

Is Infrastructure Security an Internet Market Failure?

Geoff Huston AM

APNIC Labs

#apnic58

TE WHANGANUI A TARA
WELLINGTON,
AOTEAROA NEW ZEALAND

30 August - 6 September 2024

Why is DNSSEC so hard?

- Has it taken almost 30 years for DNSSEC to have signed about 10% of domain names?
- Has it taken almost 30 years for just one third of users to perform DNSSEC validation?
- And how much longer will it take for the rest of the unprotected name space to adopt DNSSEC?
 - <insert your guess here!>



Why is Routing Security so hard?

- Is deploying a secure framework for inter-domain routing security so hard?
- Has the effort to define and adopt a framework for a secured routing system taking us decades?
 - And we are still a fair distance from completing this work!



What's going on?

- In the deregulated world of the Internet are these extended failures to adopt secure technologies instances of ***market failure?***



Is securing the infrastructure of
the Internet a market failure?

Yes!



Thanks!



#apnic58

Whoa!

Not so fast!

- What are we talking about with “**infrastructure security**”?
- What is meant by “**market failure**?”
- Is infrastructure security failing us?
- What are the implications?
- Can we improve this picture?



Infrastructure Security

- For the Internet, the common *infrastructure* we are talking about are the realms of IP addresses, Routing, Names and named Service Connections
 - Or “**names, addresses and routing**”
- And *security* is the role of protecting the integrity of these realms
 - Ensuring uniqueness of name and address tokens
 - Ensuring the integrity of the association between a holder of these resources and the resources themselves
 - Allowing a user to authenticate the validity of the use of an address or name



Market Failure and Public Good

- **Market failure** is the economic situation defined by an inefficient distribution of goods and services in the free market. In market failure, the individual incentives for rational behaviour do not lead to rational outcomes for the group (Investopedia)
- In economics, a **public good** is a good that is both non-excludable and non-rivalrous. Use by one person neither prevents access by other people, nor does it reduce availability to others. Therefore, the good can be used simultaneously by more than one person. (Wikipedia)



What's the issue here?

- The Internet operates in a largely deregulated environment which is dominated by market pressures
- There is no forcing function to ensure that *everyone* adopts a security mechanism
 - Whether you sign your domain name with DNSSEC is your choice
 - Whether you generate ROAs for your route objects is your choice
- The Internet relies on common perceptions of opportunity and cost to motivate adoption of technologies
 - These are **market forces**, not **regulatory imposts**



Heterogenous Markets

In the absence of enforced universal adoption, we normally see discretionary *partial adoption*

- Everyone reacts to prevailing market pressures in their own way
- In terms of adoption of security mechanisms, the common result is that only a proportion of infrastructure elements will have security credentials associated with them
- If the common perception of benefit to cost is high then adoption rate will be high, and when the perception of benefit to cost is generally low then adoption will stall



Adoption is variable

Where there is no clear immediate and direct benefit to adoption then a consumer is likely to defer the cost of adoption

- This is the case irrespective of the broader social benefit that may accrue from general adoption
- “It’s in our **common interest** to adopt <x> to improve the security of our digital infrastructure but **individually** the incremental cost of adoption exceeds the perception of direct benefit”



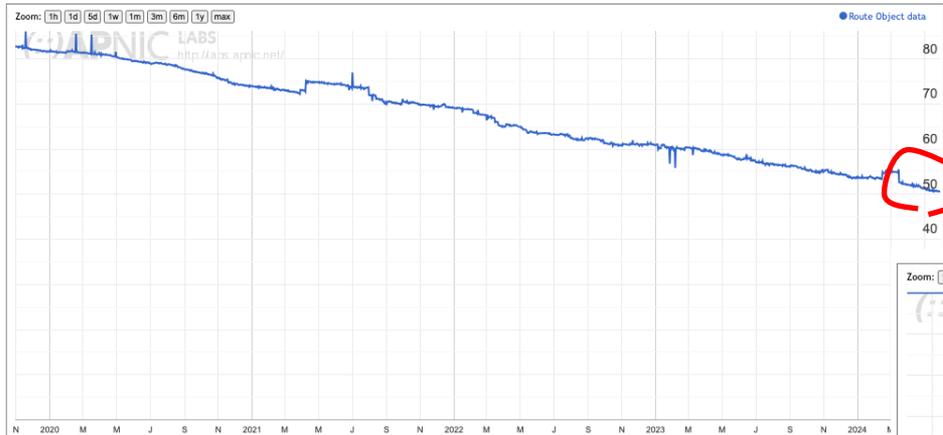
Hypothetically ...

