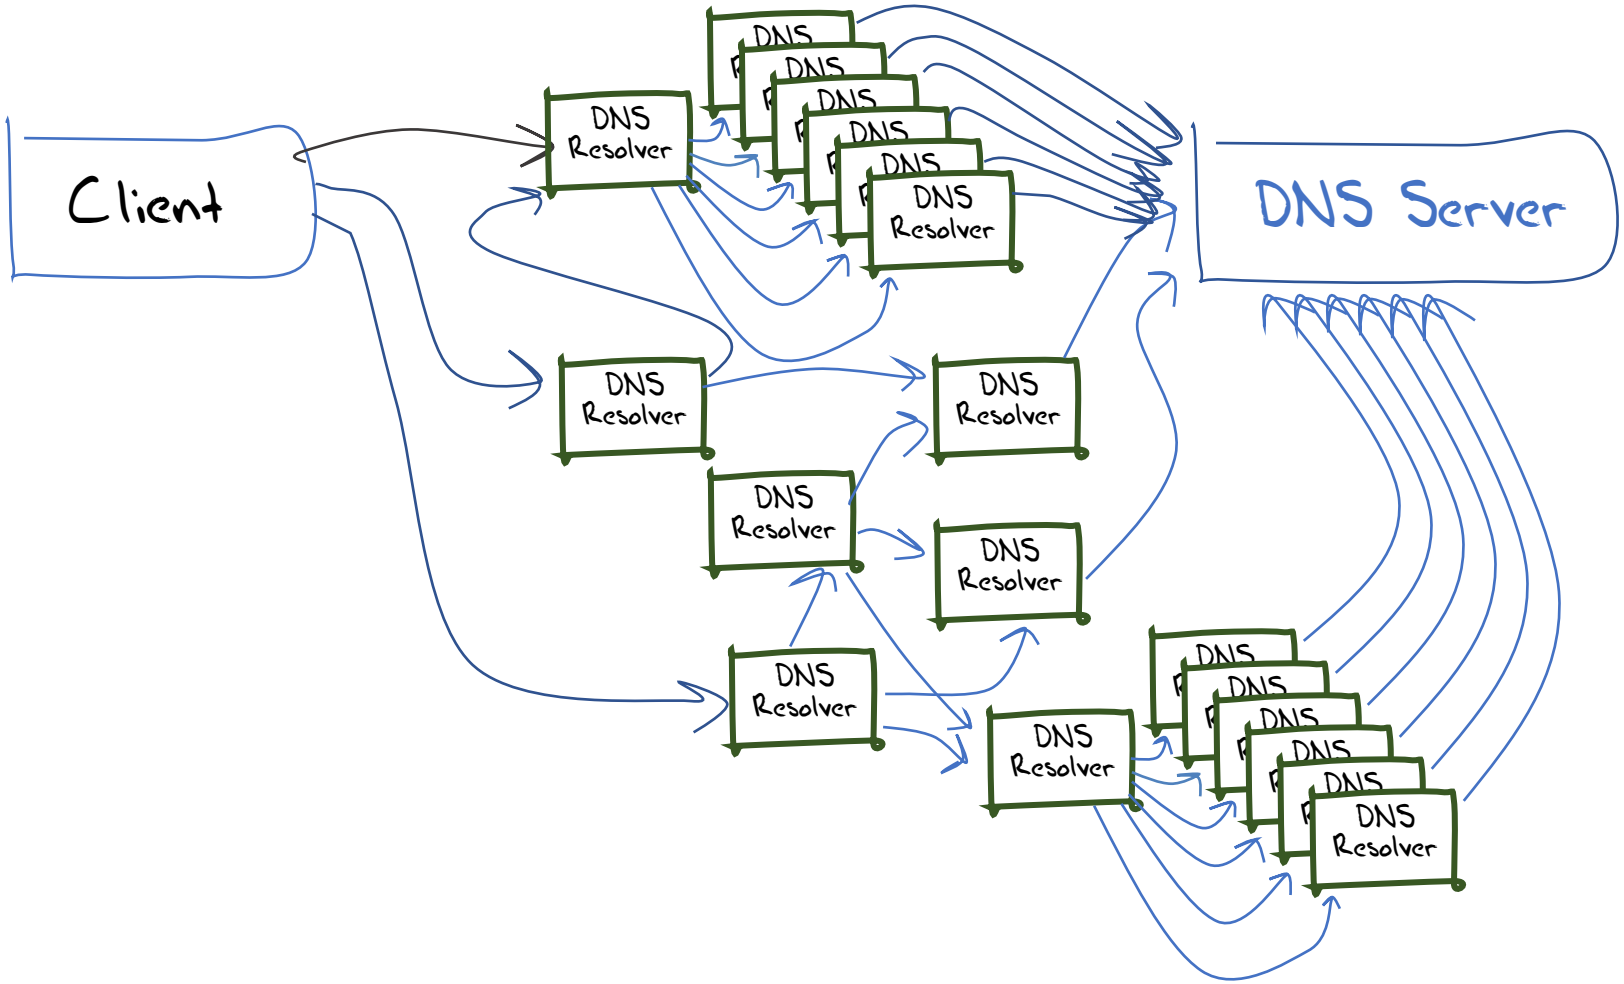

The Hidden parts of "Open"

- For an open technology the infrastructure and process of name resolution in the DNS is incredibly opaque
  - Once queries are passed to the DNS resolution infrastructure they are impossible trace. Each element does not simply forward a query, but recasts a new query using its own identity
  - There is no query tracing and no clear way to probe into the DNS infrastructure
  - There could be query chains circling in the DNS infrastructure in a hidden loop and no one would be any the wiser!
  - It is challenging to gauge the level of hidden interdependency in the DNS, nor the level of ongoing centralisation of infrastructure functions in the DNS

But all that is still not all of "the DNS"

- The DNS is more than its mapping role, and more than its infrastructure elements of resolvers and servers
- So what is the DNS?

