EPOC
Engagement and Performance Operations Center

BGP and Routing Security

Hans Addleman, addlema@iu.edu

ESnet
ENERGY SCIENCES NETWORK

INDIANA UNIVERSITY

National Science Foundation Award #1826994

Agenda

BGP and Routing Overview

Attacks against:

- Route Hijacking, leaking, spoofing

RPKI

BGP MD5 passwords

BGPSec

MANRS

Routing Working Group

EPOC
Engagement and Performance Operations Center

BGP in the wild

- Over 74,000 Autonomous Systems (ASN) in March 2022.
- Over 1,000,000 IPv4 routes advertised.
- Over 182,000 IPv6 routes advertised.

- Each Router running BGP builds its own routing table with best path information to a subset of the internet.

EPOC
Engagement and Performance Operations Center

Data from: https://bgp.he.net/report/prefixes#_prefixes

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.

EPOC
Engagement and Performance Operations Center

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.

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

EPOC
Engagement and Performance Operations Center

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/

EPOC
Engagement and Performance Operations Center

Route / Prefix Hijacking - How it works 1

- AS Path length

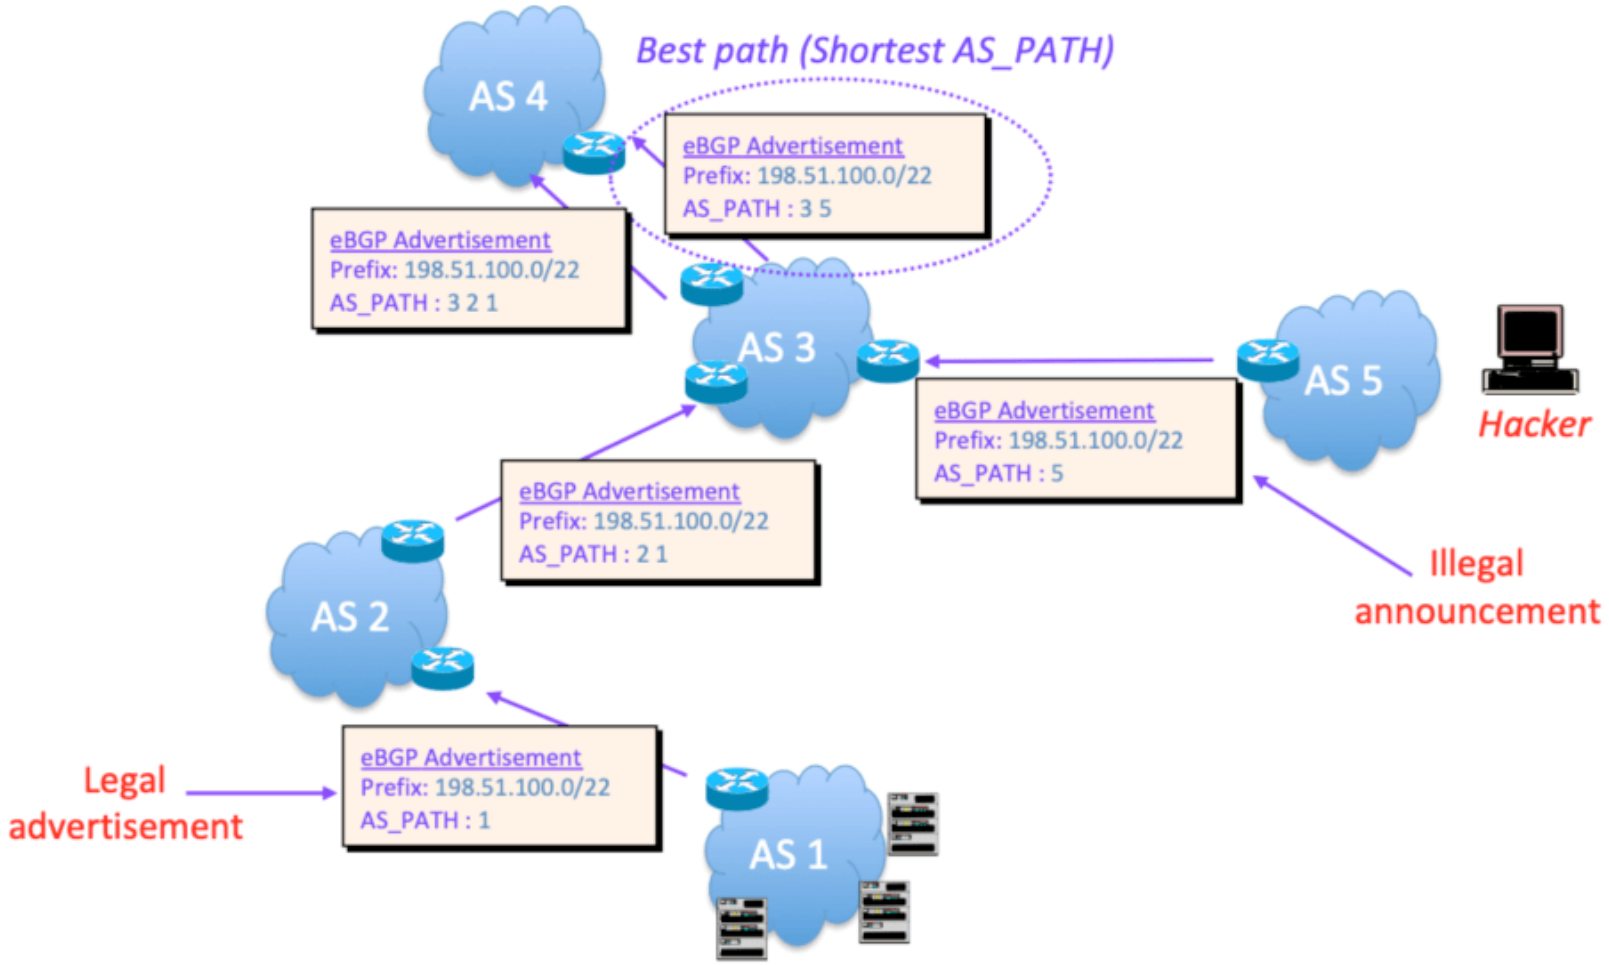

https://www.manrs.org/2020/09/what-is-bgp-prefix-hijacking-part-1/

Route / Prefix Hijacking - How it works 2

