Measurement of BGP Anycast effects – experiences in .JP

Adding an anycast node at New York

Kazunori Fujiwara <fujiwara@jprs.co.jp>
Masato Minda, Shinta Sato, Izuru Shirai, Takayasu Matsuura
Japan Registry Services Co., Ltd. (JPRS)

July 21, 2007
8th CAIDA/WIDE Measurement workshop

About JP DNS

- JP DNS - authoritative name servers of .JP (JPRS) and some in-addr.arpa zones (JPNIC)
  - Using BIND 9 (All of JP DNS are using BIND 9)

Server	Operator	Anycast	IPv6
a.dns.jp	**JPRS**	BGP Anycast	Yes
b.dns.jp	JPNIC	N/A	No
d.dns.jp	IIJ	IGP Anycast	Yes
e.dns.jp	WIDE Project	Soon	Yes
f.dns.jp	NII	N/A	Yes

a.dns.jp – Anycast status

- a.dns.jp
  - AS: 23774
  - IPv4: 203.119.1.1
  - IPv6: 2001:dc4::1
  - Located at Tokyo and Osaka
    - using BGP anycast since Feb 2004
- JPRS added one anycast node at New York.
  - for two weeks
  - It is a test run for real operation at New York

a.dns.jp - Network Topology

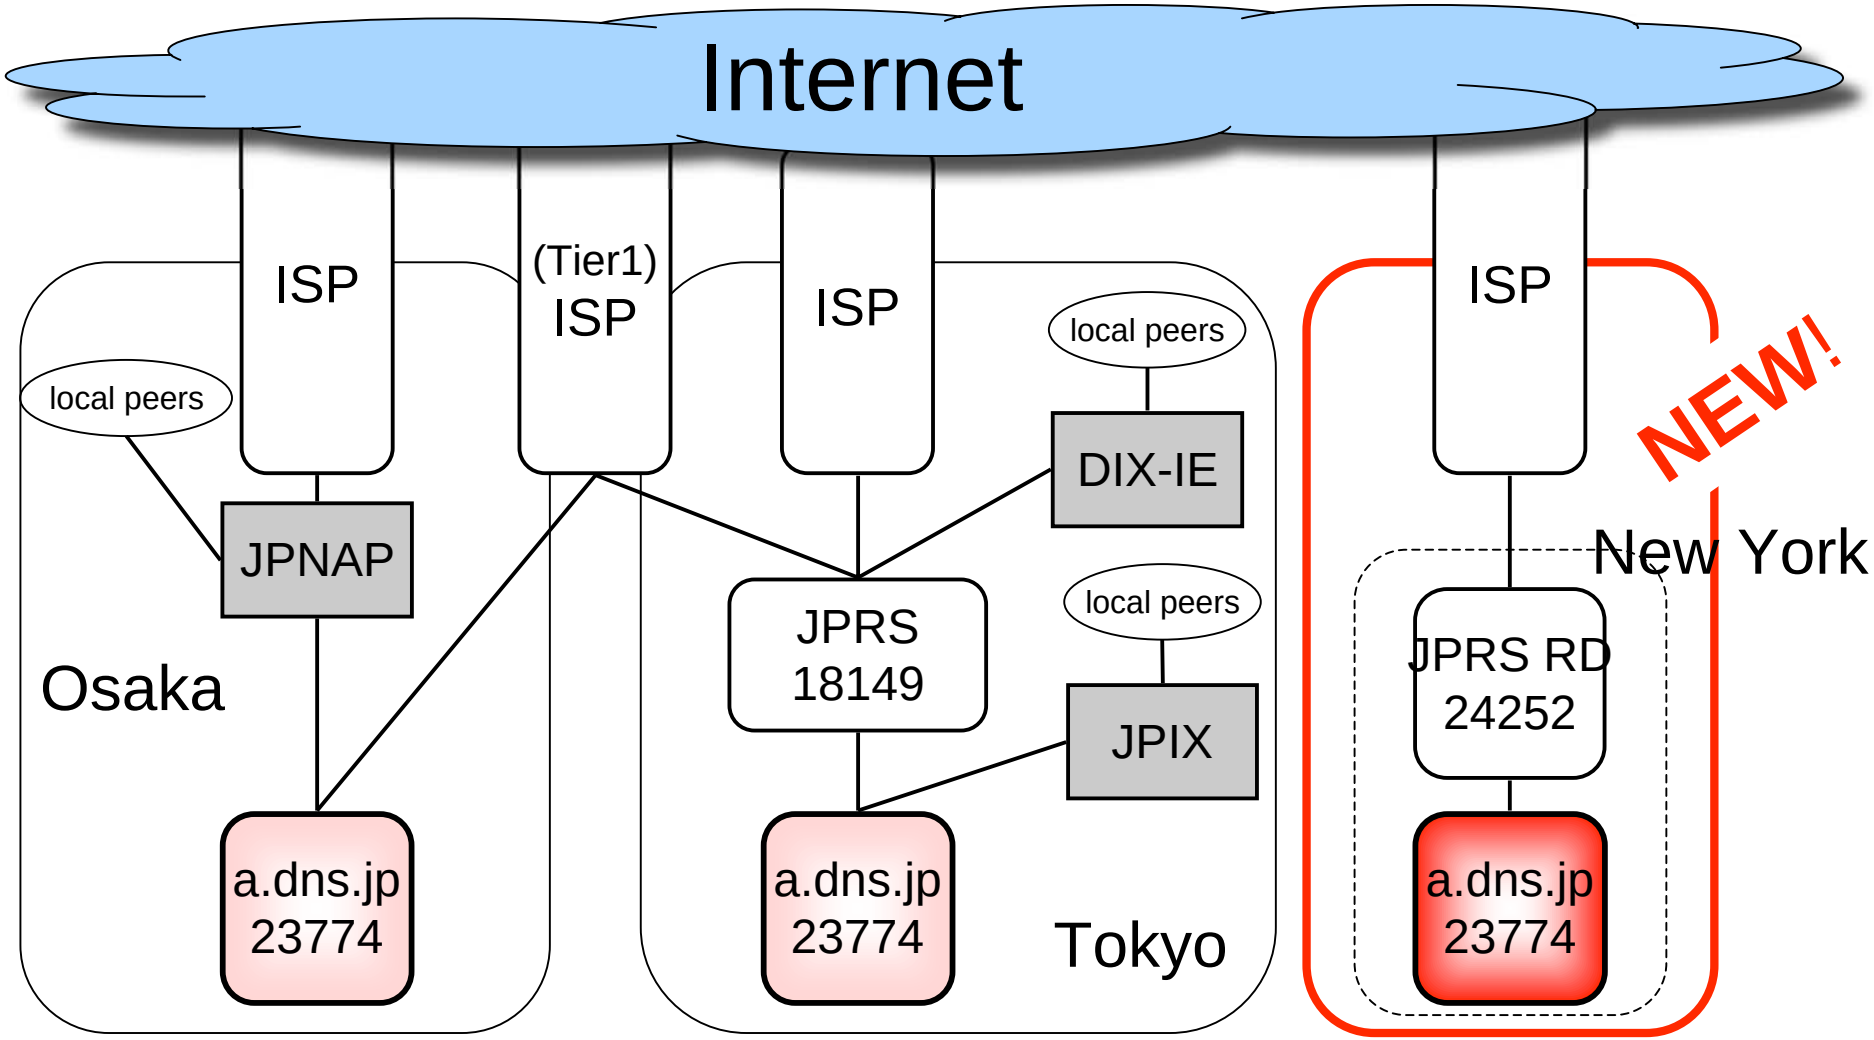

a.dns.jp - Network Topology (cont.)

- Path length between the Internet (Tier1 AS) and a.dns.jp (AS23774)
  - Tokyo: 2, Behind JPRS AS 18149
  - Osaka: 1, Direct
  - New York: 2, Behind JPRS RD AS 24252
- Osaka's BGP path length is shorter than other locations from the Internet.

