Geo-auditing RIR Address Registrations

RIPE 87

Rob Beverly <rbeverly@cmand.org>
Oliver Gasser <oliver.gasser@mpi-inf.mpg.de>
November 28, 2023

cMAND

What and Why

Regional Internet Registries (RIRs)

- Internet number allocation is *distributed* and *hierarchical*
- IANA allocates large, contiguous IP address blocks (e.g., IPv4 /8) to RIRs
- Five RIRs with regional responsibility:

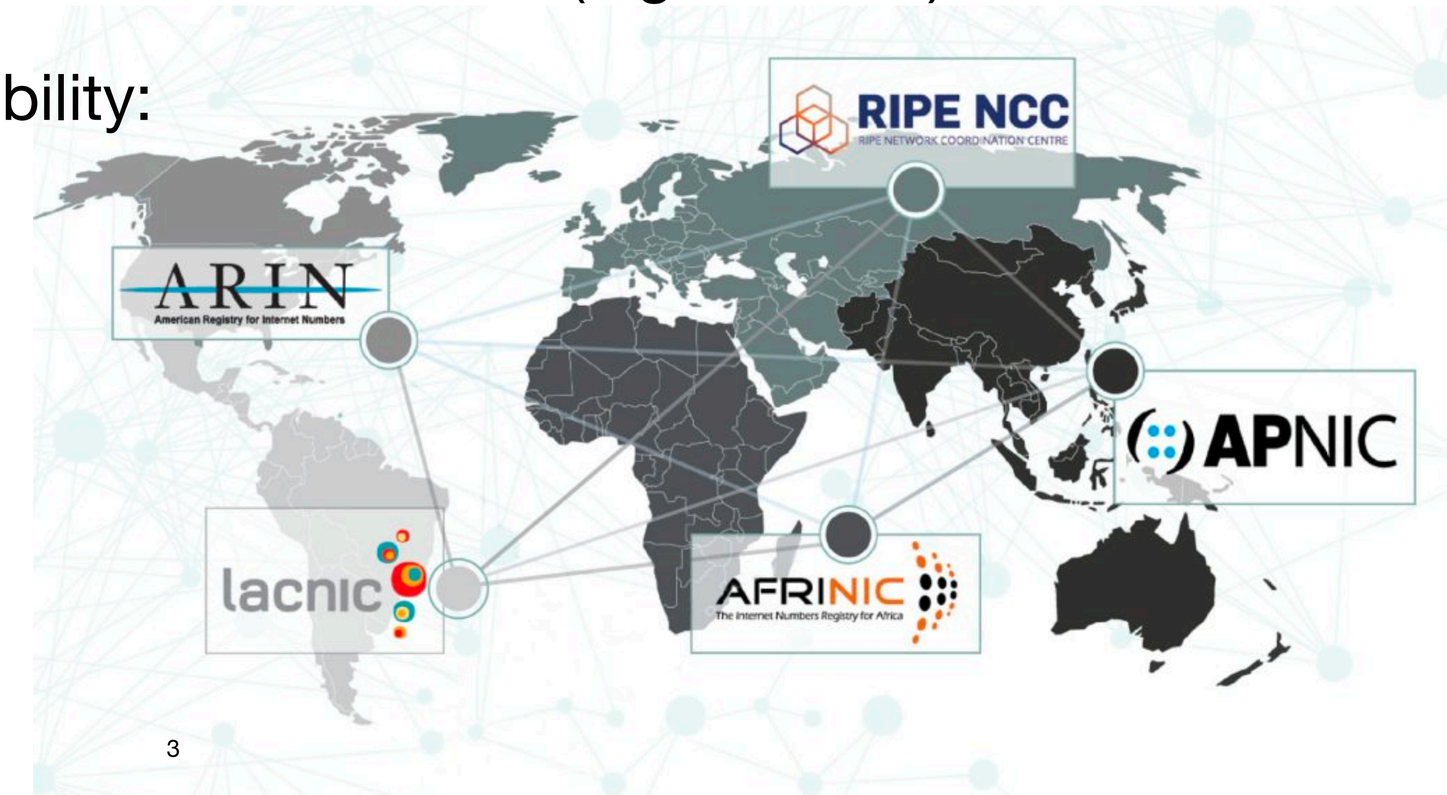

Role of RIRs

- "*The primary role of RIRs is to manage and distribute public Internet address space within their respective regions.*" [**ARIN NRPM, RIPE-738, NRO**]
- Internet numbers registry goals [**RFC 7020**]:
  - *Allocation pool management* (finite resource, uniqueness)
  - *Hierarchical allocation* (efficiency)
  - *Registration accuracy* (to meet operational needs)

Role of RIRs

- "*The primary role of RIRs is to manage and distribute public Internet address space within their respective regions.*" [**ARIN NRPM, RIPE-738, NRO**]
- Internet numbers registry goals [**RFC 7020**]:
  - *Allocation pool management* (finite resource, uniqueness)
  - *Hierarchical allocation* (efficiency)
  - *Registration accuracy* (to meet operational needs)

"A core requirement … is to maintain a registry of allocations … to provide accurate registration information of those allocations in order to meet a variety a operational requirements." RFC7020

Our Work: Geo-Auditing Prefix Registration

1. Examine IPv4 address registry information across the five RIRs
2. Active latency-based IP geolocation of allocated IPv4 prefixes
   - Where are allocated prefixes physically used?
3. Taxonomy of prefix registration geo-consistency
   - Compare physical location to RIR's service region and to registration info?
4. Geo "audit" of prefix registration consistency
   - How geo-consistent are registrations across the RIRs?

Wait! Out-of-region use is allowed!

- Not looking at inter-RIR transfers (publicly logged and vetted by RIRs):
  - Instead, out-of-region use that can only be uncovered via measurement
- Adopt a conservative view of out-of-region use:
  - If used out-of-region, is it at least consistent with the registered organization's location?
- It's complicated: different RIRs have different policies

NRO Comparative Policy Overview

https://www.nro.net/rir-comparative-policy-overview-2023-q3/

- ARIN: “*To receive resources, ARIN requests organizations to verify that it plans on using the resources within the ARIN region*”
- RIPE: “*The network that will be using the resources must have an active element located in the RIPE NCC service region*”
- APNIC: “*permits account holders located within the APNIC service region to use APNIC-delegated resources out of region*”
- LACNIC: “r*equires organizations to be legally present and have network infrastructure in the LACNIC service region to apply for and receive resources*”
- AFRINIC: “*requires organizations/persons to be legally present and the infrastructure from which the services are originating must be located in the AFRINIC service region*”

Third-Party Audit Motivation

- Increase transparency and help community understand where scarce resources are being used
- Quantify extent to which registry information is accurate and can serve operational needs (e.g., security)
- Inform ongoing discussion over “in-region” address use and policy

Third-Party Audit Motivation

- Increase transparency and help community understand where scarce resources are being used
- Quantify extent to which registry information is accurate and can serve operational needs (e.g., security)
- Inform ongoing discussion over “in-region” address use and policy

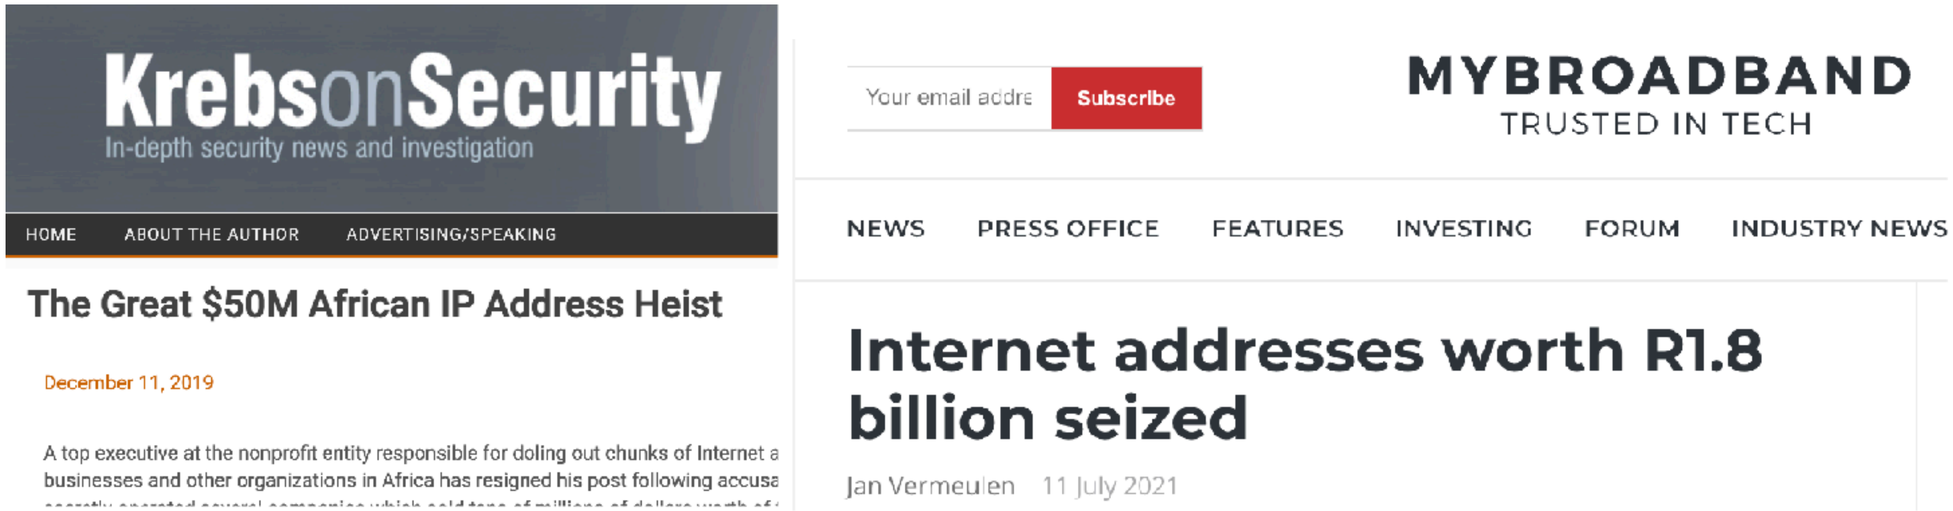

(What this talk is not)

