RIPE TTM
and divining IPv6 routing policies

Dave Wilson
DW238-RIPE

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Agenda

- Need for a routing policy
- Weaknesses compared to IPv4
- Observations and adjustments
- Conclusion

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Background – inet4

"Just pick the shortest route"
isn't good enough

- IXes and private peerings
- Research networks – Géant, Abilene
- Traffic engineering, transcontinental peerings

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Background – inet6

Problem is even more pronounced on IPv6 internet

- Tunnels and STM-16s look the same
- Some networks production, some pre-production, some testing
- No such thing as a Tier-1

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Forming a policy

Mechanisms by which one can form a routing policy

- Routing policies documented in the DBs
- Often clear distinctions between European and American ASes
- BGP Community policies well defined

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Forming a policy – inet6

These mechanisms dont exist
in the IPv6 internet

- 6bone DB not adequate, RIPE DB does not have v6 routing policy support yet
- Tunnels and software forwarding
- Iffy BGP community support

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Forming a policy – inet6

So what can we use?

- Lots of looking glasses
- WHOIS DBs, Google and mail account managers for BGP community support
- RIPE TTM project...

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Our main peers

- Géant – EU academic internet
  strict filtering, also provides transit
- Abilene – US academic internet
  no filtering, provides transit
- Tunnel to Global Crossing
  no significant filtering, provides transit
- There are others

