max planck institut informatik

RIPE 86
Rotterdam, Netherlands
22 - 26 May 2023

Hyper-Specific Prefixes: Gotta Enjoy the Little Things in Interdomain Routing

Presenter: Khwaja Zubair Sediqi

23.May.2023

Authors: Khwaja Zubair Sediqi, Lars Prehn, Oliver Gasser
zsediqi@mpi-inf.mpg.de, lprehn@mpi-inf.mpg.de, oliver.gasser@mpi-inf.mpg.de
Paper Published at: ACM SIGCOMM Computer Communication Review, Volume 52 Issue 2, April 2022

Introduction

ASes use the BGP to announce prefixes

BGP best practices recommend filtering prefixes

- more specific than /24 in IPv4 and /48 in IPv6

Plenty of /25 to /32 IPv4 and /49 to /128 IPv6 exist

hyper-specific prefixes (HSPs)

How prominent and why HSPs exist in the Internet routing ecosystem?

Related Work

In 2014 and 2015 Aben and Petrie

- announced /24, /25, and /28 IPv4 prefixes
- RIPE Atlas measurements
- HSPs visible at most 20 % of RIPE RIS peers

In 2017, Strowes and Petrie conclude

- at most one fourth of all BGP peers

In 2017, Huston analyzed different types of more-specific prefixes

1. hole punching (different origin AS),
2. traffic engineering (same origin AS, but different AS path),
3. overlay (same AS path)

Methodology

For our analysis we utilize “snapshots“ from the RC projects RIPE RIS , Routeviews, and Isolario

- From Jan.2010 to October.2021
- Quarterly, 7days per quarter
- BGP RIBs – every 24 hours
- BGP Updates – every 5 mins
- Applied filters to clean the data

