SAND2010-1523C

Measuring Availability in the Domain Name System

Casey Deccio

Prasant Mohapatra

Sandia National Laboratories

Krishna Kant
Jeff Sedayao

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

DNS Availability

- DNS is integral to Internet activity
- Name resolution is complex due to a network of dependencies
- Availability cannot be measured only by analyzing availability of servers

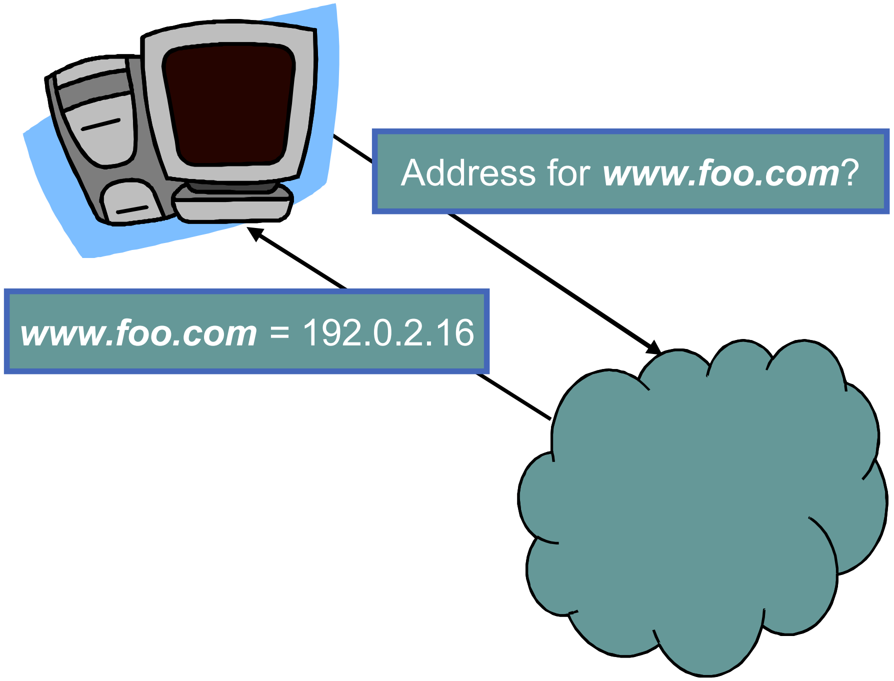

Address for ***www.foo.com***?

www.foo.com = 192.0.2.16

Objectives

- Quantify availability of a domain name
- Quantify the impact of DNS misconfigurations on availability
- Impact: security, availability, and performance

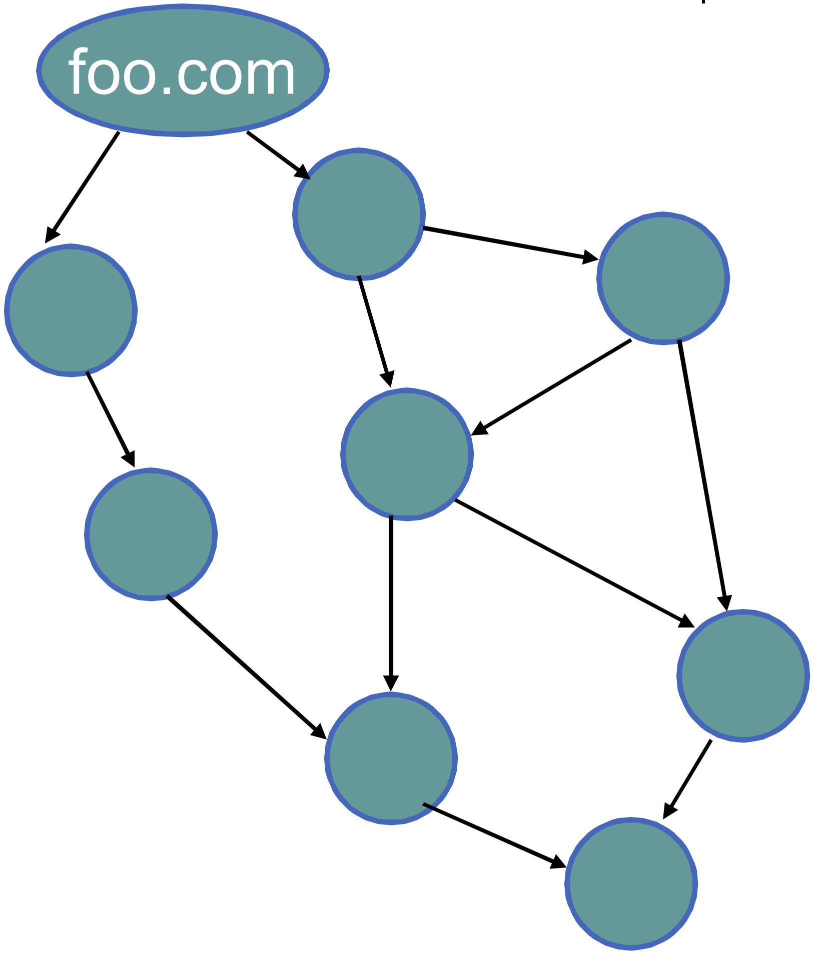

Sandia National Laboratories

DNS name dependencies

- Child-to-parent dependencies
- Alias dependencies
- NS target dependencies
  - *Names* used to designate servers authoritative for zone:

    foo.com. NS ns.foo.com.
  - Resolver needs *address* to query server:

    ns.foo.com → **192.0.2.1**
  - Resolvers must independently resolve *out-of-bailiwick* names or names without *glue* records

