DNS Poisoning: Developments, Attacks and Research Directions

Suggestions for the Idle and Curious Researcher

David Dagon[1]

[1]Georgia Institute of Technology
Atlanta, Georgia

USENIX 08 DNS Panel – July 31 2008

Objectives: Identify Research Opportunities

- More than ever, the research community is needed
- Recent DNS exploits present a broad threat, and opportunities
- These notes present an overview of new DNS poisoning techniques
- Open questions are presented **in red**
  - The panel discussion may identify interesting research topics
- The research community is urged to respond to this problem

Basic Poisoning Model

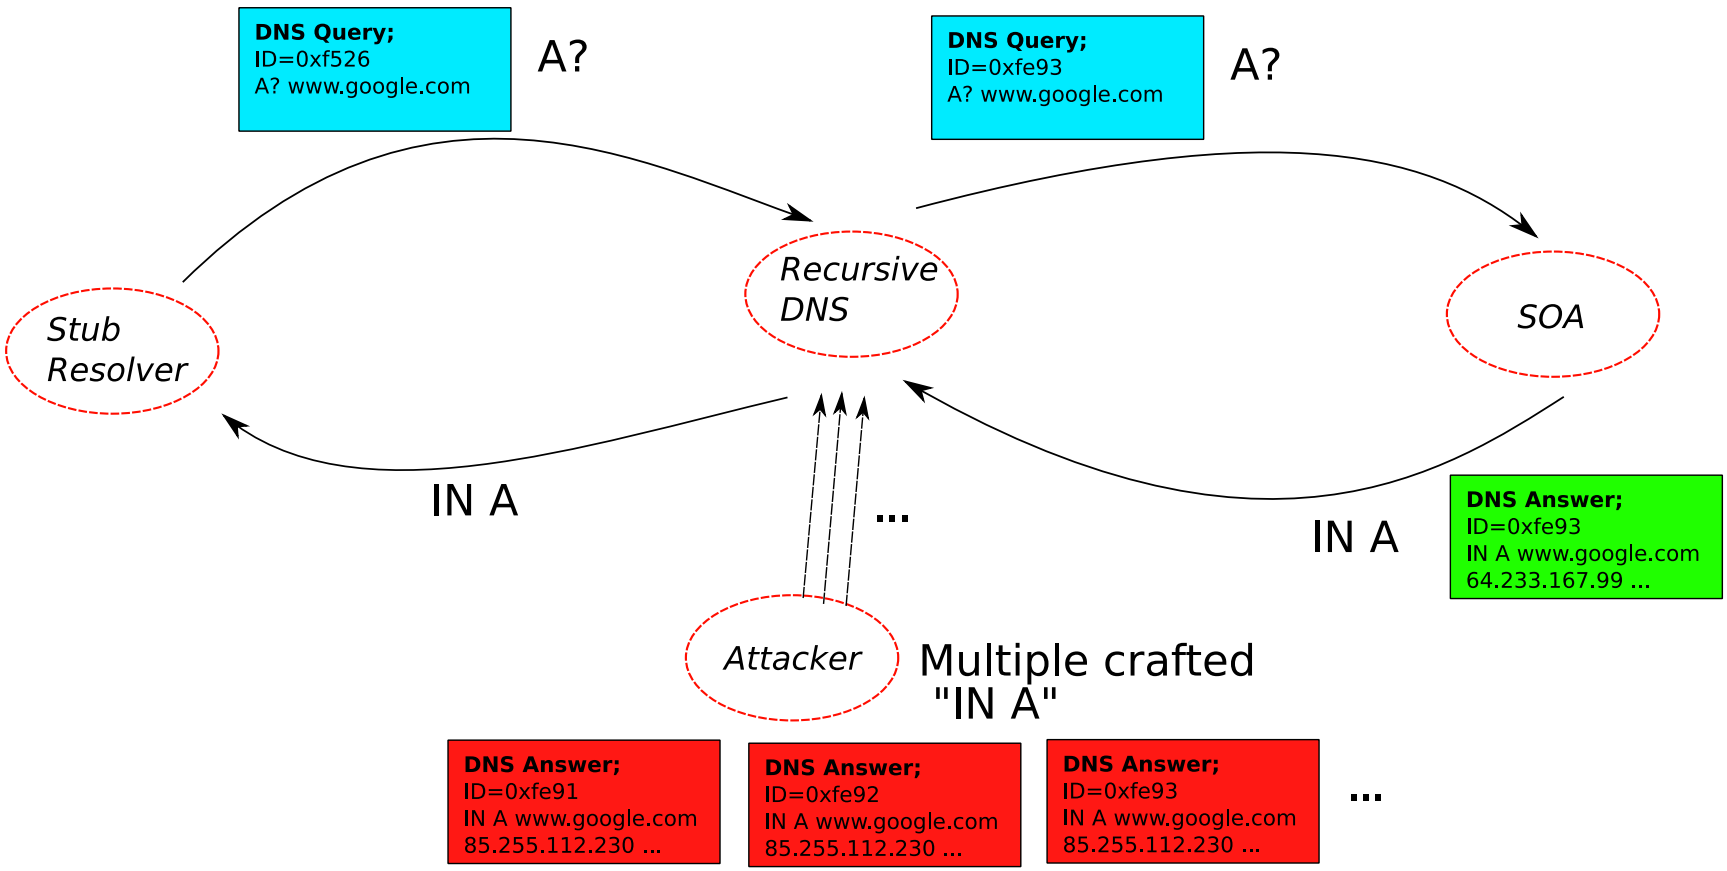

Poisoning Overview: Time-to-Success

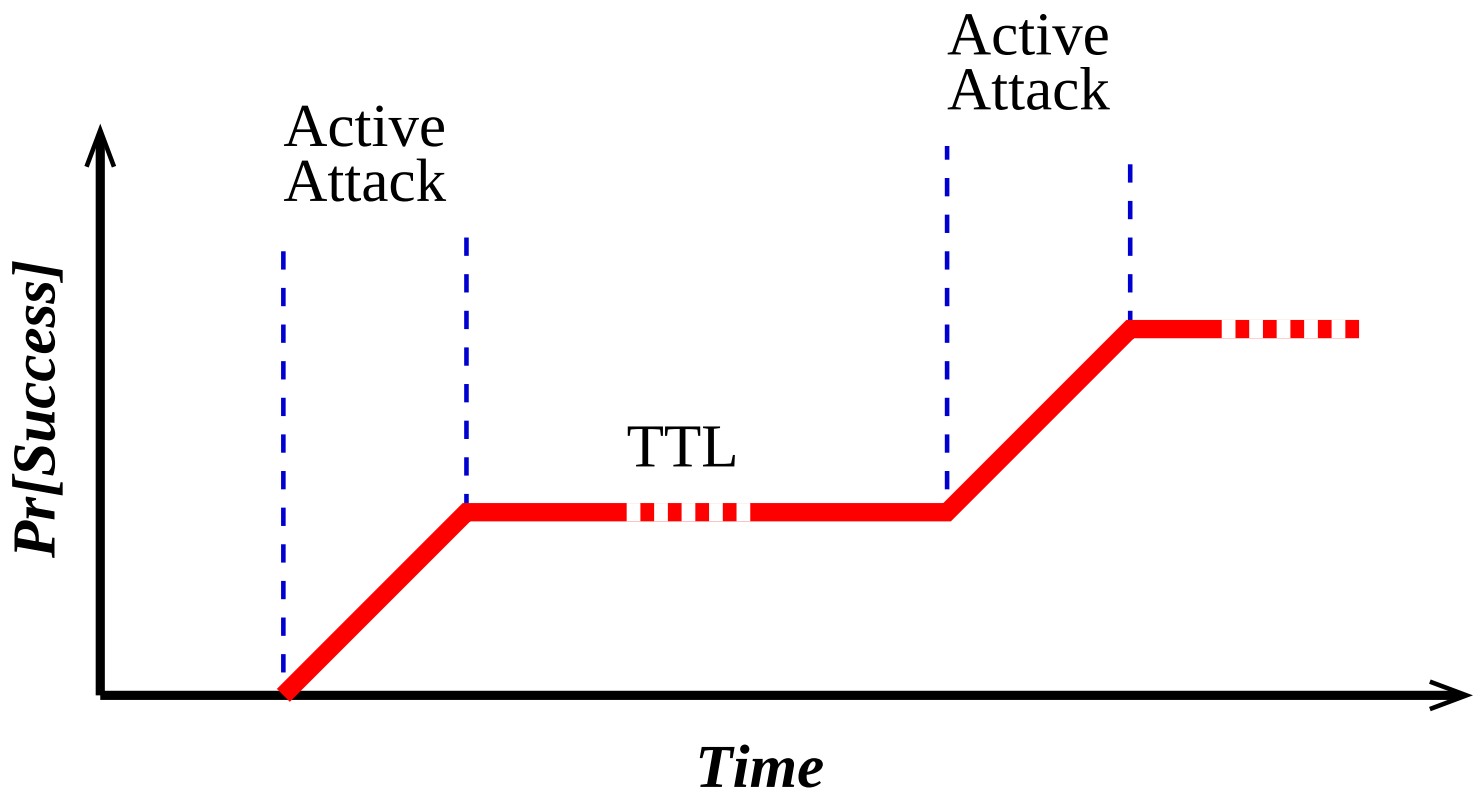

Poisoning Overview: Time-to-Success

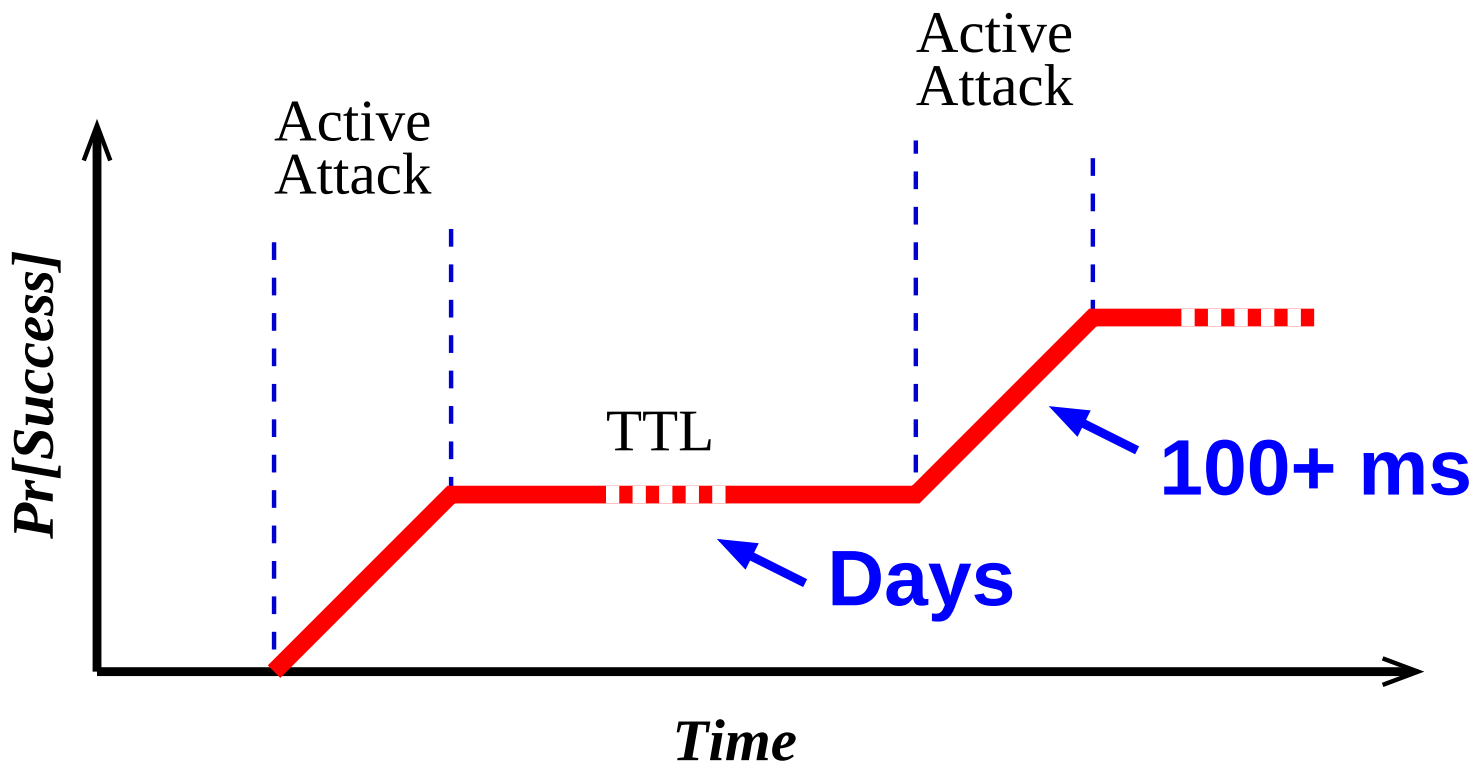

Kaminsky-Class Poisoning

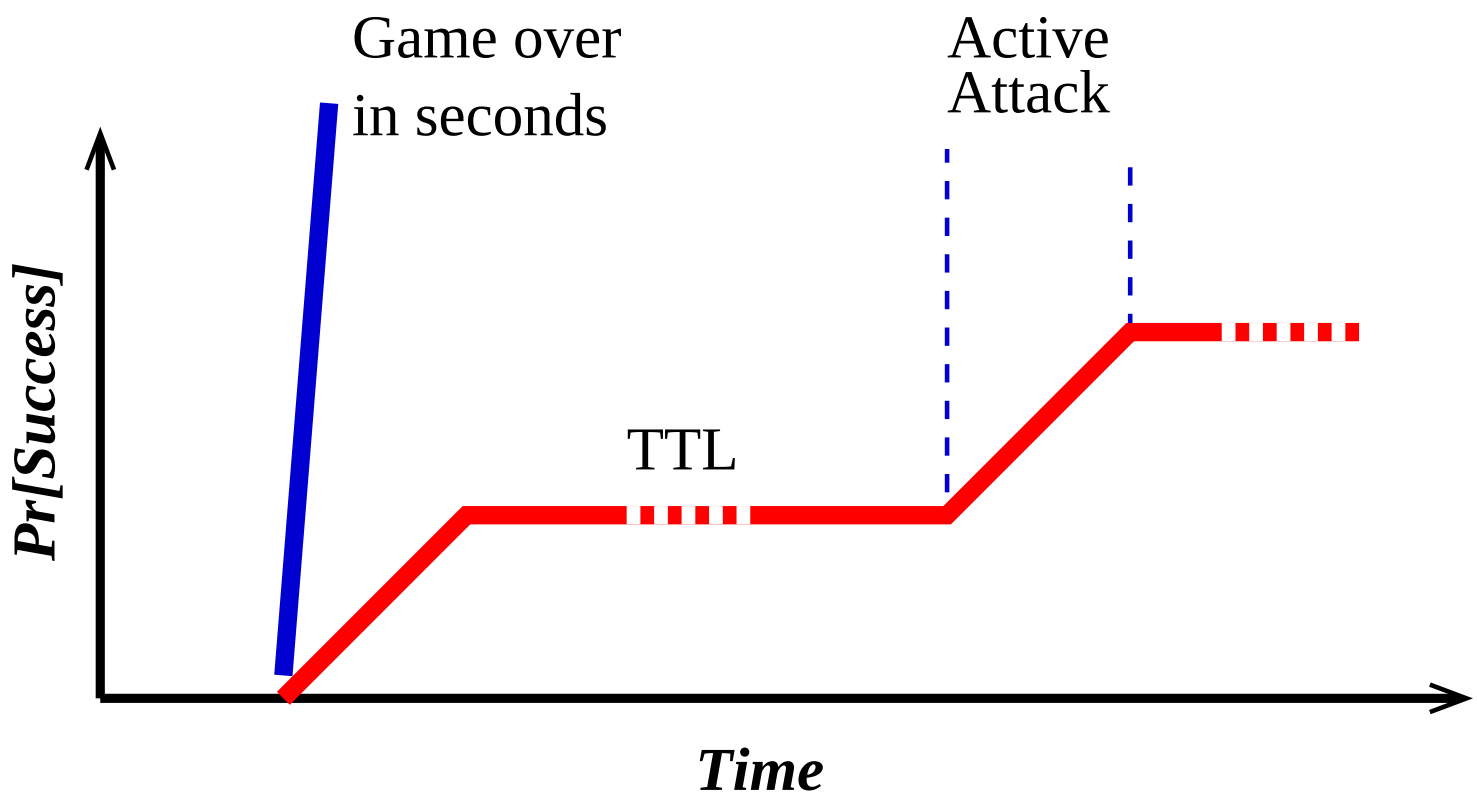

Kaminsky-Class Poisoning

- Can start anytime; now waiting for old good cached entries to expire
- No “wait penalty” for poisoning failure: TTL no longer a factor
- Generally, the attack is only bandwidth limited
- Deterministic march to cache manipulation
- Full consideration of attack dimensions at Kaminsky’s upcoming BlackHat talk

