

Peering with Route Servers

Last update: January 28, 2016

General Information

RS1: 195.246.227.1, 2001:7F8:16::1
RS2: 195.246.227.129, 2001:7F8:16::2
aut-num: 39626
macro: IPv4, IPv6: AS-SMILE
max-pref: IPv4: 130000, IPv6: 20000
description: SMILE IXP Riga Latvia
admin-c: UB2
tech-c: UB2, AB81-RIPE
noc: +371-67722211
(24x7 emergency only: +371-67801851)
noc-mail: noc@smile-ixp.net
peer-mail: bgp@smile-ixp.net
Looking-glass: <http://www.smile-ixp.net/lg>

You can accept all routes announced by SMILE Route Servers. All incoming prefixes are filtered against prefix lists built according to customer defined lists or AS Sets.

Route Server Setup

SMILE runs two independent Route Servers – RS1 and RS2. There is no iBGP peer between them.

Route Servers are located in two Data Centers where SMILE switches are located (DC-TV: RS1 and DC-P2: RS2). For customer peers connected in DC-TV (addresses 195.246.227.3 to .126) the primary Route Server is RS1 and for customer peers connected in DC-P2 (addresses 195.246.227.131 to .254) the primary Route Server is RS2. Primary RS will provide for you most optimal path. The secondary RS will provide for you alternative path as long as such exist with same AS path length, or the same path if no viable alternative exist.

We encourage you to accept MED and base your routing decision based on it. For peers connected in same Data Center as your peer you will receive prefixes with transparent MED as it is announced by announcing peer (or 0 if MED is missing). For peers connected in other Data Center you will receive MED increased by 10 (or 10, if MED is missing in incoming update). This is done to gently optimize inter-DC traffic.

We strongly recommend to peer with both RS. This will ensure you are using most optimal paths across SMILE infrastructure. Nevertheless, losing one peer with RS

will not disturb traffic flow although suboptimal routing may happen towards AS who have redundant connections in both Data Centers.

We periodically perform Route Server maintenance. Typically maintenance lasts 5 to 10 minutes per RS. We never do it on both RS at the same time. Loss of peering with one route server for up to 1 hour is acceptable according to our policy.

BGP Neighbor Configuration

For basic BGP neighbor configuration you should:

- 1) Provide via e-mail to bgp@smile-ixp.net the AS Set (AS Macro) name or Customer Defined Prefix List (CDPL), which SMILE RS will use for inbound prefix filtering.
- 2) Provide via e-mail to bgp@smile-ixp.net the MD5 password to protect BGP session. Password is not required (although preferred) if you implement TTL Security.
- 3) Declare your ability and willingness to implement TTL Security.
More information on TTL Security you can find here:
General: <http://tools.ietf.org/html/rfc5082> or Google for "TTL Security"
Cisco: http://www.cisco.com/en/US/docs/ios/12_3t/12_3t7/feature/guide/gt_btsh.html
Juniper: <http://www.juniper.net/us/en/community/junos/script-automation/library/configuration/ttl-security/>
- 4) Ensure that your router accepts paths transparently from Route Server (e.g. RS will not include AS39626 in AS Path).
For recent Cisco IOS routers you need to configure "no bgp enforce-first-as" in "router bgp" section. For Cisco IOS XR you do it per neighbor.
- 5) Ensure your outbound route filters allow only prefixes for which you provide free IP transit with good transmission quality.
- 6) Configure your peer in active mode
(e.g. do NOT use commands like "neighbor 195.246.227.1 transport connection-mode passive" or similar).

Note: AS-SET based prefix lists are built with manually started process with human validation of results. Updates are performed in average once per month. If you want to rebuild prefix list from your current AS Set or you want to provide new CDPL please write e-mail to bgp@smile-ixp.net.

Note: Under any circumstances you are **not allowed** to announce SMILE IXP prefix 195.246.227.0/24, 2001:7F8:16::/64 or any parts of it!

Route Server Control

Filtering

Route Servers performs prefix filtering in both inbound and outbound directions.

