DNS Technologies for Resiliency

Eric Ziegast - 2017-04-20
MENOG 17 - Muscat, Oman

FARSIGHT SECURITY

Worth Reviewing

(but I don’t do today)

FARSIGHT SECURITY

Practical DNS Operations

Great tutorial by John Kristof (menog-dnsops.pdf)

- DNS protocol and system overview
  - Hierarchical structure of DNS, registries vs registrars, whois
  - Caches, recursion, delegation
  - Transport (UDP *and TCP too*)
- Best Common Practices
  - Multiple nameservers, load balancing, anycast, network diversity
  - Bad stuff: Open resolvers, cache poisoning, mitigation delays, hijacking
  - Consistency, Administration, tools
- Advanced
  - PassiveDNS, Logging
  - Tools for monitoring – dsc, dnstop, zonecheck

DNS Anycast

- Great tutorial by Gaurab Raj Upadhaya @PCH (MENOG 3):
  - upadhaya-Anycast-v09.pdf
- Good overview by Martin Levy @CloudFlare (MENOG 17):
  - link-to.pdf
- Well-written 5-part blog series @DDIGuru:
  - Anycast, Static, RIP, OSPF, BGP
- Not just for authoritative servers, recursive servers as well
  - DNS is critical to operations
  - Enhanced and public DNS service providers use it
  - User-facing ISPs should investigate

Authoritative Anti-DDoS

Roots

- Many root server operators deployed anycast
  - More global bandwidth adds resiliency
  - Localized DDoS attacks
  - Reduced latency
- Some are large load-balanced nodes while others are single servers very broadly deployed
- DNS software and architecture diversity
  - Several different software back-ends
  - Different management practices
  - Different deployment strategies
- Deploy your own root? alternate ccTLD?
  - Do we have a plan for a Mirai-sized country-level DDoS?

FARSIGHT SECURITY

Authoritative Anti-DDoS

ccTLD/gTLD

- Old method:

  "Do you have a secondary I can add to my list?

- Today, in light of typical DDoS:

  "Let's add a mix of global anycast/cloud partners"

  Several have started service since 2009.

- Even then, still not enough
  - DDoS enough to knock any single provider down (Oct 21, 2016)
  - Mix of multiple providers?
  - Upstream DDoS mitigation?

End Users

- DDoS mitigation services (roll-your-own, paid, free)
- Adaptive response to DDoS (banks)

FARSIGHT SECURITY

Response Policy Zones (RPZ)

FARSIGHT SECURITY

DNS RPZ - Motivations

- "Taking back the DNS" - Paul Vixie
- Domains are cheap – hostnames are cheaper
- Cleanup of domain abuse is:
  - time-consuming
  - expensive / cost shifting
  - ineffective / too slow
  - in some cases not possible (bulletproof / registry policy)
- Criminals tend to reuse same infrastructure
  - Not just domains => global identifiers (IP, nameserver)
- Not all "crime" is equal – allow end user flexibility

FARSIGHT SECURITY

Newly observed domains

```
$ nmsgtool -C ch212 |egrep 'domain: [0-9]'
domain: 5685555.cc.
domain: 584033323.cn.
domain: 7rs5mleto3.xn--fiqs8s.
domain: 569517.cc.
domain: 569527.cc.
domain: 0452nb.cn.
domain: 4kle0j6.ddns.net.
domain: 48647536.pw.
domain: 0zhb1o842a.nom.za.
domain: 3933573.pw.
domain: 569529.cc.
domain: 8phpnr7no96.tk.
domain: 5921547.cn.
domain: 607e5d26.ngrok.io.
domain: 569296.cc.
domain: 146909rjrp3z.pw.
domain: 575297140.cn.
```

FARSIGHT SECURITY

DNS RPZ - Motivations

- “Most new domain names are crap” *(Vixie, HITB 2012)*

  Eg: “x n - - 8 0 a k 6 a a 9 2 e . c o m” → apple.com
- Many domains today are registered, abused, and abandoned within a short time (NOD)
- Desire dynamic near-real-time distribution methods
- Multiple sources of policy information
- Previous methods not scalable
  - Fakeroot
  - Proprietary software

FARSIGHT SECURITY

RPZ Constraints and Goals

- The goal of DNS RPZ is a global technology standard and market for publication/subscription of DNS reputation information
- Must be unencumbered by patents or licenses, and available in many RDNS implementations
- Must not generate new wide area DNS traffic or make RDNS more fragile / less robust / slower
- Must not directly facilitate NXDOMAIN remapping or any other form of DNS pollution

DNS Hierarchy

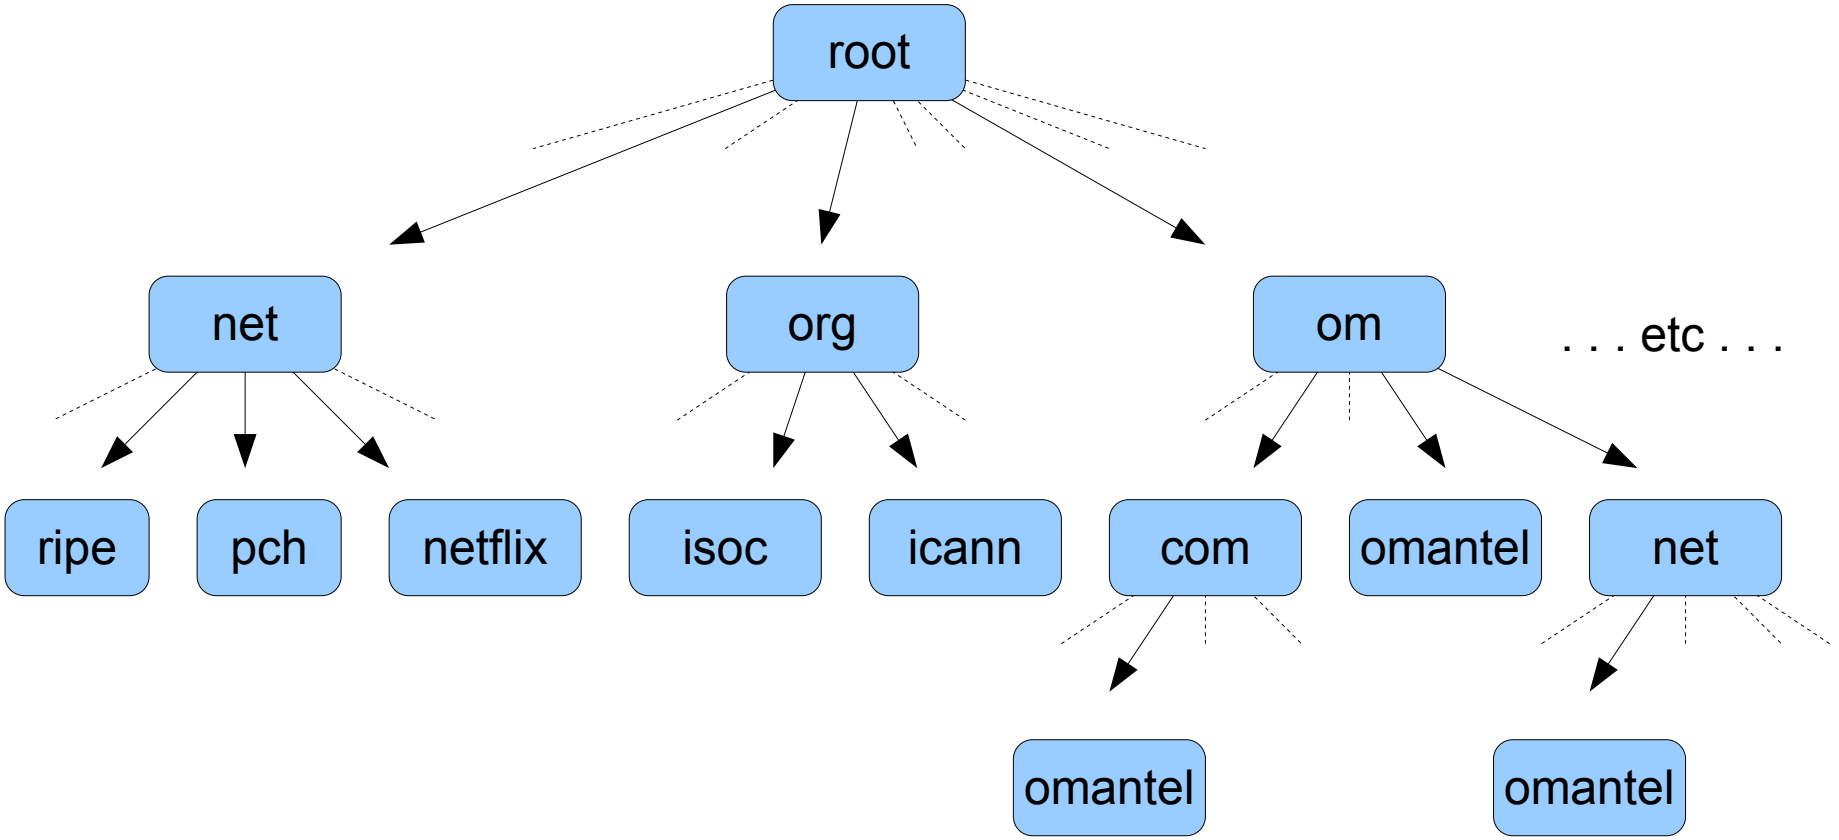

PS: Quick ”thank you” to MENOG17 sponsors

FARSIGHT SECURITY

Normal DNS

“ I’d like to visit www.malware-infected.com ”

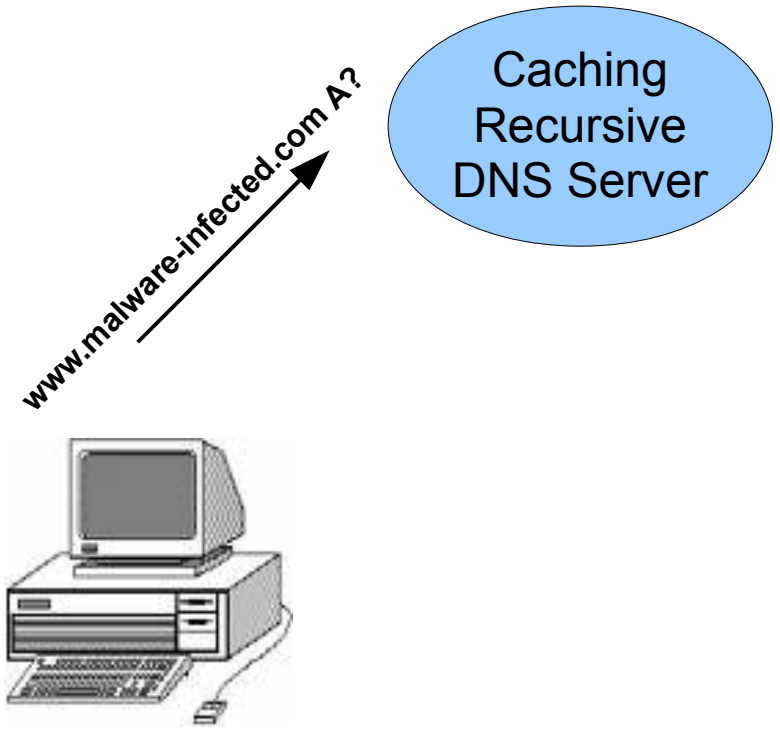

FARSIGHT SECURITY

Normal DNS

“ I’d like to visit www.malware-infected.com ”

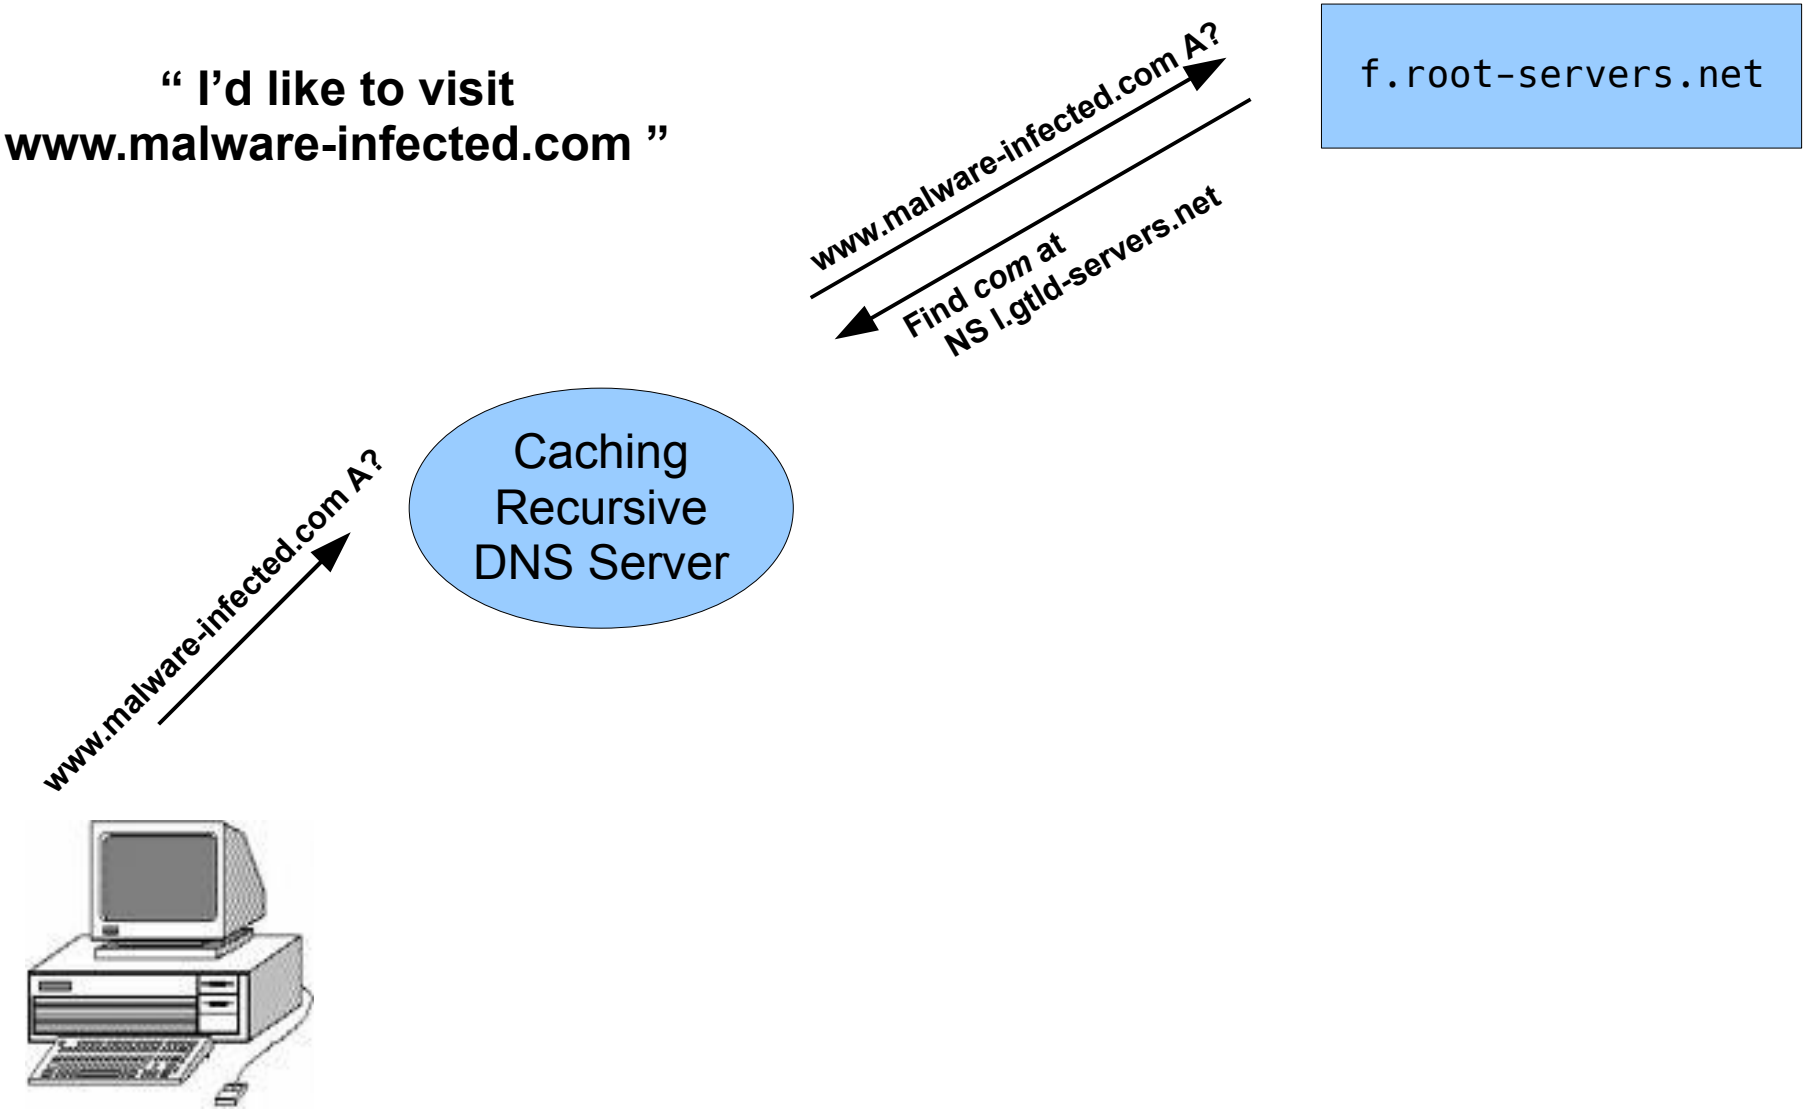

Normal DNS

“ I’d like to visit www.malware-infected.com ”

Above Recursive

`f.root-servers.net`

www.malware-infected.com A?

Find *com* at NS l.gtld-servers.net

Below Recursive

Caching Recursive DNS Server

www.malware-infected.com A?

FARSIGHT SECURITY

Normal DNS

“ I’d like to visit www.malware-infected.com ”

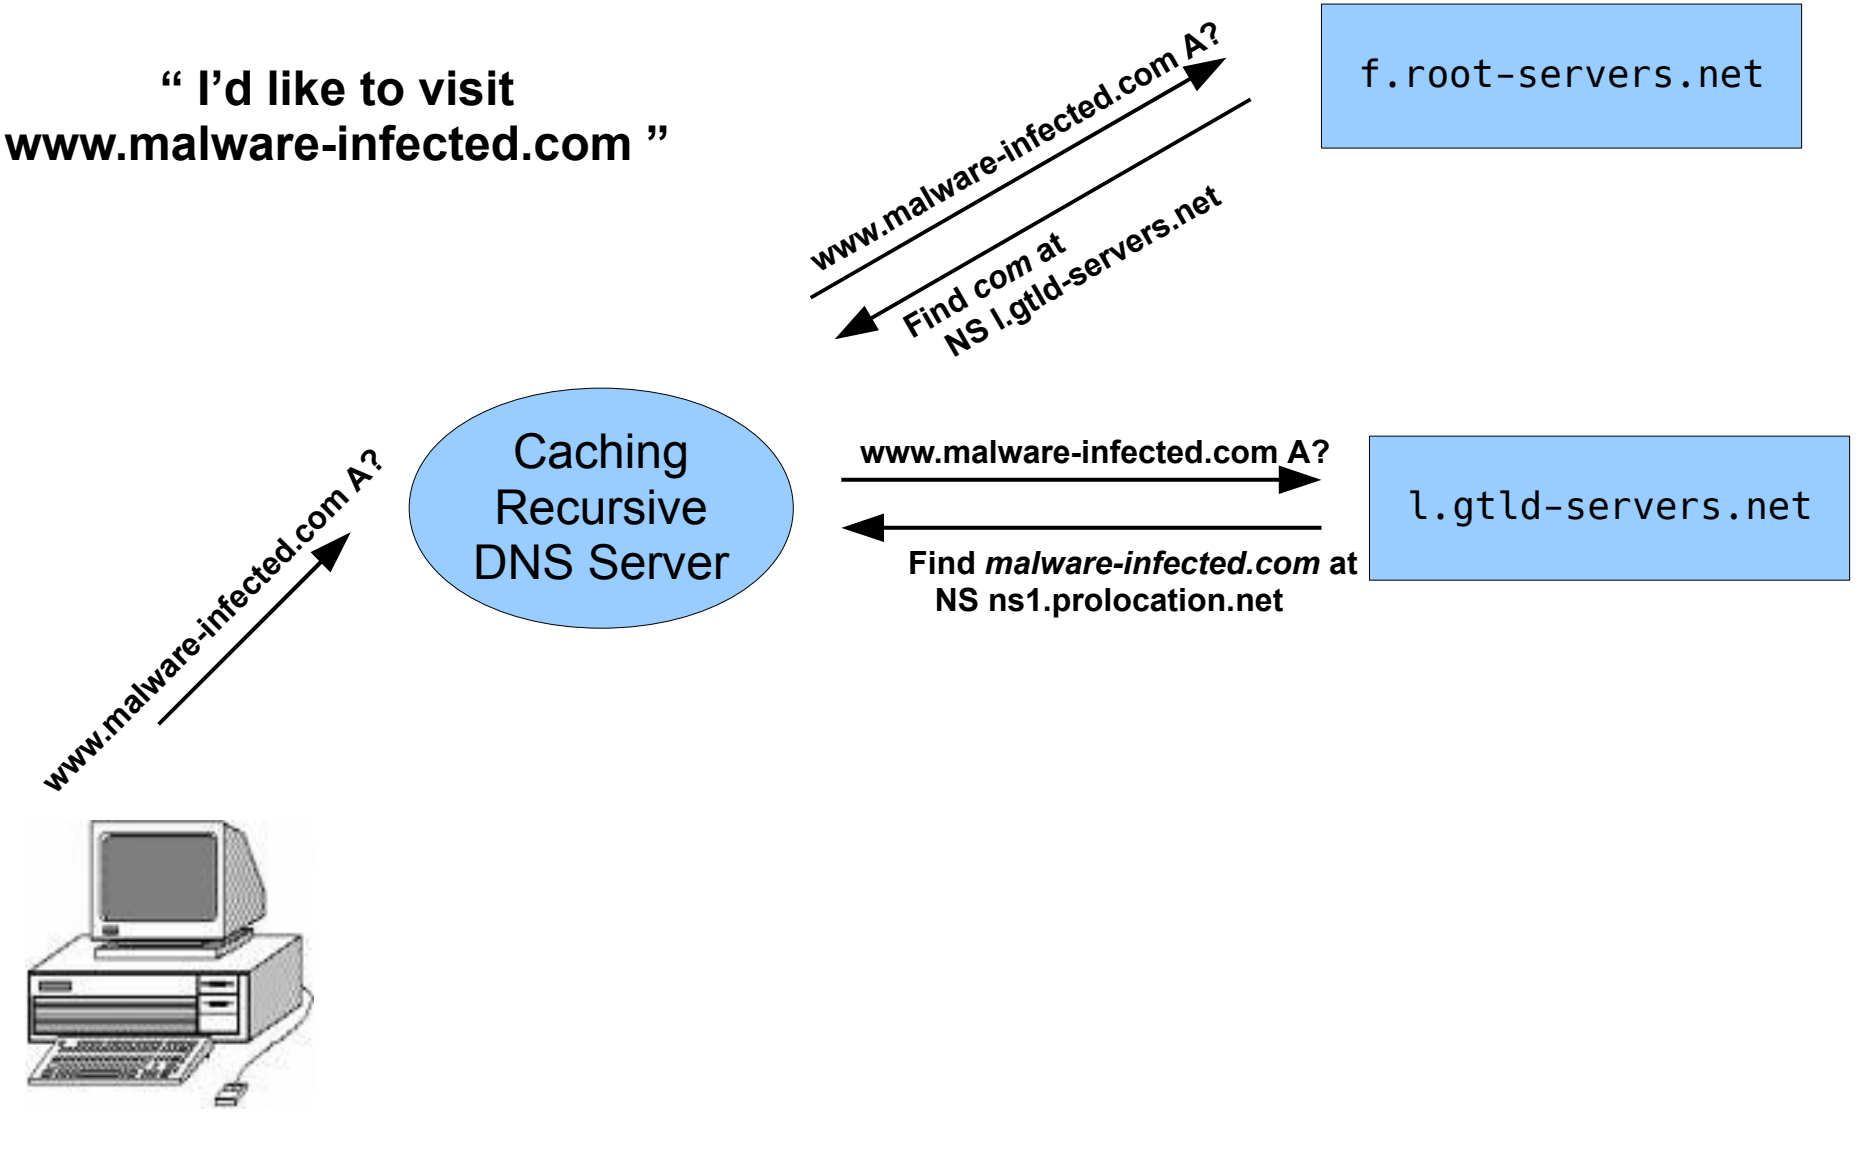

FARSIGHT SECURITY

Normal DNS

“ I’d like to visit www.malware-infected.com ”

