Lessons Learned From Using the RIPE Atlas Platform for Measurement Research

RIPE 68, Warsaw

Vaibhav Bajpai, Steffie Jacob Eravuchira, Jürgen Schönwälder
(v.bajpai|s.eravuchira|j.schoenwaelder)@jacobs-university.de

Computer Networks and Distributed Systems Group,
Jacobs University Bremen,
Bremen, Germany

May 2014

Leone Project: leone-project.eu

Background

- Partners within the Leone Project. The goal is to assess the QoE of end-users through active measurements[1]. Primarily developing new metrics for the SamKnows[2] platform.
- Started actively using the RIPE Atlas platform since the RIPE Atlas API was publicly released in 2013. Close collaboration with Daniel Karrenberg, Philip Homburg (RIPE NCC).
- This talk is a subset of a larger measurement study using the RIPE Atlas platform (recently submitted to IMC 2014). We share our experience in using the platform during this journey.

[1] ripe66.ripe.net/archives/video/1259
[2] samknows.com

Rate Limits

We have ample credits. Why can't we provision measurements?

Rate limits are setup on each account by default:

1. No more than 100 simultaneous measurements.
2. No more than 500 probes per measurement.
3. No more than $270K$ credits may be used each day.

- Although documented[3], may not be well-known.
- Limits can be lifted off by proposing the research on `atlas` mailing list[4]

---

[3] https://atlas.ripe.net/docs/udm

[4] Thanks to Vesna Manojlovic (RIPE NCC) for lifting off rate limits on our user accounts.

Probe Calibration

Firmware Variants

"Lack of calibration can lead to uncertainity of results - " [1]

- Frequency of firmware releases has increased since 2013.

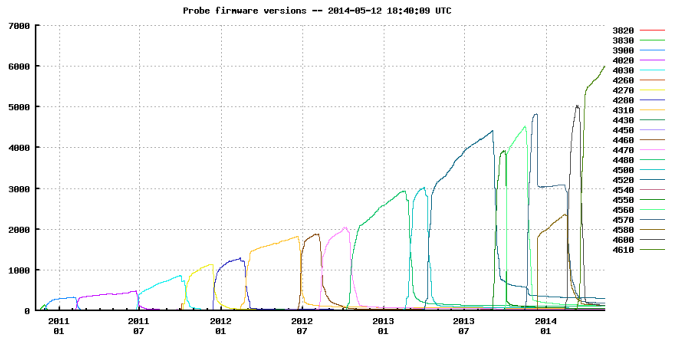

Firmware release cycle since 2011 (as of May 2014): atlas.ripe.net/results/graphs.

- Each User-Defined Measurement (UDM) tags the firmware version of the probe when reporting results.

```
{
  "prb_id": 10678,
  "type": "traceroute"
  "fw": 4560,
  ...
}
```

A snippet of a `traceroute` measurement result from a probe (as of November 2013).

Probe Calibration

Hardware Variants

- Three hardware revisions of deployed probes: v1, v2, v3
- v3 probes are more hardware capable[5] than v1/v2 probes.

8162 probes are deployed (as of May 2014): atlas.ripe.net/results/maps/network-coverage.

- Anchors are dedicated servers.
- Anchors are sources/sinks of regional measurement traffic.

56 anchors are deployed (as of May 2014): atlas.ripe.net/anchors/map.

[5] In terms of specifications

Probe Calibration

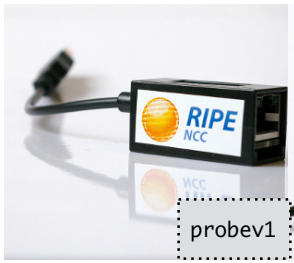

PROBE ID	HARDWARE REVISION	FIRMWARE VERSION	HARDWARE	WEBPAGE
[1, 1521)	probev1	4570	Lantronix XPort Pro	probev1.ripe.net
(2000, 5000)	probev2	4570	Lantronix XPort Pro	probev2.ripe.net
(10000, $+\infty$)	probev3	4580	TP-Link TL-MR3020	probev3.ripe.net
(6000, 6018)	anchorv1	-	Dell PowerEdge	-
(6018, 7000)	anchorv2	-	Soekris Net6501-70	anchorv2.ripe.net

Generated as of December 2013.

- Firmwares are kept in sync across hardware revisions.
- ProbeID can reveal the hardware revision[6].

