DNS Anomalies and Their Impacts on DNS Cache Servers

Katsuyasu TOYAMA, Keisuke ISHIBASHI, and Tsuyoshi TOYONO (NTT)

Masahiro ISHINO and Chika YOSHIMURA (NTT Communications)

Kazunori FUJIWARA (JPRS)

Background

- In the DNS world, these ‘players’ are related to each other
  1. application or operating system in user devices
  2. authoritative servers including root DNS servers
  3. DNS cache servers

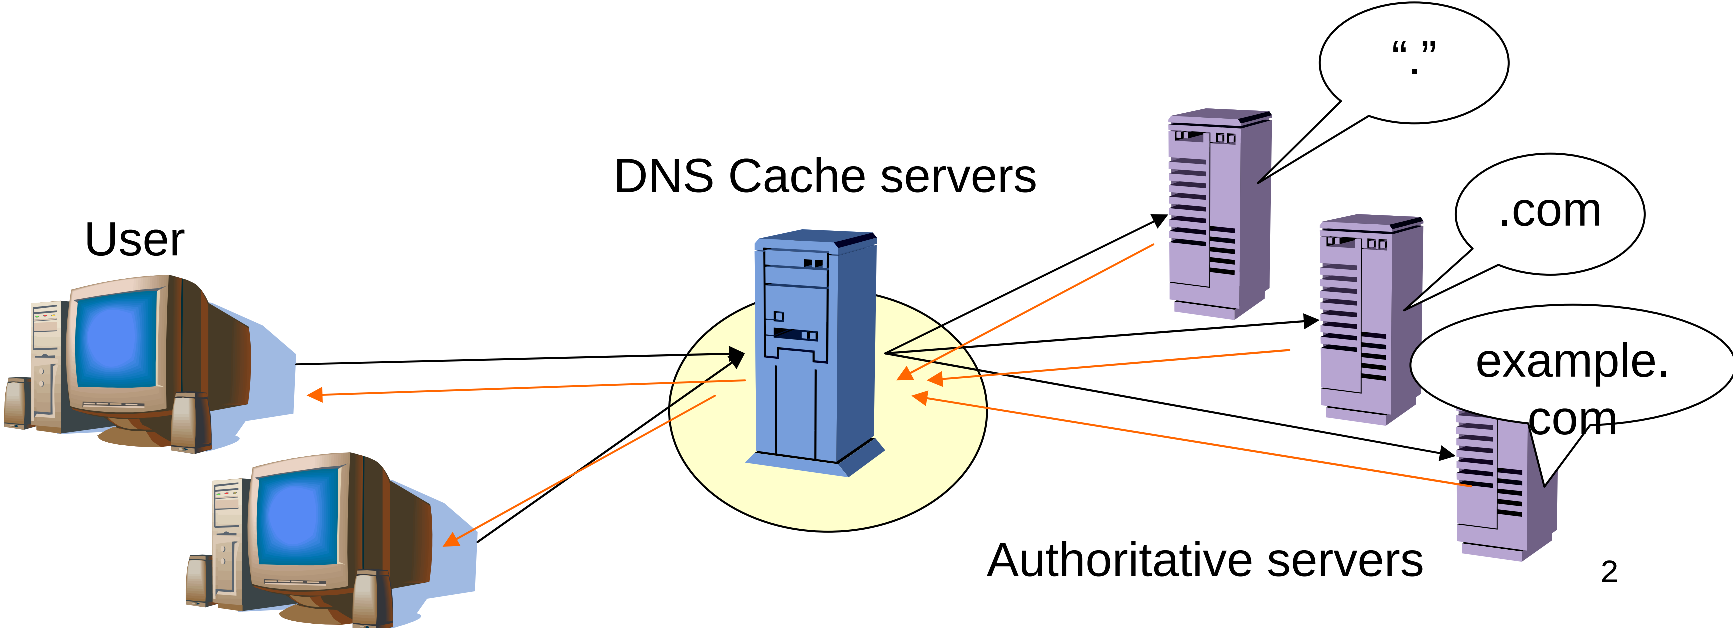

Today's topic

- The burden of DNS Cache servers caused by misconfigured DNS authoritative servers
  - Lame delegation is well known, but other factors exist as well.

- Focusing on:
  1. Virus and Worm activities

     An extreme increase in queries caused DNS cache server to become overloaded.
  2. Large RRSet and TCP filtering

     Even small number of queries could cause increase in heavy load of DNS cache server.

1. Virus and Worm

Activities of Virus and Worm

- Activities of viruses and worms sometimes cause burden on DNS cache servers.

- MyDoom
  - Attacks SCO Web site; some subspecies attack Microsoft Web site
- Antinny
  - Attacks ACCSJP Web site
    - ACCSJP: Association of Copyright for Computer Software

What is Antinny?

- Antinny is a worm that infects through famous Japanese P2P software, “Winny”.
- Some subspecies of Antinny try to send private information of the infected user to the Web site of ACCSJP once a month.
  - the first Monday of the month
    - Apr. 5th, May 3rd, June 7th…
  - to expose the user as a potential criminal??

- When trying to connect to the Web site, it resolves the FQDN “www.accsjp.or.jp”