- More specific prefix

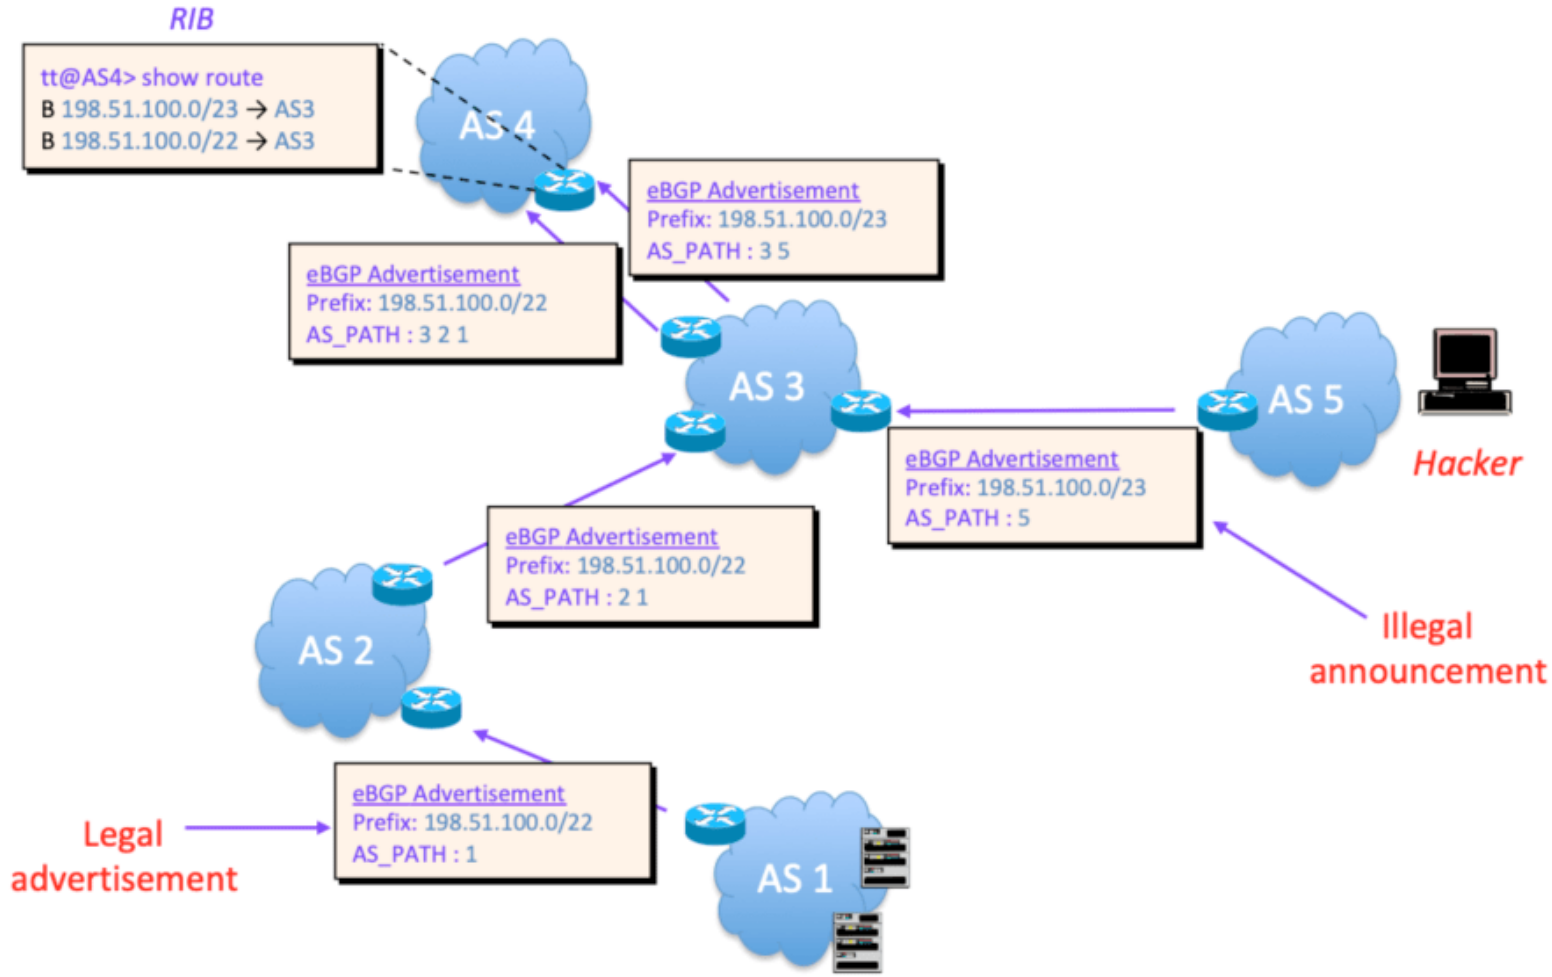

https://www.manrs.org/2020/09/what-is-bgp-prefix-hijacking-part-1/

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/

EPOC
Engagement and Performance Operations Center

Resource Public Key Infrastructure (RPKI) to the rescue

- 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

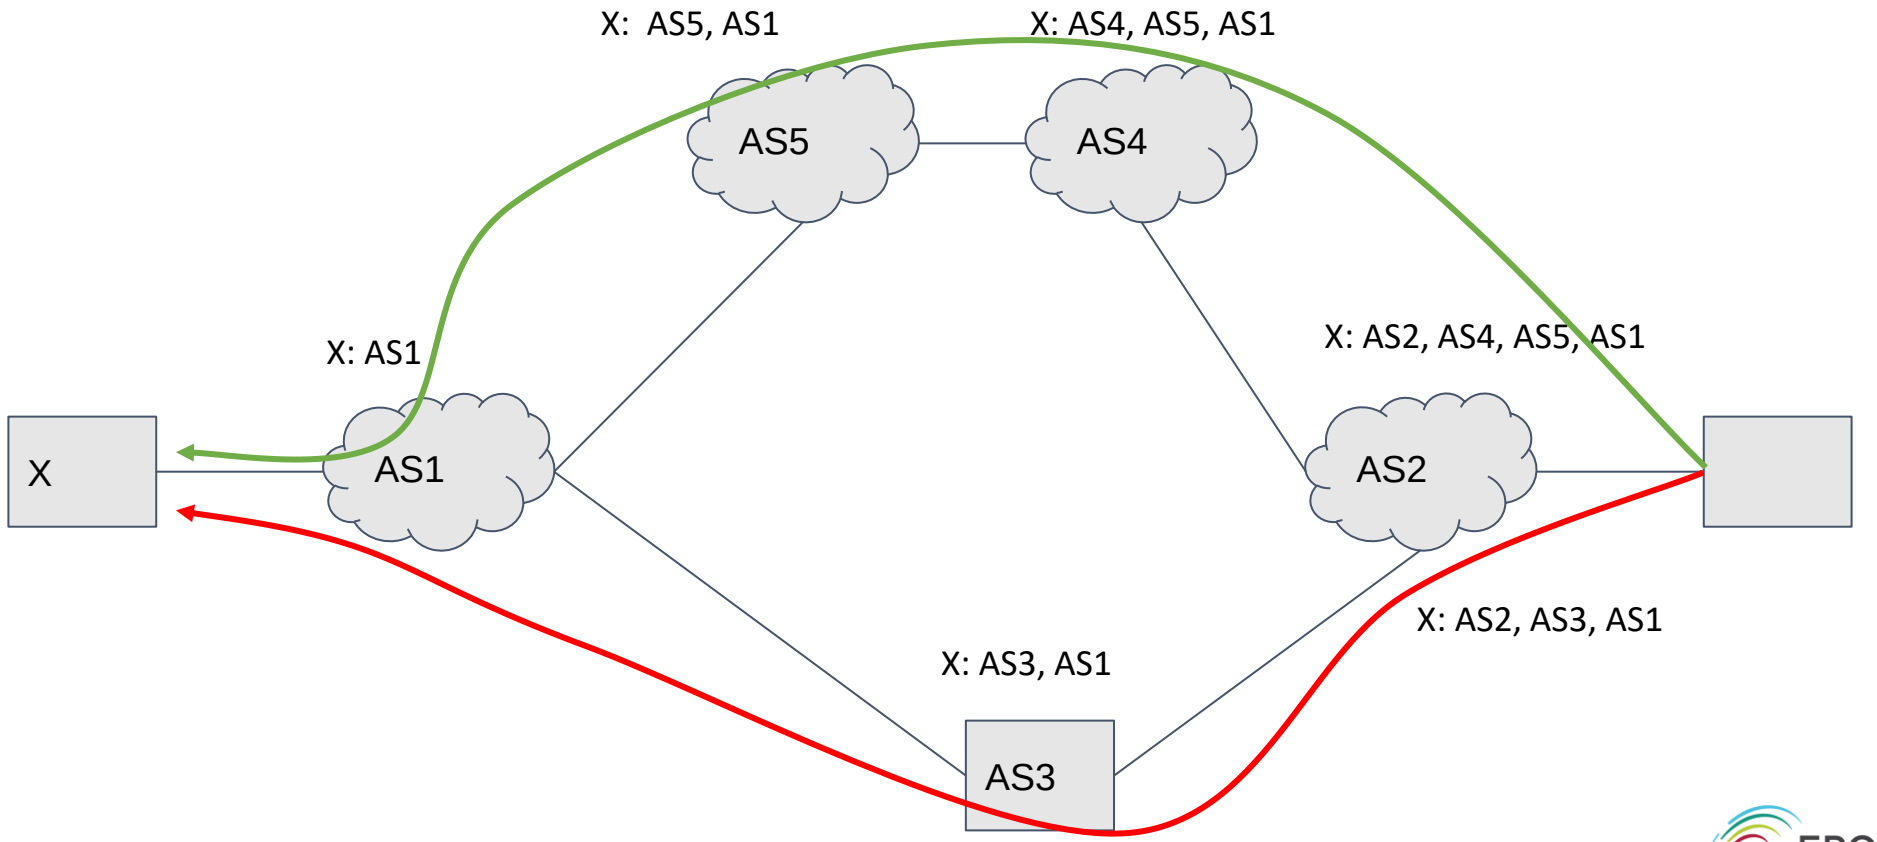

Stub network AS3 creates route leak advertising AS1 to AS2.