Inbound prefix filtering is performed based on the macro (AS-SET) you provide. This macro is resolved against whois.radb.net and appropriate prefix-list filters are built and installed. For IPv4 prefix list allows more specific prefixes up to /27 prefix

length to be accepted. Prefixes /28 and longer up to /32 are accepted only if AS-SET contains route objects exactly the same prefix length. Note that prefixes with length shorter or equal to /27 will be aggregated where possible, but prefixes /28 and longer will be processed individually.

For example following prefixes from AS-SET

```
192.0.2.0/25
192.0.2.128/25
192.0.2.48/29
```

will build following prefix list:

```
ip prefix list XXX permit 192.0.2.0/24 le 27
ip prefix list XXX permit 192.0.2.48/29
```

Outbound prefix filtering is based on Communities. In case multiple paths exist for a particular prefix each RS will announce only one best path from all paths remaining after community based filtering.

Communities

Outgoing routing information of your prefixes via RS can be controlled by you. For this you must send communities. (Cisco IOS neighbor command “send-community standard” or similar). In general if you want to announce all your prefixes to all SMILE customers you should mark all prefixes with Community 39626:39626. **[Change in SMILE Policy 2016] Prefixes received without communities 39626:* and 0:* will get added community 39626:39626 and consequently it will be advertised to all peers.** We strongly suggest adding 39626:39626 yourself and not rely on this option as your downstream ISP may add some community (e.g. 0:1), which will totally disable propagation of this prefix through RS.

If you want to announce a certain route to a certain peer AS <peer-as>, please tag the route with the community 39626:<peer-as>. If you want the route to be distributed to all peers, please tag the route with 39626:39626. You can tag one route with more than one community, though. If you want to block certain routes from being distributed to certain peers, you can do so by sending 0:<peer-as> for blocking routing advertisement by the RS to peer <peer-as> or by sending 0:39626 to block distribution to all peers. Please note that you have to tag a route with both communities 39626:39626 and 0:<peer-as> if you want to block announcement to AS <peer-as> and allow distribution to everyone else.

Evaluation order of communities is as

action	community
1. block of announcement of a route to a certain peer	0:<peer-as>
2. announcement of a route to a certain peer	39626:<peer-as>
3. block of announcement of a route to all peers	0:39626
4. announcement of a route to all peers	39626:39626

Communities not starting with 39626: or 0: are disregarded by the RS and passed transparently.

Route servers will add following communities:

- 39626:1000 – prefix received from peer connected to SMILE1 switch.
- 39626:2000 – prefix received from peer connected to SMILE2 switch.
- 39626:999 – prefix received from peer who requested MAX-PREFIX COUNT set larger than IPv4:2000 and IPv6:500. (Intended to indicate high prefix count for customers with limited hardware capabilities to accept more than 8K/16K routes).

Local Preferences

Route Server on reception mark all prefixes with Local Preference 100. In case you need to announce “last resort backup route” for some of your prefixes there is available special community 39626:80 which sets Local Preference to 80. Please do not use this community unless you have discussed with SMILE personnel about the need and implications.

Transparency

Route Servers unlike regular BGP routers provide transparency features to facilitate optimal path selection in Internet Exchange infrastructure. SMILE Route Servers provide:

- 1) AS path transparency by hiding its own ASN. This shortens AS path by one hop and allows avoid need for direct peering with other SMILE IXP customers.
- 2) Next Hop transparency. RS will provide original Next Hop address as it is received from customer. Note that you should never use RS as Next Hop for any prefix – RS will drop any traffic not destined for it.
- 3) MED transparency. This allows to optimize path in case AS have two (or more) peers with SMILE. Note that you can use MED to direct other SMILE customer traffic towards some of your interfaces. Please note MED discussion under Route Server Setup section regarding MED treatment in SMILE Route Servers.

PeeringDB

We encourage you to use PeeringDB database (<https://www.peeringdb.com>) to document your connection to SMILE Internet Exchange.

SMILE can be found by searching Exchange Points (SMILE-LV) – that’s what you have to select in your Peering Record under Public Peering Locations. Also you can search Networks by AS number 39626 – you will find info about Route Servers.

If you are not using this data base yet then please register for an account there – it is free - and then create your own record.

End of Document