Worms' impacts on networks

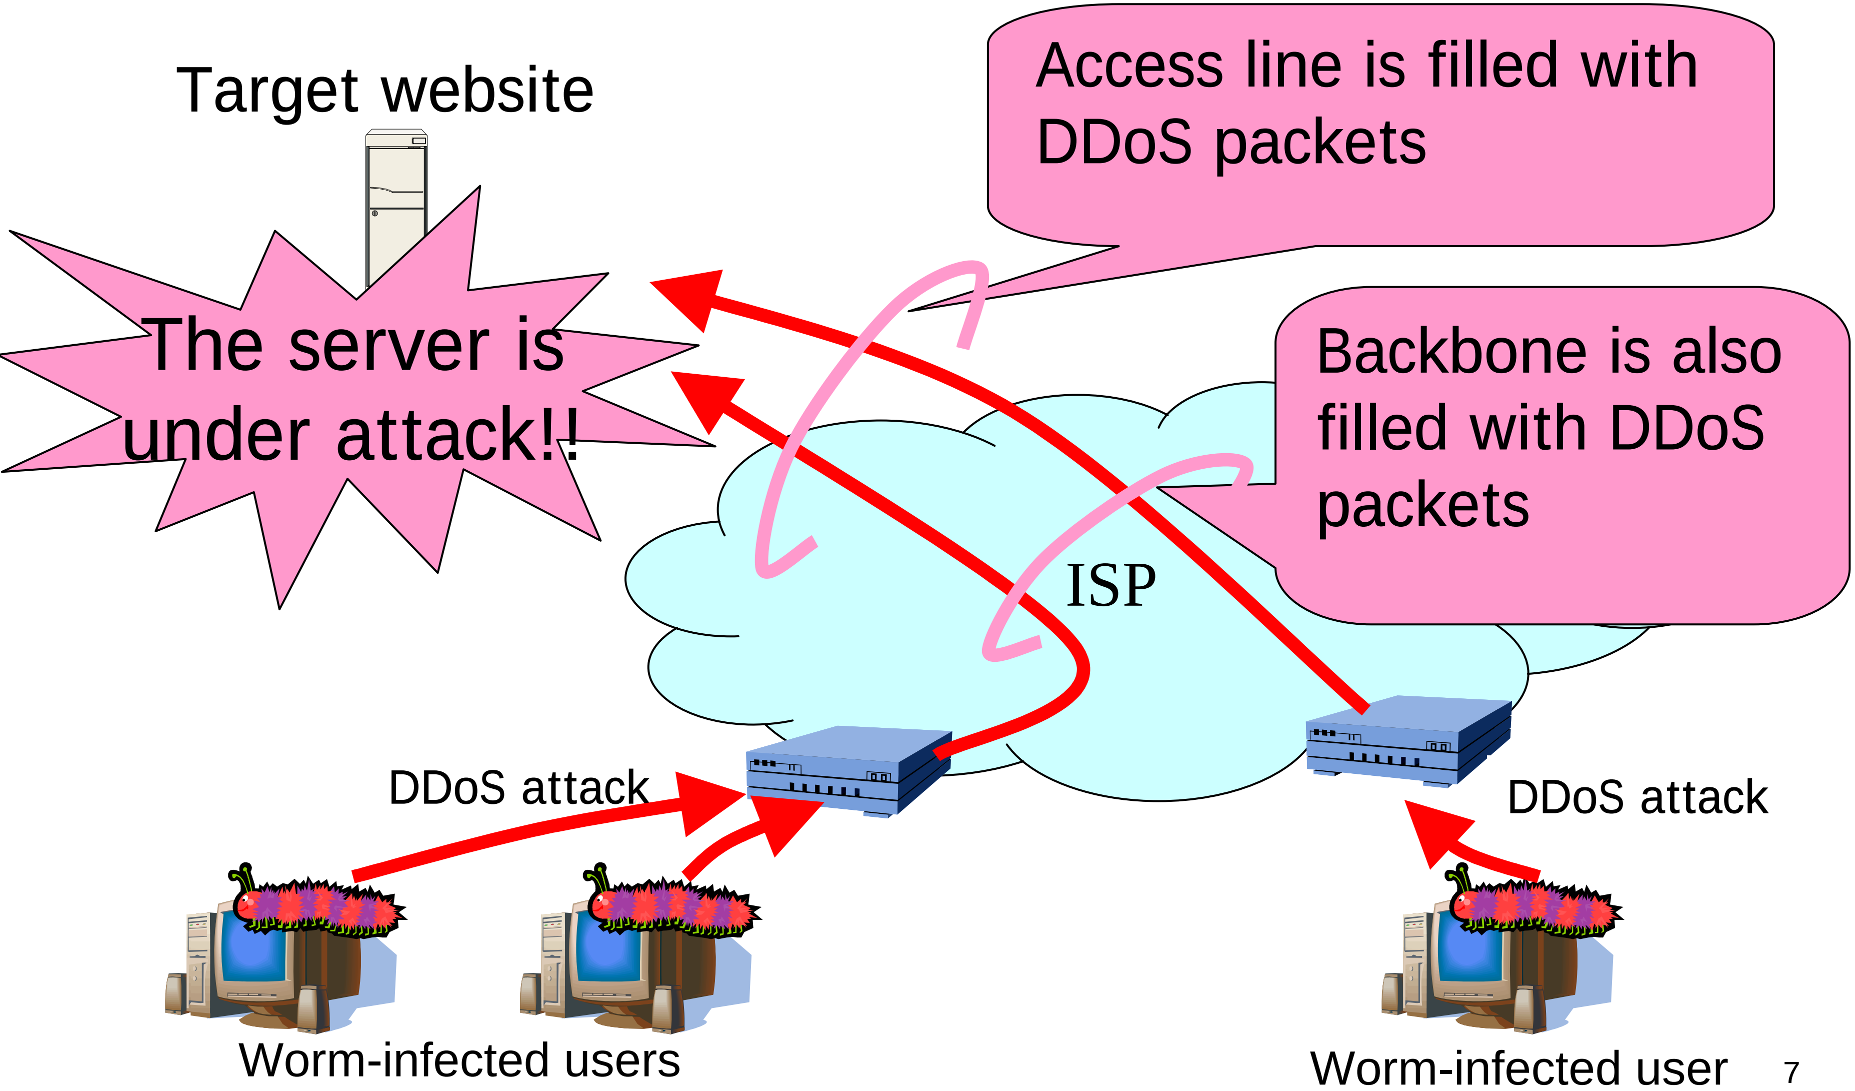
Target website
Access line is filled with DDoS packets
The server is under attack!!
Backbone is also filled with DDoS packets
ISP
DDoS attack
DDoS attack
Worm-infected users
Worm-infected user

Owner removed A RR from authoritative server!!

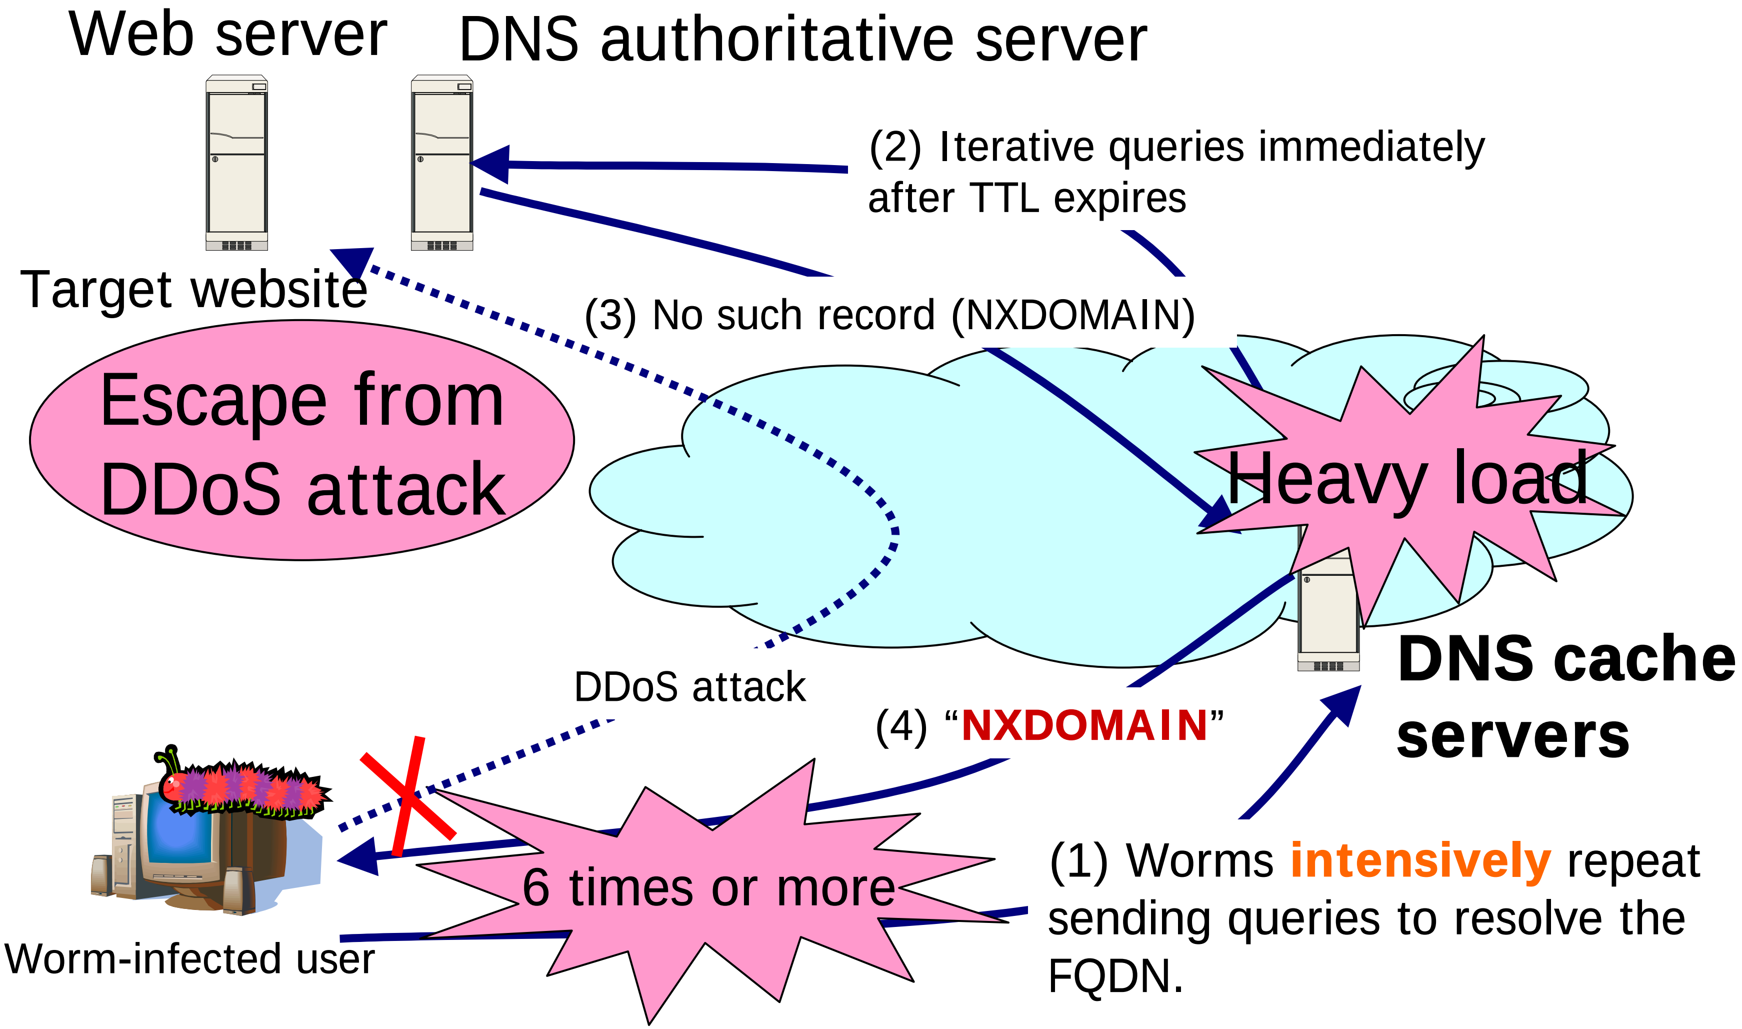

Web server

DNS authoritative server

(2) Iterative queries immediately after TTL expires

Target website

(3) No such record (NXDOMAIN)

Escape from DDoS attack

Heavy load

DNS cache servers

DDoS attack

(4) “**NXDOMAIN**”

6 times or more

Worm-infected user

(1) Worms **intensively** repeat sending queries to resolve the FQDN.

