BGP Routing Policy Best Practices

AFPIF, Accra-Ghana
August, 2023

Frank Ribeiro

About Me

Before Netflix

Ghana

AGENDA

INTERNET PEERING EVOLUTION

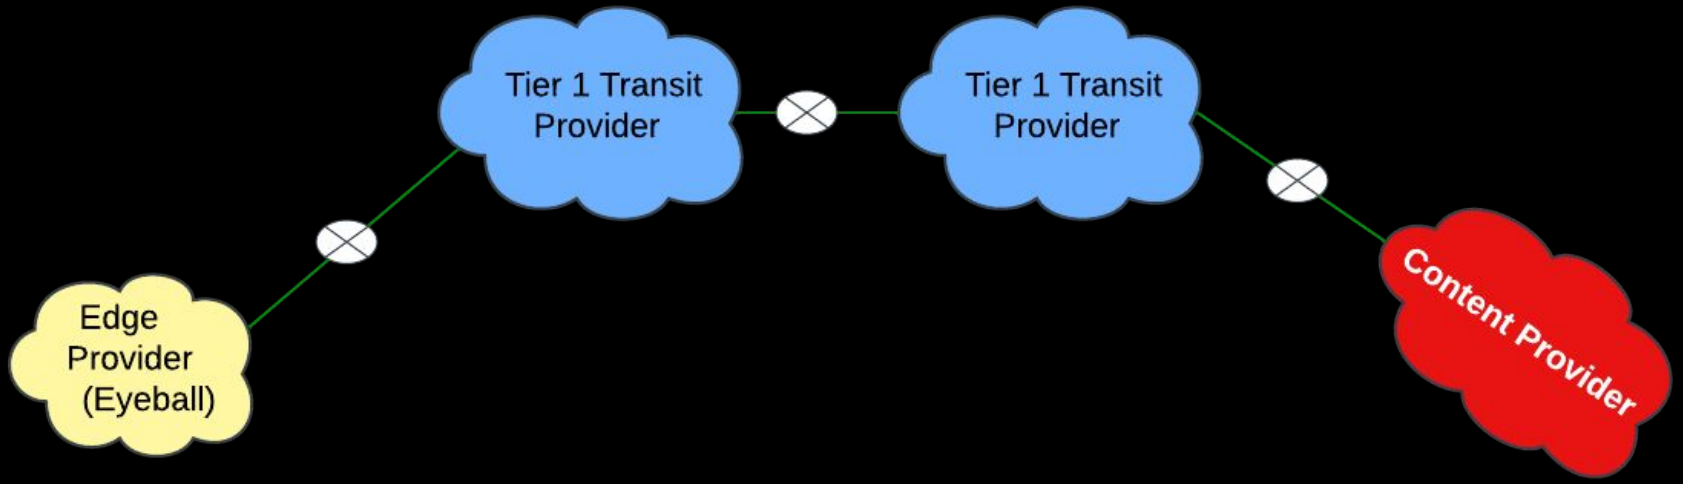

INTERNET PEERING EVOLUTION (Contd.)

