

BGP Security

Hijack and Route Leak Detection

Lefteris Manassakis | COO, Code BGP

✉ lefteris@codebgp.com

ITNOG7

May 9-10, Bologna



About me




Lefteris Manassakis

COO & co-founder | Code BGP

✉ leftieris@codebgp.com

🌐 <https://manassakis.net/>

⚠️ BGP hijacks, leaks & misconfigurations affect your network



- BGP events critically affect **reliability, security, and performance**
- Only the **tip of the iceberg** gets known

Types of BGP prefix hijacks

- **Classification by Announced AS-Path**
 - **Origin-AS (or Type-0):** The hijacker AS announces – as its own – a prefix that it is not authorized to originate. This is the most commonly observed hijack type.
 - **Type-N ($N \geq 1$):** The hijacker AS announces an illegitimate path for a prefix it does not own. The announced path contains the ASN of the victim (first AS in the path) and hijacker, e.g., {AS50414, ASx, ASy, AS1 – 212.46.55.0/24}, while the sequence of ASes in the path is not a valid route, e.g., AS50414 is not an actual neighbor of ASx.

Types of BGP prefix hijacks

- **Classification by Affected Prefix**

- **Exact Prefix Hijacking:** The hijacker announces a path for exactly the same prefix announced by the legitimate AS. Since shortest AS-paths are typically preferred, only a part of the Internet that is close to the hijacker (e.g., in terms of AS hops) switches to route towards the hijacker.
- **Sub-Prefix Hijacking:** The hijacker AS announces a more specific prefix of the prefix of the legitimate AS. Since the more specific prefixes are preferred, the entire Internet routes traffic towards the hijacker to reach the announced sub-prefix.
- **Squatting:** The hijacker AS announces a prefix owned but not (currently) announced by the owner AS.
- For a comprehensive prefix hijack taxonomy please check the [ARTEMIS paper](#).

Route leaks

- **Definition:** A route leak is the propagation of routing announcement(s) beyond their intended scope.

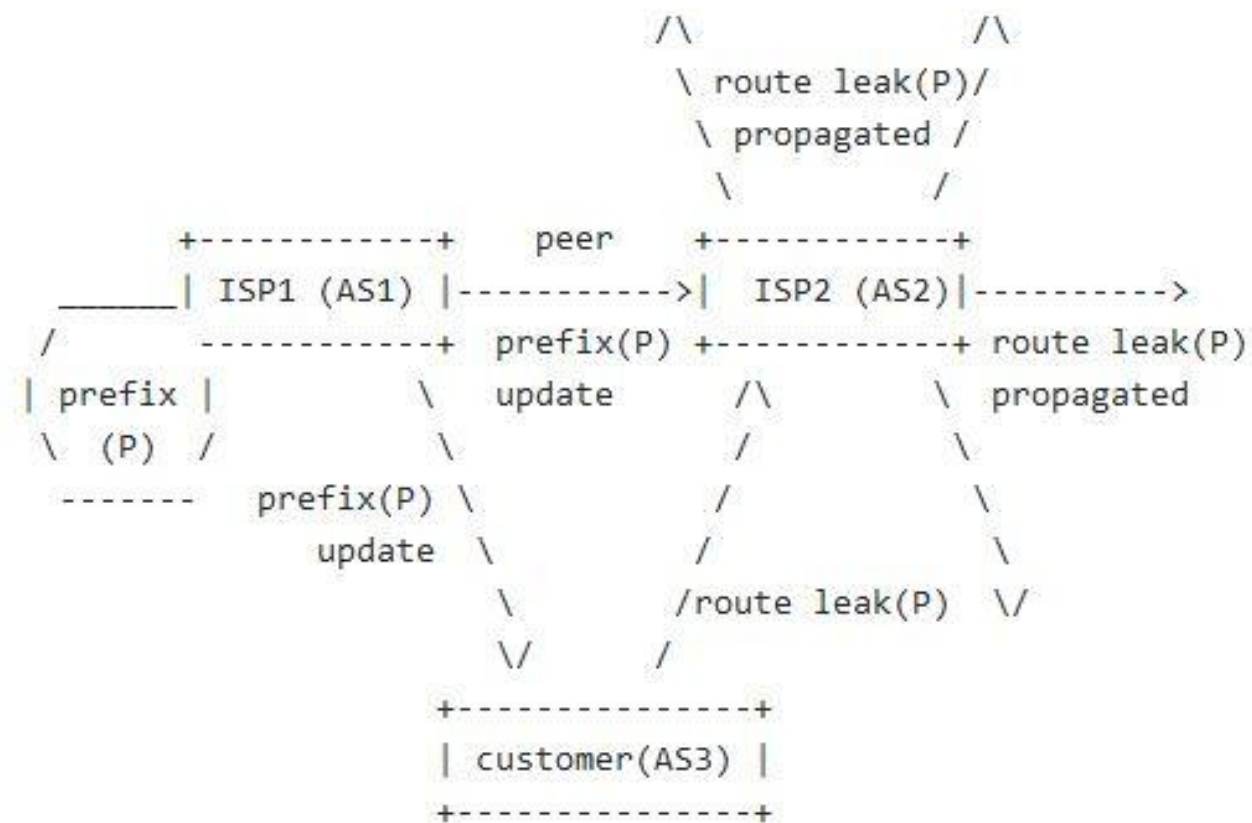



Figure 1: Basic Notion of a Route Leak

- For different types of route leaks please check [RFC 7908](#).