[6] Thanks to Robert Kisteleki (RIPE NCC) for revealing the probe ID to hardware revision mapping.

Hardware Matters

Does probe hardware revision have effect on measurement results?

- Find probes whose:
  1. Hop #1 is private [2] and
  2. Hop #2 is public.
- This ensures:
  1. Probe does not cross a wireless link[7].
  2. Probe wired directly behind a residential gateway.
- Provisioned `traceroute` measurements that lasted for a day.
- Investigated latencies to hop #1 from v2/v3 probes.

---

[7] probes themselves cannot associate to a wireless access point.

Hardware Matters

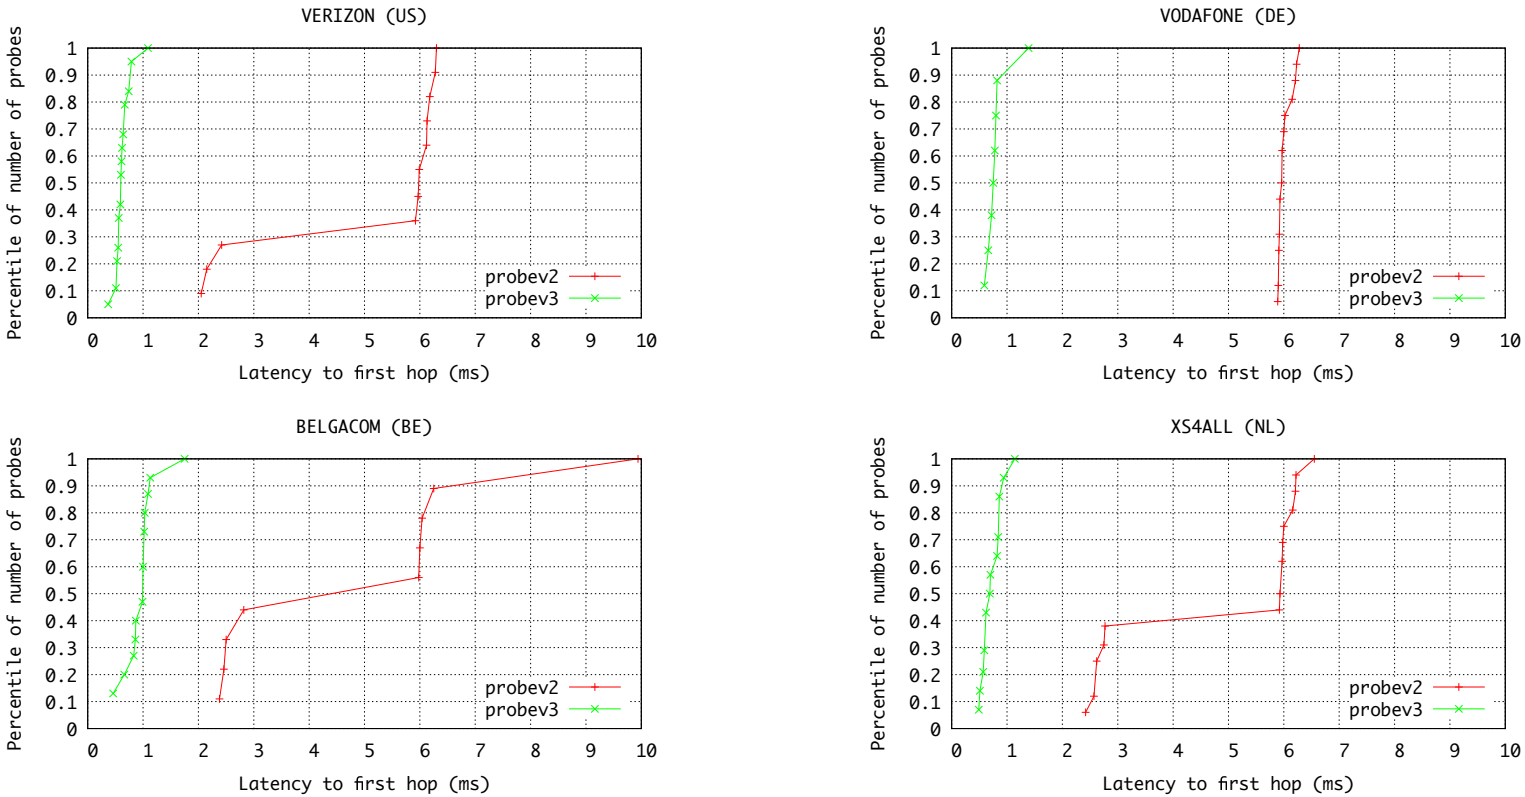

Generated as of November 2013 (RIPE Atlas)

- v3 probes show expected $<$ 1ms latencies to hop #1.
- v2 probes show surprisingly high latencies around 6ms to hop #1.

