I. Why is Securing the Internet's Routing System so hard?

Geoff Huston, APNIC

Usual BGP Disaster-Porn Clips…

Pakistan hijacks YouTube

Research // Feb 24, 2008 // Dyn Guest Blogs

Late in the (UTC) day on 24 February 2008, Pakistan Telecom (AS 17557) began advertising a small part of YouTube's (AS 36561) assigned network. This story is almost as old as BGP. Old hands will recognize this as, fundamentally, the same problem as the infamous AS 7007 from 1997, a more recent ConEd mistake of early 2006 and even TTNet's Christmas Eve gift 2004.

ars TECHNICA

BIZ & IT TECH SCIENCE POLICY CARS GAMING & CULTURE STORE

THE ACCIDENTAL LEAK —

Google goes down after major BGP mishap routes traffic through China

Google says it doesn't believe leak was malicious despite suspicious appearances.

DAN GOODIN - 11/13/2018, 6:25 PM

Google lost control of several million of its IP addresses for more than an hour on Monday in an event that intermittently made its search and other services unavailable to many users and also caused problems for Spotify and other Google cloud customers. While Google said it had no reason to believe the mishap was a malicious hijacking attempt, the leak appeared suspicious to many, in part because it misdirected traffic to China Telecom, the Chinese government-owned provider that was recently caught improperly routing traffic belonging to a raft of Western carriers though mainland China.

FURTHER READING
Strange snafu misroutes domestic US Internet traffic through China Telecom

The leak started at 21:13 UTC when MainOne Cable Company, a small ISP in Lagos, Nigeria, suddenly updated tables in the Internet's global routing system to improperly declare that its autonomous system 37282 was the proper path to reach 212 IP prefixes belonging to Google. Within minutes, China Telecom improperly accepted the route and announced it worldwide. The move by China Telecom, aka AS4809, in turn caused Russia-based Transtelecom, aka AS20485, and other large service providers to also follow the route.

According to BGPmon on Twitter, the redirections came in five distinct waves over a 74-minute period. The redirected IP ranges transmitted some of Google's most sensitive communications, including the company's corporate WAN infrastructure and the Google VPN. This graphic from regional Internet registry RIPE NCC shows how the domino effect played out over a two-hour span. The image below shows an abbreviated version of those events.

BGPMON Now part of OpenDNS

HOME BLOG ABOUT US PRODUCTS AND SERVICES CLIENT PORTAL

Today's BGP leak in Brazil

Posted by Andree Toonk – October 21, 2017 – News and Updates – No Comments

Earlier today several people noticed network reachability problems for networks such as Twitter, Google and others. The root cause turned out to be another BGP mishap.

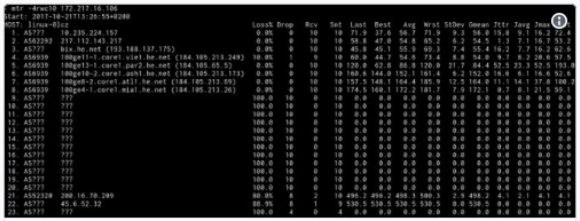

Between 11:09 and 11:27 UTC traffic for many large CDN was rerouted through Brazil. Below an example for the Internet's most famous prefix 8.8.8.0/24 (Google DNS)
At 2017-10-21 11:09:59 UTC, AS33362, an US based ISP saw the path towards Google's 8.8.8.0/24 like this:

33362 6939 16735 263361 15169

This shows the US based network AS33362, would have sent traffic to Google via 6939 (HE) to 16735 (Algar Telecom, Brazil), to 263361 infovale telecom which would have tried to delivered it to Google. The successful delivery of packets would have been unlikely, typically due to congestion which would have been the result of the increase in attracted traffic or an ACL blocking the unexpected traffic.

BGPMON Now part of OpenDNS

HOME BLOG ABOUT US PRODUCTS AND SERVICES CLIENT PORTAL

Popular Destinations rerouted to Russia

Posted by Andree Toonk – December 12, 2017 – Hijack – No Comments

Early this morning (UTC) our systems detected a suspicious event where many prefixes for high profile destinations were being announced by an unused Russian Autonomous System.

Starting at 04:43 (UTC) 80 prefixes normally announced by organizations such Google, Apple, Facebook, Microsoft, Twitch, NTT Communications and Riot Games were now detected in the global BGP routing tables with an Origin AS of 39523 (DV-LINK-AS), out of Russia.