Challenges of hijack and route leak detection

- Speed
- Accuracy
- Evasion
- Privacy



Code BGP Platform

Monitor • Detect • Protect



Data service: Code BGP Monitor

BGP Monitoring Service developed by Code BGP

- Route Reflection ([RFC 4456](#))
- BGP Add-Path ([RFC 7911](#))
- 212 full feed peerings (v4 & v6)
- 71 cities, 44 countries, 23 upstreams



Data service: RIS Live

Provides real-time JSON BGP messages via a fully filterable interactive WebSocket JSON API, and a full stream ("firehose") containing all of the messages generated by RIS. → <https://ris-live.ripe.net/>

```
{
  "prefix": null,
  "path": 50414,
  "type": null,
  "require": null,
  "moreSpecific": true,
  "lessSpecific": false,
  "host": null (all),
  "peer": null,
  "socketOptions": {
    "includeRaw": false,
    "acknowledge": true
  }
}
```

```
// Received at 09:25:59 (3.31 second delay)
{
  "timestamp": 1662877556.6,
  "peer": "2001:7f8:30:0:1:1:0:6720",
  "peer_asn": "6720",
  "id": "05-7642-108395297",
  "host": "rrc05",
  "type": "UPDATE",
  "path": [6720, 8447, 20473, 50414],
  "community": [[1120, 1]],
  "origin": "igp",
  "announcements": [
    {
      "next_hop": "2001:7f8:30:0:1:1:0:6720",
      "prefixes": [
        "2a12:bc0::/48",
        "2a12:bc0:1::/48",
        "2a12:bc0:2::/48"
      ]
    },
    {
      "next_hop": "fe80::de8c:37ff:fe6f:f612",
      "prefixes": [
        "2a12:bc0::/48",
        "2a12:bc0:1::/48",
        "2a12:bc0:2::/48"
      ]
    }
  ]
}
```

Total peerings (IPv4 & IPv6):

1448

BGP full feeds:

- IPv4: **366**
- IPv6: **401**

Code examples

Below are simple examples of using the RIS Live WebSocket interface. For a full guide, see the RIS Live manual.

