

Measurement of BGP Anycast effects – experiences in .JP

From the query to a.dns.jp

Yoshiro YONEYA <yone@jprs.co.jp>

Masato Minda, Shinta Sato, Izuru Shirai, Takayasu Matsuura

Japan Registry Services Co., Ltd. (JPRS)

IEPG Meeting@68th IETF

About JP DNS

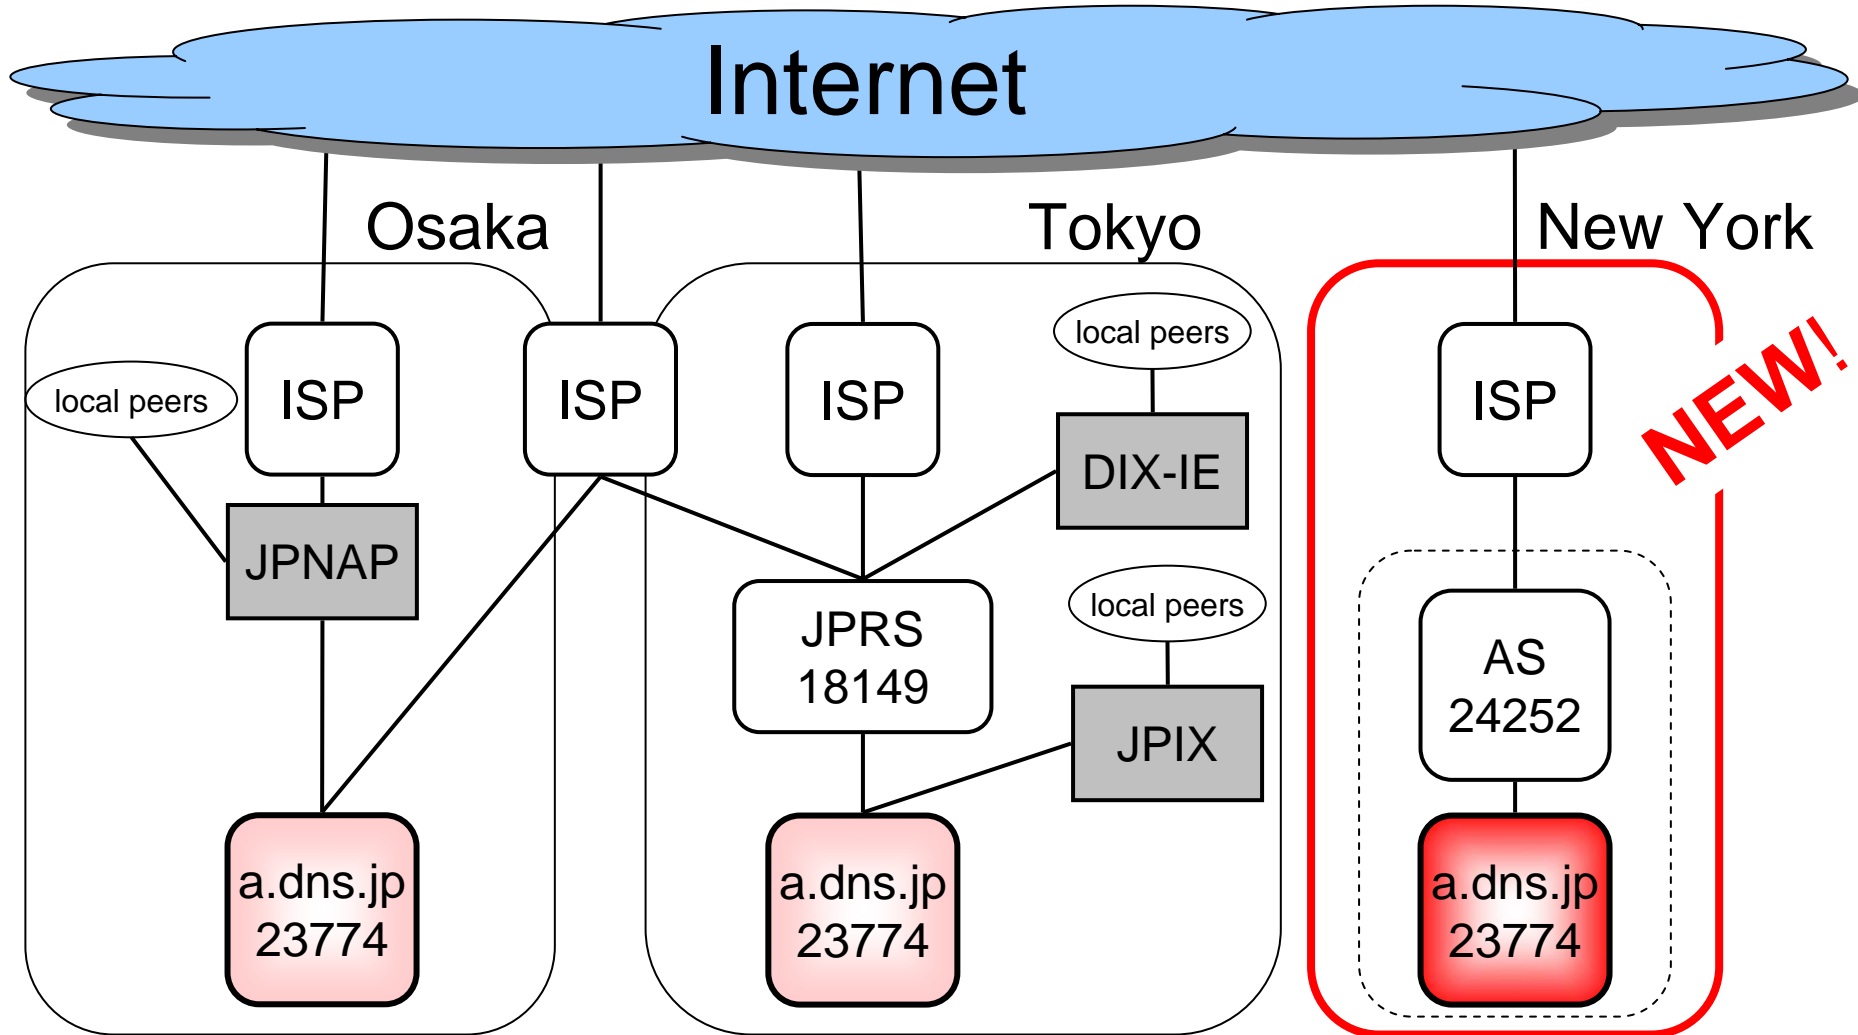
JP DNS - authoritative nameservers of .JP
a.dns.jp - f.dns.jp (except "c")

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	SINET	N/A	Yes

a.dns.jp - Anycast Status

- a.dns.jp
 - AS: 23774 IPv4: 203.119.1.1 IPv6: 2001:dc4::1
 - Using BIND 9 (All of JP DNS are using BIND 9)
- Locations of a.dns.jp
 - Tokyo and Osaka (since Feb 2004)
 - New York will up by the end of 2007 (tentatively)
- Objectives of this experiment
 - Test run and measurement before real operation at New York

a.dns.jp - Network Topology



18 Mar 2007

Copyright©2007 Japan Registry Services Co., Ltd.


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

- Path between the Internet and a.dns.jp (AS23774)
 - Tokyo: Two ASes, AS18149 and ISP
 - Osaka: One AS, ISP
 - New York: Two ASes, AS24252 and ISP

a.dns.jp – Measurement & Tools

- Query Log
By BIND logging feature (since Feb 2004)
- Maxmind GeoIP
<http://www.maxmind.com/app/ip-location>
- HELIO World
<http://www.helio.org/world/>
- DSC
<http://dns.measurement-factory.com/tools/dsc/>
- RIPE DNSMON
<http://dnsmon.ripe.net/>

a.dns.jp - Summary of queries



- Ave. 2000 qps
 - 1500 - 2500 qps
- From which country
 1. JP 50 - 60%
 2. US 10 - 20%
 3. Some Asian and European Countries ~10%
 4. Others ~1%

