National Institute of Standards and Technology
U.S. Department of Commerce

INFORMATION TECHNOLOGY LABORATORY

2018 TECHNOLOGY exchange

ORLANDO FL OCTOBER 15-18

Practical BGP Security with RPKI

If there is a will to mitigate hijacks, there is a way.

Doug Montgomery (dougm@nist.gov)

Manager, Internet and Scalable Systems Research

https://www.nist.gov/programs-projects/robust-inter-domain-routing

Establishing the Technical Basis for Trustworthy Networking

BGP Systemic Vulnerabilities

- **Faults / Accidents**

WIRED: HOW A TINY ERROR SHUT OFF THE INTERNET FOR PARTS OF THE US

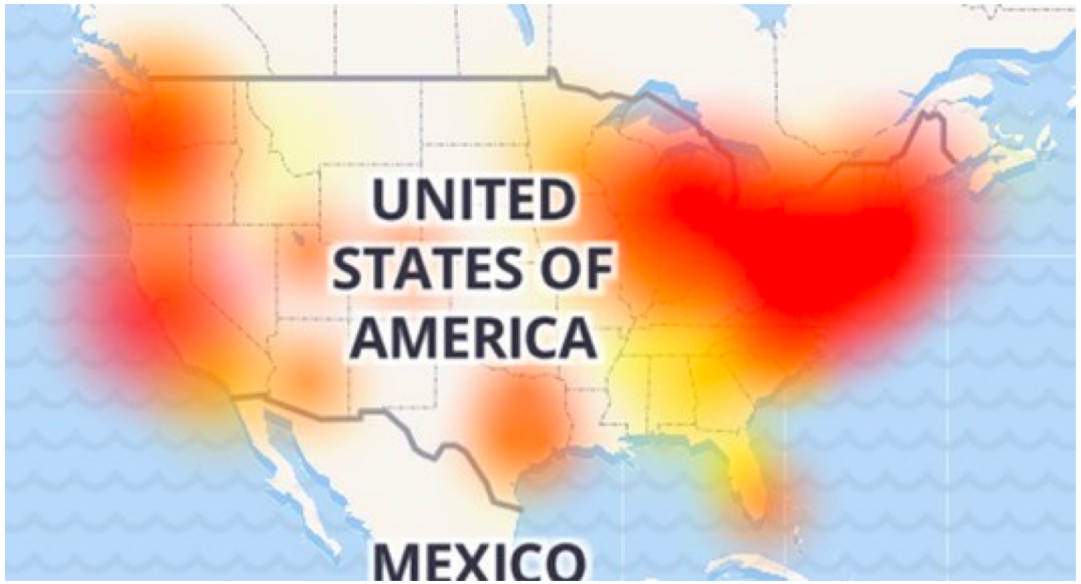

https://www.wired.com/story/how-a-tiny-error-shut-off-the-internet-for-parts-of-the-us/

- **Attacks**

Targeted Internet Traffic Misdirection

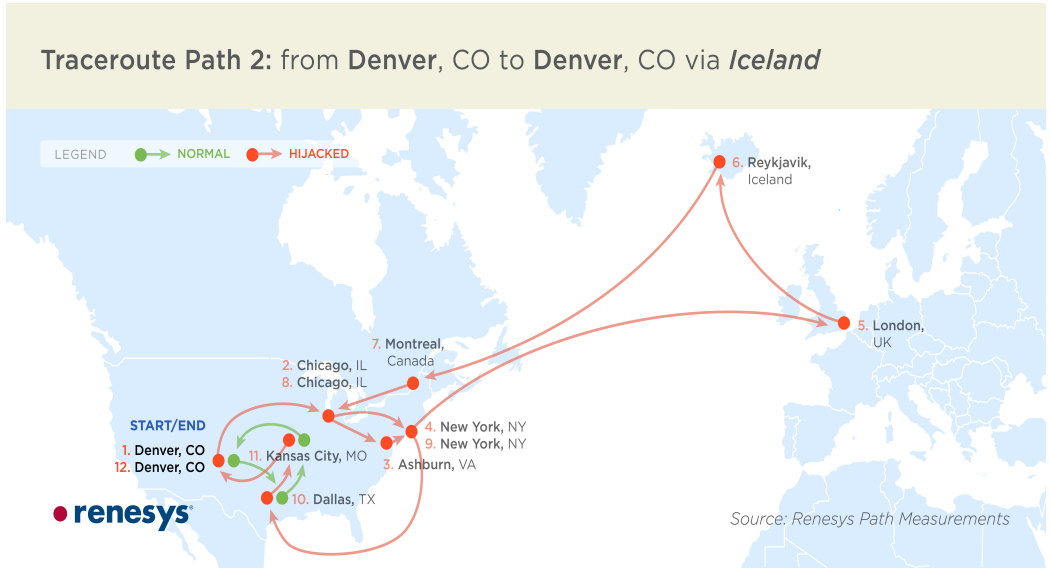

https://dyn.com/blog/mitm-internet-hijacking/

Broad Range of Threats and Motivations

- **Financially motivated attacks**


THREAT ANALYSIS

BGP Hijacking for Cryptocurrency Profit

THURSDAY, AUGUST 7, 2014
BY: JOE STEWART

https://www.secureworks.com/research/bgp-hijacking-for-cryptocurrency-profit


BIZ & IT —

Russian-controlled telecom hijacks financial services' Internet traffic

Visa, MasterCard, and Symantec among dozens affected by "suspicious" BGP mishap.

DAN GOODIN - 4/27/2017, 4:20 PM

https://arstechnica.com/information-technology/2017/04/russian-controlled-telecom-hijacks-financial-services-internet-traffic/

- **Other motivations …**


BIZ & IT —

How China swallowed 15% of 'Net traffic for 18 minutes

In April 2010, 15 percent of all Internet traffic was suddenly diverted ...

NATE ANDERSON - 11/17/2010, 2:45 PM

https://arstechnica.com/information-technology/2010/11/how-china-swallowed-15-of-net-traffic-for-18-minutes/

https://dyn.com/blog/iran-leaks-censorship-via-bgp-hijacks/

Broad Range of Risks

- **Attacks to leverage other vulnerabilities**

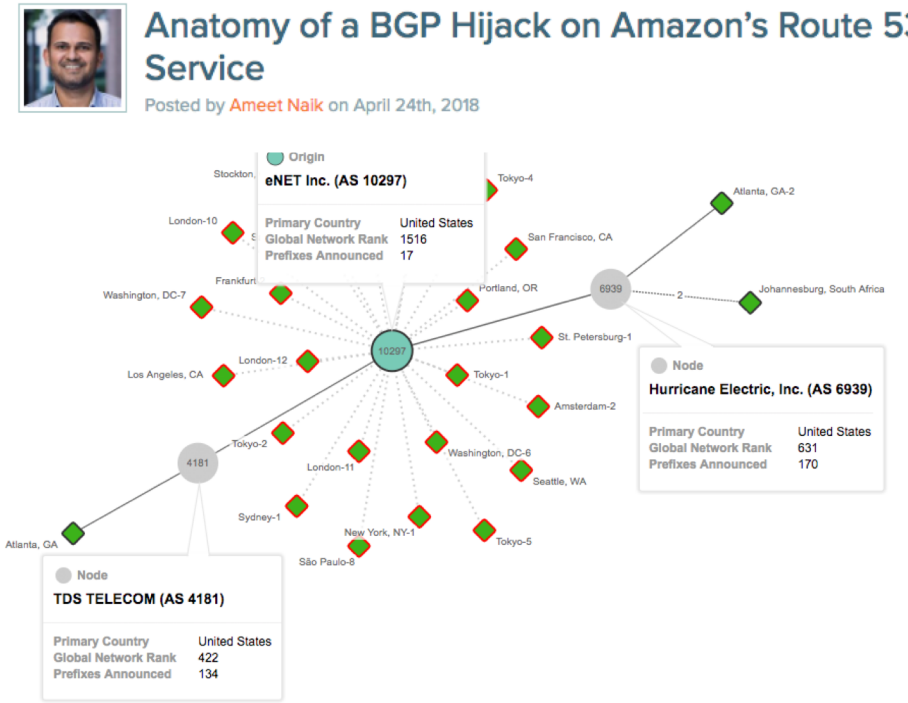

https://blog.thousandeyes.com/amazon-route-53-dns-and-bgp-hijack/

- **… and to undermine other infrastructure.**

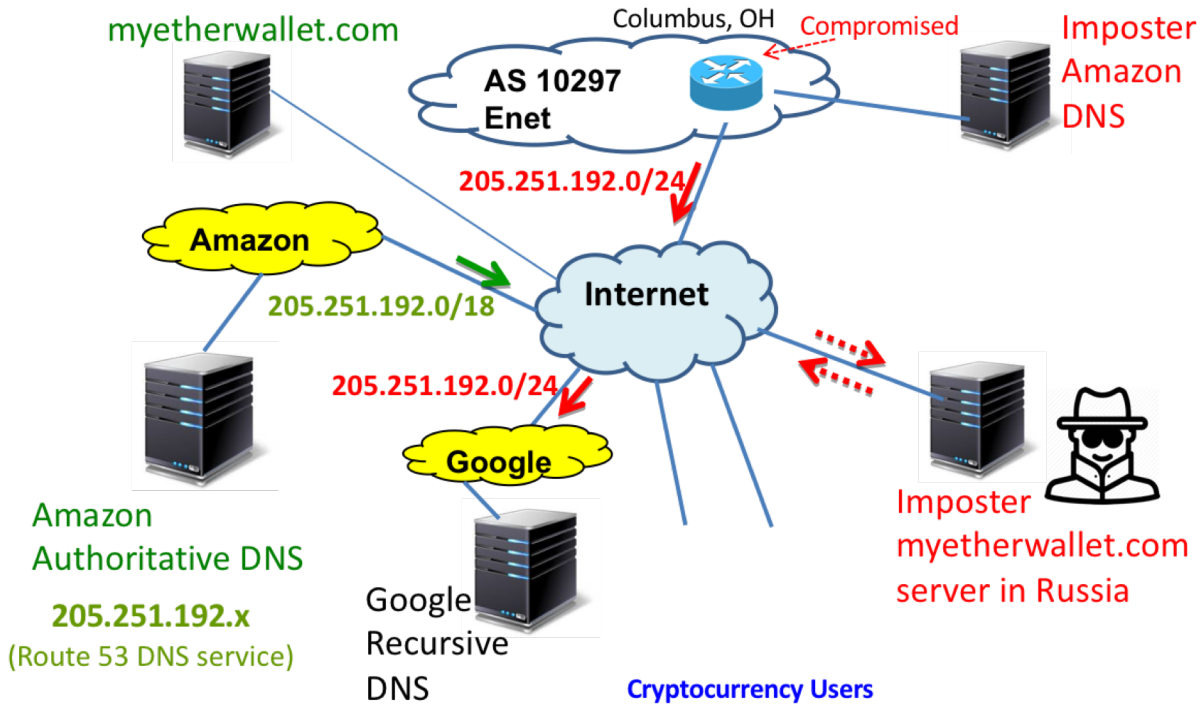

https://www.internetsociety.org/blog/2018/04/amazons-route-53-bgp-hijack/

BGP Systemic Vulnerabilities

– **Threats**
  - **Route hijacks**
    – Steers traffic away from legitimate servers
  - **Address squatting**
    – Hijacks a not-in-service prefix and sets up spam servers
  - **Route detours**
    – Modifies path causing data to flow via the attacker
  - **Route leaks**
    – Announces routes in violation of ISP policy.

