

On IPv4 and IPv6 Routing Stability

Ioana Livadariu, Ahmed Elmokashfi, Amogh Dhamdhere



[**simula** . research laboratory]

Motivation

- IPv4 addresses are running out:
 - Feb 2011: IANA allocated last /8
 - 3 RIRs (APNIC, RIPE, LACNIC) are allocating from the last /8




Figure no.1 : IPv4 Address Report *

- Problem: Slow migration to IPv6
- *Look at performance by analyzing routing stability*




Figure no.2 : IPv6 client-side adoption**

* RIR IPv4 Address Run-Down Model [<http://www.potaroo.net/tools/ipv4/>]

** IPv6 Statistics [<http://www.google.com/intl/en/ipv6/statistics.html>]

Measuring Routing Stability : metrics

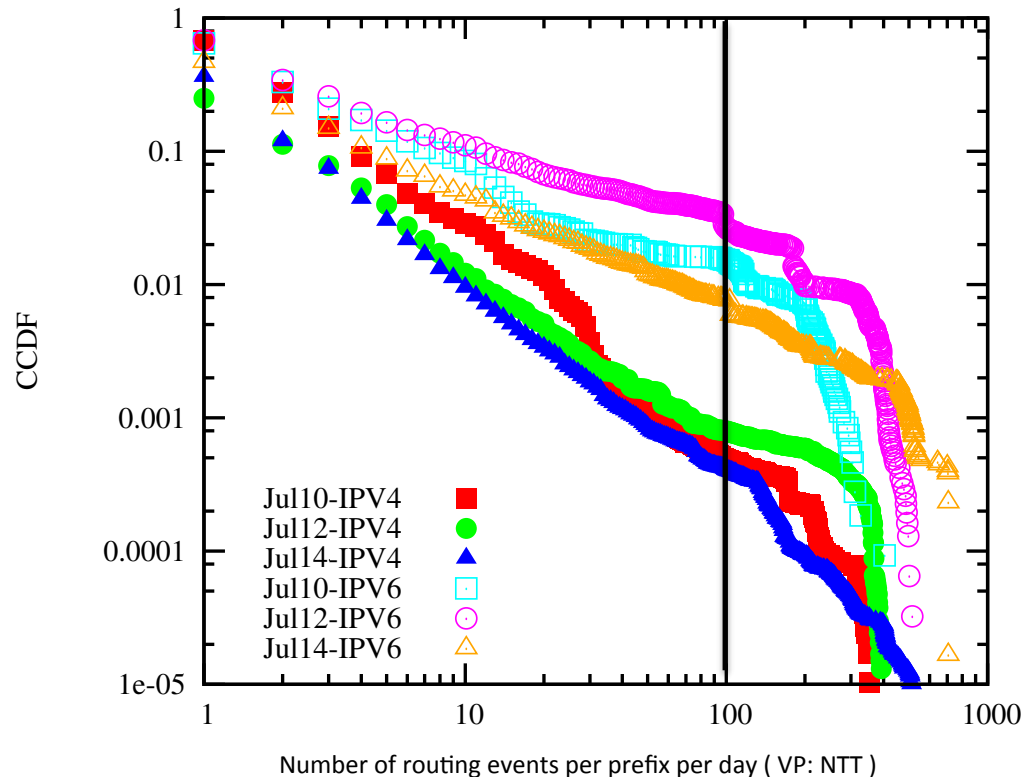
Control plane (BGP updates at 5 dual-stacked RV monitors)

- *Frequency of routing changes towards IPv4 and IPv6 prefixes*
- *Correlation of the IPv4 and IPv6 routing events*

Data plane (9 ARK monitors to probe dual-stacked targets)