Javascript Python


```
/*
```

List of Route Collectors: https://ris.ripe.net/docs/10_routecollectors.html

List of Peers: <https://www.ris.ripe.net/peerlist/all.shtml>

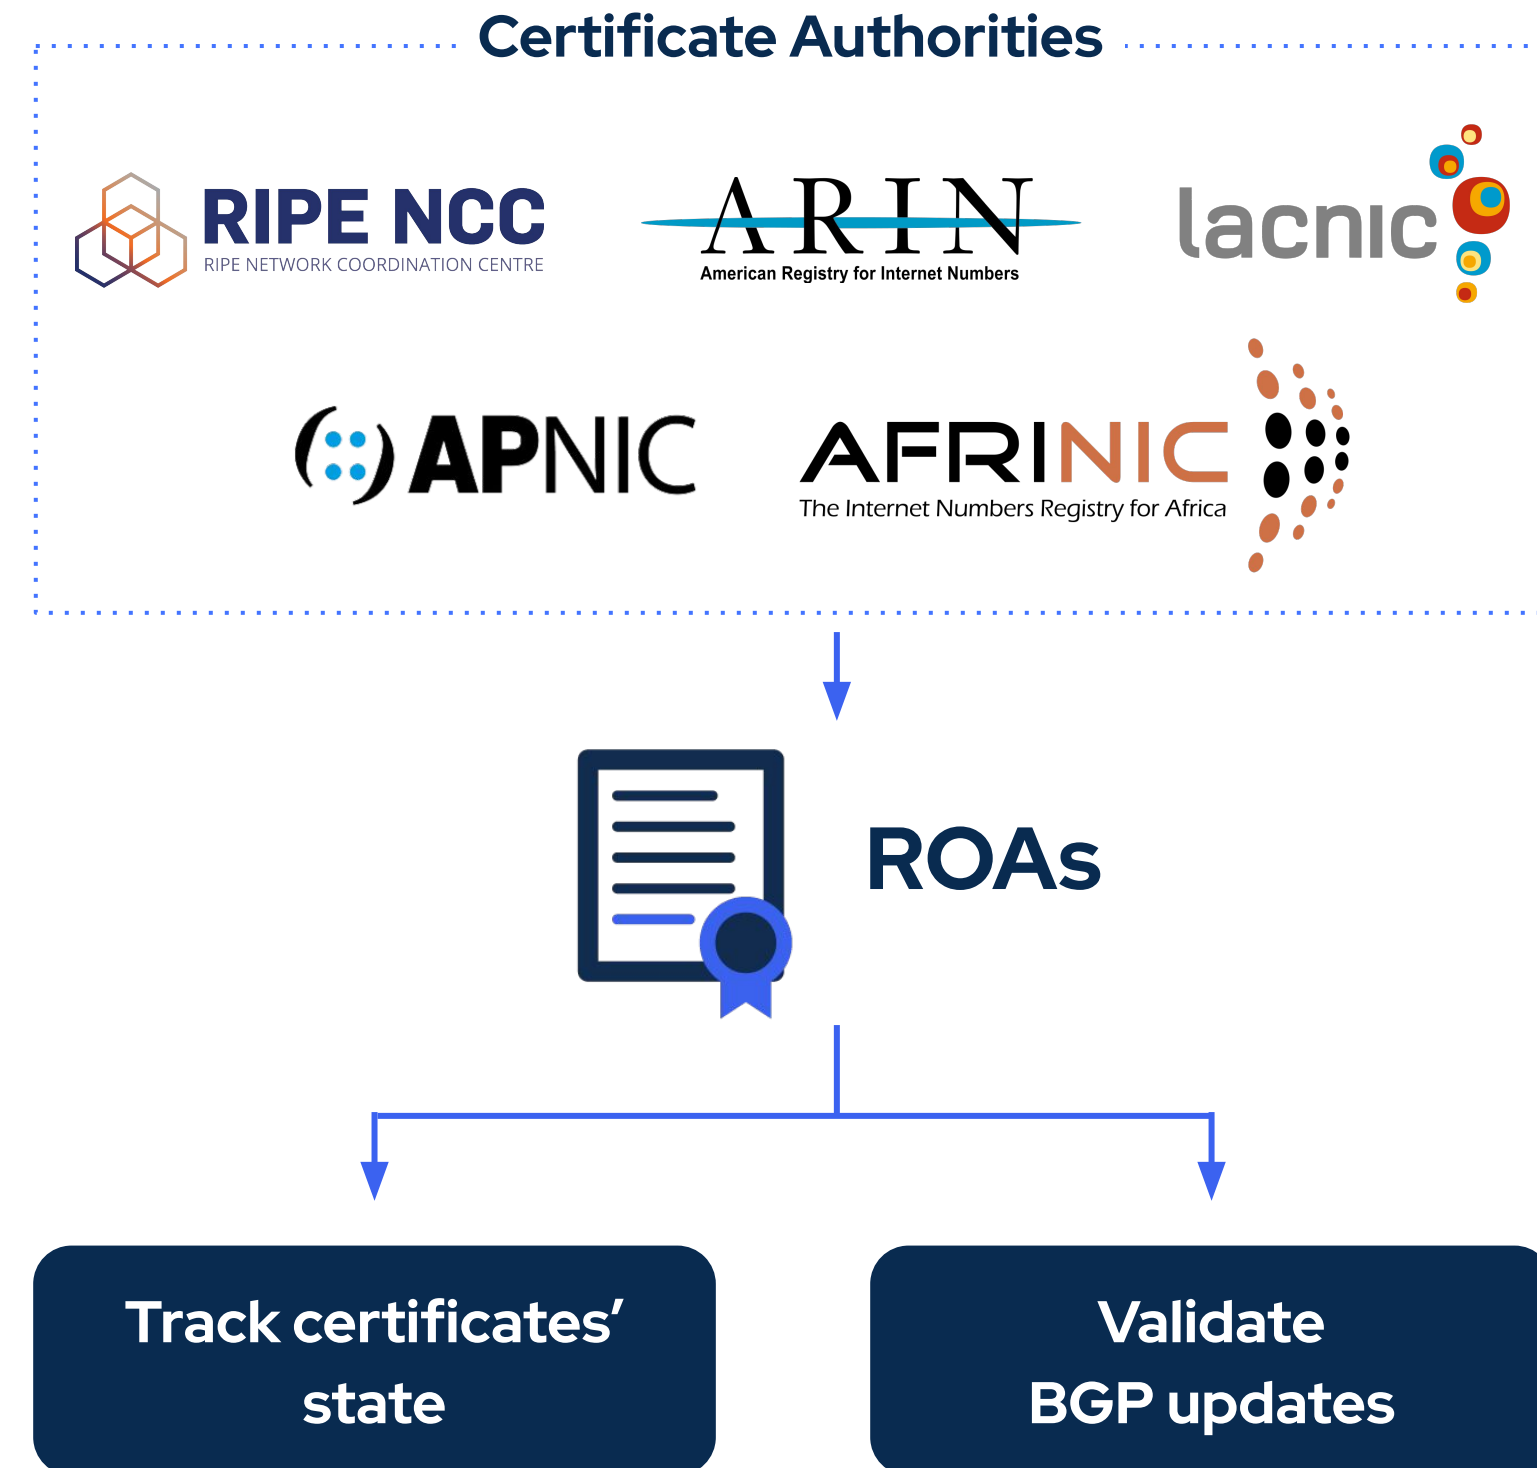
Data service: Your routers

- Multi-hop BGP sessions



Data service: RPKI

- Tracking the state of **ROA certificates**
- **Validating** BGP updates and detecting **invalids**



