

(Well-known)
"Redirect the target domain's
nameserver"
cache poisoning attacks

Kazunori Fujiwara, JPRS

fujiwara@jprs.co.jp

IEPG meeting, Toronto

2014/7/20

(Well-known) references

- Bernhard Müller, “Improved DNS spoofing using node re-delegation”, August 2008
- Wikipedia DNS spoofing
 - http://en.wikipedia.org/wiki/DNS_spoofing

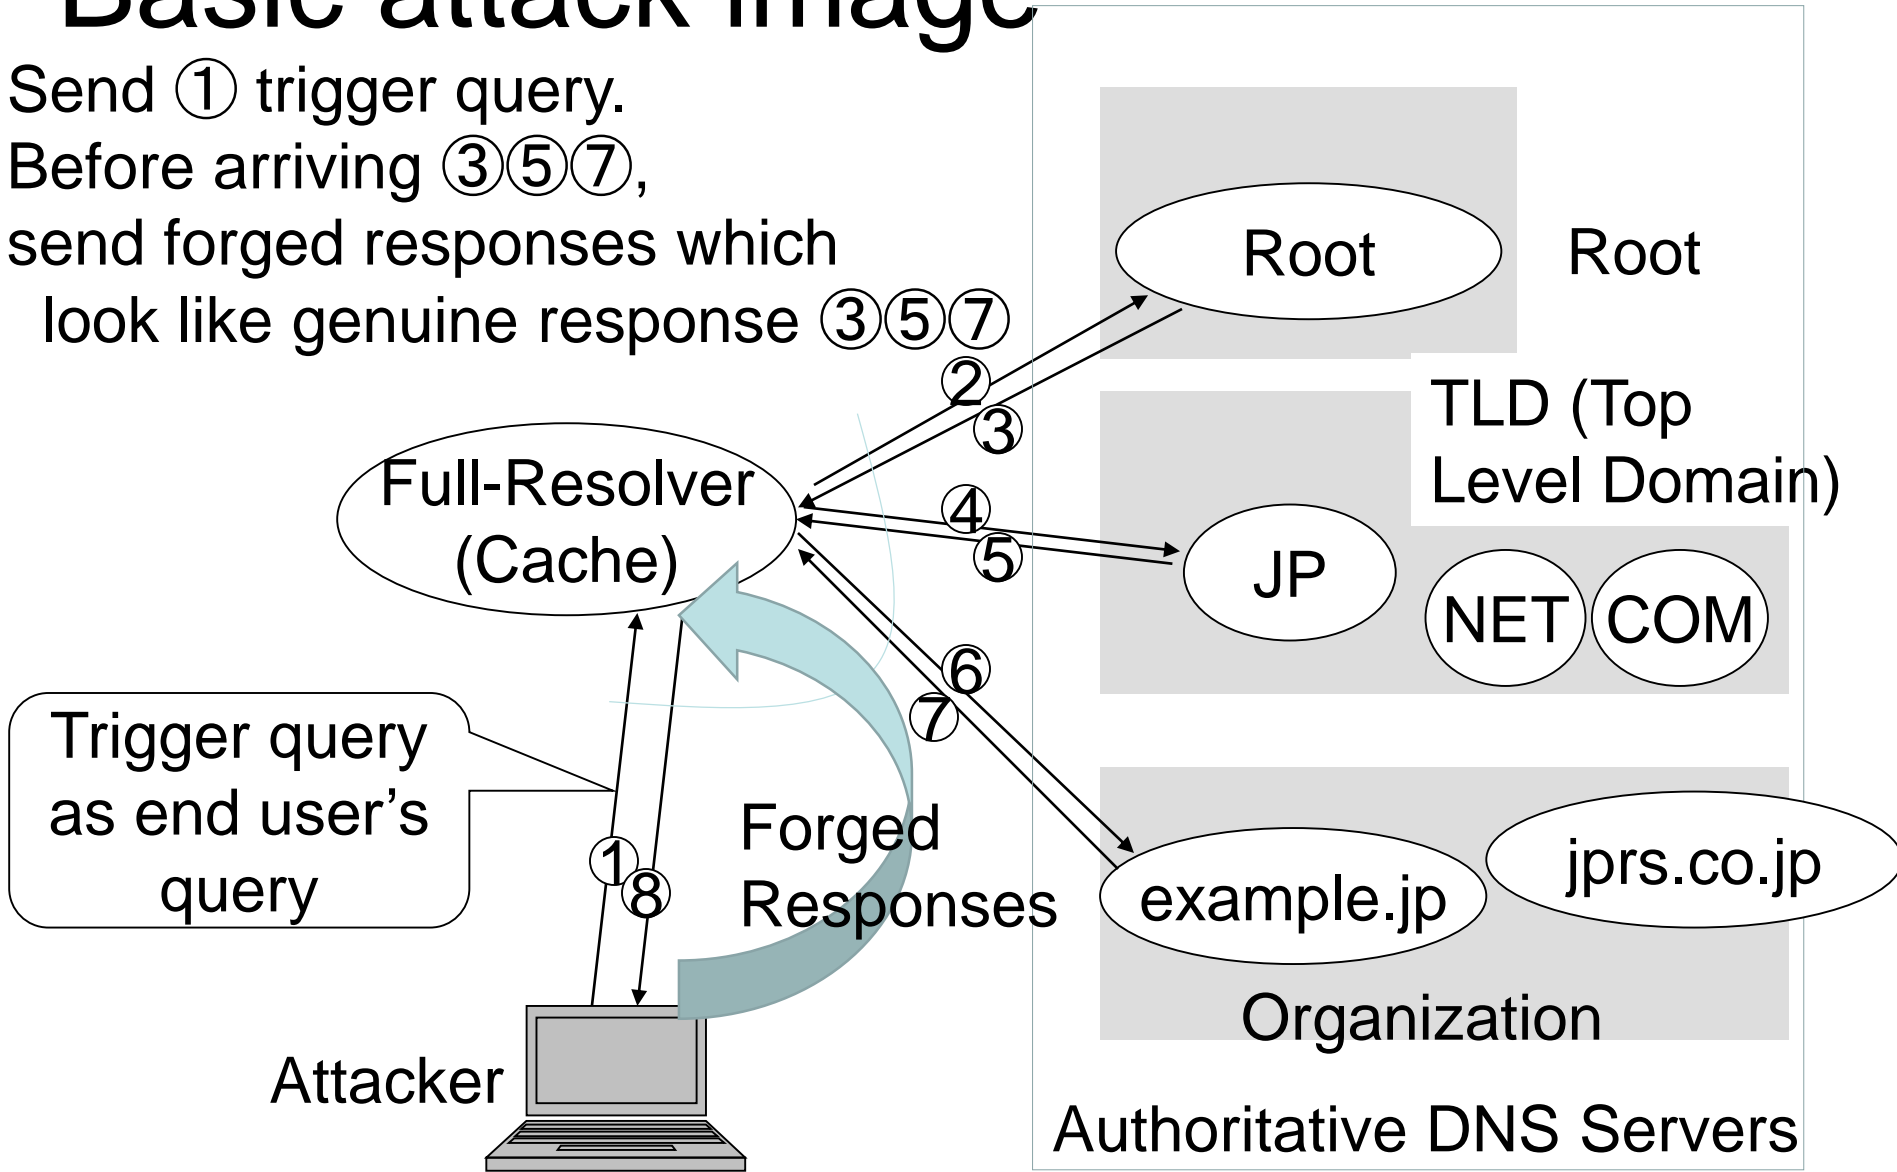
Redirect the target domain's nameserver

The first variant of DNS cache poisoning involves redirecting the nameserver of the attacker's domain to the nameserver of the target domain, then assigning that nameserver an IP address specified by the attacker.