Anycast effect measurement

- Period: 20days in February, 2007
- Method:
  - Record all queries (using BIND 9 logging function)
  - Extract query source IPv4 addresses from the query log before and after one day (24hours) of each change.
    - Note: IPv6 was not examined.
  - Analyze query source by countries

a.dns.jp - Anycast Test Run Step

1. one AS-path prepend at Osaka
2. New York: Turn ON

Today's report

3. Osaka: three AS-paths prepend at Osaka
4. Osaka: Turn OFF
5. Osaka: Turn ON with normal AS-path length
6. Osaka: Turn OFF
7. Osaka: Turn ON with normal AS-path length
8. New York: Turn OFF

Summary of a.dns.jp queries

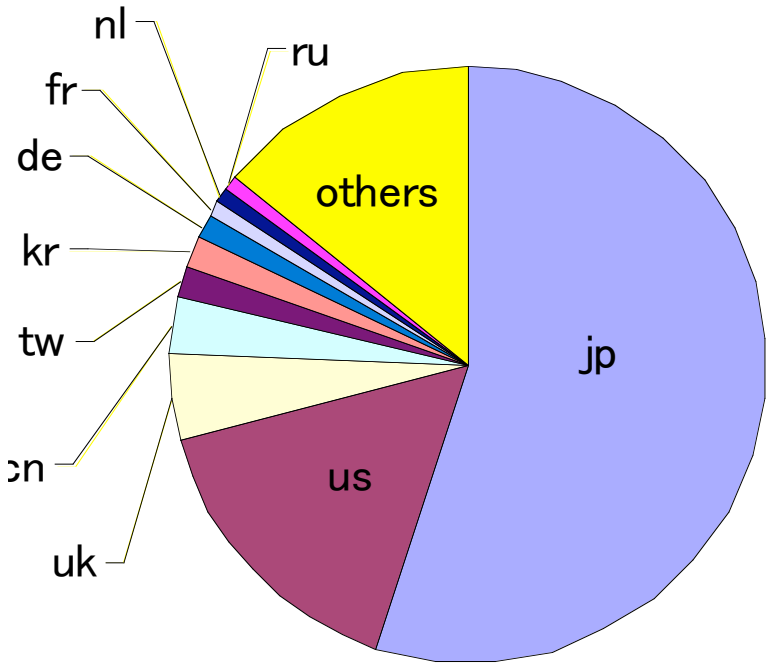

Queries/countries

- Average 1900qps
  - 1500 to 2500 qps
  - JP 83%　ARPA 17%
- Ranking

1.	JP	55%
2.	US	17%
3.	UK	5%
4.	CN	3.5%
5.	TW	2%

Anycast distribution chart

- Try to display anycast effects
- This chart was drawn by the queries ratio that reached each Anycast node per countries.
- Country's brightness indicates query rates (%) for each anycast node.
- For each country, total of all graph value is 100%.
- This example chart shows that most queries reached Osaka node for almost all countries.

