On Eliminating Root Nameservers from the DNS

Mark Allman

International Computer Science Institute

ACM SIGCOMM HotNets
November, 2019

"Any intelligent fool can make things bigger, more complex, and more violent. It takes a touch of genius and a lot of courage to move in the opposite direction."
—E. F. Schumacher

DNS Overview

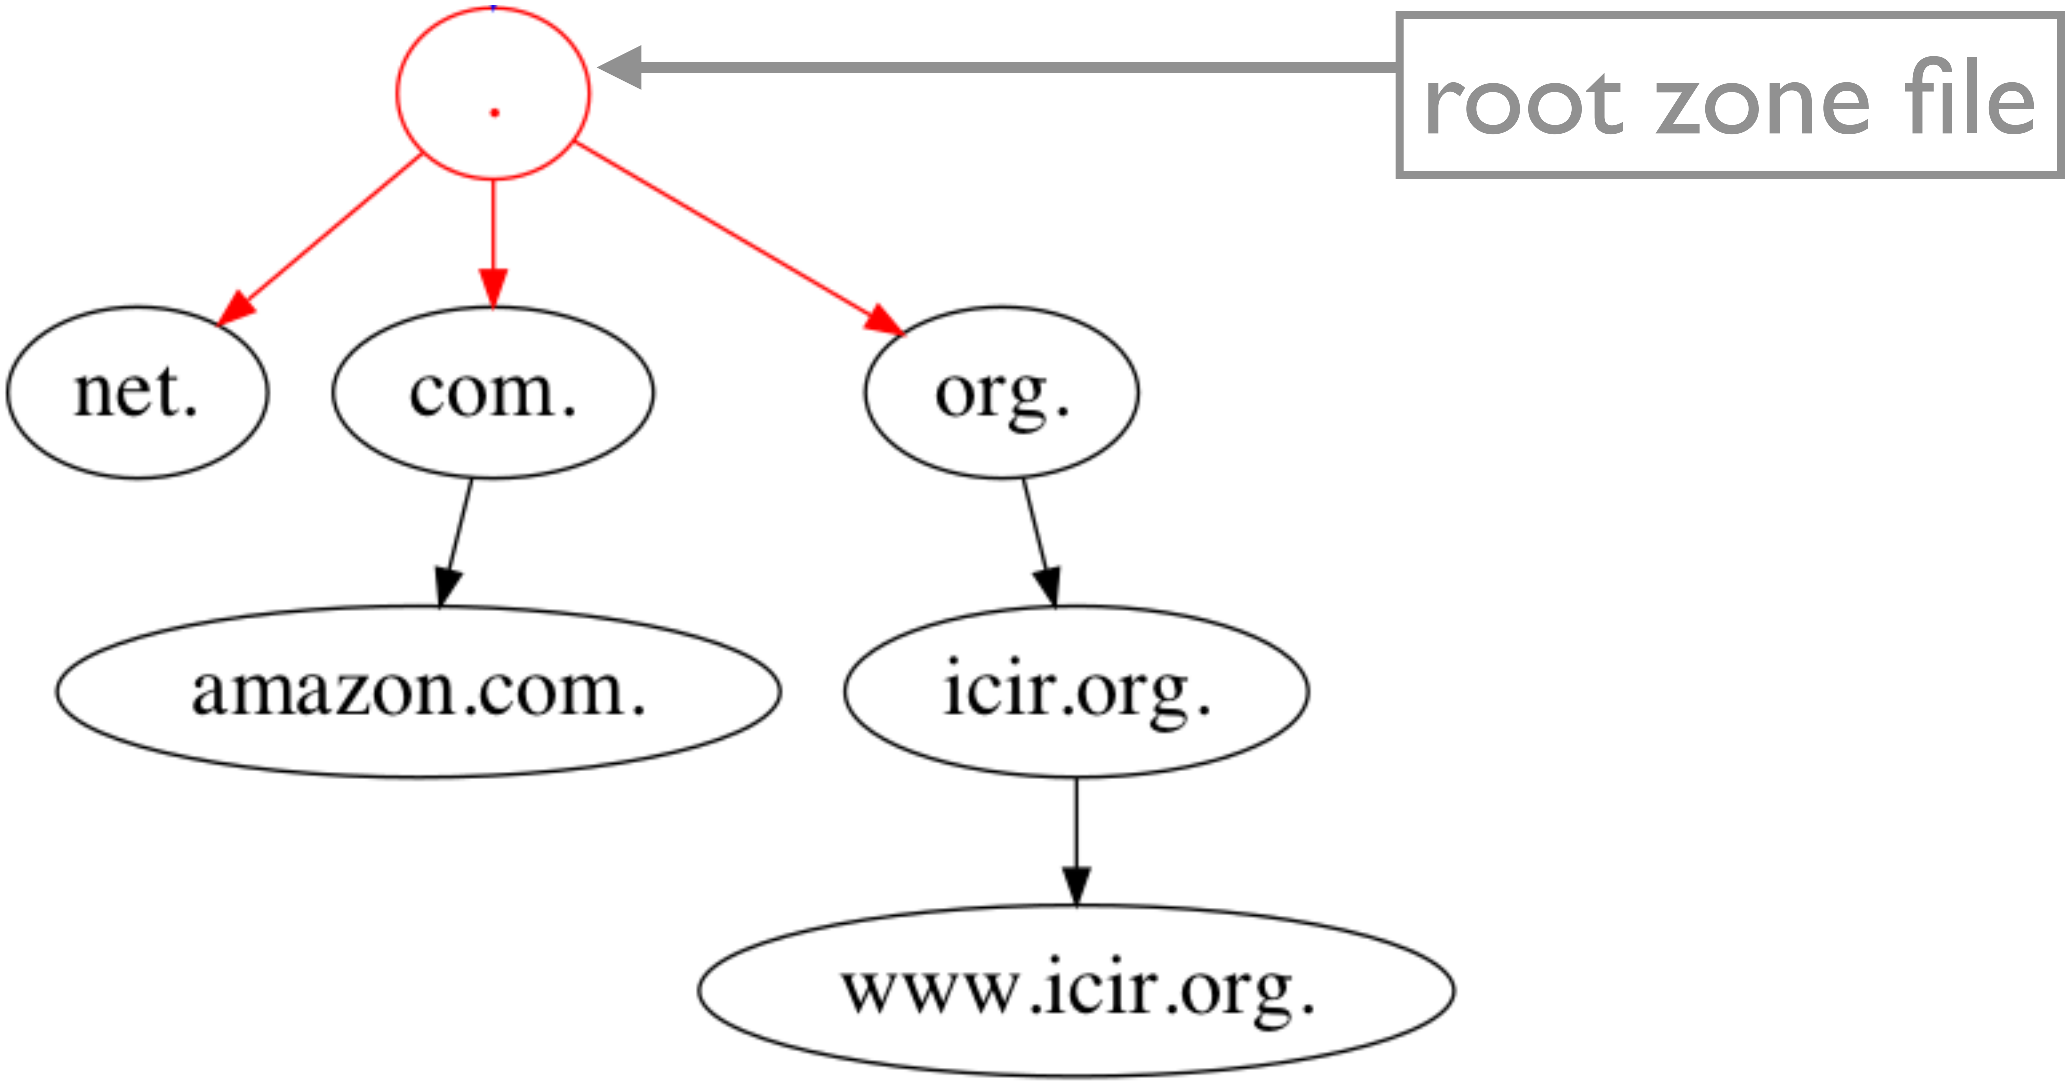

Replicating the Roots

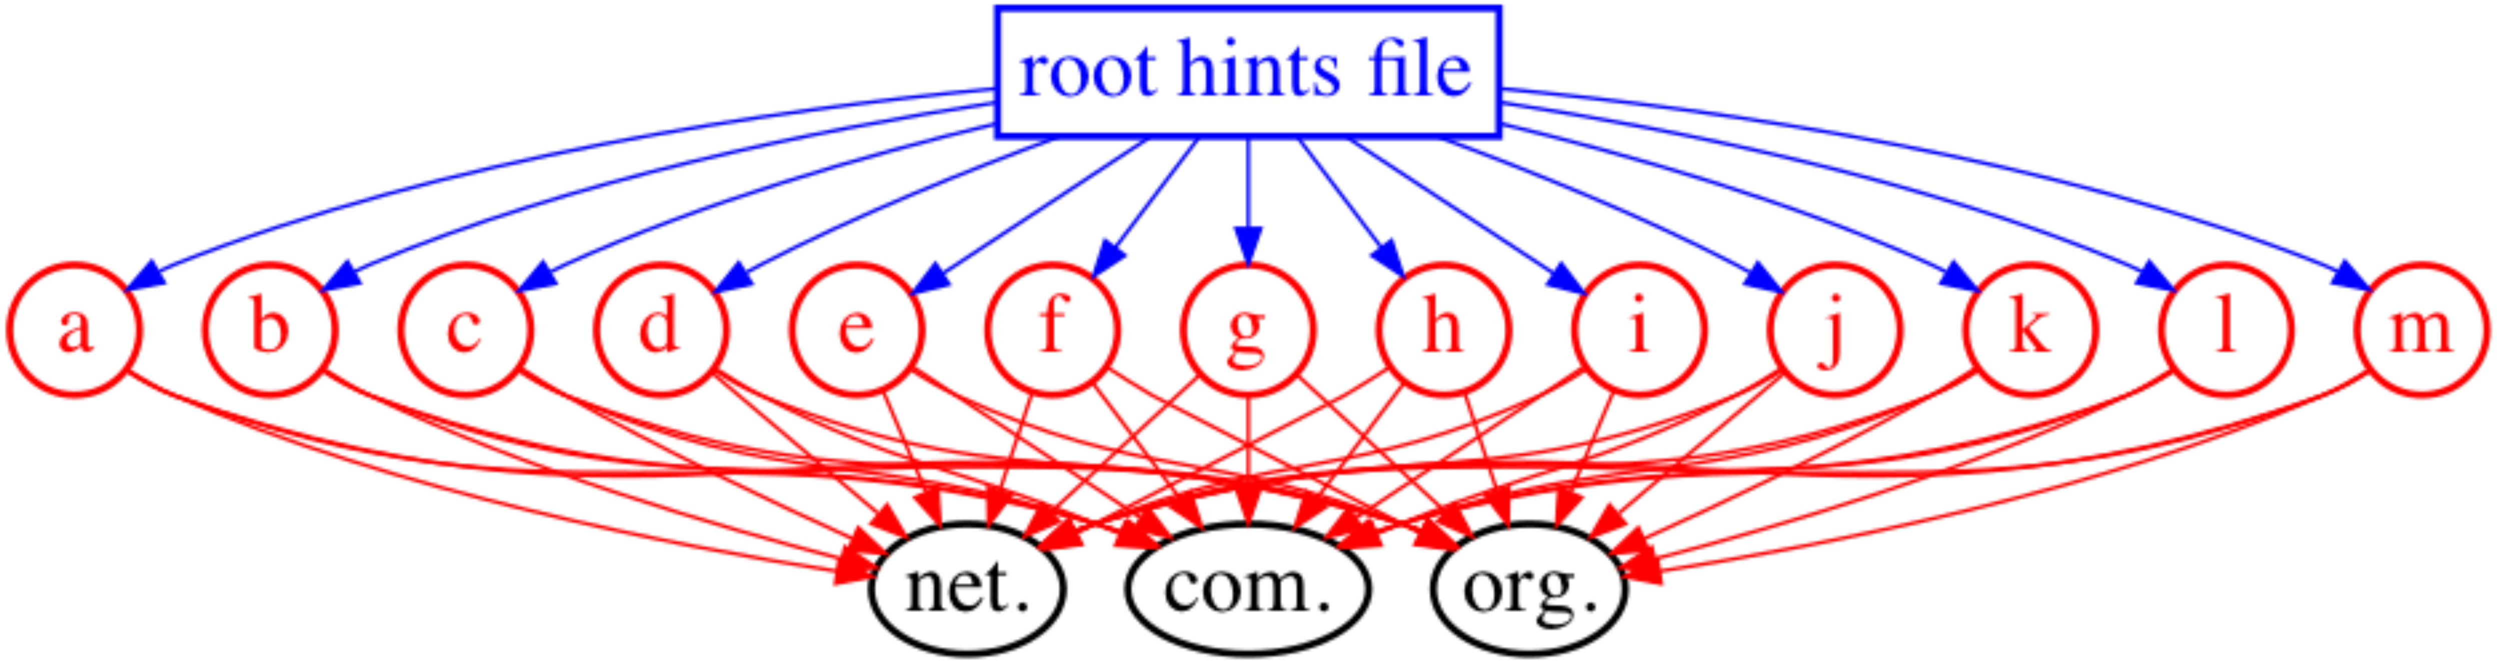

Problem: immense load

Root Server Load

	Queries/Day	Queries/Second
Total	107 B	1.2 M
Per Root	8.9 B	103 K

Data from 12 of the 13 roots on May 15, 2019 (missing g-root).