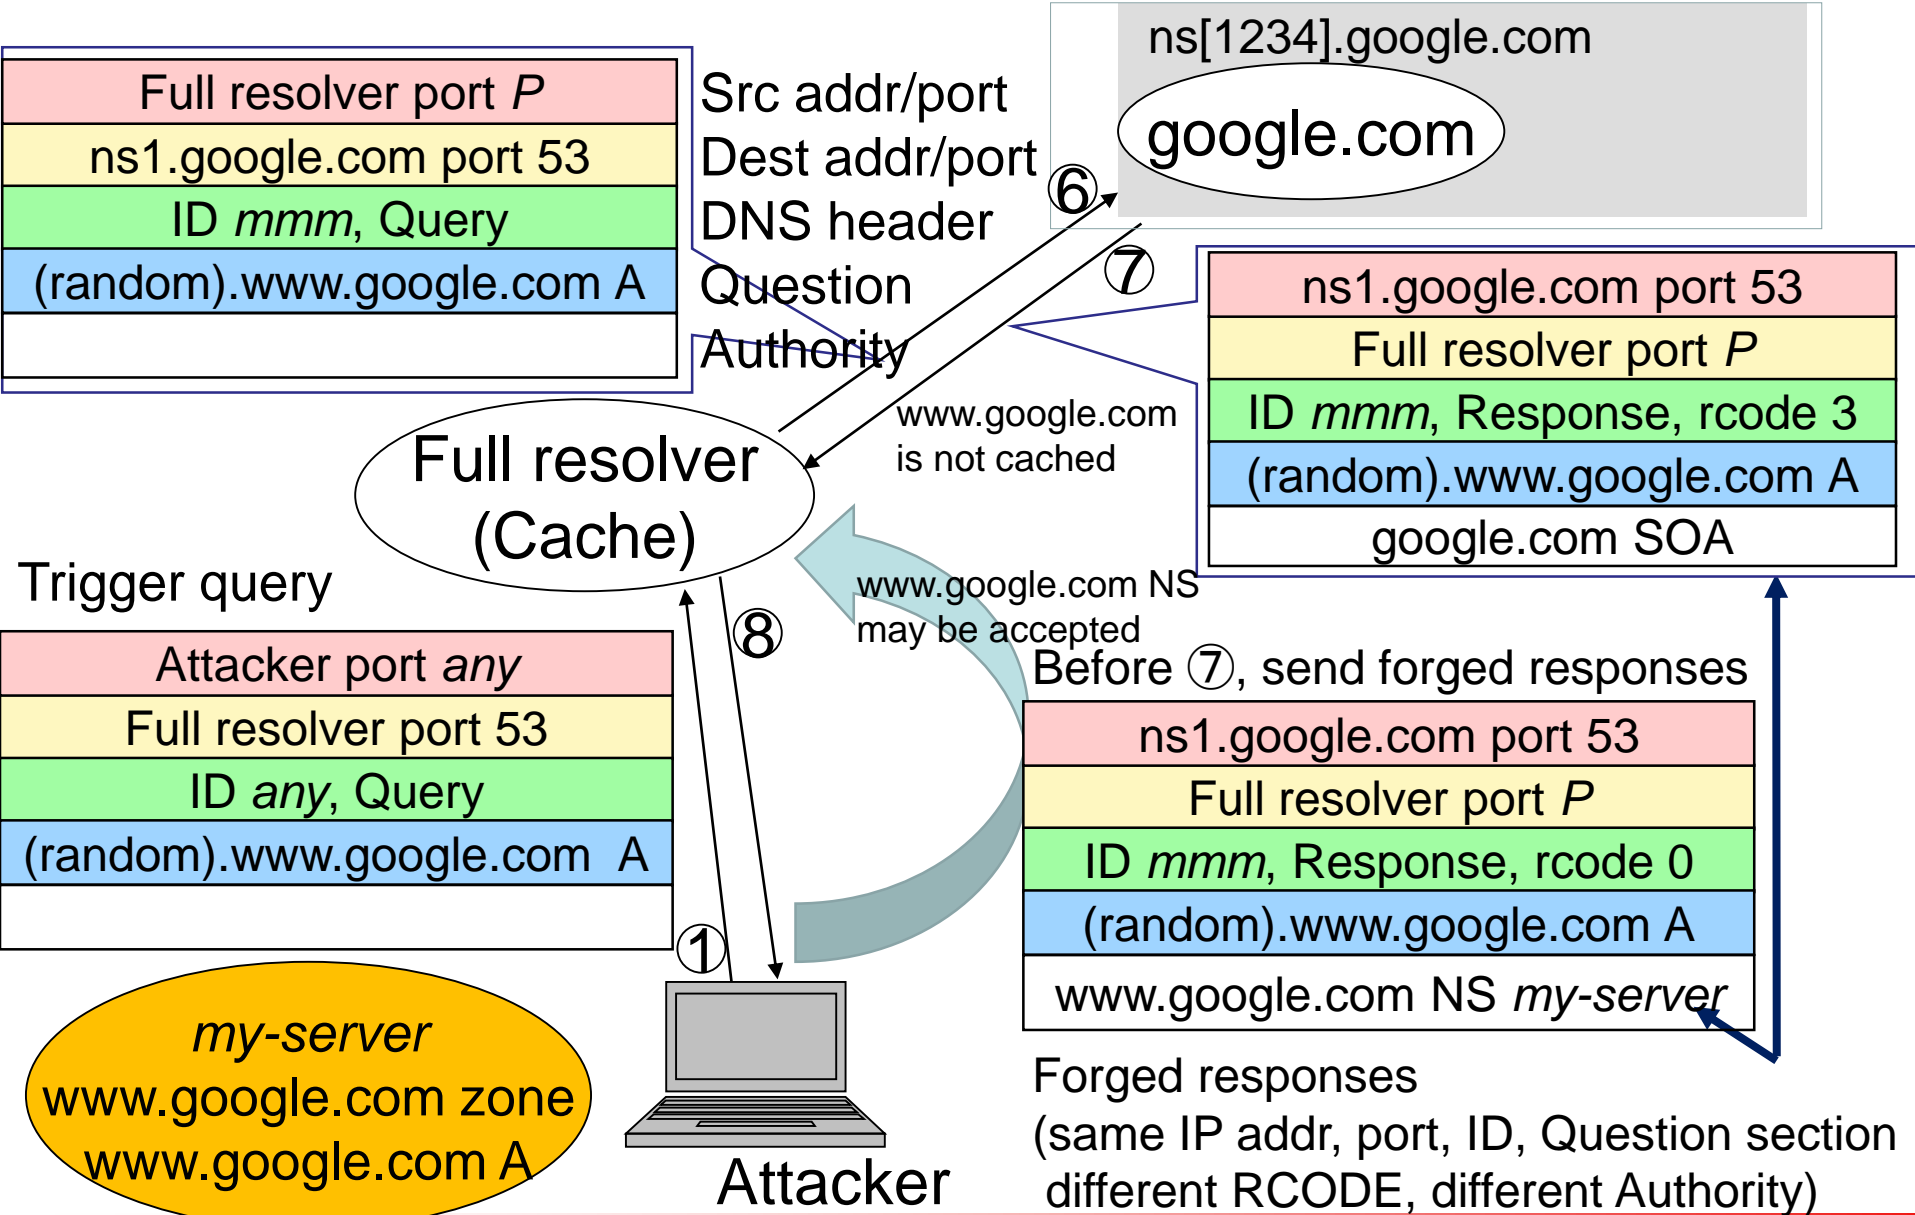
- Dan Kaminsky, “DNS 2008 and the new (old) nature of critical infrastructure”, July 2008
- RFC 3833 Threat Analysis of the Domain Name System (DNS)
 - 2.2. ID guessing and query prediction