Looking at timeline we can see two event windows of about three minutes each. The first one started at 04:43 UTC and ended at around 04:46 UTC. The second event started 07:07 UTC and finished at 07:10 UTC.

Even though these events were relatively short lived, they were significant because it was picked up by a large number of peers and because of several new more specific prefixes that are not normally seen on the Internet. So let's dig a little deeper.

One of the interesting things about this incident is the prefixes that were affected are all network prefixes for well known and high traffic internet organizations. The other odd thing is that the Origin AS 39523 (DV-LINK-AS) hasn't been seen announcing any prefixes for many years (with one exception below), so why does it all of sudden appear and announce prefixes for networks such as Google?

If we look at a few AS paths we see that 39523 is always the origin, while the next hop transit AS is always 31133 PJSC MegaFon. We also see that the announcements were picked up further and made reachable by a few large ISP's such as:

xx 6939 31133 39523 (path via Hurricane Electric)
xx 6461 31133 39523 (path via Zayo)
xx 2603 31133 39523 (path via Nordunet)
xx 4637 31133 39523 (path via Telstra)

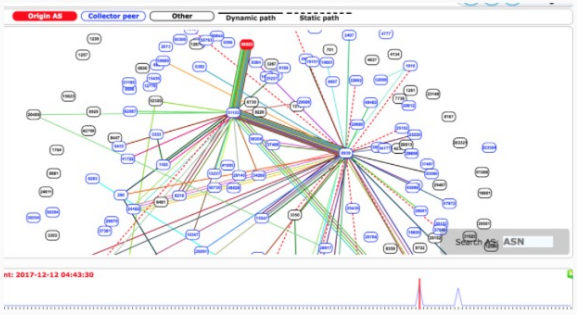

What makes this incident suspicious is the prefixes that were affected are all high profile destinations, as well as several more specific prefixes that aren't normally seen on the Internet. This means that this isn't a simple leak, but someone is intentionally inserting these more specific prefixes, possibly with the intent the attract traffic.

CommsWire
Essential daily reading for the communications industry executive
An iTWire publication www.itwire.com Editor: Stan Beer Friday 16 November 2018

TELSTRA ROUTING ERROR TAKES DOWN INTERNET

Degrees of Difficulty…

Why are some issues so challenging to solve, while others seem to be effortless?

For example: Why was the IPv4 Internet an unintended runaway success in the 90's, yet IPv6 has been a decades-long protracted exercise in industry-wide prevarication?

Routing Accidents Abound

- The majority of routing incidents are generally understood to be accidental in nature (rather than malicious)
- Because if this, we generally don't react as aggressively when there is a routing incident, and so there aren't nearly as many reputational consequences for networks which originate such routing anomalies.
- As an example, if the news had been ”<isp> maliciously steals traffic destined for X" instead of the meme of "Meh, routing is hard!" there may be more in terms of real world implications for <isp>…
  - On the other hand, there is increasing interest by governments, particularly in the area of network support of vital services, and when these fail, even through operational mishap, there can be very real consequences, both in fatalities and to the company
    - See “Optus and the 000 failure in September 2025”

Internet “Successes”

- IPv4 (and datagram packet switching)
- Network Address Translators (perversely!)
- TCP evolution and adaptation
- DNS
- Content Distribution Systems
- Streaming

Success Factors

- Piecemeal deployment without the requirement for central orchestration
- Tangible advantages to early adopters
- Economies of scale as adoption numbers increase
- Alignment of common benefit with individual benefit

Internet Non-Successes

Failures!

- SPAM
- DDOS defence
- BCP 38 deployment
- Secure end systems
- Secure networks
- Secure Routing

Failure Factors

- Need for orchestrated actions (flag days)
- Technologies that require universal or near universal adoption
- Where there are common benefits but not necessarily individual benefits
- No clear early adopter advantage

What makes a problem “hard?”

- It might be **technically challenging**: While we understand what we might want that does not mean we know how to construct a solution
- It might be **economically mis-aligned**: The costs of a solution are not directly borne by the potential beneficiaries of deploying the solution
- It might be **motivated by risk mitigation**: We are notorious for undervaluing future risk!

Why is securing the Inter-domain Routing System so hard?

- No single entity is in charge
- We can't "audit" a BGP speaker and its route set for "compliance", as we have no standard reference route set to compare it with
- We can't arbitrate between conflicting BGP information (because there is no standard reference point)
- There are no credentials that allow a BGP object to be compared against the original route injection (because BGP is a hop-by-hop protocol)
- Because BGP route propagation is based on sequences of opaque local decisions

Why should we worry?