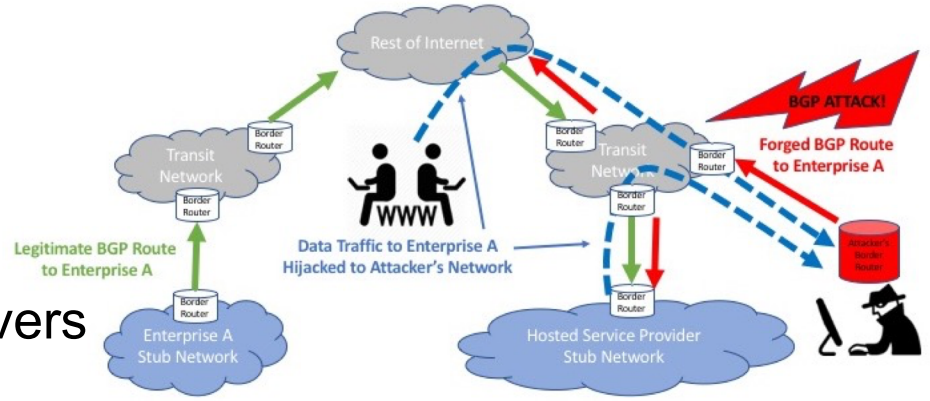

– **Ramifications**
  – Exploitations commonly result in outages, spam, misrouting of data traffic, eavesdropping on user data, DDoS, etc.

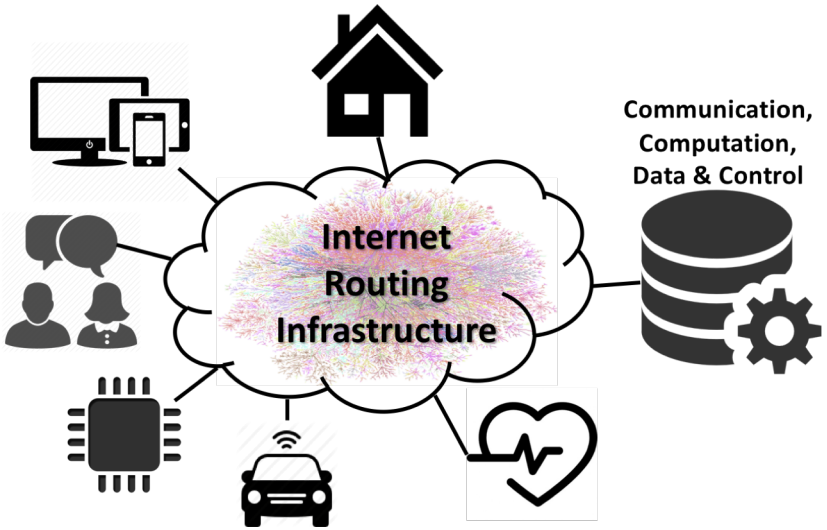

Do We Have the Will to Solve Problem?

- **These are isolated events!**
  - https://bgpstream.com/
  - https://bgpmon.net/category/hijack/
  - https://www.internetsociety.org/blog/2018/01/14000-incidents-2017-routing-security-year-review/
  - https://blog.thousandeyes.com/category/outage-reports/

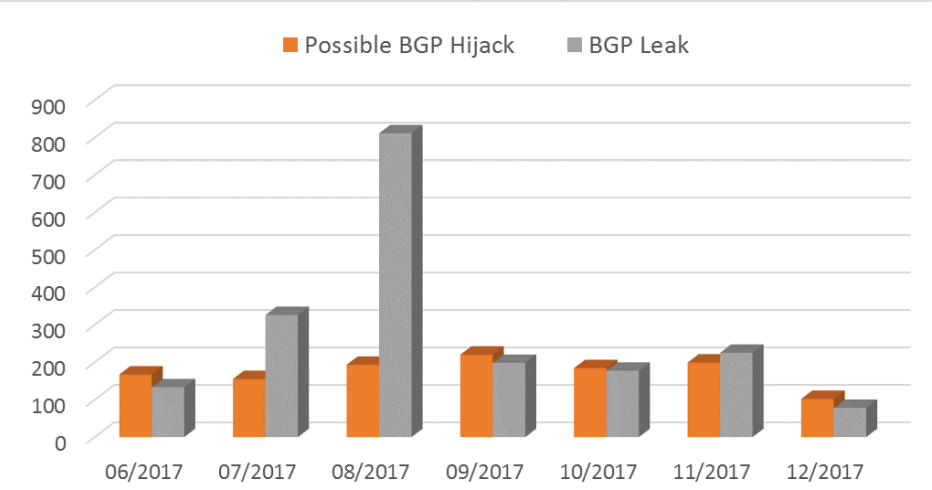

- **We watch the important routes!**
  - Important to who?
    - Your enterprise, your ISP, ISPs in other countries?
  - Measured how?
    - # of eyeballs and traffic are not the only measures of importance!
  - How does this scale?
    - ~900K prefix / origin pairs in DFZ
  - Is it proactive or reactive?
    - Is notification after the fact good enough?

Source: BGPstream

State of the Solutions Space

- **BGP Origin Validation (BGP-OV)**
  - Global public key infrastructure and protocol elements to enable BGP routers to verify that the origin AS in a BGP update, was authorized by the prefix owner to announce that route.
- **BGP Path Validation (BGP-PV)**
  - Leverages the same PKI to enable each AS to digitally sign a BGP update, proving that each AS in the PATH authorized the route announcement to its next hop.
- **BGP Route Leak Detection / Mitigation**
  - BGP protocol modifications to allow networks to detect that a BGP routed path violates typical customer-provider-peer policies for route redistribution.

