

Getting creative with IXP route collection

Ben Cartwright-Cox - PeeringDays 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:1e61: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/as/202030


Start here...



Logged in as AS206924

View

Edit



Icelandair ehf

AS Number 202030

Website <http://www.icelandair.is>

Overview

Prefixes

Connectivity

Whois

IX

Registered on
29 Apr 2014 (8
years old)

Network status
Active, Allocated under RIPE

Network type
Unknown

Prefixes Announced
5 IPv4, 0 IPv6

Upstreams

- [AS44735](#) - Nova hf
- [AS30818](#) - Advania Ísland ehf



162

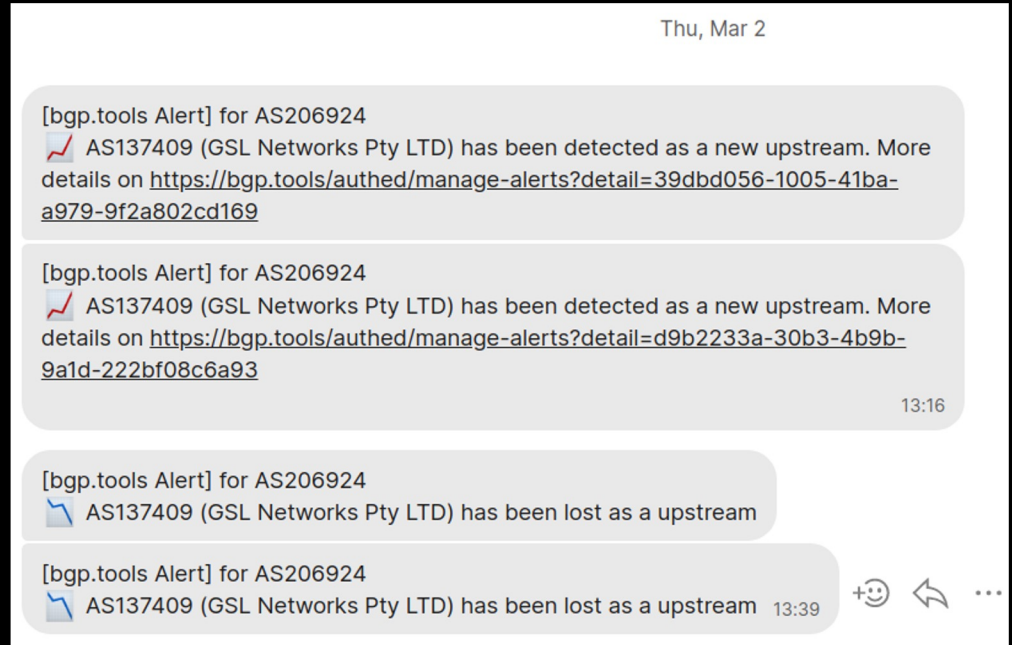
729e88e

Core points

- Built out of the frustration I had with other tools
- 910~ BGP sessions established
- Practically realtime BGP peer updates
- WHOIS is handled with more care, and in some cases is updated in near realtime
- 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)

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



ANYWAY

Sorry didn't mean to do a sales pitch



bgp.tools

16263485524785eb40804512729e88e

One of the core bits of today's adventure

Upstreams

	ASN	Description	IPv4	IPv6
	AS20473	The Constant Company, LLC	✗	✓
	AS5511	Orange S.A.	✓	✓
	AS6939	Hurricane Electric LLC	✗	✓
	AS3170	VeloxServ Communications Ltd	✓	✓
	AS44684	Mythic Beasts Ltd	✓	✓

Determining the upstreams of an ASN is a non trivial task

- You have to have loads of vantage points
- You have to detect and scrub route optimisers who lie to you