- We recognize:
  - Economic value of IP addresses
  - Need for efficient and equitable use of IP addresses
  - Operational constraints / expedience / messiness of real-world
- Goal is to shed quantitative light on IP address registration geo-consistency
  - **Not** claiming to find policy violations
  - **Not** advocating for policy changes

How

Example

```
inetnum:        79.174.20.0 - 79.174.20.255
netname:        Yunnan-Logame-Technology-Co-Ltd
country:        HK
org:            ORG-YLTC1-RIPE
admin-c:        KY603-RIPE

organisation:   ORG-YLTC1-RIPE
org-type:       OTHER
address:        37k yen chow street sham shui po
address:        unit 708 level 7 dragon center
address:        Hong Kong
```

- /24 in a /8 allocated to RIPE
- Registered owner in Hong Kong (outside of RIPE's region)
- Q: where is this /24 physically?

Example

```
inetnum:        79.174.20.0 - 79.174.20.255
netname:        Yunnan-Logame-Technology-Co-Ltd
country:        HK
org:            ORG-YLTC1-RIPE
admin-c:        KY603-RIPE

organisation:   ORG-YLTC1-RIPE
org-type:       OTHER
address:        37k yen chow street sham shui po
address:        unit 708 level 7 dragon center
address:        Hong Kong
```

- /24 in a /8 allocated to RIPE
- Registered owner in Hong Kong (outside of RIPE's region)
- Q: where is this /24 physically?
  - In RIPE's region?
  - In APNIC's region?
  - In neither RIPE nor APNIC region?

Example

```
inetnum:        79.174.20.0 - 79.174.20.255
netname:        Yunnan-Logame-Technology-Co-Ltd
country:        HK
org:            ORG-YLTC1-RIPE
admin-c:        KY603-RIPE

organisation:   ORG-YLTC1-RIPE
org-type:       OTHER
address:        37k yen chow street sham shui po
address:        unit 708 level 7 dragon center
address:        Hong Kong
```

- /24 in a /8 allocated to RIPE
- Registered owner in Hong Kong (outside of RIPE's region)
- Q: where is this /24 physically?
  - In RIPE's region?
    OK
  - In APNIC's region?
    OK
  - In neither RIPE nor APNIC region?
    INCONSISTENT

RIR Geo-consistency Taxonomy

- Given a prefix we compare:
  - RIR_{Reg}: RIR responsible for allocating the prefix
  - RIR_{CC}: RIR responsible for the country of the registered organization
  - RIR_{Geo}: RIR responsible for the inferred physical geolocation of the prefix

Result	Example RIR_{Reg}	RIR_{CC}	RIR_{Geo}
(*FC*) Fully Geo-consistent	ARIN	ARIN	ARIN
(*CC*) Country Geo-consistent	RIPE	ARIN	ARIN
(*CI*) Country Geo-inconsistent	ARIN	RIPE	ARIN
(*RI*) Registry Geo-inconsistent	ARIN	ARIN	RIPE
(*FI*) Fully Geo-inconsistent	ARIN	RIPE	APNIC