Courtesy of root-servers.org & root operators

More Root Replication

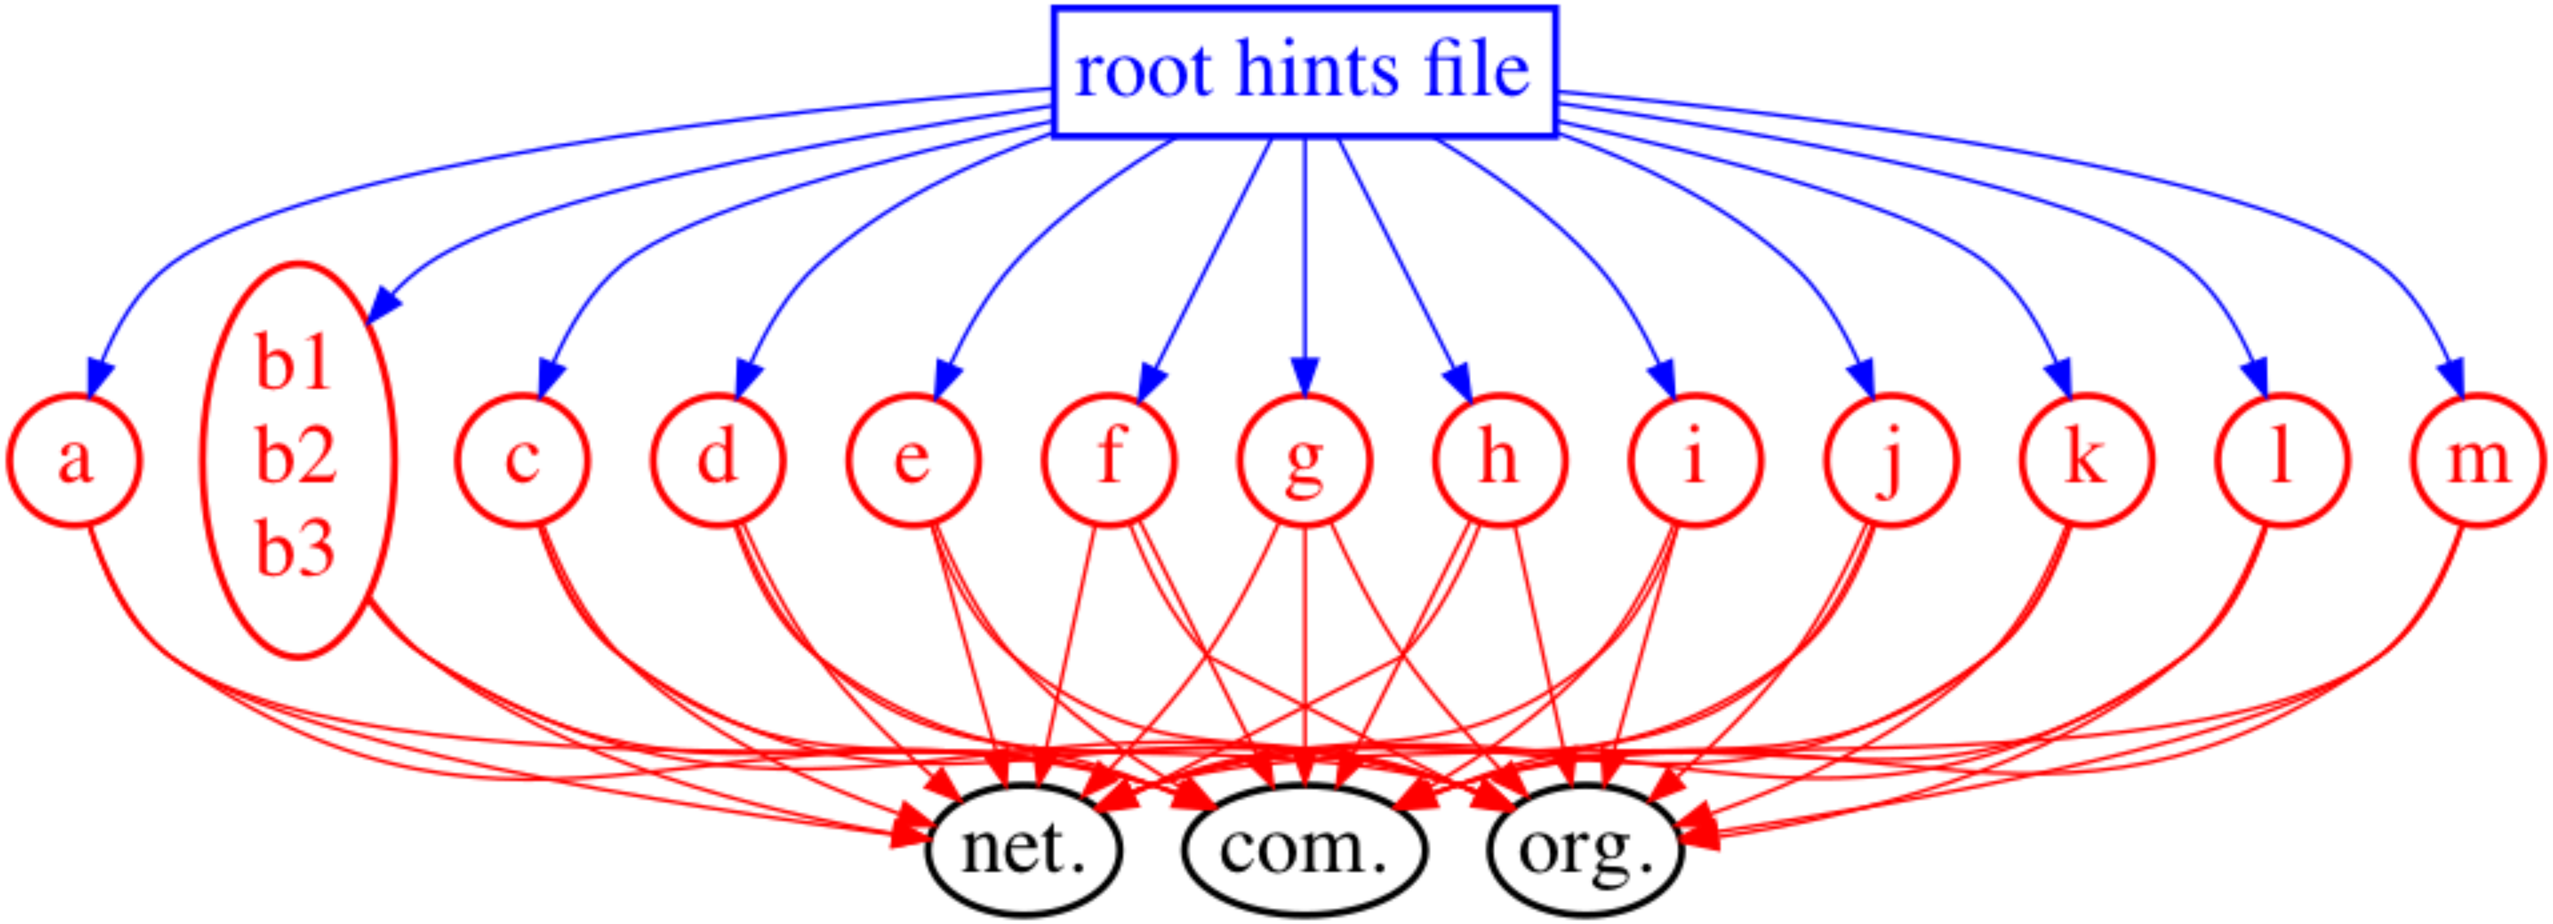

- b1, b2 and b3 have same IP address
  - anycast routing used to reach closest
- all named roots have multiple instances

