

BGP Security

Hijack and Route Leak Detection

Lefteris Manassakis | COO, Code BGP

lefteris@codebgp.com



UKNOF 51

April 4, 2023 | Manchester, UK

About me



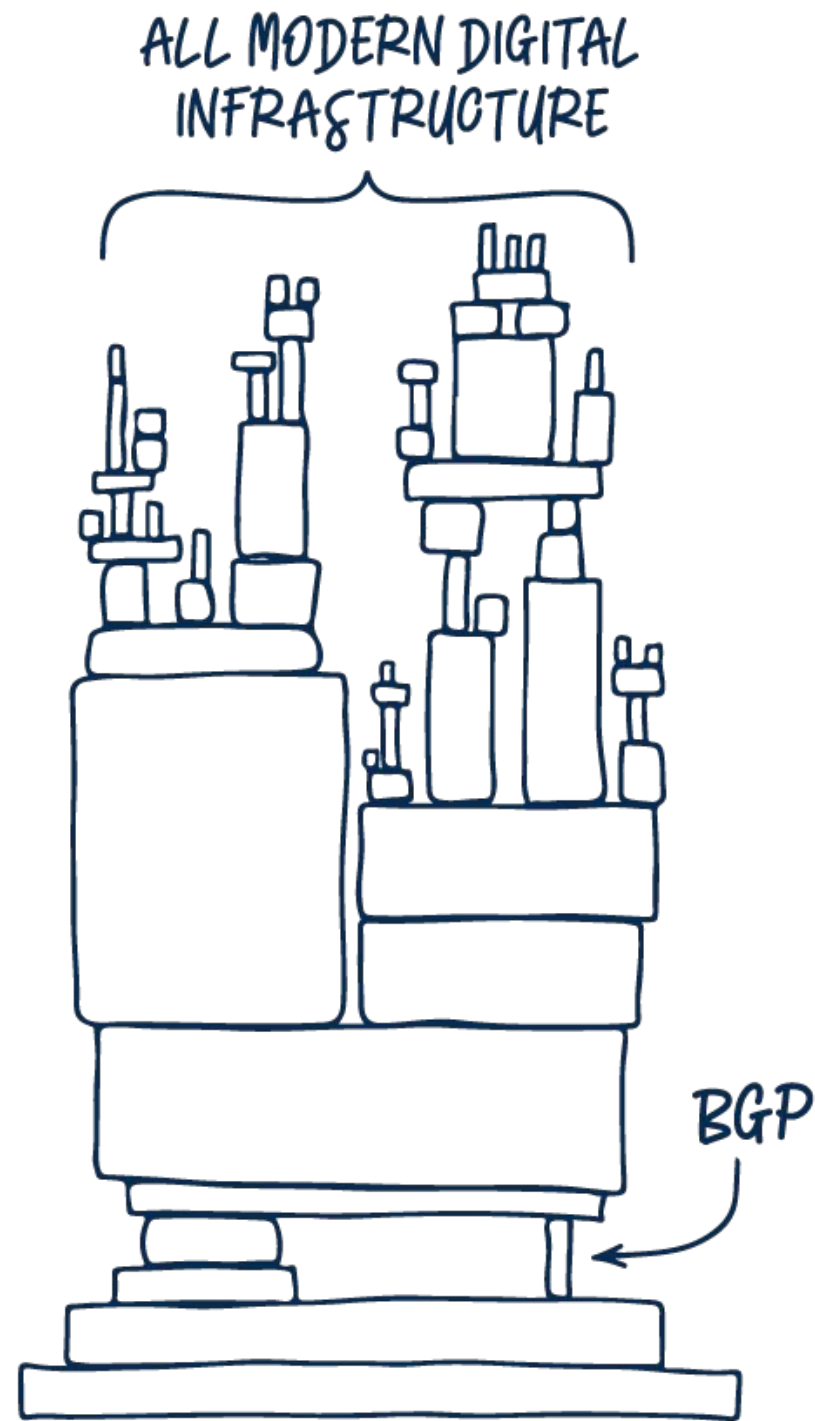
Lfteris Manassakis

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.

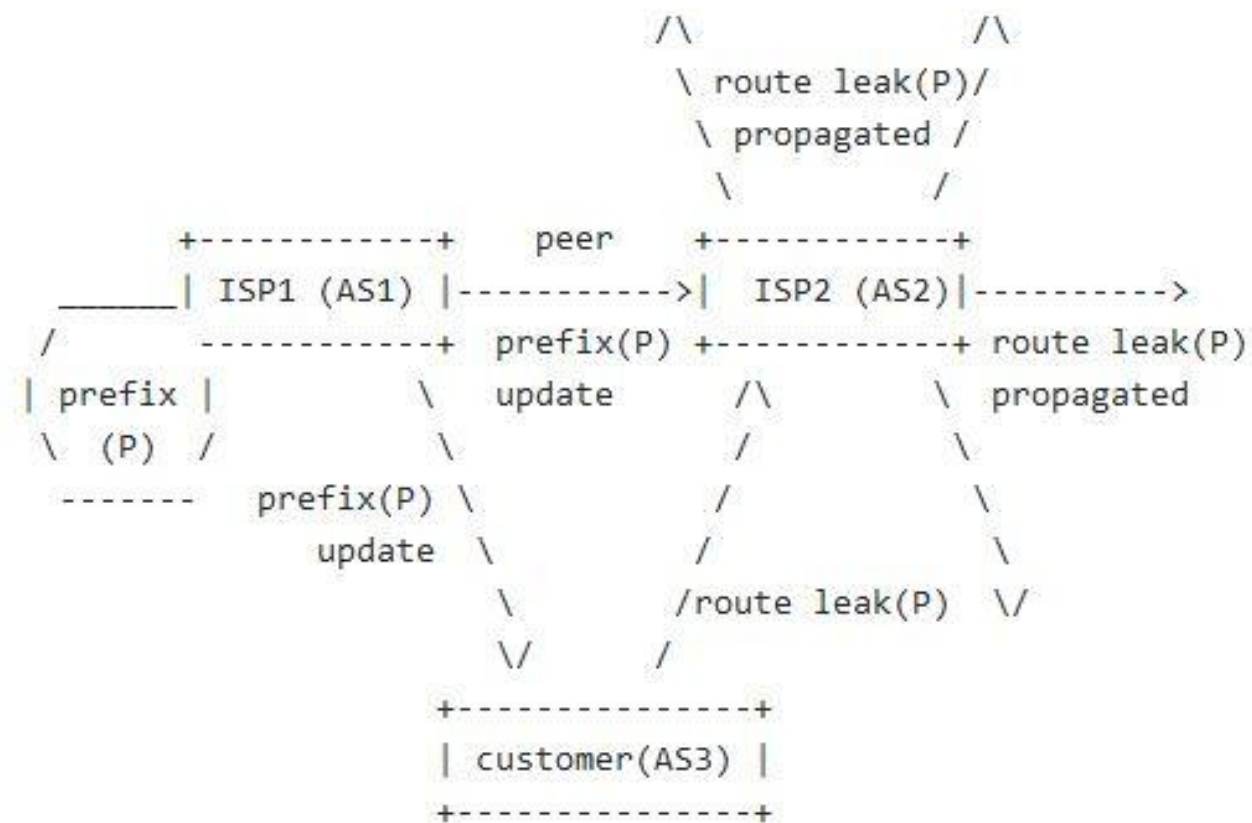



Figure 1: Basic Notion of a Route Leak

- For different types of route leaks please check [RFC 7908](https://www.rfc-editor.org/rfc/rfc7908).

