v6fix:Wiping the Slate Clean for IPv6

2006.2.15

Kenjiro Cho(WIDE Project/IIJ)

Ruri Hiromi(WIDE Project/Intec NetCore,Inc.)

Topics

- Introduction of “v6fix”, an effort to solve problems in IPv6 deployment
  - Technical analysis of specific problems found in the real world
- Tools and measurement results
  - Macro-level analysis of IPv6 deployment

IPv6 deployment status

- Majority of equipment and software are IPv6-Ready
  - Major routers, operating systems and applications
- Still, many applications and appliances work only with IPv4
- getting in various business fields
  - sensor networks, building automation, others
- Most people lack knowledge and experiences of IPv6
  - When non-experts encounter a problem, they are often clueless

Example illiteracy

- Hotel internet systems have an instruction for guests
  - Trouble shooting:

    If you have IPv6 enabled,
    please disable IPv6

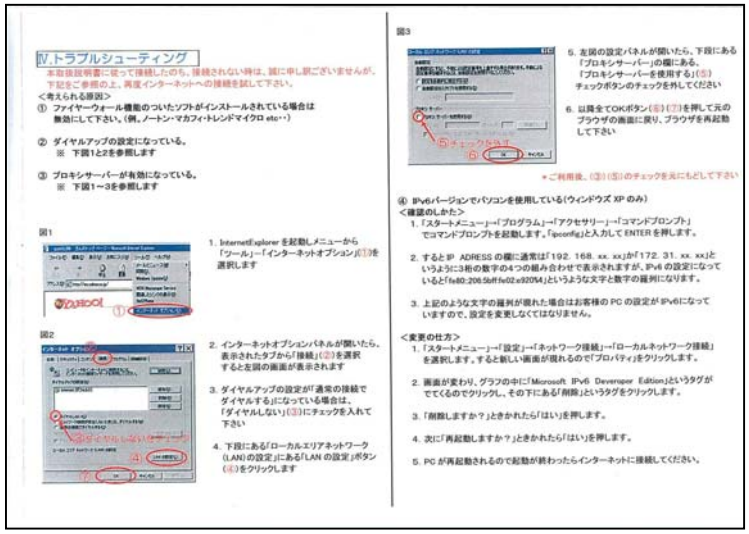
Ⅳ.トラブルシューティング

a brochure in a guest room

- The cause of the problem: a combination of
  - the DNS redirection system returns specific A for AAAA
  - client's stub-resolver accepts the A for AAAA and can't go anywhere

Wiping the slate clean for the IPv6

- Faulty behaviors are only 1% and often combinatorial, but could be fatal to acceptance of IPv6
  - slow fallback to IPv4 when encountering IPv6 errors
  - misbehaving DNS resolvers
  - filtering of ICMPv6 by firewalls
  - DNS misconfigurations (lame delegation,etc)
  - poorly configured tunnels
  - lack of peering or IPv6-capable paths
- Known problems are just the tip of an iceberg
  - We need to act now, or they would bring a negative impact to the IPv6 deployment

v6fix

IPv6 fix activities

- Technically identify/analyze/solve real-world problems in IPv6 deployment
  - Our enemy: “after trials&errors, disabling IPv6 fixed my problem!”
- Cooperation among researchers, implementers and operators
  - Overlooked problems are often found at boundaries of specification, implementation and operation

v6fix topics

- Harmful Effects of the "On-link Assumption"
- Misbehaving DNS servers and resolvers
- Slow fallback to IPv4 after a failure of TCP connection attempt
- Misconfigured firewalls
- Comparative analysis of the quality of the IPv6 Internet

Examples from reported problems

Case #1(DNS loop at hotel)

- A real story of the hotel internet system
  - DNS is intercepted and redirected to the sign-up page
  - IPv6-enabled users can’t go beyond the 1st page
  - Hotel has an instruction for disabling IPv6
- Caused by a combination of erroneous DNS redirection system and stub-resolver
  - the redirection system always returns the A for the sign-up page when receiving non-A queries
  - Client’s stub-resolver queries AAAA first for any address and always receives the same A and blindly accepts it

