BGP:201

Why use BGP and Common Missteps

John Brown, CISSP, CFI, CP-AMEL

- Been connected to the "net" since early 1980's
- Senior Security Evangelist for TEAM CYMRU
- 35+ years in software and network engineering
- Principal Technical Engineer for ICANN L-Root DNS
- Have built internet on 3 continents
- Recovering ISP Owner, had 800+ BGP Peers
- Past instructor for ISC2 (CISSP),
    - Mikrotik (MTCNA, CRE, CINE)
- Passionate about helping ISP's

What we will discuss

- Route Reflectors
- IRR's and RPKI / ROA
- Using BGP for Network Security
  - BOGONS
  - SPAMHAUS
  - RTBH (UTRS blackholes, etc)

Route Reflectors

- The Problem
- BGP Used internally is called iBGP. It is between internal routers
- As you grow the number of BGP speaking routers you need to mesh
- This creates more an more sessions to manage. N(N-1)/2
- So 6 routers will require 6(6-1)/2. or 15 iBGP sessions
- ICK!!

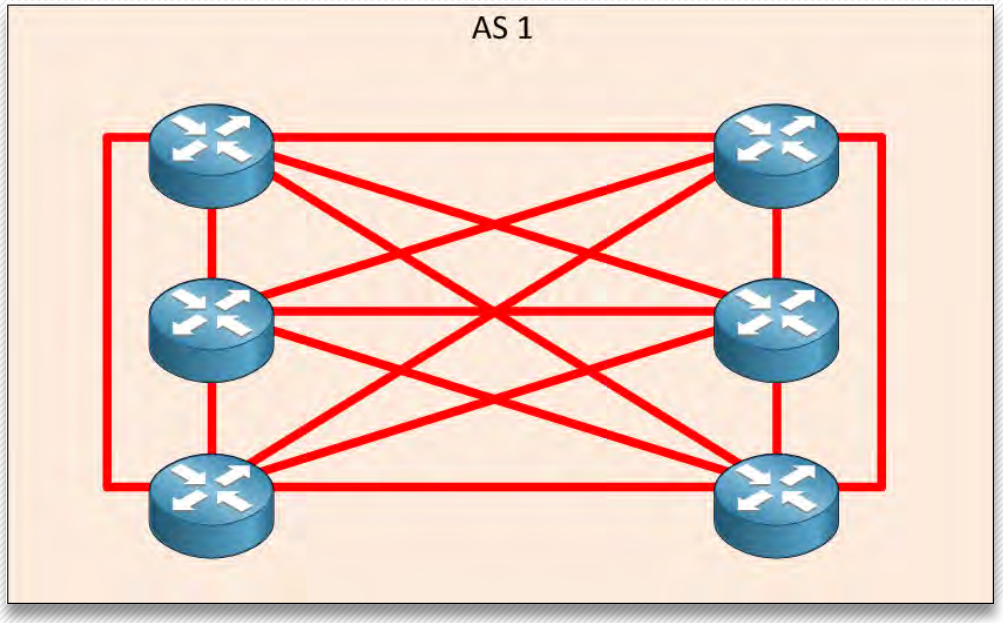

Route Reflectors

- The Solution
- Create a Route Reflector
- Both eBGP and iBGP will work with the Route Reflector
- Now you have a smaller number of sessions to manage
- SPOF can be solved by having multiple route-reflectors