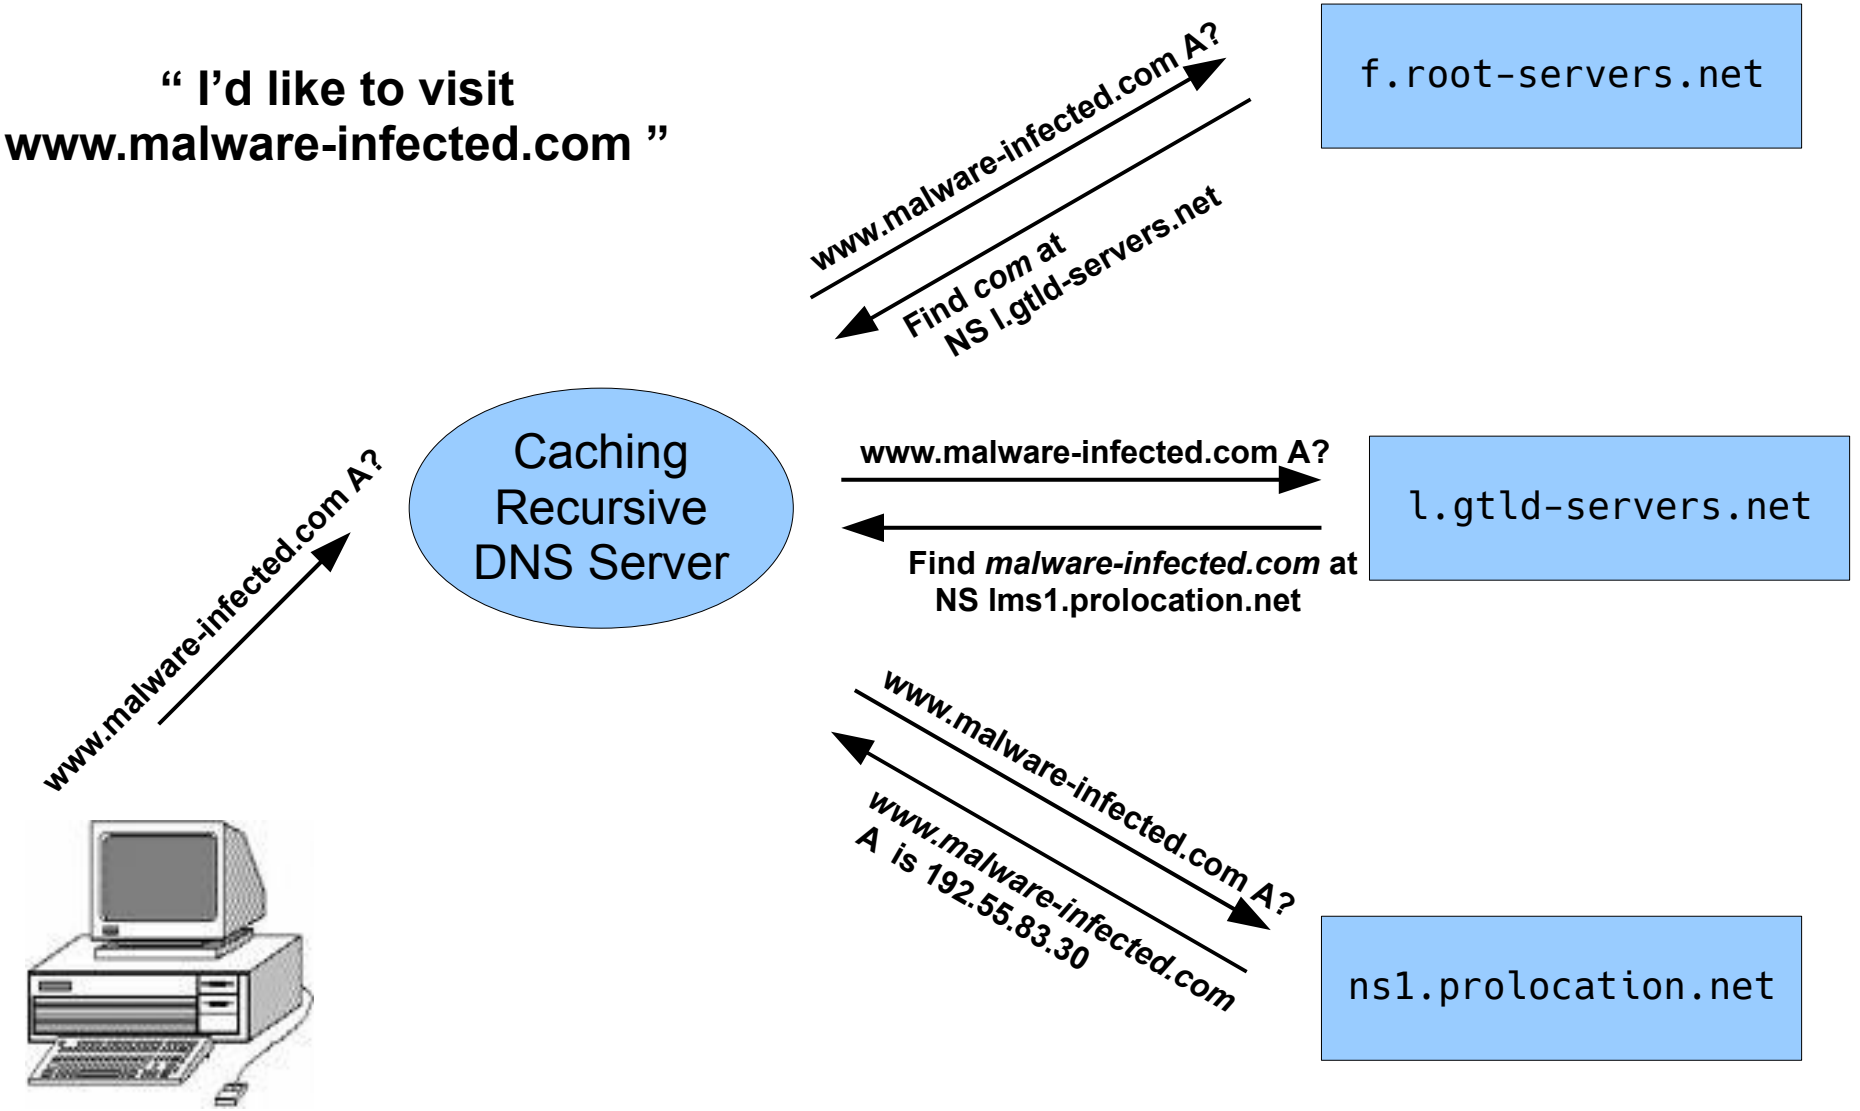

FARSIGHT SECURITY

Normal DNS

“ I’d like to visit www.malware-infected.com ”

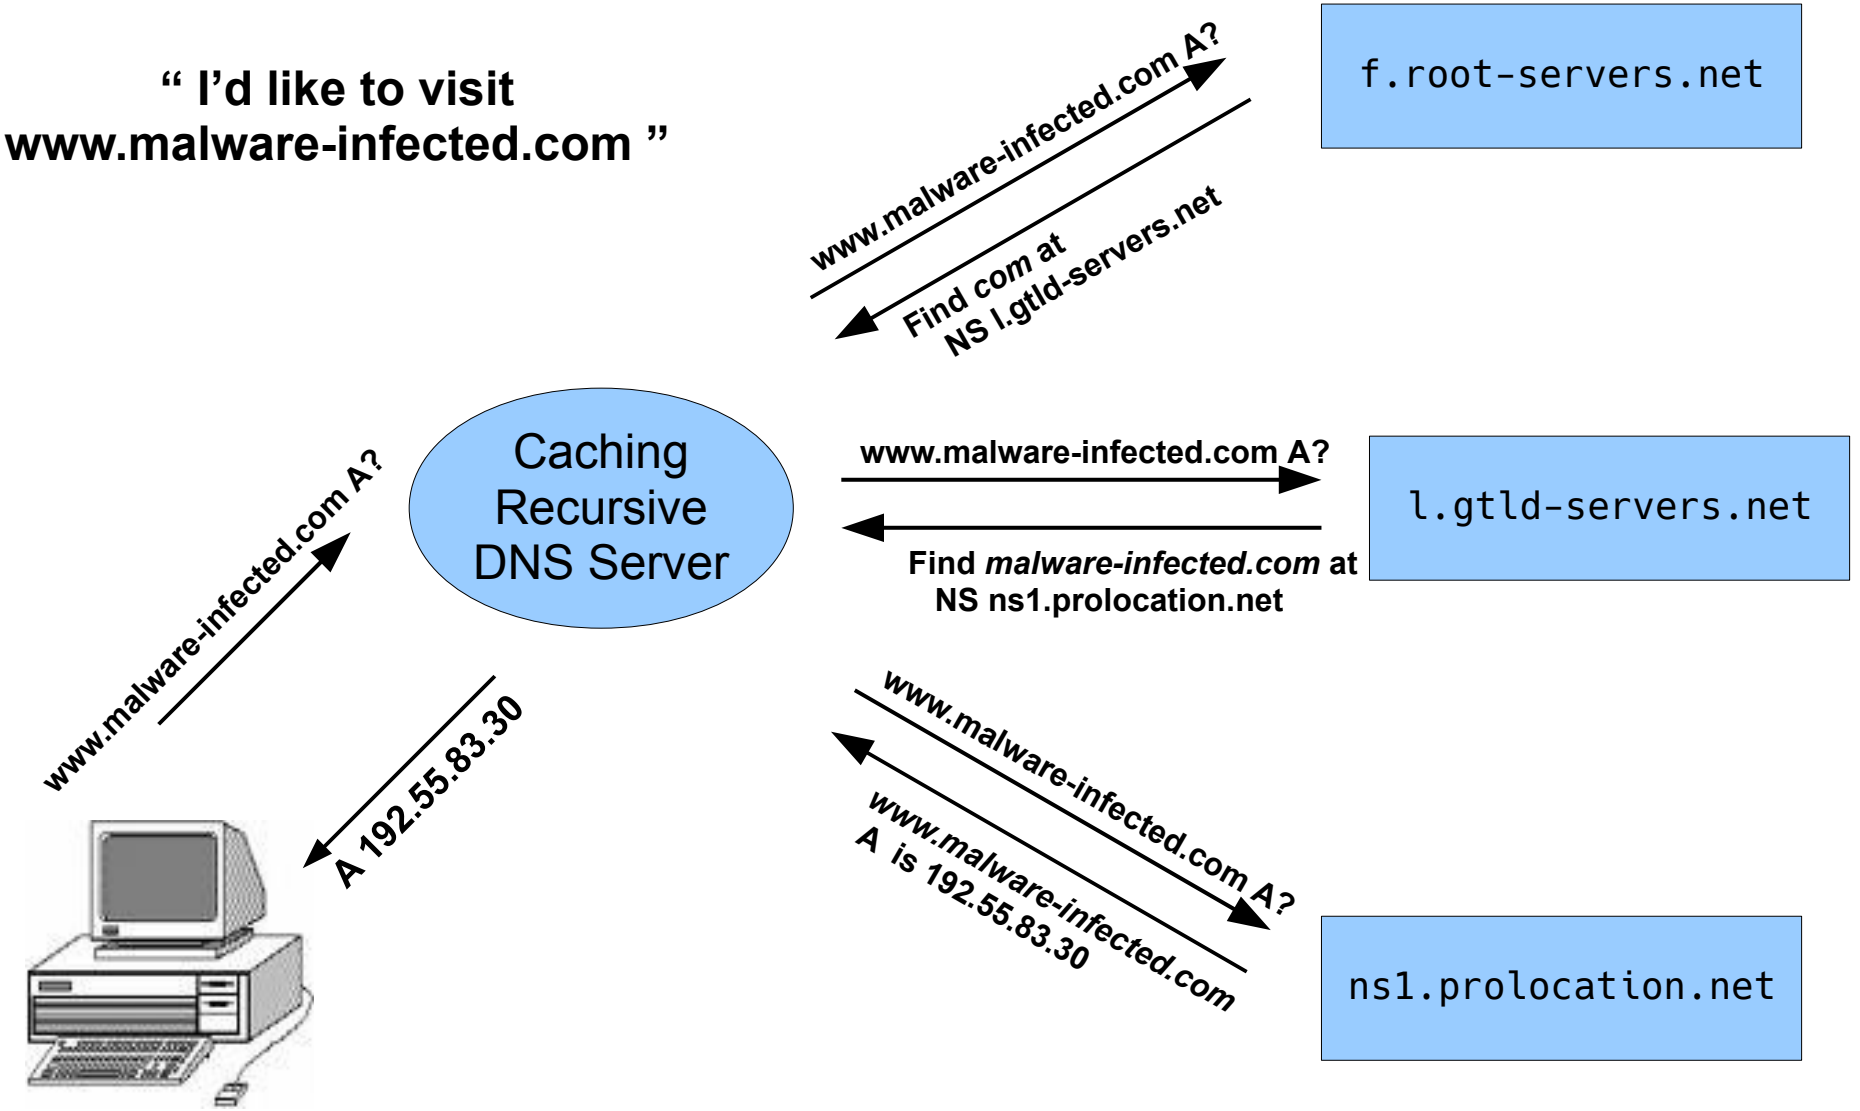

FARSIGHT SECURITY

Normal DNS

“ I’d like to visit www.malware-infected.com ”

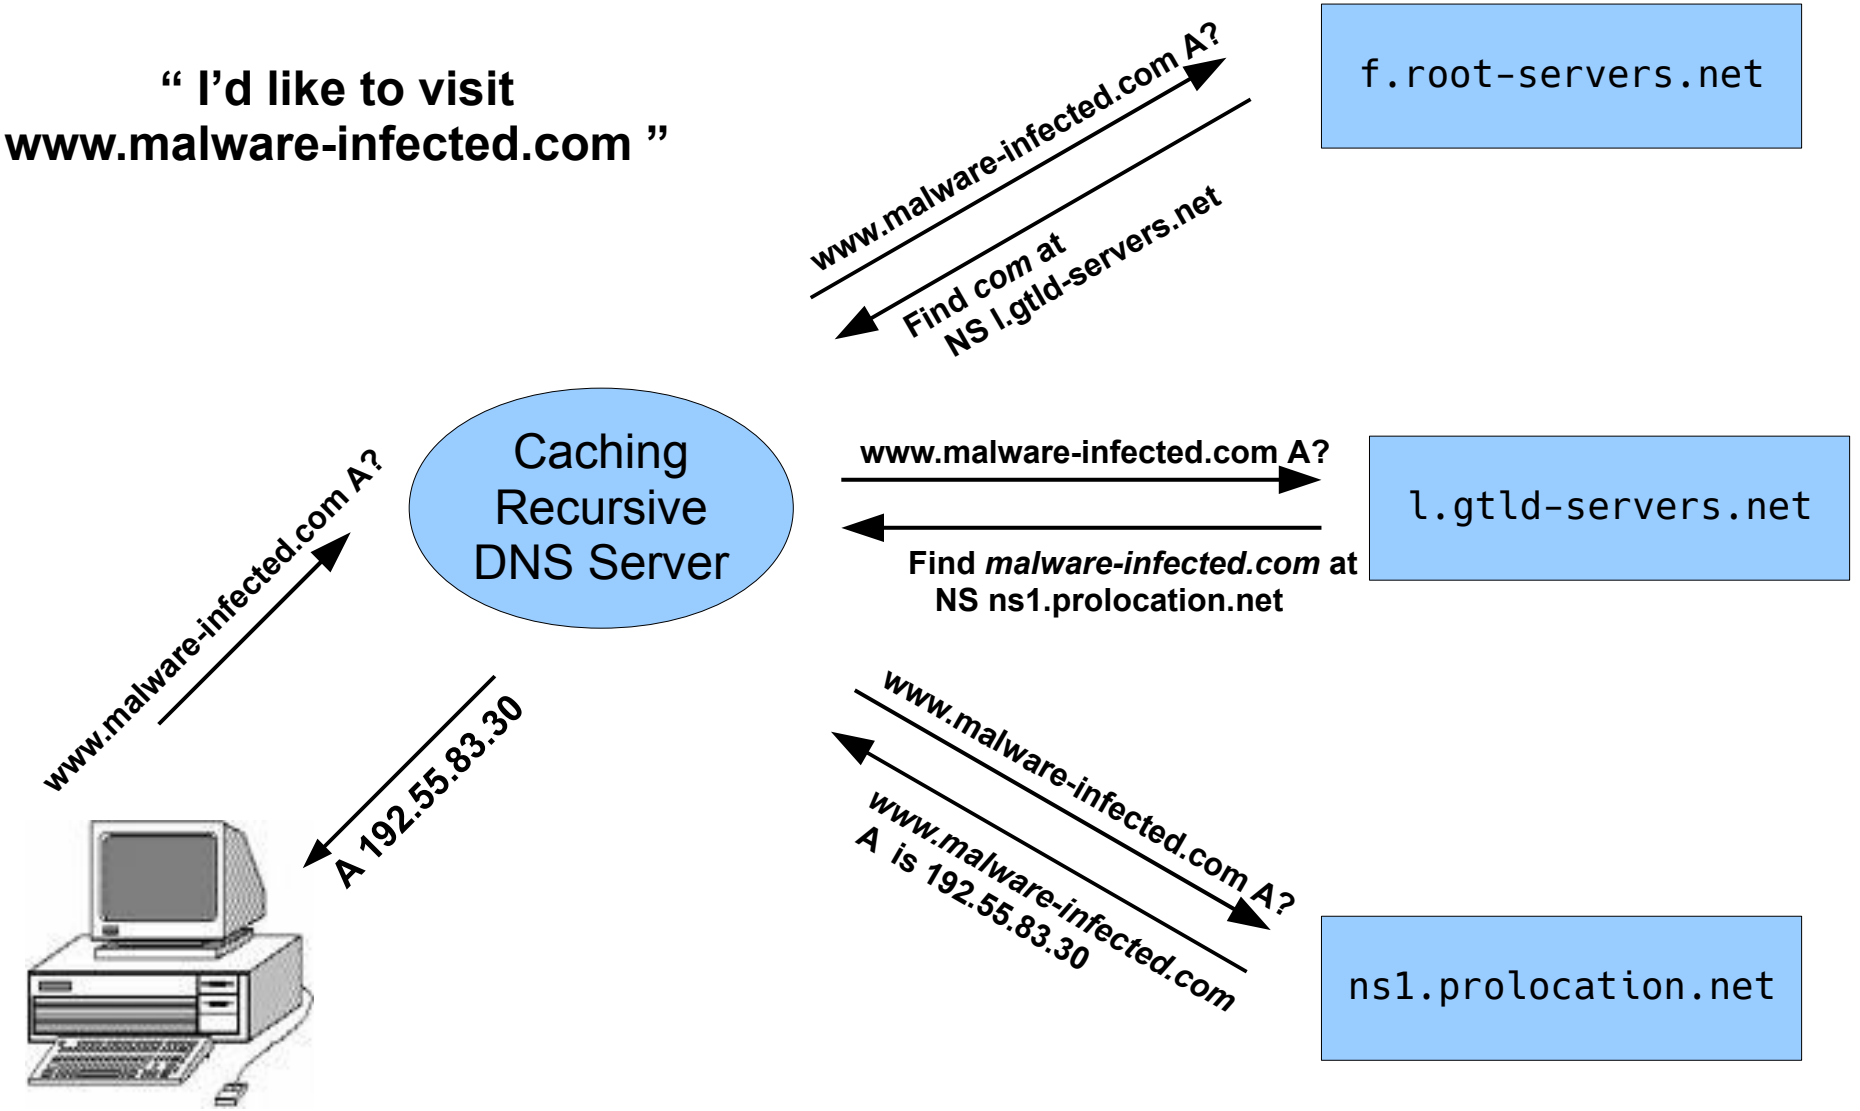

FARSIGHT SECURITY

Normal DNS

“ I’d like to visit www.malware-infected.com ”

