EPOC
Engagement and Performance Operations Center

BGP Security

Ken Miller, Jason Zurawski
ken@es.net, zurawski@es.net
ESnet / Lawrence Berkeley National Laboratory

Modern Cyberinfrastructure for Research Data Management Workshop
University of Central Florida
February 16-17, 2023

TACC TEXAS ADVANCED COMPUTING CENTER
ENERGY SCIENCES NETWORK
INDIANA UNIVERSITY

National Science Foundation Award #1826994

BGP is an OLD protocol

- Has been in use since 1994
  - https://datatracker.ietf.org/doc/html/rfc1654
- Security was not a concern and not baked into the protocol
- Believes (without help) all advertisements from peers with no checks.
- It also by default can re-advertise to other peers what it learns.

Hijacking, Leaking, and spoofing…

- MANRS reports over 10,000 routing outages or attacks in 2018*
- 40% of all incidents believed to be attacks.
- Incidents can quickly scale to global problems.

EPOC
Engagement and Performance Operations Center

*https://www.manrs.org/2019/02/routing-security-getting-better-but-no-reason-to-rest/

Route / Prefix Hijacking

- When a network advertises/originates a route that belongs to another network (without permission)
- Not always malicious can easily be caused by misconfiguration

https://www.manrs.org/2020/09/what-is-bgp-prefix-hijacking-part-1/

Route / Prefix Hijacking - How it works

- AS Path length

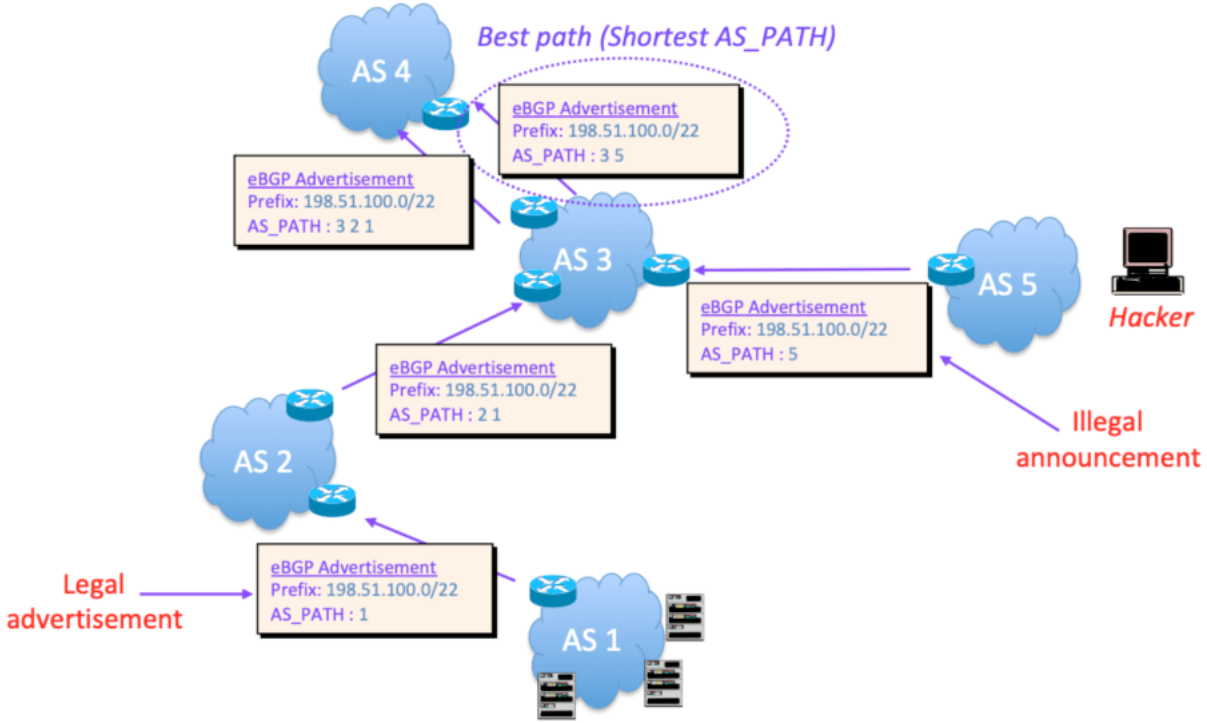

Engagement and Performance Operations Center

https://www.manrs.org/2020/09/what-is-bgp-prefix-hijacking-part-1/

Example: Youtube and Pakistan Telecom

- Before, during and after Sunday, 24 February 2008: AS36561 (YouTube) announces 208.65.152.0/22.
- Sunday, 24 February 2008, 18:47 (UTC): AS17557 (Pakistan Telecom) starts announcing 208.65.153.0/24. AS3491 (PCCW Global) propagates the announcement. Routers around the world receive the announcement, and YouTube traffic is redirected to Pakistan.
- Sunday, 24 February 2008, 20:07 (UTC): YouTube changes to announcing two /24s. Some traffic starts going back to YouTube.