More Root Replication

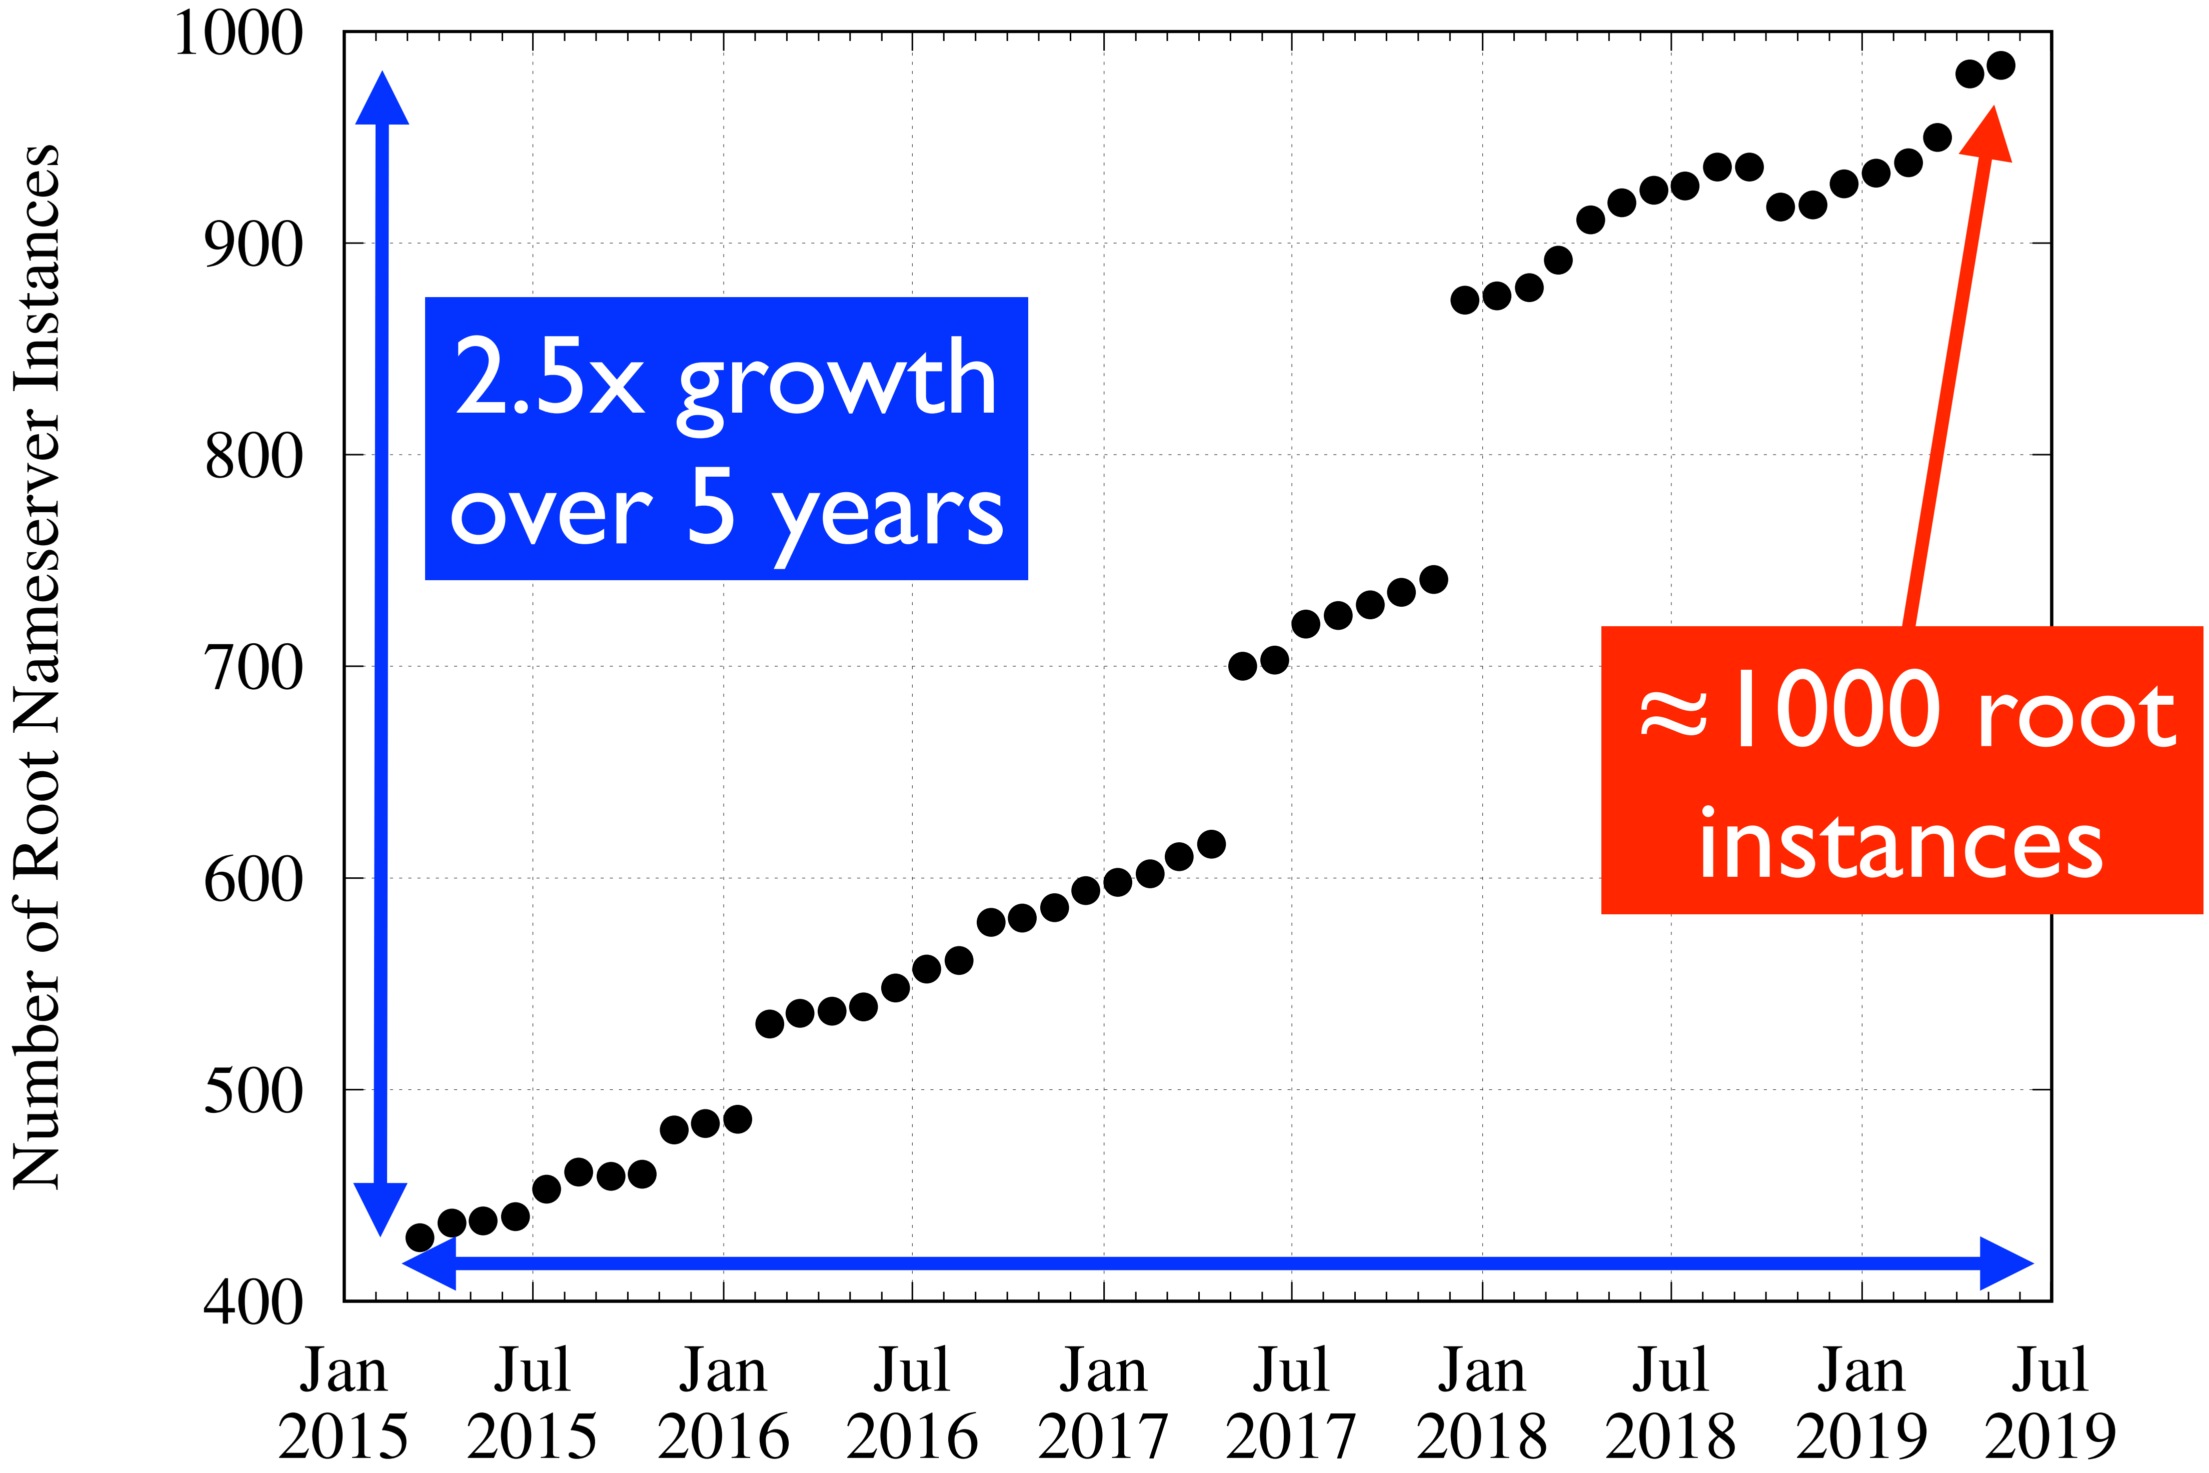

2.5x growth over 5 years

≈1000 root instances

Courtesy of root-servers.org

More Root Replication

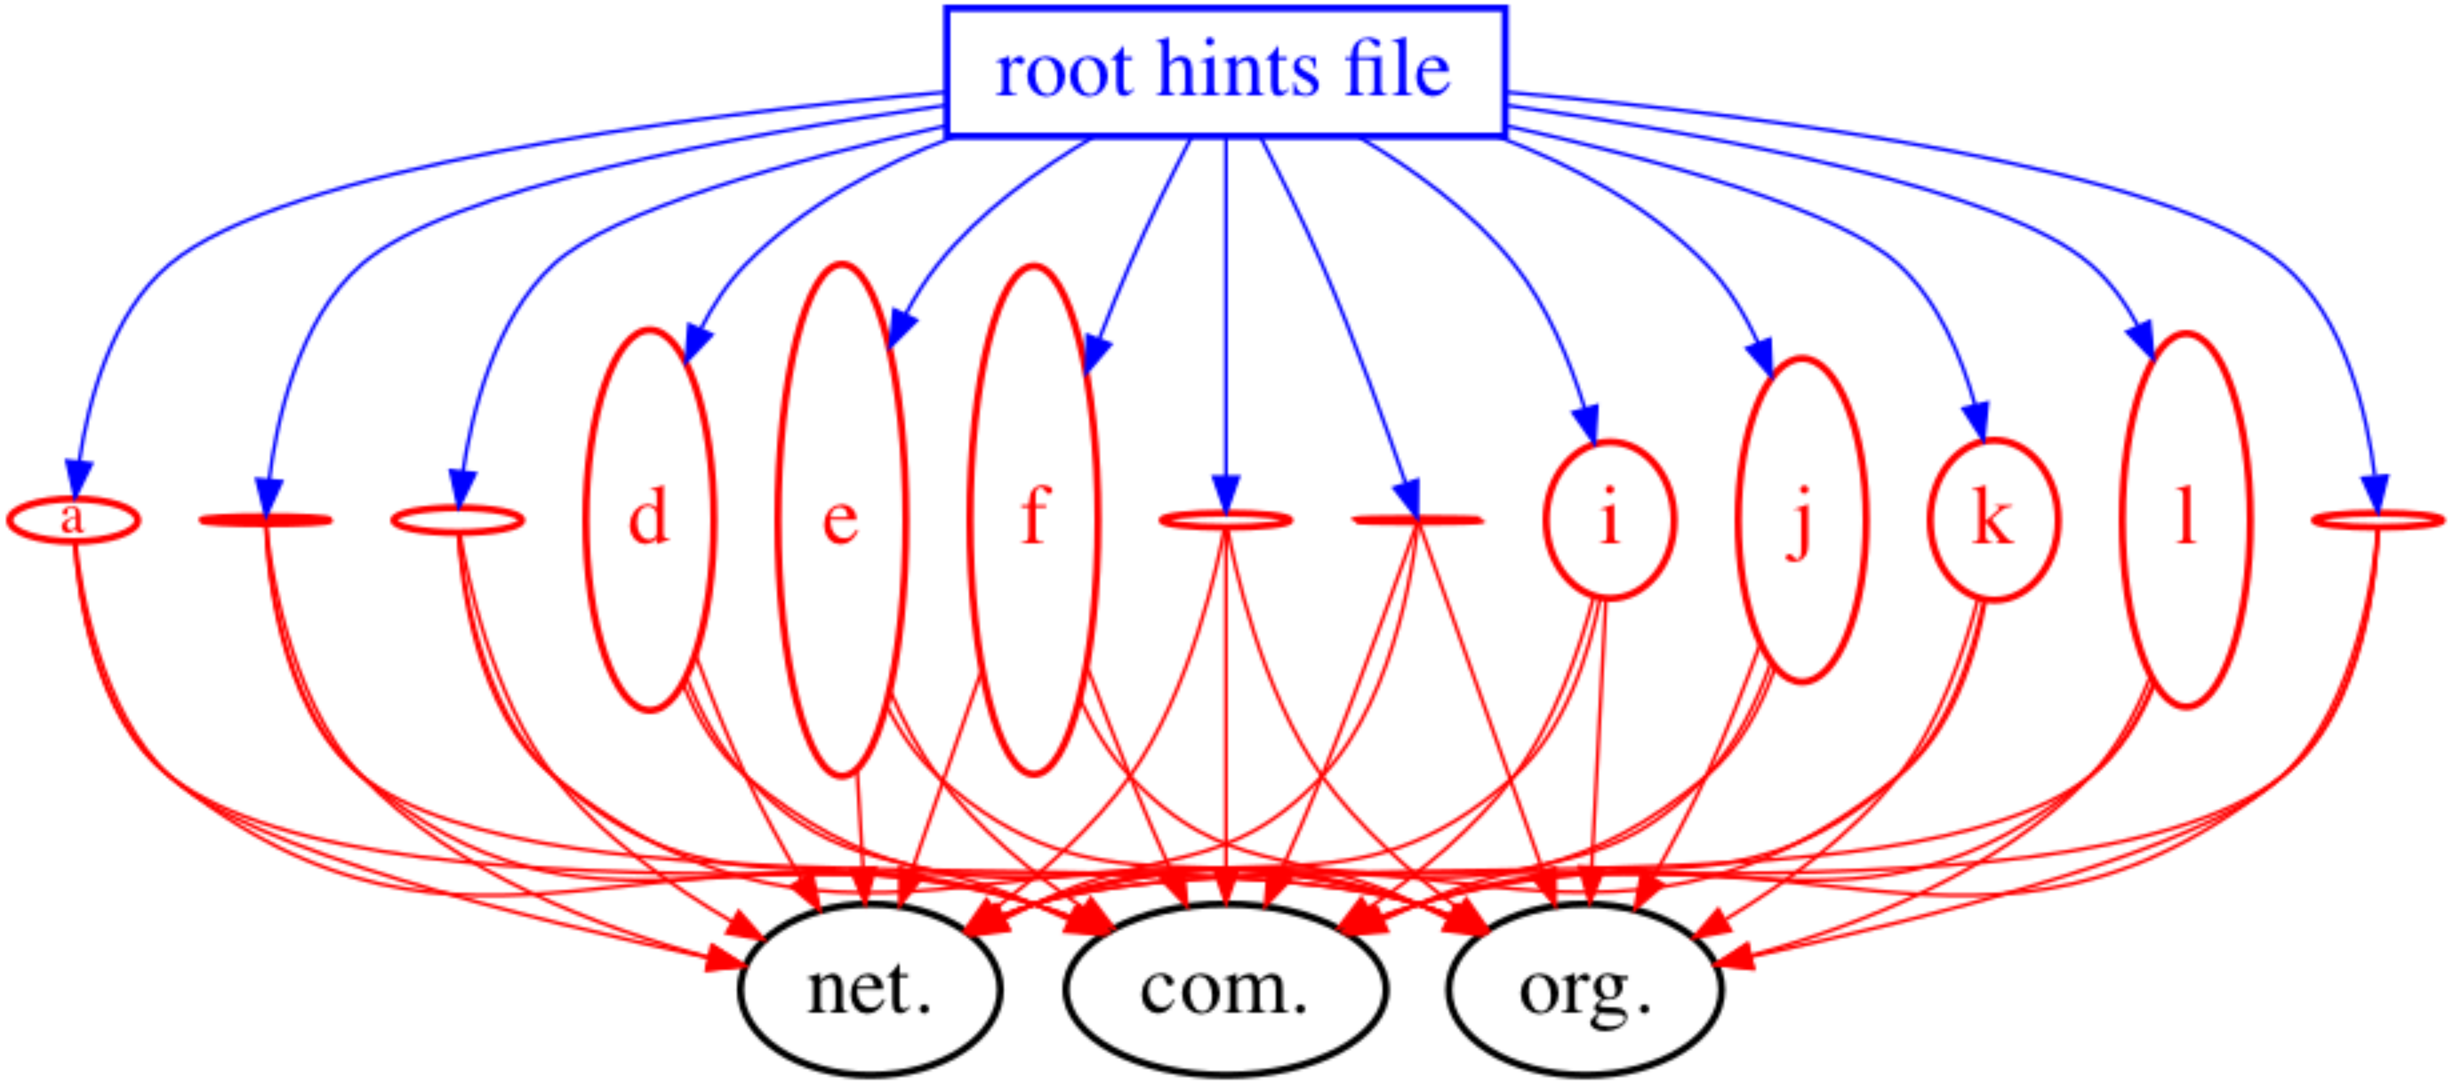

Root Server Load

	Queries/Day	Queries/Second
Total	107 B	1.2 M
Per Root	8.9 B	103 K
Per Root Instance	110 M	1.3K

Courtesy of root-servers.org & root operators

95+% Of Root Queries Are Junk

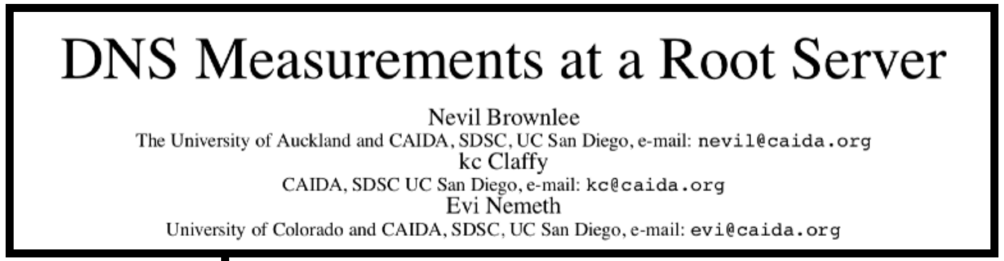

DNS Measurements at a Root Server

Nevil Brownlee
The University of Auckland and CAIDA, SDSC, UC San Diego, e-mail: nevil@caida.org
kc Claffy
CAIDA, SDSC UC San Diego, e-mail: kc@caida.org
Evi Nemeth
University of Colorado and CAIDA, SDSC, UC San Diego, e-mail: evi@caida.org

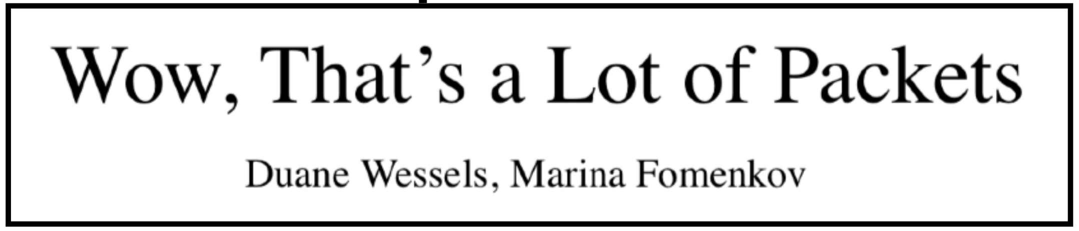

