RING-ADMINS @ NLNOG.NET

NLNOG RING UPDATE

SOME STATISTICS

- RING Nodes: **717** (+88)
- Unique ASNs: **542** (+35)
- Organisations: **527** (+32)
- Countries: **60** (+1)

- IPv4: **187** peers (+6), ~**200M** prefixes
- IPv6: **198** peers (+13), ~**60M** prefixes

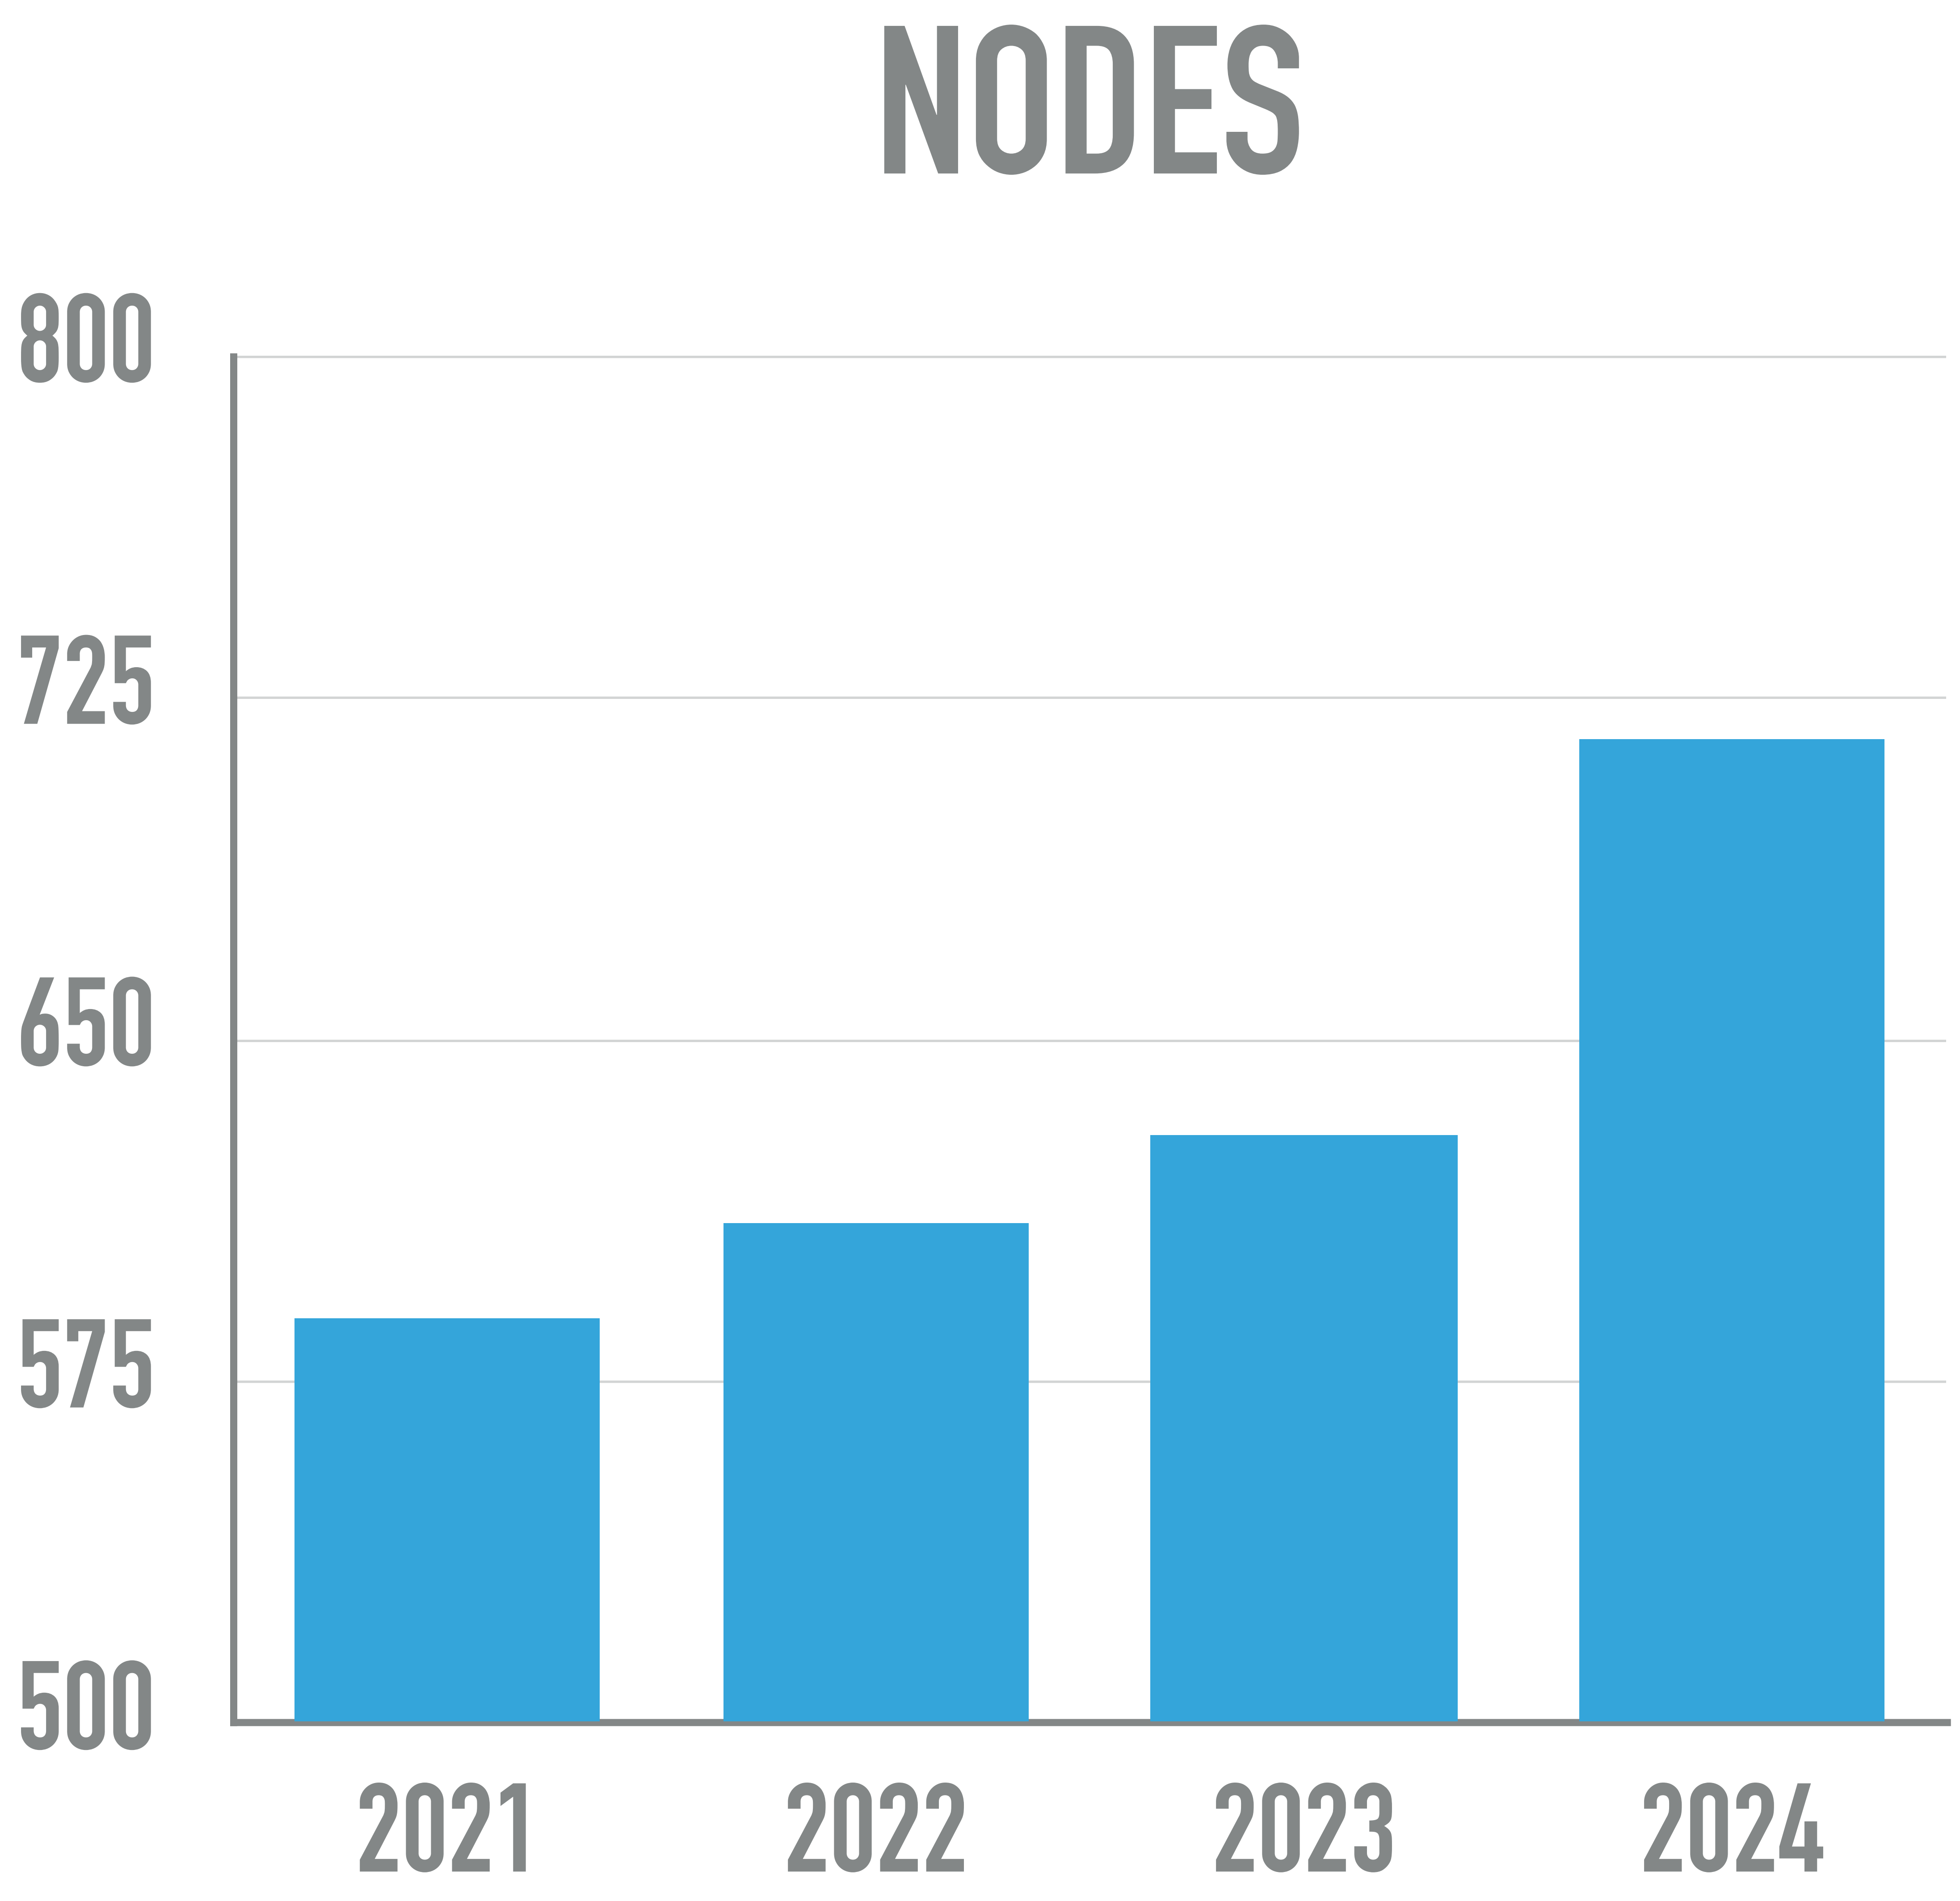

RING NODE LOCATIONS

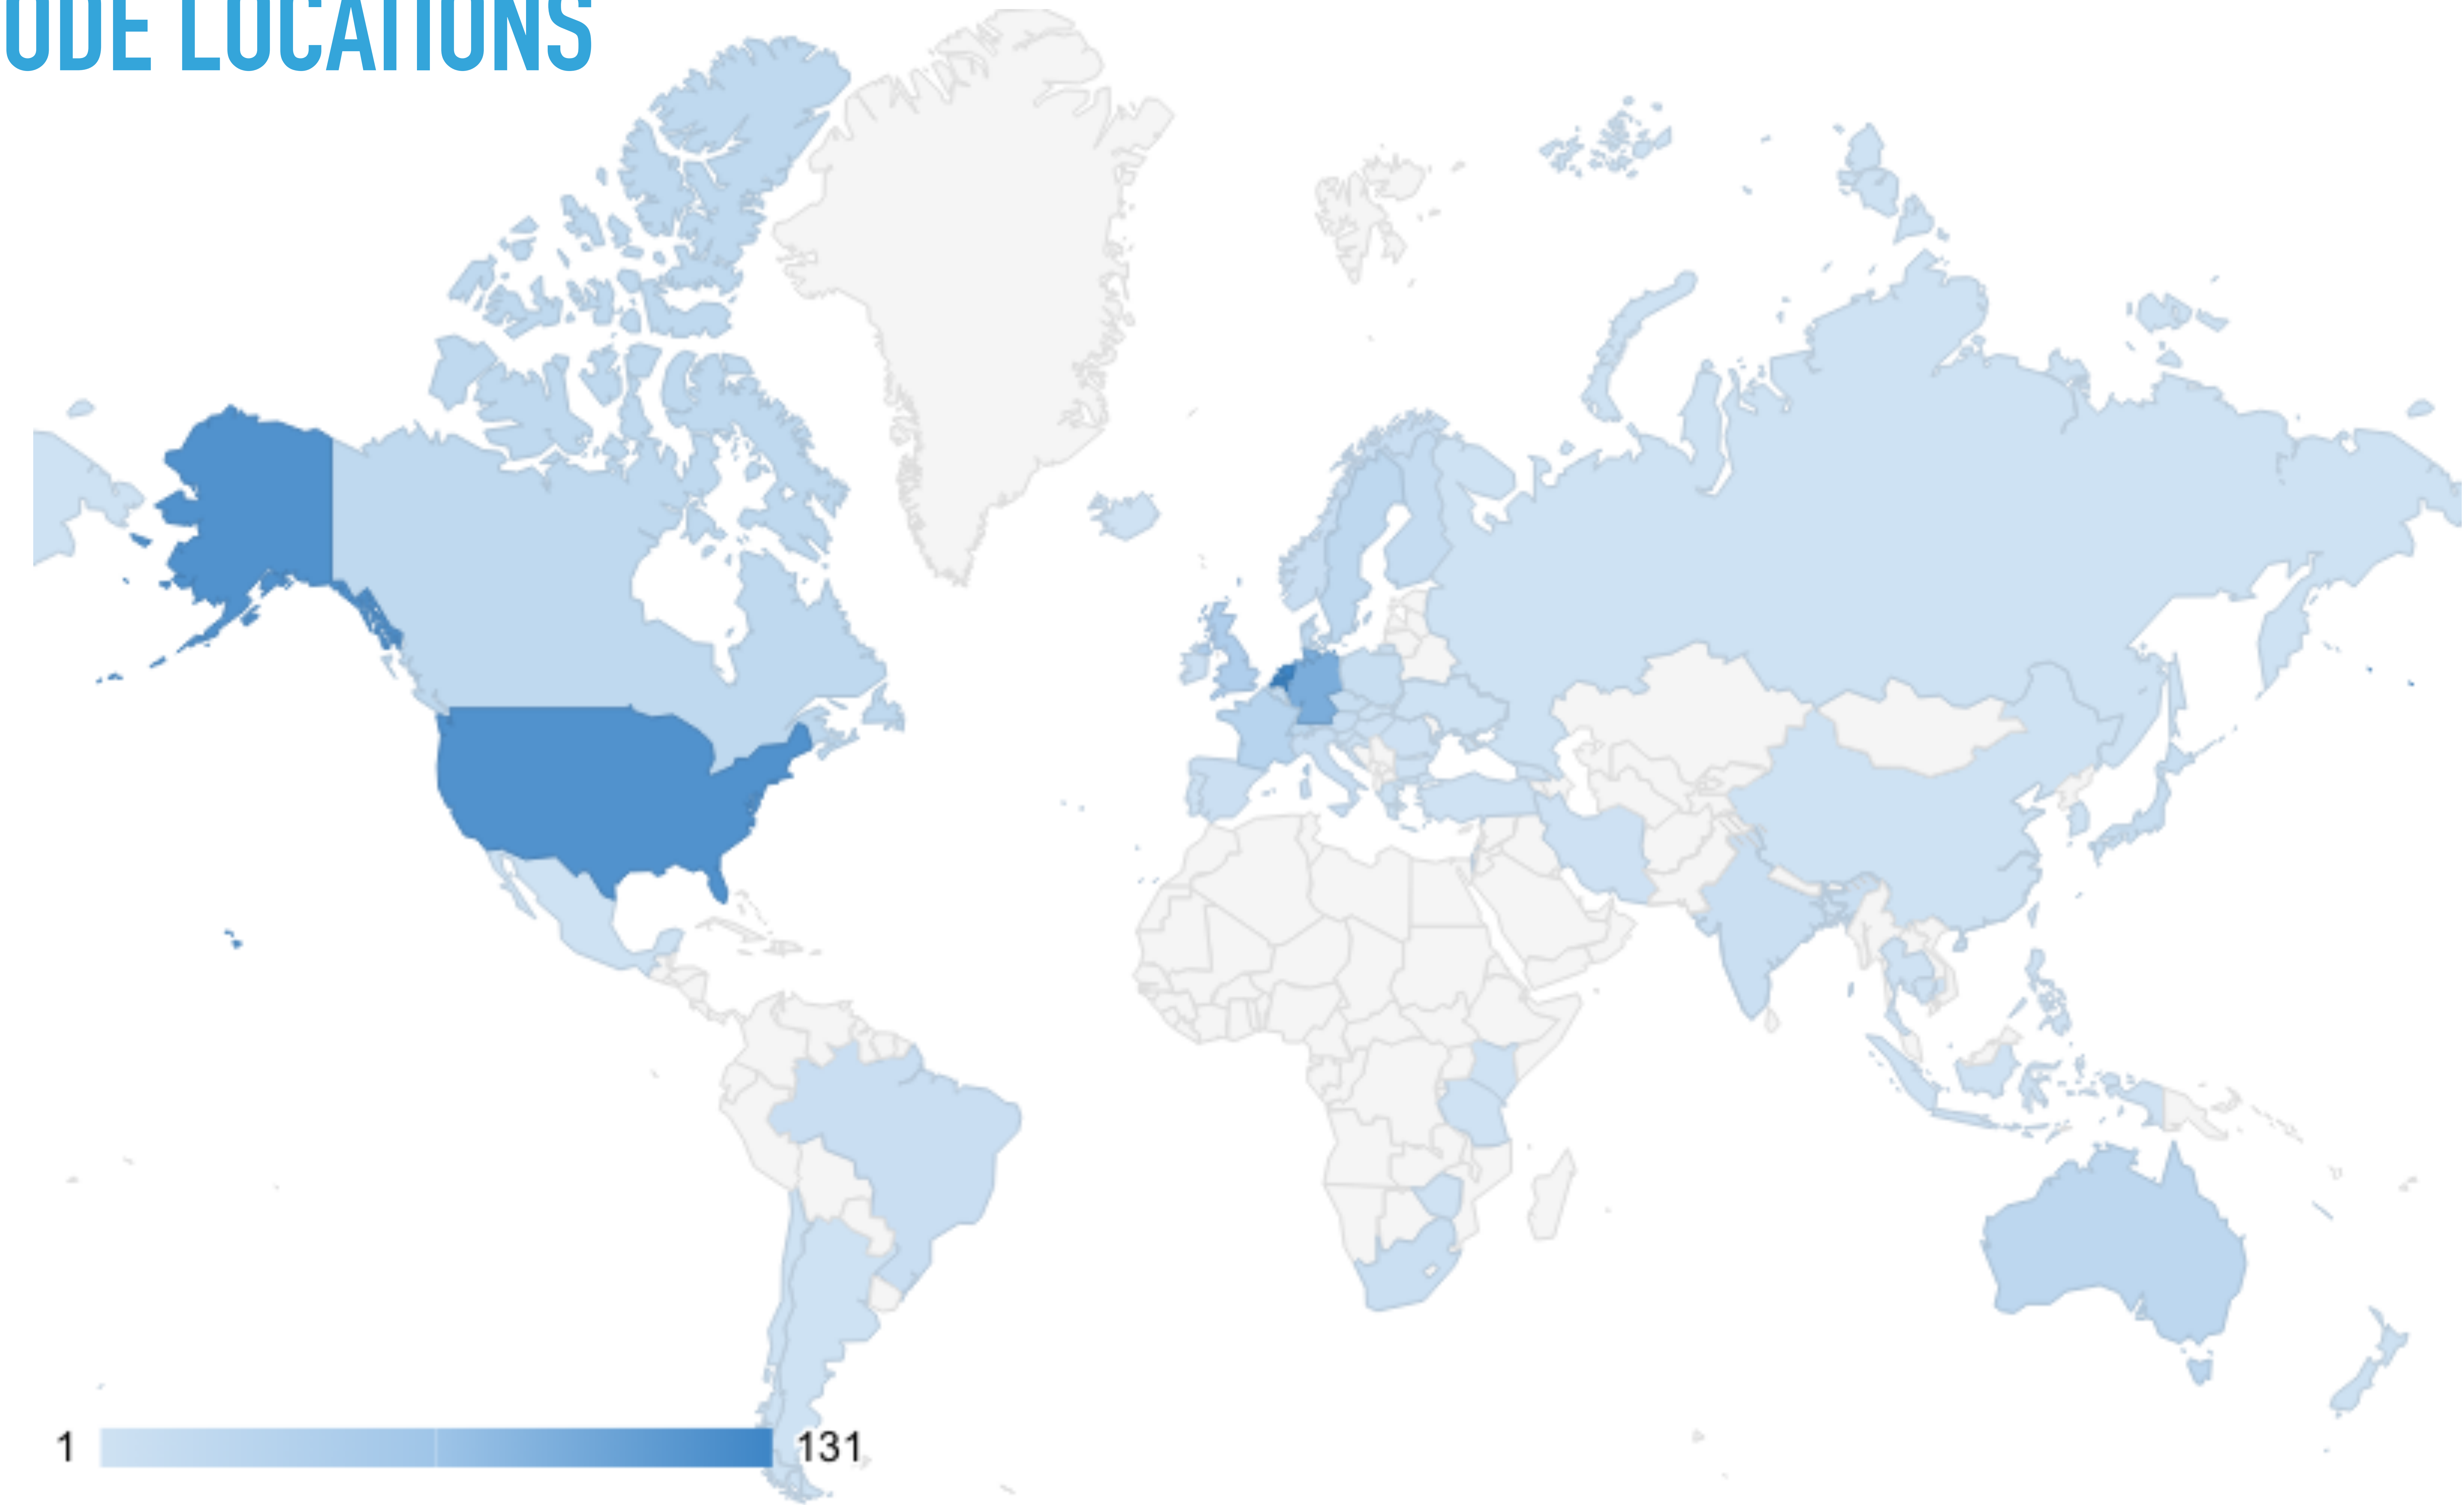

NLNOG LOOKING GLASS: BGP ROLES (RFC 9234) SUPPORT