Challenges of hijack and route leak detection

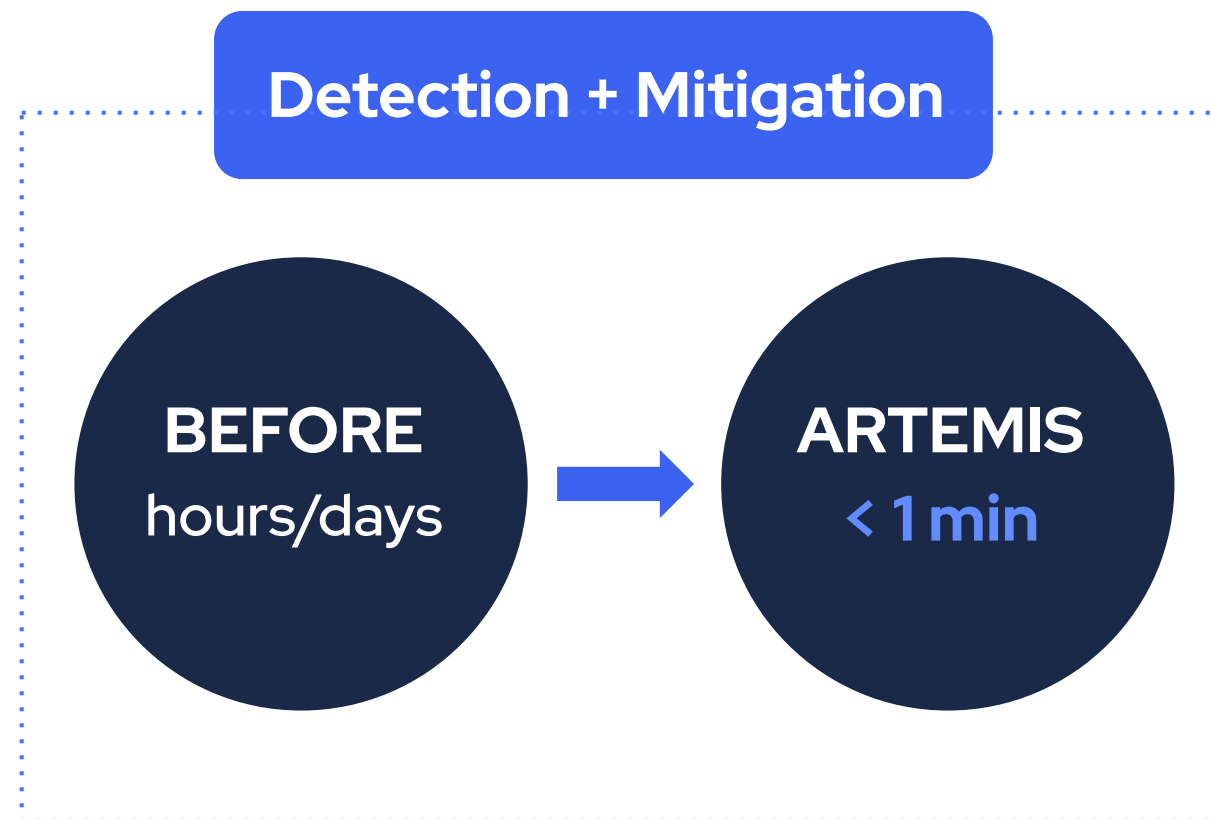
- Speed
- Accuracy
- Evasion
- Privacy and flexibility



ARTEMIS

<https://bgpartemis.org>

- On-prem **open-source** tool we developed
- We support a community of users
- Precursor of the Code BGP Platform



- The Code BGP Platform is offered as a SaaS subscription
- Both are self-operated, leveraging the contextual knowledge of the Network Operator

