

BGP.Tools and IXP Route Collection

Ben Cartwright-Cox - ITNOG7 2023



Quick overview of bgp.tools



AS206924

Browse the Internet ecosystem

Search by ASN (AS13335), or Prefix (8.8.8.0/24), or DNS (bgp.tools)

Start here...



You are connecting from

- IPv6: 2a0c:2f07:4663:4663:92e2:baff:fe61:c389
- Ben Cartwright-Cox (AS206924)
- 2a0c:2f07:4663::/48
- DNS: 185.230.223.109
- DNS: 2a0c:2f07:4896:666:216:3eff:feff:861f
- DNS: 2a0c:2f07:29:666::5353

Example Pages

- [Cloudflare \(AS13335\)](#)
- [LINX LON1](#)
- [Google DNS Prefix](#)

Recent Updates

- [March 2023 Changelog](#)
- [February 2023 Changelog](#)
- [January 2023 Changelog](#)

Why use BGP.Tools?

We offer for free:

- Near Realtime BGP Data
- User Friendly interfaces
- [Frequently updated external data](#)

We offer for paid users:

- [BGP Network Monitoring](#)
- [IRR Database Monitoring](#)

[Scripting/API](#) [Credits](#) [Pricing](#) [Contact Us](#) [Issue Tracker](#) [Contribute Data](#)

Port 179 Ltd is a company registered in England and Wales (Registration Number: 14127855)



bgp.tools


bgp.tools/as/15433

Start here...



Logged in as AS206924

View Edit



Telecom Italia San Marino S.p.A

AS Number 15433

Website <https://www.telecomitalia.sm>

Overview Prefixes Connectivity Whois

IX

Registered on
28 Jun 2000 (22 years old)

Network status
Active, Allocated under RIPE


Network type
Eyeball

Prefixes Announced
75 IPv4, 0 IPv6

Upstreams

d625 • [AS6762](#) - Telecom Italia Sparkle S.p.A a6cd164

ASN Info



BT Italia S.p.A.

AS Number **8968**

Website <http://www.bt.com/italia>

Overview Prefixes Connectivity Whois IX

Peers
57


Upstreams
4

Downstreams
31 (Cone: 37)

Network Policy

[Click here to learn more about what this graph means and what makes up a network policy](#)

Policy awesome_curie



Prefixes in selected policy

| Prefix | Description |
|----------------------------------|------------------|
| 85.20.128.0/17 | BT Italia S.p.A. |
| 213.217.128.0/18 | BT Italia S.p.A. |

How are upstreams and downstreams calculated?

Upstreams

| ASN | Description | IPv4 | IPv6 |
|------------------------|--------------------------------|------|------|
| AS1299 | Arelion (fka. Telia Carrier) | ✓ | ✓ |
| AS3356 | Lumen (Level 3) | ✓ | ✓ |
| AS5400 | British Telecommunications PLC | ✓ | ✓ |
| AS3320 | Deutsche Telekom AG | ✓ | ✓ |

Peers


| ASN | Description | IPv4 | IPv6 |
|-------------------------|---------------------------------|------|------|
| AS3303 | Swisscom (Schweiz) AG | x | ✓ |
| AS28716 | Retelit Digital Services S.p.A. | x | ✓ |
| AS1299 | Arelion (fka. Telia Carrier) | ✓ | ✓ |
| AS3356 | Lumen (Level 3) | ✓ | ✓ |
| AS3313 | BT Italia S.p.A. | ✓ | ✓ |
| AS5400 | British Telecommunications PLC | ✓ | ✓ |
| AS33942 | IREN ENERGIA S.P.A. | ✓ | x |

| | | | |
|-------------------------|---------------------------------|---|---|
| AS3303 | Swisscom (Schweiz) AG | x | ✓ |
| AS28716 | Retelit Digital Services S.p.A. | x | ✓ |
| AS1299 | T-Mobile US (Sprint) | x | ✓ |
| AS13237 | euNetworks GmbH | x | ✓ |