    com. **provides only name for** *me.baz.net*; **resolver must look up address**

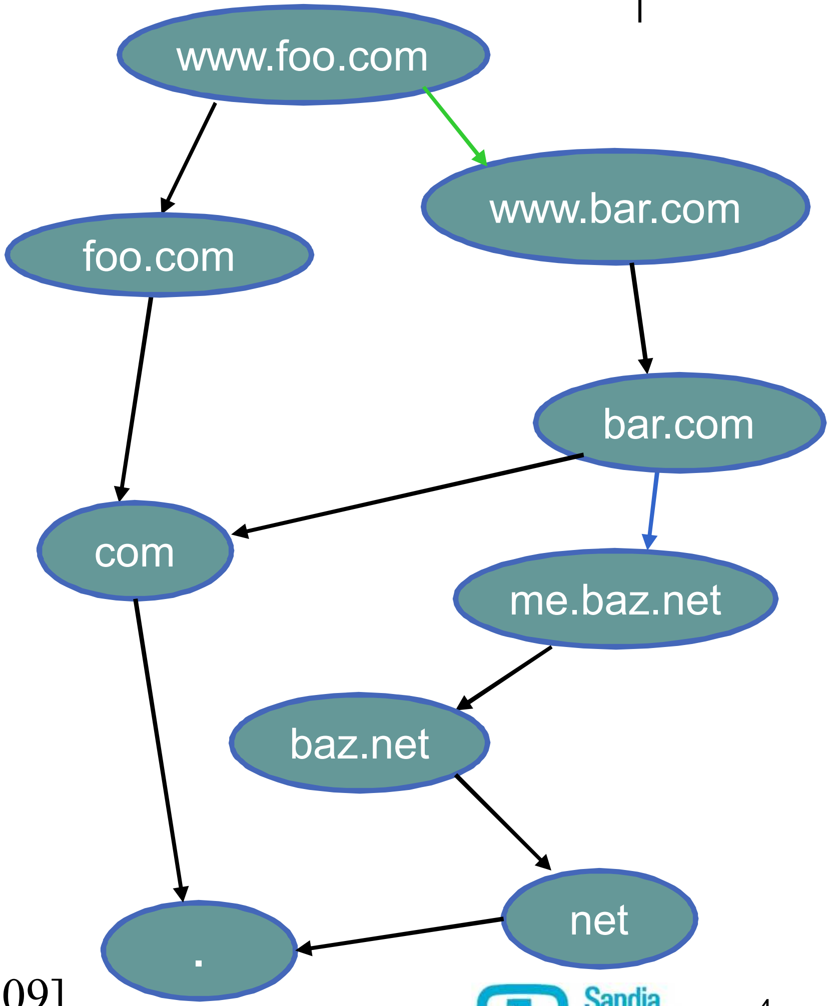

"Quality of Name Resolution in DNS" [Deccio 2009]

Adding server dependencies

- Direct server dependencies:
  - **Zone-to-server:** Dependency of zone on server whose name has in-bailiwick glue record
  - **Name-to-server:** Dependency of name on address
- Indirect server dependencies:
  - Transitivity:

    If *a* depends on *b*, and *b* depends on *c*, then *a* depends on *c*

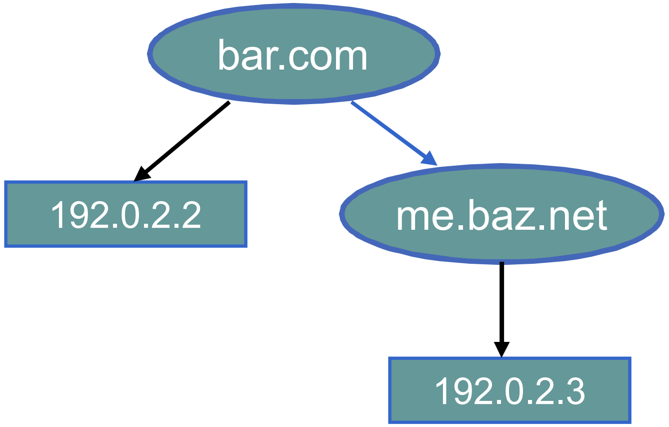

Resolver states and bootstrapping

- *Ignorant*: doesn't know names or addresses of authoritative servers
- *Knowledgeable*: knows names and addresses of authoritative servers
- *Bootstrapping*: becoming knowledgeable

Ask ***com***: ***a.gtld-servers.net*** (192.0.2.6) ***b.gtld-servers.net*** (192.0.2.7)
Ask ***bar.com***: ***me.baz.net*** (??) ***ns.bar.com*** (192.0.2.2)
Ask ***net***: ***a.gtld-servers.net*** (192.0.2.6) ***b.gtld-servers.net*** (192.0.2.7)
Ask ***baz.net***: ***ns1.baz.net*** (192.0.2.4) ***ns2.baz.net*** (192.0.2.5)
me.baz.net = 192.0.2.3

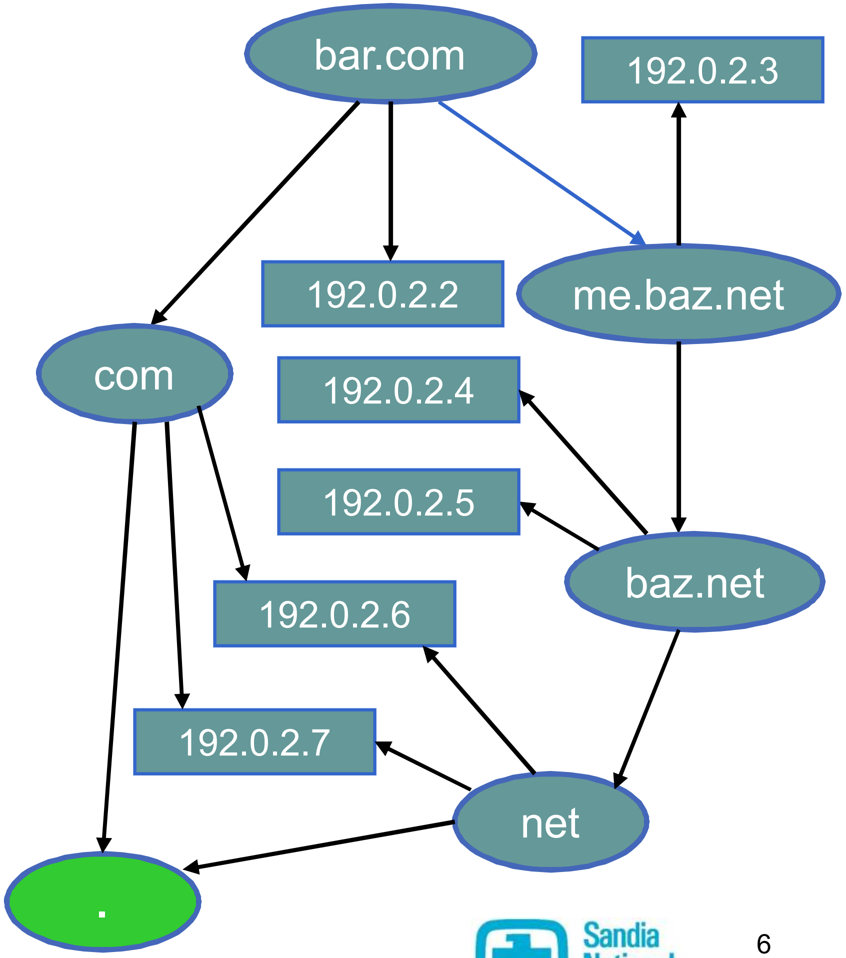

Domain name availability for knowledgeable resolvers

- When a resolver is knowledgeable about a zone, availability based on that of authoritative servers
- A resolver remains knowledgeable about a zone only until pertinent TTLs expire

