Understanding the Reachability of IPv6 Limited Visibility Prefixes

Andra Lutu*†*, *Marcelo Bagnulo*, Cristel Pelsser‡, Olaf Maennel♣*

*†IMDEA Networks Institute, *UC3M , ‡IIJ, ♣Loughborough University*

Motivation

- ***Does prefix visibility at the interdomain level have an impact on the reachability of the address space?***

- Many networks are interacting, while also defining their routing preferences
- Routing policies defined by network operators may affect the ***global visibility*** of a certain prefix, both intentionally and unintentionally
- ***Global connectivity*** issues have been reported lately in the IPv6 Internet

Outline

- The BGP Visibility Scanner for IPv6
  - Monitor prefix visibility
- Propose a Measurement Methodology
  - Test reachability of an IPv6 prefix
- Measure the Reachability of IPv6 Limited Visibility Prefixes
  - From the RIPE Atlas platform, we test the reachability of the identified IPv6 LVPs
  - Look for correlations with the visibility degree assigned to each prefix

The BGP Visibility Scanner

visibility.it.uc3m.es

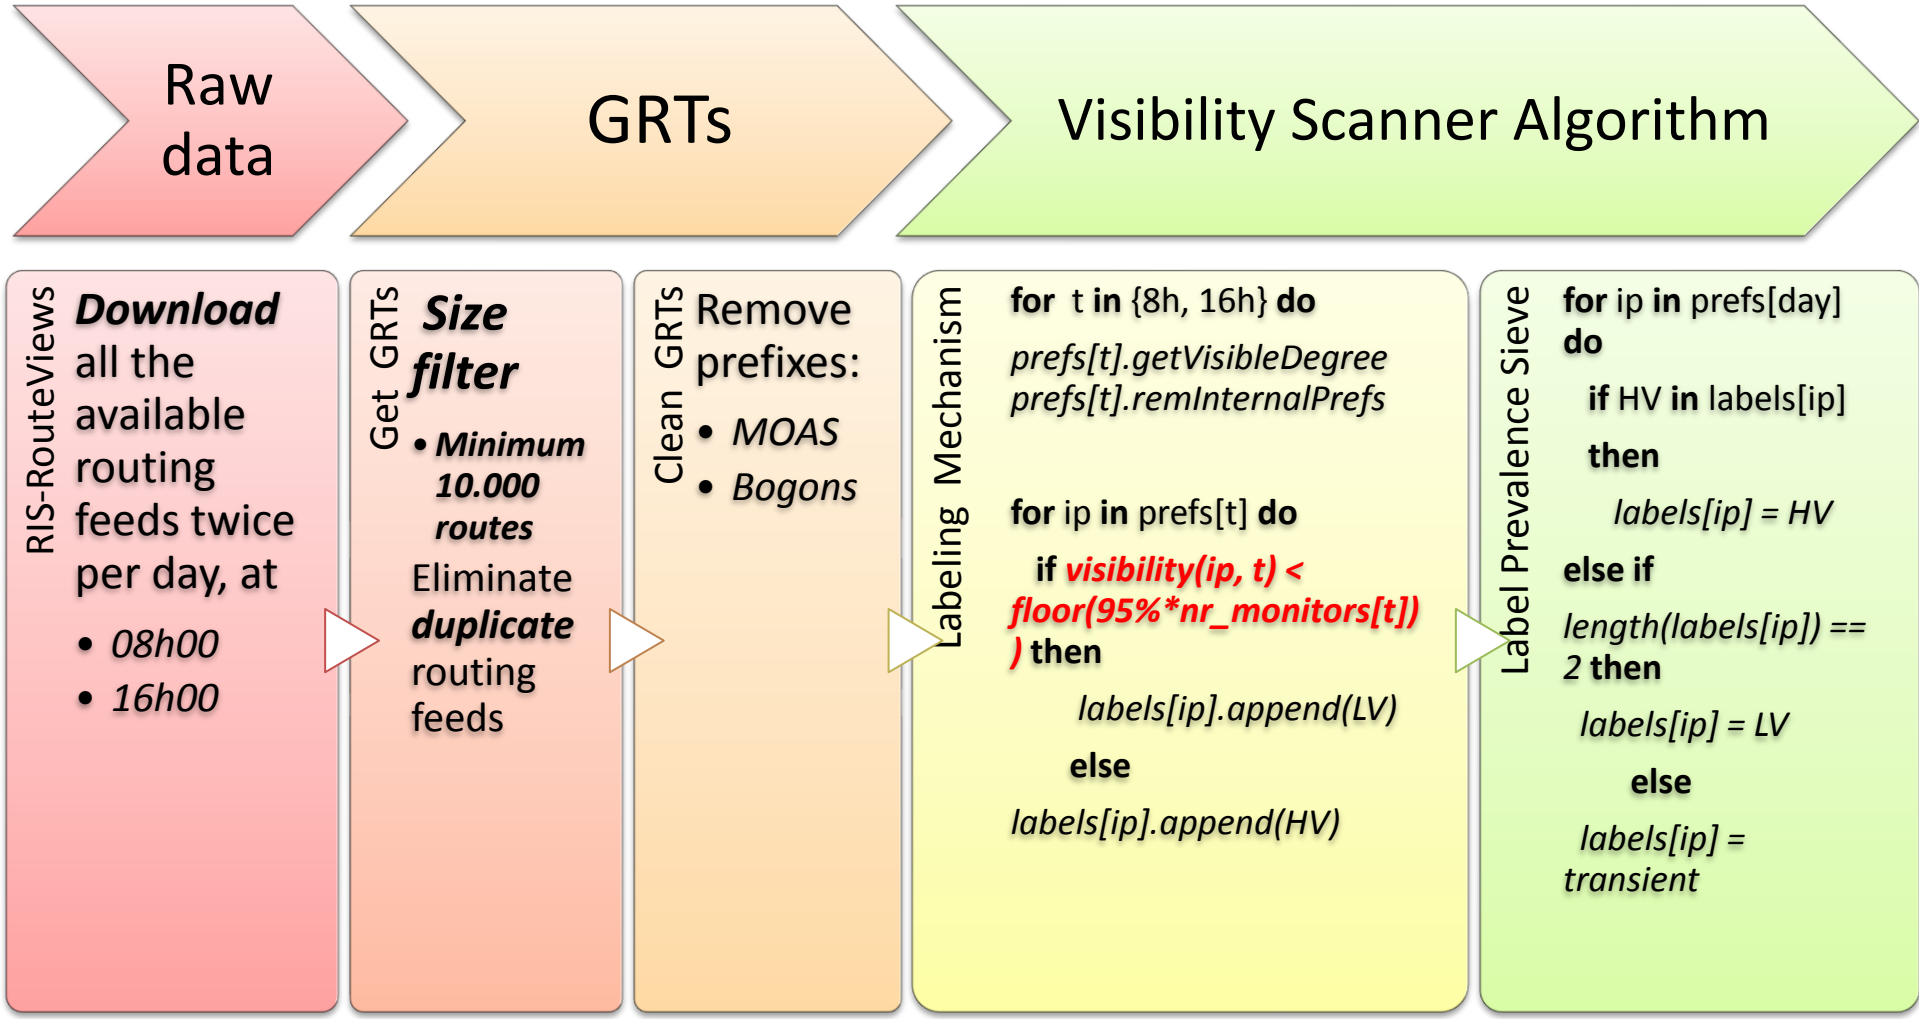

The BGP Visibility Scanner

visibility.it.uc3m.es

- 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

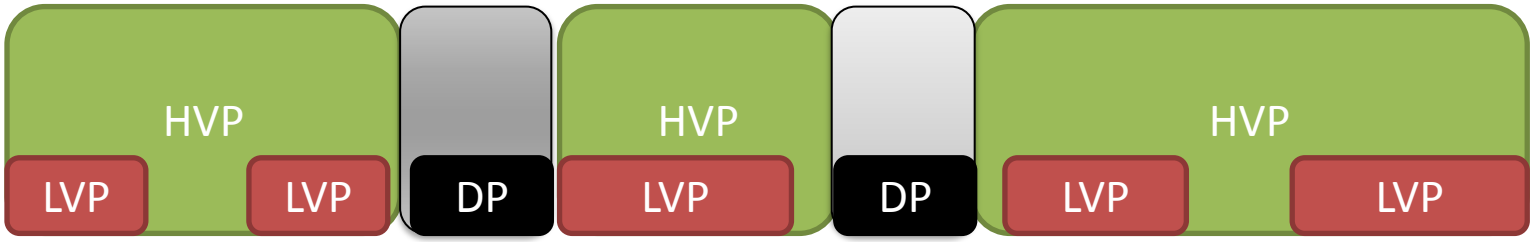

  - **DP –** limited visibility prefixes ***without*** a covering high visibility prefix

Limited Visibility Prefixes

- 110 IPv6 global routing tables
  - ~16,500 IPv6 prefixes
    - 12,500 v6HVPs
    - 3,500 v6LVPs
- ***20% of all the IPv6 prefs are LVPs***
  - 14% of the v6LVPs are ***Dark Prefixes***
  - *This is 5 times more dark address space than what we see in IPv4*
    - Only 3% of the v4LVPs are DPs
- 1,000 IPv6-active ASes inject v6LVPs (out of ~8,000 active networks in total)
  - 40% of these inject dark address space

Question: Why do LVPs emerge?

- Gathered feedback on the ***expected*** visibility status for 20,000 LVPs
  - Invite the ASes operators using the BGP Visibility scanner to fill in survey form
  - Actively interacted with operators to help debug their routing policies
  - Presented the tool in numerous venues, e.g., NANOG, ESNOG, UKNOF, RIPE Labs

Question: Why do LVPs emerge?

- ***Intended*** Limited Visibility Prefixes
- 1,400 LVPs, among which:
  - Content provider doing geographical scoping of prefix advertisements using BGP communities
