Piranha BGP

where all the fishy routes meet...

Failure Detection in BGP Networks
by Pascal Gloor

Agenda

PiranhaBGP

Agenda

PiranhaBGP

- What it is
- What it does
- What it might do
- What I need
- What you want to know

What is PiranhaBGP?

PiranhaBGP

What is PiranhaBGP

PiranhaBGP

- Route collector
  - Written in C
  - Supports address-family IPv4/IPv6
  - Attributes: next hop, community, ext community, large community, as path, origin
  - Fast dump into files, per peer, (typical file rotation, 1min)
  - Dump decoder (human readable, machine readable or JSON)

What is PiranhaBGP

PiranhaBGP

- Piranha Injector
  - Updates a database
  - State of routing table
  - Last n updates
  - (Some stats)

What is PiranhaBGP

PiranhaBGP

- Piranha UI
  - Peer Status
  - Route lookup(s)
  - Top n of:
    - Flapping Routes
    - Longest AS paths
    - Routes with invalid AS in AS paths (unallocated, private use)
    - Routes with invalid prefixes (unallocated, RFC1918, small prefixes >/24 and >/48)

What is PiranhaBGP?

PiranhaBGP

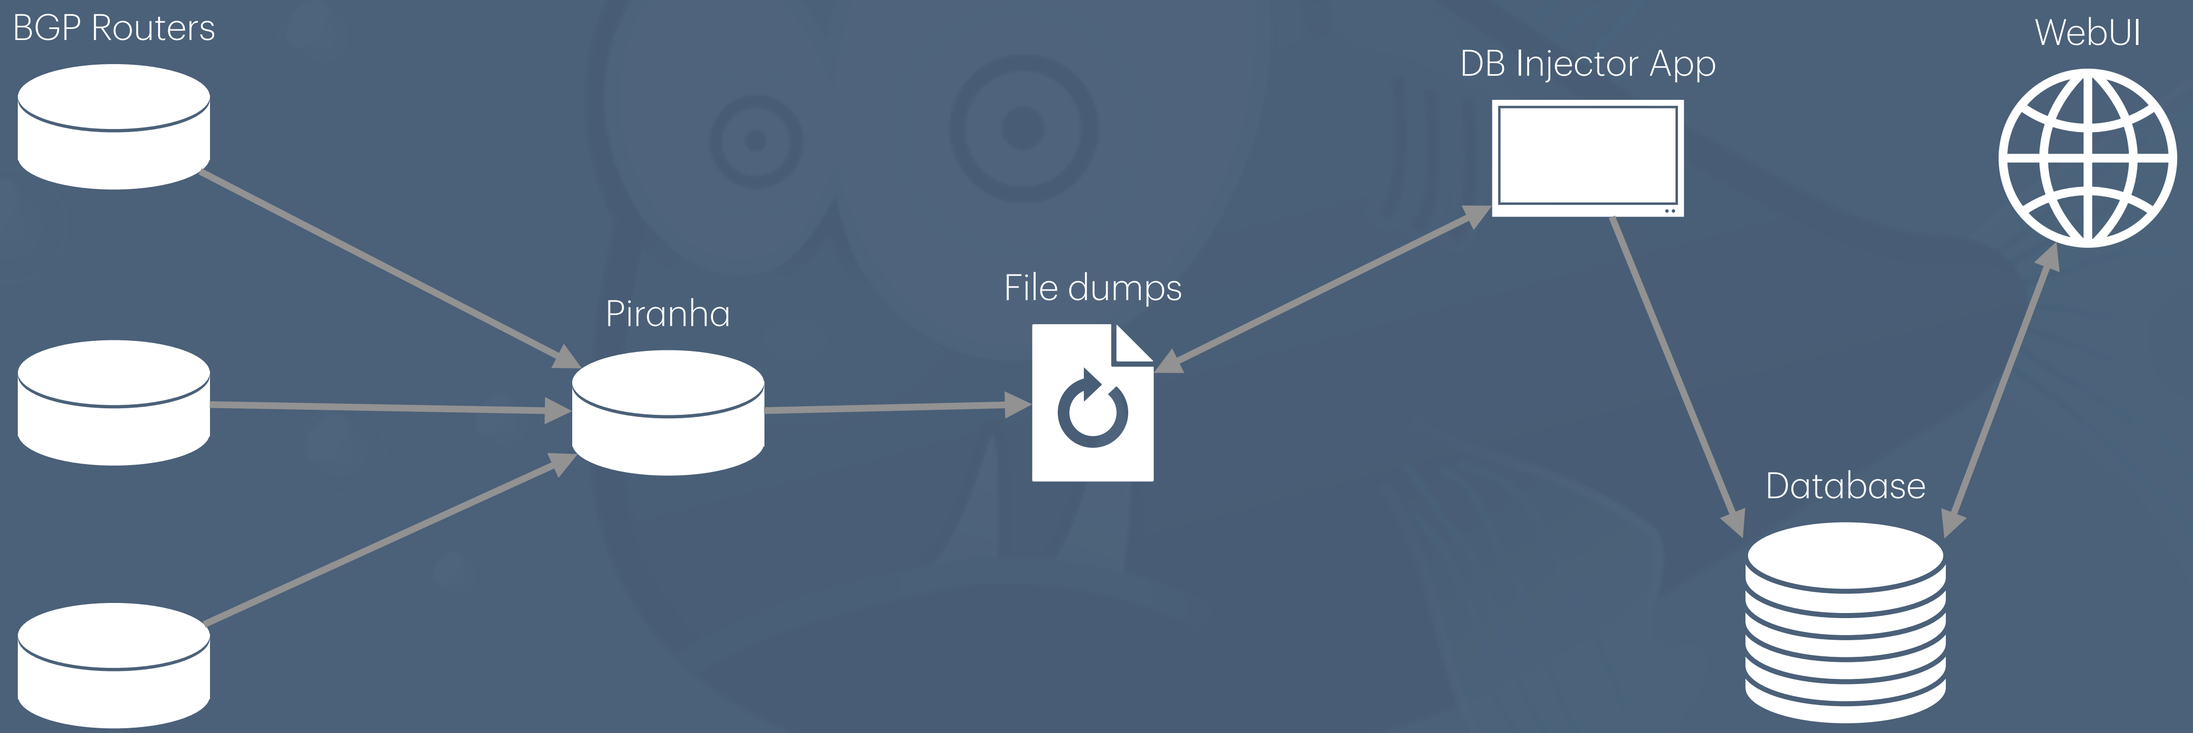

BGP Daemon

PiranhaBGP

BGP Daemon configuration

PiranhaBGP

```
# Piranha BGP Daemon configuration file
#
# [local_as]
# Local autonmous system number
local_as 65500

# [local_ip]
# Local IP Address to listen on.
# must be set in order to work.
# if you do not want to support ipv4 or ipv6
# comment the local_ipX out.
local_ip4 10.0.0.1
local_ip6 fe80::1

# [local_port] (default:179)
# Local port in which you want to listen().
local_port4 179
local_port6 179

# [export] (default: none)
# choose which route attributes to export
# in dump files
export origin
export aspath
#export community
#export extcommunity
#export largecommunity
```

```
# [export] (default: none)
# choose which route attributes to export
# in dump files
export origin
export aspath
#export community
#export extcommunity
#export largecommunity

# [bgp_router_id]
# BGP Router identifier, MUST be set to something else
# than 0.0.0.0 !

bgp_router_id 10.0.0.1

