Overview of Routing Security Landscape

for the Quilt Member Meeting, Winter 2019

Mark Beadles, CISO, OARnet mbeadles@oar.net

BGP

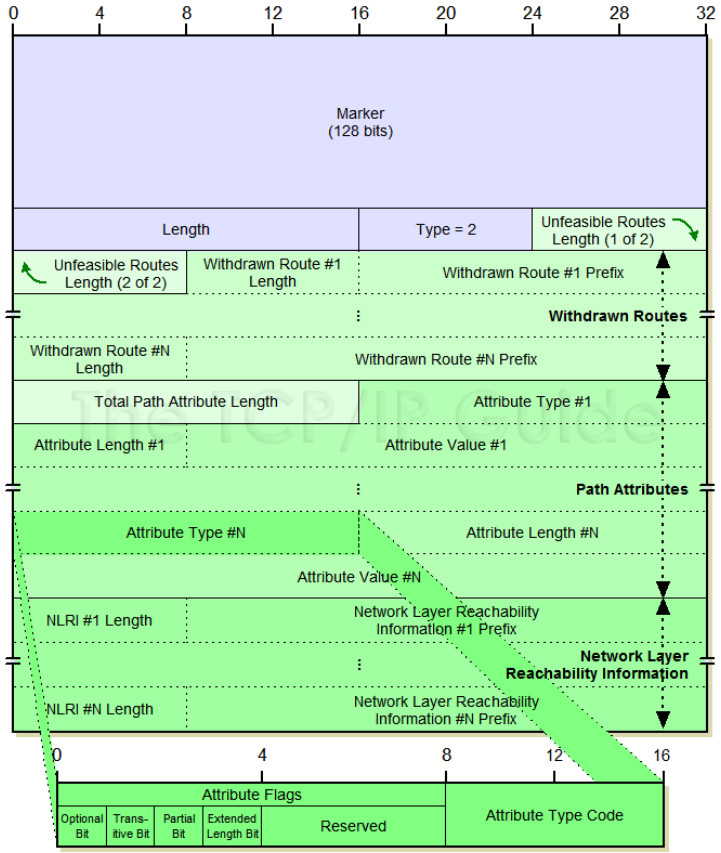
0
4
8
12
16
20
24
28
32
Marker
(128 bits)
Length
Type = 2
Unfeasible Routes Length (1 of 2)
Unfeasible Routes Length (2 of 2)
Withdrawn Route #1 Length
Withdrawn Route #1 Prefix
Withdrawn Routes
Withdrawn Route #N Length
Withdrawn Route #N Prefix
Total Path Attribute Length
Attribute Type #1
Attribute Length #1
Attribute Value #1
Path Attributes
Attribute Type #N
Attribute Length #N
Attribute Value #N
NLRI #1 Length
Network Layer Reachability Information #1 Prefix
Network Layer Reachability Information
NLRI #N Length
Network Layer Reachability Information #N Prefix
0
4
8
12
16
Attribute Flags
Optional Bit
Trans-itive Bit
Partial Bit
Extended Length Bit
Reserved
Attribute Type Code