What is "the DNS"?

- **A name space**: A collection of word-strings that are organised into a hierarchy of labels
- A distributed **name registration** framework that assigns a unique "license to use" to human-centric word-strings to entities (for money)
- A **distributed database** that maps human-centric word-strings into IP addresses
- A **protocol** used by DNS protocol speakers to "resolve" a word-string into a defined attribute (usually an IP address)
- A **signalling medium** that is universally supported across all of the Internet

Orchestration of the DNS

- If the DNS is a set of functions and a set of various actors in this space then how are their individual actions orchestrated to provide a cohesive outcome?
- How can client use the functions of “the DNS” if there is no one orchestrating all these elements of the name infrastructure?
- Why does all this work in a completely deregulated space?

- The answers lie in Markets and Market Signalling

What are DNS "Markets"?

The DNS is not a single market – it is a highly devolved framework and there are a number of discrete markets that are at best loosely coupled .

Some of these are:

- The market for new “top level” labels (gTLDs) operated by ICANN. This market is open to ICANN-qualified registry operators. A registry has an exclusive license to operate a TLD.
- The market for “registrars”, who act as retailers of DNS names and deal with clients (registrants) and register the client’s DNS names into the appropriate registry
- The market for clients to register a DNS name with a registry
- The market for DNS name certification, which is a third party that attests that an entity has control of a domain name
- The market for DNS name resolution where users direct their queries to a resolver and the resolver provides DNS “answers”
- The market for hosting authoritative name services, where “bigger is better” has driven a highly aggregated market
- The market for DNS query logs

Maybe it's more than markets

- Perhaps this is best viewed as a collection of market needs, or requirements
  - Some are well established
  - Some are emergent requirements
- Perhaps the question could be rephrased as one that asks to what extent are conventional open markets a good fit for these various DNS requirements
- And to what extent these market-based mechanisms are failing to respond to these needs?

So, what should we talk about?

If we want to talk about the DNS as a deregulated activity orchestrated through the operation of open markets then there are a few more things to bear in mind:

- The DNS is not a single market place, or even a collection of inter-dependent and tightly coupled market place
- The DNS is constructed of many elements, some of which appear to behave as tightly regulated markets, some of which are openly competitive markets, some of which are comprehensive market failures!

- Maybe we should think of the DNS using a number of themes to give some focus to this consideration of market effectiveness

Current DNS Themes

There are so many, and here are just a few:

- DNS as a control element
- DNS and privacy
- DNS and trust
- DNS and name space fragmentation
- DNS as a rendezvous tool
- DNS as a collection of markets
- DNS and market aggregation
- DNS and abuse and cyber attacks
- DNS as an economic failure
- DNS as the last remaining definition of a coherent Internet

This is now a very big agenda

- Far bigger than we have time for here!
- So I'll just take a couple of themes and develop them further

I. The Open Market for Trust

DNS and Trust

Can you trust what you learn from the DNS?

DNS and Trust

Can you trust what you learn from the DNS?

NO!

DNS and Trust

Can you trust what you learn from the DNS?

NO!

- DNS responses can be altered in various ways that are challenging to detect
- We know how to improve this situation by using digital cryptography to protect the integrity, accuracy and currency of DNS responses
- But does this capability for an improvement in the trust of the DNS correlate to a visible consumer preference?
- Is security in the DNS a market failure?

DNSSEC

- A framework to attach digital signatures to DNS responses that attest to the accuracy and currency of the DNS response
  - The method of attachment doe not alter the behaviour of the DNS protocol, nor does it require any changes to DNS servers
- A procedure for clients to follow to validate the DNS response through the processing of this digital signature

Changes:

- Zone management – add DNSSEC digital signature records through ”zone signing”
- Registry management – add DS records alongside NS records for delegated zones
- Client behaviour – perform additional DNS queries to perform digital signature validation

Is DNSSEC being used?

You might think that a change to the DNS that improved the trust in the DNS would prove to be highly popular in the DNS space

- DNS name “owners” would like their name to be trustworthy
- DNS users would like the names they use to be “genuine” and trustable

So **every** client of the DNS would want to use DNSSEC!

Right?

So let’s see if this requirement has been translated by the DNS service market into a service offering by looking at the current metrics of the use of DNSSEC

Is DNSSEC being used?

Who validates DNS responses?