https://www.ripe.net/publications/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study
https://www.cnet.com/culture/how-pakistan-knocked-youtube-offline-and-how-to-make-sure-it-never-happens-again/

EPOC
Engagement and Performance Operations Center

Example: Youtube and Pakistan Telecom 2

- Sunday, 24 February 2008, 20:18 (UTC): AS36561 (YouTube) starts announcing 208.65.153.128/25 and 208.65.153.0/25. Because of the longest prefix match rule, every router that receives these announcements will send the traffic to YouTube.
- Sunday, 24 February 2008, 20:51 (UTC): All prefix announcements originated by AS17557 (Pakistan Telecom) via AS3491 (PCCW Global), are prepended by another 17557. The longer AS path means that more routers prefer the announcement originated by YouTube.
- Sunday, 24 February 2008, 21:01 (UTC): AS3491 (PCCW Global) withdraws all prefixes originated by AS17557 (Pakistan Telecom), thus stopping the hijack of 208.65.153.0/24.

EPOC
Engagement and Performance Operations Center

Other Hijacking examples

- 2018: Amazon DNS routes hijacked and redirected to malicious DNS server: https://www.internetsociety.org/blog/2018/04/amazons-route-53-bgp-hijack/
- 2020: Rostelecom hijacks internet traffic for Google, AWS, Cloudflare, and others: https://www.zdnet.com/article/russian-telco-hijacks-internet-traffic-for-google-aws-cloudflare-and-others/

Resource Public Key Infrastructure (RPKI) to the rescue (maybe?)

- Regional Internet Registries (RIR's) certifies owners of AS numbers and IP addresses.
- They also certify route announcements
  - Route Origin Authorization (ROAs) show that you are authorized to advertise the IP addresses
- Allows you to verify addresses advertised to your router are authorized to be advertised by that entity
- Router can set the route as Valid, Invalid, or unknown
- Create route policy depending on those results
- Allows reject on wrong AS, wrong prefix, or too specific advertisement

EPOC
Engagement and Performance Operations Center

https://www.noction.com/blog/rpki-overview

Route Leak

- RFC7908 - “A route leak is the propagation of routing announcement(s) beyond their intended scope.”
- A multihomed stub network announces routes from one upstream providers routes to one or more of its other upstream providers
- Stub network becomes an inadvertent transit provider.
- Only announce AS’s and prefixes that you originate.

