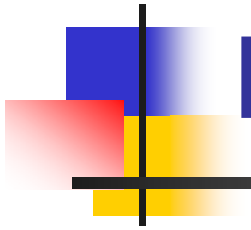

Operational Problems in IPv6: Fallback and DNS issues

Tomohiro Fujisaki, Arifumi Matsumoto,
Katsuyasu Toyama, Tsuyoshi Toyono
and Shirou Niinobe
Nippon Telegraph and Telephone Corporation

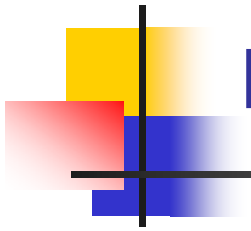

Network problems associated with IPv6

- Today, IPv6 is widely used
  - Many operating systems and routers support IPv6 without any additional software.
    - Anyone can use IPv6
- I'll introduce two problems with IPv6
  - IPv6/IPv4 fallback problem
    - Described in section 4 at http://www.v6fix.net/docs/wide-draft-v6fix.en
  - DNS cache server problem

What is ‘IPv6-to-IPv4 fallback’?

- Many current IPv6/IPv4 dual-stack operating systems start their communication using IPv6.
  - If destination has both ipv4 and ipv6 address, end host first tries ipv6. And if it fails, then tries ipv4.

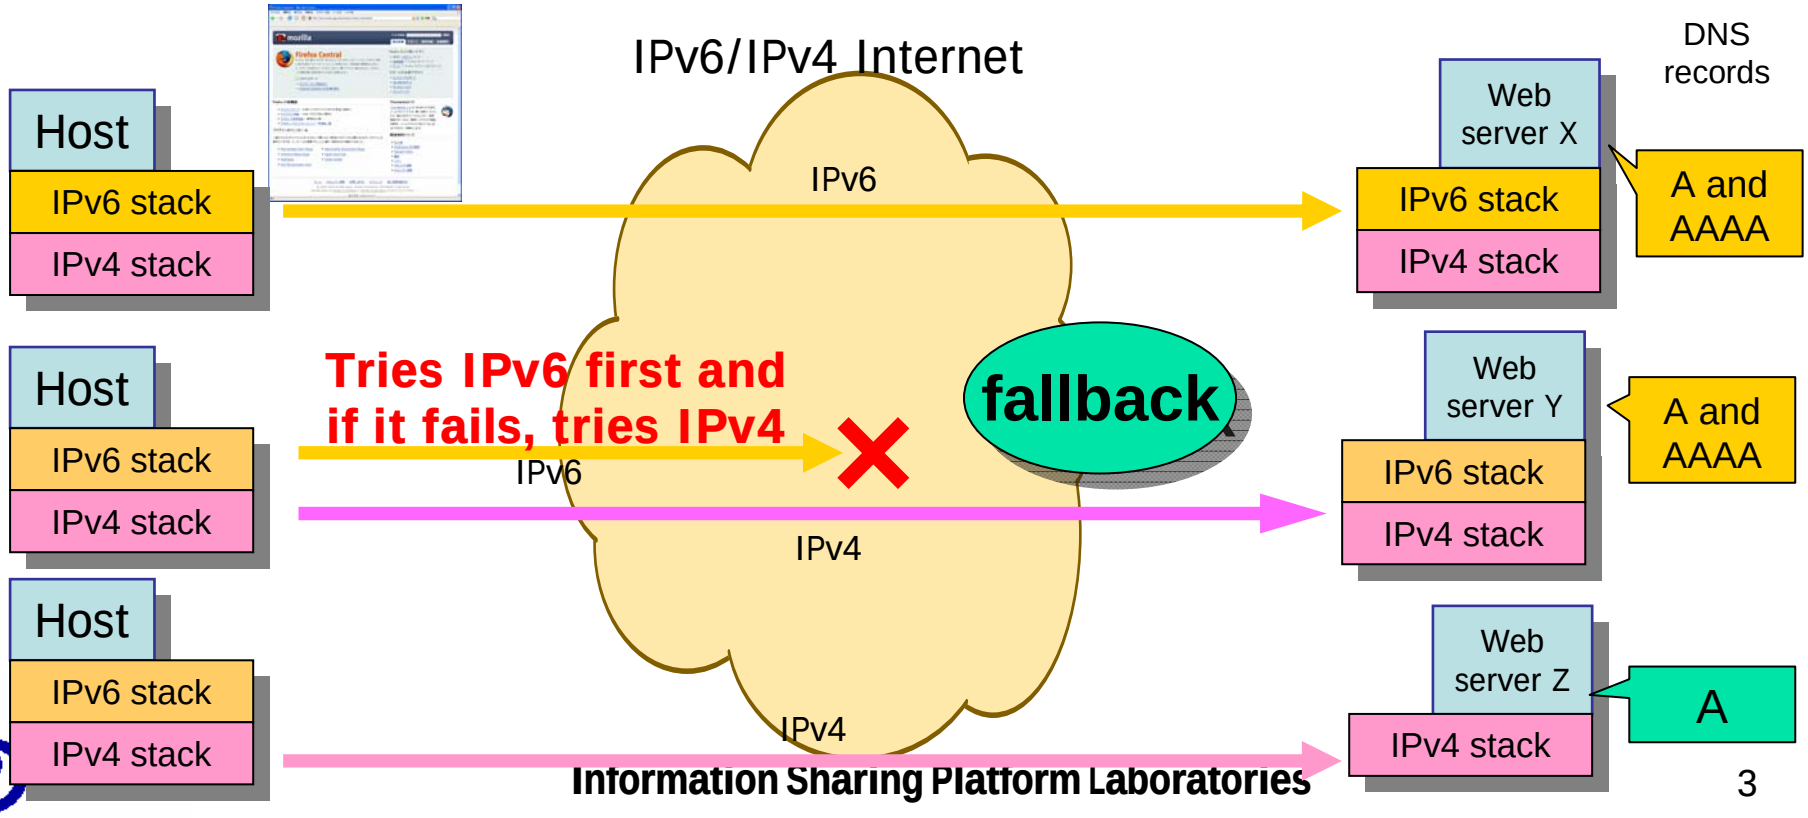

IPv6/IPv4 fallback problem

Problem: Fallback sometimes takes a long time.

- Problem is especially significant in TCP-based applications.
  - They initially need to establish communication channels
  - If destination node has multiple IP addresses, application tries them sequentially until TCP communication channel is established.

User's view

User has to wait until communication is finished until web page is displayed completely.

