

BGP FAQ

BGP Route Servers

Is it free to use the BGP-RS?

Yes, peering with the BGP Route Servers is available free of charge to all participants, as one of the many value added services at the INXes.

Why should I peer with the BGP route servers?

If you have a simple peering policy, and you're happy to peer with everyone, then there's no doubt about it; you should definitely be peering with the BGP route servers. Peering with the route servers will allow you to be able to exchange traffic from other like minded peers, without going through the hassles of setting up individual bilateral sessions. For example, there are many CDN and DNS operators at the JINX, and having a BGP session to the route servers will help to ensure that you make use of their services seamlessly.

Is it true that peering with a route server hands over your routing policy to the route server operator?

Yes, this is true. This is one of the reasons why larger networks prefer not to peer with the route servers. At the INXes, we provide you with an easy to use, BGP community based, policy tool, where you are able to choose, how to manage your routing policies. More details about this, are available on the BGP [Route Server](#) page.

That said, peering with the route server is **not** something that is dangerous either. We are particularly careful about how configuration is generated, and, **we filter strictly** and probably better than most networks do.

Should I do bilateral peering, or just multi-lateral peering? Better yet, tell me what's best practice...

Do both 😊 No, really, we mean it!

A bilateral sessions means that, if there is ever a problem with the BGP-RS, you'll be unaffected. Peering with the BGP-RS means that, until you get around to setting up bilateral sessions, you'll get the benefit of new peers that come onboard, and are already peering with the BGP-RS.

If this (peering with BGP route servers) is not considered best practice, why do you encourage it?

It's our belief that network engineers should learn and build their capacity to run their networks, so we strongly encourage you to setup bi-lateral sessions to peers that are of particular importance to you. At the same time, we recognise that operations teams are always busy and aren't able to get around to these tasks as quickly as might be possible. Multilateral peering (ie. via the route servers) will give you the benefit of being able to pickup traffic for networks you may not have direct sessions with immediately, and bilateral sessions will give you the piece of mind that if something ever happens to the IXP route servers, your peering will continue uninterrupted.

BGP Route Collectors

What is the point of the Route Collector?

It's a relatively simple way to see which networks, and which prefixes are available at the IXP via your peering sessions. We also provide simple diagnostic tools (ping, traceroute) for network debugging purposes. A web front end to each of the route collectors operated by INX is available via <https://lg.inx.net.za>

There are additional route collectors in the form of [Packet Clearing House](#) and [Routeviews.org](#)

What value do I get from peering with the Route Collector?

The collector (and its interface via the looking glass) acts as a neutral view into multiple networks. By peering with the collector, you're allowing your peers to perform simple debugging techniques (like bgp "show" commands) that can aid them in their work. In the same manner, third party networks that do this, give you the same benefit, for your network.

- [BGP Route Servers](#)
 - [Is it free to use the BGP-RS?](#)
 - [Why should I peer with the BGP route servers?](#)
 - [Is it true that peering with a route server hands over your routing policy to the route server operator?](#)
 - [Should I do bilateral peering, or just multi-lateral peering? Better yet, tell me what's best practice...](#)
- [BGP Route Collectors](#)
 - [What is the point of the Route Collector?](#)
 - [What value do I get from peering with the Route Collector?](#)
 - [Ok, but I'm curious. Why are there so many route-collectors at JINX?](#)
 - [Will the route collector advertise any prefixes back to me?](#)
 - [How safe is this? Can the service be abused by a third party?](#)
 - [I have a restrictive peering policy ...](#)
- [General BGP](#)
 - [Another IXP operator said I should set my peering local-preference to 999 for their IX.](#)
 - [This BGP thing is tough...](#)

The route collector also provides an independent tool, to ensure that what *you think* you are advertising to your peers, is actually what your peers are receiving! And since it takes a few minutes to setup, there's literally no down-side to this.

Ok, but I'm curious. Why are there so many route-collectors at JINX?

Remember that a big part of JINX is community 😊 And JINX is old. Different research groups often collect data in different ways, and we are impartial to that. We encourage and support all research.

Will the route collector advertise any prefixes back to me?

The INX route collector does not advertise any prefixes back to you. Some other route collectors might advertise a single, unused prefix, simply to let you know that the session is working as intended.

How safe is this? Can the service be abused by a third party?

Since the route collector does not advertise any prefixes back to its peers, there is no risk in peering with it. The policy of "non-advertisements" has been hard-coded into the BGP collector setup in at least two places, so there is no chance of this breaking. **It is simply not possible to hairpin peering traffic through the route collector.** That said, if you are truly afraid, simply add in a policy to not accept any routes from the INX route collector.

I have a restrictive peering policy ...

That's quite fine; we'd still welcome your peering session to the route collector, since peering with the route collector is not an indication of network policy, nor does it allow other participants to abuse your established routing policy.

General BGP

Another IXP operator said I should set my peering local-preference to 999 for their IX.

In general you want as much traffic to go across your peering interfaces as possible. But blindly setting your BGP local preference to a random high number without thinking through a complete policy, is a Really Bad Idea. It's a much better idea to plan out a BGP local preference strategy for your network, and then execute according to that. Here are some ideas to help you get started below, keeping in mind the general principle that localpref works on the basis of client_localpref should be higher than peering_localpref which should be higher than transit_localpref. Substitute the numbers for ones that make sense in your network 😊

Prefix Description	Role	Suggestion
Customer prefix	Primary prefix learnt from customer	999
Customer backup	Client backup prefix	950
Domestic peer	Peers within a metro	899
Domestic peer backup	Backup for metro peering	850
Regional peer	Peers from a defined region	799
Regional peer backup	Backup for regional peers	750
....		
Transit		90
Backup Transit		50

This is just a rough idea of system that could work. You're encouraged to spend time to think about where your entry and exit points are, how you would like traffic to flow through your network, and how to design for efficiency.

This BGP thing is tough...

There are many forums where you can get assistance with BGP. We recommend looking through some of the [BGP course materials from the African Network Operators' Group](#) training material, and [joining a local technical community](#) and asking for assistance there. Additionally, we perform **free** BGP and network design training classes periodically, so feel free to [contact us](#) for assistance.