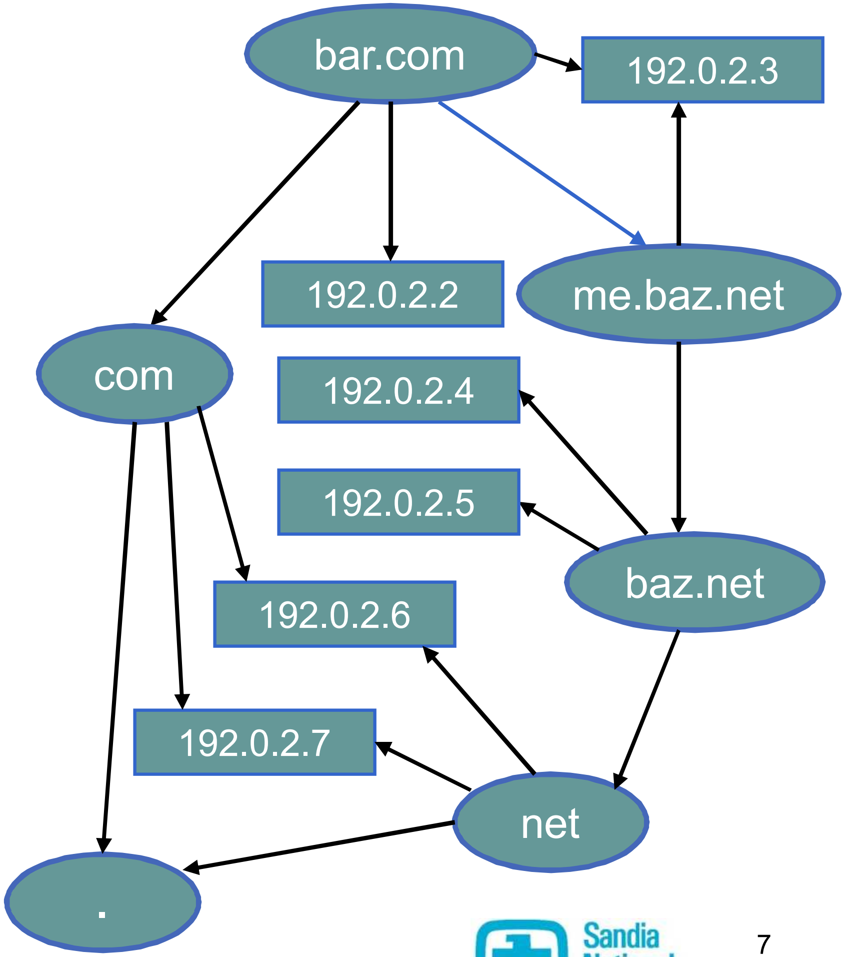

Domain name availability for ignorant resolvers

- Ignorant resolvers rely on the availability of intermediate names and servers
- Leaf nodes (addresses) represent “knowledge anchors”

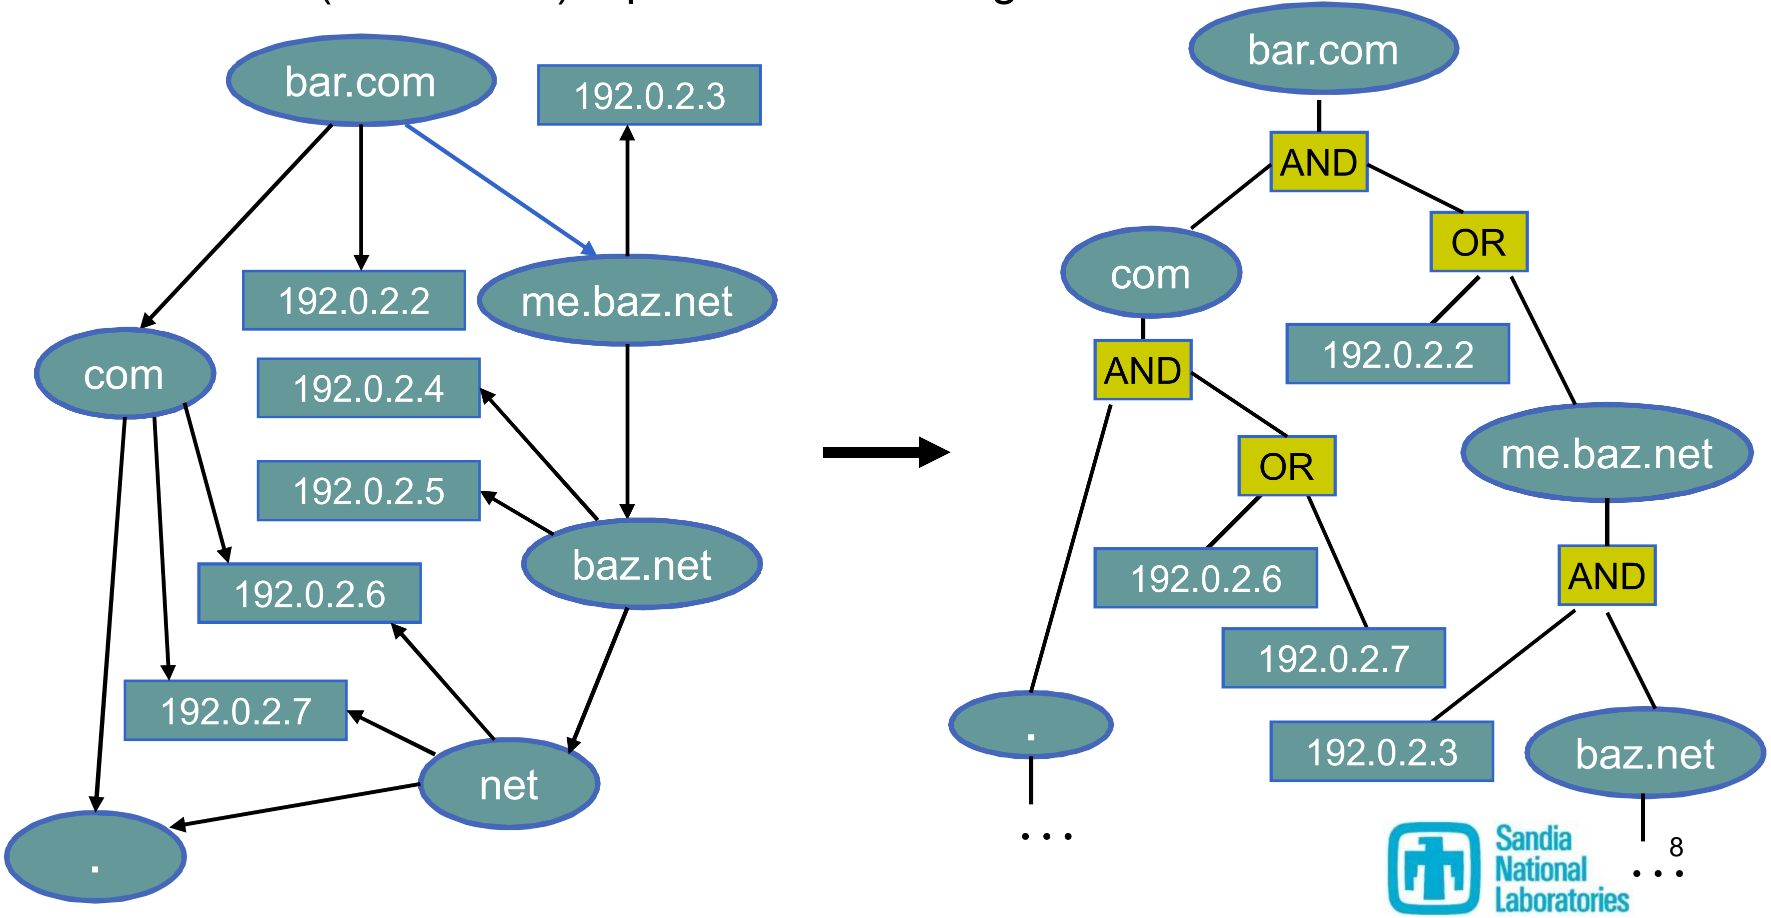

Sandia National Laboratories

Minimum servers queried (MSQ)