Downstreams

| ASN | Description | IPv4 | IPv6 |
|-------------------------|-----------------------|------|------|
| AS3313 | BT Italia S.p.A. | ✓ | ✓ |
| AS33942 | IREN ENERGIA S.P.A. | ✓ | x |
| AS47224 | Telecom Italia S.p.A. | ✓ | x |

Global Looking Glass

```
Terminal
File Edit View Search Terminal Help
recording to ~/.ssh/recordings/rec_1683493861827672619
[22:11:01] ben@metropolis:~$ ssh ifnbg@bgp.tools
Welcome 2a0c:2f07:4663:4663:92e2:baff:fe61:c389 This session is supported by:

bgp.tools> show route 185.230.223.0/24 match 1003 short
[AS1003 - andrewnet] TORv4 [1003 835 174 5511 206924] [[
174:21100 174:22012 835:10000 1003:1200 1003:1201 62513:10000 206924:666:0 20692
4:5511:0]]
[AS1003 - andrewnet] MCiv4 [1003 12186 32097 1299 3170 2
06924] [[1003:1300 1003:1302 1003:1306 12186:30000]]
bgp.tools>
```



Web UI

Terminal UI

Query all public BGP sessions connected to bgp.tools

Lookup by CIDR, only applies to sessions that have been marked to be exported publicly

185.230.223.0/24



Search Filters:


Must Contain ASN:

Query Overview:

322 Sessions Responded

451 Matching Paths Displayed

Supported by:



185.230.223.0/24 unicast [AS35487 - edge-ng-los01 0000-00-00] * (?/-) [AS206924]

Type: BGP

BGP.as_path: 35487 8849 5511 206924

BGP.community: (56630,3000) (56630,3057) (57695,13000)

unicast [AS1003 - TORv4 0000-00-00] * (?/-) [AS206924]

Type: BGP

BGP.as_path: 1003 835 174 5511 206924

BGP.community: [AS174: Route is learned from EU (Europe) non-customer.] [AS174: United Kingdom] [AS835: Source:..

Cogent Transit] (1003,1200) (1003,1201) (62513,10000)

BGP.large_community: (206924, 666, 0) (206924, 5511, 0)

unicast [AS34979 - 39D-TEL-02 0000-00-00] * (?/-) [AS206924]

Type: BGP

d6253275ab09f7afeaca74955a6cd164

IXP Info Pages

MIX-IT

 [Go to PeeringDB page](#)

 [Go to IXP-DB page](#)

Data Feeds Available:

RS Feed, Ping, MAC Address

Do you run this IX and want to help with feeds? [Contact Us!](#)

List of members (393 routers over 347 ASNs):

| ASN | Description | IPv4 | IPv6 |
|---|--------------------------------------|--------------|--------------------|
|  <input checked="" type="checkbox"/> AS16004 | MIX S.r.L. - Milan Internet eXchange | 217.29.66.1 | 2001:7f8:b:100:1d1 |
|  <input checked="" type="checkbox"/> AS5392 | TELNET S.r.l. | 217.29.66.5 | 2001:7f8:b:100:1d1 |
|  <input checked="" type="checkbox"/> AS12654 | RIPE NCC - RIS | 217.29.66.6 | 2001:7f8:b:100:1d1 |
|  <input checked="" type="checkbox"/> AS42692 | Overweb Srl | 217.29.66.7 | 2001:7f8:b:100:1d1 |
|  <input checked="" type="checkbox"/> AS25152 | RIPE - K-ROOT | 217.29.66.8 | 2001:7f8:b:100:1d1 |
|  <input checked="" type="checkbox"/> AS1267 | WIND TRE S.P.A. | 217.29.66.9 | 2001:7f8:b:100:1d1 |
|  <input checked="" type="checkbox"/> AS33891 | Core-Backbone GmbH | 217.29.66.10 | 2001:7f8:b:100:1d1 |