Supplemental datasets

caida

RAPID7

ASDB

RIPE NCC RIPE NETWORK COORDINATION CENTRE

ROUTE VIEWS 6447

Isolario Project

1. OBSERVABILITY

HSPs in Routing Ecosystem

Share of HSPs in the Interent

HSPs make ~ 14% to more than 20% of of all the prefixes

HSPs make ~ 10% of all the prefixes

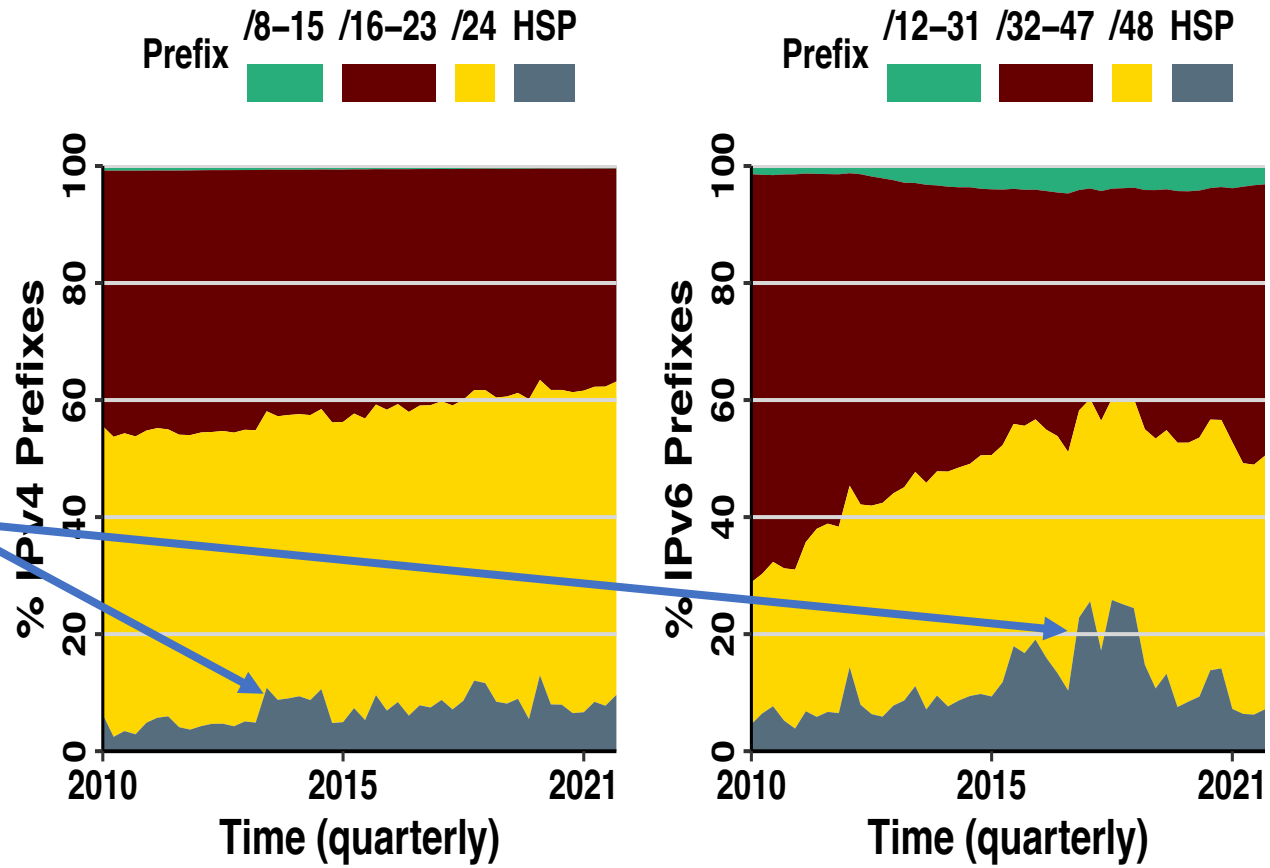

HSP Visibility and Consistency

We use one year data of BGP RIBs and updates

- to track every HSP for the whole year

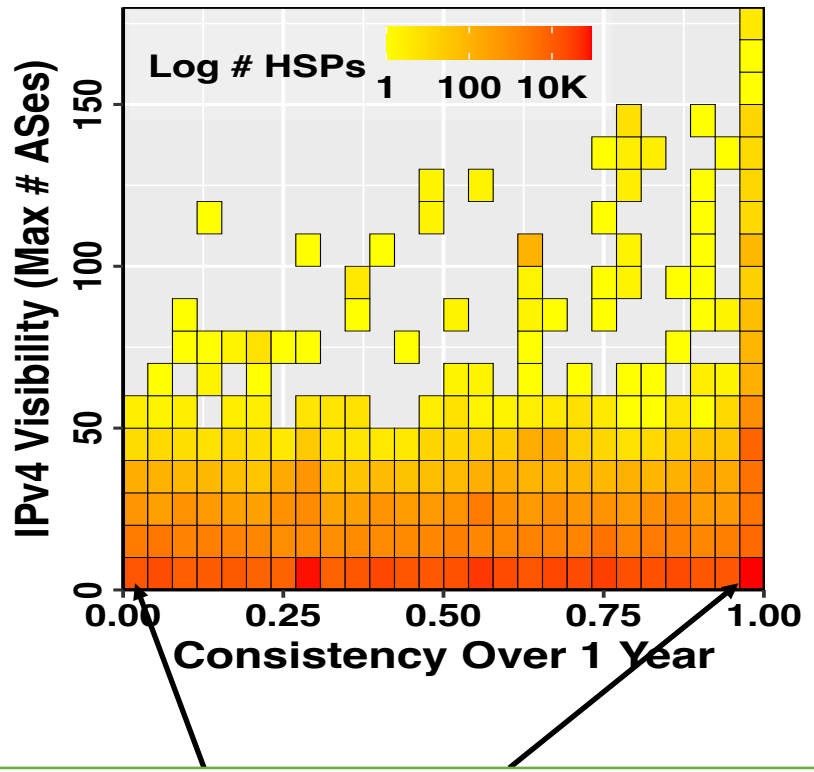

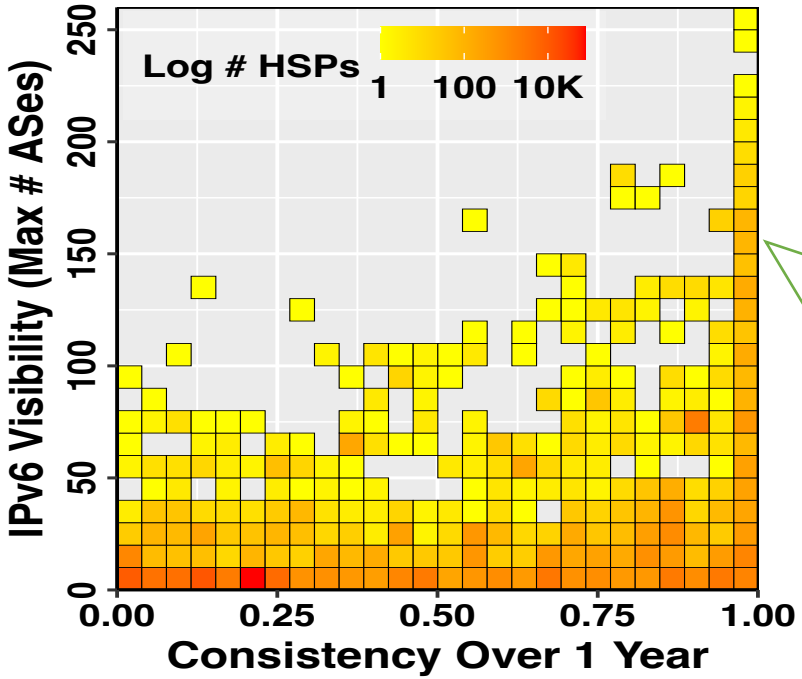

There is a correlation between consistency and visibility

HSPs have life span from days to more than a year

Many have visibility to less than 50 peer ASes

2. USE CASES & FUNCTIONS

CIDR Sizes of HSPs

CIDR sizes hint use cases

- /32 and /128 for blackholing purposes
- /30, /29 peering subnets
- /56 and /64 address block assignments
- /25 traffic engineering

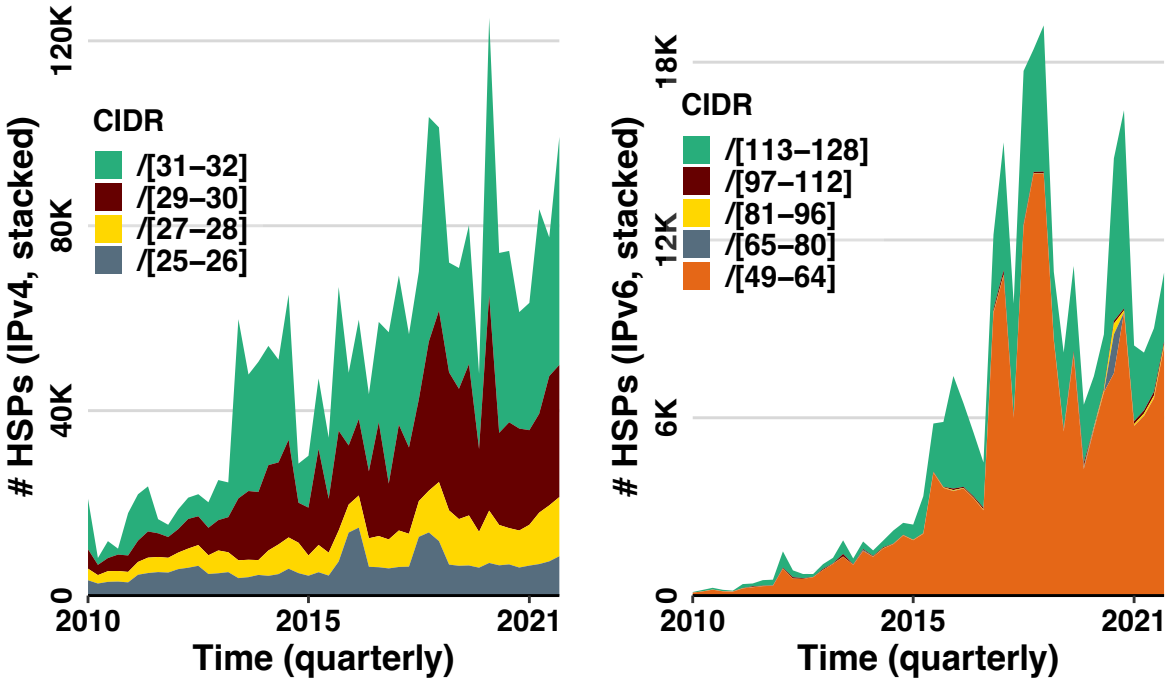

HSPs have heterogeneous use cases

Protocols on HSP IPs

We leverage Rapid7's Open Data platform

Responding hosts and total tested hosts per-protocol

Top5 Protocols:

- CWMP is only present in the IPv4-wide
- BGP is only present in the HSP

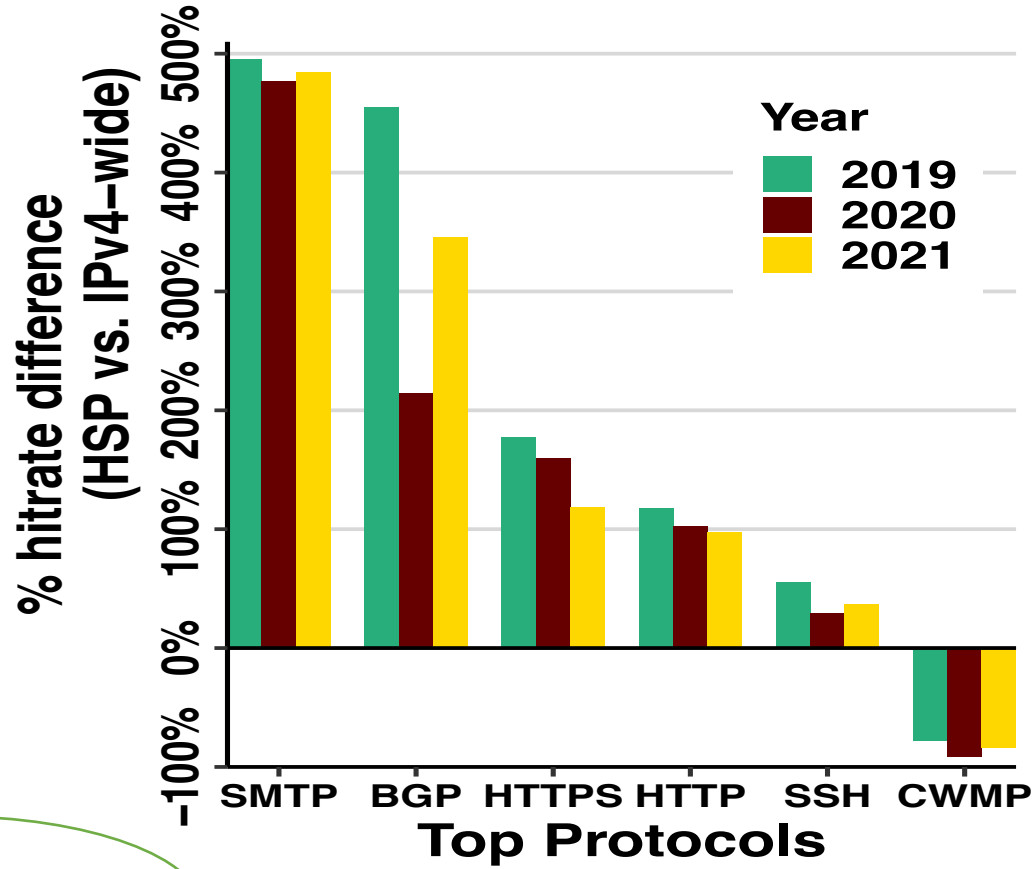

HSPs have upto 5 times higher hitrate than IPv4-wide

BGP Communities of HSPs

We examine BGP communites:

- specifically used for blackholing (BH)
- restrict route propagation (RES)

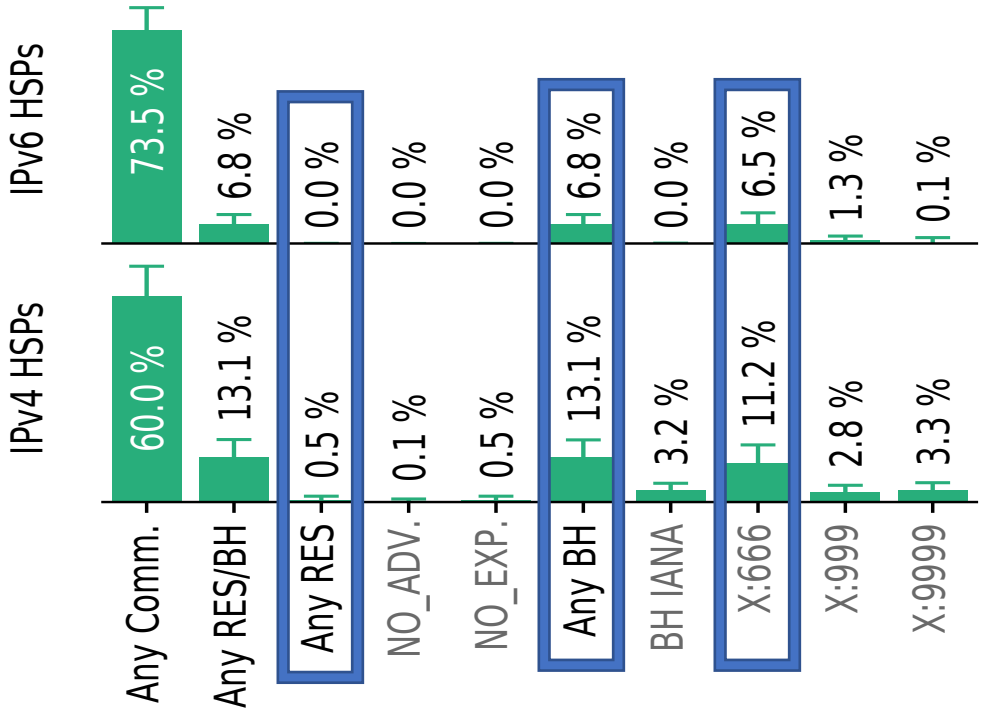

13% and 7% of IPv4 and IPv6 HSPs are Blackholing

3.INTENDED OR ACCIDENTAL USE?

HSPs Origin ASes in Public Databases

IRR has high HSP origin ASes

Many HSPs from RC/BGP have no entries in operator databases

- could be accidental announcements
- misconfigured route collector sessions
- leak of internal routes

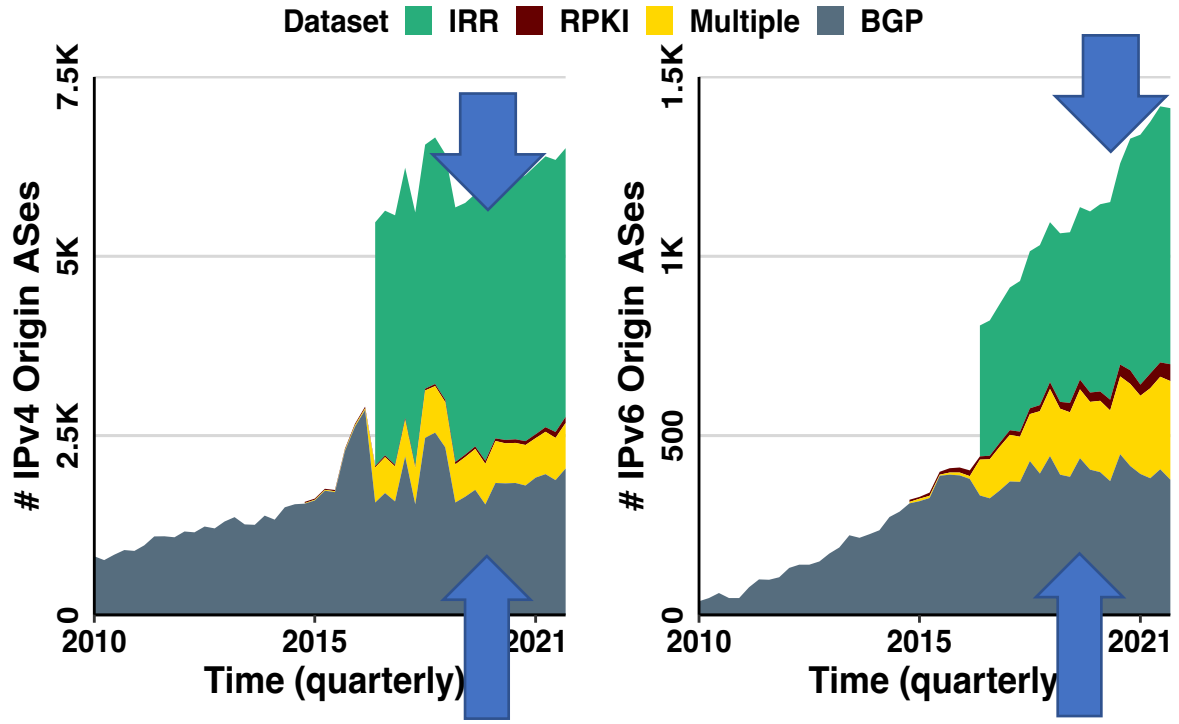

Are HSPs caused by BGP prefix hijacks?