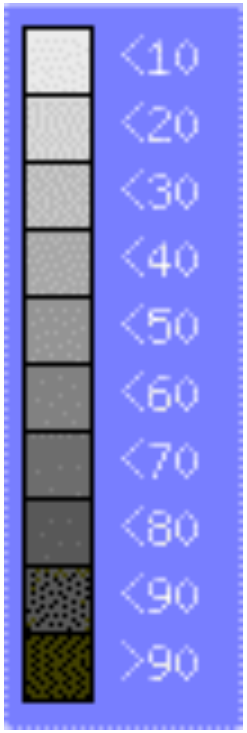

Osaka

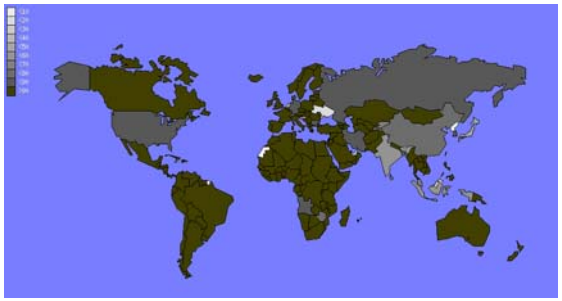

65%

Tokyo

35%

Step1: AS-path prepend at Osaka

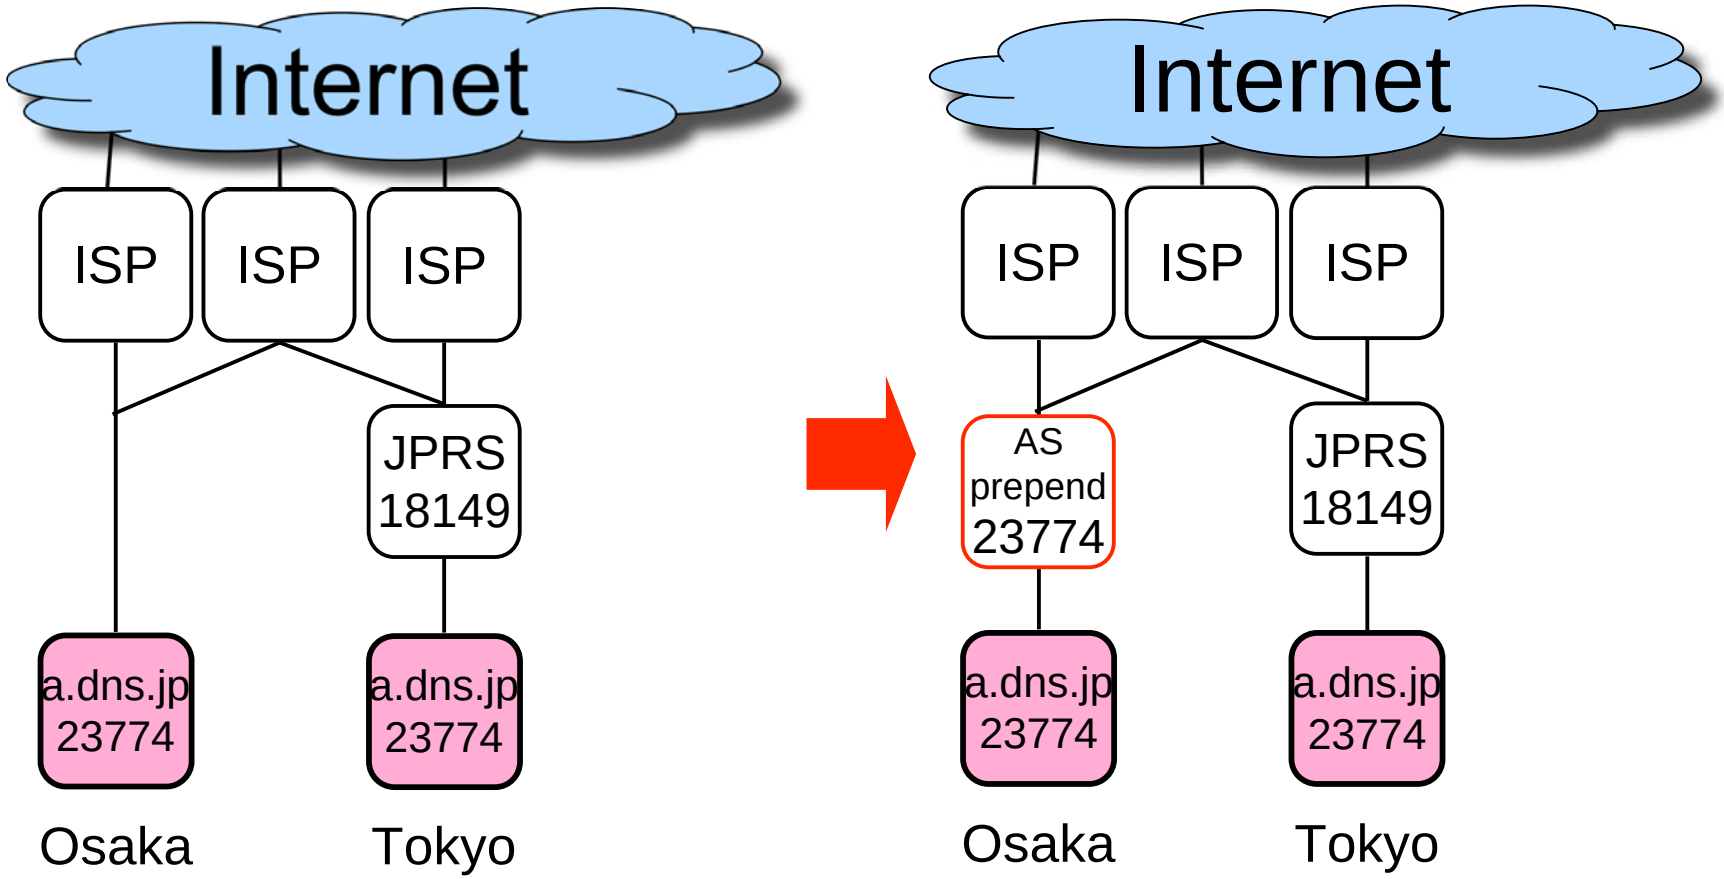

Changes between step 1 (1 hour)