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

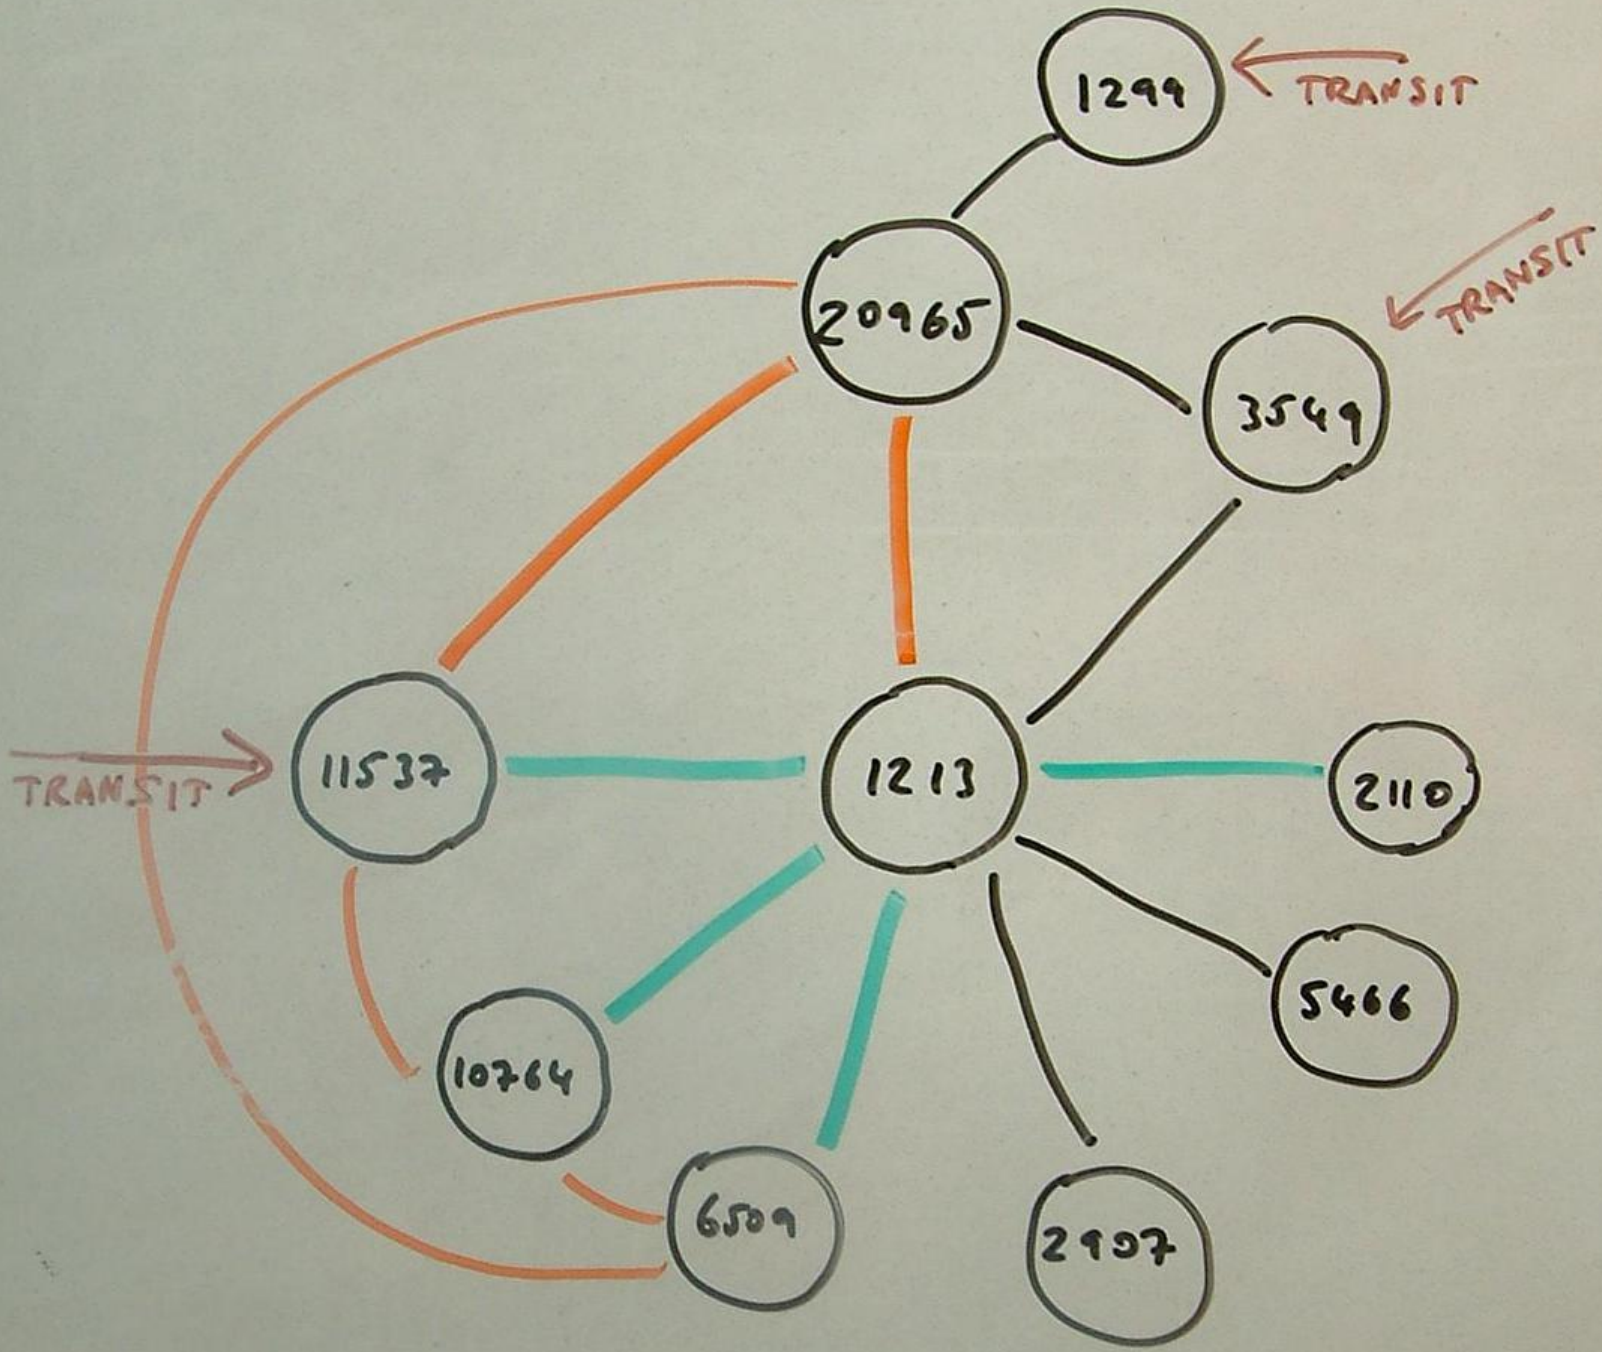
1299
TRANSIT
20965
TRANSIT
3549
TRANSIT
11537
1213
2110
5466
10764
6509
2907

IPv6 outbound traffic

Local preferences:

- 250: EU peer, native, hardware
- 240: US peer, native, hardware
- 150: EU peer, native, software
- 140: US peer, native, software
- 100: tunnel
  - ➔ including routes marked tunneled as such by Abilene/Géant

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

IPv6 inbound traffic

- Use AS-path stuffing to steer traffic toward better links
- EU/US peers, native links, hardware forwarding
  - ➔ do not prepend
- EU/US peers, native links, software forwarding
  - ➔ prepend 1213 1213
- Tunnels
  - ➔ prepend 1213 1213 1213

