Lessons Learned from May 19 China's DNS Collapse

Ziqian Liu
liuzq@chinatelecom.com.cn

2nd DNS-OARC Workshop
Beijing, Nov 5 2009

Agenda

- **What happened**
- **How did we survive**
- **How could this happen**
- **Lessons learned**

What happened- at first sight

- Time
  - 21:00—22:40, May 19, 2009
- Many provinces suffered…
  - 6 reported “almost lost Internet access” province-wide: www, emails, IM... ALL Failed!
  - Others reported “much lower Internet speed”
  - Not only CT, but other Tier-1 ISPs (Unicom, CM) in China all reported large scale Internet access failure
- Survivals remains still…
  - Some provinces see nothing wrong
- Inference 1 : backbone unstable?
  - No relevant link down, no route flapping, no configuration changed
  - So many ISPs all got involved?

What happened-a close look

- Internet users got few responses from their local DNS servers
- Operators observed their boxes way overloaded…
  - DNS servers: ~100% CPU/Memory usage
  - Layer-4 switches (if exists): the same, plus abnormally high parallel session amount
  - up-connected network nodes: several times larger Incoming traffic than normal
  - source IP addr of the incoming traffic: widely scattered
- Inference 2: DDoS attack againt DNS servers?
  - Simultaneous attacking so many servers in every ISP in China ?
  - Few malformated queries were seen
  - No prominent DDoS observed at the time

What happened-a close look (cont'd)

One of our DNS node

- memory usage
- Inbound traffic
- Parallel sessions
  （~20Kqps）

What happened-a careful look

- baofeng.com became the hottest domain
  - Top4 most-requested names all ended with baofeng.com
    ```
    videodown.baofeng.com,
    active.baofeng.com,
    live.baofeng.com,
    download.baofeng.com
    ```
- Servfail messages showed up when trying nslookup baofeng.com names
- The recursive session number reached the limit !!

```
bash-2.05$ rndc status
...
recursive clients: 49902/50000
...
bash-2.05$ more default.log.20090519-2
19-May-2009 22:21:13.186 client: warning: client 218.77.186.180#51939: recursive-clients soft limit exceeded, aborting oldest query
19-May-2009 22:21:13.213 client: warning: client 59.50.182.161#1151: recursive-clients soft limit exceeded, aborting oldest query
```

- Inference 3: something wrong with baofeng.com resolution?

Agenda

- **What happened**
- **How did we survive**
- **How could this happen**
- **Lessons learned**

How did we survive

- Filter out baofeng.com queries
- Force the cache to response something to baofeng.com queries
  - artifact ip addresses
  - still alive baofeng.com authoritative NS addr.
- Increase the TTL of baofeng.com related RRs
- Decrease the timeout threshold of each connection on layer-4 switches (if exists)
- Peace…

Outline

- **What happened**
- **How did we survive**
- **How could this happen**
- **Lessons Learned**

How could this happen – root cause

- DDoS attack!
  - Not against ISP's local DNS name servers…
  - Nor baofeng's authoritative name servers…
  - But private online game servers !

How could this happen – playback

So , who are the real victims?