MINAP Milan

 [Go to PeeringDB page](#)

 [Go to IXP-DB page](#)

Route Server ASN: [AS43369](#)











Data Feeds Available:

RS Feed, Ping, MAC Address

Top Vendors

| Vendor | % |
|--|-----|
|  Cisco Systems, Inc | 26% |
|  Juniper Networks | 24% |
|  Routerboard.com | 17% |
|  MIX s.r.l. | 6% |
|  Arista Networks | 6% |

List of members (93 routers over 82 ASNs):

| ASN | Description | IPv4 | IPv6 | Speed |
|---|------------------|-------------|--------------------------|---------|
|   <input checked="" type="checkbox"/> AS12637 | Seeweb s.r.l. | 185.1.114.2 | 2001:7f8:c5::a501:2637:1 | 10 gbps |
|   <input checked="" type="checkbox"/> AS8816 | Metrolink S.R.L. | 185.1.114.3 | 2001:7f8:c5::a500:8816:1 | 10 gbps |
|   <input checked="" type="checkbox"/> AS5392 | TELNET S.r.l. | 185.1.114.4 | 2001:7f8:c5::a500:5392:1 | 10 gbps |
|   <input checked="" type="checkbox"/> AS20836 | CDLAN SpA | 185.1.114.5 | 2001:7f8:c5::a502:0836:1 | 10 gbps |
|   <input checked="" type="checkbox"/> AS20848 | TWT S.p.A. | 185.1.114.6 | 2001:7f8:c5::a502:0848:1 | 20 gbps |



bgp.tools

Traceroutes/Looking Glass/Agents

Orange S.A.

AS Number **5511**

BGP

Select BGP Session to query:

London [IPv4] [IPv6]

Input Prefix:

80.80.80.80

Query

```
80.80.80.0/24      unicast [London 0000-00-00] * (?/-) [AS60679]
Type: BGP
BGP.as_path: 5511 3356 30247 60679
BGP.community: [AS5511: United Kingdom] [AS5511: Route received from peering partner]
[AS5511: Route received in Europe from peering] [AS5511: TUNE announce to US peers]
```

You need a bgp.tools (free) +
RIPE Atlas account for this



AS200924

View Edit Looking Glass

The Crazy Red Cat Company LTD

AS Number 34695
Website <https://www.e4a.it>

Overview Prefixes DNS Tools IX

Registered on 18 Mar 2018 (AS200924)

