DNS Inconsistency

John Kristoff <jtk@depaul.edu>
Balajee Vamanan <bvamanan@uic.edu>
Chris Kanich <ckanich@uic.edu>

Introduction

- Parent NS RRset ***p***

```
example 1D  IN NS  ns1.example.edu.
example 1D  IN NS  ns2.example.edu.
```

- Child NS RCset ***c***

```
example 1D  IN NS  ns1.example.edu.
example 1D  IN NS  ns2.example.edu.
example 1D  IN NS  ns3.example.edu.
```

Real, Badly Inconsistent Example

- Child NS RRset ***cdm.depaul.edu***

```
cdm.depaul.edu.  3600 IN NS  ns1.cti.depaul.edu.
cdm.depaul.edu.  0    IN NS  shemp.cti.depaul.edu.
cdm.depaul.edu.  3600 IN NS  ns-colo.cti.deapaul.edu.
cdm.depaul.edu.  3600 IN NS  dc-colo-cti.cti.depaul.edu.
cdm.depaul.edu.  3600 IN NS  bach.cti.depaul.edu.
cdm.depual.edu.  3600 IN NS  ellington.cti.depaul.edu.
cdm.depaul.edu.  3600 IN NS  moe.cti.depaul.edu.
cdm.depaul.edu.  3600 IN NS  mozart.cti.depaul.edu.

. . .
ns-colo.cti.depaul.edu.     AAAA 2002:d8dc:b452::dbdc:b452
dc-colo-cti.cti.depaul.edu. A    10.128.30.2
```

Conjecture

DNS infrastructure (NS parent/child RRset) inconsistency arises from **asynchronous** and **uncoordinated** NS RRset **configuration**

Methodology: .edu traversal

- Obtain all .edu names using `whois  *` hack

```
for each name in edu
    mark root_servers as visited
    get NS_RRset from an .edu NS for the name
    for each S in NS_RRset
        do_query( name, S )

do_query:
    return if already queried S for name
    mark ( name, S ) as visited
    get new_NS_RRset for name from S
    for each S* in new_NS_RRset
        do_query( name, S* )
```

NS name mapping ambiguity

name	IPv4 address
name	IPv4 address set
name	IPv6 address
name	IPv6 address set
name	IPv4 address + IPv6 address
name	IPv4 address set + IPv6 address
name	IPv4 address + IPv6 address set
name	IPv4 address set + IPv6 address set

Evaluating Inconsistency

error	bad type (e.g. CNAME)	
error	bad rdata (e.g. Ipaddr for NS)	29 (0.01)
error	TTL disagreement in NS RRset	141 (0.06)
error	DNSSEC validation failure	
error	timeout/unreachable transient (e.g. down time)	
error	timeout/unreachable permanent (e.g. misconfiguration)	1403 (6)

query_response	NOERROR	21593 (90)
query_response	NXDomain	23 (0.01)
query_response	REFUSED	679 (3)
query_response	SERVFAIL / FORMERR / NOTIMP / ...	142 (0.06)
query_response	referral after a referral	77 (0.03)
query_response	aa==0 when aa==1 expected	977 (4)
query_response	malicious or incorrect data	