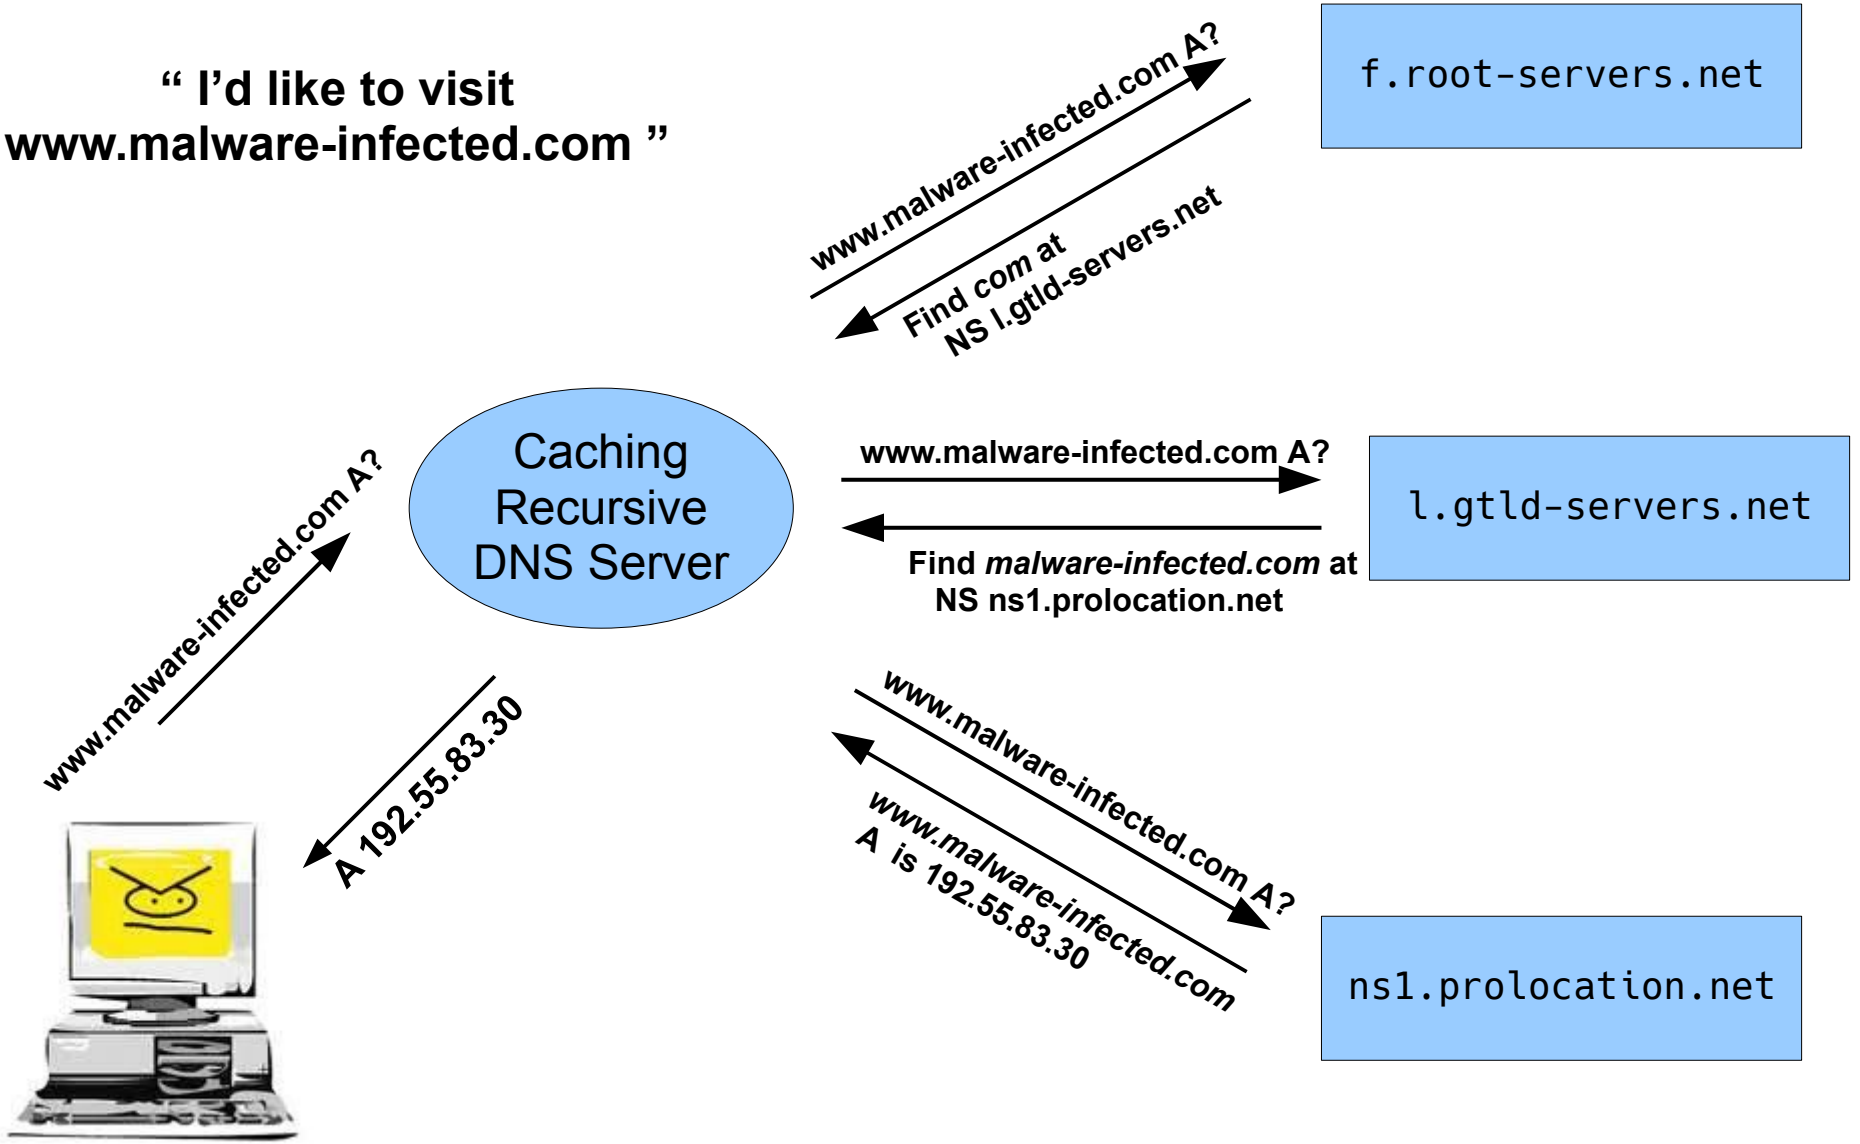

FARSIGHT SECURITY

RPZ

- “DNS firewwall”
- Subscribing RDNS servers are stealth secondary server for response policy zone(s)
- TSIG is used to control access and authenticity
- NOTIFY is used to ensure timeliness of updates
- IXFR is used to compress updates into deltas
- An RDNS can subscribe to more than one RPZ and if so they are searched in order, per query
- RDNS operators can use a mix of private and public RPZs, using search order for precedence

FARSIGHT SECURITY

RPZ

- Scalable method to modify DNS responses between recursive server and clients
- Multiple policies – ordering
- Maintained as DNS zones
- Quickly updated (dynamic updates)
- Efficiently/securely distributed
  - AXFR + IXFR updates
  - TSIG encryption
- RDNS operators can use a mix of private and public RPZs, using search order for precedence

FARSIGHT SECURITY

Multiple providers

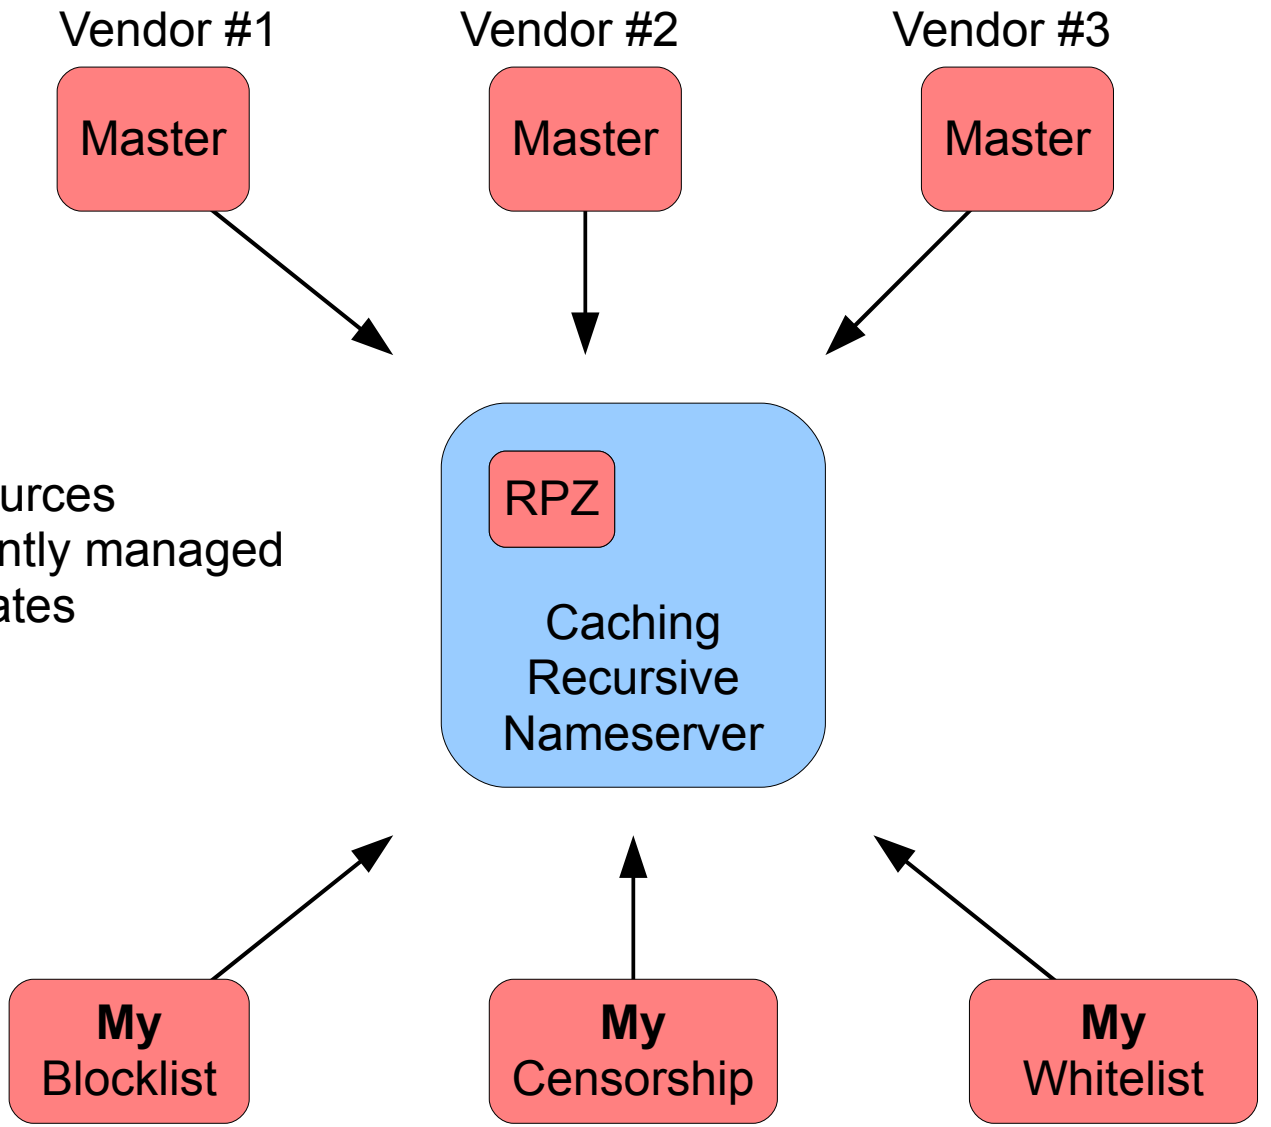

- Multiple sources
- Independently managed
- Quick updates

FARSIGHT SECURITY

RPZ Usage

- Maintained like standard DNS zone at it's own apex.
- Add zones like a secondary domain (use TSIG)
- Zone data transferred/updated like secondary domain
- RPZs are never queried and so need not be delegated by their parents nor have globally unique names
- Linkage from RDNS to RPZ is by configuration (BIND)

```
response-policy {
        zone "dns-policy.vix.com";
        zone "rpz.deteque.com";
};
```

- Read the draft:
  - https://tools.ietf.org/html/draft-vixie-dns-rpz-00

FARSIGHT SECURITY

RPZ policy actions

- To force an NXDOMAIN response:

  ```
  www.malware-infected.com.@ CNAME .
  ```

- To force a NODATA response:

  ```
  www.malware-infected.com.@ CNAME *.
  ```

- To stop processing and return the original answer:

  ```
  www.malware-infected.com.@ CNAME rpz-passthru.
  ```

- To make sure an answer is returned is returned as TCP only:

  (DDOS mitigation)

  ```
  www.malware-infected.com.@ CNAME rpz-tcp-only.
  ```

- To force no response (DROP):

  ```
  www.malware-infected.com.@ CNAME rpz-drop.
  ```

- To force a different answer:

  Use any normal RR, including CNAME:

  - `www.malware-infected.com.@ CNAME some.honeypot.server.`

FARSIGHT SECURITY

RPZ policy triggers

Rewrite answers for queried Hosts/Domains

- `host.domain.@`
- `*.domain.@`

Rewrite answers based on response IP address

- `prefix.B4.B3.B2.B1.rpz-ip.@` `(IPv4)`
- `prefix.W8.W7.W6.W5.W4.W3.W2.W1.rpz-ip.@` `(IPv6)`
- `prefix.zz.W3.W2.W1.rpz-ip.@` `("zz" is like "::")`

Rewrite all answers from a client (think "walled garden", login director)

- `prefix.zz.W3.W2.W1.rpz-client-ip.@`
- `prefix.W8.W7.W6.W5.W4.W3.W2.W1.rpz-client-ip.@`

Rewrite all answers from a particular authoritative server

- `NS.EXAMPLE.COM.rpz-nsdname.@`

Rewrite all answers from a particular authoritative server (trigger by IP address)

- `prefix.zz.W3.W2.W1.rpz-nsip.@`
- `prefix.W8.W7.W6.W5.W4.W3.W2.W1.rpz-nsip.@`

FARSIGHT SECURITY

DNS + RPZ

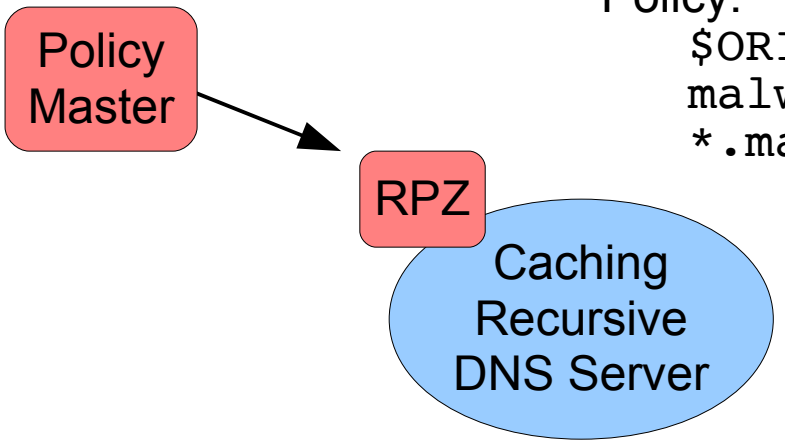

Policy:

```
$ORIGIN zone.surbl.org.
malware-infected.com.@ CNAME .
*.malware-infected.com.@ CNAME .
```

FARSIGHT SECURITY

DNS + RPZ

“ I’d like to visit www.malware-infected.com ”

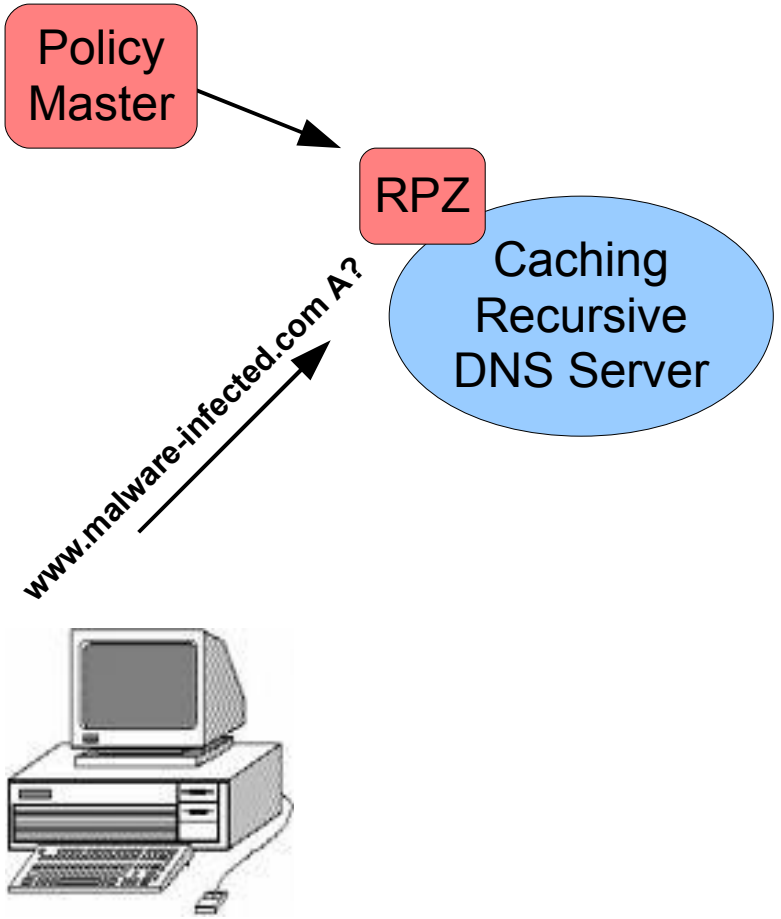

FARSIGHT SECURITY

DNS + RPZ

“ I’d like to visit www.malware-infected.com ”

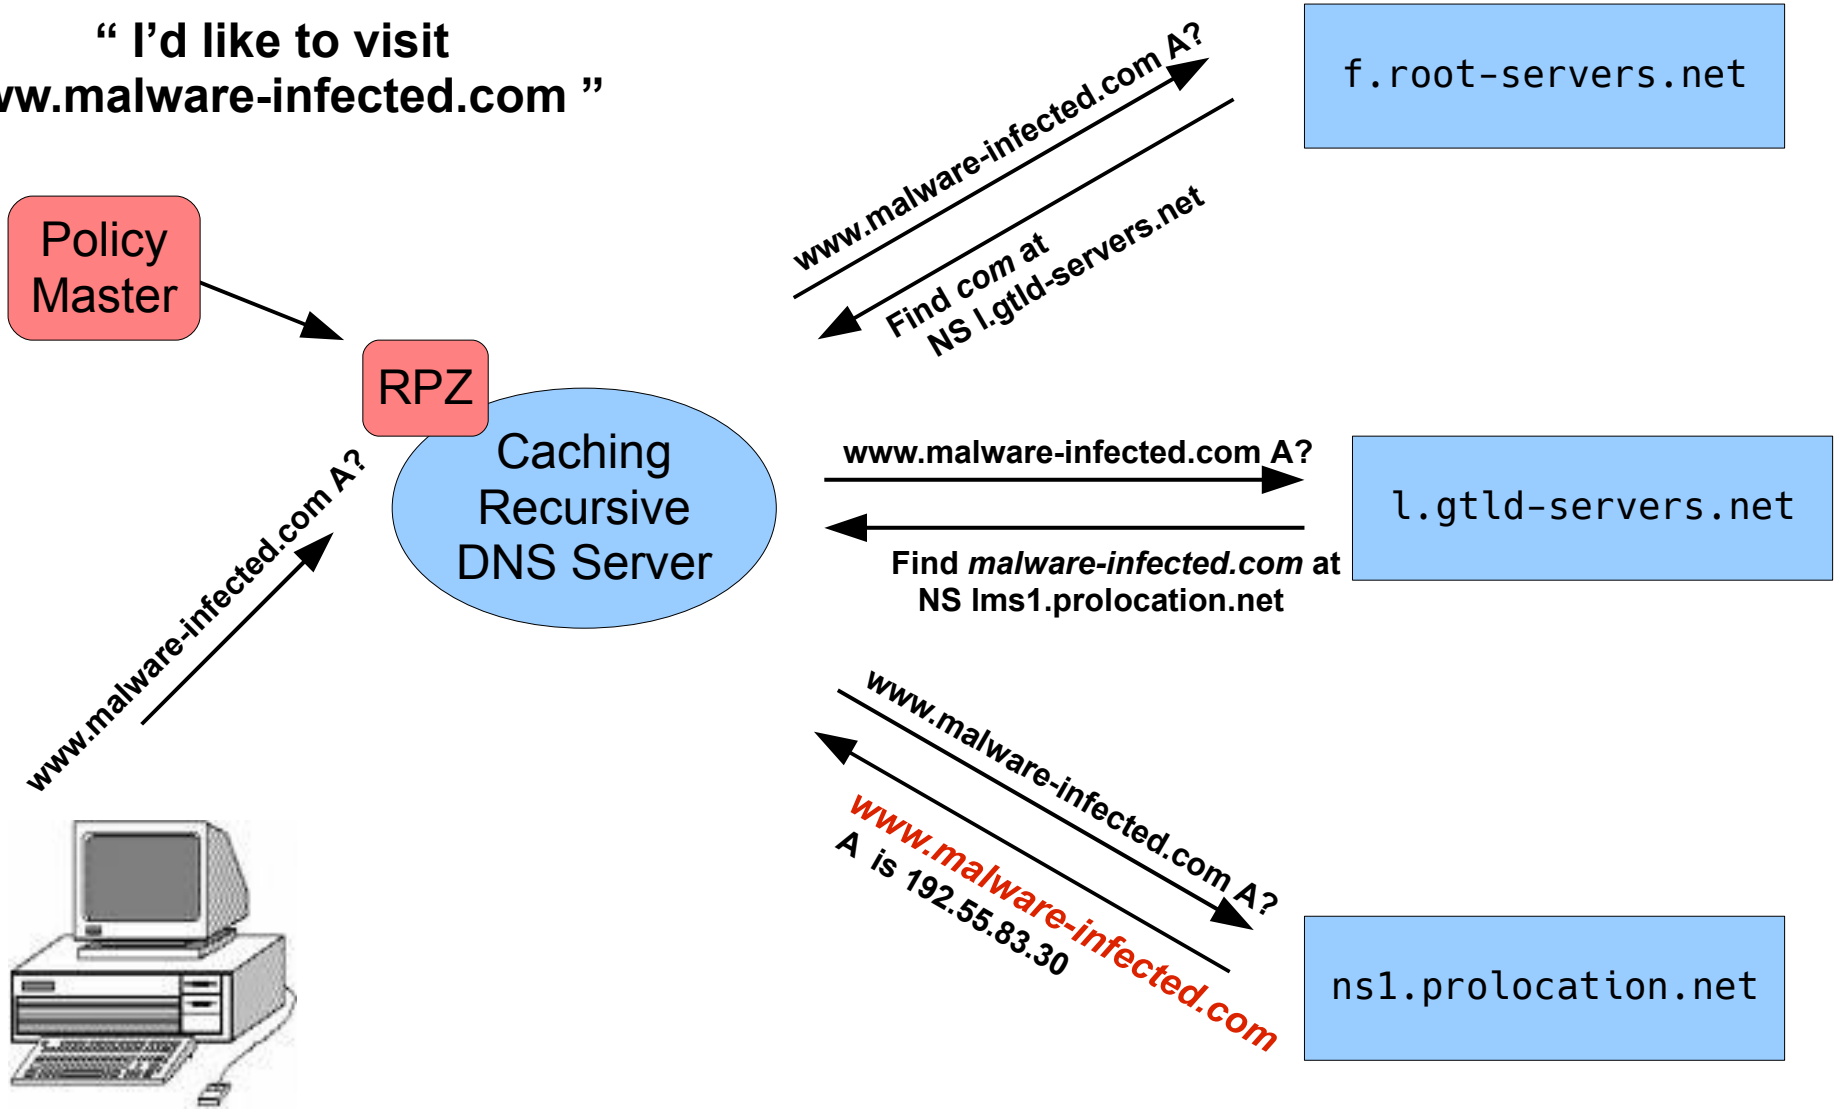

FARSIGHT SECURITY

DNS + RPZ

“ I’d like to visit www.malware-infected.com ”