# [user]
user nobody

# [neighbor]
# neighbors/peers definition
# neighbor <ip4|ipv6> <ASN> [optional password]
neighbor 10.0.0.2 65500 MyPassword
```

BGP Daemon status

PiranhaBGP

```
# ./bin/piranhactl show
/-------------------------------------------------------------------------------\
| neighbor                                asn        recv       sent  updates  status  up/down |
|-------------------------------------------------------------------------------|
| 62.220.133.33                           6893     916272       3047  3622277      up     1d2h |
| 2001:788::33                            6893     444786       3047  1399050      up     1d2h |
| 62.220.133.32                           6893     950686       3051  3701618      up     1d2h |
| 2001:788::32                            6893     529475       3044  1784871      up     1d2h |
\-------------------------------------------------------------------------------/
```

BGP Daemon dumps

PiranhaBGP

```
# find var/dump/ -type f | sort | head -10
var/dump/2001:788::32/20241030145800
var/dump/2001:788::32/20241030145900
var/dump/2001:788::32/20241030150000
var/dump/2001:788::32/20241030150100
var/dump/2001:788::32/20241030150200
var/dump/2001:788::32/20241030150300
var/dump/2001:788::32/20241030150400
var/dump/2001:788::32/20241030150500
var/dump/2001:788::32/20241030150600
var/dump/2001:788::32/20241030150700
```

BGP Daemon decoder

PiranhaBGP

```
# ./bin/ptoa
Piranha v1.1.2 Dump file decoder, Copyright(c) 2004-2017 Pascal Gloor
syntax: ./bin/ptoa -<m|j|H> <file>

-H for human readable output

-j for JSON output
One JSON elem per line.

-m for machine readable output:
timestamp|P|peer_ip|peer_as # begin of every file
timestamp|C                 # connected (Active -> Established)
timestamp|D                 # disconnected (Established -> Active)
timestamp|K                 # BGP Keepalive received
timestamp|A|network|mask|opt id|opt|opt id ...
                            # BGP Announce
timestamp|W|network|mask    # BGP Withdrawn
```

