



**RIPE NCC**

RIPE NETWORK COORDINATION CENTRE

# Internet Data Analysis with RIPE RIS

RIS

Emile Aben | 2022 | TMA PhD School

# What Is RIPE RIS?



- BGP Route Collection Project
- Run by RIPE NCC
- Goal: Provide observability of the Internet control plane BGP
- Why: Internet routing has no built in security, so the best we have is observability
  - I.e. People can observe when things go wrong (accident, on purpose) and fix it
- Secondary: We can observe trends

# Other Similar Projects



- Projects with publicly available data:

- RouteViews (U Oregon)
- PCH Route Collectors
- Isolario (now defunct)

- Common:

- Dump data in MRT format
- RFC 6396

```
[Search] [txt|html|pdf|with errata|bibtex] [Tracker] [WG] [Email] [Diff1]  
From: draft-ietf-grow-mrt-17 Proposed Standard  
Errata exist  
  
Internet Engineering Task Force (IETF) L. Blunk  
Request for Comments: 6396 M. Karir  
Category: Standards Track Merit Network  
ISSN: 2070-1721 C. Labovitz  
Deepfield Networks  
October 2011  
  
Multi-Threaded Routing Toolkit (MRT) Routing Information Export Format  
  
Abstract  
  
This document describes the MRT format for routing information  
export. This format was developed in concert with the Multi-threaded  
Routing Toolkit (MRT) from whence the format takes its name. The  
format can be used to export routing protocol messages, state  
changes, and routing information base contents.
```