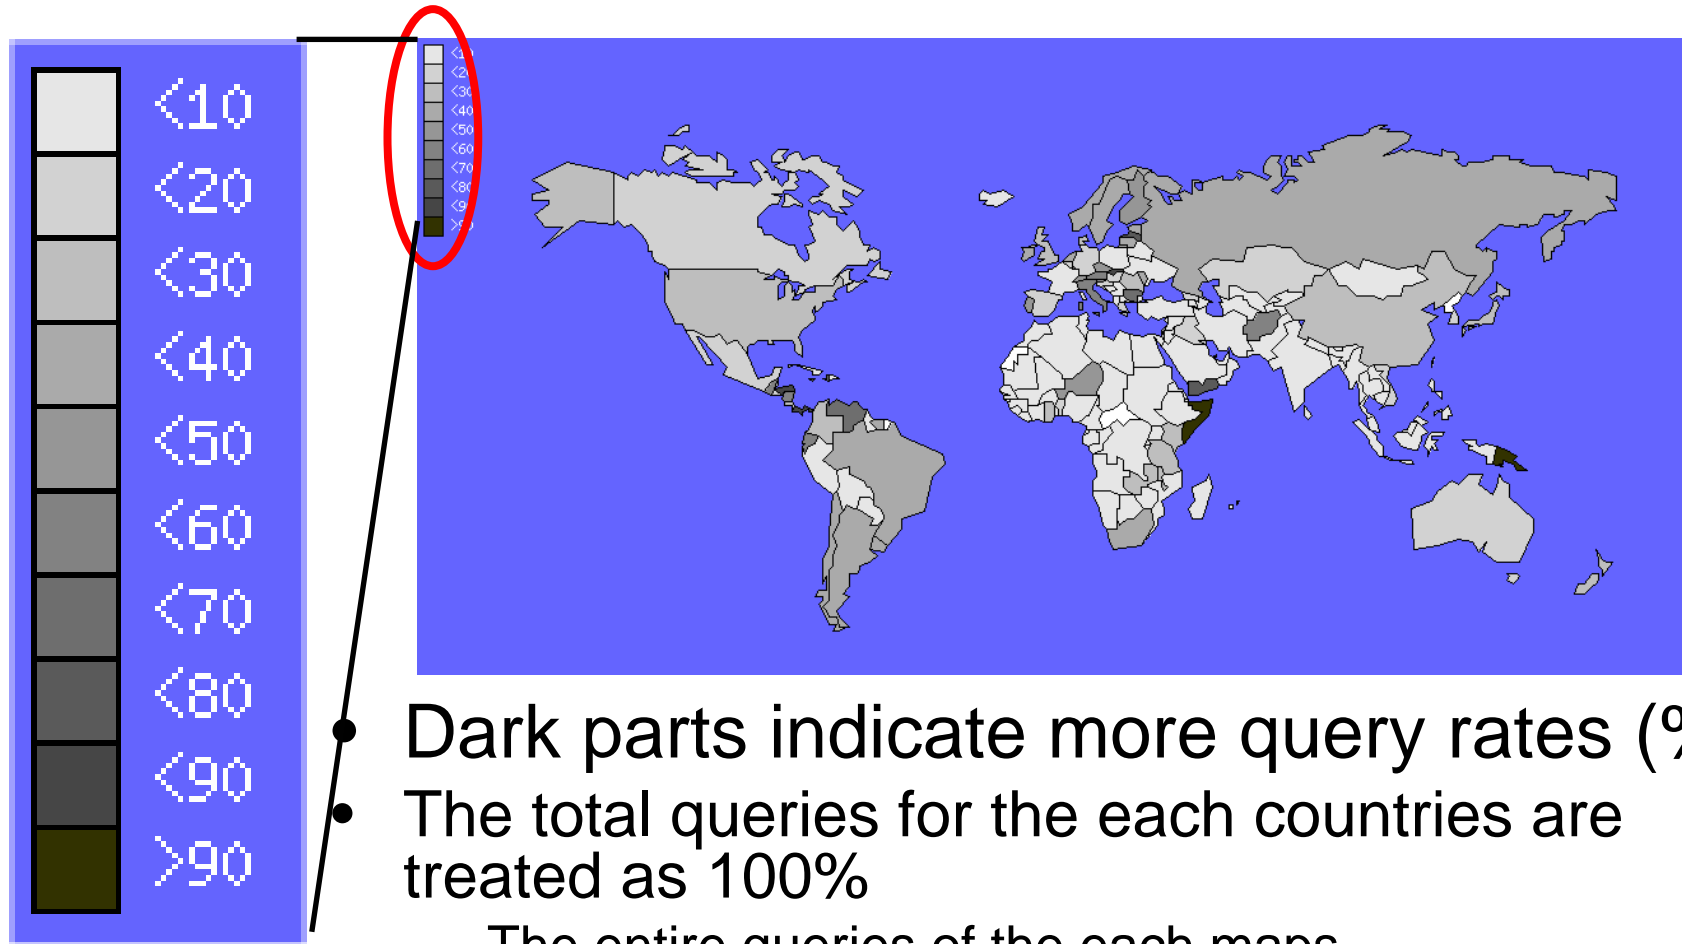
a.dns.jp - Anycast Test Run Step

1. Osaka: Add one AS-path (AS prepend)
2. New York: Turn ON
3. Osaka: Add more three AS-paths (AS prepend)
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

Schedule and Analytical Technique

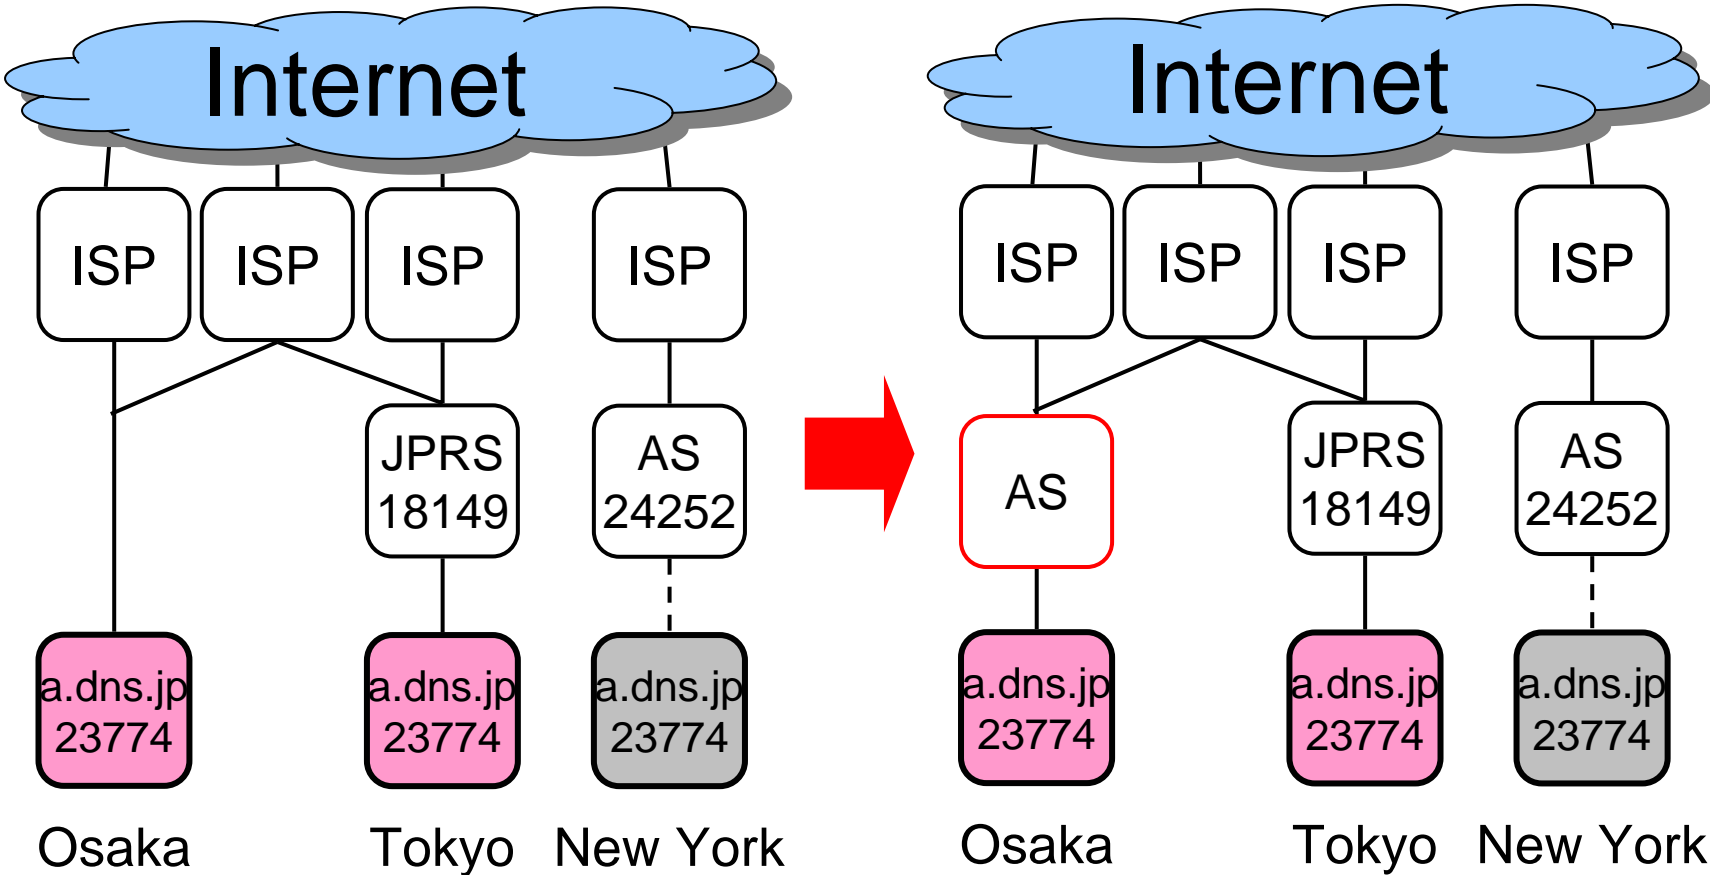
- Schedule from 2007-02-02
 to 2007-02-20
- Analytical Technique
 - Extract **IPv4** addresses (query source) from the querylog before and after one hour of each step.
 Note: IPv6 was not examined.
 - Analyze query source country by using GeoIP
 - Draw the World Map by using “HELIO World”

How to see the map




- Dark parts indicate more query rates (%)
- The total queries for the each countries are treated as 100%
 - The entire queries of the each maps do not sum up as 100%