- *MSQ*: minimum number of servers necessarily queried for resolution of a domain name
- Evaluated by reducing logical availability tree to DNF with minimum sized conjunctions

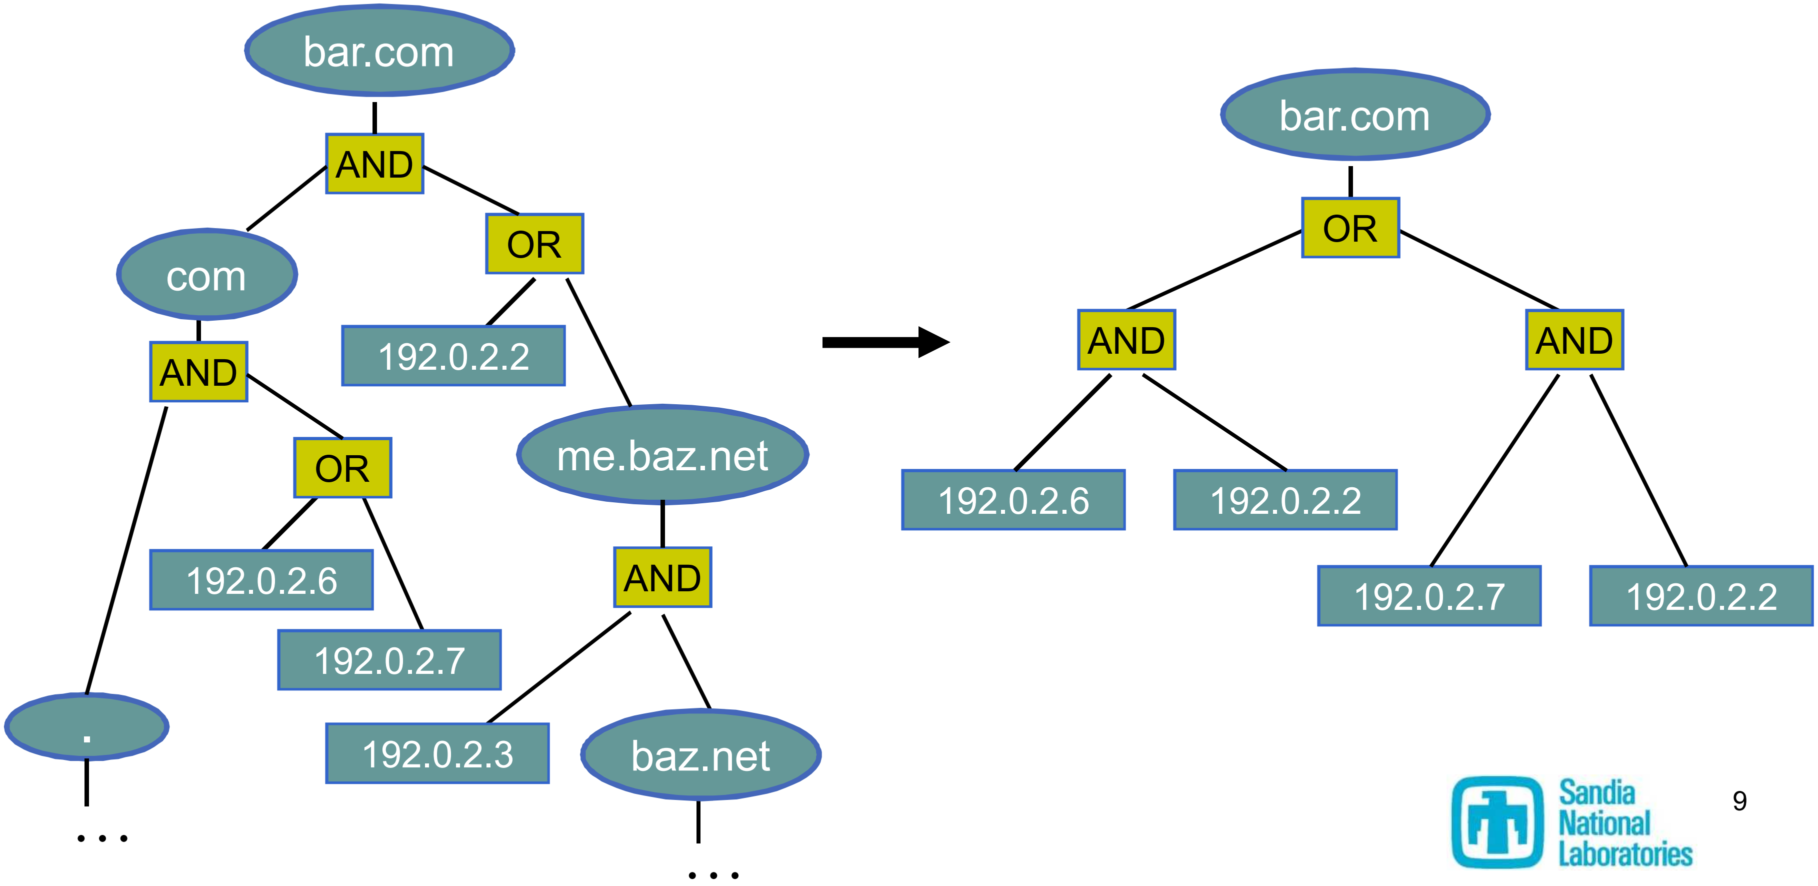

MSQ impact

- Potential impact of large MSQ:
  - Degraded performance (mitigated by caching)
  - Reduced availability (mitigated by increased redundancy)
- *Optimal MSQ*: size of conjunctions is ≤ number of ancestor zones

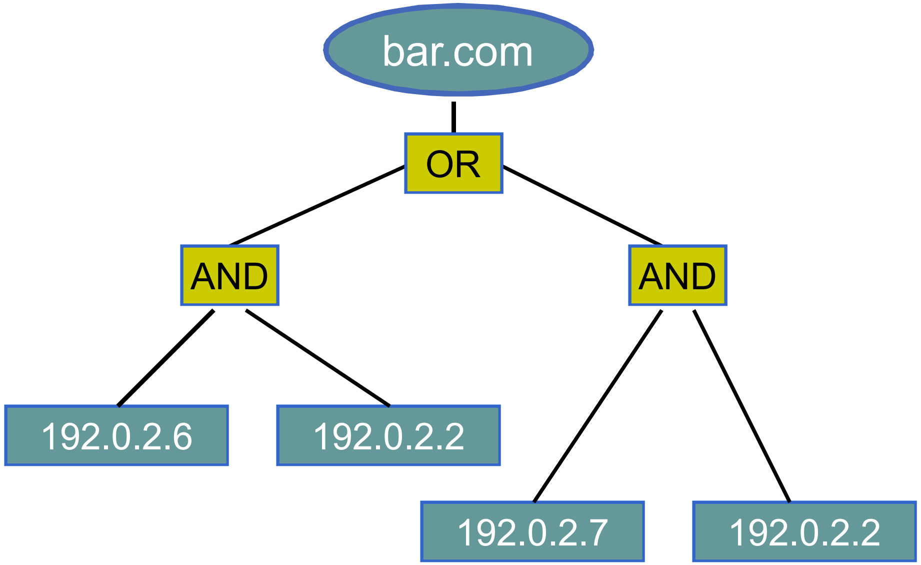

MSQ(bar.com) = 2 ≤ 2

(optimal MSQ)