Methodology

Overview

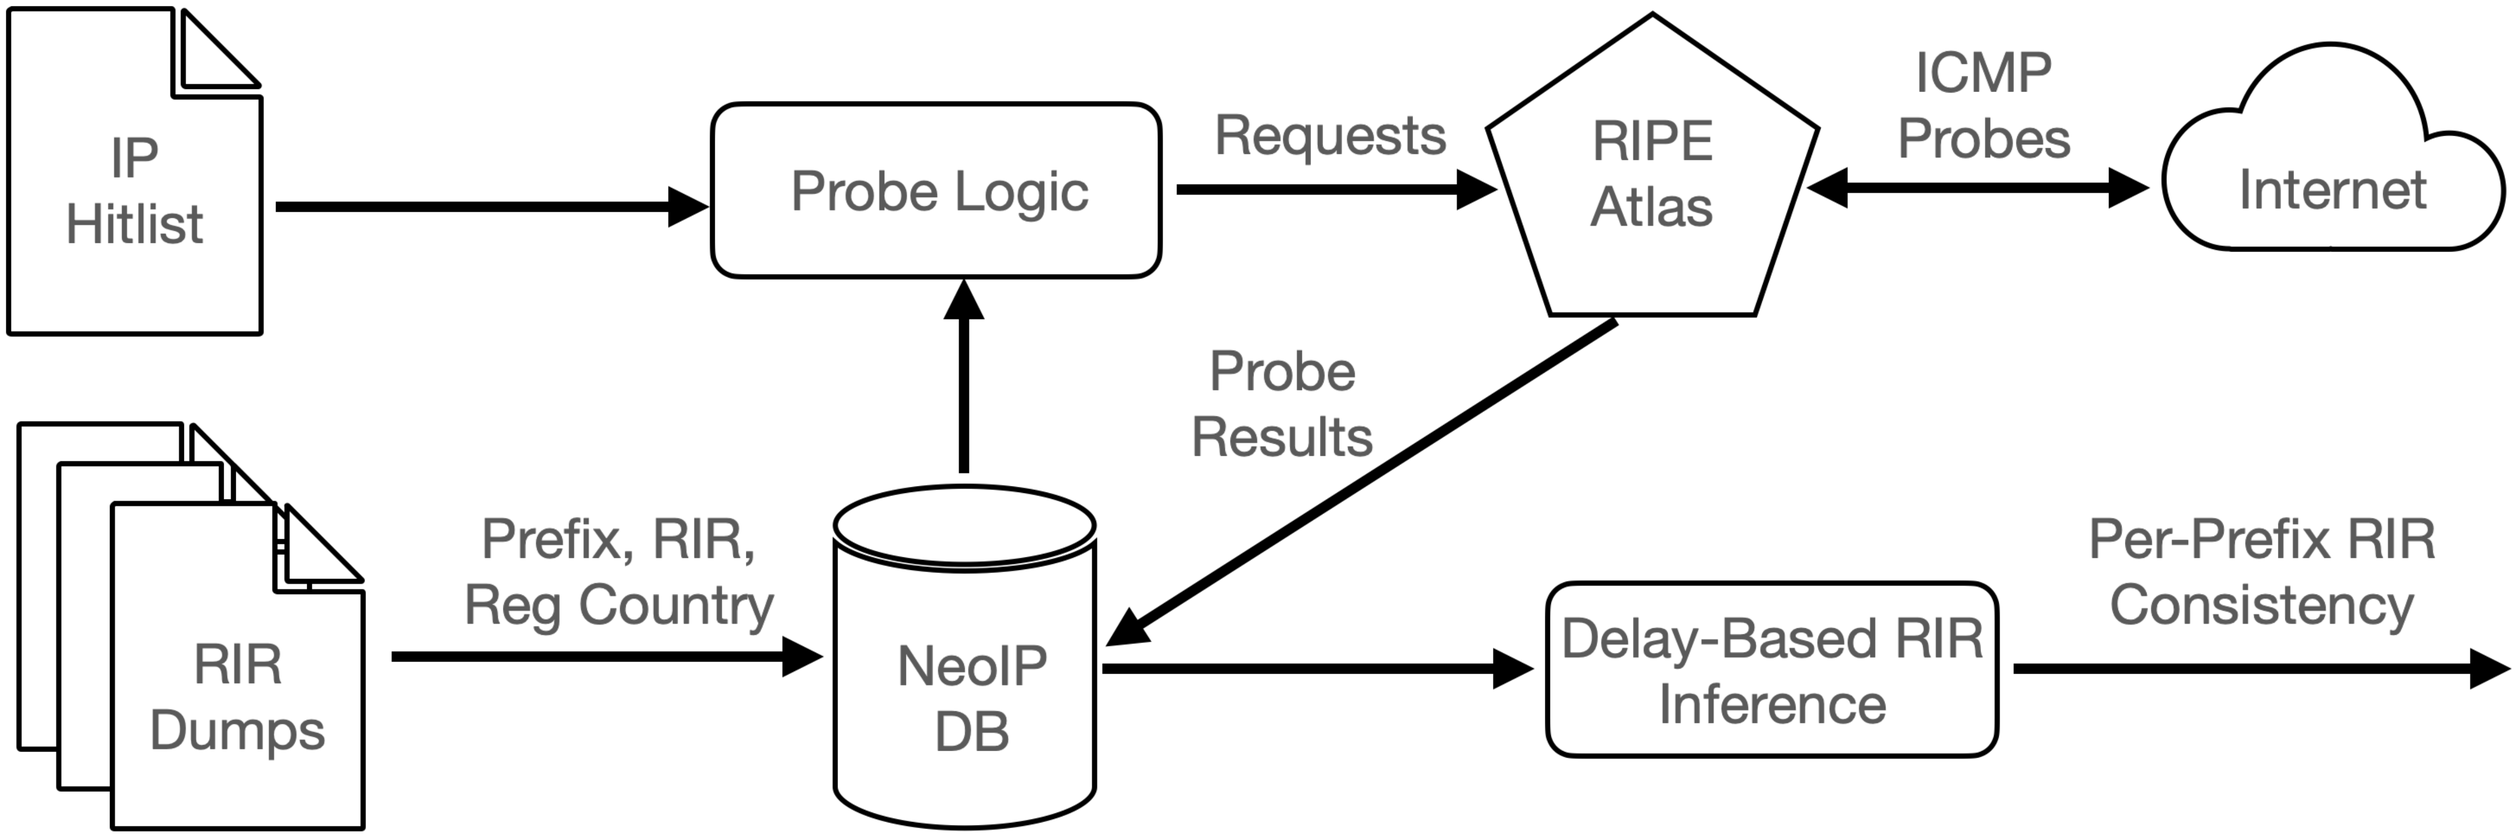

Methodology

Overview

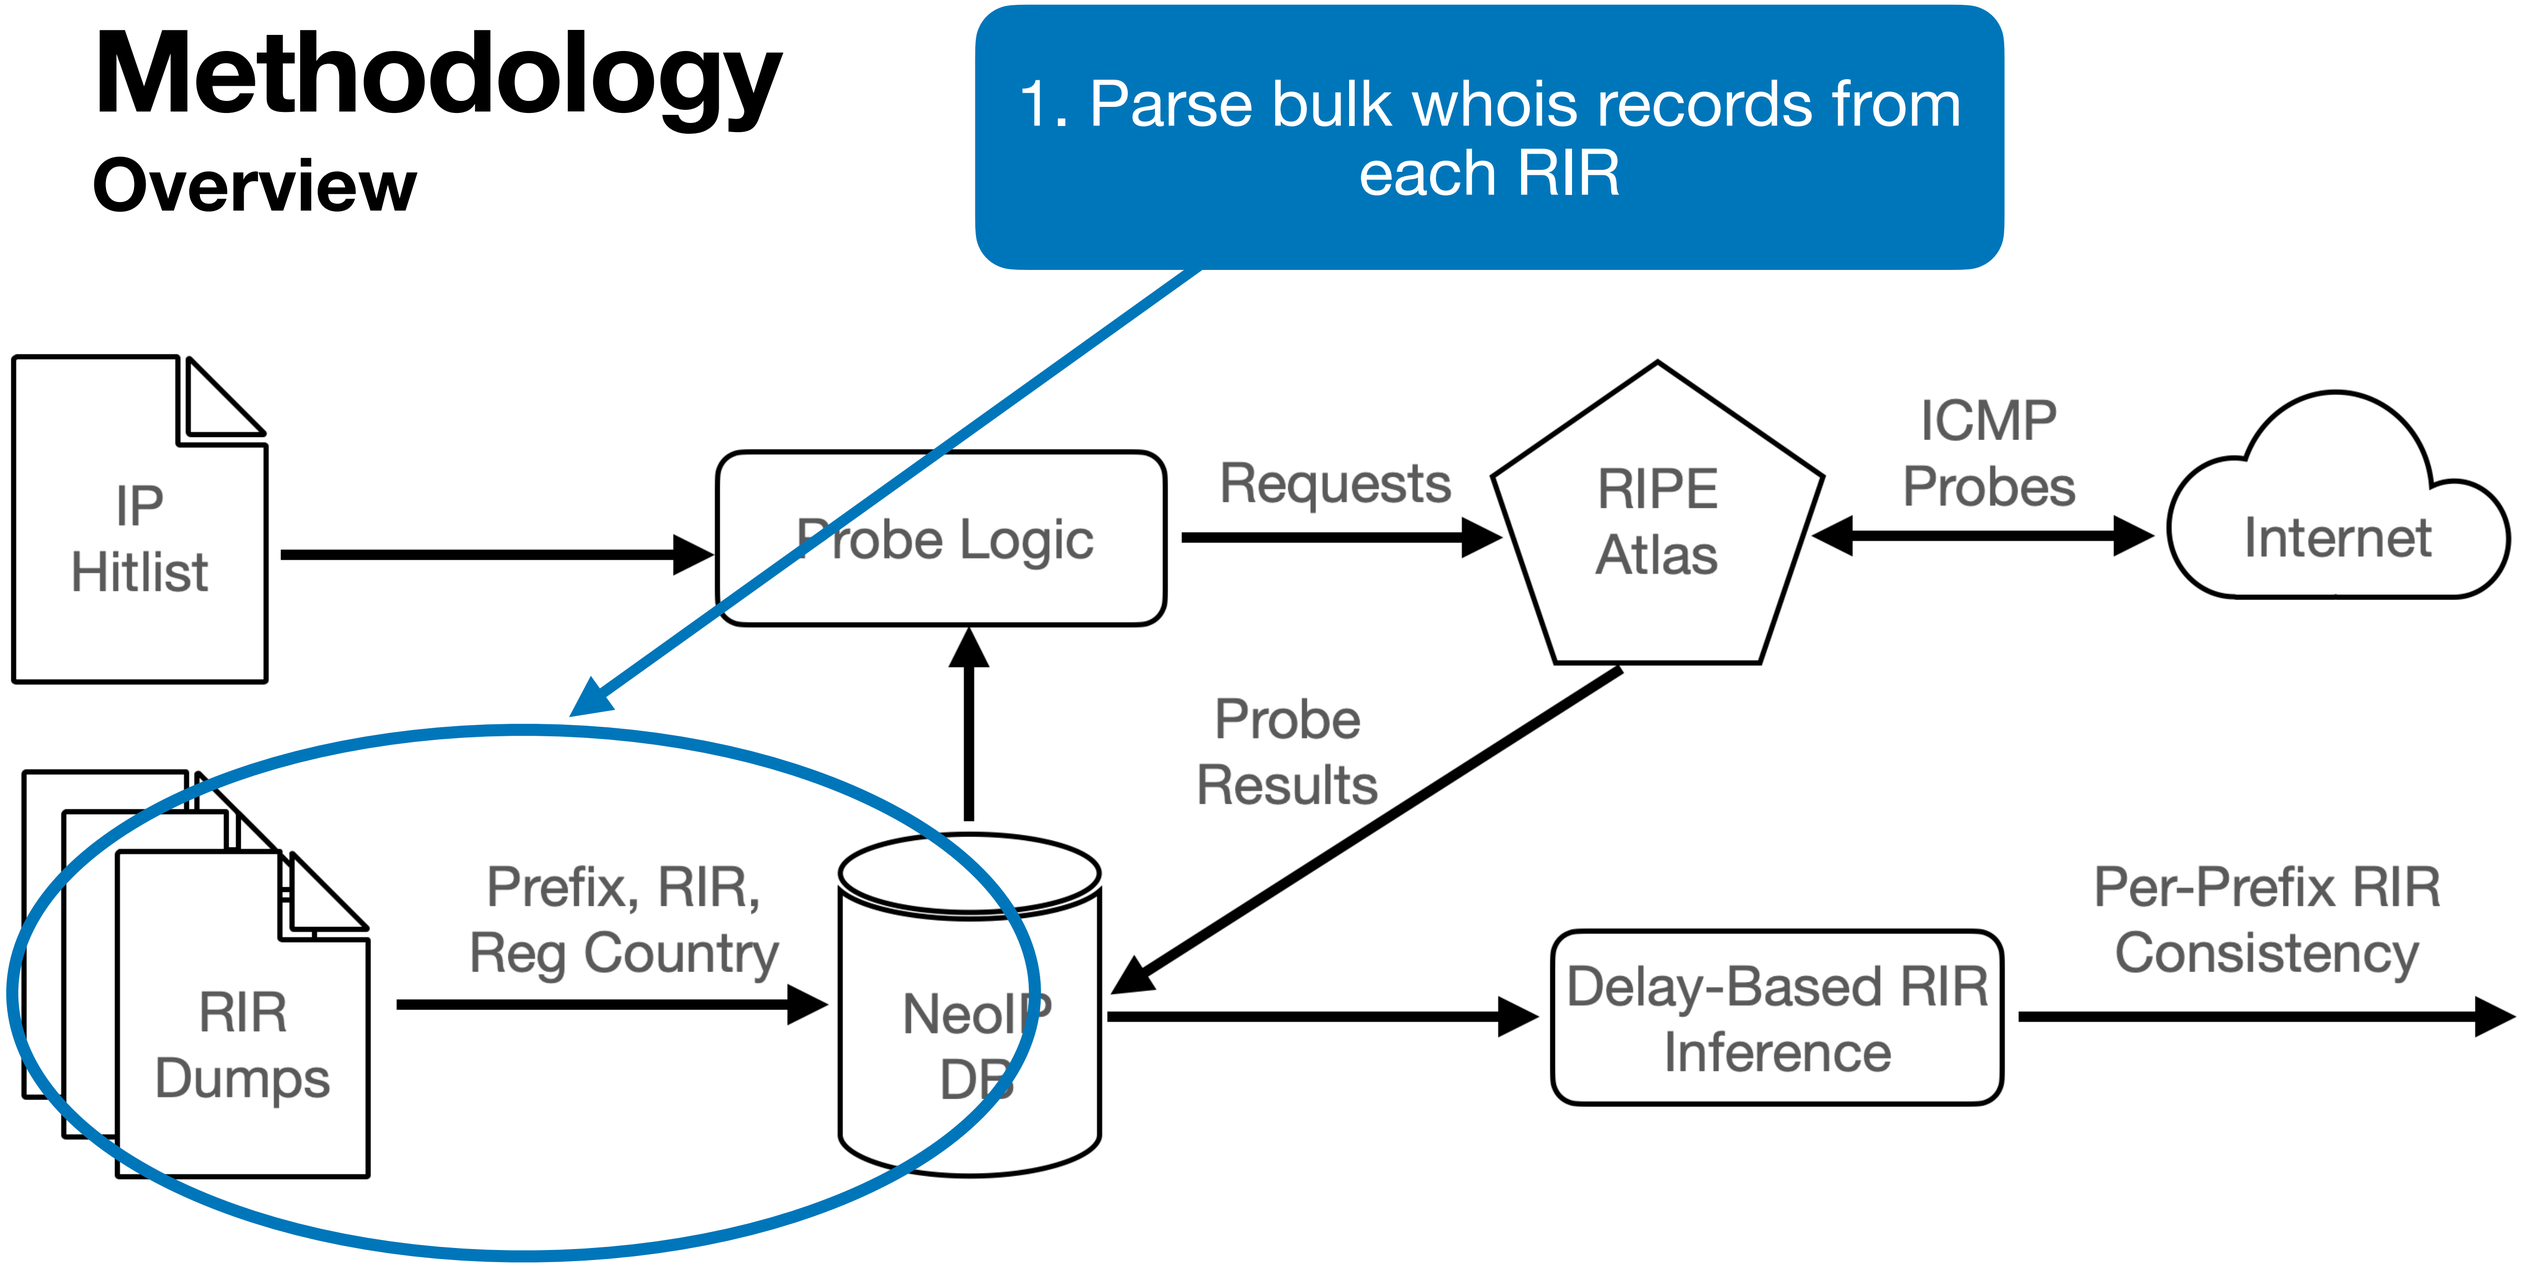

Methodology

Overview

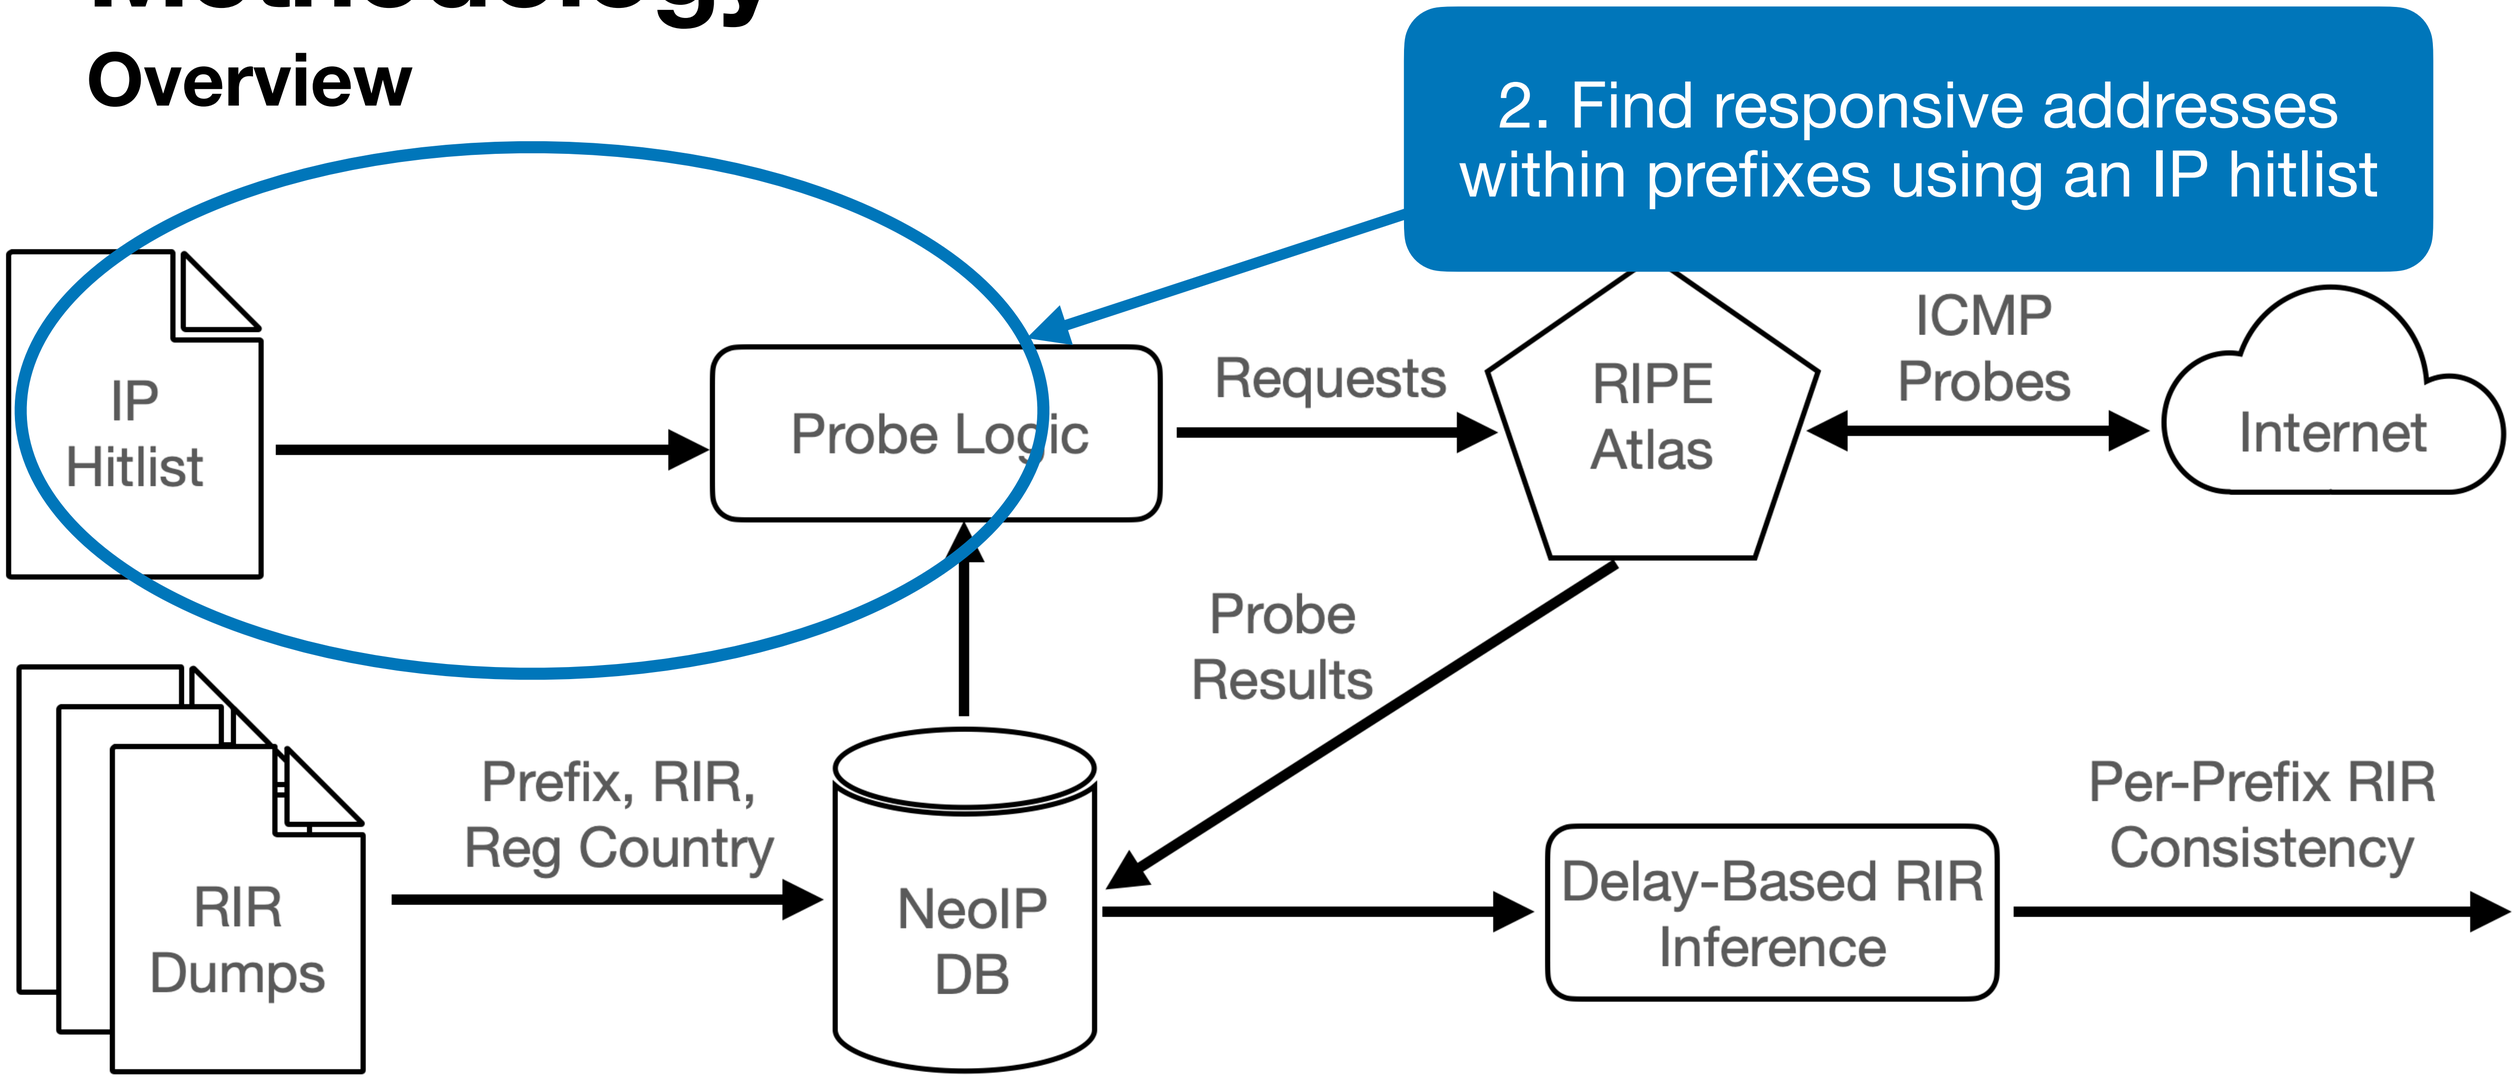

Methodology

Overview

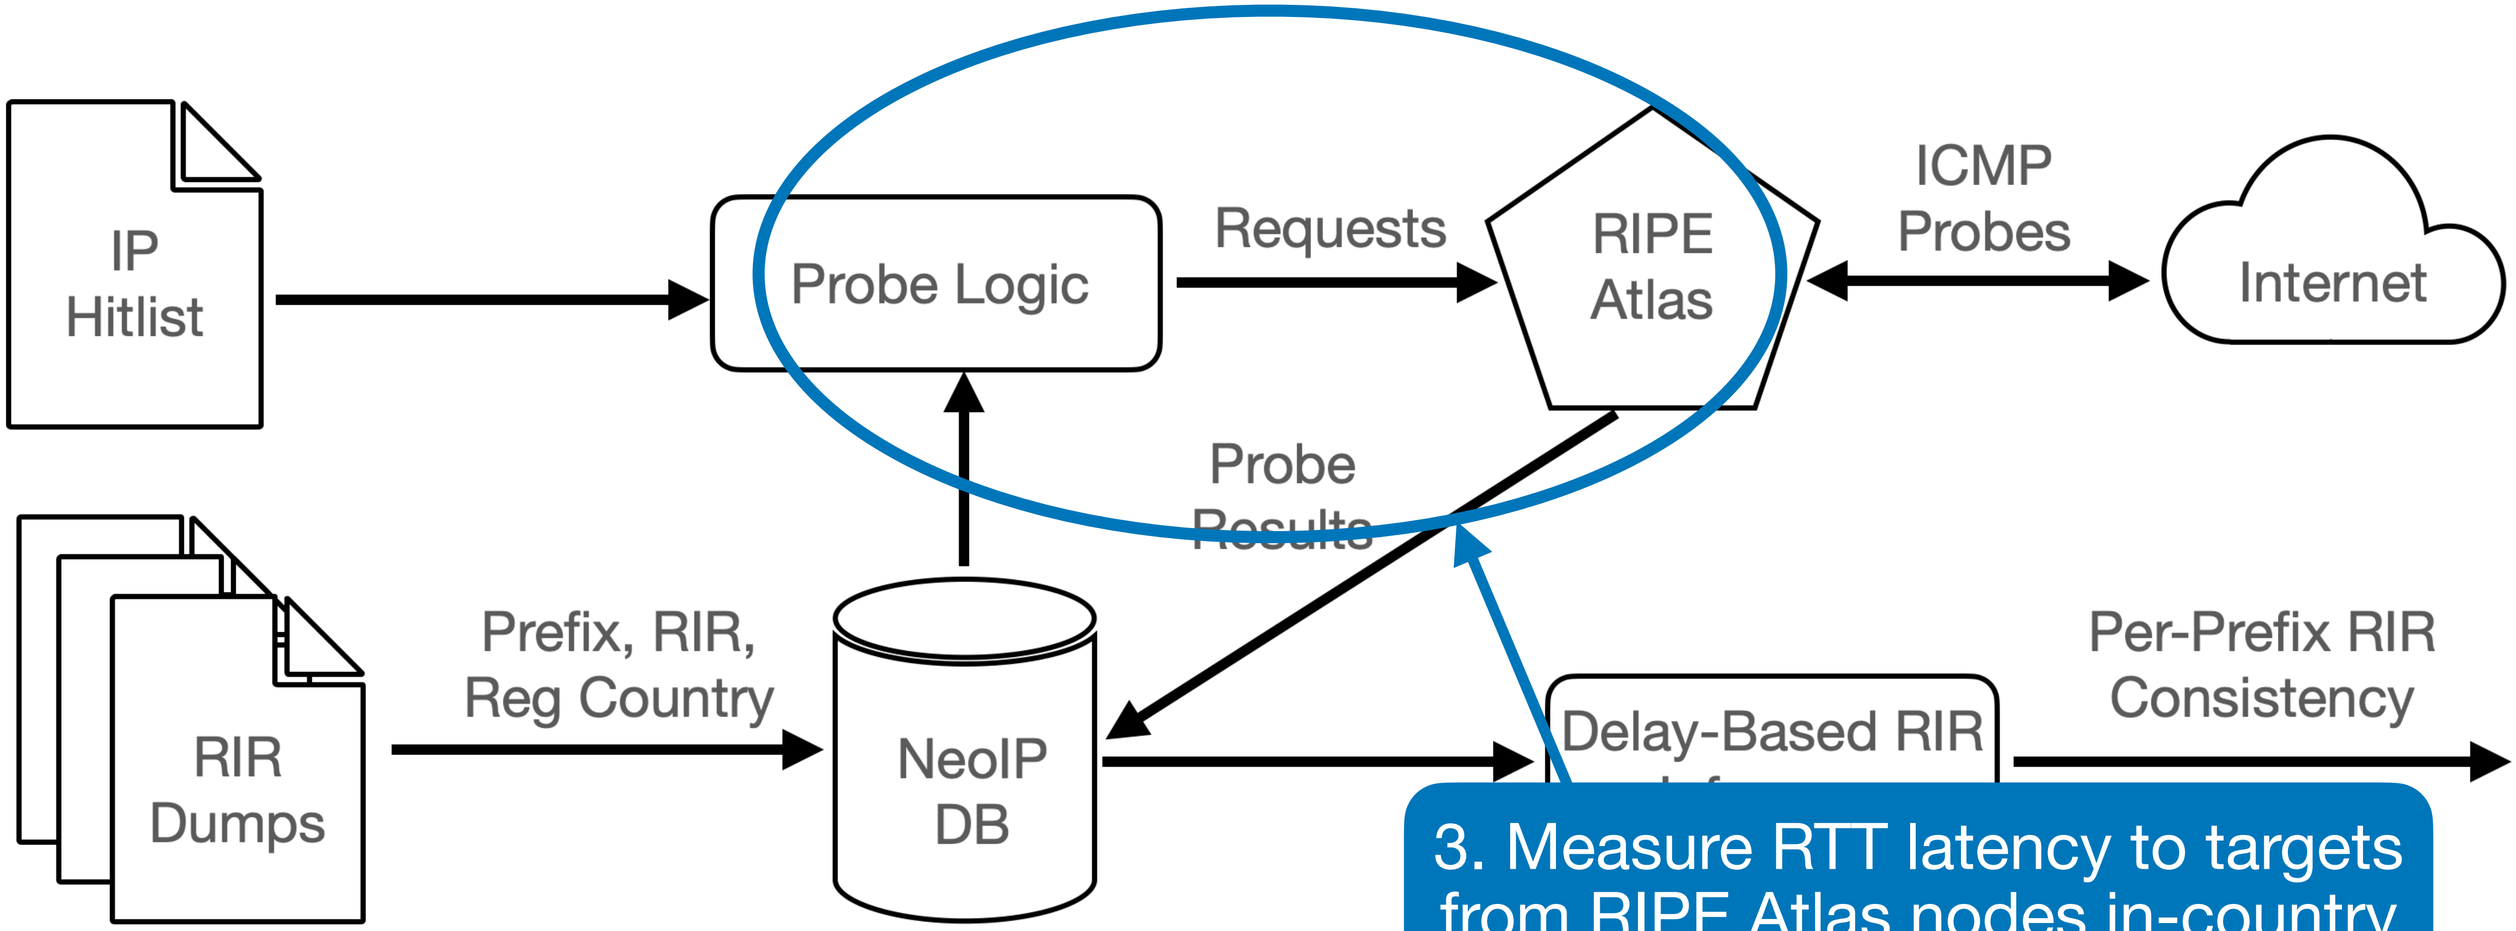

3. Measure RTT latency to targets from RIPE Atlas nodes in-country and on each continent

Methodology

Overview

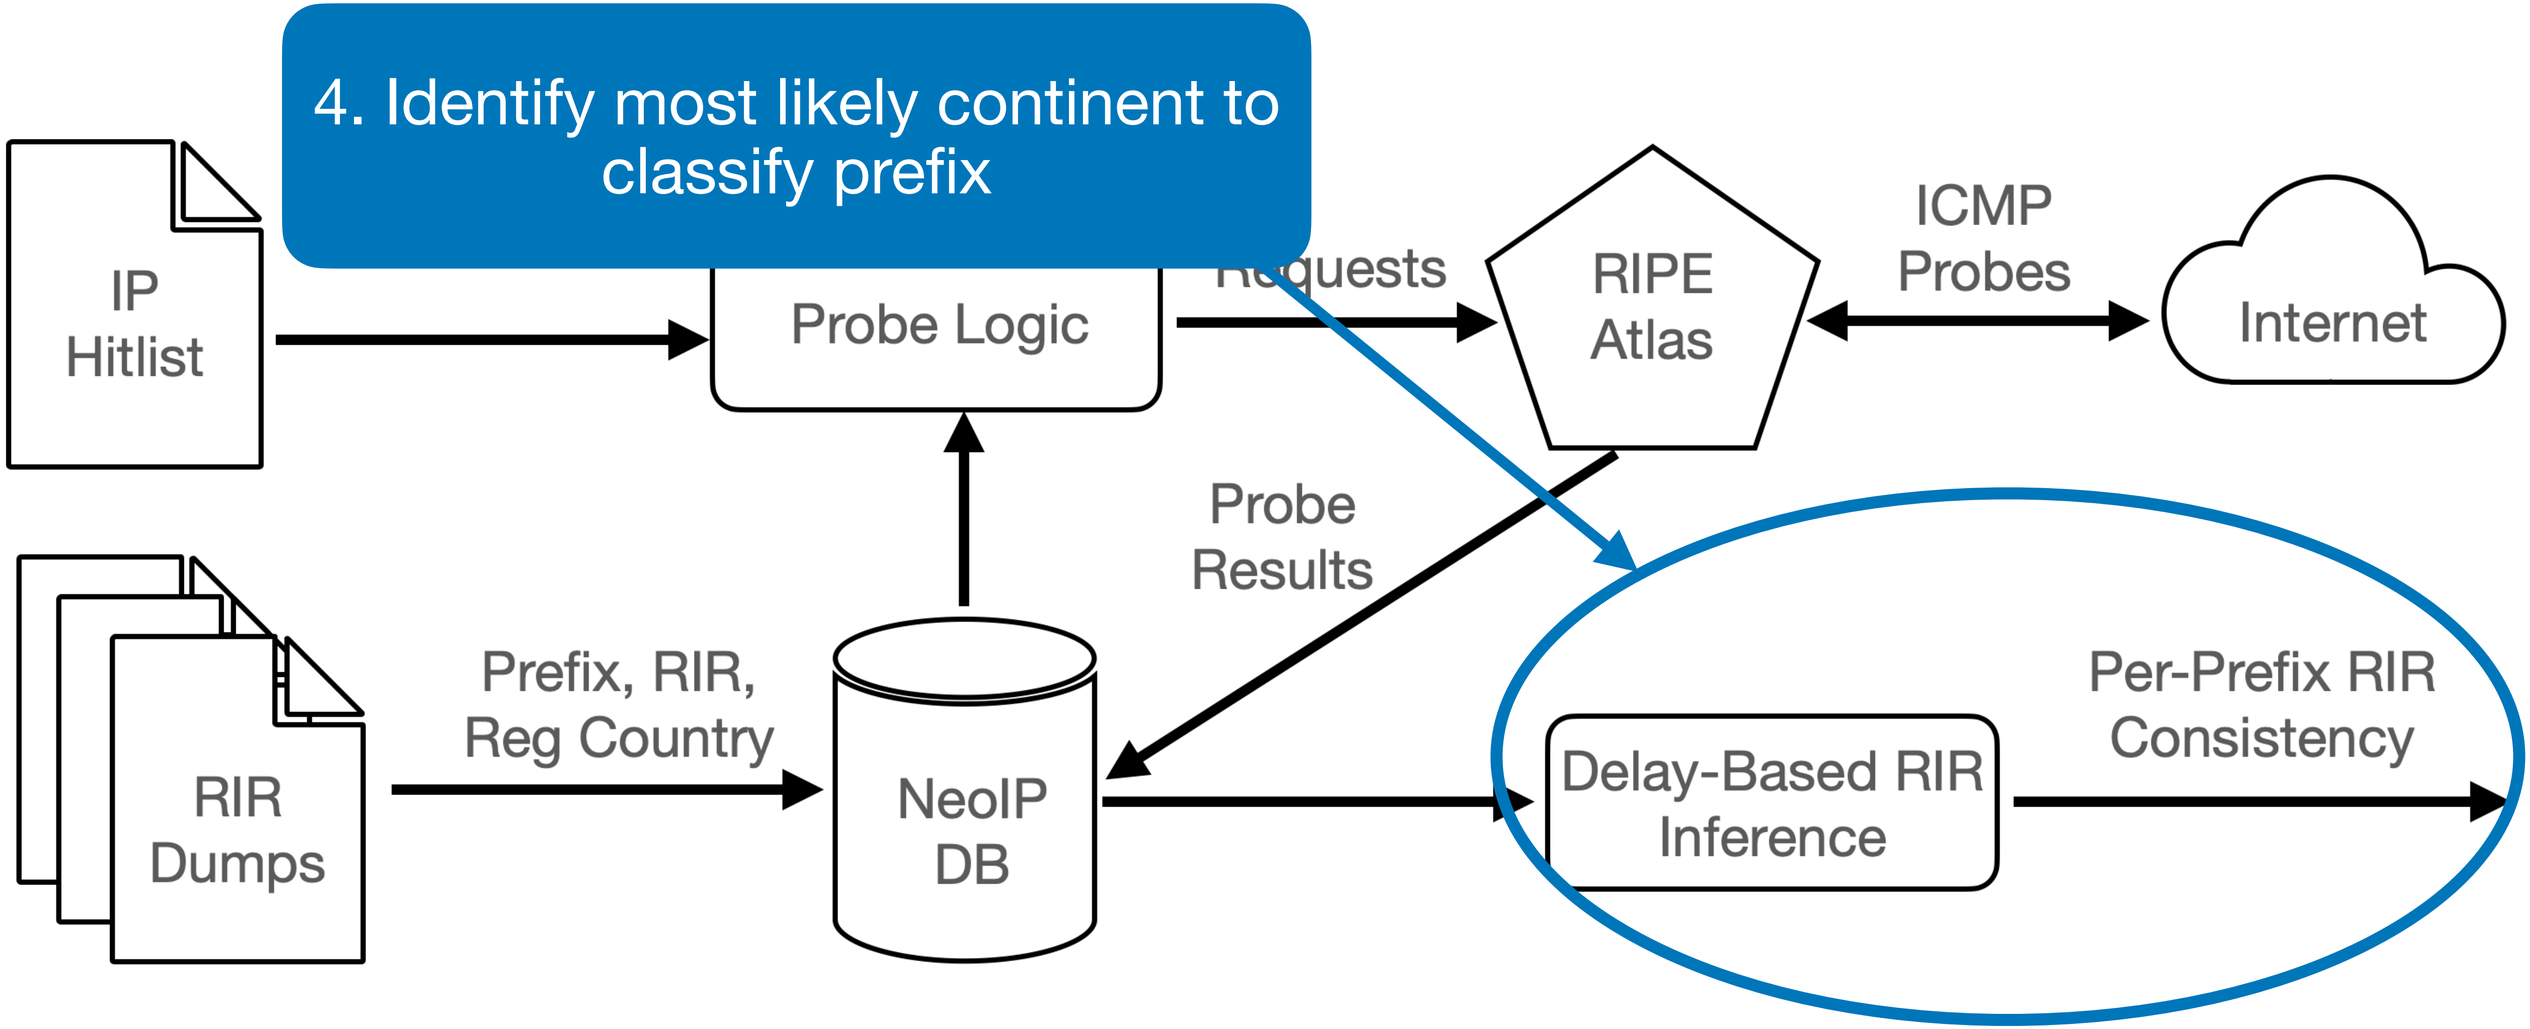

Methodology I

Bulk whois records

- Key-value pairs; different schemas for different RIRs
- Parse prefix and registered organization's mailing address
