

Fighting BGP Route Leaks with PeeringDB's new “never via route servers” flag

Theo Voss, Head of Network Engineering
ANEXIA Internetdienstleistungs GmbH

Status Quo

- Route Servers only do Origin Validation (IRR, RPKI, etc.)
- Almost no PATH validation, except for some peers using Peerlock
- Certain PATH ASNs should never be visible via RS
(e.g. Tier1 providers or networks not participating in RS)
- No single source of information for Non-RS/Tier1 networks

Result

```
tvoss@decix-router> show route receive-protocol bgp 80.81.192.157 table inet.0 aspath-regex ".*  
2914 .*" hidden
```

| Prefix | Nexthop | MED | AS path |
|------------------|--------------|-----|---|
| 45.159.56.0/24 | 80.81.193.51 | 100 | 58453 58453 40065 6453 2914 3223 18013 |
| 103.43.8.0/24 | 80.81.193.51 | 100 | 58453 58453 40065 6453 2914 137443 139640 |
| 118.184.104.0/24 | 80.81.193.51 | 100 | 58453 58453 40065 6453 2914 137443 |
| 154.210.2.0/24 | 80.81.193.51 | 100 | 58453 58453 40065 6453 2914 137443 |
| 185.181.12.0/22 | 80.81.194.69 | | 25279 2914 174 206620 206620 206620 206620 |
| 185.181.12.0/23 | 80.81.194.69 | | 25279 2914 3320 206620 206620 206620 206620 |
| 185.181.12.0/24 | 80.81.194.69 | | 25279 2914 3320 206620 206620 206620 206620 |
| 185.181.13.0/24 | 80.81.194.69 | | 25279 2914 3320 206620 206620 206620 206620 |
| 185.181.14.0/23 | 80.81.194.69 | | 25279 2914 174 206620 206620 206620 206620 |
| 185.181.14.0/24 | 80.81.194.69 | | 25279 2914 174 206620 206620 206620 206620 |
| 185.181.15.0/24 | 80.81.194.69 | | 25279 2914 174 206620 206620 206620 206620 |