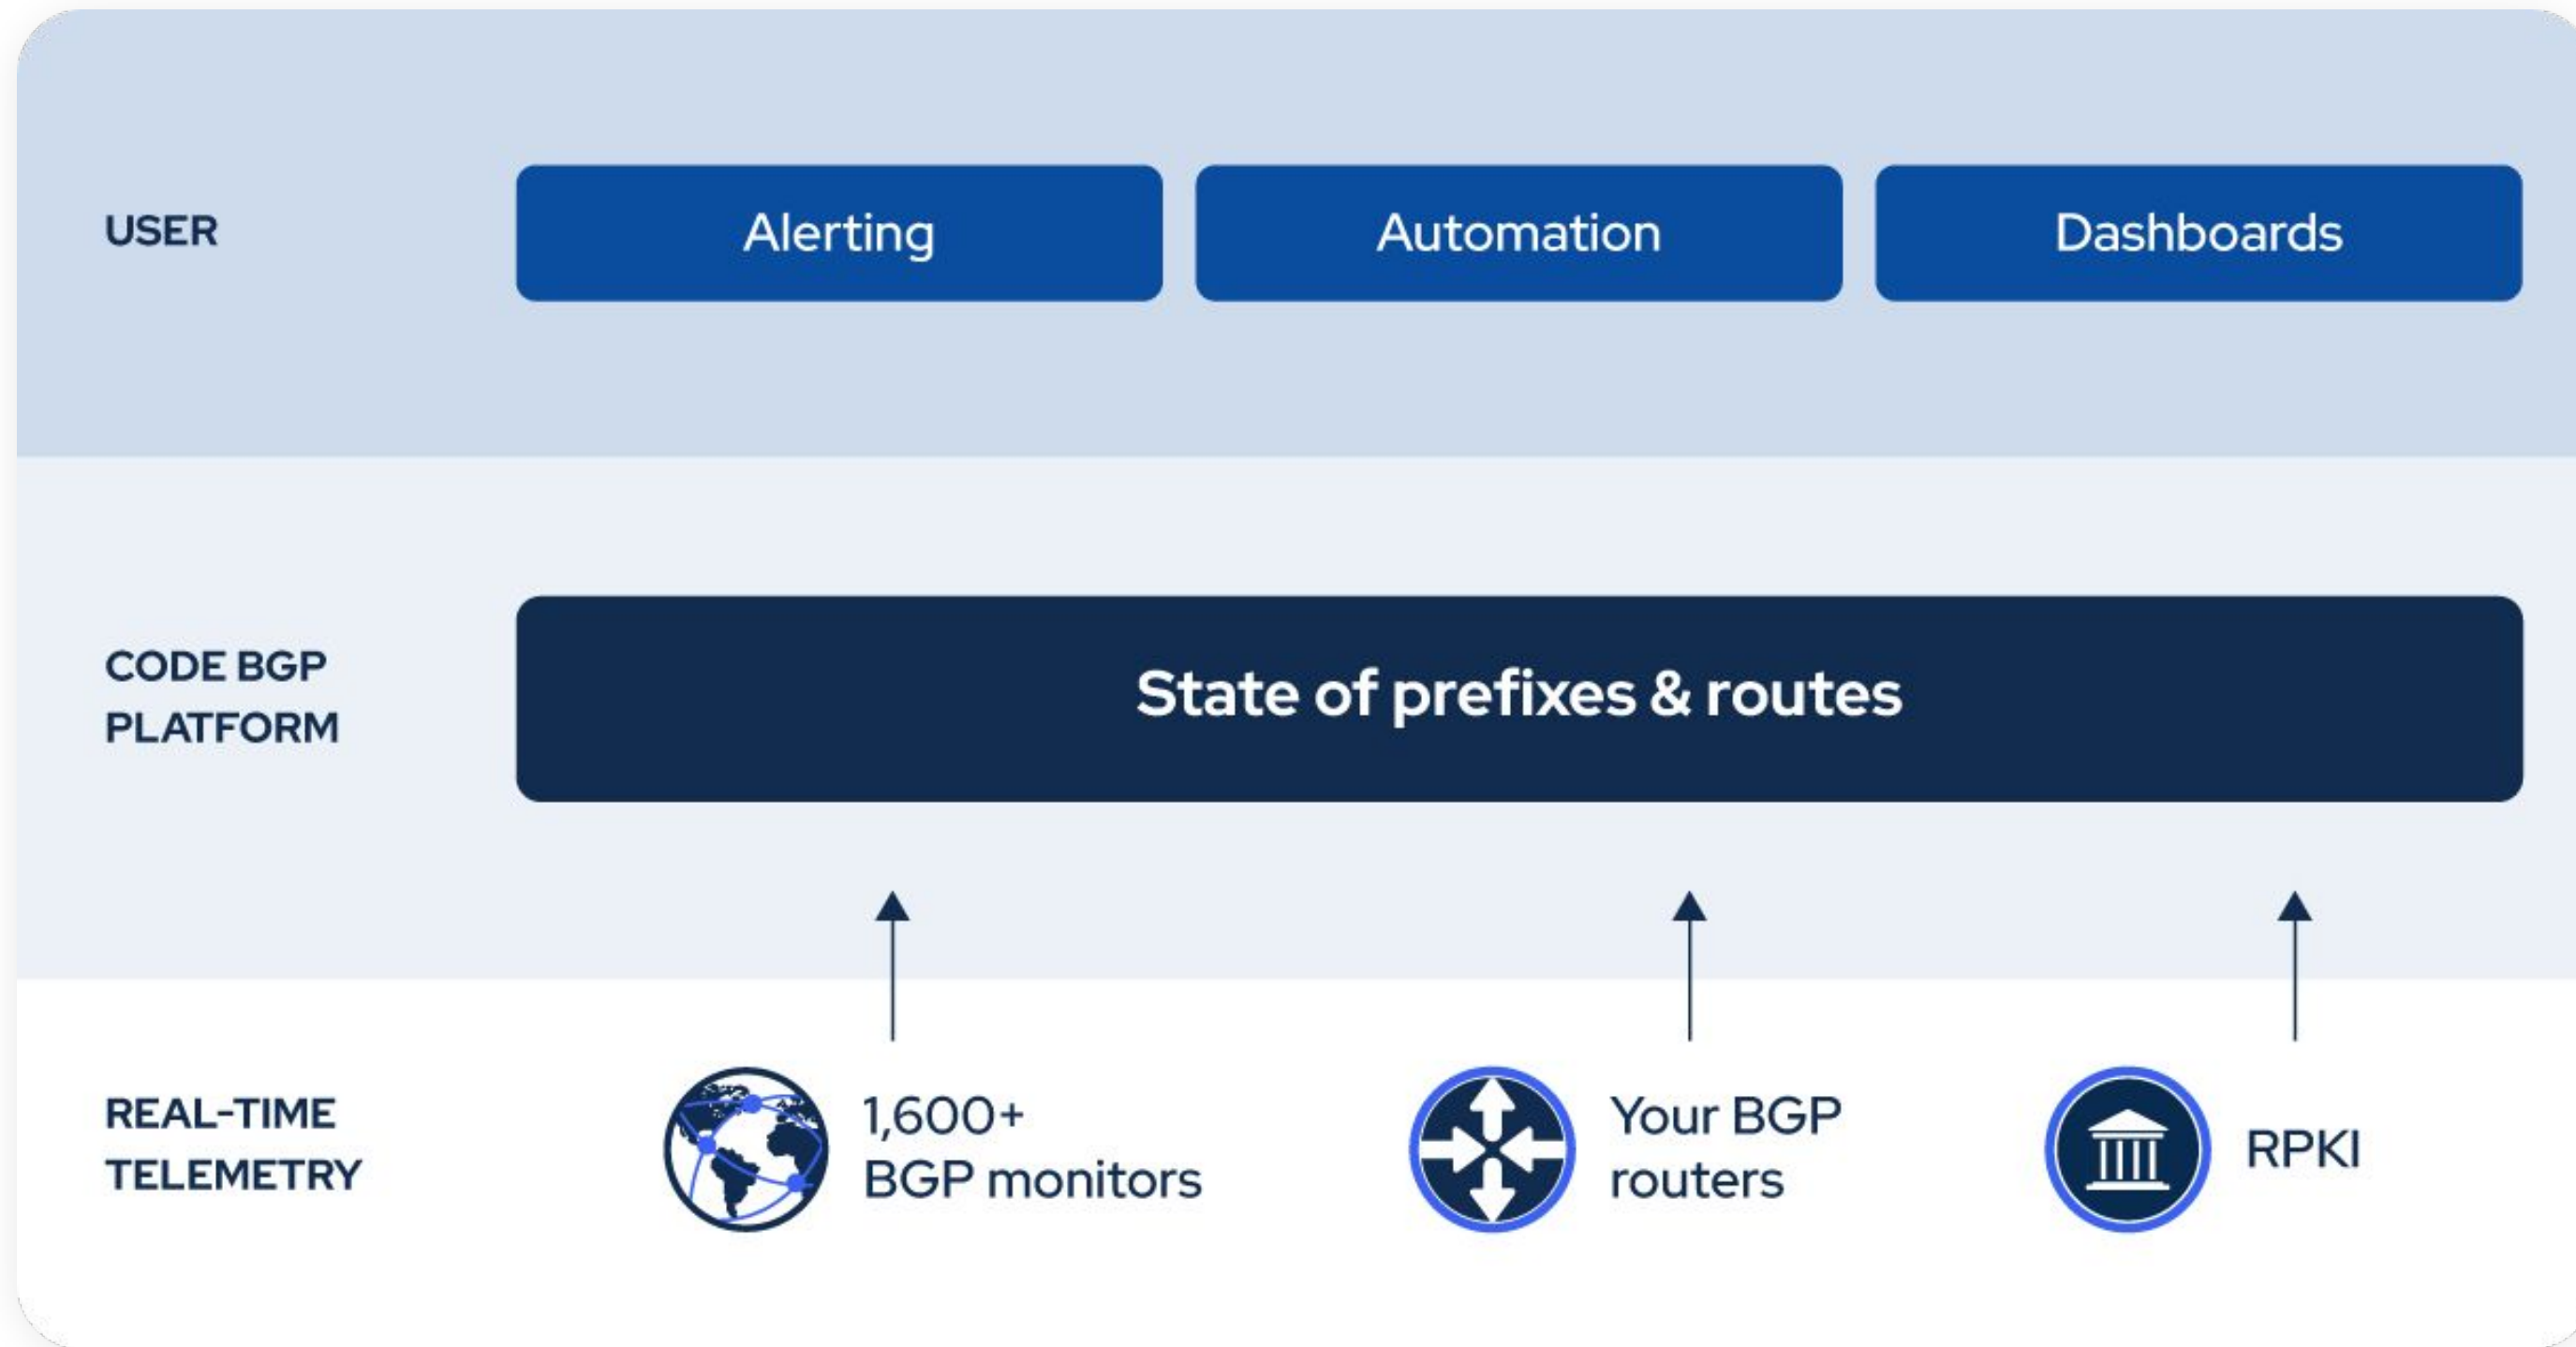


“ARTEMIS is a **fantastic** replacement for BGPmon. All around it seems like **an incredibly well-built tool** and **I use it in prod all the time**”

Chris Cummings
Network Engineer & modem.show podcast host

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](#))
- 186 full feed peerings (v4 & v6)
- 20 Upstreams
- Monitors in 37 countries, 62 cities



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": "UPDATE",
  "require": "ANY",
  "moreSpecific": true,
  "lessSpecific": false,
  "host": "null (all)",
  "peer": null,
  "socketOptions": {
    "includeRaw": false,
    "acknowledge": true
  }
}
```

Code examples

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

JavaScript Python

