BGP Route Servers

There are two BGP separate route servers on each peering LAN. It is recommended to **always** peer with both BGP Route Servers at a location, as sessions to both servers ensure that there is no disruption to the advertisement of your prefixes should it be necessary to performance maintenance on a Route Server. The Route Servers do **not** peer with each other by design, so peering with only one server is an unnecessary risk for your network!

> ⚠️ **Bi-lateral peering is considered best practice !**
>
> While the BGP Route Server service is made available as a convenience, it is strongly recommended that, in addition to any sessions you plan to establish with the BGP Route Servers, you still maintain direct bi-lateral peering sessions with peers that you feel are important to your network! BGP Route Servers should be used to pickup quick/easy/additional peers only, and **not** as a replacement for your discrete peering policy!
>
> In particular there are many peers that advertise only a subset of their prefixes to the BGP Route Server. **Always** aim for a bilateral session !

INX	ASN	Hostname	Type	IPv4	IPv6
JINX	37700	routeserver1.jinx.net.za	BIRD	196.223.14.1	2001:43f8:1f0::1
		routeserver2.jinx.net.za	BIRD	196.223.14.2	2001:43f8:1f0::2
CINX	37701	routeserver1.cinx.net.za	BIRD	196.223.22.1	2001:43f8:1f1::1
		routeserver2.cinx.net.za	BIRD	196.223.22.2	2001:43f8:1f1::2
DINX	37699	routeserver1.dinx.net.za	BIRD	196.223.30.1	2001:43f8:1f2::1
		routeserver2.dinx.net.za	BIRD	196.223.30.2	2001:43f8:1f2::2

First ASN Check

Remember that the BGP-RS service at all the INXes do not include the BGP-RS ASN in BGP update messages, as the RS is not actually a transit network. Ensure that if you do plan on peering with the BGP Route Servers, you understand that the BGP-RS does not attach its ASN to outbound BGP messages.

Please implement the IOS "`no bgp enforce-next-as`" (or IOS-XR "`enforce-first-as disable`"), or appropriate equivalent, for your platform.

Filtering policy and process

INX has **always** believed in filtering and we filter all client sessions to the BGP-RS service. We encourage peers to keep their IRR objects accurate to help us to autogenerate these filters.

- Filters are built based on IRRDB registered objects.
- We search the AfriNIC, RADB and RIPE registries (in that order).
- We permit exact match filters for both IPv4 and IPv6.
- RPKI invalids are dropped.
- Some prefixes are automatically filtered by the route servers (eg. bogons and martians).
- We do not accept BGP announcements from private ASNs

> ✅ **Max-prefix**
>
> We recommend that you set the BGP max-prefix to the BGP-RS to 100,000 prefixes for IPv4 and 50,000 prefixes for IPv6

Filtering Frequency

Filter generation happens every 4h at 4h40, 8h40, 12h40, and continues. If you need a filter update done in an emergency, please call, or email us, using the details on the INX support page.

BGP Communities for policy control

A simple set of BGP communities are made available for rudimentary policy control. These will be expanded on over time, as the BGP Route Server service is enhanced. We provide both extended and large community (RFC 8092) support. Note that if you intend to effect policy to 32bit ASNs you'll need to make use of the BGP-LC communities.

> ⓘ **Remember to use the correct ASN**
>
> Note: The communities example below applies to peers using the JINX route servers. The appropriate ASN for each INX, should be substituted when using the BGP route servers, at other INXes.

Community	Action	Explanation
0:peer-asn	deny to peer-asn	block announcement of prefix to peer-as
0:37700	block all	block announcement of prefix to all peers
37700:peer-asn	allow to peer-asn	announce prefix to specific peer-as (in conjunction with block all)
37700:37700	allow all	announce prefix to all peers (implicit default)

We honour the well-known no-export and no-advertise communities as if they were sent to us as a regular peer. If you would specifically like us to propagate these, then please tag as below:

37700: 65281	add no-export	adds the well known no-export community to all routes sent to peers
37700: 65282	add no-advertise	adds the well known no-advertise community to all routes sent to peers

BGP Large Community Support for policy control

Community	Action	Explanation
37700:0:peer-asn	deny to peer-asn	block announcement of prefix to peer-asn
37700:0:0	block all	block announcement of prefix to all peers
37700:1:peer-asn	allow to peer-as	announce prefix to specific peer-as (in conjunction with block all)
37700:1:0	allow all	announce prefix to all peers (implicit default)

Individual network filtering

> ✓ **AS-Path Stripping**
>
> The BGP route servers do not add their own ASN in the advertised path, so if you're planning on constructing a filter list to filter the BGP Route servers, do not use the BGP route servers ASN in the path!

We do not yet publish a route object for the route-servers. We will add that in the future, so, for now, please reach out to the Ops team to see how to do this most efficiently.