Network status
Active, Allocated under RIPE

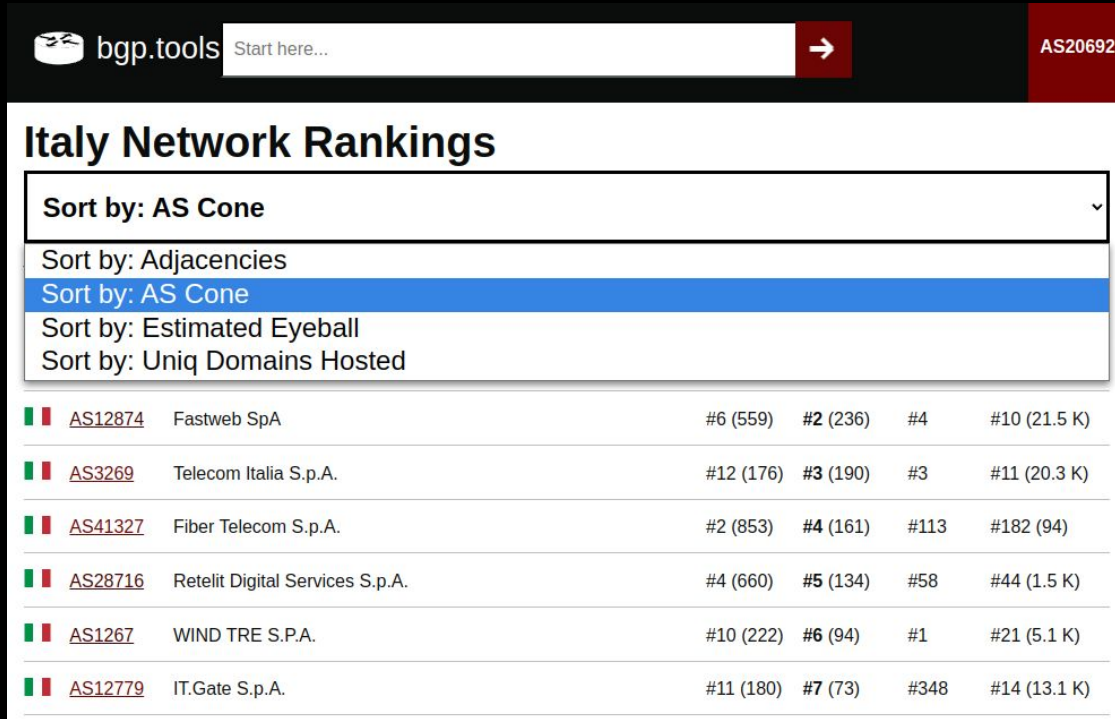
Network type
Unknown

Prefixes Announced
2 IPv4, 0 IPv6

Unstreams Rankings

This is a GIF, Sorry PDF users!

Network Ranking



bgp.tools Start here... AS206924

Italy Network Rankings







Sort by: AS Cone

Sort by: Adjacencies

Sort by: AS Cone

Sort by: Estimated Eyeball

Sort by: Uniq Domains Hosted

| | | | | | |
|--|---------------------------------|-----------|----------|------|--------------|
|  AS12874 | Fastweb SpA | #6 (559) | #2 (236) | #4 | #10 (21.5 K) |
|  AS3269 | Telecom Italia S.p.A. | #12 (176) | #3 (190) | #3 | #11 (20.3 K) |
|  AS41327 | Fiber Telecom S.p.A. | #2 (853) | #4 (161) | #113 | #182 (94) |
|  AS28716 | Retelit Digital Services S.p.A. | #4 (660) | #5 (134) | #58 | #44 (1.5 K) |
|  AS1267 | WIND TRE S.P.A. | #10 (222) | #6 (94) | #1 | #21 (5.1 K) |
|  AS12779 | IT.Gate S.p.A. | #11 (180) | #7 (73) | #348 | #14 (13.1 K) |

Can be ranked by Global or ASN Country using:

- Peer Count (*)
- AS Cone
- Eyeball Population
- Domain Records

* is improved by feeding bgp.tools BGP data

Core points

- Built out of the frustration I had with other tools
- 920~ BGP sessions established
- Practically realtime BGP peer updates
- The horrors of WHOIS is handled, and in some cases is updated in near real time
- Frequently updated (~14 days):
 - ICMP Ping data scans of IPv4 /0
 - IPv4 and IPv6 RDNS data
 - Forward DNS data (Looking what A records point to a prefix)
- Peering IXP data is provided (MAC address vendors, ping data etc)

Setting up feeds is easy

Go to (PeeringDB SSO is supported):

<https://bgp.tools/kb/setup-sessions>

You can **instantly** setup eBGP MultiHop Sessions to bgp.tools. Where you **should** export a full table.

Export to 3rd parties/Looking Glass visibility is entirely optional!



bgp.tools

New BGP Session:

Description for Router/Session: (max 16 chars)

LHR01

Select the ASN you would like us to use for you. We will only accept [AS212232 \(bgp.tools\)](#), AS206924 AS212232, and Private ASN ranges

212232

Select the ASN you are going to use with us. We will only accept AS206924 AS212232 and Private ASN ranges

206924

Select the IP you will be connecting from.

192.0.0.1 / 2001:db8::

You will get the remote (bgp.tools side) IP after you create the session.