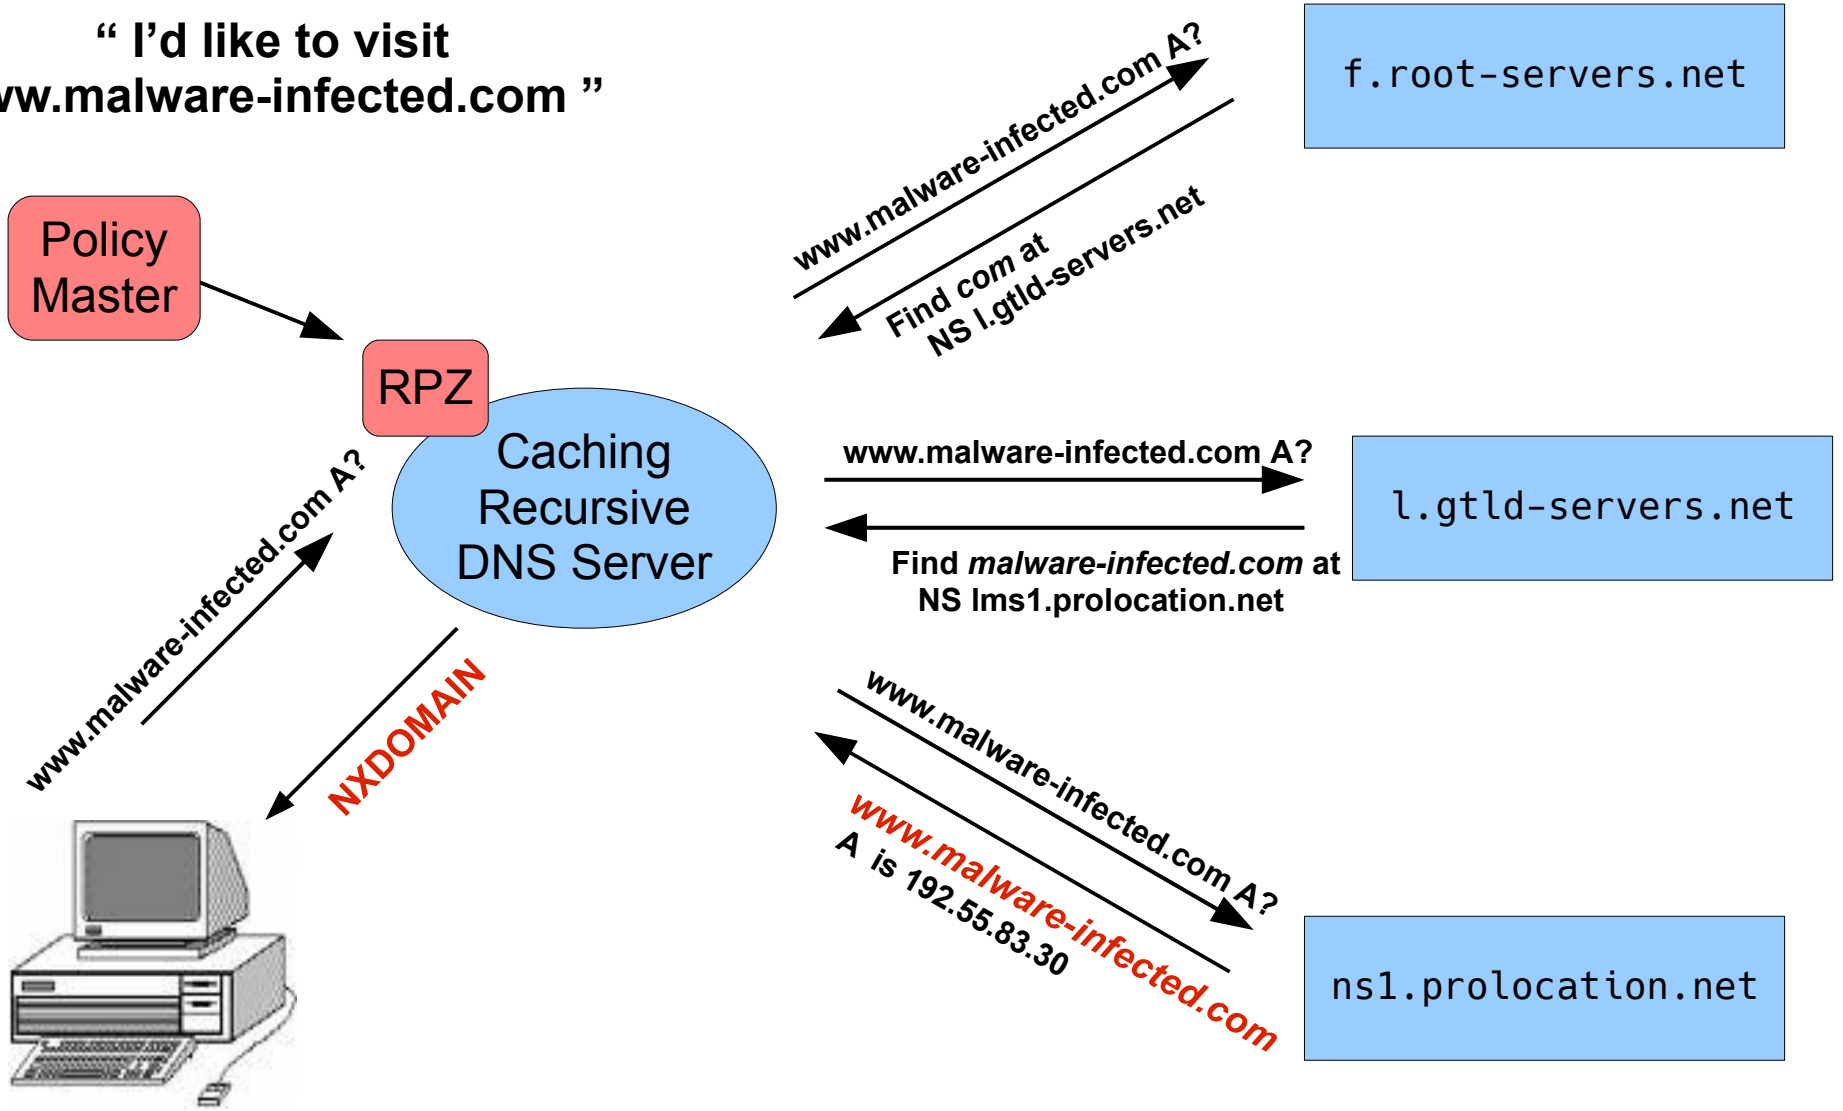

FARSIGHT SECURITY

Browser - before

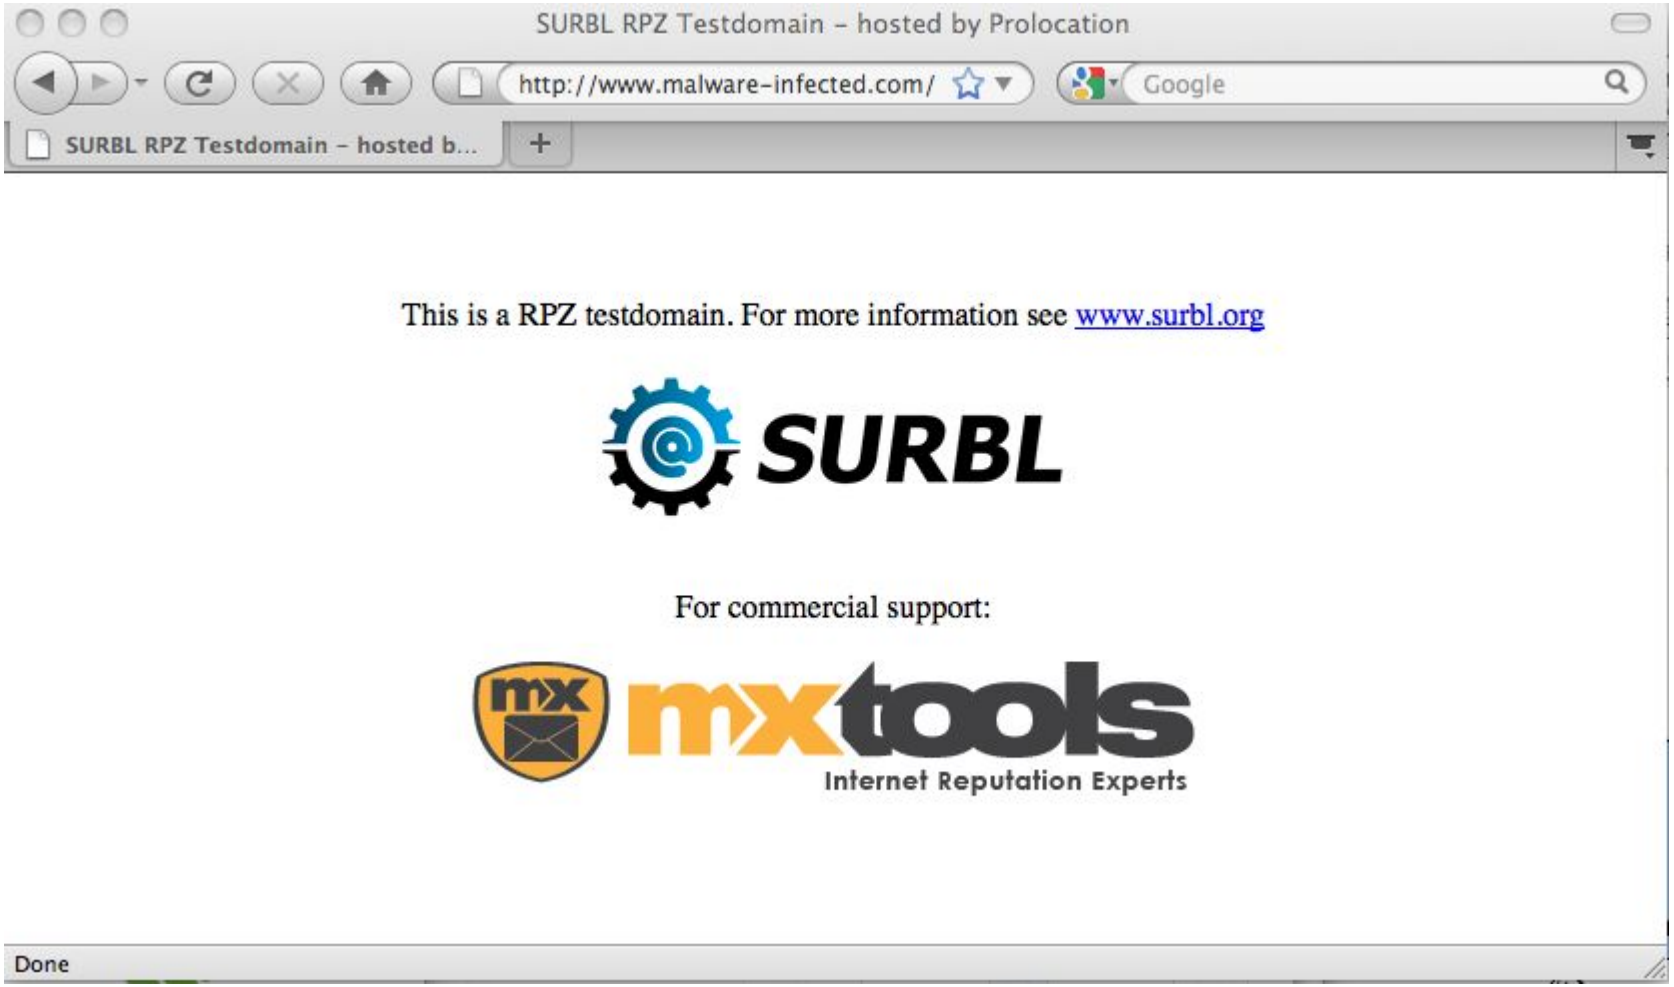

FARSIGHT SECURITY

Browser - after

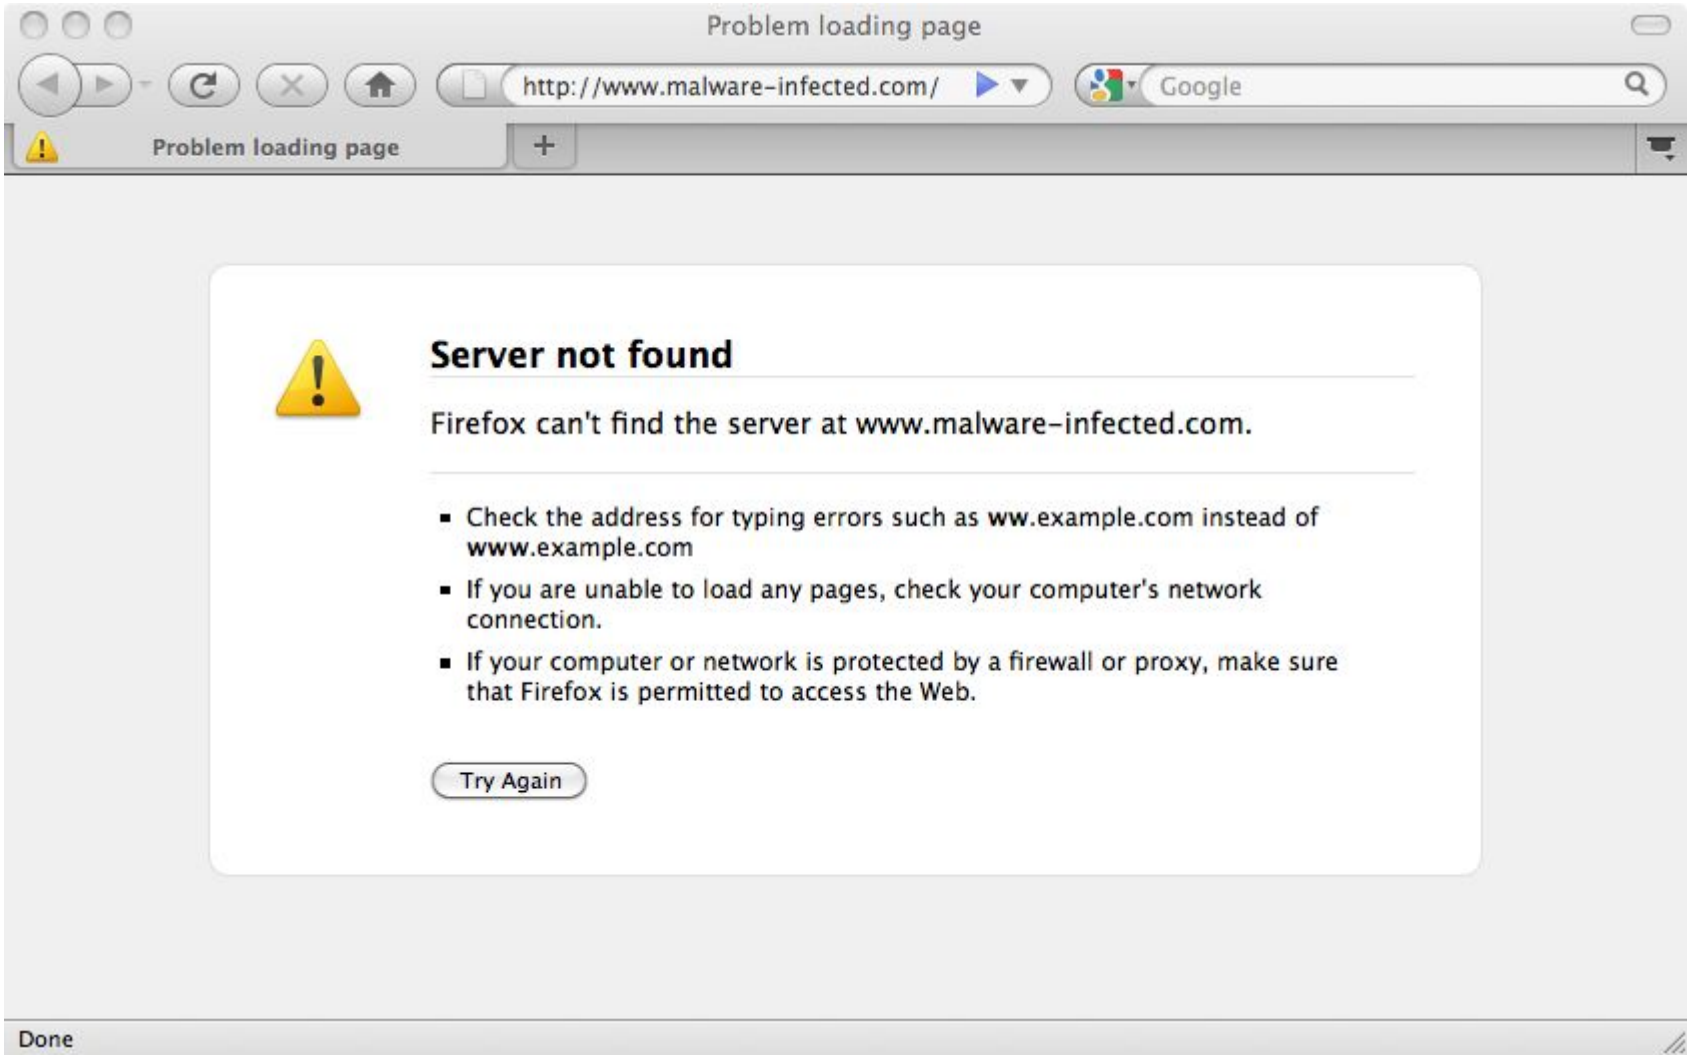

FARSIGHT SECURITY

Debugging

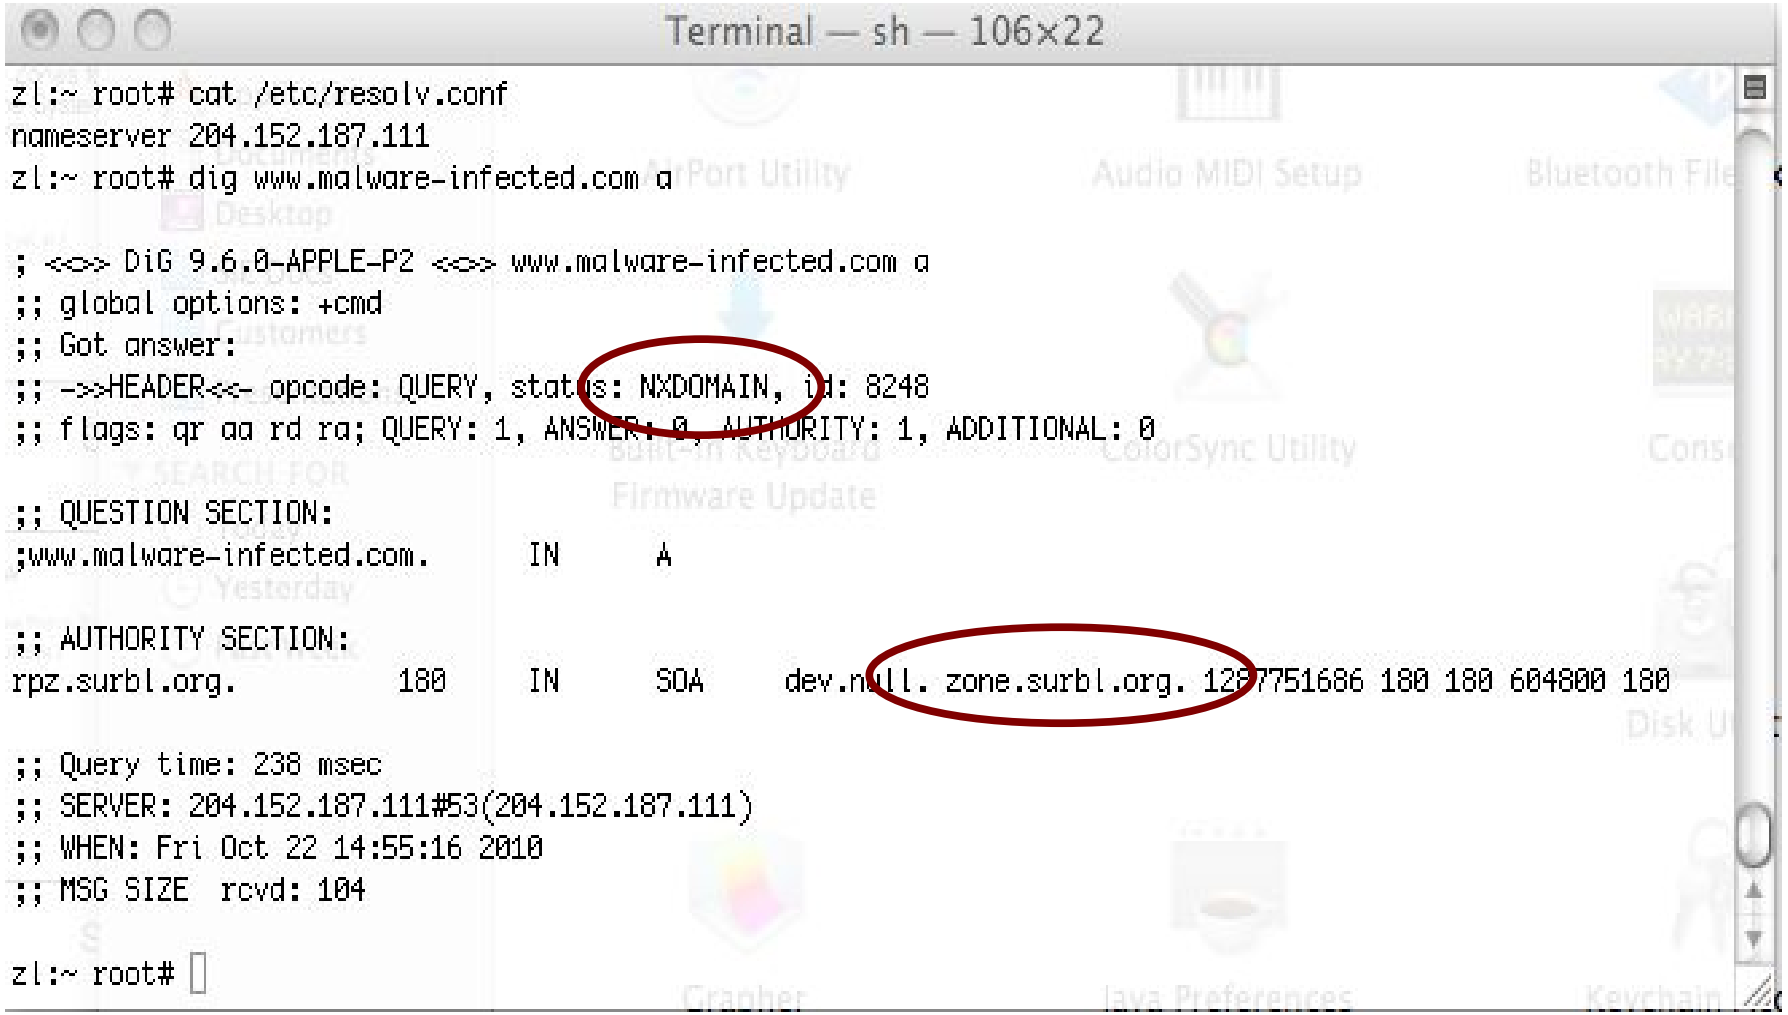

FARSIGHT SECURITY

Distribution

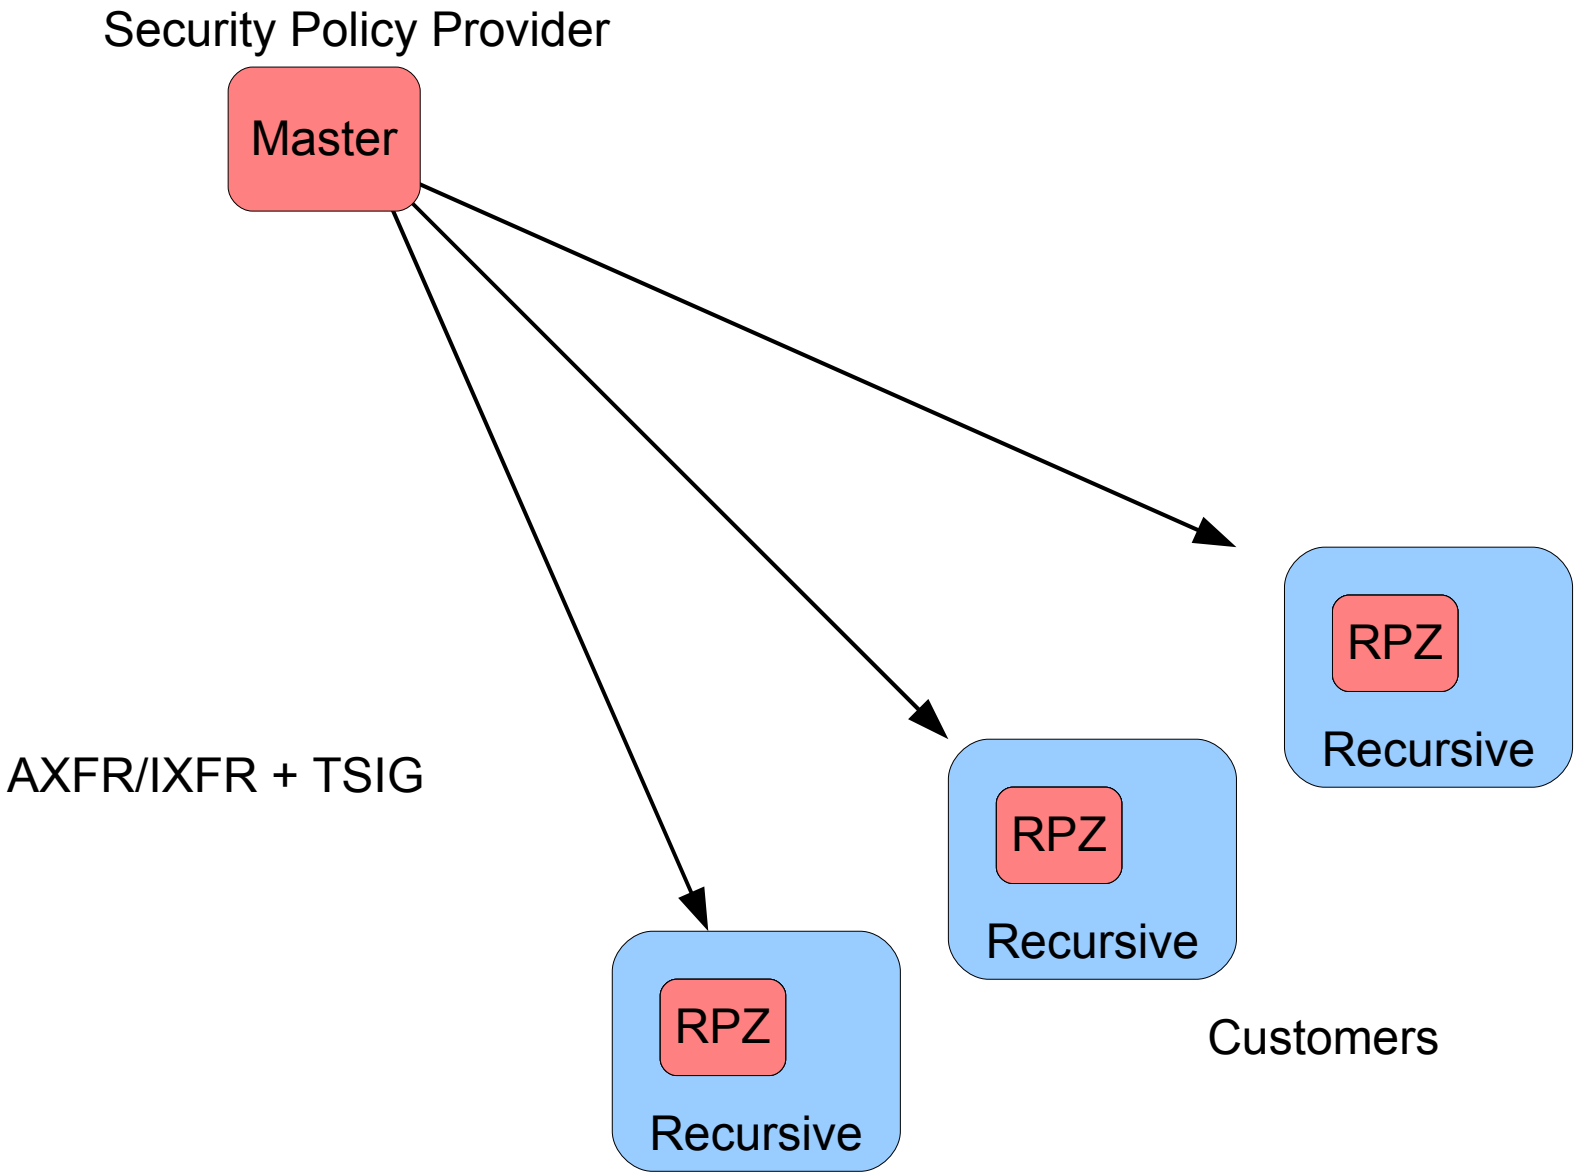

Scaling Distribution

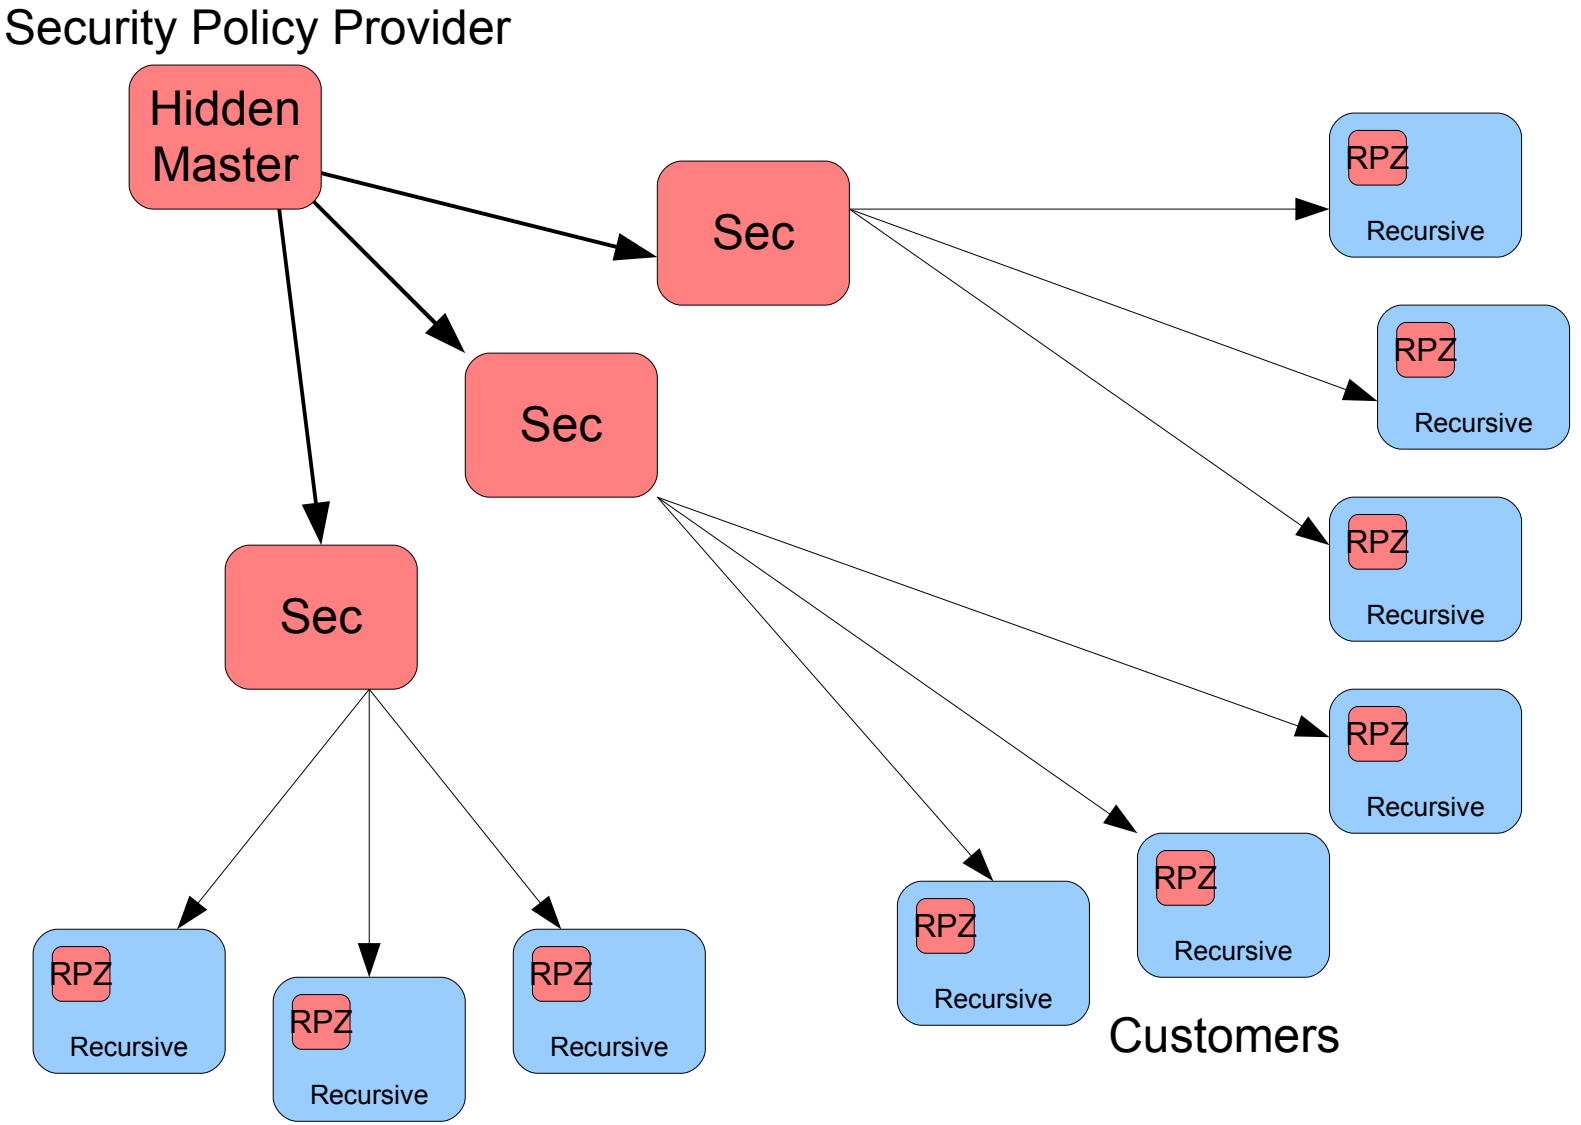

Pause

Response Rate Limiting (RRL)

http://www.redbarn.org/dns/ratelimits

FARSIGHT SECURITY

RRL - Overview

- Info: http://www.redbarn.org/dns/ratelimits