Resource Certification

RIR Hosted Resource Public Key Infrastructure

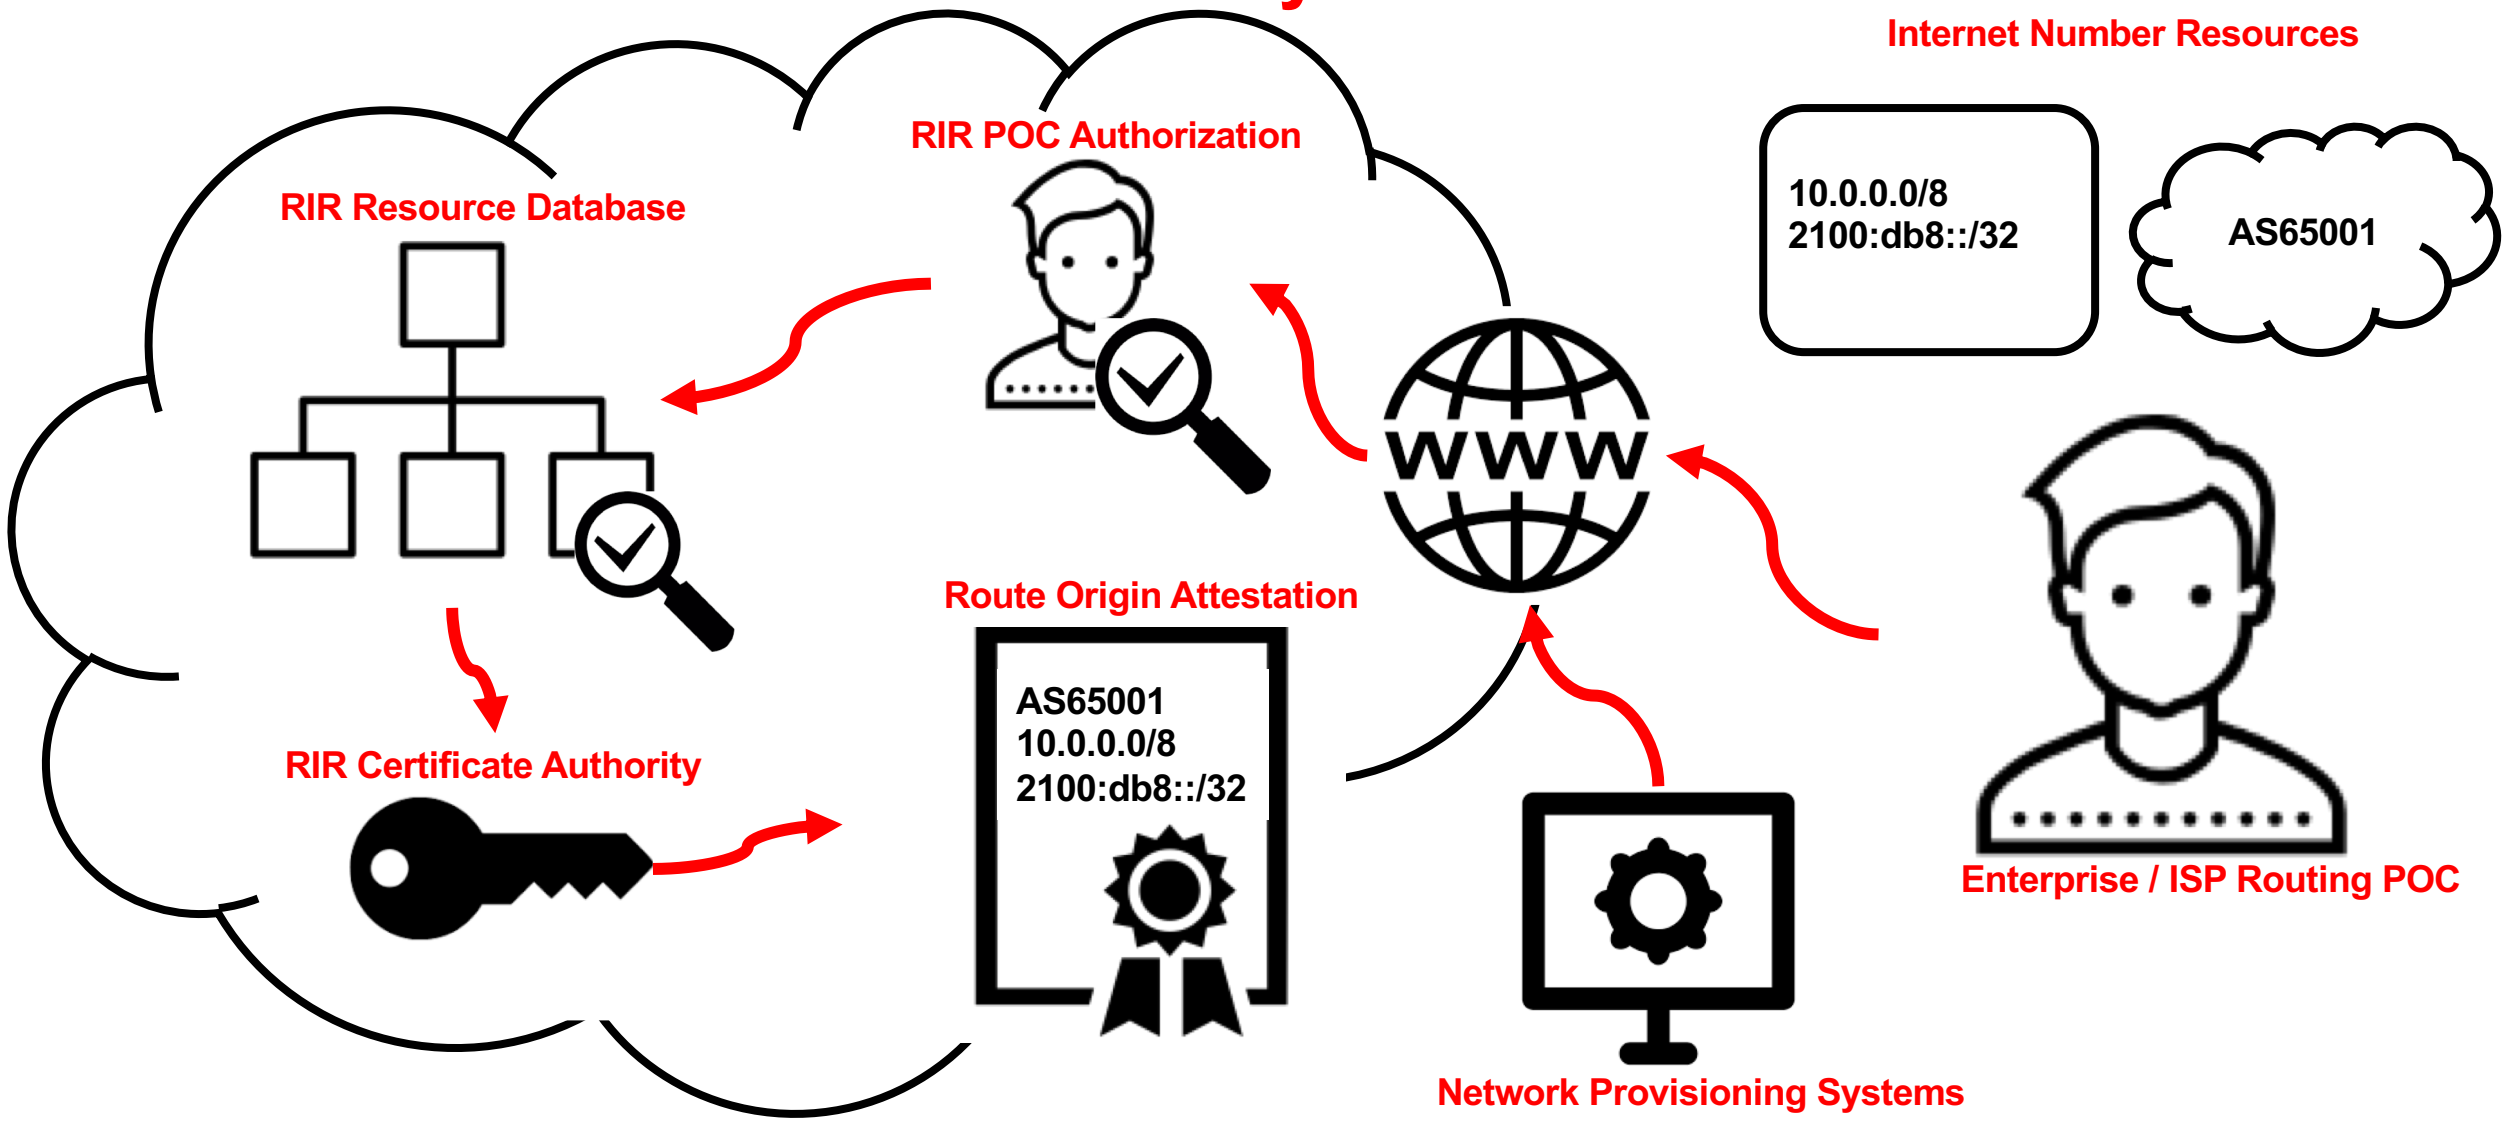

RPKI Publication & Mirroring

RIR Hosted Resource Public Key Infrastructure

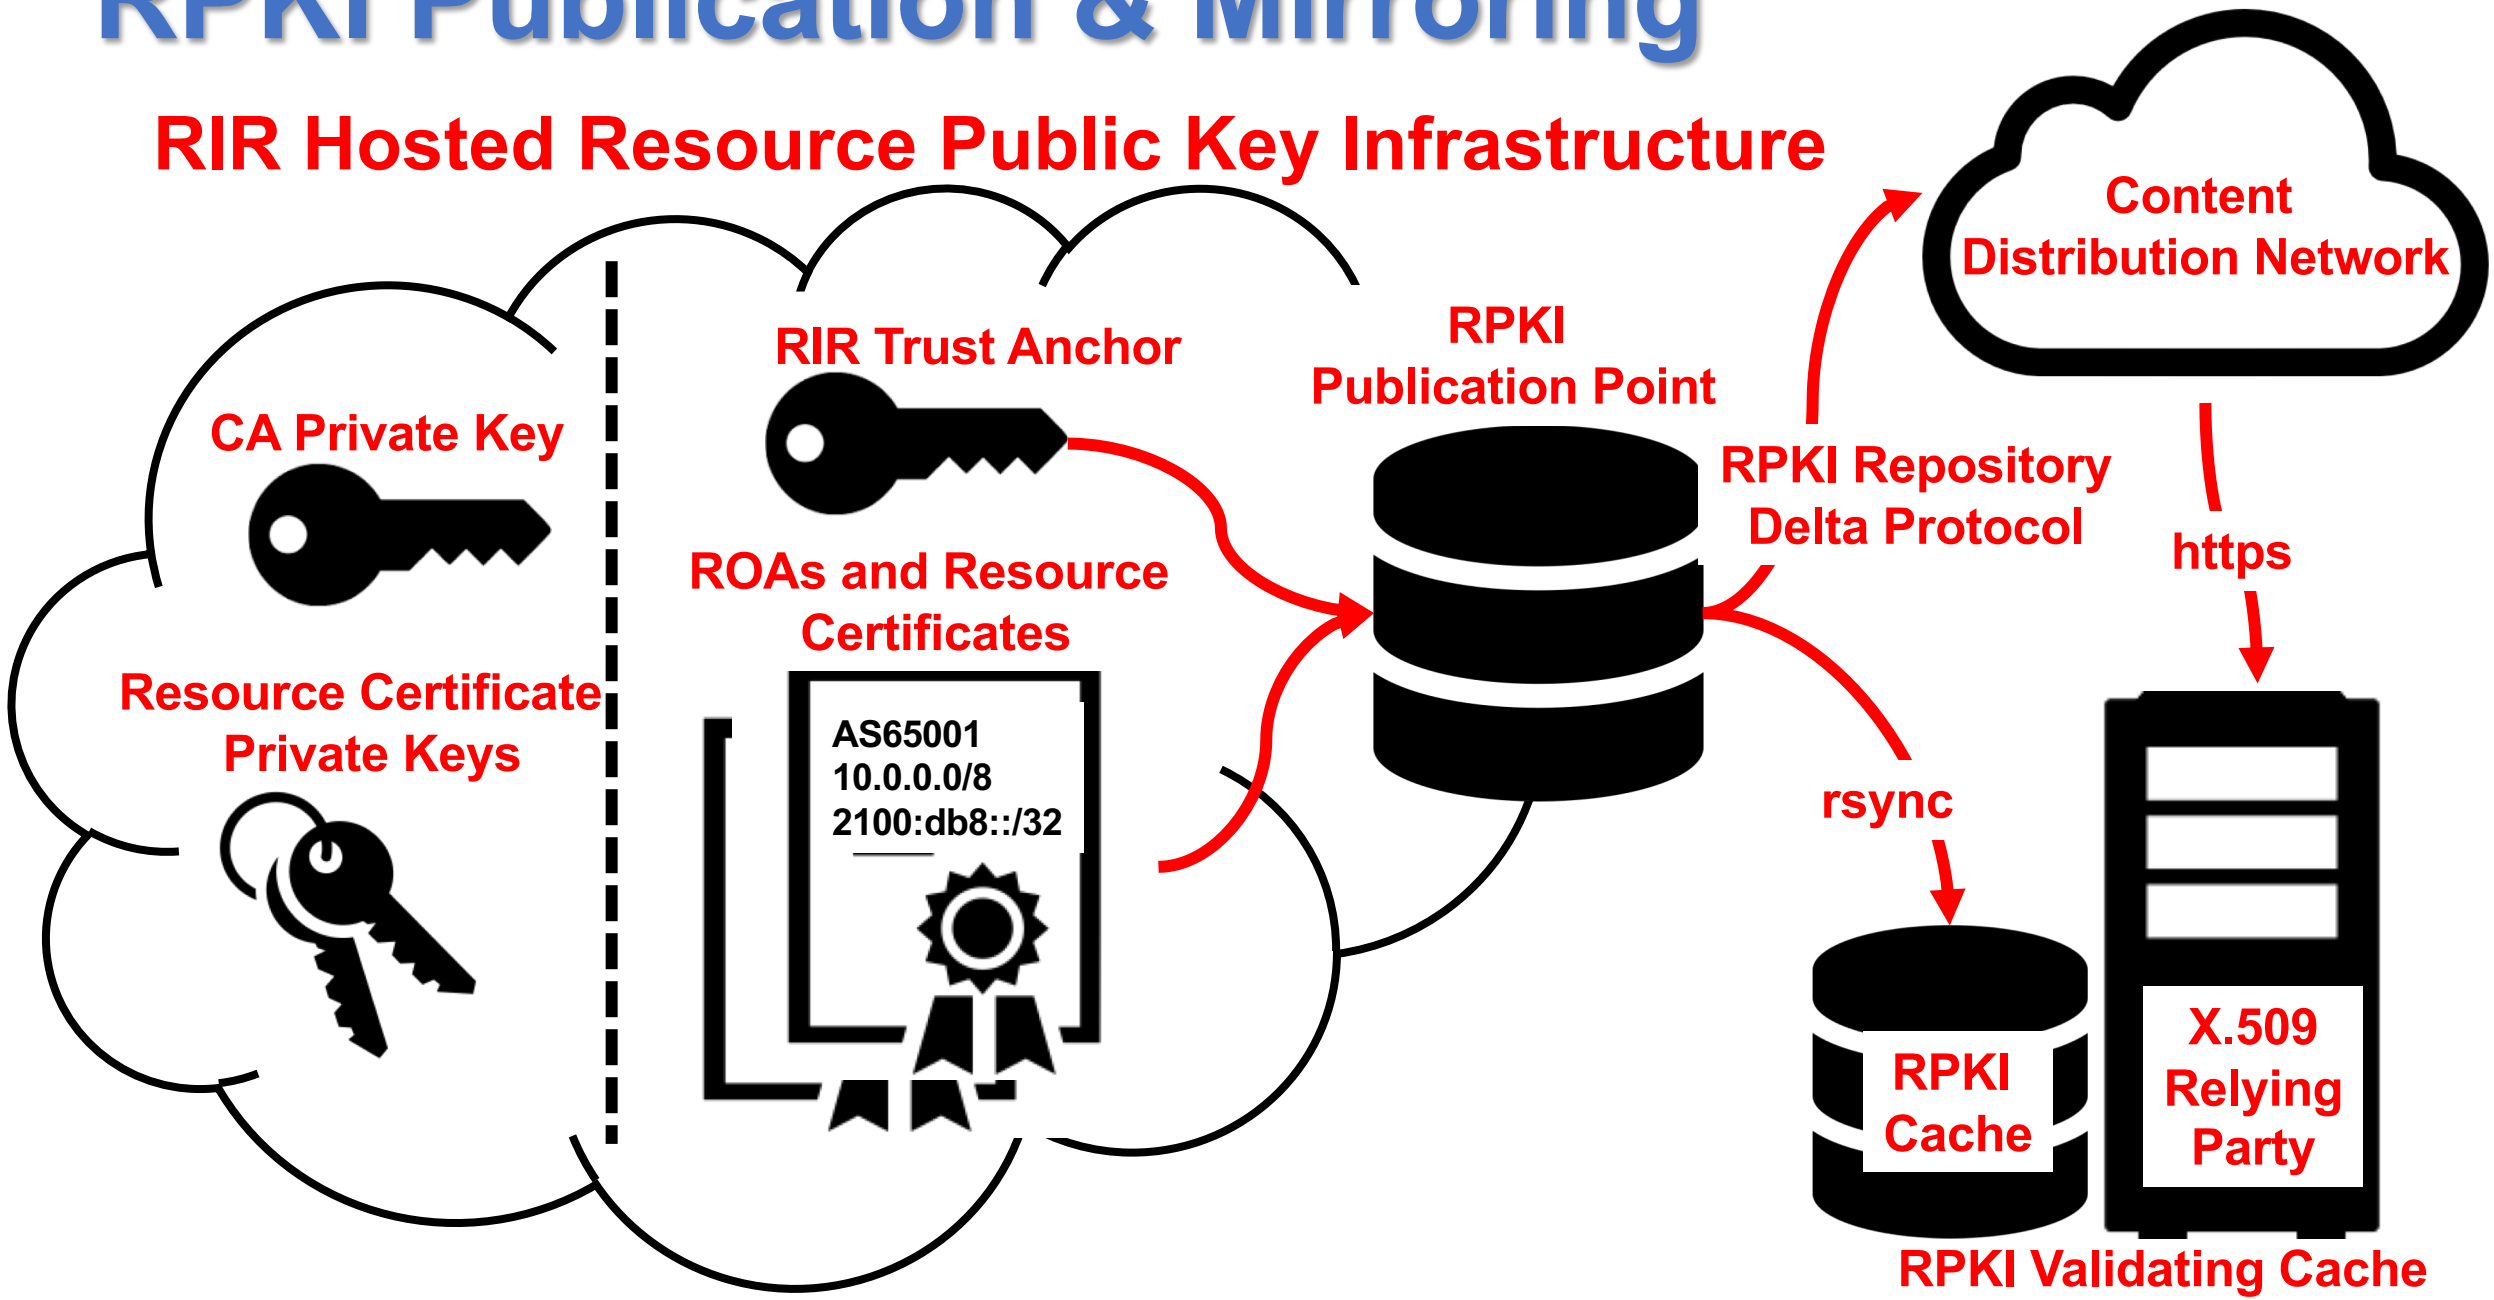

Validating RPKI Data & Passing to Router

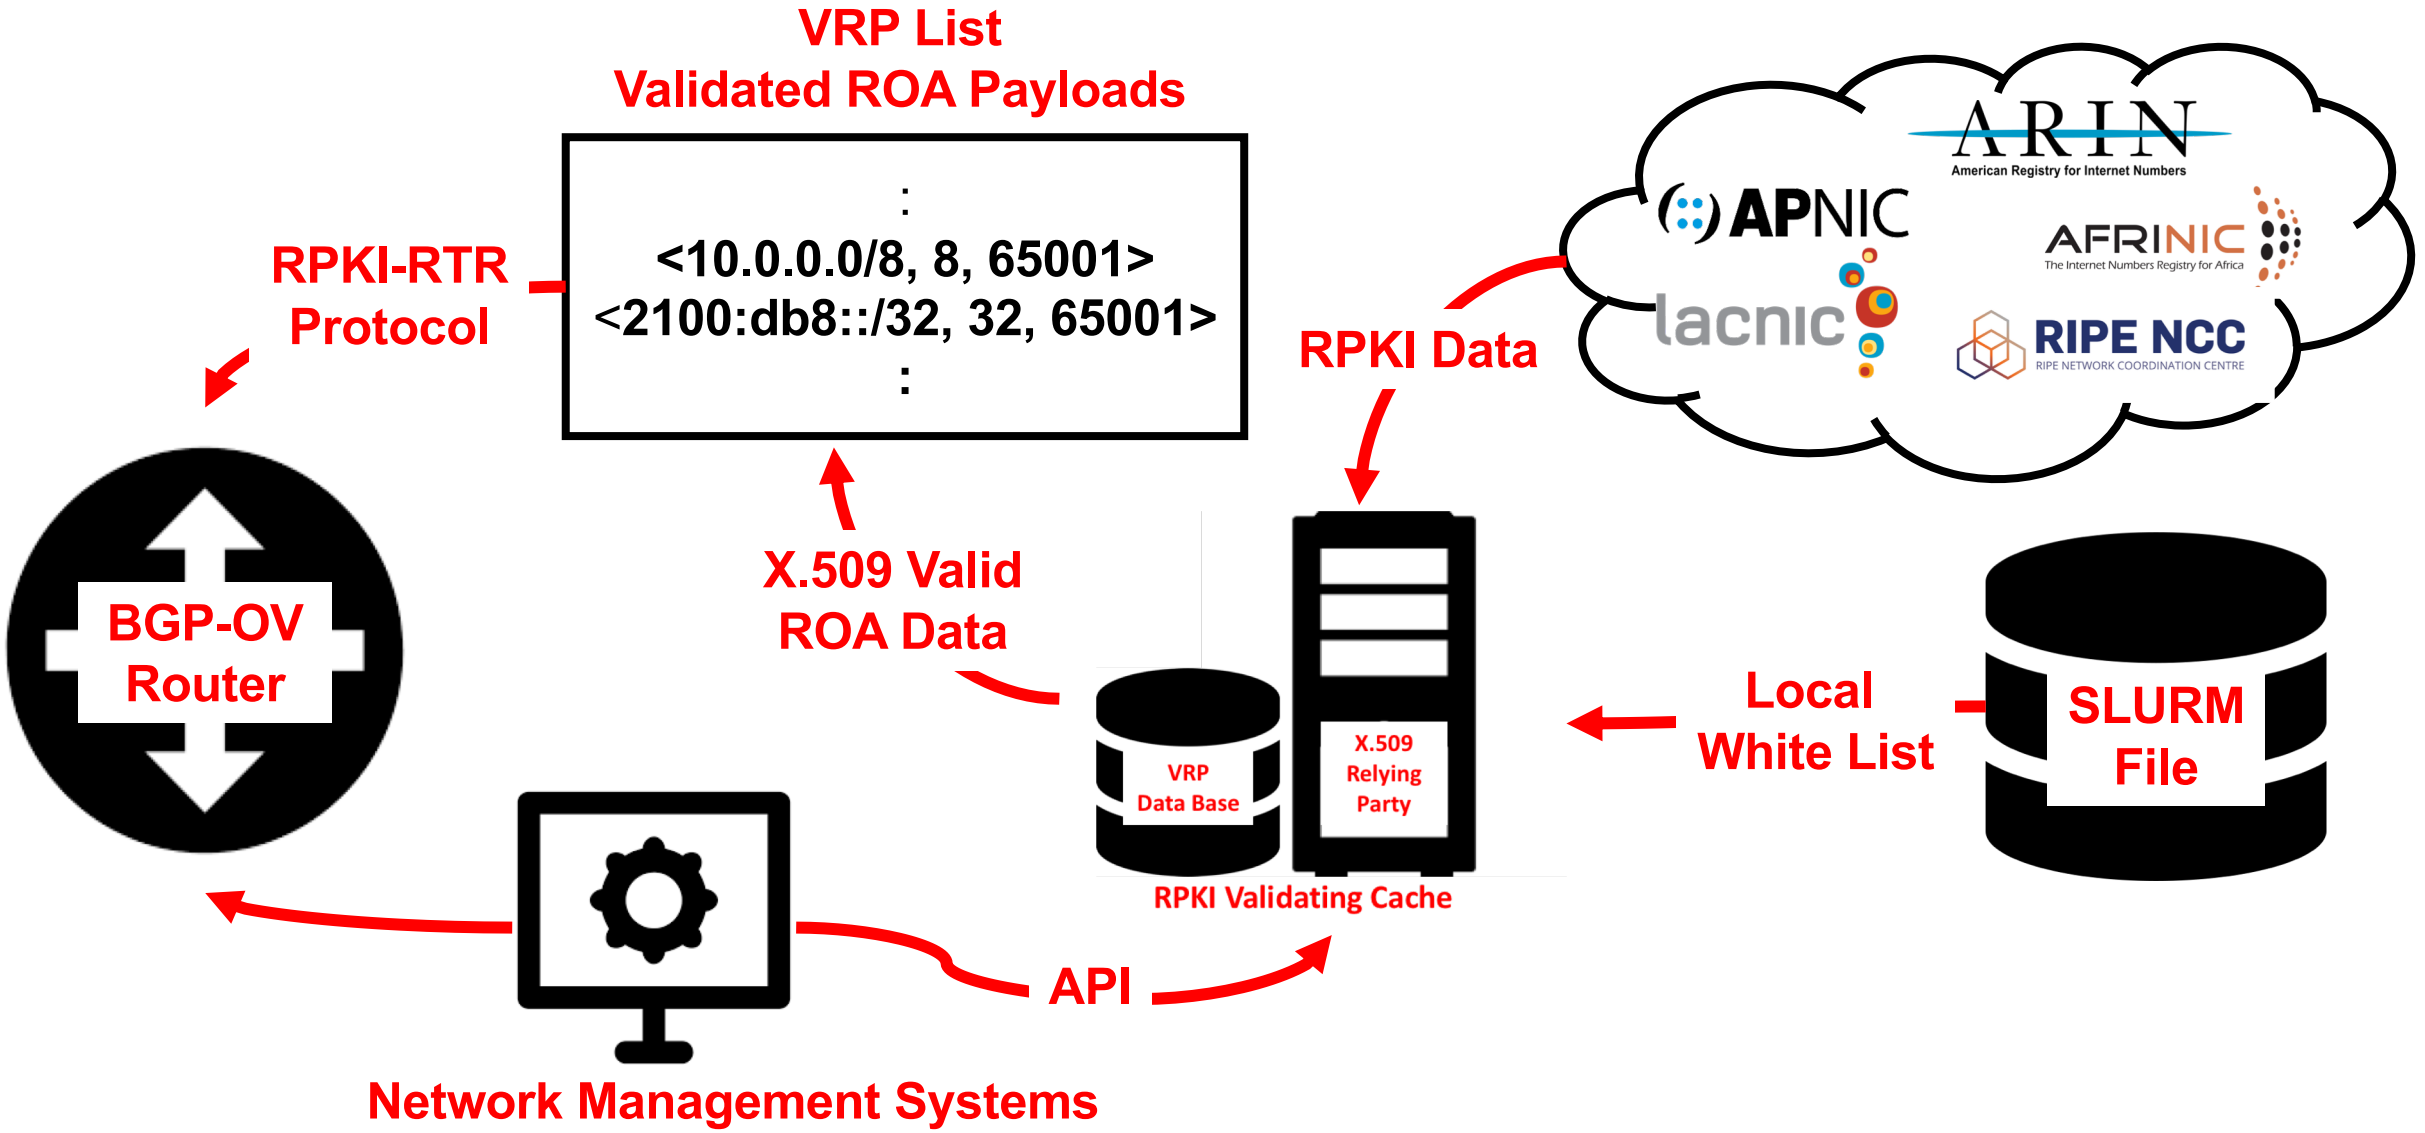

BGP Origin Validation

- **Origin Validation Algorithm RFC6811**
  - For each received eBGP and iBGP announcement.
  - Compare each received prefix+origin to VRP list.
    - **NotFound**: No VRP Covers the Route Prefix.
    - **Valid:** At least one VRP Matches the Route Prefix.
    - **Invalid:** At least one VRP Covers the Route Prefix, but no VRP Matches it.
  - **Local Policy Decisions**
    - Tag routes with OV result but take no action