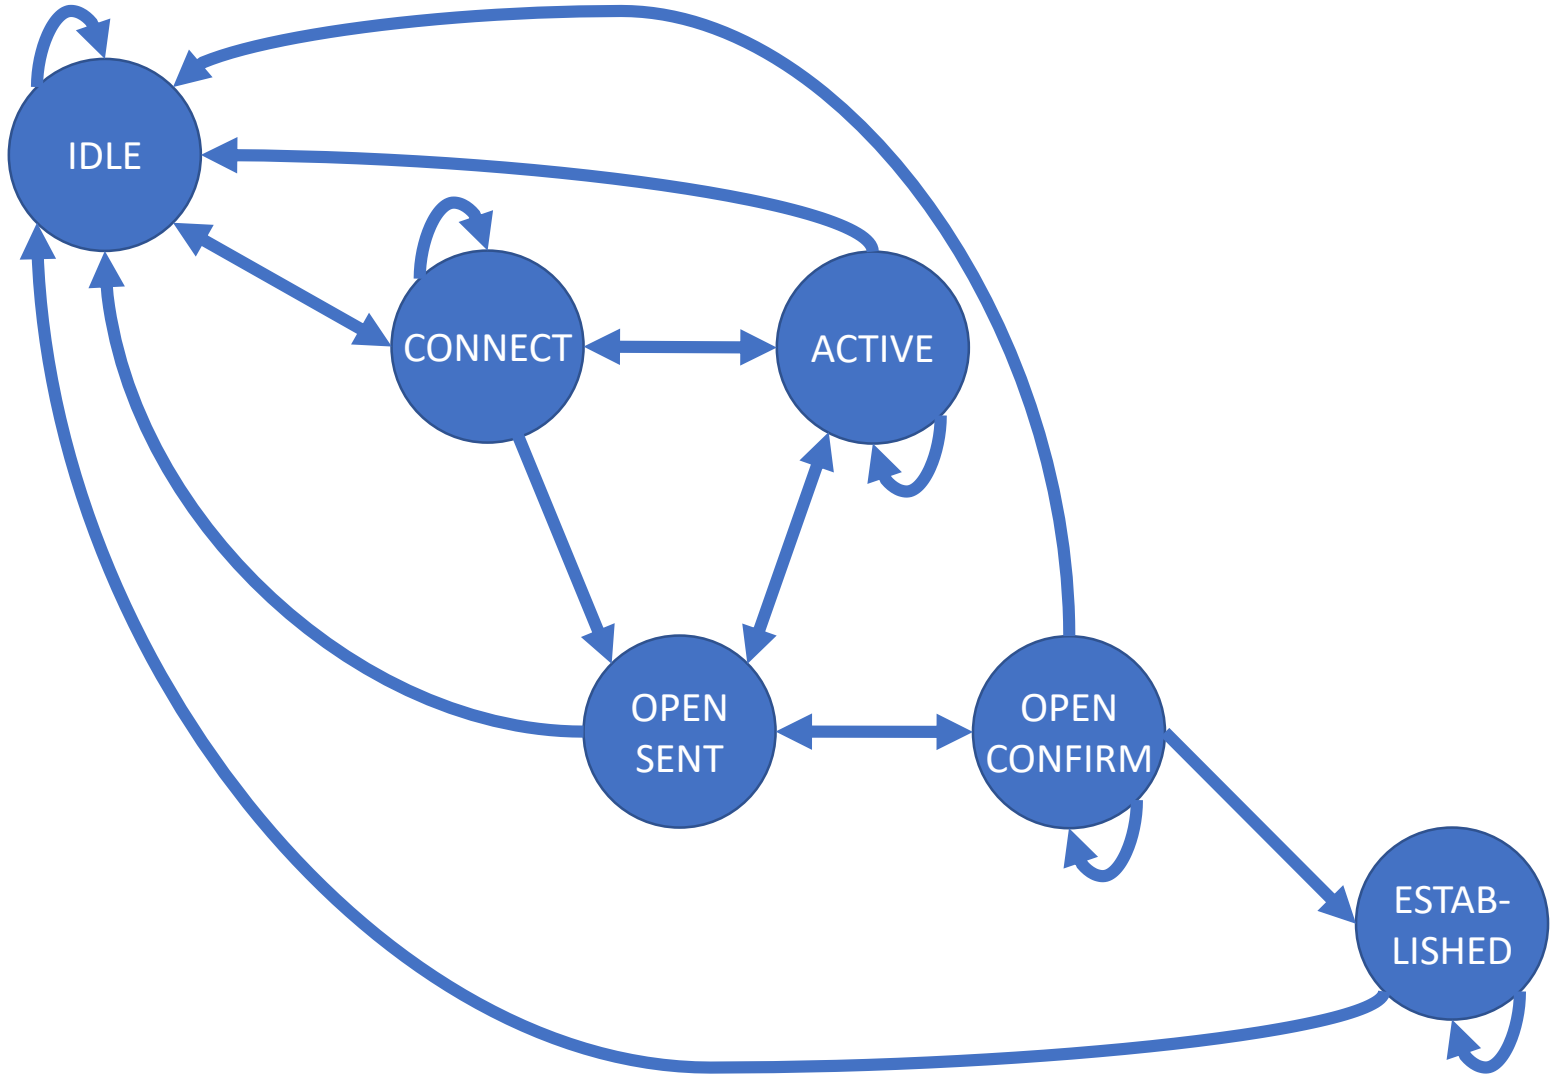
IDLE
CONNECT
ACTIVE
OPEN SENT
OPEN CONFIRM
ESTAB-LISHED

Overview of Routing Security Landscape

- Background
- Threat environment
- Current best practices
- Gaps

Background - Definitions

- BGP
  - Border Gateway Protocol, an exterior path-vector gateway routing protocol
- Autonomous System & Autonomous System Numbers
  - Collection of IP routing prefixes under control of a network operator on behalf of a single administrative domain that presents a defined routing policy to the Internet
  - Assigned number for each AS e.g. AS600

Background - The BGP Security Problem

By design, routers running BGP accept advertised routes from other BGP routers by default. (BGP was written under the assumption that no one would lie about the routes, so there's no process for verifying the published announcements.)

