



Anycast Discovery: Daily mapping the Anycast landscape for enhanced Internet resilience

RIPE FUNDS 2023 – 2 MAY 2024

RAFFAELE SOMMESE – REMI HENDRIKS

UNIVERSITY OF TWENTE

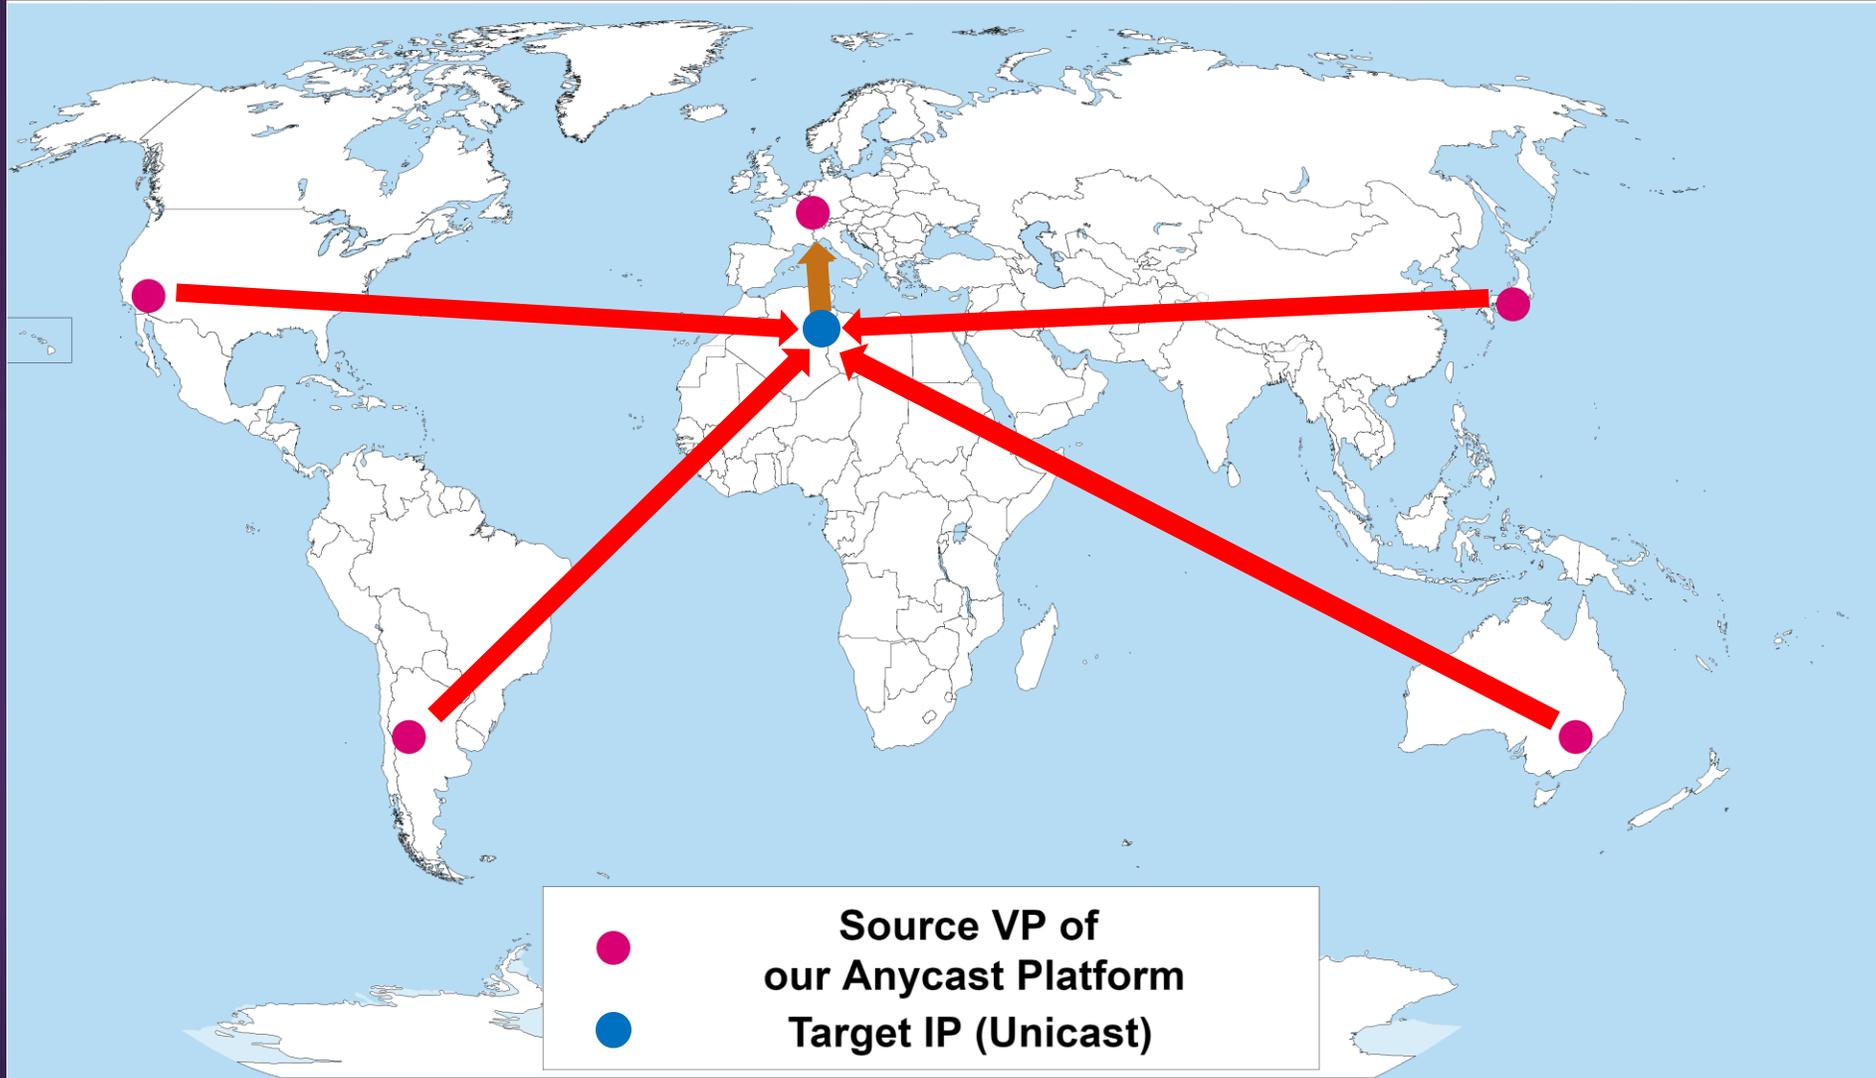
Why Anycast Census?