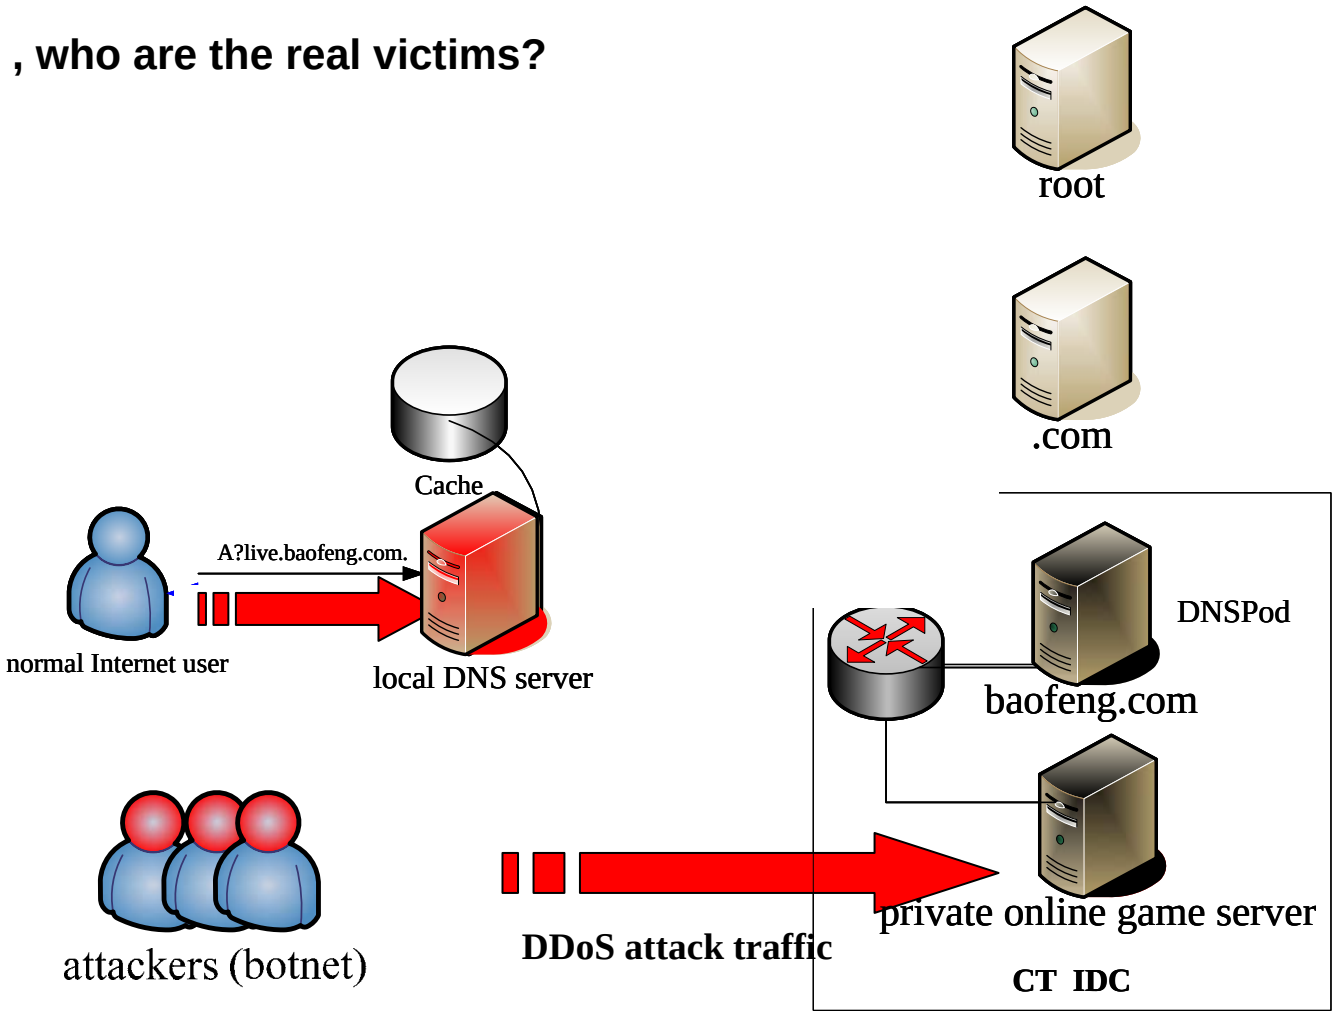

Outline

- What happened
- How did we survive
- How could this happen
- Lessons learned

Lessons Learned

- To improve the current DNS architecture to a safer one…
  - Anycast name service(domino effect?)
  - Physically separate recursive function from the cache
  - Strict access control between each boxes
  - Shadow authoritative servers needed?
  - An up-to-date copy of DNS name tree need? at least the main branches?