# How Do We Collect Data?

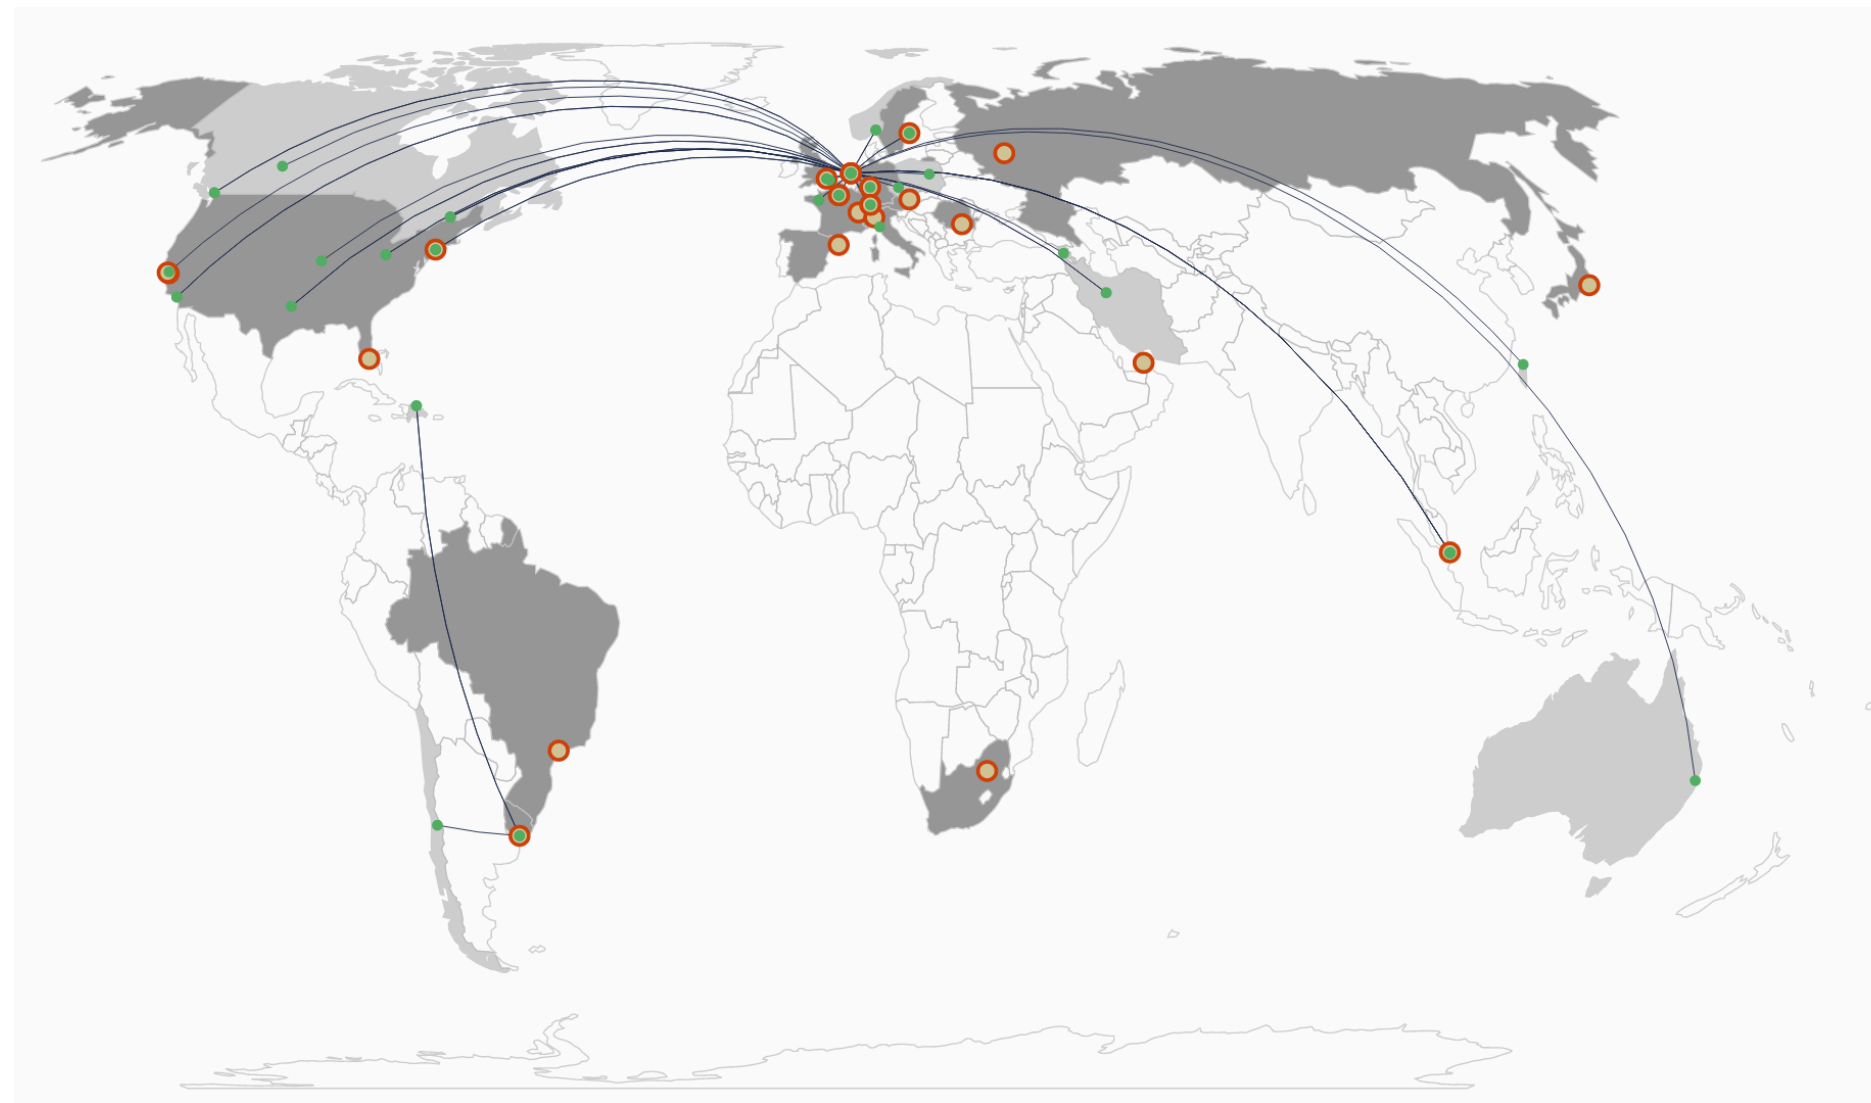


- We collect via BGP sessions to other networks
  - On Internet Exchange Points (IXPs): Limited geographic scope
  - “Multihop”: Other side of the BGP session can be anywhere in the world
- Route Collectors (RRCs): Servers that collect the data
  - Multihop collectors (rrc00, rrc24, rrc25)
  - IXP connected collectors (all other). Can be connected to multiple IXPs
- Overview: <https://www.ris.ripe.net/peerlist/all.shtml>

# RRC Locations



- Peers for multihop collectors can be anywhere
  - Prototype with location info:
    - [https://ris.ripe.net/docs/40\\_Prototypes/20\\_peer\\_metadata.html](https://ris.ripe.net/docs/40_Prototypes/20_peer_metadata.html)