- Could a ISP charge its customers a premium if it had deployed a RPKI Route Origination Validation framework?
- Could a DNS hosting service charge its customers more if it added DNSSEC front-end signing to its service?
- This inability to monetize such investments in securing network infrastructure acts as a negative incentive for adoption

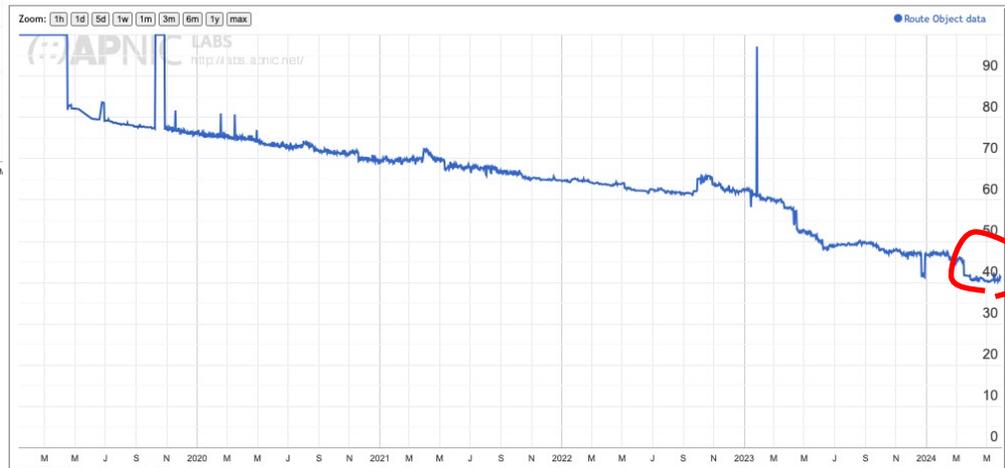


Routes without ROAs

Proportion of IPv4 Route Objects NOT covered by a ROA



Proportion of IPv6 Route Objects NOT covered by a ROA

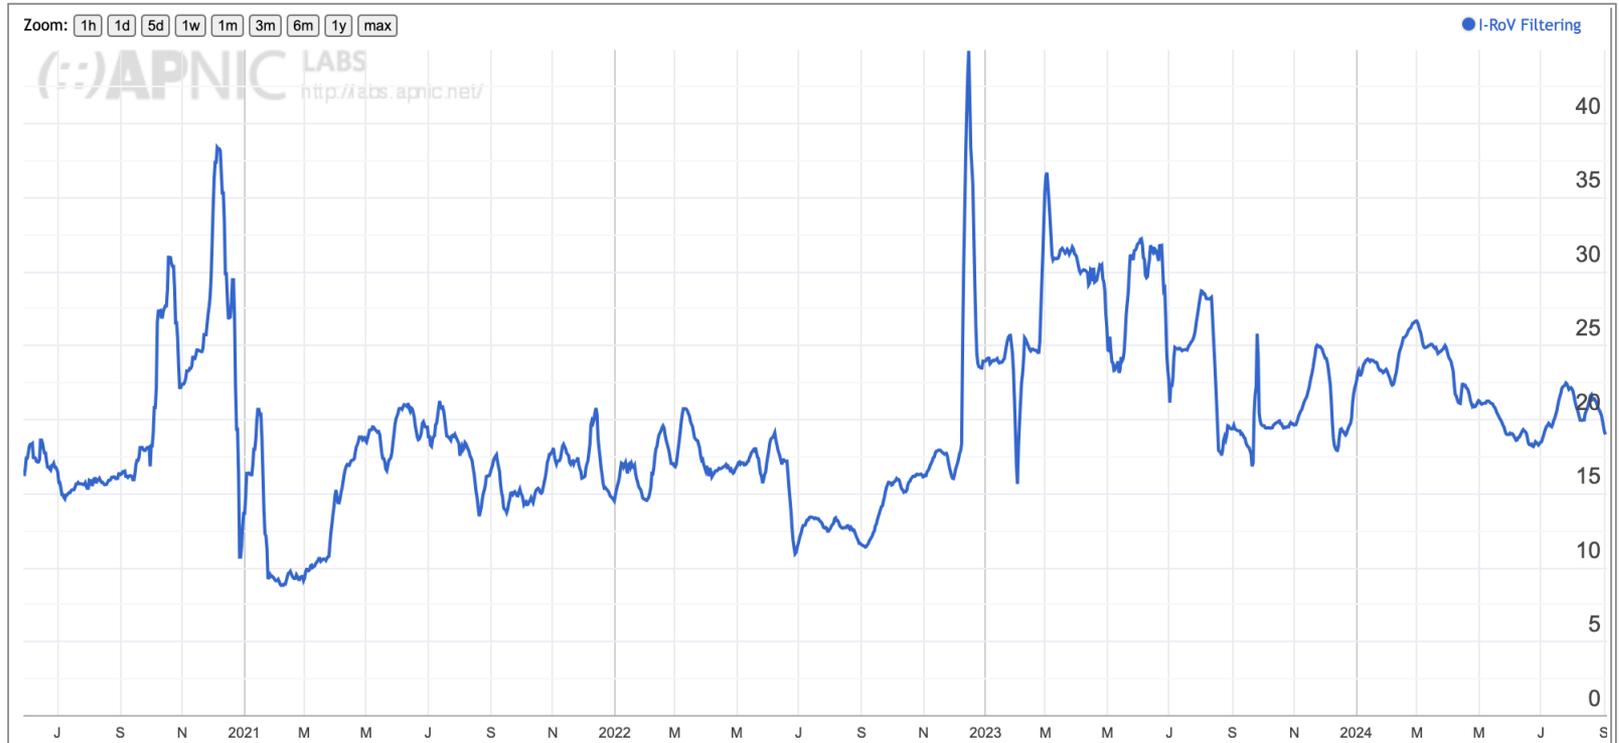


Securing Routes

- 54% of all advertised route objects in the IPv4 space are not listed in any ROA
- 48% of all advertised route objects in the IPv6 space are not listed in any ROA
- ~0% of all BGP AS Paths are protected