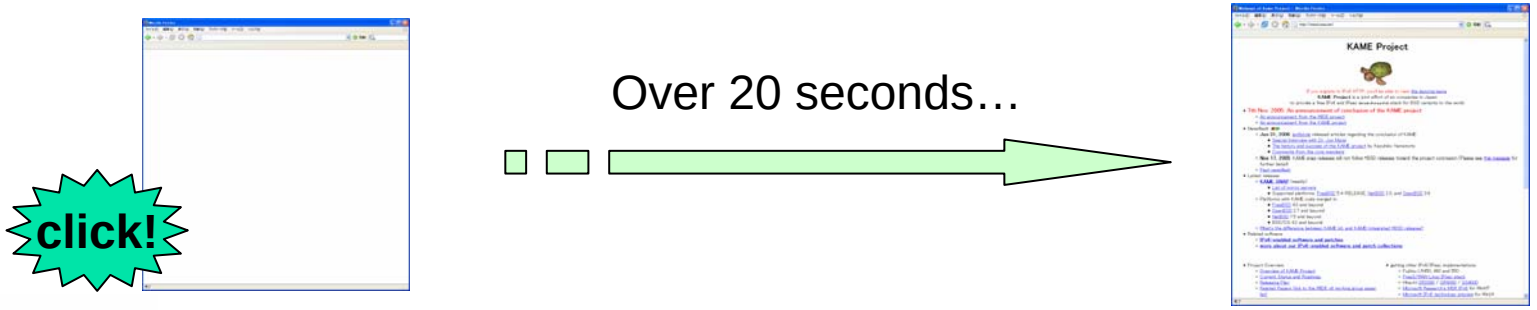

Where does fallback problem occur?

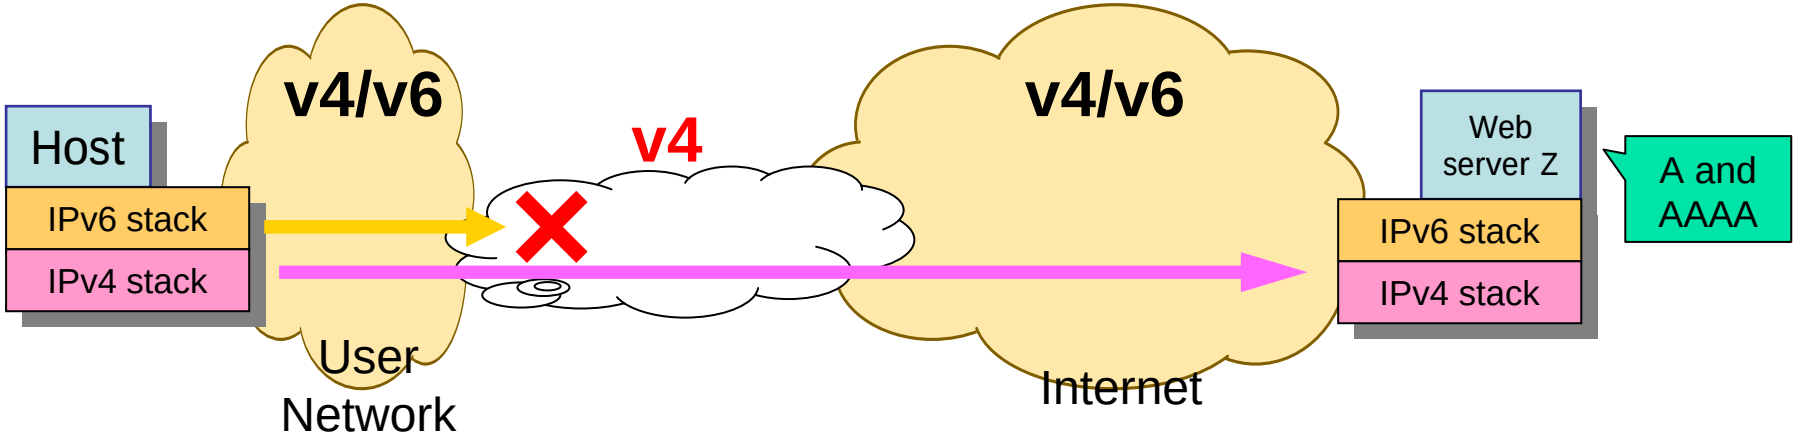

- Poor IPv6 connectivity.
  - Using unmanaged transition mechanism
  - Destination server has AAAA records, but no IPv6 connectivity
  - No IPv6 global connectivity (e.g., VPNs and networks using ULA)

Node behavior to network faults

- Networks should notify to end hosts that there is no route to the destination node, and end hosts should fall back from ipv6 to ipv4 according to the notification.
  - ICMPv6 Type1: Destination Unreachable