Name	Physical Location	Type	Scope	Raw Data
RRC00	Amsterdam, NL	multihop	global	<a href="#">data</a>
RRC01	London, GB	IXP	LINX, LONAP	<a href="#">data</a>
RRC03	Amsterdam, NL	IXP	AMS-IX, NL-IX	<a href="#">data</a>
RRC04	Geneva, CH	IXP	CIXP	<a href="#">data</a>
RRC05	Vienna, AT	IXP	VIXP	<a href="#">data</a>
RRC06	Otemachi, JP	IXP	DIX-IE	<a href="#">data</a>
RRC07	Stockholm, SE	IXP	Netnod	<a href="#">data</a>
RRC10	Milan, IT	IXP	MIX	<a href="#">data</a>
RRC11	New York, NY, US	IXP	NYIIX	<a href="#">data</a>
RRC12	Frankfurt, DE	IXP	DE-CIX	<a href="#">data</a>
RRC13	Moscow, RU	IXP	MSK-IX	<a href="#">data</a>
RRC14	Palo Alto, CA, US	IXP	PAIX	<a href="#">data</a>
RRC15	Sao Paulo, BR	IXP	PTTMetro-SP	<a href="#">data</a>
RRC16	Miami, FL, US	IXP	Equinix Miami	<a href="#">data</a>
RRC18	Barcelona, ES	IXP	CATNIX	<a href="#">data</a>
RRC19	Johannesburg, ZA	IXP	NAP Africa JB	<a href="#">data</a>
RRC20	Zurich, CH	IXP	SwissIX	<a href="#">data</a>
RRC21	Paris, FR	IXP	France-IX Paris and France-IX Marseille	<a href="#">data</a>
RRC22	Bucharest, RO	IXP	Interlan	<a href="#">data</a>
RRC23	Singapore, SG	IXP	Equinix Singapore	<a href="#">data</a>
RRC24	Montevideo, UY	multihop	LACNIC region	<a href="#">data</a>
RRC25	Amsterdam, NL	multihop	global	<a href="#">data</a>
RRC26	Dubai, AE	IXP	UAE-IX	<a href="#">data</a>

# RIS Features



- RIS Live
- RIS MRT Data
- RIS Whois
- Routing Beacons
- Interfaces via RIPEstat
  
- Prototypes
  - Per peer dumps
  - Peer metadata

# RIS Live



- Near realtime streaming RIS API
- JSON, Websockets
- <https://ris-live.ripe.net/>
- <https://ris-live.ripe.net/manual/>

```
{
  "prefix": null,
  "path": null,
  "type": null,
  "require": null,
  "moreSpecific": true,
  "lessSpecific": false,
  "host": "rrc21",
  "peer": null,
  "socketOptions": {
    "includeRaw": false,
    "acknowledge": true
  }
}
```

```
// Received at 11:32:58 (3.27 second delay)
{
  "timestamp": 1655717575.51,
  "peer": "2001:7f8:54::228",
  "peer_asn": "24482",
  "id": "21-3513-169574669",
  "host": "rrc21",
  "type": "UPDATE",
  "path": [24482, 174, 3320, 2773],
  "community": [[174, 21100], [174, 22008], [24482, 1], [24482, 12010], [24482, 12011], [24482, 20200], [24482, 64605]],
  "origin": "igp",
  "aggregator": "2773:10.10.1.1",
  "announcements": [
    {
      "next_hop": "2001:7f8:54::228",
      "prefixes": [
        "2605:b740:80::/42"
      ]
    },
    {
      "next_hop": "fe80::c203:8007:966e:500e",
      "prefixes": [
        "2605:b740:80::/42"
      ]
    }
  ]
}
```

# RIS MRT Data



- Archive of routing data RIS collects
  - First data: 1999-09
- Organised by RRC
  - Some quite big (rrc00!)
- **bview vs update:**
  - Updates are BGP messages
    - 5 min intervals
  - Dumps (bviews) are “table state”
    - Snapshot at specific time, ever 8 hours

```
https://data.ris.ripe.net/rrcXX/YYYY.MM/TYPE.YYYYMMDD.HHmm.gz
```

with:

- XX = the RRC number
- YYYY = year
- MM = month
- TYPE = the type of file, which is either **bview** (dumps) or **update** (updates)
- DD = day
- HH = hour
- mm = minute

Currently dumps are created every 8 hours, and updates are created every 5 minutes