- Ignore transferred / non-managed records
- Map mailing address countries to the RIR responsible for that country
- Gives RIR_{Reg} and RIR_{CC}

```
NetHandle:      NET-104-148-63-0-1
OrgID:          C05266659
Parent:         NET-104-148-0-0-1
NetName:        WEB-OMEGA-DO-BRASIL
NetRange:       104.148.63.0 - 104.148.63.255
```

```
inetnum:        195.24.192.0 - 195.24.223.255
netname:        CM-CAMTEL-970403
descr:          Data communication and
international
descr:          telecommunication of Cameroon
country:        CM
```

```
inetnum:        185.135.75.0 - 185.135.75.255
netname:        NON-RIPE-NCC-MANAGED-ADDRESS-
BLOCK
descr:          Japan
country:        JP
```

Bulk whois macro stats

RIR	Prefixes (k)	Out-region Prefixes (k)	Addresses (/24s)	Out-Region Addresses (/24s)
ARIN	3,109.8	77.3 (2.5%)	5,491,682	128,546 (2.3%)
RIPE	3,556.7	29.8 (0.8%)	2,925,866	50,579 (1.7%)
APNIC	1,150.8	2.7 (0.2%)	9,136,159	14,327 (0.2%)
LACNIC	66.5	0.3 (0.5%)	251,088	651 (0.3%)
AFRINIC	148.5	21.1 (14.2%)	486,456	23,601 (4.9%)
Total:	8,032.3	131.3	18,291,251	217,705

- April 2023 raw dumps from all five RIRs
- Approximately 8M IPv4 prefix registrations

Inter-RIR region registration is common

- Addresses obtained / registered to organizations outside of the RIR's service region may be explicitly **allowed**:
  - "*ARIN registered resources may be used outside the ARIN service region... provided that the applicant has a real and substantial connection with the ARIN region.*"
  - "The RIPE NCC delegates or registers resources to organizations and individuals that have a need in its service region. The network that will be using the resources must have an active element located in the RIPE NCC service region."