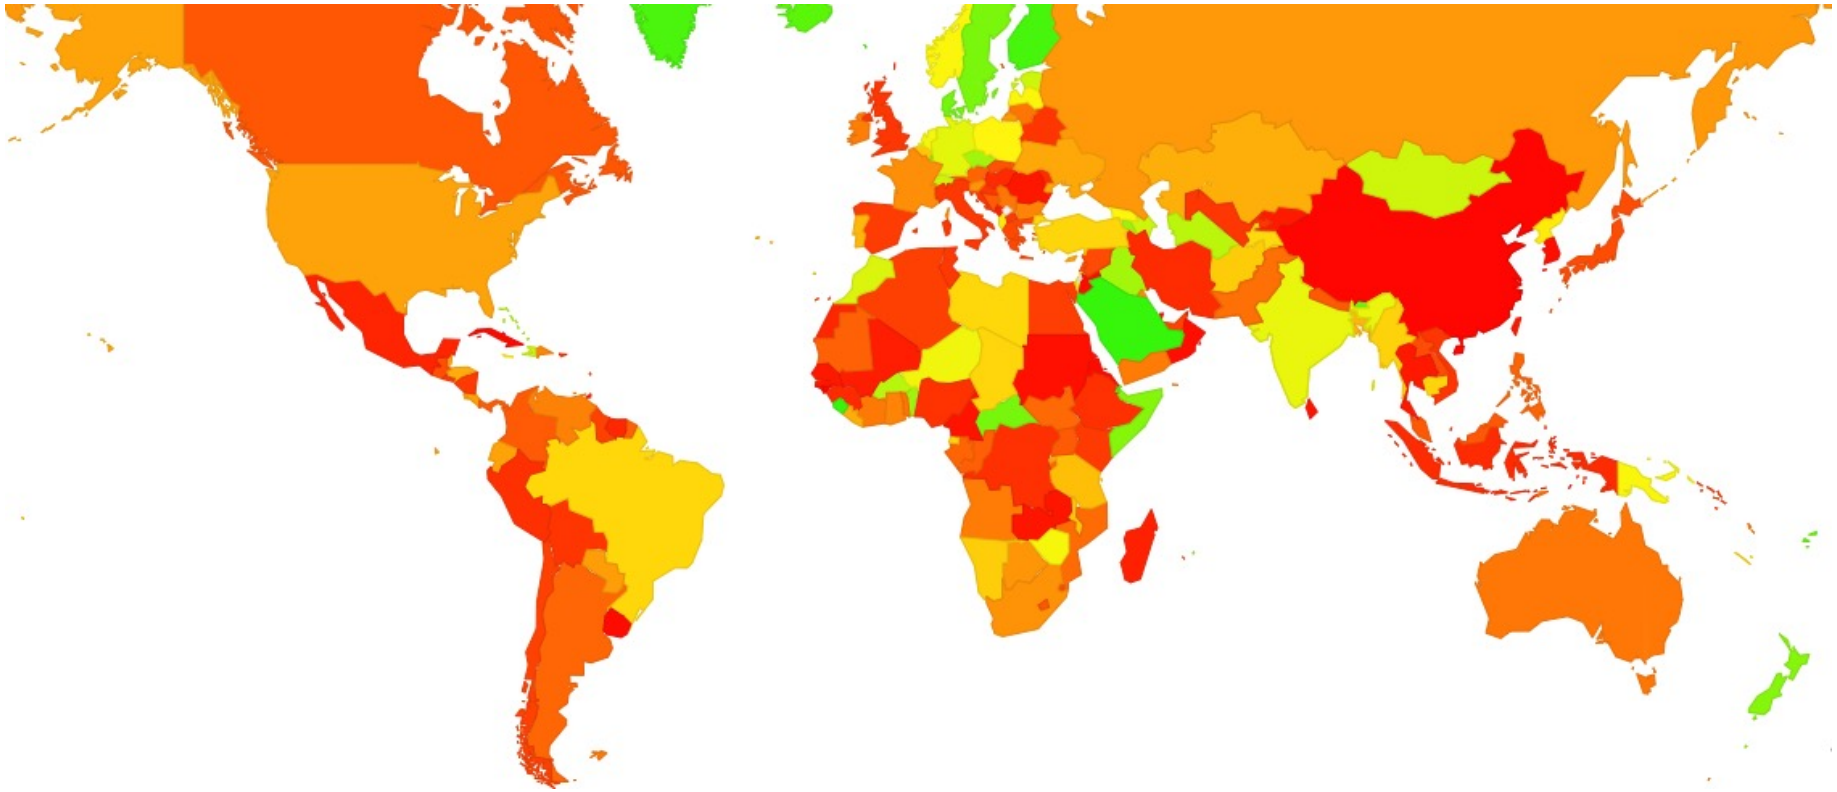

Is DNSSEC being used?

