The BGP Visibility Scanner

Andra Lutu[1,2] , ***Marcelo Bagnulo***[2] *and* ***Olaf Maennel***[3]

Institute IMDEA Networks[1]*, University Carlos III Madrid*[2] *, Loughborough University*[3]

Problem Statement

- The routing preferences are designed to accommodate various operational, economic, and political factors

- ***Problem:***
  - Only by configuring a routing policy, the origin AS cannot also ensure that it will *achieve the anticipated results*
  - The implementation of routing policies is a complicated process, involving subtle tuning operations that are *error-prone*
- Operators need to complement their internal perspective on routing with the information retrieved from external sources

Internet Prefix Visibility

- ***Prefix visibility*** as an expression of policy interaction
  - Not all the routes make it to every routing table (RT) in the interdomain
- ***Limited Visibility Prefixes*** – prefixes that are not in every RT
- ***High Visibility Prefixes*** – prefixes which are in almost all the RT

- ***The BGP Visibility Scanner:***
  - Analyze all BGP routing data from ***RouteViews*** and ***RIPE Routing Information Service (RIS)*** projects
    - All together there are 24 different RT collection points
    - More than 130 different ASes periodically dump their *entire* routing tables

- **Limited Visibility Prefixes** (**LVPs**)
  - *Intentional/Deliberate*
  - *Inflicted by third parties*
  - *Unintentional/Accidental*

Next…

- Data manipulation – methodology
- Study case: example of applying the methodology
- Characteristics of the prefixes with limited visibility
- Presenting the tool and its capabilities
- Use cases

The BGP Visibility Scanner

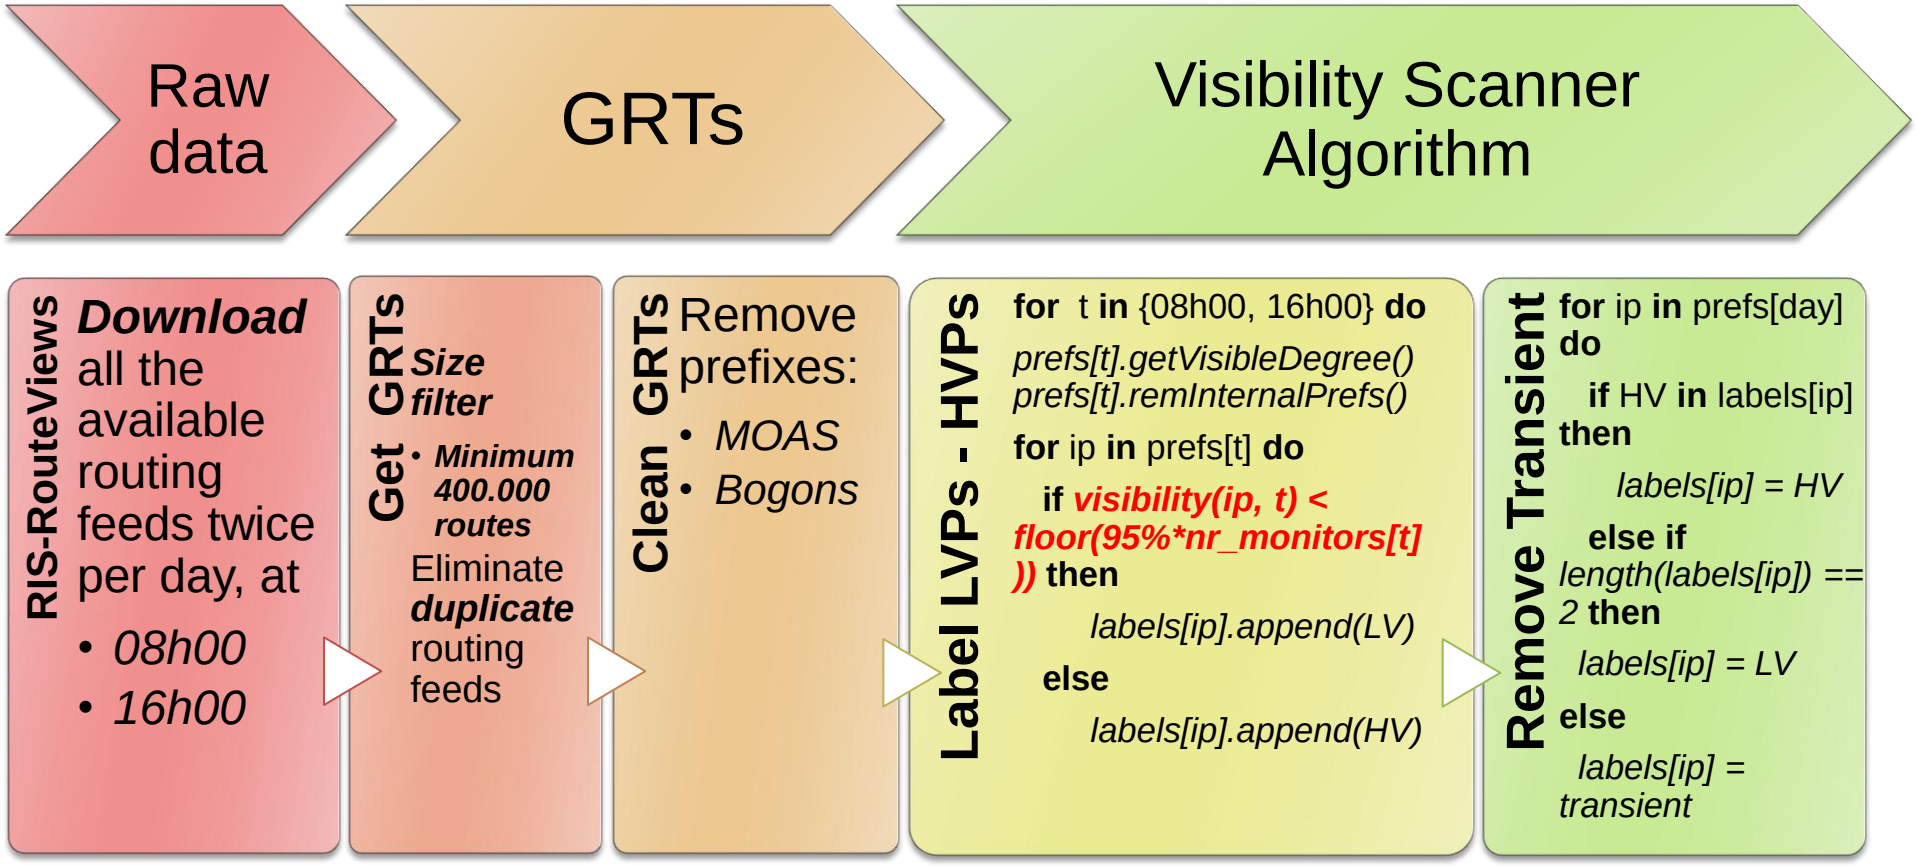

The BGP Visibility Scanner

Raw data

RIS-RouteViews

Download all the available routing feeds twice per day, at

- *08h00*
- *16h00*

The BGP Visibility Scanner

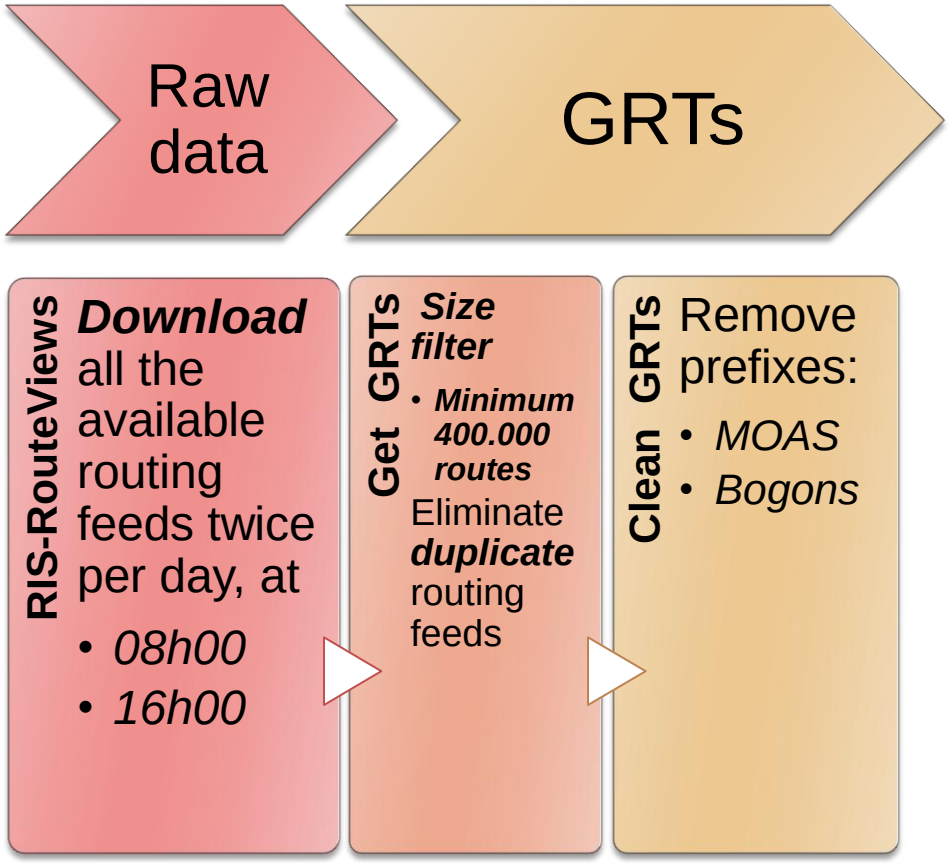

- ***Example:*** sampling time 23.10.2012
- ***Global Routing Table*** - contains almost all the prefixes injected in the interdomain
  - 129 GRTs from RIPE RIS and RouteViews
    - 9/129 ASes in LACNIC
    - 14/129 in APNIC
    - 37/129 in ARIN
    - 68/129 in RIPE NCC
- Polishing the full routing tables for our study
  - No bogons/martians present
    - Discard 500 bogon prefixes
  - No MOAS prefixes
    - Filter out approx. 4,500 MOAS prefixes

