CSE 127: Computer Security

Network Security

Kirill Levchenko

November 30, 2017

DNS

- **Domain Name System (DNS):** maps host names to addresses
  - Large distributed database
- Critical Internet infrastructure
  - Control of DNS allows attacker to impersonate any site (in the absence of end-to-end host authentication using SSH or TLS)
- Target of attacks

DNS

- **DNS Record:** Unit of information in DNS
  - **Type:** type of data it contains
  - **TTL:** time to live
- **A record:** IP address for a host name
- **NS record:** Name server to contact for a domain
- **MX record:** SMTP (mail) server for domain

Authorities

- A DNS server has a set of records for which it is the authoritative source

```
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 30439
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 3, ADDITIONAL: 6

;; QUESTION SECTION:
;bob.ucsd.edu.          IN  A

;; ANSWER SECTION:
bob.ucsd.edu.       3600    IN  A   132.239.80.176

;; AUTHORITY SECTION:
ucsd.edu.       3600    IN  NS  ns0.ucsd.edu.
ucsd.edu.       3600    IN  NS  ns1.ucsd.edu.
ucsd.edu.       3600    IN  NS  ns2.ucsd.edu.
```

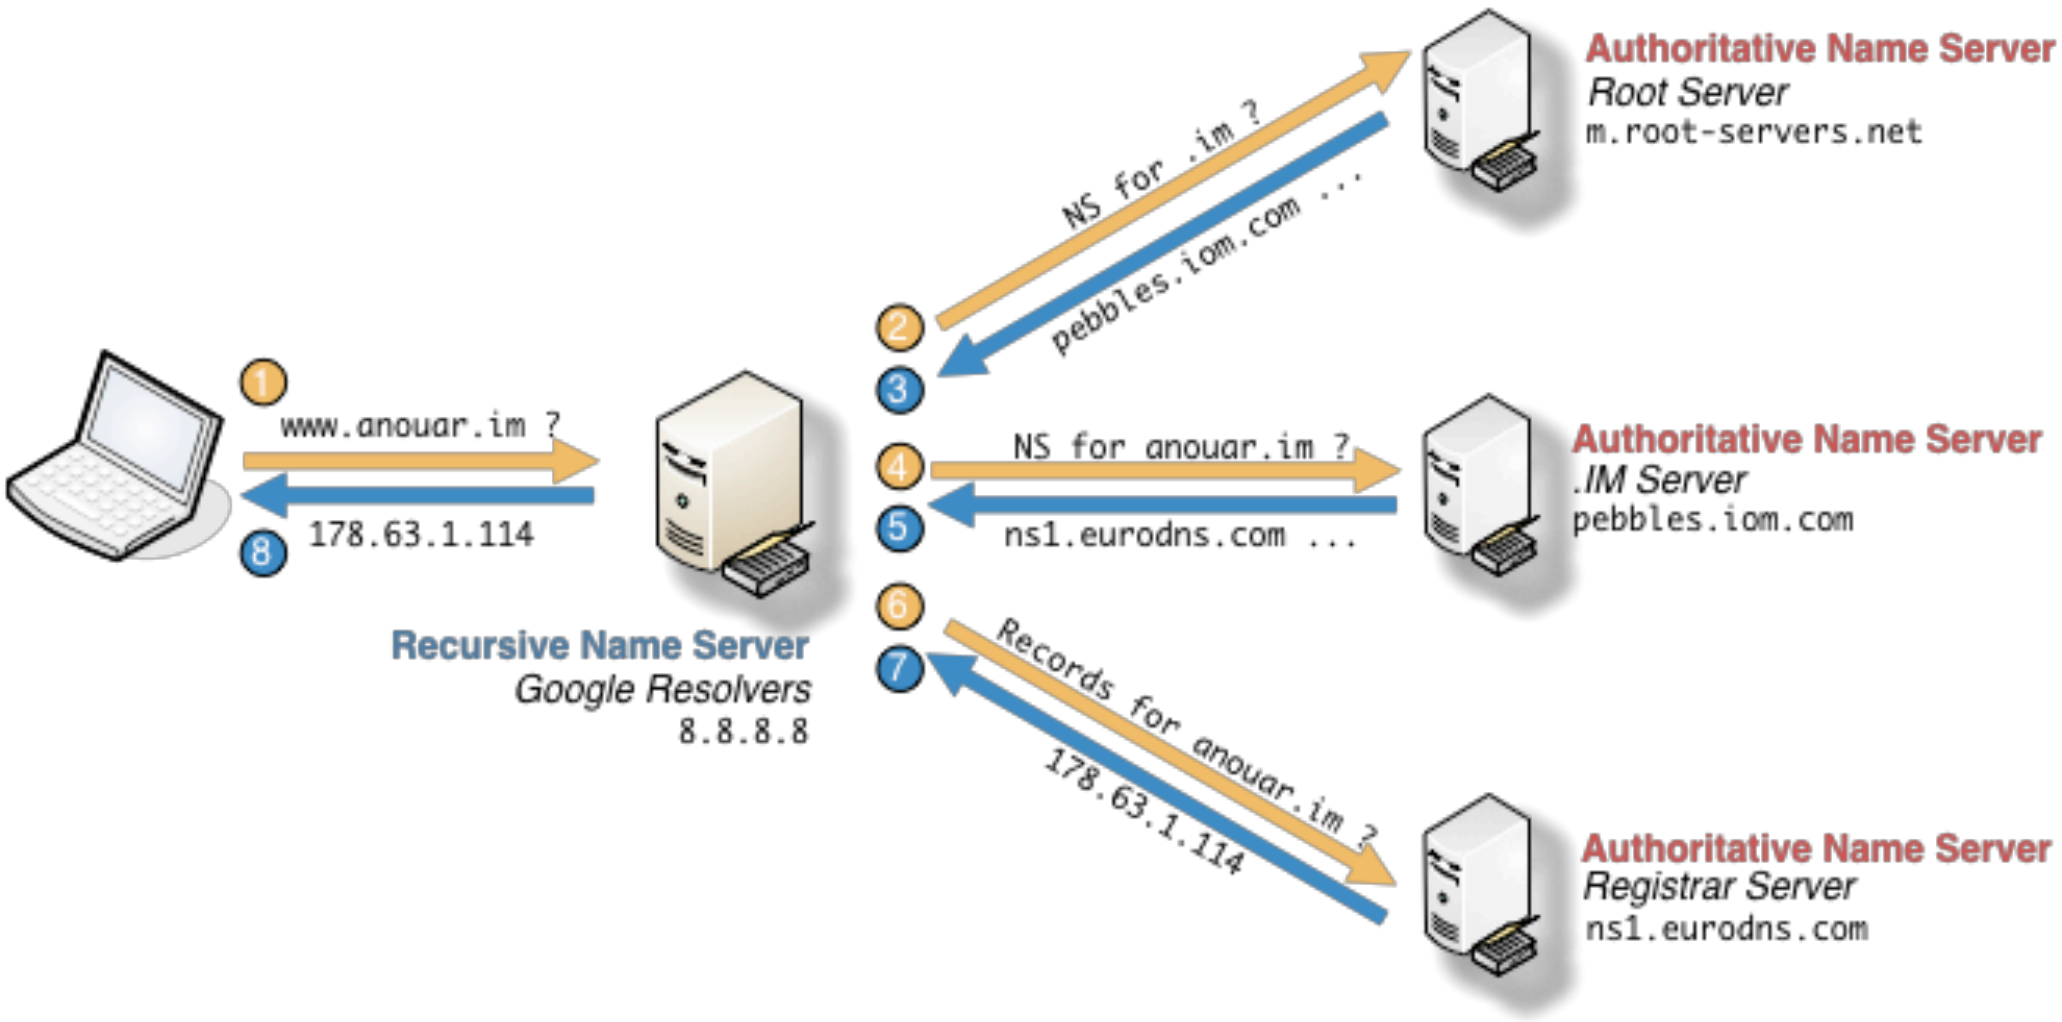

source: http://anouar.adlani.com/2011/12/useful-dig-command-to-troubleshot-your-domains.html

DNS Roles

- **Authoritative server:** provides authoritative information for a set of domains
  - Does not handle queries about other domains
- **Recursive resolver:** provides recursive resolution of a domain to return requested record to client
  - Handles queries about all domains
- Same protocol for both types of servers
  - Distinction is in intended purpose only

Security of DNS

- Basic DNS uses UDP without any authentication
- Man-in-the-middle attacks are possible
  - Forging response to observed query is trivial
- Off-path attacks require guessing two parameters
  - Query ID from response (16 bits)
  - Source port (approx. 15 bits)

