Measurement of Anycast Effects

- from the experience on .JP anycast deployment -

Shinta Sato <shinta @ jprs.co.jp>
Japan Registry Services Co., Ltd. (JPRS)
RIPE 54 DNS Working Group
Thursday 10 May 2007

Contents

- JP DNS Overview
- DNS Measurement from ccTLD View
- Measurement Results

JP DNS Overview

Charasteristics of .JP TLD

- Domain names registered in both 2LDs and 3LDs
- About 900,000 domain names registered
- Registrant's postal address in Japan required
- Zone data is updated and advertised every 15 minutes
- About 1 billion DNS queries handled per day

JP DNS – authoritative nameservers of .JP

- a.dns.jp – f.dns.jp (except “c”)
- dns.jp zone served by the same DNS servers
- 339 of in-addr.arpa zones (managed by JPNIC) also served by the same DNS servers
- All (except “a”) are operated on a voluntary basis

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

JP DNS Server Locations

- Active Sites (7 Japan, 2 US)
- Planned Sites (2 US, 1 Europe)

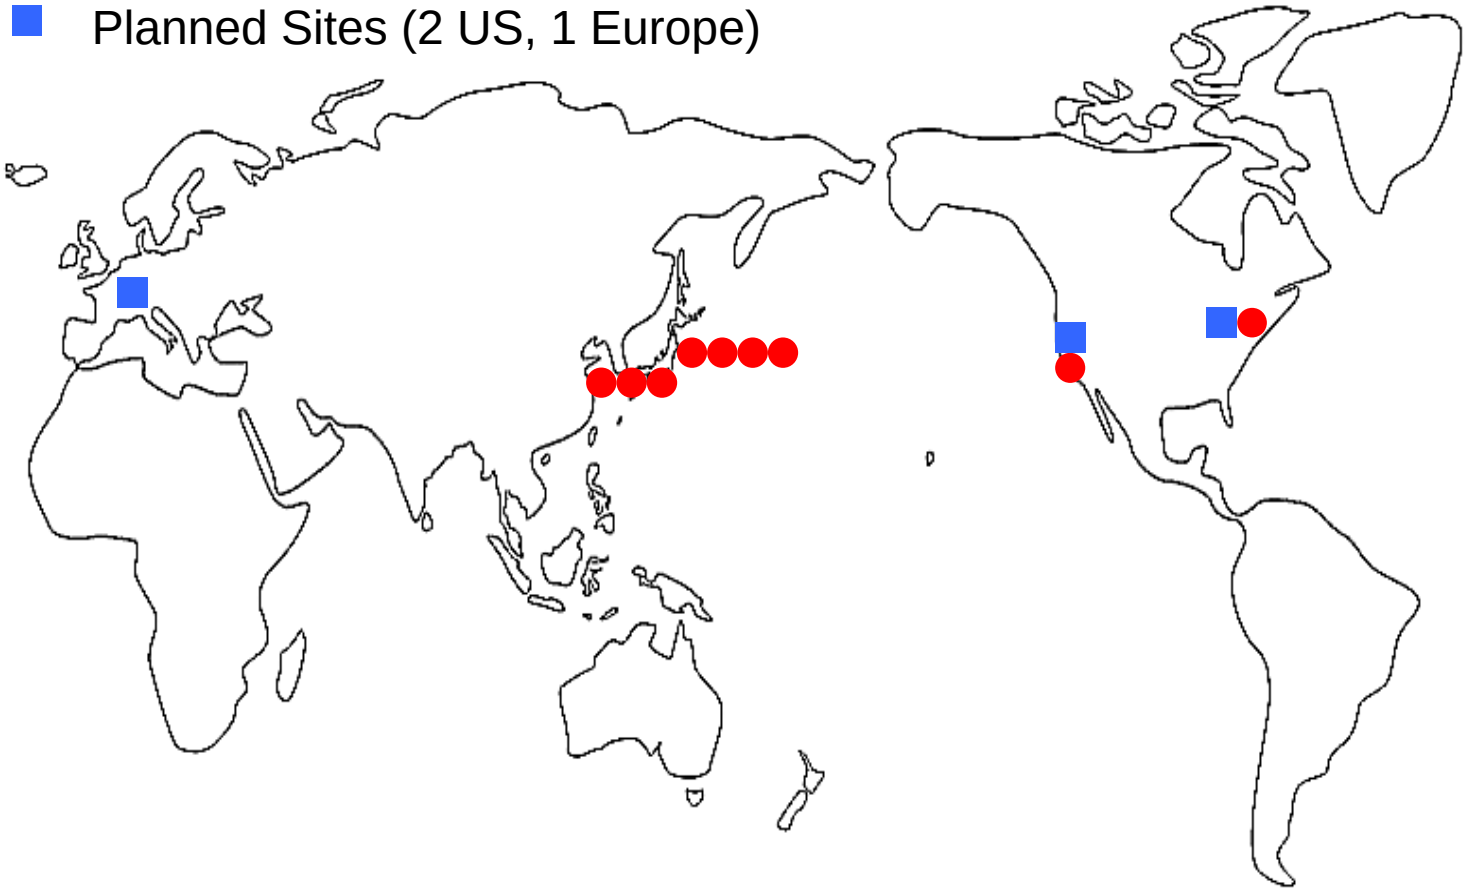

Summary of queries (a.dns.jp)

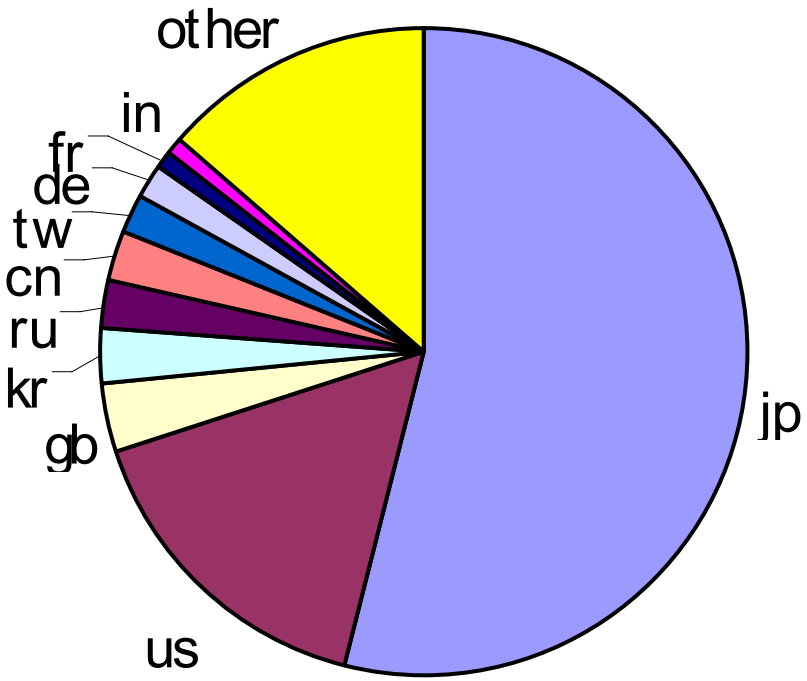

- Ave. 2000 qps
  - 1500 - 2500 qps
- Originating from
  1. JP 50 - 60%
  2. US 10 - 20%
  3. Some Asian and European Countries ~10% each
  4. Others ~1% each

DNS Measurement from ccTLD View

Reasons for the Measurements

Confirming how our IP Anycast motivations are satisfied

- Improvement of the DNS overall response
- DDoS mitigation
  - How effective each location is
  - Which queries go to which nodes?

Knowing the trends for the future system plans

- Any other conditions of anycast nodes

Any other measurements

- For techies to be prepared for the future?

What we do (for a.dns.jp)

What we collect

- Traffic information
- Query logs for all DNS servers
- Continuous SOA answers from all DNS servers
- RTT measuring from various probes (DNSMON)

What we look into

- DNS availability
- Data synchronization delay
- Query sources on country basis
- RTT distribution

Measurement Tools of JP (1/2)

- Original monitoring tools
  - Data sync checker
    - Collect SOA data from all server nodes every 10 seconds
    - Check data every hour and report if there is 120sec or bigger sync delay
- Query log / tcpdump
  - Using BIND logging feature
  - Recording all DNS queries for a.dns.jp since Feb 2004
- DSC
  - http://dns.measurement-factory.com/tools/dsc/

Measurement Tools of JP (2/2)