    - Pref or depref routes based upon OV result
    - Ignore routes based upon OV result.
  - **Be Careful in Choosing Policies**
    - Not all options provided are effective.

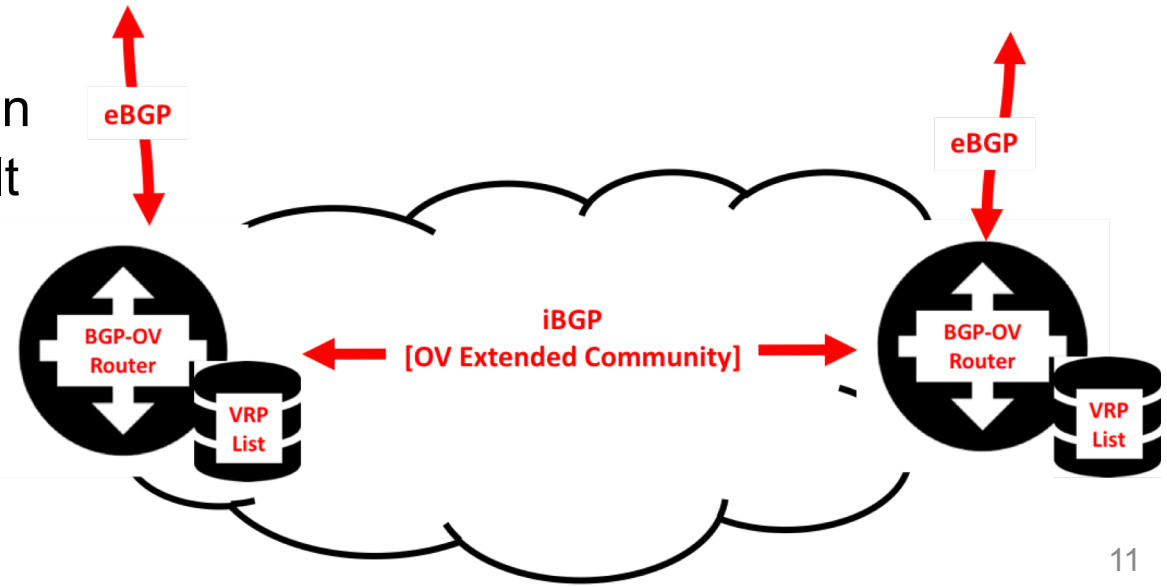

BGP Origin Validation - Status

- **Core specifications are complete**
  - IETF SIDR Working Group - https://datatracker.ietf.org/wg/sidr/documents/
  - IETF SIDR Ops WG - https://datatracker.ietf.org/wg/sidrops/documents/

- **Commercial implementations and production services exists.**
  - All 5 RIRs operate production RPKI services.
  - Multiple independent implementations of RPI Validating Caches exist.
    - More are on the way.
  - Commercial and opensource routers support RPKI-based origin validation.

- **Pilot and operational deployments.**