- Motivated by participation of authoritative DNS servers in reflectie DDoS attacks
  - isc.org/ANY & ripe.net/ANY
- RRL Limits the number of ***unique responses*** returned by a DNS server to an IPv4 /24, or IPv6 /48
  - Not just random drops of queries
  - Implemented in NSD, BIND, Knot, PowerDNS, Microsoft, more...

Reflective DDoS

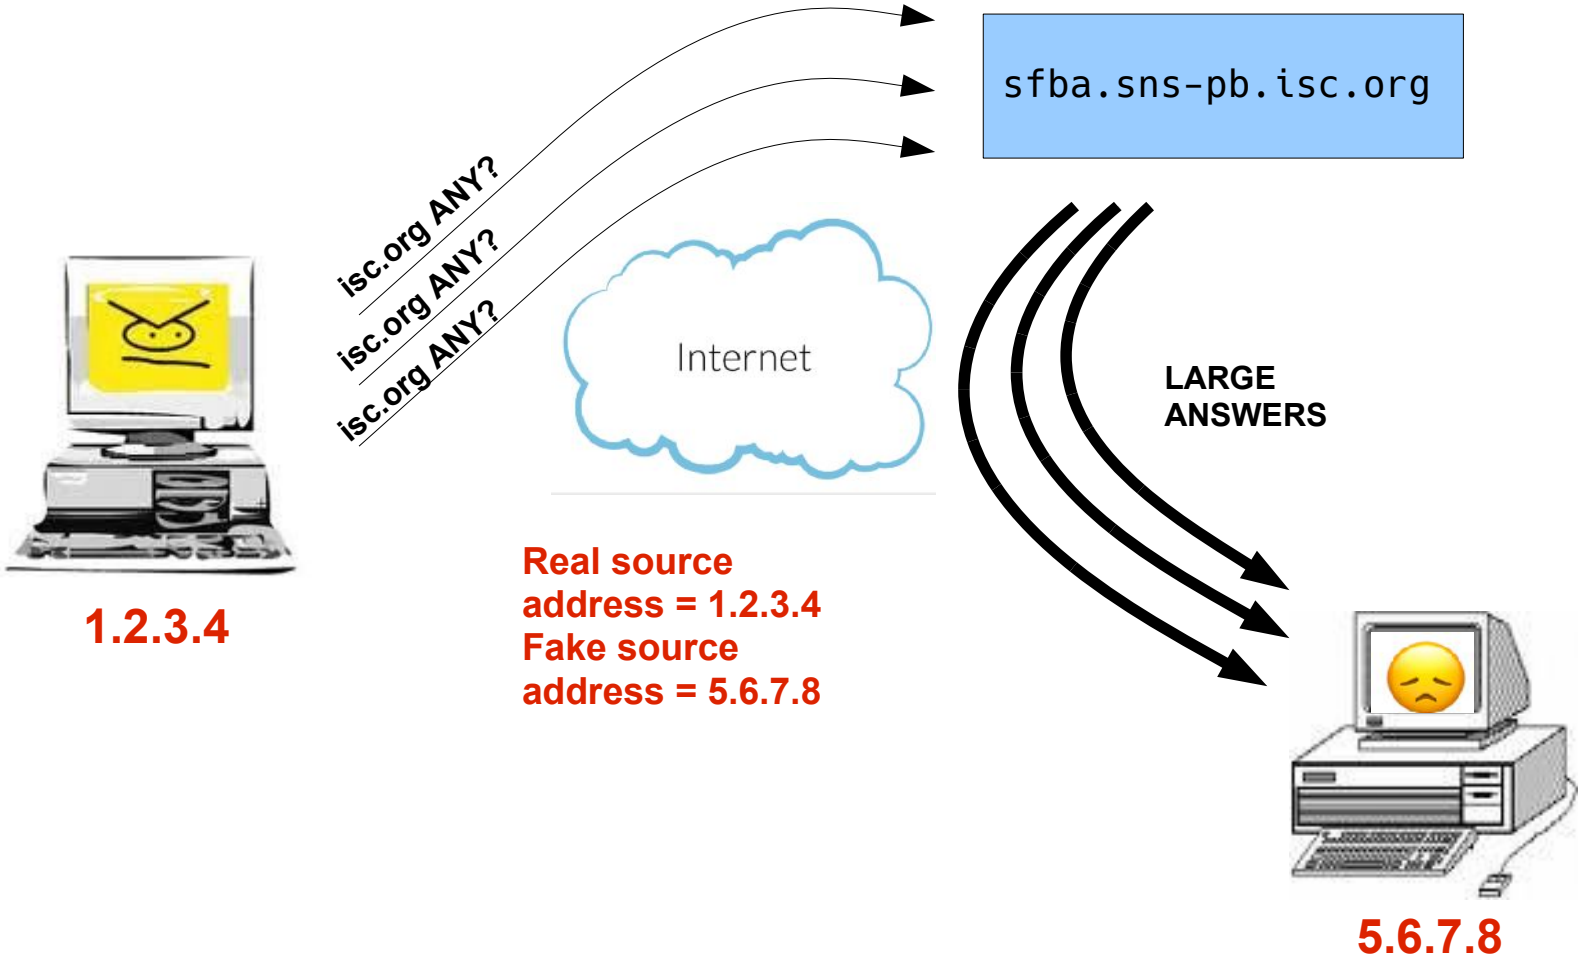

Net effect of RRL

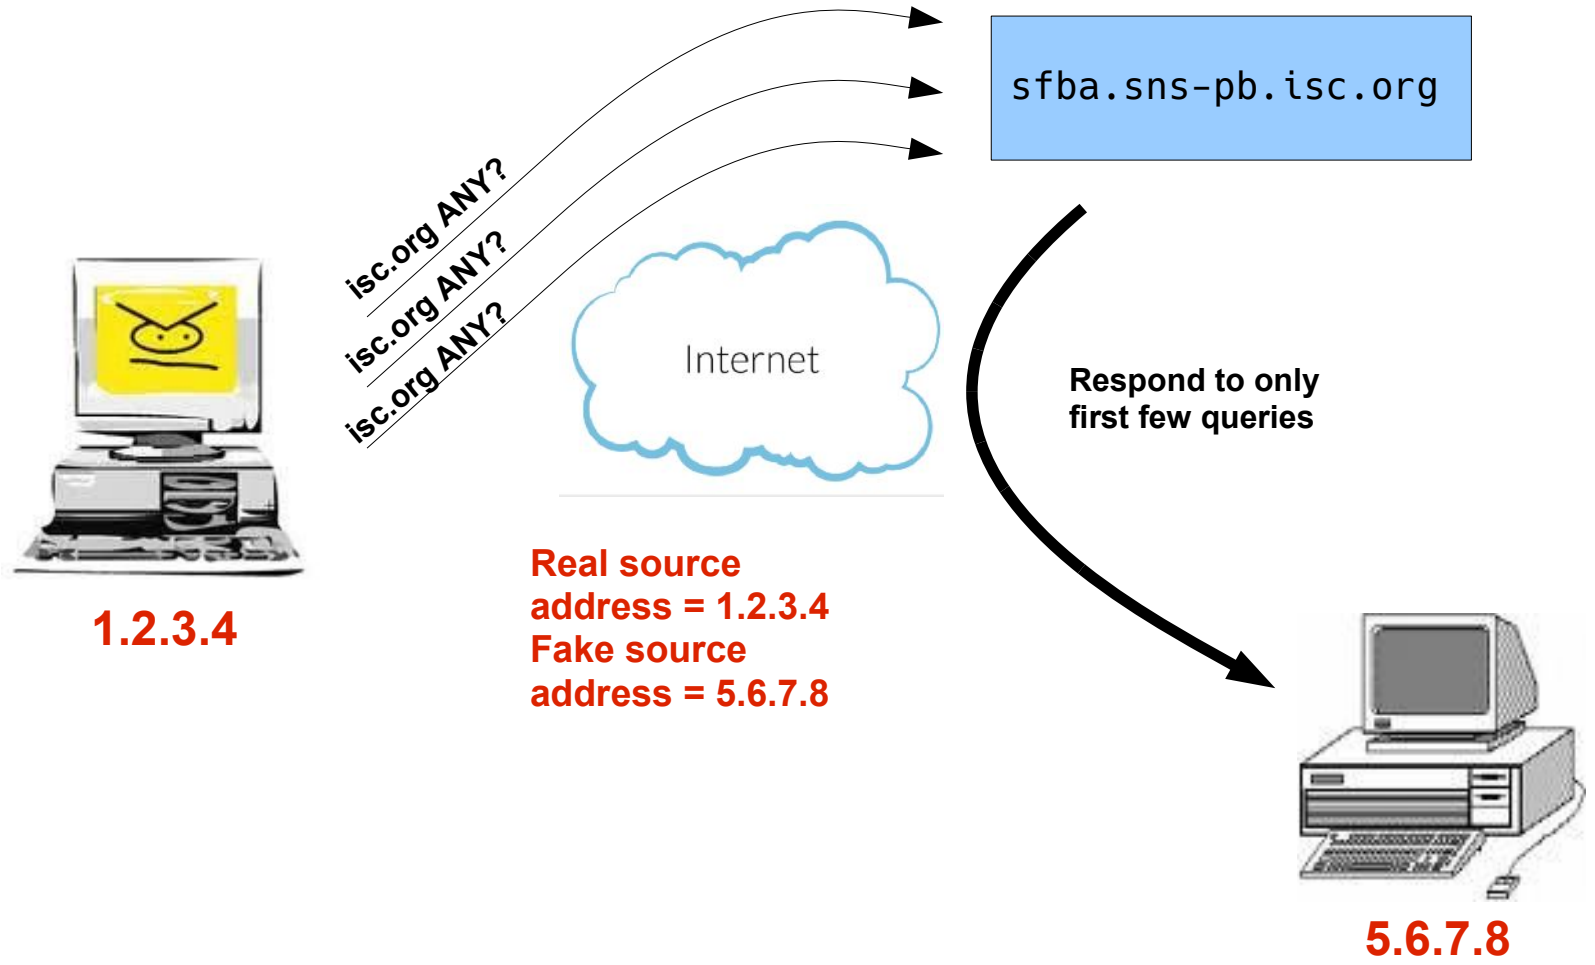

Slide courtesy Eddy Winstead @ ISC (LISA 14)

ISC F-Root

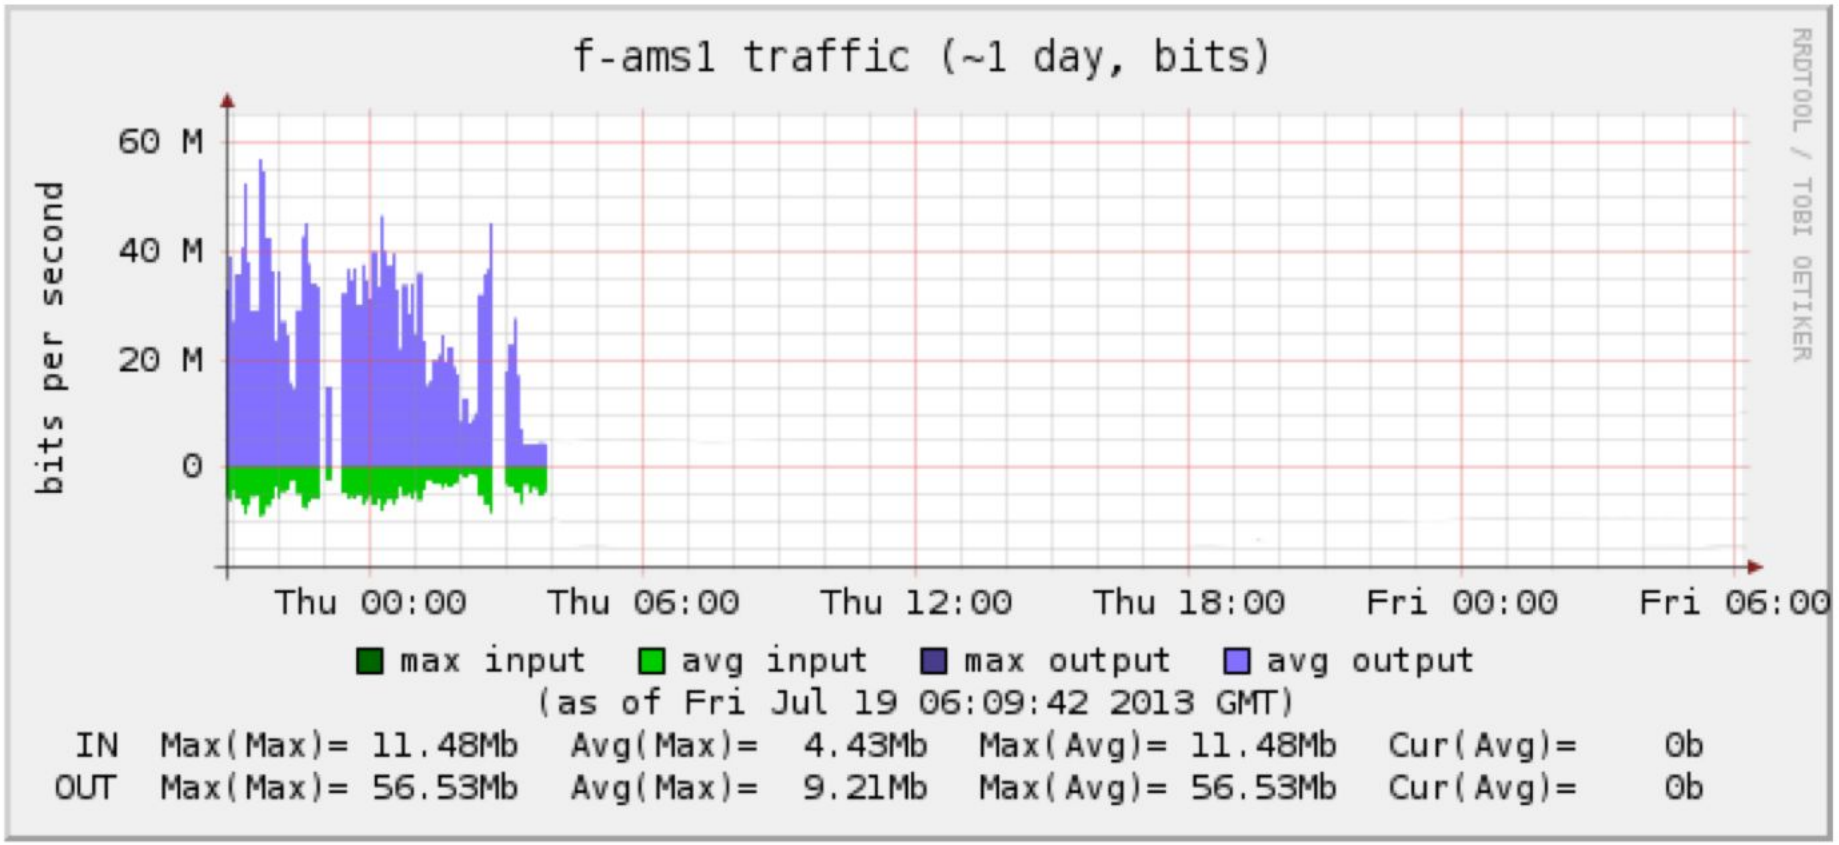

Slide courtesy Eddy Winstead @ ISC (LISA 14)