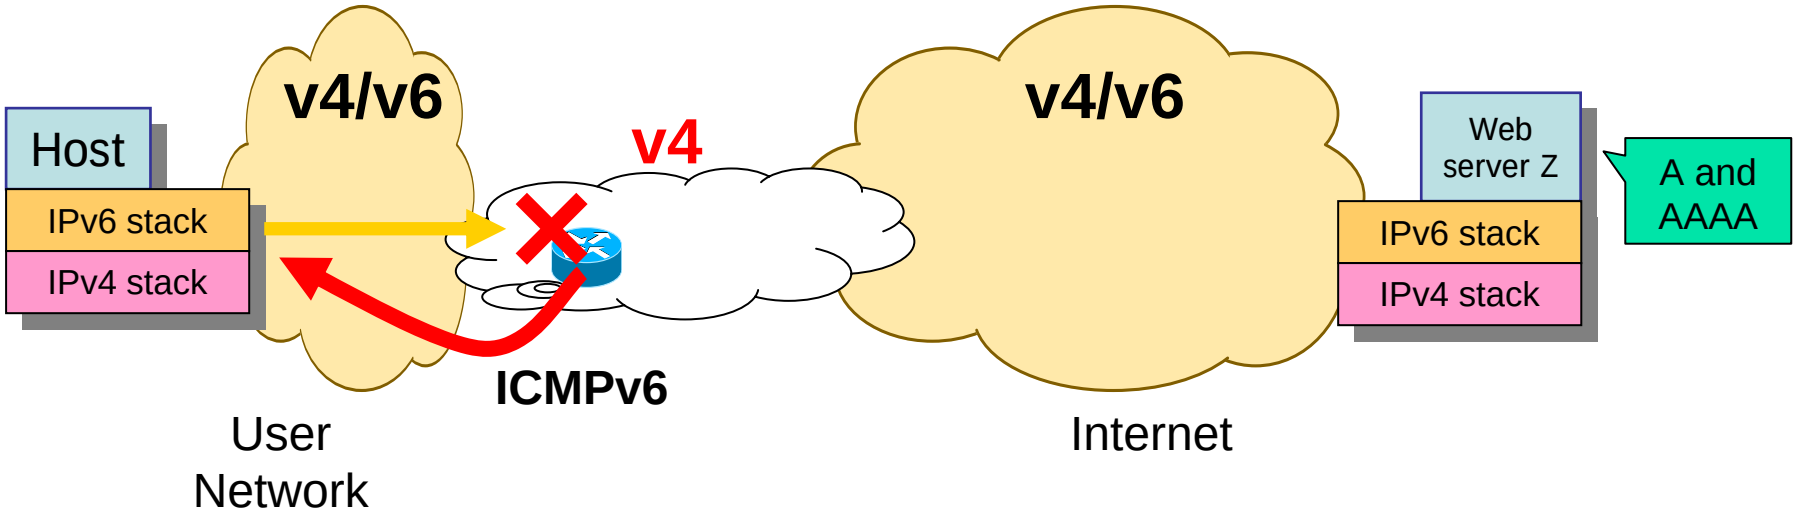

Fallback experiment

- We tested various operating systems' fallback behavior under following situations:
  - No errors are reported by the network
    - (1) Timeout of connection
  - Errors are reported by the network
    - (2) ICMP errors
      - Mainly `ICMP destination unreachable' message
  - Force to fall back
    - (3) TCP RST
      - not legitimate solution??

Fallback experimental setup

- Measuring time required to fall back from IPv6 to IPv4 at client PC when:
  - No response from network
  - ICMPv6 destination unreachable returned
    - no route to dest
    - administratively prohibited
    - address unreachable
    - port unreachable
  - TCP RST returned

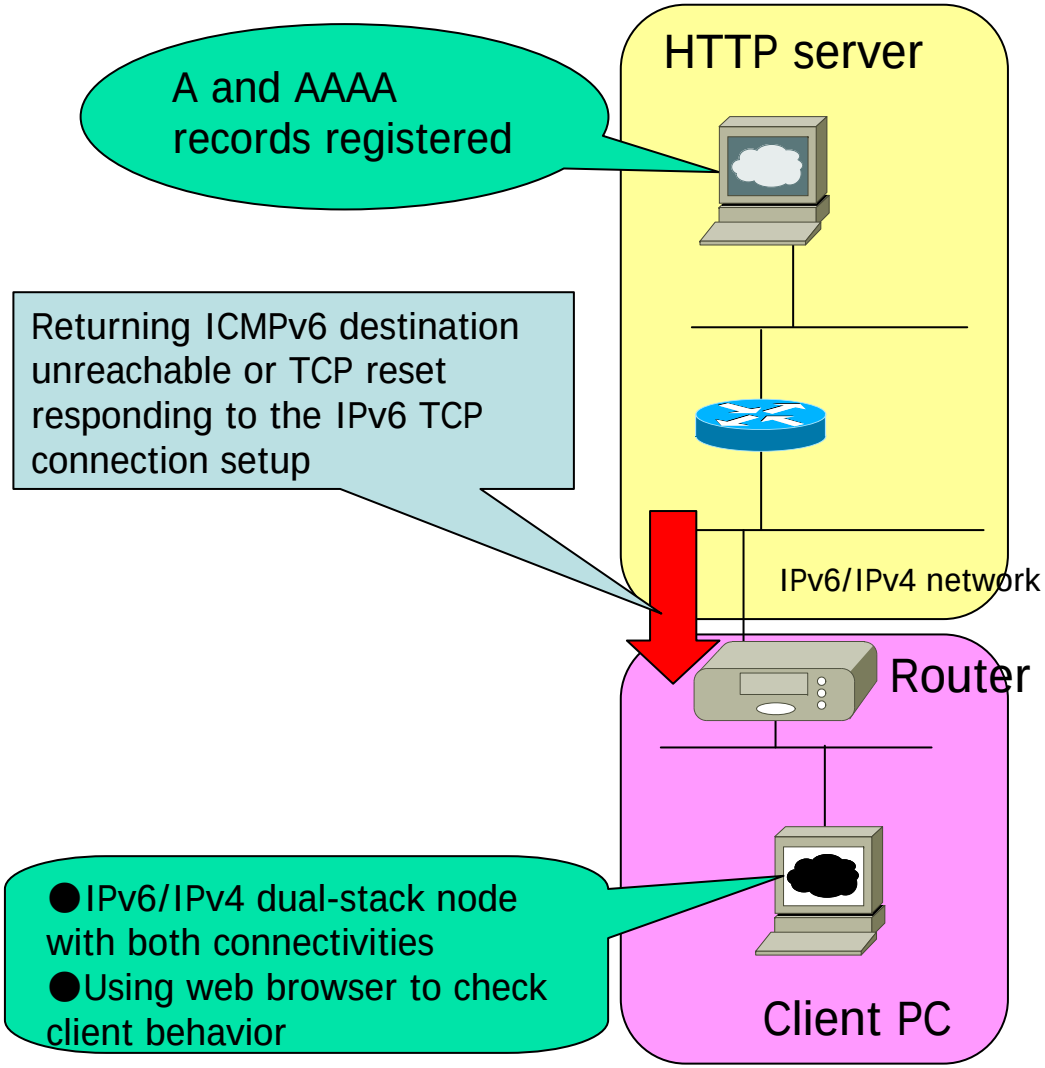

Results of experiments

unit : seconds

OS	Browser	No Error	Type=1(Destination Unreachable)							TCP reset
			Code=0	Code=1	Code=2	Code=3	Code=4	Code=5	Code=6	
Windows Vista Home Basic	IE	19.99	21.00	20.99	21.00	20.99	20.99	20.99	20.99	1.01
	FireFox	21.00	21.00	21.00	21.00	20.99	20.99	20.99	20.99	1.01
Windows Vista Enterprise	IE	21.06	21.00	20.99	21.00	20.99	21.00	21.01	21.00	1.01
	FireFox	20.99	21.00	20.99	20.99	21.00	21.00	20.99	21.00	1.00
Mac OSX (10.4.8 8L2127)	Safari	74.79	11.80	11.83	17.37	11.68	11.75	74.86	74.89	0.01
	FireFox	74.91	11.61	11.73	11.70	No fallback	11.63	74.79	74.77	0.01
FreeBSD (R6.2-#p1)	FireFox	74.99	12.61	12.61	12.69	No fallback	12.61	74.99	74.99	0.01
Fedora Core 6 (kernel-2.6.20)	FireFox	188.98	0.01	0.01	0.01	0.01	0.0.1	No fallback	No fallback	0.01

Time between first IPv6 TCP SYN packet and IPv4 TCP SYN packet immediately after fallback occurs.