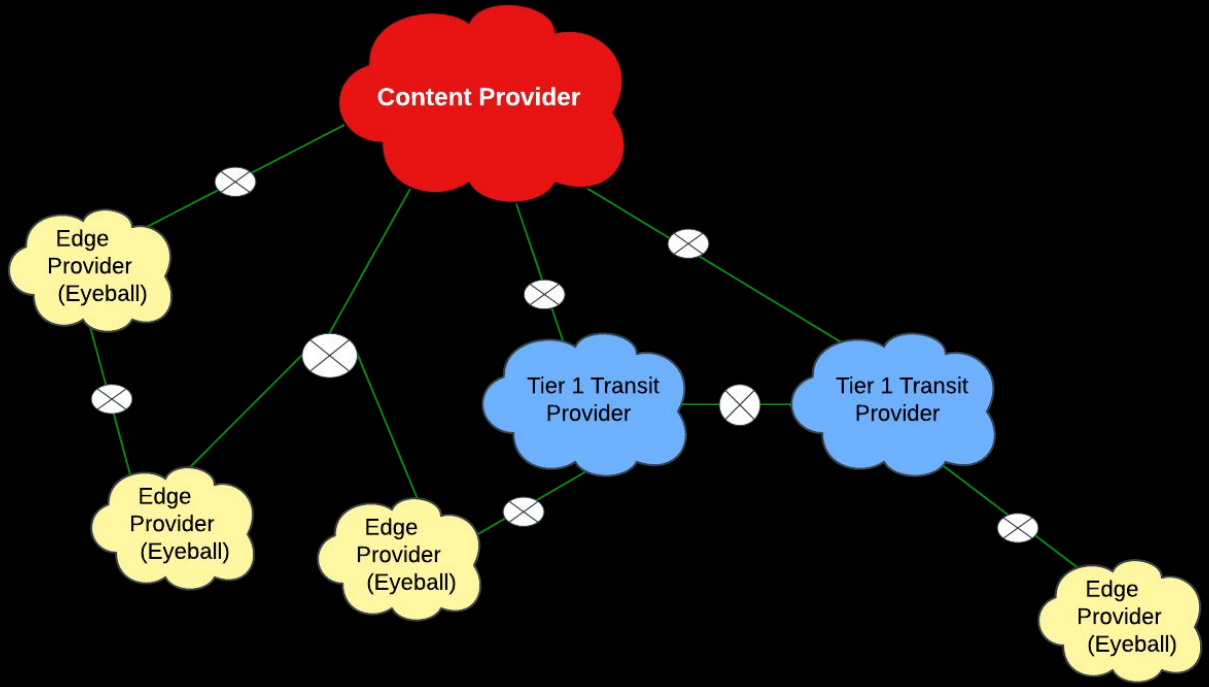

Why Adhere to BGP Best Practices?

- Enhance stability of the internet by reducing outages due to misconfiguration or security breaches

- Better Network Management at Scale due to uniform or structured policies deployed across your network. Eg. group policies for customers, peers and transit

- Security, stability and predictability of your own network and minimizes costs

BGP Policy Best Practices

- RFC8212 defines the default behavior of a BGP speaker when there is no Import or Export Policy associated with an External BGP session.

```
router bgp 2906
  neighbor 192.0.2.2 remote-as $asn_a
  neighbor 192.0.2.5 remote-as $asn_b
!
```

- Configuring BGP peering without using filters means:
  - All best paths on the local router are passed to the EBGP neighbour
  - All routes announced by the neighbour are received by the local router and can cause severe consequences.

- Desired behavior : Deny all rather than fail open

BGP Policy Best Practices

A typical BGP Peering session would have the policy applied in two directions:

Ingress BGP Policy: To control all prefixes which enter your AS

Egress BGP Policy: To control all prefixes which are sent to the AS's peer and the Internet.

The Egress BGP Policy would normally mirror the Ingress BGP Policy of their peer.

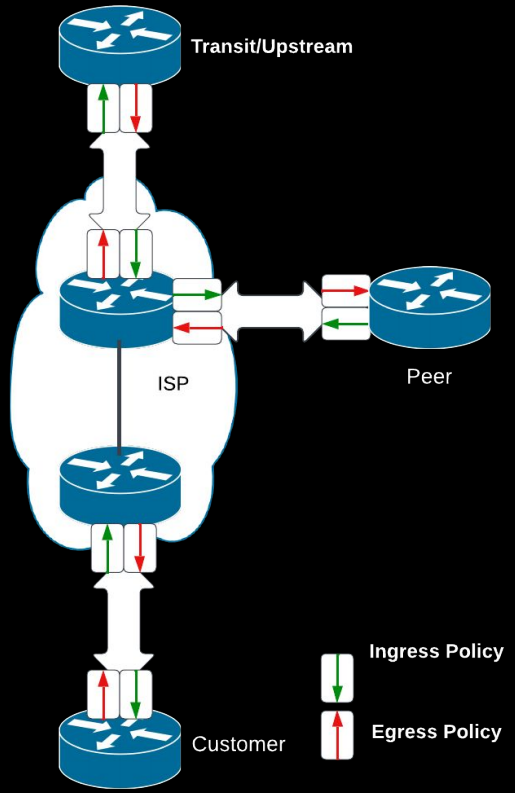

BGP Policy Best Practices - Transit/Upstream Policy

Upstream/Transit Provider is an ISP who you pay to give you transit to the WHOLE Internet.

Receiving prefixes from them is not desirable unless really necessary

- Unless you multihome, full routes are not required, accept only default
- Don’t accept your own prefixes
- Don't accept Special use prefixes for IPv4 and IPv6, BOGONS
- For IPv4:
  - Don’t accept prefixes longer than /24
- For IPv6:
  - Don’t accept prefixes longer than /48

BGP Policy Best Practices - Policy for Peers

What do you **announce** to other networks?

- Your prefixes.
- Customer's Provider Independent (PI) prefixes
- More specific customers prefixes (customers who are multihoming)

What do you **NOT** send to other networks?

- Special use prefixes eg. RFC 5735
- Bogons
- Do not advertise other peer's routes.
- Prefixes longer than /24, i.e, /25 to /32 for IPv4
- Prefixes longer than /48 for IPv6 , ie /64 etc..