Step1: Osaka: Add one AS-Path

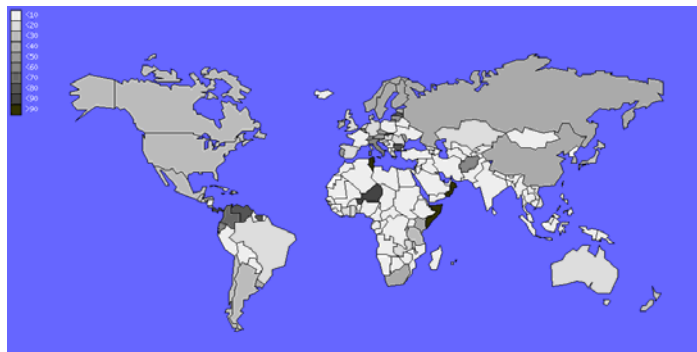


Step1: results

Osaka

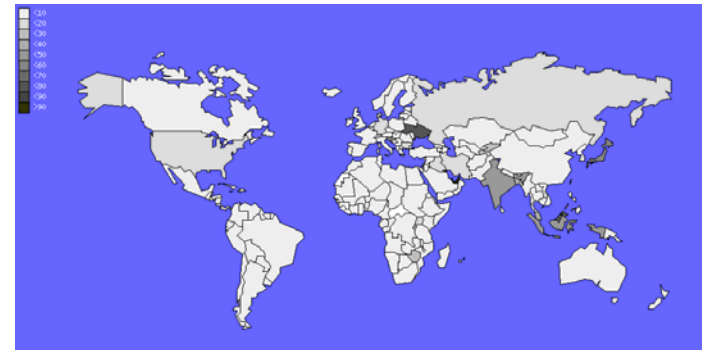


62% (1300qps)

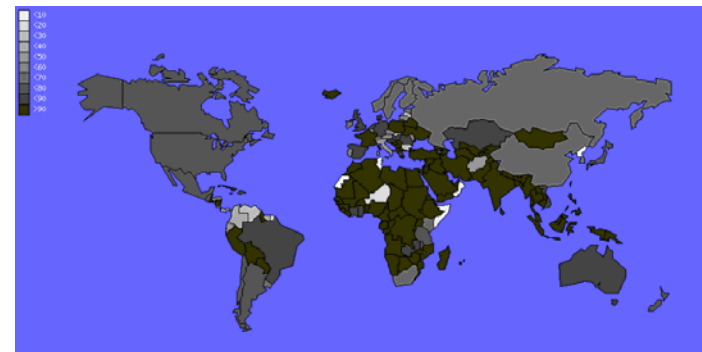


23% (460qps)

Tokyo

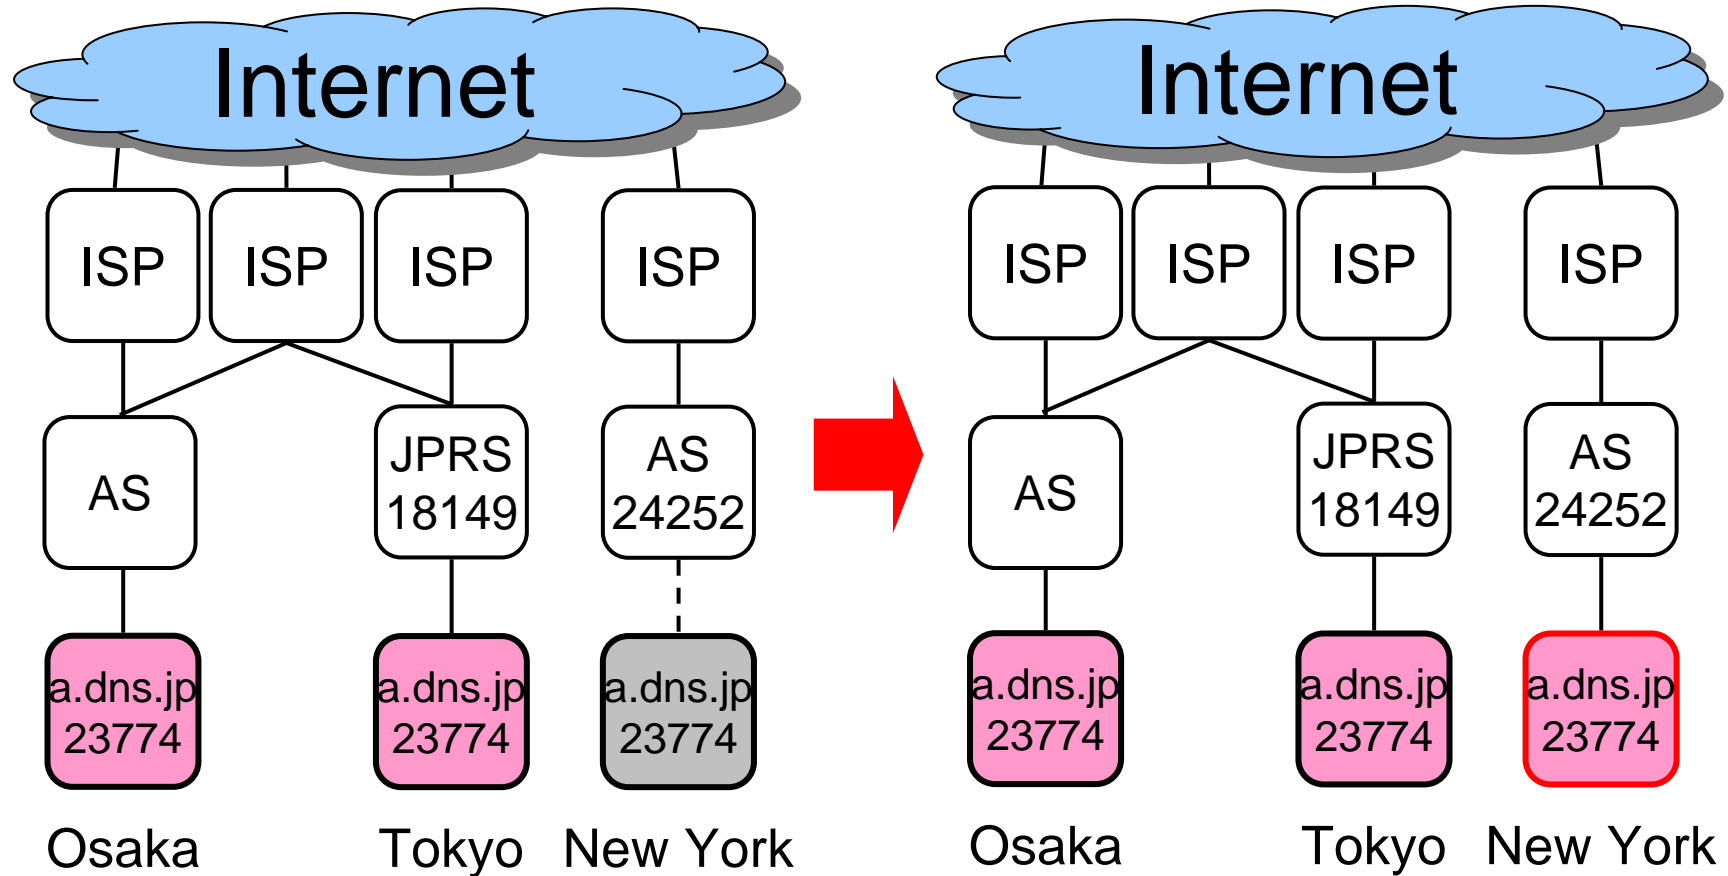


38% (800qps)




77% (1500qps)

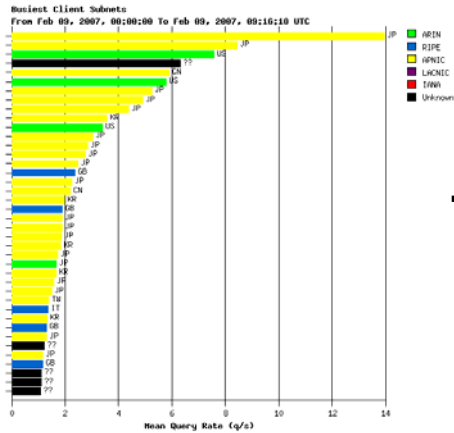
Step2: New York: Turn ON



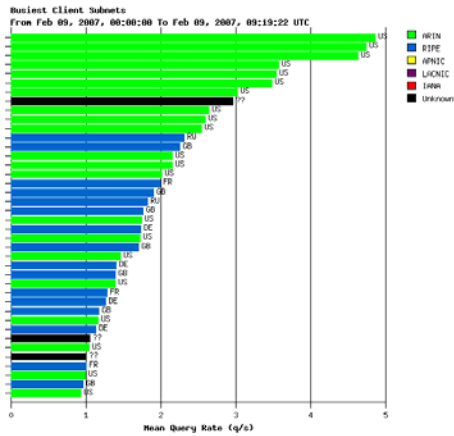
Step2: results (1)



Step2: results (2)