Case #2(DNS server slowdown)

- A real story of a Japanese ISP
  - ISP upgraded a DNS cache server to BIND9, and received complaints about slowdown
  - Recompiling BIND9 with “--disable-ipv6” fixed the problem, and reported it at JANOG
- Caused by older BIND9 w/o IPv6 connectivity
  - The server without IPv6 connectivity always tries to talk over IPv6, and ends up with falling back to IPv4 after timeouts
  - Fixed in BIND 9.2.5 and 9.3.1

Common factors

- The problems appear only with specific combinatorial conditions
- The implementers and operators didn’t notice the problems until someone reported the problems
- Even for the professionals, it wasn’t easy to track down the problems

tools

v6fix tools and measurement results

- Goal: to understand the macro-level IPv6 healthiness
- Current methodologies
  - Wide-area measurement of behaviors of 2nd/3rd-level DNS servers
  - Comparative paths analysis with IPv4/IPv6 dual-stack tools

DNS server measurement

- Active measurement of AAAA responses under .jp domains

behavior of JP domains/servers with AAAA

	2004		2005	
	domain names	DNS servers	domain names	DNS servers
faulty	0.04%	0.11%	0.04%	0.13%
okay	82.16%	84.39%	85.08%	84.26%
unknown	17.80%	15.50%	14.88%	15.61%

  - Breakdown of problems from recognized as “faulty” server

problem breakdown

	20004	2005
ignoring AAAA RR queries	4.70%	5.80%
responding with RCODE 3("Name Error")	4.70%	1.70%
responding other than No Error and Name Error	8.50%	4.20%
responding with broken answers	0.00%	0.00%

dual-stack path analysis

- measurement techniques specifically designed for dual-stack
  - take measurements for both IPv4 and IPv6 at the same time
  - compare IPv6 results to IPv4 results
  - extract problems which exist only in IPv6
- methodology
  - dual-stack node discovery
    - create a dual-stack node list by monitoring DNS AAAA replies
  - dual-stack ping
    - run ping/ping6 to target dual-stack nodes
    - select a few representative nodes per site (/48) by RTT ratios
  - dual-stack traceroute
    - run traceroute/traceroute6 to the selected nodes
    - visualize IPv6 Path MTU to observe macro-level paths/IPv6 tunnels
    - visualize IPv4/IPv6 paths to neighboring sites for comparative analysis
- results from Dec 2005

distribution of IPv6/IPv4 RTTs

(4,000 ping targets: IPv4 on x-axis, IPv6 on y-axis)

- individual nodes far above the unity line: leaf issues
- clusters above the unity line: backbone issues
  - as observed in the left graph

