BGP operations and security

draft-ietf-opsec-bgp-security-01.txt

Jerome Durand

Gert Doering

Ivan Pepelnjak

Goals

- Describe BGP security best practices for the Internet
- Synthesis of many existing pieces available (Cymru, RIPE, many IETF docs, some well known pages…)
- Help smaller AS'es build secure and stable BGP networks
- Have consistent recommendations / best practices
- IP version agnostic (IPv4 and IPv6)

What's covered: the basics

- Control-plane protection (ACL or CoPP)
- BGP session protection (TTL, MD5, TCP-AO), reference to KARR

What's covered: prefix filters

- Default routes
- Special addresses per IPv4 and IPv6 special purpose address registries
- Unallocated addresses (IANA and RIR-based)
- RPKI
- Too specific prefixes (descriptive)
- IXP subnets (with examples)

IXP LAN: don't accept more specifics

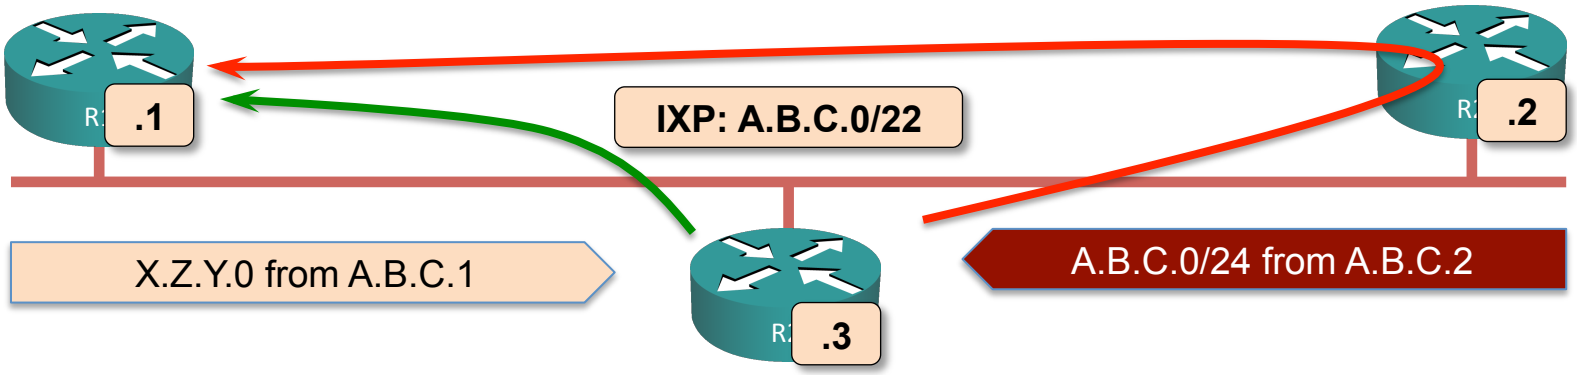

- More-specific IXP prefixes misdirect traffic and destroy EBGP sessions
- A router MUST NOT accept more specific prefixes for IXP LAN prefix

IXP LAN prefix with pMTUd and uRPF

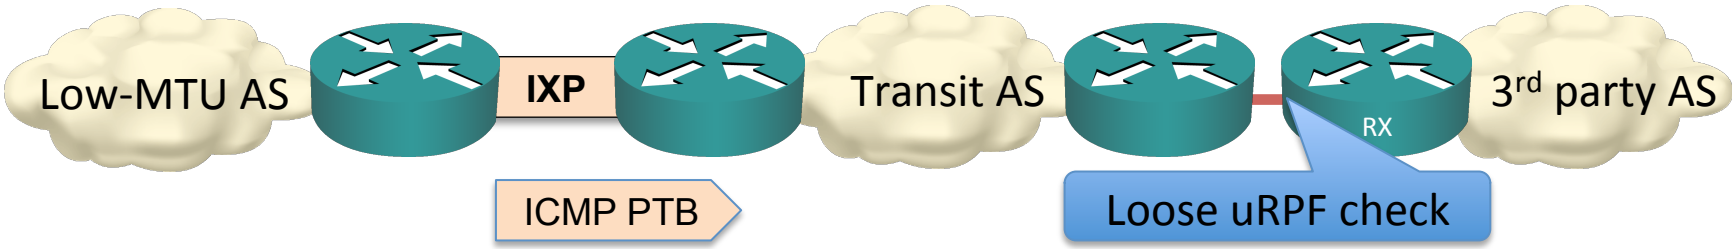

- ICMP packet sourced from IXP LAN address
- uRPF check might drop ICMP packet ➔ IXP LAN prefix SHOULD be advertised
- Downstream AS might perform strict RIR filter ➔ IXP prefix SHOULD pass RIR filter
- Solution: IXP AS advertises IXP LAN prefix