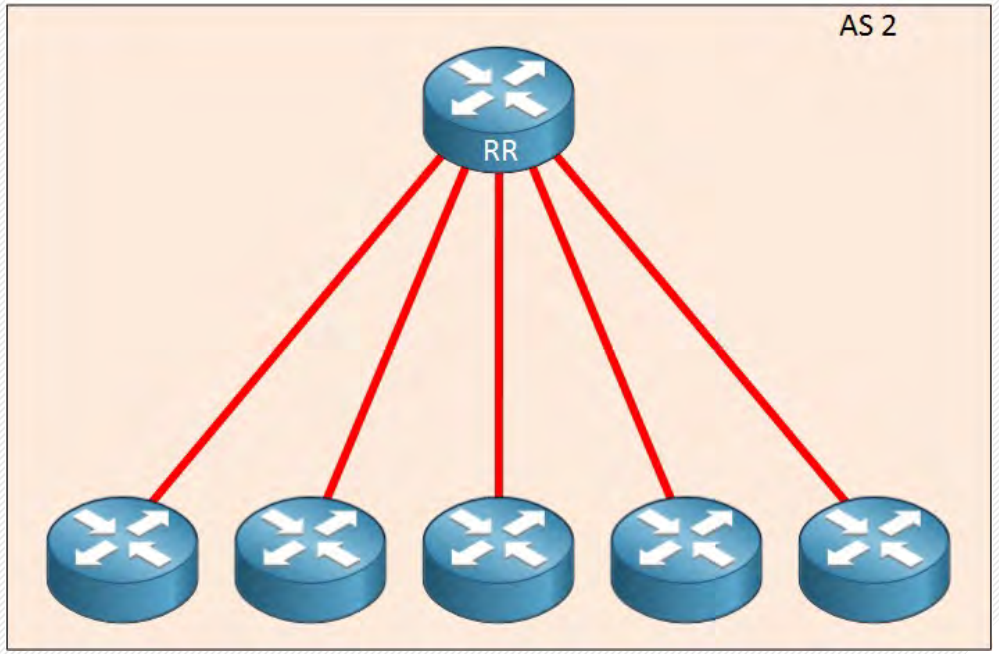

IRR's and RPKI / ROA

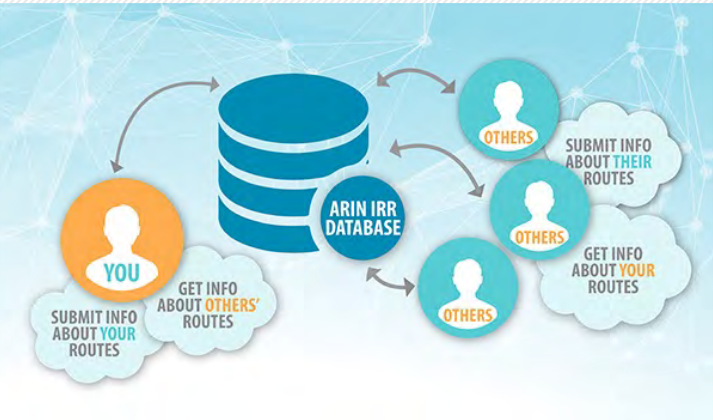

- Routing Security and Authentication
- BGP natively does not have any authentication of routes
- In the early days it became apparent that we needed a way to say "This is my route, and it should be advertised by AS XXX"
- In addition, providers wanted a way to communicate some policy information for both advertised and received routes
- Thus, the IRR (Internet Route Registry) concept was born.

IRR's and RPKI / ROA

- **What does a IRR dB store / hold**
  - There are multiple "objects" that can be stored in a IRR's dB
    - route – This is a specific route. 192.168.1.0/24
    - route6 – This is a specific IPv6 route 1234:456::/32
    - route-set – This is a object of route objects. Used to group routes
    - aut-num – This is a specific ASN. ASN65535
    - as-set – This is an object of AS objects, used to group ASN together

IRR's and RPKI / ROA

- **How and Who uses IRR data**
  - Large Transit providers use IRR data to dynamically create filters.
  - The query the IRR to learn what they should expect from you.
  - Internet Exchanges are now using IRR data to also create filters on what should be received FROM you.
  - These providers will use automated scripts to pull down your objects and create input / import filters. These typically update once a day
  - If your IRR data is missing, or incorrect, then your routing may be broken!!

IRR's and RPKI / ROA

- **Who provides IRR services**
  - There are Commercial and Free IRR providers
  - Commercial
    - RADDB (Merit) One of the original IRR's
    - $450 or so per year to be a member.
    - Not as strict about object validation
  - Non-Commercial (RIR's)
    - ARIN
    - LACNIC
    - RIPE
    - APNIC

IRR's and RPKI / ROA

- **What is RPKI**
  - RPKI == Resource Public Key Infrastructure
  - Uses Cryptography to digitally sign route objects
  - This signed object attests that the route can be advertised by a specific ASN.
  - Defined in RFC 6480
  - Significantly reduces the threat of hijacked routes.
  - Most major providers (HE, COGENT, etc) are now requiring RPKI
  - If your route(s) are not signed, more and more providers will REFUSE your prefixes.