Hardware Matters

We next investigated RIPE Atlas source code[8].

- Measurement tools are adaptation of `busybox` utilities.
- Measurements modified to run in an evented manner using `libevent`.
- Measurements do not spawn new processes (instead invoked as function calls).
- Circumvents absence of MMU in v1/v2 probes.
- A single evented loop handles all measurement requests.

[8] https://atlas.ripe.net/get-involved/source-code

Hardware Matters

```
static struct trtbase *traceroute_base_new (
  struct event_base *event_base
) {
...
event_assign(&base->event4, base->event_base,
  base->v4icmp_rcv, EV_READ | EV_PERSIST,
  ready_callback4, base);
}

static void ready_callback4 (
  int __attribute((unused)) unused,
  const short __attribute((unused)) event,
  void *s
) {
...
ms=(now.tv_sec-state->xmit_time.tv_sec)*1000 +
   (now.tv_usec-state->xmit_time.tv_usec)/1e3;
}
```

`traceroute` code snippet from v4570 running on v2 probes as of November 2013.

- RTT time-stamping is performed in user-space in the evented callback.
- A probe loaded with multiple measurements will witness time-stamping delays.
- Delays more pronounced on probes with hardware constraints: v1/v2.

Proper Statistics Matter

Can per-hop averaging of RTT significantly vary the results?

RIPE Atlas:

1. `evtraceroute` sends 3 ICMP queries (default) to each hop.
2. RTT from each ICMP response is separately made available.

SamKnows:

1. `mtr` sends 3 ICMP queries (default) to each hop.
2. RTT are averaged over each hop[9].

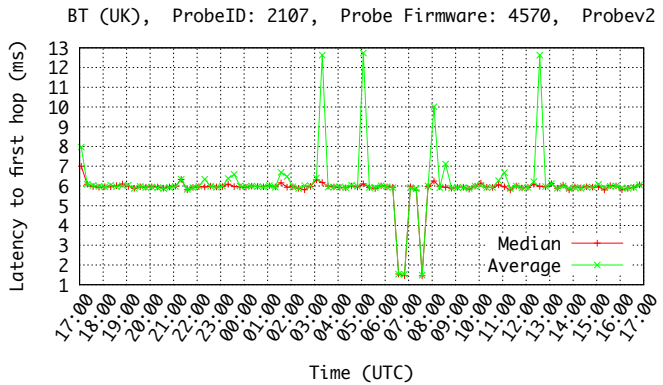

Generated as of November 2013 (RIPE Atlas)

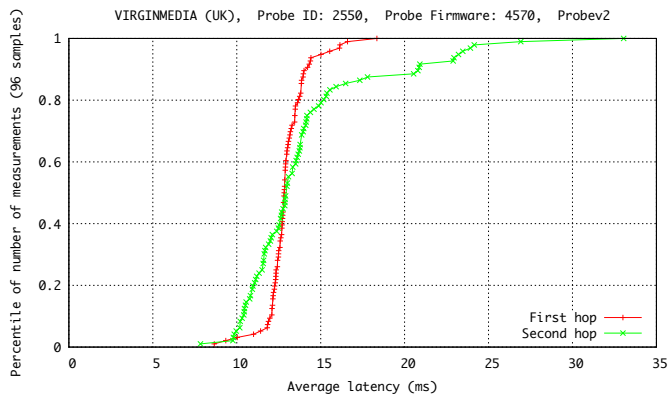

Generated as of November 2013 (RIPE Atlas)

[9] We replaced the `traceroute` test within SamKnows to expose results without aggregation.

Heavy-Tailed AS-based Probe Distribution

- 56% (2307) of probes fall within AS ranks <= 200.
- ASes with Rank > 200 have less than 10 probes.

- AS with Rank #1 contributes 4% of all probes.
- ASes with Rank >= 10 contribute 19% of all probes.