What's covered: prefix use cases

Use cases

- Full routing networks: filters with peers, upstreams and customers
- Leaf networks

Filters described:

- Inbound and outbound filters
- Loose or strict filters

What's covered: AOB

- BGP Route Flap Dampening (don't)
- Maximum prefixes per peering/BGP neighbor (just recalling it's an observed best practice without recommending any particular thing)
- AS-path filters (including customer-facing filters) Clarification that exceptions could occur upon ASN migration
- Next-hop filters (or next-hop enforcement)
- BGP community scrubbing

BGP next hop filters

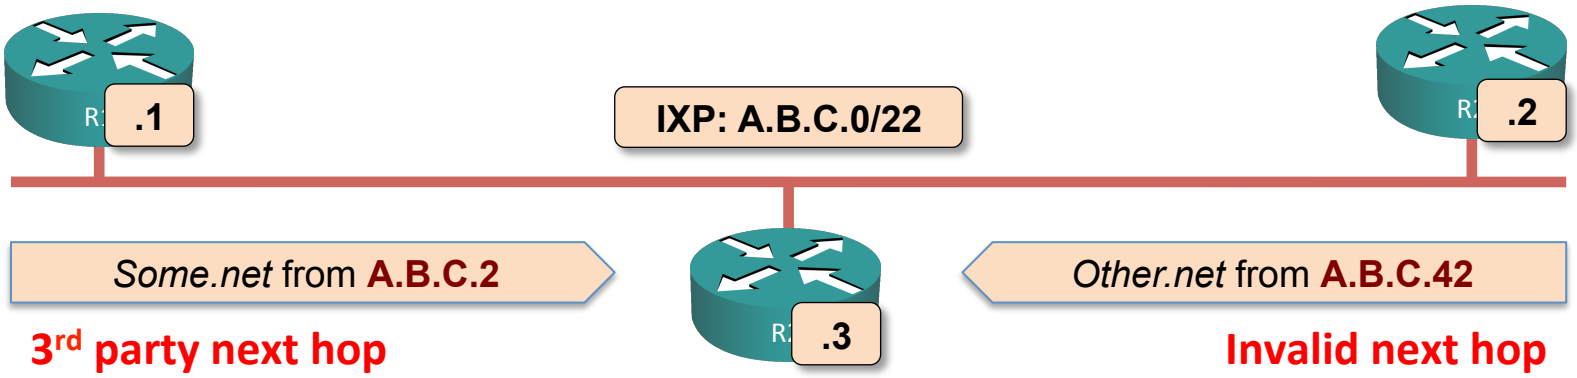

- BGP updates can have 3rd party next hop
- Good for optimal traffic flow, bad on IXP LAN
  - Problem#1 – Traffic redirection
  - Problem#2 – Blackholing (invalid next hop)
- **Solution:** change BGP next-hop to peer's IP address with inbound policy

Changes between opsec-01 and -00

- Obsolete RFC2385 moved from normative to informative reference
- Clarification of preference of TCP-AO over MD5
- Mentioning KARP efforts in TCP session protection section
- Removing reference to SIDR working-group but instead give reference documents
- Better dissociating origin validation and path validation to clarify what's potentially available for deployment
- Adding that SIDR mechanisms should be implemented in addition to the other ones mentioned throughout this document
- Added a paragraph about ASN renumbering for AS-PATH filtering
- Change of security considerations section to clarify what’s not covered
- Added the newly created IANA IPv4 Special Purpose Address Registry instead of references to RFCs listing these addresses

Conclusion and next steps

- Great feedback received so far!
  - Lot of support and many contributions received
  - THANK YOU !!
  - Read, Review & Comment!
- Read the document @ https://datatracker.ietf.org/doc/draft-ietf-opsec-bgp-security/
- Conclusion of WGLC
  - Need to spread the widest possible audience
  - Quick presentations in SIDR, GROW and IDR WG during this IETF to gather maximum feedback
- Discuss on IETF OPSEC WG mailing list @ https://www.ietf.org/mailman/listinfo/opsec

➔Questions ?