HSPs in the RPKI Database

Invalid (Length) - largest group

Invalid (Origin) - a minor fraction

Invalid (Origin) and Invalid (Both):

- not entered sibling ASes
- DDoS Protection Service (DPS)

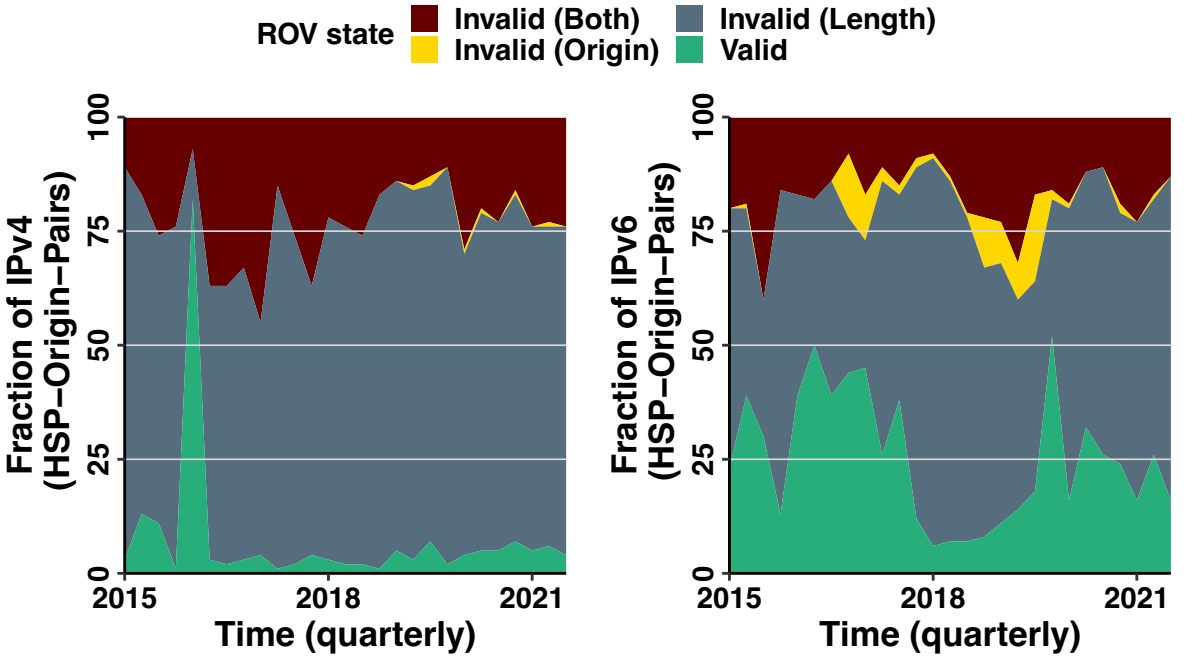

legitimate ASes announce 75 % of HSPs

4. THE FUTURE OF HSPS

Discussion: Research Community

RC projects play a vital role in awareness

HSP dashboard https://hyperspecifics.io

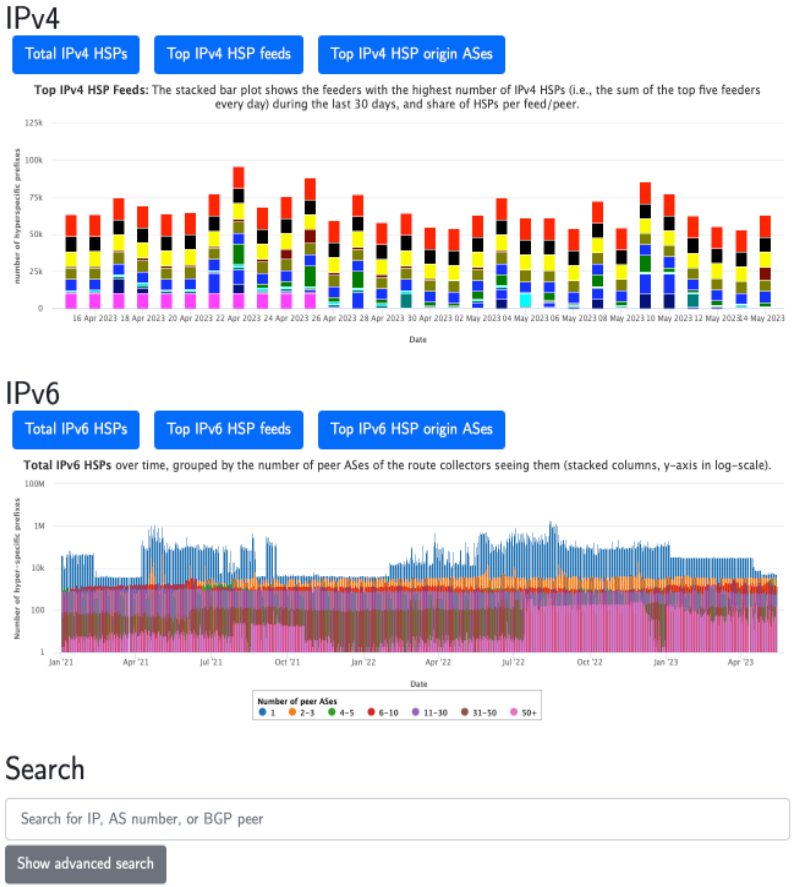

Discussion: Operator Community

Discussing with thirteen operators

- cutomer requests
- traffic engineering

Question: Should operators filter HSPs in the first place?

- for IPv6, Yes, no shortage of IPv6, avoid large routing table size
- for IPv4, shifting filters by a few CIDR sizes (e.g., /26 or /28)

How do you handle HSPs in your network/work ?