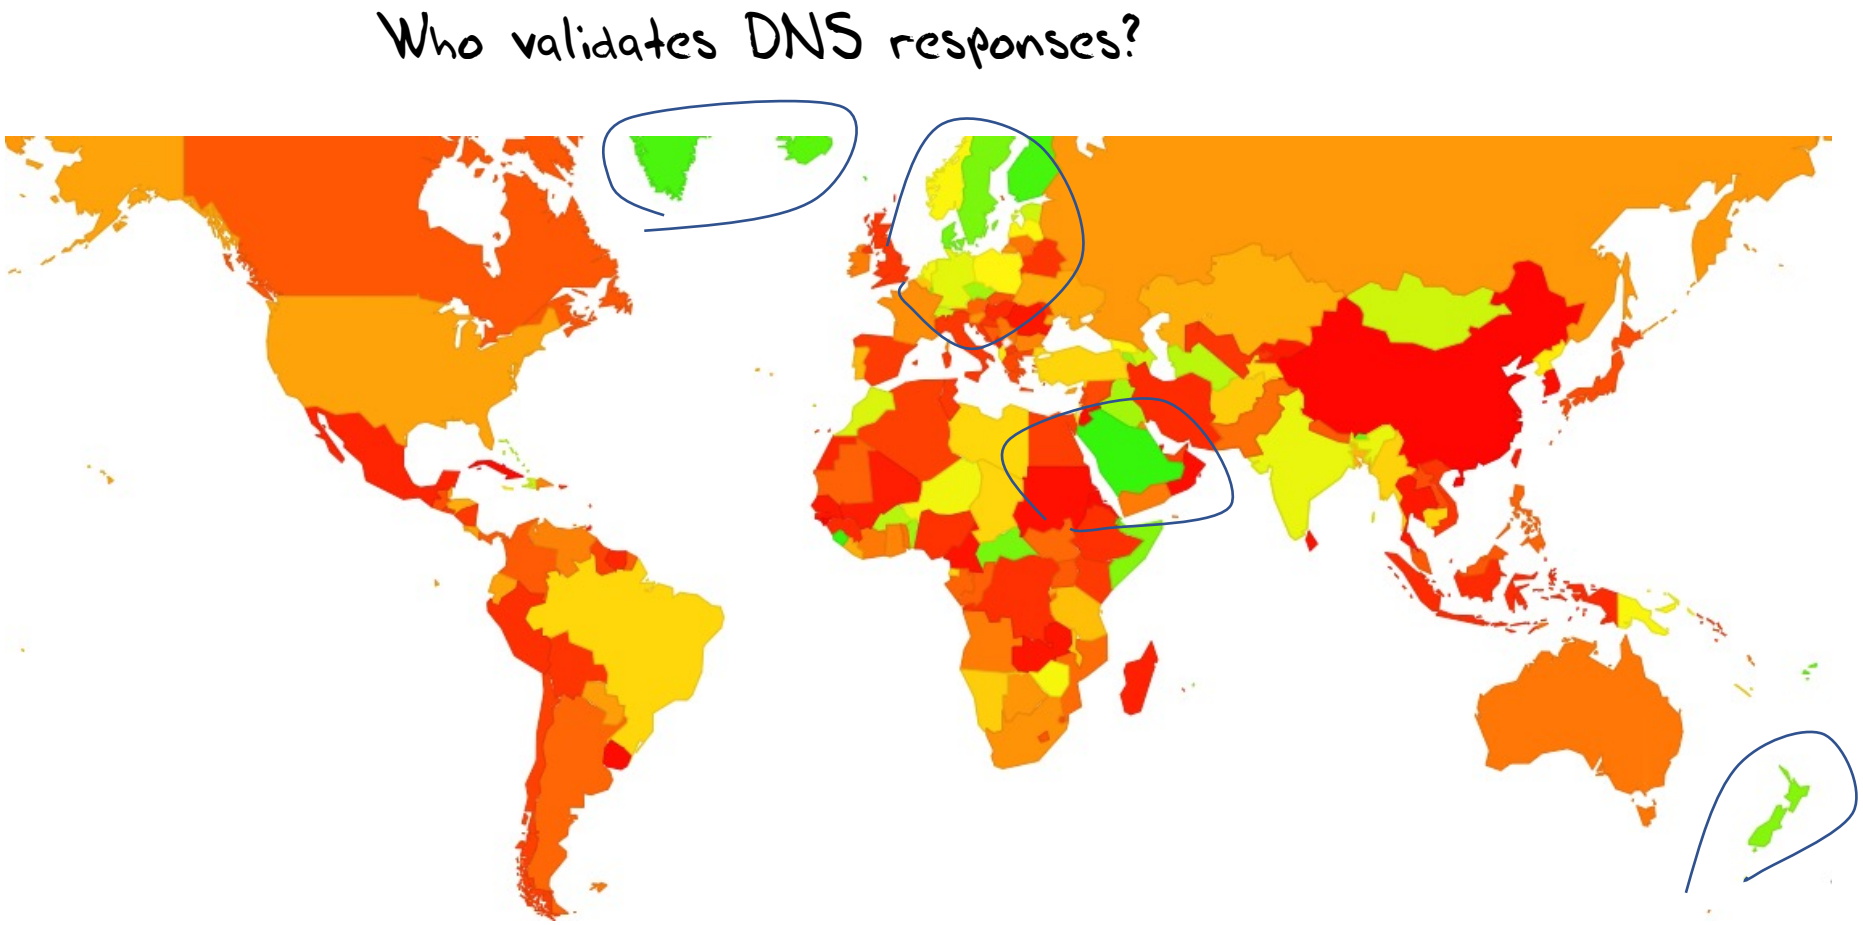

Is DNSSEC being used?