ISC F-Root

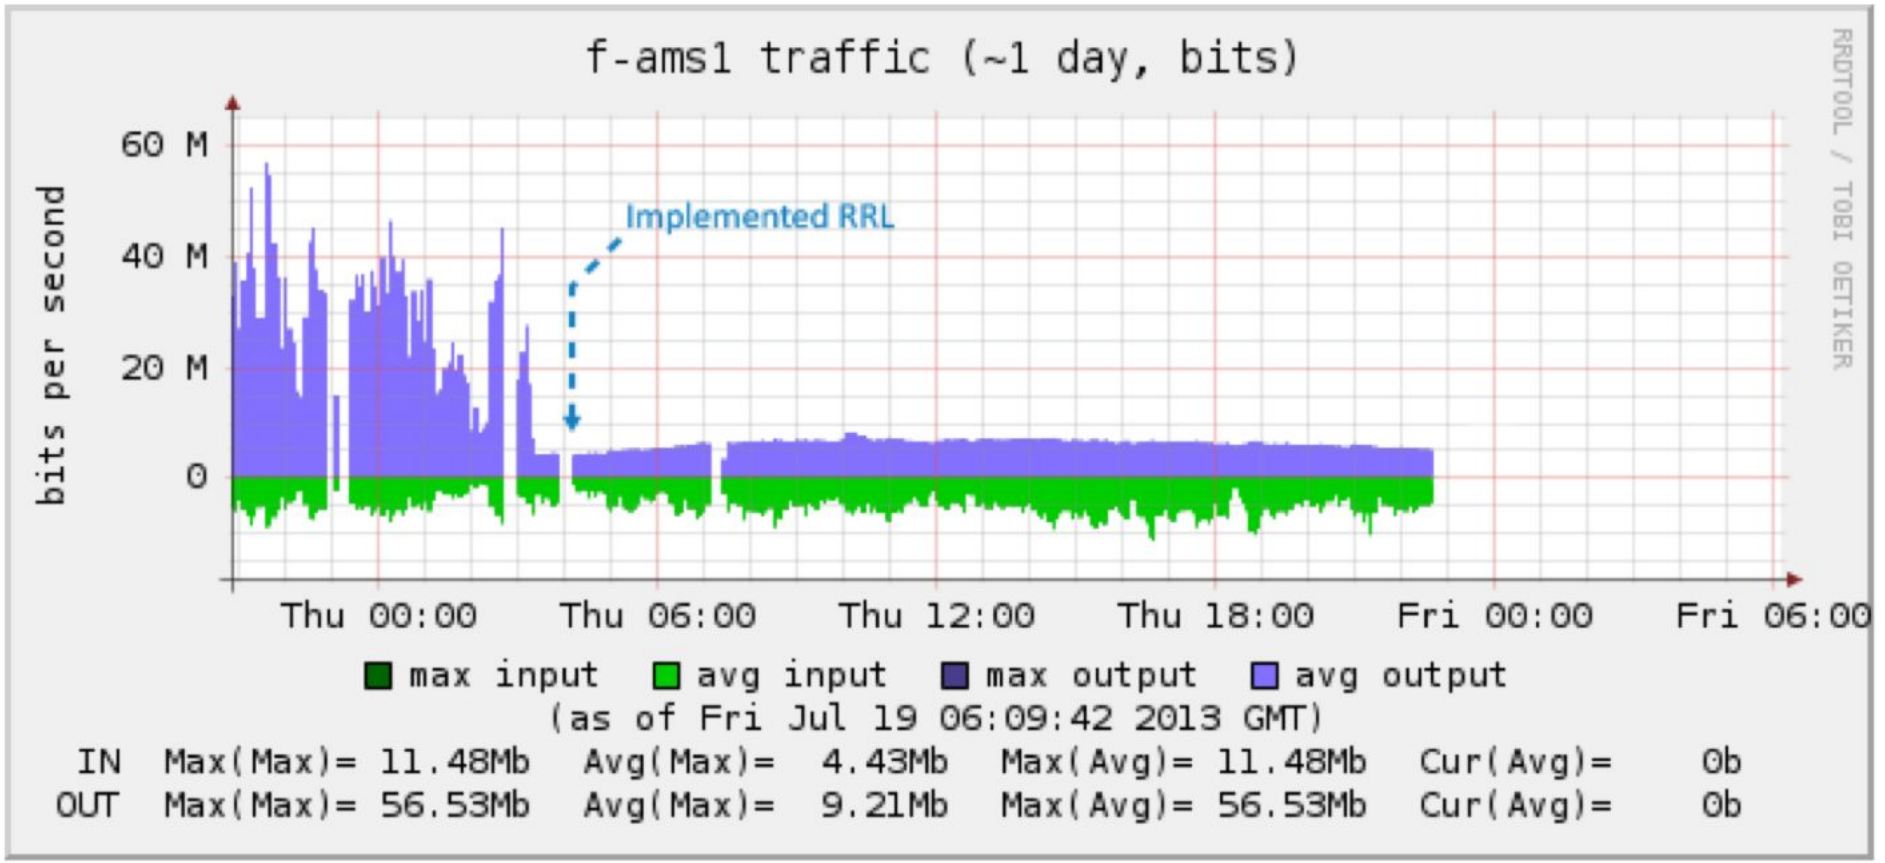

Slide courtesy Eddy Winstead @ ISC (LISA 14)

ISC F-Root

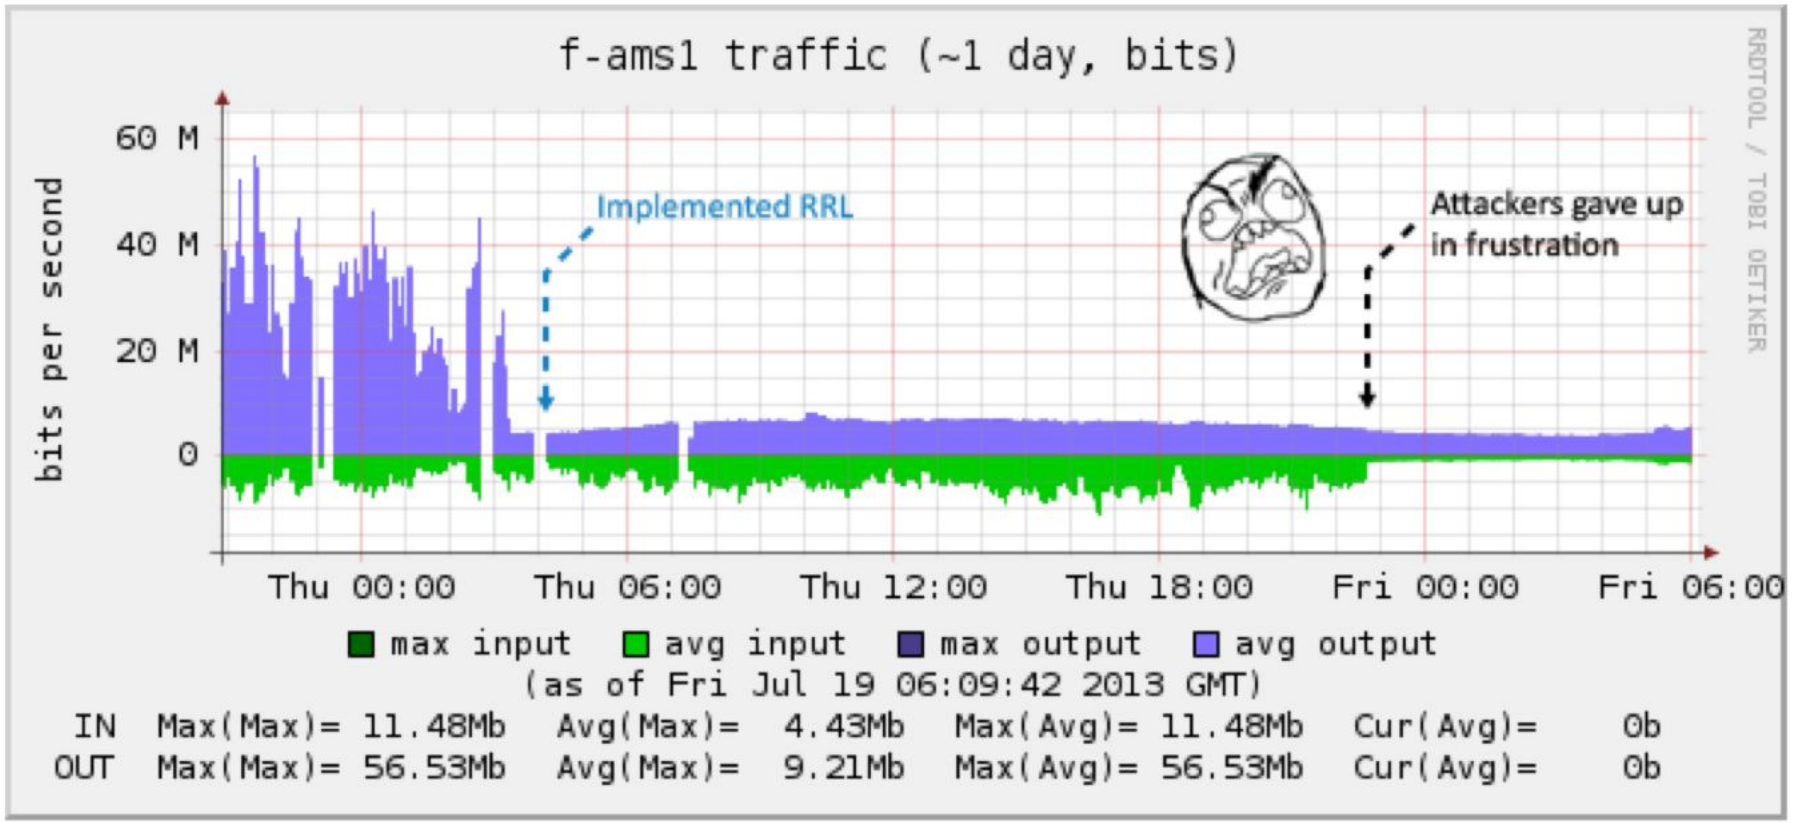

Advantages of RRL

- Improved efficiency
  - Ability to deflect attacks
  - Reduces traffic
- Brand protection
  - Less likely used as part of attack (softer targets)
- Better service
  - Servers less loaded
  - Minimal impact on traffic (compared to filtering)

Common configuration

- Responses per second & window seconds
  - How many identical requests from the same subnet need to be seen before RRL turns on? (for example 15 requests in 5 seconds)
- SLIP or TruncateRate (try “2”)
  - What ratio of responses should be truncated?
  - Common malformed response signals real clients to retry request with TCP to minimize disruption
- Start conservative & authoritative only

Recursive Rate Limiting (the other RRL)

FARSIGHT SECURITY

Recursive rate limiting

- Response Rate Limiting designed for authoritative servers
- With Mirai/IoT, NTP, Chargen, other DdoS methods, Open Recursive Servers, DNS is not at forefront, but is still used in attacks.
- RRL alone is not enough; but it's still a good idea
- DNS servers have context that IP filters won't understand
- Investigate recursive server rate limiting. For example:
  - BIND (fetches-per-*)
  - Unbound (ratelimit-*)
- PRSD attack

Slide courtesy Eddy Winstead @ ISC (LISA 14)

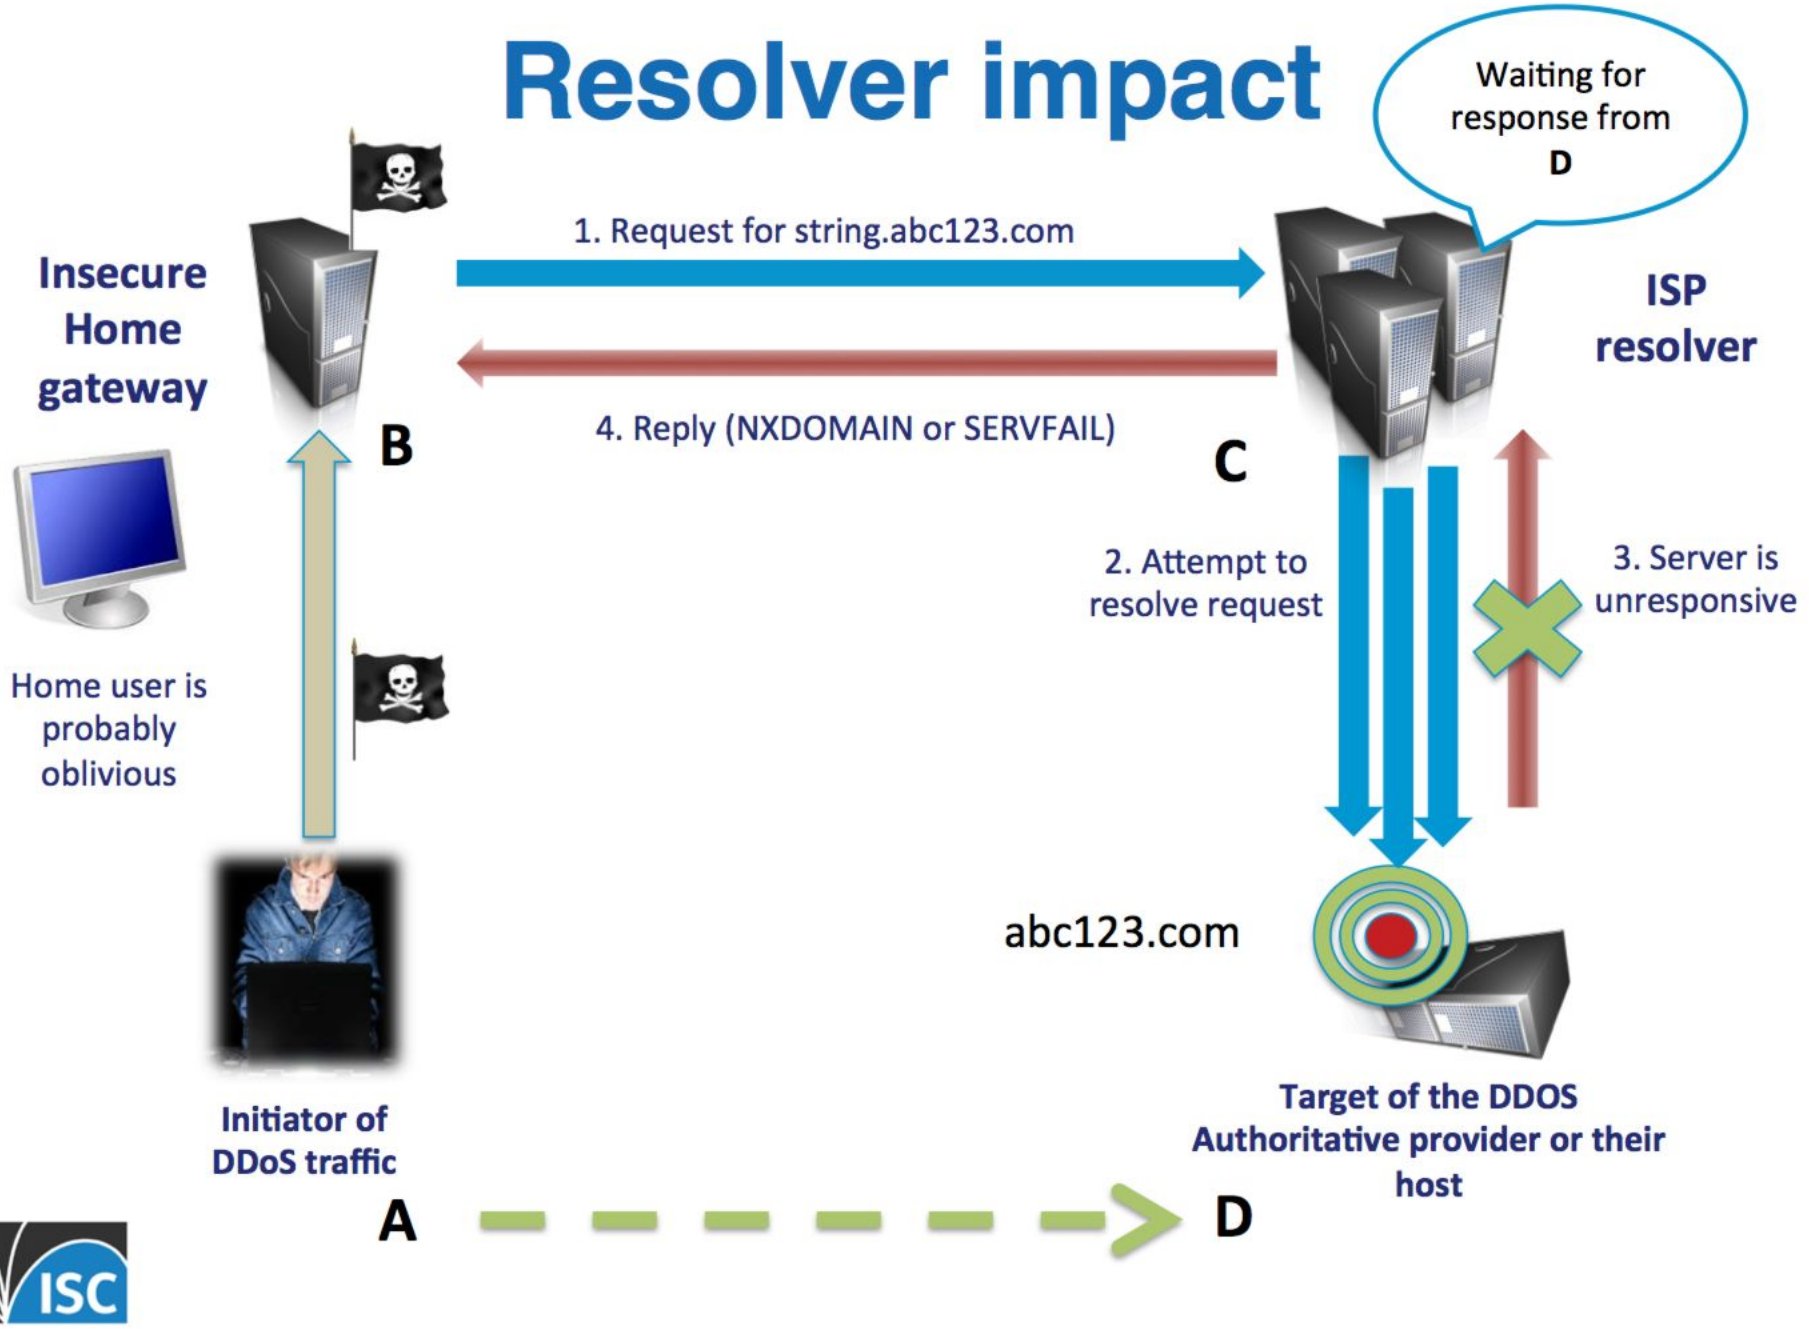

Pause

DNSTAP

FARSIGHT SECURITY

DNSTAP

- http://dnstap.info/
- Built/designed by Robert Edmonds
- Several slides courtesy of Jeroen Massar (APWG 14)