Conclusion

We analyzed HSPs in routing ecosystem for the last decade

Most HSPs visible by a few RC peers, still plenty propagate to hundreds of RC peers

IPv4 HSPs: blackholing and infrastructure announcements

IPv6 HSP: related to address block reassignments

Though, hundreds of networks use HSPs intentionally, we attribute even more cases to the accidental “leakage” of internal routes

HSP dashaboard and the paper

https://hyperspecifics.io

Backup Slides

Users of HSPs

Comparing all BGP-visible Ases to HSP origin ASes

- ISP(Transit) originate more HSPs
- 12 to 15 of the total 19 Tier 1's originate HSP
- most hypergiants do not originate HSPs

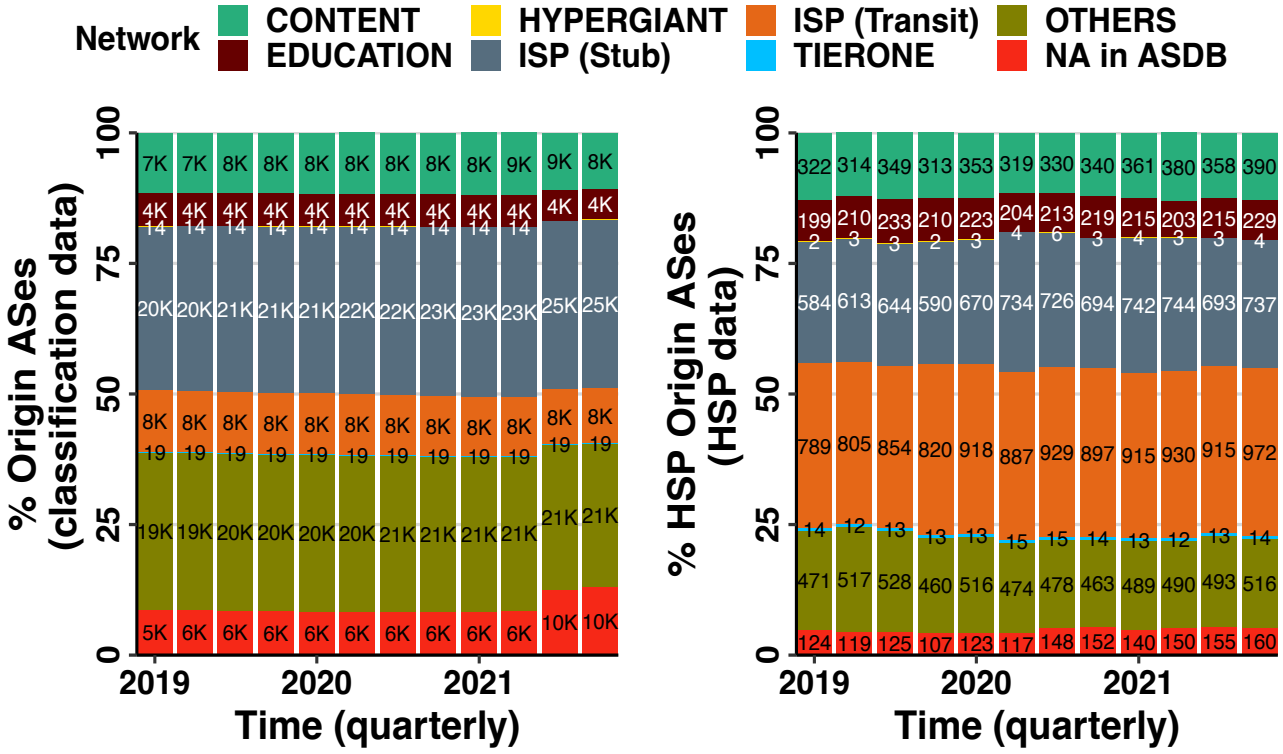

Growth of HSPs Over Time

presence of HSPs increased

one-tenth of all the prefixes

in IPv4 the increase in HSPs is driven by an increment in feeder ASes

IPv6 we see an increase also for a constant set of feeder ASes

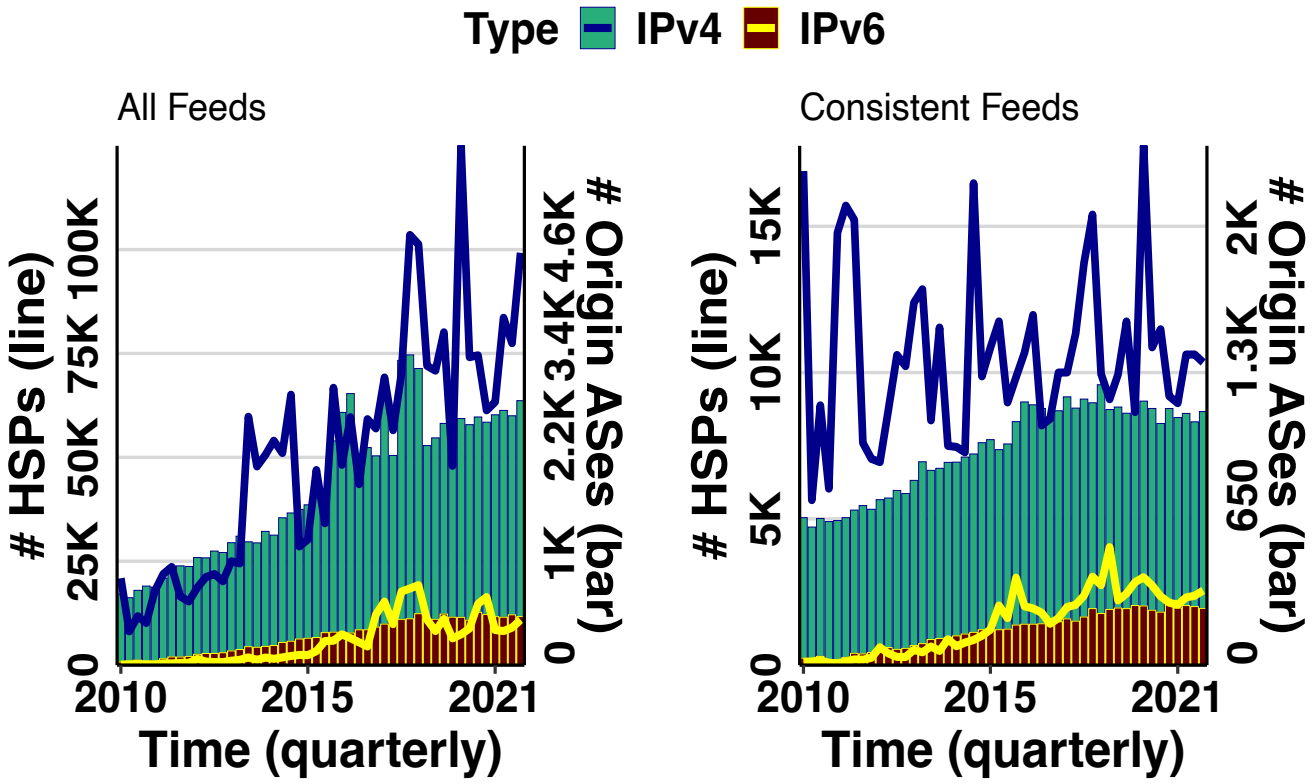

HSP Aggregation

Analyse anchor-prefixes:

- /24 in IPv4
- /48 in IPv6

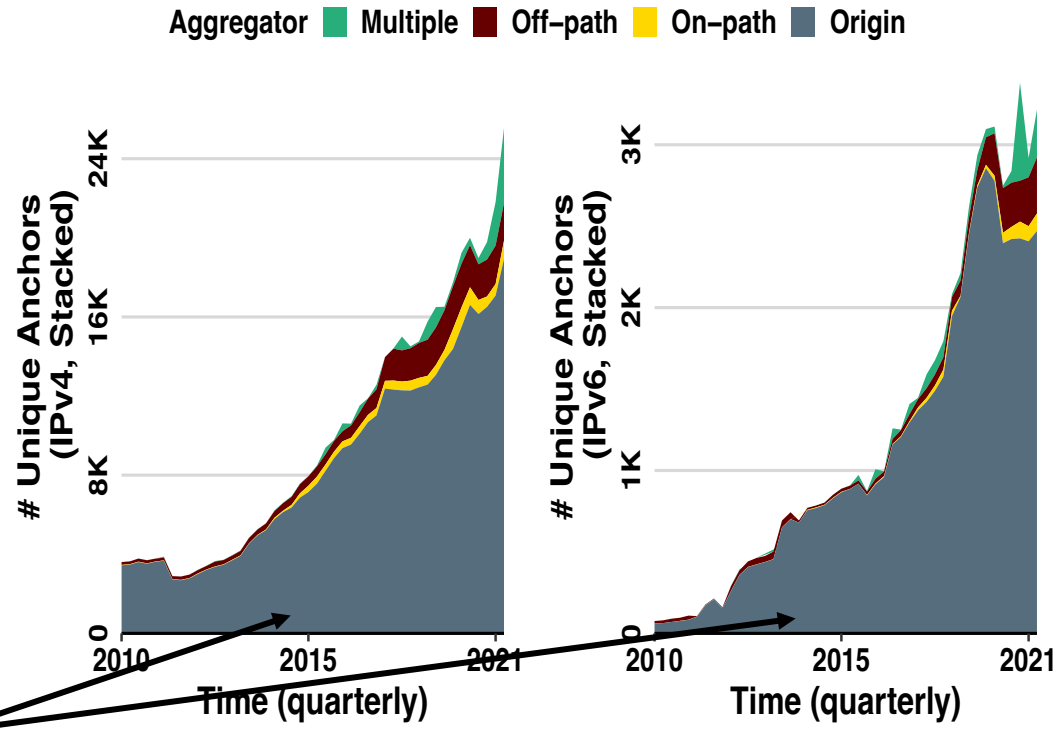

majority of HSPs are aggregated at the origin – BGP confederation

How Far HSPs Propagate?

IPv6 HSPs have better visibility than IPv4 HSPs

Most of HSPs are visible on less than 10 peers

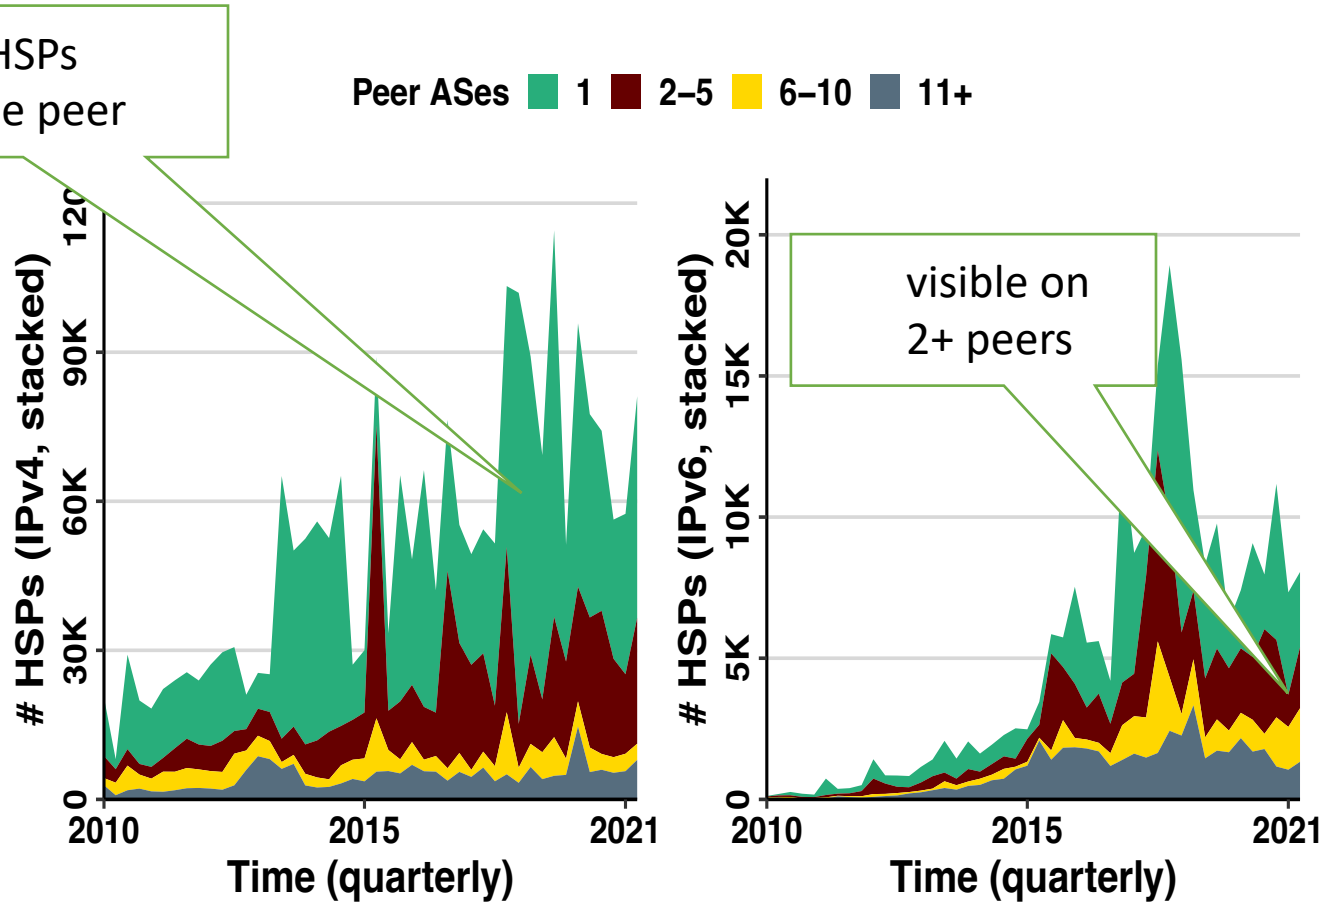

HSP Anchors in Various Datasets

Observations:

- Current RC infrastrucure misses 1/3 of anchors potentially contain HSP
- less noisy, linear increase in the number of anchor prefix for which HSPs

