TTM AS-level Traceroutes

Matching IPs to ASes

René Wilhelm

New Projects Group
RIPE NCC
<wilhelm@ripe.net>

Motivation

- TTM performs frequent traceroutes to find closest IP route for delay measurements
- Many small changes in routes, due to load balancing and rerouting inside provider networks
  - minimal changes in delays
- Want a handle on bigger changes in routing, when different upstream or backbone provider is used

→ transform collected IP traces to AS level traces

Traceroute from TBx to TBy

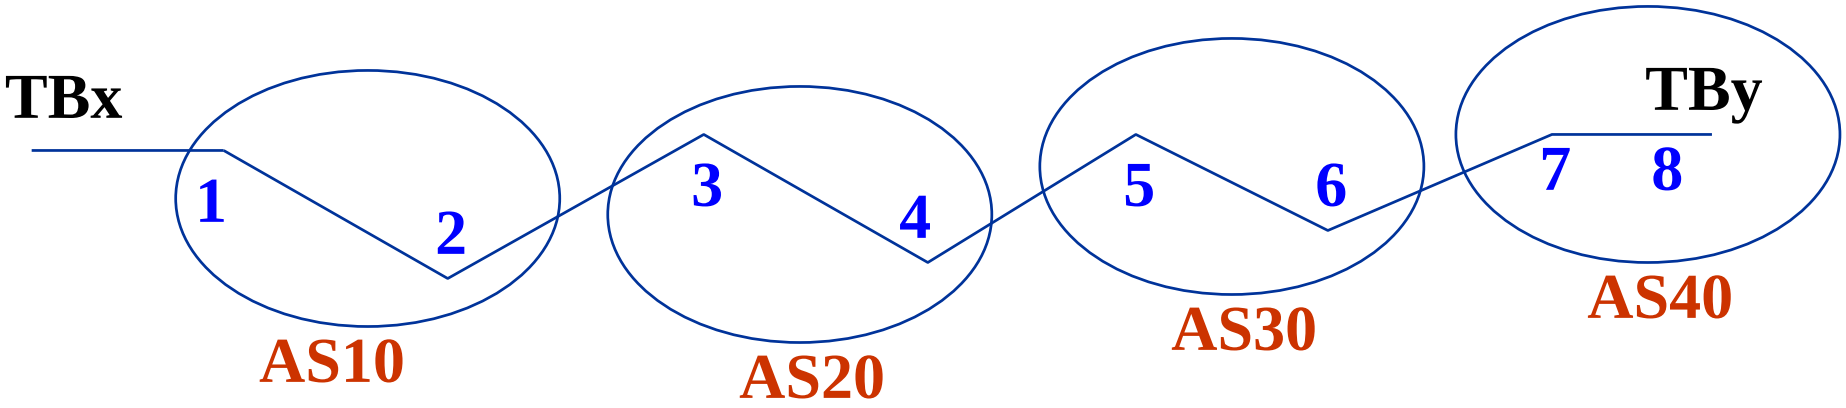

8 Hops recorded in the traceroute

4 Autonomous systems traversed

Matching IPs to ASnums

- Internet Routing Registry
  - Autonomous systems register the prefixes they originate
  - find best match
- Global Routing table
  - prefixes and aspaths, last as in path is seen as originator
  - find best match
- Currently, TTM uses IRR
  - speed, matching done offline, traceroutes on boxes simple
  - general feeling IRR not up to date; check and quantify

Internet Routing Registry

- No single authoritative registry:

Internet Routing Registry (2)

- Not practical to query everyone of them
  - time consuming, remote locations, slow connections
- Compromise:
  - whois.ripe.net mirrors RADB, APNIC, ARIN, CW, VERIO
  - local to TTM analysis machine, fast response
  - retrieve all route objects matching the IP
  - find longest prefix match, return AS number of the object
- May return more than 1 AS number
  - route registered in multiple registries with different AS
  - route registered more than once in same registry

Global Routing Table

- Not one single global routing table
  - aggregation and filtering lead to different views at different locations
- RIPE NCC ‘s RIS project
  - collects BGP updates at 10 locations, from 302 peers
  - daily dumps of Routing Table stored in per collector databases
  - accessible from one central machine

RIS Routing Tables

- Join data from all RIS collectors in one table
  - prefix + origin AS
  - 132995 prefixes (july 29th)
- Map IP to longest matching prefix
- Can return more than 1 AS number
  - multiple origin AS in one route collector
  - different origin AS in different route collectors (e.g. due to aggregation)

Results for TTM

- Typical day in the life of TTM
  - 63 active boxes, 6 traces per hour to each other box
  - 13,322,691 IP addresses, but many duplicates
  - 3618 unique IP addresses to match with AS

- IRR: 2856 IPs in 251 prefixes (79%) matched
- RIS: 3584 IPs in 297 prefixes (99%) matched

- RIS does a better job, but look at the differences

IRR vs. RIS

- 51 prefixes not found by IRR are in RIS
  - 4 of these have multiple origin AS in RIS
    check ARIN,RIPE,APNIC database → exchange points
  - 80% of missing entries are with 5 large ISPs
- 17 prefixes have different AS in RIS and IRR
  - objects not updated after mergers, prefixes once announced by 2 or more ASes, now by single AS
  - objects not maintained, outdated
- 11 prefixes represented by aggregates
  - either RIS or IRR has a more specific, with different AS

IRR vs RIS (2)

- 5 prefixes not found by RIS are in Routing Registry
  - exchange points
- 9 prefixes multiple objects in IRR, more than 1 AS
  - only 1 AS in RIS → outdated objects?
- 17 IPs (from 7 /24s) not found in RIS nor in IRR
  - RFC1918 addresses
  - exchange points
  - internal infrastructure

RIS vs. IRR matches: summary

- For TTM traceroutes IRR is ~80% correct
- Value of IRR would increase If large ISPs would register and maintain route objects
- Currently, routing tables are the best approach.
  - RIS route collectors provide a publicly available view from different vantage points (US, Europe, Japan) with 302 total peers

TTM AS traces vs. BGP aspath

- Two clear differences
  - exchange points
    - traceroute detects IP of interface on exchange, if a match is found it maps to the exchange AS, not the peer AS
      e.g. from RIPE NCC: 3333 1200 1103 …
  - unknown AS in the traceroute aspath
    - 1 or more consecutive IP hops not matched or not responding
    - can't tell if they belong to preceding AS, next AS or a different AS altogether; flag it with AS number 0.
- In depth comparison planned
  - for sites which both host a testbox and peer with RIS

Conclusions & Future

- RIS provides a good means to match IP to AS
- IRR could use better commitment by ISPs

- TTM will switch to using RIS for IP-AS mapping in the traceroute database
- Expand code to also handle IPv6

- RIPE-NCC will set up an IP-AS mapping service
  - Derived product: traceroute -A with AS from RIS

Questions / Discussion