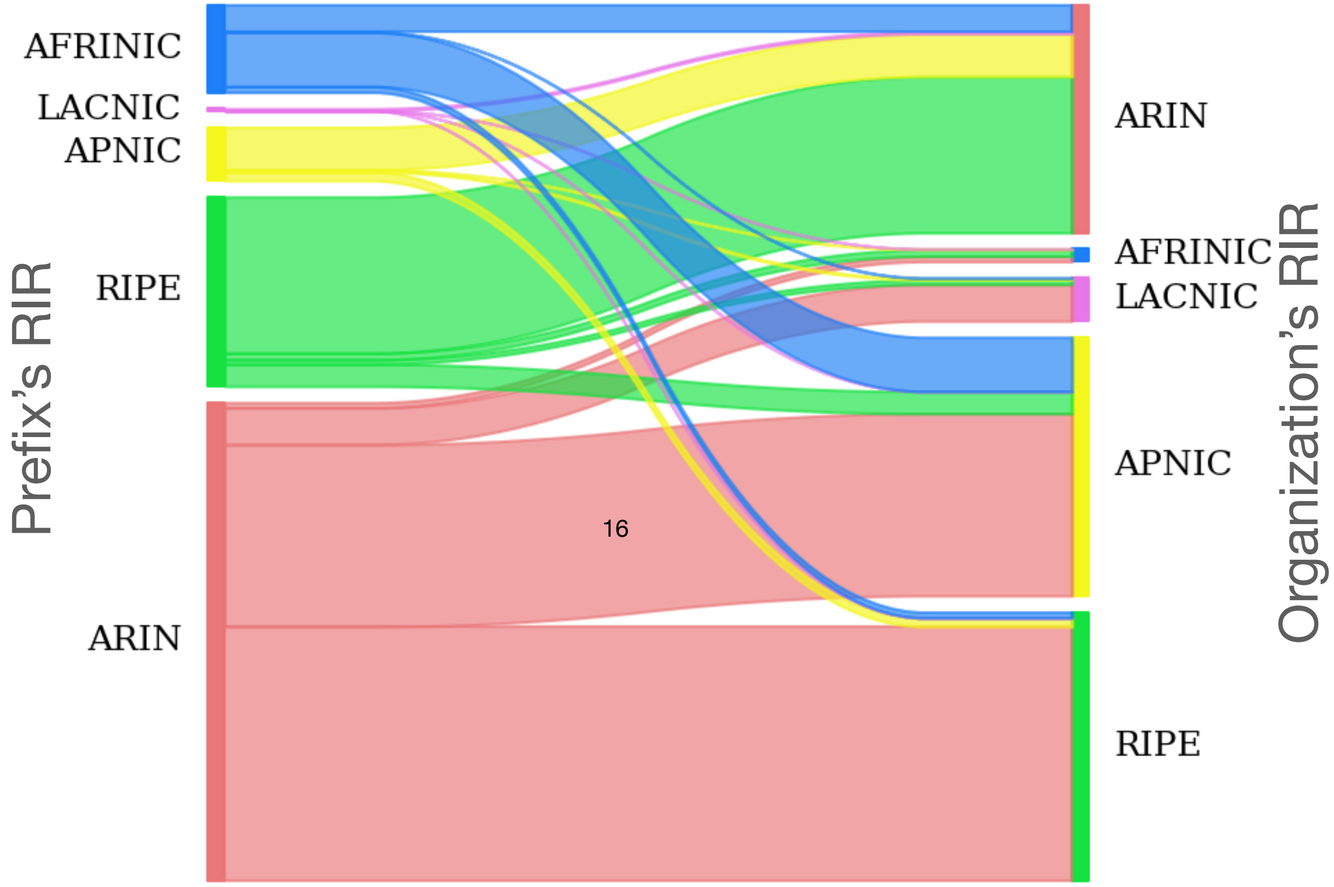

Methodology II
IPv4 Hitlist

- Utilize a “hitlist” of known / likely-responsive IPv4 addresses
- Longest-prefix match hitlist addresses to RIR prefix
  - Ignore prefixes without any responsive addresses
  - Ignore anycast prefixes
- Randomly sample 10k non-anycast prefixes with responsive targets from each RIR (50k total prefixes)

RIPE Atlas

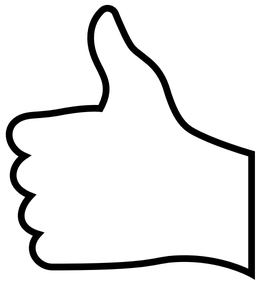

- A big thanks — Atlas is a valuable resource to the community!
- Atlas is essential to our research:
  - Extensive physical coverage, especially in-country
  - Sane and usable API
  - Persistent and FAIR (findable, accessible, interoperable, reusable) measurement results:
    - #cmand, #neo-ip, #neo-ip-20230927

Methodology III
Delay-based IP Geolocation

- Utilize 20 RIPE Atlas nodes to send 3 ICMP probes to a target prefix address
- Select Atlas nodes:
  - 3 nodes within each RIR (15 total vantage points)
  - 5 nodes within the registered country
- RIR_{Geo} is RIR responsible for RIR node returning minimum RTT

Limitations

- Prefix bias:
  - Randomly select 10k from each RIR
  - No ICMP-responsive target in prefix
  - No Atlas probes within the prefixes' registered country
- Geolocation
  - Atlas node location may be incorrect
  - Registration country may be a corporate headquarters elsewhere
  - Inconsistent prefixes

Limitations

- Prefix bias:
  - Randomly select 10k from each RIR
  - No ICMP-responsive target in prefix
  - No Atlas probes within the prefixes' registered country
- Geolocation
  - Atlas node location may be incorrect
  - Registration country may be a corporate headquarters elsewhere
  - Inconsistent prefixes

Initial work; select equal number of prefixes from each RIR

Meaningful coverage, with in-country nodes: 43k nodes in 87% of all countries

5 nodes in-country; 3 nodes on each continent. Refinement round.

Use registred country as a "second chance" to be consistent; work stands if we only look at RIR and geolocation

Current/Future work

Why latency-based geolocation?

- BGP and AS origin information can obscure true location
- IP Geolocation databases (e.g., MaxMind) known to contain inaccuracies, and use whois themselves
- Latency-based geolocation relies on physical signal propagation constraints
- Minimizing error:
  - Latency-based geolocation known accurate at continent and country granularity
  - Sound in proving geo-consistency (cannot manipulate speed-of-light constraint)
  - If any geo-inconsistency found, we select a new set of 20 nodes and repeat

Results

Case Study

RIPE: `X.X.X.X/16: Big CDN Corp UK`

Case Study

RIPE: `X.X.X.X/16: Big CDN Corp UK`

5x Atlas UK Nodes:
min(RTT) = 129ms

Case Study

RIPE: `X.X.X.X/16: Big CDN Corp UK`

5x Atlas UK Nodes:
min(RTT) = 129ms

RIPE Atlas Nodes:
min(RTT) = 149ms

Case Study

RIPE: `X.X.X.X/16: Big CDN Corp UK`

5x Atlas UK Nodes:
min(RTT) = 129ms

RIPE Atlas Nodes:
min(RTT) = 149ms

African Atlas Nodes:
min(RTT) = 258ms

Case Study

RIPE: `X.X.X.X/16: Big CDN Corp UK`

5x Atlas UK Nodes:
min(RTT) = 129ms

RIPE Atlas Nodes:
min(RTT) = 149ms

African Atlas Nodes:
min(RTT) = 258ms

ARIN Atlas Nodes:
min(RTT) = 71ms

Case Study

RIPE: `X.X.X.X/16: Big CDN Corp UK`

5x Atlas UK Nodes:
min(RTT) = 129ms

RIPE Atlas Nodes:
min(RTT) = 149ms

African Atlas Nodes:
min(RTT) = 258ms

ARIN Atlas Nodes:
min(RTT) = 71ms

Further refinement with Atlas nodes in ARIN region constrain to a Phoenix, AZ node with 7ms RTT. RIPE registry, RIPE organization, ARIN location => “registry geo-inconsistent”