```
/*
...
*/
```

```
// 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:

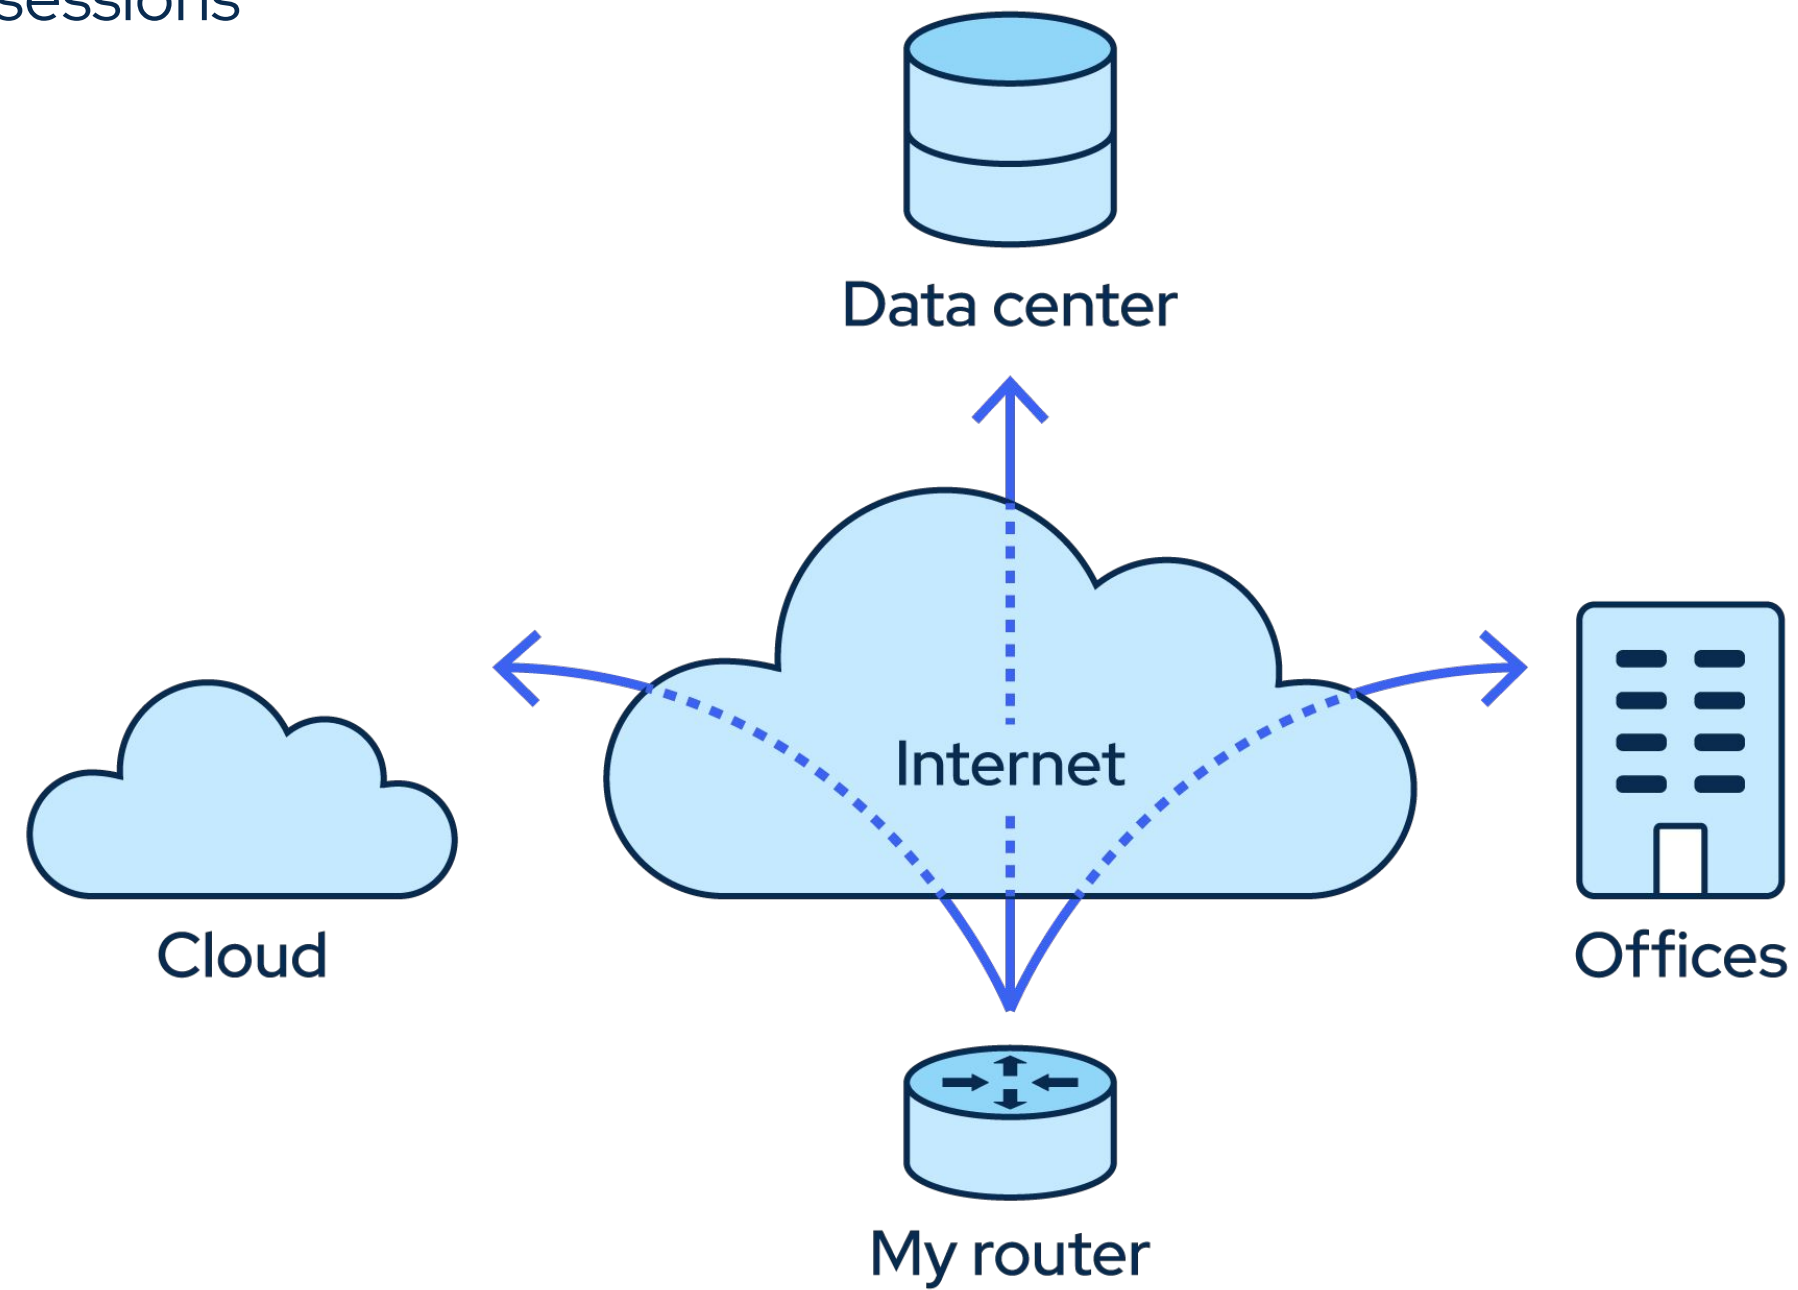
- IPv4: **366**
- IPv6: **401**

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

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

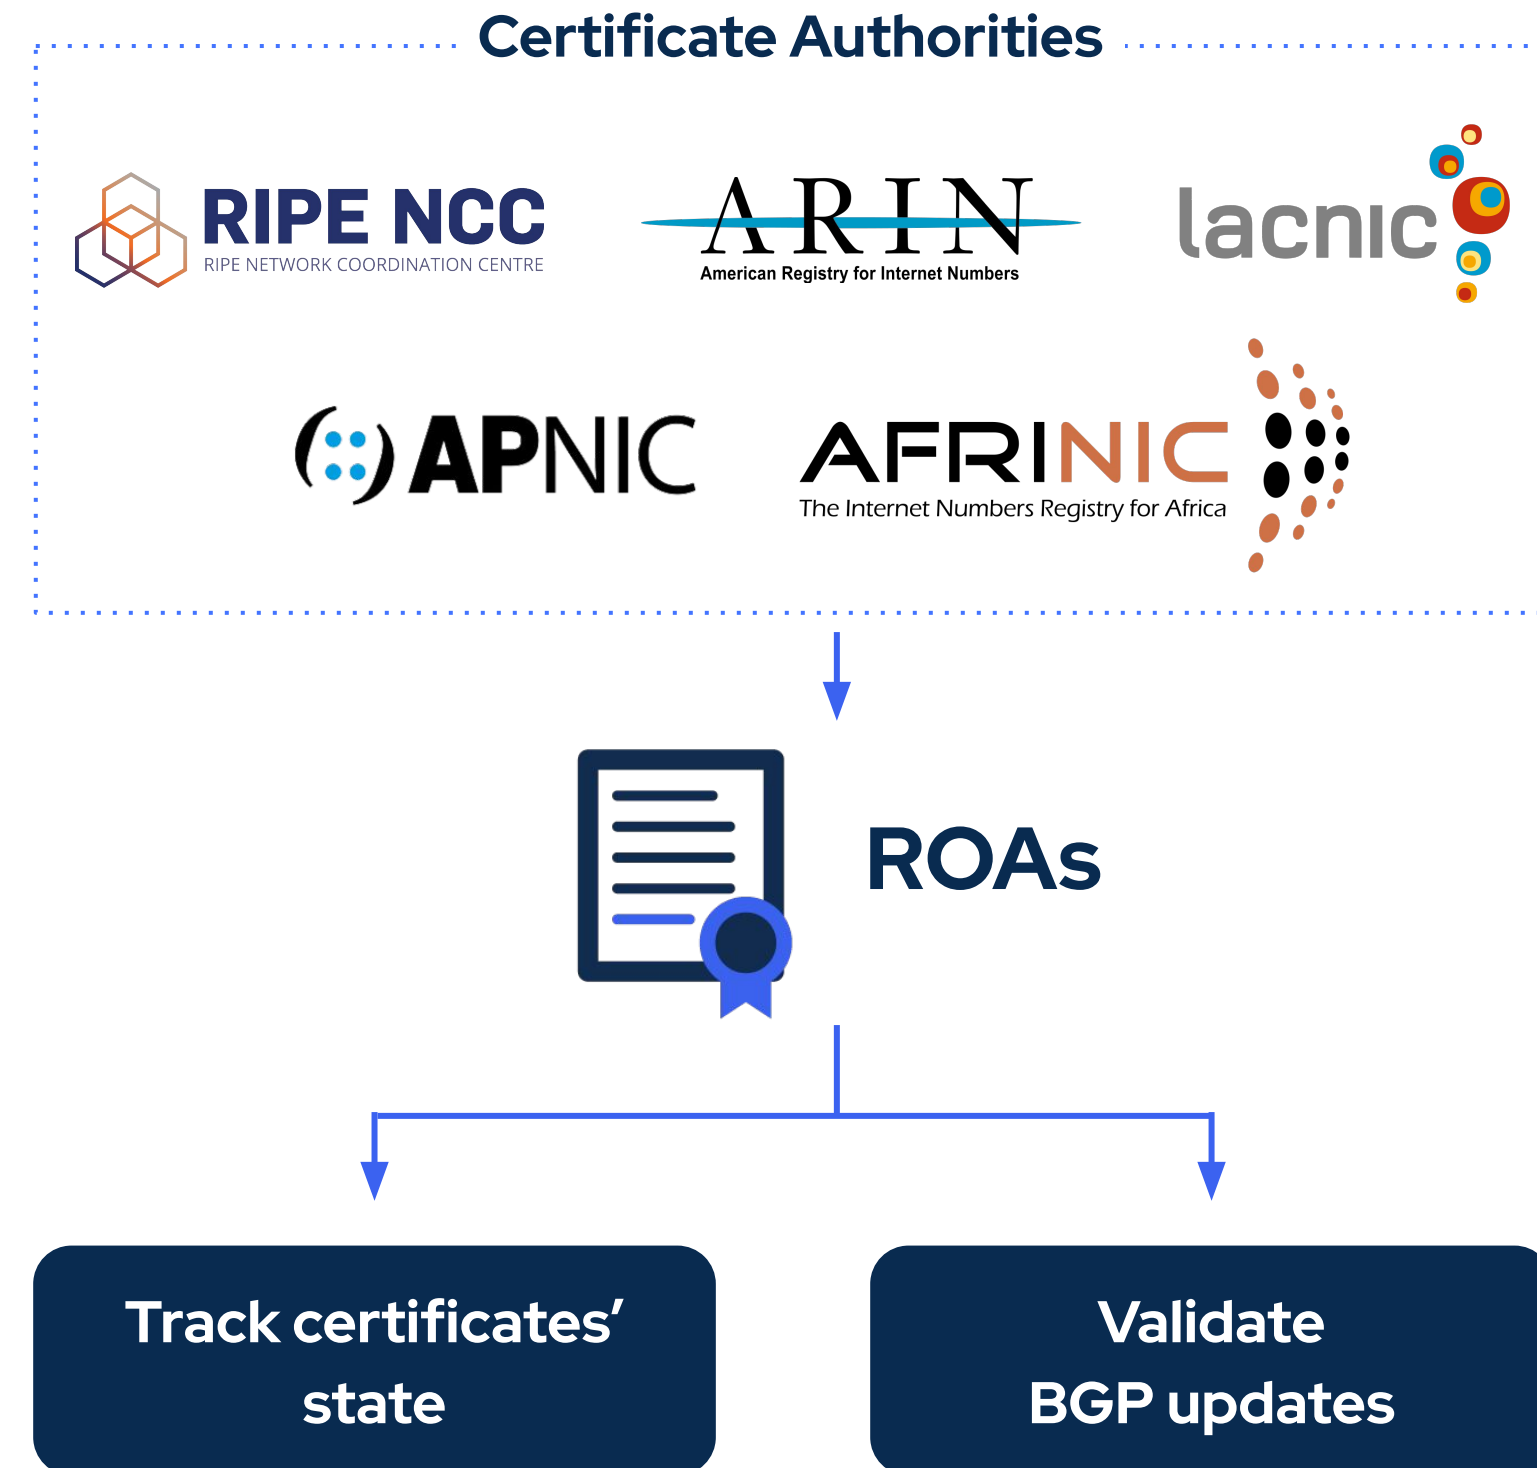
Data Service: Your routers

- **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

GraphQL basics



- **What it is**
 - Query language for APIs
 - Runtime for fulfilling queries with existing data
- **Features**
 - Ask exactly the data you need
 - Get many resources in single request
 - Single endpoint + type system: organized in terms of types and fields, not endpoints
 - No-version API evolution
 - Integration with own data + code
 - Supports subscriptions