  - Prefixes injected only to some peers, and not providers

Question: Why do LVPs emerge?

- ***Unintended*** Limited Visibility Prefixes
- 18,500 LVPs, among which:
  - Large ISP accidentally announcing 4,000 internal routes to peers because of misconfigured outbound filters
  - ISP with Dark Prefixes because of misconfiguration in its provider's routing policies
  - Prefixes without an object defined in the Regional Registry's database got filtered

Methodology

- Traceroute to a random IP address within the prefix
- The target IPv6 prefix is reachable if:
  - The traceroute probe traverses the network to which the prefix has been allocated.
  - The traceroute probe traverses the second-last AS along the source's BGP AS-Path for the target prefix.

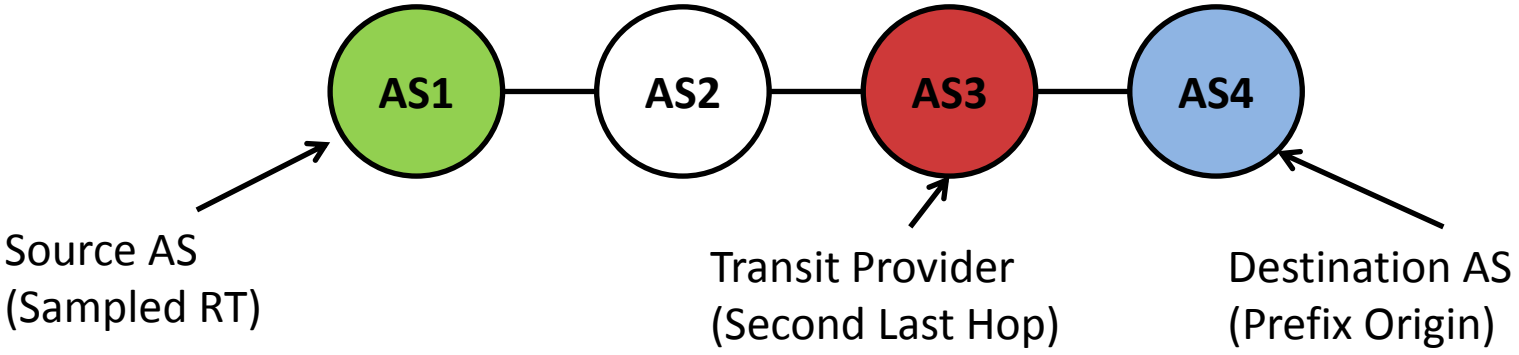

Local Reachability Measurements

- Local reachability measurements
  - Check prefix visibility from the point of view of the Japanese ISP
- We test three different groups of prefixes, from a single source, for which *we also have the BGP routing information*:
    - Data from 8th of August, 2013
  - 13,195 HVPs [prefixes present in the RT]– 92% reachable
  - 2,359 LVPs [have a covering HVP] – 94% reachable
  - **511 DPs [don't have a covering HVP] – <5% reachable**

Global Reachability Measurements

- We test the reachability of the globally-defined v6DPs using ***100 active probes*** within the RIPE Atlas platform

Global Reachability Measurements

- ***Target Prefixes***:
  - ***473 IPv6 DPs*** after analyzing 110 GRTs
  - ***3,200 v4DPs*** after analyzing 154 GRTs
  - Data from the 8th of August, 2013
- Perform one-off ICMP traceroute measurements from each Atlas source probe towards a random address within each v6DP

Methodology

- A destination prefix is reachable if:
  - The traceroute probe reaches the network to which the prefix has been allocated.
  - ~~The traceroute probe traverses ***the second-last AS*** along the BGP AS-Path for the target prefix.~~
  - *The traceroute probe traverses **any of the probable second-last ASes to the origin AS** of the target prefix.*

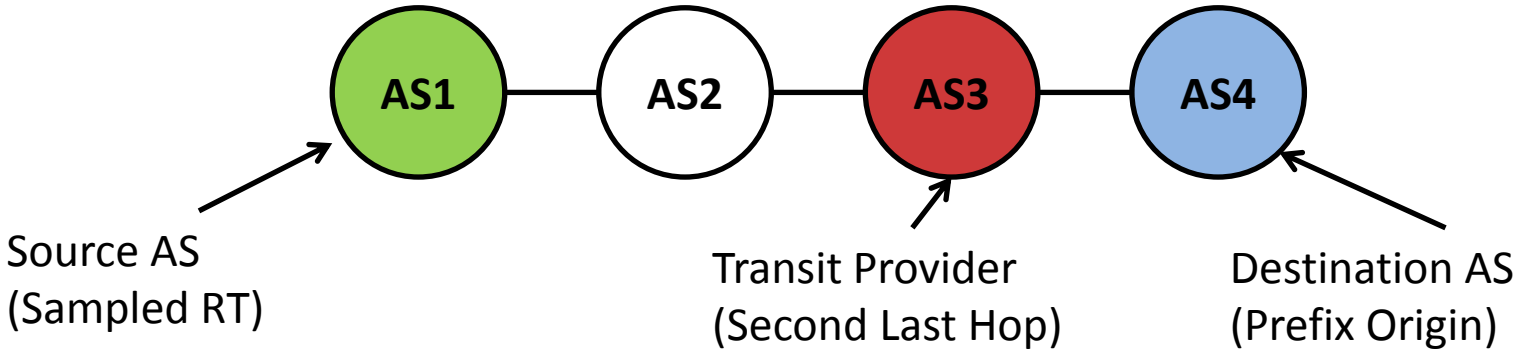

Results

- Average reachability degree for a v6DP is of **46.5%**
- Average reachability degree for v4DPs is of **17.4%**

