JPRS' DNS server/service evaluation
--- user side evaluation ---

Kazunori Fujiwara, JPRS
fujiwara@jprs.co.jp

March 13, 2010

Contents

- Introduction of how JPRS' DNS server/service is evaluated before the DNS software/service will be used as JP DNS servers

- Motivation
- Evaluation
- Result

Motivation

- TLD DNS servers MUST always answer correct DNS responses
- JP zone is a complex zone compared to gTLDs and root zones. Because of this complexity, JPRS is more heavily affected by DNS software bugs than other organizations
- Then, JPRS evaluates DNS server software / service extremely deeply before using them as JP DNS server

JP zone's characteristics

- JP domain name structure consist of multiple type of domains
  - General use domain name: Second level domains: like gTLD
    - example.jp, jprs.jp
  - Organizational domain name: Third level
    - jprs.co.jp, wide.ad.jp, u-tokyo.ac.jp, kantei.go.jp
  - Geographic domain names: Third or forth level domains
    - metro.tokyo.jp, city.chiyoda.tokyo.jp, pref.nara.jp, city.nara.nara.jp
- JP zone is one zone
  - No delegations on co.jp, ad.jp, ac.jp, go.jp, tokyo.jp, chiyoda.tokyo.jp, nara.jp, nara.nara.jp, …
  - There are many empty non-terminals.

JP zone example

```
$ORIGIN JP.
@ IN SOA
        IN NS …
JPRS            IN NS …
JPRS.CO         IN NS …
WIDE.AD         IN NS …
; CO, AD are empty non-terminals
METRO.TOKYO     IN NS …
CITY.CHIYODA.TOKYO IN NS …
; TOKYO, CHIYODA.TOKYO are empty non-terminals
```

JP zone's update

- JP DNS server uses both AXFR and IXFR to transfer JP zone
  - AXFR: once a day
    - Useful for changing DNSSEC parameters
    - To avoid possible IXFR bugs (did not confronted yet)
  - IXFR: normal update, every 15 minutes

Evaluation history

- When JPRS had chosen secondary DNS service.
- When JPRS Introduced DNSSEC
  - BIND 9.4.3 to 9.7.1
  - DNSSEC evaluation itself was another work
- Version up of DNS server software
  - BIND 9.7.1 to 9.7.3
  - Secondary DNS service's software update

(planned)

- When JPRS will use another DNS server software: BIND 10, NSD, ...

Evaluation steps

1. Define current running software as a reference
2. Read new software documents carefully
3. Use the target software for small zones
4. Perform zone transfer test (JP zone)
5. Perform DNS response performance test (JP zone)
6. Perform DNS response test (JP zone)

1. Define reference version

- Writing a reference DNS response generator is best solution, but it is hard and comparison with current running version seems to be useful.

- When JPRS has chosen secondary DNS service.
  - Current running DNS server as a reference:
    BIND 9.4.3 or 9.7.1 was a reference
- When JPRS Introduced DNSSEC
  - BIND 9.4.3 to 9.7.1: 9.4.3 was a reference
- Version up of DNS server software
  - BIND 9.7.1 to 9.7.3: Reference was 9.7.1

2. Read documents carefully

- It is obvious
- Changes, manuals tell us a lot of information
  - BIND 9's CHANGES may contain important bug fixes (After a new version released, security advisories were sometimes open to the public)
- Read with extra caution by noting the following points
  - Differences from reference DNS server
  - Changes of default settings and paths
  - Changes of configuration syntax
  - Bugs or fixes after the reference version released

3. Use the target software for small zones

- To collect operational practices

- I used the new version on JPRS' lab network and my private environment

4. Zone transfer test

- Set up the test target as JP slave
  - IXFR test: every 15 minutes
  - AXFR test: daily
- Test tool sends JP SOA query every second to the master and targets, collects and parses responses. (timing is configurable)
- After zone data will be in sync, compare transferred zone data with the master's zone data using AXFR.