Alert types

Supported Alert Types	Description
Exact Prefix Hijack	Illegal origin ASes that announce configured prefixes.
Sub-Prefix Hijack	Illegal origin ASes that announce subprefixes of configured prefixes.
Route Leak	Unexpected prefixes in the list of prefixes that are announced by configured ASes.
New Neighbor	New neighbors that appear to peer with configured ASes. Possible AS path manipulation.
Neighbor Leak/Hijack	New neighbors that not only appear to peer with configured ASes, but also propagate their prefixes.
Squatting	Illegal origin ASes announcing prefixes that are not currently announced by configured ASes.
Presence in AS Path	Presence of ASes in paths towards configured prefixes.
Invalid AS Path Pattern	Violation of valid pattern by AS paths towards configured prefixes.
Long AS Path	Paths towards configured prefixes exceed a specified length threshold.
Prefix Visibility Loss	Visibility of prefix falls below a configured data source count threshold.
Peering Visibility Loss	Visibility of peering falls below a configured data source count threshold.

Supported Alert Types	Description
RPKI-Invalid Detection	RPKI-Invalid announcements of configured prefixes by other ASes.
RPKI-Invalid Announcement	RPKI-Invalid announcements by configured ASes.
RPKI-Invalid Propagation	RPKI-Invalid routes propagated by configured ASes.
RPKI-NotFound Propagation	RPKI-NotFound routes propagated by configured ASes.
Bogon (Exact-)Prefix	Announcements of bogon prefixes by configured ASes.
Bogon (Sub-)Prefix	Announcements of bogon subprefixes by configured ASes.
Bogon AS	In-path presence of bogon ASes, in routes towards configured prefixes.
AS Path Comparison	Discrepancies in AS paths towards the same prefix, comparing between different Data Services, up to a terminating (end) AS.
Prefix Comparison	Discrepancies in prefixes announced by configured ASes, comparing between different Data Services.
Custom	User-defined

Root DNS Servers

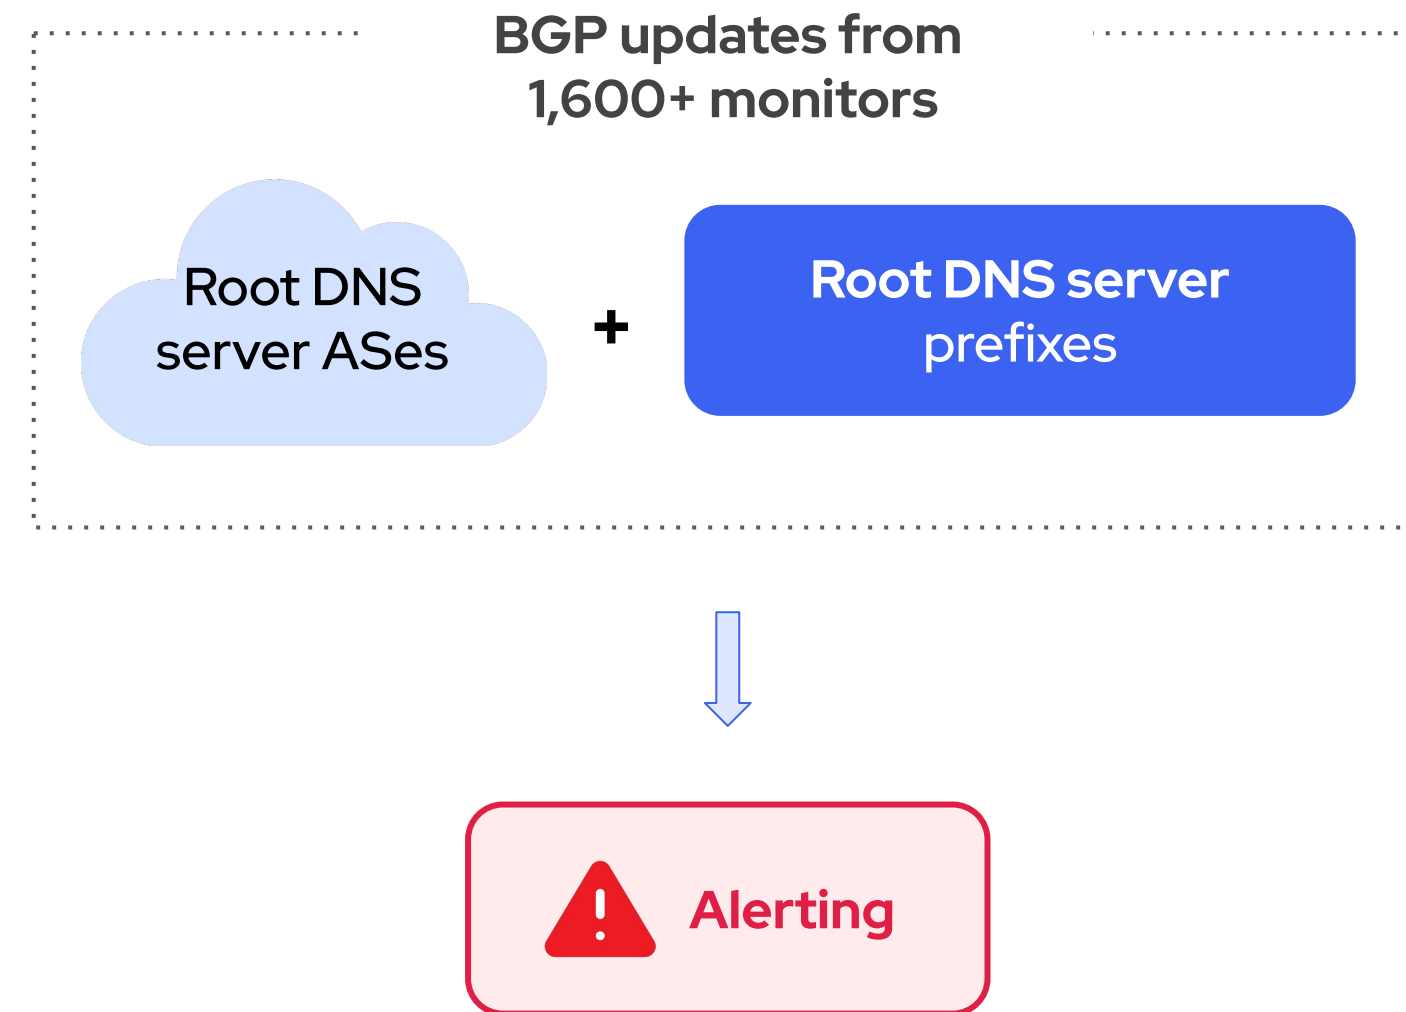
- The authoritative name servers that serve the DNS root zone