# RIS MRT Data Parsers



- Many options
  - BGPDump: Oldest, written in C, command line
  - BGPKit: Newest(?), written in Rust, python binding, command line
  - BGPStream: Package, written by researchers, C , python binding, command line
  - [https://ris.ripe.net/docs/20\\_raw\\_data\\_mrt.html#tooling](https://ris.ripe.net/docs/20_raw_data_mrt.html#tooling) :

- [bgpdump](#) : One of the first MRT parsers, written in C
- [CAIDA bgpstream](#) : Software suite, consisting of commandline, bindings to python, C library
- [bgpscanner](#) : MRT parser in C, written for speed. Part of the now-defunct Isolario project
- [microbgp suite](#) : MRT parser in C, written for speed
- [BGPKit](#) : MRT parser written in Rust
- [java MRT](#) : MRT parser written in Java
- [mrt parser](#) : A minimal, experimental MRT parser in Rust

# RISWhois / dumps



- WHOIS interface for prefix-to-origin mappings
  - Data dumps also available
- Documented here:
  - [https://ris.ripe.net/docs/27\\_riswhois.html](https://ris.ripe.net/docs/27_riswhois.html)

```
$ whois -h riswhois.ripe.net 193.0.0.1
% This is RIPE NCC's Routing Information Service
% whois gateway to collected BGP Routing Tables, version2.0
% IPv4 or IPv6 address to origin prefix match
%
% For more information visit http://www.ripe.net/ris/riswhois.html
%
% Connected to backend ris-whois15.ripe.net


route:      193.0.0.0/21
origin:     AS3333
descr:      RIPE-NCC-AS Reseaux IP Europeens Network Coordination Centre (RIPE NCC), NL
lastupd-frst: 2021-12-08 16:11Z 80.81.192.175@rrc12
lastupd-last: 2022-06-09 05:48Z 86.104.125.170@rrc22
seen-at:    rrc00,rrc01,rrc03,rrc04,rrc05,rrc06,rrc07,rrc10,rrc11,rrc12,rrc14,rrc15,rrc16,rrc18,rrc19
num-rispeers: 393
source:     RISWHOIS
```

23969	1.4.222.0/24	384
23969	1.4.223.0/24	384
23969	1.4.224.0/20	418
23969	1.4.240.0/21	420
23969	1.4.248.0/23	387
23969	1.4.250.0/23	1
23969	1.4.252.0/22	420
4725	1.5.0.0/16	361
9583	1.6.0.0/22	402
9583	1.6.1.0/24	401
9583	1.6.4.0/22	402

# RIS Beacons



- Periodically announced and withdrawn prefixes
- From specific RRCs
  - Together with an anchor prefix
  - RIS peers may or may not accept this prefix (don't assume they will!)
- Useful to study dynamics of announce/withdraw on the Internet
  - Withdraw takes much longer, due to “path hunting”



# RIS prototypes




- Location data for individual peers:
  - [https://ris.ripe.net/docs/40\\_Prototypes/20\\_peer\\_metadata.html](https://ris.ripe.net/docs/40_Prototypes/20_peer_metadata.html)
- Per peer dump files:
  - [https://ris.ripe.net/docs/40\\_Prototypes/10\\_per\\_peer\\_dumps.html](https://ris.ripe.net/docs/40_Prototypes/10_per_peer_dumps.html)
  - <https://www.ris.ripe.net/dumps-per-peer-rest/prototype/>
  - Smaller files and hourly granularity

```
{  
  "rrc": 25,  
  "asn": 212744,  
  "ipv4": "103.158.223.243",  
  "ipv6": "2a0f:5707:aa80:2744::1",  
  "country": "DE",  
  "city": "FRA",  
  "feed-type": ""  
},
```

# RIS Data in RIPEstat



- RIPEstat: one-stop shop for info on ASNs and prefixes
  - Contains a lot of data derived from RIS
  - Example: BGPlay



# How Representative Is RIS for the Internet?



- RIS is a opportunistic sampling of the Internet
  - We don't know how representative it is
  - Quantitative observation on RIS != Quantitative observation on the Internet



# Hands On Exercise



- Pick your poison!
- Try to get one of the MRT parsers working for you
  - [https://ris.ripe.net/docs/20\\_raw\\_data\\_mrt.html#tooling](https://ris.ripe.net/docs/20_raw_data_mrt.html#tooling)
- What worked for me:
  - Bgpdump
  - Bgpstream
  - Bgpkit
- Tell us your experiences. What can be improved?