Osaka

62%

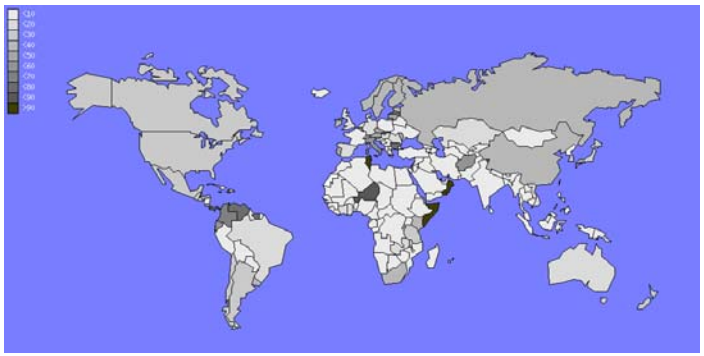

23%

Tokyo

38%

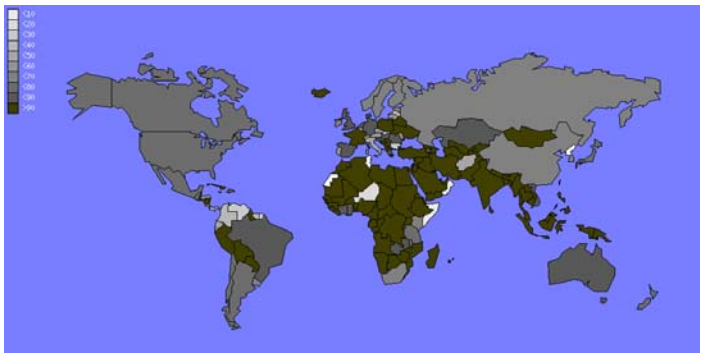

77%

Changes between step 1 (24hours)

Osaka

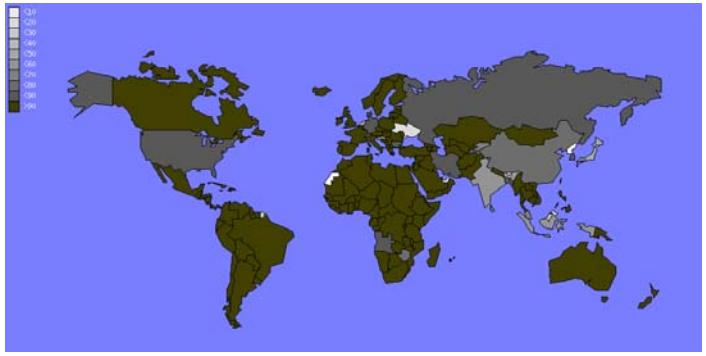

65%

22%

Tokyo

35%

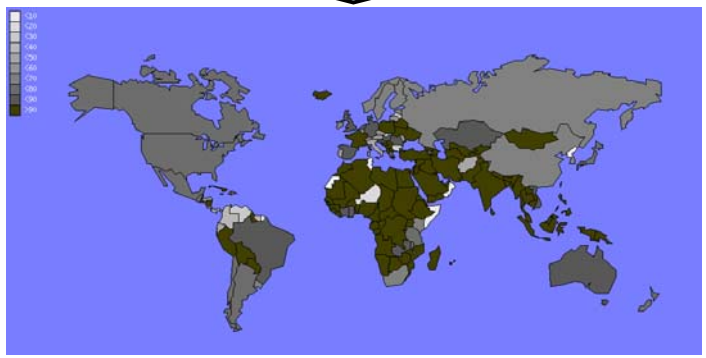

78%

Result of AS-path prepend

- Traffic trend was changed immediately
  - Many queries to Osaka were moved to Tokyo
- AS-path prepend works well for traffic control.