- Aggregated class only contains on-path aggregated anchor prefixes

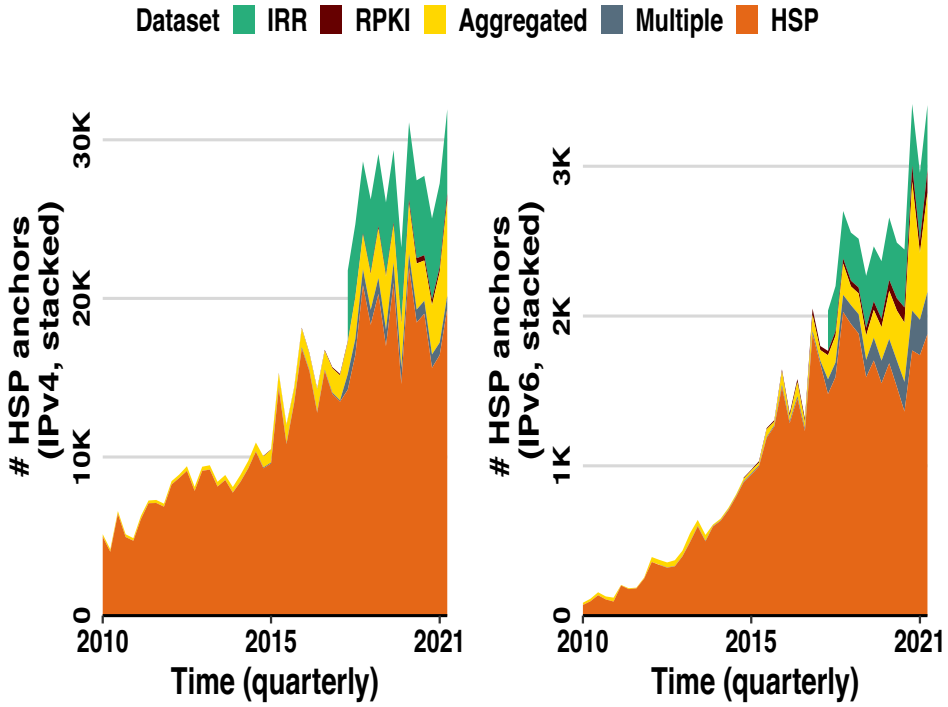

HSP Originators Across Datasets

Observations

- HSP origins has more than doubled for IPv4
- For IPv6, the growth rate of more than 25x
- little overlap between the individual data sets

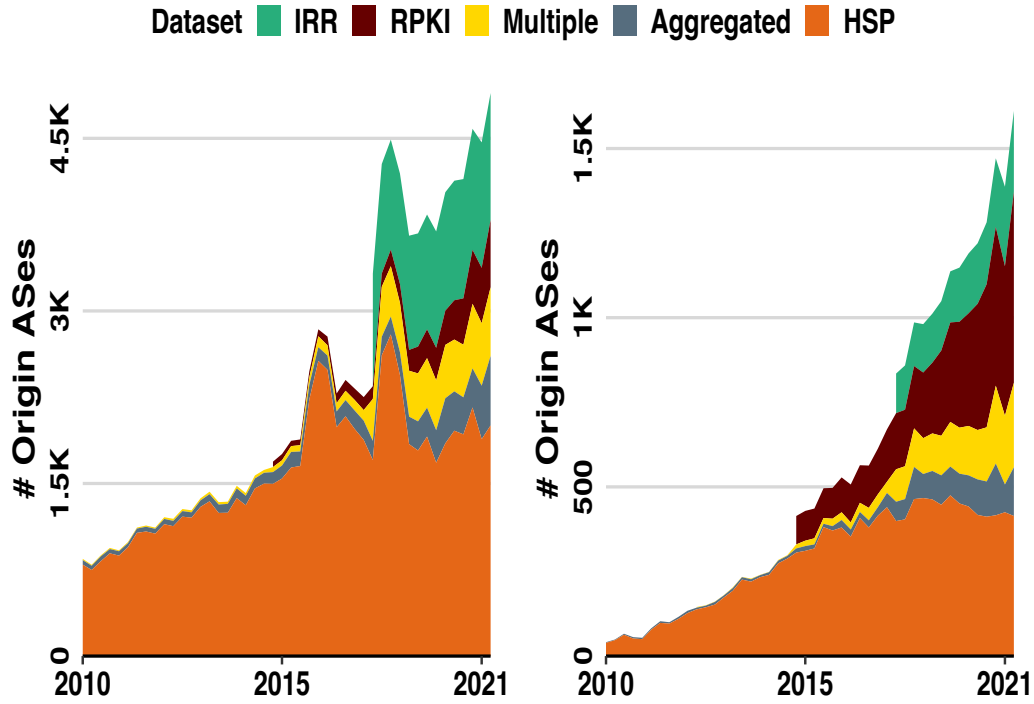

Methodology

- Route Collectors‘ Data
- 11+ years (2010-2021)
- BGP RIBs + updates
- From 3 Projects

Passive Measurement

- IRRs Snapshots
- RPKI Snapshots
- AS Relationships Inferences
- AS Classification Inferences
- ASDB

Supplemental data sets

- Advertise our own HSPs to the Internet and contuct experiment.

Active Measurement

Cleaning Noisy Data

Rule1:

- Misconfigured Peer ASes
- Abnormal Prefixes
- Private IP ranges
- Private Origin ASes
- Multicast and IPv4 class E

Rule2:

Testable HSP

- For all HSPs, check if it was announced via a route that crossed at least one additional AS then “testable“.

HSP Propagation Pattern

We use:

AS triplets (three consecutive ASes)

AS Relatship Inferences of CAIDA

- No single occurrence of P2P relationships
  - ASes strongly filter the routes they send to peers
- for IPv4 almost all ASes redistribute HSPs “upwards”
  - Customers pay their providers to reannounce their prefixes

HSPs are only propagated “vertically” and never “horizontally”.

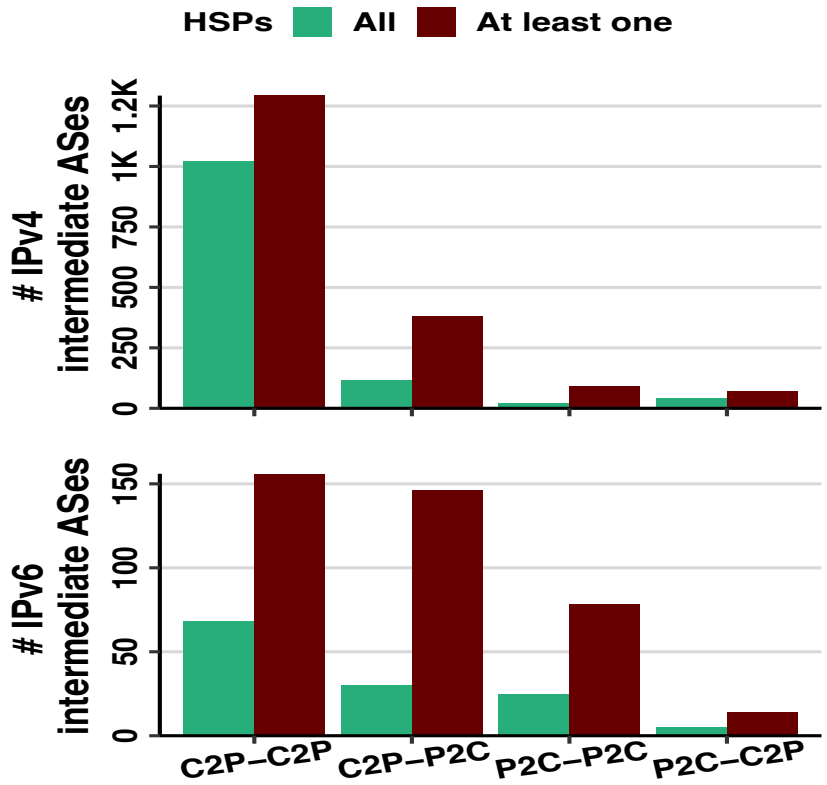

Route Collector Data

For our analysis we utilize “snapshots“ from the RC projects Isolario , RIPE RIS , and Routeviews