Zone Transfer Test

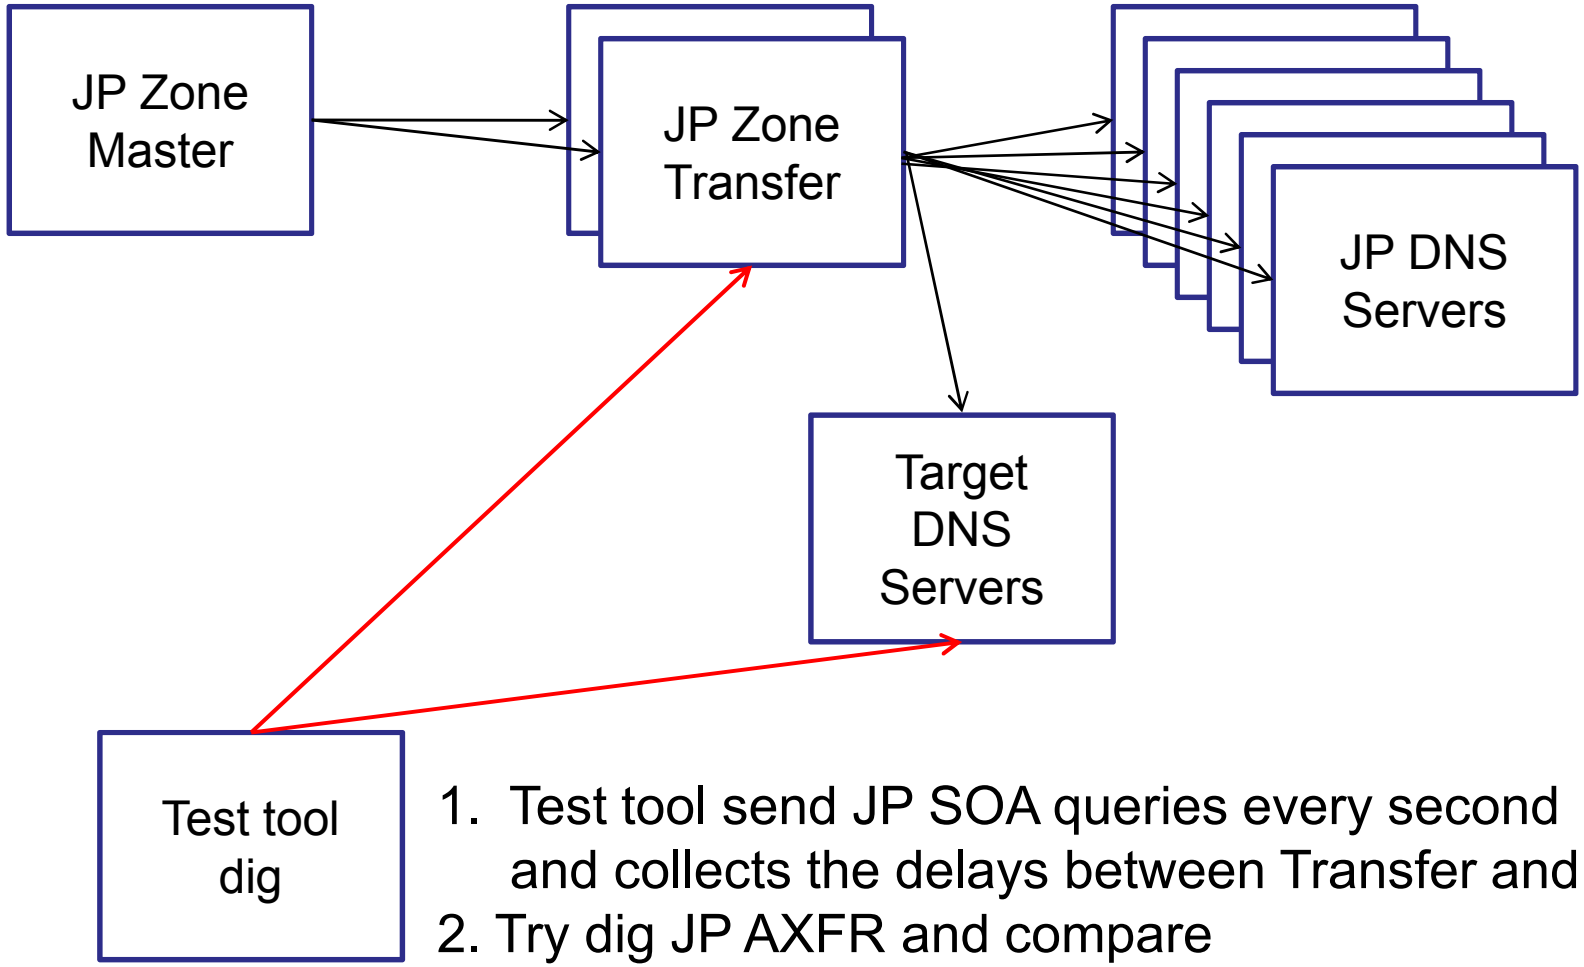

1. Test tool send JP SOA queries every second and collects the delays between Transfer and Target
2. Try dig JP AXFR and compare

Some result of Zone Transfer Test

- If the DNS server is located oversea, AXFR transfer may take large time.
  - It sometimes takes over 15 minutes
- If the DNS server's connectivity is poor, the test tool sometimes cannot detect SOA changes

- On my test, I found old BIND 9 (prior to 9.7.1) stops responding queries while it is dumping zone backup file immediately after AXFR.