Route Origination Validation



Route Origination Validation

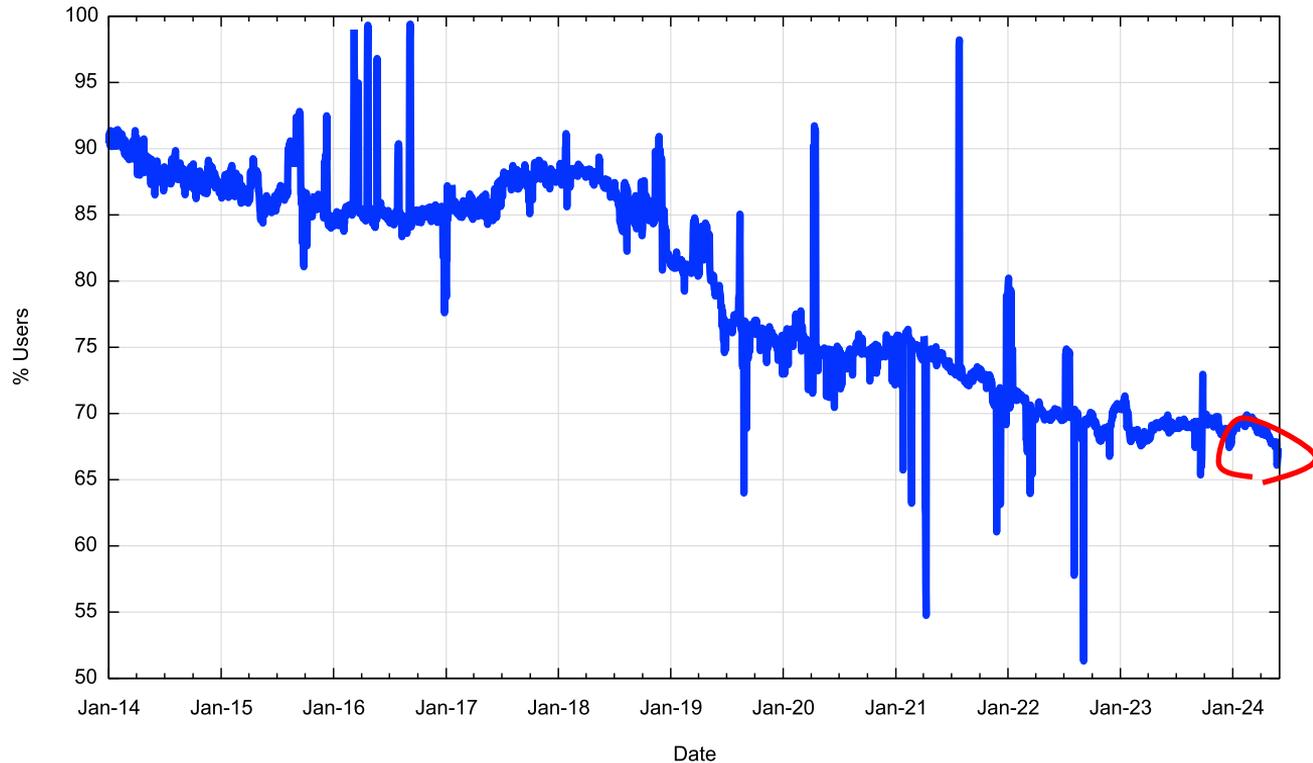
- All those ROAs appear be to achieving

Very little!



No DNSSEC Validation

Proportion of users who will accept a badly signed DNS result



Securing Names with DNSSEC

- Some 70% of users are located behind recursive resolvers where there is no validation of DNSSEC credentials
- Some 90% of domain names are not signed with DNSSEC
 - (Tranco Top 1M names)



Securing Addressing

- 0% of address database entries are digitally signed
 - Why do I constantly need to refer to an unsigned address registry to “prove” that this is my IP address?
 - Why can’t I associate a public/private key pair with my addresses and use digitally signed attestations to assert my holding of addresses?
 - If this “works” for ROAs in a routing context, then why can’t I use the same address credentials in other contexts?



What should an application do?