Domain name redundancy

- *Redundancy*: minimum number of redundant servers in required resolution path of a domain name (i.e., “availability bottleneck”)
- Evaluated by reducing logical availability tree to CNF with minimum sized conjunctions

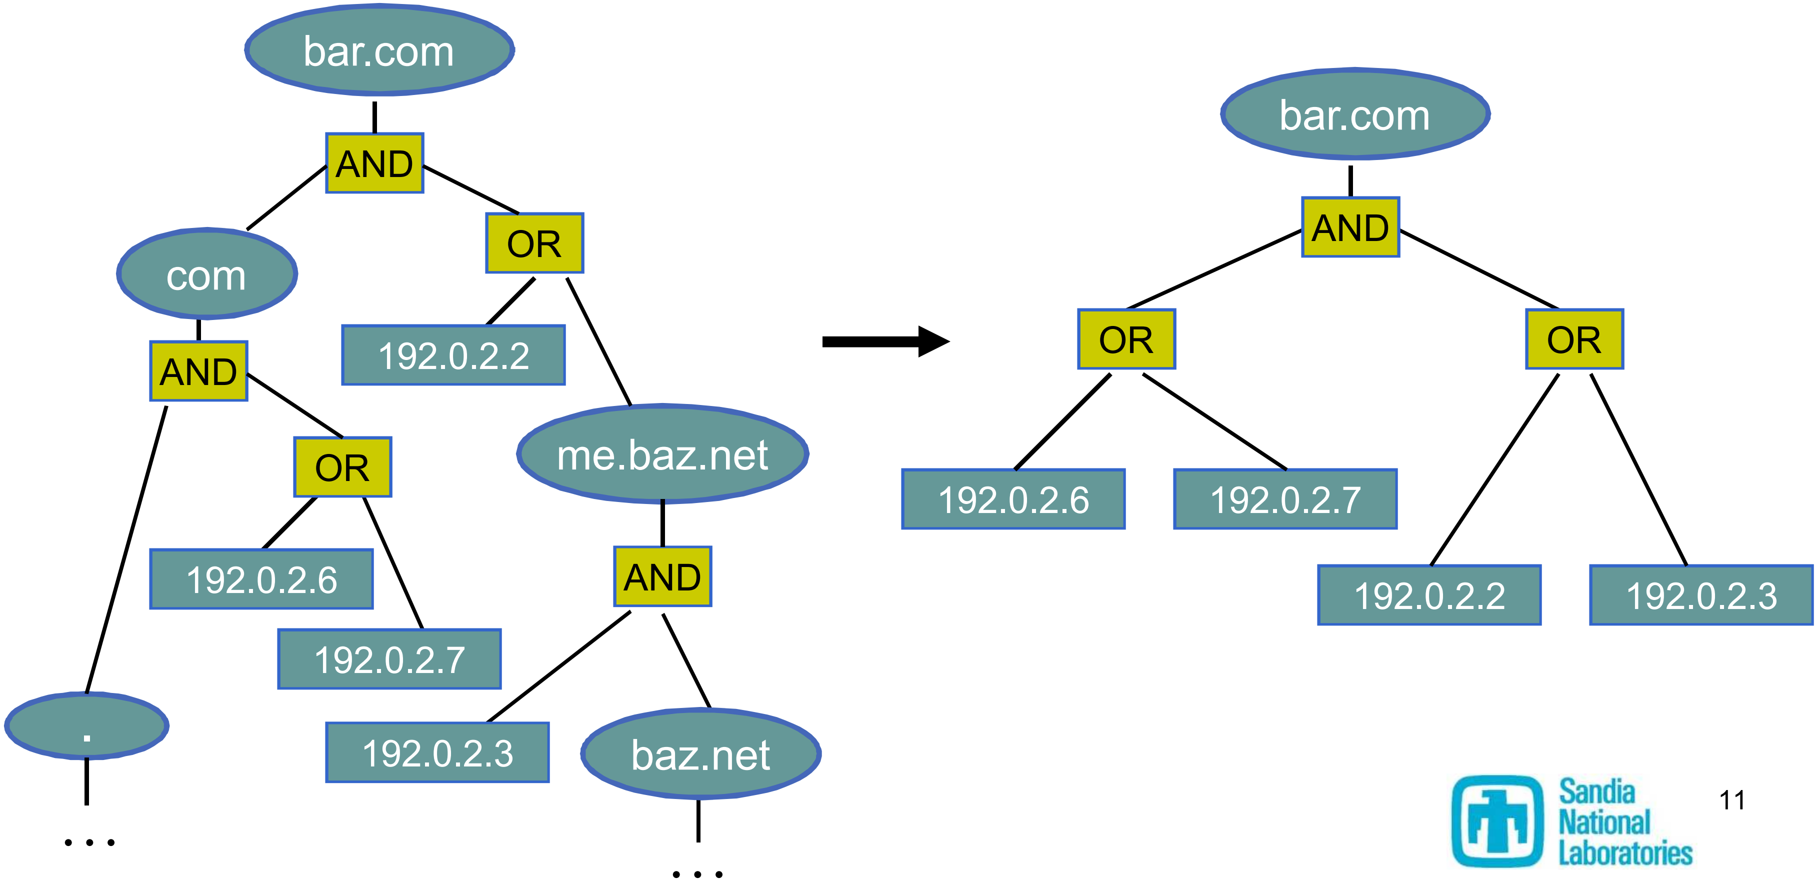

Redundancy impact

- If all servers fail in any one set of redundancies, the domain name becomes unavailable
- *False redundancy*: size of disjunctions is < number of NS targets for domain name

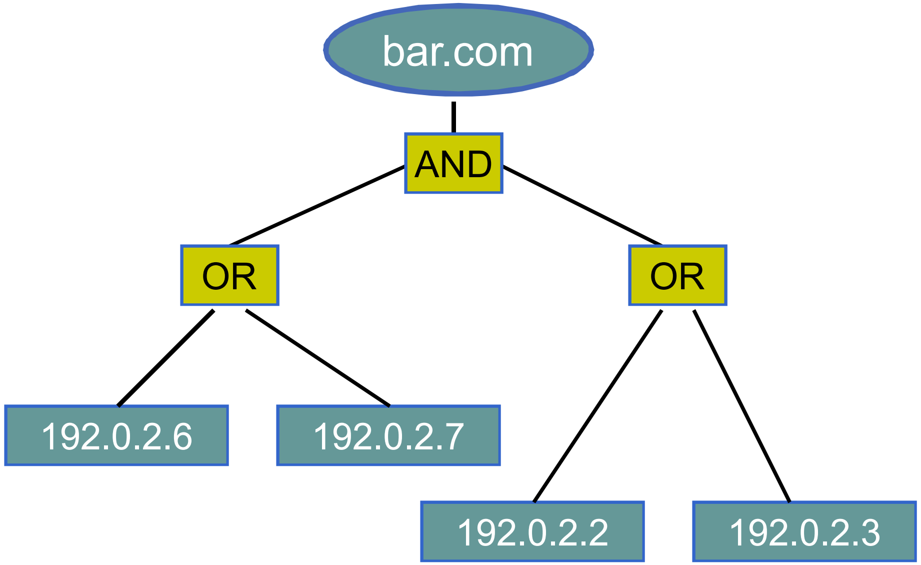

Redundancy(bar.com) = 2

(*not* false redundancy)