Basic attack image

Send ① trigger query.
Before arriving ③⑤⑦,
send forged responses which
look like genuine response ③⑤⑦



Detailed attack image (www.google.com)



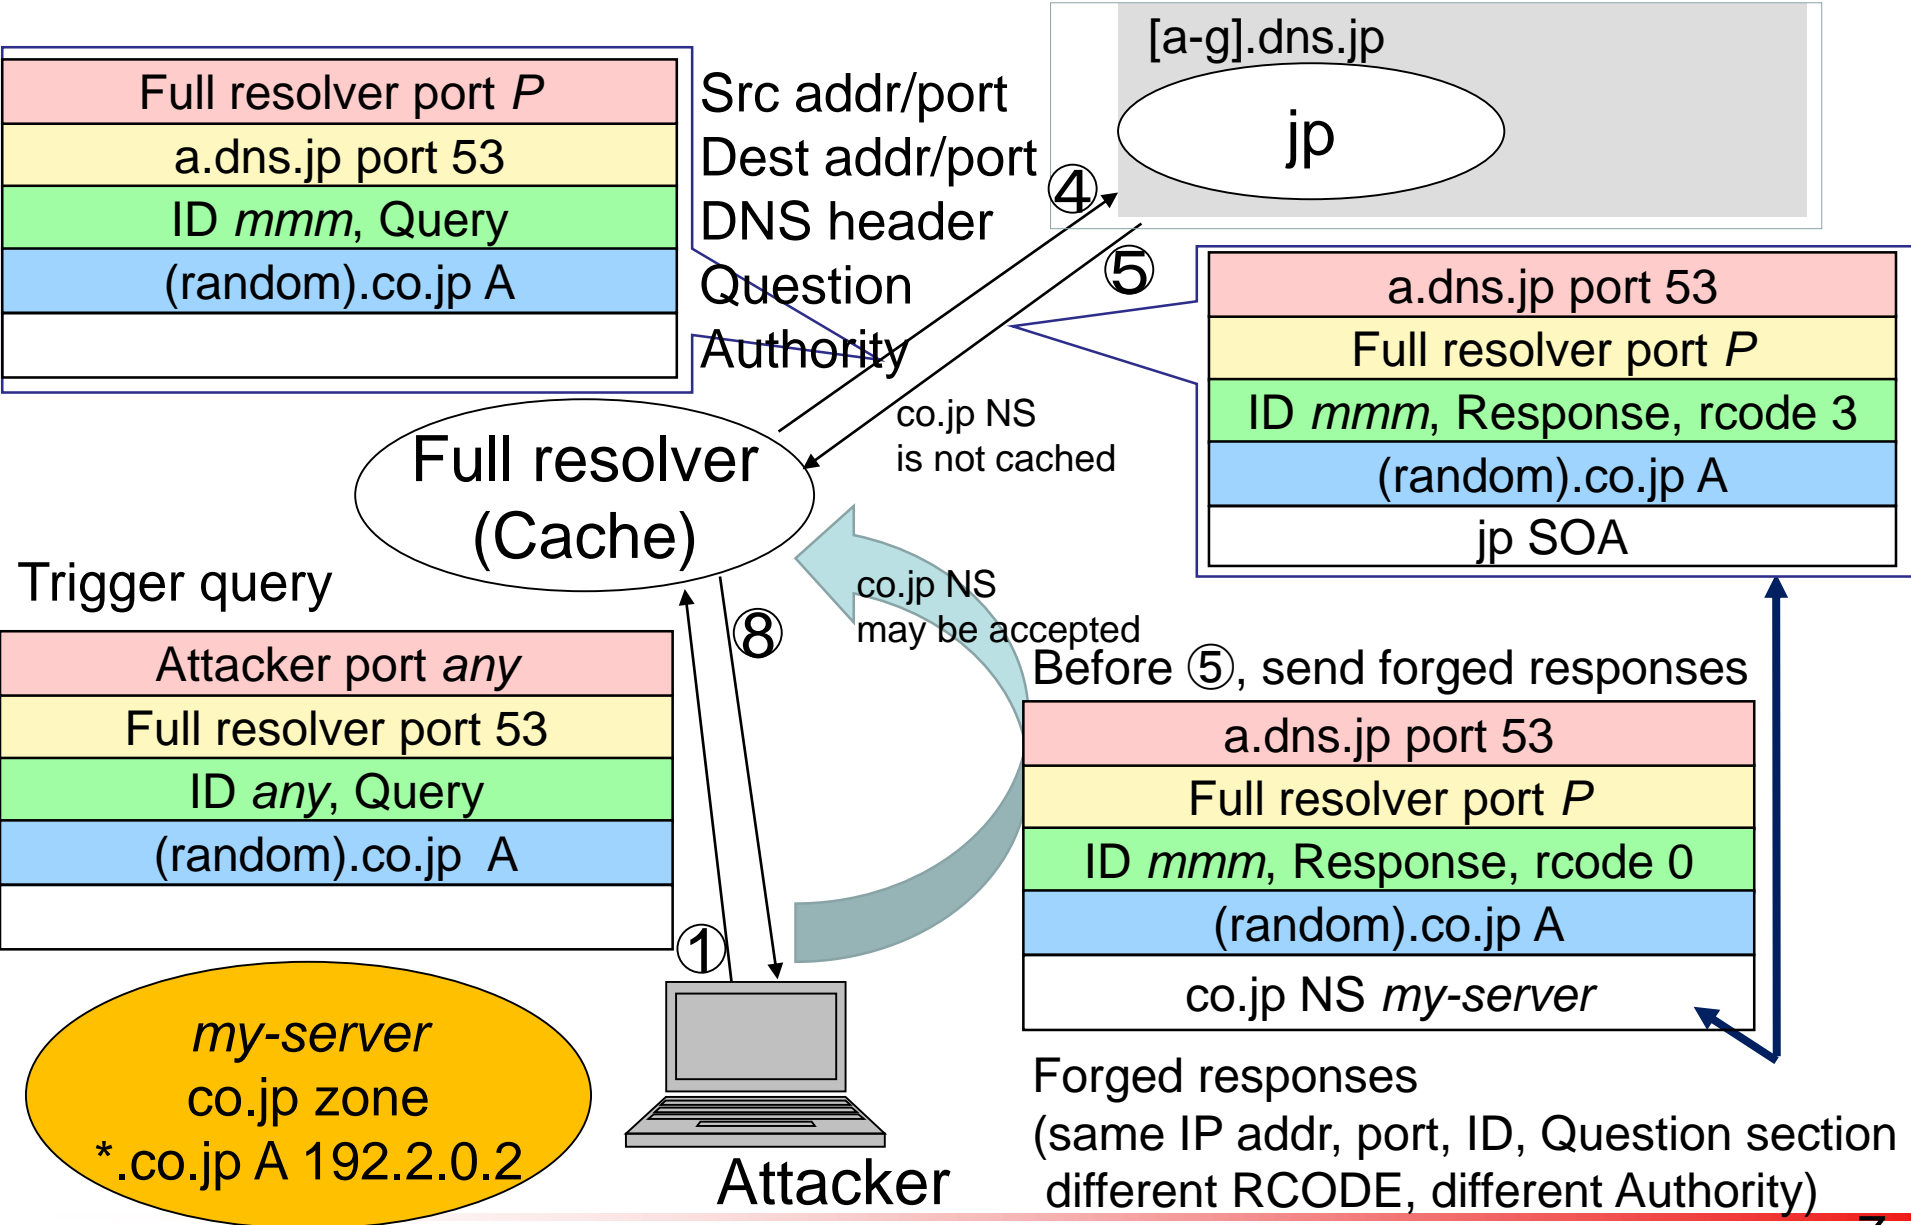
Attack details

- Choose a trigger domain name (\$Trigger)
- Send a trigger query “(random).\$Trigger A”
 - Authoritative server returns name error (rcode 3) for the query name
 - The trigger domain name is not cached by the full resolver
- Send forged responses
 - IP src = authoritative server **addresses** of Trigger domain
 - IP dst = full resolver’s address
 - src port = 53
 - dst port = full resolvers port (static, or **random number**)
 - DNS header: Rcode 0, response, ID = **random number**