- Defines the relationship between BGP peers to prevent route leaks
- Introduces BGP attribute **only-to-customer** (OTC)
- NLNOG LG advertises itself as **customer**
- Peers should have role **provider**
- Roles are shown in peer details
- OTC attribute is shown in route details
- Not all router OS'es support it yet ☹️

Peer details for 'MASSAR-v4'

Remote ASN	57777 NLNOG RING node: massar01 (CH)
Remote IP	185.173.128.2
BGP State	Established
Last up/down	1d00h55m
Prefixes received	2,748,625
Capabilities	4-byte ASN refresh ~~enhanced refresh~~ BGP roles
Local BGP role	customer
Remote BGP role	provider
Protocols	IPv4 unicast
Add-path	bidir IPv4 unicast

NLNOG LOOKING GLASS: BGP COMMUNITIES

- Many new BGP community definitions
- Added support for **draft-ietf-grow-yang-bgp-communities** by Martin Pels
  - Definition of BGP communities in a standardised format
  - RIPE NCC has implemented this for AS197000 and AS25152
  - https://github.com/rodecker/draft-yang-bgp-communities
  - https://datatracker.ietf.org/doc/draft-ietf-grow-yang-bgp-communities/

NLNOG LOOKING GLASS: BGP COMMUNITIES

```json
{
  "name": "PREPEND-TO-NA",
  "description": "Example Regular: Prepend X times to North American peers",
  "globaladmin": 65001,
  "localadmin": {
    "fields": [
      {