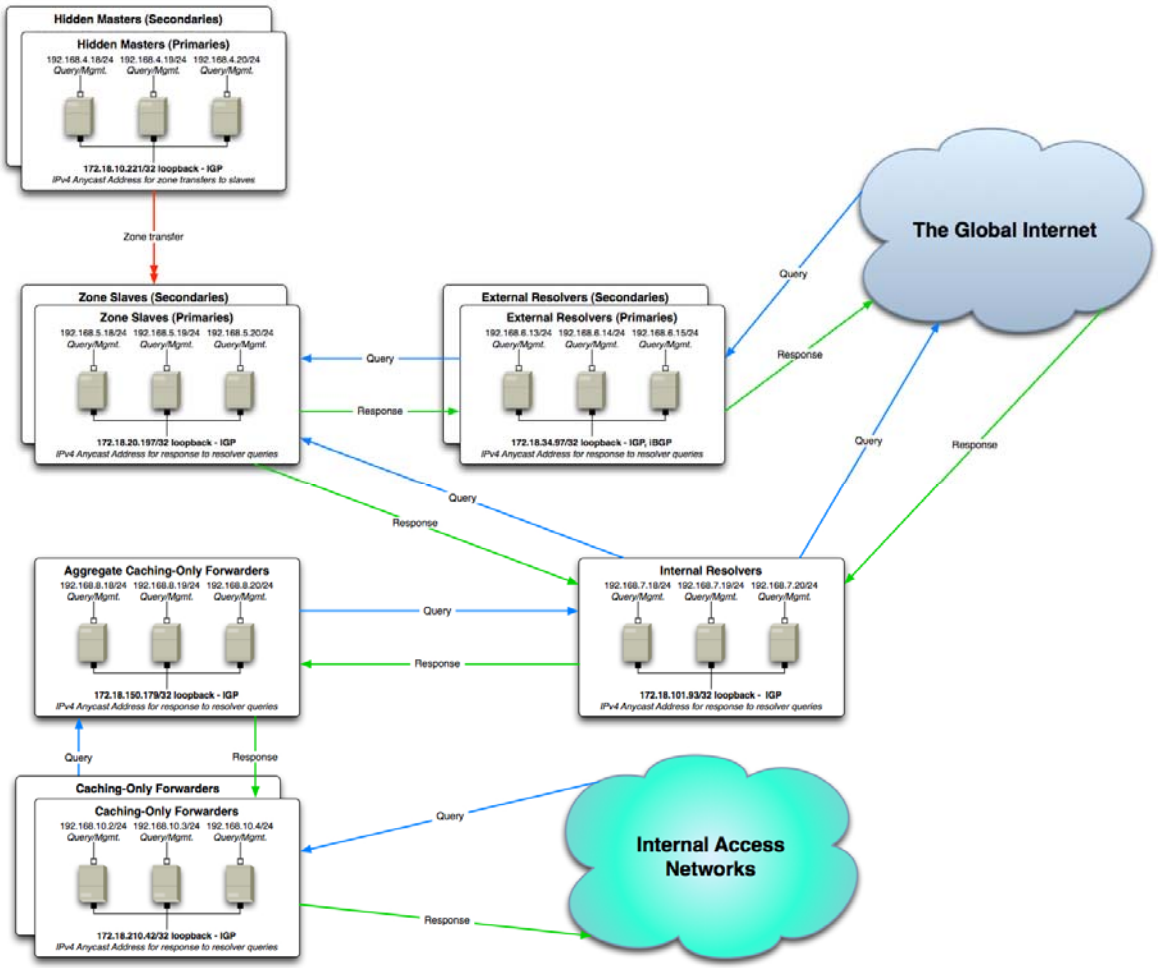

Cited from R. Dobbins Apnic 28 slide.

Lessons Learned (cont'd)

- Always keep an eye on your box!
- High performance boxes may help (to a limit extend…)
- ISPs should seriously concern about not only the DNS systems but
  - Killer Apps
  - Name resolution abusers
  - DDoS attack (against DNS servers)
    - Not only the DNS community but the Internet Community at large
- Caution: **NEVER** blackhole DNS servers unless you are fairly clear about the outcomes !

Thank you!

& comments?