Name	IPv4	IPv6	Operator
A-Root	198.41.0.4	2001:503:ba3e::2:30	Verisign, Inc.
B-Root	199.9.14.201	2001:500:200::b	USC, Information Sciences Institute
C-Root	192.33.4.12	2001:500:2::c	Cogent Communications
D-Root	199.7.91.13	2001:500:2d::d	University of Maryland
E-Root	192.203.230.10	2001:500:a8::e	NASA (Ames Research Center)
F-Root	192.5.5.241	2001:500:2f::f	Internet Systems Consortium, Inc.
G-Root	192.112.36.4	2001:500:12::d0d	US Department of Defense (NIC)
H-Root	198.97.190.53	2001:500:1::53	US Army (Research Lab)
I-Root	192.36.148.17	2001:7fe::53	Netnod
J-Root	192.58.128.30	2001:503:c27::2:30	Verisign, Inc.
K-Root	193.0.14.129	2001:7fd::1	RIPE NCC
I-Root	199.7.83.42	2001:500:9f::42	ICANN
M-Root	202.12.27.33	2001:dc3::35	WIDE Project

Why Monitoring Root DNS Server Prefixes

- Critical Internet infrastructure, worth protecting
- These prefixes are heavily anycasted
 - BGP anomalies (e.g. exact prefix hijacks) will go largely unnoticed, due to their limited impact on the data plane

We provide access for free to a Code BGP Platform instance which monitors the root DNS prefixes



Suspicious route detected for root DNS prefix - Apr. 28

- AS 137661 announced prefix 199.7.83.0/24 which belongs to ICANN and is the IPv4 prefix of the "[L-Root](#)" domain server (AS 20144). Eight (8) days later the event is ongoing.
- Blog post: <https://www.codebgp.com/blog/suspicious-route-against-a-root-dns-prefix/>

The screenshot displays the Code BGP Platform interface. The top navigation bar includes the Code BGP logo, the user name 'Lefteris Manassakis' (editor | publicdemo), and a profile icon. The left sidebar contains navigation options: Overview, Setup, Looking Glass (selected), API, Alerts, and Dashboards. The main content area is titled 'Looking Glass' and has tabs for Prefixes, Autonomous Systems, Peerings, Routes (selected), and RPKI ROAs. A filter 'Origin AS: 137661' is applied. The main table shows a single route entry:

Prefix	AS Path	RPKI Status	First Detected ↓	Last Update
199.7.83.0/24	57695 60068 9498 9829 (7) 137661	NotFound	Apr 28, 2023, 14:15:06	Apr 28, 2023, 14:15:06

Below the main table, there is a section titled 'Data Sources of Route 199.7.83.0/24 - 57695 60068 9498 9829 (7) 137661'. This section contains a table with the following data:

Data Service	IP	ASN	City	Country	Continent	Last Update ↓
Code BGP Monitor	194.156.163.203	57695	Singapore	Singapore	Asia	Apr 28, 2023, 14:15:06

At the bottom of the interface, there are pagination controls showing 'Rows per page: 10' and '1-1 of 1'.