Delegation consistency

- NS RRs and glue records for a zone maintained *separately* in parent zone
- Potential problems:
  - NS RR mismatches - 587,865 zones (20%)
  - Missing glue records - 901 (0.024%)
  - Incorrect glue records - 108,737 zones (3.6%)

(Parent)

```
$ORIGIN com.
bar.com. NS ns.bar.com.
bar.com. NS me.baz.net.
ns.bar.com. A 192.0.2.2
```

(Authoritative)

```
$ORIGIN bar.com.
bar.com. NS ns.bar.com.
bar.com. NS me.baz.net.
ns.bar.com. A 192.0.2.2
```

MSQ

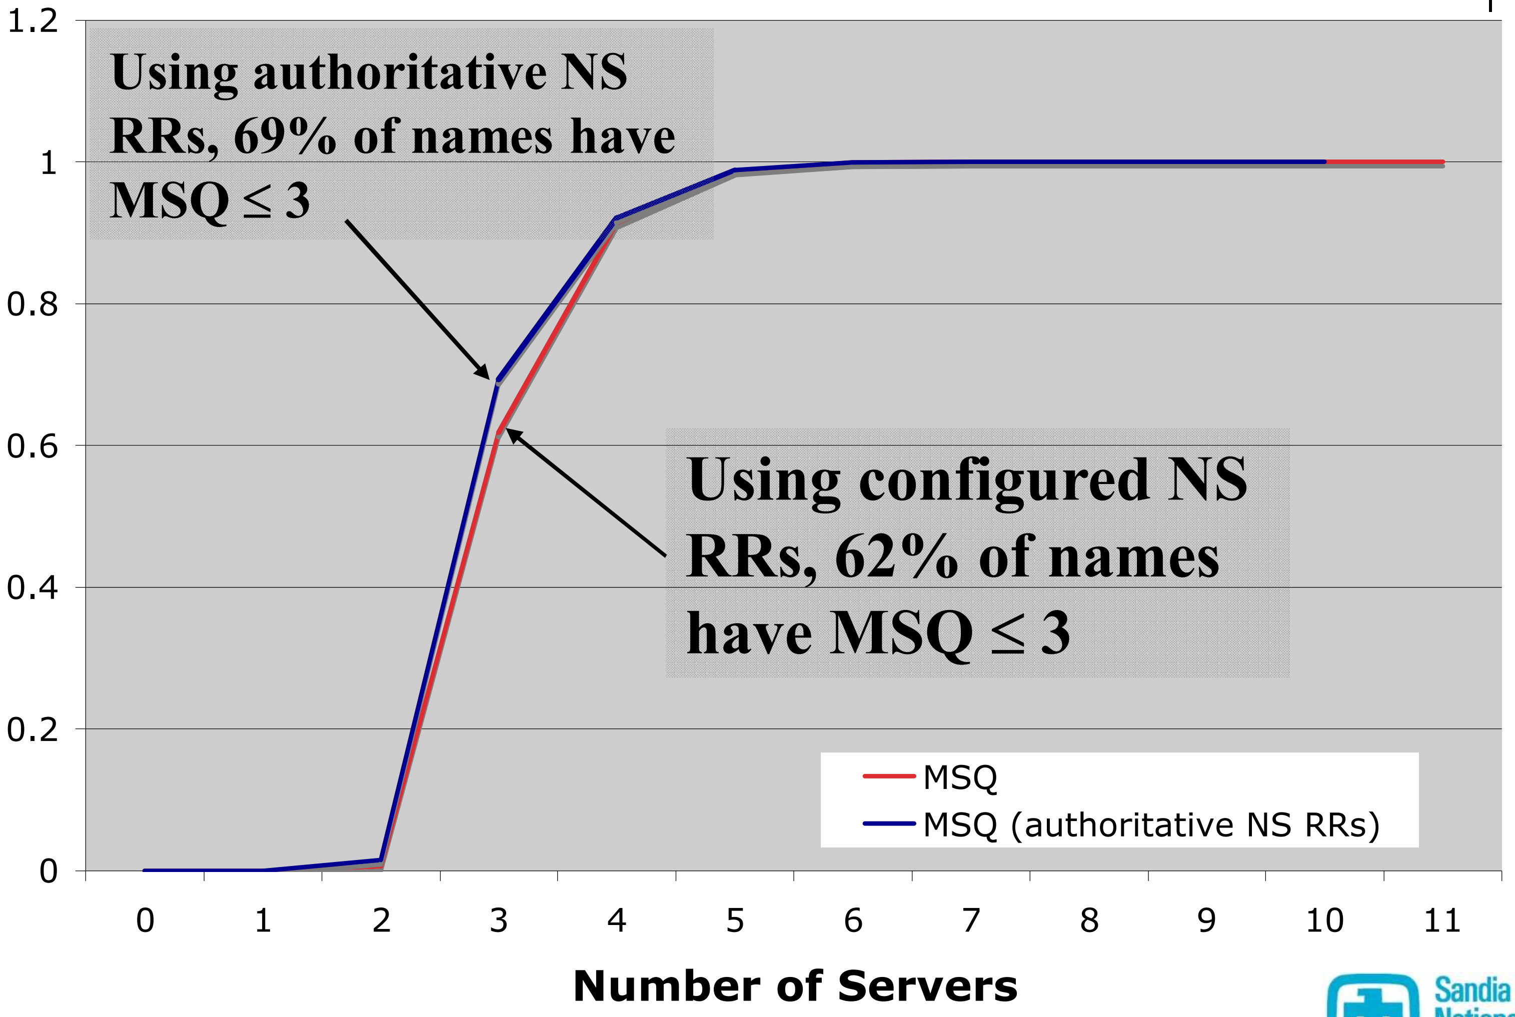

Using authoritative NS RRs, 69% of names have MSQ ≤ 3

Using configured NS RRs, 62% of names have MSQ ≤ 3

Redundancy

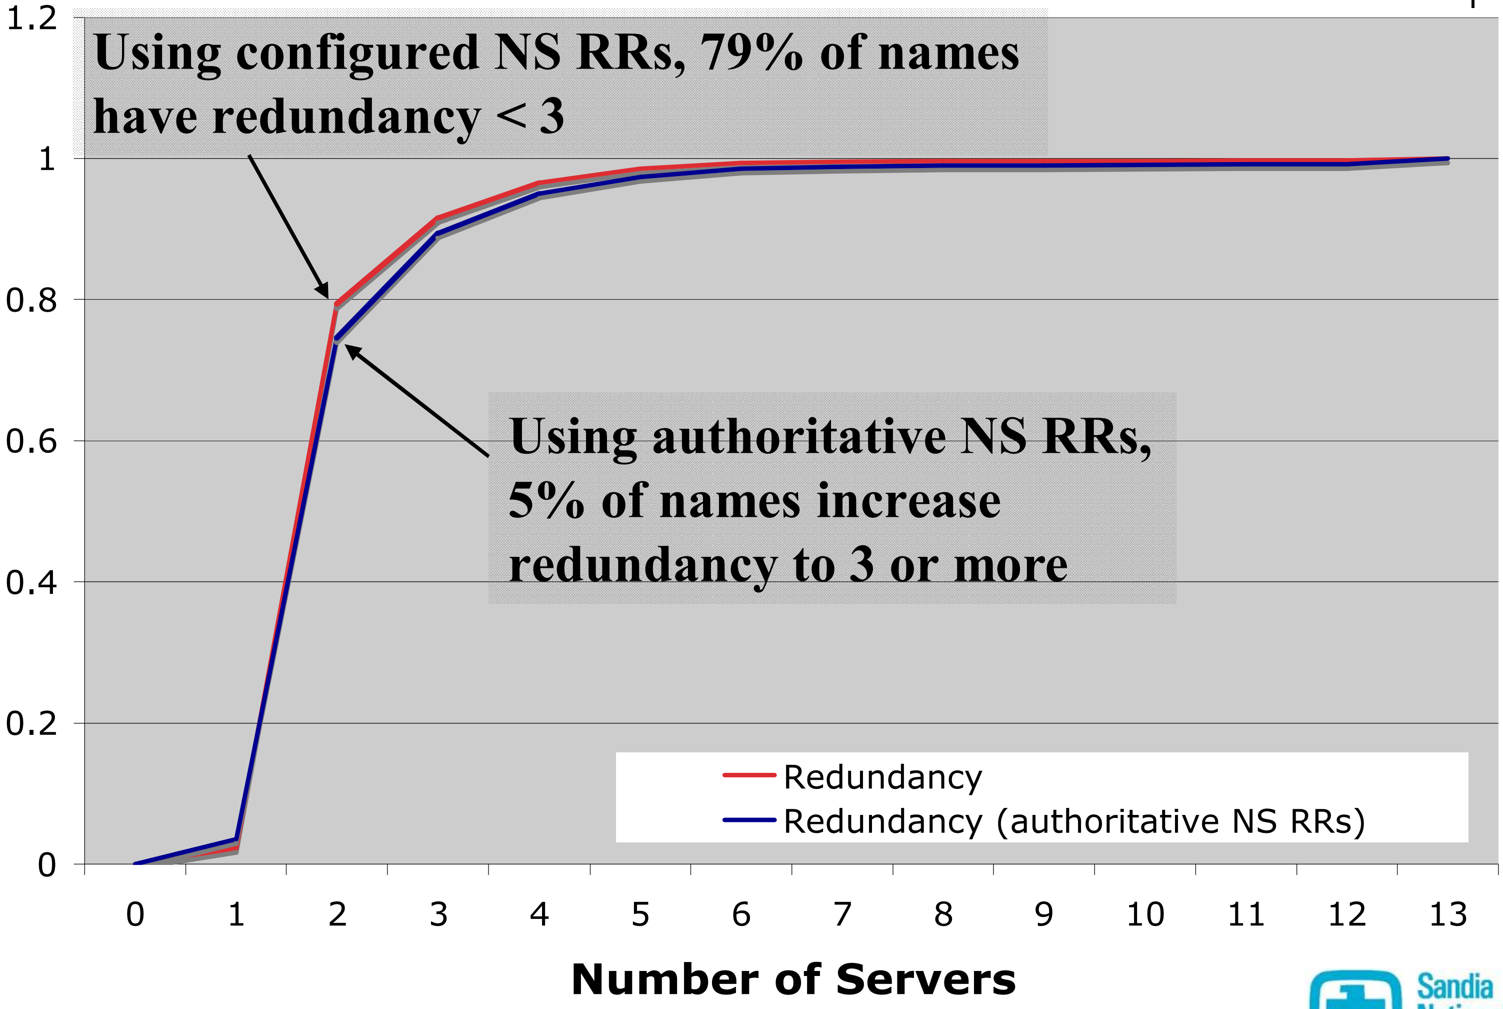

Using configured NS RRs, 79% of names have redundancy < 3

Using authoritative NS RRs, 5% of names increase redundancy to 3 or more