Because its too easy to be “bad” in routing

What's the Risk?

- User traffic gets diverted enabling a Man-in-the-Middle attack on a service

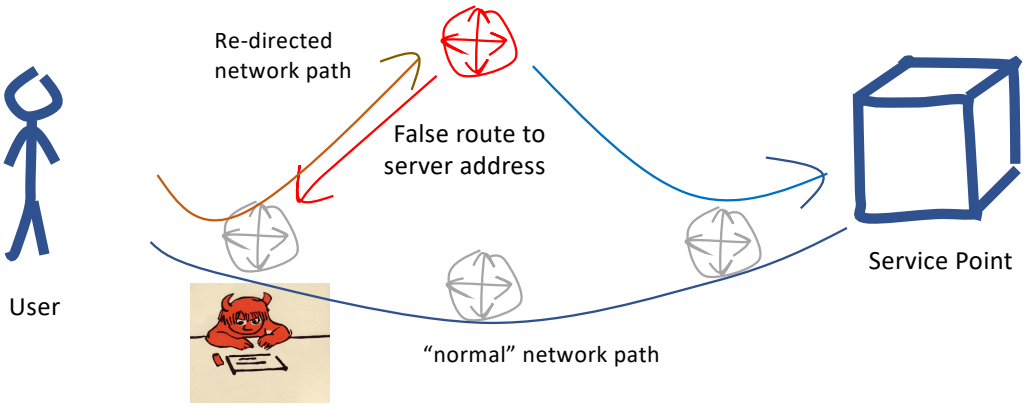

What's the Risk?

- Mount a Denial Attack:
  - Divert the traffic to a sinkhole, thereby denying access to the service
    - Crude, but effective!

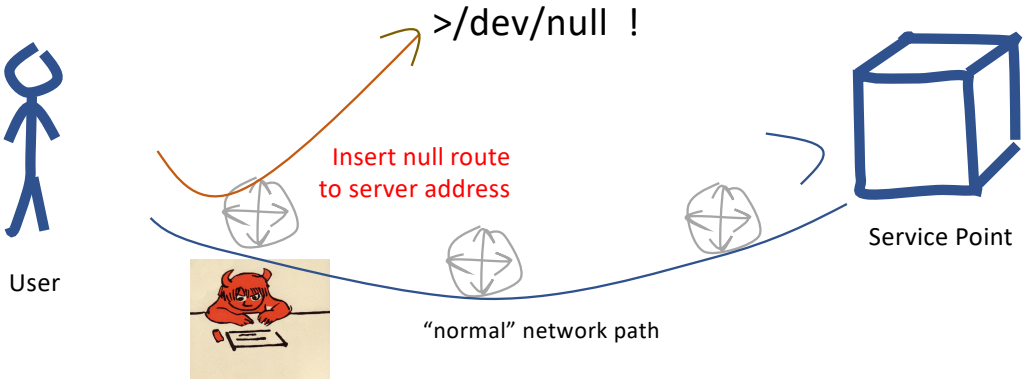

What's the Risk?

DNS Attacks via Routing

- Divert DNS traffic to fake DNS servers and provide fake answers
  - Very few domains are DNSSEC-signed and nowhere near enough resolvers perform DNSSEC validation
  - The faked answer can pass unchallenged

What's the Risk?

Server Attacks Attacks via Routing

- Divert TCP traffic to fake servers and provide fake answers
  - The attacker can collect user credentials while shadowing the actual site