 - Question section = “(random).\$Trigger A”
 - Authority = what you want to inject : “Target NS my-server”
- Prepare Target zone in my-server
 - *.Target IN A 192.2.0.2

Target domain names

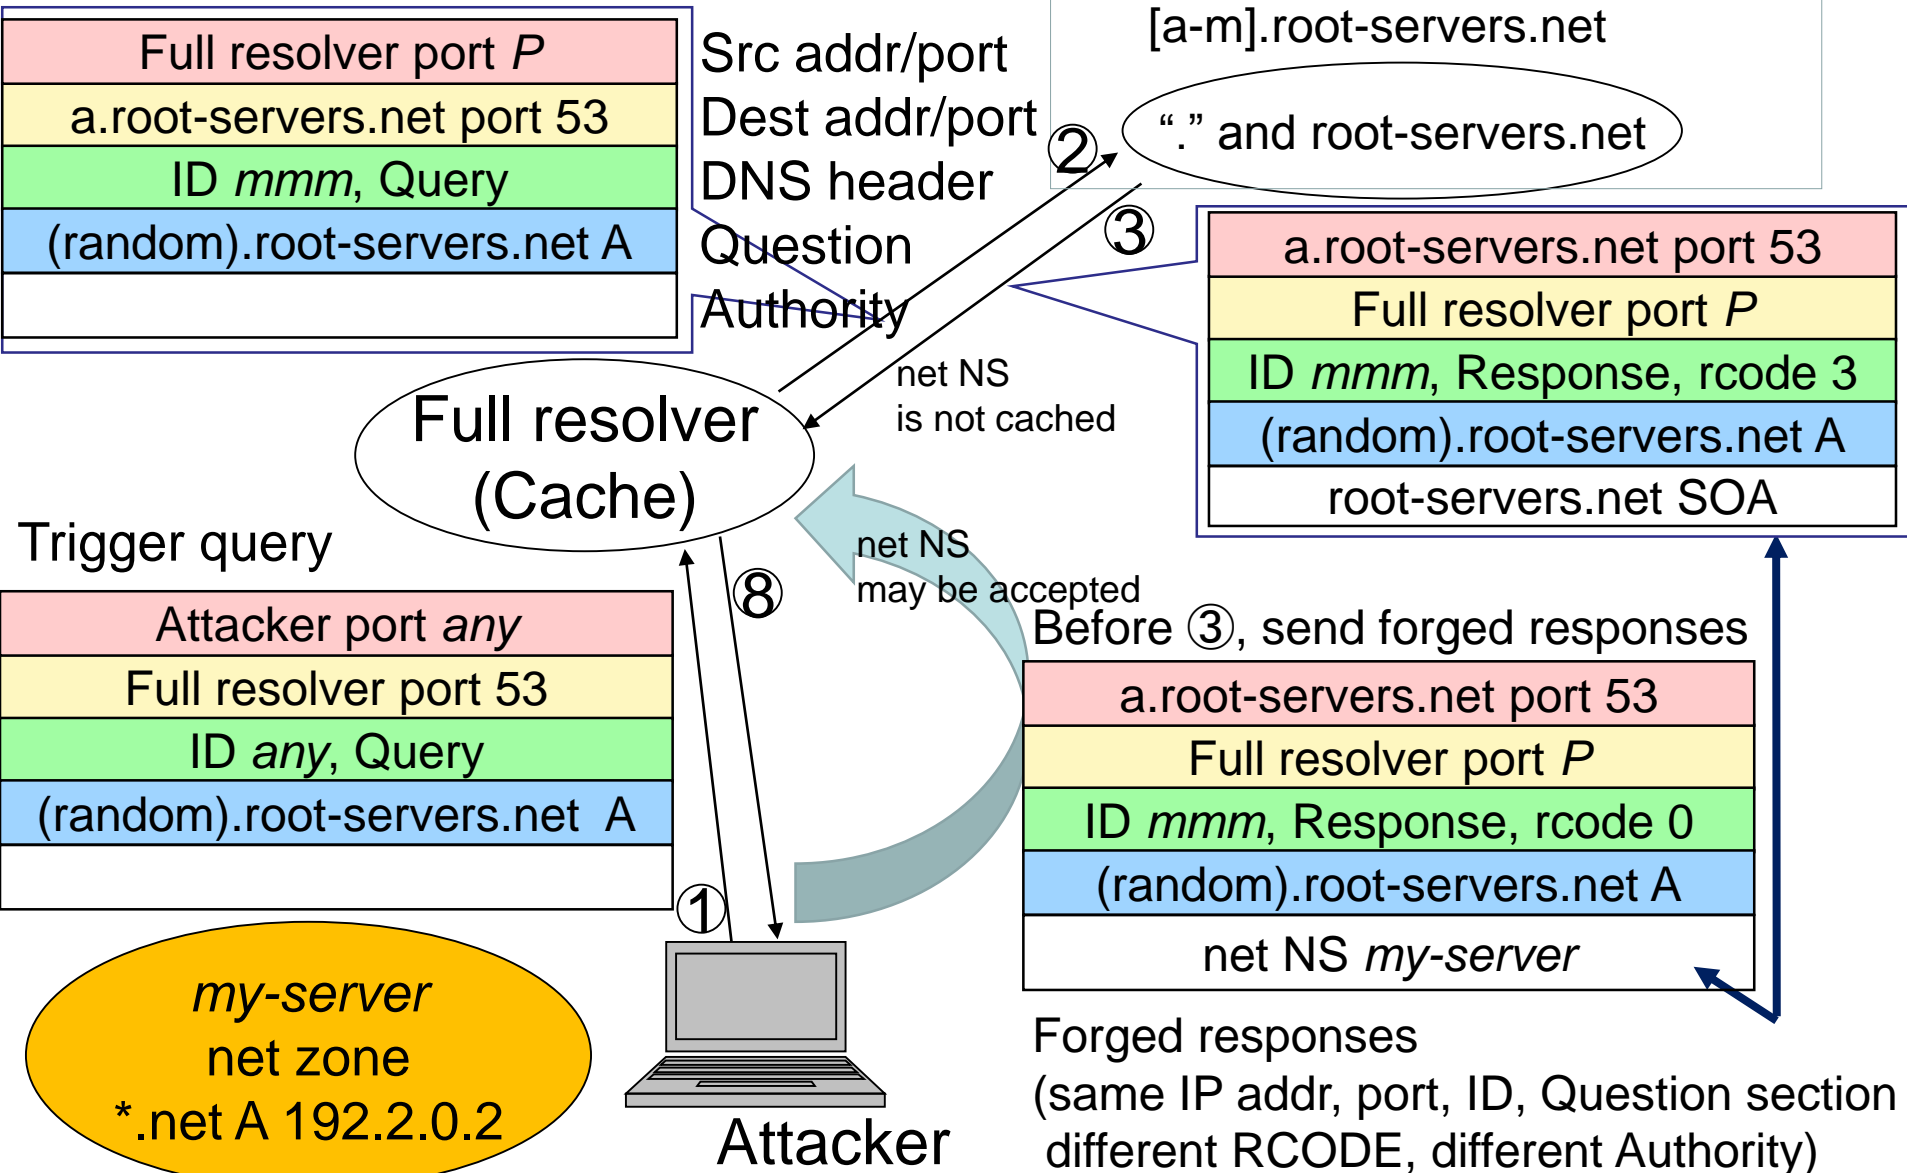
1. Domain names which does not have NS
 - **Host names** (www.google.com, www.iepg.org, ...)
 - Some **non-terminals** (co.jp, tokyo.jp, ...)
2. Some zone cuts
 - When an authoritative server serves both a zone and its descendant zones
 - The shallowest zone apex is not a target
 - If a deepest name is a zone cut, it is not a target
 - Example 1: “**net**” and “**root-servers.net**”
 - Root servers serve both “.” and “root-servers.net”
 - Example 2: “**co.uk**”
 - uk servers serve both “uk” and “co.uk”