BGP visibility scanner

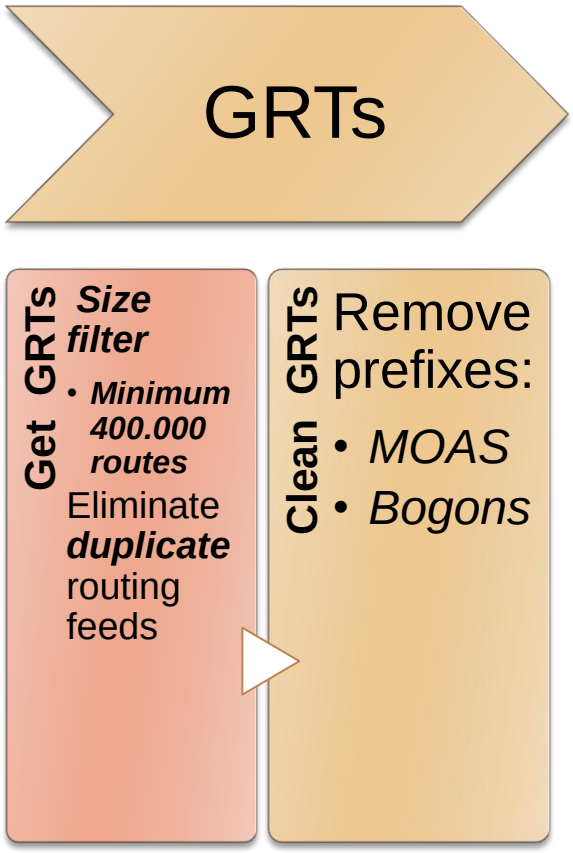

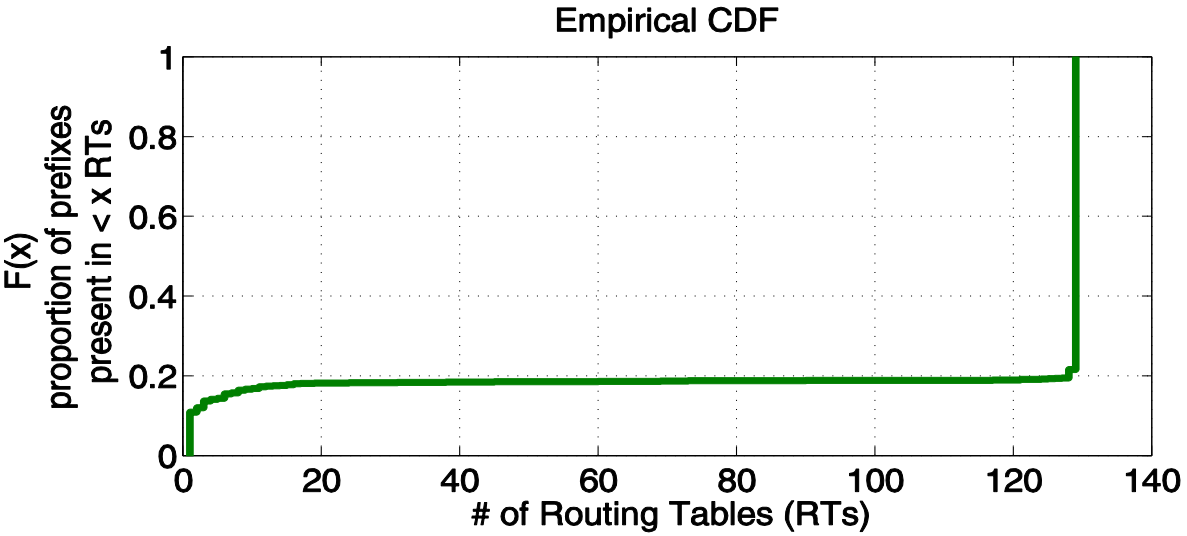

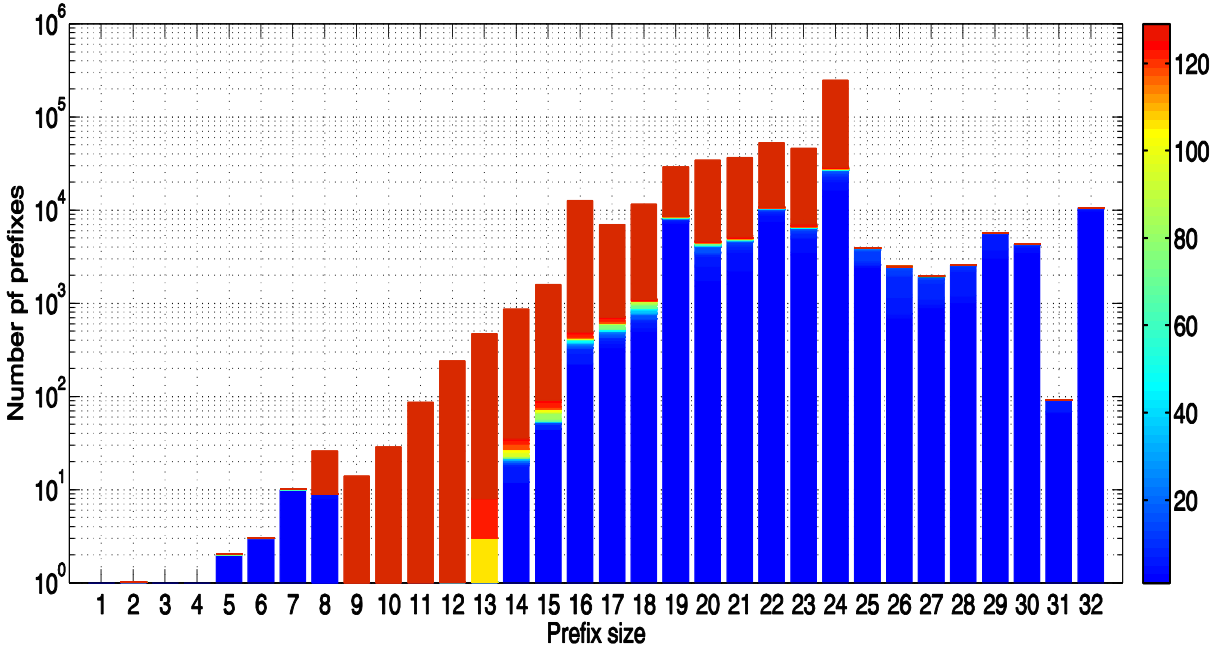

BGP Visibility Scanner

- We find that not all the GRTs identified contain **all** the prefixes injected in the interdomain
- Expression of policies which may have backfired
- Sample from 23.10.2012 – 08h00

The BGP Visibility Scanner

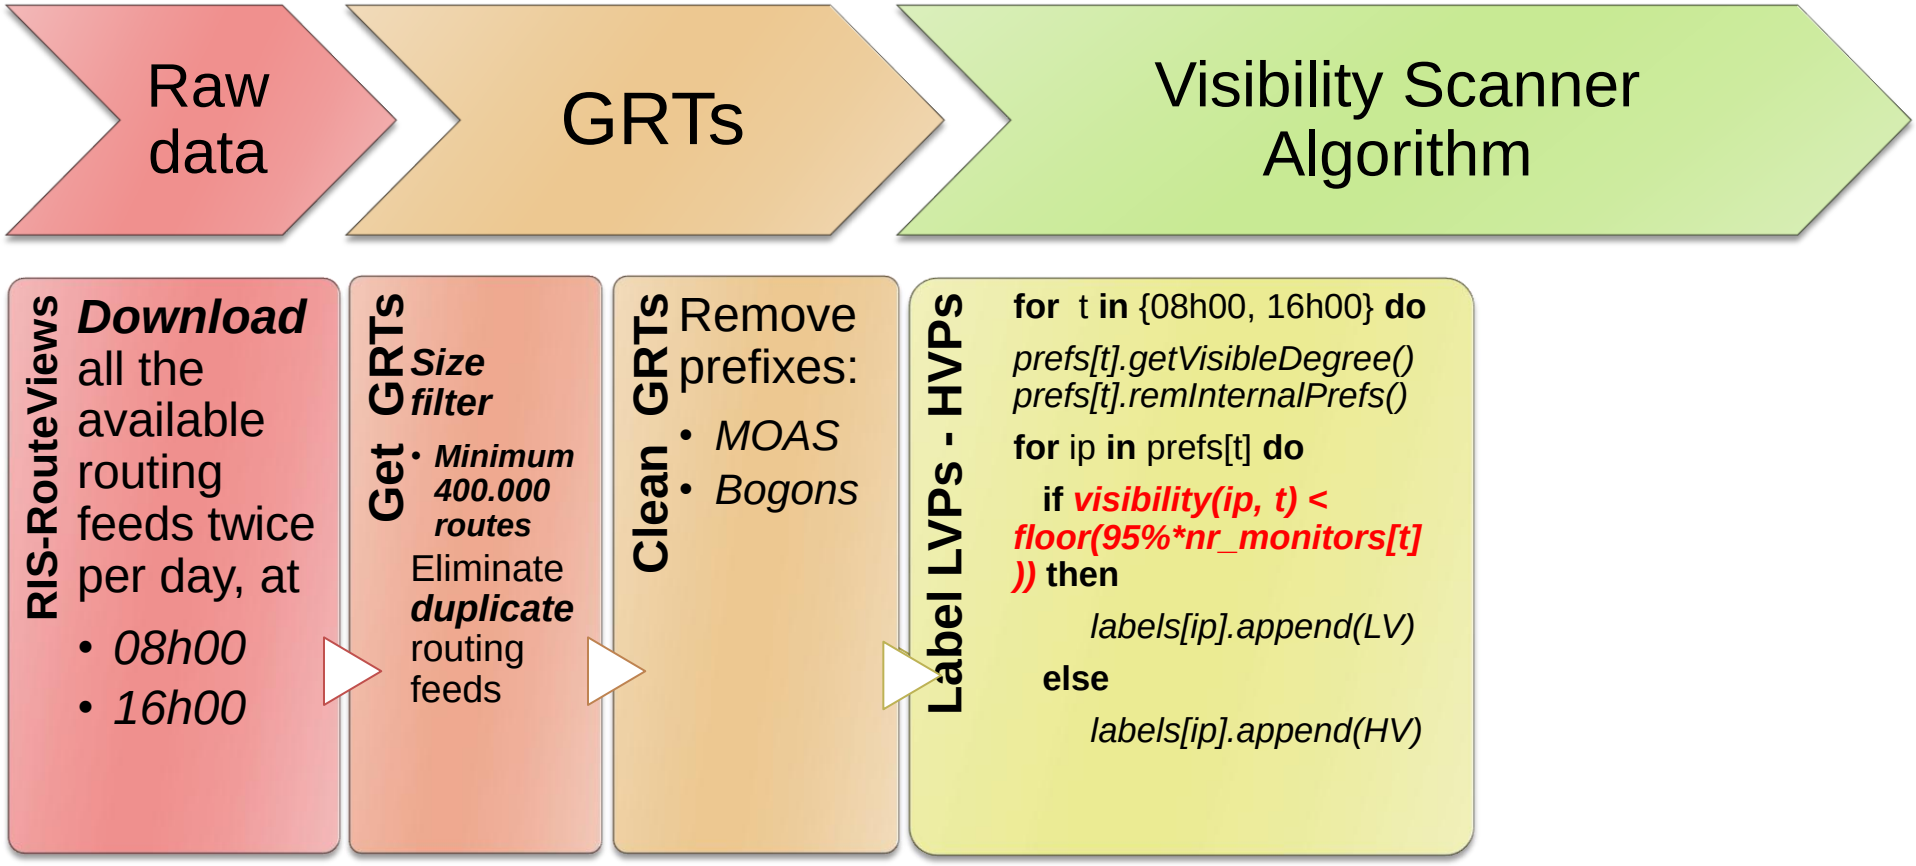

- ***Filter internal routes***
  - Not considering prefixes only present in *1 RT with an AS-Path of length 1*
  - ***23.10.2012:*** filter out 10.500 internal routes