This allows for automatic and decentralized routing of traffic across the Internet, but it also leaves the Internet potentially vulnerable to accidental or malicious disruption, known as BGP hijacking.

Due to the extent to which BGP is embedded in the core systems of the Internet, and the number of different networks operated by many different organizations which collectively make up the Internet, correcting this vulnerabilityis a technically and economically challenging problem.

Background – BGP Terminology

- Bogons
  - Objects (addresses/prefixes/ASNs) that don't belong on the internet
- Spoofing
  - Lying about your address. A major cause of Distributed Denial-of-Service (DDoS) attacks.
- BGP Leak
  - Propagation of routing announcements beyond their intended scope
- BGP Hijack
  - Illegitimate takeover of prefixes by AS's
    - Announcing prefix you don't own
    - Announcing more-specific prefix of someone else's
    - Announcing a shorter route than actually exists

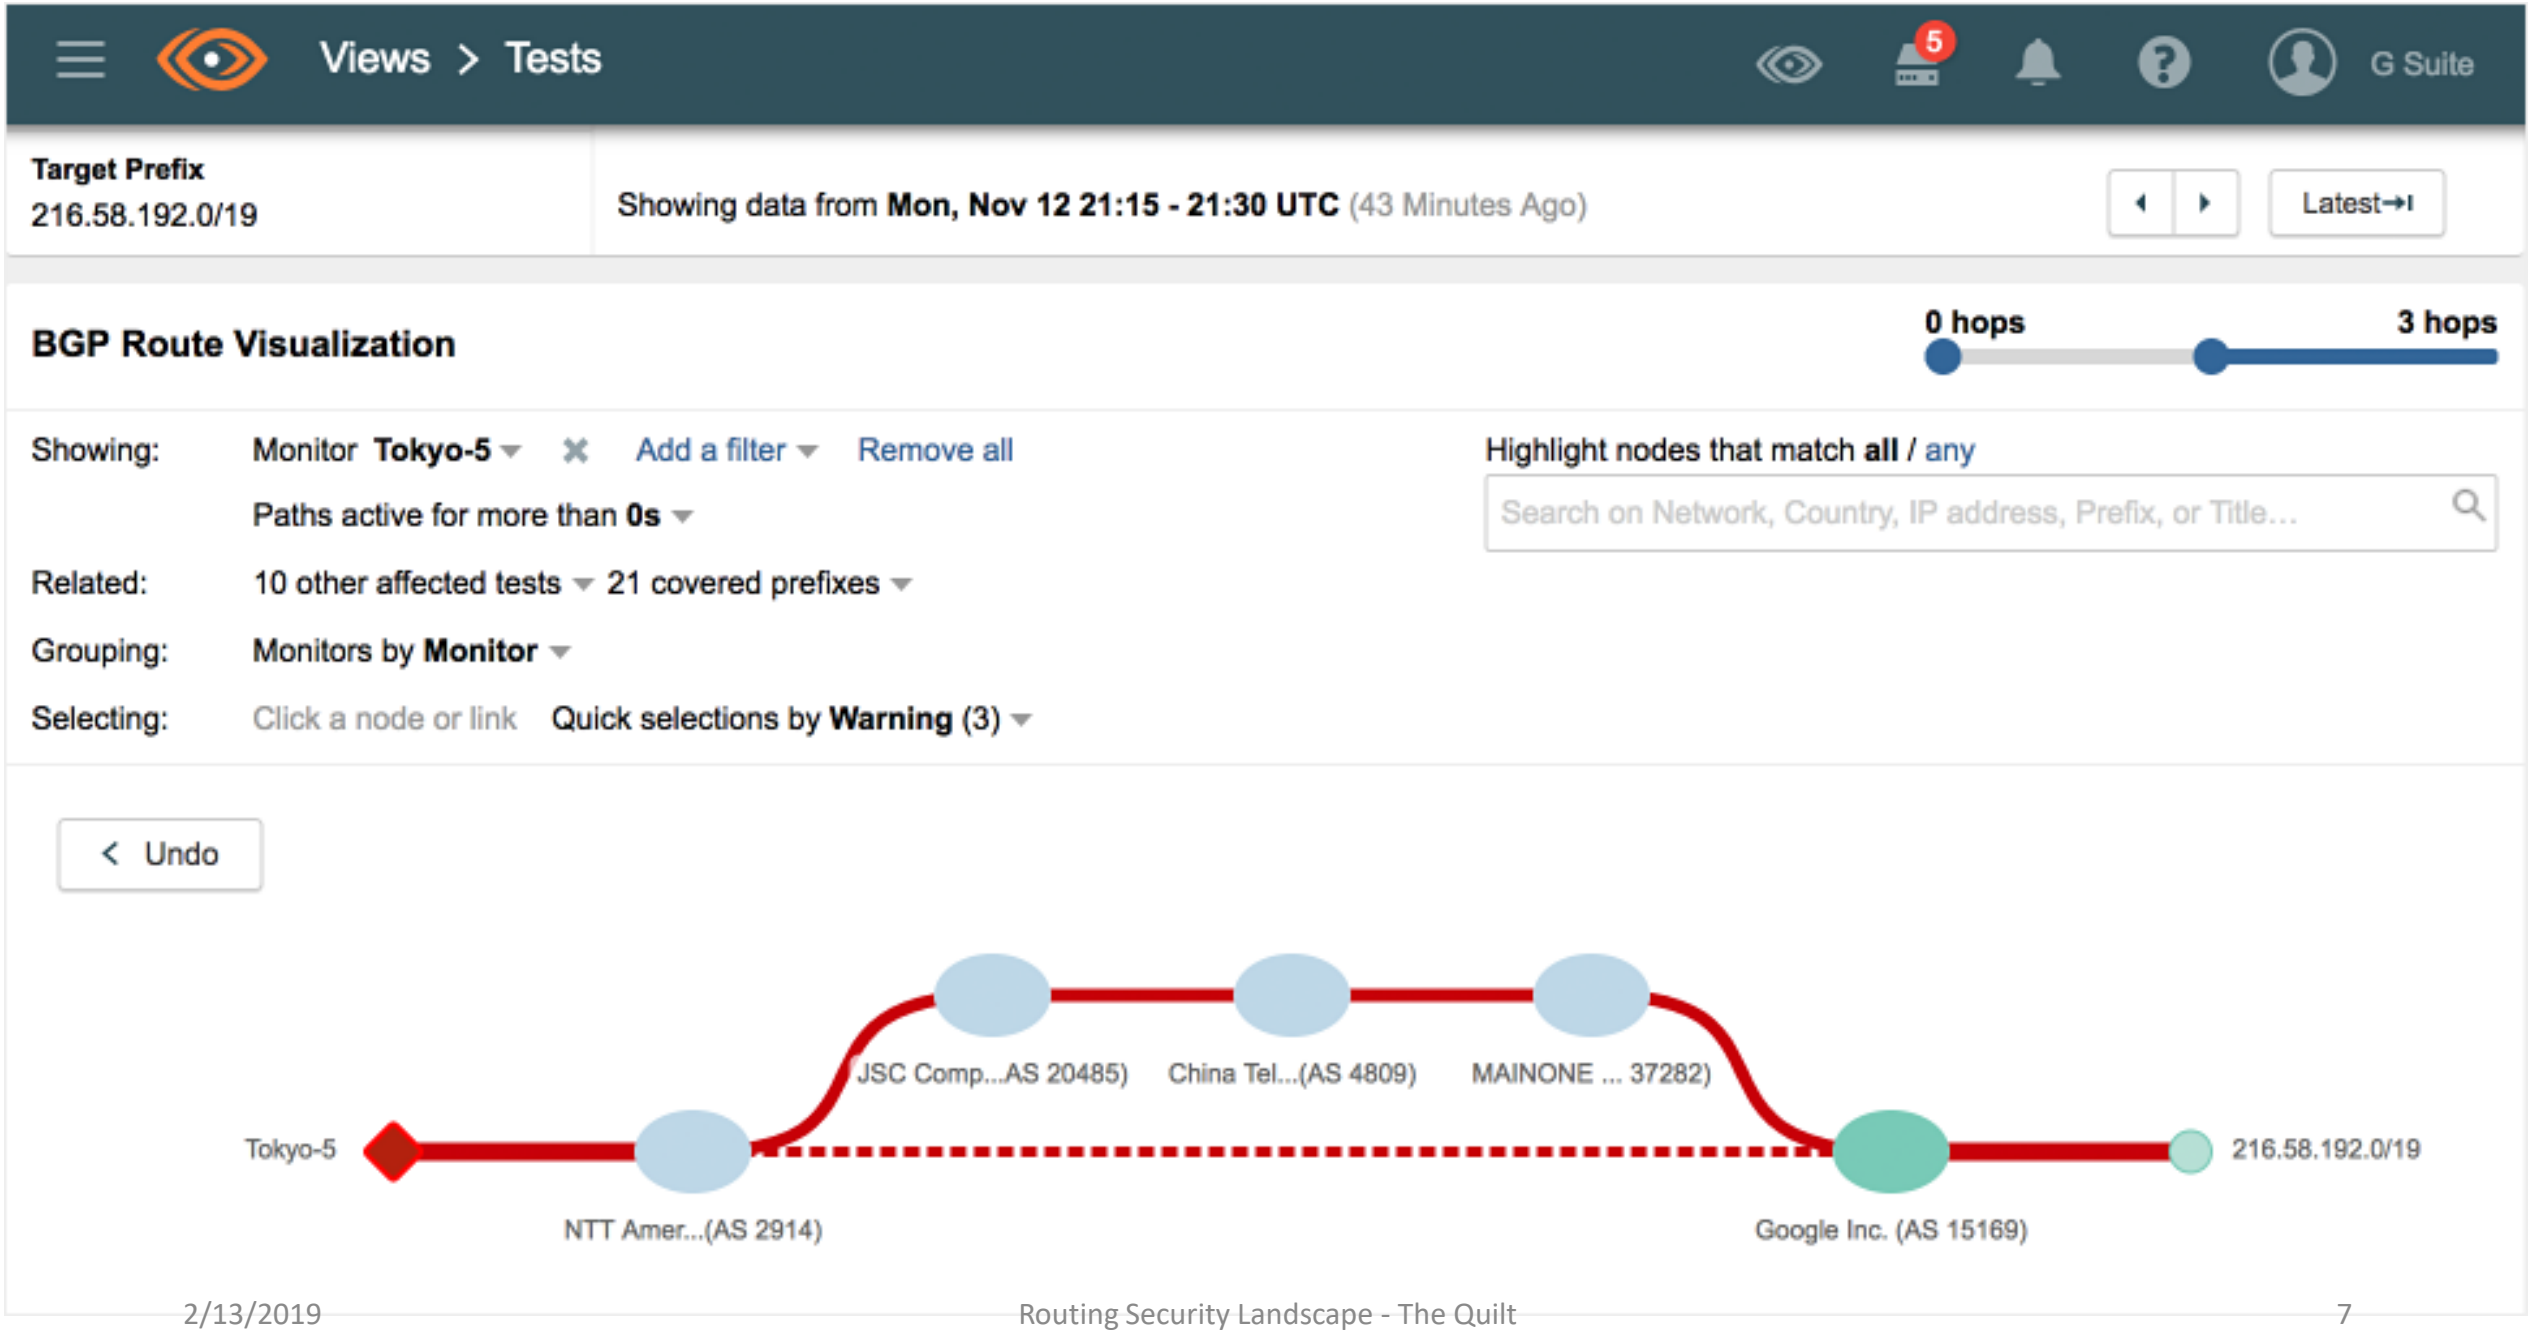
Views > Tests
5
G Suite
Target Prefix
216.58.192.0/19
Showing data from Mon, Nov 12 21:15 - 21:30 UTC (43 Minutes Ago)
Latest
BGP Route Visualization
0 hops
3 hops
Showing:
Monitor Tokyo-5
Add a filter
Remove all
Paths active for more than 0s
Highlight nodes that match all / any
Search on Network, Country, IP address, Prefix, or Title…
Related:
10 other affected tests
21 covered prefixes
Grouping:
Monitors by Monitor
Selecting:
Click a node or link
Quick selections by Warning (3)
Undo
JSC Comp...AS 20485)
China Tel...(AS 4809)
MAINONE ... 37282)
Tokyo-5
216.58.192.0/19
NTT Amer...(AS 2914)
Google Inc. (AS 15169)