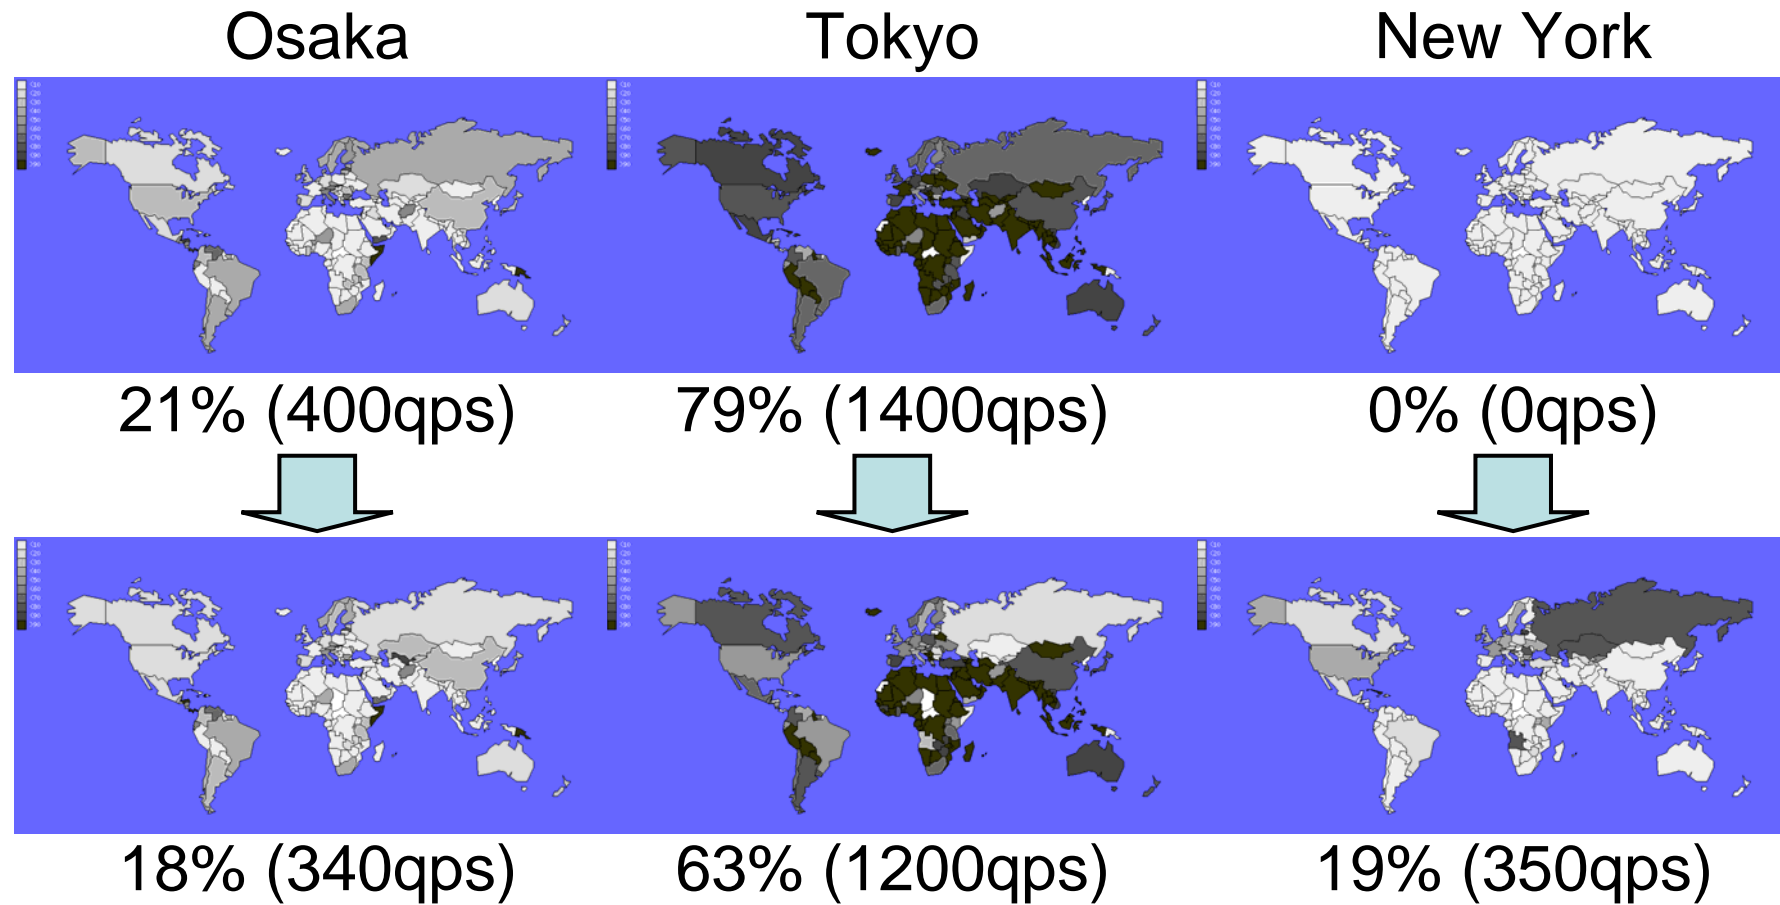


Tokyo

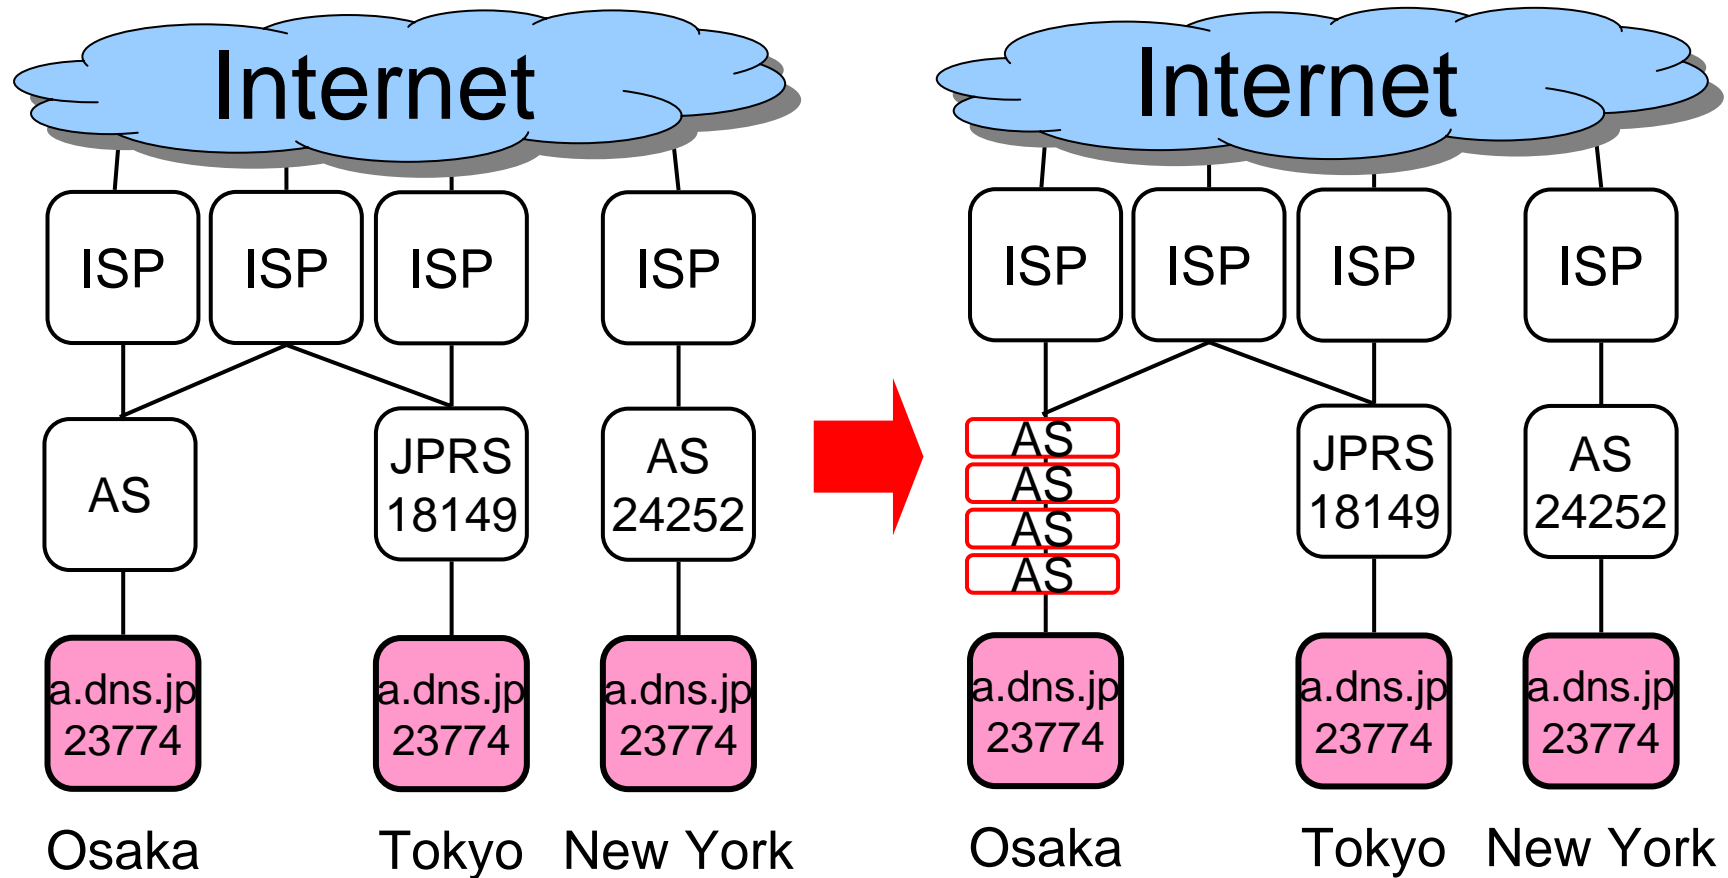


New York

Step2: results (3)