Please send **Full tables** rather than just your peering routes/customer routes. bgp.tools may automatically switch your sessions to only import your peering routes to save RAM, but allow us to figure that out for future flexibility!

We support (and encourage) BGP AddPath, and MultiProtocol/MultiFamily BGP

If you absolutely need a MD5 Password on the session, please enter the desired MD5 password

Export this data into publicly available MRT files (also enables the public looking glass)

Also allow commercial products to use those MRT files

Send notifications if session is down for more than 2 hours

Create BGP Session

Possibly interesting recent improvements to the site

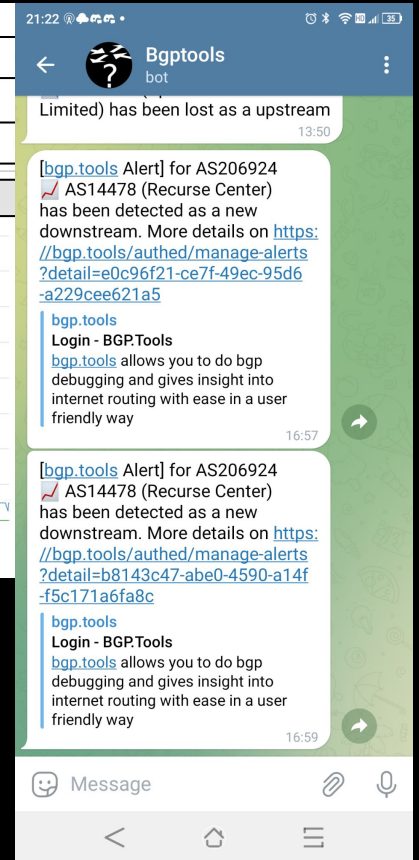

- The site now has a place to document BGP community values
- AS-SET decoding has been implemented
- Client side Agent software is available for testing (for traceroutes etc)
- RPKI ASPA support has been added
- Historical Graph data is being recorded for monitoring users
- Table dumps (not MRTs) are being produced for CAIDA
- "Super" looking glass is now live, on the website and over SSH
- IX-F feeds are used over PeeringDB where available

Monitoring services

- Goal is to fund the site (and myself) by providing nice paid tools like:
- Rapid BGP/RPKI/IRRdb monitoring
- Managed looking glasses
- BGP Session data recording (and replay for incident review)
- Logo placement on Looking Glass



bgp.tools



d6253275ab09f7afeaca74955a6cd164

ANYWAY

Sorry didn't mean to do a sales pitch



bgp.tools

d6253275ab09f7afeaca74955a6cd164

Peering RIS/Routeviews style

Peering LAN first

- RIS/RV collectors live on IXPs
- Kind of a mess of IXP peers just sending their customer cone, not their full table
- Previously (Maybe still currently?) RIS has been limiting new sessions due to capacity concerns?

| Name | Physical Location | Type | Scope | Raw Data |
|-------|-------------------|----------|---------------|----------------------|
| RRC00 | Amsterdam, NL | multihop | global | data |
| RRC01 | London, GB | IXP | LINX, LONAP | data |
| RRC03 | Amsterdam, NL | IXP | AMS-IX, NL-IX | data |
| RRC04 | Geneva, CH | IXP | CIXP | data |
| RRC05 | Vienna, AT | IXP | VIXP | data |
| RRC06 | Otemachi, JP | IXP | DIX-IE | data |
| RRC07 | Stockholm, SE | IXP | Netnod | data |
| RRC10 | Milan, IT | IXP | MIX | data |
| RRC11 | New York, NY, US | IXP | NYIIX | data |
| RRC12 | Frankfurt, DE | IXP | DE-CIX | data |
| RRC13 | Moscow, RU | IXP | MSK-IX | data |
| RRC14 | Palo Alto, CA, US | IXP | PAIX | data |
| RRC15 | Sao Paulo, BR | IXP | PTTMetro-SP | data |
| RRC16 | Miami, FL, US | IXP | Equinix Miami | data |
| RRC18 | Barcelona, ES | IXP | CATNIX | data |
| RRC19 | Johannesburg, ZA | IXP | NAP Africa JB | data |