FARSIGHT SECURITY

DNS lookups

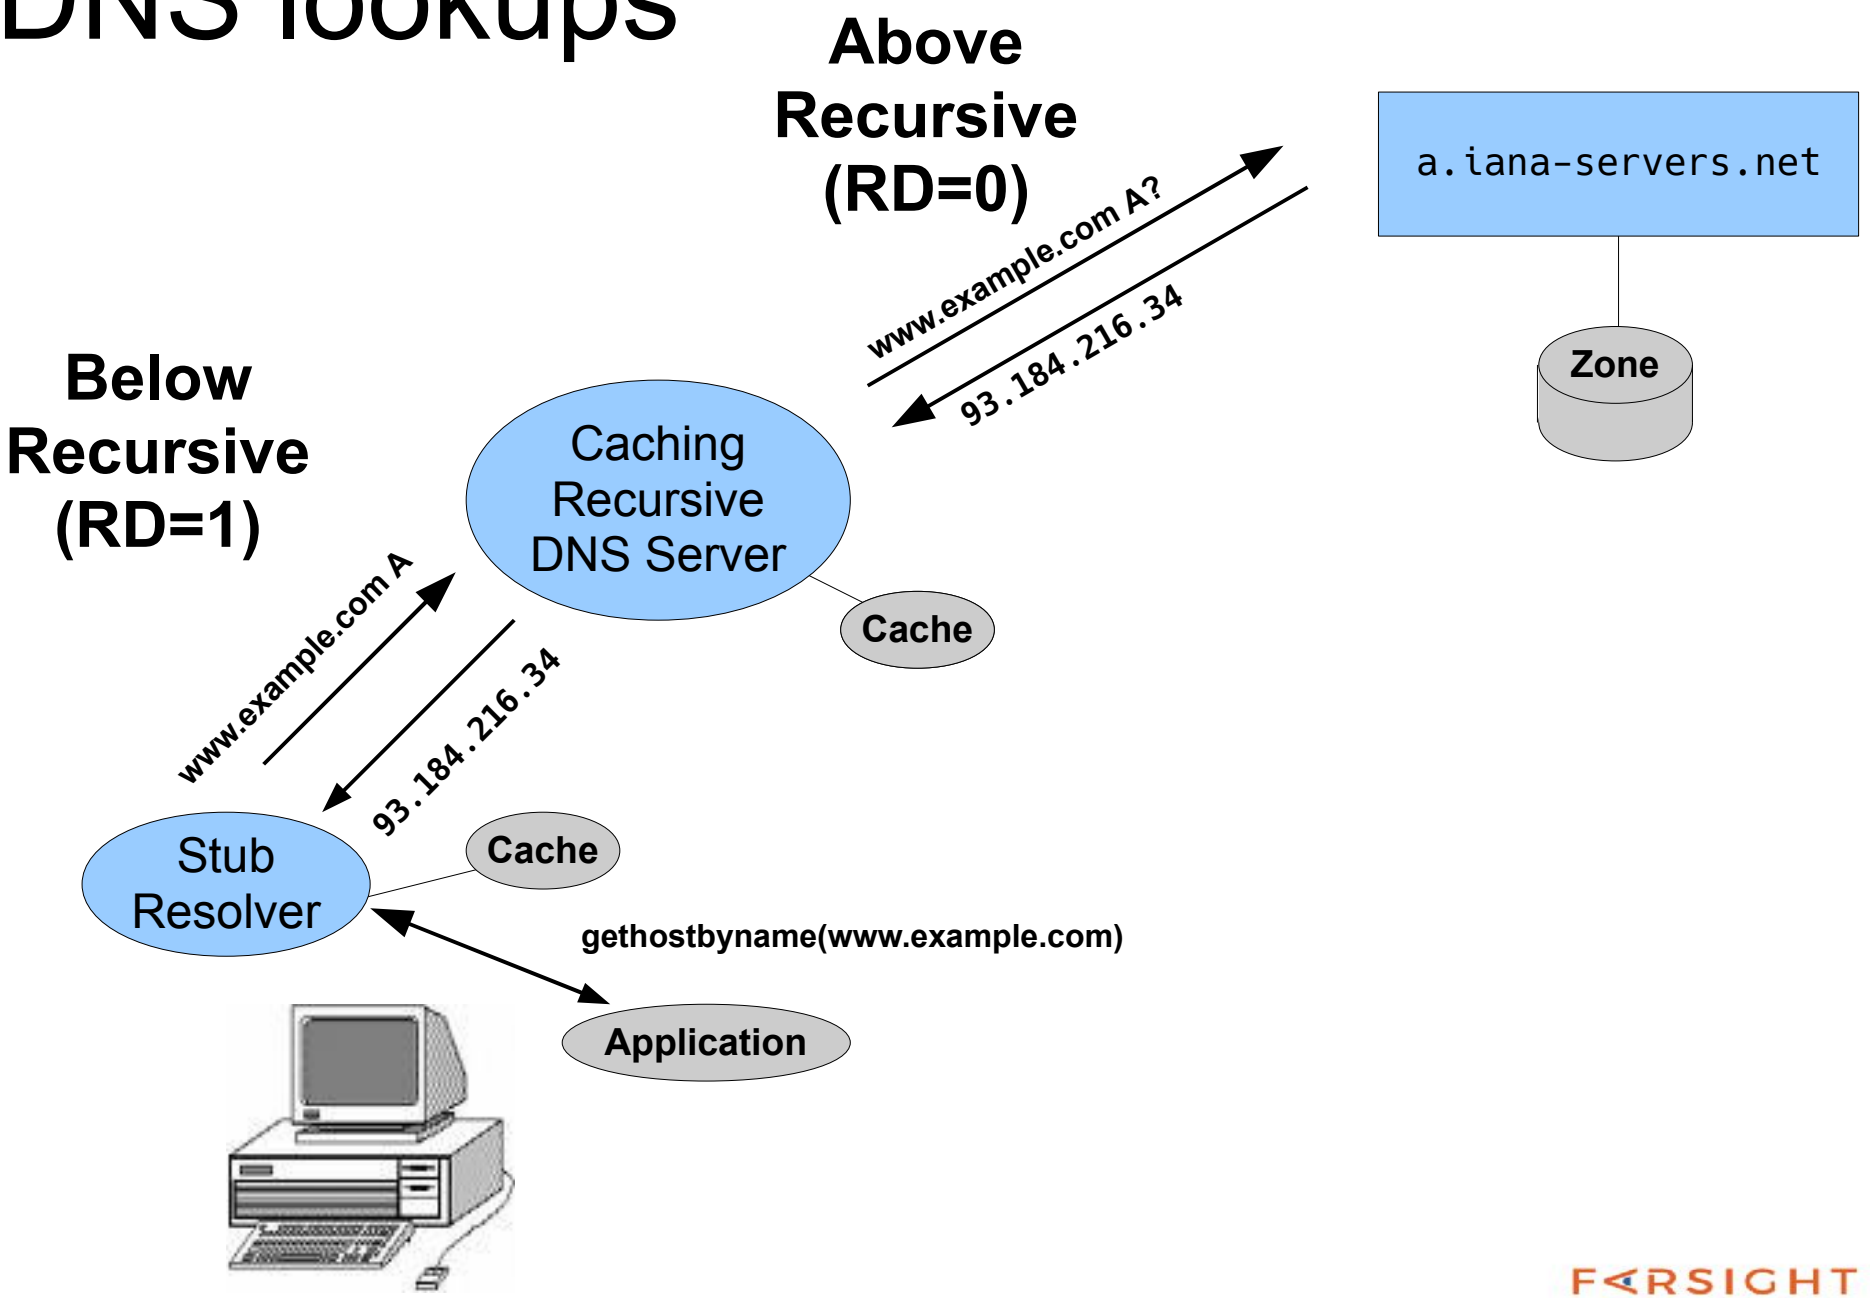

FARSIGHT SECURITY

Simplified view

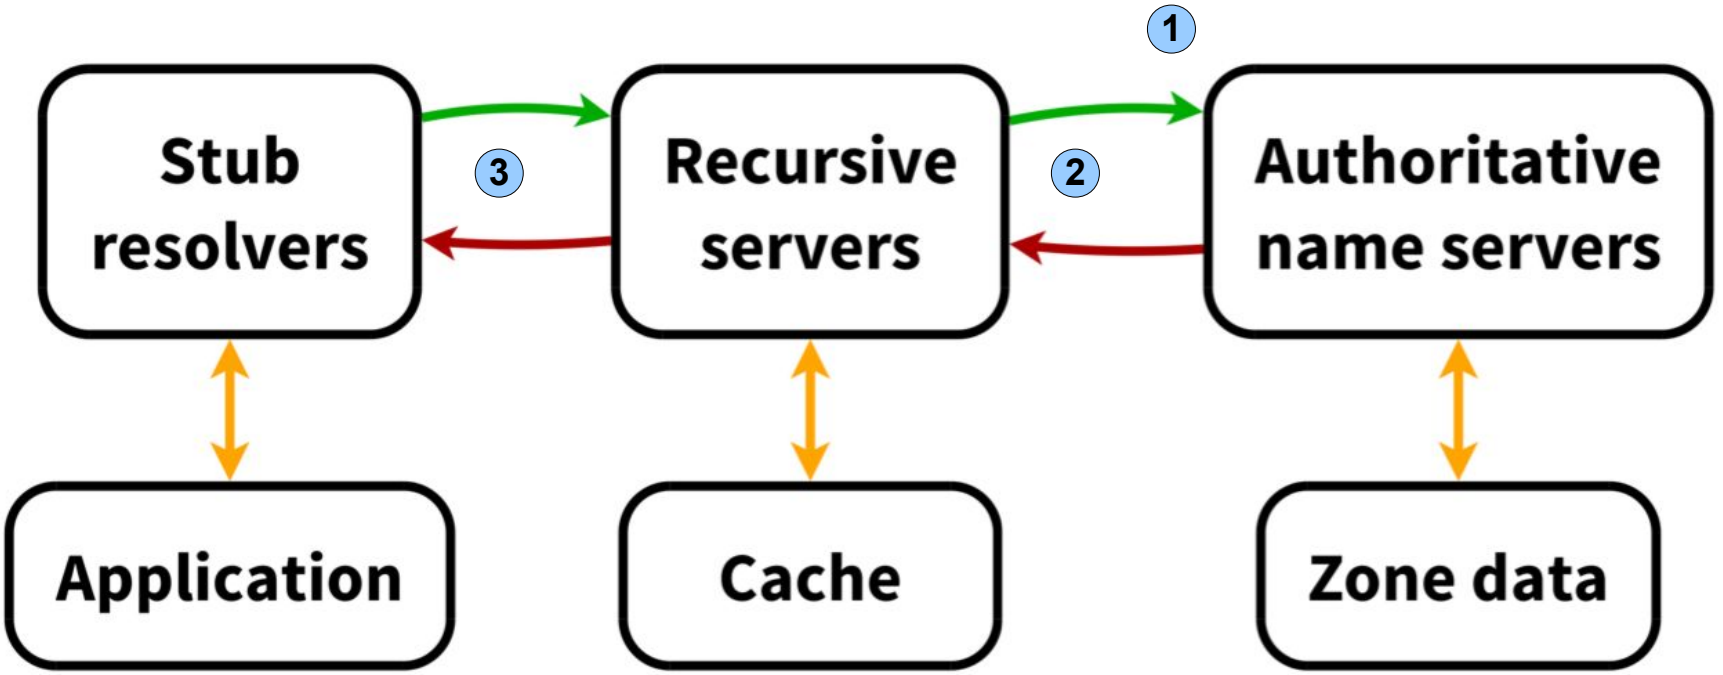

1. Query logging (eg: DSC)
2. PassiveDNS Replication
3. Query logging (eg: IDS)

FARSIGHT SECURITY

Logging methods

- Auth queries
  - Wire: DSC, dnscap
  - Server: query logging (inefficient)
- PassiveDNS
  - tcpdump, dnscap, nmsgtool
  - Issues: no TCP, hardening, bailiwick reconstruction
- Client query logging
  - Server: query logging (inefficient)
  - Network: tcpdump, IDS (some TCP)

DNSTAP monitoring types

"SQ" STUB_QUERY

"SR" STUB_RESPONSE

"CQ" CLIENT_QUERY

"CR" CLIENT_RESPONSE

"RQ" RESOLVER_QUERY

"RR" RESOLVER_RESPONSE

"AQ" AUTH_QUERY

"AR" AUTH_RESPONSE

"FQ"/"FD" FORWARDER_QUERY/RESPONSE (same as RQ/RR, but RD=1)

"TQ/"TD" TOOL QUERY/RESPONSE

FARSIGHT SECURITY

DNSTAP types

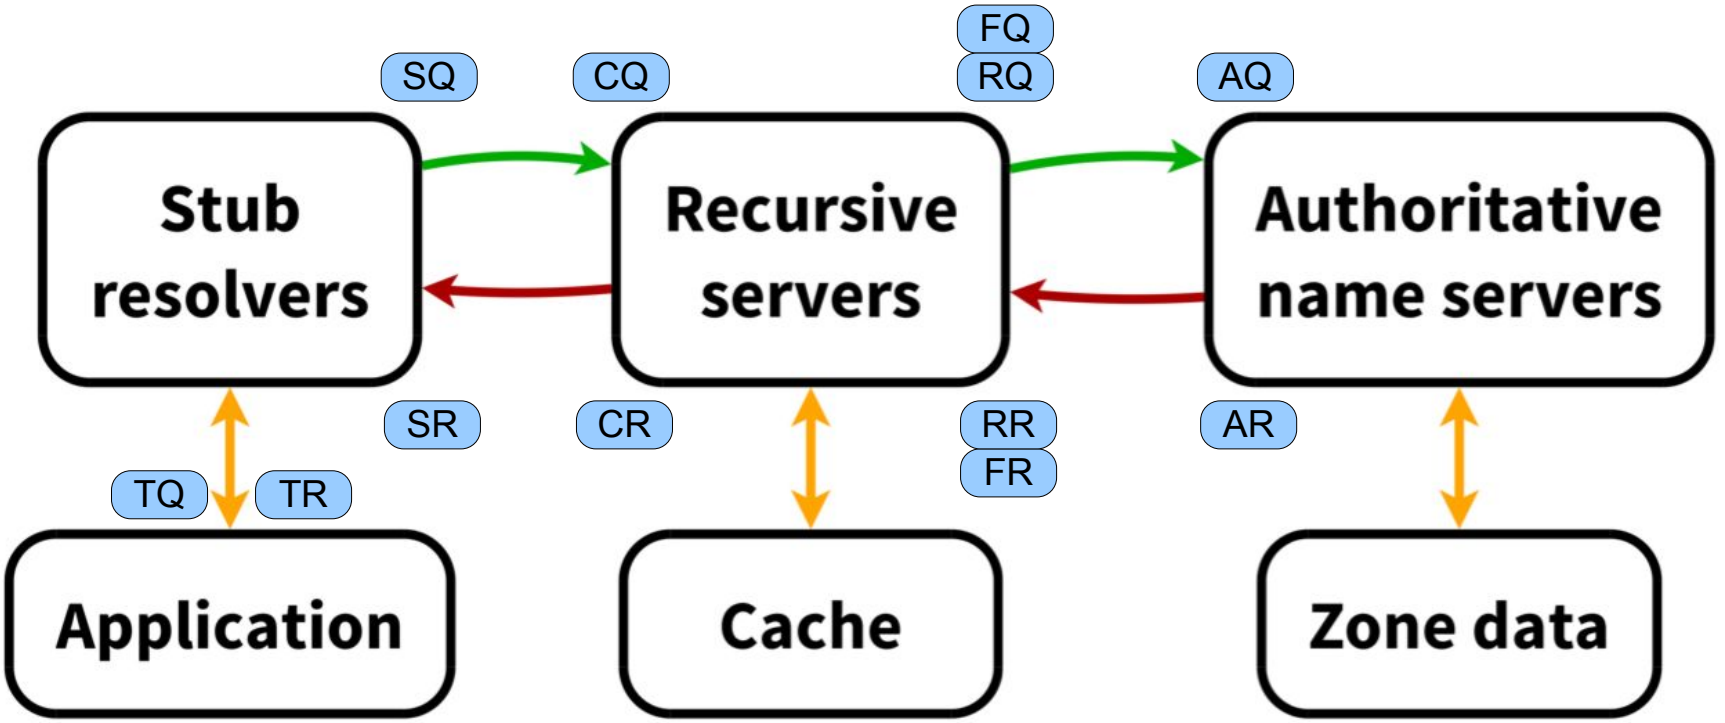

DNSTAP non-blocking

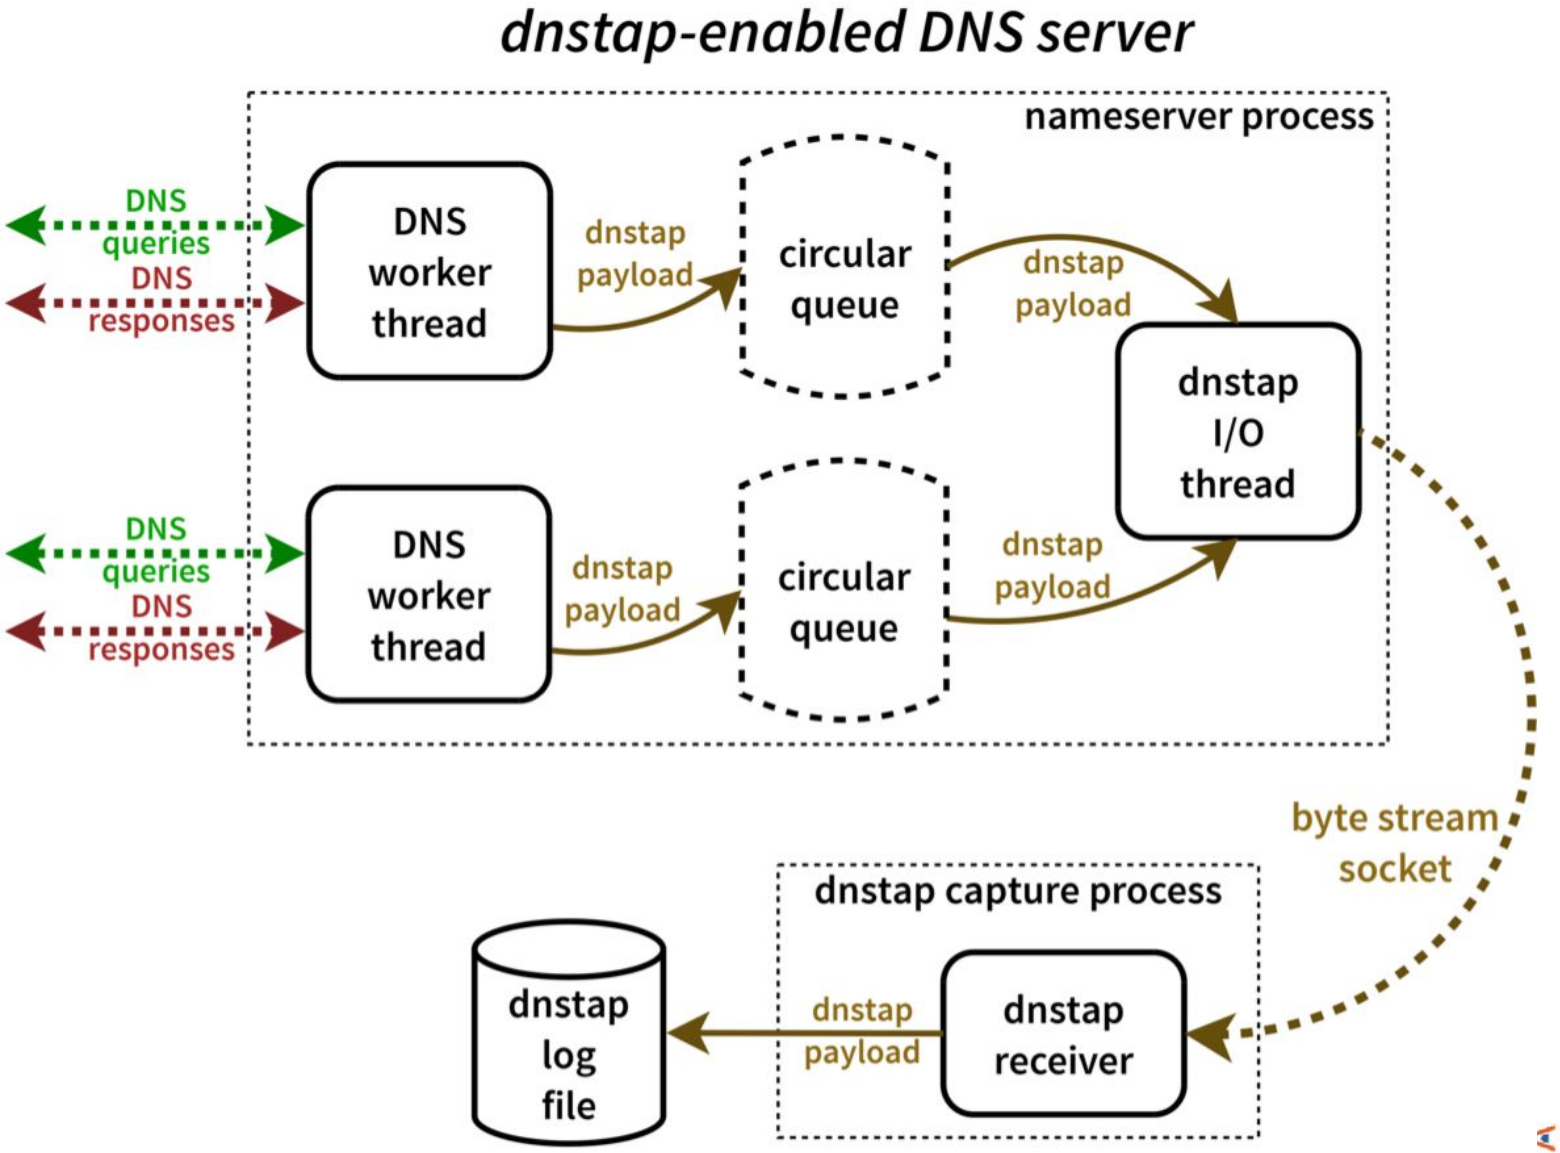

DNSTAP architecture

- Supported by most major nameservers:
  - BIND, Unbound, Knot, Akamai
- nameserver writes to Unix socket
- fstrm reads from socket, dumps to file

```
fstrm_capture -u /var/run/unbound/dnstap.sock \
    -s 60 --gmtime -t protobuf:dnstap.Dnstap \
    -w /DIR/FILE.%Y%m%d-%H%M%S.dnstap
```

  - other options in future (dnstap-nmsg)
- dnstap-read (BIND) reads from file
- Google Protocol Buffers binary format

Reading DNSTAP data

```
# dnstap-read dump.20170411-174346.dnstap
11-Apr-2017 13:43:45.863 RR 199.7.83.42 UDP 866b www.akamai.com/IN/A
11-Apr-2017 13:43:45.911 RR 2001:7fd::1 UDP 852b f.gtld-servers.net/IN/AAAA
11-Apr-2017 13:43:45.917 RR 192.31.80.30 UDP 789b f.gtld-servers.net/IN/AAAA
11-Apr-2017 13:43:45.917 RR 192.41.162.30 UDP 789b m.gtld-servers.net/IN/AAAA
11-Apr-2017 13:43:45.917 RR 192.41.162.30 UDP 789b d.gtld-servers.net/IN/AAAA
11-Apr-2017 13:43:45.924 RR 192.31.80.30 UDP 771b av2.nstld.com/IN/AAAA
11-Apr-2017 13:43:45.924 RR 192.26.92.30 UDP 771b av3.nstld.com/IN/A
11-Apr-2017 13:43:45.924 RR 192.31.80.30 UDP 771b av1.nstld.com/IN/AAAA
11-Apr-2017 13:43:45.924 RR 192.41.162.30 UDP 771b av4.nstld.com/IN/AAAA
11-Apr-2017 13:43:45.928 RR 192.228.79.201 UDP 852b h.gtld-servers.net/IN/AAAA
11-Apr-2017 13:43:45.931 RR 192.82.134.30 UDP 286b av3.nstld.com/IN/A
11-Apr-2017 13:43:45.931 RR 192.82.134.30 UDP 286b av1.nstld.com/IN/AAAA
```