Wow, That's a Lot of Packets

Duane Wessels, Marina Fomenkov

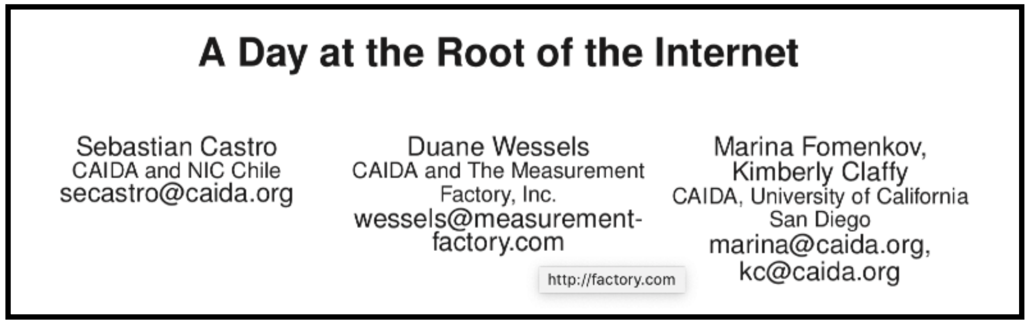

A Day at the Root of the Internet

Sebastian Castro
CAIDA and NIC Chile
secastro@caida.org

Duane Wessels
CAIDA and The Measurement Factory, Inc.
wessels@measurement-factory.com

http://factory.com

Marina Fomenkov,
Kimberly Claffy
CAIDA, University of California San Diego
marina@caida.org,
kc@caida.org

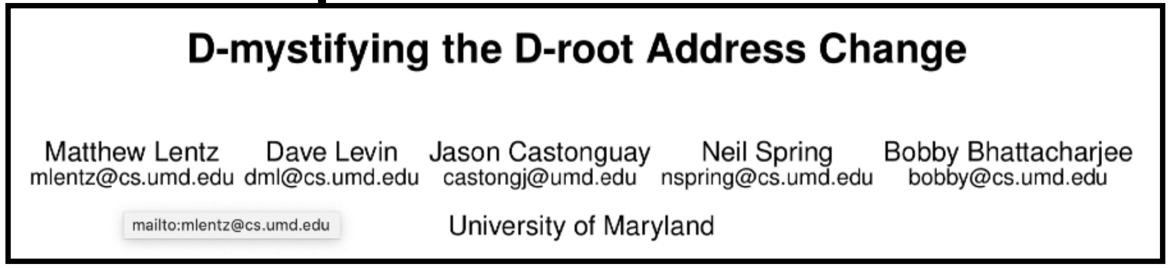

D-mystifying the D-root Address Change

Matthew Lentz
mlentz@cs.umd.edu

Dave Levin
dml@cs.umd.edu

Jason Castonguay
castongj@umd.edu

Neil Spring
nspring@cs.umd.edu

Bobby Bhattacharjee
bobby@cs.umd.edu

mailto:mlentz@cs.umd.edu

University of Maryland

We find 96.7-99.5% of the j-root queries are junk.