But I can't deny IXP collection does not work

- RIPE has ~1535 BGP sessions online,
 - 372 / 407 Full IPv4/IPv6 tables
 - (by their own calculations)
 - Some of these sessions have issues about to be mentioned, **some are immensely useful views of the internet though!**

But I can't deny IXP collection does not work

- RIPE has ~1535 BGP sessions online,
 - 372 / 407 Full IPv4/IPv6 tables
 - (by their own calculations)
 - Some of these sessions have issues about to be mentioned, **some are immensely useful views of the internet though!**
- bgp.tools is 99% eBGP Multihop only
 - 910~ Sessions online
 - 610 / 983 Full IPv4/IPv6 tables (!!)

Problems with IXP Route Collection

- Really expensive if you don't have friends
 - IXP Membership fees + XC fees + colo fees
 - IXP membership alone can be more than the last two
 - <https://peering.exposed>
- **Huge** bias to AS6939
 - They are on almost all of the large IXPs, and provide you 180k+ of peered v4 routes that will likely be preferred over transit, hiding transit paths from the collector

Solving for XC Fees / Colo

- What is the cheapest, smallest, most insane thing we could ship to a *willing* IXP?



Solving for XC Fees / Colo

- What is the cheapest, smallest, most insane thing we could ship to a *willing* IXP?



bgp.tools

<https://blog.benjojo.co.uk/post/smart-sfp-linux-inside>

d6253275ab09f7afeaca74955a6cd164

Solving for XC Fees / Colo

- What is the cheapest, smallest, most insane thing we could ship to a *willing* IXP?



- No XC, The switch is the power supply, you can hitch backhaul either via someone friendly on the IXP, or relaying via a VPS or something
- Cheap, Around 150 USD all in
- Single core ARMv7, with 512M of RAM running Debian Jessie
- **Completely crazy.** Everyone is going to look at you like you lost the plot!
- Made by a Russian company who is totally now a Dubai company since the Russian Invasion of Ukraine



bgp.tools

<https://blog.benjojo.co.uk/post/smart-sfp-linux-inside>

d6253275ab09f7afeaca74955a6cd164

Looking in the right places



bgp.tools

d6253275ab09f7afeaca74955a6cd164

Looking in the right places



- 80~ USD from FS.com
- Runs a 400Mhz 32bit MIPS core, 64MB of RAM
- Literally running OpenWRT out of the packaging
- Also, here are the SSH credentials to all of them
 - user **ONTUSER**
 - password **7sp!1wUBz1**
- The constrained RAM and MIPS CPU μ Arch makes this a challenge to program for
- Thankfully Zig lang has a mostly working MIPS target!
- Can change the vendor to w/e with just ``sfp_i2c -i 0 -s "CISCO"```
- To use as a generic "Linux box" you must perform *some kernel changes*