Effect of A RR removal

- Web site owner is happy because:
  - it is very easy to remove A RR
  - their link has become very quiet

- But… DNS cache servers are overloaded!

Why is DNS cache overloaded?

- ‘A RR’ was removed from authoritative server.
  - NXDOMAIN was returned, and its TTL was short (60 sec.)

- Worms repeatedly sent queries even if they could not resolve the name.
  - They never gave up…
  - The highest was approximately 700 queries per second from an IP address!

- Negative cache did not seem to be effective in some Operating Systems or applications
  - Negative cache is disabled at any time?

Quick fix: Return blackhole IP address by each cache server

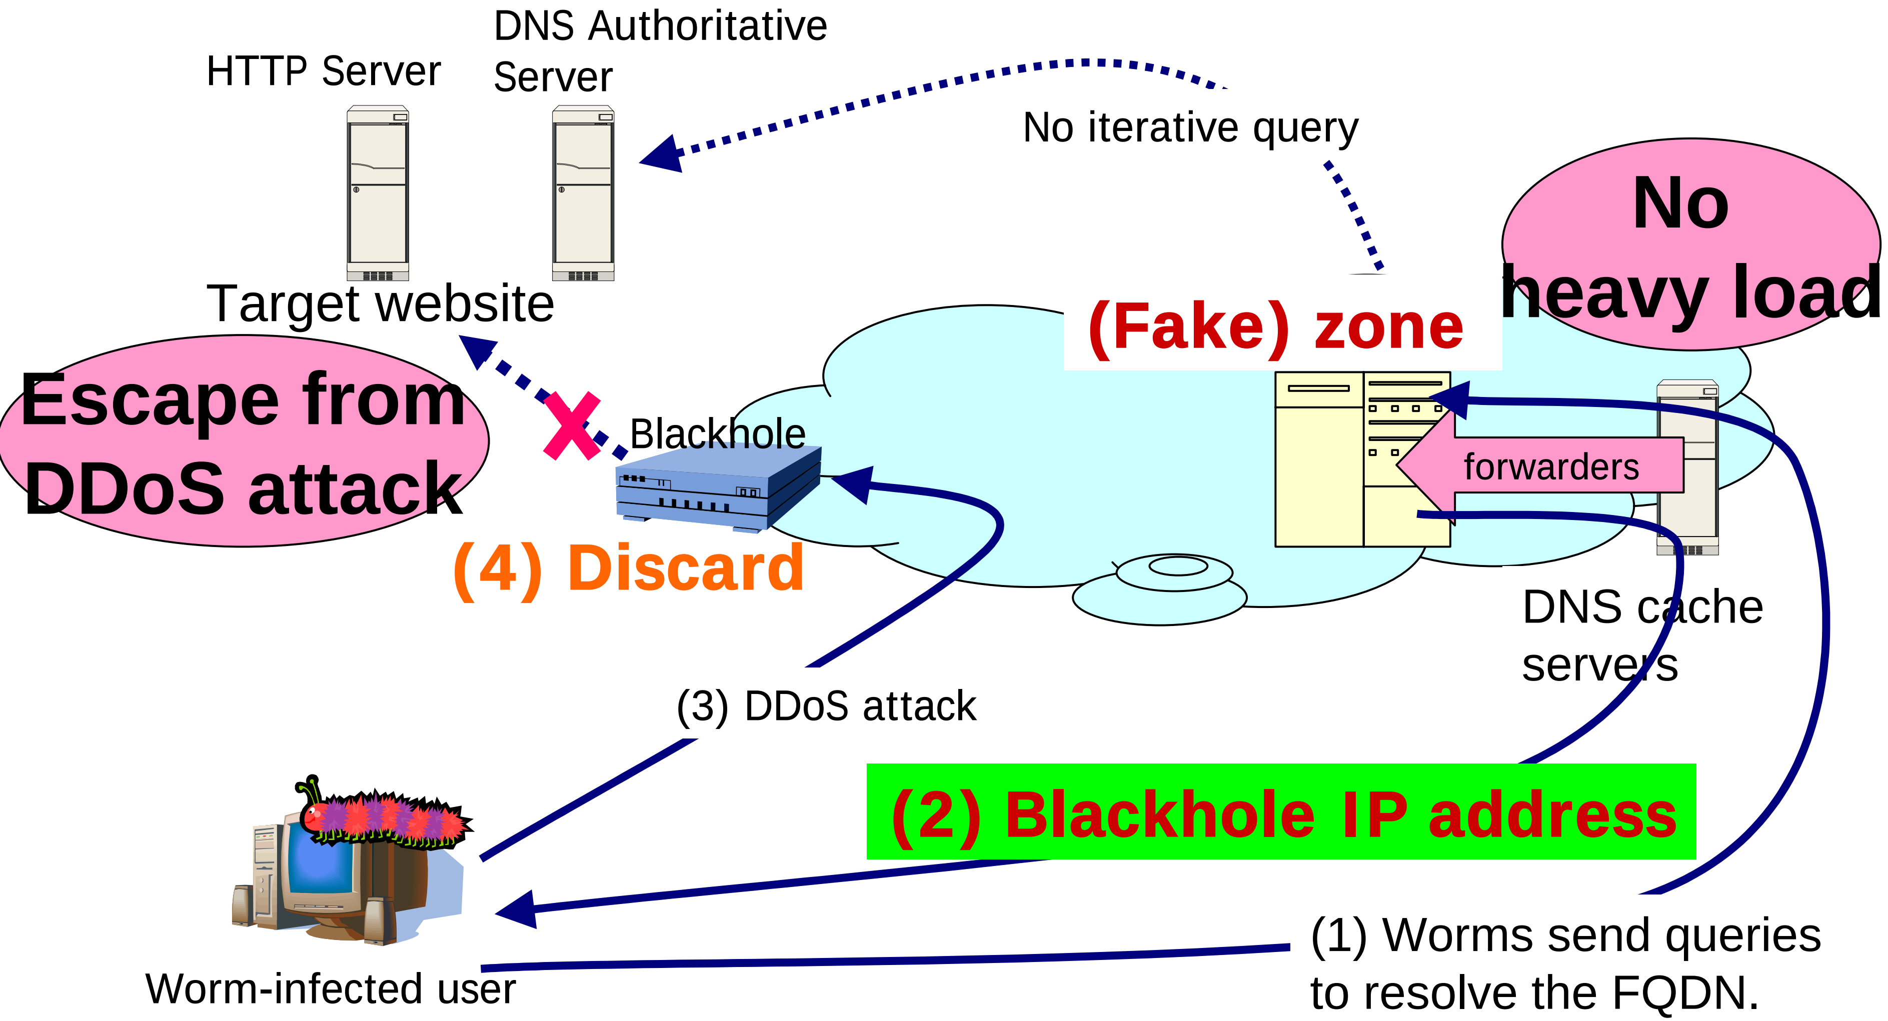

HTTP Server

DNS Authoritative Server

No iterative query

No heavy load

Target website

(Fake) zone

Escape from DDoS attack

Blackhole

forwarders

(4) Discard

DNS cache servers

(3) DDoS attack

(2) Blackhole IP address

Worm-infected user

(1) Worms send queries to resolve the FQDN.