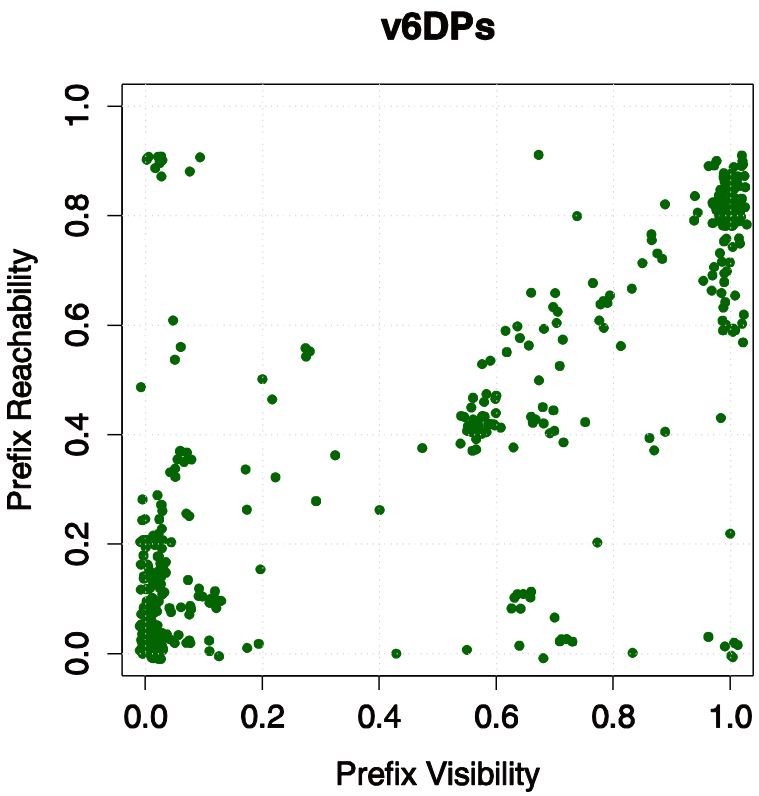

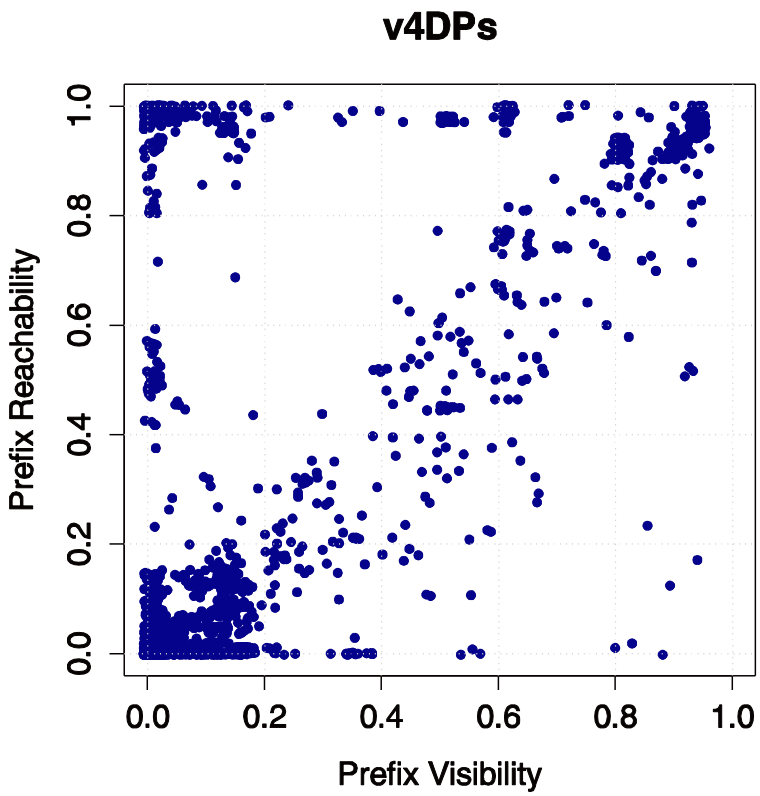

Results

- Average reachability degree for a v6DP is of **46.5%**
- Average reachability degree for v4DPs is of **17.4%**