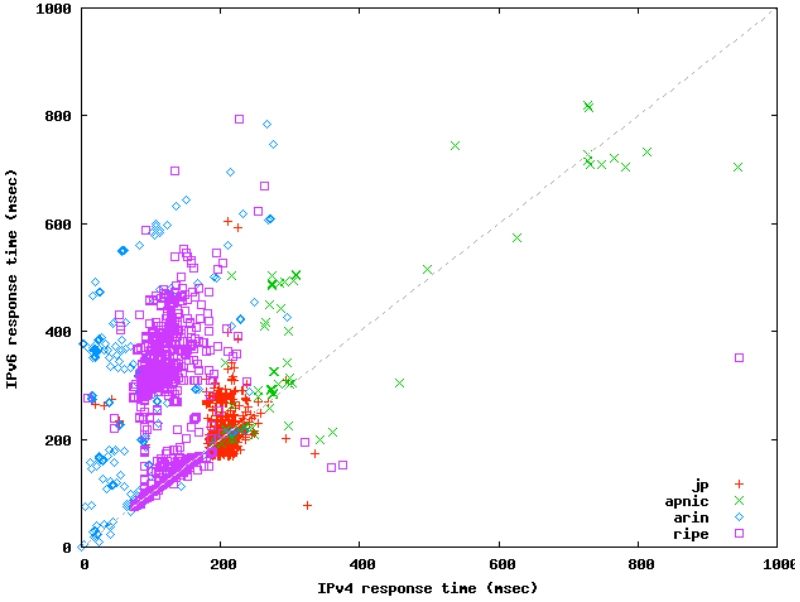

from NYSERNET

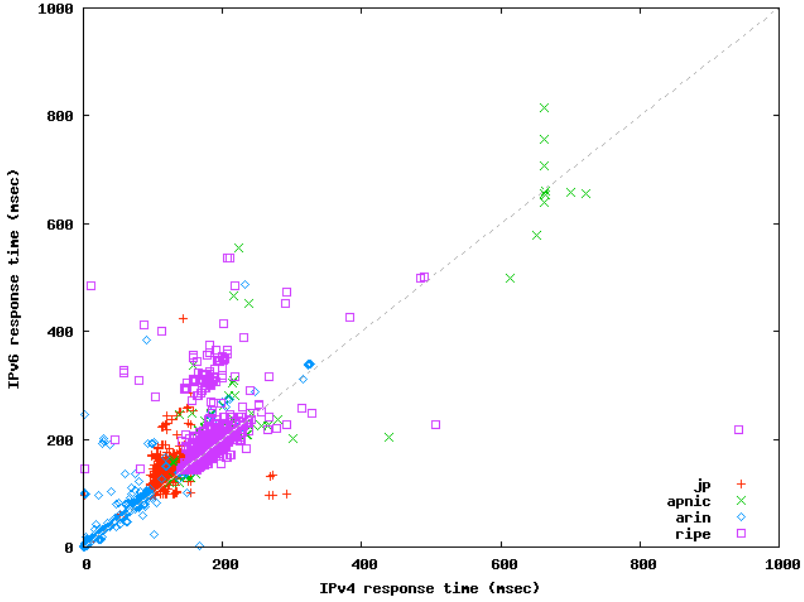

from ISC

Paths and PMTU visualization

(from NYSERNET to ARIN sites)