Result

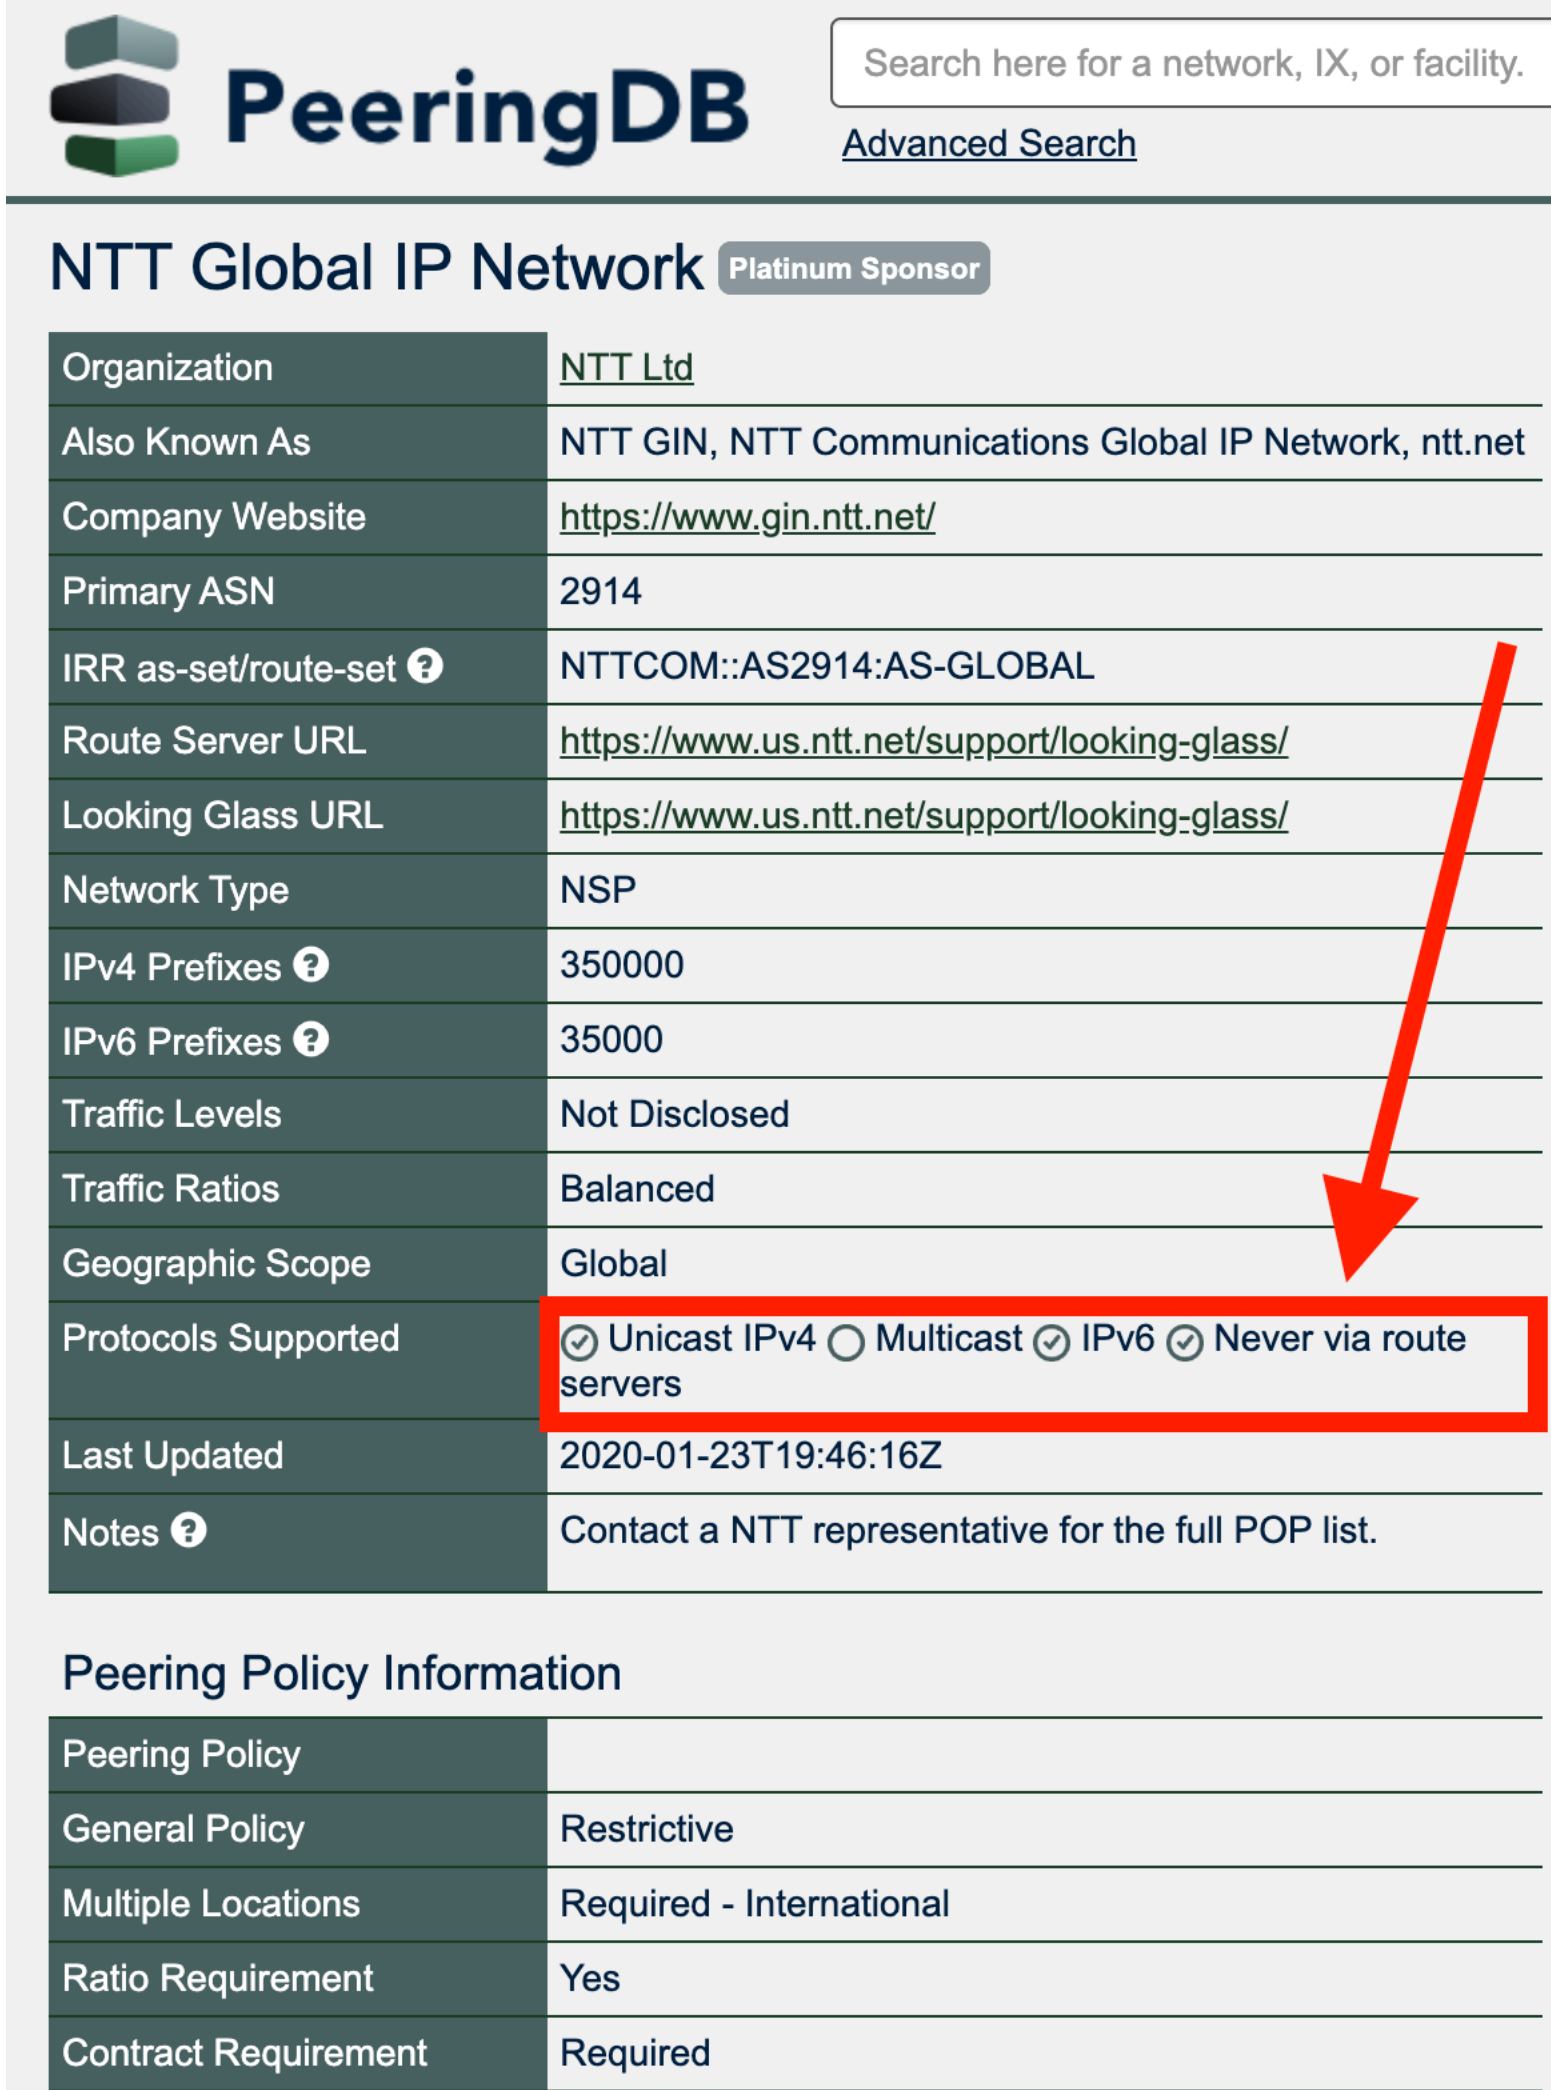
```
tvoss@decix-router> show route receive-protocol bgp 80.81.192.157 table inet.0 as-path-regex ".*  
2914 .*" hidden
```