Increased total queries at DNS cache

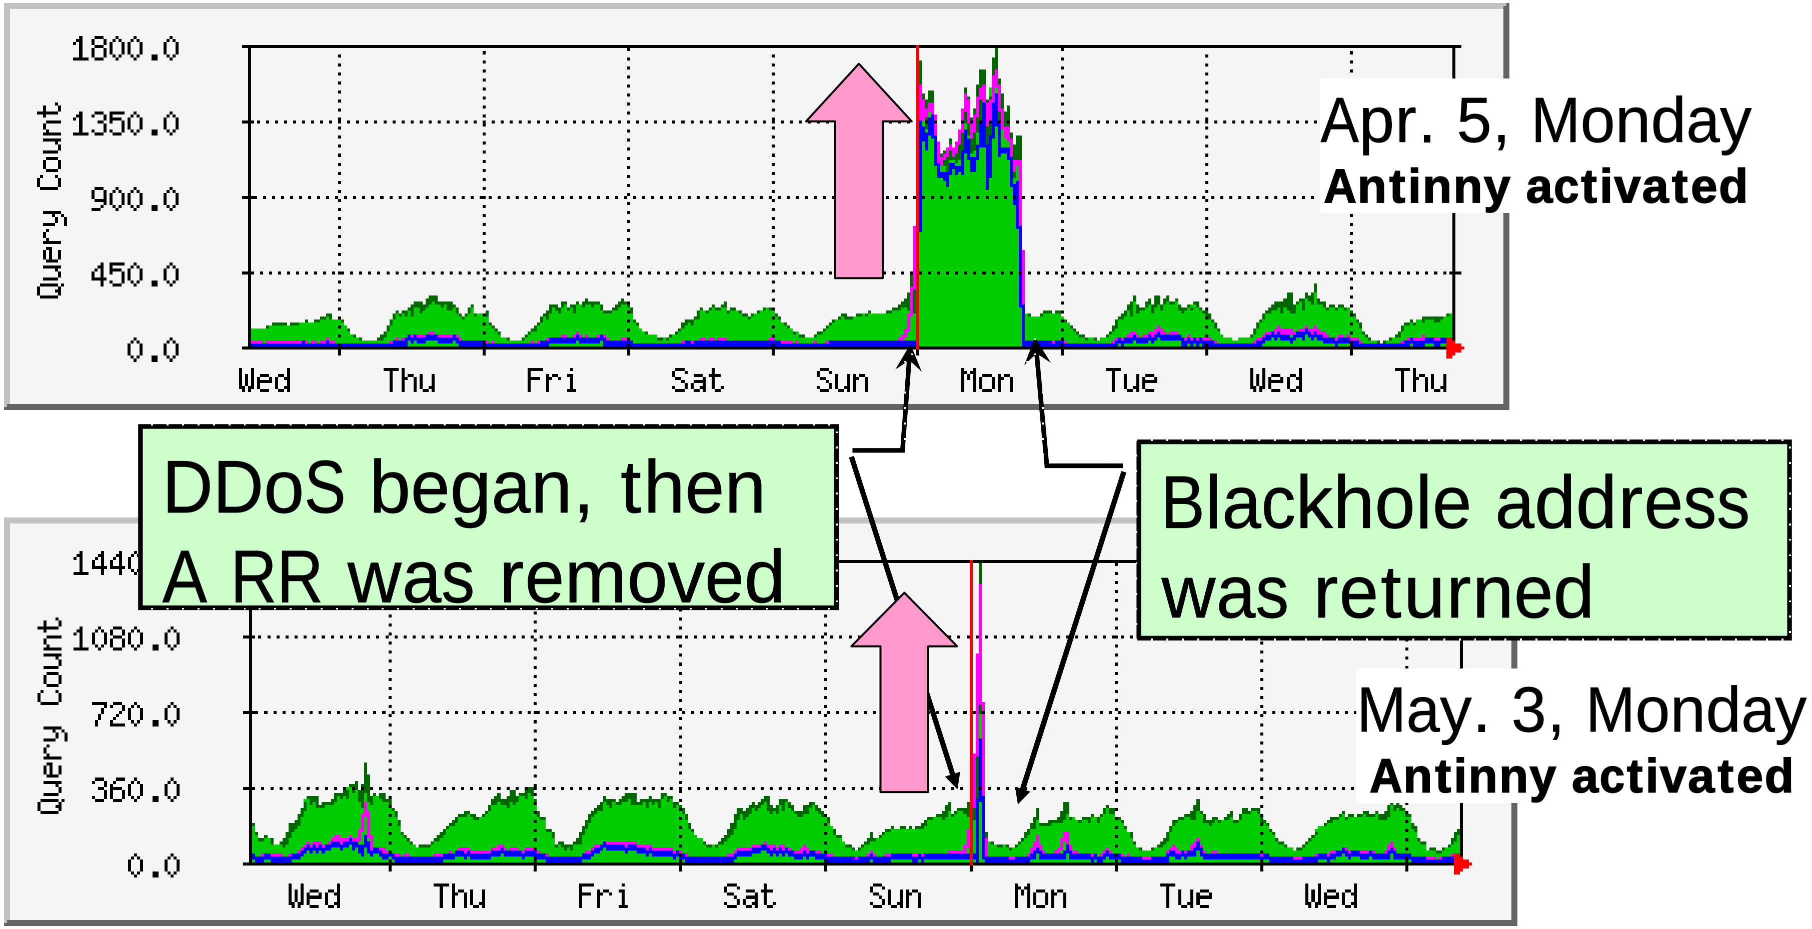

Apr. 5, Monday
Antinny activated

DDoS began, then A RR was removed

Blackhole address was returned

May. 3, Monday
Antinny activated

Green: Number of total queries
Blue: Number of NXDOMAIN

Number of queries resolving FQDN of target Web server (May 1st – 7th, per 5 min.)