Step3: Osaka: Add three more AS-paths

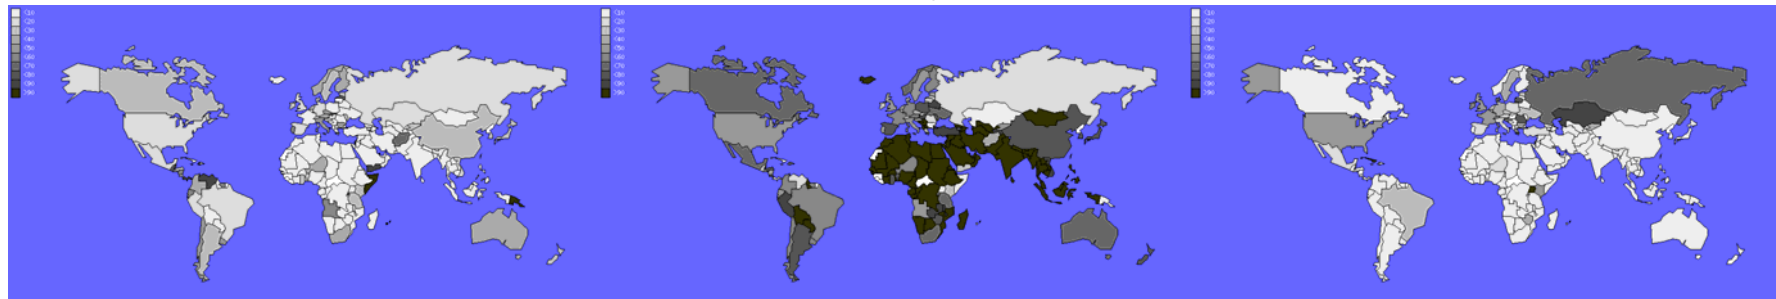


Step3: results

Osaka

Tokyo

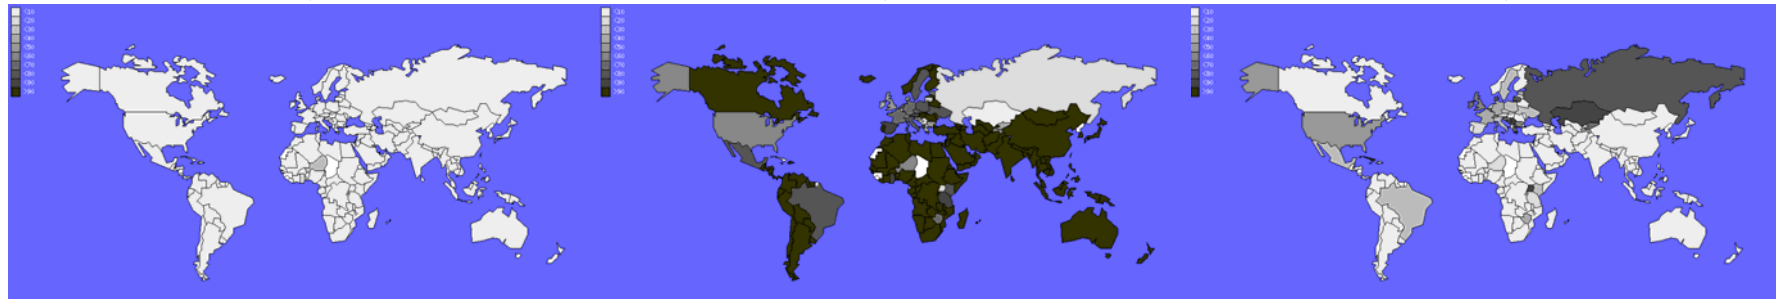
New York



18% (290qps)

64% (1000qps)

18% (290qps)

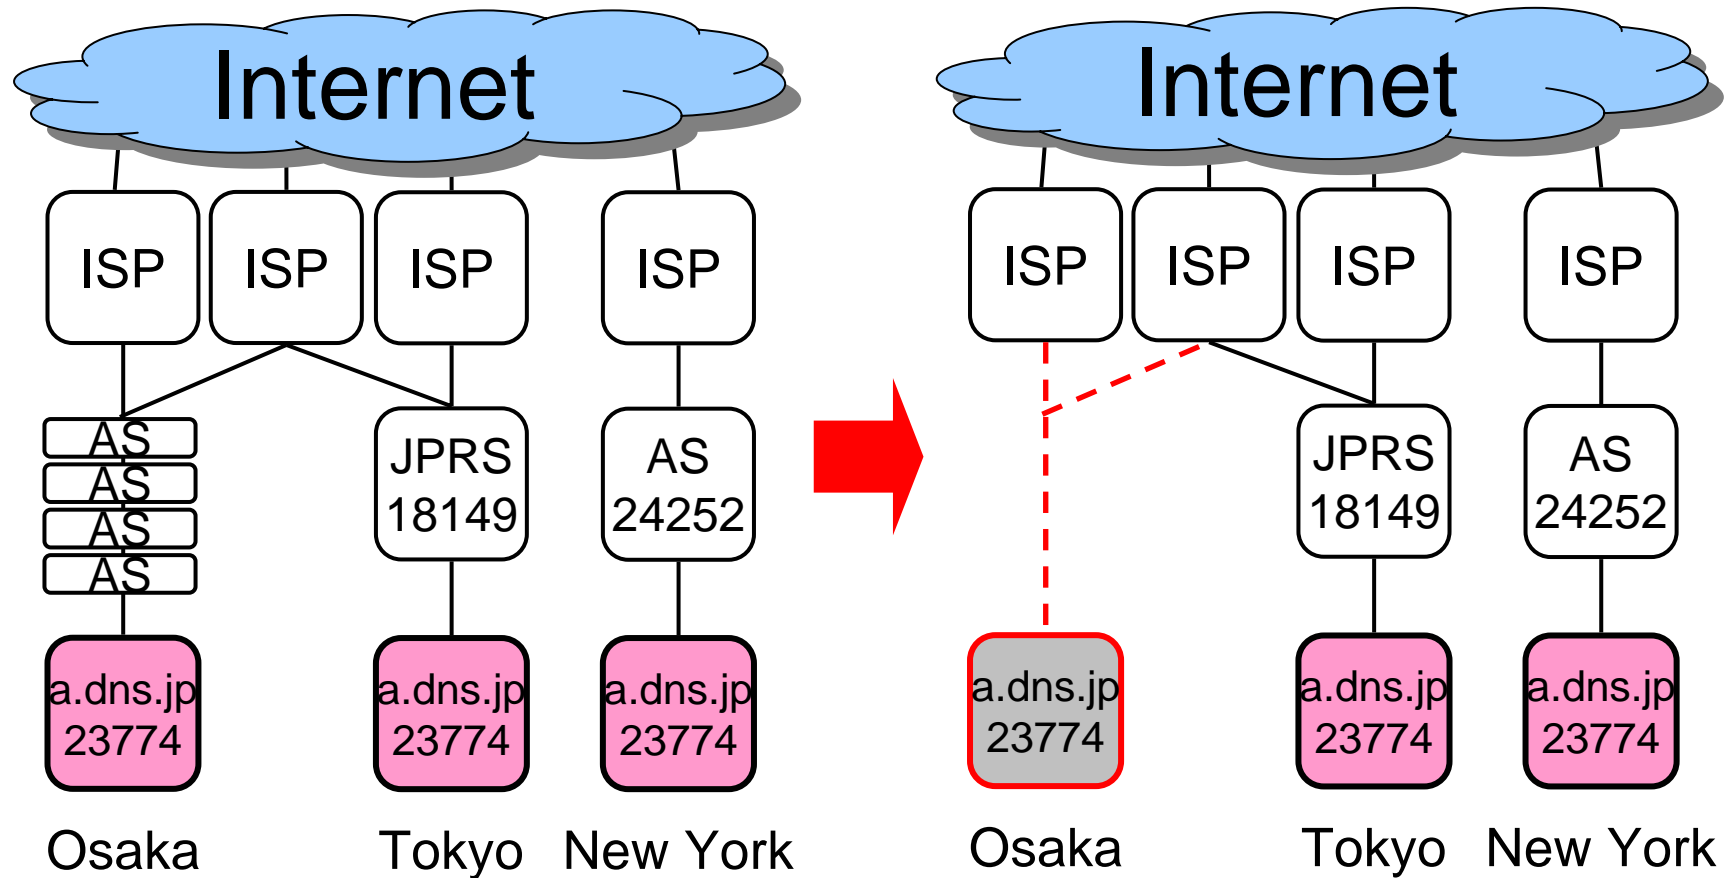


3% (60qps)