IRR's and RPKI / ROA

- **How to Sign a Route**
  - There are two ways, Hosted or Delegated.
  - Delegated
    - Requires YOU to stand up a server that will sign objects
    - Your server must be reachable by others that want to query
    - Takes several days to get working and multiple skill sets.
    - If it breaks, you have to fix it ☺
  - Hosted
    - The RIR (ARIN for example) maintains the infra-structure
    - You use a web interface to sign your route
    - Takes like 15 minutes or less to sign / publish your route
    - Doesn't cost anything, ARIN members get it as part of the services!
    - https://www.arin.net/resources/manage/rpki/hosted/

IRR's and RPKI / ROA

- AREN'T YOU DONE ALREADY WITH GETTING YOUR RPKI SETUP ??
- By the time this panel is done yammering about BGP your RPKI WOULD BE DONE ☺ ☺ ☺ ☺
- If you have questions GOTO the ARIN Help Desk here at the show.

Using BGP For Network Security

- **We can use BGP to enhance our network security.**
  - Routers forward packets based on what is in the FIB (Forwarding Information Base).
  - The FIB is populated by the RIB (Routing Info Base).
  - Each protocol (OSPF, BGP, STATIC) puts entries into the RIB and these are digested down into the FIB based on best path.
  - We can leverage BGP to "inject" routes to places we DO NOT want to send packets to.
  - We can leverage BGP to tell OTHERS to not send us packets addressed to a specific prefix.

Using BGP For Network Security

- **BOGONS**
  - BOGON's are route prefixes that SHOULD NOT be in the public routing table.
  - Examples:
    - RFC-1918 and other SUA (Special Use Addresses)
    - Prefixes that have not been allocated by a RIR / IANA
    - Prefixes that have been reclaimed by an RIR
  - You should not send traffic to any of the above prefix classes.
  - Team Cymru has a FREE service that maintains a BGP feed of BOGON addresses.
  - https://www.team-cymru.com/bogon-networks

Using BGP For Network Security

- **BOGONS**
  - Simple to configure.
  - Create a input filter that will take all routes learned via the TC BOGON feed and set the next-hop to /dev/null.
  - On some routers (Mikrotik) this is setting the route to blackhole
  - Poof, Now your router will automatically forward all traffic with a BOGON as a DST ADDR to the bit bucket and you won't be polluting the internet ;)

Using BGP For Network Security

- **SPAMHAUS**
  - Well known Anti-Abuse organization.
  - They have BGP based feeds that you can use to drop traffic TO well known malicious actors.
  - Some of the feeds are free, some cost a little.
  - ALL OF THEM ARE WORTH IT.
  - Same concept as BOGONS, we inject bad addresses into BGP and then the router drops the traffic to that bad address
  - https://www.spamhaus.org/blocklists/network-protection/

Using BGP For Network Security

- **UTRS (Unwanted Traffic Removal Service). DDOS Mitigation**
  - Team Cymru has a FREE service you can leverage to tell many many other ISP's around the globe you don't want traffic from them.
  - https://www.team-cymru.com/ddos-mitigation-services
  - Very useful if you are being DDOS'd. Advertise the vicitim IP via BGP to the TC Servers, and we will validate, redistribute to the other members. POOF, traffic dramatically drops.
  - This leverages RTBH methods.

Using BGP For Network Security

- **S/RTBH (Source Remote Triggered Black Hole)**
  - Leverages uRPF (Unicast Reverse Path Forwarding)
  - A form of validation on where the packet came from and is the source valid for that path.
  - Can be used to prevent Source Spoofed packets from leaving your network
  - Can also be used to DROP traffic from outside your network that you don't want.
  - All injected via BGP.

Contact Information

- **Team Cymru General Contact Info**
  - Email: outreach@cymru.com
- **Me Directly**
  - Team Cymru Email: jbrown@cymru.com
  - LinkedIn: https://www.linkedin.com/in/john-brown-cissp-020135

THANK YOU

Dennis Burgess

- Been in the WISP industry since 2001
  - Started a WISP in 2001 using SmartBridge's
  - Link Technologies, Inc since 2006
  - TowerCoverage.com since inception
  - MikroTik Certified Trainer
  - Author of two "Learn RouterOS" Books
  - Been consulting since 1997
    - RouterOS / BGP / OSPF / Firewalls, etc.
    - Stop by Link Technologies, Inc. Booth

What you will learn

- **Common Items**
  - Hold Time
  - Keepalive Time
  - Route filters
  - Max-Prefix's
  - BFD
- **Communities**
  - AS:666 – Blackhole
  - AS:60-120 – set local pref to
  - Others

Lets Start Common Items

- **Common Items**
  - Hold Time
  - Keepalive Time
  - Route filters
  - Max-Prefix's
  - BFD