DNS Packet

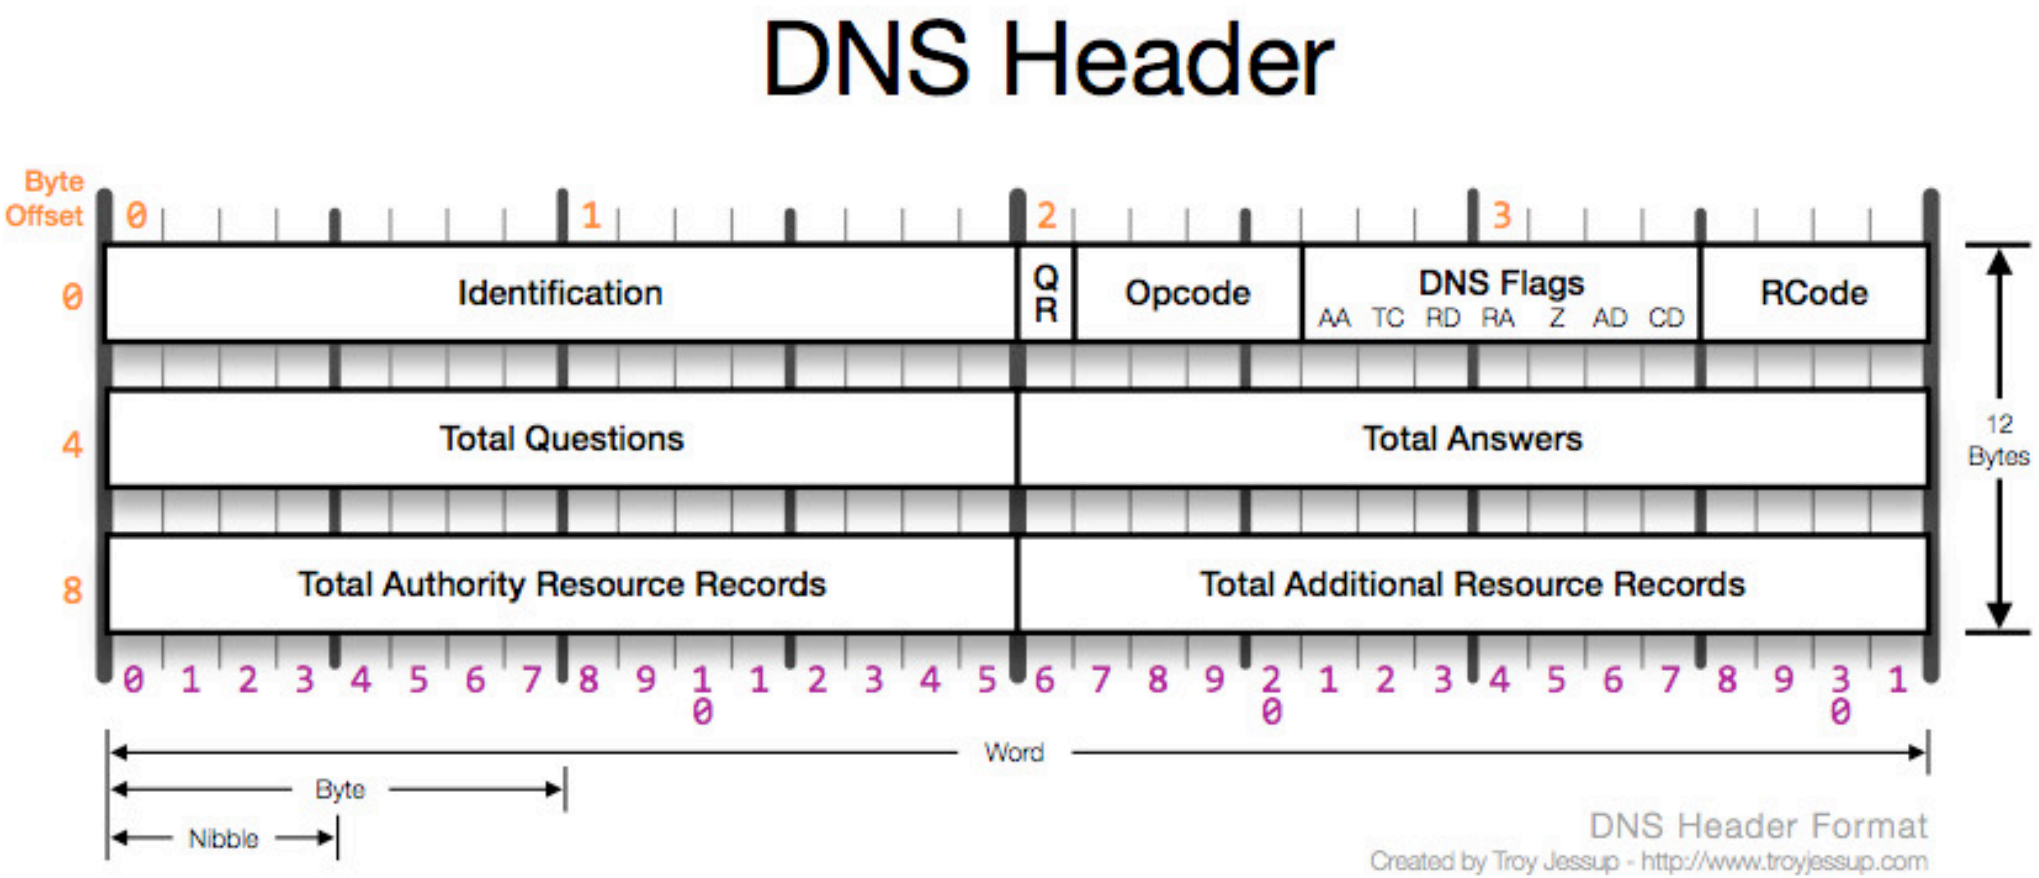
DNS Header
Byte Offset
0
1
2
3
0
Identification
QR
Opcode
DNS Flags
AA TC RD RA Z AD CD
RCode
4
Total Questions
Total Answers
8
Total Authority Resource Records
Total Additional Resource Records
12 Bytes
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31
Word
Byte
Nibble
DNS Header Format
Created by Troy Jessup - http://www.troyjessup.com

Caching

- Recursive resolvers cache records to avoid repeating recursive resolution process for each query
- Lifetime of record determined by record TTL
  - Could also be evicted from cache because of limited memory
- Injecting spoofed records into a resolver's cache is called DNS cache poisoning
  - No protocol-defined way for client to refresh cached record

Forging DNS Replies

- For performance reasons, some DNS resolvers (e.g. BIND) re-used the same socket for all queries
  - **If source port is same:** can be determined by attacker by having recursive resolver query attacker-controlled authoritative server
- 16-bit query ID now only thing need to guess
  - Non-random query ID generators may make it easier

Naive Cache Poisoning

1. Query resolver for target domain `www.ucsd.edu`
   - If name already in cache, nothing you can do, wait for it to expire
   - Some resolvers may only serve particular hosts, need insider
2. Recursive resolver issues query to `.edu` authoritative server to get authoritative server for `ucsd.edu`
3. Generate a flood of forged replies with different query IDs appearing to be from `.edu` authoritative server
   - Some chance of guessing query ID and winning packet race

Naive Cache Poisoning

- Requires guessing query ID and source port
- Throttled by TTL (discussed later)

Additional Record Injection

- DNS query results include Additional Records section
  - Provide records for anticipated next resolution step
- Early servers accepted and cached all additional records
  - What is wrong with this?