Validation Rate of Signed DNS responses

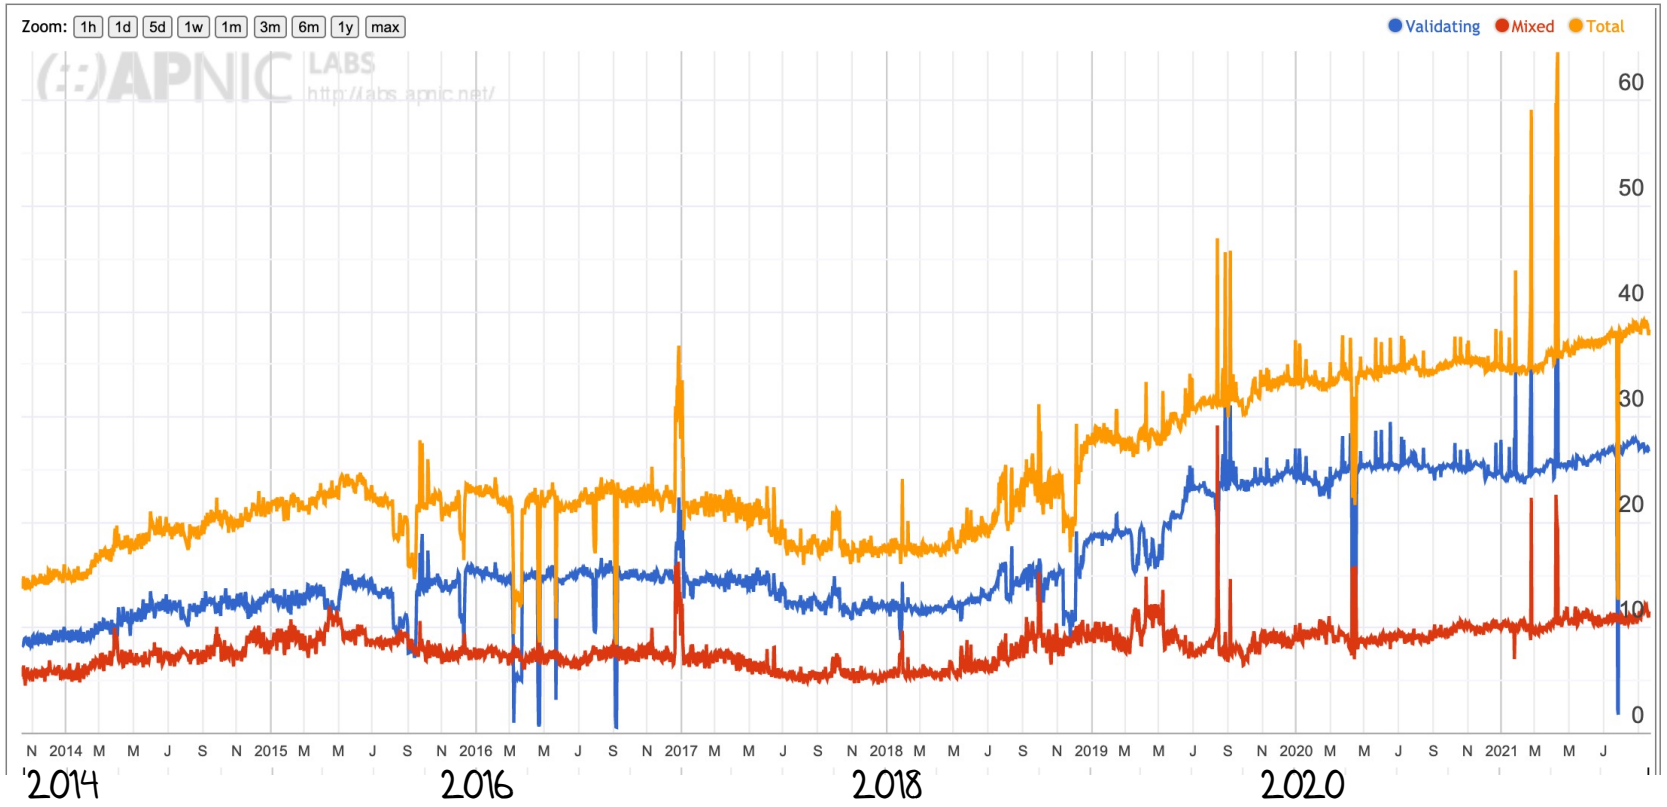

25÷ of users are behind DNSSEC-validating resolvers who will not resolve a badly signed DNS name

Is DNSSEC being used?