- **New specifications and implementations are emerging for:**
  - Deployment optimizations and implementation clarifications
  - Extensions to support additional security services.

Resource Certification Status

- **Production RPKI services in RIRs:**
  - AFRINIC:
    - http://afrinic.net/en/initiatives/rpki-certification
  - APNIC:
    - http://www.apnic.net/services/services-apnic-provides/resource-certification
  - ARIN
    - https://www.arin.net/resources/rpki/
  - LACNIC:
    - https://rpki.lacnic.net/rpki/
  - RIPE NCC:
    - http://www.ripe.net/certification/

Resource Certification Measurement

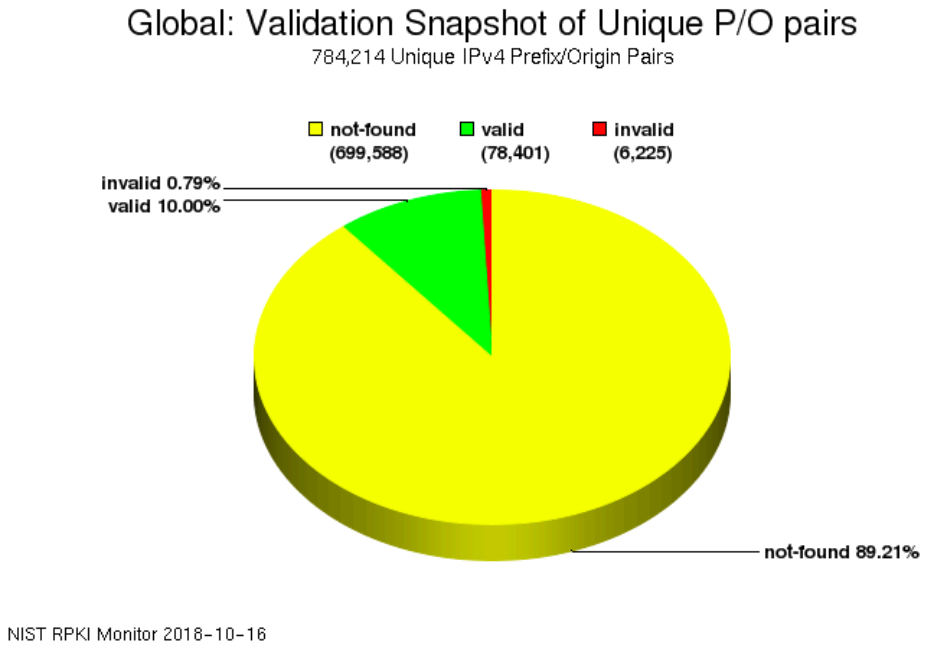

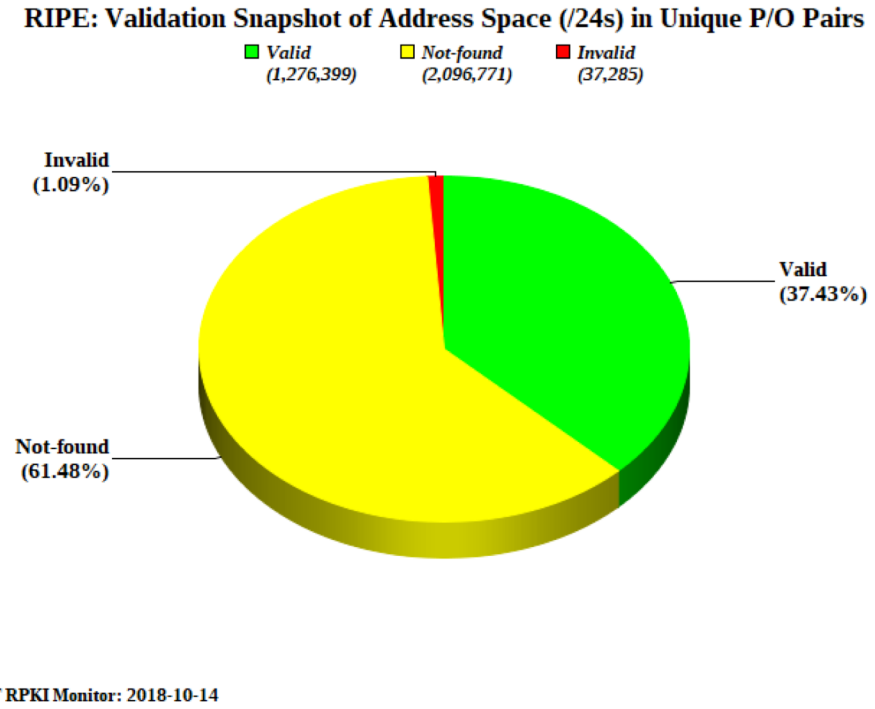

NIST RPKI Monitor: https://rpki-monitor.antd.nist.gov/

Who are the Early Adopters?