  - Because dumping of JP zone takes 5 seconds and my tool detected 5 seconds' no response.
  - It is fixed in BIND 9.7.1 and NSD does not stop responding immediately after AXFR.

5. Response performance test

- Using queryperf
  - Two test cases: NO_ERROR case and NAME_ERROR case
  - to the target DNS server

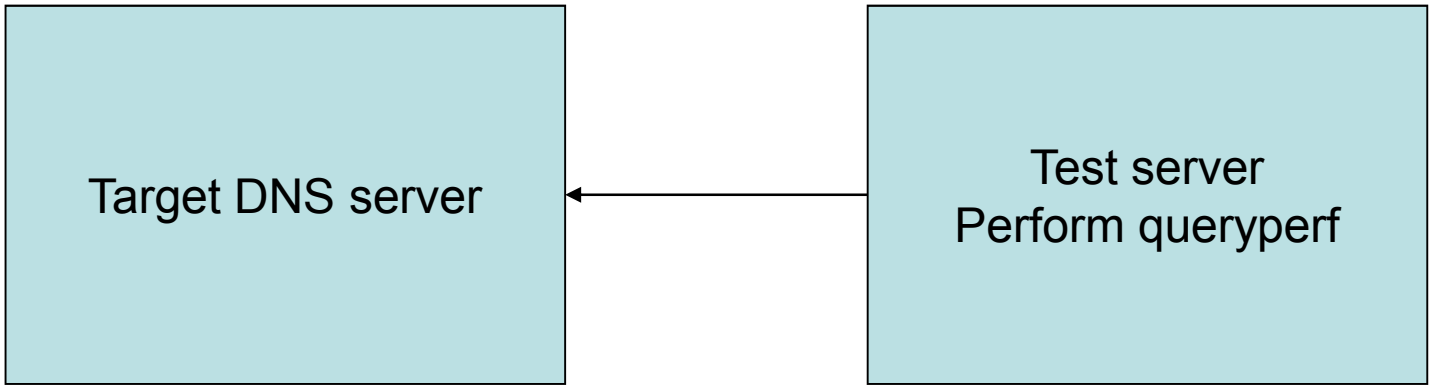

6. DNS Response test

- Goal: The software answers all queries correctly
- Setup both the test target and the reference as JP slave
- Send all possible queries to both reference DNS server and target DNS server
- Compare all responses