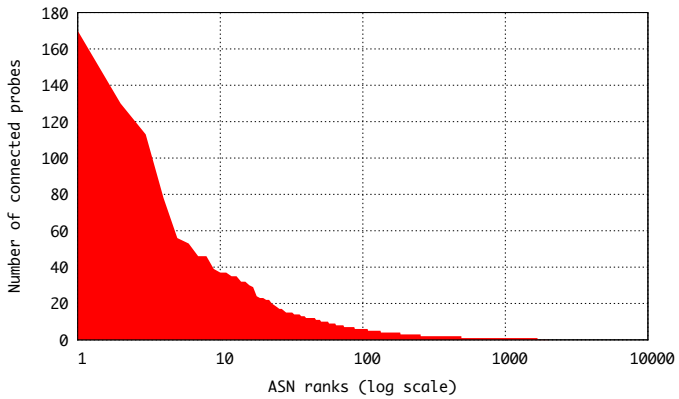

Distribution of public, connected and non-anchored probes (4133) as of October 2013

AS Rank	AS (ASN)	Connected
01	COMCAST (7922)	170
02	DTAG (3320)	130
03	LGI-UPC (6830)	113
04	PROXAD (12322)	78
05	ZIGGO (9143)	56
06	XS4ALL (3265)	53
07	VIRGINMEDIA (5089)	46
08	UUNET (701)	46
09	KABELDEUTCHLAND (31334)	39
10	UNITYMEDIA (20825)	37

Generated as of October 2013.

Metadata is (Changing) Data

How is the probe connected to the Internet?

- The connection speed
- Network Type: (Core, Research, IXP, Access, Home) Network.
- WAN Type: DSL, Cable, Fibre to the X.

Ability to track changes and API access to metadata history.

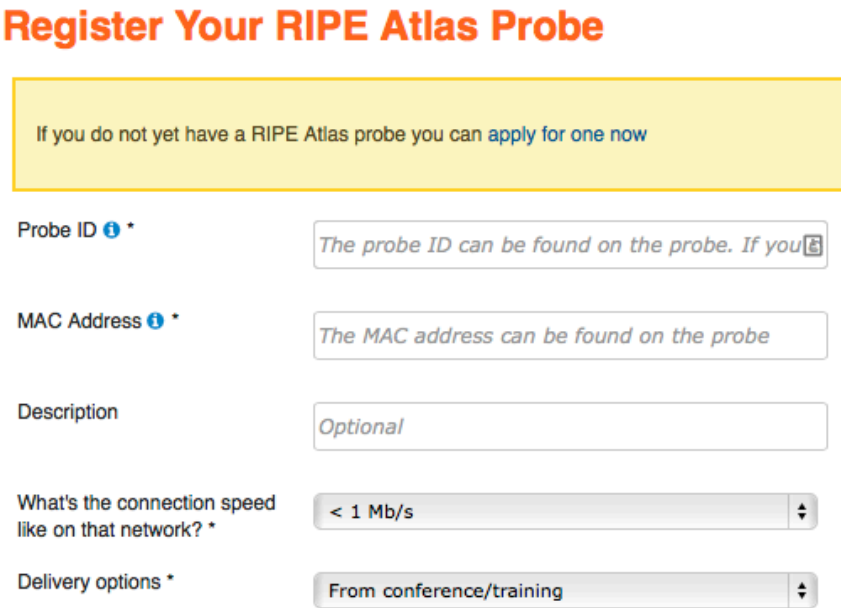

Probe registration page: atlas.ripe.net/register

Conclusion

1. Rate Limits
2. Probe Calibration
3. Hardware Matters
4. Proper Statistics Matters
5. Heavy-tailed AS-based Probe Distribution
6. Metadata is Data

References

[1] P. Eardley, M. Mellia, J. Ott, J. Schönwälder, and H. Schulzrinne, "Global Measurement Framework (Dagstuhl Seminar 13472)," *Dagstuhl Reports*, vol. 3, no. 11, pp. 144–153, 2014. [Online]. Available: http://drops.dagstuhl.de/opus/volltexte/2014/4440

[2] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear, "Address Allocation for Private Internets," RFC 1918 (Best Current Practice), Internet Engineering Task Force, Feb. 1996, updated by RFC 6761. [Online]. Available: http://www.ietf.org/rfc/rfc1918.txt