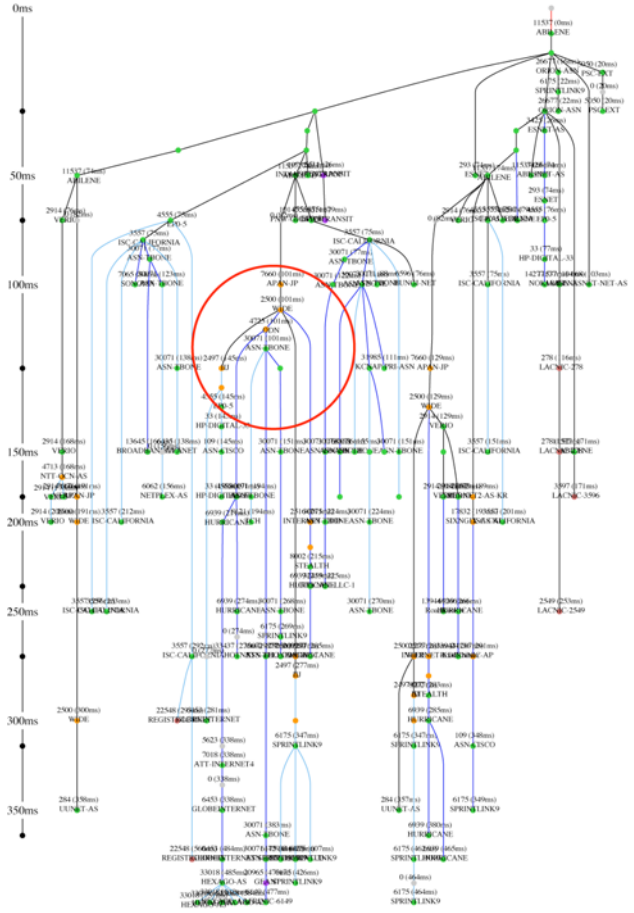

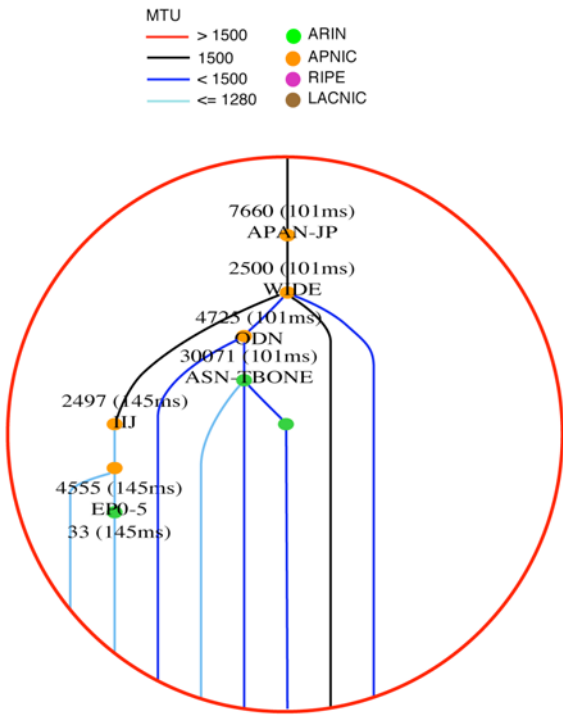

many ARIN paths via jp!
(lack of peering)

Paths and PMTU visualization

(from ISC to ARIN sites)