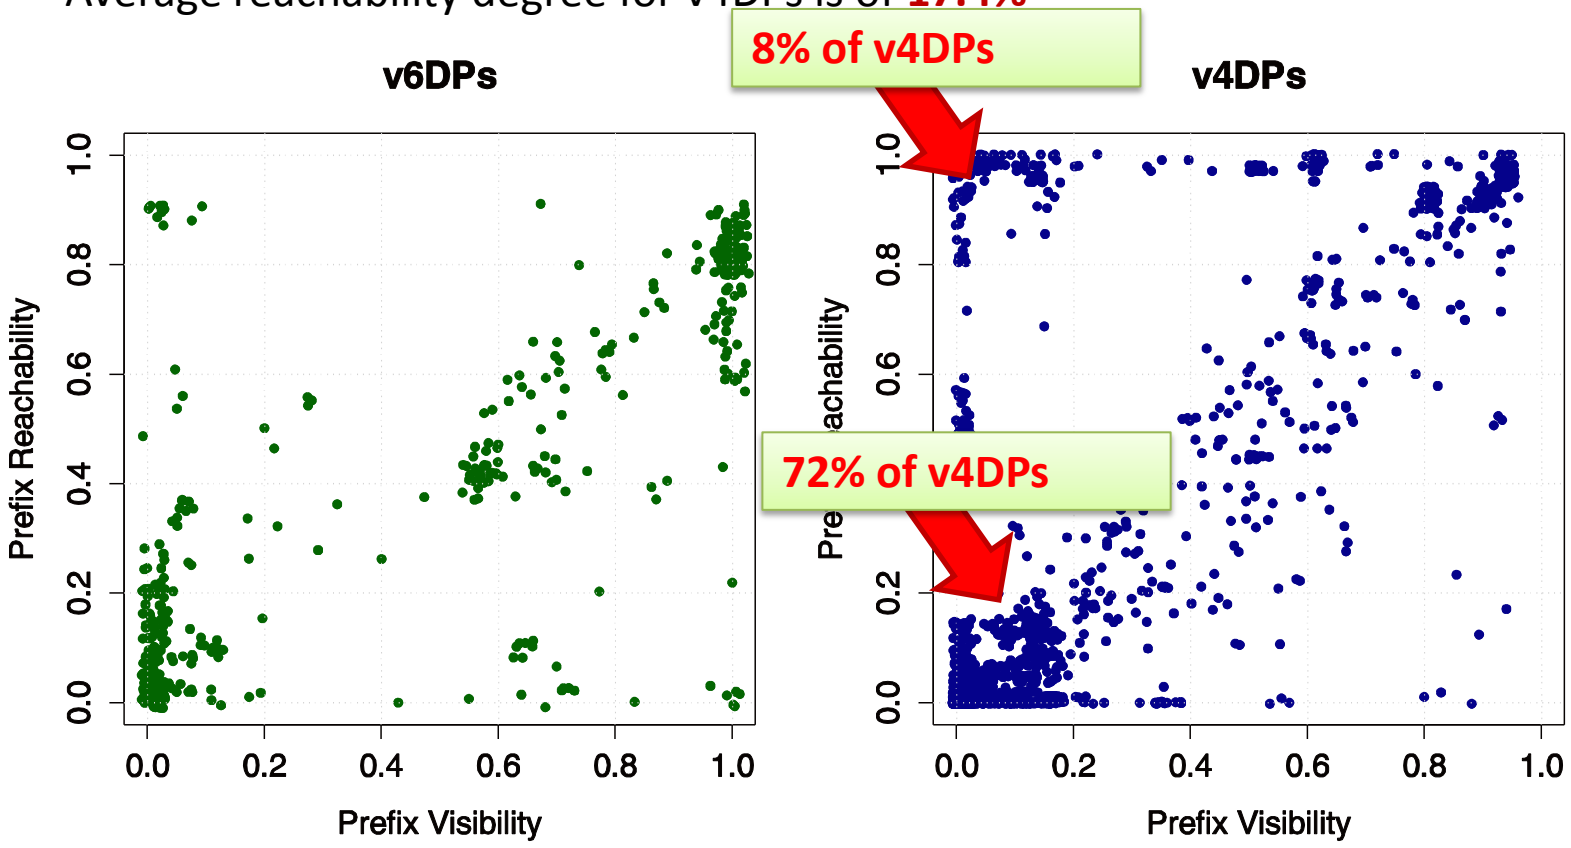

Conclusions

- While the ratio of LVPs is similar for IPv4 and IPv6, we see *5 time more DPs in IPv6 than in IPv4*
- Strong correlation between visibility and reachability for v6DPs
  - The lack of visibility may signal more important problems in IPv6, namely the lack of global connectivity
- While the v4DPs may be largely explained as long-lived route leaks or mistakes, we believe this is not the case for the v6DPs!
  - Side-effect of early stages of IPv6 deployment

Help us to help you!

- Go to **visibility.it.uc3m.es**
- Check if the prefixes of an AS are LVPs/DPs– monitor the global visibility of your prefixes!
- … and tell us why the prefixes discovered have limited visibility in the first place: intended/unintended behaviour?

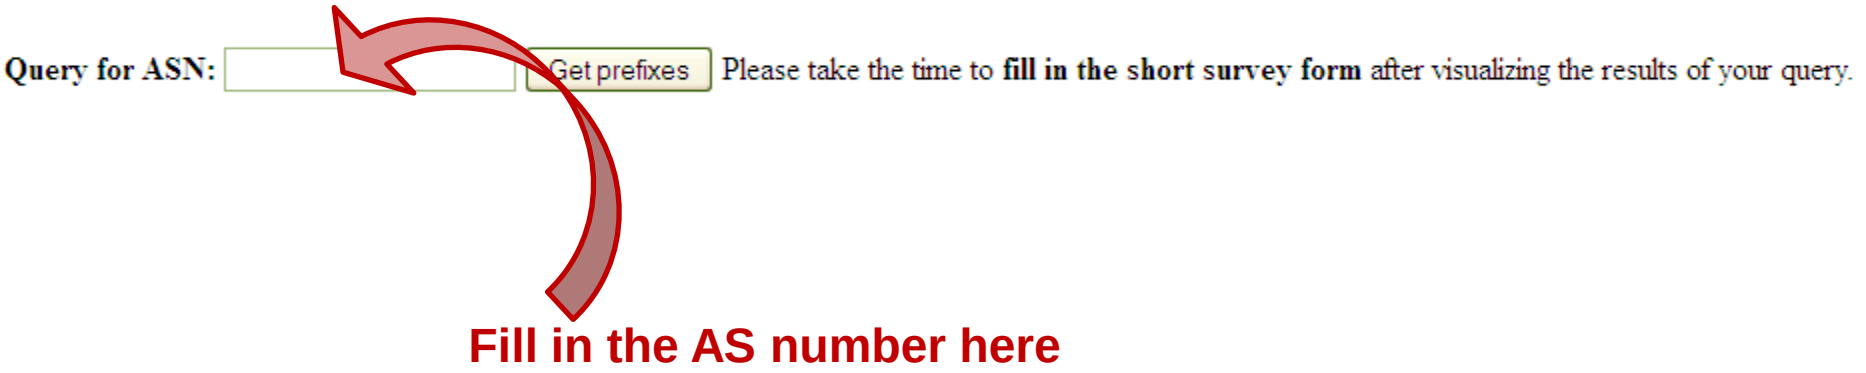

Help us to help you!

- For questions/feedback use the **FORM** at the end of the query!
- …or e-mail us!

andra.lutu@imdea.org

marcelo@it.uc3m.es

cristel@iij.ad.jp

O.M.Maennel@lboro.ac.uk

Questions?