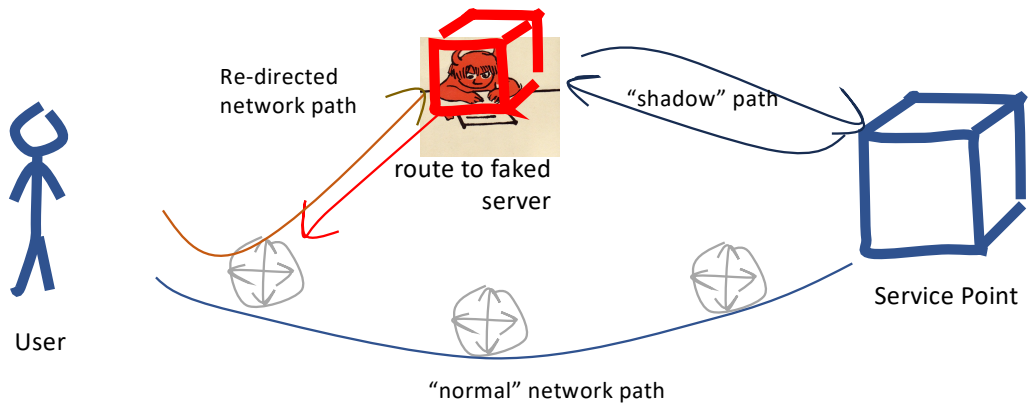

21

An attack vector on HTTPS…

- Let's say you can find an online trusted CA
  - that uses the DNS as proof-of-possession of a DNS name in order to mint a domain name certificate
  - And the DNS name is not DNSSEC protected
- You can mint a fake domain name certificate by:
  - Mount a routing attack on the DNS infrastructure with a fake DNS responder
  - Answer everything correctly except for *.victim ACME DNS challenge from the CA
  - And for the *.victim challenge queries respond with your own answer
  - Which means you can answer the CA's DNS challenge
- Now you have a trusted fake domain name certificate

- You are now able to pull off a MITM attack on a TLS 'protected' service

A similar attack can be mounted in a web challenge environment by hijacking the IP address of the web service

“I want a Pony” Routing Security wish list

1. Identify whether an address is “bogus” or not
2. Assure that the address holder has given their permission for an address to be announced into the routing system
3. Identify which AS(s) have been given this permission
4. Identify if the AS Path is consistent with the ‘correct’ operation of BGP
5. Identify if the AS Path is consistent with the routing policies of the each of the Ases
6. Identify when routing information is being ‘incorrectly’ withheld

II. Background - How we got to here

Route Registeries

- First used in the early 1990's as the Route Arbiter Database (RADB) as part of the NSFNET program
- Describes route origination and inter-AS routing policies
- An explicit declaration of intent in routing
- Route Registries can be used by others to filter BGP announcements, filtering out route advertisements that are not described in the route registry
  - Primary value in preventing neighbor route leaks
  - Can be used to prevent hijacks

RPSL and IRR Tools

Route Policy Specification Language

- A formal language to allow network operators to describer routing policies and a set of tools to help operators generate routing filters

Tools

- IRRToolSet, RtConfig, rpsltool
- Gather routing policy data from route registries and transform the data into a route filter set
- The route filter set is intended to filter route objects that do not match the declared route policy

Route Registry Issues

- Poor Authority Model (or the complete lack of one in many cases!)
  - How can a user know that a RR entry is genuine and current?
  - How can a user know that a RR entry is maintained by an entity who is the authoritative “owner” of an IP address or ASN?
  - How can a user tell the difference between a current RR entry and a lapsed historical RR entry?
- Too many Route Registries
  - If two different RRs contain conflicting information, what are users meant to do?
- Incomplete Data
  - If a route is not described in a Route Registry is it just the registry that is missing data or is the route itself invalid?
- Scaling issues
  - No realistic way to apply IRR filters to upstreams
- RPSL got too geeky!
  - The Route Policy Language used by Route Registries got overly expressive and complex to use

What's missing in that picture?

- If we want to improve the usefulness of route registries, we probably need a robust authority model

How about Digital Signatures?

- The signatures can provide currency and authenticity
- The authority model can allow RR entries to be seen as explicit authorities or permissions from address holders to network operators and from network operators to other networks

From trying to Secure Routing to trying to Secure the Routing Protocol - sBGP

sBGP – an effort to secure the "correct" operation of the BGP protocol

- Secure UPDATE messages to protect them from inflight tampering, assure the receiver of the authenticity of the sender, assure the receiver that the UPDATE was correctly addressed, and the sender was authorized to send the UPDATE
- Requires X.509 certificates and a PKI for address and ASN ownership
- Route origination required a signed authority by the address owner to permit the AS operator to originate the prefix
- Route propagation required the route advertiser to sign over the route, it's AS Path and the AS of the peer that is receiving the route update
- Cannot validate route withdrawals
- Does not incorporate routing policy constraints

And an Alternative - soBGP

- sBGP required intensive processing capability on routers, as route propagation requires a new signature for every propagation step for every prefix
- **soBGP** used a “lighter” approach that combined route origination attestations with AS Adjacency (topology) attestations
  - sBGP is able to *verify* the AS Path of an update, while soBGP can conclude that the AS Path represents a *plausible* path through the inter-AS topology
- Neither can validate route withdrawals
- Neither incorporates routing policy constraints