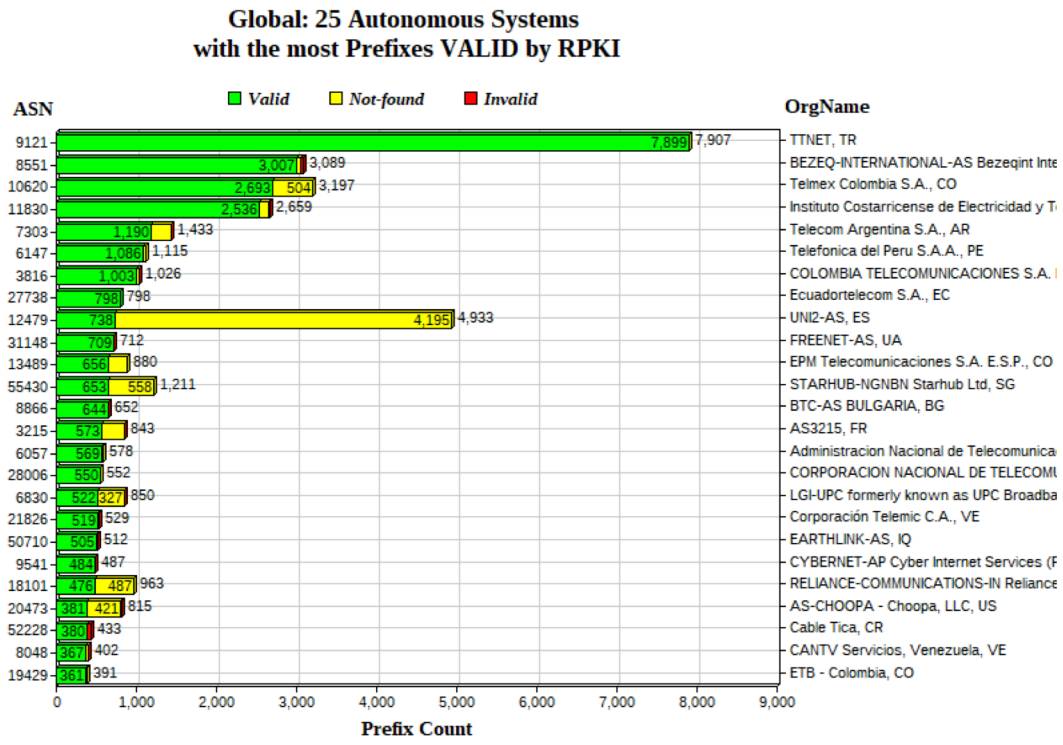

NIST RPKI Monitor: 2018-10-14

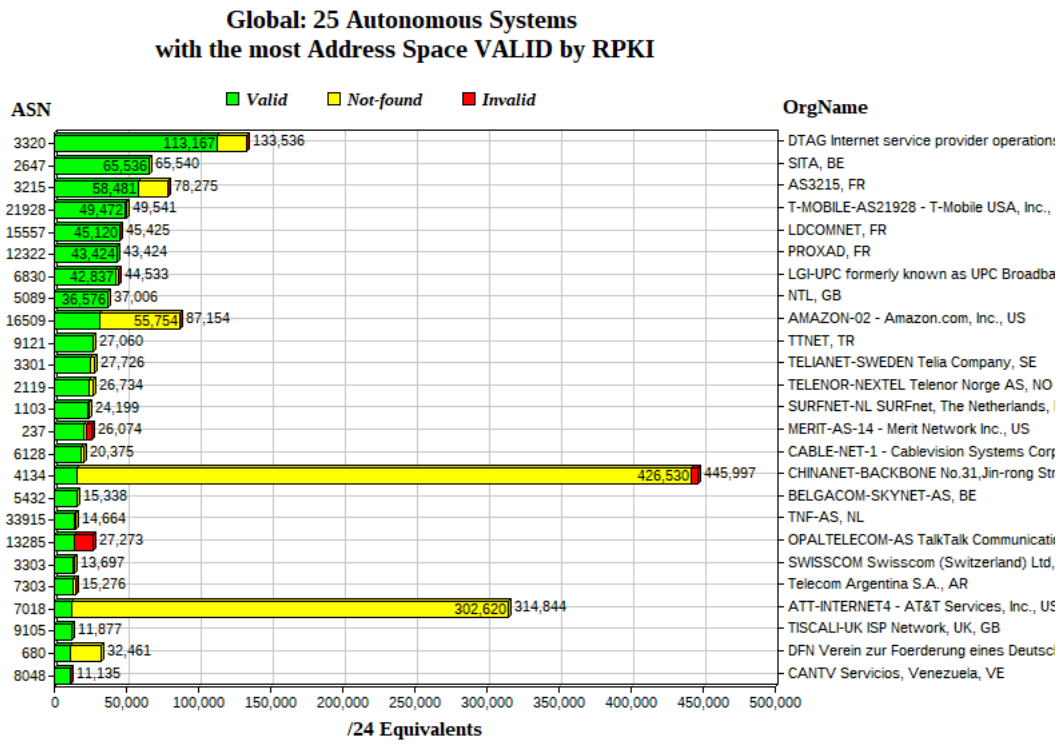

NIST RPKI Monitor: 2018-10-14

Current Invalid Routes?

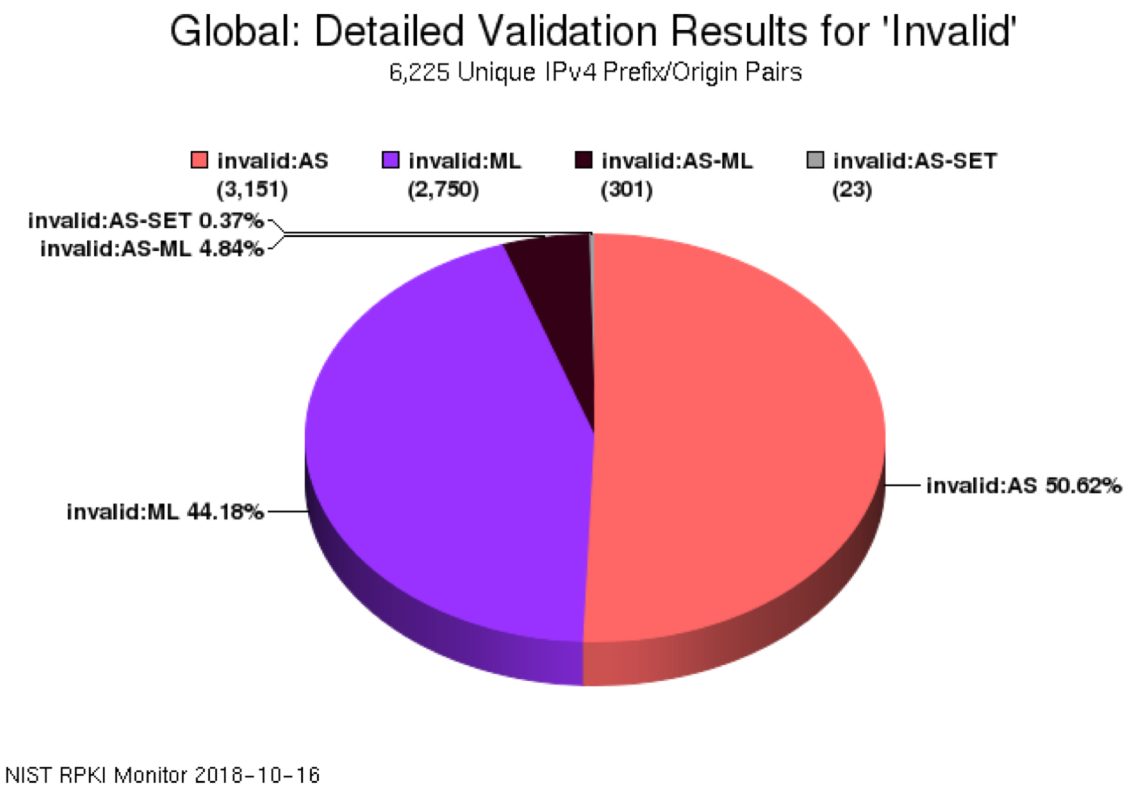

- **Is this an issue?**
  - 68% of current INVALID routes are covered by route that is VALID (97%) or NOTFOUND (3%).

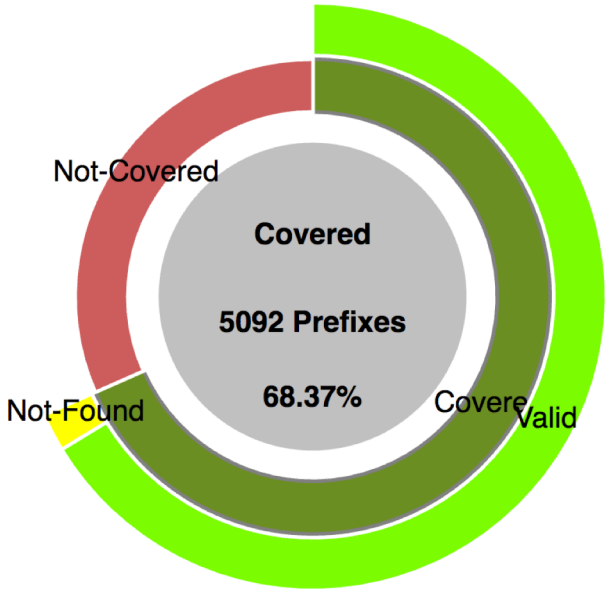

Who is Filtering Based Upon RPKI?

- **Tough Measurement Problem**
  - See: *Towards a Rigorous Methodology for Measuring Adoption of RPKI Route Validation and Filtering*, ACM SIGCOMM Computer Communication Review Volume 48 Issue 1, January 2018.
  - *See live measurements using the above methodology at: https://rov.rpki.net/*
    - *~65 ASes – most using ROV@AMS-IX route server.*