- Maxmind GeoIP
  - http://www.maxmind.com/app/ip-location
  - Country level resolution is precise enough in most scenes
- HELIO World
  - http://www.helio.org/world/
  - For mapping of analyzed data
- DNSMON
  - http://dnsmon.ripe.net/
  - Need more probes out of Europe

Measurement Results

Experiences from Test Run of NY Site

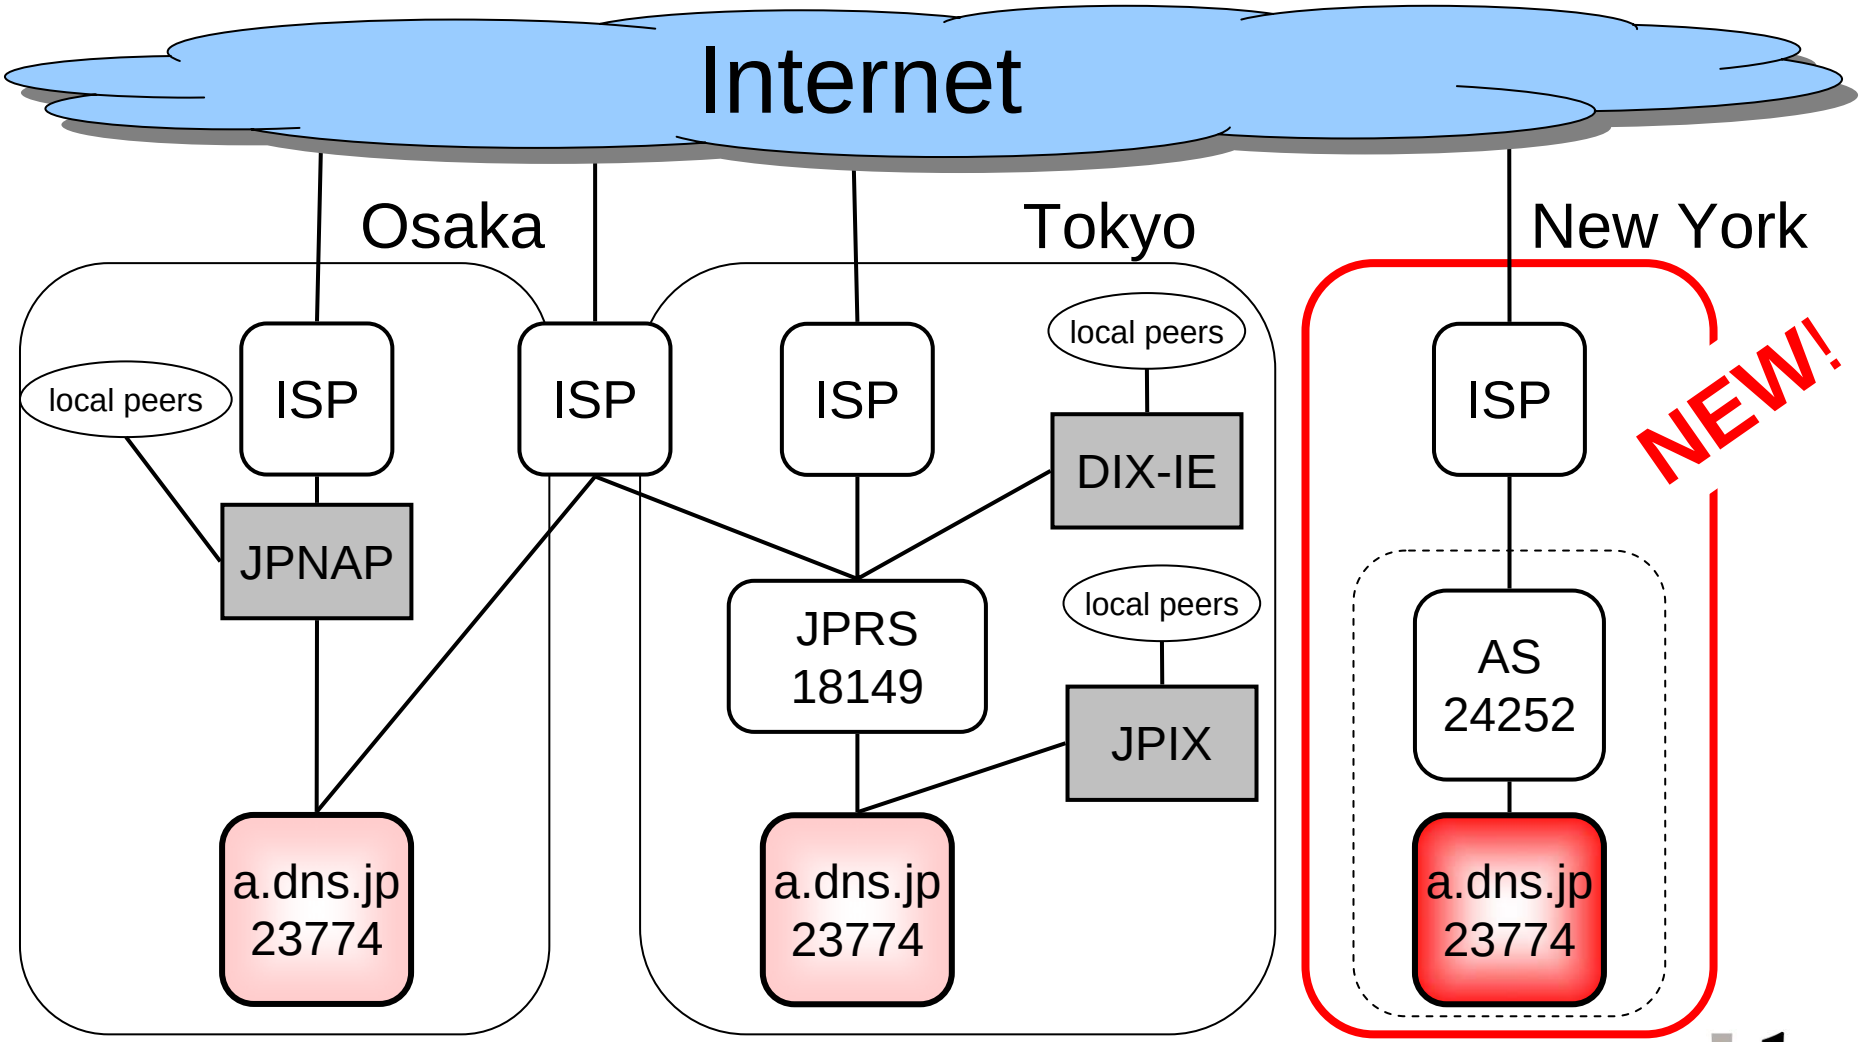

Test Run Steps

BGP operations produced

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

Effect of the new node

- Colors of the country show the percentage of queries from the country going to relevant node
  - Estimation of the DDoS mitigation