IE Version : 7.0.6000.16386
Firefox version : 2.0.0.1
Safari version : 2.0.4

Code=0: no route to destination [RFC2463]
Code=1: communication with destination administratively prohibited [RFC2463]
Code=2: beyond scope of source address [RFC4443]
Code=3: address unreachable [RFC2463]
Code=4: port unreachable [RFC2463]
Code=5: source address failed ingress/egress policy [RFC4443]
Code=6: reject route to destination [RFC4443]

TCP stack behavior when ICMP errors are returned

- When node receives ICMP error packets, TCP stack behavior is defined in RFC 1122 (for IPv4 ICMP only).
  - When node receives an ICMP hard error, TCP aborts connection immediately.
  - When node receives ICMP soft error, TCP must not abort connection.

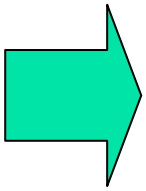

Currently, ICMPv6 destination-unreachable error handling is different in each OS.

- The “draft-ietf-tcpm-tcp-soft-errors” proposes IPv6 version of ICMPv6 soft error handling….

Summary of IPv6/IPv4 fallback problem

- In IPv6-and-IPv4 enabled network, quality of each network affects the user's nodes.
  - If quality (e.g. reachability) of the network is clearly bad, administrators should control communication by:
    - Changing protocol preference using RFC3484 mechanism
    - Returning TCP RST to minimize fallback time

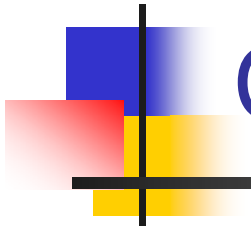

DNS cache server issue:
Increase in Number of DNS AAAA Queries

This topic was presented by T. Toyono
at 2006 OARC Member Workshop

http://public.oarci.net/oarc/workshop-2006/agenda/

Focus on

- User-Cache DNS queries, not on Cache-Authoritative queries
- increase in number of queries between users and cache servers caused by
  - 1. IPv6 support
    - Number of AAAA queries same as that of A queries
  - 2. Domain name completion
    - Domain name completion by operating system (API), and by applications
  - 3. These Combinations
    - Sequence of queries

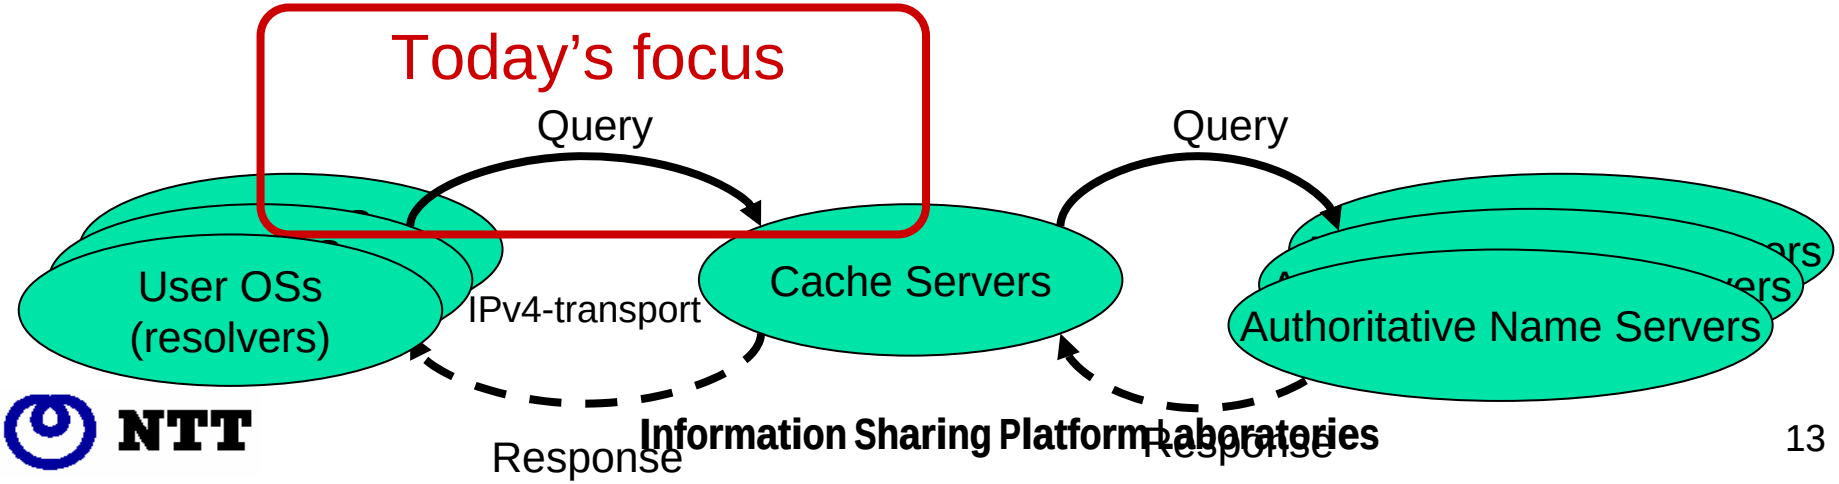

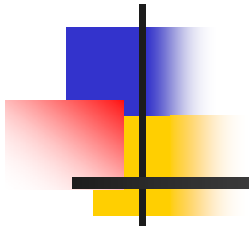

(1) IPv6-enabled OS increases DNS queries

1. IPv6 and OS Resolver

- IPv6-enabled OSs ask for both A and AAAA resource records
  - “A” query = IPv4 name resolution
  - “AAAA” query = IPv6 name resolution

- Sends both A and AAAA queries for every name resolution
  - Currently, almost no application specifies “DNS Query Type”; therefore, OS sends both

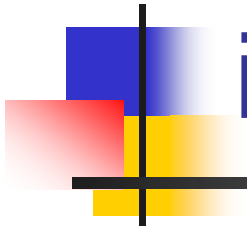

(2) Domain name completion increases DNS queries

2. Domain Name Completion

- When a name resolution fails, both OS and APP automatically resolve the domains with prefix/suffix completion
  - e.g., when name resolution of “host” failed → host.com → host.org → host.net ...

- OS using these domains to complete:
  - FreeBSD: specified by “search” “domain” in /etc/resolv.conf and distributed via DHCP
  - Windows: configured in control panel and distributed via DHCP

- Applications:
  - Mozilla: retries name resolution for a domain by adding “www. ” domain prefix
  - IE6: using MSN search, then adds a domain suffix ".com" ".net" ".org" and ".edu"

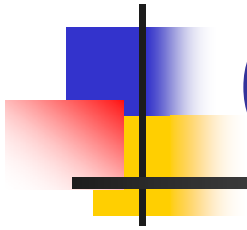

(3) Combination of (1) and (2)

Combination in FreeBSD

- Sequence
  - Sends A query first, then AAAA query
- Domain Completion
  - Tries domain completions for every set of “A+AAAA”
- IPv6 address
  - Sends AAAA queries even if it doesn’t have an IPv6 address

(Ex) User Query: noexist-example.com

A noexist-example.com

AAAA noexist-example.com

A noexist-example.com.com

AAAA noexist-example.com.com

A noexist-example.com.net

AAAA noexist-example.com.net .....