- ▶ Anycast is one of the most effectively distribution and resilience technique used in the world.
- ▶ The working principle of Anycast relying on BGP makes it opaque to the rest of the internet.
- ▶ Mapping the evolution of the adoption of Anycast at scale is fundamental for the analysis of the development of the global Internet.

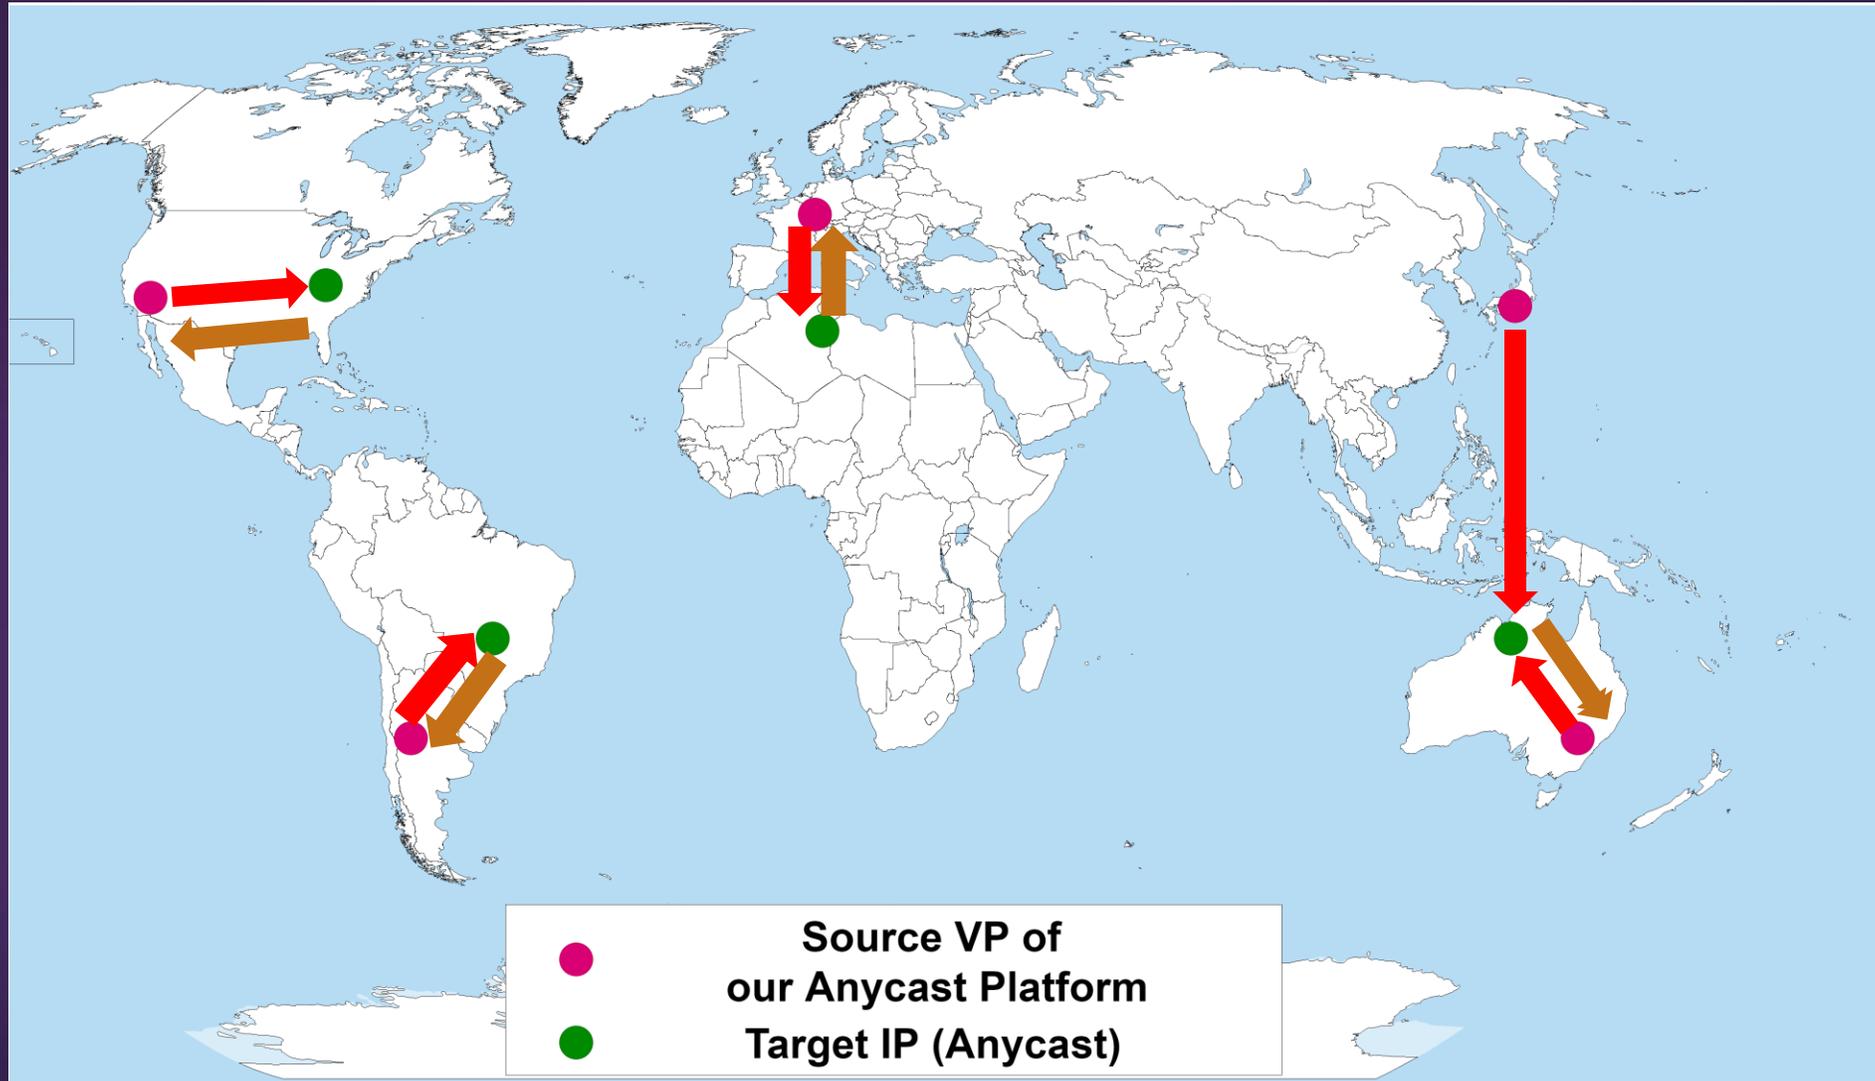
A «new» approach to measure Anycast: MAnycast2

- ▶ Developed in an IMC2020 submission.
- ▶ We leverage the concept of using anycast to measure anycast:
 - ▶ Pinging a unicast address from an anycast network results in packet to be routed always to a SINGLE node, regardless of the source anycast site.
 - ▶ But pinging an anycast address from an anycast network results in packet routed to different MULTIPLE nodes, depending on the source site.
 - ▶ We leverage this behaviour to identify anycast networks.