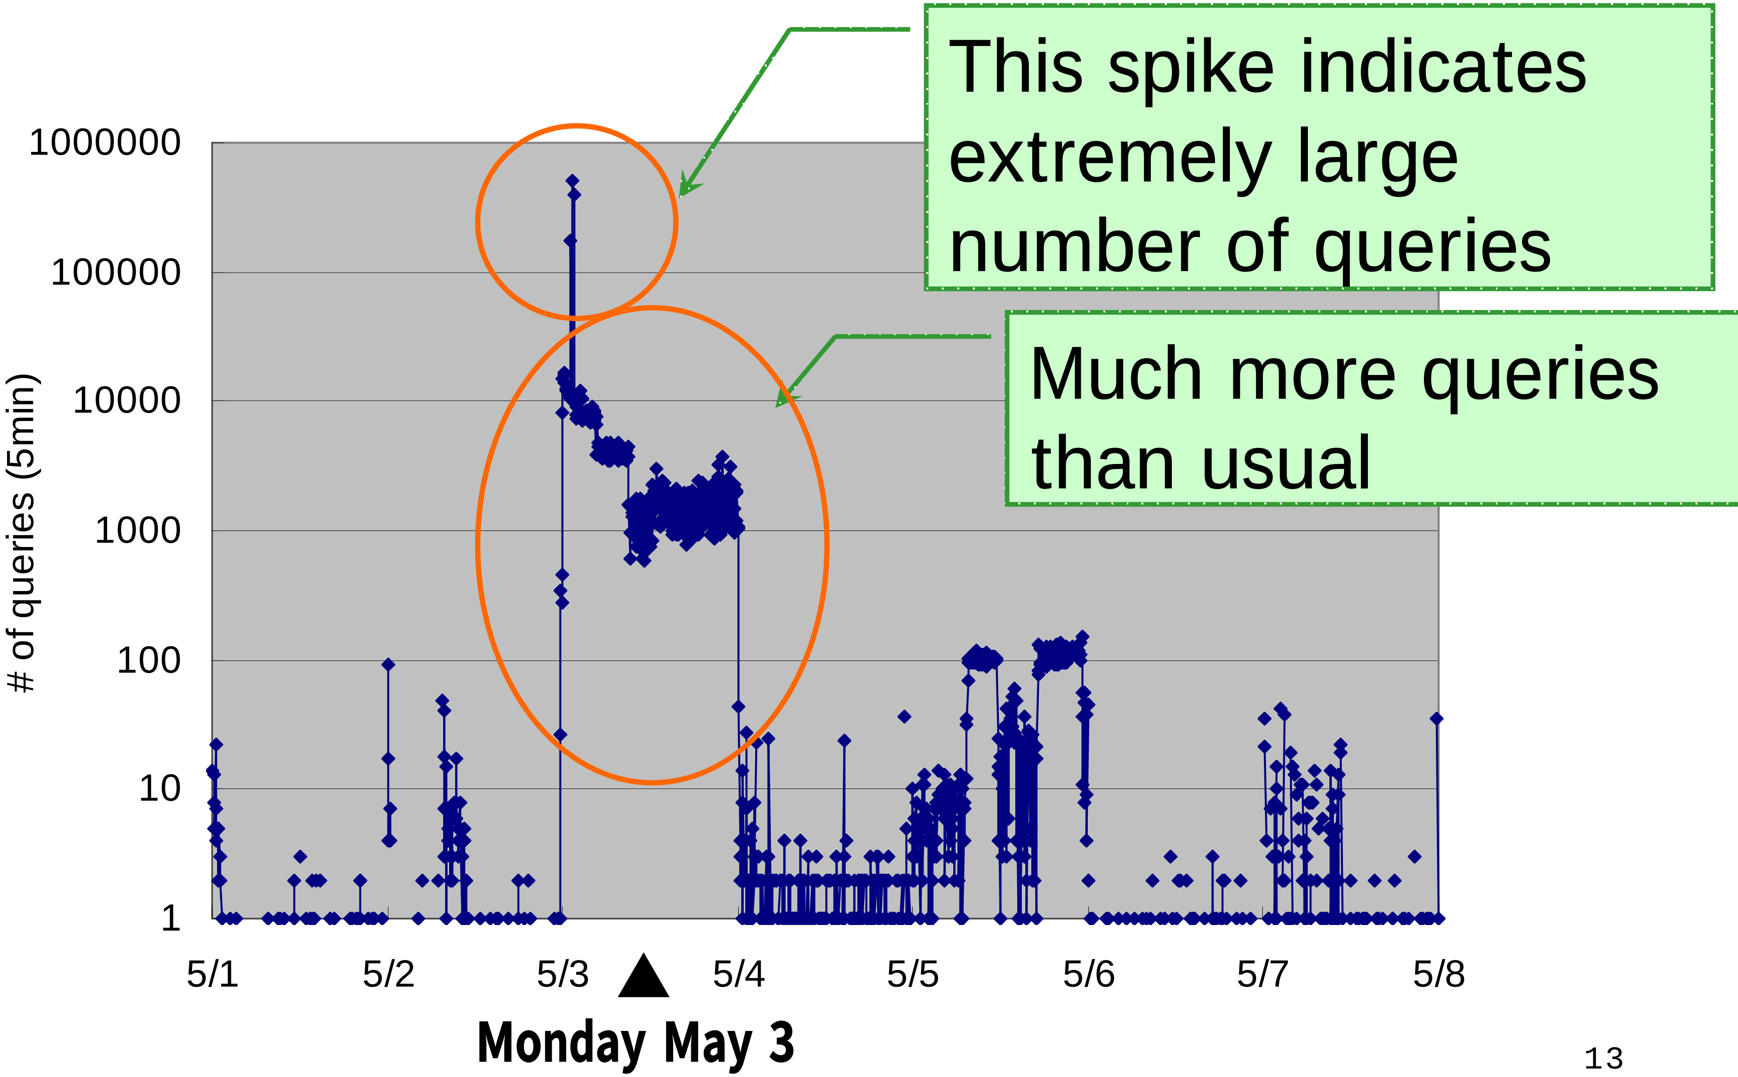

This spike indicates extremely large number of queries

Much more queries than usual

Number of queries resolving FQDN of target Web server (May 3rd, 01:00 – 03:00, per seconds)

- Queries extremely increased only when NXDOMAIN returned

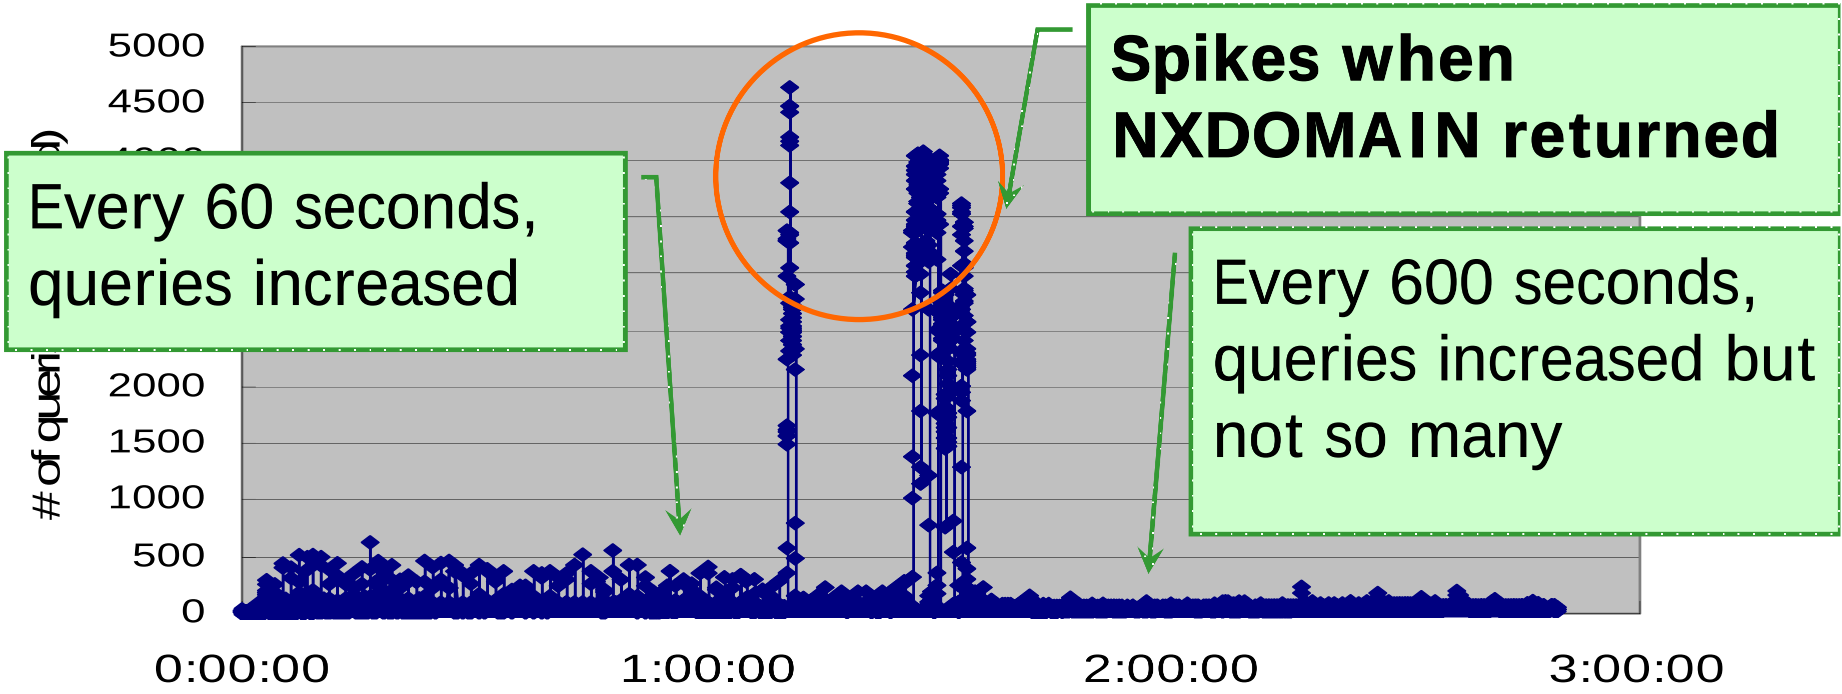

Spikes when NXDOMAIN returned

Every 60 seconds, queries increased

Every 600 seconds, queries increased but not so many

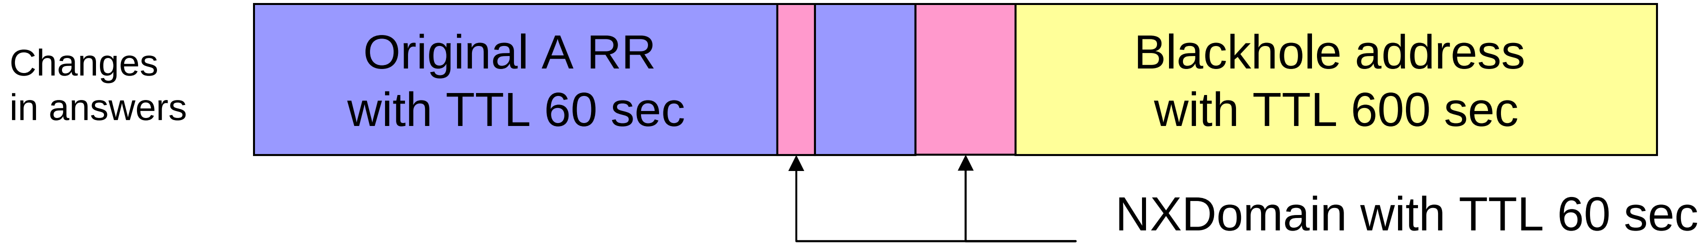

Original A RR with TTL 60 sec

Blackhole address with TTL 600 sec

NXDomain with TTL 60 sec

Number of unique IP address

(May 3rd, 01:00 – 03:00, per second)

- # of unique IP address increased only 2-3 times in “NXDomain period”, while queries increased 10 times or more!