Low Visibility of Suspicious Route

- Route currently not visible by RIS Live. However, it is picked up by our platform and other LGs.

The screenshot shows the NLNOG Looking Glass interface. The search bar contains the IP or prefix '199.7.83.0/24' with 'Exact match' and 'on' selected. The search results show three entries for the route 199.7.83.0/24, all originating from AS 3214 via next-hop 185.255.55.19. The first entry is highlighted in red and labeled 'XTOM1-v4'. The second entry is labeled 'XTOM1-v4' and the third 'XTOM1-v4'. The interface shows various validation states and communities for each entry.

The screenshot shows the Alice - The friendly BGP looking interface. The search bar contains the IP or prefix '199.7.83.0/24'. The search results show a list of routes for the prefix 199.7.83.0/24, including the origin AS, next-hop IP, and the receiving AS. The routes are listed in a table format. The first entry is highlighted in red and labeled 'Invalid Origin-AS'. The second entry is labeled 'Invalid Prefix Object'. The interface also shows a network diagram on the right side.

Prefix	Next Hop	AS	Origin AS	Receiving AS
199.7.83.0/24	195.66.226.95	37468	Angola Cables	RS3.LON1 (IPv4)
199.7.83.0/24	195.66.225.60	5564	Brightsolid Online Technology Ltd.	RS3.LON1 (IPv4)
199.7.83.0/24	195.66.227.126	2018	Tenet	RS3.LON1 (IPv4)
199.7.83.0/24	195.66.225.145	35313	Infonias Ltd	RS3.LON1 (IPv4)
199.7.83.0/24	195.66.225.217	23028	Team Cymru, Inc.	RS3.LON1 (IPv4)
199.7.83.0/24	195.66.225.115	33920	aq Limited	RS3.LON1 (IPv4)
199.7.83.0/24	195.66.225.82	25091	IP-Max	RS3.LON1 (IPv4)
199.7.83.0/24	195.66.225.134	8473	Bahnhof	RS3.LON1 (IPv4)
199.7.83.0/24	195.66.226.41	37100	SEACOM	RS3.LON1 (IPv4)
199.7.83.0/24	195.66.224.205	3216	VimpelCom PJSC	RS3.LON1 (IPv4)
199.7.83.0/24	195.66.226.204	9498	Bharti Airtel Ltd	RS3.LON1 (IPv4)
199.7.83.0/24	195.66.224.237	200612	GBI HQ Cooperatief UA	RS3.LON1 (IPv4)
199.7.83.0/24	195.66.224.105	29668	vTream	RS3.LON1 (IPv4)
199.7.83.0/24	195.66.227.231	262589	Internexa	RS3.LON1 (IPv4)
199.7.83.0/24	195.66.227.22	37100	SEACOM	RS3.LON1 (IPv4)
199.7.83.0/24	195.66.226.204	9498	Bharti Airtel Ltd	RS1.LON1 (IPv4)
199.7.83.0/24	195.66.225.60	5564	Brightsolid Online Technology Ltd.	RS1.LON1 (IPv4)
199.7.83.0/24	195.66.224.105	29668	vTream	RS1.LON1 (IPv4)
199.7.83.0/24	195.66.225.134	8473	Bahnhof	RS1.LON1 (IPv4)
199.7.83.0/24	195.66.227.29	29119	Aire Networks del Mediterraneo S.L.U	RS1.LON1 (IPv4)

AS Path Prepending of Suspicious Route

- This route is the only one which is prepended so excessively (7 times!) by AS 9829