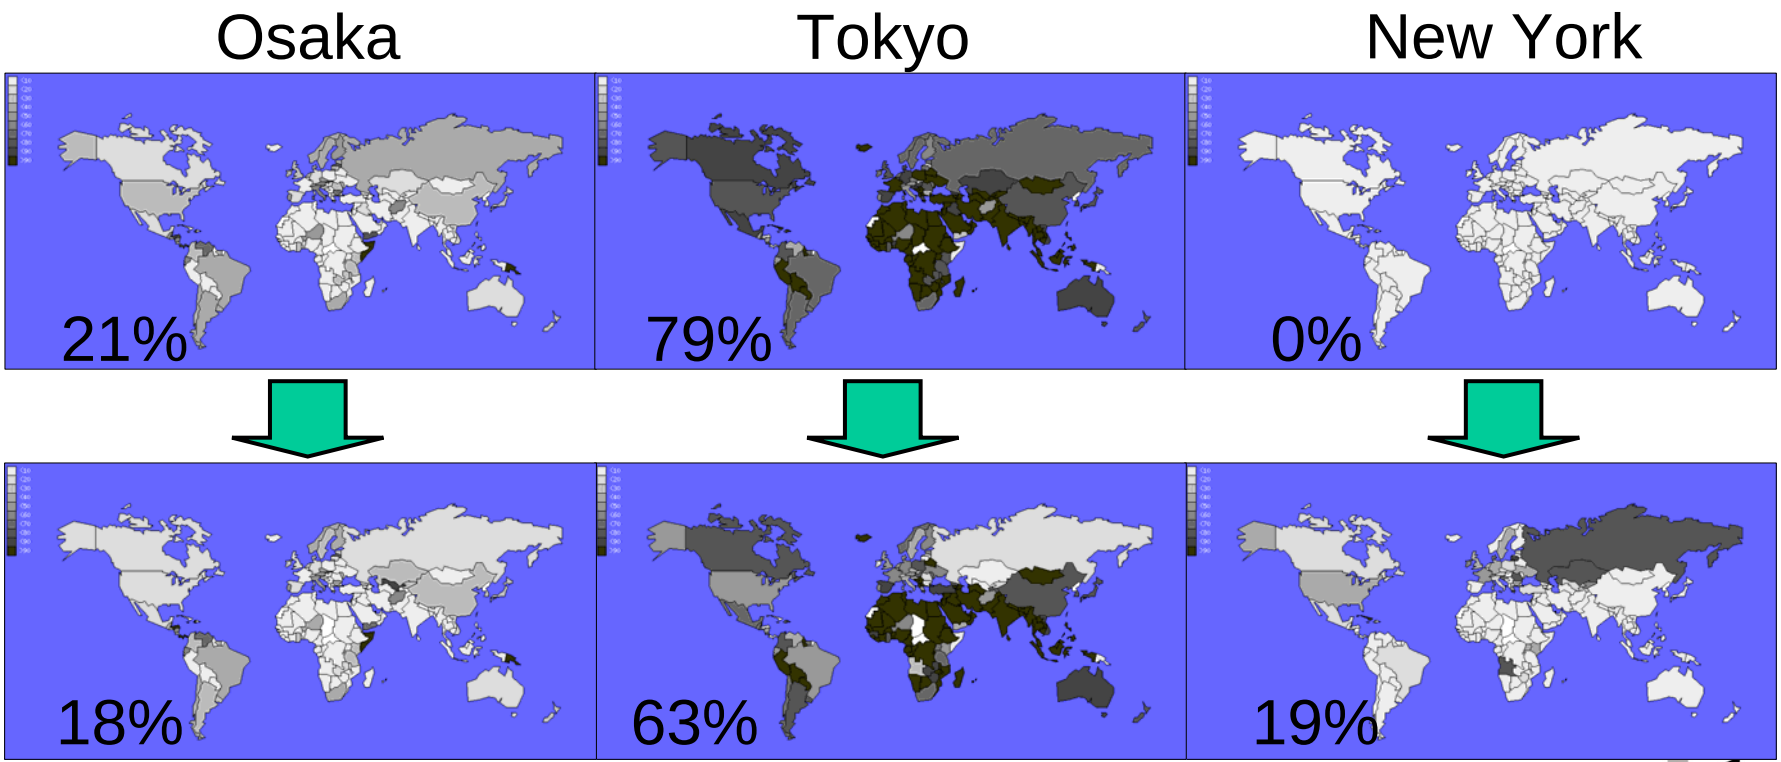

Decrease of the RTT (from DNSMON)

Regional Distributions (from DSC)

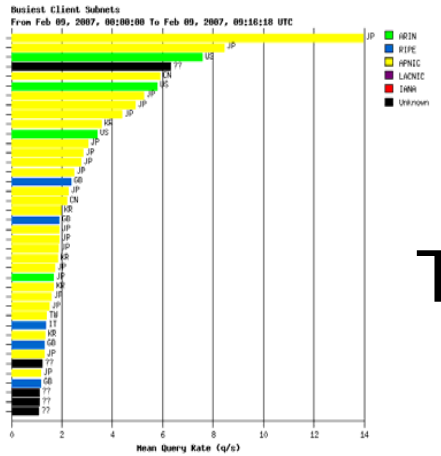

Tokyo

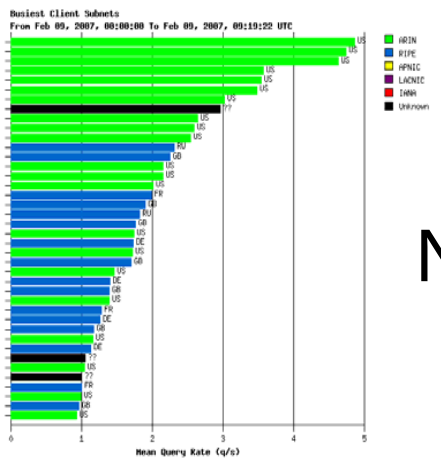

New York

ARIN
RIPE
APNIC
LACNIC
IANA
Unknown