Early Attack Strategy

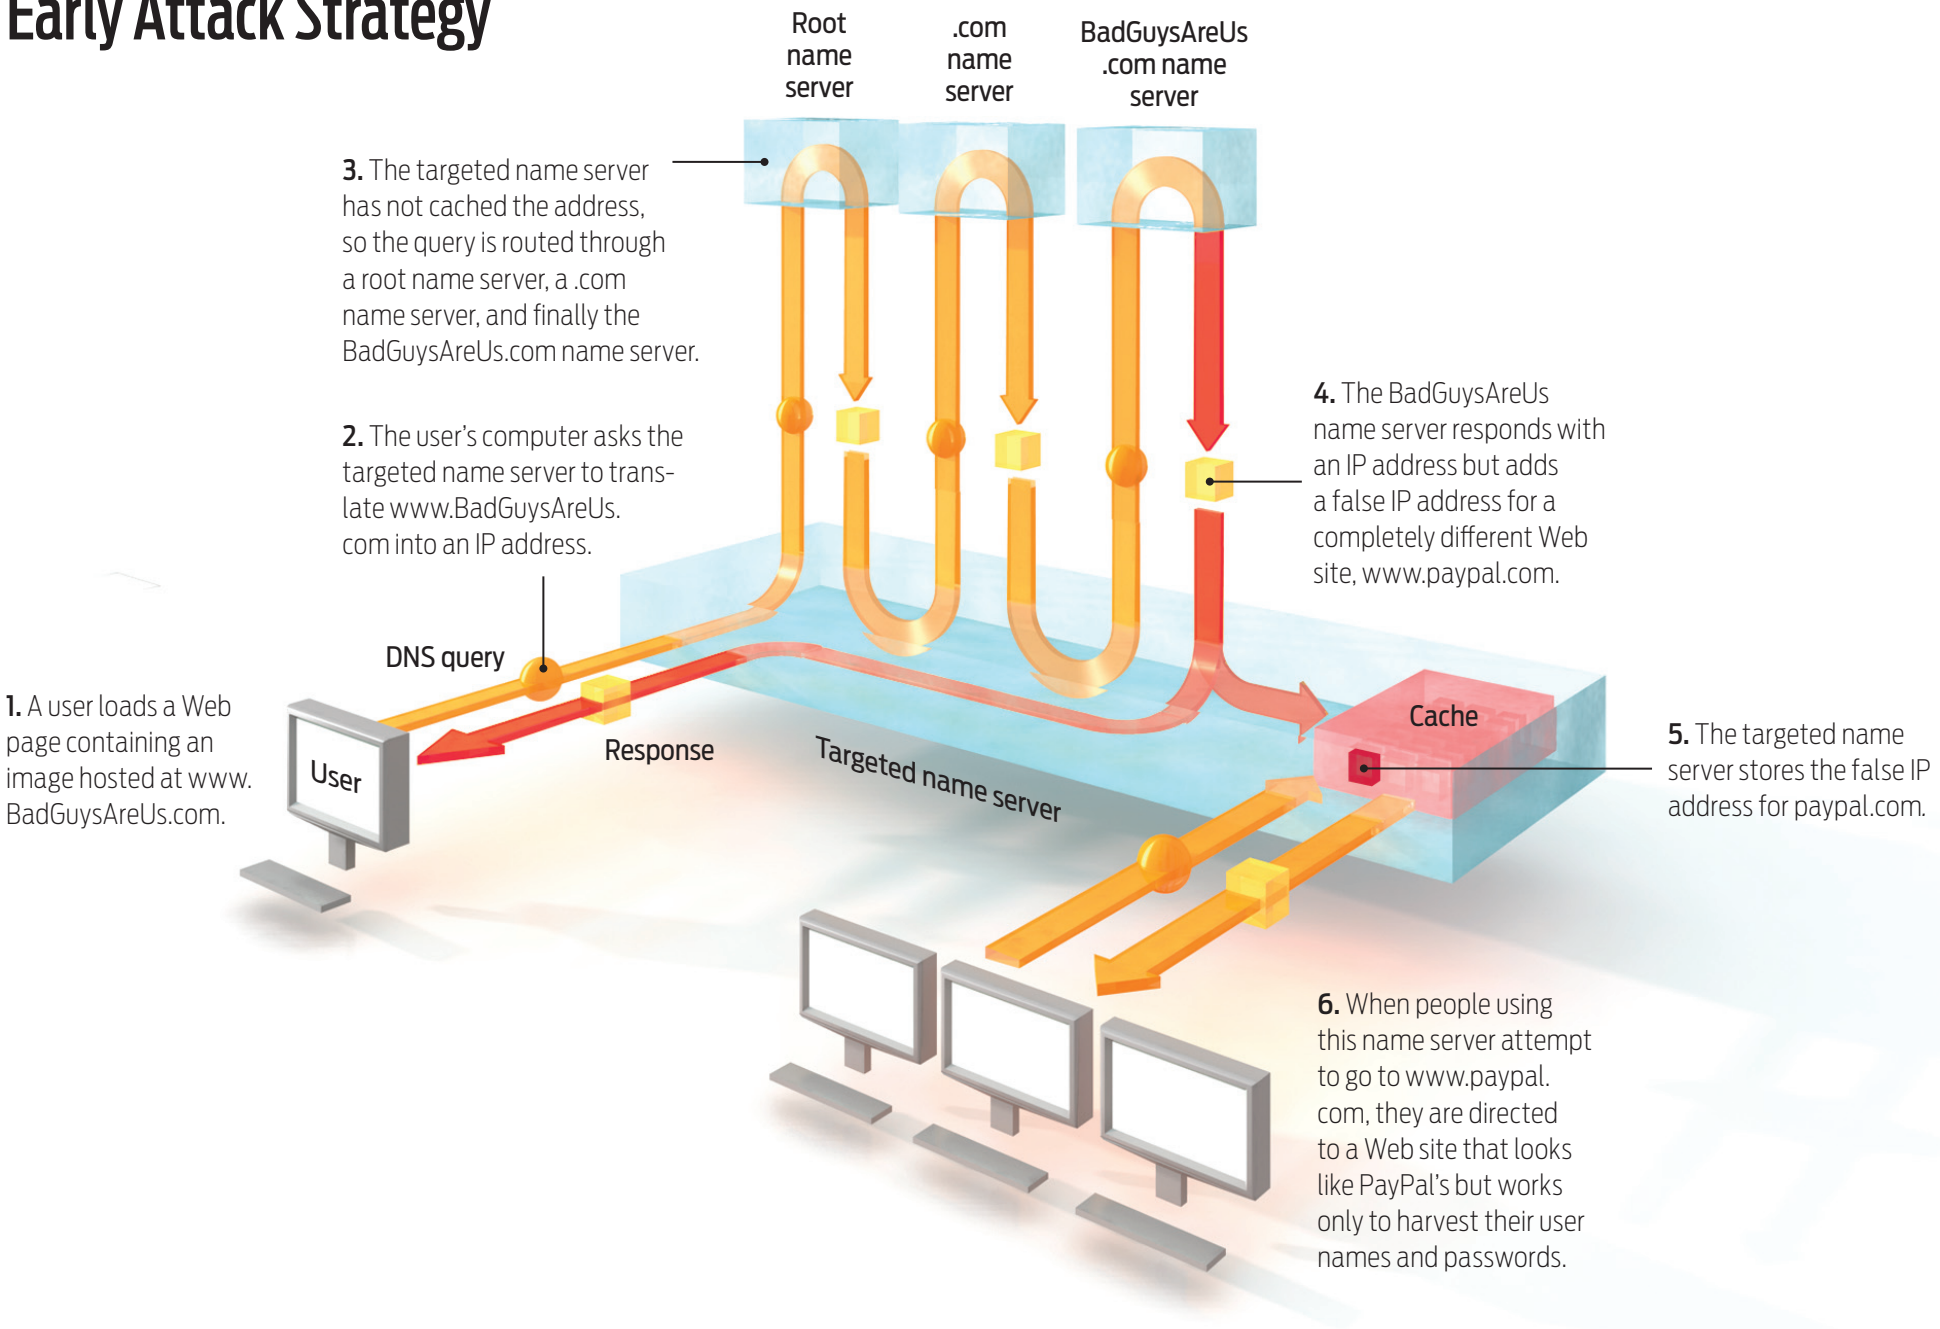

(source: *IEEE Spectrum*, Oct 2008)

Additional Record Injection

- DNS query results include Additional Records section
  - Provide records for anticipated next resolution step
- Early servers accepted and cached all additional records
  - What is wrong with this?
- Can we just stop using additional section?
  - Only accept answers from authoritative servers?

DNS Glue Records

- Can we just stop using additional section?
  - Only accept answers from authoritative servers?
- **Glue records:** non-authoritative records necessary to contact next hop in resolution chain
  - Necessary given current design of DNS

```
; <<>> DiG 9.6-ESV-R4-P3 <<>> @192.5.6.30 ucsd.edu
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 12781
;; flags: qr rd; QUERY: 1, ANSWER: 0, AUTHORITY: 3, ADDITIONAL: 4
;; WARNING: recursion requested but not available

;; QUESTION SECTION:
;ucsd.edu.                      IN      A

;; AUTHORITY SECTION:
ucsd.edu.               172800  IN      NS      ns1.ucsd.edu.
ucsd.edu.               172800  IN      NS      ns2.ucsd.edu.
ucsd.edu.               172800  IN      NS      ns0.ucsd.edu.

;; ADDITIONAL SECTION:
ns1.ucsd.edu.           172800  IN      A       128.54.16.2
ns2.ucsd.edu.           172800  IN      A       132.239.1.52
ns0.ucsd.edu.           172800  IN      A       132.239.1.51
ns0.ucsd.edu.           172800  IN      AAAA    2607:f720:100:100::231
```

edu authority

Names of ucsd.edu authoritative servers

Glue records for authoritative servers

What should be the policy for accepting additional records?