Looking Glass Info

Prefixes Autonomous Systems Peerings **Routes** RPKI ROAs

Origin AS: 137661

Prefix	Origin AS	Neighbor AS	AS Path	RPKI Status	First Detected ↓	Last Update
> 199.7.83.0/24	137661	9829	57695 60068 9498 9829 (7) 137661	NotFound	Apr 30, 2023, 15:02:02	Apr 30, 2023, 15:02:55
> 103.122.36.0/24	137661	55824	57695 137409 64049 55836 9885 55824 137661	Valid	Apr 30, 2023, 15:01:14	Apr 30, 2023, 15:01:43
> 103.122.36.0/24	137661	9829	57695 60068 6762 6453 9829 137661	Valid	Apr 30, 2023, 15:01:14	Apr 30, 2023, 15:01:43
> 103.122.36.0/24	137661	55824	6939 64049 55836 9885 55824 137661	Valid	Apr 30, 2023, 15:01:14	Apr 30, 2023, 15:01:43
> 103.122.36.0/24	137661	9829	57695 48024 137443 174 9829 137661	Valid	Apr 30, 2023, 15:01:14	Apr 30, 2023, 15:01:43
> 103.122.36.0/24	137661	55824	44684 64049 55836 9885 55824 137661	Valid	Apr 30, 2023, 15:01:14	Apr 30, 2023, 15:01:43
> 103.122.36.0/24	137661	55824	60011 3356 64049 55836 9885 55824 137661	Valid	Apr 30, 2023, 15:01:14	Apr 30, 2023, 15:01:43
> 103.122.36.0/24	137661	9829	61317 6762 6453 9829 137661	Valid	Apr 30, 2023, 15:01:14	Apr 30, 2023, 15:01:43
> 103.122.36.0/24	137661	55824	57695 9009 64049 55836 9885 55824 137661	Valid	Apr 30, 2023, 15:01:14	Apr 30, 2023, 15:01:43
> 103.122.36.0/24	137661	55824	42473 8529 64049 55836 9885 55824 137661	Valid	Apr 30, 2023, 15:01:14	Apr 30, 2023, 15:01:43

Rows per page: 10 51-60 of 379

Announcement of super-prefix

- The super-prefix 199.7.82.0/23 belongs to ICANN and is also being announced by AS137661!

The screenshot shows the Twelve99 Looking Glass interface. The browser address bar displays "lg.twelve99.net/?type=bgp&router=sngc-b5&address=199.7.82.0/23". The main content area displays BGP routing information for the prefix 199.7.82.0/23. The output shows multiple paths, with the path through AS137661 highlighted in blue. The highlighted path includes the following details: Origin IGP, metric 0, localpref 200, IGP metric 0, weight 0, tag 0; Received 5d17h ago, valid, external, ECMP head, ECMP, AS Origin not found; Communities: 1299:431 (RPKI state Unknown); and the path itself: 1299:1000 1299:37000 1299:37200 9498:1 9498:33 9498:91 9498:9829 34111:9498 34911:9498 40512:9498.

The screenshot shows the Alice - The friendly BGP looking interface. The browser address bar displays "alice-rs.linx.net/search?q=199.7.82.0/23". The main content area displays a table of BGP routes for the prefix 199.7.82.0/23. The table has columns for Prefix, Next Hop, AS, and Origin. The path through AS137661 is highlighted in blue. The highlighted path includes the following details: Prefix: 199.7.82.0/23, Next Hop: 195.66.226.204, AS: 9498, Origin: Invalid Origin-AS; Received 5d17h ago, valid, external, ECMP head, ECMP, AS Origin not found; Communities: 1299:431 (RPKI state Unknown); and the path itself: 1299:1000 1299:37000 1299:37200 9498:1 9498:33 9498:91 9498:9829 34111:9498 34911:9498 40512:9498.

Example of good routing hygiene

- DE-CIX also receives both routes, but does not propagate them

SEARCH ON ALL ROUTE SERVERS

199.7.83.0/24

Go to: **Filtered** Accepted

Found 157 received and 8 filtered routes.
Query took 7146.61 ms to complete.
Routes cache was built an hour ago and will be refreshed in 43 minutes.

ROUTES FILTERED Showing all of 8 routes

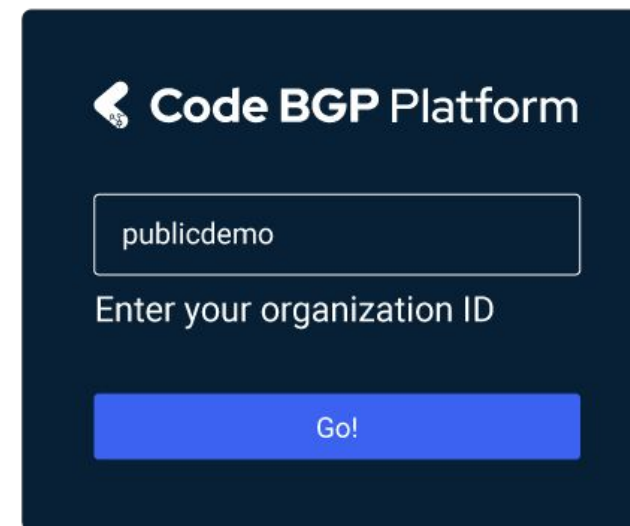
Network	Next-Hop	AS Path	ASN	Neighbor	RS
199.7.83.0/24 IRRDB lookup: Unable to resolve prefix for origin AS	80.81.194.250	9498 9829 9829 9829 9829 9829 9829	9498	Bharti Airtel Limited	rs1.fra.de-cix.net (IPv4)
199.7.83.0/24 IRRDB lookup: Unable to resolve prefix for origin AS	80.81.196.112	9498 9829 9829 9829 9829 9829 9829	9498	Bharti Airtel Limited	rs1.fra.de-cix.net (IPv4)
199.7.83.0/24 IRRDB lookup: AS-SET does not include origin AS	80.81.195.242	47734 41494 39107 20144	47734	Telecom IT Solutions	rs1.fra.de-cix.net (IPv4)
199.7.83.0/24 IRRDB lookup: Unable to resolve prefix for origin AS	185.1.47.7	9498 9829 9829 9829 9829 9829 9829	9498	Bharti Airtel Limited	rs2.mrs.de-cix.net (IPv4)
199.7.83.0/24 IRRDB lookup: Unable to resolve prefix for origin AS	185.1.47.7	9498 9829 9829 9829 9829 9829 9829	9498	Bharti Airtel Limited	rs1.mrs.de-cix.net (IPv4)
199.7.83.0/24 IRRDB lookup: Unable to resolve prefix for origin AS	80.81.194.250	9498 9829 9829 9829 9829 9829 9829	9498	Bharti Airtel Limited	rs2.fra.de-cix.net (IPv4)
199.7.83.0/24 IRRDB lookup: Unable to resolve prefix for origin AS	80.81.196.112	9498 9829 9829 9829 9829 9829 9829	9498	Bharti Airtel Limited	rs2.fra.de-cix.net (IPv4)
199.7.83.0/24 IRRDB lookup: AS-SET does not include origin AS	80.81.195.242	47734 41494 39107 20144	47734	Telecom IT Solutions	rs2.fra.de-cix.net (IPv4)