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?)

EPOC
Engagement and Performance Operations Center

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

EPOC
Engagement and Performance Operations Center

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/

EPOC
Engagement and Performance Operations Center

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-trends/

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

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

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

EPOC
Engagement and Performance Operations Center

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)

EPOC
Engagement and Performance Operations Center

BGP MD5 / GTSM

- MD5
  - Configure a password shared with your BGP peer
  - Each router checks the MD5 hash. If none or no match the router drops packets
- Generalized TTL Security Mechanism (GTSM)
  - sets time to live (TTL) of BGP packets to 255
  - Receiving router checks for 255 TTL and drops if different
  - This stops attacks from coming from several hops away because each routed hop decrements the TTL
  - NOTE: This does not work for multi-hop BGP.

EPOC
Engagement and Performance Operations Center

https://www.noction.com/blog/bgp-security-md5-password-gtsm

Mutually Agreed Norms for Routing Security (MANRS)

- https://www.manrs.org/
- MANRS improves the security and reliability of the global Internet routing system, based on collaboration among participants and shared responsibility for the Internet infrastructure.
- Programs, training, tutorials, best practices, and more for Network Operators, Exchange Point Operators, and Content Delivery Networks.

EPOC
Engagement and Performance Operations Center

https://www.manrs.org/wp-content/uploads/2020/09/About-MANRS-Presentation-PDF-Sep-2020.pdf

MANRS Actions for Network Operators

Filtering

Prevent propagation of incorrect routing information

Ensure the correctness of your own announcements and announcements from your customers to adjacent networks with prefix and AS-path granularity

Anti-spoofing

Prevent traffic with spoofed source IP addresses

Enable source address validation for at least single-homed stub customer networks, their own end-users, and infrastructure

Coordination

Facilitate global operational communication and coordination between network operators

Maintain globally accessible, up-to-date contact information in common routing databases

Global Validation

Facilitate validation of routing information on a global scale

Publish your data so others can validate

https://www.manrs.org/wp-content/uploads/2020/09/About-MANRS-Presentation-PDF-Sep-2020.pdf

Blue shading = Mandatory Action

MANRS Implementation Guide and Tutorials

- Based on Best Current Operational Practices deployed by network operators around the world
- Recognition from the RIPE community by being published as RIPE-706
  - https://www.manrs.org/bcop/
- Tutorials based on information in the Implementation Guide
  - https://www.manrs.org/tutorials

EPOC
Engagement and Performance Operations Center

International Routing Working Group

Established to create a community to work on international routing issues to improve R&E performance. The goal is to engage network owners and NRENs to not only reactively discuss and address ineffective routes, but will work proactively across the community to systematically create policies to prevent them from occurring.

- Co Chairs:
  - Brenna Meade , Indiana University (meadeb@iu.edu)
  - Warrick Mitchell, AARNET (Warrick.Mitchell@aarnet.edu.au)
  - Hans Addleman, Indiana University (addlema@iu.edu)

EPOC
Engagement and Performance Operations Center

Routing Working Group 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

Join the Routing Working Group

- Mailing list routing-wg@gna-g.net
- Contact Brenna to be added meadeb@iu.edu
- Slack
  - APAN Slack Instance, Channel: Routing
- Web
  - https://www.gna-g.net/join-working-group/gna-g-routing-wg/

EPOC
Engagement and Performance Operations Center

Interesting Routing / BGP Resources

- https://twitter.com/bgpstream
- https://bgpstuff.net/
  - https://twitter.com/bgp4_table
  - https://twitter.com/bgp6_table
- http://www.routeviews.org/routeviews/

EPOC
Engagement and Performance Operations Center

EPOC Takeaways - Questions, Comments, Jaded Opinions?

- EPOC is an NSF-funded operations center to help scale science engagement and problem resolution
- Single point of contact for end-to-end performance issues
  - epoc@iu.edu
- We're ready to believe (and help) you!
- More about EPOC:
  - http://epoc.global
- Hans Addleman - addlema@iu.edu

EPOC
Engagement and Performance Operations Center

EPOC is funded by US NSF award #1826994