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

HEAnet->Géant

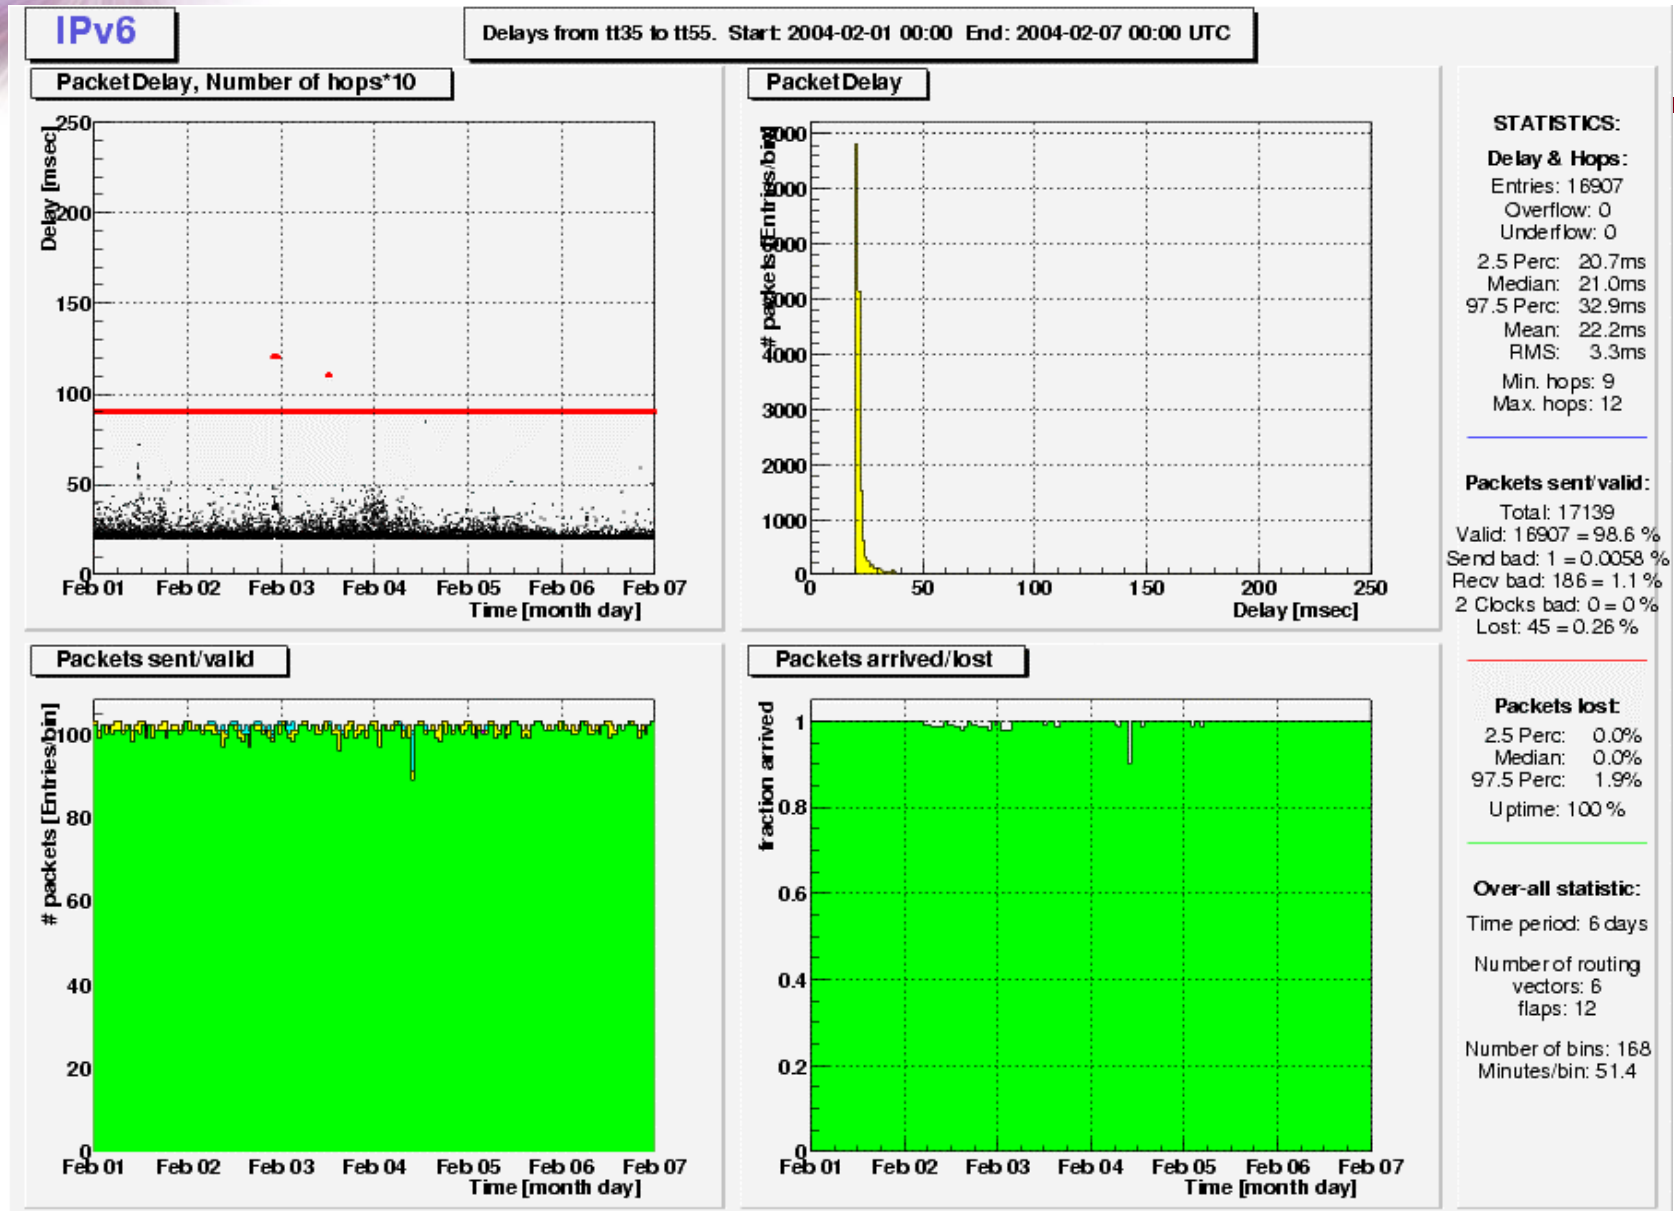

HEAnet->Tunnel