DNS Response Test

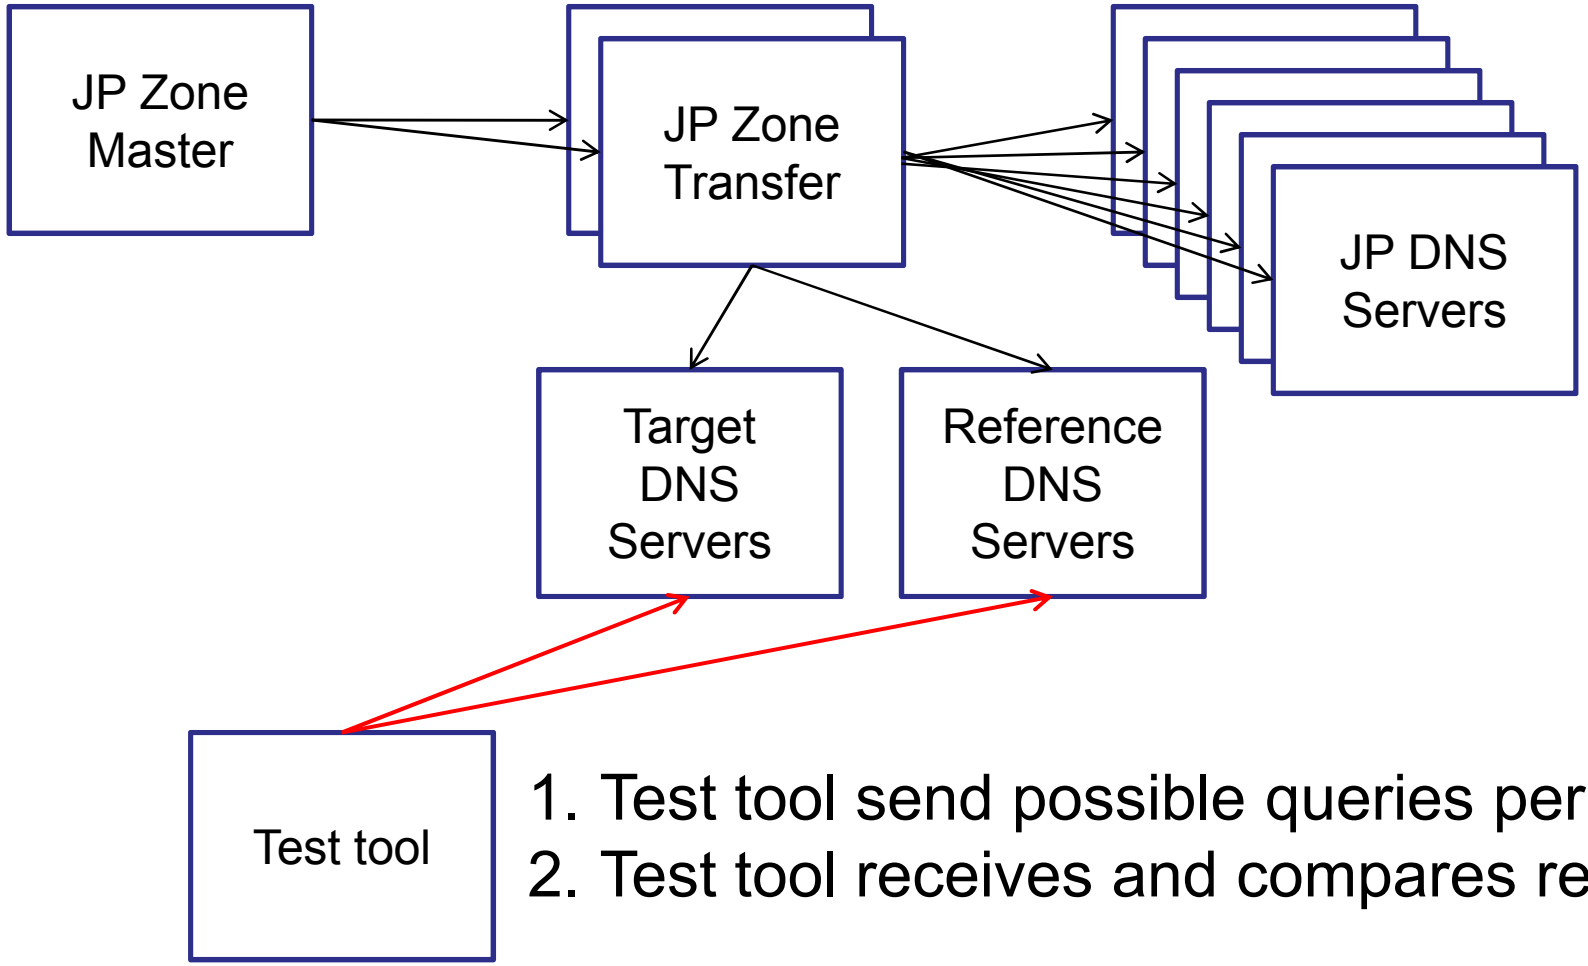

1. Test tool send possible queries periodically
2. Test tool receives and compares responses

Possible queries are

- Owner names from JP zone as $dom
  - Registered domain names
  - Glue host names
  - Non-existing name (xx-yy.jp)
  - Empty non-terminals (co.jp, ad.jp, …)
- 28 patterns of domain name and query type
  - $dom {A,AAAA,MX,NS,CNAME,SPF,TXT,NAPTR,DS,RRSIG,NSEC,DNSKEY,NSEC3,NSEC3PARAM}
  - noexistence.$dom (for maybe non-existence name) {A,AAAA,MX,NS,CNAME,SPF,TXT,DS,RRSIG, NSEC,DNSKEY,NSEC3,NSEC3PARAM}
  - _sip._udp.$dom SRV
- Three attributes: noEDNS0, EDNS0, DO=1

Total queries

- JP zone has about 1,300,000 owner names (registered domain name and glues)
- Times 28 patterns
- Times 3 attributes
- Times 2 servers
- Makes 218,400,000 queries
- Test tool send the queries specified time steps
  - 1 milli-second step case, it sends 500 queries/sec for both servers
  - The test takes 218,400 seconds: about 3 days

Comparison on DNS response test

- There are different DNS responses but they are correct DNS responses
  - Ordering in the sections
  - Additional section may contain glue RRs
  - Authority section may contain zone's NS RRs
  - EDNS0 payload size differences

- Correct differences need to be treated as no-problem
  - If I find a difference, I evaluate it is OK or not.
  - If Ok, I need to update the comparison program not to report the differences
  - I don't know how to automate the step

Some findings of DNS Response Test

- When I found some bugs, I reported and they were fixed (Or didn't use the software)
- BIND 8 was old
  - It put NS RRs in answer section at delegation
  - Recent DNS servers put NS RRs in authority section
- BIND 9 sometimes changed the response patterns
  - Recent BIND 9 does not add authority section in DS or DNSKEY answer to minimize DNS packets

6': DNSSEC and Non-DNSSEC response test

- Prepare test signed JP zone and load it into the test target
  - Added some DS RRs and signed
- Prepare reference DNS server with traditional (non-DNSSEC) JP zone
- Sent all query patterns to both servers
- Compare responses
  - Ignored differences of DNSKEY, RRSIG, DS, NSEC3
- This test resulted that resolvers are not affected by JP zone signing if the resolvers does not perform DNSSEC validation.

Conclusion

- Trying all of possible query patterns are very useful for DNS server evaluation even if it takes very long time.

- There were many bugs
- We are trying to avoid bugs on our DNS servers

- DNS software / service evaluation is important for JPRS

We would like to know

- Do you evaluate DNS server software / services on a user's point of view ?

Comments and Questions ?