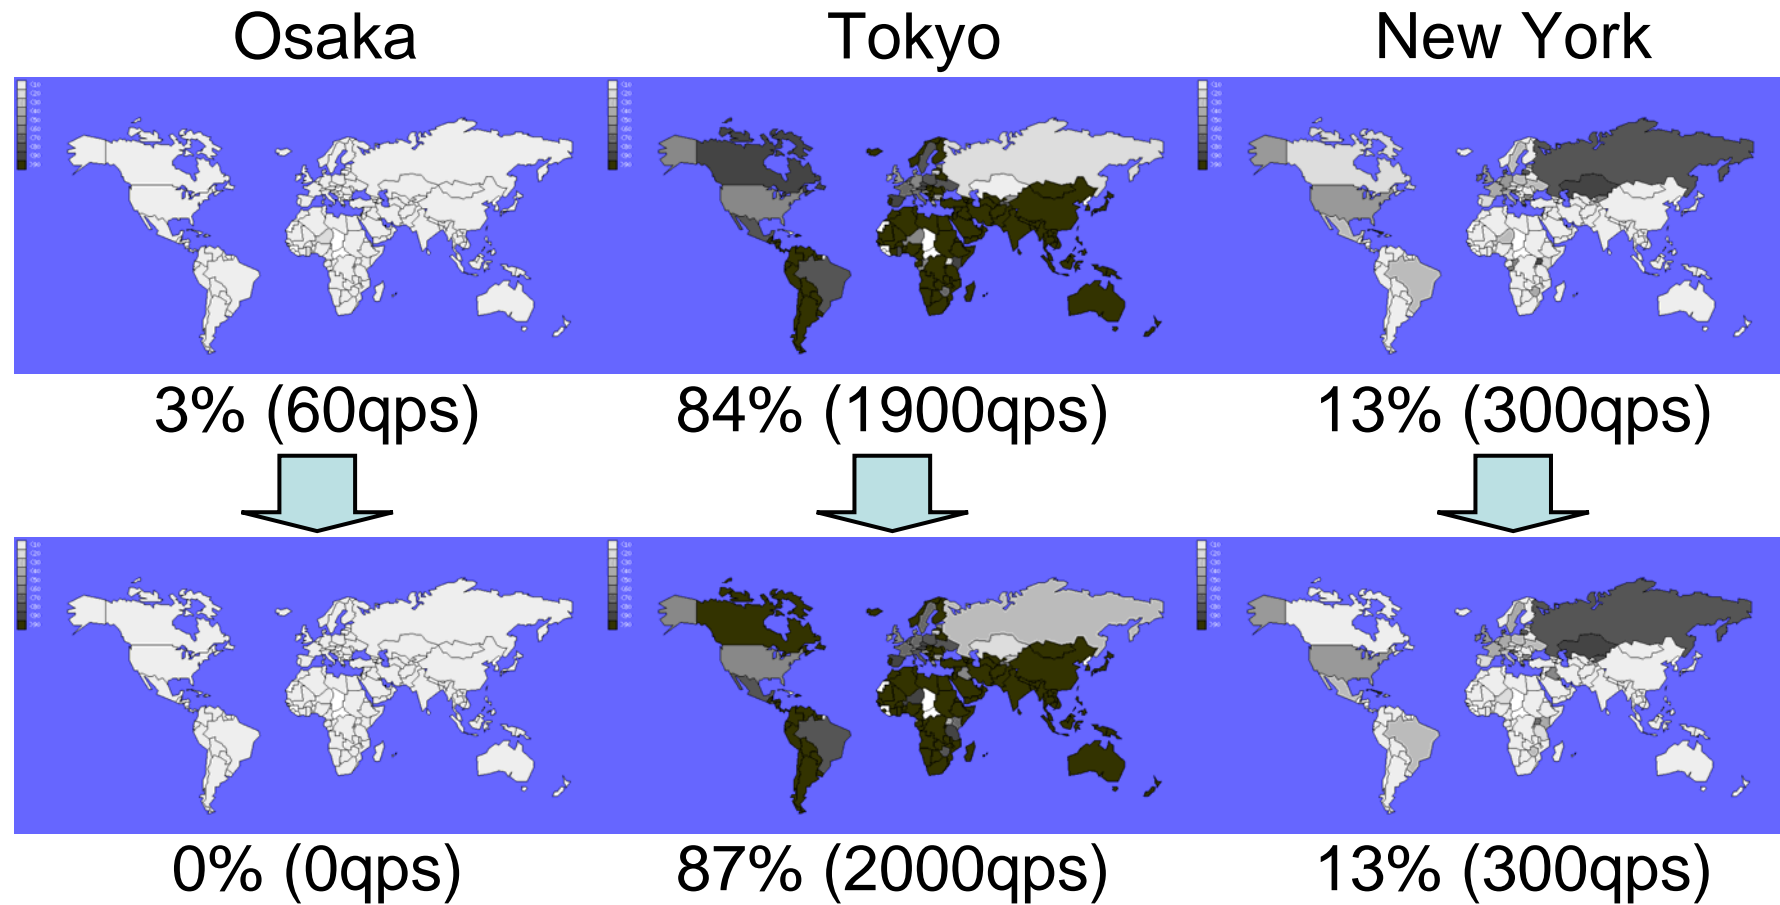
84% (1900qps)

13% (300qps)