BGP Daemon decoder (machine)

PiranhaBGP

```
# ./bin/ptoa -m var/dump/2001:788::32/20241030145900
1730300340.880126|P|2001:788::32|6893|i
1730300342.984359|W|2605:9cc0:c00::|48
1730300342.984364|A|2a02:ac80:50::|48|O|I|NH|2a02:2178:3:6b::1|AP|
  29075 6939 25145|C|0:0 29075:6700
..
1730300400.456128|E
```

BGP Daemon decoder (human)

PiranhaBGP

```
# ./bin/ptoa -H var/dump/2001:788::32/20241030145800
2024-10-30 14:58:00.824 peer ip 2001:788::32 AS 6893 TYPE ibgp
2024-10-30 14:58:00.824 prefix announce 2406:840:e23a::/48 origin IGP
nexthop 2001:1700:2b00:4::1 aspath 6730 6939 15353 140915 community
6730:6200 6730:6210 6730:6212 6893:11301
2024-10-30 14:58:05.240 prefix withdrawn 2a06:de01:8cb::/48
..
2024-10-30 14:59:00.880 eof
```

BGP Daemon decoder (json)

PiranhaBGP


```
{
    "timestamp": 1730300340.880126,
    "type": "peer",
    "msg": {
        "peer": {
            "proto": "ipv6",
            "ip": "2001:788::32",
            "asn": 6893,
            "type": "ibgp"
        }
    }
},
{
    "timestamp": 1730300342.984359,
    "type": "withdrawn",
    "msg": {
        "prefix": "2605:9cc0:c00::/48"
    }
}
```

```
{
    "timestamp": 1730300342.984364,
    "type": "announce",
    "msg": {
        "prefix": "2a02:ac80:50::/48",
        "origin": "IGP",
        "nexthop": "2a02:2178:3:6b::1",
        "aspath": [
            29075,
            6939,
            25145
        ],
        "community": [
            "0:0",
            "29075:6700"
        ]
    }
}
```

Current Features

PiranhaBGP

Piranha UI

PiranhaBGP

- Fast access to routing table state across all your BGP speakers
- Lookup any route or AS (including withdrawn routes)
- RDAP Lookup (cached in the database)
- See “Top 100”
  - Longest AS paths
  - Flapping routes
  - Invalid ASN
  - Invalid prefixes

Peer Status

PiranhaBGP

Peer status

Peer Status

Name	lsn-prc-MX24-01
ASN	6893
IP	2001:788::33
Password	Yes
State	Up
Last connect	2024-10-29 15:54:56.941
Last disconnect	never
Last update	2024-10-30 14:58:30.791
Delay	1m14.256s
Valid routes	241384
Invalid routes	56665

Netmask distribution

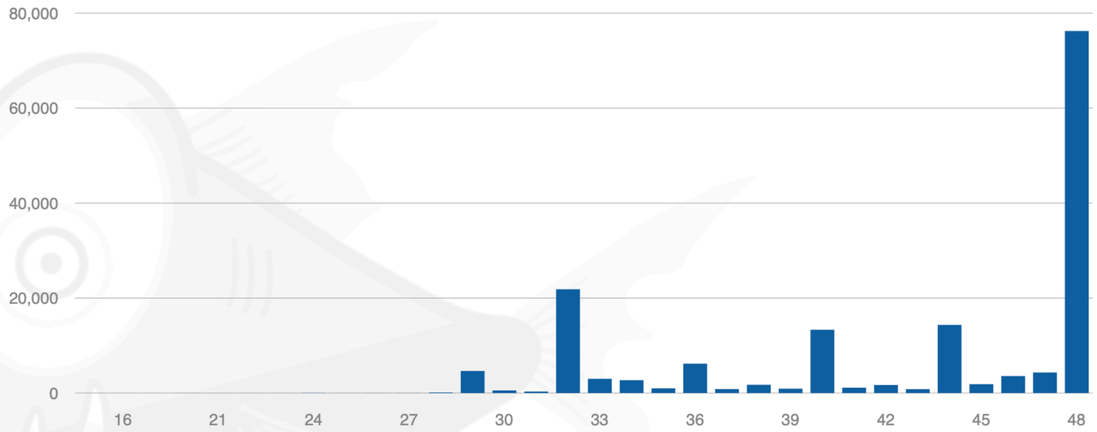

Peer Statistics (1/2)