Dramatis Personae

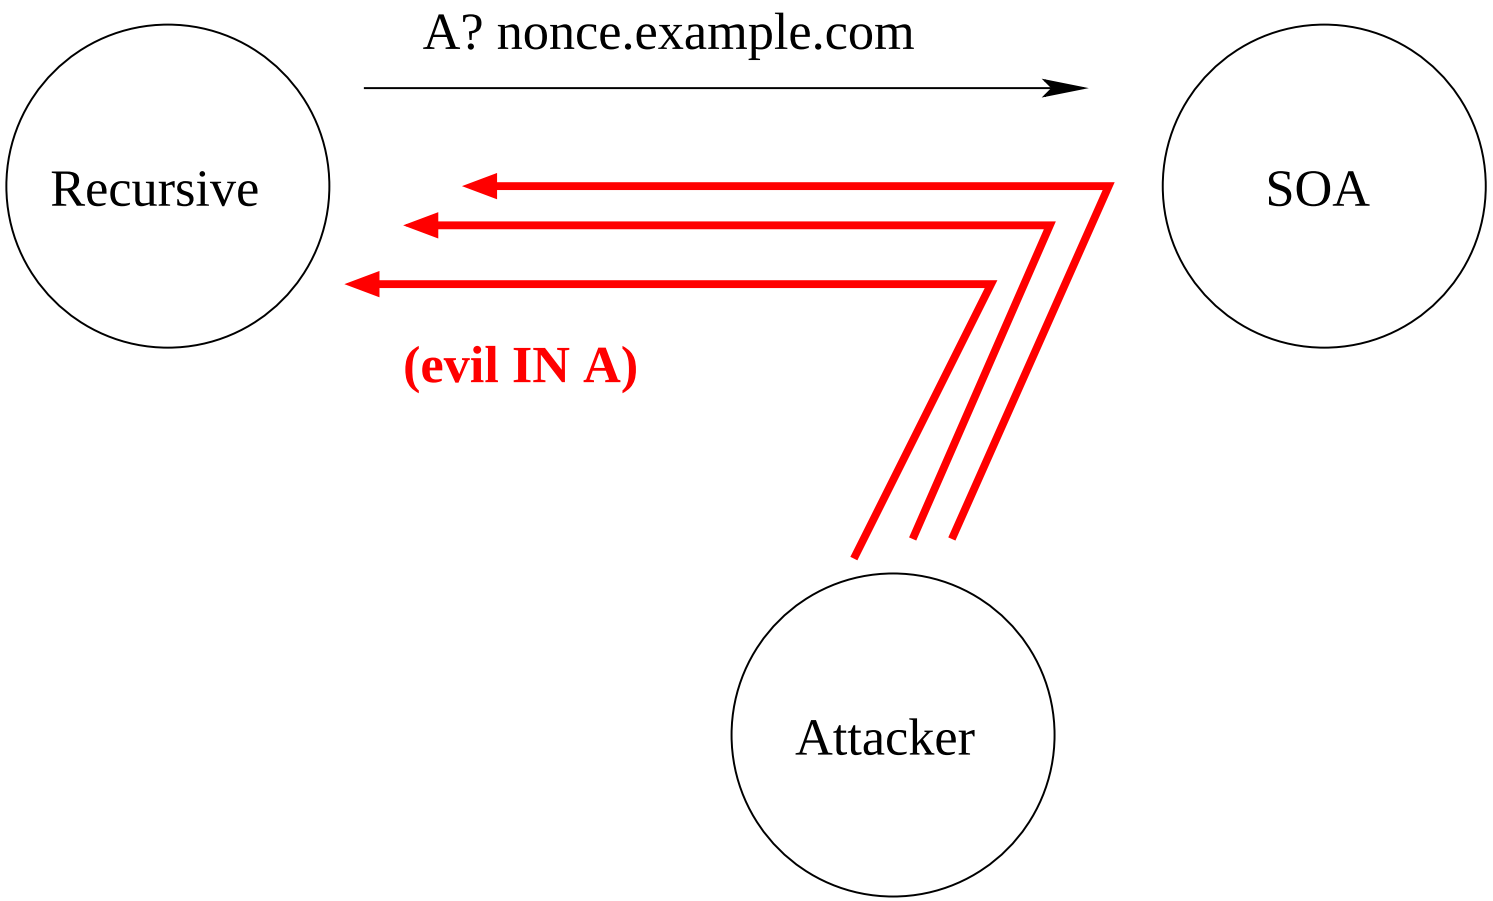

Dramatis Personae

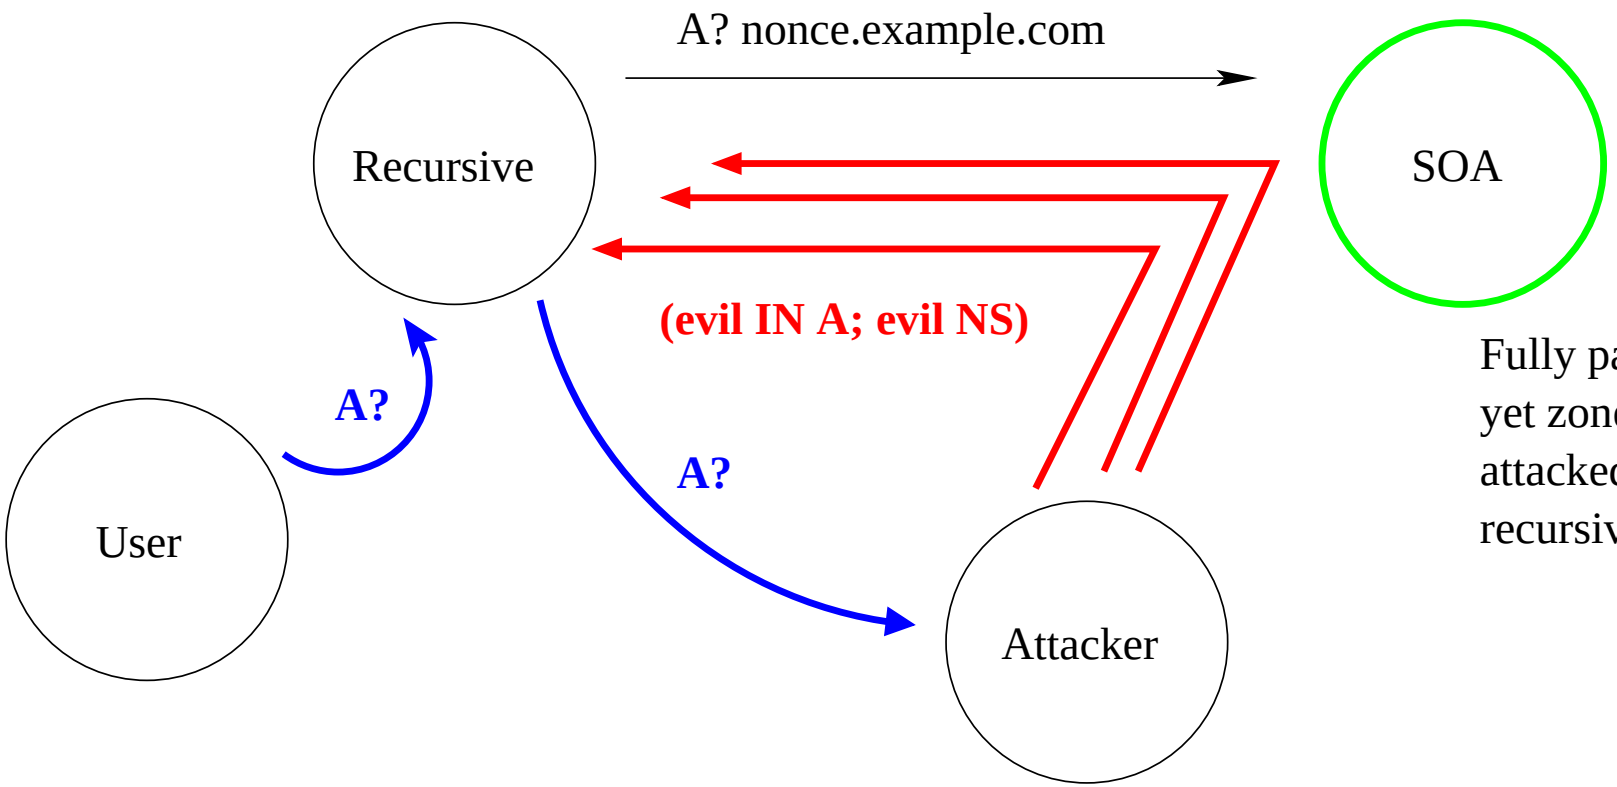

Dramatis Personae: Implications

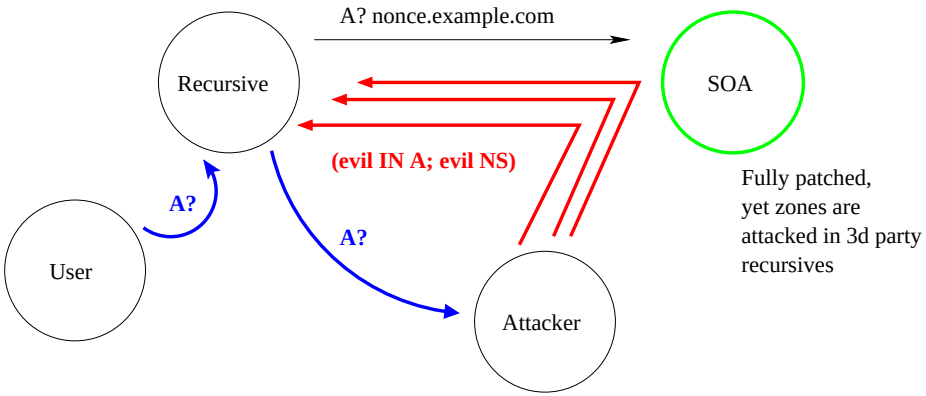

- Note the diverse threat:
  - Recursive's risk: DNS reputation; integrity of service
  - Authority's risk: Visitors at risk, domain brand
  - User's risk: all DNS-aware applications

Attack Scenarios

- Why, then, is this an important exploit?
- There are countless trivial exploits built on top of this single vulnerability disclosure.
- One example now being seen:
  - Message interception
- High risk/high yield: HTTP to HTTPS sites
- Numerous other scenarios... see Dan Kaminsky's talk for more.