- ***Labeling Mechanism*** – each prefix gets a visibility label based on the *95% minimum visibility threshold rule*
  - **HV –** high visibility if present in more than 95% of routing tables
  - **LV –** limited visibility if present in less than 95% of routing tables

BGP visibility scanner

Visibility Scanner Algorithm

Label LVPs - HVPs

```
for  t in {08h00, 16h00} do
prefs[t].getVisibleDegree()
prefs[t].remInternalPrefs()
for ip in prefs[t] do
   if visibility(ip, t) <
floor(95%*nr_monitors[t]
)) then
        labels[ip].append(LV)
   else
        labels[ip].append(HV)
```

The BGP Visibility Scanner

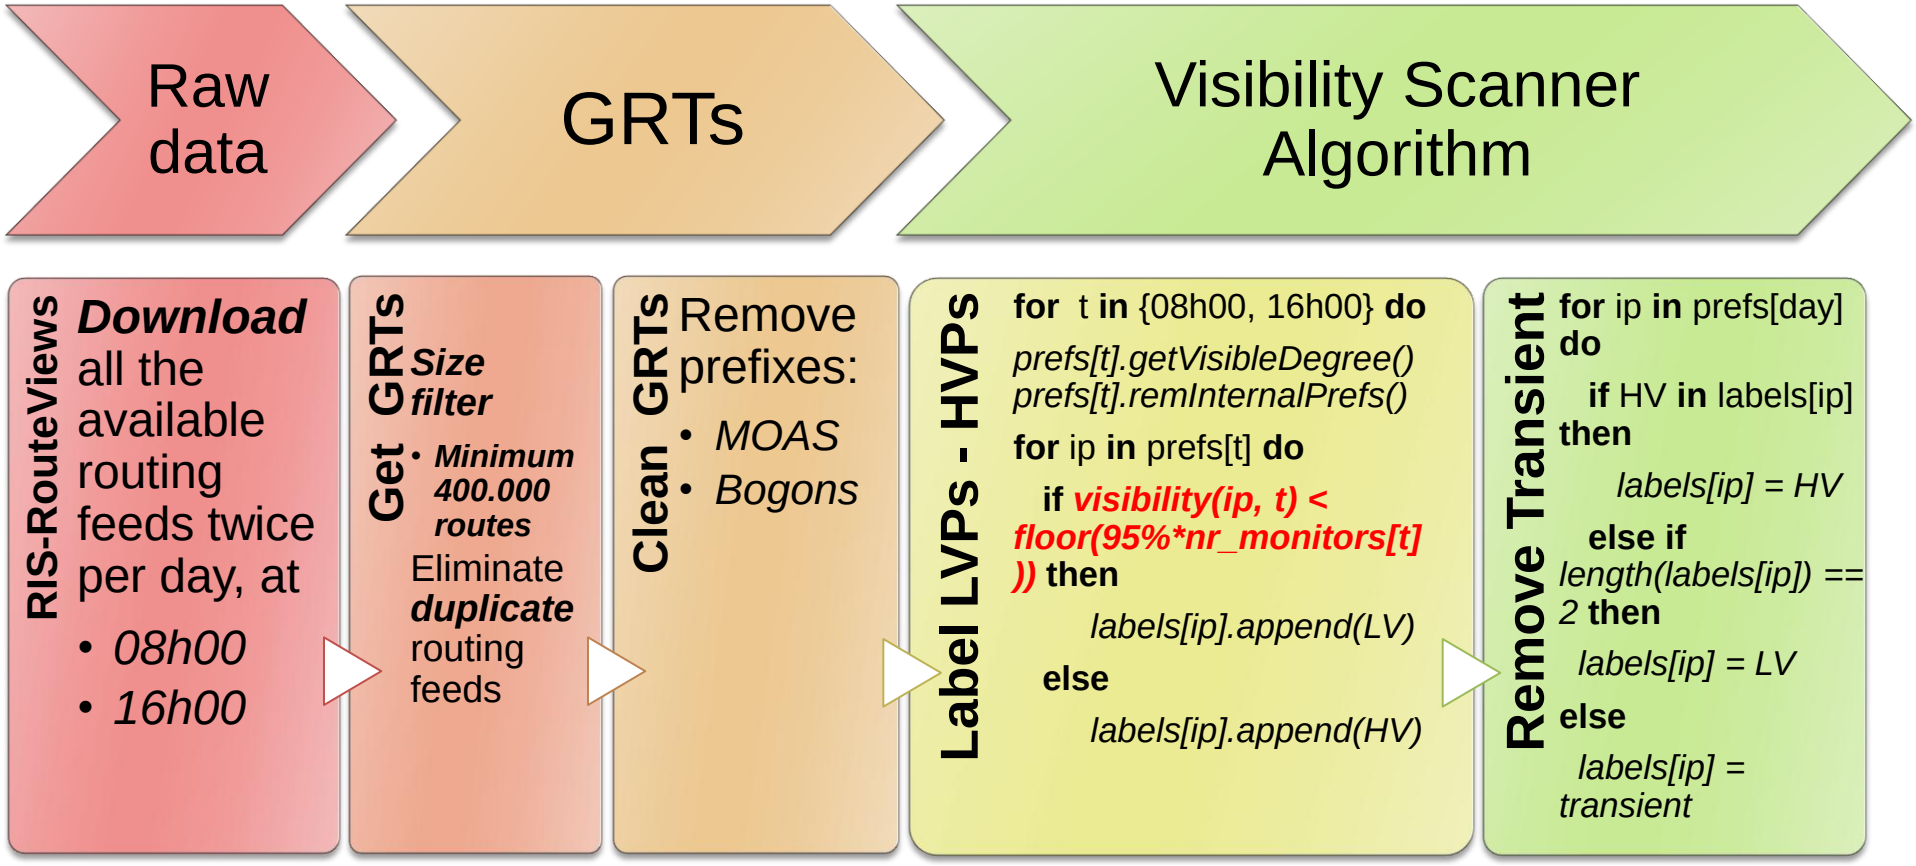

- ***Label Prevalence Sieve*** – rule of prevalence for the visibility labels tagged on each prefix
- Filter transient routes
  - Filter the prefixes that are not consistently appearing in the two samples analyzed
  - Discard 7,800 prefixes
- A total of 512.000 prefixes identified
  - 415.576 High-Visibility prefixes (**HVPs**)
  - 98.253 Limited-Visibility prefixes (**LVPs**)

BGP visibility scanner

Visibility Scanner Algorithm

Remove Transient

```
for ip in prefs[day] do
  if HV in labels[ip] then
    labels[ip] = HV
  else if length(labels[ip]) == 2 then
    labels[ip] = LV
  else
    labels[ip] = transient
```

Dark Prefixes

- Dark Prefixes (DP) are the ***LV prefixes that are not covered by any HV prefix***

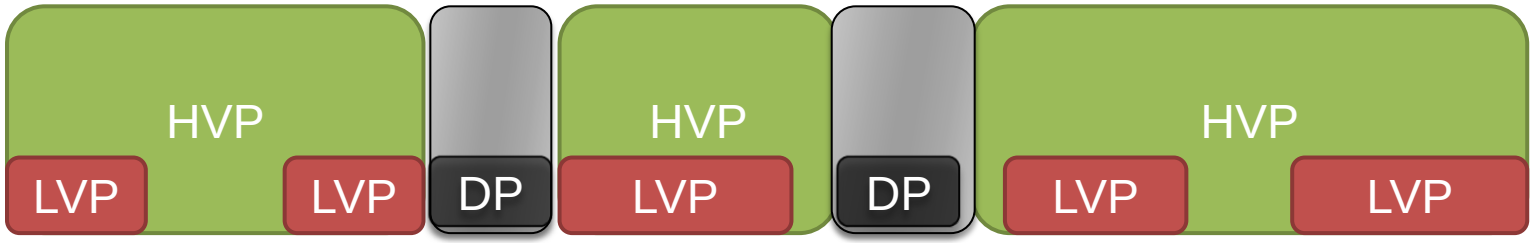

- This would constitute address space that may not be globally reachable (in the absence of a default route)
  - In 2012.10.23 there were ~2.400 dark prefixes in the LV prefix set