        "name": "Continent",
        "length": 3,
        "pattern": "200",
        "description": "North America"
      },
      {
        "name": "Prepend-count",
        "length": 1,
        "pattern": "[1-5]"
      }
    ]
  }
}
```

```
teun@ishamael:~/src/draft-yang-bgp-communities                      ⌥⌘4
 ~/src/draft-yang-bgp-communities  main *1 !1 ?4
› cat bgp-communities.txt
65001:2001
65001:2003
65001:3001
65002:2001

 ~/src/draft-yang-bgp-communities  main *1 !1 ?4
› scripts/parser.py bgp-communities.txt bgp-communities.json
65001:2001 - PREPEND-TO-NA (65001:Continent=North America,Prepend-count=1)
65001:2003 - PREPEND-TO-NA (65001:Continent=North America,Prepend-count=3)
65001:3001 - Unknown
65002:2001 - Unknown

 ~/src/draft-yang-bgp-communities  main *1 !1 ?4
›
```

NLNOG RING AS A RESEARCH TOOL

- Tobias Fiebig asked to use the RING to perform DNS measurements:
  - co-location of DNS root servers
  - DNS root server integrity
- The paper will be published on ACM
- While measuring DNS root server integrity they also found… **bitflips** 🤔
- Read the paper on https://fiebig.nl/imc-24.pdf

The Roots Go Deep: Measuring '.' Under Change

Florian Steurer
Max Planck Institute for Informatics
Saarland University
Saarbruecken, Germany
fsteurer@mpi-inf.mpg.de