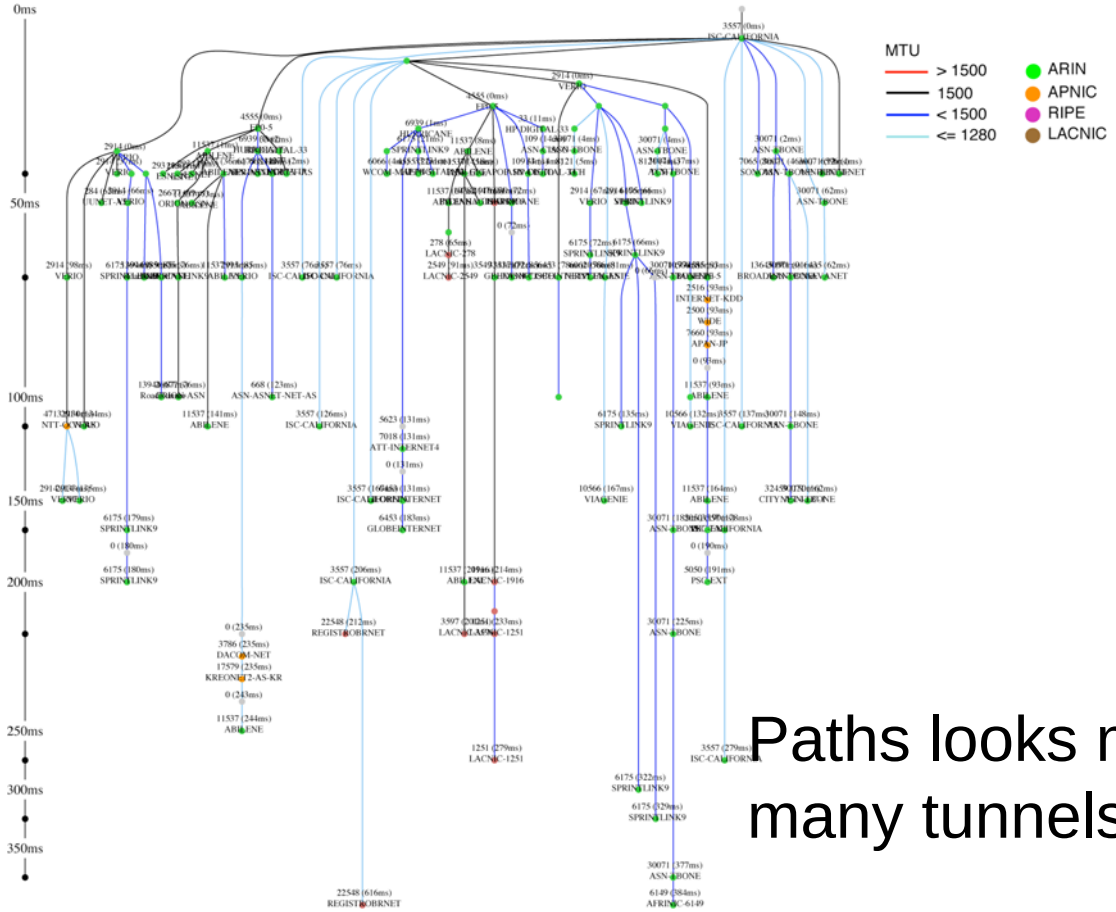

Paths looks much better but many tunnels in backbone

Abilene case: a well-known problem

- Abilene has been trying to encourage IPv6 adoption
  - NO-AUP, tunnel services for IPv6
- but ended up with horrible IPv6 paths, mostly with tunnels
  - ISPs are reluctant to move to paid IPv6 connectivity
- Abilene is thinking about suspending its relaxed AUP for IPv6

- out tool tries to illustrate such issues
  - hoping to help convince operators

dual-stack traceroute

to ABILENE from WIDE (IPv4:upper, IPv6:bottom)