Prefix visibility – distribution on prefix length

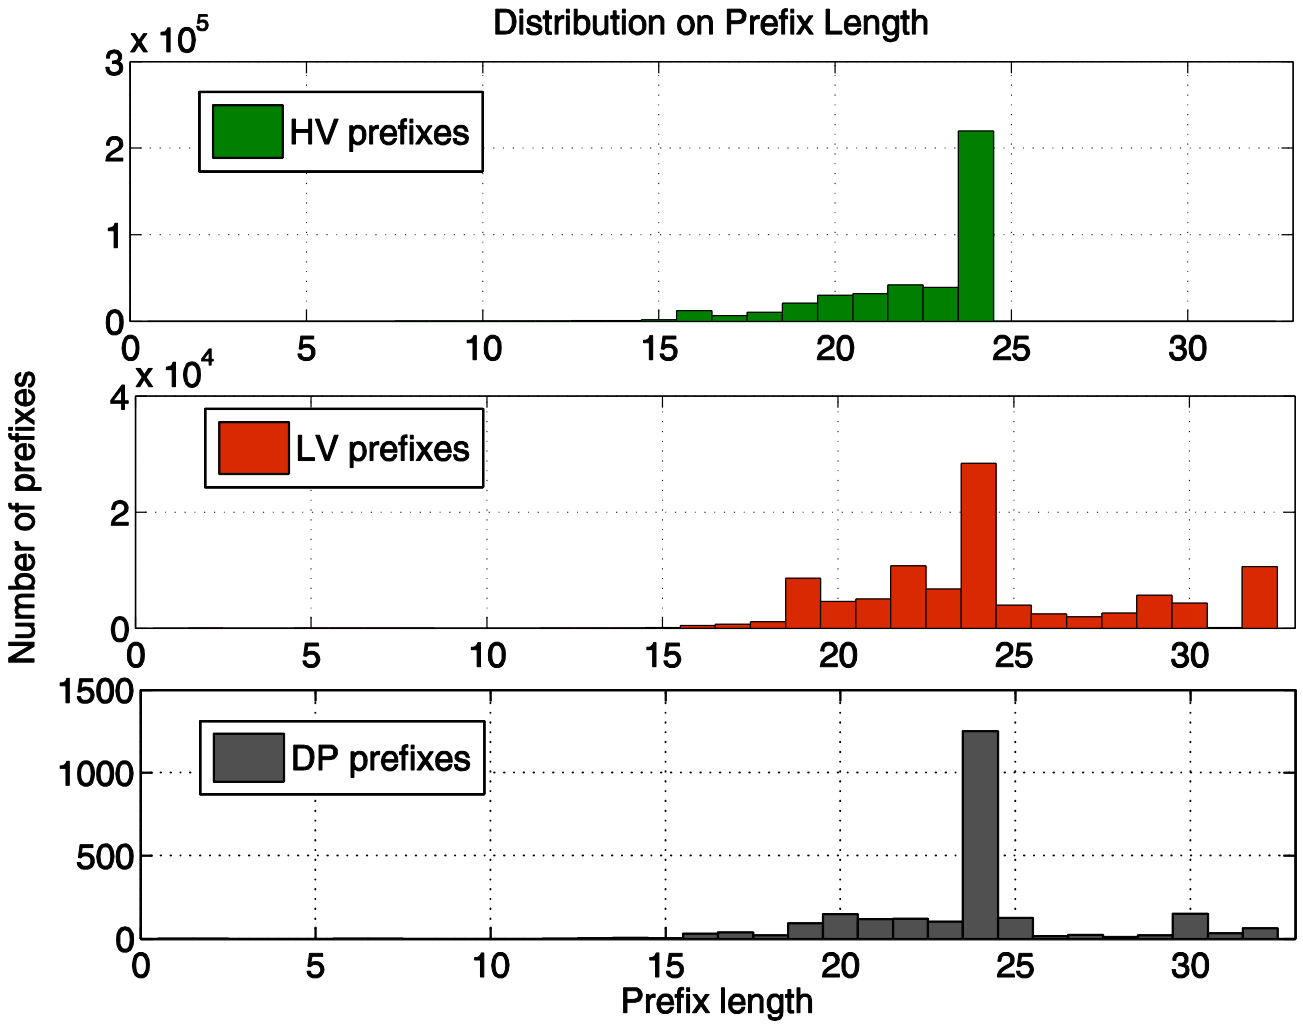

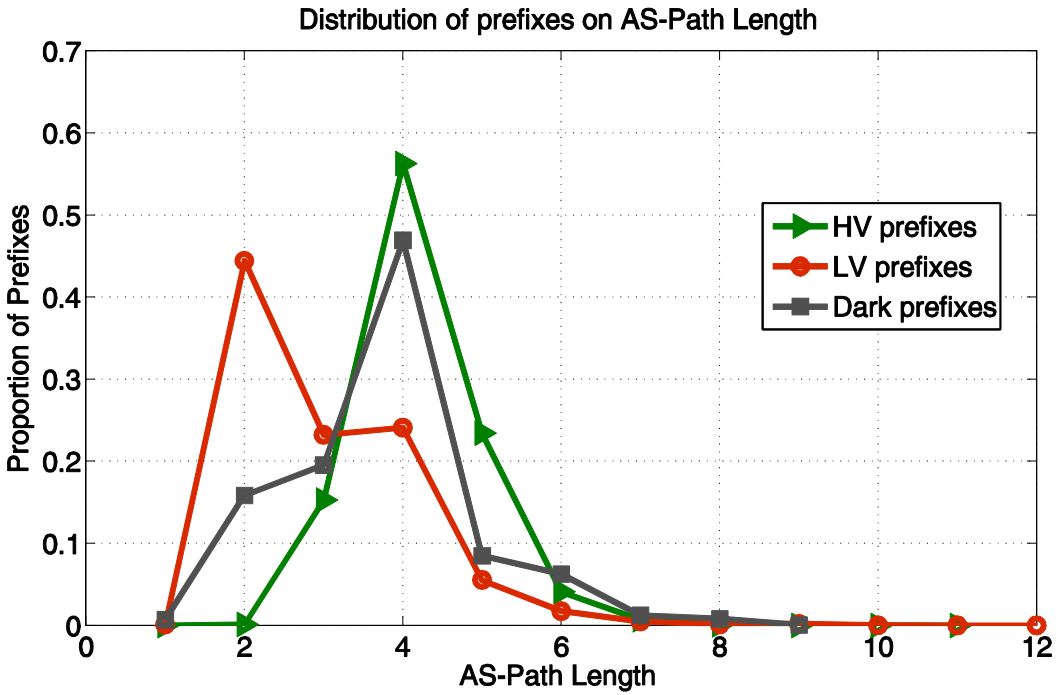

AS-Path length

- The per set ***mean*** AS-Path length (no prepending considered):
  - LV prefixes – 3.02
    - **Mode = 2**
  - HV prefixes – 4.16
    - **Mode = 4**
  - Dark prefixes – 3.75
    - **Mode = 4**

Prefix visibility as of 23.10.2012

- Visibility distribution: # of LV prefixes present in *n* monitors, where n = 1, … 129
  - Low sensitivity to the visibility threshold included in the Labeling Mechanism