Background - BGP Attack Simulation

- https://www.youtube.com/watch?time_continue=5&v=WgX9LiIDaVE

Threat Landscape – Scale of the problem

- Scale:
  - About 90,000 allocated ASN's
    - https://www-public.imtbs-tsp.eu/~maigron/RIR_Stats/RIR_Delegations/World/ASN-ByNb.html
  - About 63,000 active ASN's
    - https://www.cidr-report.org/as2.0/
  - About 750,000 prefixes in the global routing table
    - Largest AS by prefixes has 5000+ prefixes (AS8151 UNINET Mexico)
    - Largest AS by addresses has 115,000,000+ addresses (AS4134 CHINANET backbone)

Threat Landscape – History of Notable Incidents

- April 1997: The "AS 7007 incident"
- December 24, 2004: TTNet in Turkey hijacks the Internet
- May 7, 2005: Google's May 2005 Outage
- January 22, 2006: Con-Edison hijacks big chunk of the Internet
- February 24, 2008: Pakistan's attempt to block YouTube access within their country takes down YouTube entirely.
- November 11, 2008: The Brazilian ISP CTBC - Companhia de Telecomunicações do Brasil Central leaked their internal table into the global BGP table. It lasts over 5 minutes. Although, it was detected by a RIPE route server and then it was not propagated, affecting practically only their own ISP customers and few others.
- April 8, 2010: Chinese ISP hijacks the Internet - China Telecom originated 37,000 prefixes not belonging to them in 15 minutes, causing massive outage of services globally.
- July 2013: The Hacking Team aided Raggruppamento Operativo Speciale (ROS - Special Operations Group of the Italian National Military police) in regaining access to Remote Access Tool (RAT) clients after they abruptly lost access to one of their control servers when the Santrex IPv4 prefix *46.166.163.0/24* became permanently unreachable. ROS and the Hacking Team worked with the Italian network operator Aruba S.p.A. (AS31034) to get the prefix announced in BGP in order to regain access to the control server.
- February, 2014: Canadian ISP used to redirect data from ISPs. - In 22 incidents between February and May a hacker redirected traffic for roughly 30 seconds each session. Bitcoin and other crypto-currency mining operations were targeted and currency was stolen.
- January 2017: Iranian pornography censorship.
- April 2017: Russian telecommunication company Rostelecom (AS12389) originated 50 prefixes for numerous other Autonomous Systems. The hijacked prefixes belonged to financial institutions (most notably MasterCard and Visa), other telecom companies, and a variety of other organizations.
- December 2017: Eighty high-traffic prefixes normally announced by Google, Apple, Facebook, Microsoft, Twitch, NTT Communications, Riot Games, and others, were announced by a Russian AS, DV-LINK-AS (AS39523).
- April 2018: Roughly 1300 IP addresses within Amazon Web Services space, dedicated to Amazon Route 53, were hijacked by eNet (or a customer thereof), an ISP in Columbus, Ohio. Several peering partners, such as Hurricane Electric, blindly propagated the announcements July 2018: Iran Telecommunication Company (AS58224) originated 10 prefixes of Telegram Messenger.[21]
- November 2018: US-based China Telecom site originated Google addresses.

Threat Landscape - Trends

- Categories of incidents
  - Human error
  - Denial of service
  - State actors implementing censorship or economic manipulation
  - Theft (From small scale up to organized crime)
    - 1. BGP Hijack traffic aimed for DNS servers
    - 2. Fake web certificate man-in-the-middle
    - 3. Profit?

Current Best Practices

- RPKI
  - Resource Public Key Infrastructure
  - https://www.arin.net/resources/rpki/
- MANRS
  - Mutually Agreed Norms for Routing Security (ISOC)
  - https://www.manrs.org/isps/
- NSF
  - Draft SP 1800-14, "route origin validation (ROV) by using RPKI in a manner that mitigates some misconfigurations and malicious attacks associated with route hijacking"
    - https://www.nccoe.nist.gov/projects/building-blocks/secure-inter-domain-routing
  - Draft SP 800-189, "Secure Interdomain Traffic Exchange: BGP Robustness and DDoS Mitigation"
    - https://csrc.nist.gov/publications/detail/sp/800-189/draft

MANRS (Mutually Agreed Norms for Routing Security)

1. Prevent propagation of incorrect routing information.
2. Prevent traffic with spoofed source IP addresses.
3. Facilitate global operational communication and coordination between network operators.
4. Facilitate validation of routing information on a global scale.