Step 2: Starting New York

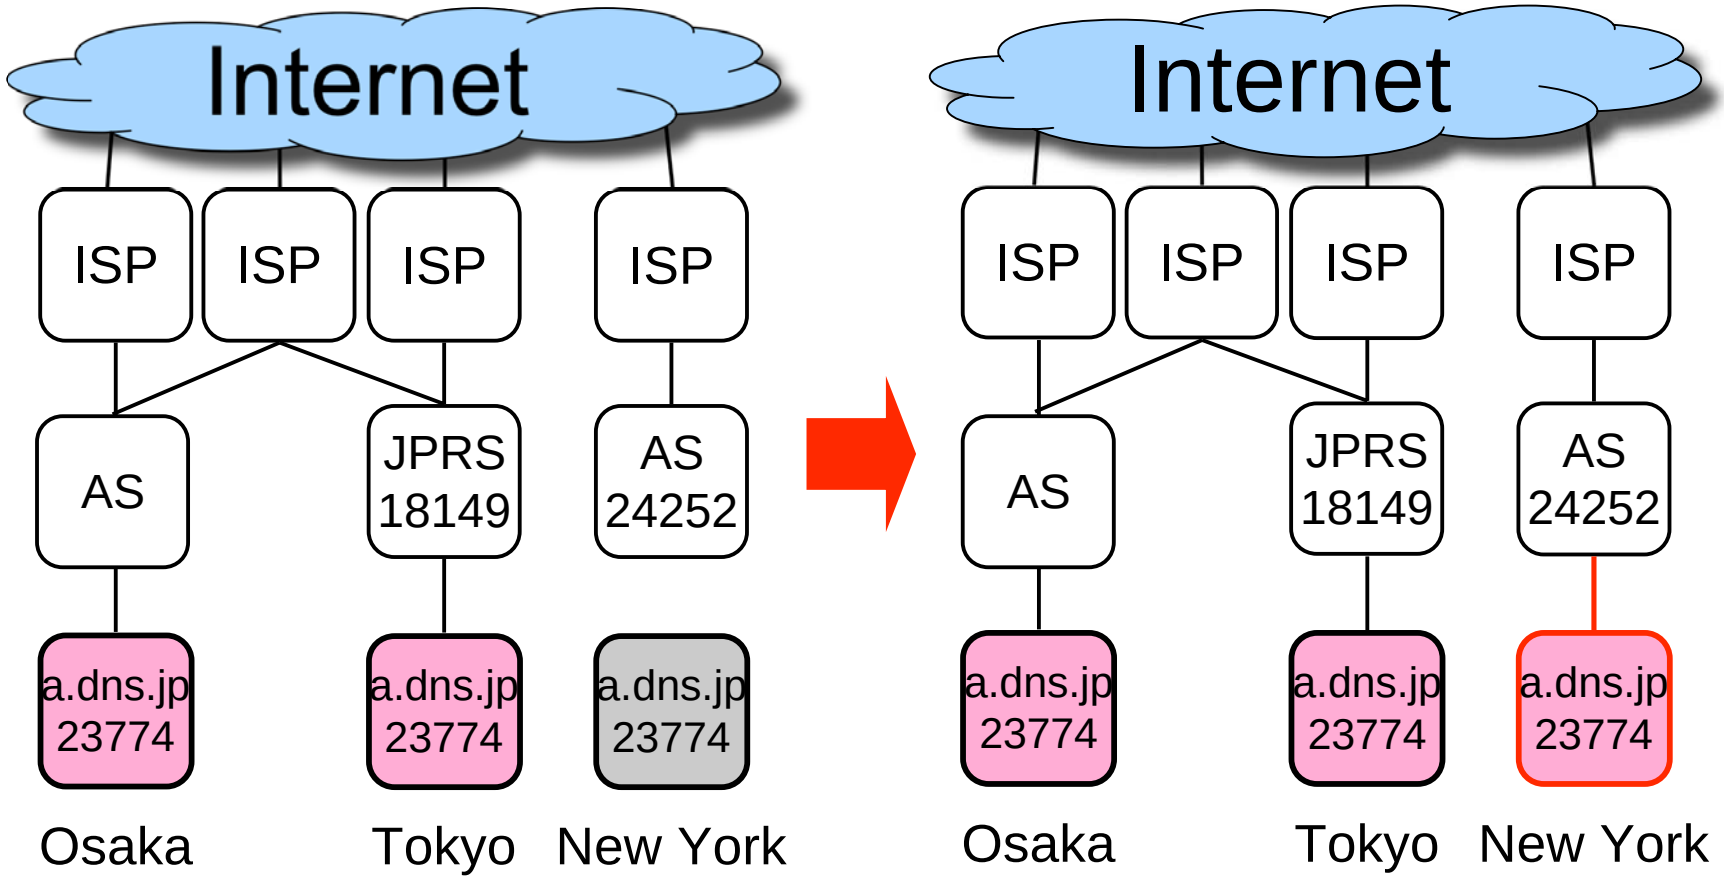

Result of starting New York (1): from RIPE DNSMON

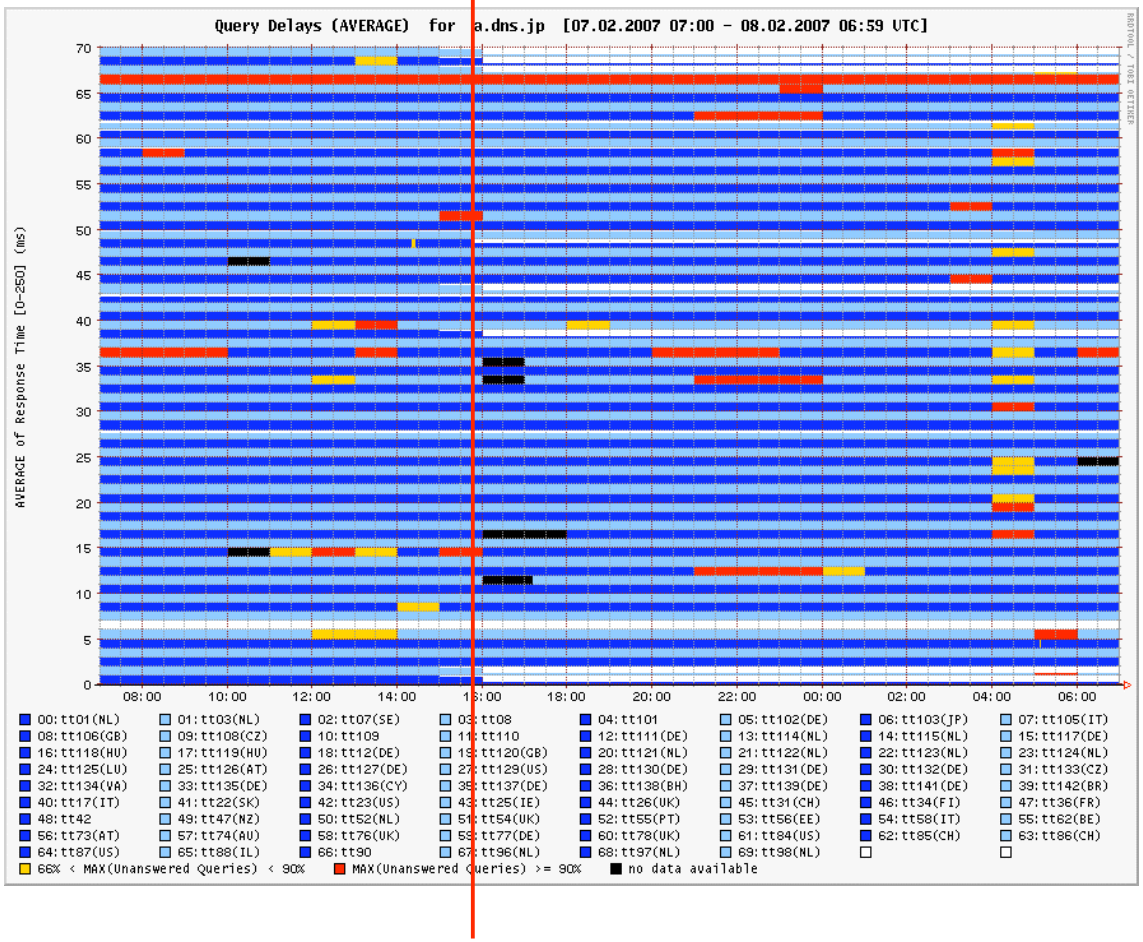