Attack example: **co.jp**



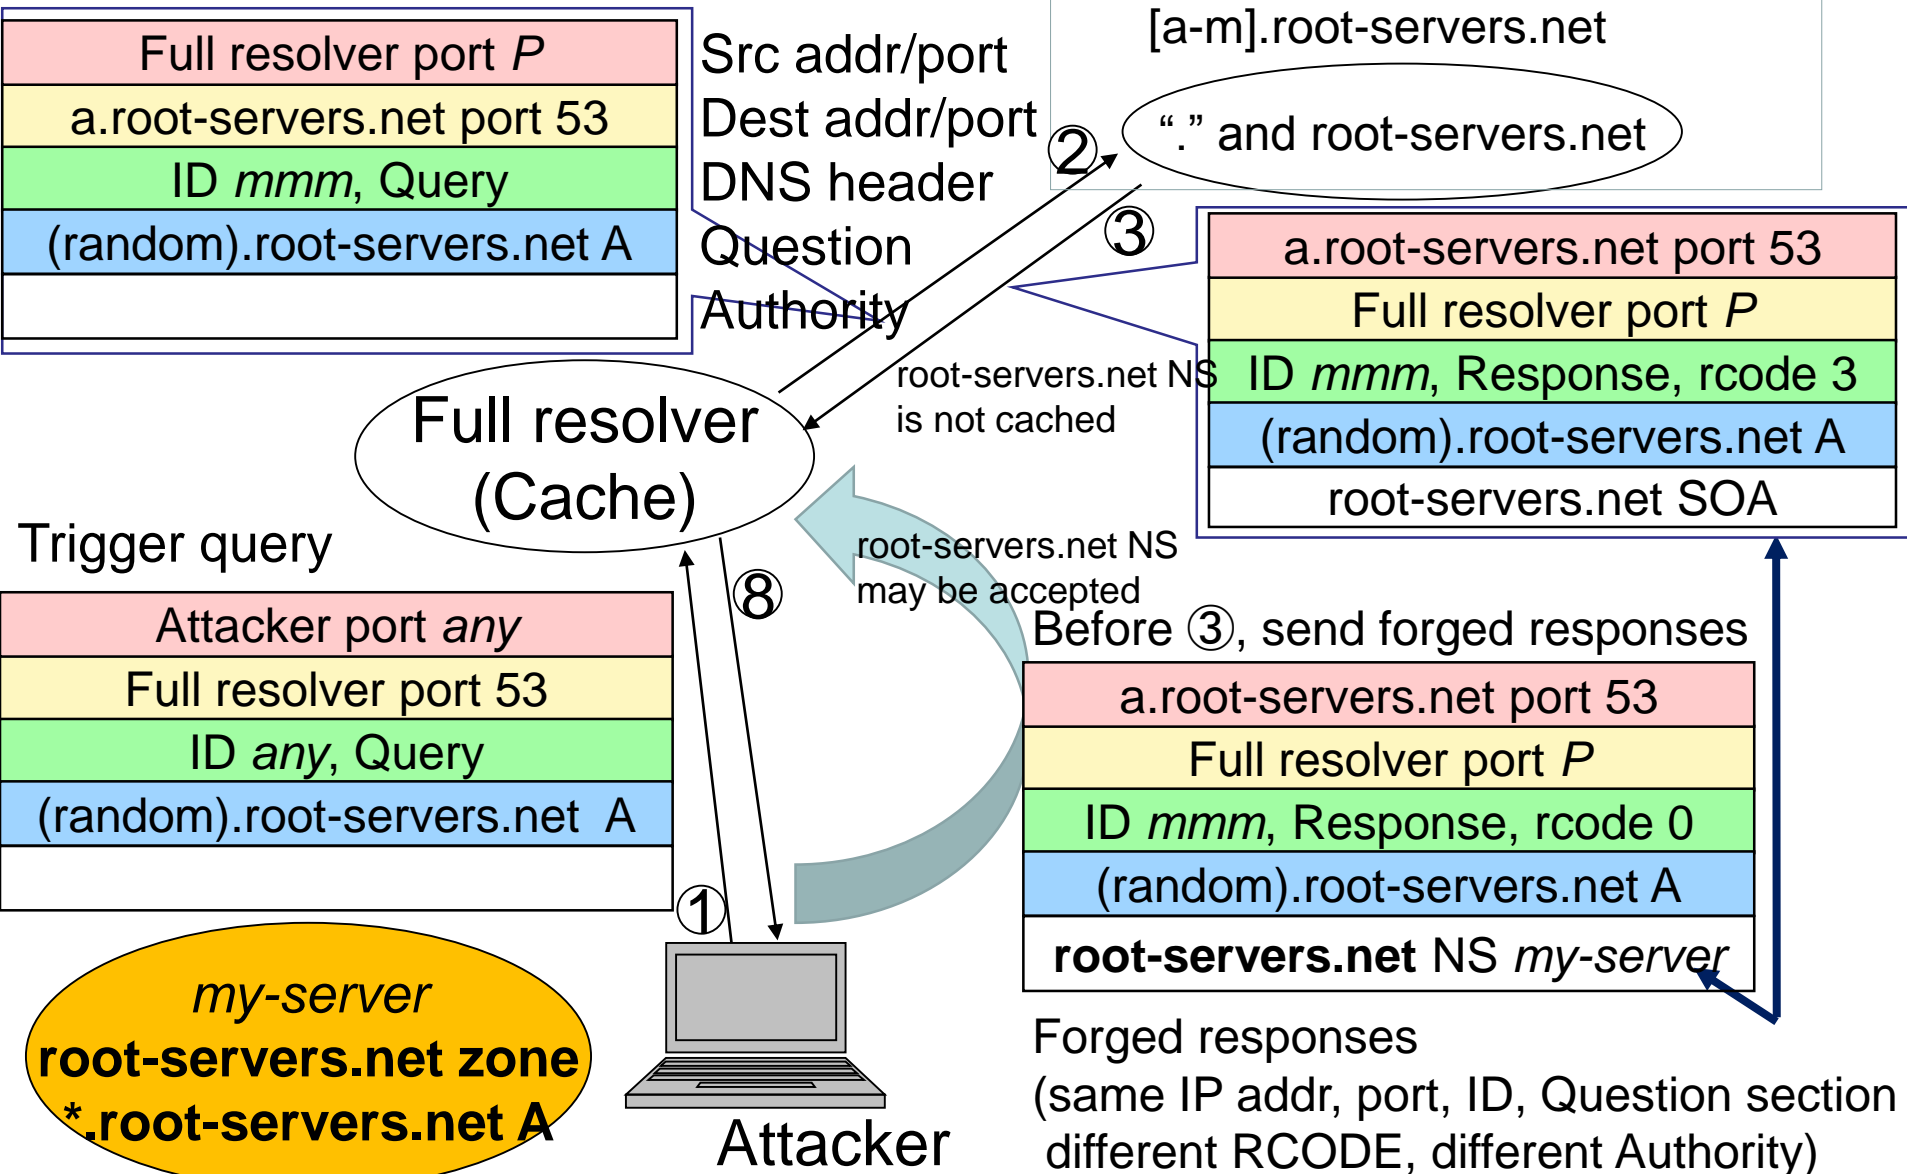
Attack example:net

Only when "net NS" is not cached by victim full resolvers

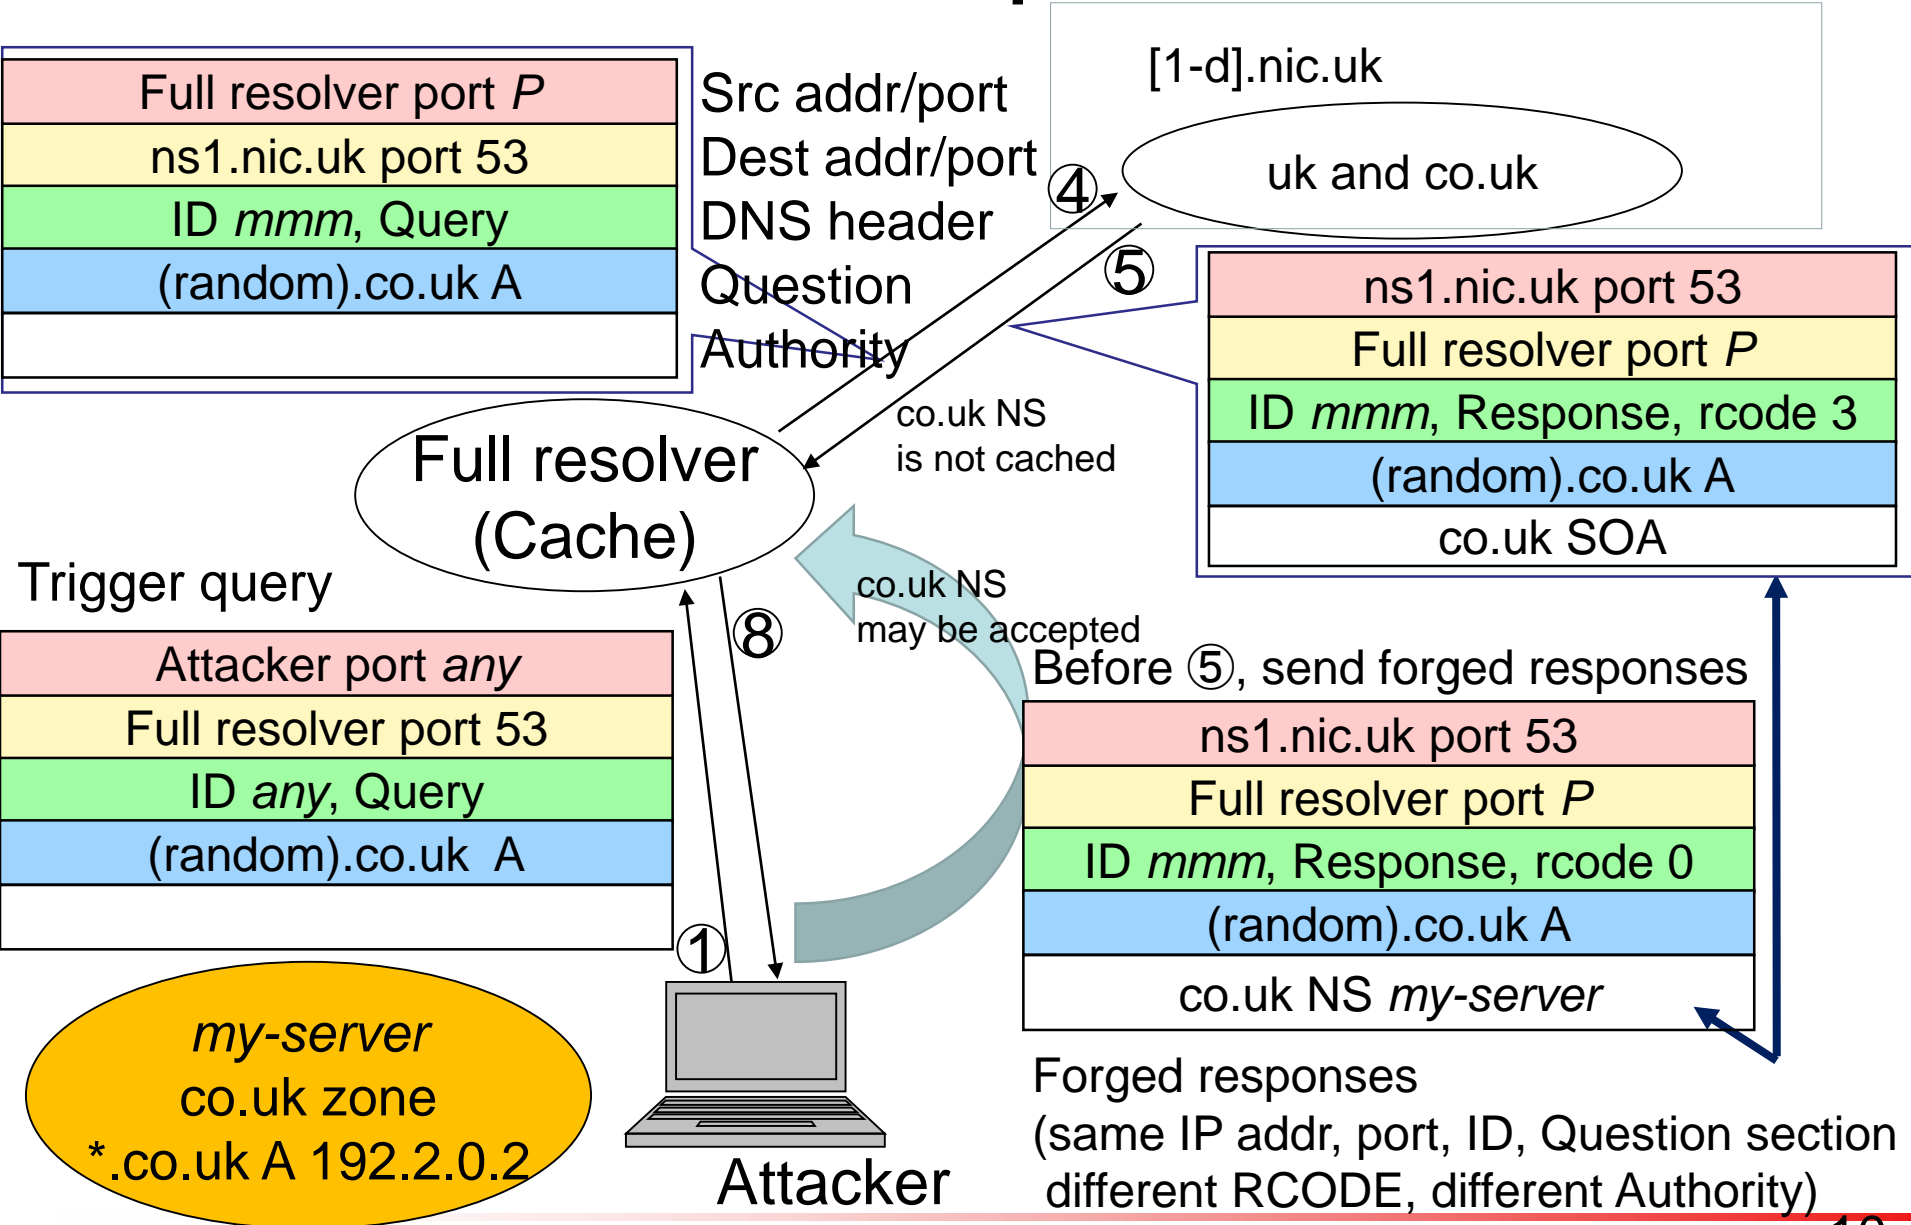


Attack example: root-servers.net JPRS JAPAN REGISTRY SERVICES

~~Only when "net NS" is not cached by victim full resolvers~~



Attack example: **co.uk**



Attack success probability and expected value of attack time

	Variable	Unit	Values	Example values
Num of resolver's ports	N_{port}		$1 \sim 2^{16}$	1 or 64,000
Num of IDs	N_{qid}		2^{16}	65,536
Num of auth servers	N_{ns}		$1 \sim 13$	4
RTT to auth servers	T_{auth}	second	$0.001 \sim 0.2$	0.039