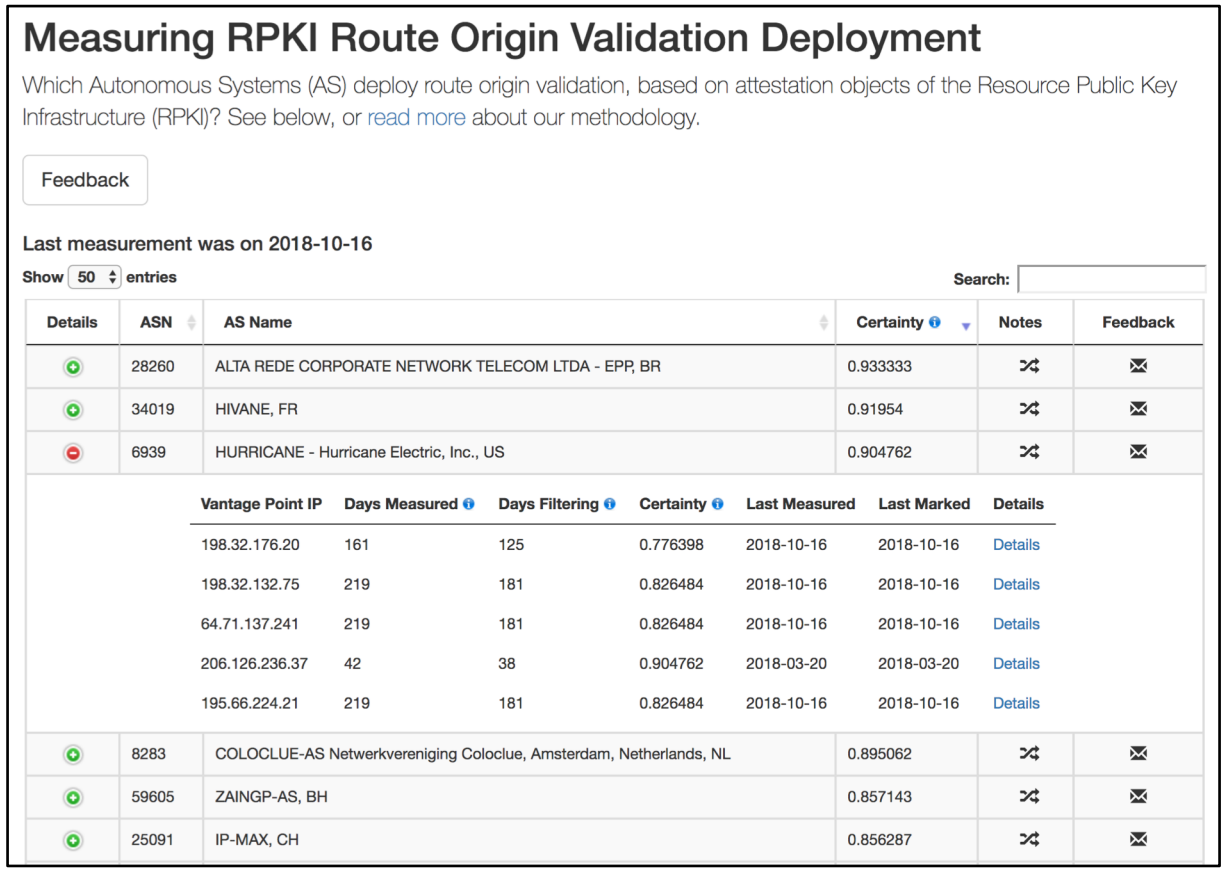

Measuring RPKI Route Origin Validation Deployment

Which Autonomous Systems (AS) deploy route origin validation, based on attestation objects of the Resource Public Key Infrastructure (RPKI)? See below, or read more about our methodology.

Feedback

Last measurement was on 2018-10-16

Show 50 entries Search:

Details	ASN	AS Name	Certainty	Notes	Feedback
⊕	28260	ALTA REDE CORPORATE NETWORK TELECOM LTDA - EPP, BR	0.933333	⤮	✉
⊕	34019	HIVANE, FR	0.91954	⤮	✉
⊖	6939	HURRICANE - Hurricane Electric, Inc., US	0.904762	⤮	✉

Vantage Point IP	Days Measured	Days Filtering	Certainty	Last Measured	Last Marked	Details
198.32.176.20	161	125	0.776398	2018-10-16	2018-10-16	Details
198.32.132.75	219	181	0.826484	2018-10-16	2018-10-16	Details
64.71.137.241	219	181	0.826484	2018-10-16	2018-10-16	Details
206.126.236.37	42	38	0.904762	2018-03-20	2018-03-20	Details
195.66.224.21	219	181	0.826484	2018-10-16	2018-10-16	Details

⊕	8283	COLOCLUE-AS Netwerkvereniging Coloclue, Amsterdam, Netherlands, NL	0.895062	⤮	✉
⊕	59605	ZAINGP-AS, BH	0.857143	⤮	✉
⊕	25091	IP-MAX, CH	0.856287	⤮	✉

BGP-OV Implementations

- **RPKI Infrastructure**
- RIPE validator 2
  - https://www.ripe.net/manage-ips-and-asns/resource-management/certification/tools-and-resources
- RIPE validator 3
  - https://github.com/RIPE-NCC/rpki-validator-3/wiki
- Routinator - Open source Relying Party - NLnet Labs
  - https://github.com/NLnetLabs/routinator
- RPKI.net Open Source Implementation of RPKI Tools
  - https://github.com/dragonresearch/rpki.net/
- RPSTIR - BBN Validation Tool
  - https://sourceforge.net/projects/rpstir/

- **Router Implementations**
  - Cisco
    - https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/command/irg-cr-book/bgp-m1.html#wp3677719851
  - Juniper
    - https://www.juniper.net/documentation/en_US/junos12.2/topics/topic-map/bgp-origin-as-validation.html
  - Nokia
    - https://infoproducts.alcatel-lucent.com/cgi-bin/dbaccessfilename.cgi/9300731102_V1_7750%20SR%20OS%20Router%20Configuration%20Guide%2012.0.R4.pdf
  - Quagga / FRR, BIRD
    - https://www.nist.gov/services-resources/software/bgp-secure-routing-extension-bgp-srx-prototype
    - http://rtrlib.realmv6.org/
  - Go BGP
    - https://github.com/osrg/gobgp

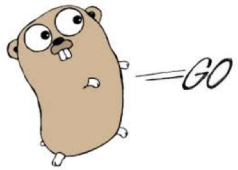

Emerging Community Guidance w/ RPKI

- Mutually Agreed Norms for Routing Security (MANRS) Implementation Guide.
  - https://www.manrs.org/isps/guide/
- Report to the President on Enhancing Resilience Against Botnets.
  - https://www.commerce.gov/page/report-president-enhancing-resilience-against-botnets
- Secure Interdomain Traffic Exchange, Draft NIST Special Publication 800-189.
  - To be released Nov 2018.
- Draft: NIST SP 1800-14, *Protecting the Integrity of Internet Routing: Border Gateway Protocol (BGP) Route Origin Validation*.
  - https://www.nccoe.nist.gov/projects/building-blocks/secure-inter-domain-routing

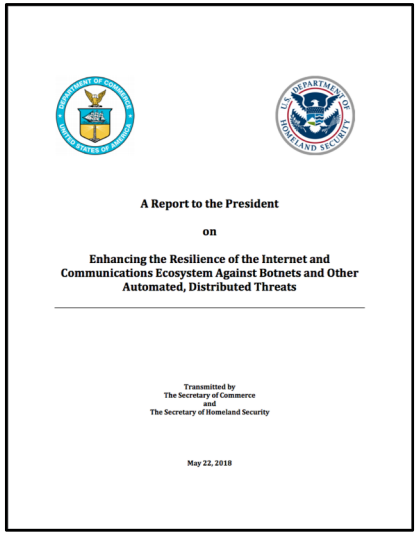

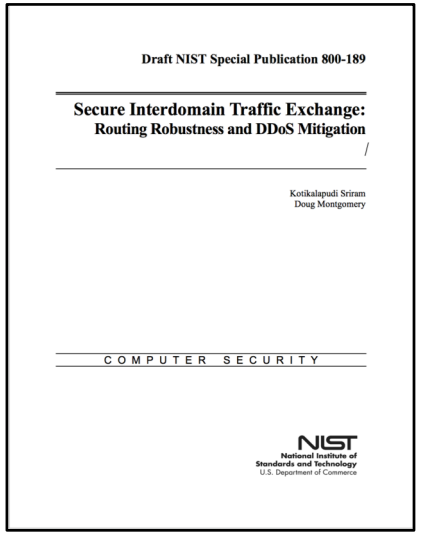

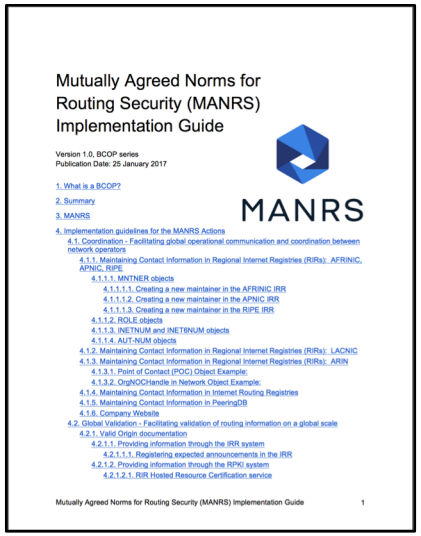

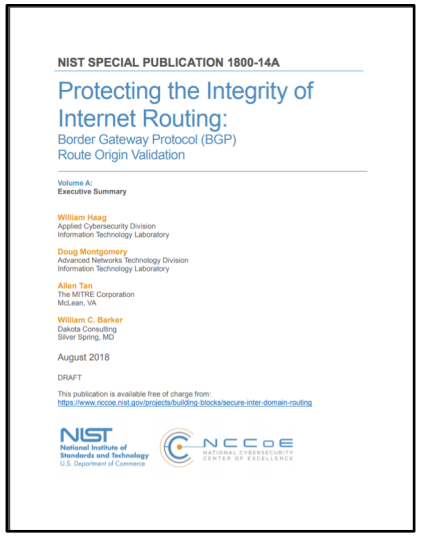

NCCoE BGP-OV Demonstration Project

- **Validating Caches**
  - 3 distinct implementations
- **Routers**
  - 2 customer edge class commercial hardware routers.
  - 2 commercial VM routers.
  - 1 software router.
- **Basic Functionality Tests.**
  - Complete
- **Robustness / Scaling Tests.**
  - Ongoing

- **Interesting Issues**
  - Cache lacking support for delta update by serial query. Router required to request full tables each time.
  - SLURM file size restrictions, and software bugs related to SLURM use when not connected to RPKI publication points.
  - Interoperability issues in iBGP signaling of ROV status in extended community string.
  - Inconsistent handling of AS_SETs across vendors.