similar RTTs/hops for IPv4/IPv6: native dual-stack paths

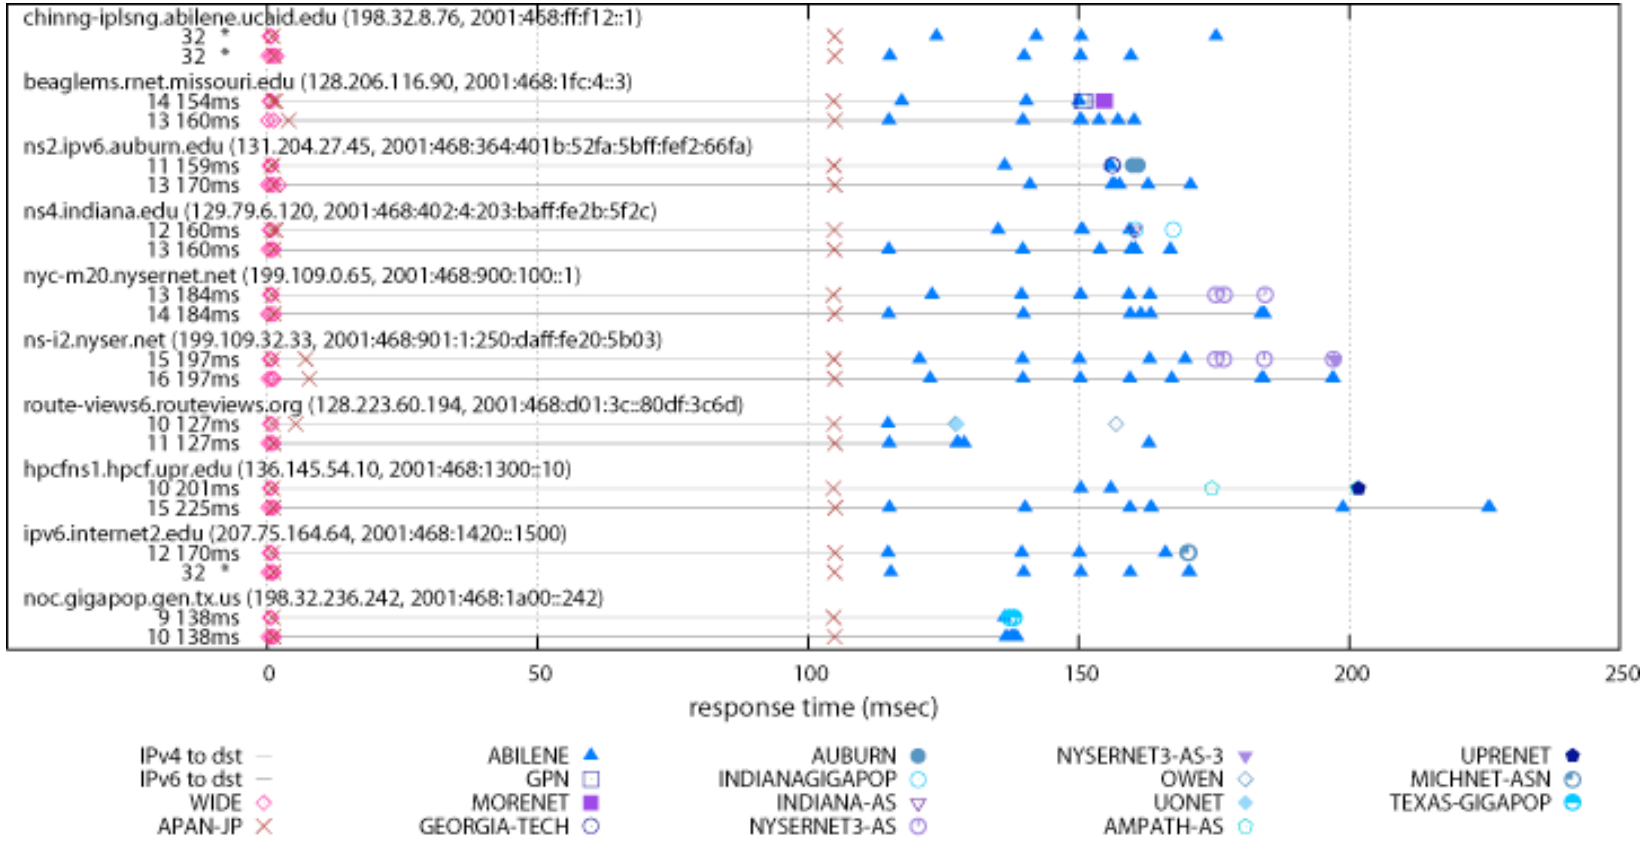

dual-stack traceroute

to ABILENE from IIJ

similar RTTs but different paths: currently more common

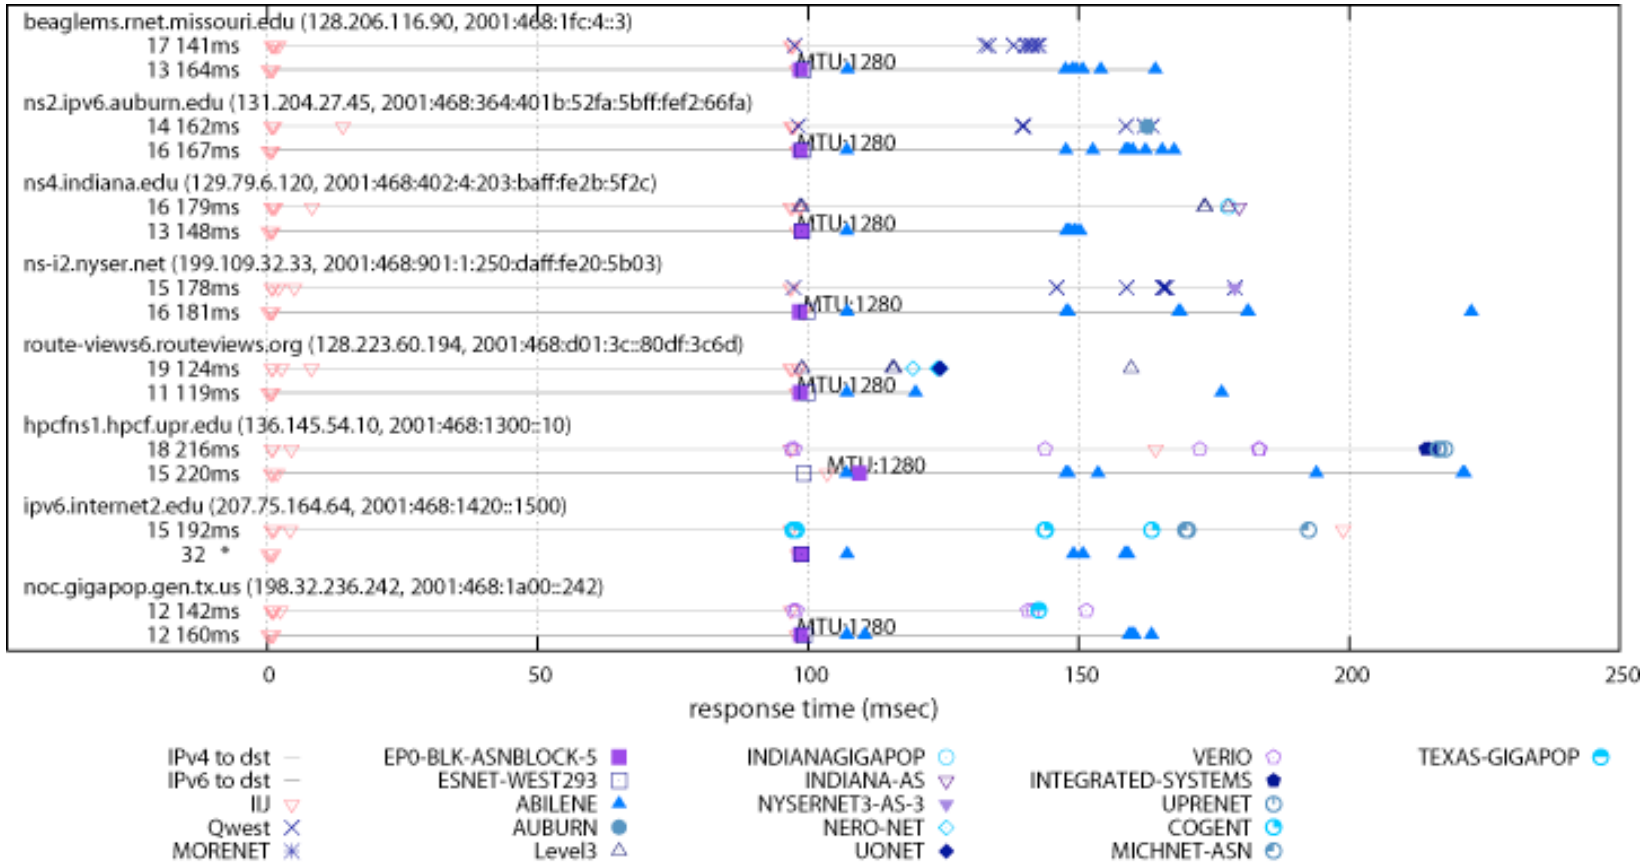

dual-stack traceroute

to ABILENE from ES

IPv6 RTTs larger than IPv4: roundabout tunnels

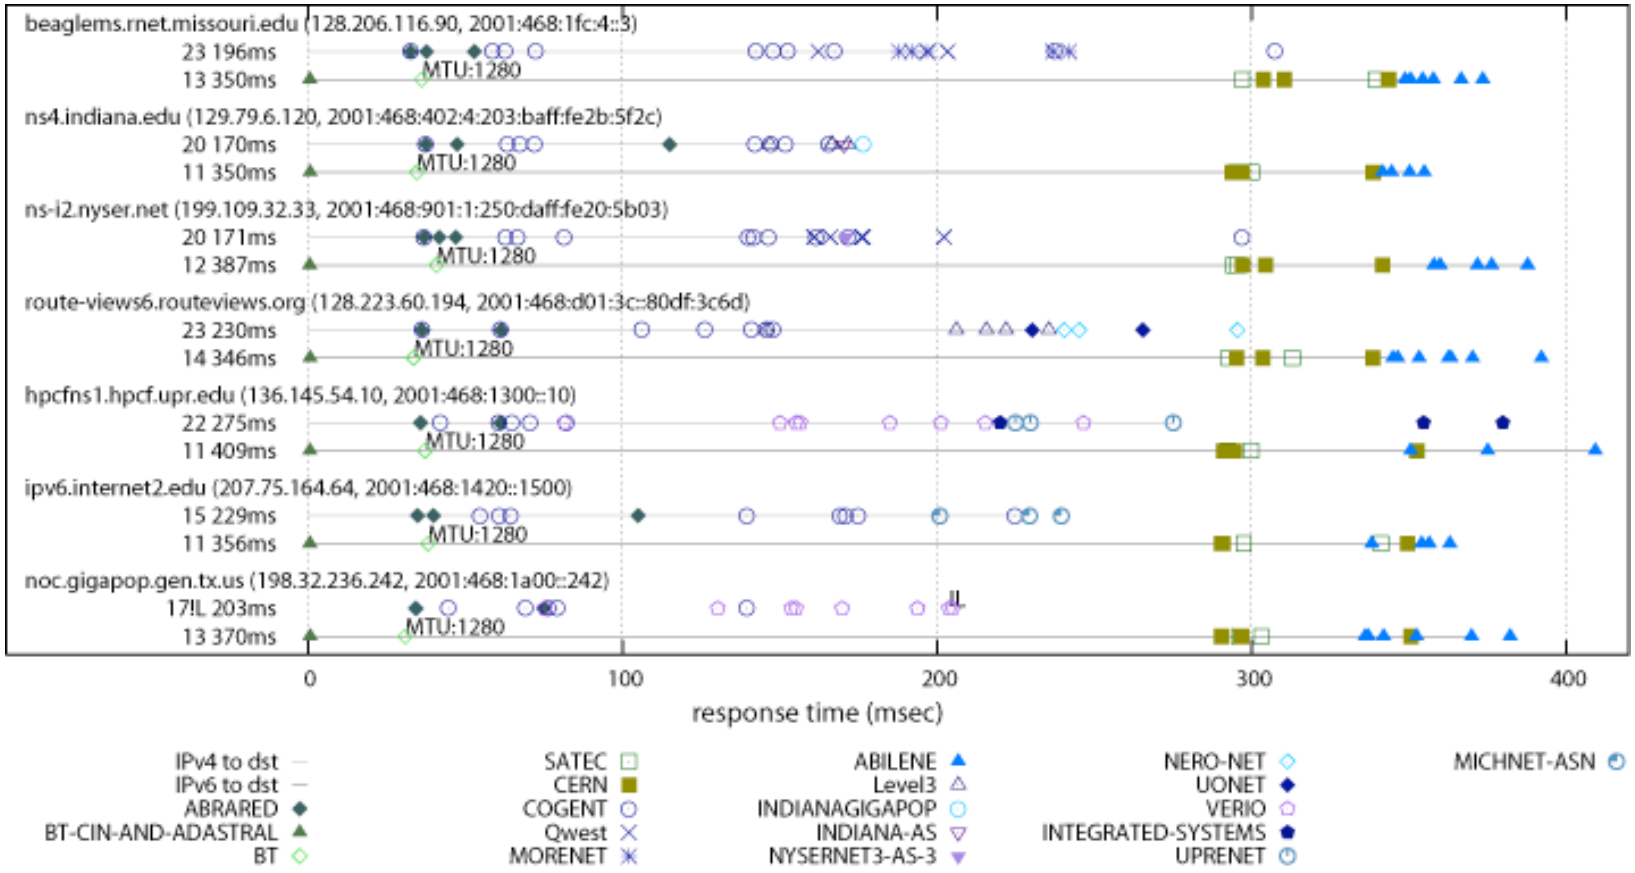

Conclusion

- Faulty behaviors are only 1% and often combinatorial, but could be fatal to acceptance of IPv6
  - slow fallback to IPv4 when encountering IPv6 errors
  - misbehaving DNS resolvers
  - filtering of ICMPv6 by firewalls
  - DNS misconfigurations (lame delegation,etc)
  - poorly configured tunnels
  - lack of peering or IPv6-capable paths
- Knowledge sharing
  - we need to realize the dangers of harmful adoption of the IPv6
  - cooperation among researchers, implementers and operators
  - we need to act now, or they would bring a negative impact to the IPv6 deployment

acknowledgments

v6fix members

- Kazu Yamamoto, Tatsuya Jinmei, Koichi Kunitake, Toshio Shimojo, Sohgo Takeuchi, Hideaki Yoshifuji, other WIDE project members

Collaborators for tool development

- Matthew Luckie (U. Waikato), Bradley Huffaker (CAIDA)

Measurement cooperation

- JPRS, CAIDA, NYSERNET, ISC, consulintel

This project is partly funded by

- Japanese Ministry of Internal Affairs and Communications
- WIDE project

Documents / Information

- Web pages

  http://v6fix.net/

  – Overview, documents and fact database

- E-Mail

  contact@v6fix.net

  – reports are welcome