visibility.it.uc3m.es

andra.lutu@imdea.org

marcelo@it.uc3m.es

cristel@iij.ad.jp

O.M.Maennel@lboro.ac.uk

Use Case – Internet Root Servers

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

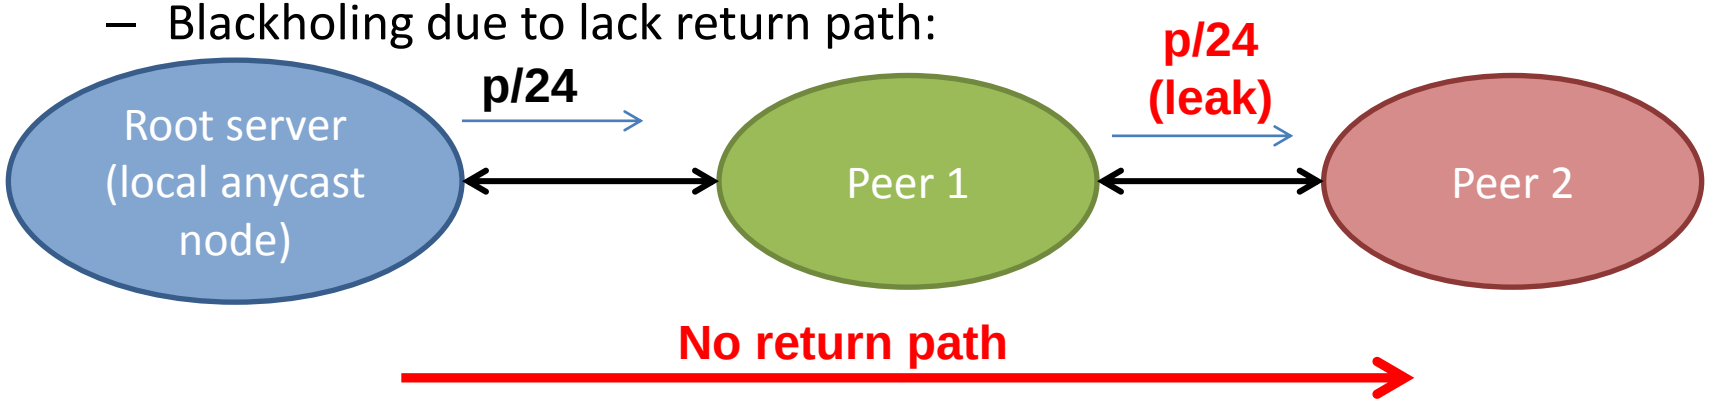

Use Case – Internet Root Servers

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