Validation Rate of Signed DNS responses

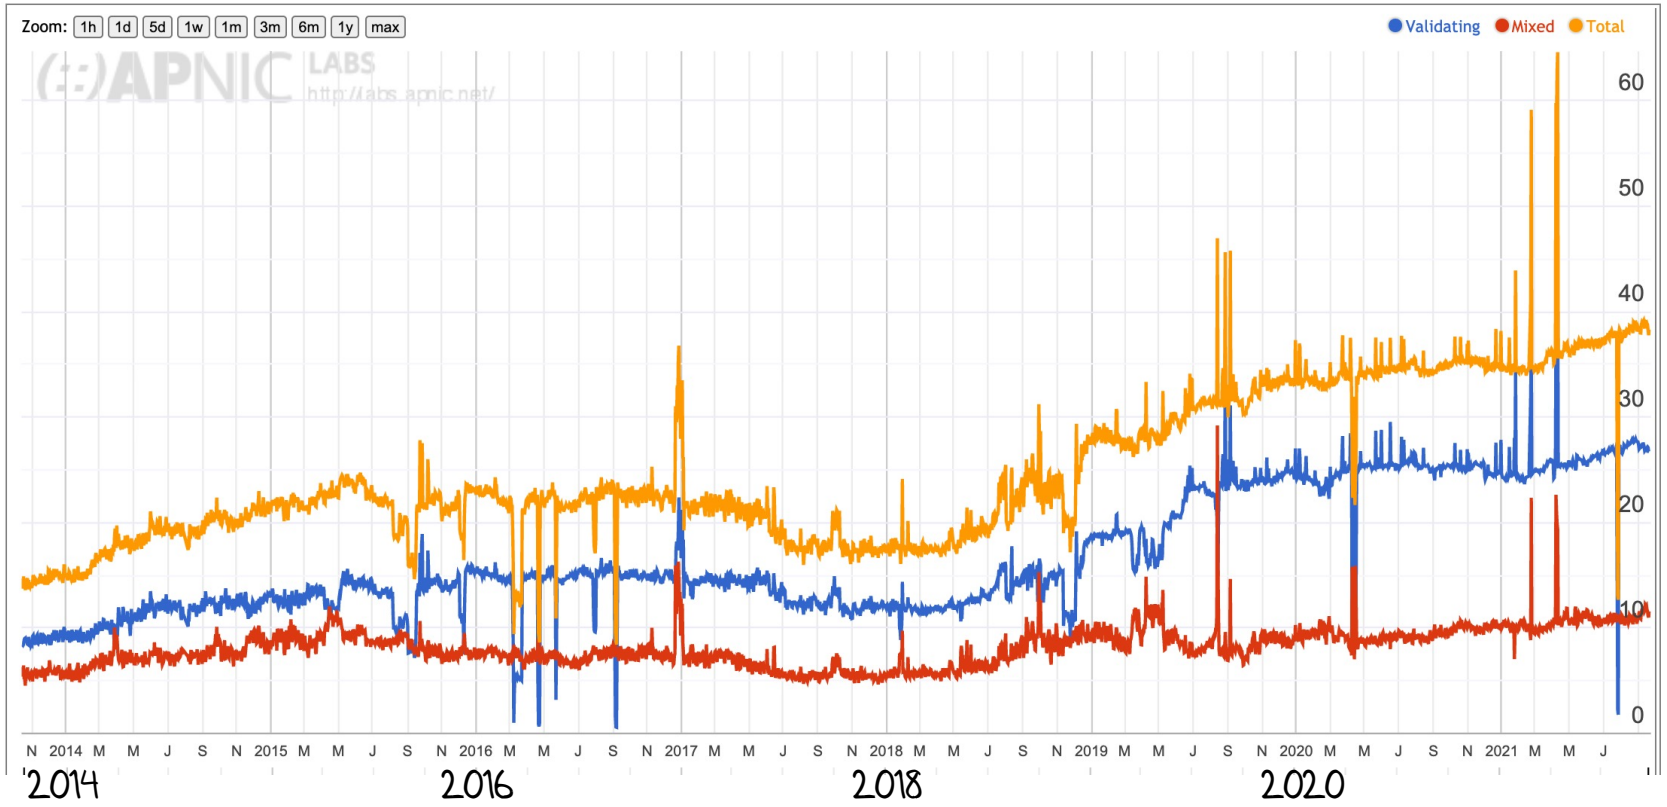

Three quarters of the Internet's user base can't or won't check if a DNS response is genuine

Problems with DNSSEC

- Large DNS responses cause robustness issues for DNS
  - Getting large responses through the network has reliability issues with UDP packet fragmentation and timing issues with signalled cut-over to TCP
  - The validator has to perform a full backtrace query sequence to assemble the full DNSSEC signature chain
  - So the problem is that DNSSEC validation may entail a sequence of queries where each of the responses may require encounter UDP fragmentation packet loss

All this adds to the time to resolve a signed name

And nobody is tolerant of delays in today's Internet

Some More Problems with DNSSEC

- Cryptographically “stronger” keys tend to be bigger keys over time, so the issue of cramming more data into DNS transactions is not going away!
- The stub-to-recursive hop is generally not using validation, so the user ends up trusting the validating recursive resolver in any case

The current DNSSEC framework represents a lot of effort for only a a tiny marginal gain

DNSSEC is a Market Failure!

- Users don’t pay for queries
  - Users have no leverage with recursive resolvers in terms of expressing their preference for authenticity in DNS responses
- Users don’t have a choice in what they query for
  - Users have no ability to express a preference for only using domain names that are signed
- The benefits of a signed DNS zone and validating resolvers are indirect
  - Cost and benefit are totally mis-aligned in this space!

DNSSEC is a Market Failure!