PiranhaBGP

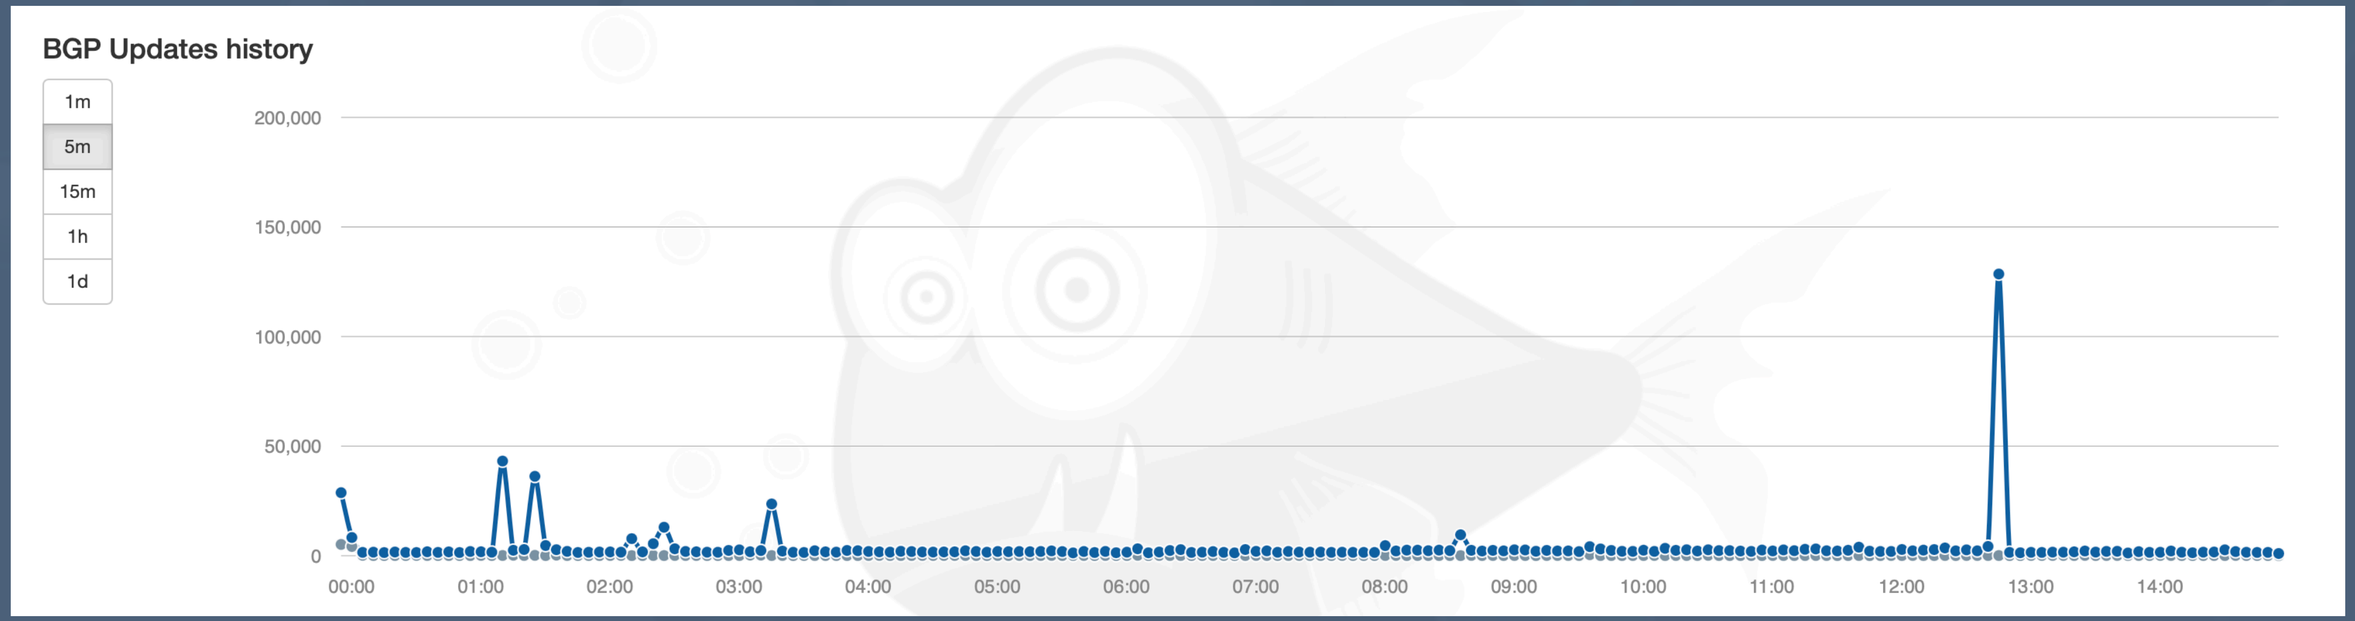

Peer Statistics (2/2)

PiranhaBGP

BGP Events

Timestamp	Event
2024-10-29 15:54	Connect

BGP Updates last 100 entries

The 100 most recent updates in the rolling buffer.