Attack Scenario: Mail Intercept

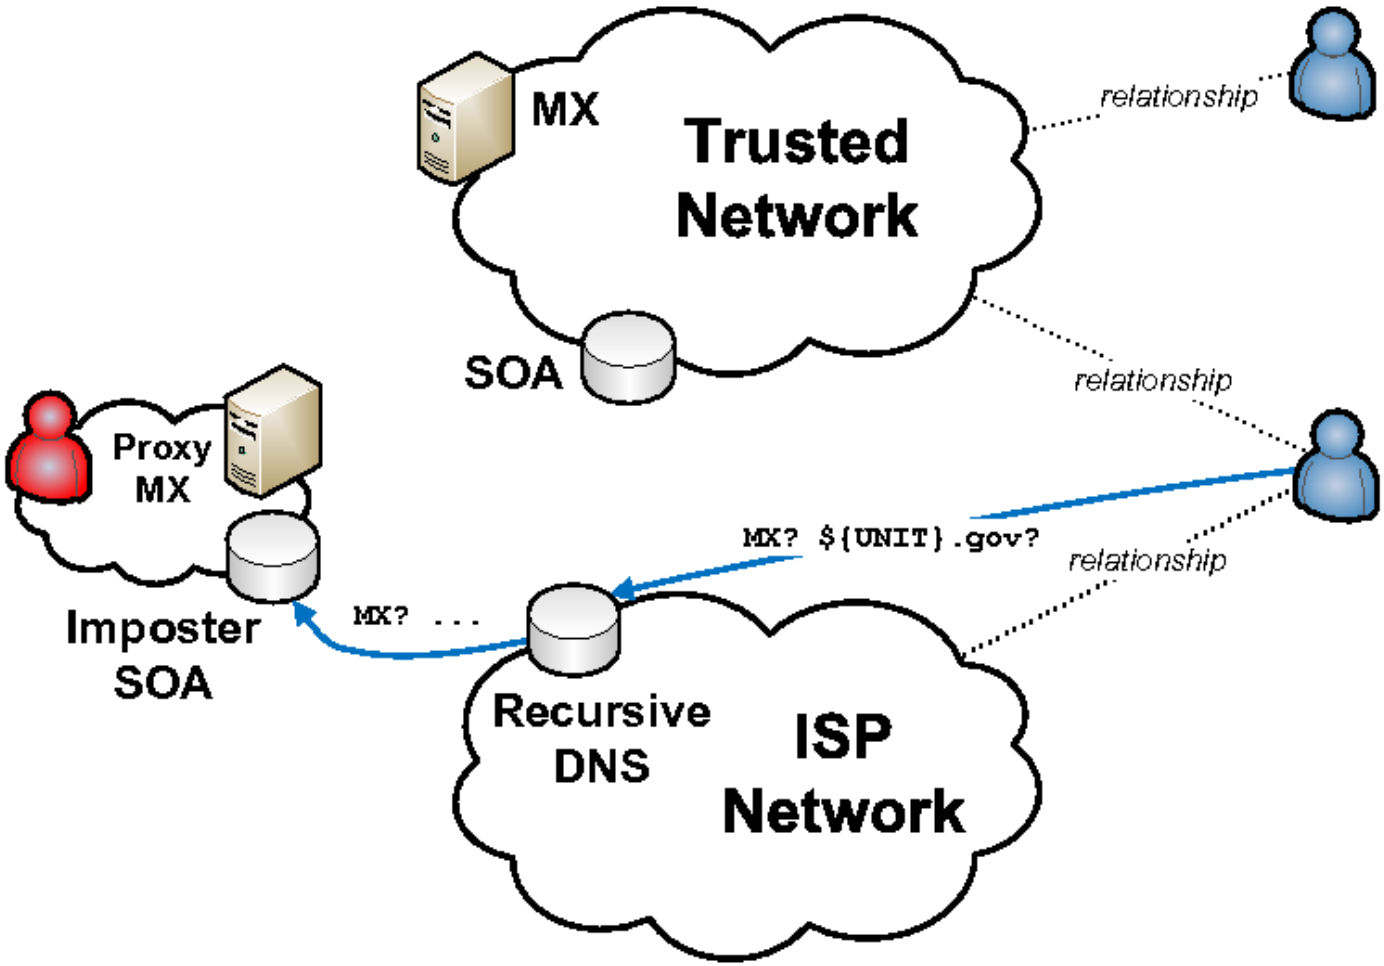

Attack Scenario: Mail Intercept

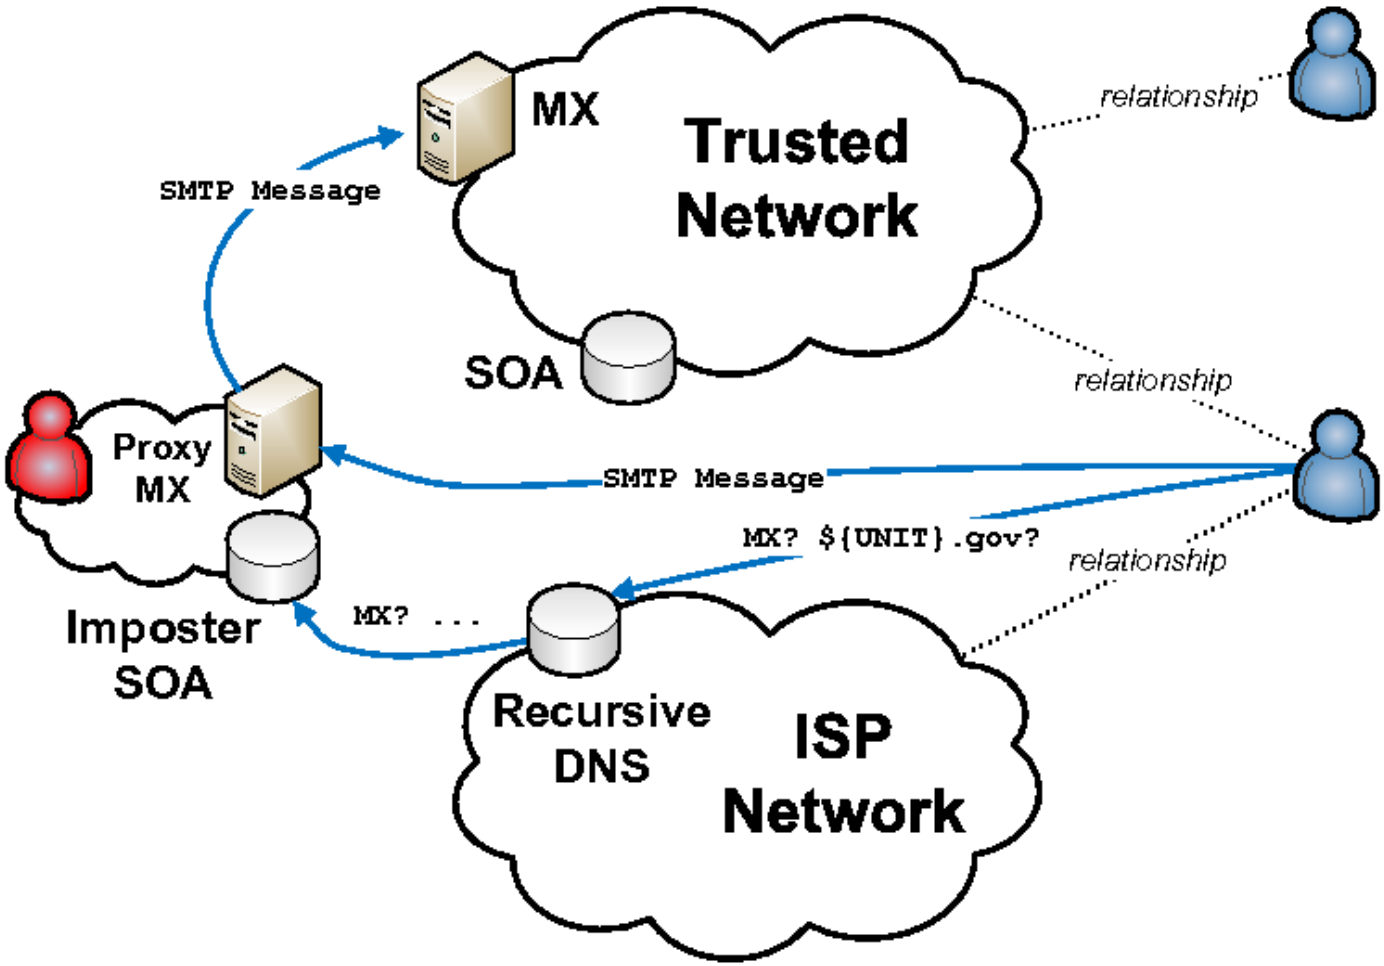

DNS Patch Rates Over Time

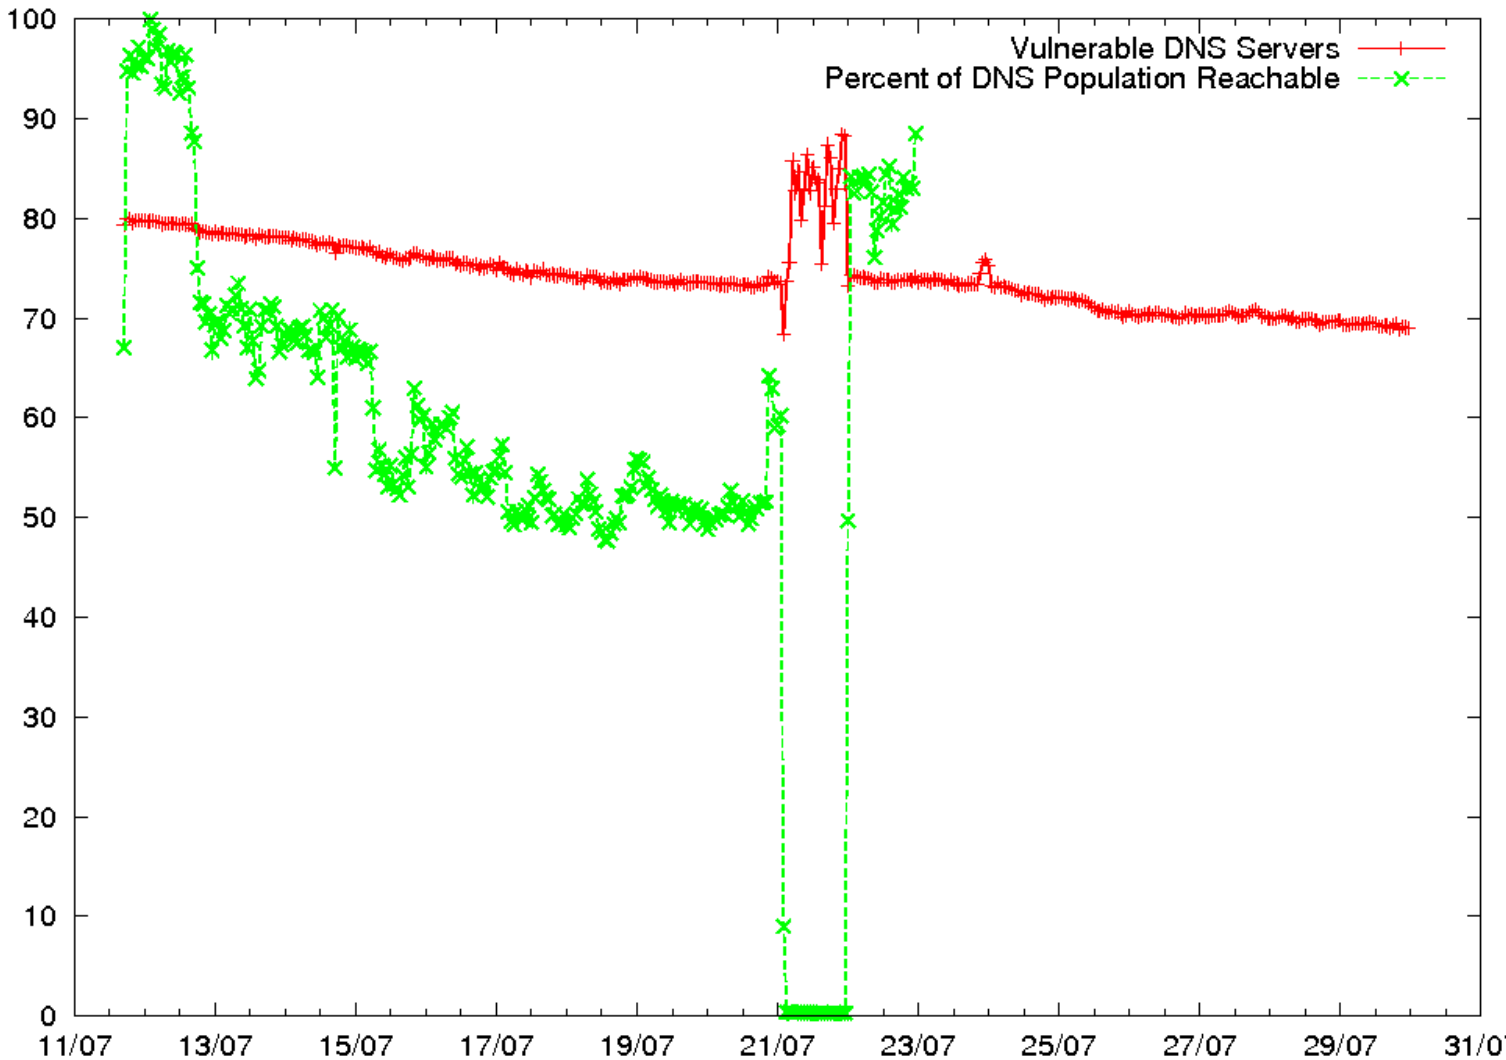

DNS Patch Rates Over Time

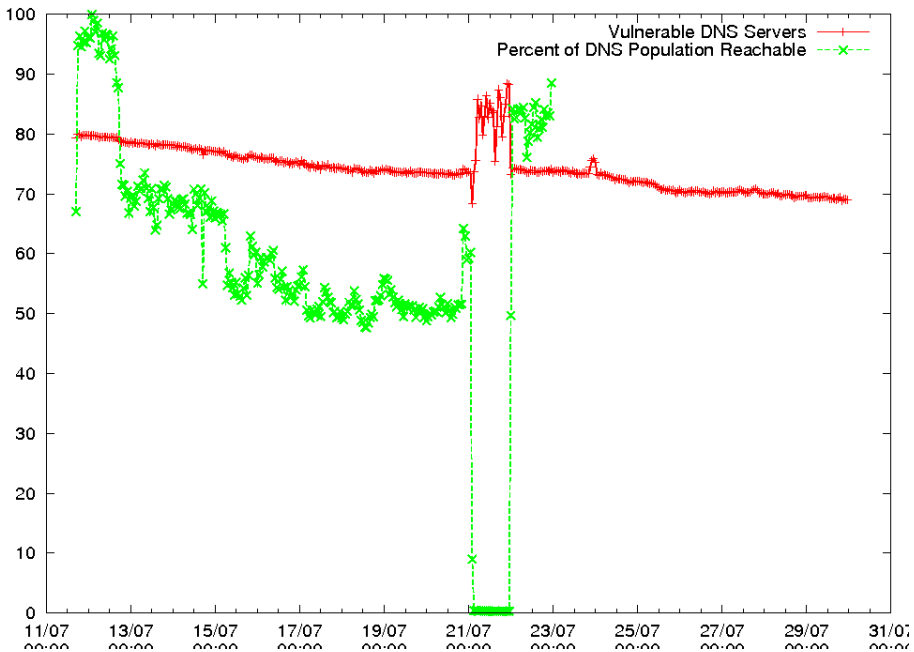

Salient points:

- Some 398,270 unique DNS servers probed over days, post `VU-800113`
- Slight decrease in vulnerable rate
- ... however, we also fail to reach many of the DNS servers originally identified (as much as half)
- Why? Many are dynamic hosts

DNS Patch Rates: Current

- Current rates of patching:
- Based on a subsample of tens of thousands of DNS resolvers
- 50% by number are unpatched
- 40% by popularity remain unpatched
- **Research Need:** Understand the patch agility of networks, applications, and services.