| Prefix | Nexthop | MED | AS path |
|------------------|--------------|-----|---|
| 45.159.56.0/24 | 80.81.193.51 | 100 | 58453 58453 6453 2914 3223 18013 |
| 103.43.8.0/24 | 80.81.193.51 | 100 | 58453 58453 40065 6453 2914 137443 139640 |
| 118.184.104.0/24 | 80.81.193.51 | 100 | 58453 58453 40065 6453 2914 137443 |
| 154.210.2.0/24 | 80.81.193.51 | 100 | 58453 58453 40065 6453 2914 137443 |
| 185.181.12.0/22 | 80.81.194.69 | 100 | 25279 2914 174 206620 206620 206620 206620 |
| 185.181.12.0/23 | 80.81.194.69 | 100 | 25279 2914 3320 206620 206620 206620 206620 |
| 185.181.12.0/24 | 80.81.194.69 | 100 | 25279 2914 3320 206620 206620 206620 206620 |
| 185.181.13.0/24 | 80.81.194.69 | 100 | 25279 2914 3320 206620 206620 206620 206620 |
| 185.181.14.0/23 | 80.81.194.69 | 100 | 25279 2914 174 206620 206620 206620 206620 |
| 185.181.14.0/24 | 80.81.194.69 | 100 | 25279 2914 174 206620 206620 206620 206620 |
| 185.181.15.0/24 | 80.81.194.69 | 100 | 25279 2914 174 206620 206620 206620 206620 |