BGP Policy Best Practices-Policy for Customer Peering

Service Providers should only accept assigned or allocated prefixes from their downstream peer/customer.

- If the RIR has assigned Eg. /19 to your customer, accept only that.
- If your customer is multihomed, then accept specific prefix assigned to them by the other ISP, if required.
- Filter Special use prefixes eg. RFC 5735.
- Filter Bogons.

Special Use Addresses RFC 5735

0.0.0.0/8 and 0.0.0.0/32—Default and broadcast

127.0.0.0/8 - Host loopback

192.0.2.0/24 - TEST-NET for documentation

198.51.100.0/24 - TEST-NET-2 for documentation

203.0.113.0/24 - TEST-NET-3 for documentation

10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16

169.254.0.0/16 - End node auto-configuration

Also RFC 6598

100.64.0.0 to 100.127.255.255 - CGNAT Range

Bogons

- Bogons are defined as public addresses that have not been allocated or assigned by a Regional internet registry (RIR) or the Internet Assigned Numbers Authority.
- Bogon ASNs;
  - 0
  - 23456
  - 64496-131071
  - 4200000000-4294967295
- Dynamic ways of receiving Bogon Updates
  BGP Peering (Bogon Route Service Project)

Designing BGP Route Policies - Filtering

01	What	• A. Prefix Filters
		  ○ Prefix List
		  ○ AS PATH Filters
		  ○ BGP Community Filters
		• B. RPKI

02	Where	• Ingress • Egress
03	Peers Policy	• A. Transit / Upstream • B. Peers • C. Customers

BGP Policy Best Practices - RPKI Filtering

Resource Public Key Infrastructure (RFC3779) is a public key infrastructure framework designed to secure the Internet's routing infrastructure

Mainly to prevent prefix hijacking by bad actors on the internet.

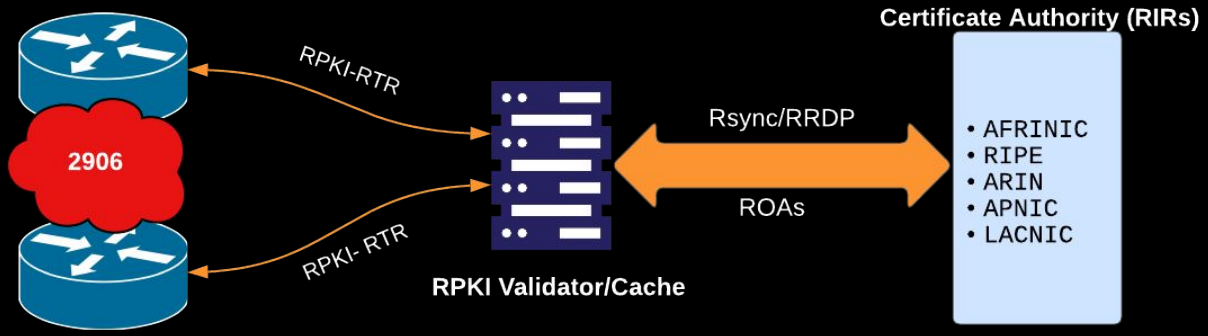

Other Best Practices

1. BGP Max Prefixes Tracking

2. Aggregation

3. Keeping Data Up to Date

Keeping Data Up to Date

- To develop robust routing policies in our networks, we need to have reliable network provider data for engineers to develop appropriate policies where needed.

- Many major transit providers and several content and ip-geo providers use what is contained in the Internet Routing Registry and PeeringDB.

- Best practice for any internet provider is to document their routing policy in the IRRs.
  - At least a route-object

Summary

- Standardise network configuration based on interconnection type and filtering best practices.

- Implement fundamental routing hygiene practices (MANRS)

Want to Discuss More?

If you want to discuss special cases, traffic engineering and more, please join our workshop on Friday morning

Date and Time: Friday 25 August from 8:30 am - 1:00 pm
Where: Accra City Hotel
Tano Hall

Coffee and Lunch will be served

Please Register with Salam or Frank

References/Further Reading

Default External BGP (EBGP) Route Propagation Behavior without Policies: https://datatracker.ietf.org/doc/html/rfc8212
IPv6 Special Use Addresses RFC 5156 : https://datatracker.ietf.org/doc/html/rfc5156
IPv4 Special Use Addresses RFC 5735 :https://datatracker.ietf.org/doc/html/rfc5735
IANA Reserved Shared IPv4 Reserved Space : https://www.rfc-editor.org/rfc/rfc6598.html

Thank You.

Frank Ribeiro
frankr@netflix.com