



Route Server Rules

Version 1.3 Technical

Table of Contents

GENERAL DESCRIPTION	1
USAGE	1
INGRESS FILTERS	2
BGP COMMUNITIES	2
NOTES ON COMMUNITIES	2
FILTERING POLICY	3

General Description

Route Servers (**RS**) simplifies multilateral interconnection. A single **BGP** session with each **RS** allows all participants to see all prefixes.

RSs retain *next-hop* and other prefix attributes. **RSs** never route traffic; they simply facilitate the exchange of prefixes.

Usage

We provide two **RSs** for resiliency and ask all participants to establish a **BGP** session with each of them. Having two **RSs** also allows us to perform secure upgrades and maintenance without interrupting service.

InteRed members can establish **BGP** sessions with the following information:

- **RS1** - ASN 10391 - IPv4: 179.0.201.1/24 - IPv6: 2801:17:c000:1::1/64.
- **RS2** - ASN 10391 - IPv4: 179.0.201.2/24 - IPv6: 2801:17:c000:1::2/64

Our Routers operate as "**BGP** Route Servers", so they do not insert our **ASN** in the "**AS Path**". For this reason, you need to disable the verification of the first **ASN** of the **AS Path** in your router. For example, on Cisco routers, you need to enter the command "no bgp enforce-first-as" under the **BGP** configuration.

Reference:

http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xr-16/irg-xe-16-book/configuring-a-bgp-route-server.html#d160487e514a1635

To enable **BGP** sessions with the **RSs** it is necessary for **InteRed** staff (soporte@intered.org.pa) to allow your **BGP** session beforehand.

Ingress Filters

InteRed RS implements prefix filtering using Resource Public Key Infrastructure (**RPKI**) and Internet Routing Registry (**IRR**) data from the various **IRR databases** (RIPE, RADB, LACNIC, ARIN, etc.) to allow connected members to announce only publicly registered prefixes. If your prefix has a valid **RPKI ROA**, it will be accepted.

If the result of the **RPKI ROA** check is not found (you have not yet configured a **ROA**), we will test with **IRRDB**. If a prefix does not have a valid **ROA** (invalid status) or is not correctly registered in the **IRRDB**, it will not be accepted.

Use the "Filtered Prefixes" tab and the **RS Looking Glass** in the member portal (portal.interred.org.pa) to determine if any of your prefixes are being filtered or if you are getting to the maximum prefix limit (250 by default). If a higher maximum value is required, please contact the **InteRed** team.

BGP Communities

RS participants must filter their announcements so that they are not sent to other participants. This is useful if you want to prevent your prefixes from reaching IP transit clients through **InteRed RS**, or in other situations. The filtering logic is expressed with the use of **BGP** communities. Large Communities **LCs** also allow prefixes to use "prepend" when advertising to defined participants.

Action	Standard Community	Large Community
Send all prefixes to all other RS participants (default)	10391:10391	10391:1:0
Send a prefix to a RS participant with a specific ASN	10391: ASN	10391:1: ASN
Do not send a prefix to an RS participant with a specific ASN	0: ASN	10391:0: ASN
Do not send a prefix to any RS participant	0:10391	10391:0:0
Perform a prepend to a specific ASN		10391:101: ASN
Prepend to a specific ASN twice		10391:102: ASN
Prepend to a specific ASN three times		10391:103: ASN

Notes on Communities

1. If no communities are specified, the default behavior is to announce all prefixes to all participants.

Example: to announce a prefix only for AS65001 and AS64500, tag the prefix with the communities 0:10391, 10391:65001 and 10391:64500.

Example: to announce a prefix to all members except AS64500, tag the prefix with community 0:64500.

2. Large communities are evaluated before standard **BGP** communities.
3. Most members will want to send their prefixes to **RSs** and tag the community 10391:10391 (or 10391:1:0).
4. To avoid limitations when using standard communities with 32-bit **ASNs**, we recommend using only large communities (RFC 8092) if your router supports it.

5. The known community (*Well-Known BGP Communities*) of RFC 1997 NO_EXPORT (65535: 65281) is not processed by the **RS**, but is passed transparently to the participants. The NO_EXPORT configuration specifies that its prefixes will be advertised to **InteRed** members, but that they should not advertise them to downstream **ASNs** or other external **ASNs**.
6. The logic of using communities is identical for **IPv4** and **IPv6**.

Filtering Policy

InteRed RS filtering policy is based on the **BIRD2** configuration developed by **IXP Manager**:

- Discard prefixes smaller than **/24 for IPv4** and greater than **/48 for IPv6**
- Discard all known martians and bogons (private and reserved addresses defined by **RFC 1918**, **RFC 5735** and **RFC 6598**)
- Ensure that there is at least one (1) **ASN** and less than 64 **ASNs** in the **AS Path**
- Ensure that the **ASN** of the participant is the first **ASN** in the **AS Path**
- Discard any prefix where the IP address of the next hop is not the same as the participant's IP address. This prevents prefix hijacking
- Discard any prefix with an **ASN** from a known transit network in the **AS Path**
- Ensure that the origin **ASN** is in the member's **ASN** set is in **IRRDB AS-Set**. If a set is not specified, all prefixes must originate from the member's **ASN**
- If the prefix is evaluated as a **valid RPKI**, accept it
- If the prefix is evaluated as a **invalid RPKI**, discard it
- If the prefix is evaluated as **RPKI not found** (no ROA exists), revert to standard **IRRDB** prefix filtering:
 - All **origin ASN** must be listed as members: in **AS-Set: (AS Macro)** in the **IRRDB**
 - There must be a **route:** or **route6:** object with a correct **origin:ASN** for the prefix to be accepted
- **InteRed RSs** accept the more specific **IRRDB** prefixes, however, we recommend a *route:object* for each prefix you intend to advertise

To protect against configuration errors, a maximum prefix limit is also enforced, which is set to 250 by default. **RS** will terminate the **BGP** session if the maximum prefix limit is exceeded. If you need to increase this value, please contact **InteRed staff** (soporte@intered.org.pa).