EPOC
Engagement and Performance Operations Center

https://datatracker.ietf.org/doc/html/rfc7908

Simple Campus/Institution Route Leak Example

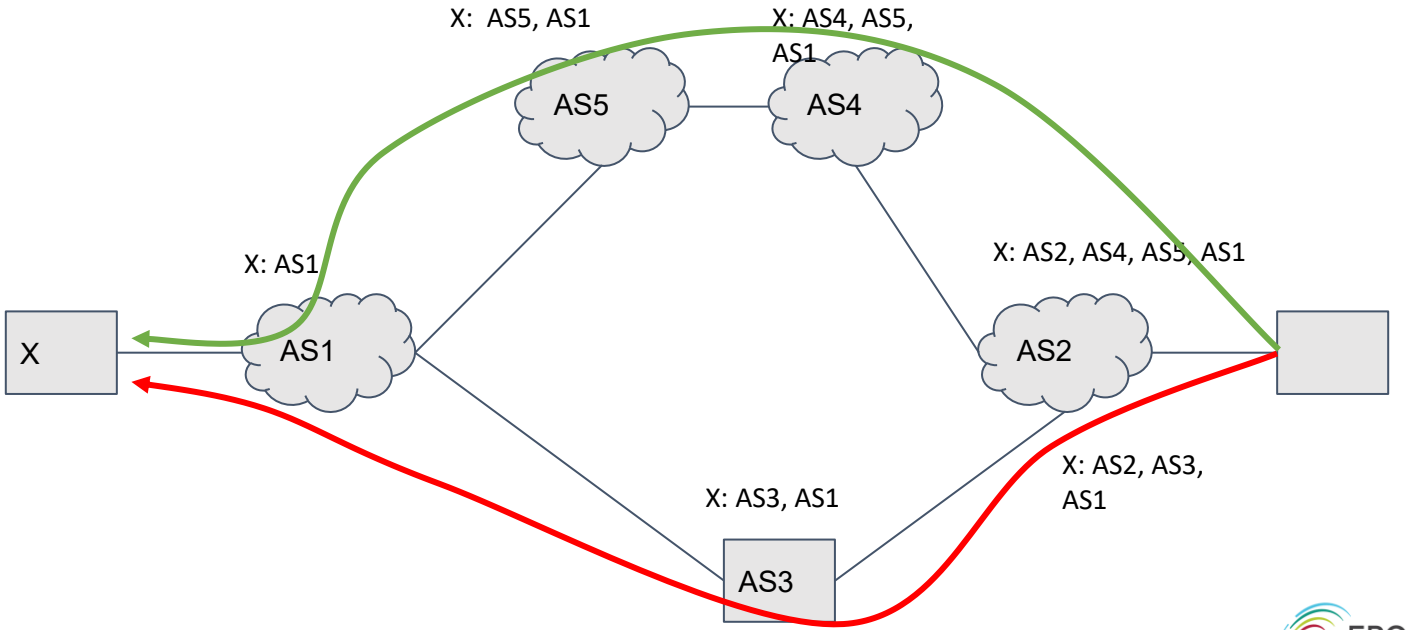

Stub network AS3 creates route leak advertising AS1 to AS2.

EPOC
Engagement and Performance Operations Center

Route Leak Example

- 2017: Rostelecom Route Leak Targets E-Commerce Services: https://www.thousandeyes.com/blog/rostelecom-route-leak-targets-ecommerce-services
  - Confirmation that traffic destined for those E-Commerce sites went through the leakers network (possible inspection?)

Route Policy to fix Leaks - Overview

- BGP Operations and Security RFC: https://datatracker.ietf.org/doc/html/rfc7454
  - Includes lots of great best practices for AS and prefix filtering
- Good Primer: https://www.noction.com/wp-content/uploads/2019/08/BGP-Filtering-Best-Practices.pdf

Route Policy to fix Leaks - Inbound

- Loose Inbound Filtering Highlights include:
  - Don’t accept your own prefixes from a peer.
  - Filter Bogons (Addresses not assigned)
    - IPv4 not so much anymore but IPv6 YES
- Be careful of more specific prefixes
  - IPv4: more specific than a /24
  - IPv6: more specific than a /48
- Strict Filtering: use scripts or tool to validate incoming prefixes against route registries.
  - https://www.irr.net/

Route Policy to fix Leaks - Outbound

- If you are a multihomed only advertise what you originate.
- Don’t advertise private space (RFC1918)
- Prefixes used on your internal networks
- Default route

EPOC
Engagement and Performance Operations Center

IP Spoofing

- Attacker creates and send IP packets with false source address
- Commonly used in Distributed Denial of Service (DDOS) attacks
  - DNS, memcached, NTP, UDP - lots of vulnerabilities

- November 2021: Microsoft detects and mitigates a 3.47Tbps (340 million packets per second) 15 minute long DDOS attack using UDP reflection.

EPOC
Engagement and Performance Operations Center

https://azure.microsoft.com/en-us/blog/azure-ddos-protection-2021-q3-and-q4-ddos-attack-

Source Address Validation and IP Spoofing

- Unicast Reverse Path Forwarding (uRPF)
  - Router checks it's forwarding information table (FIB) for source address in each packet.
    - Strict: Source Address must be reachable via incoming interface (strict) or in the FIB (loose) or packet is dropped.
- Can be done with ACL's as well but can require a lot of manual configuration.
- Best Current Practices (BCP) 38
  - http://www.bcp38.info/index.php/Main_Page
  - https://datatracker.ietf.org/doc/html/rfc2827

EPOC
Engagement and Performance Operations Center

https://learn.nsrc.org/bgp/urpf

https://blog.apnic.net/2022/02/07/source-address-validation-use-cases-and-gap-

BGPSec

- RPKI doesn't validate the entire ASPATH of a prefix.
- BGPSec intended to verify the full path.
- https://datatracker.ietf.org/doc/html/rfc8205 and more
- IETF working groups moving forward (https://datatracker.ietf.org/wg/sidrops/about/)
- No commercial implementations yet.
- few open source projects (https://github.com/usnistgov/NIST-BGP-SRx)

Takeaways!