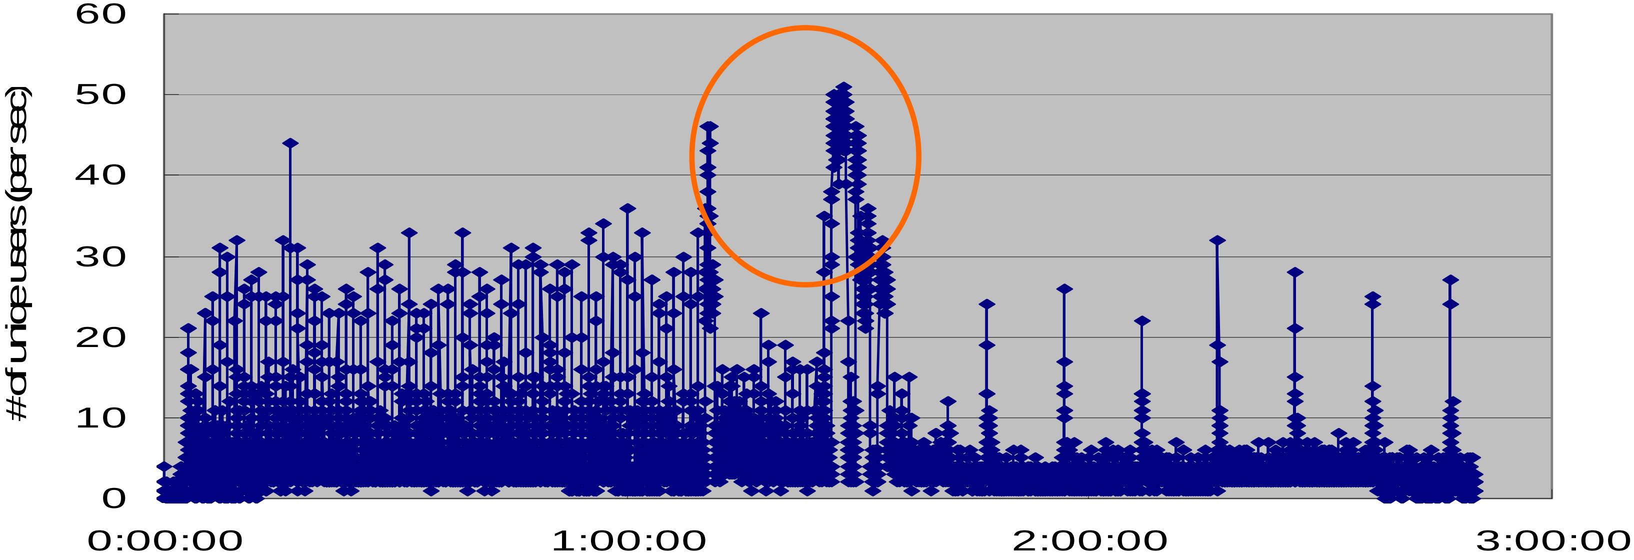

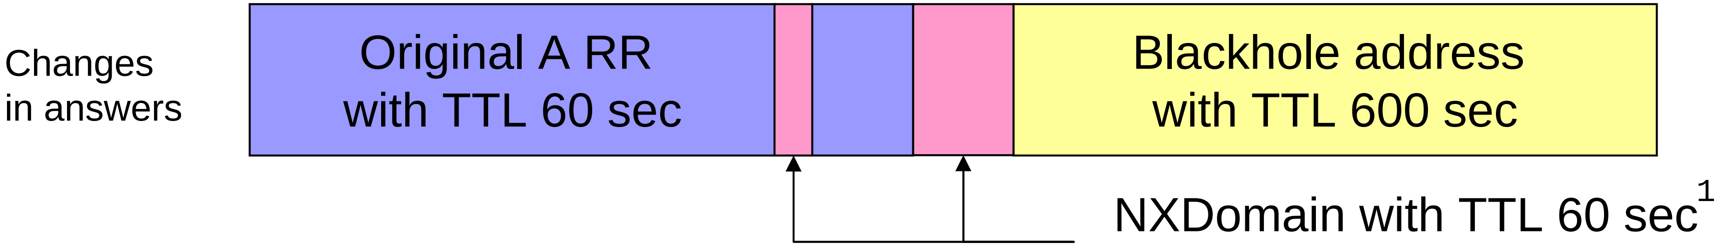
Changes in answers

Original A RR with TTL 60 sec

Blackhole address with TTL 600 sec

NXDomain with TTL 60 sec

Solution:
Return Blackhole IP Address by authoritative server

- In June, ISPs collaborated to defend the attack
- Asked the administrator of the authoritative server to return a blackhole address