Bailiwick Rules

- **Bailiwick rule:** defines what response records a recursive resolver will accept
- **Bailiwick:** set of domains about which a server is has direct or indirect authority to speak
  - Bailiwick determined by the initiator of query
- Answer should be relevant (in response to request)
- Answer should be in bailiwick

Bailiwick Checking Rule from BIND

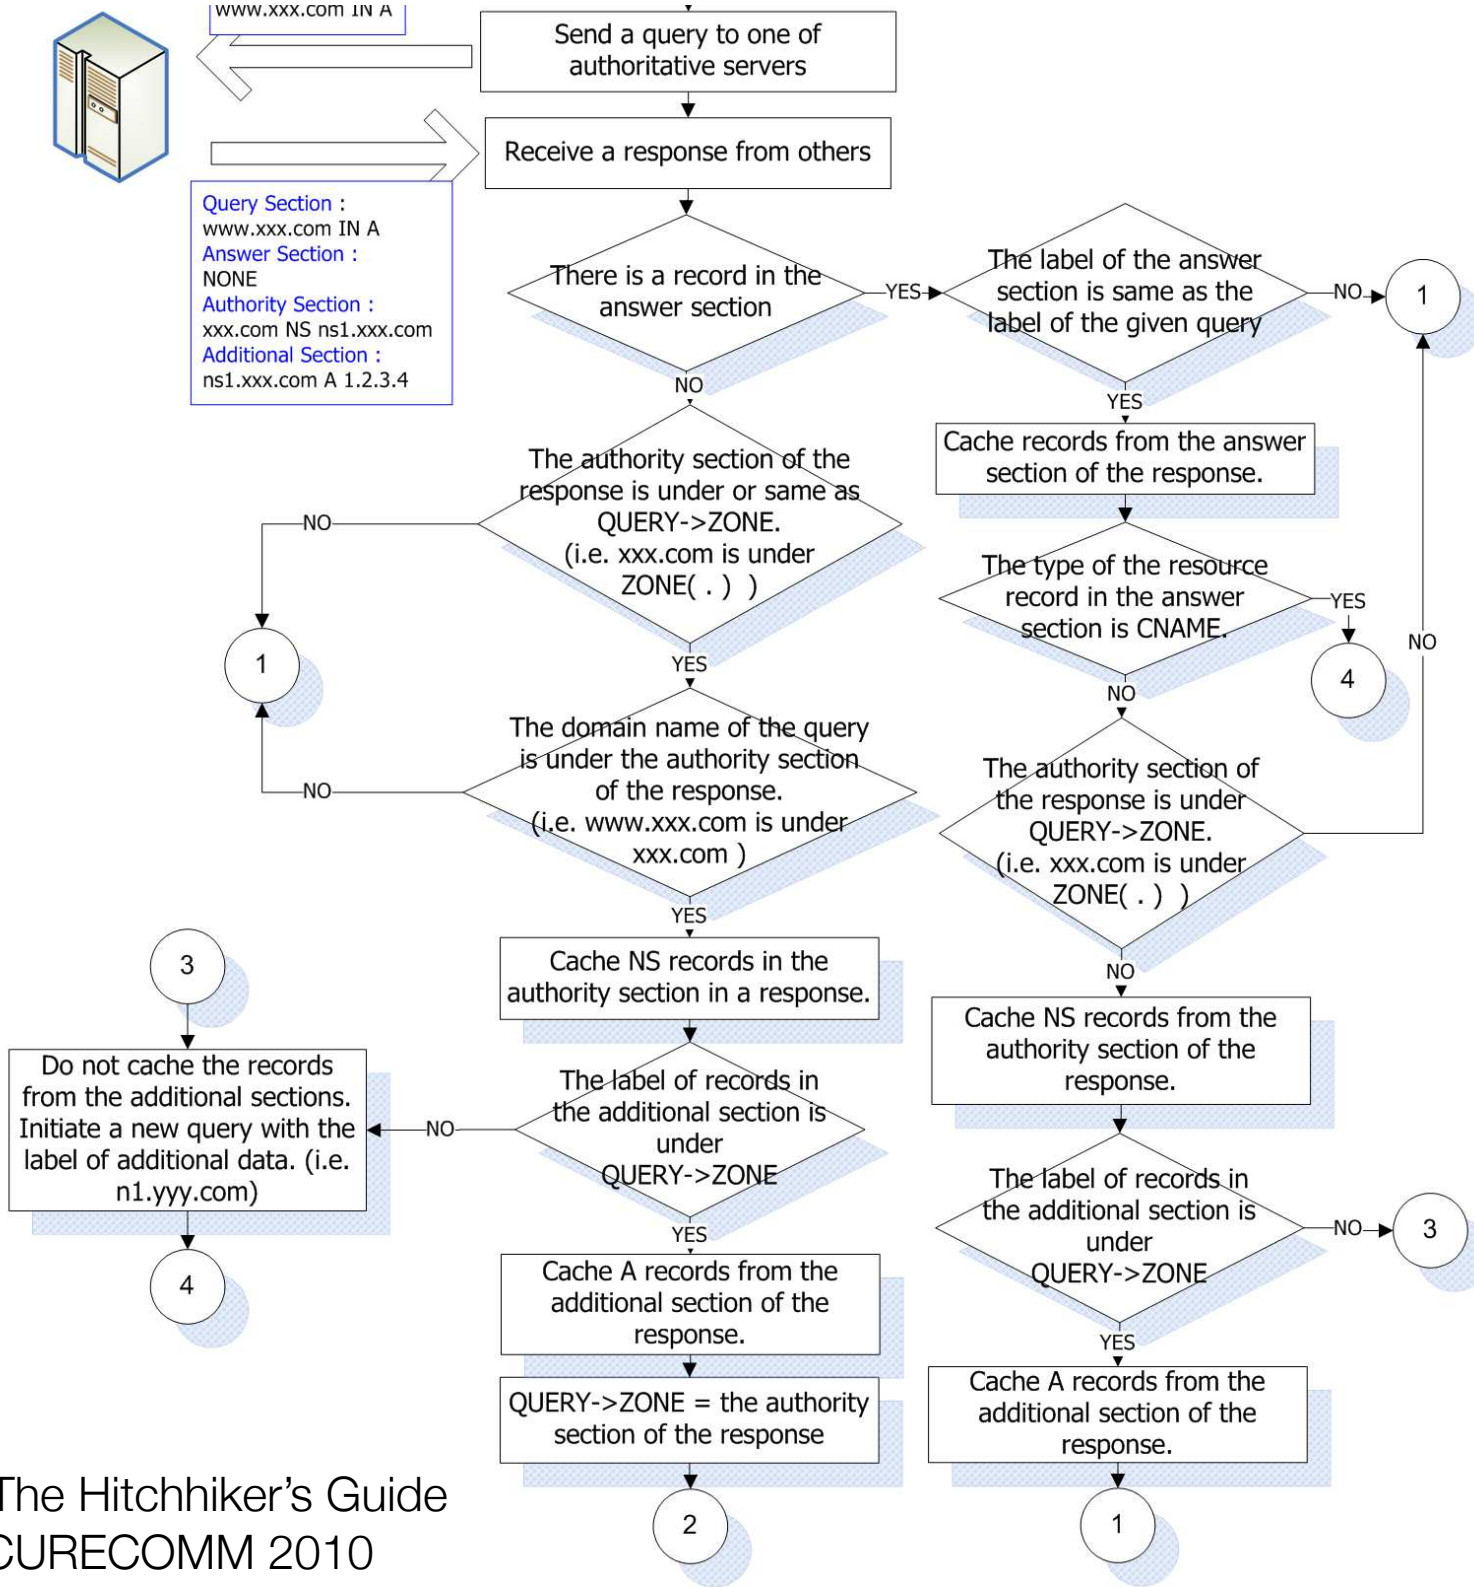

source: Son and Shmatikov, “The Hitchhiker’s Guide to DNS Cache Poisoning” SECURECOMM 2010

BIND Bailiwick Rule (roughly)

- Authorities must be for queried domain
  - `ns0.csd.edu` accepted as authority for `ucsd.edu` only when initiating query was for subdomain of `ucsd.edu`
- Additional records must be *in bailiwick* for query
  - A record for `ns0.ucsd.edu` accepted because `edu` server has indirect authority over `ns0.ucsd.edu`

DNS should have been designed with addresses, not names, in NS records and MX records. The “additional section” of DNS responses should have been eliminated. RFC 1035 observes correctly that NS indirection and MX indirection “insure [sic] consistency” of addresses; however, this indirection should have been handled by the server, not the client.

— *Daniel J. Bernstein*

source: http://cr.yp.to/djbdns/notes.html

Kaminsky's Attack

- Naive attacks on DNS are throttled by TTL
- Each result of query cached for TTL duration
  - Future attack queries answered from cache
- **Attack:** inject forged *additional records* when non-existent domains are queried
- Bypasses TTL throttling