Deciding is hard

- sBGP and soBGP represented different trade-offs between processing load and levels of assurance
- A new Working Group (RPSEC) was formed to list the set of functional requirements that a secure routing framework should address
  - The working group reached agreement on the require to add signed credentials to validate the function of route origination
  - It stalled on the topic of path validation, with no consensus between path validation and path plausibility
  - The group did not get around to considering adherence to policy constraints

Authority Injection

- Both sBGP and soBGP had the issue of how to create a testable association of an address and AS number resource with the holder/operator of the resource
- An effort was constructed on the foundation of RFC3779, with the active engagement of the Regional Internet Registries as Certification Authorities to build and maintain a PKI for IP number resources

RFC3779: X.509 Public Key Certificates for IP addresses and AS Numbers

- An X.509 Public key certificate that includes a set of IP addresses and AS numbers
- If a certificate can be validated against a trust anchor, then it indicates that:
  - The IP addresses and/or AS numbers have been validly allocated
  - The holder of the subject key pair is the current holder of the IP addresses and/or AS numbers
  - Attestations validly signed using the private key can be considered as genuine authorities that cannot be repudiated

Route Origination Authority

- An address holder can convey a ‘permission’ for an AS to originate a BGP route for the address by signing a permission authority (ROA) using a signing key associated with a valid public key address certificate
- This authority:
  - can be validated by any interested party
  - is dated, so currency is known
  - cannot be repudiated

If we all used ROAs then:

1. Identify whether an address is “bogus” or not
2. Assure that the address holder has given their permission for an address to be announced into the routing system
3. Identify which AS(s) have been given this permission
4. Identify if the AS Path is consistent with the ‘correct’ operation of BGP
5. Identify if the AS Path is consistent with the routing policies of the each of the Ases
6. Identify when routing information is being ‘incorrectly’ withheld

Is 3 out of 6 good enough to get a pony?

NO!

- The hijack can reproduce the origin and if the ROA is sloppy then it can use a more specific
- Even if the ROA is tight the conflicting routes can still support a desired attack profile

From ROAs to a fully secure BGP - BGPSEC

ROAs are good, but probably not enough to stop a determined routing attacker

- The attacker simply needs to replicate the BGP origination in the AS path to be accepted as “good”

So we really need to secure the BGP AS Path as well

We can do this with RPKI certs!

- Every eBGP speaker has a key that is certified by the AS
- When an update is passed to a neighbor AS, the router signs across the existing AS Path signature and the neighbor AS
- A BGPSEC speaker validates a received update by checking that
  - there is a current ROA to describe the address and origin AS
  - The received AS Path can be validated as a sequence of sign-over-sign operations by the AS keys

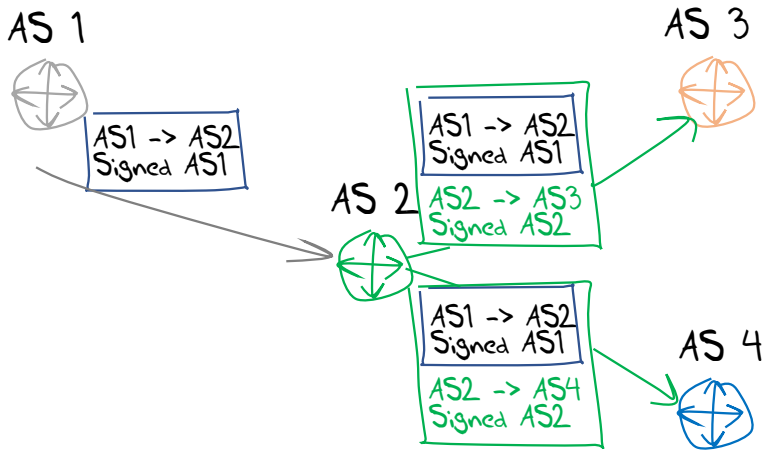

But this form of ASPath protection is hard…

- This is a replay of sBGP: BGPSEC cannot cope with partial adoption
  - It cannot jump across non-participating networks
- It has a high crypto overhead for session restarts
- It does not define how to promulgate the collection of certificates required to validate the digital signatures
- It does not necessarily identify and prevent route leaks

- Which means that BGPSEC is not looking like its going to be deployed everywhere