DNS Server BL Listing Periods

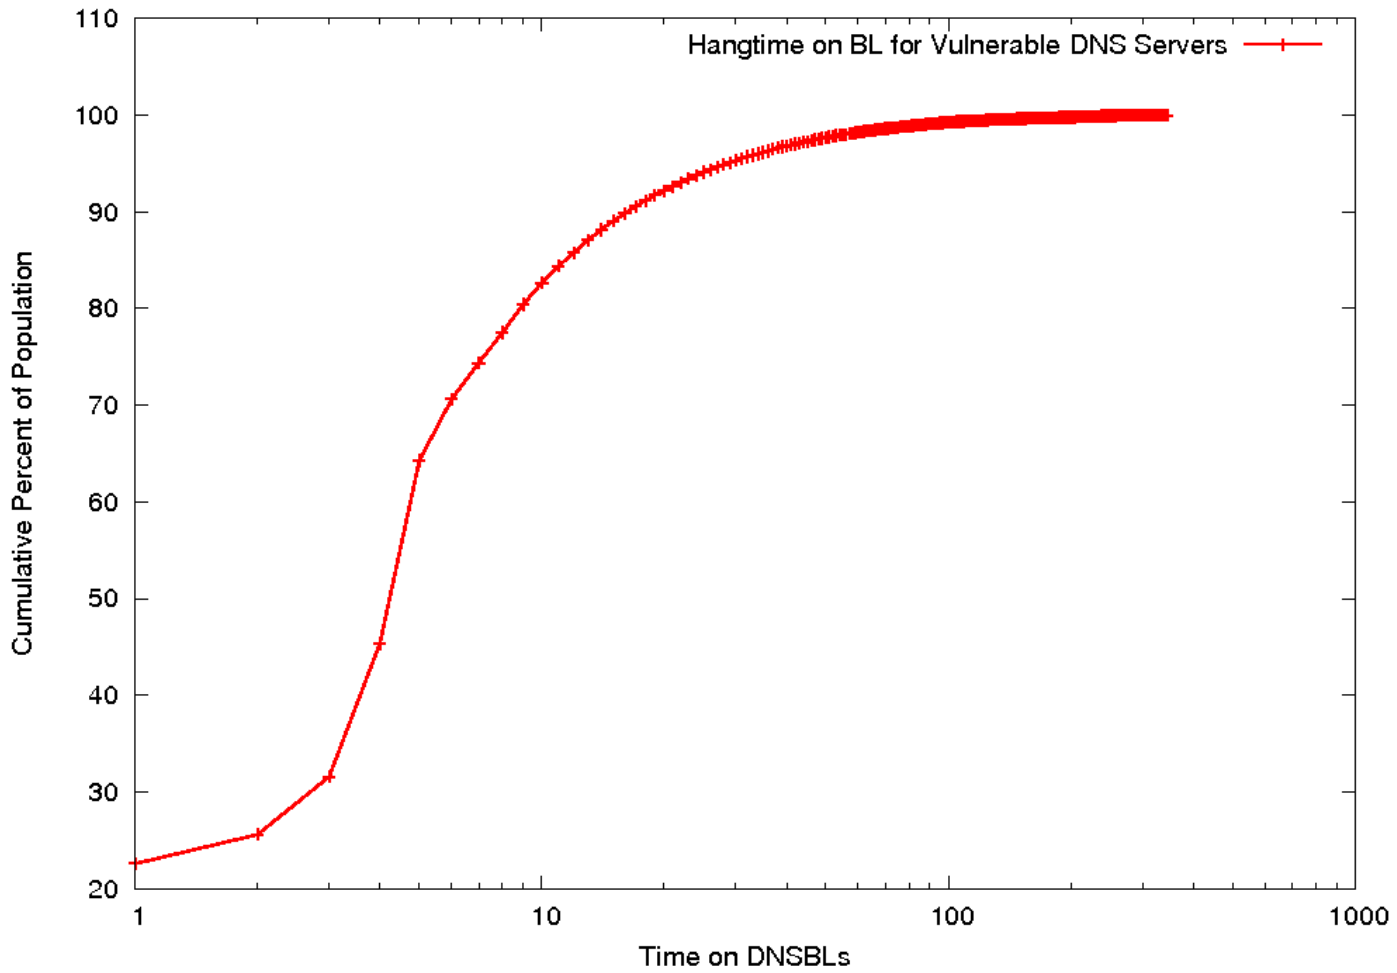

DNS Server BL Listing Periods

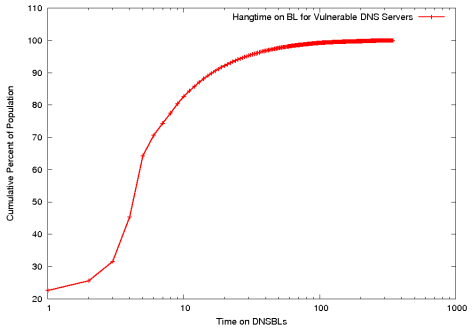

Salient points:

- Why were so many DNS servers no longer reachable after the initial probes?
- Some (34K, or $\approx$ 9%) are listed in the XBL, suggesting: an open recursive SOHO device, NAT'ing traffic for infected hosts at home (diurnal pattern masked where epoch is 86400).
- The graph shows the portion of population that persisted on the XBL, in days.
  - Thus, 80% of the "infected" hosts running DNS servers remained infected for 10 days or less. 20% are highly recidivist.

Vulnerable DNS Server Profiles

- How far behind in patches are the vulnerable servers?