In today's Internet **no** application can count on the presence of a fully secured infrastructure

- DNS queries generate unreliable answers
- IP packets may be routed to unintended destinations



What should an application do?

- The safest course of action in the face of partial adoption is to assume that these infrastructure services - the secured resolution of domain names to IP addresses, and the secured forwarding of packets to destination addresses is **not reliably available**
- The application is unable to derive a consistent benefit from partial deployment of secured infrastructure



What should an application do?

Therefore, a well-engineered application has to invest in its own efforts to validate the authenticity of the service destination to which it is connecting and invest in protecting the integrity of the service transaction that is it undertaking

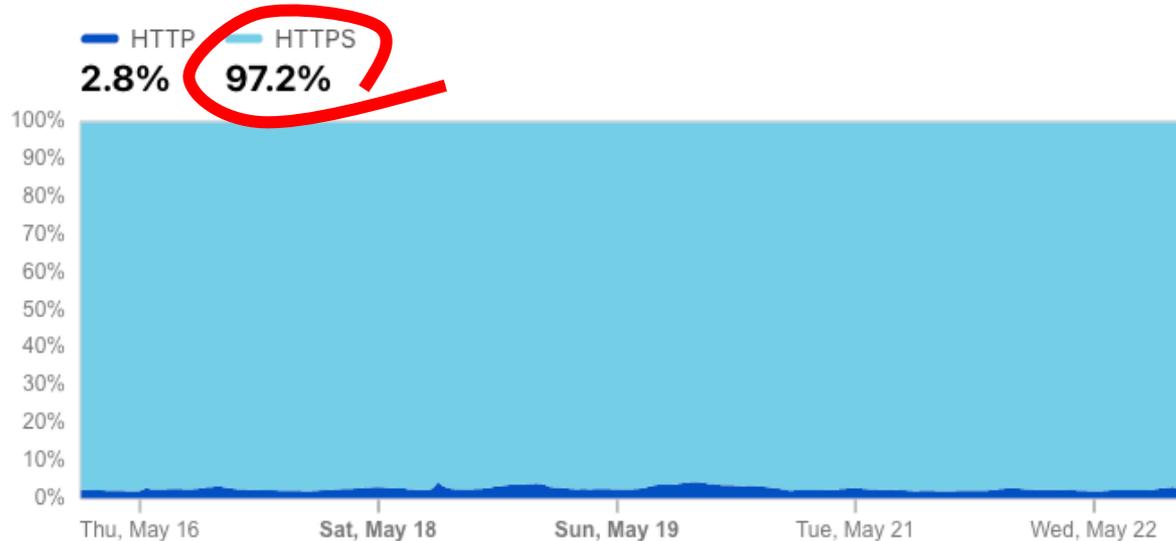
Which is a widespread feature of today's Internet!