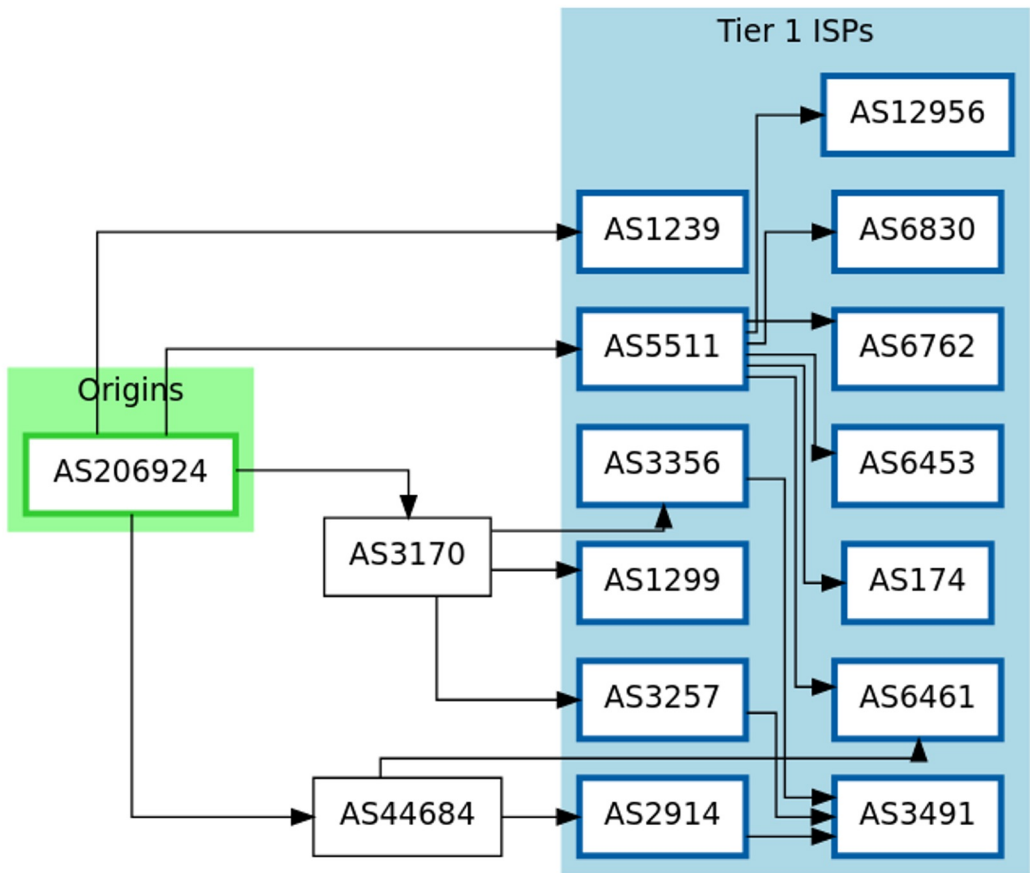
Determining the upstreams of an ASN is a non trivial task

- You have to have loads of vantage points
- You have to detect and scrub route optimisers who lie to you

- "Upstreams" is not really a truly well defined concept, but I view it as someone who sends you eventually up to a T1 network