GraphQL subscriptions



- Subscriptions are a **GraphQL feature** that allows a server to send data to its clients when a specific event happens. They are implemented with WebSockets, and the server maintains a steady connection to its subscribed client. This also breaks the “Request-Response-Cycle” that were used for all previous interactions with the API.
- Instead, the client initially opens up a **long-lived connection** to the server by sending a subscription query that specifies which event it is interested in. Every time this particular event happens, the server uses the connection to push the event data to the subscribed client(s).

```
GraphQL API | Editor ▶ Prettify History Explorer Docs  
1 Subscription AutonomousSystemNumbers {  
2   autonomousSystems(order_by: {number: asc}) {  
3     number  
4   }  
5 }  
6
```


Insert Alert Rules using the UI

Add Alert Rule

1 — 2 — 3
Alert Rule Configuration Preview

Type
Exact Prefix Hijack

Custom Name
Internet Systems Consortium

[Cancel](#) **Next**

Add Alert Rule

✓ — 2 — 3
Alert Rule Configuration Preview

Internet Systems Consortium (Exact Prefix Hijack)

Parameters

Valid Origin ASes
3557 ASes

Originated Prefixes
192.5.4.0/23 192.5.5.0/24
2001:500:2e::/47 2001:500:2f::/48

Prefixes

Notifications

Type
Email

Email Address
lfteris@codebgp.com

[Cancel](#) [Previous](#) **Next**

Add Alert Rule

✓ — ✓ — 3
Alert Rule Configuration Preview


Summary for alert rule:
Internet Systems Consortium (Exact Prefix Hijack)

Parameter	Value
ASes	3557
Prefixes	192.5.4.0/23, 192.5.5.0/24, 2001:500:2e::/47, 2001:500:2f::/48

Notifications:
Email (lfteris@codebgp.com)

[Cancel](#) [Previous](#) **Add Alert Rule**





How we use GraphQL Subscriptions for Alert Rules

- **Example** of a subscription query (which is entered to the system as a mutation) to detect exact prefix hijacks for prefixes belonging to Code BGP (AS 50414).

```
mutation MutationExactPrefixHijack {
  insertAlertSubscription(object: {name: "Exact Prefix Hijack", query: "subscription IllegalOriginsFromWhichExactPrefixesAreAnnounced($asns:
[bigint!] = [], $prefixes: [cidr!] = []) { routes(where: {originAutonomousSystem: {number: {_nin: $asns}}, prefix: {network: {_in: $prefixes}}}}
order_by:
{as_path: asc, prefix: {network: asc}, originAutonomousSystem: {number: asc}}) { originAutonomousSystem { number } prefix { network } as_path
}}", vars: {asns:[50414],
prefixes:["212.46.55.0/24","2a12:bc0::/48","2a12:bc0:1::/48","2a12:bc0:2::/48","2a12:bc0:3::/48","2a12:bc0:4::/48","2a12:bc0:5::/48"]},
fire_alert_regex: "^.*/routes.*as_path.*$", type: "as_origin_violation_exact", severity: "critical", description: "Illegal origin ASes that
announce configured prefixes."}) {
  id
  name
  query
  vars
  fire_alert_regex
  type
  severity
  description
}
}
```

Root DNS Servers

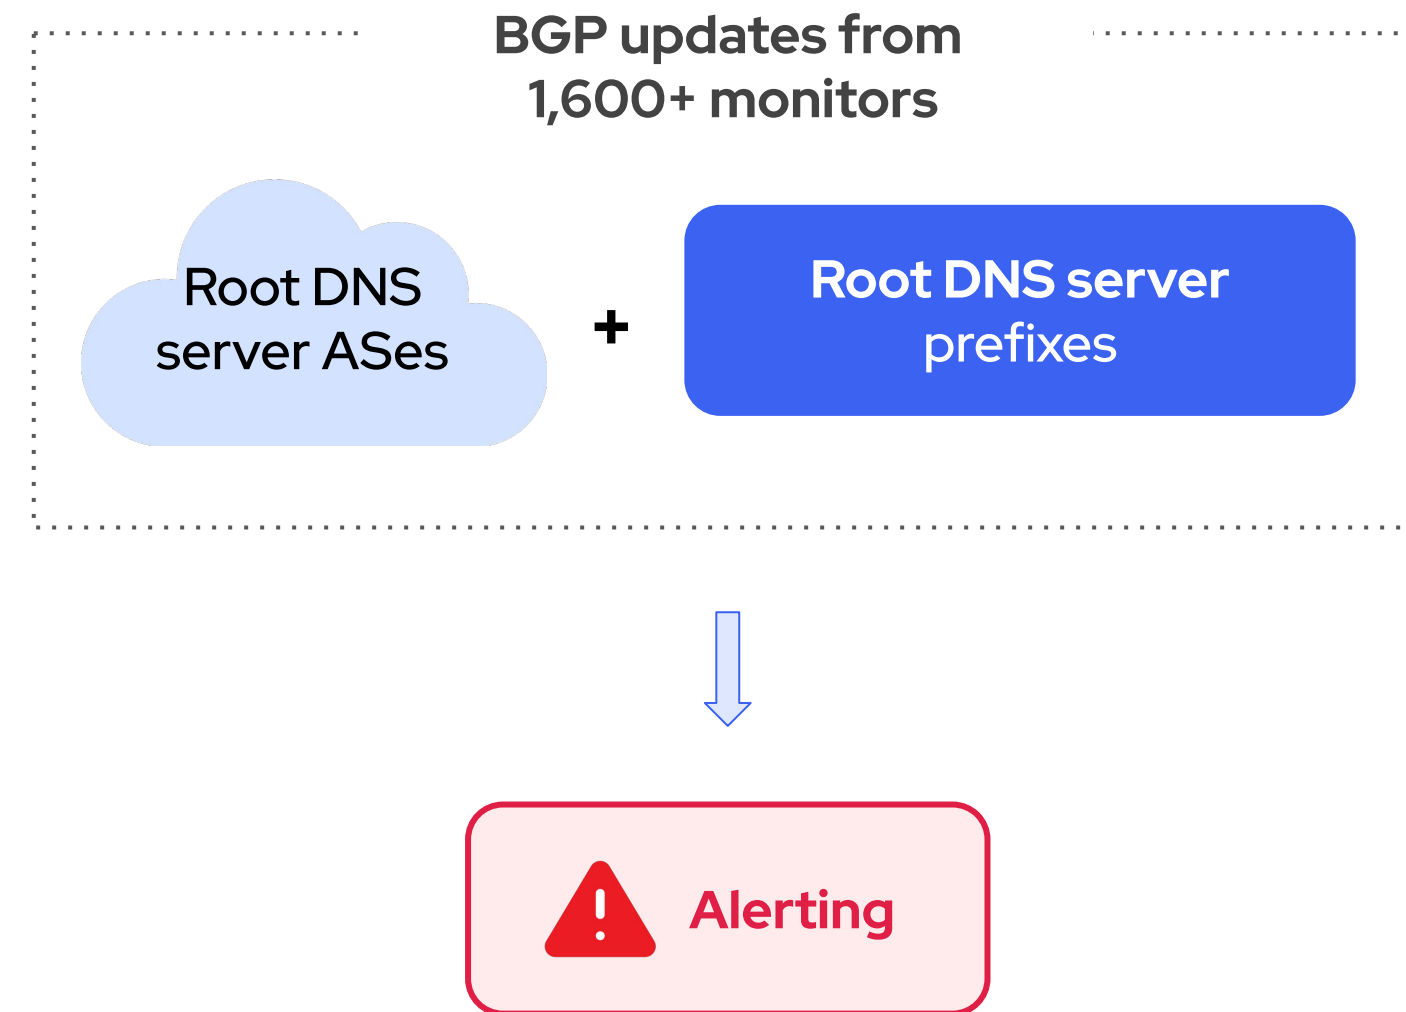
- The authoritative name servers that serve the DNS root zone