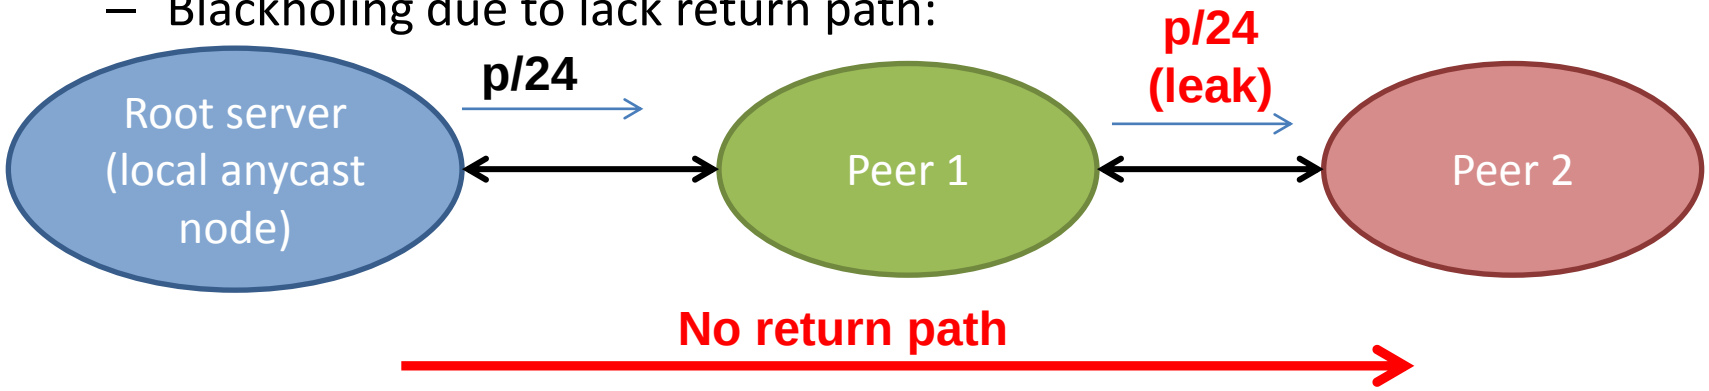

  - No full transit at the IXP => tag with NO EXPORT

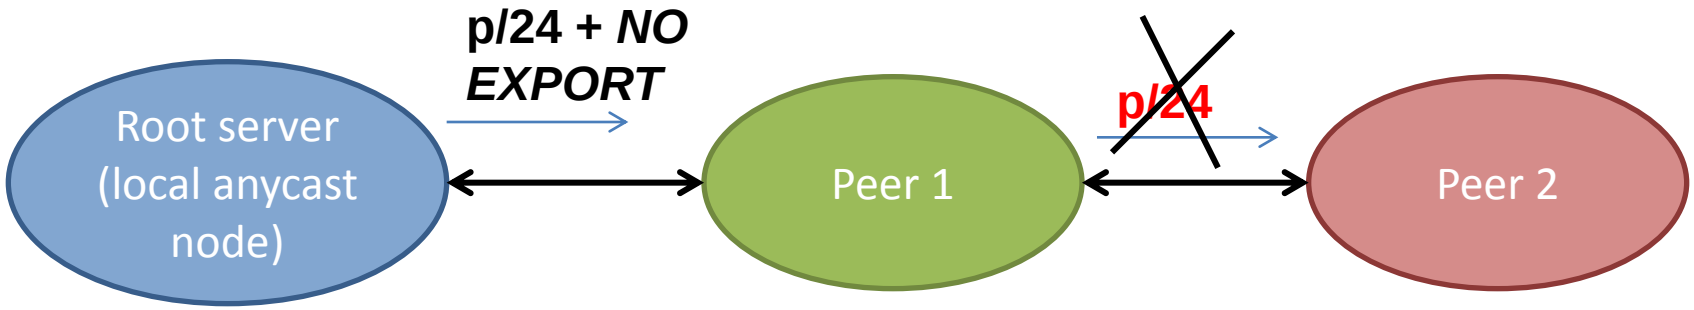

Use Case – Internet Root Servers

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

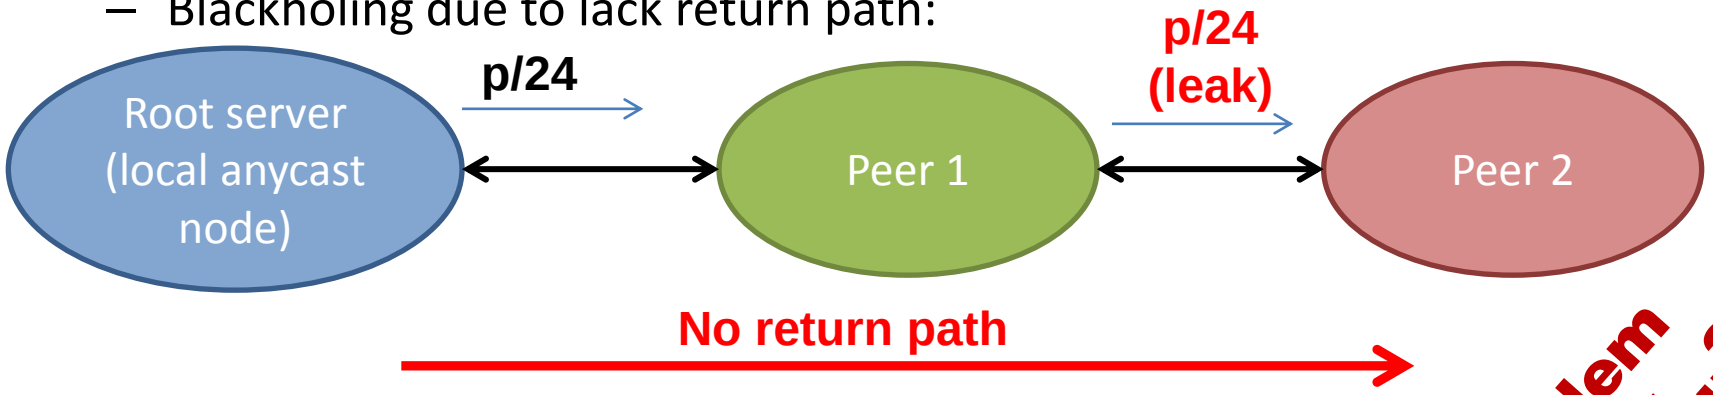

  - No full transit at the IXP => tag with NO EXPORT