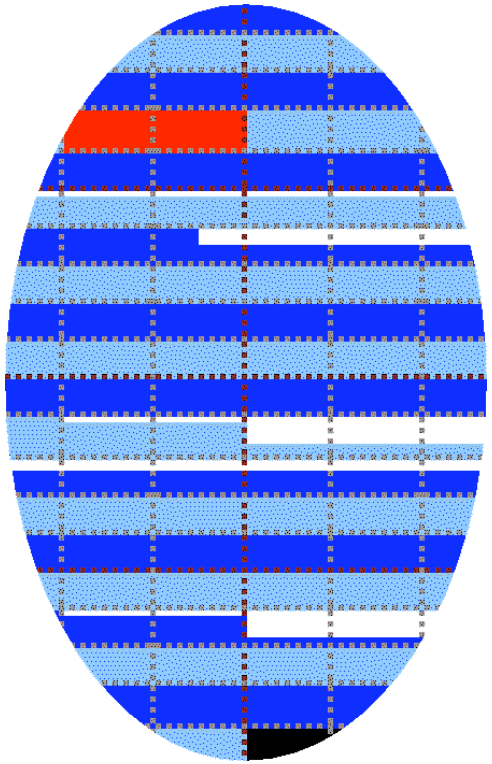

Magnified

Result of starting New York(2): Anycast distribution chart