Kaminsky DNS Attack

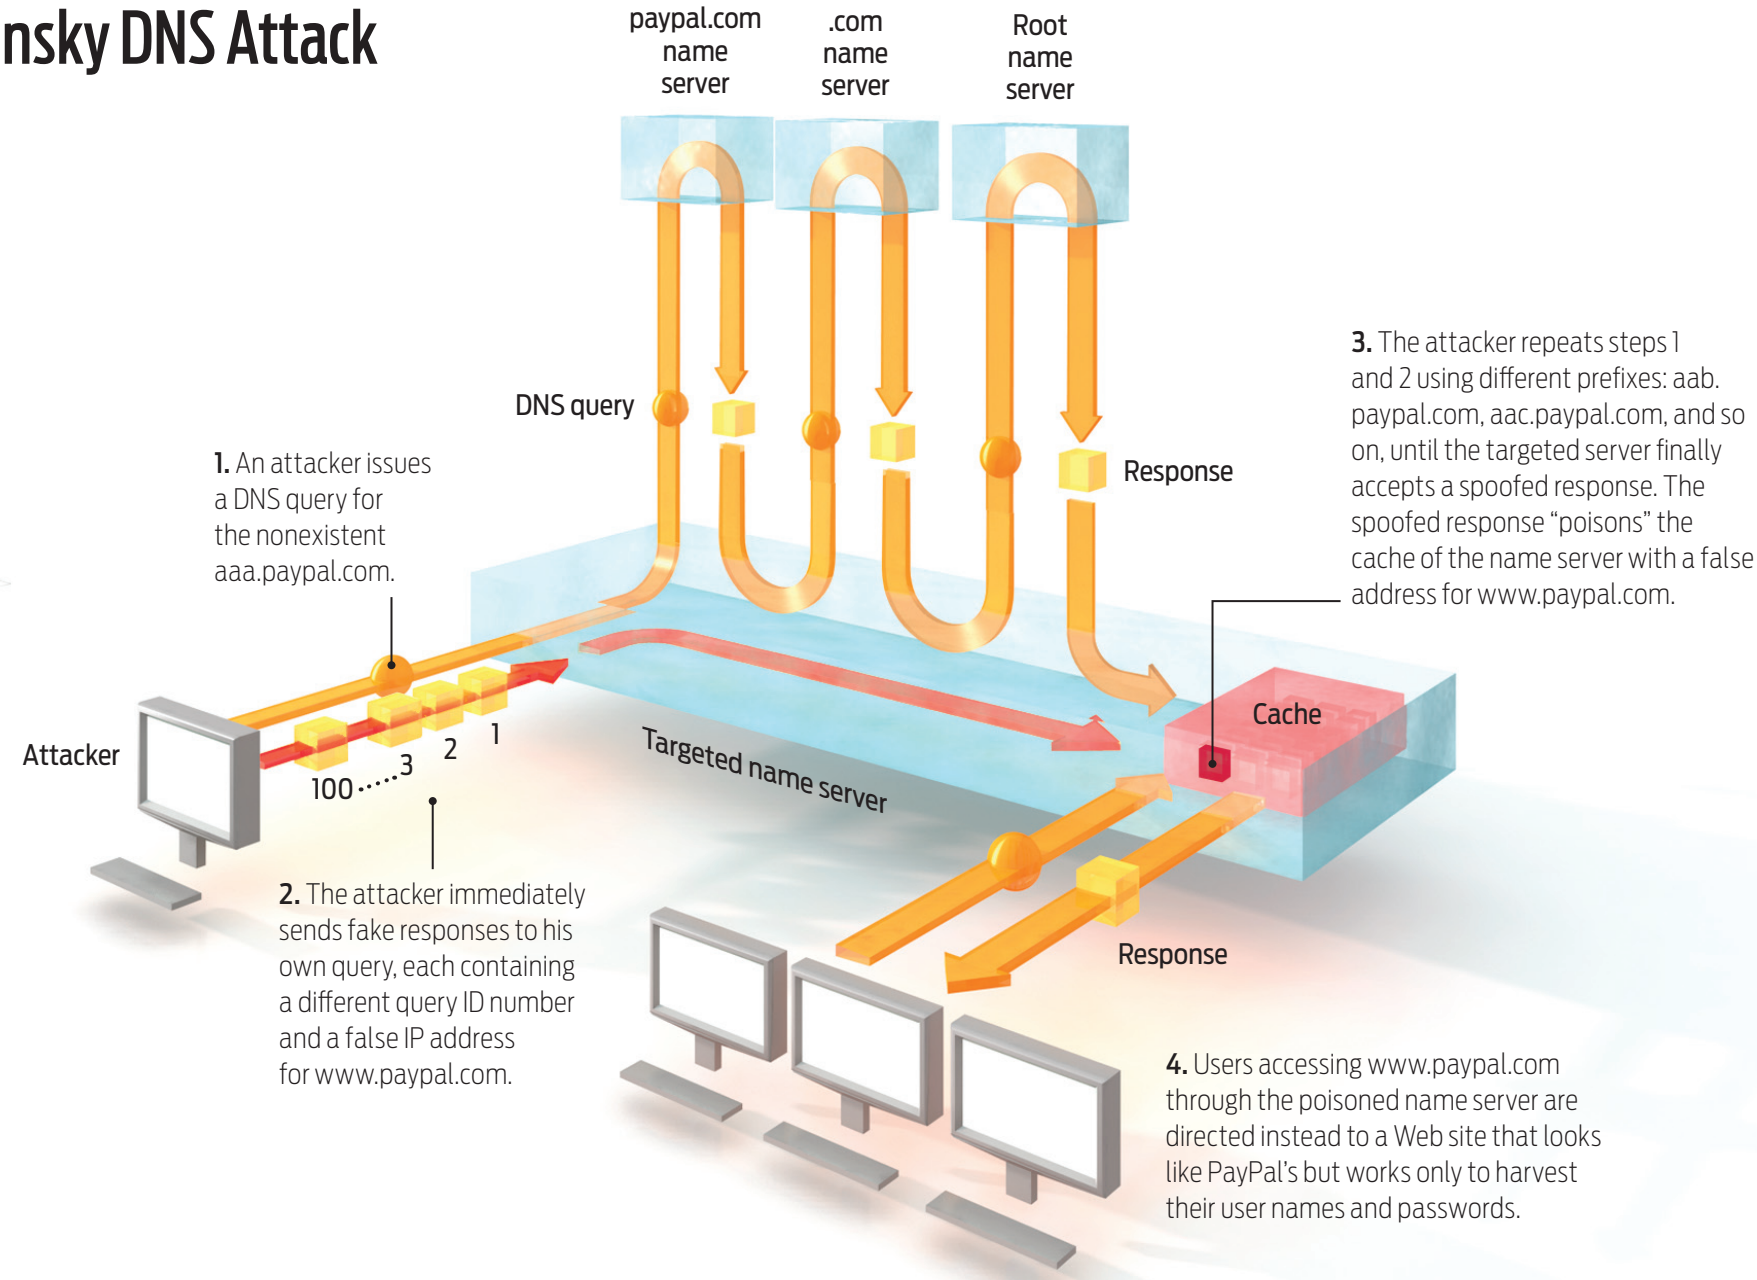

BRYAN CHRISTIE DESIGN

0x20 Encoding

- **Goal:** Increase entropy in DNS
- **Idea:** Vary capitalization of queried domain
  - DNS is case insensitive, so this is okay
- DNS server's response must use same capitalization
  - All known servers happen to do this
- One additional bit of entropy per letter
  - Attack must guess capitalization also

DNSSEC

- Cryptographically sign DNS records
- Chain of trust form root to subdomains, etc.
- DNSKEY record: a public key
- RRSIG record: digital signature of data
- NSEC record: negative response
- DS record: delegation record