2001

2010

2019

Position

Current:

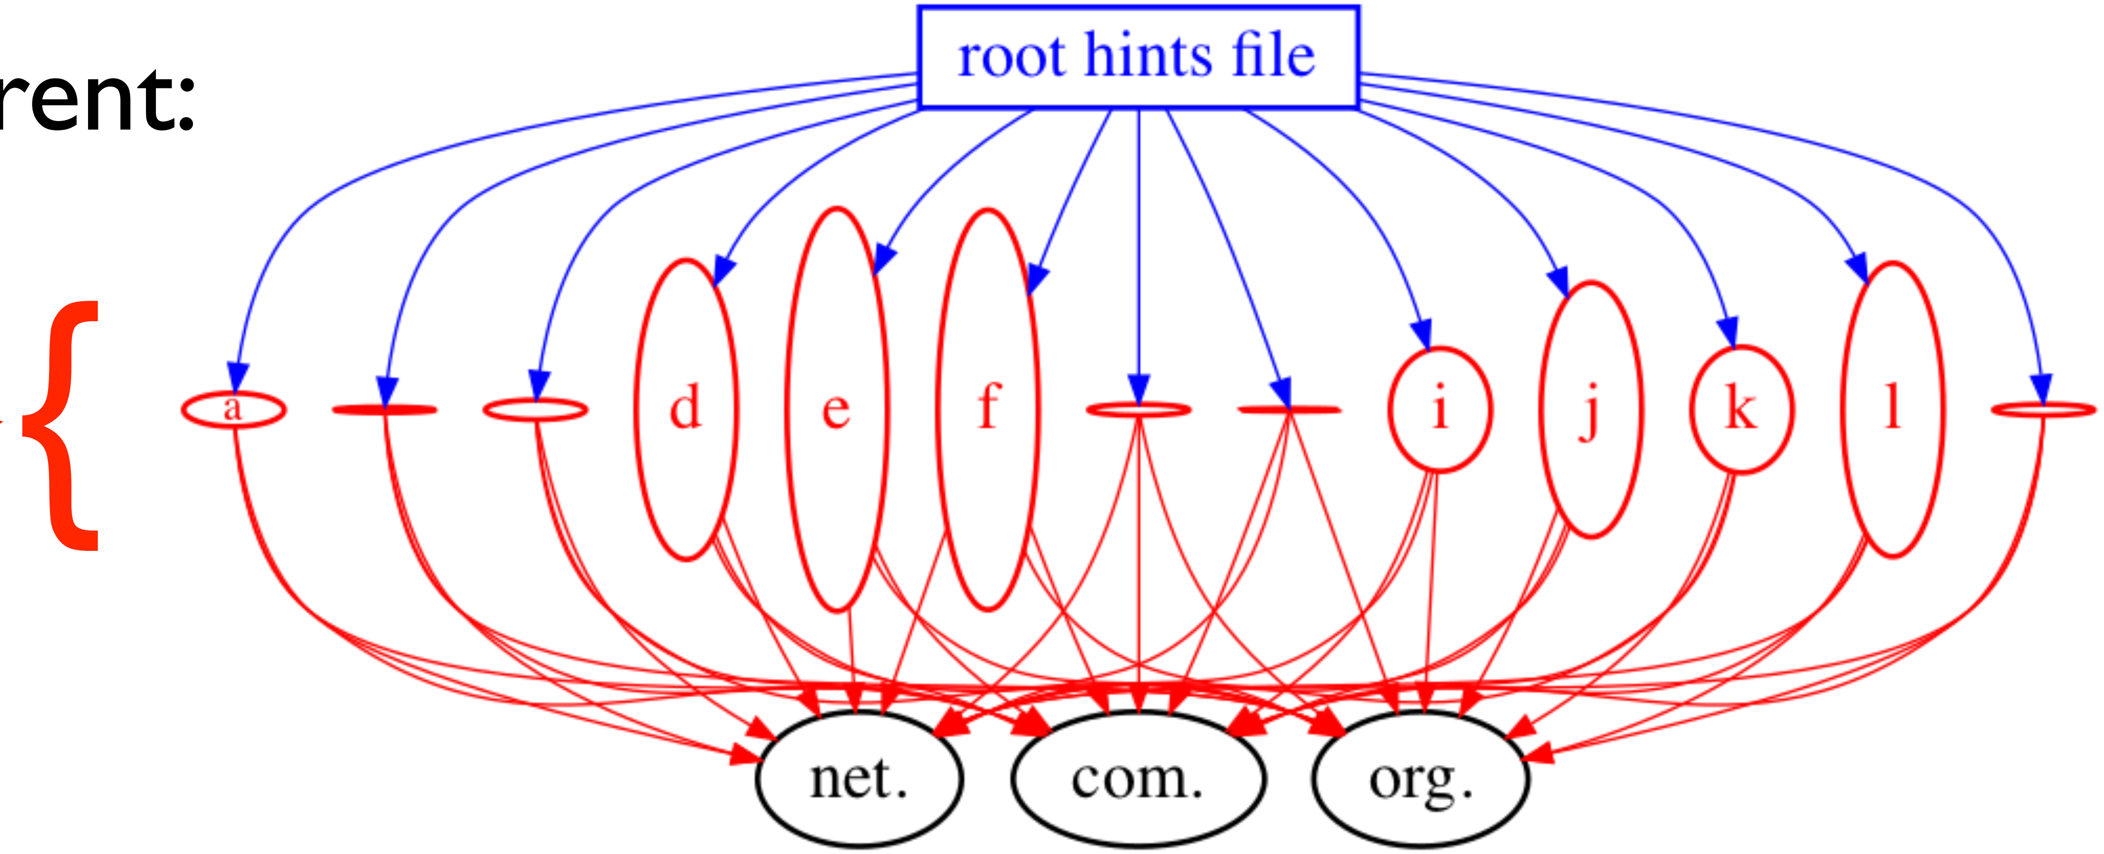

Position:

- bunch of infrastructure
- doing lots of mostly useless work

How To Realize?

Record Integrity

```
org.         172800 IN  NS  a0.org.afilias-nst.info.
org.         172800 IN  NS  a2.org.afilias-nst.info.
org.         172800 IN  NS  b0.org.afilias-nst.org.
org.         172800 IN  NS  b2.org.afilias-nst.org.
org.         172800 IN  NS  c0.org.afilias-nst.info.
org.         172800 IN  NS  d0.org.afilias-nst.org.
org.         86400  IN  DS  9795 7 1 364DFAB3DAF254CAB477B5675B107
org.         86400  IN  DS  9795 7 2 3922B31B6F3A4EA92B19EB7B52120
org.         86400  IN  RRSIG  DS 8 1 86400 20191117050000 2019110
org.         86400  IN  NSEC   organic. NS DS RRSIG NSEC
org.         86400  IN  RRSIG  NSEC 8 1 86400 20191117050000 20191
```