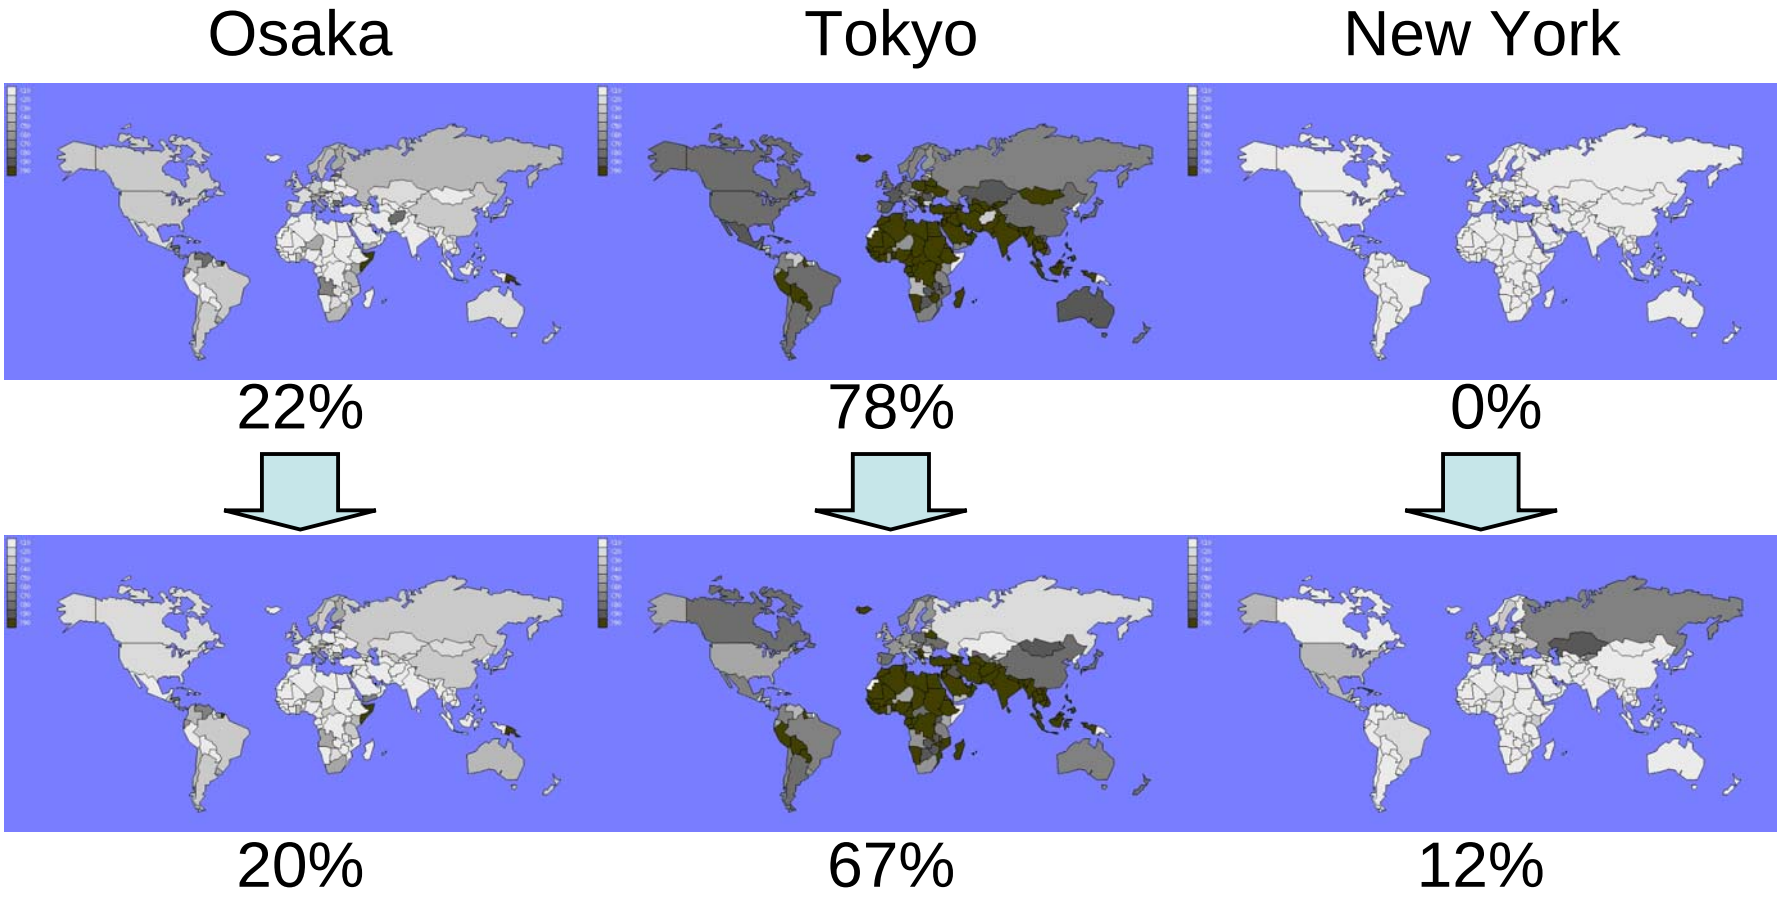

Result of starting New York(3): Asia

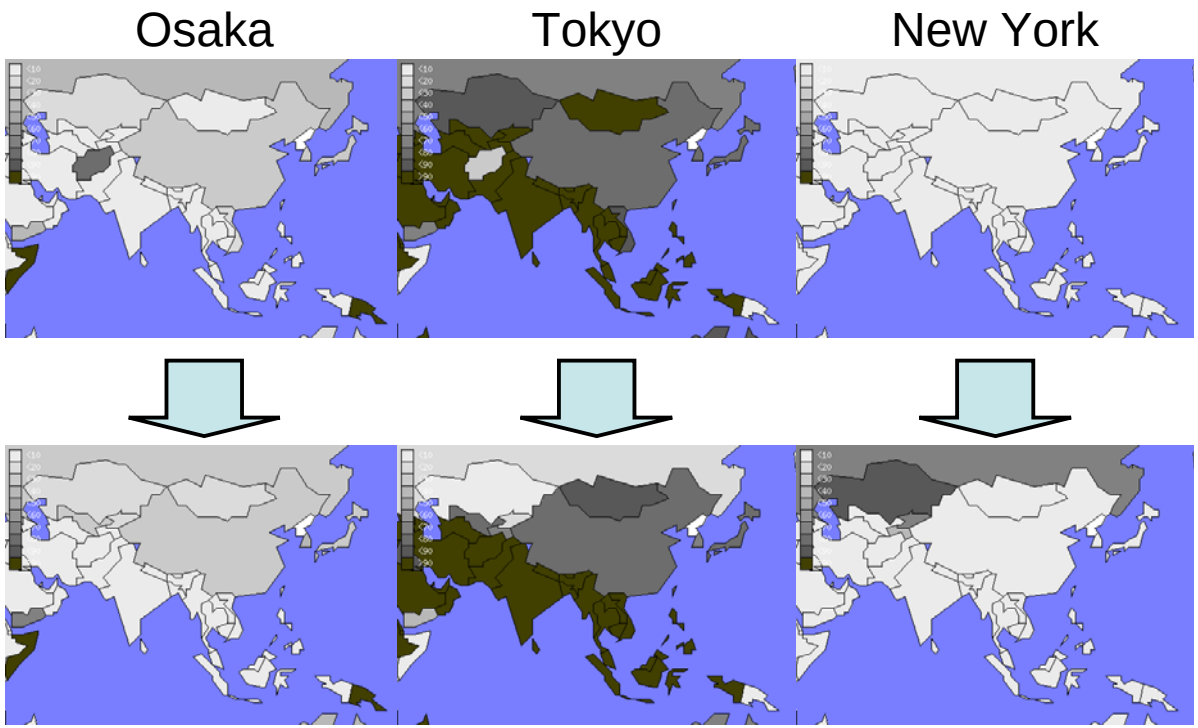

- not changed in Asia.

  (except in Russia)