NCCoE BGP-OV Demonstration Project

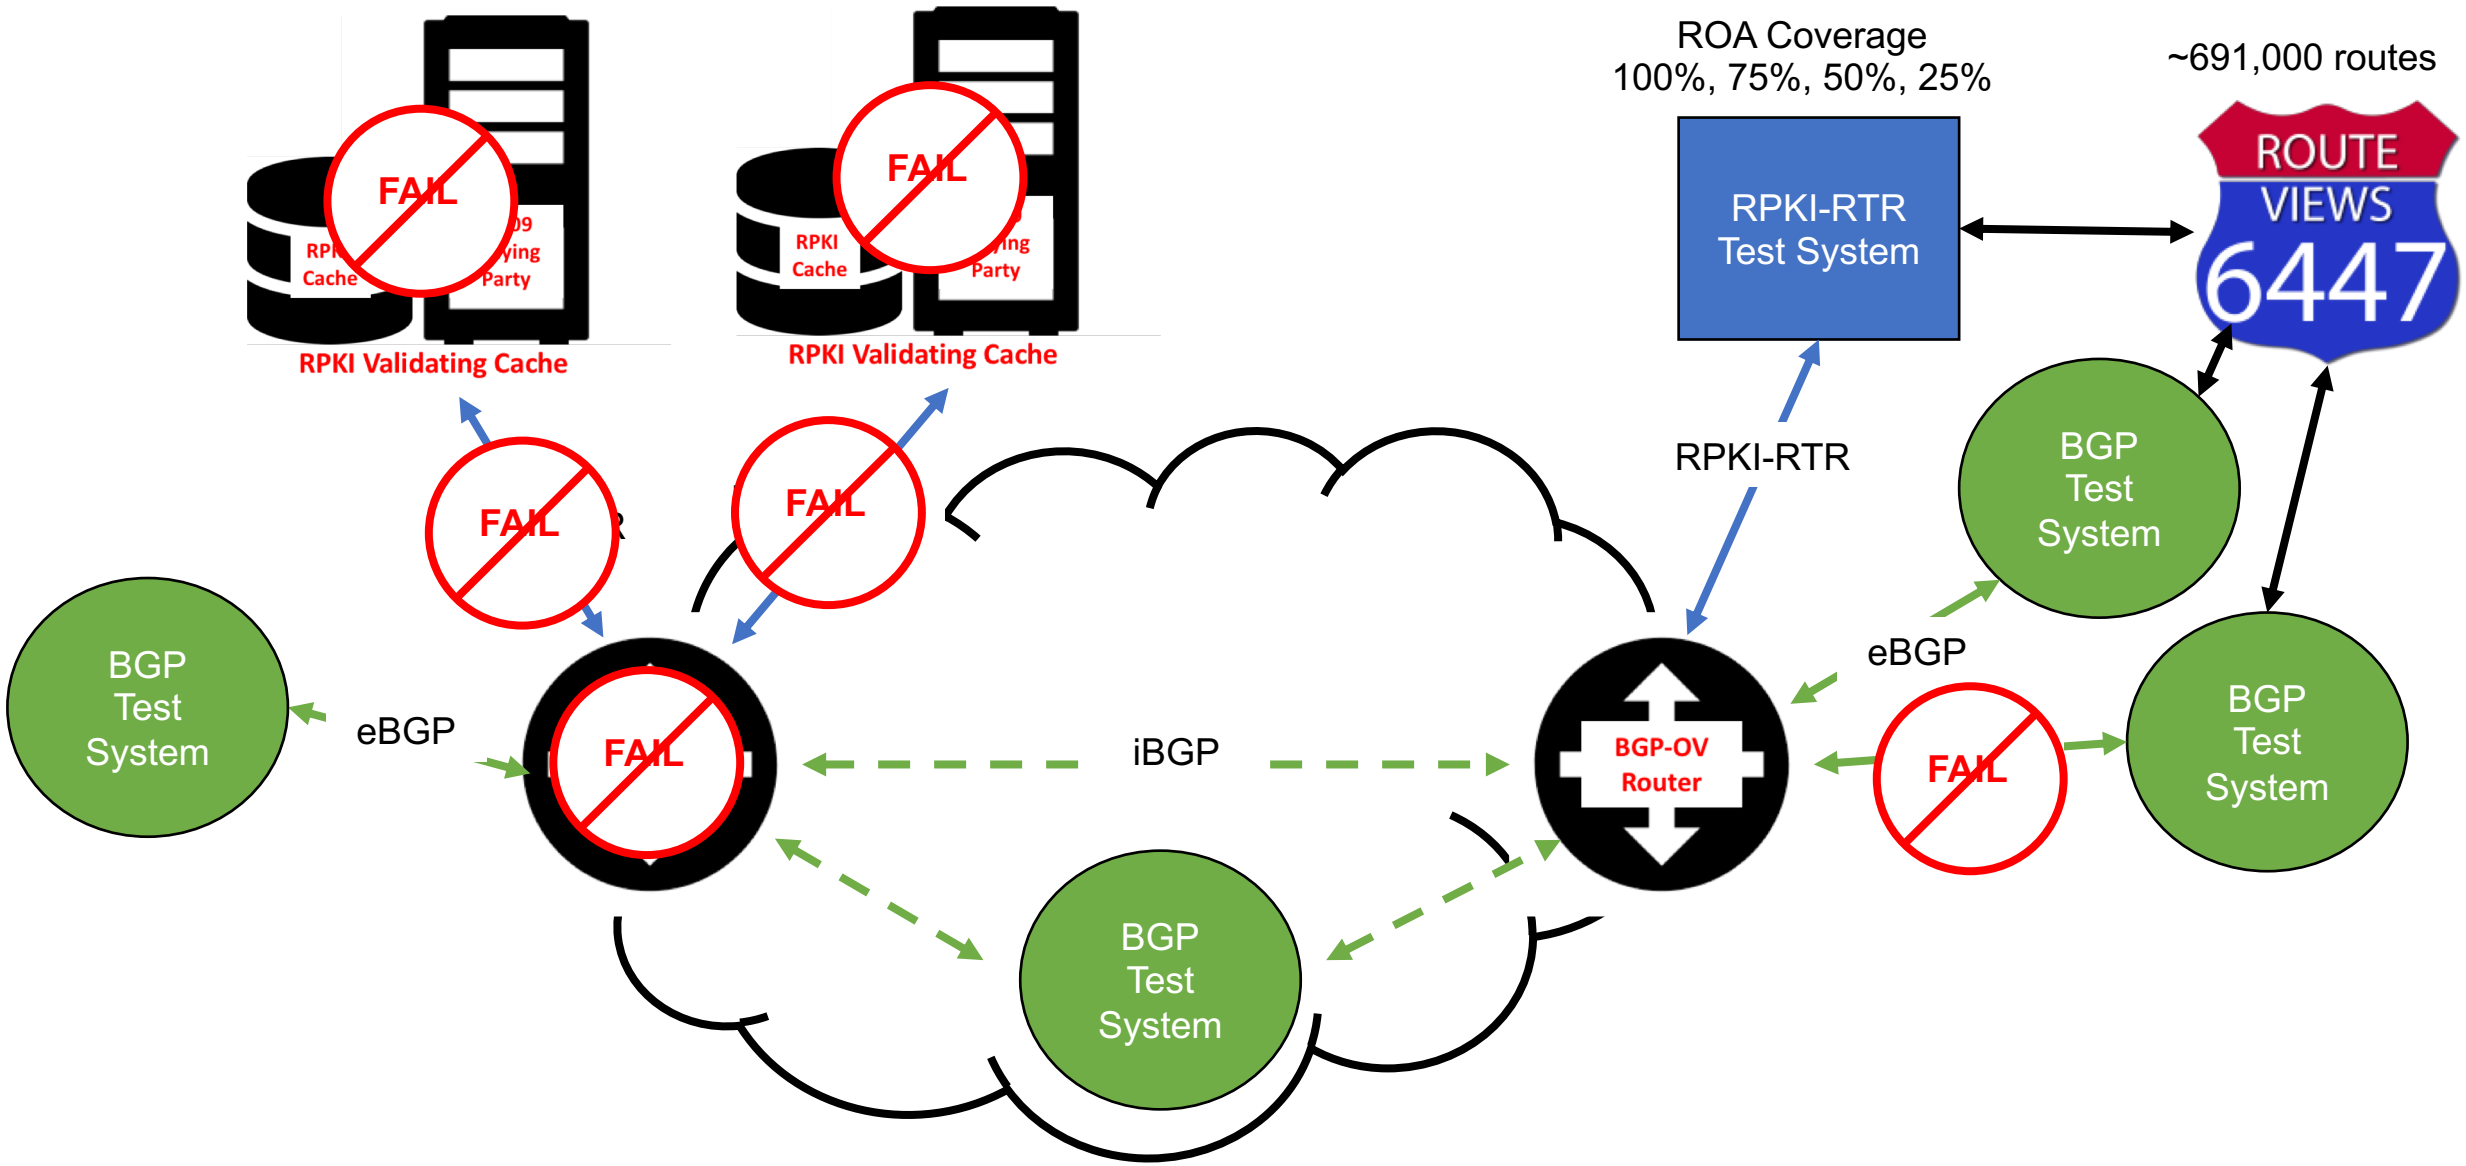

NCCoE Demonstration Project

- **Issues ….**
  - Not all implementations performed ROV on local routes.
    - Some treated all such routes as VALID.
  - Some implementations favored ROV status in community string over local validation.
  - ROV validation signaling can impact prefix packing in some rare scenarios.
  - Routers vary in when they react to the loss of a validation cache.
- **Scaling**
  - Measured CPU, memory, convergence time impacts of large VRP lists.
  - ~700,000 routes. ROA coverage of 100%, 75%, 50%, 25%.
    - i.e., at 100% RPKI-RTR delivery of ~700K VRPs.
  - Impact on convergence time on session reset of full routing table.
    - Consistently between 4-7% slower convergence.
    - 16-20% increase in RAM used. ~150MB.
    - 1.4% increase in CPU utilization.

RPKI Origin Validation Trends

- **Innovative Tools & Techniques**
  - AMS-IX Implementation
    - https://www.ripe.net/support/training/ripe-ncc-educa/presentations/use-cases-stavros-konstantaras.pdf
  - Cloudflare Architecture & Tools
    - https://blog.cloudflare.com/rpki-details/

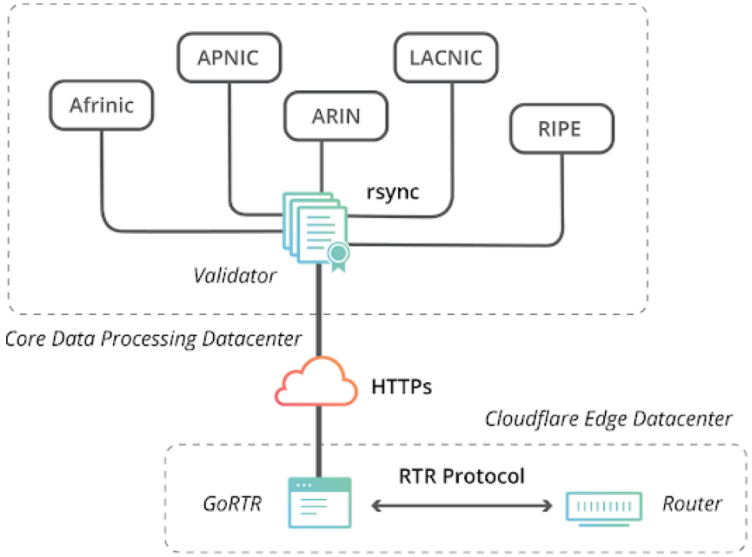

- **Ongoing BGP-OV Specifications**
  - Autonomous System Provider Authorization
    - https://datatracker.ietf.org/doc/draft-azimov-sidrops-aspa-profile/
    - https://datatracker.ietf.org/doc/draft-azimov-sidrops-aspa-verification/
  - Customer Cone AS-Sets
    - https://ripe76.ripe.net/presentations/3-RIPE76-ASCones.pdf
    - https://datatracker.ietf.org/doc/draft-ss-grow-rpki-as-cones/
  - Drop Invalid if Still Routable
    - https://datatracker.ietf.org/doc/draft-sriram-sidrops-drop-invalid-policy/

Questions and Discussion ?

- **For more information:**
  - NIST Robust Inter-Domain Routing Project
    - https://www.nist.gov/programs-projects/robust-inter-domain-routing
  - NIST National Cybersecurity Center of Excellence (NCCoE)
    - https://www.nccoe.nist.gov/

- **DISCLAIMER** Certain commercial entities, equipment, products, or materials may be identified in this document in order to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement by NIST or NCCoE, nor is it intended to imply that the entities, equipment, products, or materials are necessarily the best available for the purpose.