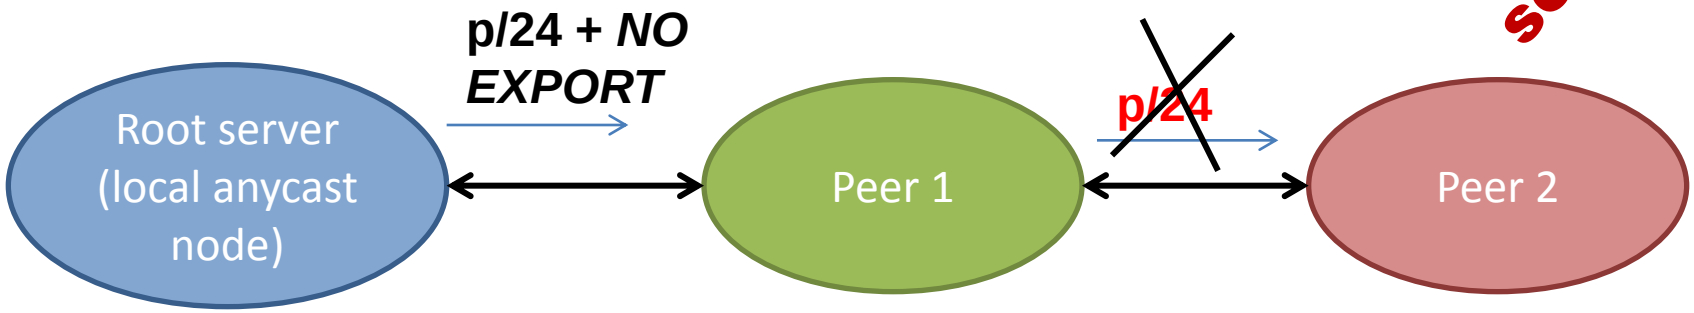

Use Case – Internet Root Servers

– Blackholing due to no announcement

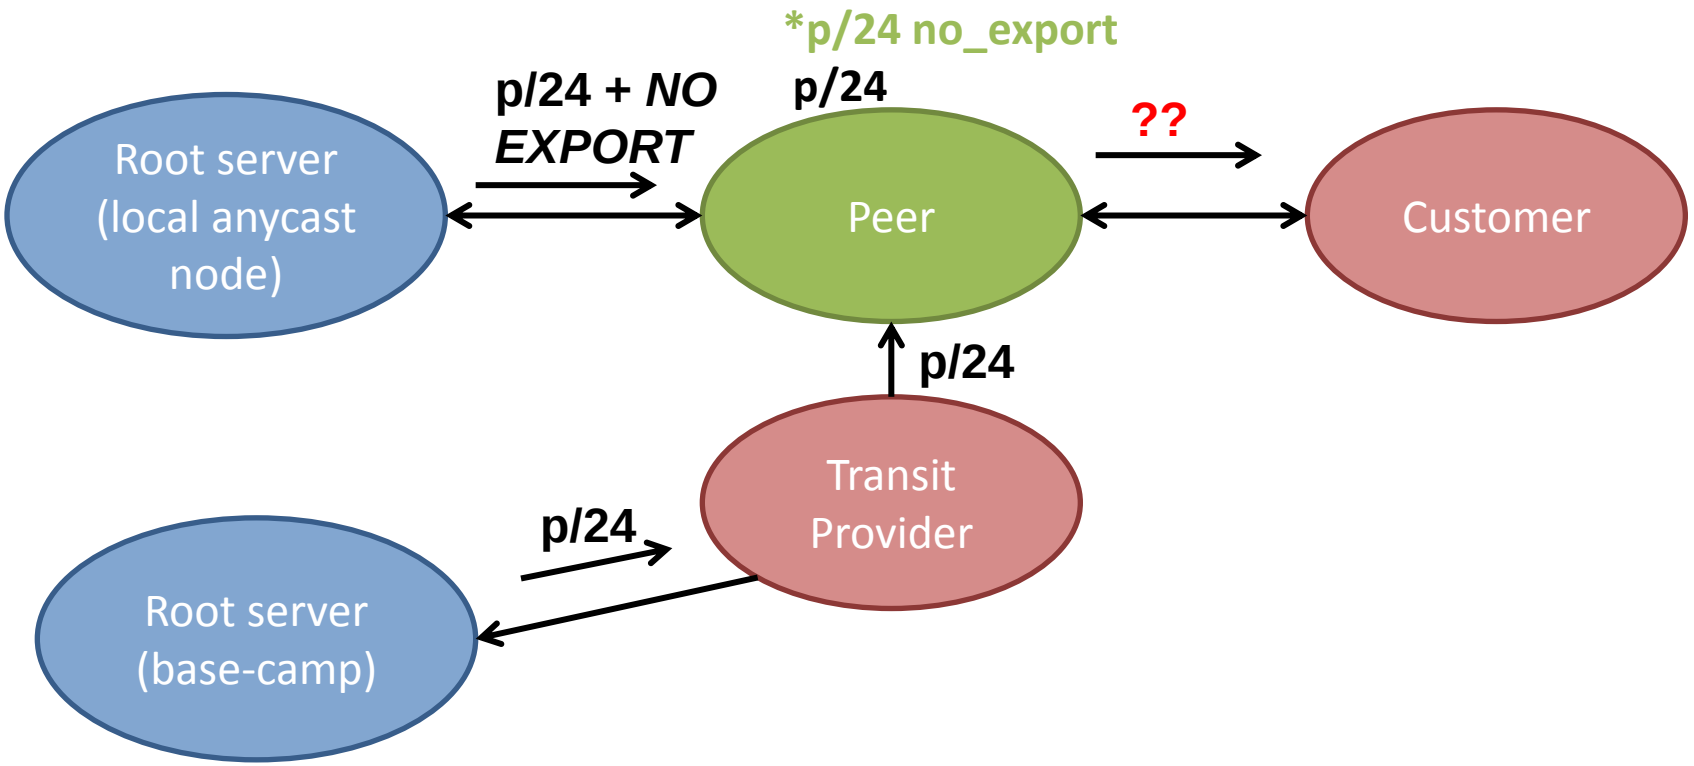

Use Case – Internet Root Servers

- Blackholing due to no announcement

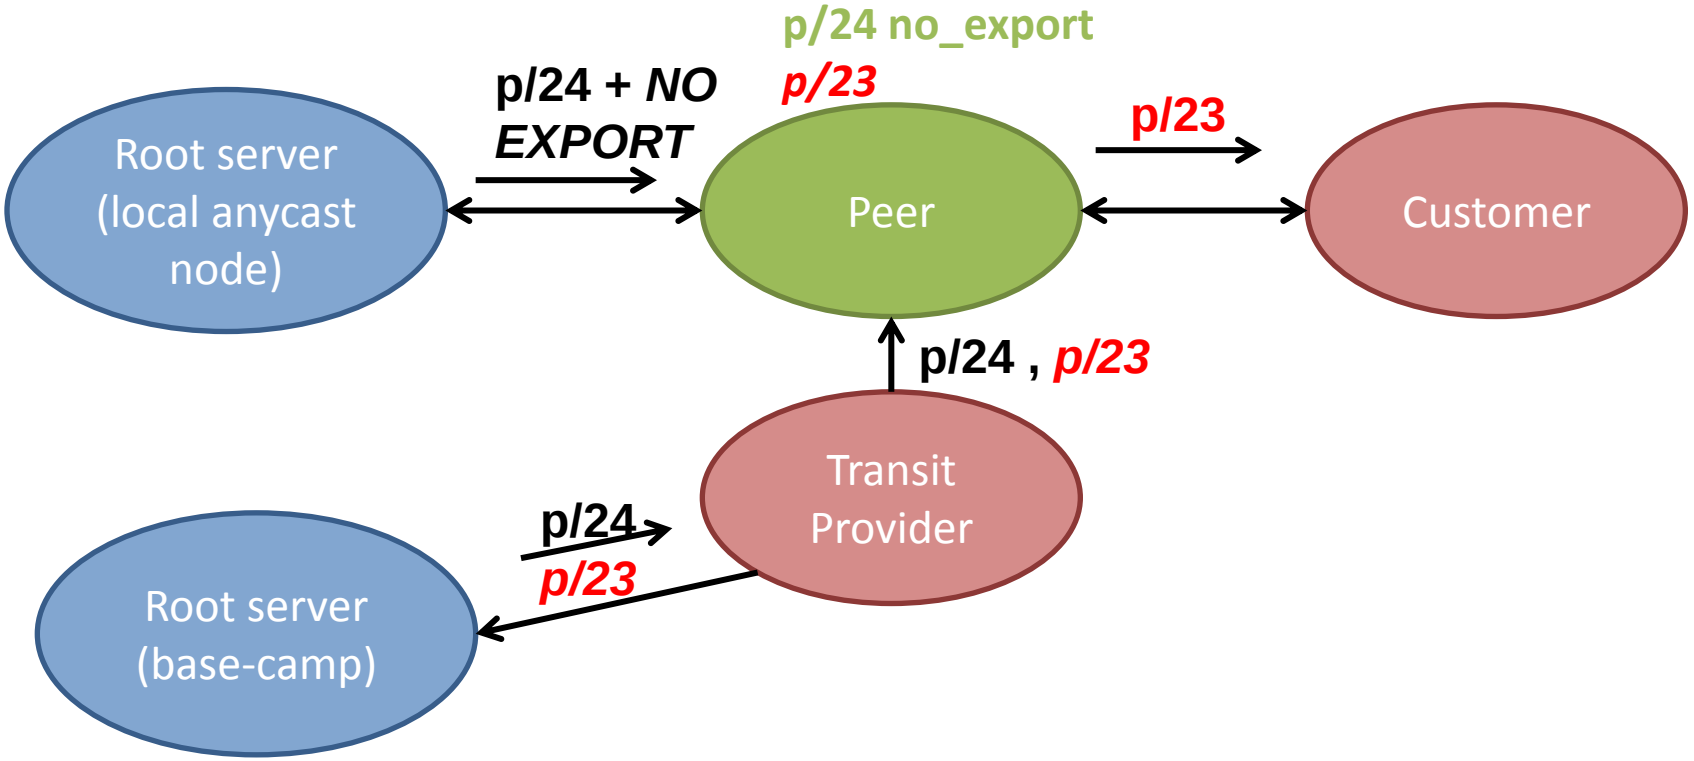

Use Case – Internet Root Servers