TLS as an Application-level security Response

HTTP vs. HTTPS

Distribution of HTTP vs. HTTPS requests [?](#) [🔗](#)



Which means ...

- When an application is using TLS to assure itself of the authenticity of the service to which it is connecting, then what is the incremental value it places on the partial presence of infrastructure-based address and name security?



Which means ...

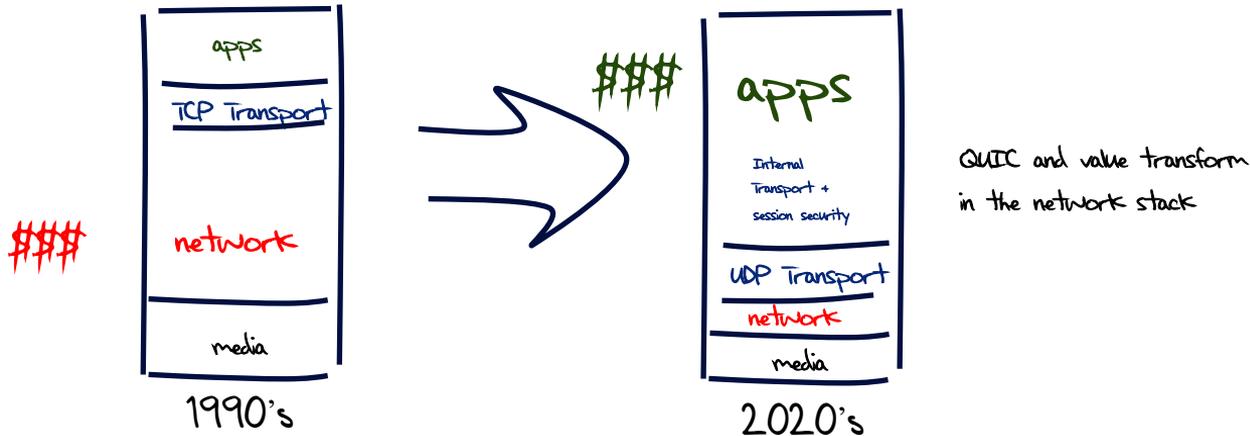
- When an application is using TLS to assure itself of the authenticity of the service to which it is connecting, then what is the incremental value it places on the partial presence of infrastructure-based address and name security?

Zero!



Internet Economics

- The economic picture of the last 2 decades of the Internet is the shift of value UP the protocol stack



Internet Economics

- The economic picture of the last 2 decades of the Internet is the shift of value UP the protocol stack
- The “value” of the Internet now lies in content and services, and the application frameworks that deliver them
- Applications are now taking over many of the roles that were performed by the lower layers of the protocol stack
 - QUIC is a good example of moving transport, flow control, and session integrity up from the common platform into the application itself
- The lower layers of the stack, including common infrastructure services, are being **commoditized** and **stripped of intrinsic value**



What does this mean?

- 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, and performing service authentication directly within the application
 - QUIC (and MASQUE) is a visible part of this process of wrapping up traffic in encryption and redirection



What does this mean?

- For the network operator there is little left to do at the common infrastructure level, and little money to do it with in any case!
- For the infrastructure operator the same story of marginalisation applies
- Who is left to invest in infrastructure security?
- And who funds them?