  Which means that there is little value in deploying it anywhere

Reviving Path Plausibility

- We were pretty convinced about the value of RPKI certificates and digital signatures
  - Because we really have nothing better to offer in their place
- But the AS Path protection elements of BGPSEC are a critical problem!
- We revived a variant of soBGP's AS connection certificate: ASPAs
  - Combines part of inter-AS topology and path inter-AS routing policy
  - An AS lists all of its "upstream" providers in a signed object
- This can provide a form of path plausibility together with a set of policy filters


Should have bought me
that pony

Additional Material

- Some measurement snapshots
- What can you do today
- soBGP redux
- Some General Comments

1 – Measurement Snapshots

- ROA creation is VERY widespread

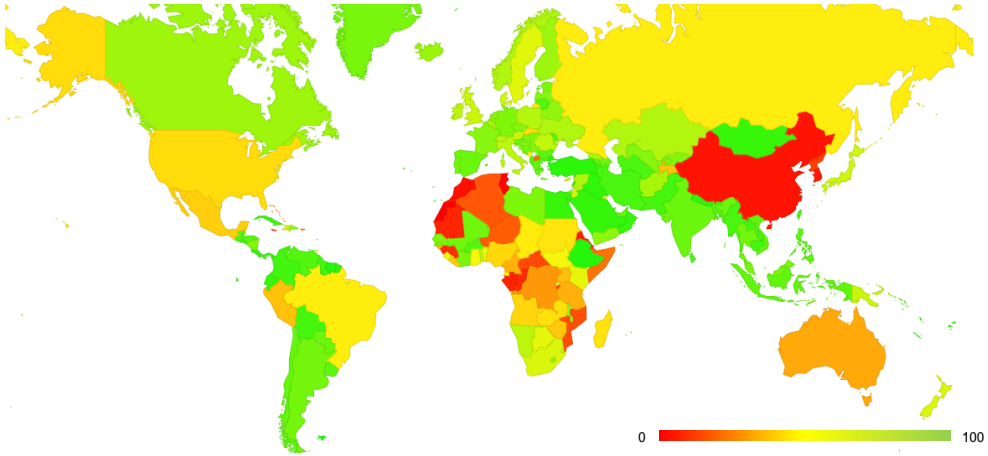

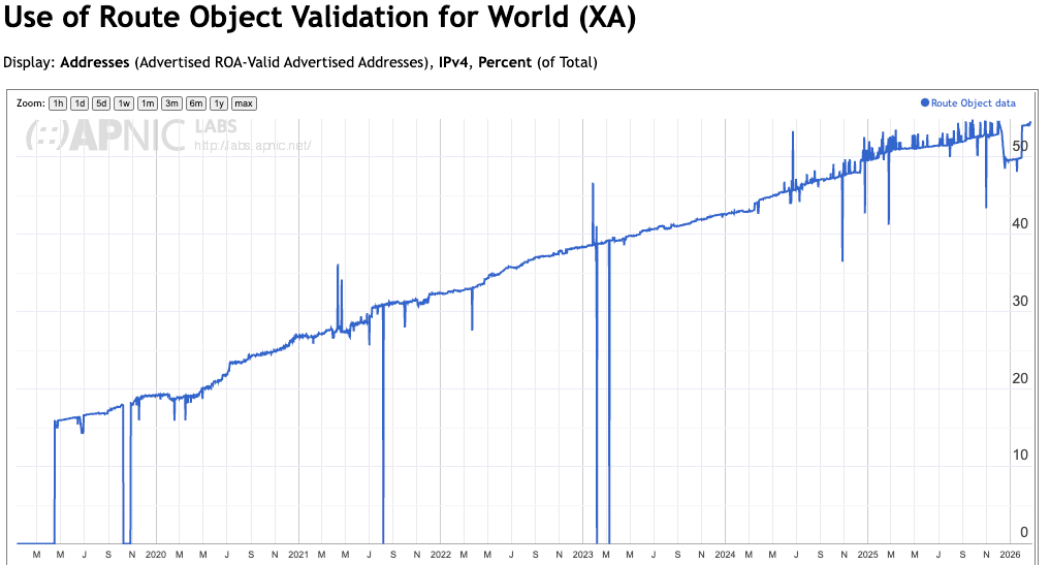

https://stats.labs.apnic.net/roas

Measurement Snapshots

- But applying ROAs to routing updates across all Ases has never really been adopted!