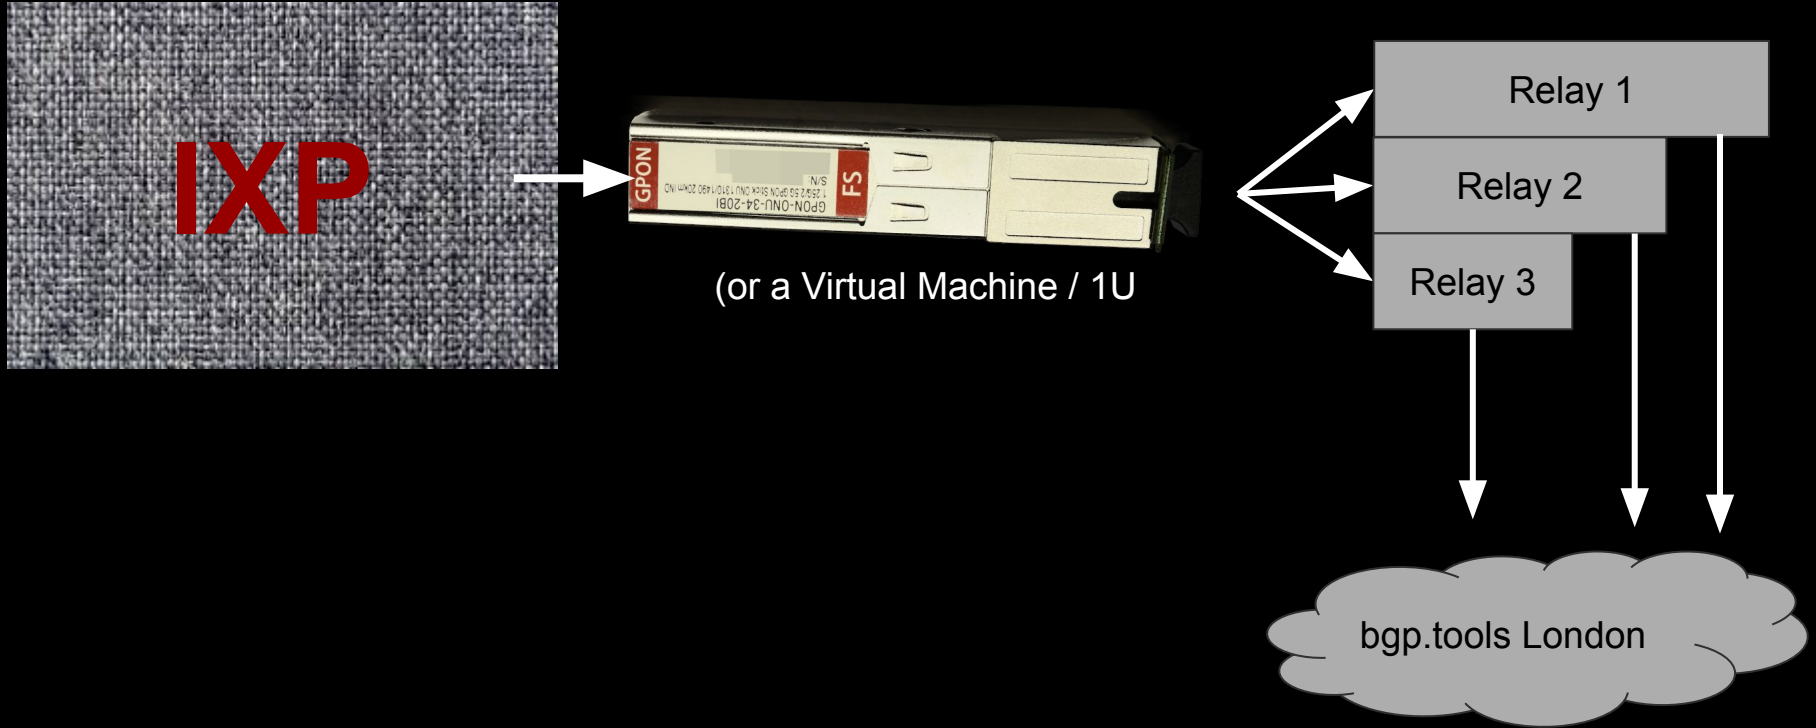
Dialing down the insanity

- A lot of IXPs have reseller programs, and with friends you can pick up very cheap ports and maybe a tagged vlan on a virtual machine.
- Downside is that even with this, a lot of the IXPs still require you to be a full member to be present on the LAN

All roads lead back to London

- You have have noticed it isn't really possible to store a *modern* full internet table on 64MB of RAM.
- Instead of storing sessions locally, the local collector will "rehost" the BGP session back in London where all of the infra is.
- This is because with how bgp.tools is designed, all BGP data has to be within 3ms~ of the web server to ensure a enjoyable experience

Overall



Current expanding plans

- NL-IX is up
 - South African IXP relays are in progress (JINX, CINX, DINX)
 - Got (friendly) contacts at AMS-IX / FranceIX? I'd love to know them
 - Looking for cheap ways to get to IX.BR
-
- Some IXPs are setting up eBGP multihop sessions from their route servers! Very nice of them!