If IPv4 address is resolved, stop here.

Combination in Linux

- Tries AAAA queries for all domain completions, then A queries with domain completions
- IPv6 address
  - Sends AAAA queries even if it doesn't have an IPv6 address

(Ex) User Query: noexist-example.com

AAAA noexist-example.com
AAAA noexist-example.com.com
AAAA noexist-example.com.net
A noexist-example.com
A noexist-example.com.com
A noexist-example.com.net

Even if domain has IPv4 addresses,
first, AAAA queries are sent.

Combination in Windows Vista (before β2 Build5270)

- Tries AAAA queries for all domain completions, then tries A queries with domain completions
- Same as Linux (kernel 2.6.15) behavior
- IPv6 address
  - Sends AAAA queries even if it doesn't have an IPv6 address

(Ex) User Query: noexist-example.com

```
AAAA noexist-example.com
AAAA noexist-example.com.com
AAAA noexist-example.com.net
A    noexist-example.com
A    noexist-example.com.com
A    noexist-example.com.net
```

Even if domain has IPv4 addresses, first, AAAA queries are sent.

Windows Vista (β2 Build5270)+IE7.0(at the time)

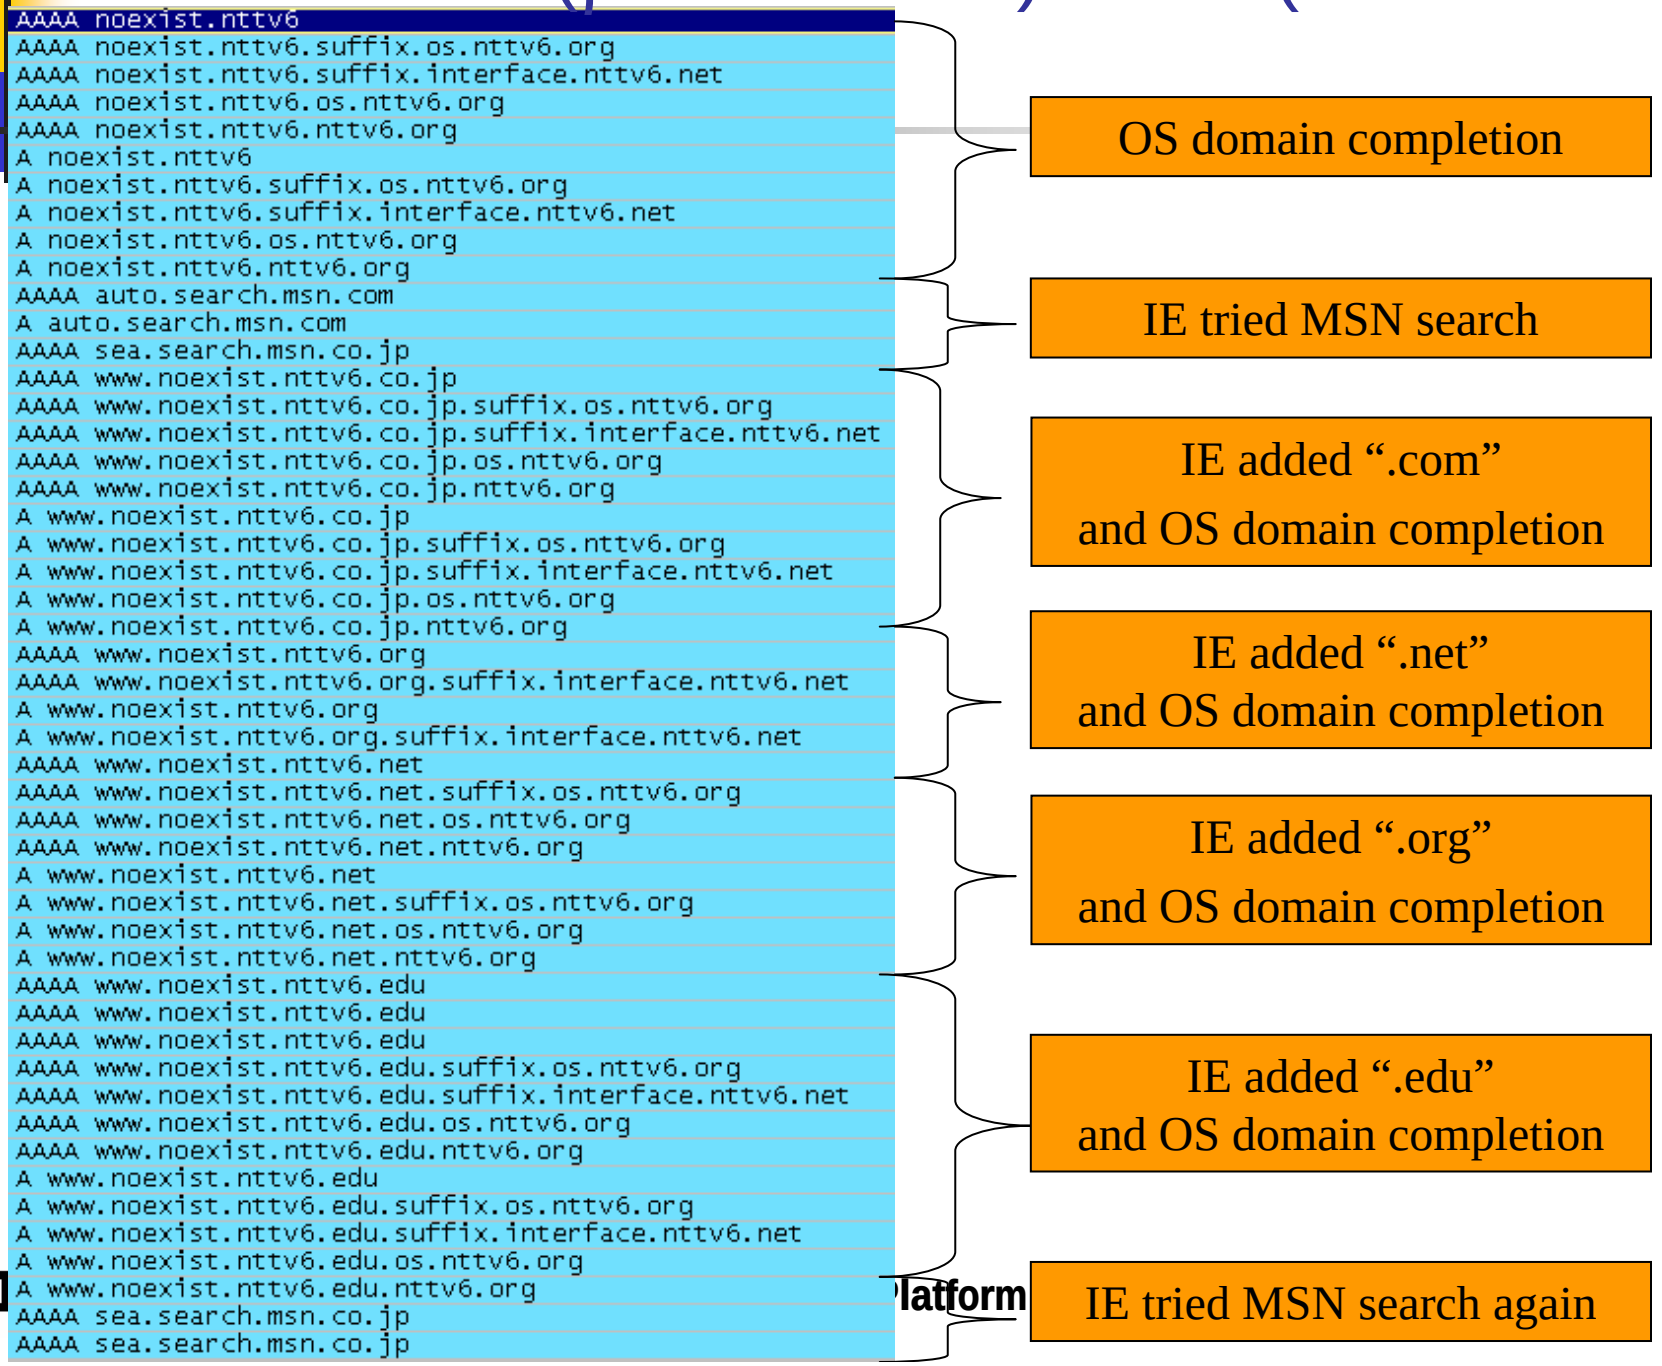

Our alert report and MS response

- NANOG36 (2006/2)
  - We reported this behaviors (Vista β) and alerted increase in number of DNS queries
- NANOG38 (2006/10)
  - Abolade Gbadegesin@Vista Internet Protocols team
  - “The NetIO Stack in Windows Vista: Functionality and Deployment”
    - “NTT Labs: NANOG36 report with preliminary analysis based on Windows Vista”
- In his slides:
  - “Deployments of new behavior are best undertaken as joint efforts between host software vendors and public network operators”