- From Jan.2010 to October.2021
- Quarterly, 7days per quarter
- BGP RIBs – every 24 hours
- BGP Updates – every 5 mins

seven-day window allows us to achieve a consistency of 97 % and 98 % for IPv4 and IPv6, respectively.

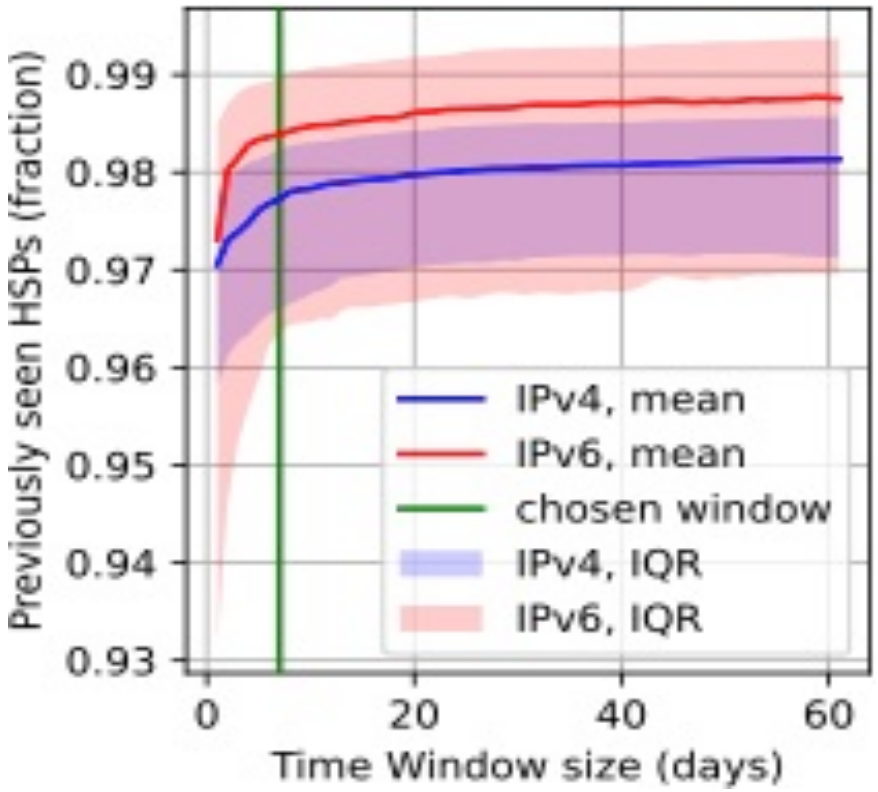

Real World Experimentaiton

The PEERING testbed

- 180 IPv4 and 152 IPv6 neighboring ASes
- 8 IPv4 and 9 IPv6 neighboring ASes redistributed HSPs

Used Prefixes

- IPv4:184.164.240.0/23
- IPv6:2804:269c:4::/46

RIPE Atlas probes

- To maximize AS coverage - one probe per AS
- prefer dual-stack probes
- Highest stable

Experimemt design

- announce HSP and anchors
- wait convergence
- run paris-traceroutes from all probes
  - simultaneously issue ICMP, TCP, and UDP probing
- withdraw prefixes
- map traceroutes to AS Paths using bdrmapit

How Far HSPs Propagate?

We did experiment by advertising anchor + HSPs to the Internet

- conduct traceroute from probes
- check it in RC‘s peer ASes

Current RC‘s infrastructure underestimates data plane reachibility

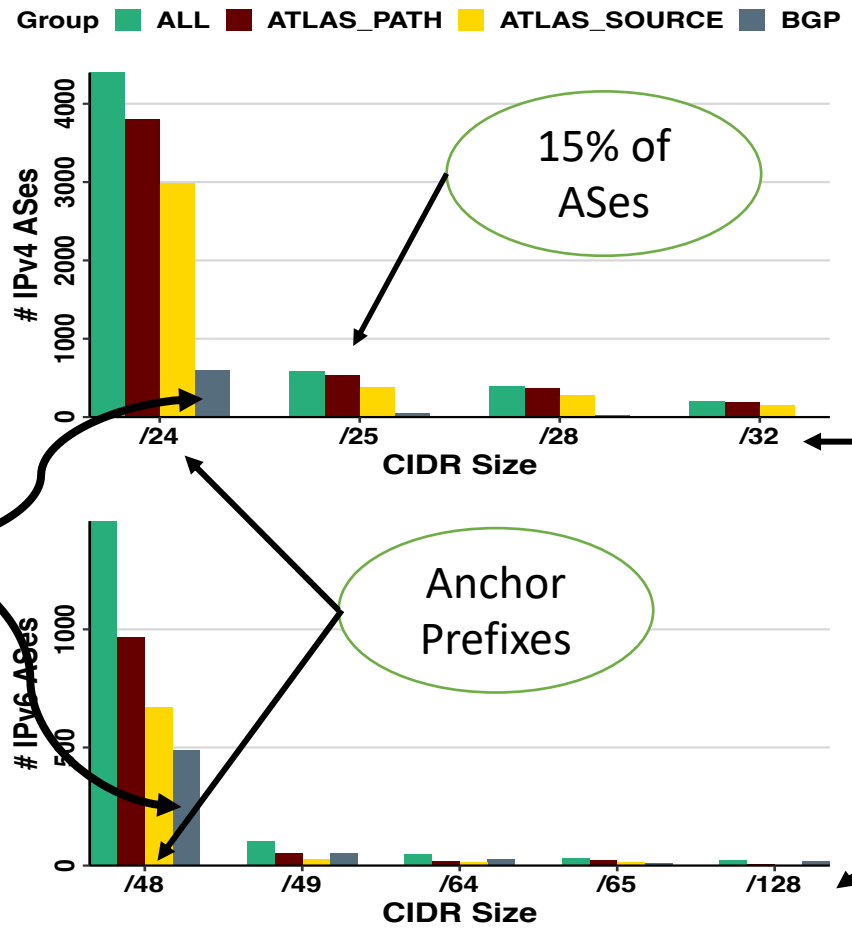

The more specific the CIDR size, the less propagation chances.