Display options: Peer IP | Peer ASN | Nexthop | Valid | Updated | Prefix | AS Path | Communities | Announce | Withdrawn

Valid	Updated	Prefix	AS Path	Communities
✓	2024-10-30 14:58:30.875	2404:3fc0:3::/48	6893 25091 9002 32590	6893:11303 25091:22416 65400:51706 65401:416
✓	2024-10-30 14:58:30.875	2404:4200::/32	6893 29075 1299 136106 24536	0:0 1299:37000 29075:7050
✓	2024-10-30 14:58:30.875	2401:2ce0::/33	6893 25091 6453 58453 58453 58453 134823 150881	6453:86 6453:3000 6453:3600 6453:3601 25091:15305
✓	2024-10-30 14:58:30.875	2401:43e0:2::/48	6893 25091 6762 17557 149021	6762:1 6762:30 6762:40 6762:13930 25091:15306
✗	2024-10-30 14:58:30.875	2404:76c0::/36		
✗	2024-10-30 14:58:30.875	2404:76c0:3000::/36		
✗	2024-10-30 14:58:30.875	2404:76c0:e000::/36		
✗	2024-10-30 14:58:30.875	2404:76c0:4000::/36		

Flapping routes (1/2)

PiranhaBGP

Top 100

Search options

Top 100 Flapping routes | All peers

IPv4 | IPv6

All | Only valid | Only invalid

Display options

Peer IP | Peer ASN | Nexthop | Valid | Updated | Prefix | AS Path | Communities | Announce | Withdrawn

Peer IP	Peer ASN	Valid	Updated	Prefix	AS Path	Announce(s)	Withdrawn(s)
62.220.133.32	6893	✖	2024-10-30 12:26:03.280	223.196.46.0/24	6893 25091 3257 174 55644 55644 55644 55644 55644 55644	3800	408
62.220.133.32	6893	✖	2024-10-30 12:45:48.631	223.196.68.0/24	6893 29075 174 55644	4059	128
62.220.133.32	6893	✖	2024-10-30 12:45:19.658	23.199.236.0/22	6893 174 55644	3887	249
62.220.133.32	6893	✖	2024-10-30 12:45:49.742	112.79.45.0/24	6893 25091 1273 55410 38266	3667	126
62.220.133.33	6893	✖	2024-10-30 12:45:21.990	223.196.68.0/24	6893 29075 174 55644	3467	128
62.220.133.33	6893	✖	2024-10-30 12:26:03.280	223.196.46.0/24	6893 25091 3257 174 55644 55644 55644 55644 55644 55644	3183	407

Flapping routes (2/2)

PiranhaBGP

Top 100

Search options

Top 100 Flapping routes | All peers

IPv4 | IPv6

All | Only valid | Only invalid

Display options

Peer IP | Peer ASN | Nexthop | Valid | Updated | Prefix | AS Path | Communities | Announce | Withdrawn

Peer IP	Peer ASN	Valid	Updated	Prefix	AS Path	Announce(s)	Withdrawn(s)
62.220.133.32	6893	✓	2024-10-30 12:45:49.738	203.145.74.0/24	6893 174 3491 10118 17794	2057	0
62.220.133.32	6893	✓	2024-10-30 12:45:49.738	203.145.78.0/24	6893 174 3491 10118 17794	2056	0
62.220.133.32	6893	✓	2024-10-30 12:45:48.629	202.45.88.0/24	6893 174 3491 10118 17794	2022	0
62.220.133.33	6893	✓	2024-10-30 12:45:21.991	203.145.74.0/24	6893 174 3491 10118 17794	1859	0
62.220.133.33	6893	✓	2024-10-30 12:45:21.991	203.145.78.0/24	6893 174 3491 10118 17794	1858	0

Longest AS paths

PiranhaBGP

Top 100

Search options

Top 100 Longest AS paths ▾ | All peers ▾

IPv4 | IPv6 — All | Only valid | Only invalid

Display options

Peer IP | Peer ASN | Nexthop | Valid | Updated | Prefix | AS Path | Communities | Announce | Withdrawn