Daniel Wagner
DE-CIX
Frankfurt, Germany
daniel.wagner@de-cix.net

Danny Lachos
BENOCS GmbH
Berlin, Germany
dlachos@benocs.com

Anja Feldmann
Max Planck Institute for Informatics
Saarbruecken, Germany
anja@mpi-inf.mpg.de

Tobias Fiebig
Max Planck Institute for Informatics
Saarbruecken, Germany
tfiebig@mpi-inf.mpg.de

Abstract

As the entry point to the DNS hierarchy, the DNS root zone, served by the DNS root server system, is essential for the Internet. It consists of 13 deployments managed by 12 independent root server operators. Due to its importance, the root zone deserves special scrutiny, which it has received from researchers and operators alike.

In this study, we measure all root servers over a period of 174 days from 675 vantage points in 523 networks and 62 countries using IPv4 and IPv6. Using this data, we first investigate the *co-location* between root servers, finding that almost 70% of clients observe co-location of at least two servers. Second, we monitor the *integrity* of zone transfers, finding rare issues like bitflips or stale zone files. Finally, by enriching our data with passive ISP and IXP data, we quantify the *role of IPv6* for performance and behavior under change, finding that even seemingly similar subsets of root servers can differ considerably.

CCS Concepts

• **Networks → Network measurement.**

Keywords

DNS, Root Zone, Anycast

ACM Reference Format:
Florian Steurer, Daniel Wagner, Danny Lachos, Anja Feldmann, and Tobias Fiebig. 2024. The Roots Go Deep: Measuring '.' Under Change. In *Proceedings of the 2024 ACM Internet Measurement Conference (IMC '24), November 4–6, 2024, Madrid, Spain.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3646547.3689008

1 Introduction

The root zone is the top of the DNS hierarchy, containing the delegations to the top-level domains. DNS root servers MUST answer queries for the root zone [6], providing a crucial function for DNS and the Internet. RSSAC037 [16] reflects this importance, defining stability, reliability, and resilience goals for root server operations.

To remain fast and reliable in a growing Internet, the scale of the root server system (RSS) steadily grew. As of 2023-12-24, the RSS consists of 1750 instances, operated by 12 independent operators, and serving tens of billions of queries per day [40].

However, such a large deployment may lead to co-location of servers, as it is attractive to deploy instances at locations with good (local) connectivity, such as IXPs. Co-location and the reuse of last hop infrastructure may reduce the redundancy of the system and consequently, negatively affect stability and reliability. Thus, we examine: *How much server co-location exists in the RSS? (RQ1).*

Using active traceroute measurements, we find that co-location is prevalent with almost 70% of clients observing co-location of two or more root servers and some clients being routed to sites with 12 root servers present. While *not* questioning the reliability of the system as a whole, our results indicate that diversifying last-hop infrastructure at certain sites may be worthwhile.

As one of the first systems to deploy IP anycast, and due to the availability of rich data sources [11], the RSS became one of the most popular systems to study the behavior of anycast in practice. Existing studies have investigated performance [20, 38], routing stability [20, 31] or how resolvers react to changes in the RSS [24].

However, existing studies of the root servers' anycast deployment focus on IPv4. To the best of our knowledge, there is no study which comprehensively examines these characteristics for *all* root servers using IPv6. It remains unclear whether results obtained via IPv4 are applicable to IPv6. This work aims to close this gap, answering the question: *What are the differences in the root servers' performance and behavior between IPv4 and IPv6 (RQ2).*

Utilizing data from a large scale active measurement, we show that clients of individual servers are up to 40% (`g.root`) more likely to experience changes of the contacted anycast site when querying via IPv6. While we observe that the overall geographical distance from clients to the contacted anycast sites is comparable to IPv4, we find differences in the experienced RTTs based on the clients location. For example, even though `i.root` and `l.root` have a similar number of replicas deployed in South America, clients experience