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

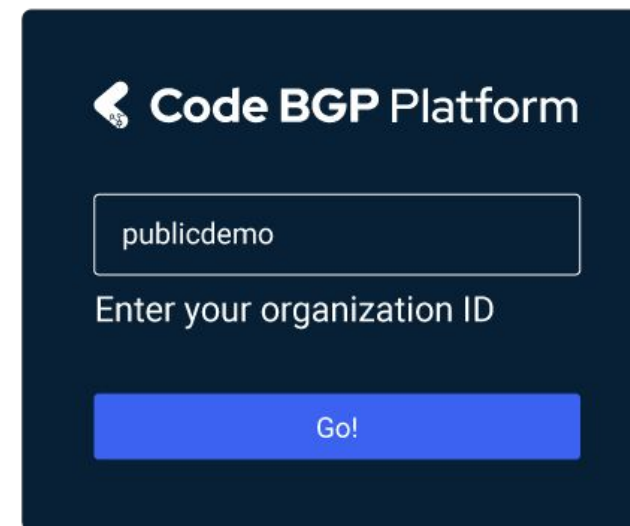
- Critical Internet infrastructure, worth protecting
- These prefixes are heavily anycasted
 - BGP anomalies (e.g. exact prefix hijacks) will go largely unnoticed, due to their limited impact on the data plane

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

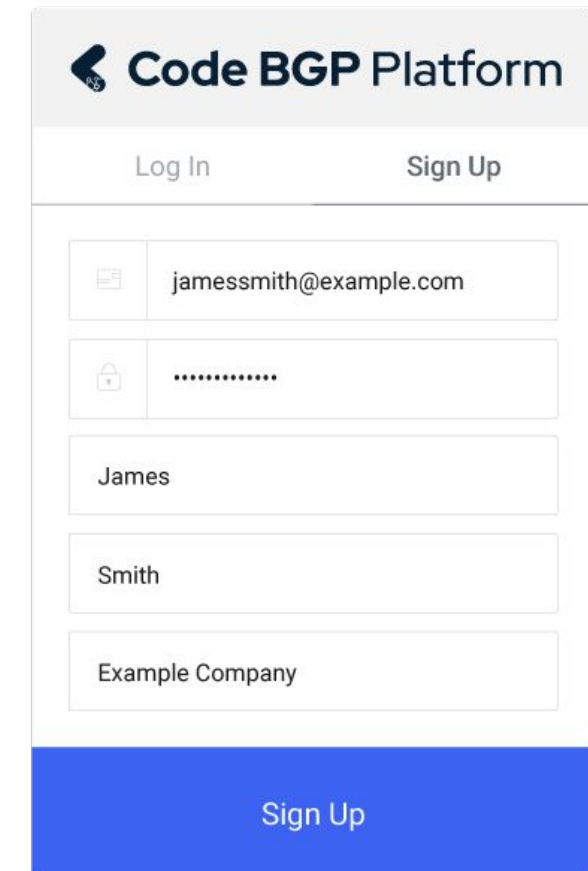


How to get access to the Route DNS monitoring instance


- 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 interface for the Code BGP Platform. At the top, there is a back arrow icon and the text "Code BGP Platform". Below this is a text input field containing the value "publicdemo". Underneath the input field, the text "Enter your organization ID" is displayed. At the bottom of the form is a blue button labeled "Go!".



The screenshot shows a light-themed sign-up form for the Code BGP Platform. At the top, there is a back arrow icon and the text "Code BGP Platform". Below this are two links: "Log In" and "Sign Up". 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". At the bottom of the form is a large blue button labeled "Sign Up".



Exact Prefix Hijack detected for root DNS prefix - Jan 27

- AS 24028 announced prefix 2001:500:2f::/48 which belongs to [ISC](#), and serves as the IPv6 prefix of the “[F-Root](#)” domain server (AS 3557)
- Seen only by one source, which happens to be a neighbor of the offending network. The limited propagation is possibly due to RPKI ROV

The screenshot shows the Code BGP Platform interface. The top navigation bar includes the Code BGP Platform logo and the user profile for Lefteris Manassakis (editor | tenant7). The left sidebar contains navigation options: Overview, Setup, AS Filters, Prefix Filters, Alert Rules, Data Services, State (selected), and API.

The main content area is titled "State" and has tabs for Prefixes, Autonomous Systems, Peerings, Routes (selected), and RPKI ROAs. A filter "Origin AS: 24028" is applied. The "Routes" table shows a single entry for the prefix 2001:500:2f::/48, originating from AS 24028 and advertised by neighbor AS 38001. The RPKI Status is "Invalid", and it was first detected on Jan 27, 2023, at 11:48:20.

Below the route table, a section titled "Data Sources of Route 2001:500:2f::/48 - 38001 24028" displays a table of data sources:

Data Service	Route Collector	IP	ASN	City	Country	Continent	Last Update ↓
RIS Live	RRC00	2406:f400:8:34::1	38001	Singapore		Asia	Jan 27, 2023, 11:48:18

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

Exact Prefix Hijacks detected for root DNS prefixes – Feb. 25

- AS 7639 announced prefix 2001:500:a8::/48 which belongs to NASA and is the IPv6 prefix of the “[E-Root](#)” domain server (AS 21556)
- At the exact same time, the same AS 7639 announced prefix 2001:500:2f::/48 which belongs to F-Root (ISC AS 3557)

Code BGP Platform

Lefteris Manassakis editor | publicdemo

State Info

Prefixes Autonomous Systems Peerings Routes RPKI ROAs

Origin AS: 17639

Prefix	Origin AS	Neighbor AS	AS Path	RPKI Status	First Detected ↓	Last Update
> 2001:500:a8::/48	17639	1239	1239 17639	NotFound	Feb 25, 2023, 01:47:48	Feb 25, 2023, 01:47:46
> 2001:500:a8::/48	17639	396998	396998 17639	NotFound	Feb 25, 2023, 01:47:43	Feb 25, 2023, 01:47:35
> 2001:500:a8::/48	17639	1239	205148 9002 1239 17639	NotFound	Feb 25, 2023, 01:47:43	Feb 25, 2023, 01:47:36
> 2001:500:2f::/48	17639	396998	396998 17639	Invalid	Feb 25, 2023, 01:47:43	Feb 25, 2023, 01:47:35
> 2001:500:a8::/48	17639	1239	9002 1239 17639	NotFound	Feb 25, 2023, 01:47:43	Feb 25, 2023, 01:47:35
> 2001:500:2f::/48	17639	137409	57695 137409 17639	Invalid	Feb 25, 2023, 01:47:35	Feb 25, 2023, 01:47:35
> 2001:500:a8::/48	17639	137409	57695 137409 17639	NotFound	Feb 25, 2023, 01:47:34	Feb 25, 2023, 01:47:35

Rows per page: 10 1-7 of 7

Exact Prefix Hijacks detected for root DNS prefixes - Feb. 25

- The "E-Root" 2001:500:a8::/48 prefix is not covered by a RPKI ROA. The event lasted 2 days

Code BGP Platform

Prefix: 2001:500:a8::/48 Origin AS: 17639

Prefix	Origin AS	Neighbor AS	AS Path	RPKI Status	Last Update
2001:500:a8::/48	17639	1239	1239 17639	NotFound	Feb 25, 2023, 01:47:46