Questions?

Want to feed bgp.tools?

go to bgp.tools and go to bottom link "Contribute Data"

More complex queries:

IRC: Benjojo-bgptools (terahertz) / benjojo (everything else)

Or email: admin@bgp.tools




bgp.tools

Bonus slides (just in case)

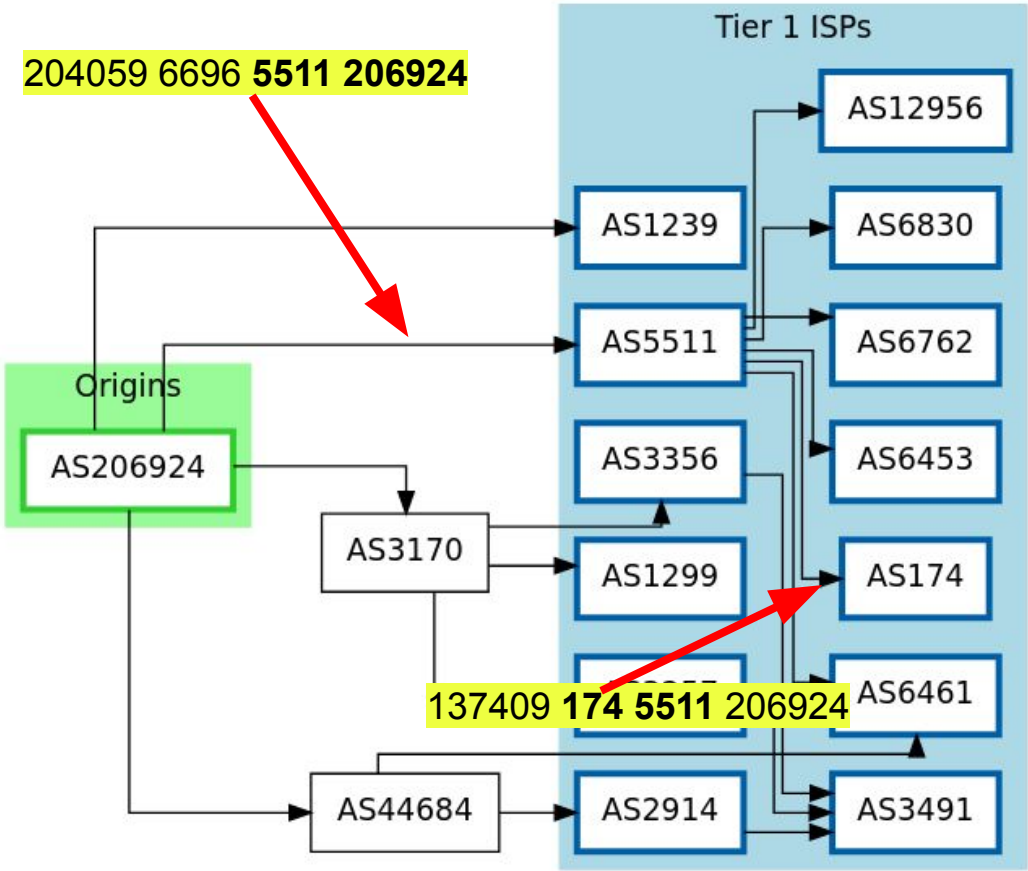
Determining the upstream

- You have to have loads of
- You have to detect and so
- "Upstreams" is not really who sends you eventually



Determining the upstream

- You have to have loads of data
- You have to detect and sort
- "Upstreams" is not really a concept, it's just who sends you eventually



EOF