- We can use ‘`fpdns -f`‘ to estimate...

```
40%  ISC BIND 9.2.3rc1 -- 9.4.0a0
15%  No match found
10%  TIMEOUT  id unavailable
9%  ISC BIND 9.2.0rc7 -- 9.2.2-P3
0.6%  Microsoft Windows DNS 2000
...
1 instance: ‘‘Dan Kaminsky nomde
                DNS tunnel’’
```

- **Research Need:** How to better characterize the *quality* of DNS service?

Where Are The Vulnerable Servers?

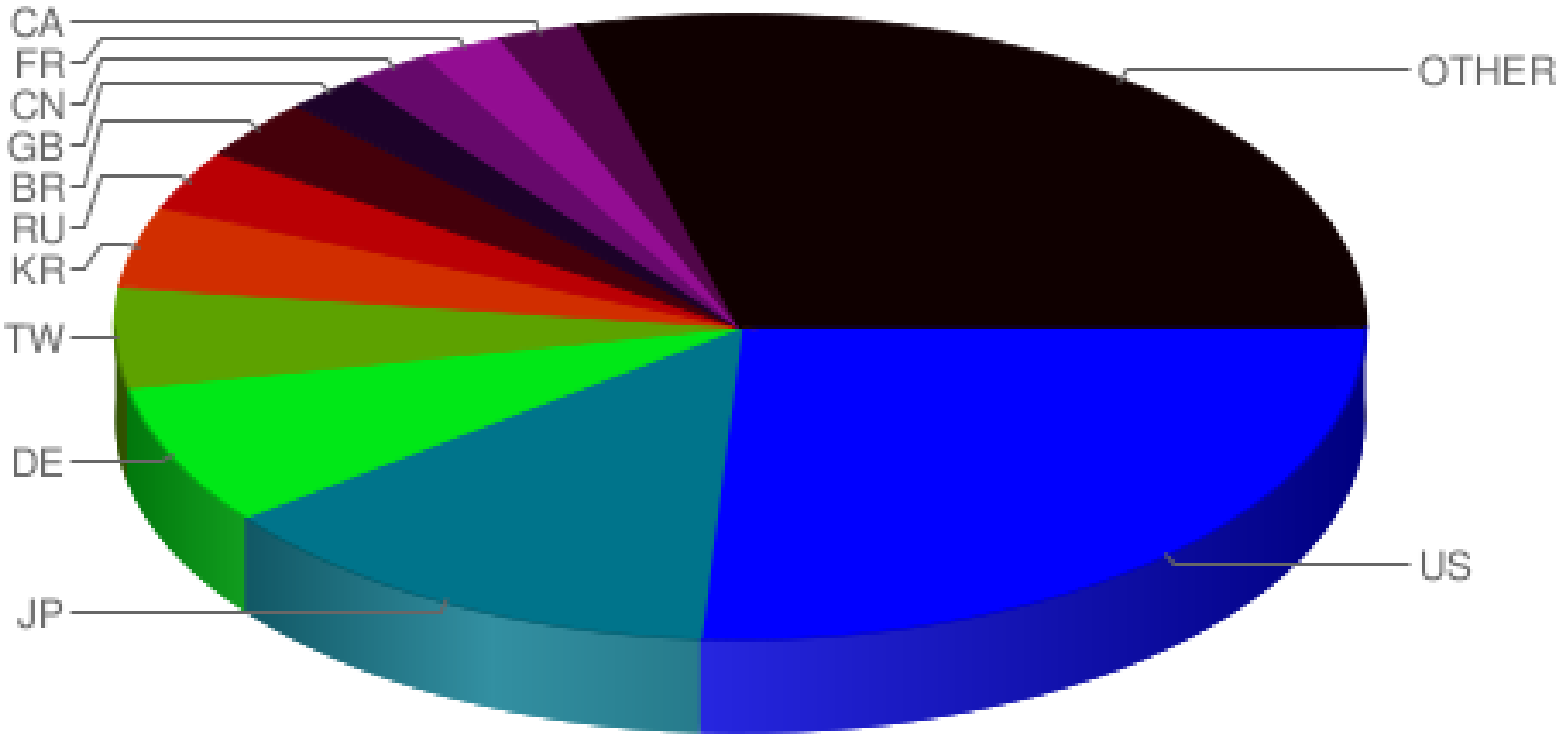

Where Are The Vulnerable Servers?

- Every security talk ever given has a pie chart showing IPs by country.
- This is that slide.

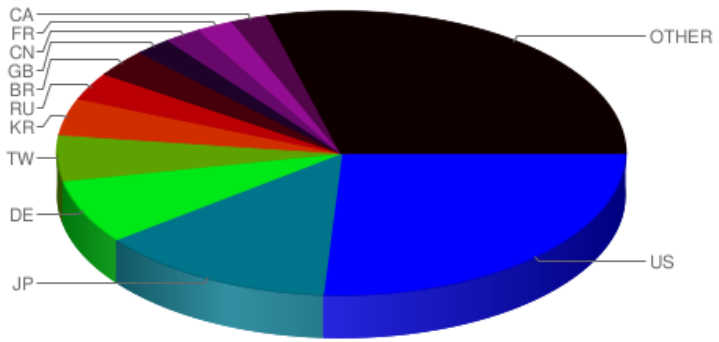

- What's the take-away?
- **Research Need:** A more relevant, actionable, insightful analysis of IP reputation systems.