- *Reachability of the dual-stacked probed targets*
- *RTT variations*


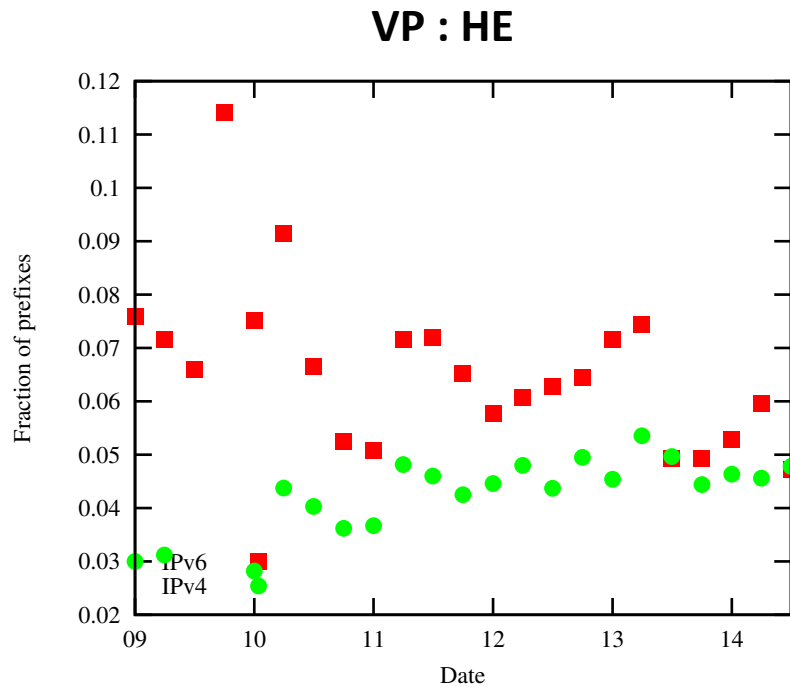
IPv6 routing system exhibits more routing changes than IPv4



- Identified IPv4 and IPv6 routing events
- 0.1% of the IPv4 versus 2% of the IPv6 prefixes experience more than 100 events per day

Zooming in : Active prefixes

- *Active prefix* : a prefix that experiences a routing change at least once per day

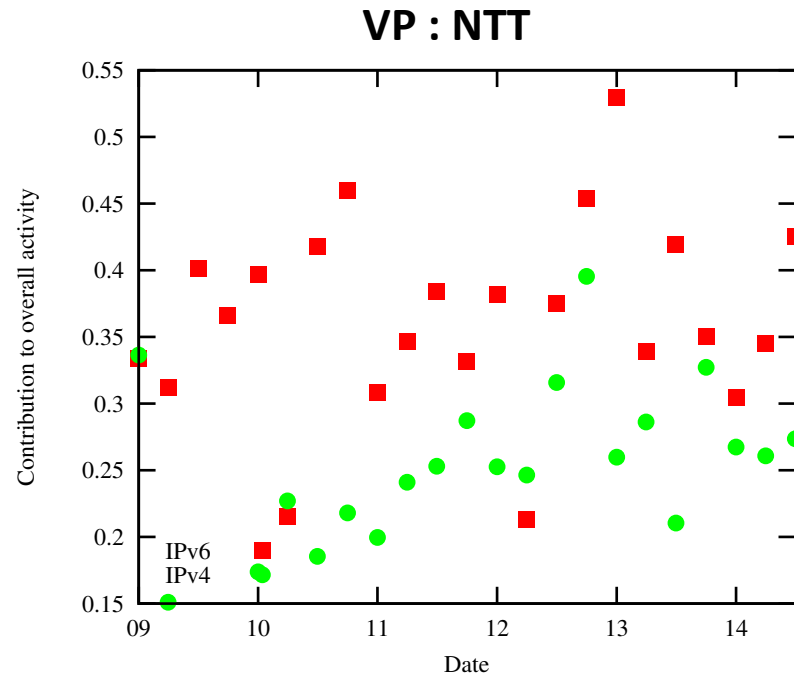
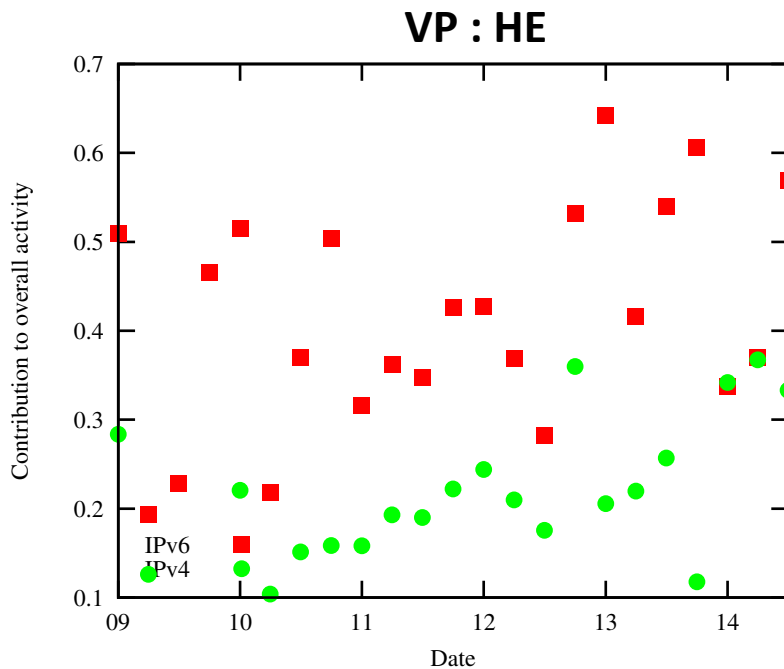