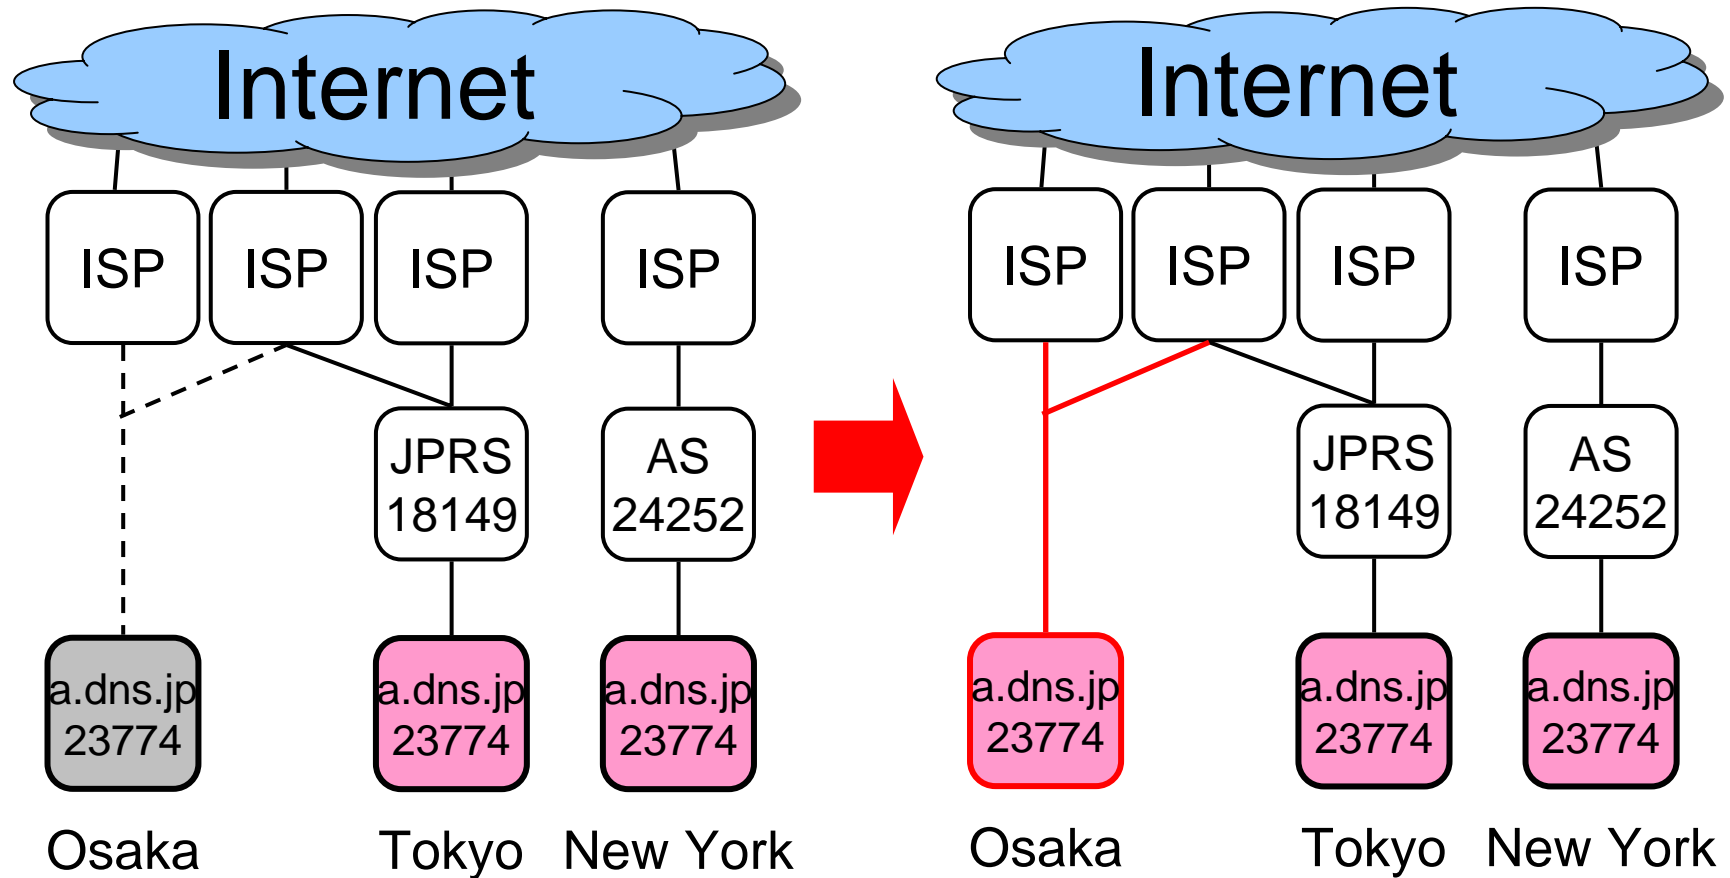
Step4: Osaka: Turn OFF



Step4: results



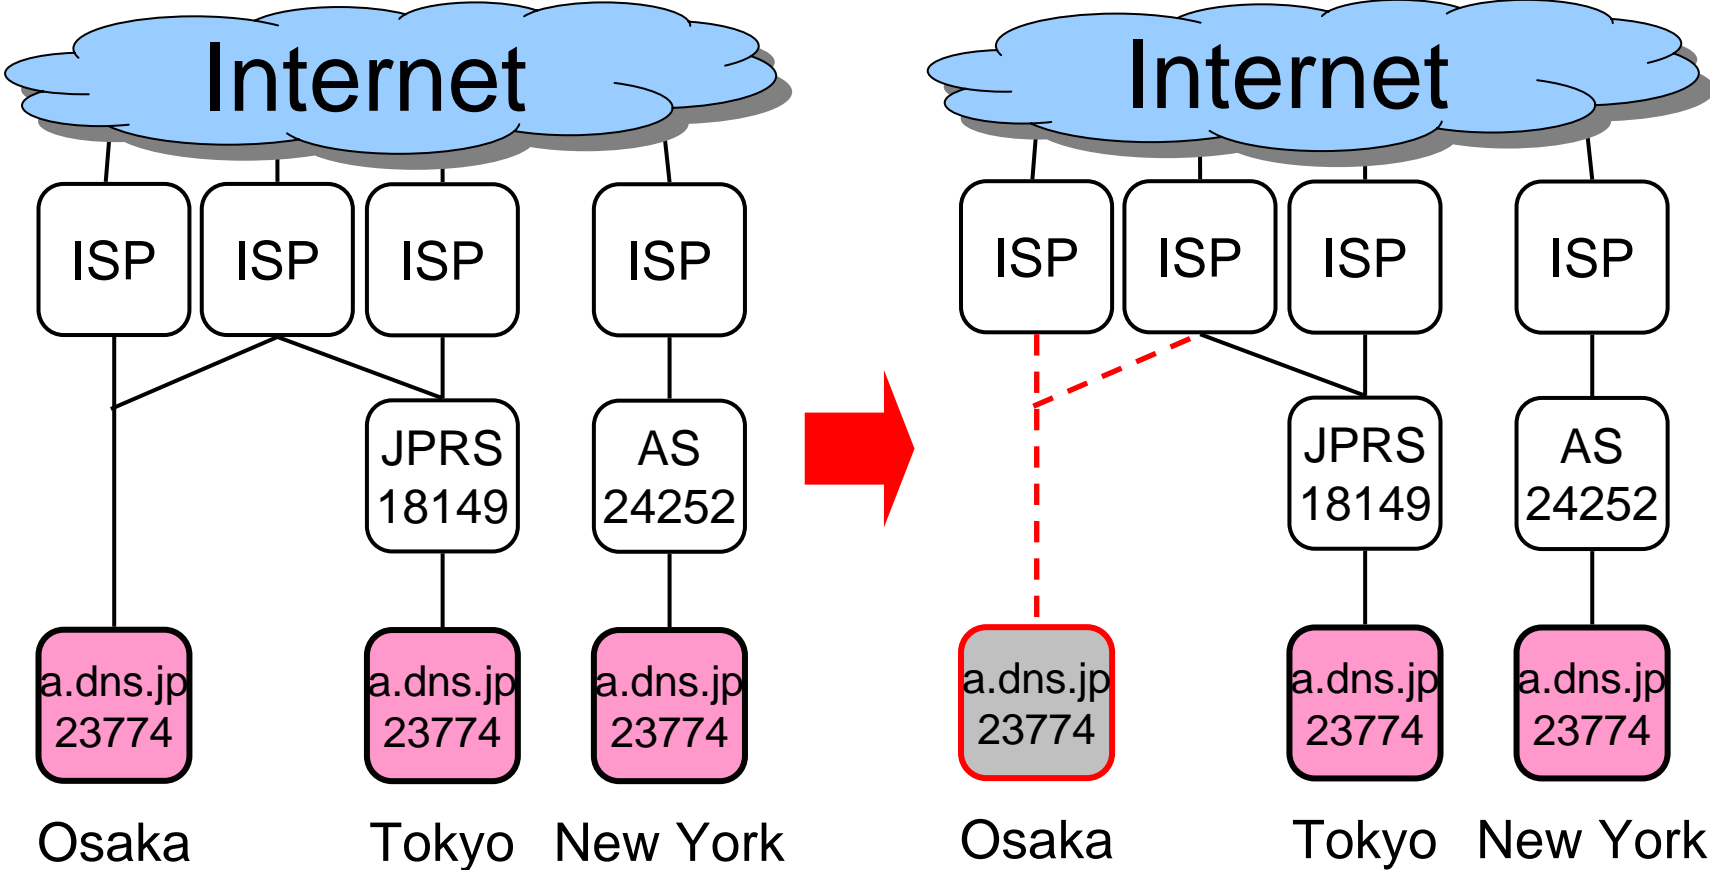
Step5: Osaka: Turn ON



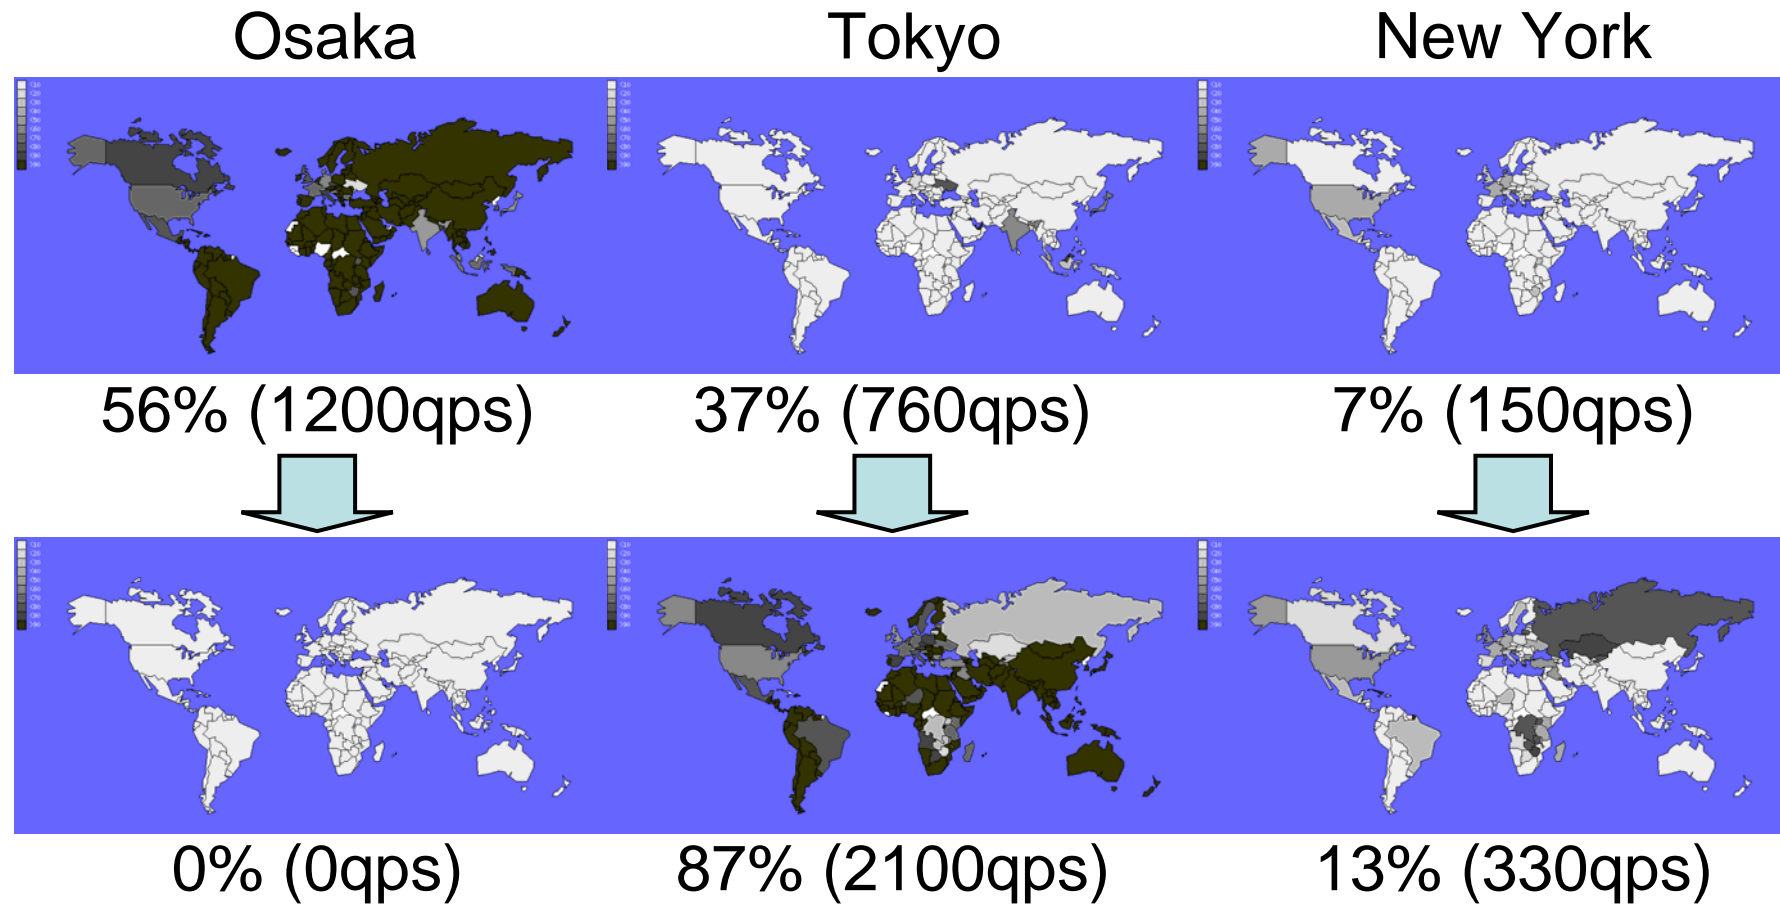
Step5: no results

- The logfile was missed
- Over 2GB logged at New York
 - We believed that 64bit version of BIND was used. However, it was 32bit version.