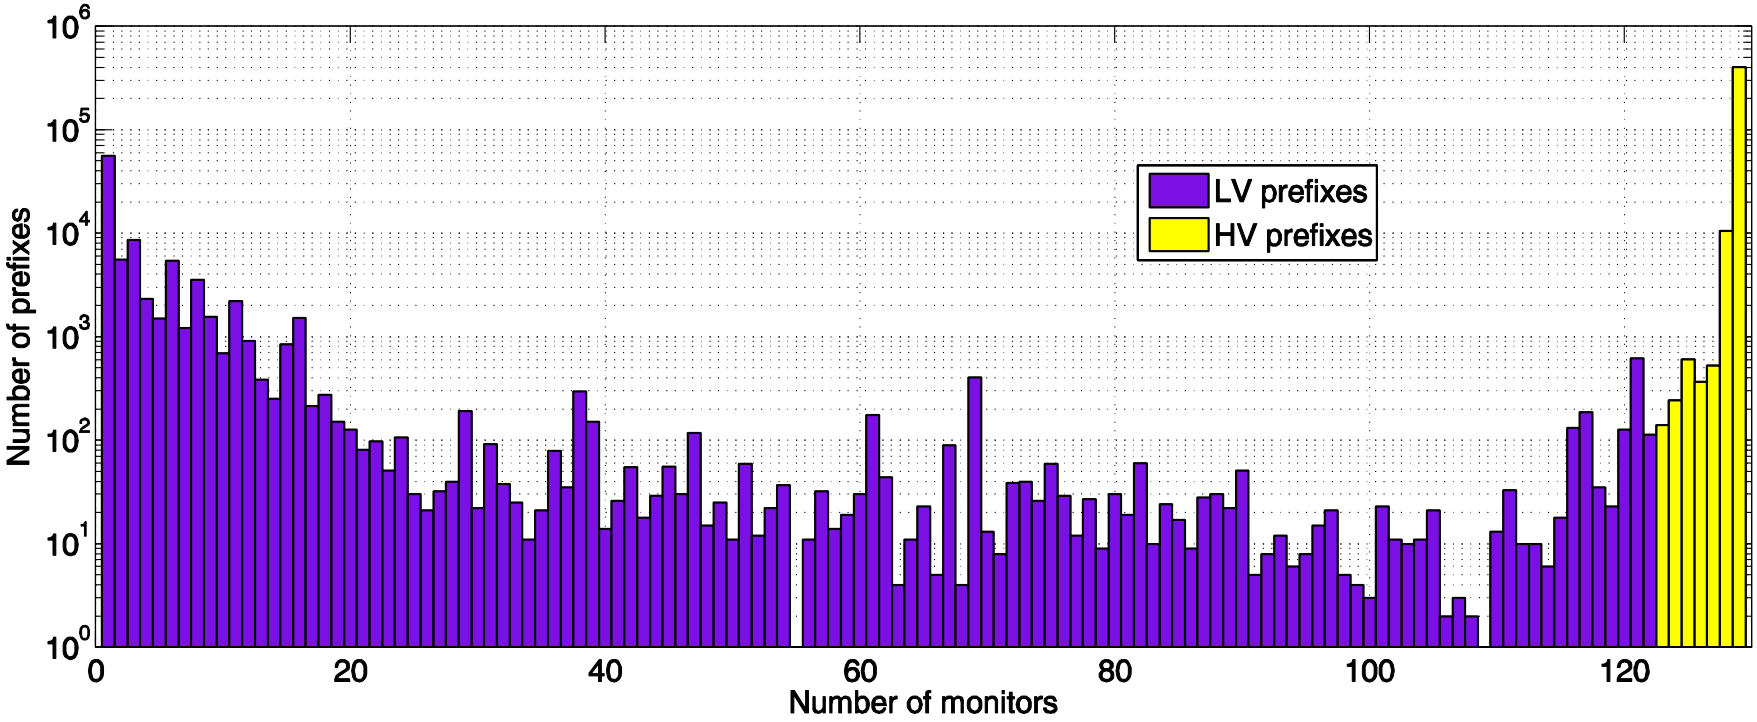

Prefix Label Stability in 2012.10

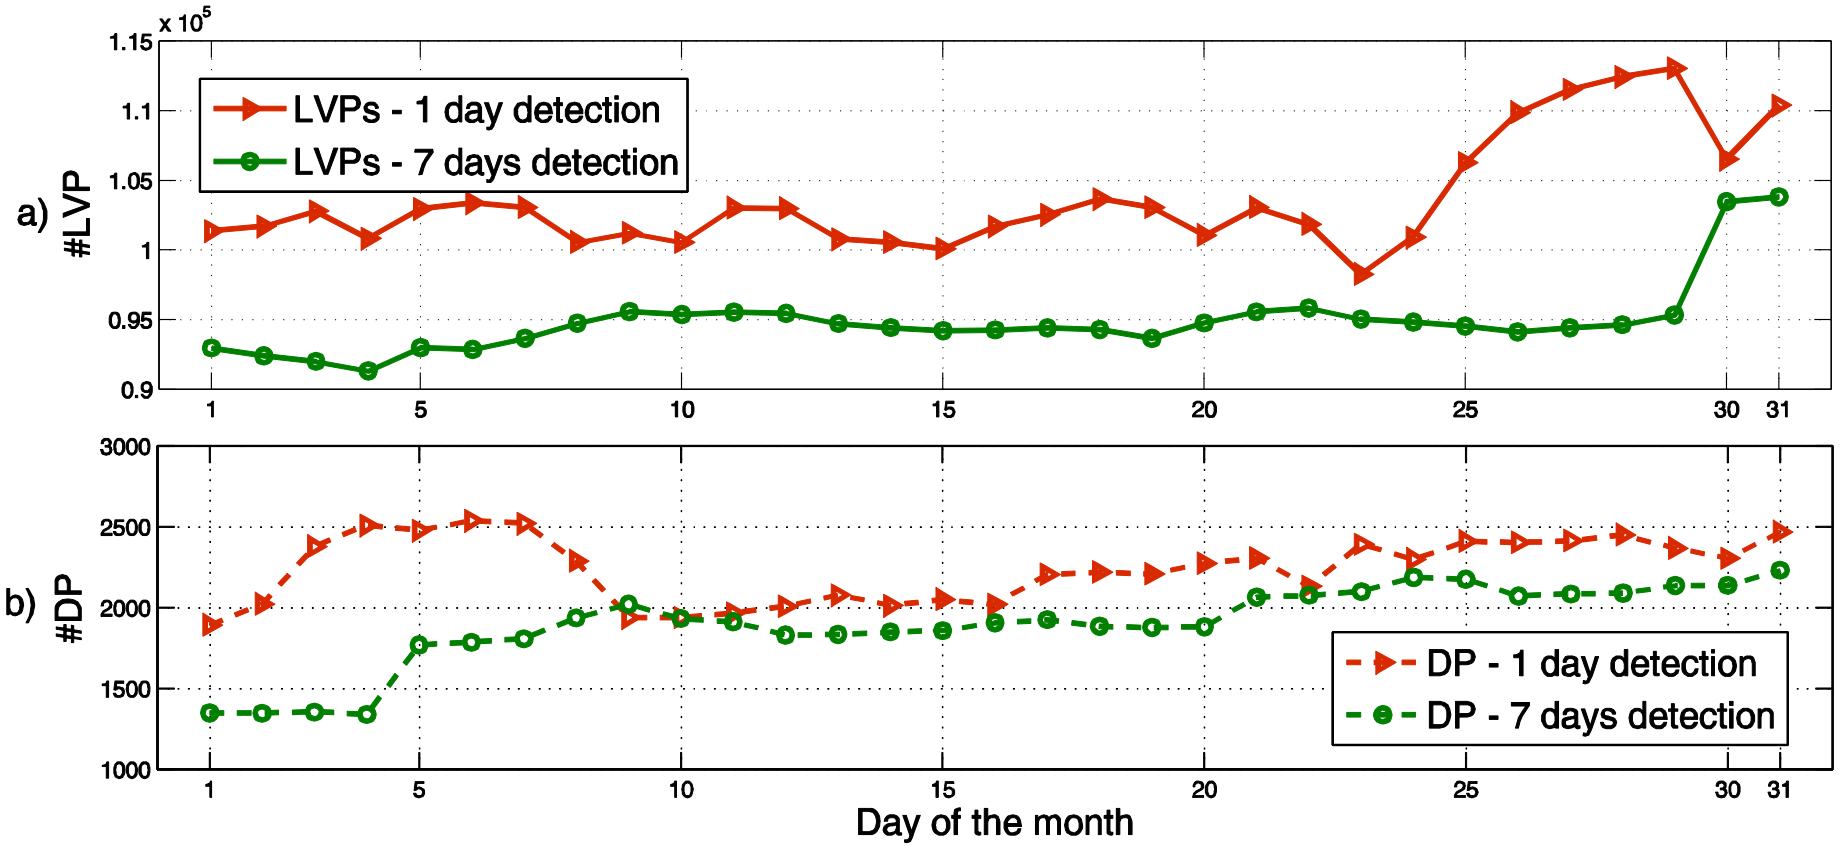

Origin ASes for the LV prefixes

- Identified 3.570 different ASes originating the LV prefixes identified on 2012.10.23:
  - 14% in LACNIC (~493 ASes)
  - 30.5% in APNIC (~1.081 ASes)
  - 30.1% in RIPE (~1.068 ASes)
  - 22.4% in APNIC (~795 ASes)
  - 1.1% in AFRINIC (~42 ASes)

What are these prefixes?

- We are looking to explain this phenomena:
  - Is it something the origin AS intended or is it something that the AS is suffering?
- All the results of this study are made available online

visibility.it.uc3m.es

  - *Up to date information on LV announced by each AS*
  - *Check to see if your AS is originating LV prefixes*
  - *Retrieve those prefixes and see if there are any Dark Prefixes within that set*
- Please provide feedback!
  - Short form that you can fill in and send

How does it work?

visibility.it.uc3m.es

Limited Visibility Prefixes

Retrieve the limited visibility prefixes per origin AS

By inputing an AS number, you can retrieve the limited visibility IPv4 prefixes injected by that particular network, according to the data we have observed during our study. You can also check if the LV prefix retrieved is a dark prefix (marked with DP) or simply limited visibility (marked with LV).

Please also take the time to fill in a short form after visualizing the results of your query.

Query for ASN: Get prefixes

How does it work?

visibility.it.uc3m.es

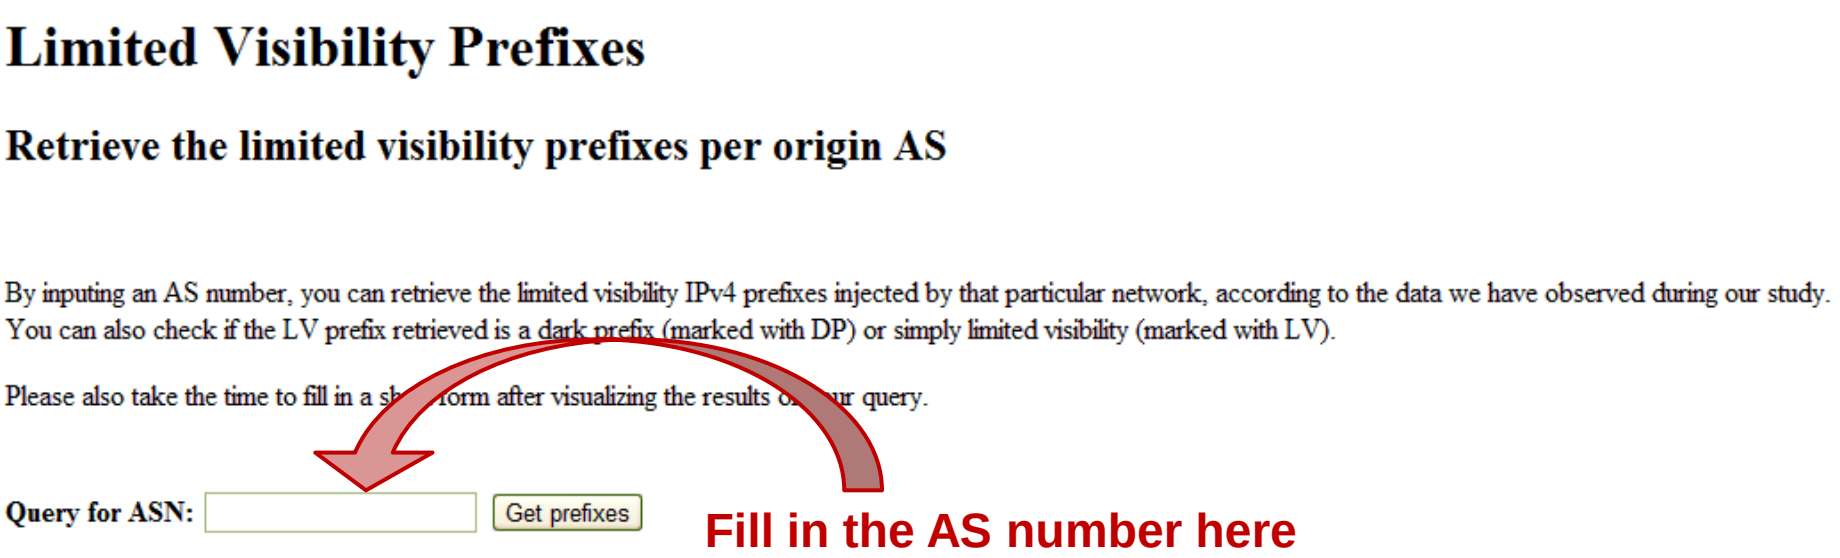

How does it work?

visibility.it.uc3m.es

- Example of output:

Please take the time to fill in the form concerning the prefixes listed below!