Average fraction of daily active prefixes

- The overall fraction of active prefixes is higher in IPv6 than IPv4

Zooming in (more) : Highly active prefixes

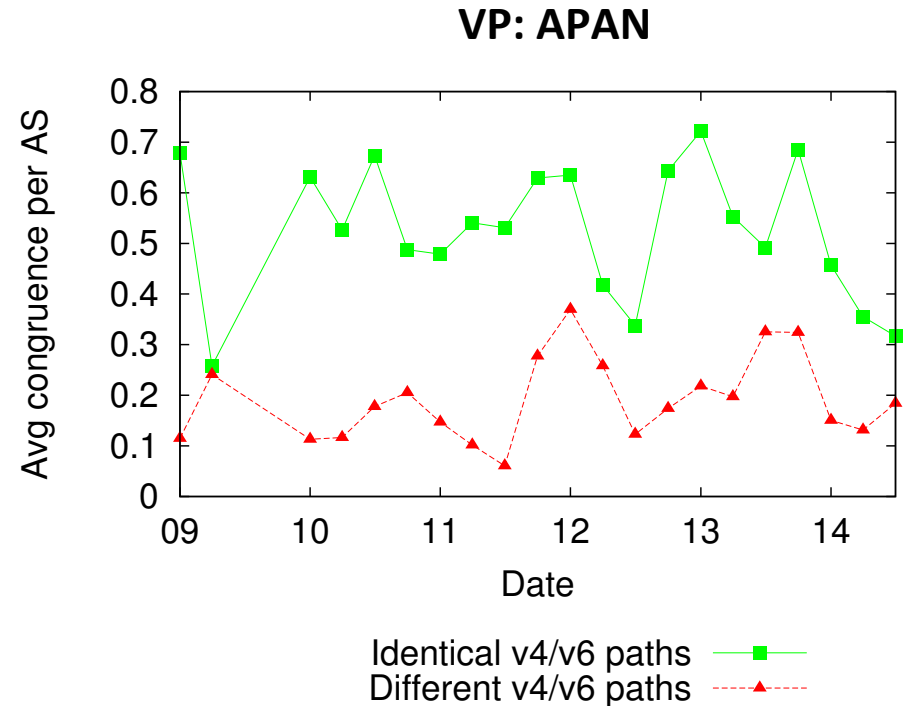
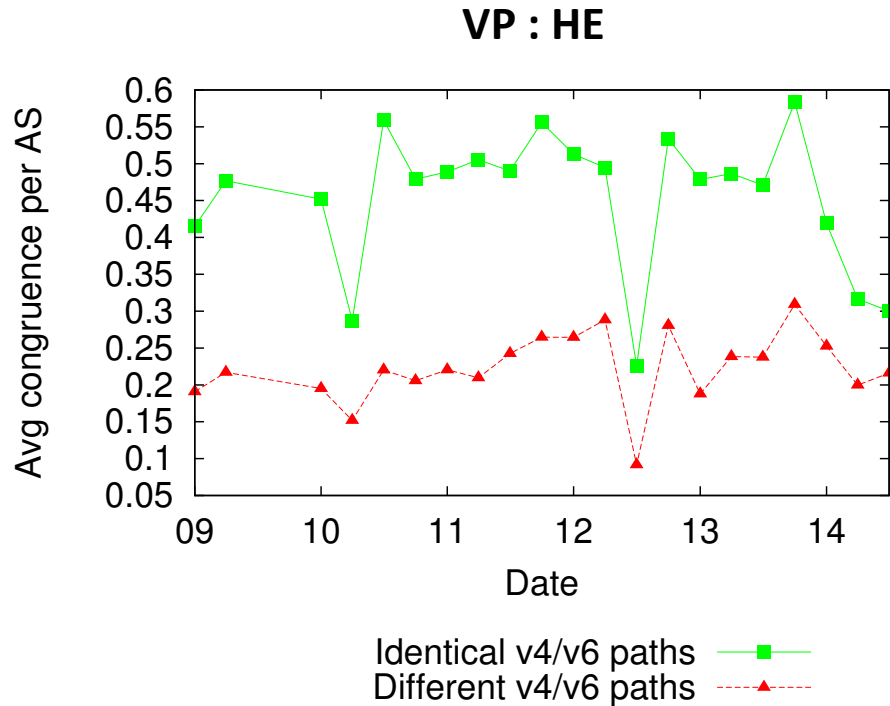
- *Highly active prefixes* : top 1% of the active prefixes in terms of contribution to the BGP dynamics