Findings

Result	ARIN	RIPE	APNIC	LACNIC	AFRINIC
Fully Geo-consistent	94.7%	98.1	98.1%	97.0%	81.3%
Country Geo-consistent	1.2%	1.1%	0.5%	0.8%	7.6%
Country Geo-inconsistent	0.8%	0.2%	0.2%	0.0%	0.0%
Registry Geo-inconsistent	3.2%	0.4%	1.1%	2.1%	10.2%
Fully Geo-inconsistent	0.1%	0.2%	0.1%	0.0%	0.9%

- Overall, 96% of prefixes are fully consistent
- Primary contributor to fully inconsistent RIPE prefixes are prefixes geolocated to North America (CA, MX, US)
- AFRINIC has largest fraction of registry geo-inconsistencies (dominated by Europe and China)

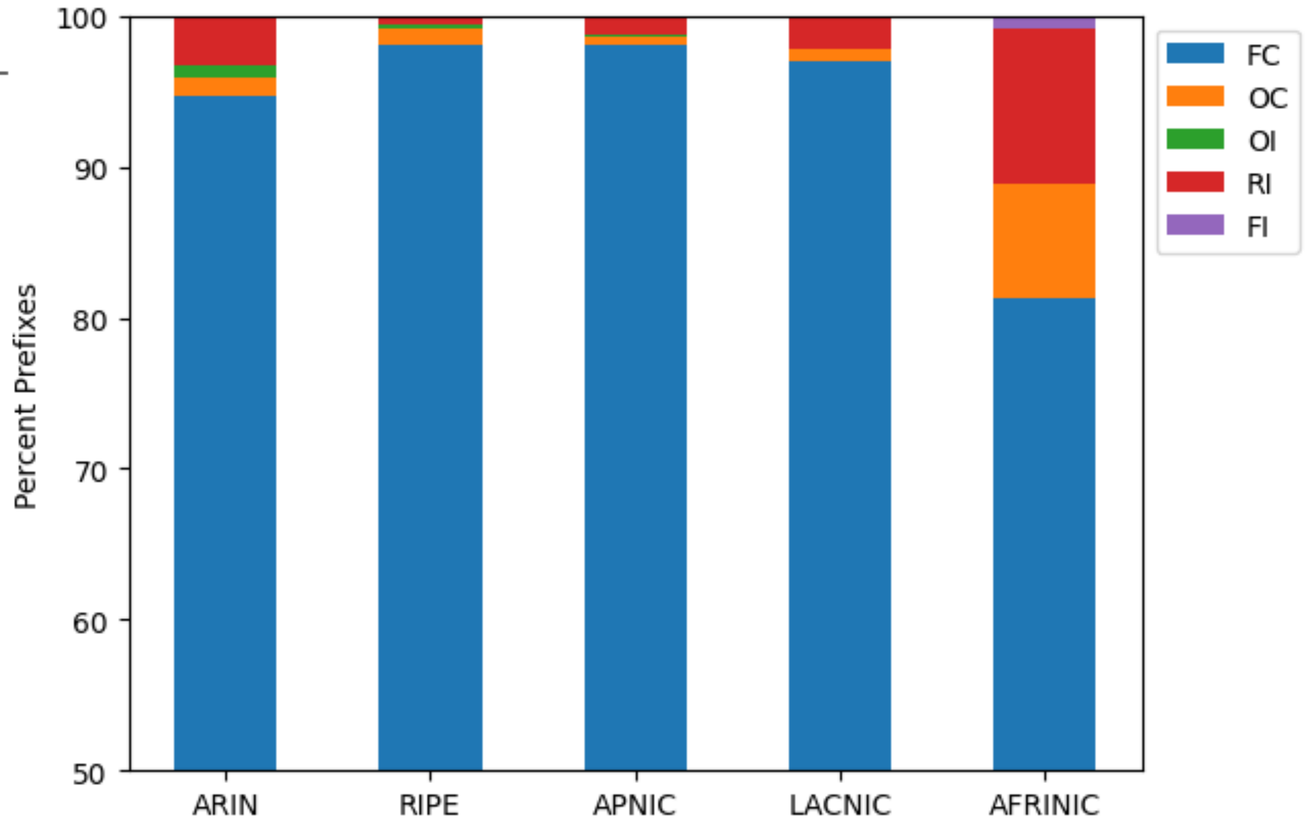

Current Work

- Intra-Prefix Inference Consistency
- IPv6 prefix registration audit
- Correlations with: registration age, prefix length, status attributes
- ASes responsible for most inconsistencies
- Validation with RIRs

Intra-Prefix Inference Consistency

Two IPv4 targets within large gaming provider's `X.Y.Z.0/20`

```
MeasID: XXXXXXXX
         US ARIN 82.095ms
        *NL RIPE 12.003ms
         NZ APNIC 298.107ms
         NL RIPE 15.385ms
         US ARIN 88.577ms
         ZA AFRINIC 174.829ms
         CA ARIN 112.797ms
         US ARIN 141.072ms
         ZA AFRINIC 188.942ms
         CA ARIN 92.787ms
         US ARIN 99.045ms
         MX LACNIC 145.058ms
         MX LACNIC 140.563ms
         CA ARIN 147.034ms
         NZ APNIC 298.748ms
         NZ APNIC 279.816ms
```

```
MeasID: YYYYYYYY
         DE RIPE 135.782ms
         BR LACNIC 148.534ms
         US ARIN 40.190ms
         BR LACNIC 172.870ms
         BR LACNIC 166.621ms
         DE RIPE 129.862ms
         MA AFRINIC 146.560ms
         CA ARIN 43.771ms
         MA AFRINIC 157.716ms
         CA ARIN 62.010ms
         CA ARIN 59.697ms
         US ARIN 49.252ms
         US ARIN 86.423ms
        *US ARIN 26.833ms
         NZ APNIC 158.495ms
         NZ APNIC 162.620ms
```

Intra-Prefix Inference Consistency

Two IPv4 targets within large gaming provider's `X.Y.Z.0/20`

```
MeasID: XXXXXXXX
        US ARIN 82.095ms
       *NL RIPE 12.003ms
        NZ APNIC 298.107ms
        NL RIPE 15.385ms
        US ARIN 88.577ms
        ZA AFRINIC 174.829ms
        CA ARIN 112.797ms
        US ARIN 141.072ms
        ZA AFRINIC 188.942ms
        CA ARIN 92.787ms
        US ARIN 99.045ms
        MX LACNIC 145.058ms
        MX LACNIC 140.563ms
        CA ARIN 147.034ms
        NZ APNIC 298.748ms
        NZ APNIC 279.816ms
```

```
MeasID: YYYYYYYY
        DE RIPE 135.782ms
        BR LACNIC 148.534ms
        US ARIN 40.190ms
        BR LACNIC 172.870ms
        BR LACNIC 166.621ms
        DE RIPE 129.862ms
        MA AFRINIC 146.560ms
        CA ARIN 43.771ms
        MA AFRINIC 157.716ms
        CA ARIN 62.010ms
        CA ARIN 59.697ms
        US ARIN 49.252ms
        US ARIN 86.423ms
       *US ARIN 26.833ms
        NZ APNIC 158.495ms
        NZ APNIC 162.620ms
```

Intra-Prefix Inference Consistency

Two IPv4 targets within large gaming provider's `X.Y.Z.0/20`

BGP: `X.Y.N/23`

```
MeasID: XXXXXXXX
        US ARIN 82.095ms
       *NL RIPE 12.003ms
        NZ APNIC 298.107ms
        NL RIPE 15.385ms
        US ARIN 88.577ms
        ZA AFRINIC 174.829ms
        CA ARIN 112.797ms
        US ARIN 141.072ms
        ZA AFRINIC 188.942ms
        CA ARIN 92.787ms
        US ARIN 99.045ms
        MX LACNIC 145.058ms
        MX LACNIC 140.563ms
        CA ARIN 147.034ms
        NZ APNIC 298.748ms
        NZ APNIC 279.816ms
```

Inconsistency explained by subnets advertised in BGP

BGP: `X.Y.M/23`

```
MeasID: YYYYYYYY
        DE RIPE 135.782ms
        BR LACNIC 148.534ms
        US ARIN 40.190ms
        BR LACNIC 172.870ms
        BR LACNIC 166.621ms
        DE RIPE 129.862ms
        MA AFRINIC 146.560ms
        CA ARIN 43.771ms
        MA AFRINIC 157.716ms
        CA ARIN 62.010ms
        CA ARIN 59.697ms
        US ARIN 49.252ms
        US ARIN 86.423ms
       *US ARIN 26.833ms
        NZ APNIC 158.495ms
        NZ APNIC 162.620ms
```

Take-aways

- Different RIRs have different out-of-region address use policies
  - But limited visibility of where resources used, especially post-allocation
- RIR allocations are largely geo-consistent, with some notable exceptions
- Geo-inconsistencies raise operational and security concerns that suggest registration information should be updated
- RIR whois records use inconsistent schemas, complicating data analysis (RDAP will hopefully fix this!)

Thanks!

- First quantitative geo-audit of RIR IP registry information
  - Technical draft paper: https://arxiv.org/abs/2308.12436
  - All RIPE Atlas data open and public for transparency
- Future work: expand measurements, extend to IPv6, and engage with RIRs
- We welcome feedback/flames!

Rob Beverly <rbeverly@cmand.org>
Oliver Gasser <oliver.gasser@mpi-inf.mpg.de>