Responses on a single VP for unicast target



Responses on multiple VPs for anycast target



Project goals

- ▶ Expanding the coverage beyond ICMP probing, encompassing both service-aware, non-service-aware measurements and IPv6.
- ▶ Building a stable anycast infrastructure to facilitate MAnycast² measurements.
- ▶ Conducting a daily anycast census by combining both methodologies and making data publicly available to the community

Project timeline

Month 1:

- ▶ Project kick-off

Month 2-3:

- ▶ Develop enhancements to support TCP/UDP/IPv6 scanning and "service-aware" scans

Month 3-5 (current):

- ▶ Develop a pipeline for the daily census of the IPv4 and IPv6 space

Month 6-10:

- ▶ Enable data storage and sharing mechanisms via public repository and S3-compatible storage.
- ▶ Testing and validation

Month 11-12:

- ▶ Engage with researchers, network operators, and stakeholders to gather feedback
- ▶ Academic publication submission-ready on the results collected from the process and the improved methodology.

Our Team @ DACS



Remi Hendriks
PhD Candidate
University of Twente



Raffaele Sommese
Assistant Professor
University of Twente



Mattijs Jonker
Assistant Professor
University of Twente



Roland Van Rijswijk
Full Professor
University of Twente

Existing Partners

The original MAnycast² scanning pipeline has been implemented with the collaboration of:

- ▶ SIDN Labs
- ▶ CAIDA
- ▶ ISI/USC

Benefits for the RIPE community

- ▶ Reduced burden on the RIPE Atlas infrastructure: avoiding costly latency-based measurements to detect anycast
- ▶ Advancing Anycast Detection
 - Researchers will benefit from enhanced insights into anycast deployments enabling more accurate assessments and optimization of resilience techniques.
 - Network operators within the RIPE community can use the knowledge gained from the shared data to improve the resilience and performance of critical network services

Current Status

Improved measurement tool

- ▶ Measuring tool designed for MAnycast2 measurements
 - Scalable
 - Efficient
 - Robust
 - Responsible
 - Increased precision (reduced number of FPs)
 - Increased coverage (IPv4, IPv6, TCP, DNS, ICMP, ..)

MAnycast2

▶ Pros:

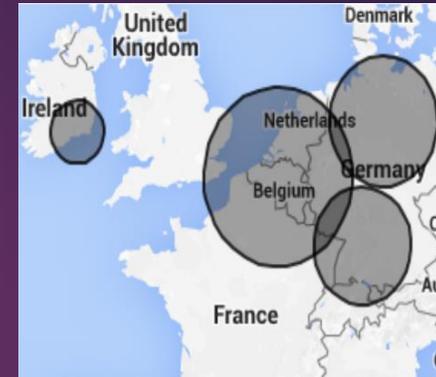
- High recall (low FN rate compared to previous work)
- Scalable, efficient, low probing cost

□ Cons:

- Has FPs (unicast responding to multiple sites)

iGreedy

- ▶ Latency based GCD technique
- ▶ Pros:
 - Detection, enumeration, geolocation of anycast
 - Low FP rate
- ▶ Cons:
 - Performance dependent on # and distribution of VPs
 - Large probing cost
- ▶ Works great, but too costly
 - Not suitable for daily measurements at Internet scale



Daily census

(combining methodologies)

▶ Hitlist

- IPv4: ISI/USC hitlist ~5.6 million /24s
- IPv6: TUM, OPENIntel -> ~800k /48s

▶ Entire hitlist

- MAnycast2 census using TANGLED anycast platform
- iGreedy census using TANGLED

▶ MAnycast2 target hitlist (~30k v4, ~6k v6)

- Intensive iGreedy follow-up measurement
 - Ark
 - RIPE Atlas (?)

Conclusion

- ▶ This project will help the Internet community to gain a fine-grained insight on the deployment of Anycast on the global Internet
- ▶ Thanks to RIPE Funding, we are transforming and extending an academic publication in a sustained long-term effort to provide a valuable dataset to the research and operational community
- ▶ We hope other operators and researchers will use this data towards assessment of the resilience and development of Internet.

Questions?

- ▶ remi.hendriks@utwente.nl
- ▶ r.sommese@utwente.nl
- ▶ <https://github.com/UT-DACS/Anycast-Census>