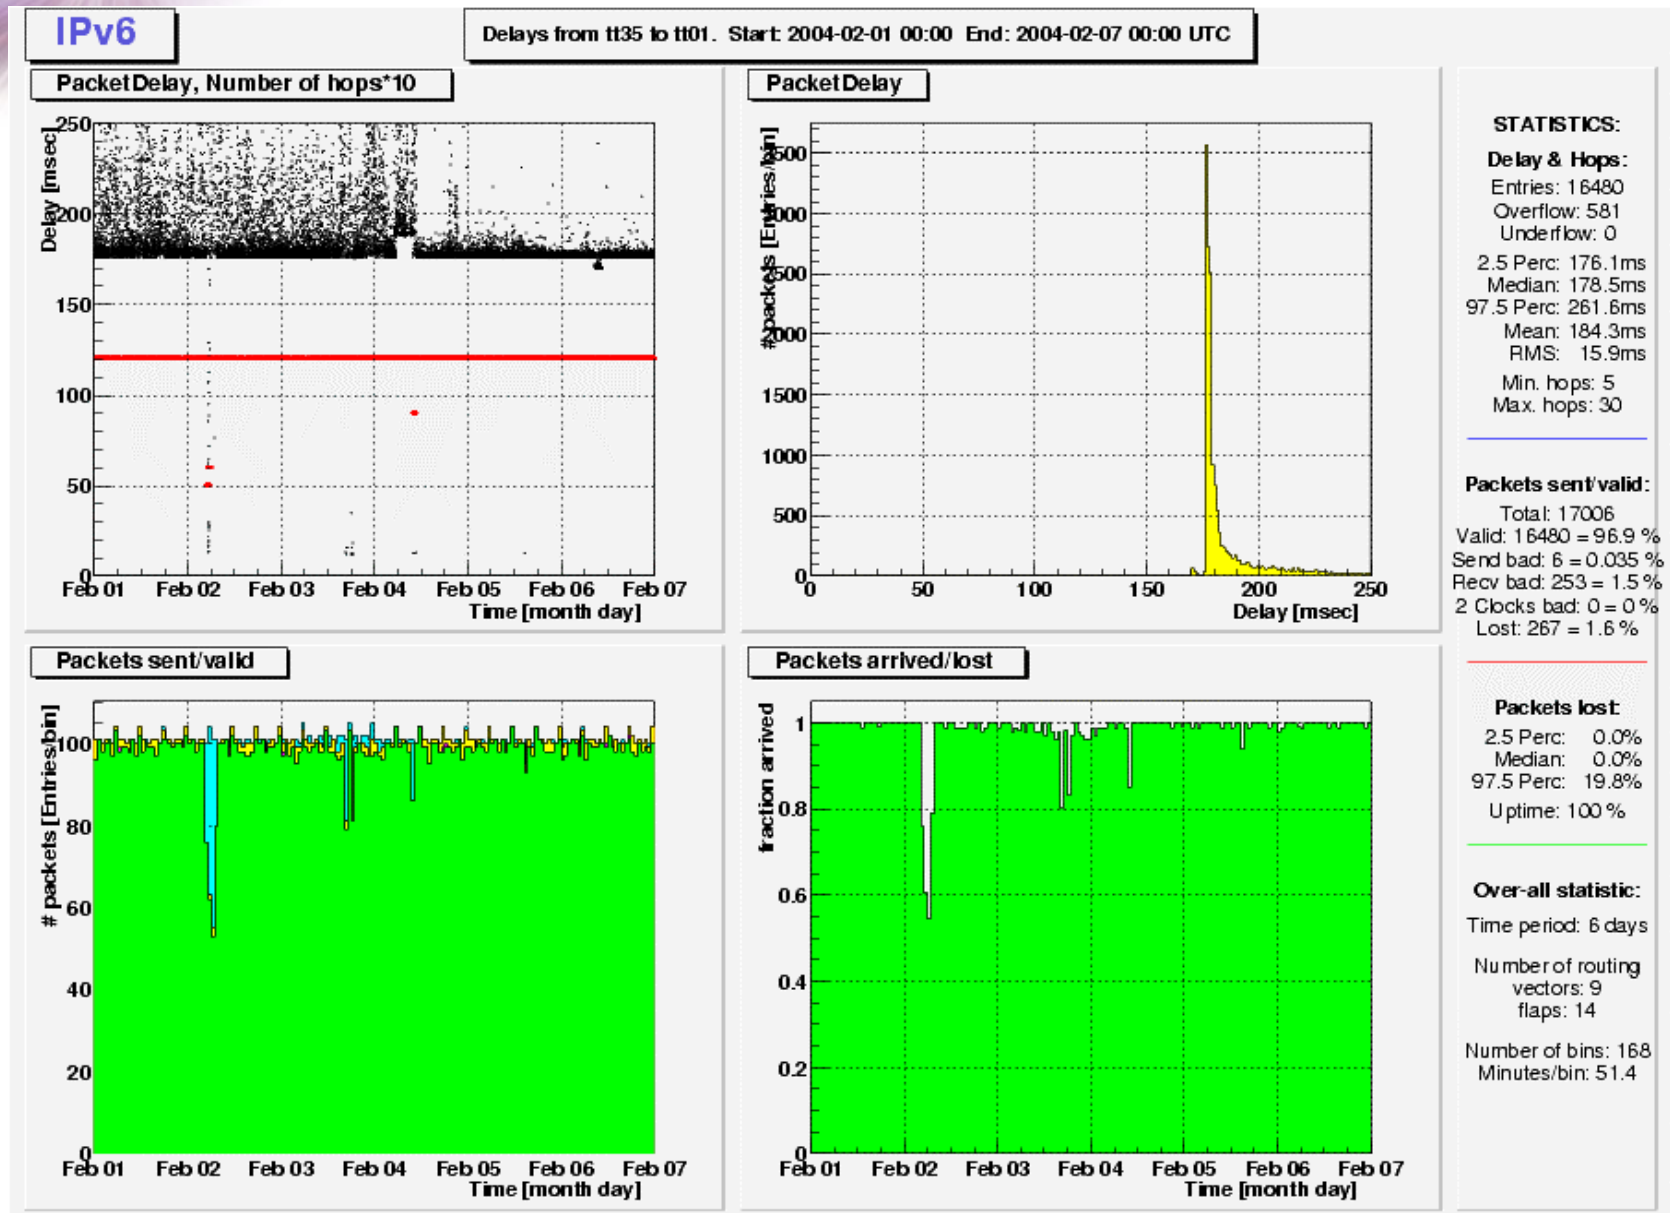

Some got worse!