- Most traffic may pass direct or suitable links to Japan

Result of starting New York(4): Europe

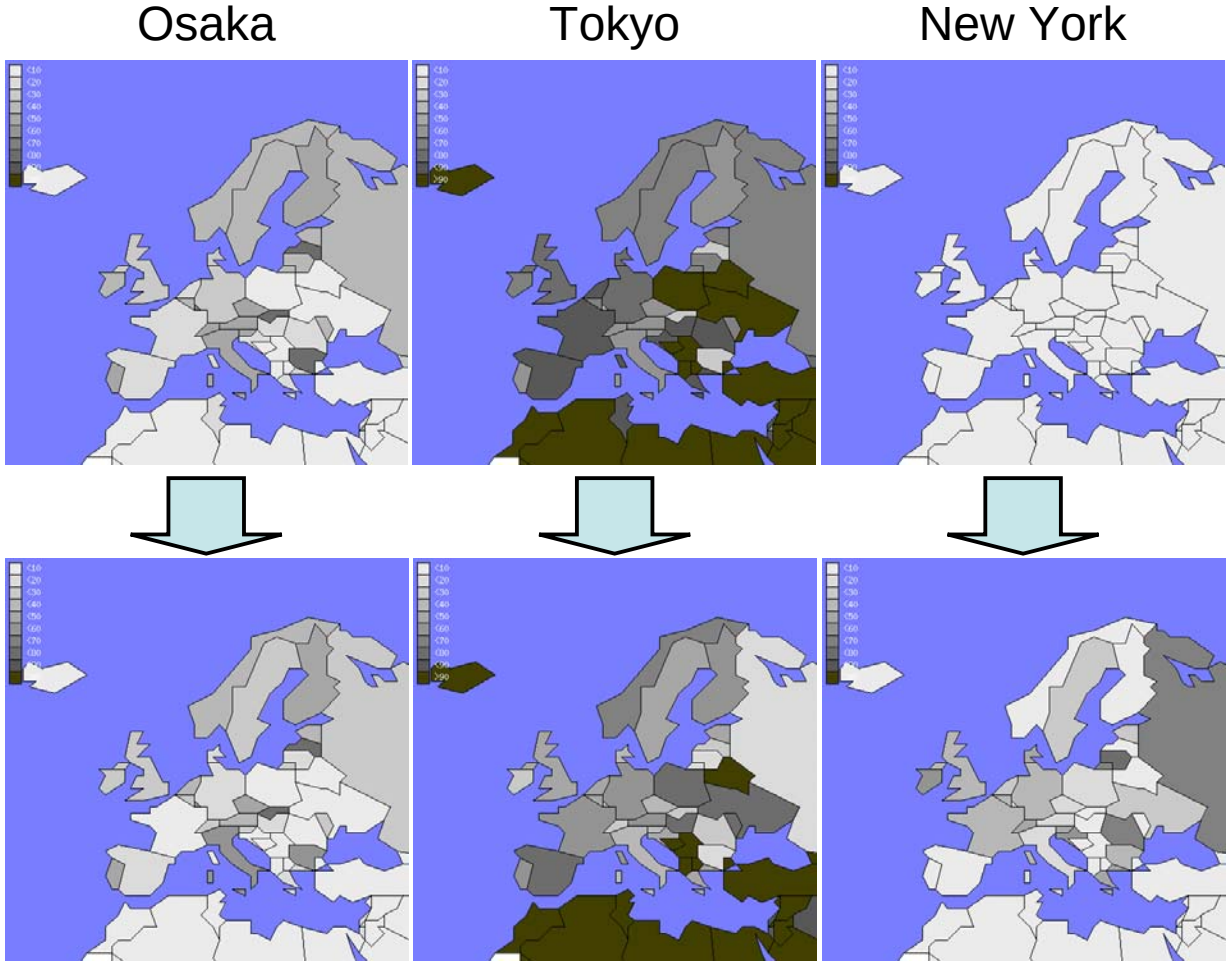

- Queries to Tokyo were partially reduced.
- A decrease in Tokyo was moved to New York.

Result of starting New York(5): South America

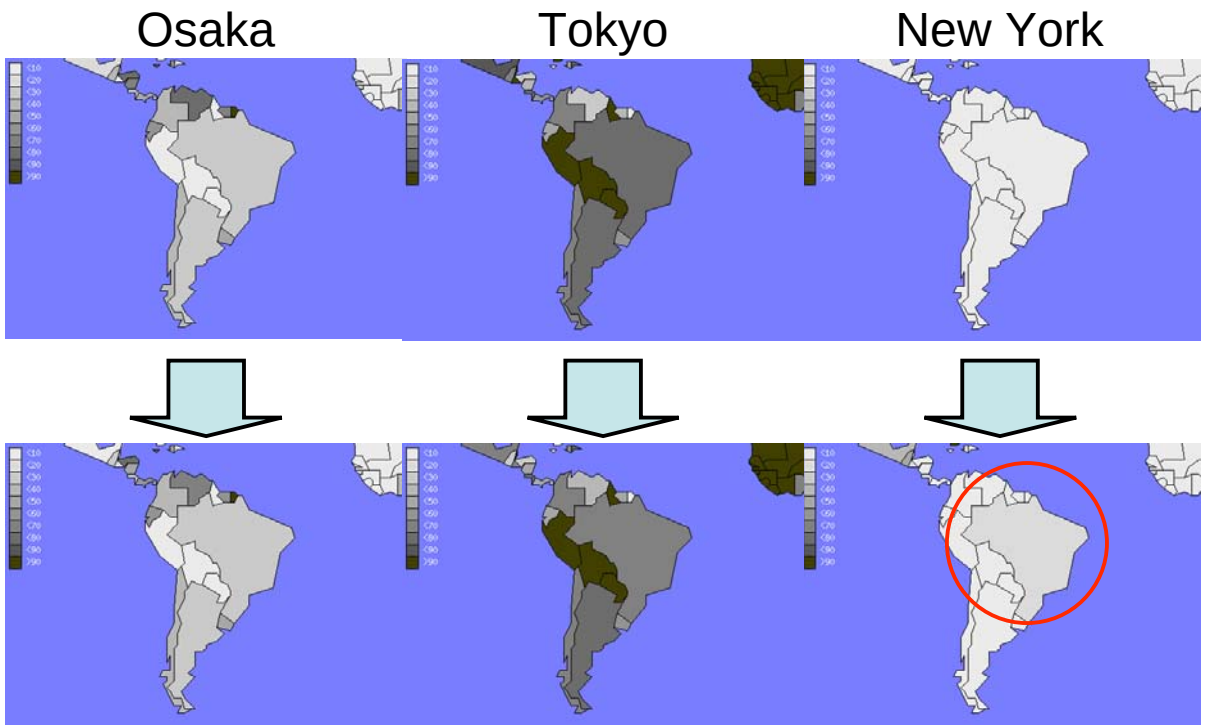

- No large change, but some queries were moved to New York.

Result of starting New York (6)

- Though a.dns.jp is connected only one ISP and 1 more AS path at New York, 12% queries were moved to New York immediately.
- From RIPE DNSMON, latencies from some European probes are remarkably reduced.

Conclusion/Comments?

- BGP anycast is useful for traffic control and reducing latency.
  - AS-path prepend works well for IP Anycast traffic control
  - Anycast node change is immediately reflected.
- Anycast distribution chart was proposed to display anycast effectiveness.

Tools for this analysis

- BIND 9 query log function
- Maxmind GeoIP

  http://www.maxmind.com/app/ip-location
- HELIO World

  http://www.helio.org/world/
- RIPE DNSMON

  http://dnsmon.ripe.net/