The contribution of the top 1% active prefixes to the BGP dynamics

- The highly active prefixes are responsible for between 40-50% of IPv6 updates compared to 20-30% of IPv4 updates

Correlating IPv4 and IPv6 routing events



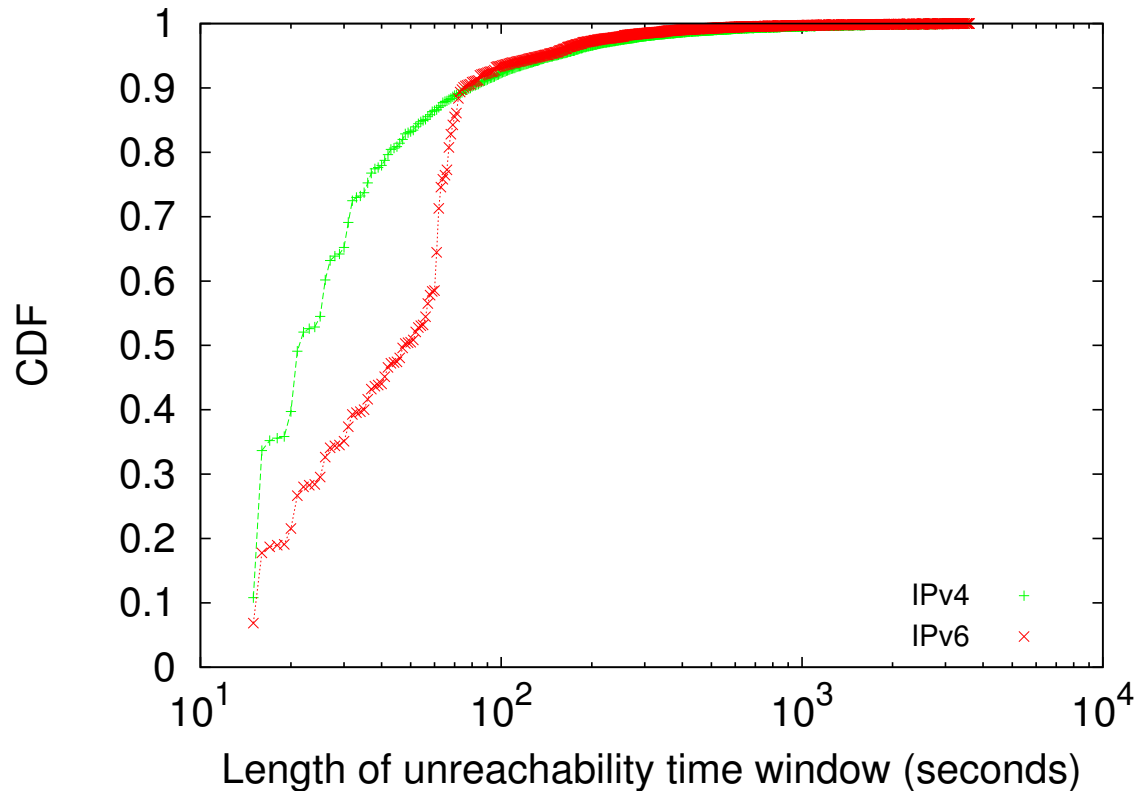
- Computed the fraction of overlapping IPv6 and IPv4 events
- Higher correlation for identical paths than for different paths