Peer IP	Peer ASN	Valid	Updated	Prefix	AS Path	Announce(s)	Withdrawn(s)
62.220.133.33	6893	✔	2024-10-29 15:55:22.860	170.84.40.0/22	6893 174 16735 16735 16735 265020	1	0
62.220.133.32	6893	✔	2024-10-29 15:55:23.828	170.84.40.0/22	6893 174 16735 16735 16735 265020	1	0
62.220.133.33	6893	✔	2024-10-30 01:13:29.586	103.235.66.0/23	6893 29075 1299 38158 55701 63490 63490 63490 63490 63490 63490 63490 63490 63490 63490 63490 63490 63490 63490 63490	7	0

Invalid routes

PiranhaBGP

Top 100

Search options

Top 100 Invalid Global Prefixes | All peers

IPv4 | IPv6 | All | Only valid | Only invalid

Display options

Peer IP | Peer ASN | Nexthop | Valid | Updated | Prefix | AS Path | Communities | Announce | Withdrawn

Peer IP	Peer ASN	Valid	Updated	Prefix	AS Path	Announce(s)	Withdrawn(s)
2001:788::33	6893	✓	2024-10-29 15:54:59.314	::/0	6893	1	0
2001:788::32	6893	✓	2024-10-29 15:54:58.328	::/0	6893	1	0
2001:788::32	6893	✗	2024-10-29 23:12:41.272	2001:418:1401:4::/64	6893 25091 3257 20940	1	1
2001:788::32	6893	✗	2024-10-29 23:12:41.272	2001:418:1401:7::/64	6893 25091 3257 20940	1	1

Invalid ASN

PiranhaBGP

Top 100

Search options

Top 100 Invalid ASN | All peers

IPv4 | IPv6 | All | Only valid | Only invalid

Display options

Peer IP | Peer ASN | Nexthop | Valid | Updated | Prefix | AS Path | Communities | Announce | Withdrawn

Peer IP	Peer ASN	Valid	Updated	Prefix	AS Path	Announce(s)	Withdrawn(s)
2001:788::32	6893	✓	2024-10-29 15:54:58.220	2001:788::13/128	6893 65100 65313	1	0
2001:788::32	6893	✓	2024-10-29 15:54:58.220	2001:788::14/128	6893 65100 65314	1	0
2001:788::32	6893	✓	2024-10-29 15:54:58.220	2001:788::15/128	6893 65100 65315	1	0
2001:788::32	6893	✗	2024-10-30 08:05:14.246	2001:788::17/128	6893 65100 65104 65103 65111	2	1
2001:788::32	6893	✓	2024-10-30 08:05:14.246	2001:788::27/128	6893 65100 65104 65103	3	0
2001:788::32	6893	✓	2024-10-29 15:54:58.220	2001:788::39/128	6893 65100 65201	1	0
2001:788::32	6893	✗	2024-10-30 08:05:14.246	2001:788::45/128	6893 65100 65104	2	1

Lookup routes (1/2)

PiranhaBGP

Peer filter: All peers

Prefix search: 8.8.8.8

Origin AS search: 23456

Optional filter(s): Any | IPv4 | IPv6 — All | Only valid | Only invalid

Display options: Peer IP | Peer ASN | Nexthop | Valid | Updated | Prefix | AS Path | Communities | Announce | Withdrawn

Input help

`address`
Search for any prefix the IP belongs to, no matter the netmask.

`address/netmask`
Search for **exact match** prefixes.

`address/netmask-netmask`
Same as `address` but restrict the search to the range of given netmasks.

Peer IP	Peer ASN	Valid	Updated	Prefix	AS Path	Communities	Announce(s)	Withdrawn(s)
62.220.133.33	6893	✓	2024-10-30 00:00:22.596	8.0.0.0/9	6893 25091 3356	3356:0 3356:2 3356:100 3356:123 3356:501 3356:2065 6893:11303 25091:11301	2	0
62.220.133.32	6893	✓	2024-10-30 00:00:50.697	8.0.0.0/9	6893 25091 3356	3356:0 3356:2 3356:100 3356:123 3356:501 3356:2065 6893:11303 25091:11301	2	0
62.220.133.33	6893	✓	2024-10-30 00:00:22.596	8.0.0.0/12	6893 25091 3356	3356:0 3356:2 3356:100 3356:123 3356:501 3356:2065 6893:11303 25091:11301	2	0
62.220.133.32	6893	✓	2024-10-30 00:00:50.697	8.0.0.0/12	6893 25091 3356	3356:0 3356:2 3356:100 3356:123 3356:501 3356:2065 6893:11303 25091:11301	2	0
62.220.133.33	6893	✓	2024-10-30 12:44:53.743	8.8.8.0/24	6893 25091 15169	6893:11303 25091:21403	2	0
62.220.133.32	6893	✓	2024-10-30 12:45:01.891	8.8.8.0/24	6893 25091 15169	6893:11303 25091:21403	2	0