FARSIGHT SECURITY

Performance

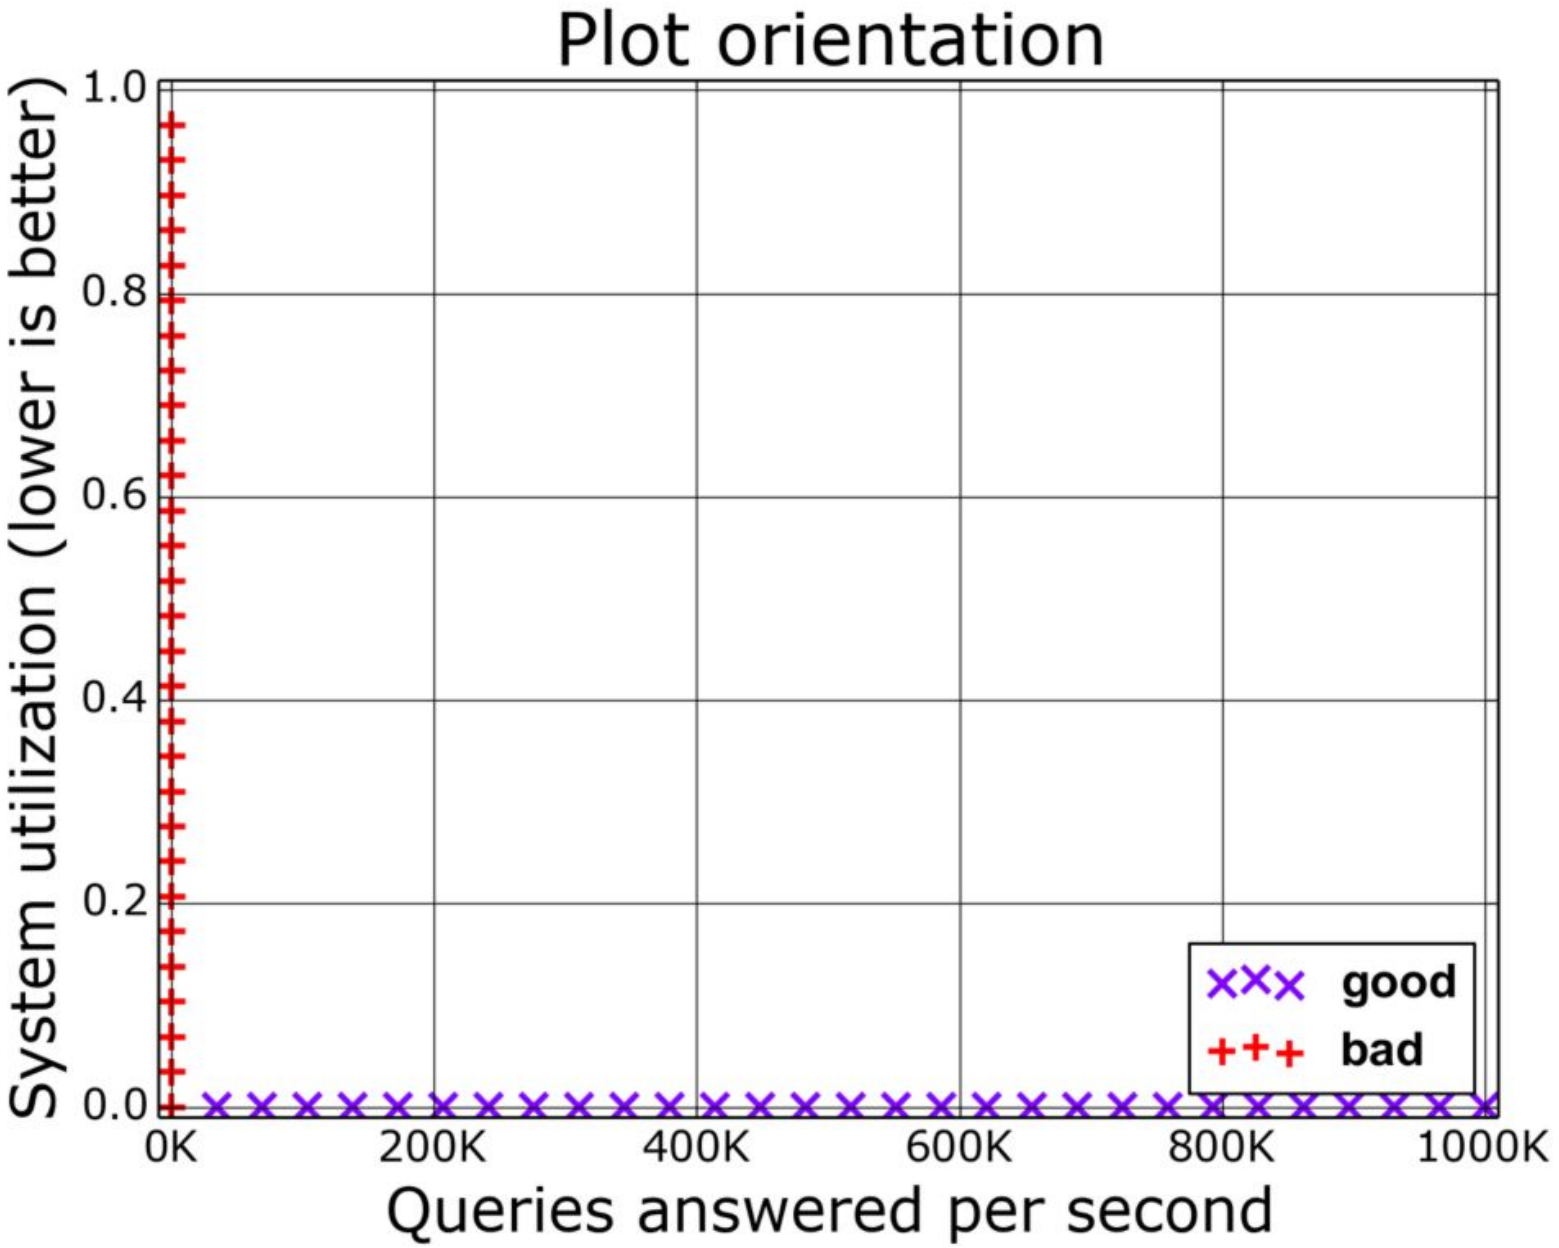

Performance

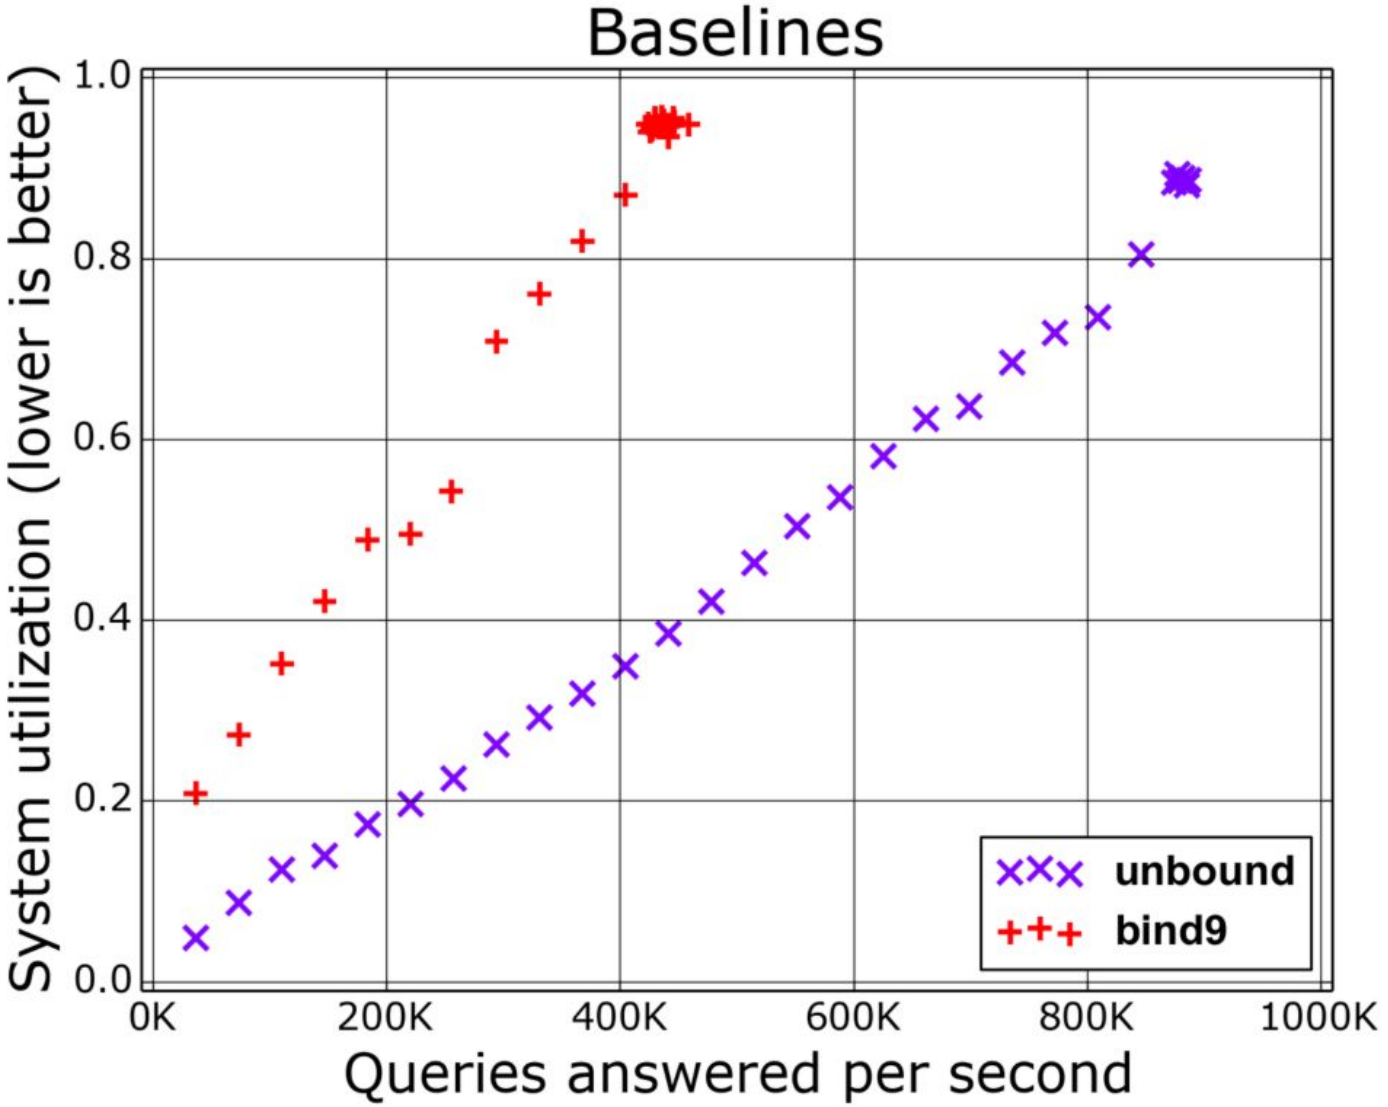

FARSIGHT SECURITY

Performance

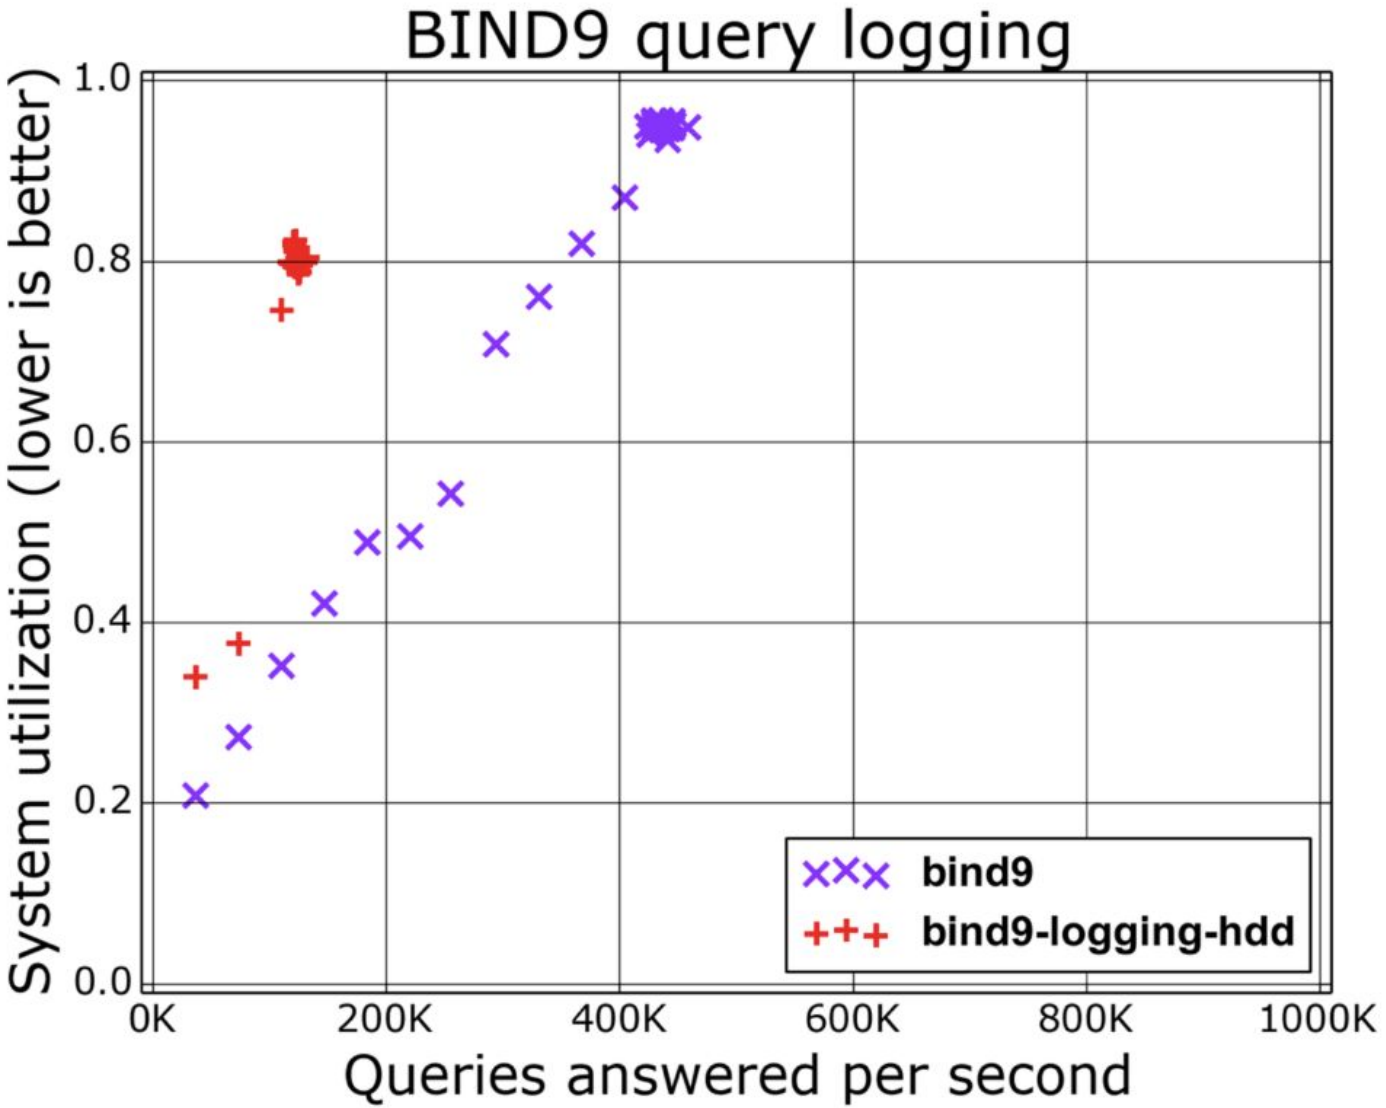

FARSIGHT SECURITY

Performance

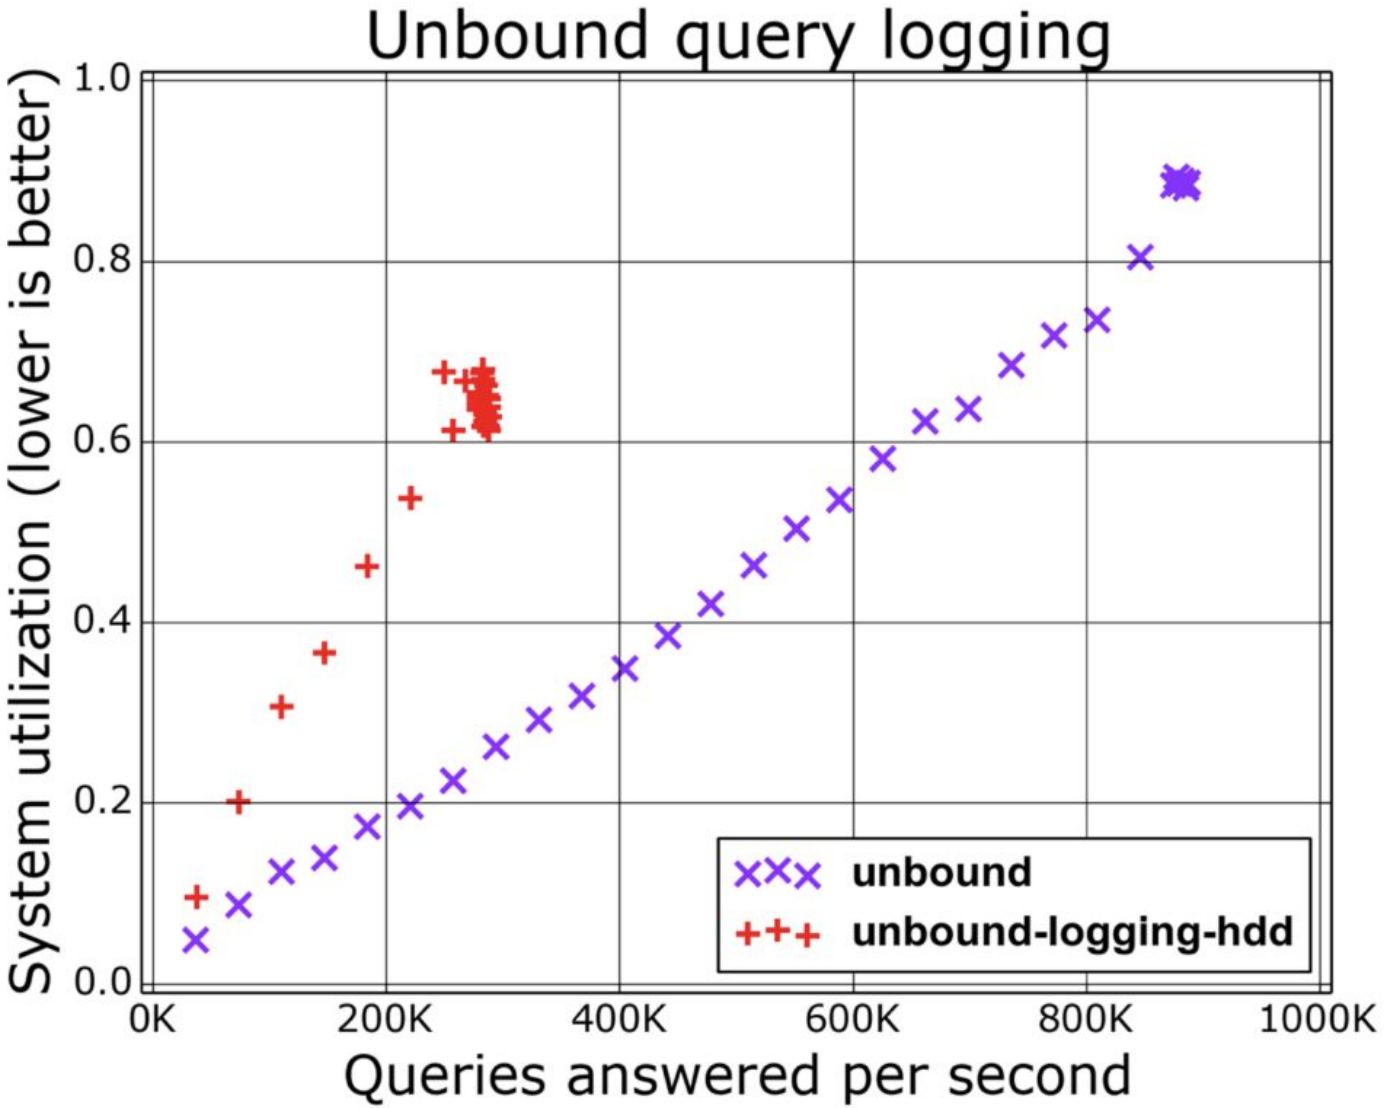

FARSIGHT SECURITY

Performance

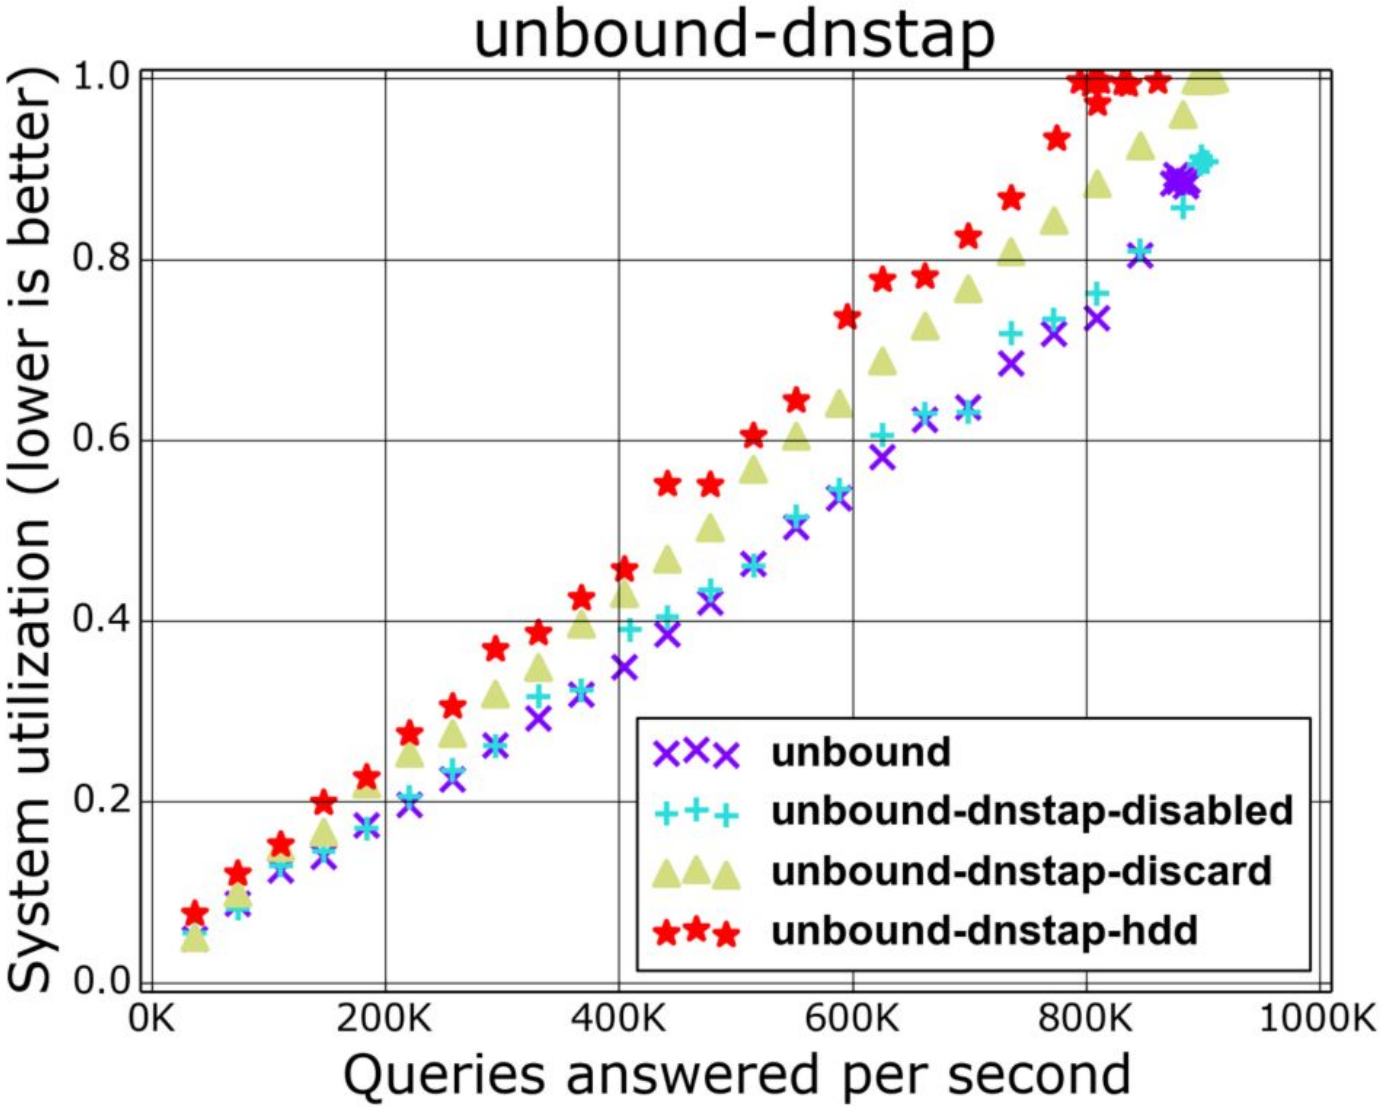

FARSIGHT SECURITY

Performance update

- two-thread recursive nameserver under PRSD attack: 200% system CPU
- tcpdump 10%+ of system CPU
- dnstap <1% of system CPU
- Look to DNS-OARC 26 for recent performance comparison.

DPRIVE

- Encrypted TLS/tcp client for DNS queries between stub resolver and recursive servers
- https://datatracker.ietf.org/doc/rfc7858/
- Already implemented – moving monitoring to clients and nameservers
- IDS vendors may need to adapt

FARSIGHT SECURITY