Communities returned for filtered routes

If your prefix is filtered by the BGP-RS, we'll return one of the BGP communities below, that should help aid in the debugging process.

Filtered community List

```
PREFIX_LEN_TOO_LONG     = ( routeserverasn, 1101, 1  )
PREFIX_LEN_TOO_SHORT    = ( routeserverasn, 1101, 2  )
BOGON                   = ( routeserverasn, 1101, 3  )
BOGON_ASN               = ( routeserverasn, 1101, 4  )
AS_PATH_TOO_LONG        = ( routeserverasn, 1101, 5  )
AS_PATH_TOO_SHORT       = ( routeserverasn, 1101, 6  )
FIRST_AS_NOT_PEER_AS    = ( routeserverasn, 1101, 7  )
NEXT_HOP_NOT_PEER_IP    = ( routeserverasn, 1101, 8  )
IRRDB_PREFIX_FILTERED   = ( routeserverasn, 1101, 9  )
IRRDB_ORIGIN_AS_FILTERED = ( routeserverasn, 1101, 10 )
PREFIX_NOT_IN_ORIGIN_AS = ( routeserverasn, 1101, 11 )
RPKI_UNKNOWN            = ( routeserverasn, 1101, 12 )
RPKI_INVALID            = ( routeserverasn, 1101, 13 )
TRANSIT_FREE_ASN        = ( routeserverasn, 1101, 14 )
TOO_MANY_COMMUNITIES    = ( routeserverasn, 1101, 15 )
```

Prefixes auto-filtered by the Route Servers

For the overall safety and security of our participants, we actively filter the following prefixes at the Route Servers. That is, advertisements from peers, containing the following networks, will be dropped, and not onward announced.

IPv4 prefixes filtered by the BGP-RS (RFC6890)

```
        martians = [
                10.0.0.0/8+,
                              100.64.0.0/10+,
                169.254.0.0/16+,
                172.16.0.0/12+,
                192.0.0.0/24+,
                192.0.2.0/24+,
                192.168.0.0/16+,
                198.18.0.0/15+,
                198.51.100.0/24+,
                203.0.113.0/24+,
                224.0.0.0/4+,
                240.0.0.0/4+,
                0.0.0.0/32-,
                0.0.0.0/0{25,32},
                0.0.0.0/0{0,7}
        ];
```

IPv6 prefixes filtered by the BGP-RS

```
        martians = [
                ::/0,                    # Default (can be advertised as a
route in BGP to peers if desired)
                ::/96,                   # IPv4-compatible IPv6 address -
deprecated by RFC4291
                ::/128,                  # Unspecified address
                ::1/128,                 # Local host loopback address
                ::ffff:0.0.0.0/96+,      # IPv4-mapped addresses
                ::224.0.0.0/100+,        # Compatible address (IPv4 format)
                ::127.0.0.0/104+,        # Compatible address (IPv4 format)
                ::0.0.0.0/104+,          # Compatible address (IPv4 format)
                ::255.0.0.0/104+,        # Compatible address (IPv4 format)
                0000::/8+,               # Pool used for unspecified,
loopback and embedded IPv4 addresses
                0200::/7+,               # OSI NSAP-mapped prefix set
(RFC4548) - deprecated by RFC4048
                3ffe::/16+,              # Former 6bone, now decommissioned
                2001:db8::/32+,          # Reserved by IANA for special
purposes and documentation
                2002:e000::/20+,         # Invalid 6to4 packets (IPv4
multicast)
                2002:7f00::/24+,         # Invalid 6to4 packets (IPv4
loopback)
                2002:0000::/24+,         # Invalid 6to4 packets (IPv4
default)
                2002:ff00::/24+,         # Invalid 6to4 packets
                2002:0a00::/24+,         # Invalid 6to4 packets (IPv4
private 10.0.0.0/8 network)
                2002:ac10::/28+,         # Invalid 6to4 packets (IPv4
private 172.16.0.0/12 network)
                2002:c0a8::/32+,         # Invalid 6to4 packets (IPv4
private 192.168.0.0/16 network)
                fc00::/7+,               # Unicast Unique Local Addresses
(ULA) - RFC 4193
                fe80::/10+,              # Link-local Unicast
                fec0::/10+,              # Site-local Unicast - deprecated
by RFC 3879 (replaced by ULA)
                ff00::/8+,               # Multicast
                ::/0{49,128}             # Filter small prefixes
        ];
```

Filtering of the Route Servers.

We publish IRR record to show networks that peer with the BGP-RS service. Peers are are so inclined, may use this to create their own filters that they can then elect to use to filter the BGP-RS in question. These are available separately for each INX as AS-SETs for:

- CINX: AFRINIC::AS-CINX-RS
- DINX: AFRINIC::AS-DINX-RS
- JINX: AFRINIC::AS-JINX-RS

This is also published at PeeringDB