Distributing the Root Zone

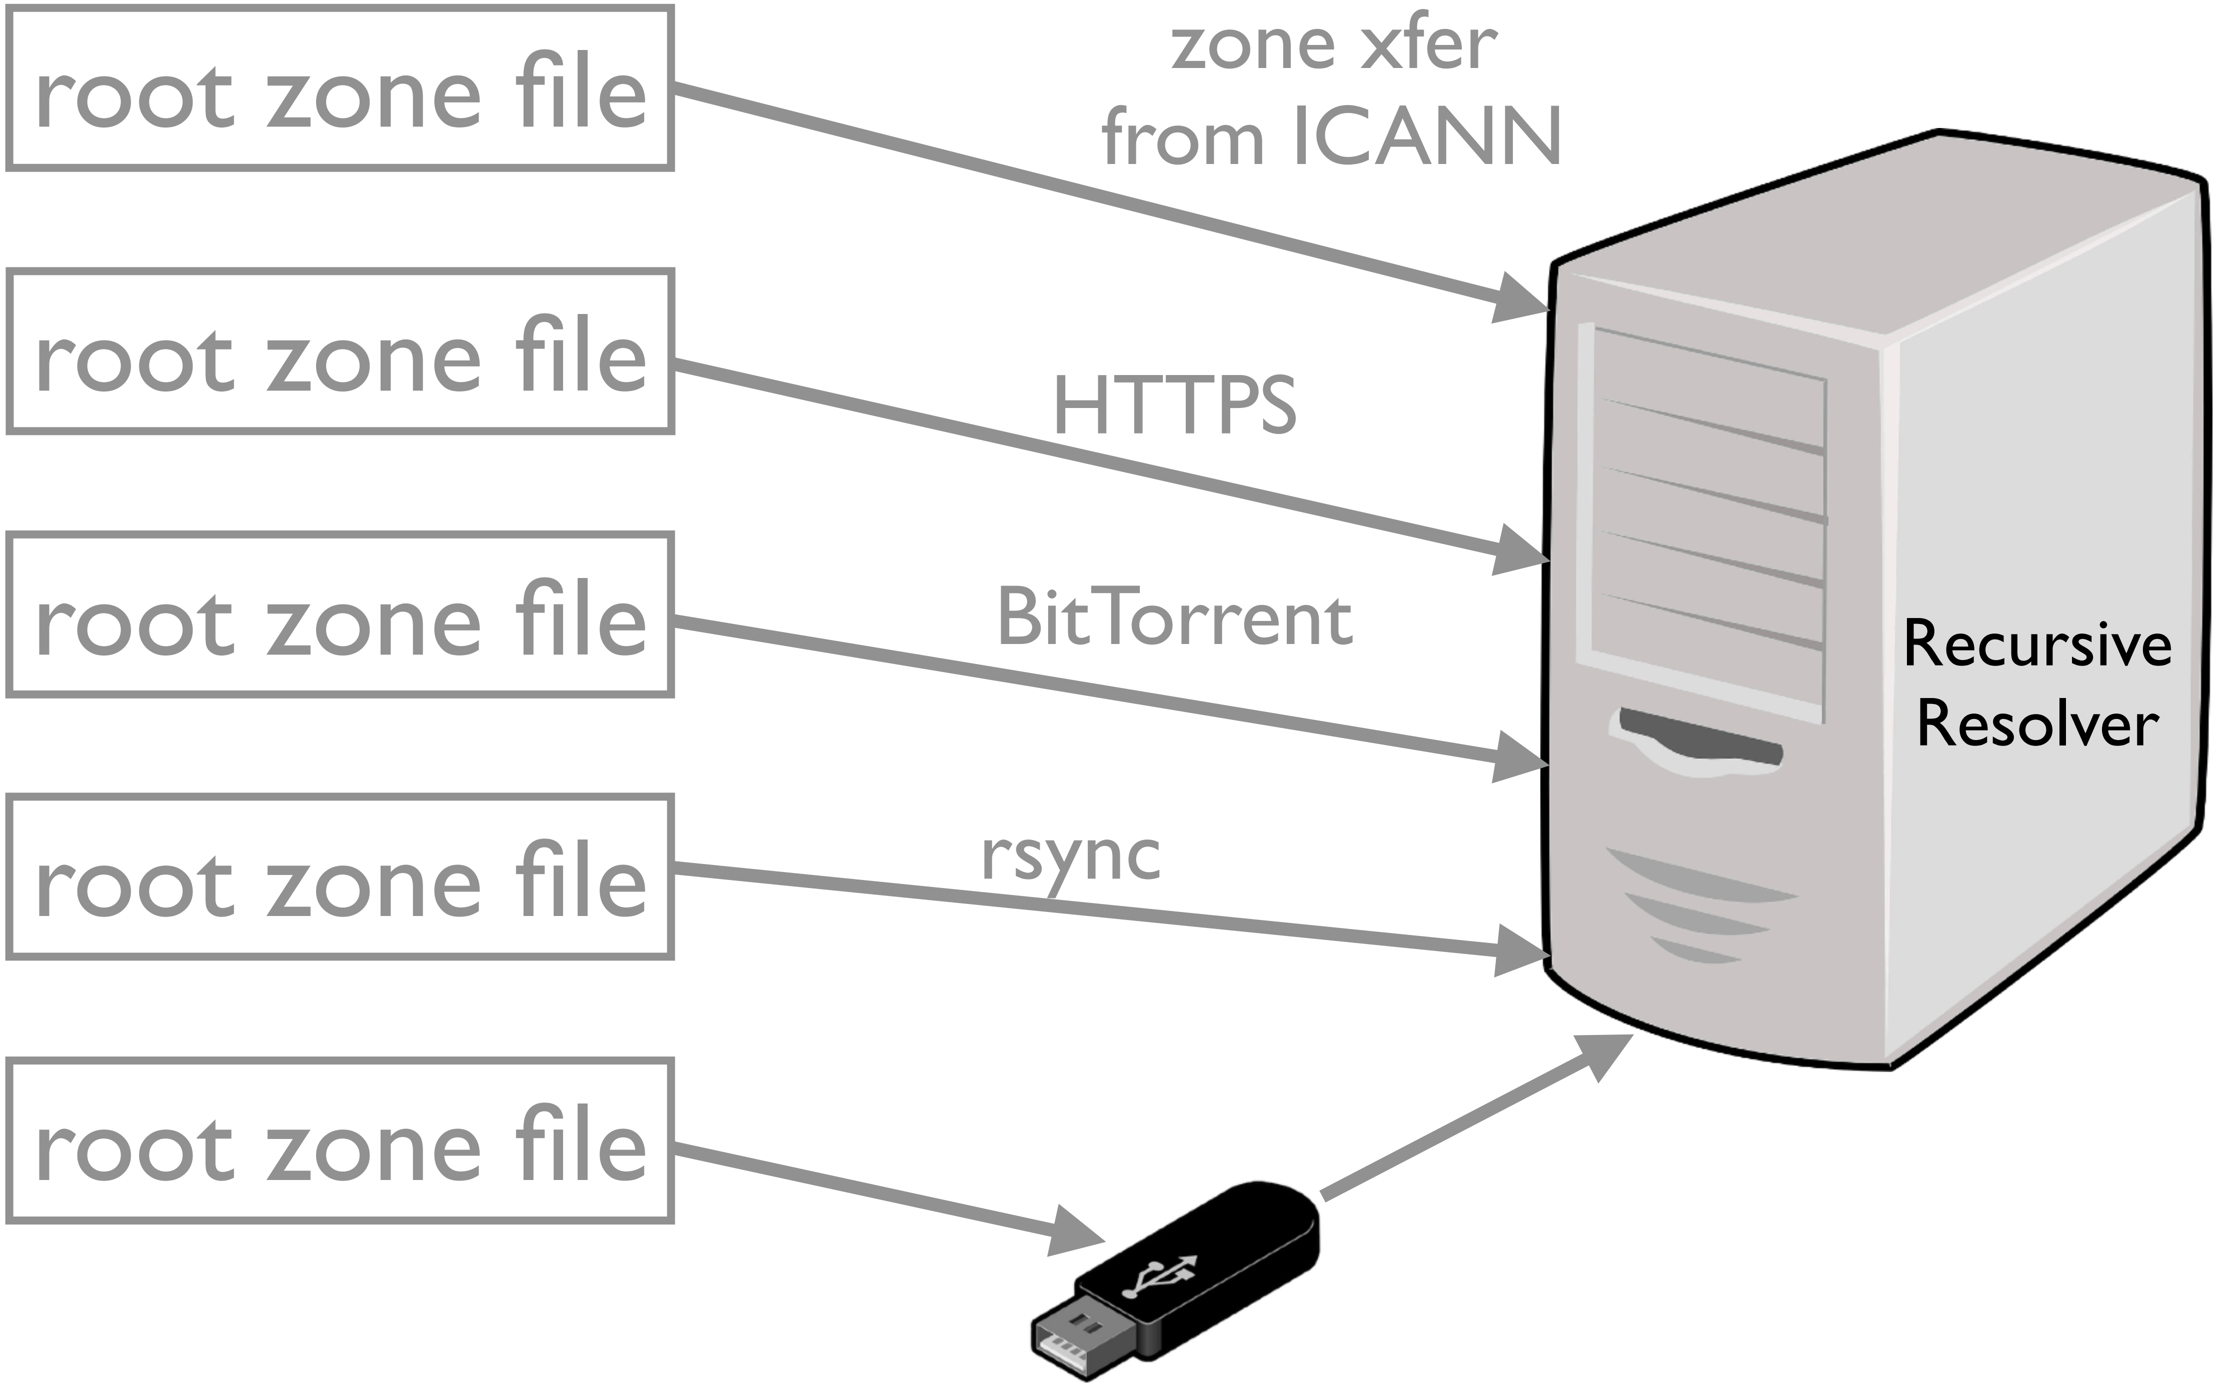

root zone file

zone xfer from ICANN

root zone file

HTTPS

root zone file

BitTorrent

root zone file

rsync

root zone file

Recursive Resolver

Using the Root Zone

- Using root zone files:
  - pre-load RRs into cache
  - demand-load RRs into cache
  - load RRs from file w/o cache
  - leverage database
  - etc.