Reference:
NANOG36 “Clear and Present Increase of AAAA Queries”
NANOG38 “The NetIO Stack in Windows Vista: Functionality and Deployment”

Combination in Windows Vista

- Status in Vista
  - “DNS sends A query first, follows up with AAAA only to servers that have some info, then stops”
  - “Vista doesn’t send AAAA queries if the only global IPv6 addresses

We appreciate this change by Microsoft!

Reference:
NANOG38 “The NetIO Stack in Windows Vista: Functionality and Deployment”

Windows Vista + IE 7.0.5744.16384

- NX-Domain
  - Sends A query first, and answer is “NX-Domain”, stops sending AAAA query
  - Doesn’t try domain name completions

User query “none.nttv6.net”(NX)

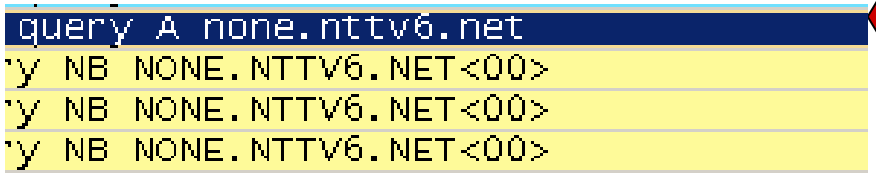

Recieves “NX Domain”,
so doesn’t send AAAA query

Windows Vista (RC2 Build5744) + IE 7.0.5744.16384

- No Answer
  - sends A query first, and answer is “No Answer”, then sends AAAA query
  - Doesn’t try domain name completions
  - Application displays “Not found: nttv6.net” page

User query “nttv6.net”(NoAnswer)

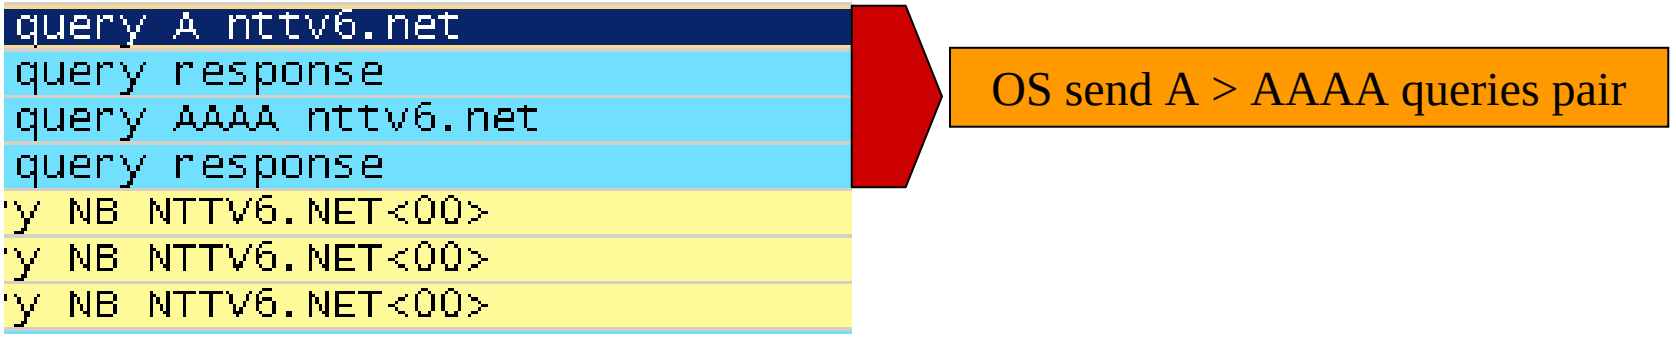

Results

	FreeBSD	Linux	MacOS X	Vista (β)	Vista
A & AAAA query sequence order	A first	AAAA first	A first	AAAA first	A first
When does domain name completion occur?	After A+AAAA	All AAAA completion first, then A	After A+AAAA	All AAAA completion first, then A	No completion
Send AAAA queries even if no IPv6 addresses assigned	Yes	Yes	No	Yes	No

- Linux sends AAAA queries first
- Linux sends all suffix completions of AAAA first, then A

- FreeBSD, Linux and old Vista send AAAA queries even if don't have IPv6 reachability