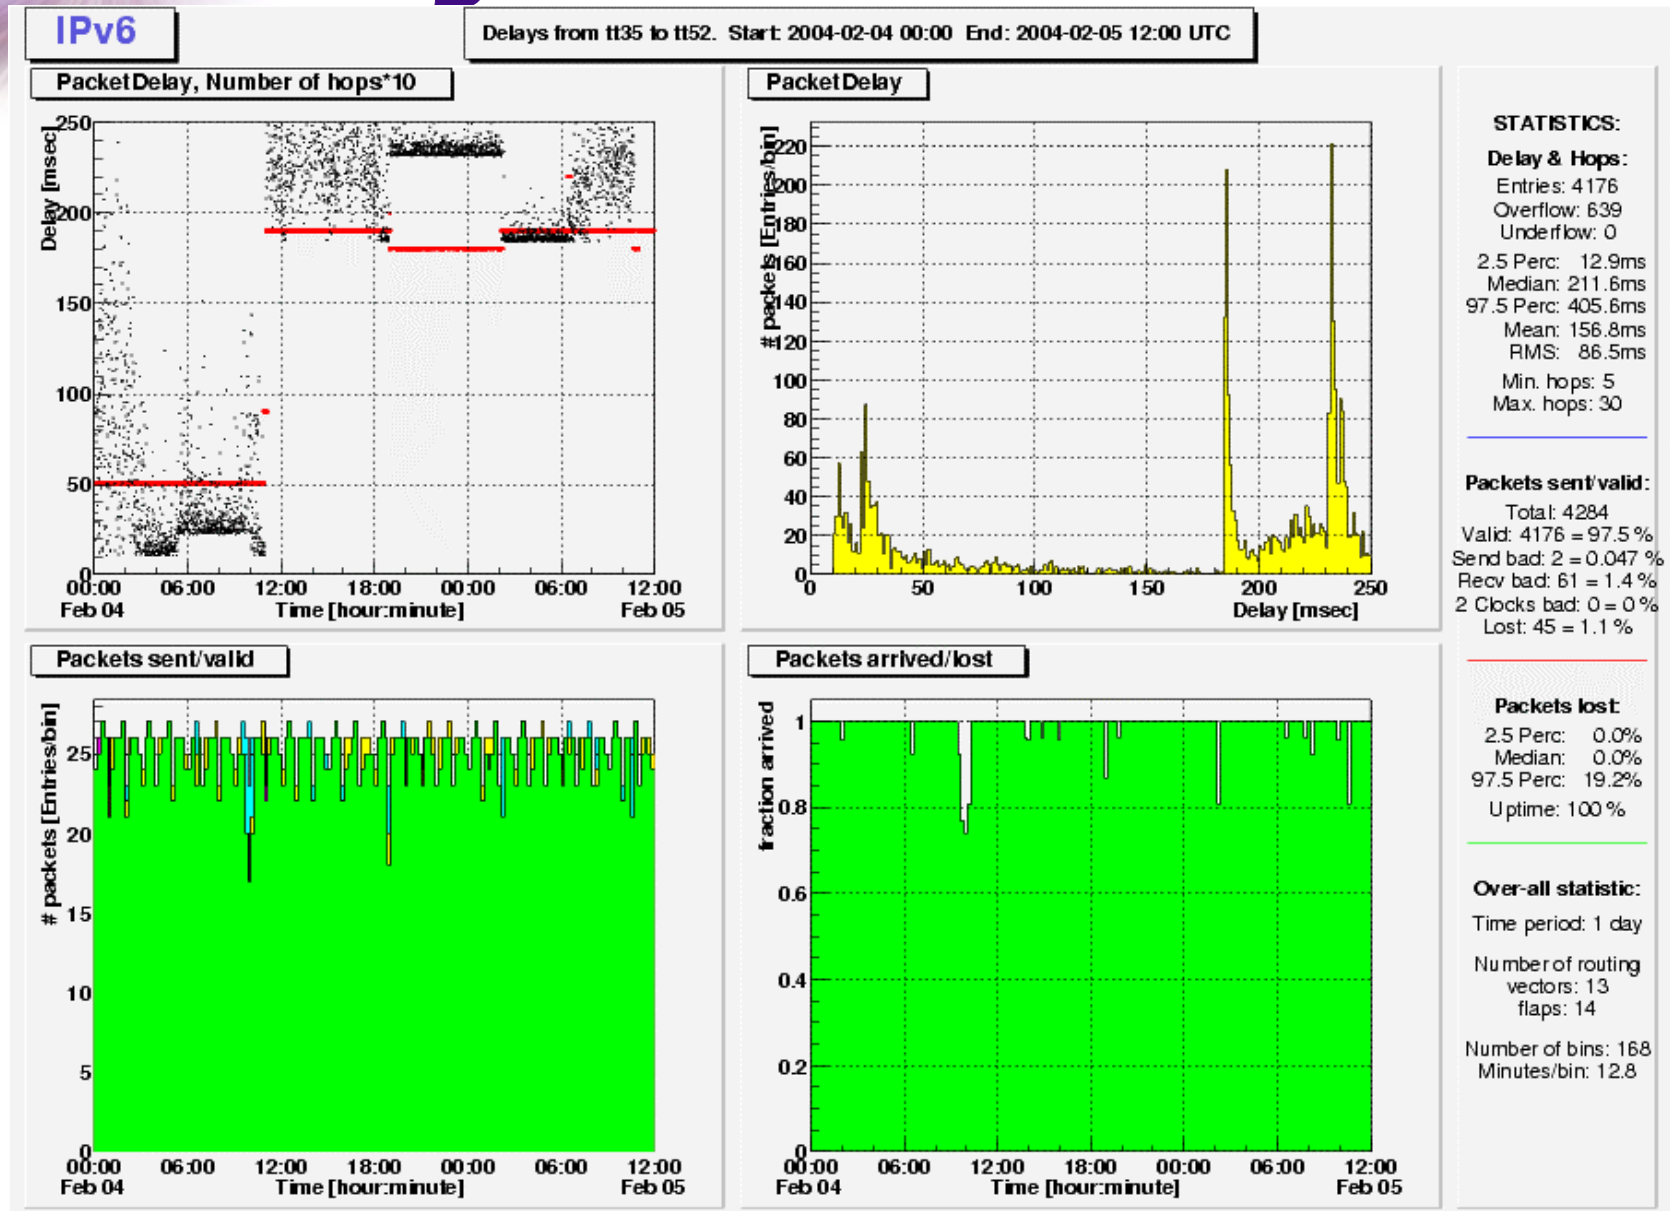

Why the change?

- Checked traceroutes from TTM (results came in the following day)
- Some routes which originally used our tunnel to Global Crossing in Amsterdam now used Abilene
- But the performance went down – why?

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

Why the change?

- Intra-Abilene routes are excellent, as are most of their onward connections
- Abilene mark tunnels with 11537:600 – these had lower localpref assigned
- But some international routes (ITN) from Abilene already come from Europe – lowered the localpref on these also