Remedy #1: Source Port Randomization

- Vendor patches are available
- General direction: port randomization with cache logic enhancements
- **Research Need:** High performance techniques to randomize ports, without impacting resources

Remedy #2: DNS-0x20

- DNS messages often preserve the query formatting
- DNS-0x20 is an anti-poisoning technique that uses mixed-case queries
- We can run the dig command to test:

```
  dig @a.iana-servers.net www.EXamPLE.cOM
; <<>> DiG 9.5.0-P1 <<>> @a.iana-servers.net www
...
;; Got answer:
...
;; QUESTION SECTION:
;www.EXamPLE.cOM.                    IN      A
....
;; ANSWER SECTION:
www.EXamPLE.cOM.          172800    IN      A
```

Remedy #2: DNS-0x20

- DNS signaling preserves qname formatting.
- This allows provides addition bits for transaction processing.
- Thus, attacks have to guess not only the ID and src port, but also the *DNS-0x20 encoding* of the qname

DNS-0x20 Conceptual Use

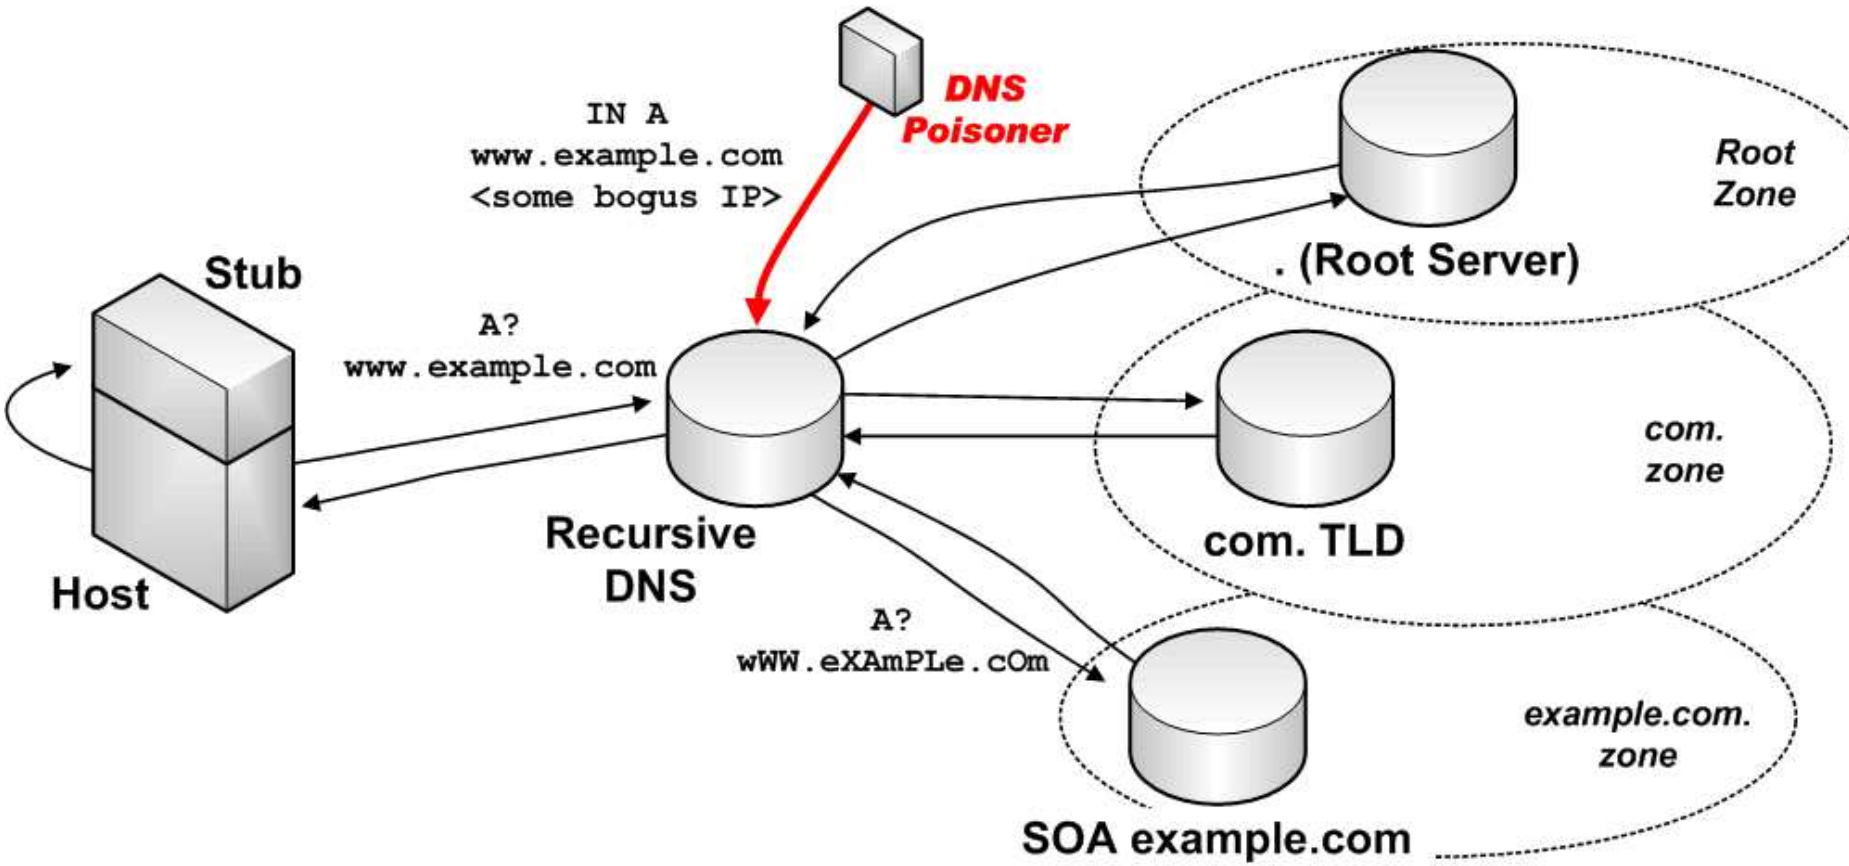

Where to learn more:

- IETF draft at:
  `http://www.ietf.org/internet-drafts \`
  `/draft-vixie-dnsext-dns0x20-00.txt`
- Unbound DNS server uses; BIND releasing shortly
- 99.7% of authority servers can handle DNS-0x20; the non-compliant zones can be forwarded
- Note: Kaminsky-class poisoning can affect TLDs, which have few 0x20-capable bits
- **Research Need**: Additional light-weight enhancements for forgery resistance

Remedy #3: You and Your Lab

- **Research Needs:**
- The research community needs to describe this poisoning attack, scenarios, and create a new threat model.
  - Enumerate capabilities of attackers and defenders
  - Identify invariants
  - Rethink old assumptions (e.g., domain whitelistings)
- Merging of IDS/DNS technologies
- Better understanding of control-plane monitoring
  - How to determine if “wrong” DNS answers are malicious
  - Engineering problems; data collection problems; anonymity problems

Remedy #4: SIE Monitoring

- SIE (among other things) is a large, diverse DNS traffic sharing system.
  - ISC hosted; 60 Mb/s traffic
  - Numerous US ISPs, Universities contributing
  - Data sharing; data amplification
  - More details at `https://sie.isc.org/`
- **Research Need:** Leveraging SIE data for analysis of traffic reveals attacks, trends, flux, bots, spam, ... everything.