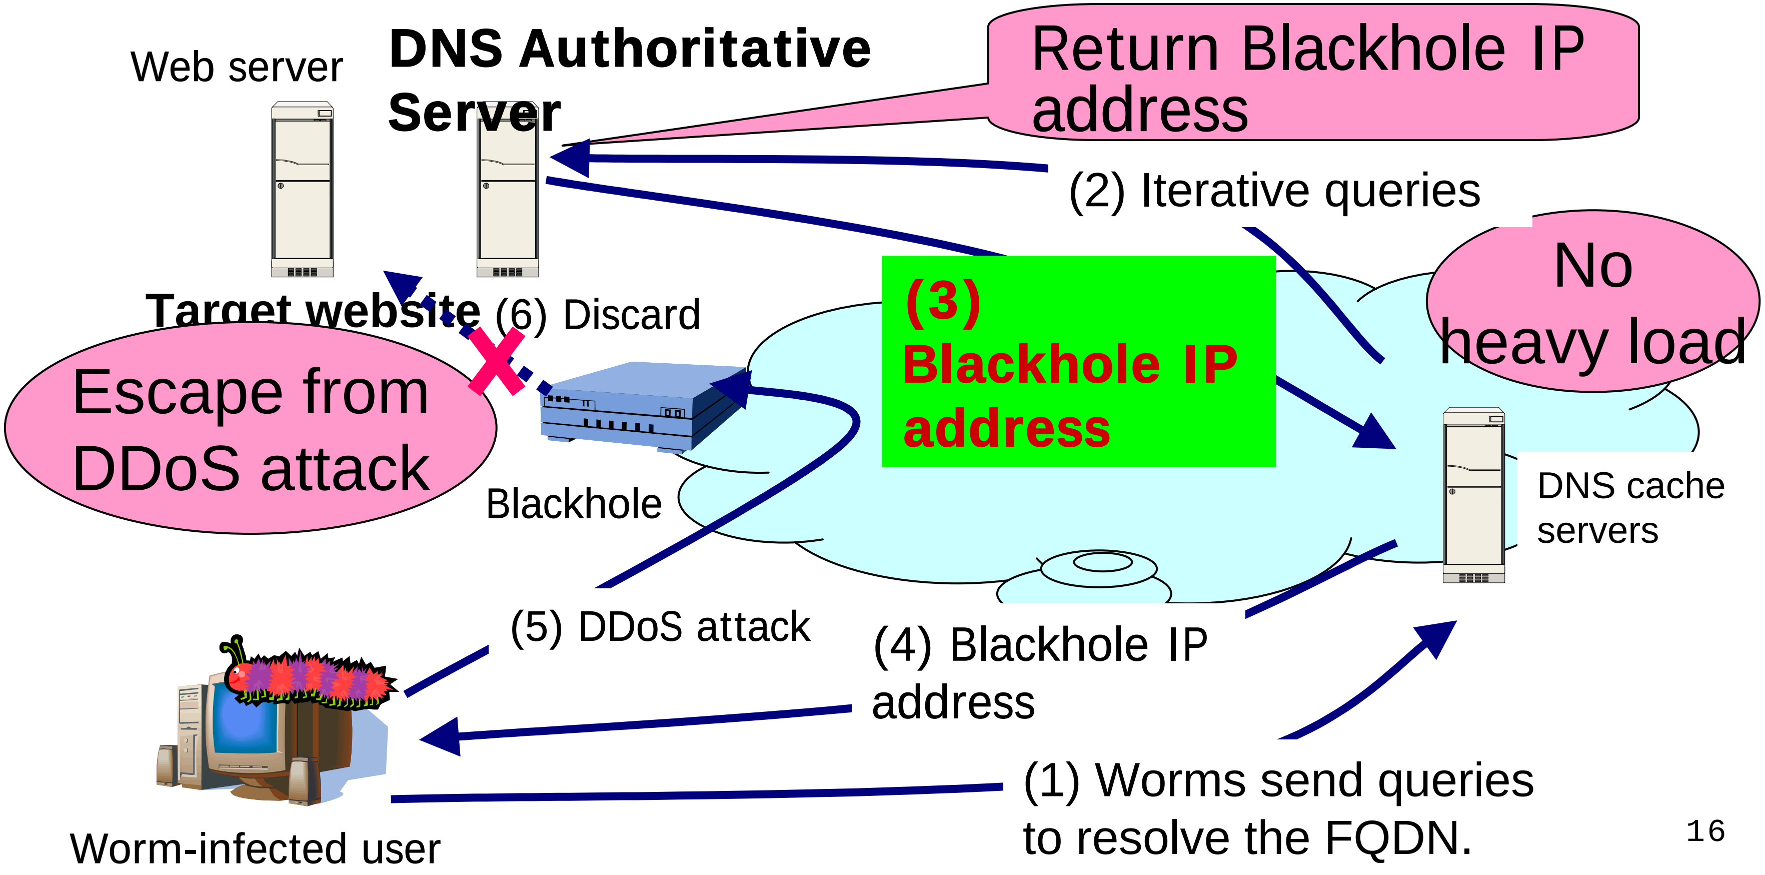

Lessons from the attack

- Removing A RR is not a good method
  - Behavior of worms, resolvers at user, and authoritative servers sometimes causes DNS cache server to collapse.
- Collaboration between victims and ISPs is required.
  - This time blackhole IP address worked very well.

- **Generic blackhole addresses** are necessary
  - If ISPs set the blackhole addresses to be discarded at their routers, then an authoritative server can return the addresses to escape from DDoS attacks.
  - For instance, TEST-NET (192.0.2.0/24, RFC3330)

2. Large RRSet and TCP filtering

Overload of DNS cache server

- Increasing load of DNS cache server, but..
- Number of total queries was normal
- We detected increase in the number of TCP queries

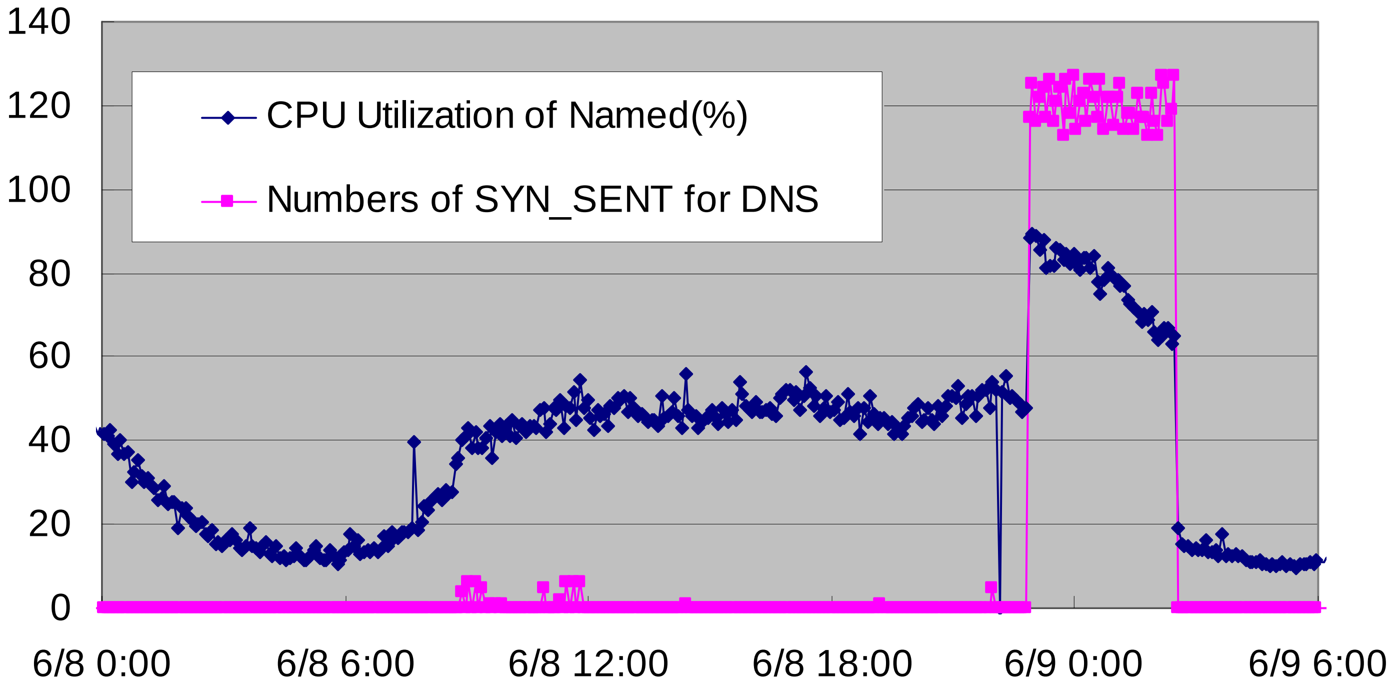

Why TCP queries were increased

- The number of truncated answers and that of TCP packets are synchronized
- We guessed truncated answers caused TCP queries.

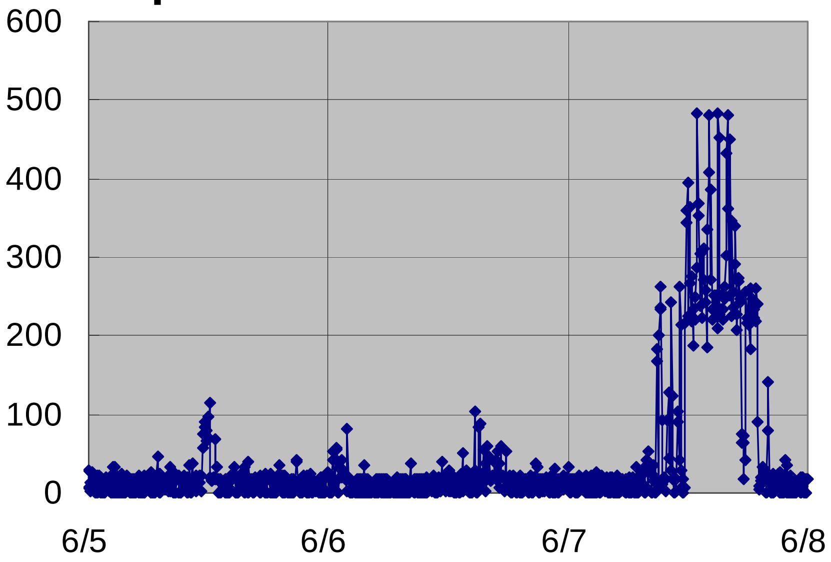

Number of TCP packets sent to authoritative servers (per 5 min.)

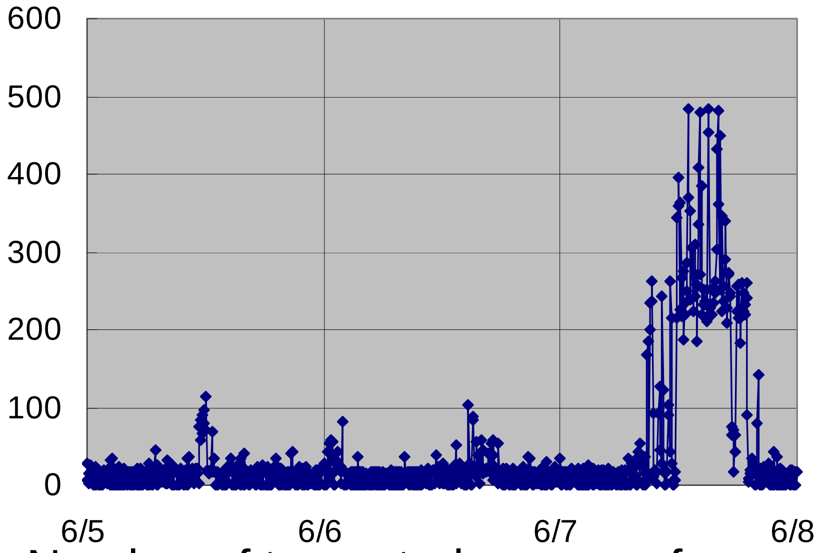

Number of truncated answers from authoritative servers (per 5 min.)