Fill Form Back

Prefix	Origin AS	Dark Prefix (DP) / Limited Visibility (LV)	Prefix visibility (#RTs out of the sample)
140.212.21.0/24	7018	LV	61/75

How does it work?

visibility.it.uc3m.es

- Example of output:

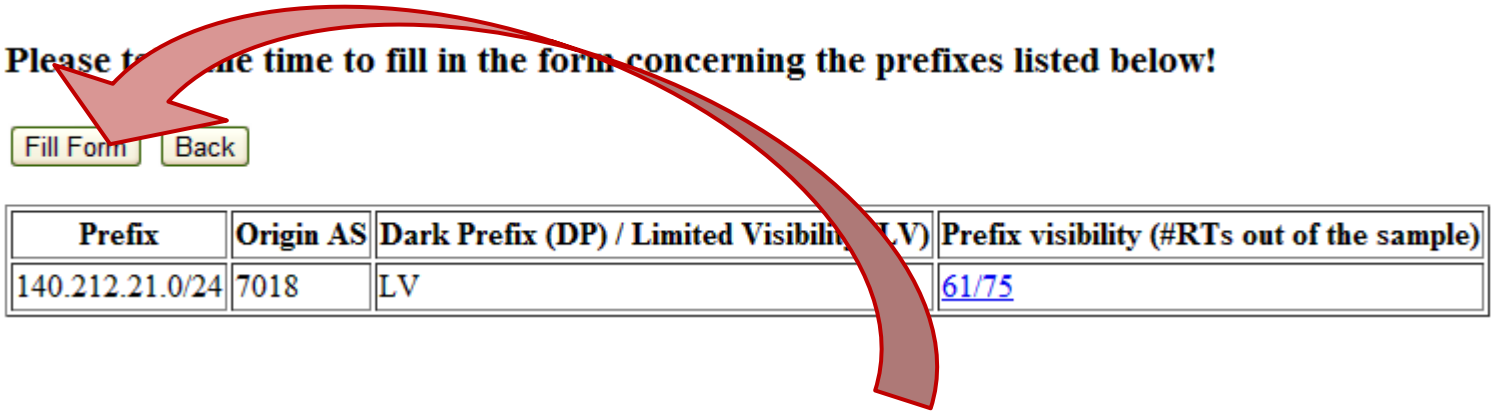

Please t[illegible]e time to fill in the form concerning the prefixes listed below!

Fill Form | Back

Prefix	Origin AS	Dark Prefix (DP) / Limited Visibili[illegible]LV)	Prefix visibility (#RTs out of the sample)
140.212.21.0/24	7018	LV	61/75

Next step: fill in form!

How does it work?

Limited Visibility Prefixes Survey

Webpage disclaimer:

All the infomation provided in this survey is used to generate anonymized aggregated reports regarding the Limited Visibility prefixes observed during our study. All the questions are optional. We appreciate any level of infomation that you wish to provide.

1. **Are you aware that the prefixes retrieved in the previous table have limited visibility?** ○ **Yes** ○ **No**
2. **Were any of these prefixes intended to have full global visibility?** ○ **Yes** ○ **No**
3. **Are some of these prefixes accidentaly leaked outside the network?** ○ **Yes** ○ **No**
4. **Could you point out some of the reasons for which these prefixes are not visible everywhere?**
   - ☐ Scoped advertisements
   - ☐ Use of Communities
   - ☐ Advertised only to peers
   - ☐ Partial transit
   - ☐ Leaked prefixes
   - ☐ Filtered by other AS
   - ☐ Other:

Type your answer here.

Submit!!

Use Cases

- Different use case:
  - Intended Scoped Advertisements
    - Inject prefixes only to peers
  - Intended Scoped Advertisements: *Content provider*
    - Geographical scoping of prefix
  - Config errors: *Large ISP*
    - Outbound filters mistakes in configuration
    - Leaking routes to direct peers
  - Third-party inflicted: *Internet root servers*
    - Tackle problems rising from the interaction between Ases
      - Blackholing due to lack of return path
      - Blackholing due to no announcement

Use Cases – Internet Root Servers

- Observe two prefixes: p/24 -LVP and p/23 – HVP
  - Blackholing due to lack return path:

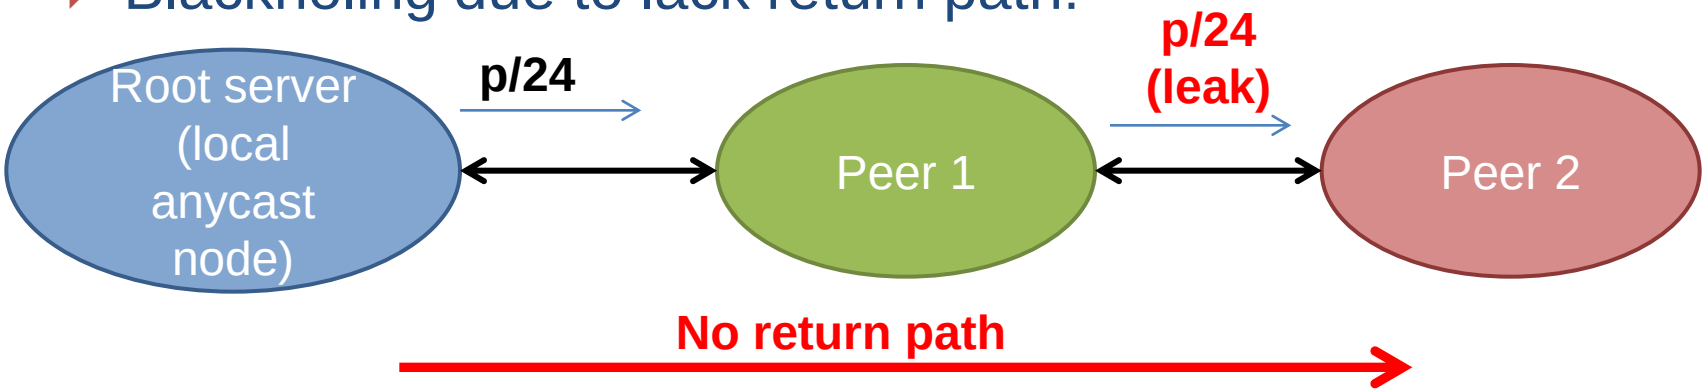

Use Cases – Internet Root Rervers

- Observe two prefixes: p/24 -LVP and p/23 – HVP
  - Blackholing due to lack return path:

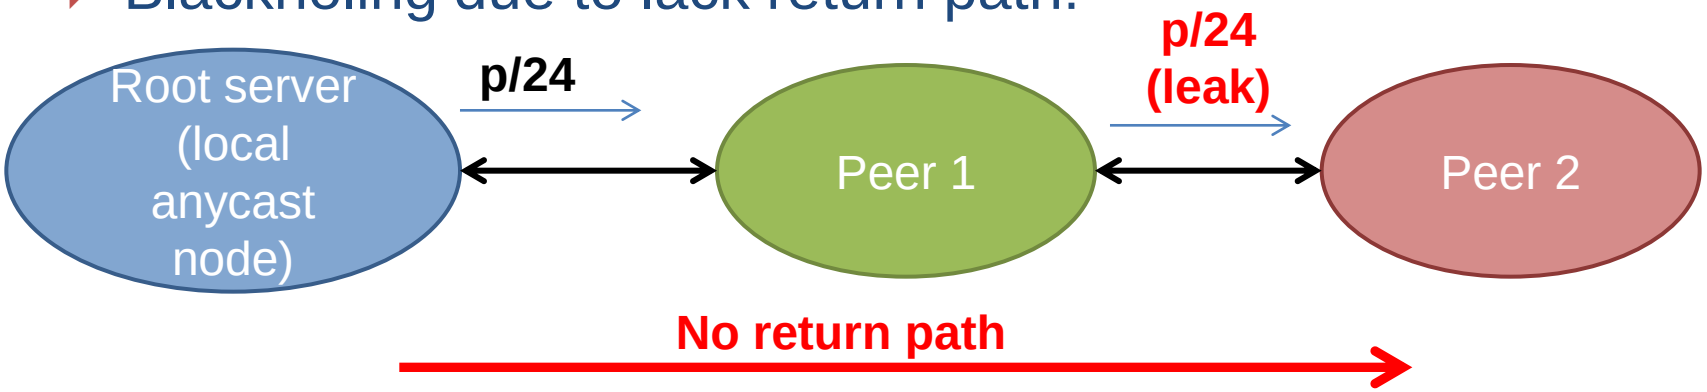

  - No full transit at the IXP => tag with NO EXPORT

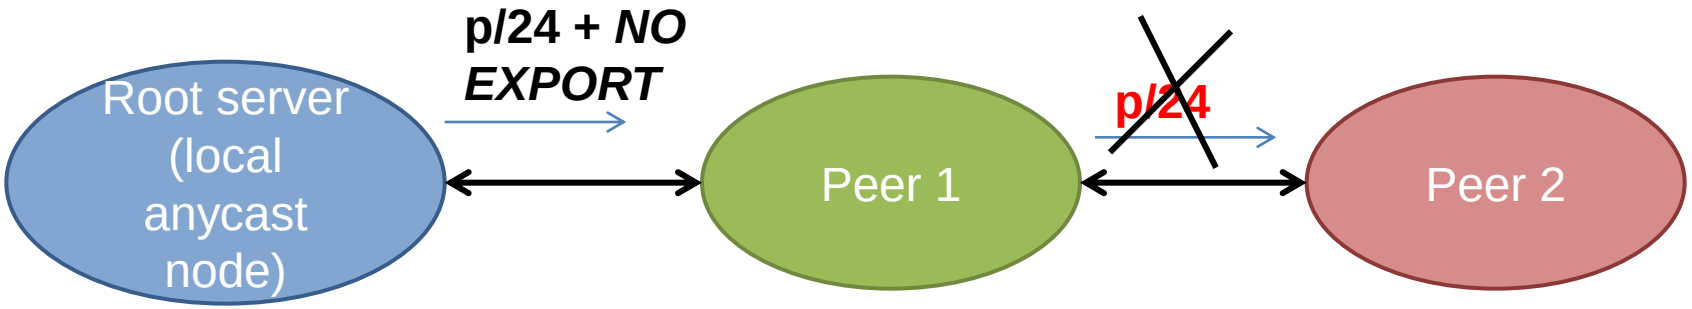

Use Cases – Internet Root Servers

- Observe two prefixes: p/24 -LVP and p/23 – HVP
  - Blackholing due to lack return path:

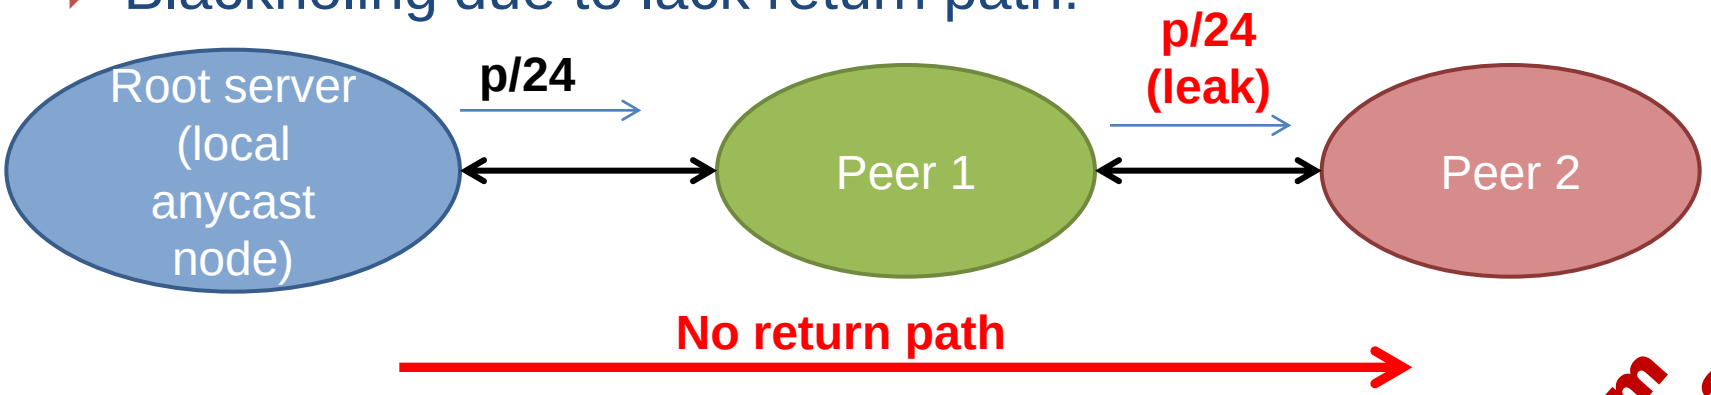

  - No full transit at the IXP => tag with NO EXPORT

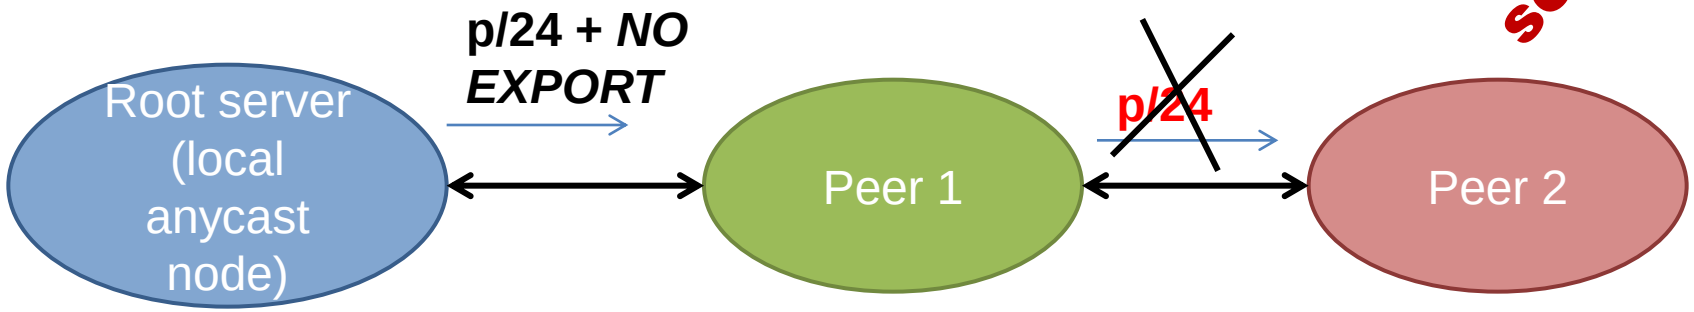

Use Cases – Internet Root Server

- Blackholing due to no announcement

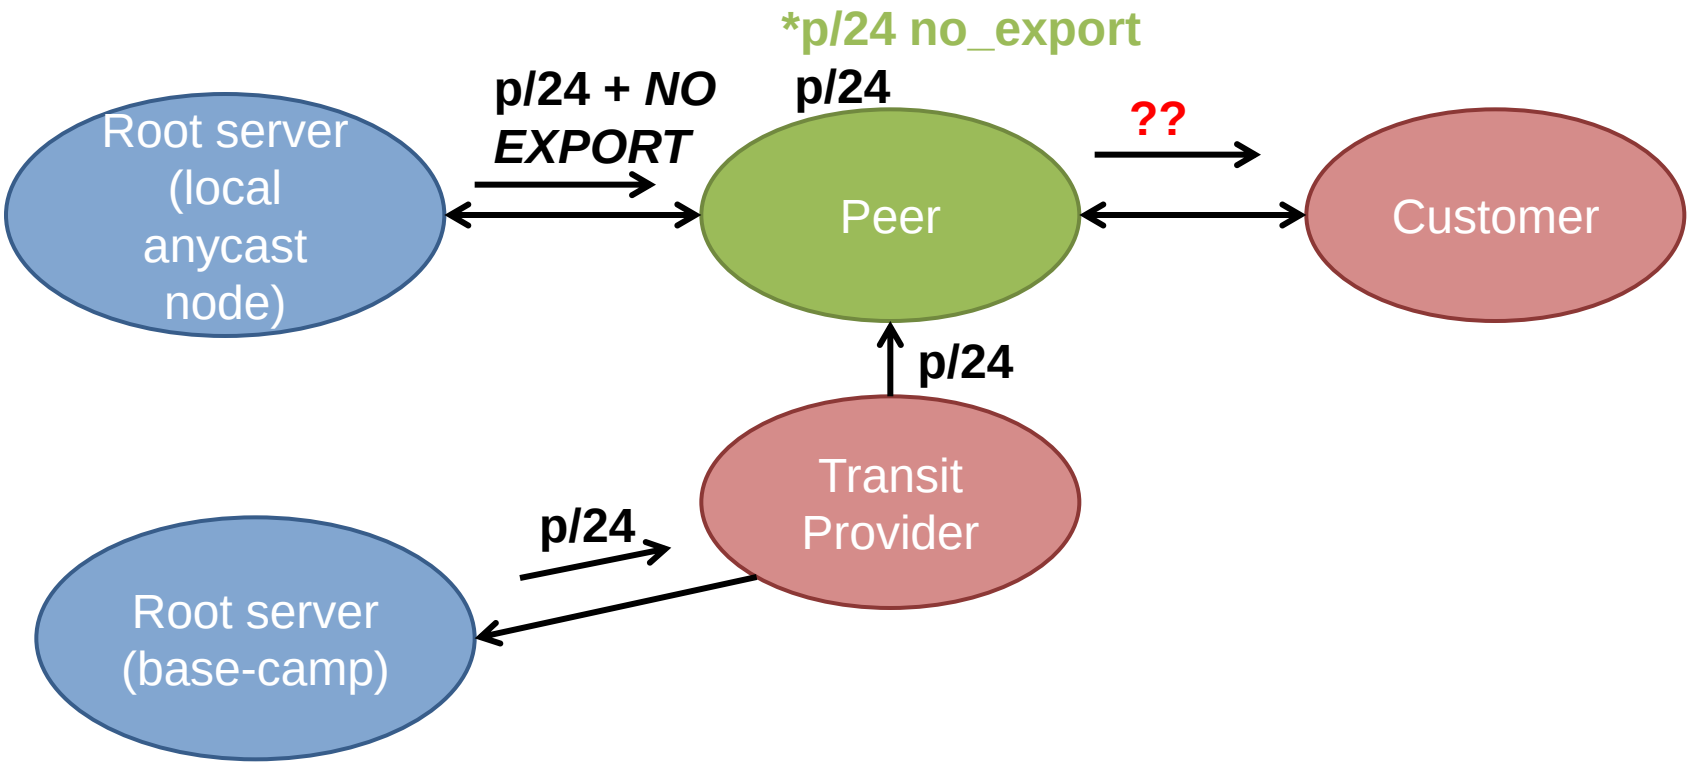

Use Cases – Internet Root Server

- Blackholing due to no announcement