Eliminating Root Servers

- Resolvers can switch independently
- No flag days

- Decommissioning root nameservers can happen gradually

Benefits

Less Infrastructure

- No longer need 1K servers to deal with 100+B (worthless) requests per day

- Less coordination effort
  e.g., DNS Root Server System Advisory Committee

Performance

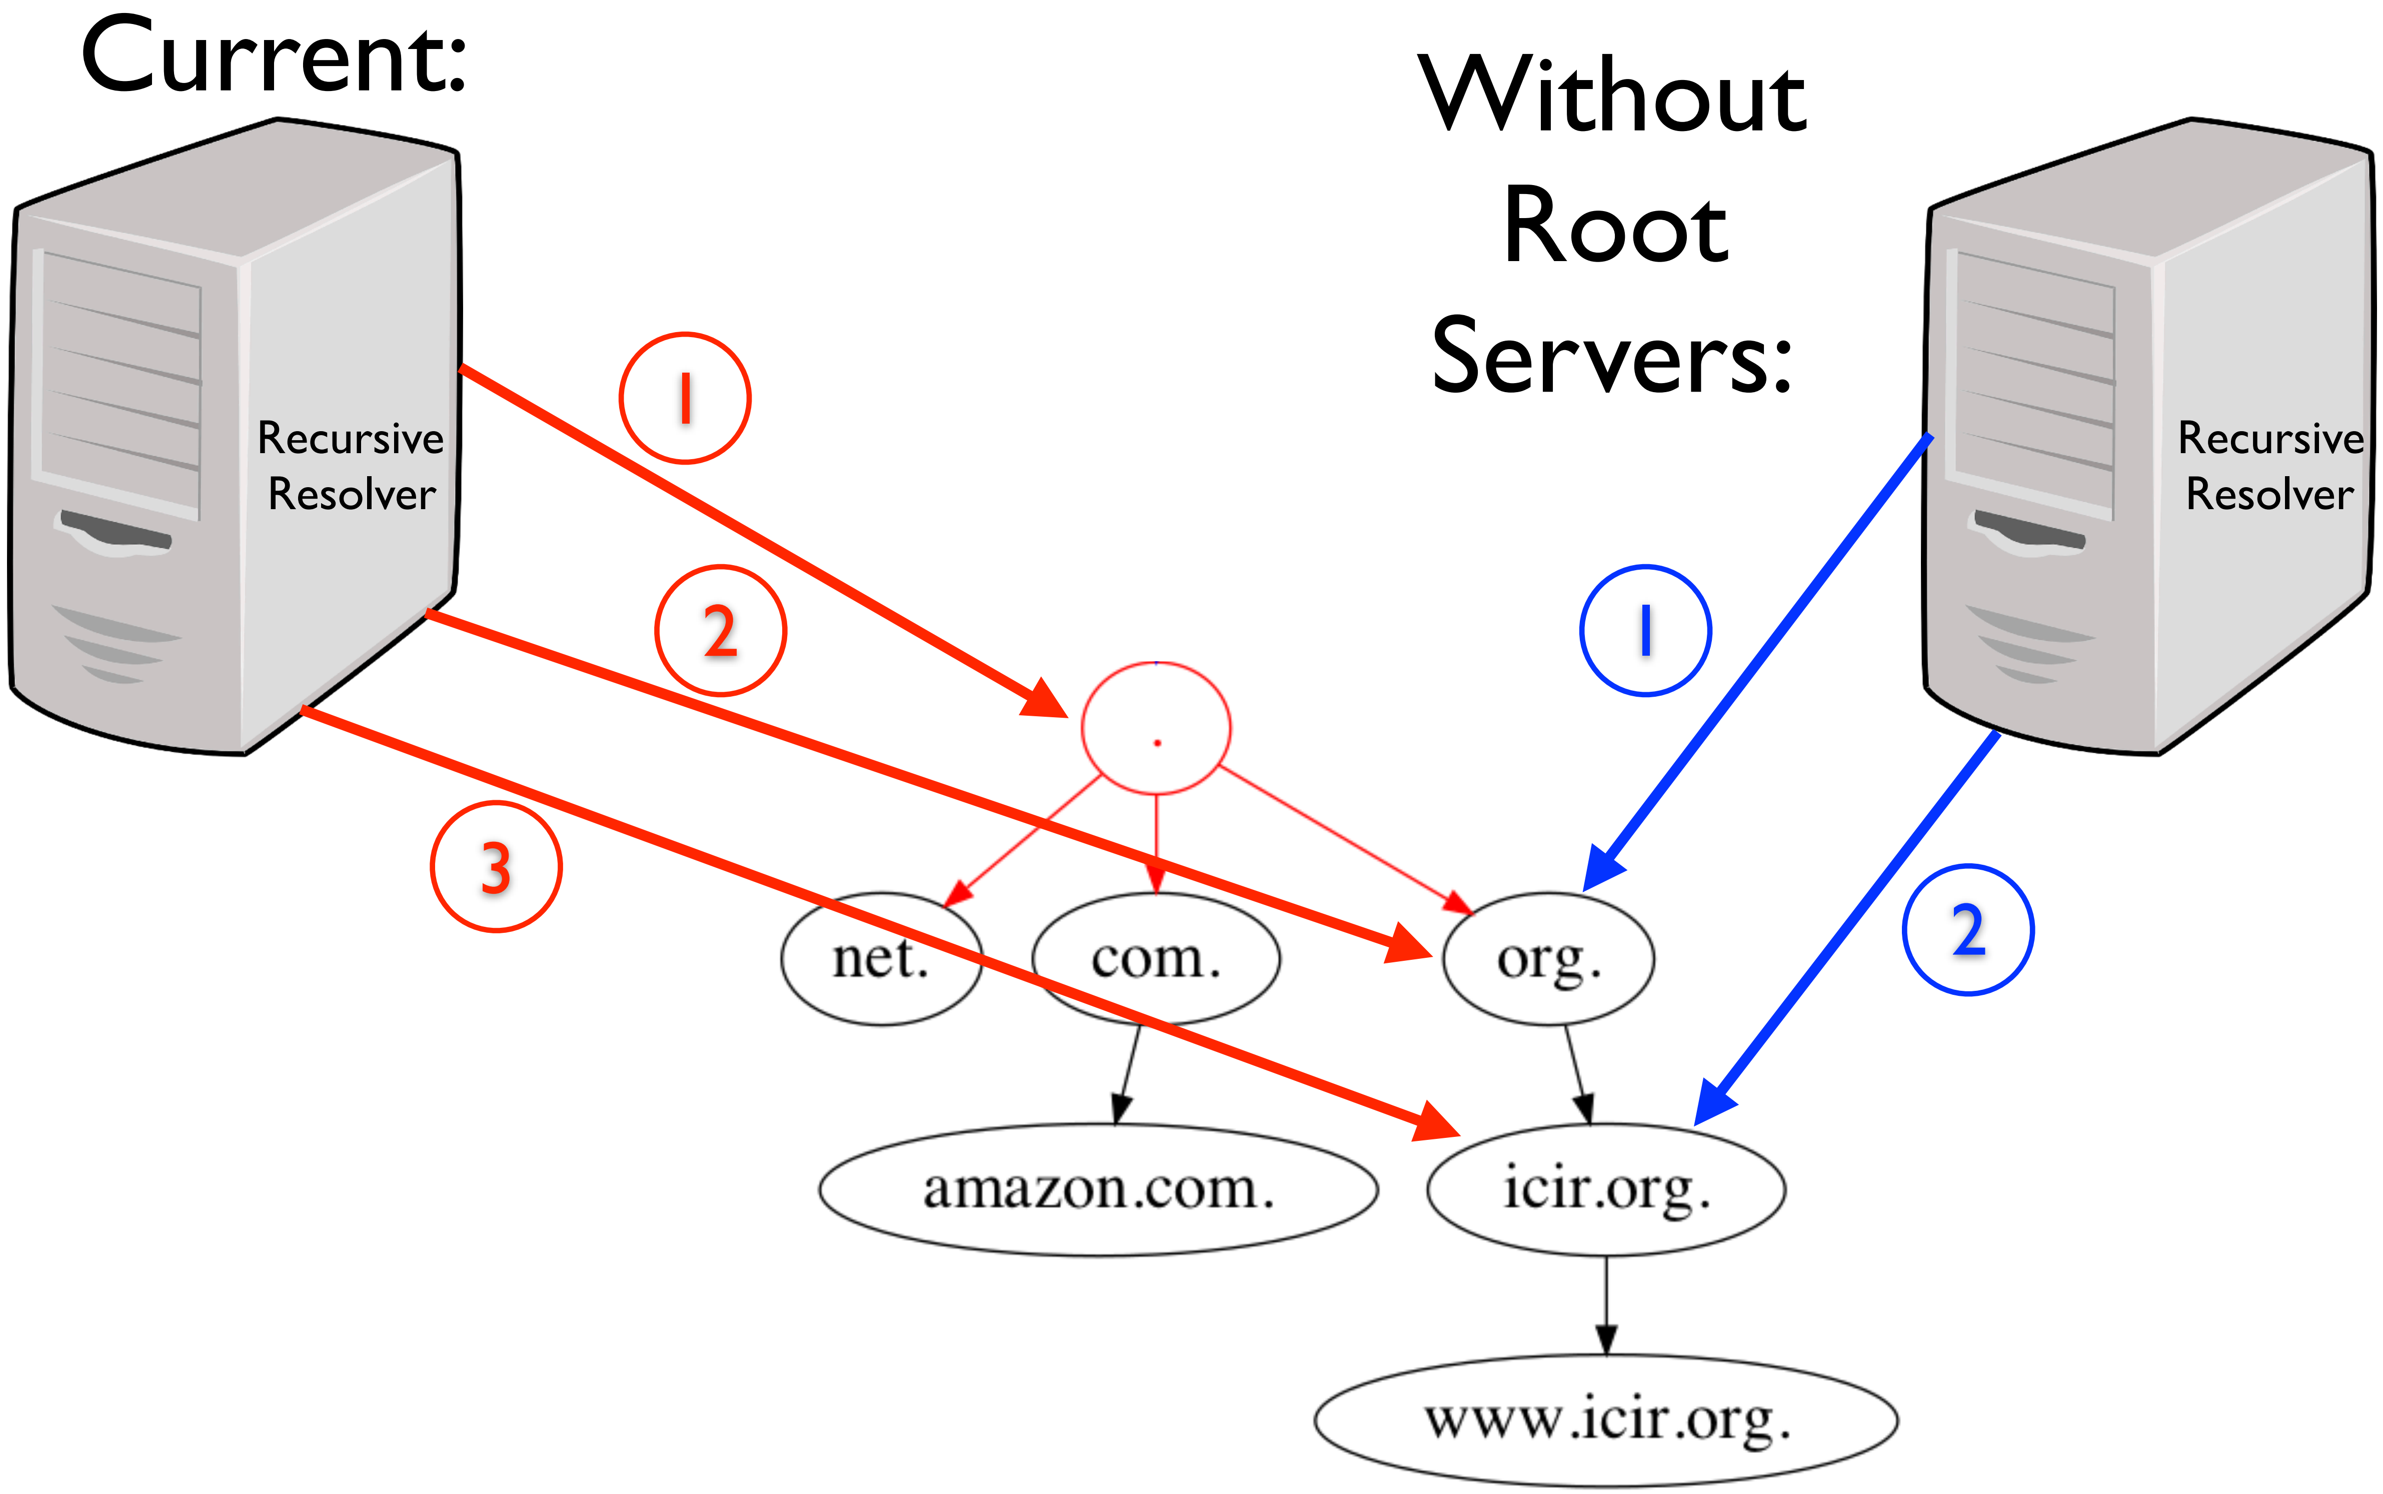

Security

- No root nameservers eliminates risk of …
  - DDoS attacks

- Eliminating some transactions lessens risk of …
  - man-in-the-middle attacks
  - cache poisoning attacks
  - censorship

Privacy

- Potentially sensitive lookups not revealed to …
  - the root nameservers
  - network monitors