Lame delegation

- Symptoms
  - Non-responsive server designated as authoritative - 187,023 servers (2.5%)
  - Non-authoritative server designated as authoritative - 90,745 servers (1.2%)
- Causes
  - Delegation inconsistency
  - Misconfiguration on affected authoritative server
  - Outdated zone data
- Impact: potentially increases MSQ and decreases redundancy, decreasing availability and performance

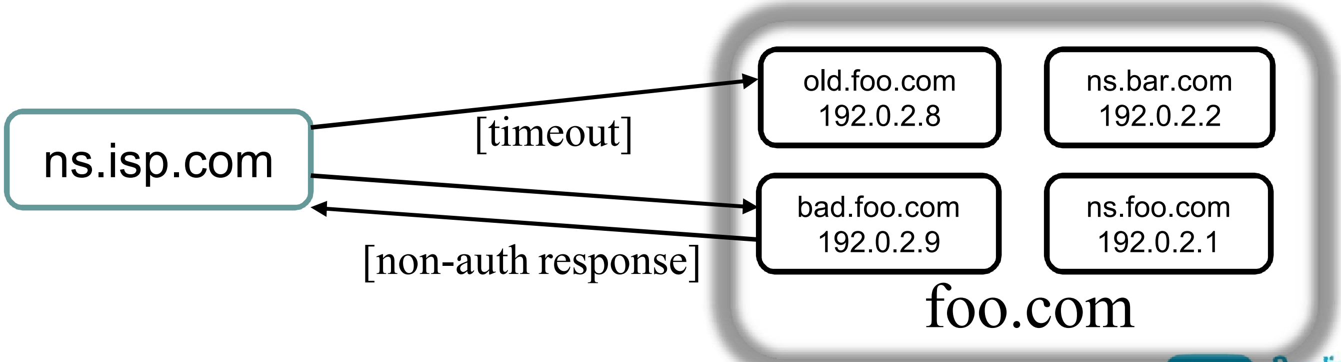

“Impact of Configuration Errors on DNS Robustness” [Pappas 2004]

Cyclic dependencies

- Symptom
  - Cycle in dependency graph - 2,835 zones (0.095%)
- Causes
  - Missing glue record - 76% of affected zones
  - Other cyclic dependency - 24% of affected zones
- Impact: Potentially decreases availability