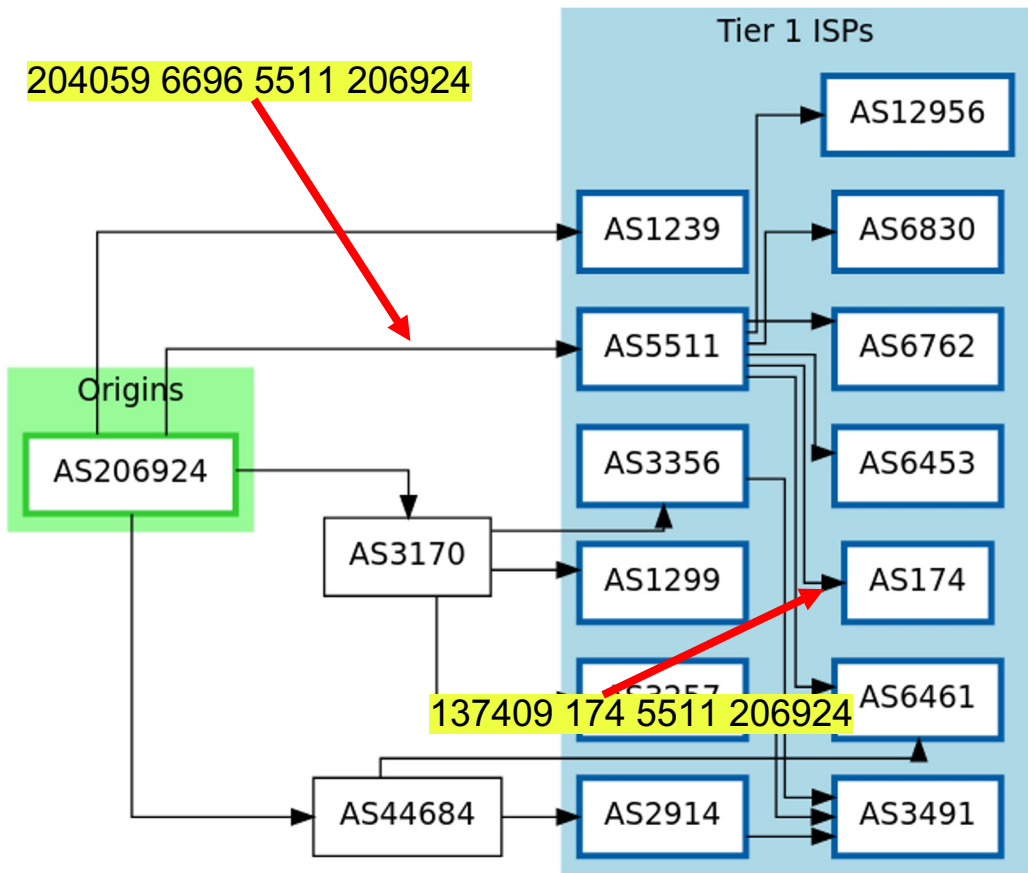
Determining the upstream

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



Determining the upstream

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



Determining the upstreams of an ASN is a non trivial task

- You have to have loads of vantage points
- You have to detect and scrub route optimisers who lie to you

- "Upstreams" is not really a truly well defined concept, but I view it as someone who sends you eventually up to a T1 network

Determining the upstreams of an ASN is a non trivial task

- You have to have loads of vantage points
- You have to detect and scrub route optimisers who lie to you

- "Upstreams" is not really a truly well defined concept, but I view it as someone who sends you eventually up to a T1 network

- It is not clear that this is a good definition in some regions
- ISPs like Zayo and Sprint peer a *lot* on Route Servers, so you have to apply logic to them to avoid everyone on a IXP having Zayo and Sprint as a upstream

The two problems this talk will explore

You have to have
loads of vantage points

It is not clear that this
is a good definition in
some regions

What is upstreaming?

Current status quo

- Bgp.tools considers any path that crosses a T1 a path where upstreams can be learned.
- Current calculation goes:
 - 917 137409 174 5511 206924 112
 - 917 -> 137409 -> 174 is ignored, since the feed could easily be lying about this to juice their upstream stats (or incompetence)
 - 174->5511 are considered peers, since they are T1s
 - 5511->206924 results in 206924 having 5511 as a upstream (and 206924 being a downstream of 5511)
 - 206924->112 results in the same effect, 112 as a upstream of 206924
 - All of this with a single BGP ASPATH!

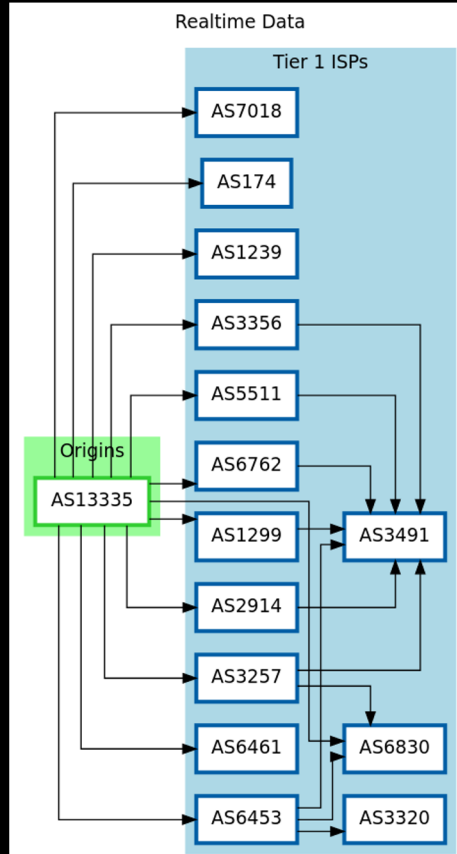
Current status quo (problems #1)