Why TCP queries increased (cont'd)

- What happened was…

(1) The cache server sent iterative query to the authoritative servers.

(2) The authoritative servers answered with a truncated bit; over 512 octets answer

(3) Then the cache server tried to resend query by TCP.

The authoritative servers did not respond to TCP query.

- The authoritative servers didn't support EDNS0
- TCP packets seemed to be filtered before authoritative server process
- The cache server was waiting 75 seconds for acknowledgement from the authoritative servers, and pending queries were increased.

⇒ the cache server overloaded

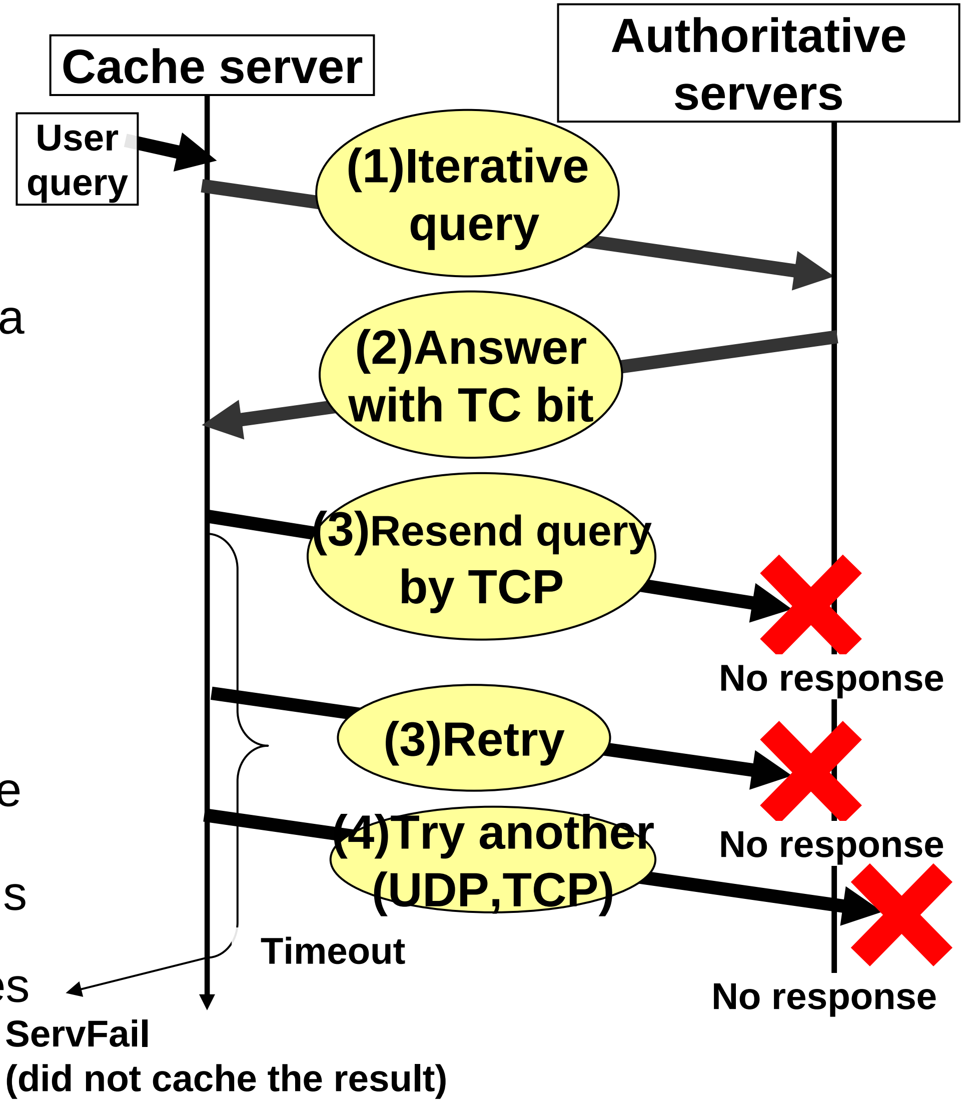

Dump example

```
% dig @yyy.yyy.yyy.yyy -x zzz.zzz.zzz.zzz
```

(1) Iterative query

```
21:23:14.604332 IP xxx.xxx.xxx.xxx.51705 > yyy.yyy.yyy.yyy.53:  24775+ PTR?
      zzz.zzz.zzz.zzz.in-addr.arpa. (42)
```

TC bit

(2) Answer with TC bit

```
21:23:14.744362 IP yyy.yyy.yyy.yyy.53 > xxx.xxx.xxx.xxx.51705:  24775| 15/0/0 PTR
      redirect.***.co.in., PTR onlyoriginal.com.my., PTR redirect.***.com.cn., PTR
      redirect.china.***.com., PTR redirect.jp.***.com., PTR ***.co.jp., PTR redirect.***.co.jp., PTR
      ***.ne.jp., PTR redirect.***.co.kr., PTR redirect.kt***.co.kr., PTR redirect.***.com.my., PTR
      redirect.***.com.ph., PTR redirect.***.com.sg., PTR redirect.***.co.th., PTR
      redirect.***.com.tw. (488)
```

(3) Resend query by TCP (SYN packet)

```
21:23:14.744748 IP xxx.xxx.xxx.xxx.60035 > yyy.yyy.yyy.yyy.53: S 2600114847:2600114847(0)
      win 65535 <mss 1460,nop,nop,sackOK,nop,wscale 1,nop,nop,timestamp 1455660 0>

21:23:17.736935 IP xxx.xxx.xxx.xxx.60035 > yyy.yyy.yyy.yyy.53: S 2600114847:2600114847(0)
      win 65535 <mss 1460,nop,nop,sackOK,nop,wscale 1,nop,nop,timestamp 1455960 0>

21:23:20.937001 IP xxx.xxx.xxx.xxx.60035 > yyy.yyy.yyy.yyy.53: S 2600114847:2600114847(0)
      win 65535 <mss 1460,nop,nop,sackOK,nop,wscale 1,nop,nop,timestamp 1456280 0>
```

Workaround

1. Asked the administrators of the authoritative server
   - to change the settings to accept TCP connection, or
   - to decrease the size of their records to fit them into a UDP packet.

--> But the administrators will not change the settings…

2. Denied the users sending the query
   - by using BIND blackhole setting

Another possible approach

- Patching BIND, not to query by TCP when truncation occurs.
  - Quick hacking could reduce TCP sessions, but it may violate RFC.
  - Proper modification would be:
    - Cache the following information for the RR during the TTL
      - TC = 1
      - Does not accept TCP query
    - Return ServFail without iterative query if above cache exists
  - …this should be proposed to IETF

Lessons through this phenomenon

- Strongly recommend that administrators check the configuration of authoritative servers.
  - Answer TCP queries (mandatory)
    - RFC 1123
      DNS servers must be able to service UDP and **should** be able to service TCP queries … it **should not refuse to service a TCP query** just because it would have succeeded with UDP.
  - and either
    - Set the size of answer to be smaller than 512 octets, or
    - Support EDNS0 option

Summary

To be ‘friendly’ to DNS cache servers,

- Administrators should check and modify the settings of authoritative servers.
  - Some generic blackhole address should be used, instead of removing A RR in case of DDoS.
  - Configuration should be consistent.
    - Oversized RRSet with no EDNS0 and closed TCP port is not good but often seen.

Contacts

- Chika YOSHIMURA (NTT Communications)
    - yosimura@ocn.ad.jp
- Katsuyasu TOYAMA (NTT)
    - toyama.katsuyasu@lab.ntt.co.jp

- Special Thanks to:
  - Ichiro Mizukoshi (NTT Communications)
  - Haruhiko Ohshima (NTT Communications)
  - Yasuhiro Morishita (JPRS)
  - Hirotaka Matsuoka (NTT)