Loop time	T_{loop}	second		0.020
Response send rate	R_{ans}	packets/sec		100,000

- Success probability of first time trial

$$P = \frac{T_{auth} * R_{ans}}{N_{qid} * N_{port} * N_{ns}}$$

Example:
0.0148
Or
0.00000023

- Success probability under continuation attack is 1
- Expected attack time

$$T = \frac{N_{qid} * N_{port} * N_{ns}}{R_{ans}}$$

Example:
2.62 seconds
Or
167,772 (2 days)

My attack tool

- Written in C, 500 lines, standard library only
 - Using select() to control timing and sockets
 - Send forged responses using raw socket
- Tool's parameters
 - Victim resolver's IP address and port number
 - Trigger domain name
 - NS RR information
 - Authoritative servers' IP address list
- Usage example
 - `./a.out 192.2.0.2 20001 www.google.com
www.google.com ns.dnslab.jp
216.239.32.10/216.239.34.10/216.239.36.10/216.2
39.38.10`

Diagram of the attack tool

Initialization
prepare two sockets
UDP and Raw
Initialize timer
Send first trigger query

Loop
select
If readable,
receive response
if success --> Termination
if fail --> Send trigger query
After Tloop, Send trigger query
If writable to raw socket,
send forged responses

Termination
Display the results

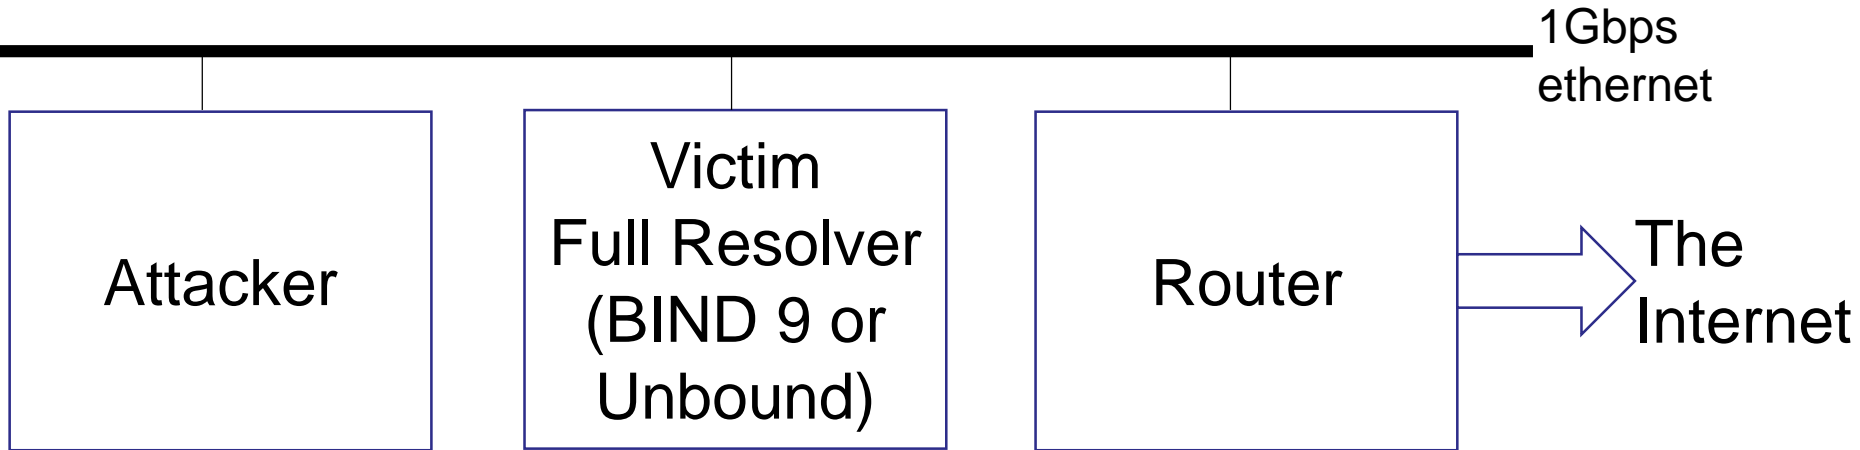
Send trigger query
generate (random) label
send "(random).\$Trigger
A" queries to victim resolver