Approach to study Data plane Stability

- *Goal* : Study **reachability** and **performance**
- *Measurement setup*: Use nine monitors from the ARK infrastructure to ping dual-stacked targets* every 5 seconds (~105 targets per monitor) for 1 ½ months (August – September 2014);
- *Limitations* : Data could be influenced by the availability of the webservers and random losses;
- *Data Filtering* : Filter out very long or very short periods of unreachability([15s , 3600s])

* - Alexa [http://www.alexa.com/topsites]


Network reachability over IPv4 than IPv6



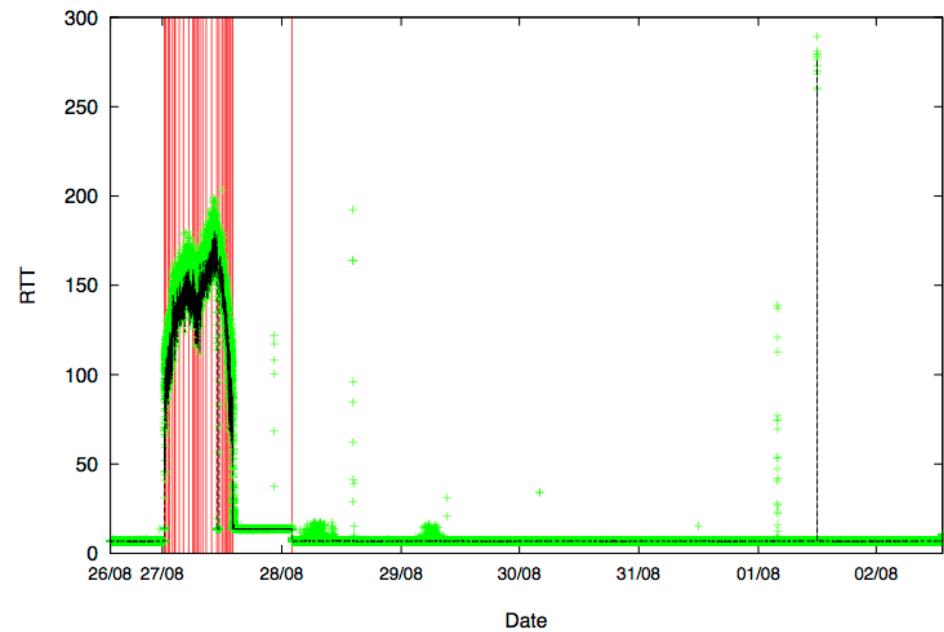
- Longer unreachability intervals over IPv6 than IPv4
- Paired measurements : The difference in unreachability towards the same AS can be up to 15%

Performance : Similar RTT time series

IPv4




IPv6




Monitor : ams-nl; Target : AS 197043

Performance : Discrepancies in the RTT time series

IPv4



IPv6



Monitor : ams-nl; Target : AS 15982

Conclusions & Future work

Conclusions (so far)

- Control plane:
 - IPv6 prefixes are less stable than IPv4
 - Most IPv6 routing dynamics are generated by a small fraction of pathologically unstable prefixes
- Data plane:
 - IPv6 unreachability intervals longer and more frequently than IPv4 unreachability intervals

Ongoing work

- Use the same measurement setup to collect traces of the paths between the ARK monitors and the probed webservers
- Correlate the control and data plane events