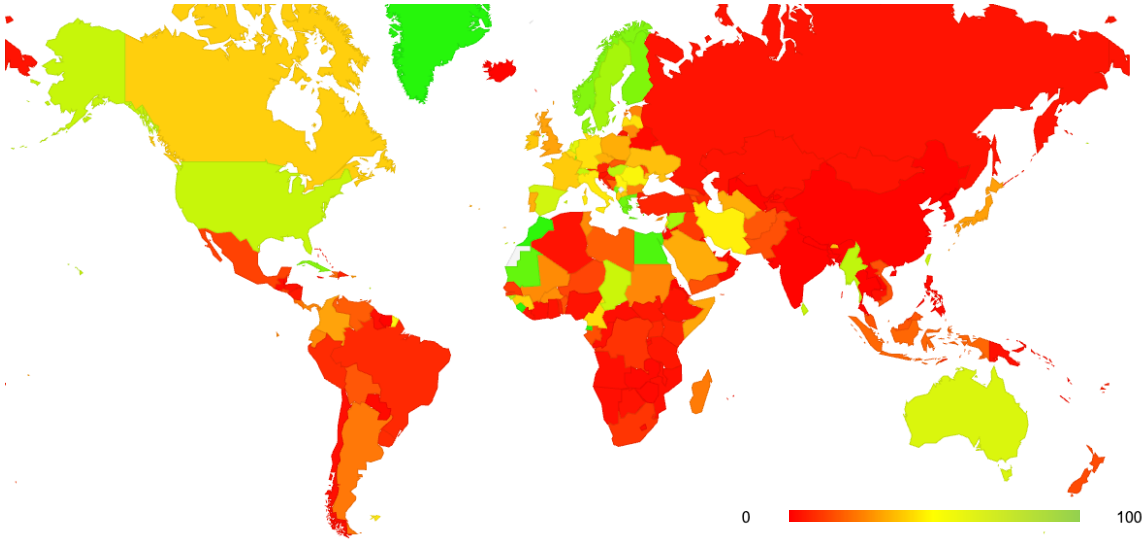

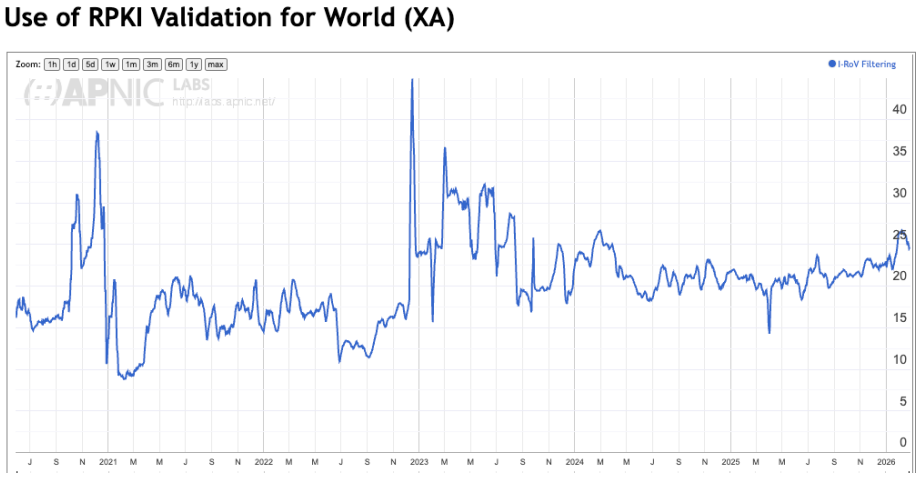

https://stats.labs.apnic.net/rpki

2. What can you do today?

“Don’t let the perfect be the enemy of the good!”

- You could just wait for a complete routing security framework to be invented
- Or you could do something practical right now that might be helpful

What can you do today?

You might want to take some steps to make routing attacks easier to detect and easier to deflect

- ❑ BCP38 filters can help
  - UDP DOS attacks are very common
- ❑ Generating ROAs can help
  - Maybe they won't help a lot today, but as more networks filter on ROAs then they will be more effective to protect against simple address hijacking
- ❑ Route Registry objects can help
  - www.irr.net
  - Again this is not a complete answer, but its better than nothing
- ❑ You should really should filter your customers
  - Filter customer routing updates according to BCP38, ROAs and IRR profiles
- ❑ Consider signing up to MANRS
  - https://www.manrs.org/ (Even spending a few minutes thinking about routing security is better than not thinking about it at all)
- ❑ DNSSEC-sign your domain name
- ❑ Validate DNS responses