- If it wasn't for the lead taken by Google and their adoption of DNSSEC validation in their Public DNS Service then its likely that DNSSEC would still be a complete failure
- And only was possible because of the significant level of market share in the use of Google's Public DNS Service
- So perhaps we should look at the market for resolvers

II. The Open Market for Resolution

Who resolves Names?

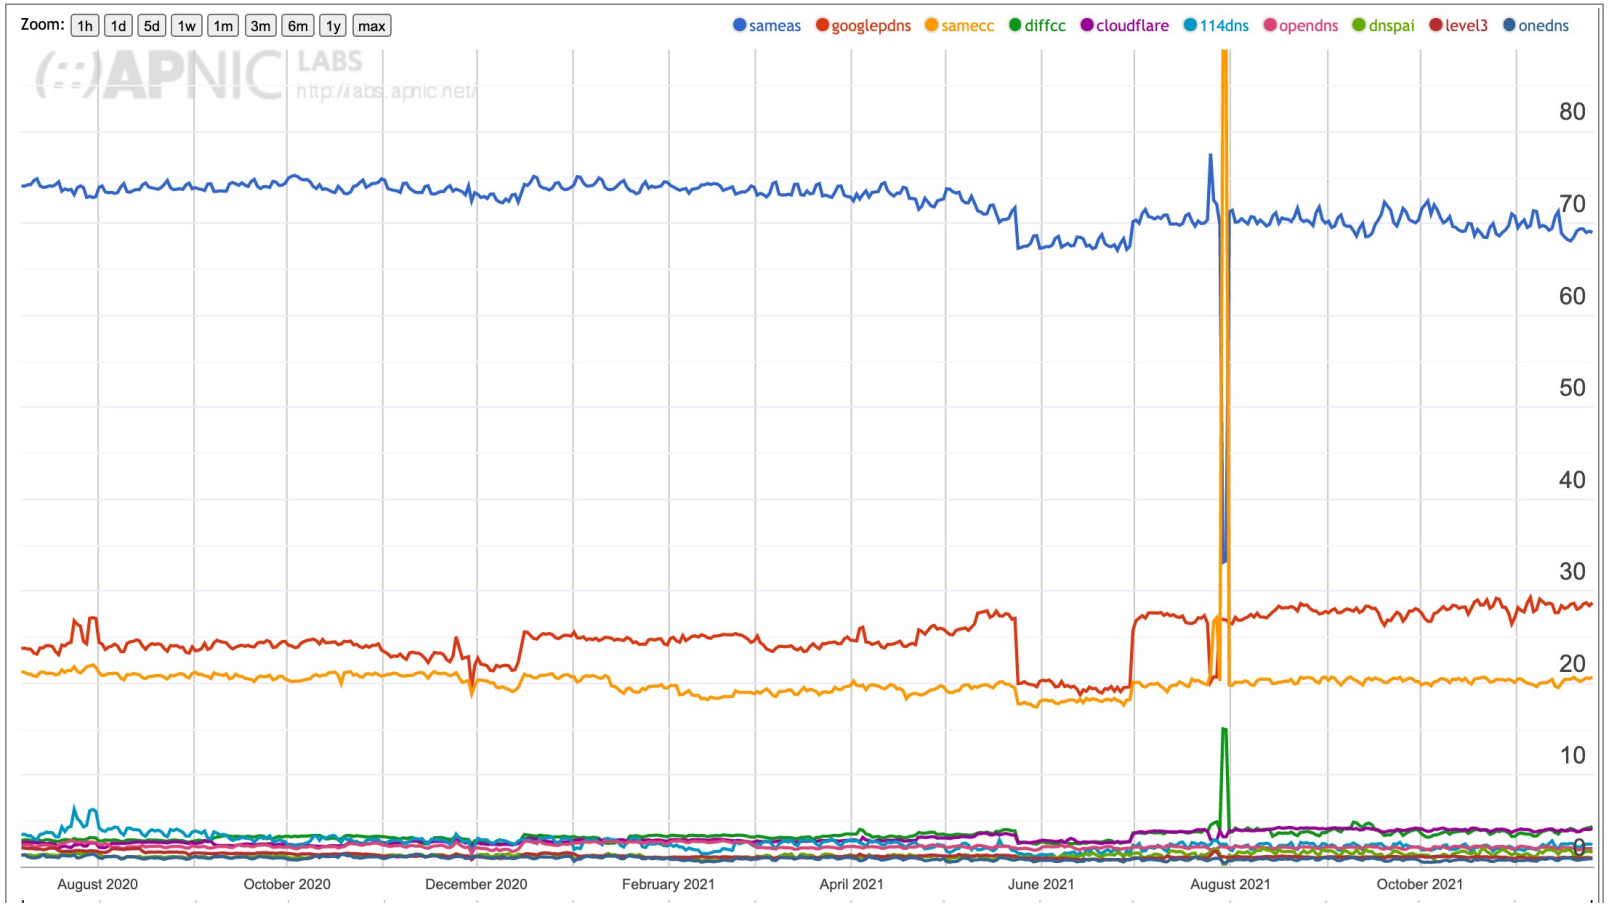

My ISP

Google

Same country, different network (likely to still be my ISP)

Everyone else

Who resolves Names?