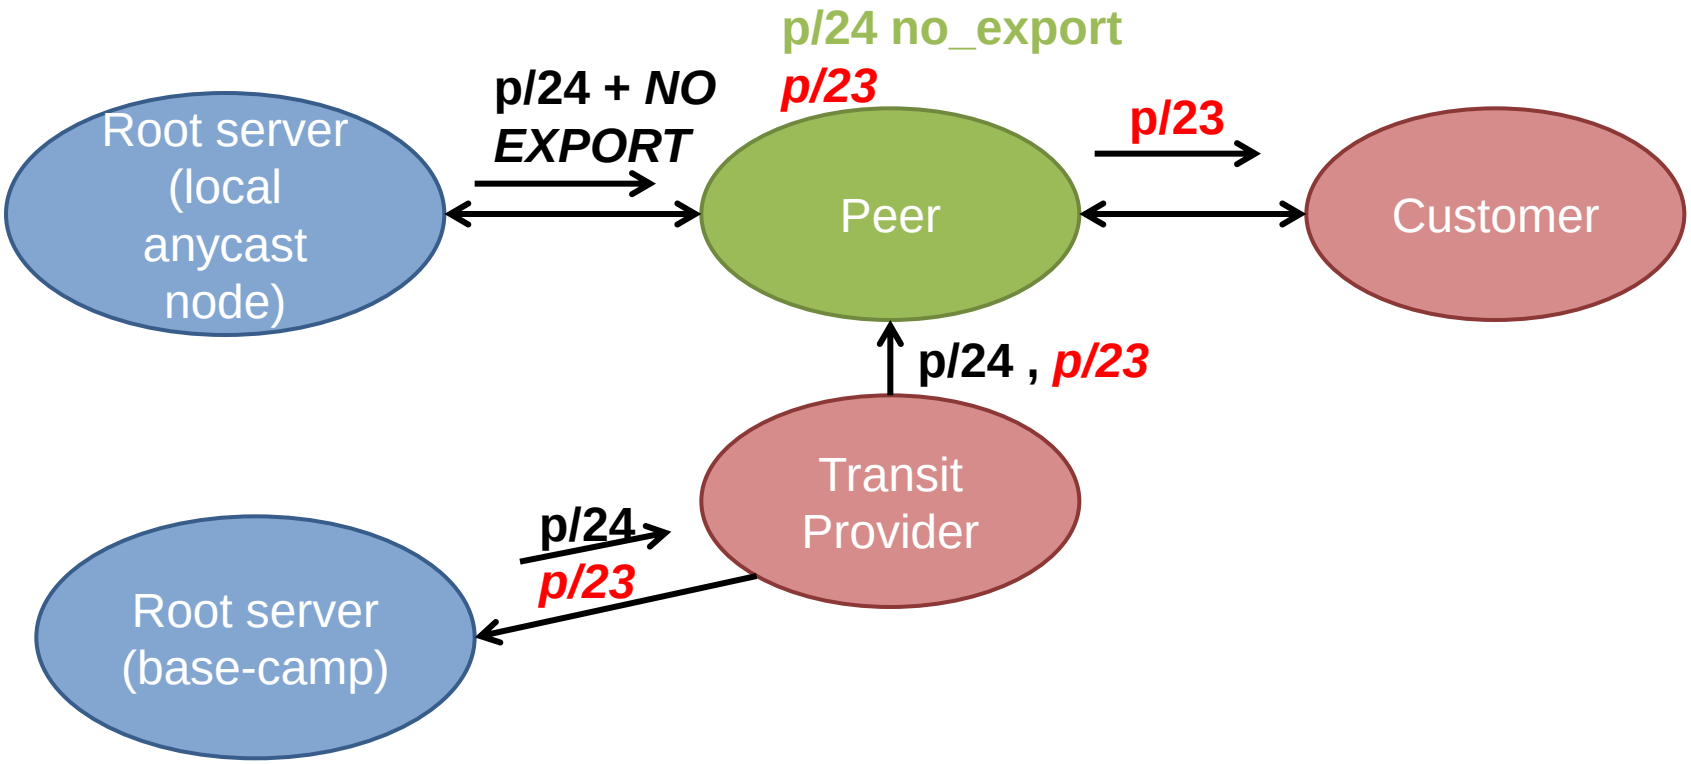

Use Cases – Internet Root Server

- Blackholing due to no announcement

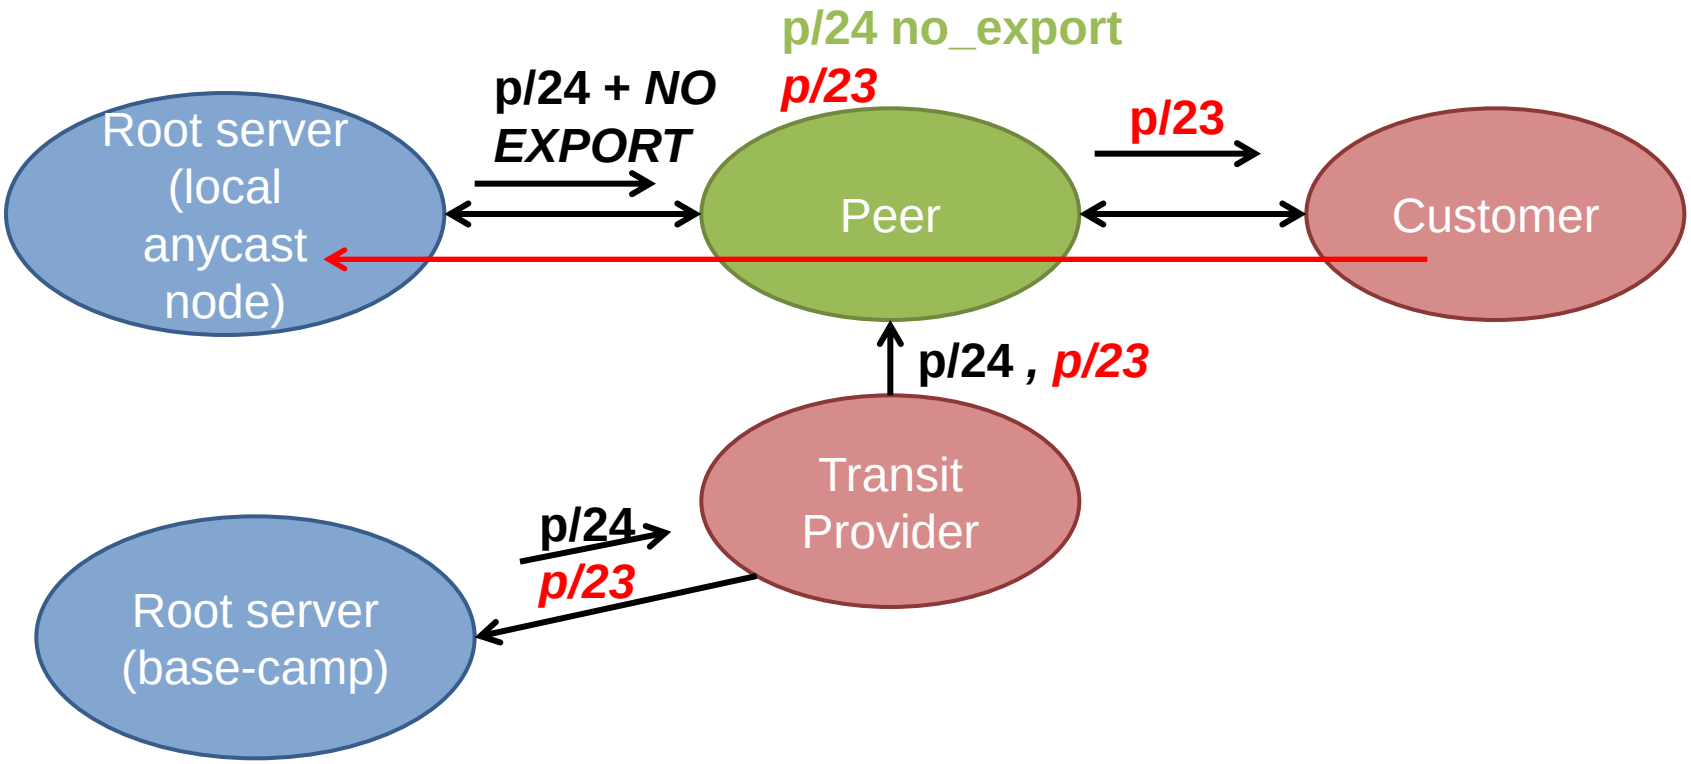

Conclusions

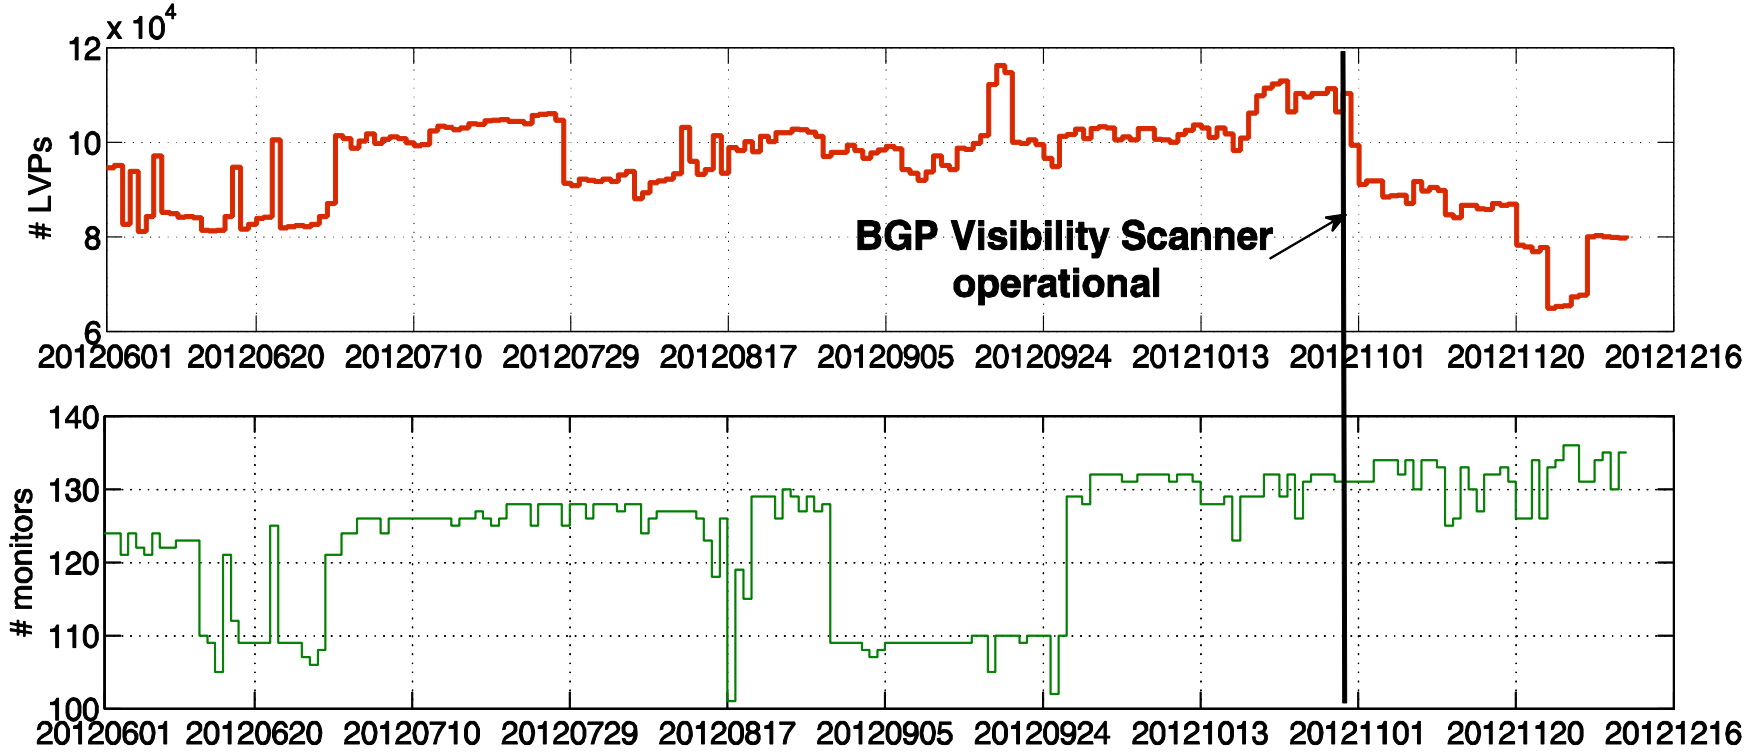

visibility.it.uc3m.es

Questions?

andra.lutu@imdea.org

marcelo@it.uc3m.es

O.M.Maennel@lboro.ac.uk

The paper (GI'13):

The BGP Visibility Scanner