- This does not consider that you may be a peer of the T1.
 - If I see **137409 174 206924** It could (but very unlikely) be that AS206924 is a peer of cogent
 - This is more likely with (A) Zayo, (B) Sprint since (A) peers a lot, (B) imports peers on the route servers of IXPs
 - Though sprint won't be a problem for too much longer I guess...
 - Zayo and Sprint have a requirement of a "double hop", meaning we **must** observe two T1 ASN hops before considering it a upstream route
 - Since that would imply that route is in their customer cone
 - This trick works well, but it does not scale to all T1's because...

Current status quo (problems #2)

My code

AT&T, Liberty Global, ✨, 🍊, TATA, GTT, NTT, Lumen, Cogent and 🦁 are upstreams



My (own brain) interpretation

AT&T might be peering???
Cogent is almost defo transit
Sprint, Zayo are likely peering

🍊, Lumen, ✨, 🦁, NTT, GTT, TATA are all double hops. Assumed upstreams

AT&T, Liberty are likely not upstreams



New idea

- A upstream is someone who (*with consent*) carries your routes to their peers/upstreams
- Basically, if it's in your cone, you are a upstream for them

- The basic way of implementing this is to keep doing the T1 logic, but also import in IXP route servers, since they are a great insight into the cone of networks in a pure-ish way

- Downside: You must get feeds from loads of different IXP route servers

The two problems this talk will explore

You have to have
loads of vantage points

It is not clear that this
is a good definition in
some regions

RIS and Routeviews

Peering LAN first

- Most 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 ~1500 BGP sessions online,
 - 374 / 403 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 eBGP Multihop only
 - 910~ Sessions online
 - 522 / 800 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>

16263485524785eb40804512729e88e

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 deranged.** 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>

16263485524785eb40804512729e88e

Looking in the right places



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"```

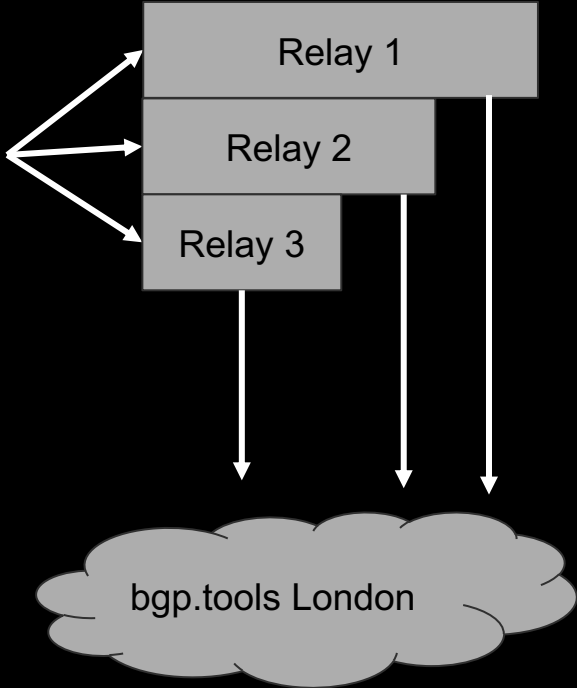

Dialing down the insanity

- 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 1ms~ of the web server to ensure a enjoyable experience

Overall



Current expanding plans

- NL-IX is coming up via a reseller port
- African IXP relays are slowly being figured out via two different carriers
- Got contacts at AMS-IX / DE-CIX / FranceIX / MIX? I'd love to know them
- IX.BR is home to a lot of interesting networks that don't wish to eBGP multihop peer, I am slowly working out a reseller port setup there!
 - Same with some African networks

- 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)

Email: admin@bgp.tools