**Same situation with other large operators,
Tier-1s and non-RS participants**

Never via route servers

- New feature „Never via route servers“
- Indicates if routes should be visible via RS
- Available since PeeringDB 2.18.0
- 8 networks currently using it on IXPs
we (47147, 42473) are connected to



PeeringDB

Search here for a network, IX, or facility. [Advanced Search](#)

NTT Global IP Network Platinum Sponsor

| | |
|------------------------|--|
| Organization | NTT Ltd |
| Also Known As | NTT GIN, NTT Communications Global IP Network, ntt.net |
| Company Website | https://www.gin.ntt.net/ |
| Primary ASN | 2914 |
| IRR as-set/route-set ? | NTTCOM::AS2914:AS-GLOBAL |
| Route Server URL | https://www.us.ntt.net/support/looking-glass/ |
| Looking Glass URL | https://www.us.ntt.net/support/looking-glass/ |
| Network Type | NSP |
| IPv4 Prefixes ? | 350000 |
| IPv6 Prefixes ? | 35000 |
| Traffic Levels | Not Disclosed |
| Traffic Ratios | Balanced |
| Geographic Scope | Global |
| Protocols Supported | <input checked="" type="checkbox"/> Unicast IPv4 <input type="checkbox"/> Multicast <input checked="" type="checkbox"/> IPv6 <input checked="" type="checkbox"/> Never via route servers |
| Last Updated | 2020-01-23T19:46:16Z |
| Notes ? | Contact a NTT representative for the full POP list. |

Peering Policy Information

| | |
|----------------------|--------------------------|
| Peering Policy | |
| General Policy | Restrictive |
| Multiple Locations | Required - International |
| Ratio Requirement | Yes |
| Contract Requirement | Required |

PeeringDB API Calls

```
% curl -s -X GET --header 'Accept: application/json' 'https://peeringdb.com/api/net/13902' | jq  
'data[].netixlan_set[]'  
{  
  "id": 36695,  
  "ix_id": 31,  
  "name": "DE-CIX Frankfurt: DE-CIX Frankfurt Peering LAN",  
  "ixlan_id": 31,  
  [...]  
}
```

```
% curl -s -X GET --header 'Accept: application/json' 'https://peeringdb.com/api/net/13902' | jq  
'data[].netixlan_set[].ixlan_id'  
31  
50  
60  
1270  
[...]
```

PeeringDB API Calls

```
% curl -s -X GET --header 'Accept: application/json' 'https://peeringdb.com/api/ixlan/31' | jq '.data[].net_set[] |  
{ "asn": .asn, "never_via_rs": .info_never_via_route_servers } | select(.never_via_rs == true)'  
{  
  "asn": 6908,  
  "never_via_rs": true  
}  
{  
  "asn": 2914,  
  "never_via_rs": true  
}  
{  
  "asn": 8455,  
  "never_via_rs": true  
}  
{  
  "asn": 36459,  
  "never_via_rs": true  
}  
[...]
```

BGP Filters

- API output is used to generate AS_PATH filters for our edge routers
- Filters are being updated every night and pushed via NETCONF

Update Policy

Version *
Specifies syntax parsing, currently only version 1 (default)

Type *
Manual = User can edit via UI but not API, API managed = Policy can only be changed via API

1

API managed

Name *

NANOG-DEMO-NON-RS-ASN

Config

```
set policy-options as-path-group NON-RS-ASN as-path AS11164 ".* 11164 .*"  
set policy-options as-path-group NON-RS-ASN as-path AS2914 ".* 2914 .*"  
set policy-options as-path-group NON-RS-ASN as-path AS3491 ".* 3491 .*"  
set policy-options as-path-group NON-RS-ASN as-path AS36459 ".* 36459 .*"  
set policy-options as-path-group NON-RS-ASN as-path AS6830 ".* 6830 .*"  
set policy-options as-path-group NON-RS-ASN as-path AS6908 ".* 6908 .*"  
set policy-options as-path-group NON-RS-ASN as-path AS8455 ".* 8455 ."
```


**Don't peer via RS?
Please enable it, today.**

**Don't filter yet?
Do it, soon :-)**

Questions?