→ Now, if Vista doesn't have IPv6 address, they don't send AAAA queries

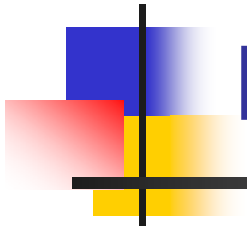

Network environment factors

Network environment factors for query increase

- Number of end users who have IPv6 addresses
- Some OSs send AAAA queries even if they don't have IPv6 reachability
  - Such as FreeBSD, Linux
- Others factors
  - If the answer was "NX-Domain"
  - Has A Resource Record, but doesn't have AAAA Resource Record
  - Domain suffix distribution to users by DHCP

Number of end users who have IPv6 addresses

- IPv6 environment in Japan
  - Many ISPs already provide IPv6 connectivity services to end users
    - e.g., NTT, Yahoo, IIJ, KDDI, and nifty, for example (The market share of these large ISPs is about 60-70% of all broadband users)
  - ISPs use IPv6 for their streaming services and IP-phone services, for example.
  - Global IPv6 addresses are assigned to end hosts
    - Vista will send AAAA queries

Share of large ISP's DNS cache queries, from users (2006/10 one day total)

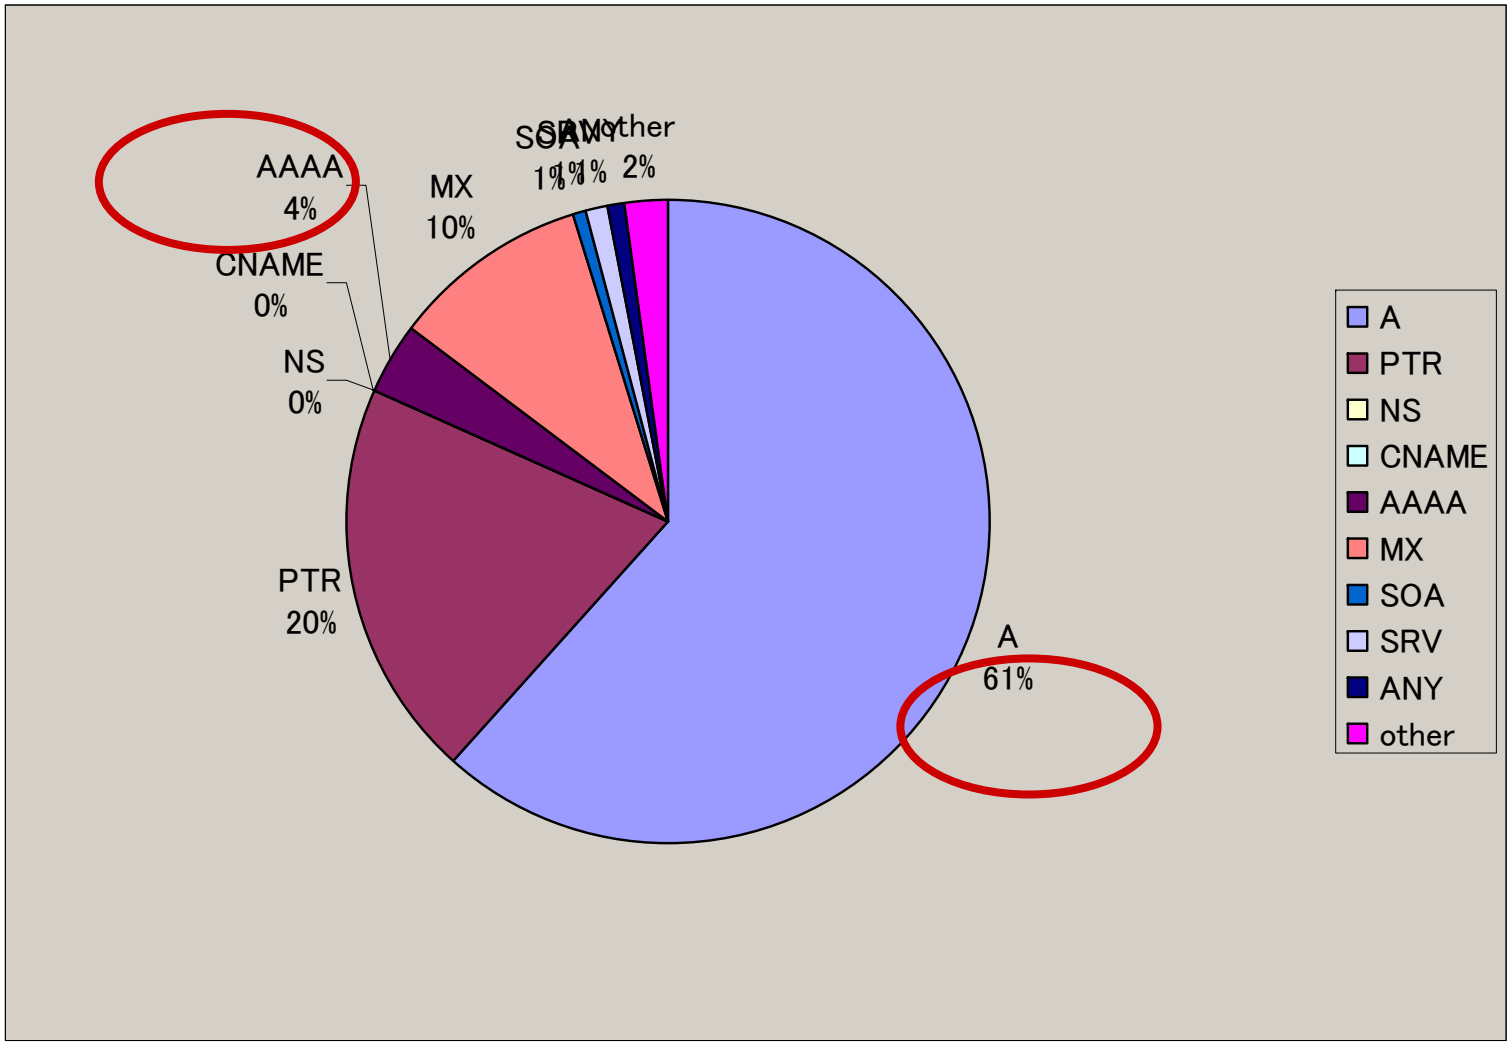

Expected increase in number of user queries

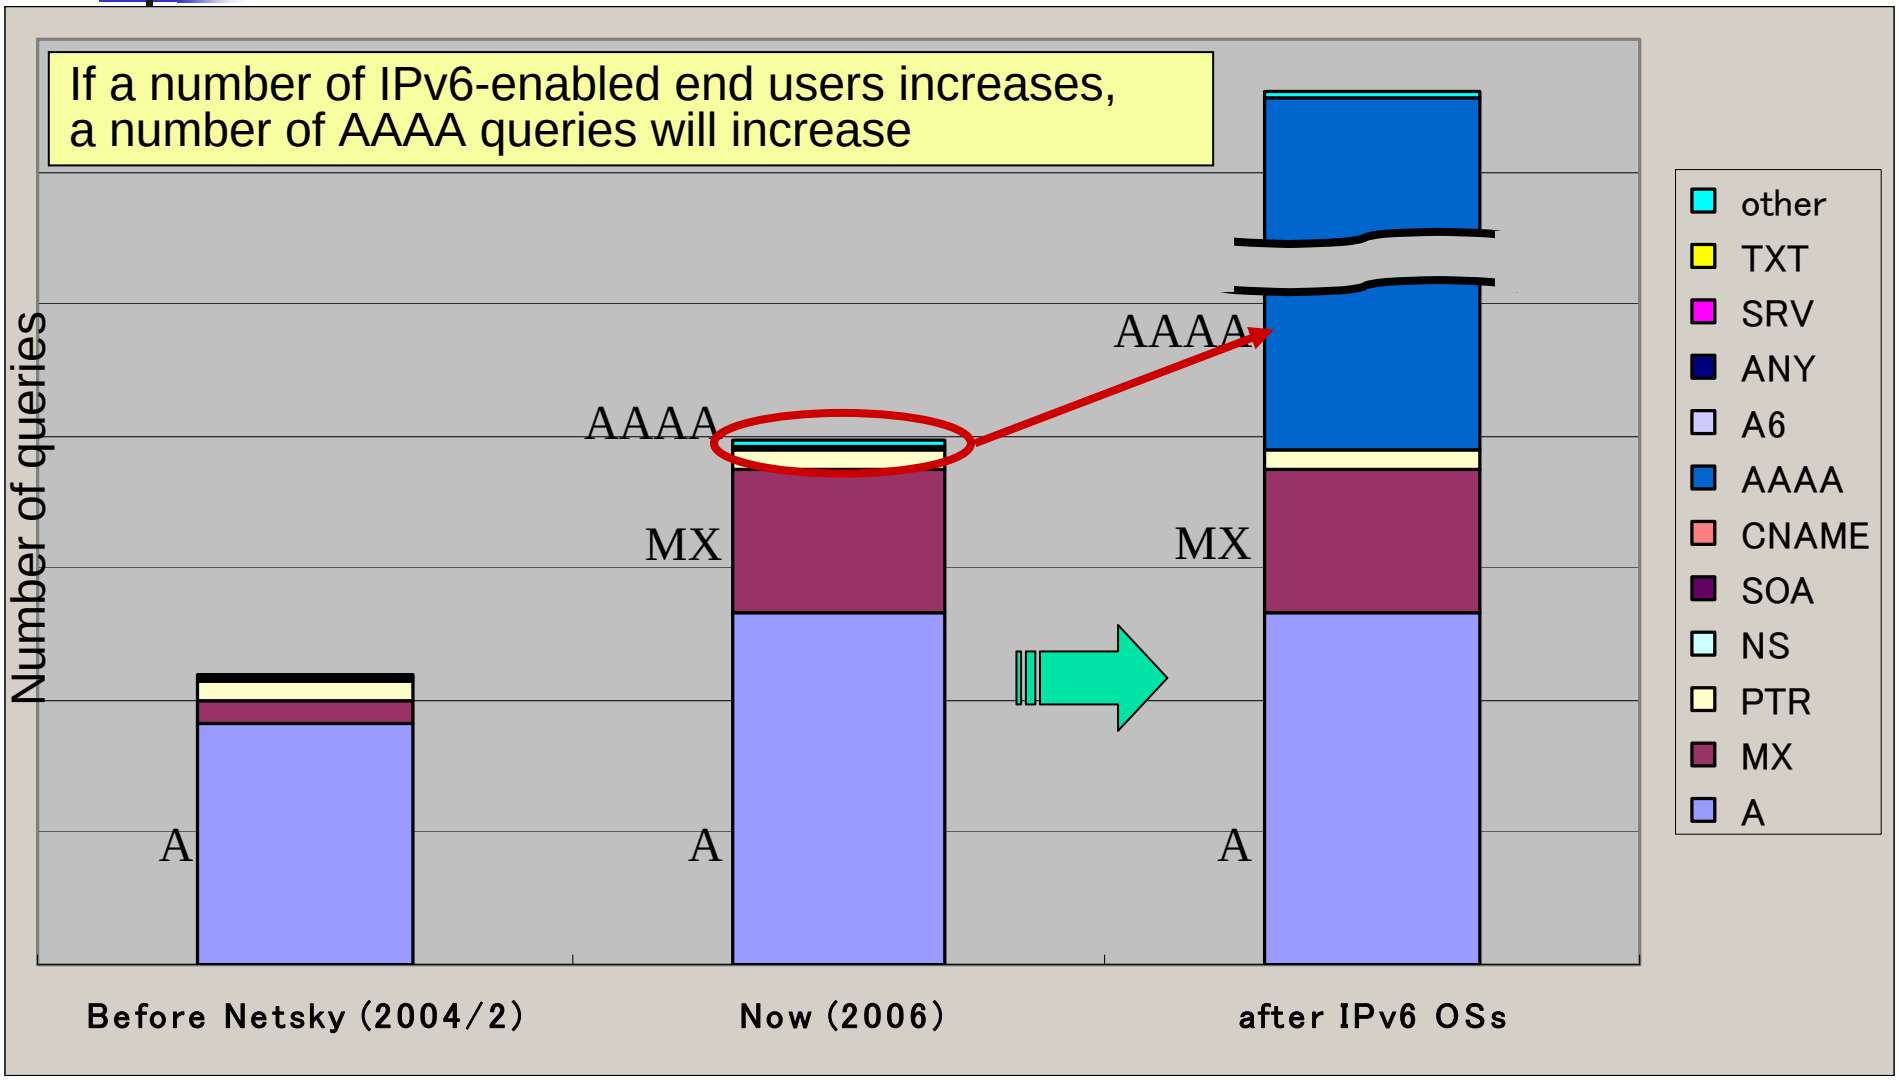

Conclusion

- If a number of IPv6-enabled end users increases, a number of AAAA queries will increase
  - The number of query increase depends on IPv6-enabled OSs and Applications implementation
- Some OSs send AAAA queries even if hasn't IPv6 reachability
  - As for Vista, the impact was minimized
- We have to prepare increase in number of DNS queries
  - Cache servers should be prepared for those increases
    - Large ISPs Cache servers (that use load balancing) would be better off preparing for those increases
    - Preparing authoritative servers for increases would be better
  - Is current search order of resolvers & applications appropriate?
  - Should IPv6 transport DNS be used?

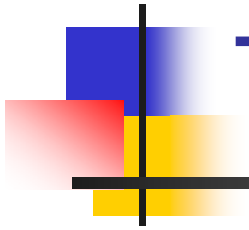

Thank you.