If this is a market then it appears to be highly skewed:

- Google has a market share of just under 30% of all users
- The next largest open resolvers are Cloudflare (4%), OpenDNS (2%) and Quad9 (0.5%)
- Everything else is just the default of the local ISP performing the name resolution service for their customers

Why has this happened?

How has Google succeeded and other open DNS resolver services appear to struggle to get significant adoption by clients?

The DNS Name Resolution Economy

- In the public Internet, end clients don't normally pay directly for DNS name resolution services
- Which implies that outside of the domain of the local ISP, DNS resolvers are essentially unfunded by the resolver's clients
- And efforts to monetise the DNS with various forms of funded misdirection (such as NXDOMAIN substitution) are generally viewed with extreme disfavour
- Open Resolver efforts run the risk of success-disaster
  - The more they are used, the greater the funding problem to run them at scale
  - The greater the funding problem the greater the temptation to monetise the DNS resolver function in more subtle ways

Google is Special

- Search is a critical asset for GOOGLE
  - Search drives eyeballs and profiles
  - Eyeballs and profiles drive advertising
  - Advertising drives revenue
- If we could use the DNS as a defacto search engine then Google has a problem
  - But that's what is happening when DNS resolvers don't pass back NXDOMAIN but instead pass back a reference to a search engine, or even better use a search engine to pass back the resolution response that a hidden search engine obtained when it processed the queried DNS name
- So it's in Google's interests to have a fast, accurate and unaltered DNS resolution service
  - And Google's real target market for this service is not individual users, but ISPs
  - Because the real threat model for Google is NXDOMAIN substitution and other forms of DNS response manipulation by ISPs as a local revenue source
- No other Open DNS resolver service shares Google's motivation

The DNS Name Resolution Economy

- The default option is that the ISP funds and operate the recursive DNS service, funded by the ISP's client base
  - >70% of all end clients use their ISPs' DNS resolvers
- However the fact that it works today does not mean that you can double the input costs and expect it to just keep on working tomorrow
- For ISPs the DNS is a cost department, not a revenue source
  - We should expect strong resistance from ISPs to increase their costs in DNS service provision

The resistance to change in the DNS

- The quality of an ISP's DNS service does not appear to be a significant competitive discriminatory factor in the consumer market
- So the ISP does not generally devote many resources to tuning their DNS infrastructure for high performance, resiliency and innovation
- Most users don't change their platform settings from the defaults and CPE based service provisioning in the wired networks and direct provisioning in mobile networks will persist
- So current innovations such as improved DNS privacy (DNS over TLS, DNS over HTTPS) are looking like being another mainstream market failure in the DNS space

The resistance to change in the DNS

- The quality of an ISP's DNS service does not appear to be a significant competitive discriminatory factor in the consumer market
- So the ISP does not generally devote many resources to tuning their DNS infrastructure for high performance, resiliency and innovation
- Most users don't change their platform settings from the defaults and CPE based service provisioning in the wired networks and direct provisioning in mobile networks will persist
- So current innovations such as improved DNS privacy (DNS over TLS, DNS over HTTPS) are looking like being another mainstream market failure in the DNS space
- But maybe that's not the full story …

Fragmenting the DNS

- Is appears more likely that **applications** who want to tailor their DNS use to adopt a more private profile will hive off to use DNS over HTTPS to an application-selected DNS service, while the platform itself will continue to use libraries that will default to DNS over UDP to the ISP-provided recursive DNS resolver
- That way the application ecosystem can fund its own DNS privacy infrastructure and avoid waiting for everyone else to make the necessary infrastructure and service investments before they can adopt DNS privacy themselves
- The prospect of **application-specific naming services** is a very real prospect in this scenario

Fragmenting the DNS

- Is appears more likely that **applications** w… or their DNS use to adopt a more private p… NS over HTTPS to an appli… platform itself… over UDP to the I…
- That … privacy infras… else to make the neces… vice investments before they can adopt … themselves
- The prospect of **application-specific naming services** is a very real prospect in this scenario

Those parts of the Internet space with sufficient motivation and resources will simply stop waiting for everyone else to move. They will just do what they feel they need to do!

It's life Jim, but not as we know it!*

- The overall progression here is an evolution from network-centric services to platform-centric services to today's world of application-centric services
- It's clear that the DNS is being swept up in this shift, and the DNS is changing in almost every respect
- The future prospects of a single unified coherent name space as embodied in the DNS, as we currently know it, for the entire internet service domain are looking pretty poor right now!

* With apologies to the Trekkies!

Is the DNS "Open"?

- Yes, today.

Although the cracks are clearly evident with the use of various DNS customizations to tailor each DNS response to the querier, the interest in lifting the DNS into the application space as an application service and not a common infrastructure, the strong levels of centralisation and the destruction of open competition in this space, the continuing erosion of trust and the strong commercial impetus to use the DNS as a profiling tool in the surveillance economy.

Will it continue to be "open"?

- Probably not!

It’s just getting too hard to keep it together and keep it open and the market-based pressures continue to tear the DNS apart and haul it away from a unified open space into a set of close and opaque markets with a fragmented name space.

And the common benefit of a single, cohesive open space looks like falling victim to the Tragedy of the Commons