Send forged responses
generate random ID
choose authoritative server address
from address-list
Question Section is the same as
sent trigger query
Authority Section is forged NS RR
from auth port 53 to victim resolver
send packet using Raw socket

Attacker address port <i>any</i>
Full resolver port 53
ID <i>any</i> , Query
(random).\$Trigger A

Auth server address port 53
Full resolver port <i>P</i>
ID <i>random</i> , Response, rcode 0
(random).\$Trigger A
Injecting NS RR

Environment of attack experiments



- Two VMs (Attacker and Victim)
 - 2.5GHz Xeon, 2 cpu, 3GB memory
 - On same link (1Gbps)
 - Victim full resolver sends queries to the Internet
- Victim full resolver
 - BIND 9.9.5 (without source port randomization)
 - (Unbound 1.4.21 without source port randomization)

Two attack experiments

- `www.google.com`
 - `% sudo ./a.out 203.178.129.2 20001 www.google.com www.google.com ns.dnslab.jp 216.239.32.10/216.239.34.10/216.239.36.10/216.239.38.10`
 - The injection took 46 seconds,
 - 28 qps queries and 66,726 pps forged responses
- `net`
 - “**rndc flush**” on victim full resolver
 - `sudo ./a.out 203.178.129.2 20001 root-servers.net net ns.dnslab.jp 198.41.0.4/192.228.79.201/...(root server addresses)`
 - The injection took 8 seconds,
 - 30 qps queries and 69,158 pps forged responses

Expected attack time and results

	Var.	Unit	Range	www. google. com attack		net attack	
				no	yes	no	yes
Port randomization				no	yes	no	yes
Num of resolver ports	Nport		1~ 2 ¹⁶	1	64,000	1	64,000
Num of IDs	Nqid		2 ¹⁶	65536	←	65536	←
Num of auth servers	Nns		1~13	4	←	13	←
RTT to auth servers	Tauth	sec	0.001 ~0.2	0.036~ 0.094	←	0.006~ 0.272	←
Loop interval	Tloop	sec		0.036	←	0.033	←
Response send rate	Rans	pps		66726	←	69158	←
First time trial Success possibility				0.00916	0.00000 014	0.00267	0.00000 004
expected attack time		sec		3.9	251,434	12.3	788,425
Attack time (result)		sec		46	×	8	×

If you detect poisoning

1. Flush poisoned RRSet
2. Resolve the poisoned RRSet
3. Compare the poisoned RRSet with other full resolvers
 - Your other full resolvers
 - Or, public DNS services (for example:Google)

 - Because cache flush may cause another poisoning

Measures

- Source port randomization
- Monitoring
 - Attacks may increase victim full resolver's load
 - Traffic
 - Number of packets of input and output are usually balanced on DNS
 - Victim full resolver receives many unmatched responses
- Many countermeasures are proposed/implemented
 - harden-referral-path on Unbound
 - nonce prefix on Google Public DNS
 - Use of TCP transport
 - DNS cookies
- See also
 - draft-fujiwara-dnsop-poisoning-measures-00

Detection by DNSSEC validation

- Cache poisoning is possible on BIND 9 and Unbound validators
- However, DNSSEC validation works well and they return validation error as “Server Failure”
- NSEC3 Opt-Out issue
 - Complex structured TLDs may have empty non-terminals
 - Empty non-terminals with NSEC3 are able to be validated because NSEC3 RRs say no NS
 - Some empty non-terminals do not have NSEC3 RRs
 - When there is no secure delegation for the name
 - One example: saitama.jp
 - They were not able to be validated
 - So, JPRS added TXT RRs to all empty non-terminals in JP zone
 - BIND 9’s dnssec-signzone generates NSEC3 RRs for **non-empty** non-terminals

Conclusion

- "Redirect the target domain's nameserver" cache poisoning attack is easy
 - Most of domain names are conquerable
 - However, it is easy to trace attackers because it needs authoritative DNS servers
- Measures
 - Source port randomization works well
 - DNSSEC validation is effective to detect
 - Detection and correspondence of attacks are important

details

Success probability of first time trial

- Success if ID, port, address are matched
- Probability of first time trial P is probability within $Tauth$

$$P = \frac{Tauth * Rans}{Nqid * Nport * Nns}$$

- However, $Tauth$ is not controllable
- Use $Tloop$ instead of $Tauth$
under $Tloop < Tauth$

$$P = \frac{Tloop * Rans}{Nqid * Nport * Nns}$$

Success probability under continuation attack

- Success probability of not succeeding by $n - 1$ st time, but succeeding n -th time

$$P_n = (1 - P)^{n-1} * P \quad (1)$$

- Success probability of succeeding by n -th time Q_n is the sum of P_n

$$Q_n = \sum_{i=1}^n (1 - P)^{i-1} P \quad (2)$$

$$Q_n = 1 - (1 - P)^n \quad (3)$$

- n is brought close to infinity

$$Q_n \rightarrow 1$$

Expected attack number of tries

- Success probability of not succeeding by $n - 1$ st time, but succeeding n -th time

$$P_n = (1 - P)^{n-1} * P \quad (1)$$

- Expected attack number of tries by n -th time E_n is a sum of the multiply of value i and P_n

$$E_n = \sum_{i=1}^n i(1 - P)^{i-1} P \quad (2)$$

- Then

$$E_n = \frac{1}{P} - \frac{(1+nP)(1-P)^n}{P} \quad (3)$$

- n is brought close to infinity

$$E_n \rightarrow \frac{1}{P}$$

Expected attack time

$$\begin{aligned} T &= T_{loop} * E \\ &= \frac{T_{loop}}{P} \\ &= \frac{N_{qid} * N_{port} * N_{ns}}{R_{ans}} \end{aligned}$$

Termination of attack

- Prepare “*.\$Trigger IN A 192.2.0.1” on forged authoritative DNS server
- If the NS injection succeed, a response of a trigger query returns rcode 0, “(random).\$D IN A” RR