Routing will not take care of itself

- Old routes may not work well with new networks
- New routes may not work as planned

How do we address routing anomalies as a community?

The Routing Working Group!

EPOC
Engagement and Performance Operations Center

How we used to solve this:

I think I’ve found a bad route. How do I address this?

Start a conversation! “Hello, it appears that there’s a lot of traffic going between your institution and X that probably shouldn’t

Can we work together to improve it?"

EPOC
Engagement and Performance Operations Center

What do we do now?

Routing Working Group case process

1. Cases are submitted via the mailing list, slack or at the monthly meeting
2. Teams are selected to assist with the case by the chairs
3. The case is added to the master case list (open access)
   - RWG Master Case List
4. A folder is created for the specific case and team members are given access to the documentation
   - Most cases are worked on via email or slack
5. Case updates are given at each RWG meeting

EPOC
Engagement and Performance Operations Center

Routing Working Group - What are the goals?

- **Engineering focus**
  - Document possible erroneous routes
  - Identify teams to address them
  - Check in together as we work through them
- **Policy Focus**
  - Detail routing policies for paths
    - Including preferred backup paths!
  - Verify if policy is being followed

EPOC
Engagement and Performance Operations Center

Routing Challenges We've Observed

Asymmetrical routing - meaning a source to a destination takes one path and takes a different path when it returns to the source

R&E data takes a less efficient route around the world - affecting performance

- Europe to Asia routes traversing the US
- Africa to Europe routes traversing the US

R&E data takes a commodity route when an R&E path is available

New R&E links are removed or added but routing does not adjust appropriately

Leaking of Private ASN's into the global routing table by R&E networks

IP blocks advertised with a Bogon Origin ASN's within R&E routing

ITB <> Starlight

Summary: Asymmetrical routing is preventing access the starlight ESnet DTN.

- Route from ITB to ESnet Chicago is using R&E networks via TEIN, TransPAC and Pacific Wave
- Chicago ESnet to ITB is using the PCCW and commercial paths.

Resolution: Worked with ESnet to update accepted routes from TransPAC to use R&E path. Trace routes are symmetrical now, We received confirmation from ITB engineers that they are able to access the DTN.

Team: Simon Peter Green (SingAREN) , Basuki Suhardiman (ITB) , Brenna Meade (IU)

Singapore to New Zealand

Summary: Traffic from Singapore to New Zealand traversing the US, and is asymmetrical

Resolution: Asymmetric routing has been resolved by moving traffic to the Singapore <> Guam link and confirmed via PS tests for both IPv4 and IPv6

Routes before changing the routing:

SingAREN@SG -> SingAREN@LA - > Internet2@LA -> REANNZ

SingAREN@SG -> APAN-JP -> Transpac@SEA -> REANNZ

Current symmetrical routing (using Guam <> Singapore link) :

SingAREN@SG -> GOREX-> REANNZ

Team: Brenna Meade (IU) , Dylan Hall (REANNZ) , Francis Lee (SingAREN), Simon Green (SingAREN)

Taiwan to Indonesia

Perfsonar tests indicate asymmetric routes

Taiwan to Indonesia (16 hops)

ASGCNET <> APAN <> SINET <> TEIN2 <> BANDUNG-NET

Indonesia to Taiwan (11 hops)

BANDUNG-NET <> TIEN2 <> ASGCNET

Resolution: Peering was changed at ASGCNET to prefer BANDUNG-NET

NetSage

Questions?

Transfer Performance problems? EPOC is here to help!

- epoc@tacc.utexas.edu
- https://epoc.global/

NSF Award: 1826994

EPOC
Engagement and Performance Operations Center

EPOC
Engagement and Performance Operations Center

BGP Security

Ken Miller, Jason Zurawski
ken@es.net, zurawski@es.net
ESnet / Lawrence Berkeley National Laboratory

Modern Cyberinfrastructure for Research Data Management Workshop
University of Central Florida
February 16-17, 2023

TACC TEXAS ADVANCED COMPUTING CENTER
ENERGY SCIENCES NETWORK
INDIANA UNIVERSITY

National Science Foundation Award #1826994