Lookup routes (2/2)

PiranhaBGP

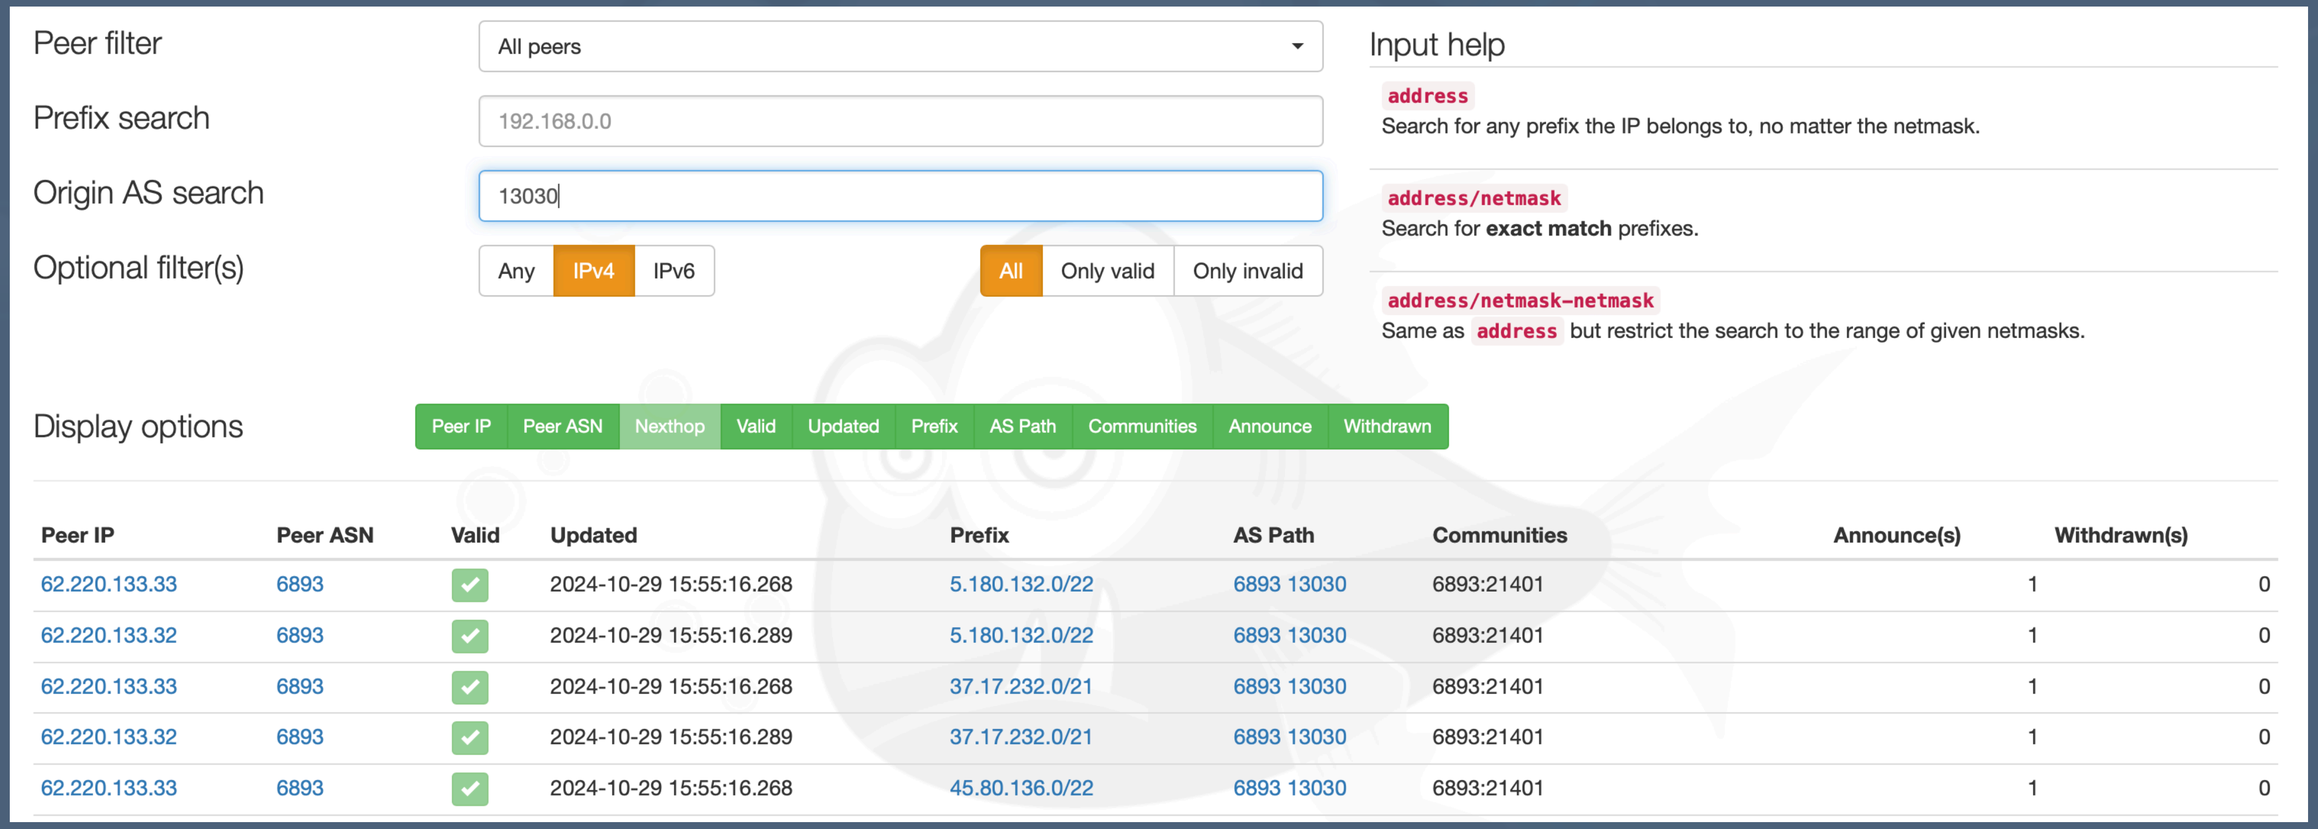

Peer IP	Peer ASN	Valid	Updated	Prefix	AS Path	Communities	Announce(s)	Withdrawn(s)
62.220.133.33	6893	✓	2024-10-29 15:55:16.268	5.180.132.0/22	6893 13030	6893:21401	1	0
62.220.133.32	6893	✓	2024-10-29 15:55:16.289	5.180.132.0/22	6893 13030	6893:21401	1	0
62.220.133.33	6893	✓	2024-10-29 15:55:16.268	37.17.232.0/21	6893 13030	6893:21401	1	0
62.220.133.32	6893	✓	2024-10-29 15:55:16.289	37.17.232.0/21	6893 13030	6893:21401	1	0
62.220.133.33	6893	✓	2024-10-29 15:55:16.268	45.80.136.0/22	6893 13030	6893:21401	1	0

RDAP Lookup (whois)

PiranhaBGP

2001:788::32	6893	✓	2024-10-30 14:58:34.441	2405:7e40:f910::/48	6893 25091 6762 58717 137515 137515 65901
2001:788::32	6893	✓	2024-10-30 14:58:34.441	2405:7e40:f911::/48	6893 25091 6762 58717 137515 137515 65901
2001:788::32	6893	✓	2024-10-30 14:58:44.147	2800:2b03:100::/40	6893 25091 6762 7303 64500 64501
2001:788::32	6893	✗	2024-10-30 14:03:40.403	2803:7710:2515::/48	6893 25091 6762 174 23520 272112 272110 65551
2001:788::32	6893	✓	2024-10-30 14:58:47.488	2804:2438:a000::/48	6893 25091 6762 53013 53158 264209 64500
2001:788::32	6893	✓	2024-10-29 15:54:59.304	2a0c:a641::/32	6893 25091 3257 8781 42298 65540

✖ ‹ 1/2 › 2800:2b03:100::/40 Telecom Argentina S.A. Piso 11 Dorrego Buenos Aires 1425 AR +54 1149684975#0000

Future Features

PiranhaBGP

Future features

PiranhaBGP

- Detect IPv4/IPv6 AF routing inconsistencies (loops, ghost routes)
- RPKI validation
- Collect VPNv4/VPNv6 routes -> searchable & detect inconsistencies
- Collect EVPN routes -> searchable & detect inconsistencies

Caveats

PiranhaBGP

- MySQL was a poor choice (especially for IPv6 storage), must be replaced (PgSQL?)
- extcommunities and largecommunities supported by BGP daemon but not by database/UI

What's next?

PiranhaBGP

What's next?

PiranhaBGP

- Is it worth it?
- Test deployments
- Need more contributors

Q&A

PiranhaBGP

PiranhaBGP

Q&A

Questions?

https://piranhabgp.net