Lets Start Common Items

- **Common Items**
  - **Hold Time**
    - Default hold time on MikroTik is 120 seconds
      - This means 180 seconds must pass before we consider the peer dead. Up to 180 seconds of downtime, but it will be lower than that.
      - Once the peer is deemed dead, we will remove all routes learned from that peer.
      - I recommend 30 seconds
      - This should only be changed if you have more than one peer. One peer does not matter ..
  - Keepalive Time
  - Route filters
  - Max-Prefix's
  - BFD

Lets Start Common Items

- **Common Items**
  - Hold Time
  - **Keepalive Time**
    - Keepalive is a timer that keeps sending data every so often.
    - Cisco's uses 1/3 of the hold timer, so typically 60 seconds
    - **I recommend 10 seconds if the hold timer is 30.**
  - Route filters
  - Max-Prefix's
  - BFD

Lets Start Common Items

- **Common Items**
  - Hold Time
  - Keepalive Time
  - **Route filters**
    - Route filters are the key to BGP
    - They tell the router what to do with inbound and outbound prefixes.
    - Example:
      - if (dst-len in 1-24 && afi ipv4) {set blackhole yes; set comment BOGON; set bgp-communities 888:889; accept}
    - This does:
      - **Sets prefixes learned as Blackholes**
      - **Sets Comment to BOGON**
      - **Sets BGP-Communites 888:889**
      - **Finally Accepts the prefix**
    - As long as
      - **The dst-length is between 1-24 and it's a IPV4 prefix.**
  - Max-Prefix's
  - BFD

Lets Start Common Items

- **Common Items**
  - Hold Time
  - Keepalive Time
  - **Route filters**
    - Can run multiple filters in one filter rule in v7
    - Here we are running Ipv4 and IPv6 rules in one filter rule
  - Max-Prefix's
  - BFD

```
if (dst-len in 1-24 && afi ipv4) {set blackhole yes; set comment BOGON; set bgp-communities 888:889; accept}
if (dst-len in 1-128 && afi ipv6) {set blackhole yes; set comment BOGON; set bgp-communities 888:889; accept}
```

Lets Start Common Items

- **Common Items**
  - Hold Time
  - Keepalive Tim
  - **Route filters**
    - Note that if you have multiple Filter rules in one chain, you do NOT need a reject
    - In RouterOS v7, all chains are implicit REJECT at the bottom.
    - Many customers with OSPF have OSPF filters setup, but never used them, since they have them setup they upgrade to v7 and all of sudden OSPF stops working (well to them.
      - This is caused by in V6 it was an implicit ACCEPT at the bottom, in V7 its an implicit REJECT
  - Max-Prefix's
  - BFD