ROUTES ACCEPTED Showing all of 157 routes

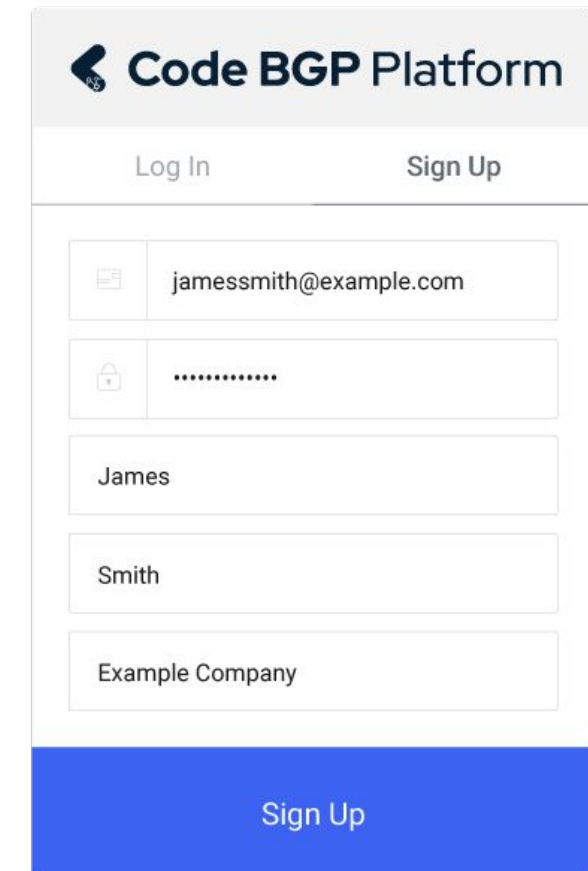
Network	Next-Hop	AS Path	ASN	Neighbor	RS
199.7.83.0/24	149.112.27.3	6939 20144	6939	Hurricane Electric	rs2.phx.de-cix.net (IPv4)
199.7.83.0/24	185.1.172.4	6939 20144	6939	Hurricane Electric	rs2.seecix.net (IPv4)
199.7.83.0/24	185.1.48.9	34984 20144	34984	Superonline Iletisim Hizmetleri A.S.	rs1.list.de-cix.net (IPv4)
199.7.83.0/24	185.1.48.16	6939 20144	6939	Hurricane Electric	rs1.list.de-cix.net (IPv4)
199.7.83.0/24	185.1.48.7	25091 20144	25091	IP-Max SA	rs1.list.de-cix.net (IPv4)
199.7.83.0/24	80.81.193.13	8218 20144	8218	Zayo Infrastructure France SA	rs1.fra.de-cix.net (IPv4)
199.7.83.0/24	80.81.194.197	6697 20144	6697	Beltelecom	rs1.fra.de-cix.net (IPv4)

How to get access to the Route DNS monitoring instance


- Go to <https://cloud.codebgp.com/> and in the Organisation ID type "publicdemo"
- Sign up
- Docs: <https://docs.codebgp.com/>



The screenshot shows a dark-themed form titled "Code BGP Platform". It features a text input field containing the text "publicdemo". Below the input field is the label "Enter your organization ID" and a blue "Go!" button.



The screenshot shows a light-themed sign-up form titled "Code BGP Platform". It has "Log In" and "Sign Up" links at the top. The form contains several input fields: an email field with "jamesmith@example.com", a password field with masked characters, a first name field with "James", a last name field with "Smith", and a company name field with "Example Company". A large blue "Sign Up" button is at the bottom.



Questions



✉ lefteris@codebgp.com

🌐 codebgp.com