MANRS

1. Prevent propagation of incorrect routing information.
    1. Use Internet Routing Registries (IRRs) and require the customers to register route objects
    2. Use RPKI and require customers to create Route Origin Authorizations (ROAs)
    3. Use an internal database with the information provided as part of the provisioning process
2. Prevent traffic with spoofed source IP addresses.
    1. uRPF (strict/loose/feasible)
    2. Dynamic user ACLs (RADIUS)
    3. IETF SAVI (DHCP snooping)
    4. Source Verification (vendor-specific)
    5. Explicit ACL's at PE-CE boundary
3. Facilitate global operational communication and coordination between network operators.
    1. Maintaining Contact Information in Regional Internet Registries (RIRs)
    2. IRR's
    3. PeeringDB https://www.peeringdb.com
    4. Operator's own web site
4. Facilitate validation of routing information on a global scale.
    1. Valid origin information in IRR's
    2. Cryptographic validation using RPKI

MANRS

- 132 current network operator participants
  - https://www.manrs.org/isps/participants/

RPKI

- **Resource certificates**: These certificates digitally verify that a resource has been allocated or assigned to a specific entity
- **Route Origin Authorizations (ROAs)**: Digital statements specifying which Autonomous System may originate a specific IP address or range
- **Trust Anchor Locator (TAL):** File used to allow relying parties to retrieve the data within ARIN's RPKI validator (via rsync) and base routing decisions upon that data. ARIN's TAL contains two things: The URL of ARIN's published RPKI repository, and ARIN's PEM-encoded public key.
- Two aspects to adoption:
  - Using RPKI as a Relying Party: Obtaining information about routes and using RPKI as a relying party (to make routing decisions for your network). You need to download the ARIN Trust Anchor Locator (TAL) and use it with an RPKI validator.
  - Providing Certification for Your Resources: Certify that you have authority over routes that originate from your resources by creating certificates and Route Origin Authorizations (ROAs).

Tools…so many tools…

- IRR's and RPSL
  - http://www.irr.net/
  - ARIN, RIPE, APNIC, AFRINIC, LEVEL3, RADB...
- PeeringDB
  - https://www.peeringdb.com/
- RADB
  - https://www.radb.net/
- Hurricane Electric BGP Toolkit https://bgp.he.net/
- Looking Glass servers… https://tools.keycdn.com/bgp-looking-glass
  - e.g. http://merry.netsys.more.net/lg/index.cgi
- CAIDA (Center for Applied Internet Data Analysis)
  - http://www.caida.org/, AS-RANK (http://as-rank.caida.org/)
- University of Oregon Route Views & BGPPlay
  - http://www.routeviews.org/, http://bgplay.routeviews.org/
- Qrator
  - https://radar.qrator.net
- ARIN RPKI
  - https://www.arin.net/resources/rpki/
- Routinator (RPKI Validator)
  - https://www.nlnetlabs.nl/projects/rpki/routinator/

Gaps & Unknowns

- Adoption
  - Network operators – current programs are opt-in, enforcement is lacking, percentage-wise adoption is low
  - Customers – many or even most customers take no active role in managing routing
  - Benefits may not accrue until majority of networks adopt the measures
- Disparate incentives
  - AS deploying a prefix filter does not have particularly strong incentives to do so, other than protecting *the rest of the Internet* from attacks by its *own customers*
  - Choosing economic factors over security factors
- Availability threats
  - Adding new operational elements (e.g. RPKI) increases points of failure
  - Adding new computational requirements impacts performance (e.g., BGPSEC)
  - Lack of proven, diverse operational software implementations
- Staleness
  - Routing-related data stores are highly distributed, informal, and poorly maintained by customers
- Incomplete solutions
  - RPKI does not prevent all kinds of attacks (route leaks, announcing short bogus paths)
  - RPKI validates origins, not paths. BGPSEC is an attempt at cryptographic path validation.
- Impact of growth of “cloud” & “content” providers
- Reference: https://queue.acm.org/detail.cfm?id=2668966 (Sharon Goldberg, BU)

Questions / Discussion (and please join MANRS)