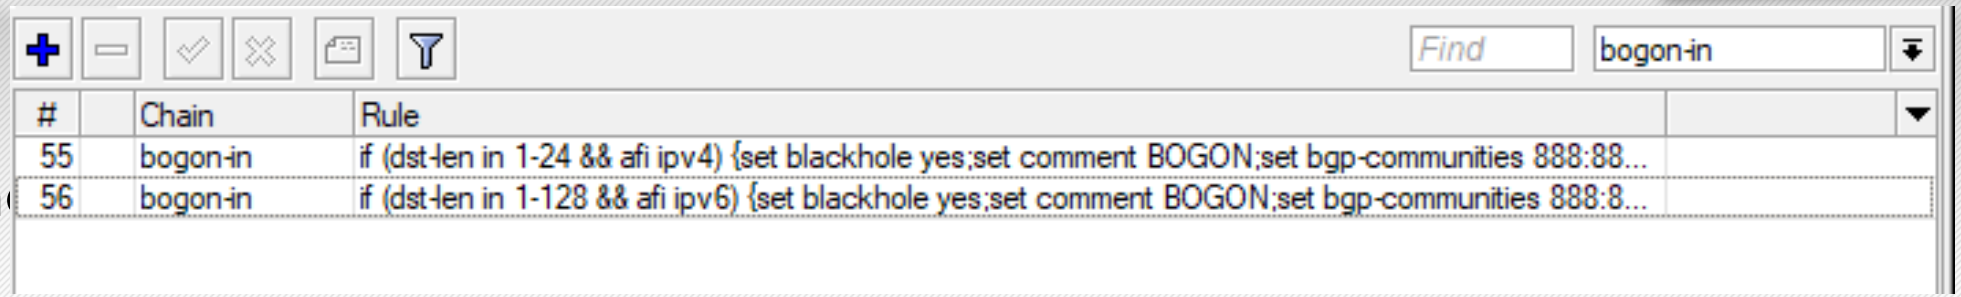

Lets Start Common Items

- **Common Items**
  - Hold Time
  - Keepalive Time
  - **Route filters**
    - if (chain allowed.in) {set distance 5; set comment IBGP_IN_ISP; set bgp-local-pref 120;accept}
      - In this case we are specifying a chain, allowed.in, if it matches those rules then we do the action.
      - **Action: Set DISTANCE to 5, Set Comment to IN_ISP, and BGP-Local-Pref to 120, accept**
  - Max-Prefix's
  - BFD

Lets Start Common Items

- **Common Items**
  - Hold Time
  - Keepalive Time
  - Route filters
  - **Max-Prefix's**
    - Once the peer would reach this limit, the peer would be dropped.
    - Usually a Max-prefix-reset-timer
    - Changed to input.limit-process-routes-ipv4/6.
      - Same function
      - The reset timer says how long to wait till reset
      - On Internet Exchanges this is a good thing
  - BFD

Lets Start Common Items

- **Common Items**
  - Hold Time
  - Keepalive Time
  - Route filters
  - Max-Prefix's
  - **BFD**
    - **Bi-Direction Forwarding Detection**
    - Both OSPF and BGP Support BFD
      - By default sends BFD messages 5 times a second, or every .2 seconds.
      - Failures can occur in sub-second range
      - These timers are calculated based on the average latency in each direction.
      - This is not good for Links that are bumping against their limit
        - As increases in latency (even though the protocol calculates that in automatically) can cause false positives.
        - This is for both fiber/wireless links
      - This can be used to fail VoIP calls over another link.

Other more uncommon BGP Features

- **Communities**
  - AS:666 – Blackhole
  - AS:60-120 – set local pref to
  - Others

Other more uncommon BGP Features

- **Communities**
  - **FIRST DISCLAIMER**
    - Larger providers, Cogent, Hurricane Electric, AT&T, L3, etc, have BGP guides, that should LIST all of the BGP communities that they support.
    - Most small providers or regional DO NOT SUPPORT or offer a listing of what they support.
      - Can they support it, sure, do they, its up the people there
    - There is ALSO NO STANDARDS set, (not that I know of)
    - So 174:60 would be to tell Cogent to set their local-pref to 60 for that prefix
      - But if you did the same thing of HE, or ATT it may not do anything or do something else!
      - BE SURE to get a guide or document that lists what they are wiling to do.
  - AS:666 – Blackhole
  - AS:60-120 – set local pref to
  - Others

Other more uncommon BGP Features

- **Communities**
  - **AS:666 – Blackhole**
    - This is a common BGP community :666, this tells the remote host that this prefix should be blackholed.
      - If you advertise a /32, or an individual IP, in v4 or up to a /64 on v6, with the upstream AS number:666 , it will blackhole
      - Useful to drop all traffic coming into your network to a specific IP, like a DDOS
      - This does take that customer offline that is or was using that IP.
  - AS:60-120 – set local pref to
  - Others

Other more uncommon BGP Features

- **Communities**
  - AS:666 – Blackhole
  - AS:60-120 – set local pref to
    - This is **AS:60-120**
      - This is commonly used to set their Local-pref
      - In eBGP you cannot set their local-pref.
        - Commonly used to set backup routes or routes of last resort.
        - You will send their AS:60 to set to the lowest backup, or 120 for the highest.
        - Note that this will NOT go out their upstreams, unless they have a guide that you can setup a different community on.
  - Others

Other more uncommon BGP Features

- **Communities**
  - AS:666 – Blackhole
  - AS:60-120 – set local pref to
  - **Others**
    - Other common communities that I have used in the past
      - AS:3000 – DO NOT ANNOUCE TO PEERS
      - AS:3001 – PREPEND AS 1 time to peers
      - AS:3002 – PREPEND AS 2 Times to peers
      - AS:962 – Set localpref to 10 upon enter AU regain from other regions of Cogent Backbone
    - You can also learn where the routes came from if the upstream gives you communities
      - EX: 174:21001 – Cogent learned the route from NA internal or customer router.
      - 174:22009 – Cogent learned the route from Italy
- Many of these can allow you to path inbound routing to your desired connection

THANK YOU