Data Sources of Route 2001:500:a8::/48 - 1239 17639

Data Service	Route Collector	IP	ASN	City	Country	Continent	Last Update
RIS Live	RRC01	2001:7f8:4::4d7:1	1239	London	United Kingdom	Europe	Feb 25, 2023, 01:47:46
RIS Live	RRC12	2001:7f8:4d7:0:1	1239	Frankfurt	Germany	Europe	Feb 25, 2023, 01:47:44

Rows per page: 10 1-2 of 2

Prefix	Origin AS	Neighbor AS	AS Path	RPKI Status	Last Update
2001:500:a8::/48	17639	396998	396998 17639	NotFound	Feb 25, 2023, 01:47:35

Data Sources of Route 2001:500:a8::/48 - 396998 17639

Data Service	Route Collector	IP	ASN	City	Country	Continent	Last Update
RIS Live	RRC11	2001:504:1::a539:6998:1	396998	New York	USA	North America	Feb 25, 2023, 01:47:35

Rows per page: 10 1-1 of 1

Prefix	Origin AS	Neighbor AS	AS Path	RPKI Status	Last Update
2001:500:a8::/48	17639	1239	205148 9002 1239 17639	NotFound	Feb 25, 2023, 01:47:36

Data Sources of Route 2001:500:a8::/48 - 205148 9002 1239 17639

Data Service	Route Collector	IP	ASN	City	Country	Continent	Last Update
RIS Live	RRC00	2a0d:f407:101:dead::1	205148				Feb 25, 2023, 01:47:36

Rows per page: 10 1-1 of 1

Share your suggestions with our team!

Send us an Email

Acknowledgments: RIPE NCC

Code BGP Platform

Prefix: 2001:500:a8::/48 Origin AS: 17639

Prefix	Origin AS	Neighbor AS	AS Path	RPKI Status	Last Update
2001:500:a8::/48	17639	1239	1239 17639	NotFound	Feb 25, 2023, 01:47:46
2001:500:a8::/48	17639	396998	396998 17639	NotFound	Feb 25, 2023, 01:47:35
2001:500:a8::/48	17639	1239	205148 9002 1239 17639	NotFound	Feb 25, 2023, 01:47:36
2001:500:a8::/48	17639	1239	9002 1239 17639	NotFound	Feb 25, 2023, 01:47:35

Data Sources of Route 2001:500:a8::/48 - 9002 1239 17639

Data Service	Route Collector	IP	ASN	City	Country	Continent	Last Update
RIS Live	RRC11	2001:504:1::a500:9002:1	9002	New York	USA	North America	Feb 25, 2023, 01:47:35
RIS Live	RRC12	2001:7f8::232a:0:1	9002	Frankfurt	Germany	Europe	Feb 25, 2023, 01:47:35
RIS Live	RRC07	2001:7f8:dff::157	9002	Tobolsk	Russia	Europe	Feb 25, 2023, 01:47:35
RIS Live	RRC01	2001:7f8:4::232a:1	9002	London	United Kingdom	Europe	Feb 25, 2023, 01:47:35

Rows per page: 10 1-4 of 4

Prefix	Origin AS	Neighbor AS	AS Path	RPKI Status	Last Update
2001:500:a8::/48	17639	137409	57695 137409 17639	NotFound	Feb 25, 2023, 01:47:35

Data Sources of Route 2001:500:a8::/48 - 57695 137409 17639

Data Service	Route Collector	IP	ASN	City	Country	Continent	Last Update
Code BGP Monitor		103.170.233.249	57695	Tokyo	Japan	Asia	Feb 25, 2023, 01:47:35
Code BGP Monitor		194.156.163.203	57695	Singapore	Singapore	Asia	Feb 25, 2023, 01:47:35

Rows per page: 10 1-2 of 2

Share your suggestions with our team!

Send us an Email

Acknowledgments: RIPE NCC

Exact Prefix Hijacks detected for root DNS prefixes – Feb. 25

- The “F-Root” 2001:500:2f::/48 prefix is covered by a RPKI ROA. The event lasted 18 hours

The screenshot displays the Code BGP Platform interface. The top navigation bar includes the Code BGP logo, the user name 'Lefteris Manassakis', and a profile icon. The left sidebar contains navigation options: Overview, Setup, State (selected), API, and Alerts. The main content area is titled 'State' and has tabs for Prefixes, Autonomous Systems, Peerings, Routes (selected), and RPKI ROAs. A search filter is applied: Prefix: 2001:500:2f::/48 and Origin AS: 17639. The main table shows route information with columns for Prefix, Origin AS, Neighbor AS, AS Path, RPKI Status, and Last Update. Two routes are listed, both with an 'Invalid' RPKI status. The first route has AS Path 396998 17639, and the second has AS Path 57695 137409 17639. Below each route, a 'Data Sources' table provides details on the route collector, IP, ASN, city, country, and continent. The first route's data source is RIS Live from RRC11 in New York, North America. The second route's data sources are Code BGP Monitor from Tokyo, Asia and Code BGP Monitor from Singapore, Asia. A footer in the sidebar encourages users to share suggestions via email.

Prefix	Origin AS	Neighbor AS	AS Path	RPKI Status	Last Update		
2001:500:2f::/48	17639	396998	396998 17639	Invalid	Feb 25, 2023, 01:47:35		
Data Sources of Route 2001:500:2f::/48 - 396998 17639							
Data Service	Route Collector	IP	ASN	City	Country	Continent	Last Update
RIS Live	RRC11	2001:504:1::a539:6998:1	396998	New York	USA	North America	Feb 25, 2023, 01:47:35
2001:500:2f::/48	17639	137409	57695 137409 17639	Invalid	Feb 25, 2023, 01:47:35		
Data Sources of Route 2001:500:2f::/48 - 57695 137409 17639							
Data Service	Route Collector	IP	ASN	City	Country	Continent	Last Update
Code BGP Monitor		103.170.233.249	57695	Tokyo	Japan	Asia	Feb 25, 2023, 01:47:35
Code BGP Monitor		194.156.163.203	57695	Singapore	Singapore	Asia	Feb 25, 2023, 01:47:35



Prefix Hijacking Demo

Setup (Start Node)

Basic network and activate the user interface using the instructions in the terminal window. This is a basic network for a small network with a few nodes and links.

Network is a configuration of nodes and links. The network is a graph with nodes and links.

Network Name	Nodes	Links
Network 1	1000	1000
Network 2	1000	1000
Network 3	1000	1000
Network 4	1000	1000
Network 5	1000	1000
Network 6	1000	1000
Network 7	1000	1000
Network 8	1000	1000
Network 9	1000	1000
Network 10	1000	1000
Network 11	1000	1000
Network 12	1000	1000
Network 13	1000	1000
Network 14	1000	1000
Network 15	1000	1000
Network 16	1000	1000
Network 17	1000	1000
Network 18	1000	1000
Network 19	1000	1000
Network 20	1000	1000
Network 21	1000	1000
Network 22	1000	1000
Network 23	1000	1000
Network 24	1000	1000
Network 25	1000	1000
Network 26	1000	1000
Network 27	1000	1000
Network 28	1000	1000
Network 29	1000	1000
Network 30	1000	1000
Network 31	1000	1000
Network 32	1000	1000
Network 33	1000	1000
Network 34	1000	1000
Network 35	1000	1000
Network 36	1000	1000
Network 37	1000	1000
Network 38	1000	1000
Network 39	1000	1000
Network 40	1000	1000
Network 41	1000	1000
Network 42	1000	1000
Network 43	1000	1000
Network 44	1000	1000
Network 45	1000	1000
Network 46	1000	1000
Network 47	1000	1000
Network 48	1000	1000
Network 49	1000	1000
Network 50	1000	1000

Back

Submit

Questions



✉ lefteris@codebgp.com

🌐 codebgp.com