Step6: Osaka: Turn OFF (again)



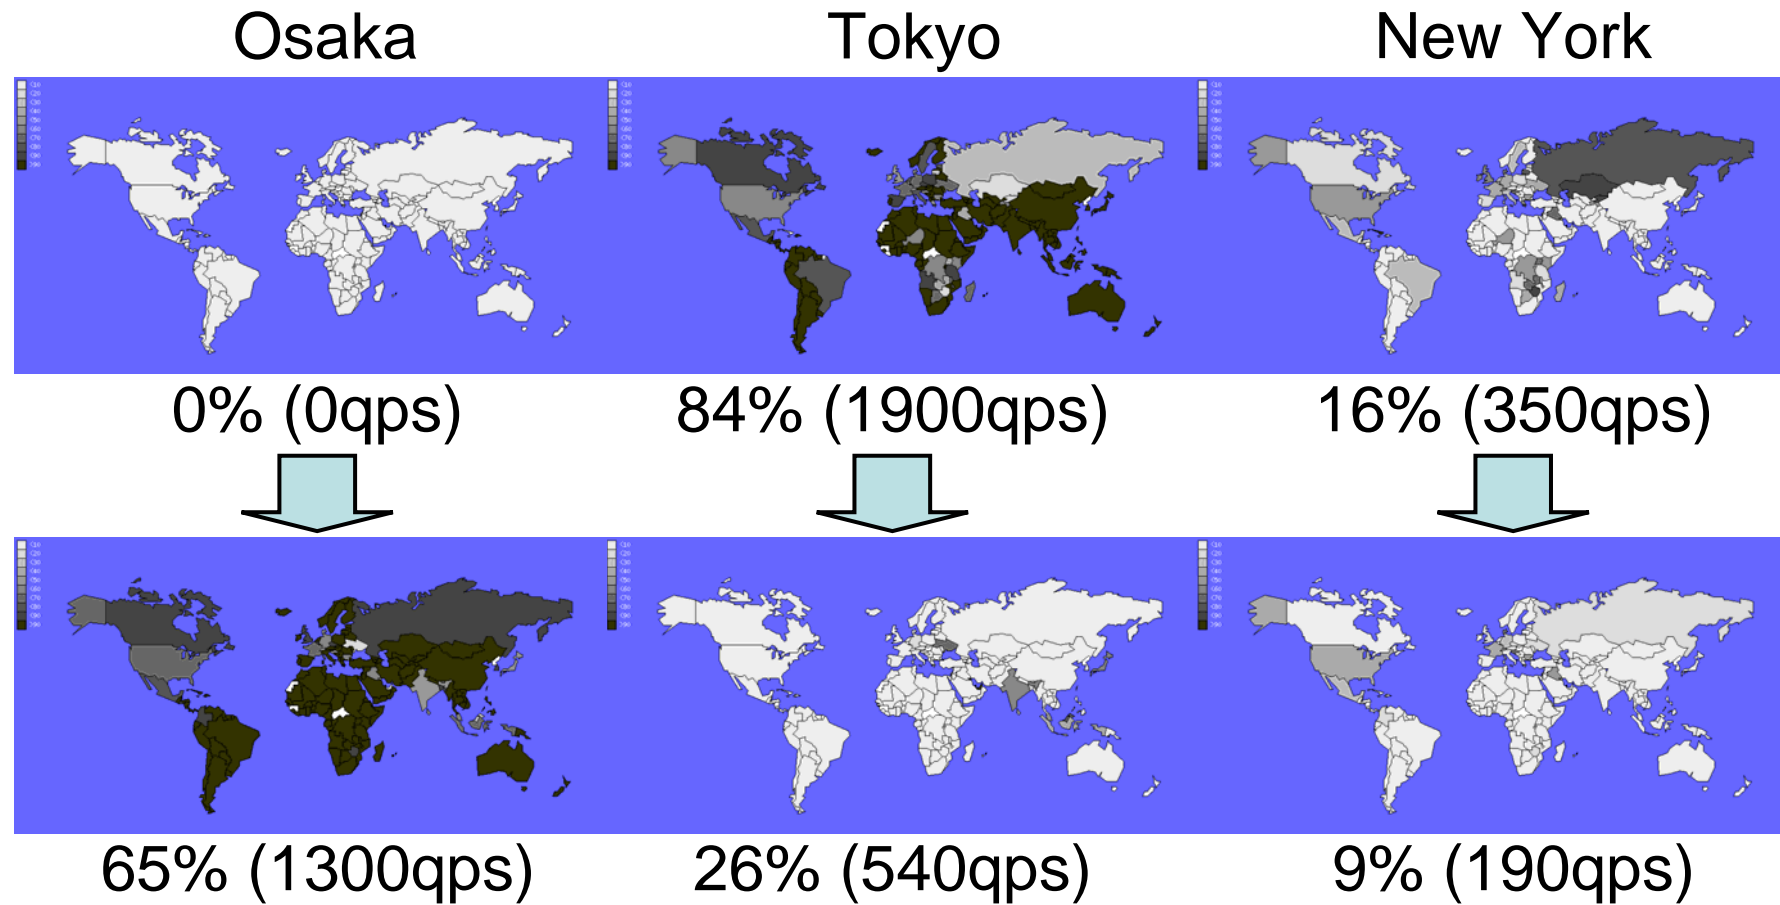
Step6: results



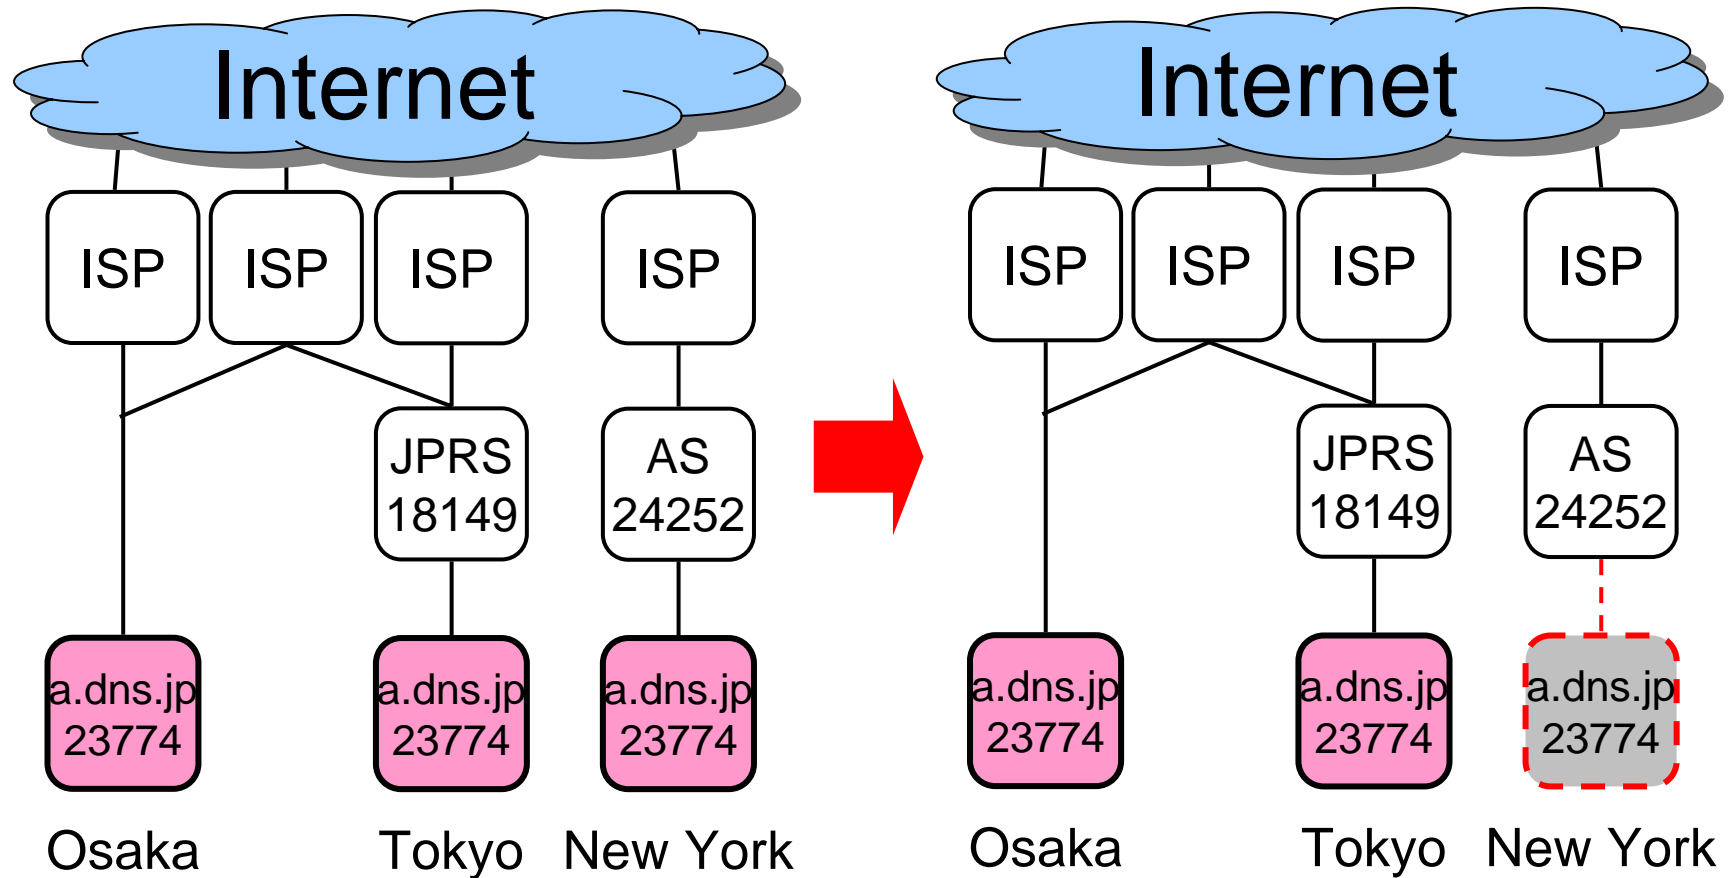
Step7: Osaka: Turn ON



Step7: results



Step8: New York: Turn OFF



Step8: results

Osaka

Tokyo

New York



67% (1100qps)

23% (400qps)

10% (170qps)



62% (1100qps)

38% (700qps)

0% (0qps)

Conclusion

- In anycast operation, AS-Path prepend is useful for load balancing.
 - The setting of present Tokyo and Osaka is just good.
- European countries are close to New York.
 - Especially, It is interesting for Japan and Russia to be distant.

Future Work

- We will analyze...
 - Origin AS Number
 - Raw data of RIPE DNSMON
 - IPv6
 - etc...

Q and A