- Blackholing due to no announcement

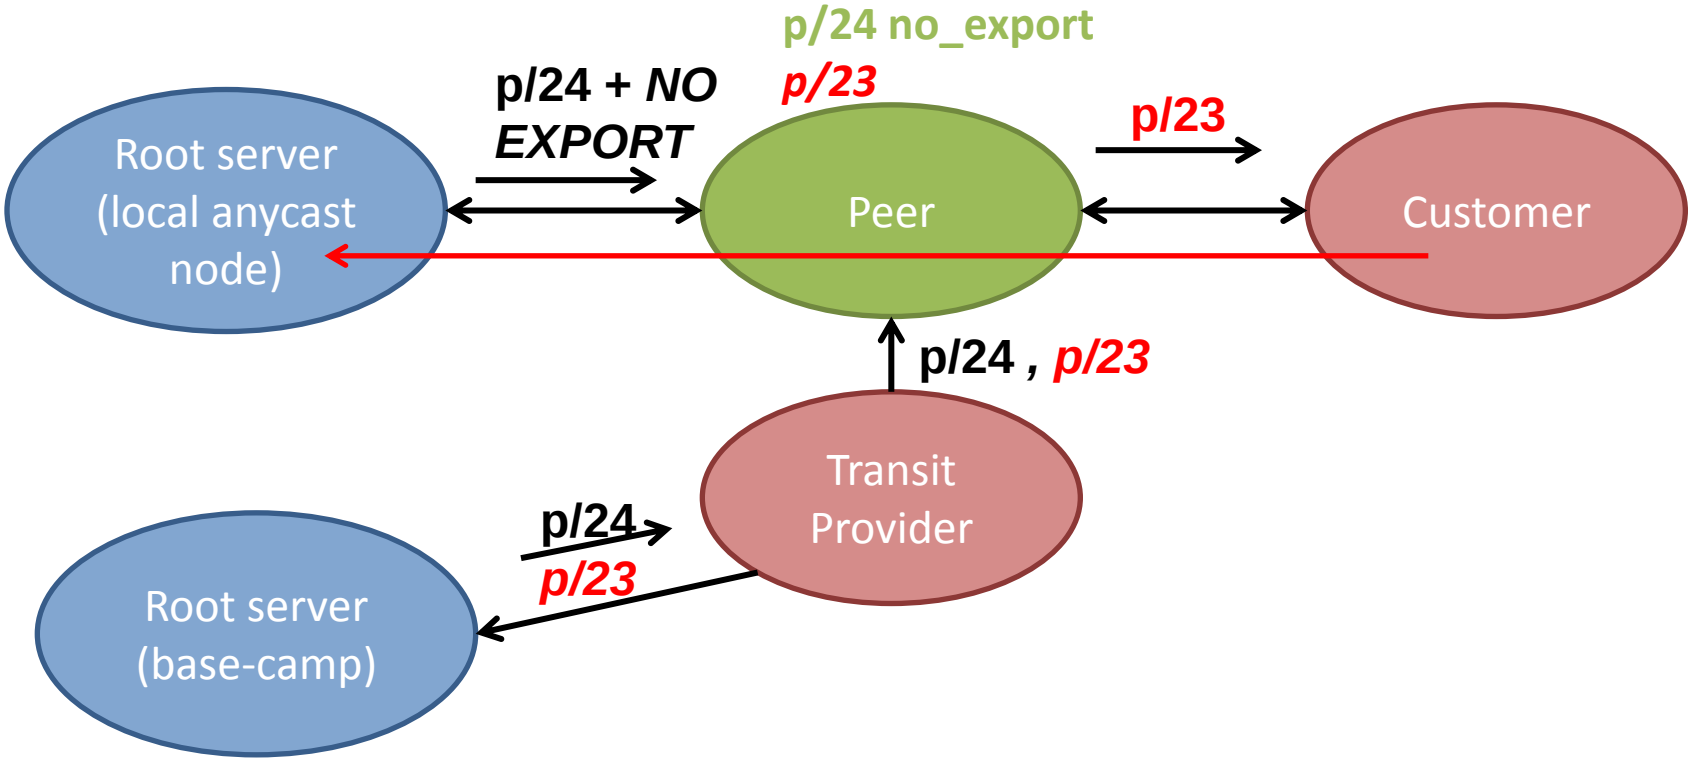

BGP Prefix Visibility – IPv4

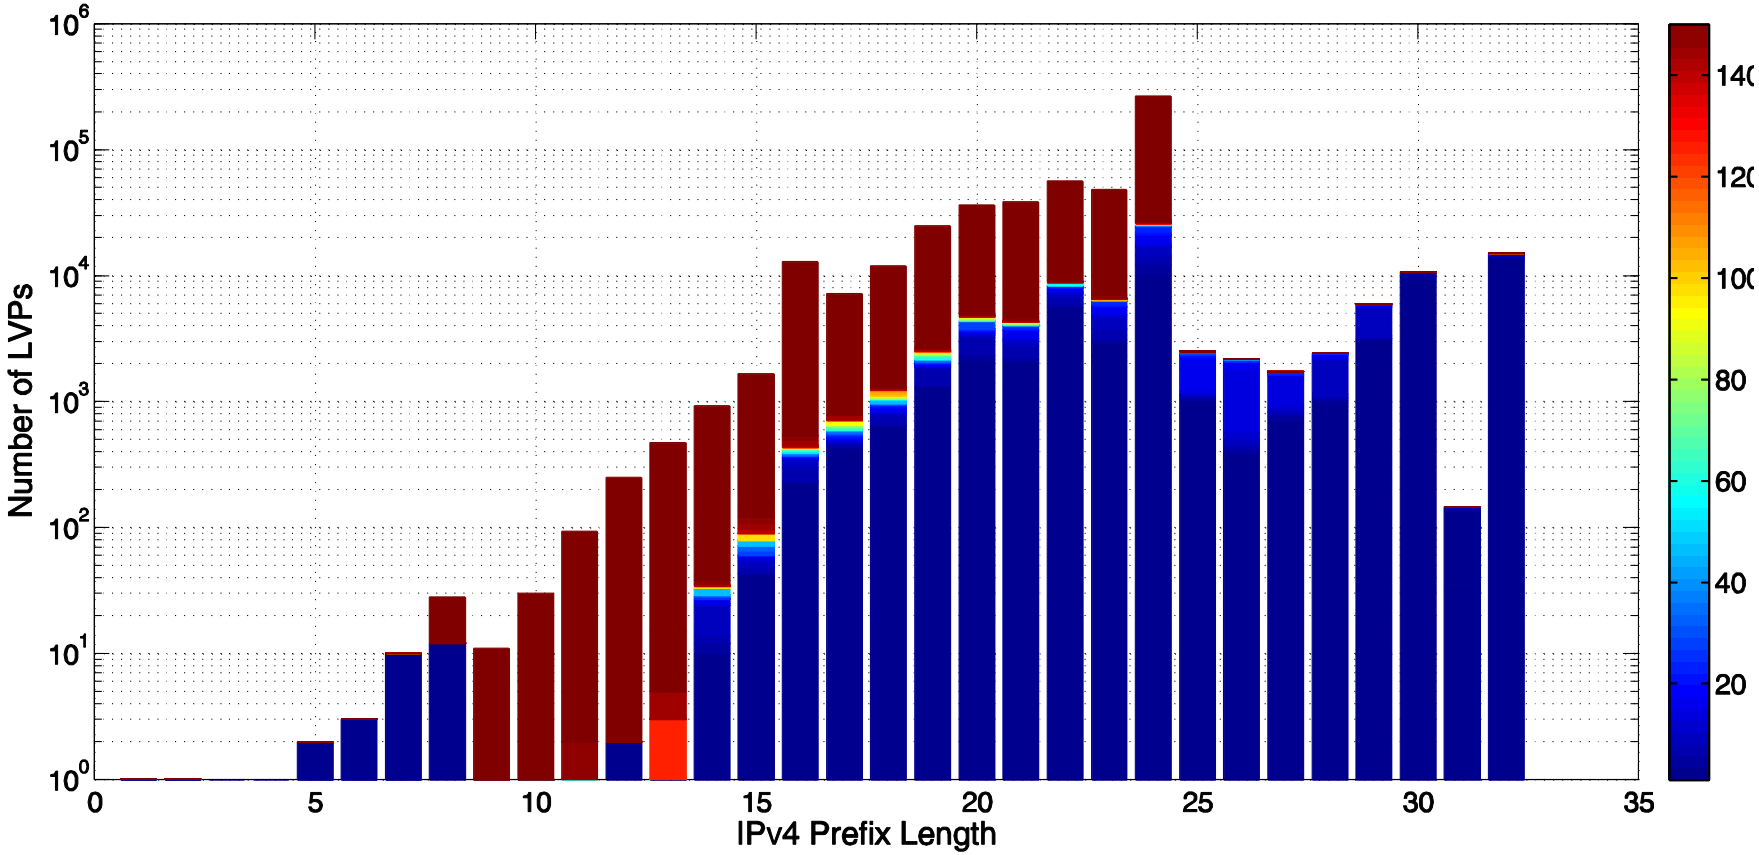

BGP Prefix Visibility – IPv6

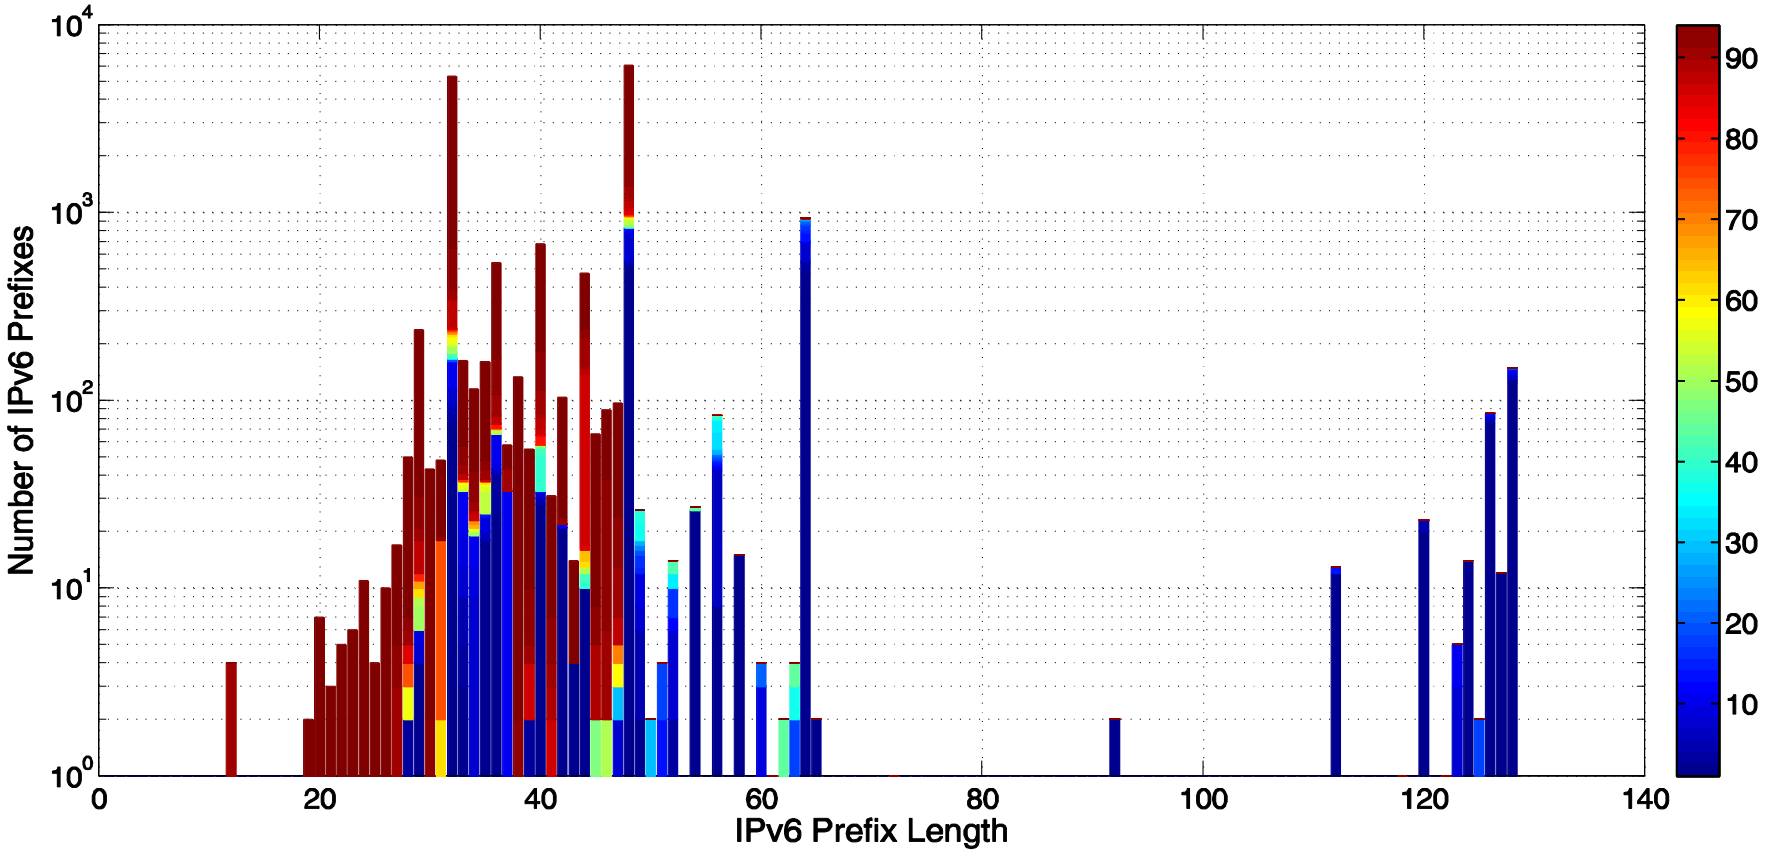

Prefix visibility as of 23.10.2012

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