What should we do?

Should we invest more money and effort into DNSSEC?

- Who funds such an effort and where would the funding be directed?
 - DNSSEC-signing a zone entails higher operational effort by the zone operator - i.e. an increase in cost
 - Validating a DNSSEC-signed response also invokes higher levels of operational effort on the part of the DNS resolver operator
- Or is this a case of misdirected effort with an assumption of funding that simply does not exist in the first place?



What should we do?

Invest more money and effort into the routing space and push for universal deployment of BGPSEC?

- Who funds such an effort and where would the funding be directed?
- Should we even care about routing security at all when most CDNs are now adjacent to most of the Internet user access networks?
 - Most content and service delivery in today's Internet is NOT routed!
- What are we trying to protect here in routing security and who has an economic interest in the outcomes of such measures?
 - Or is TLS doing an adequate job for content and service operators?



Perhaps that's too extreme

Let's try and rephrase the problem statement...



Improving the situation

How can applications improve the assurance of the authenticity of the content and services that they access

- Third-Party CAs are the weak point of the entire TLS architecture – can we use a DNSSEC-like rooted trust hierarchy to remove the inherent vulnerability of so many trust points?
- Should we make more use of obfuscation approaches such as MASQUE to make more use of “semi-trusted” intermediaries that further obscure content and service transactions from the network?



What is happening?

- Moore's Law has a lot to answer for!
- The original communications industry was an industry that was about rationing a scarce common public resource via pricing signals
- The Internet replaced human endpoints with digital services
- The constant opportunities of ever faster and cheaper digital processing meant that it was cheaper to bring a copy of the service closer to each user than it was to use the network to bring each user to the service
- The role of the common network is being devalued as a common infrastructure resource and the value of content aggregation and distribution platforms have risen in its place



Why is this relevant?

- Because the failure of common infrastructure security mechanisms to achieve universal adoption is reflective of a broader technology shift, and not a failure on the part of individual actors
- Which means that the conventional response to market failures – regulatory imposts – become regressive imposts rather than helpful interventions when applied to common infrastructure



Such as...



Media Contact:
MediaRelations@fcc.gov

For Immediate Release

FCC CHAIRWOMAN PROPOSES INTERNET ROUTING SECURITY REPORTING REQUIREMENTS

Broadband Providers Would Need BGP Security Plans and Largest Providers Would File Quarterly Reports

WASHINGTON, May 15, 2024—FCC Chairwoman Jessica Rosenworcel today proposed requiring the largest broadband providers to file confidential reports on Border Gateway Protocol (BGP) security so the FCC and its national security partners can for the first time collect more up-to-date information about this critical internet routing intersection. BGP is the technology used for routing information through the physical and digital infrastructure of the internet.

National security experts have raised concerns that, by accessing vulnerabilities in BGP, bad actors can disrupt critical services that rely on the internet and result in misdirection, interception, inspection, or manipulation of data. A bad network actor may deliberately falsify BGP reachability information to redirect traffic. Russian network operators have been suspected of exploiting BGP's vulnerability for hijacking in the past. "BGP hijacks" can expose Americans' personal information, enable theft, extortion, state-level espionage, and disrupt otherwise-secure transactions.

"It is vital that communication over the internet remains secure," **said Chairwoman Rosenworcel.** "Although there have been efforts to help mitigate BGP's security risks since its original design, more work needs to be done. With this proposal, we would require broadband providers to report to the FCC on their efforts to implement industry standards and best practices that address BGP security."



#apnic58

What should we do?

Press on for the next 30 years, doing what we've been doing for the past 30 years?

- What do we need to do to motivate every domain name to be DNSSEC-signed?
- What do we need to do to make every network operator generate a comprehensive set of ROA credentials and run BGPSEC?
- What do we need to do to get the registries to digitally sign every resource element in their registry?



Or should we be asking:

- How can we improve service transactions and content access within the application space to assure users of the authenticity and accuracy of their transactions?



Thanks!



#apnic58



APNIC

58

[#apnic58](#)

**TE WHANGANUI A TARA
WELLINGTON,
AOTEAROA NEW ZEALAND**

30 August - 6 September 2024