3. soBGP: an alternative to BGPSEC

- Instead of the high overhead of AS Path validation we can look at secure origin BGP (soBGP) from 2003
- soBGP looked at the AS Path as a topology vector composed of a number of paired AS adjacencies
  - An AS publishes a signed adjacency attestation for all of its neighbors
  - If a signing AS appeared in an AS Path then its neighbors in the AS Path must also by described in the adjacency attestion

soBGP and AS Adjacencies

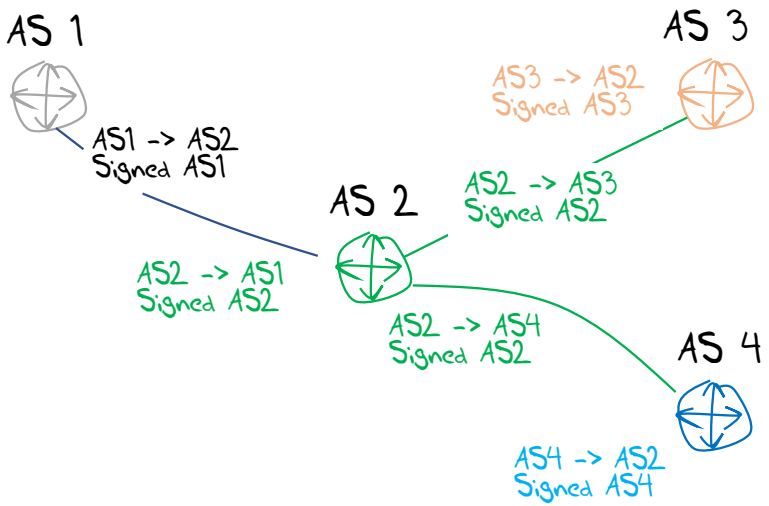

AS Path Processing using AS Adjacency 'hints'

AS1 -> AS2 -> AS3	plausible
AS1-> AS3 -> AS2	implausible
AS1-> AS2 -> AS3 -> AS4	implausible

soBGP compared to BGPSEC

- Lower crypto overhead
- Can be used in scenarios of partial adoption
- Does not prevent a network from learning false information, but prevents a network being used in a falsified AS path
  - Unless you also include the AS's peers
  - And so on
  - Incremental deployment generates incremental benefit
- Can include directionality in the AS adjacency attestation
  - As a simple "policy" filter

soBGP compared to ASPAs

- ASPAs add the aspect of a policy assertion into a eBGP peering session – namely that one network is an “upstream” provider to the other
  - Routes learned from one provider cannot be propagated to another provider
- Can ASPAs provide robust assurances in scenarios of partial deployment?
- Is a unilateral assertion of BGP adjacency an adequate protection against rogue AS’s generating bogus assertions?

4. Generic Concerns over PKIs

Is a ***trust hierarchy*** the best approach to use?

- The concern here is **concentration of vulnerability**

  If validation of routing information is dependent on the availability and validity of a single root trust anchor then what happens when this single digital artifact is attacked?

- But is there a viable alternative approach?

  Can you successfully incorporate robust diversity of authentication of security credentials into a supposedly highly resilient secure trust framework?

This is a very challenging question about the nature of trust in a diverse networked environment!

Web trust – 1,500 CAs vs DNSSEC trust – 1 key

Which is 'better'?

Generic Concerns over universality

A major issue here is that of ***partial use and deployment***

- This security mechanism has to cope with partial deployment in the routing system
  - The basic conventional approach of "what is not certified and proved as good must be bad" will not work in a partial deployment scenario
- In BGP we need to think about both origination and the AS Path of a route object in a partial deployed environment
  - AS path validation is challenging indeed in an environment of piecemeal use of secure credentials, as the mechanism cannot tunnel from one BGPsec "island" to the next "island"
- A partially secured environment may incur a combination of high incremental cost with only marginal net benefit to those deploying BGPsec

Generic Concerns: Prevention vs Detection

Is certification the ***only way*** to achieve useful outcomes in securing routing?

- Is this form of augmentation to BGP to enforce "protocol payload correctness" over-engineered, and does it rely on impractical models of universal adoption?
- Can various forms of *routing anomaly detectors* adequately detect the most prevalent forms of typos and deliberate lies in routing with a far lower overhead, and allow for unilateral detection of routing anomalies?
- Or are such anomaly detectors yet another instance of "cheap security pantomime" that offer a thinly veiled placebo of apparent security that is easily circumvented or fooled by determined malicious attack?