Costs

Root Zone File Size

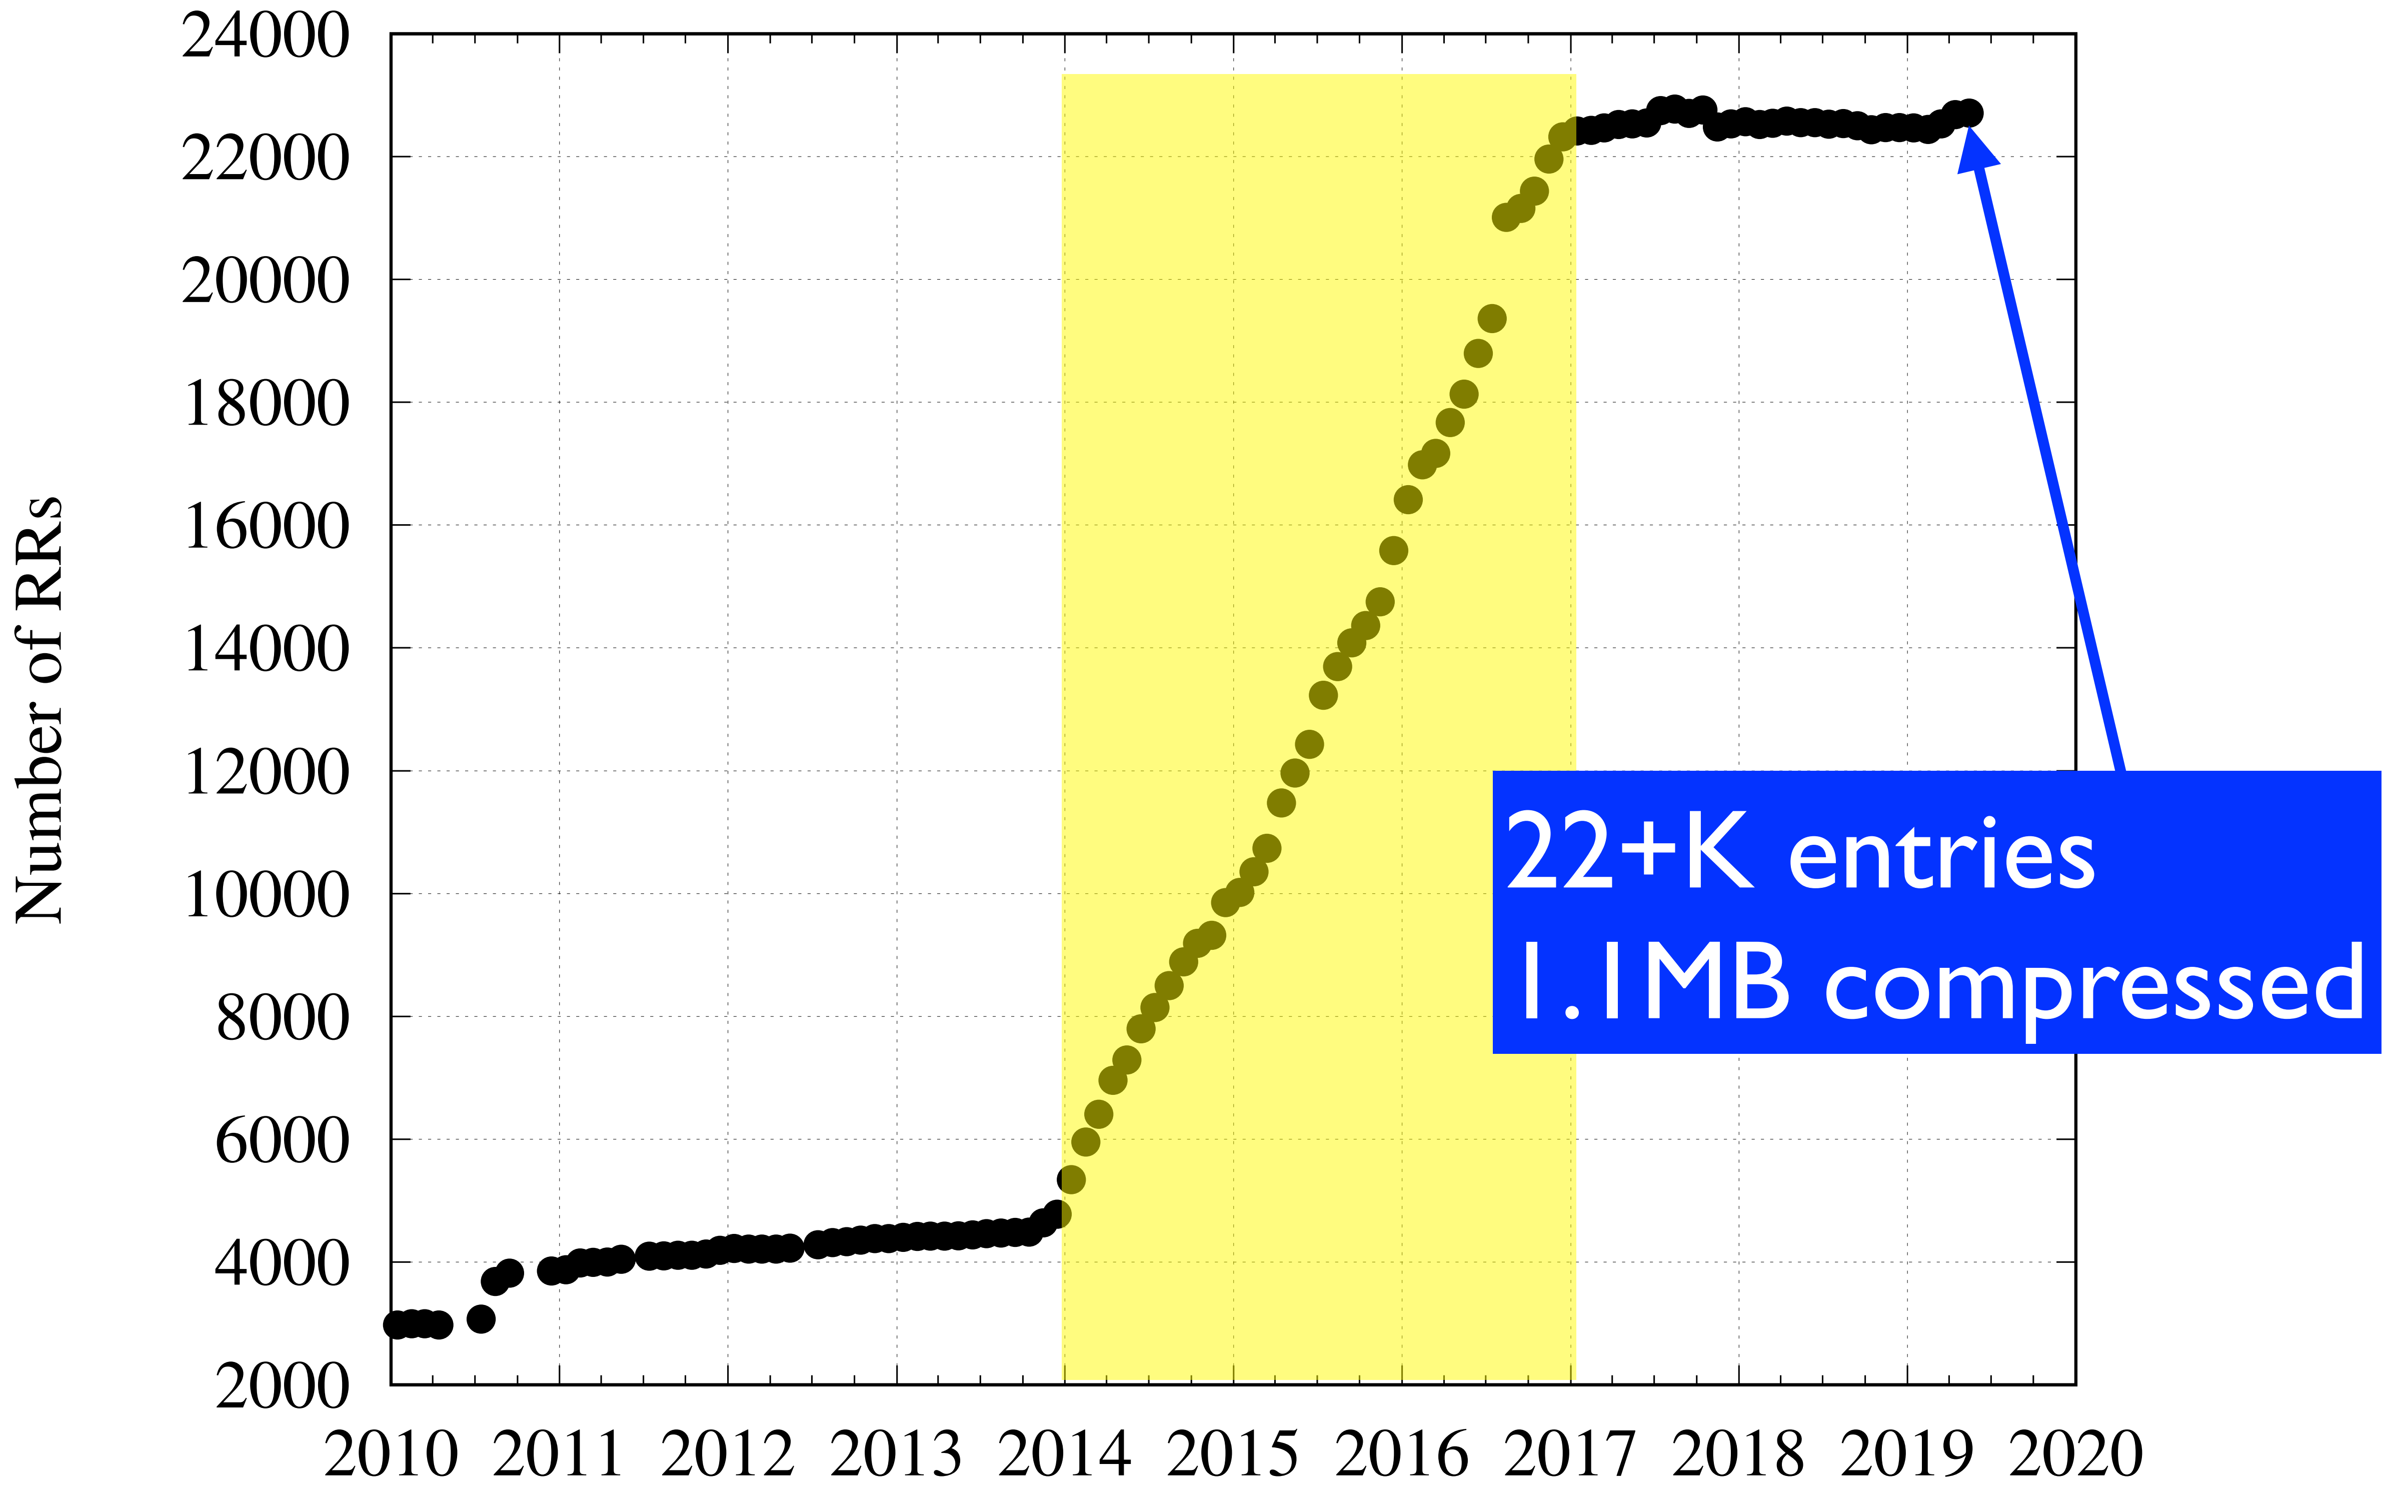

22+K entries
1.1MB compressed

How Onerous is Distribution?

- 1.1MB transfer every 2 days per resolver (the current TTL in root zone file)

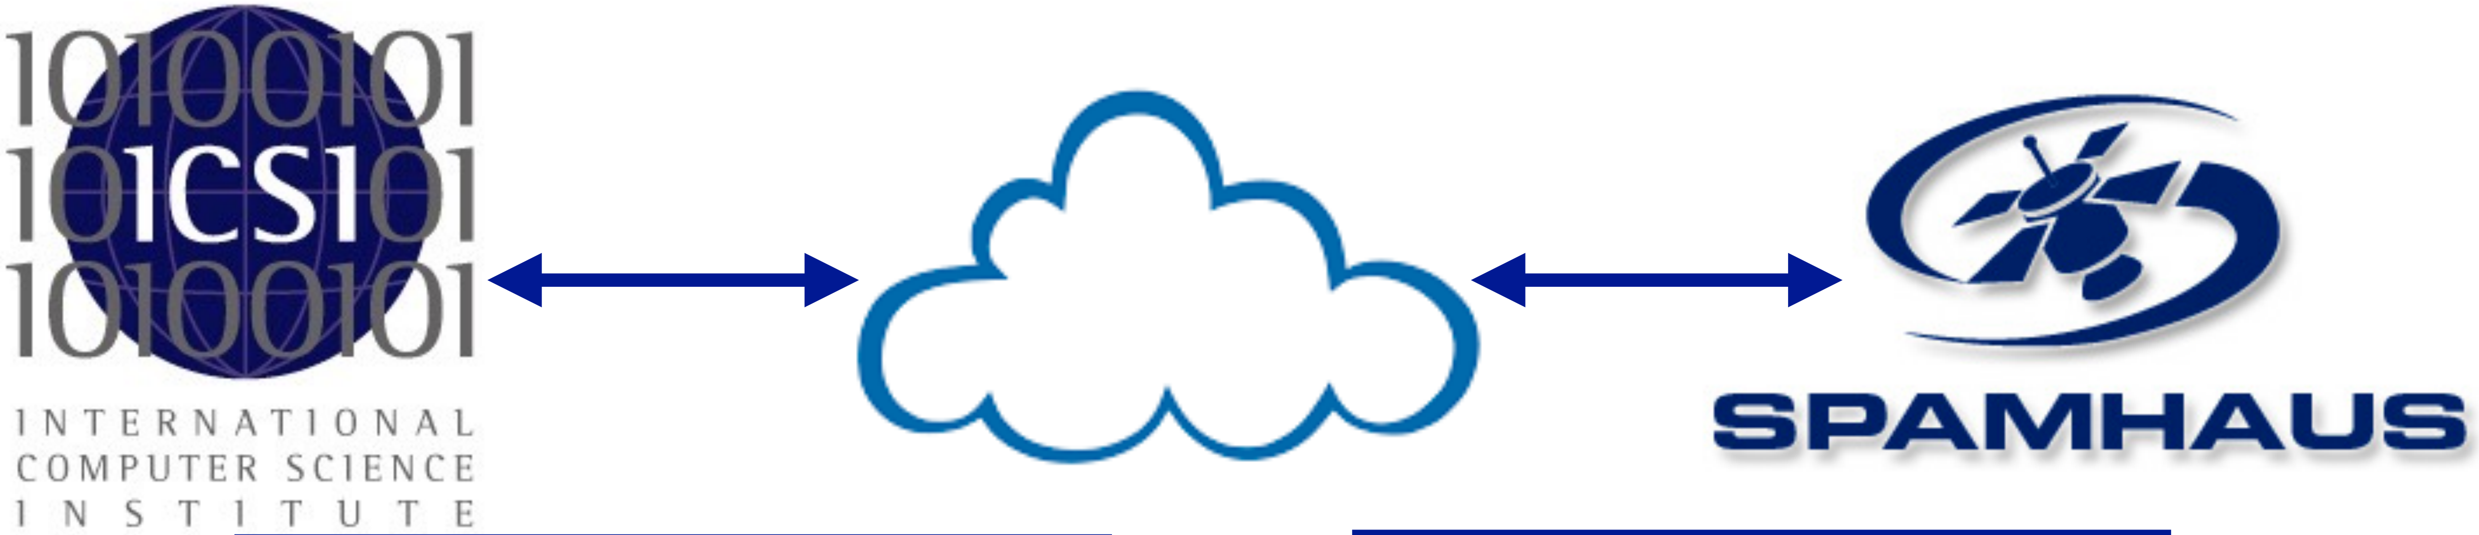

rsync once/minute

total: 3.1GB / day

- The root zone file is fairly static and the TTL could be increased to lower the distribution load

Summary

- DNS root infrastructure is …
  - … large
  - … and growing
  - … busy
  - … but doing relatively little useful work

- We ***can*** get rid of the infrastructure
- We ***should*** re-organize to a system without root nameservers

Questions? Comments?

Mark Allman, *mallman@icir.org*
https://www.icir.org/mallman/
@mallman_icsi