Parent/Child NS RRset Consistency

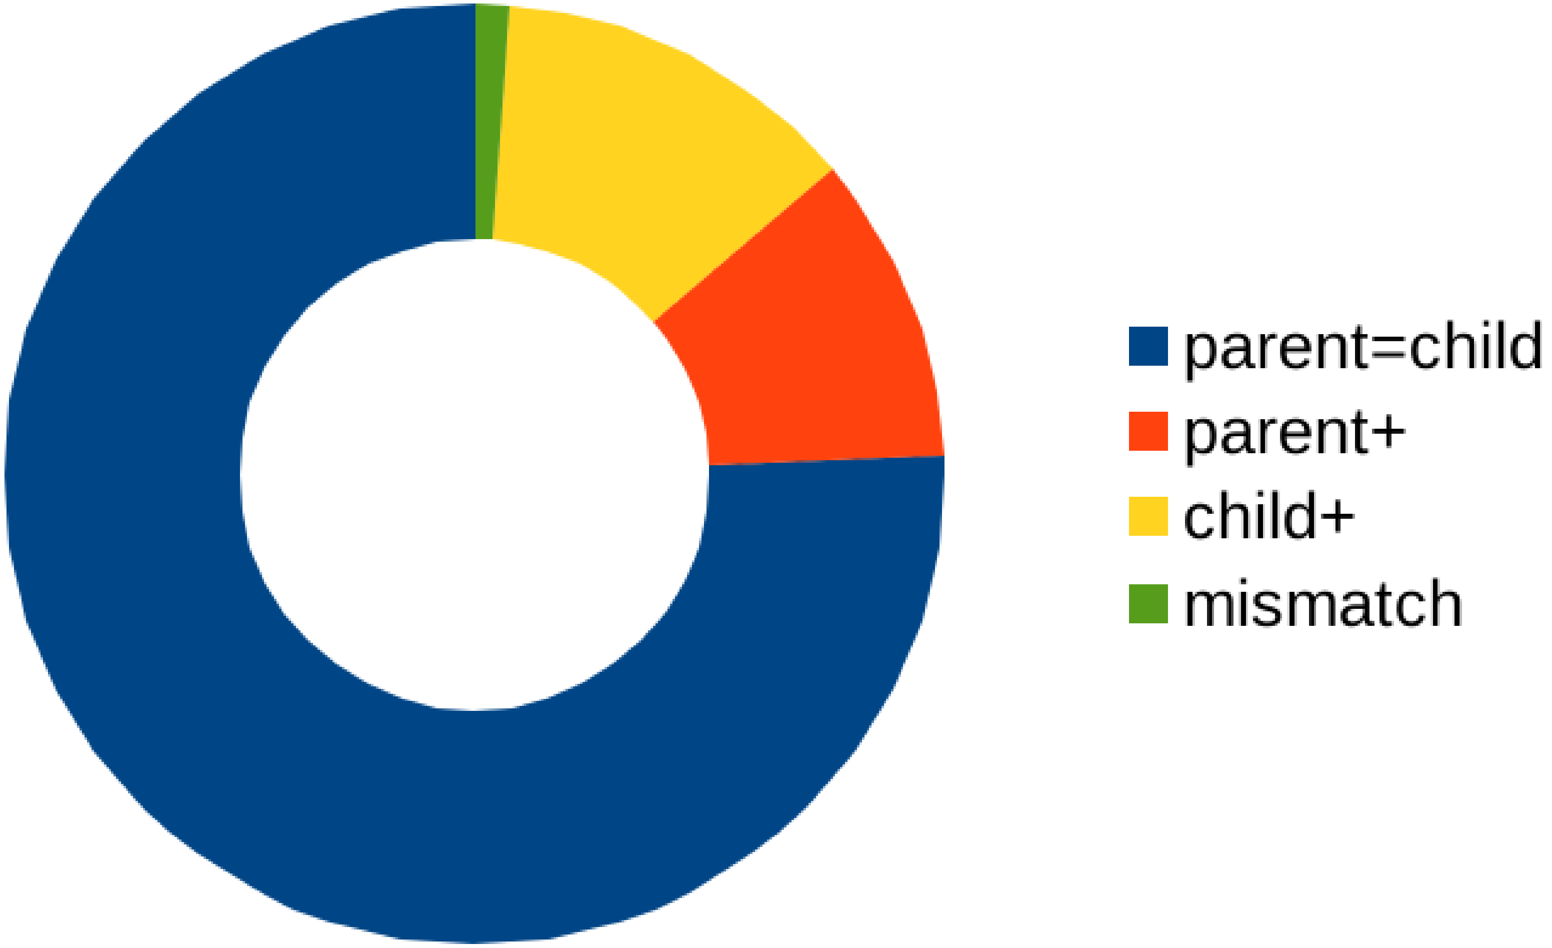

Namespace != Infrastructure Graph

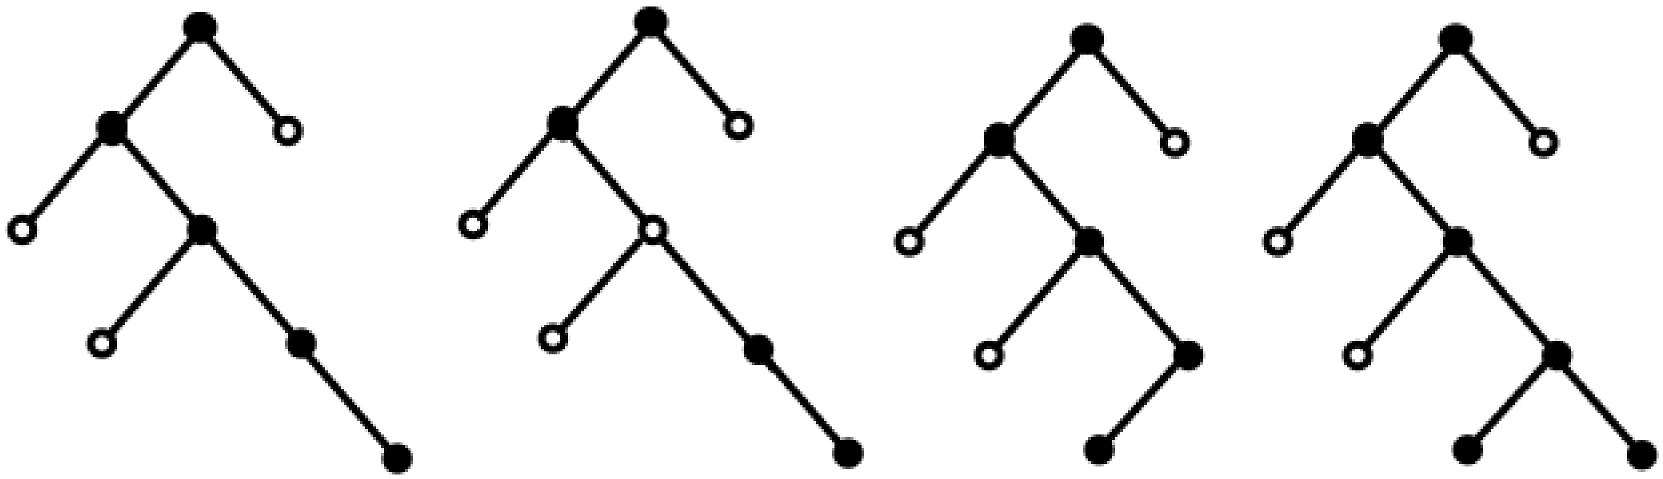

Resolver (in)Stability

	distribution	avoidance	recovery
BIND	proportional	no	< 1 sec
PowerDNS	spike dist.	no	3 min
Unbound	uniform	yes	15 min
DNSCache	uniform	no	< 1 sec
WindowsDNS	uniform	yes	1 sec

- Source: Yu et al., Authority server selection in DNS cachine resolvers, ACM SIGCOMM CCR 2012
- NOTE: negative caching => bursts of repeated failures

Discussion

- Inconsistency increases down the name space
- Inconsistency could exacerbate security threats
- Inconsistency may affect performance
- Inconsistency may lead to non-determinism

Questions

- Are some NS infrastructure graphs unknowable?
- Should consistency be encouraged? If so, how?
- There is no up/down sync, should there be?
- Should minimal-responses be preferred?
- Should repeated failures influence retry algorithms?
- Should NS RRs have had IPaddrs as RDATA?
- Is inconsistency worth studying further?