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK

...and got better again

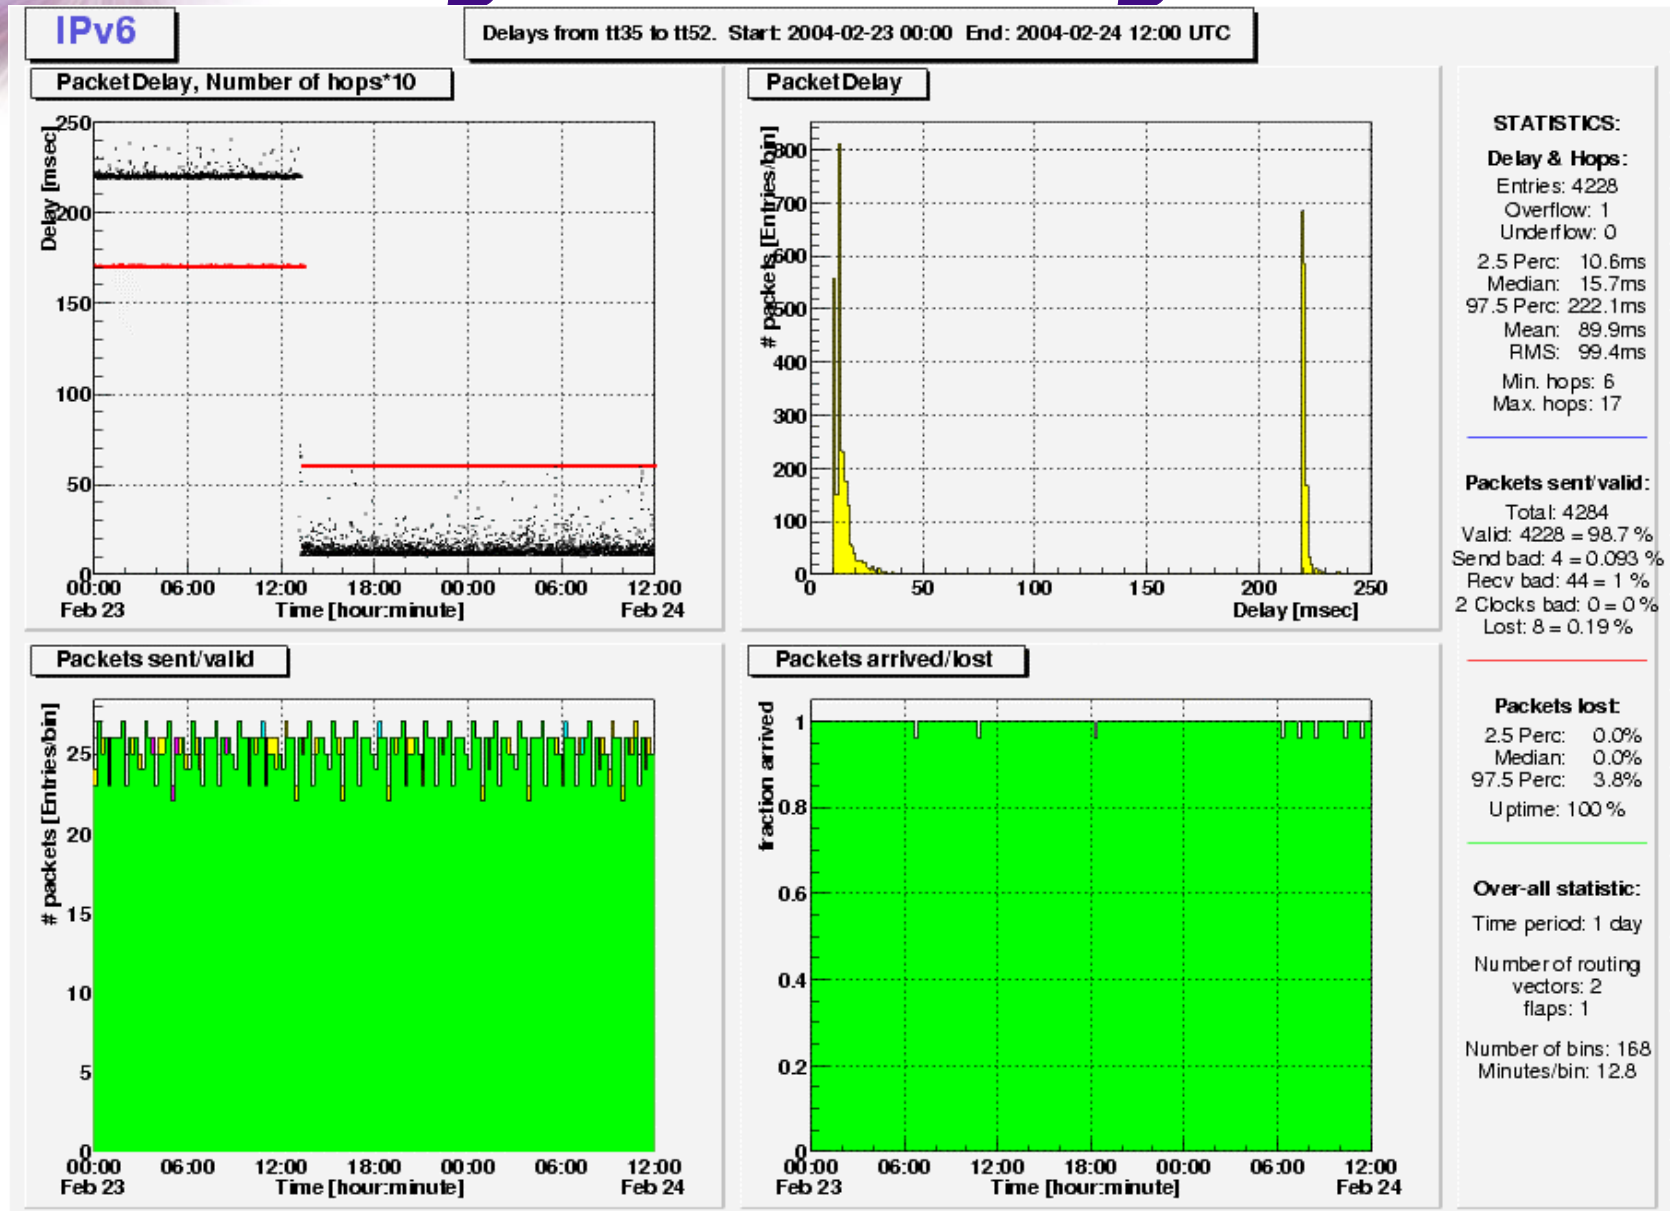

Summary

- If some peers filter and some don't, traffic may not go the way you expect
- Finding what to expect is very difficult in the absence of routing registries
- Sometimes a tunnel is better than a native link, if the peers of your peer are good
- If you have a TTM box, much closer attention can be paid *by others* to your connectivity

.

HEAnet
IRELAND'S NATIONAL EDUCATION & RESEARCH NETWORK