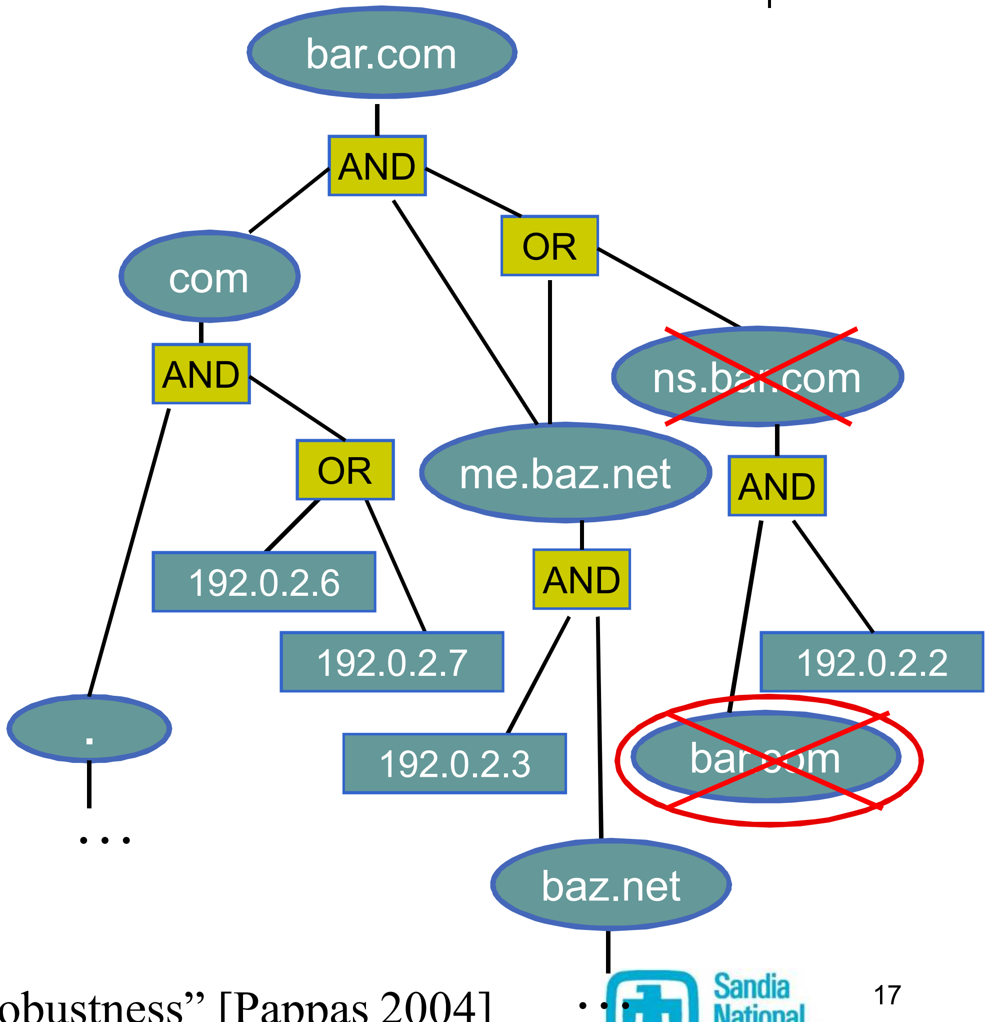

"Impact of Configuration Errors on DNS Robustness" [Pappas 2004]

Summary

- DNS availability model
  - Quantifies availability of domain names using:
    - Minimum servers queried (MSQ)
    - Redundancy
  - Quantifies impact of misconfigurations on availability

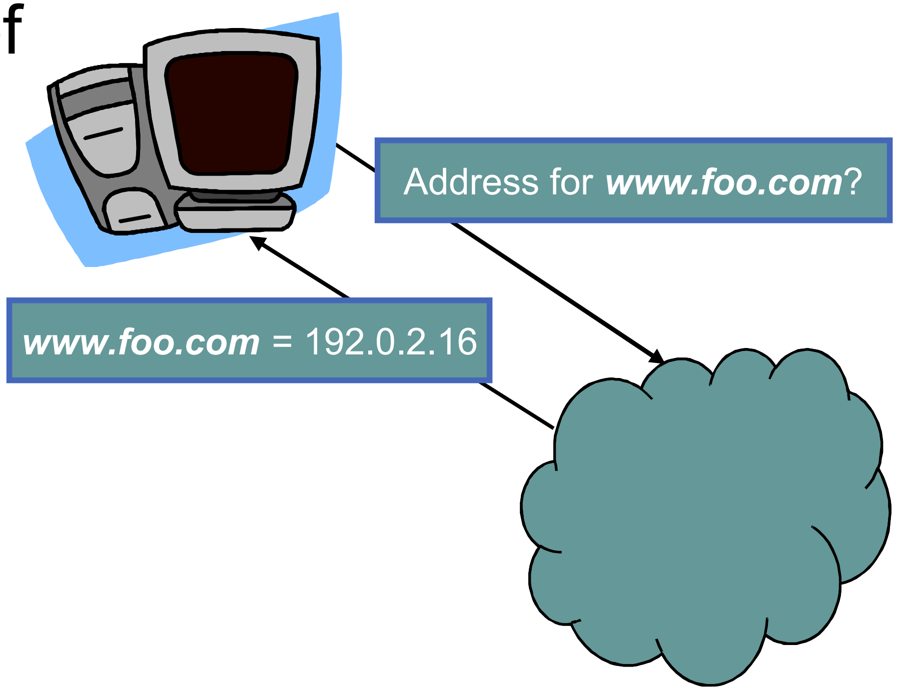

Address for ***www.foo.com***?

www.foo.com = 192.0.2.16

Data collection

- Extracted ~3 million names from Open Directory Project (dmoz.org)
- Collected additional 100,000 names from SC08
- Crawled dependencies of each name
- Resulting graph:
  - 8.4 million nodes
  - 22.3 million edges

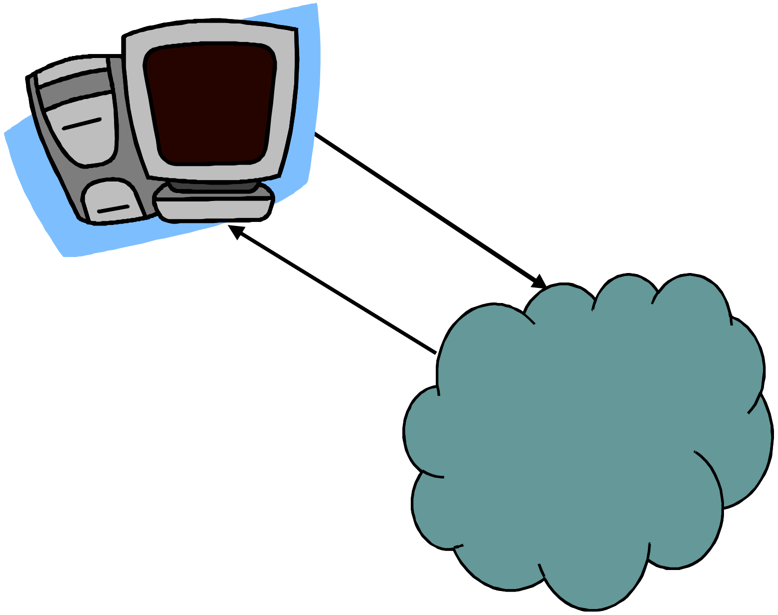

Future work

- Current availability model assumes that if a server provides an answer, then answer is correct
- Availability can be extended to include possibility for compromise
- DNSSEC availability
  - DNSSEC in “early adopter” deployment phase
  - Signed zone data has limited lifetime, requires regular maintenance, synchronization
  - Increased reliance on dependencies

Sandia National Laboratories