

# WHEN PARENTS AND CHILDREN DISAGREE: DIVING INTO DNS DELEGATION INCONSISTENCY

Raffaele Sommese<sup>1</sup>, Giovane C. M. Moura<sup>2</sup>, Mattijs Jonker<sup>1</sup>,  
Roland van Rijswijk-Deij<sup>1</sup>, Alberto Dainotti<sup>3</sup>, K Claffy<sup>3</sup>, Anna Sperotto<sup>1</sup>

<sup>1</sup>University of Twente, <sup>2</sup>SIDN Labs, <sup>3</sup>CAIDA/UCSD



The Domain Name System (DNS) is one of the most critical components of the Internet



DNS is a distributed, hierarchical database



DNS maps hosts, services and applications to IP addresses and various other types of records.

# INTRODUCTION



A key mechanism that enables the DNS to be hierarchical and distributed is delegation



The DNS hierarchy is organized in parent and child zones typically managed by different entities



Different zones need to share common information (NS records) about which are the authoritative name servers for a given domain.

# DNS AND DELEGATIONS



RFC1034 states that the NS records at both parent and child **should** be “consistent and remain so”



Is this in practice the case?

IS COMMON INFORMATION  
CONSISTENT?

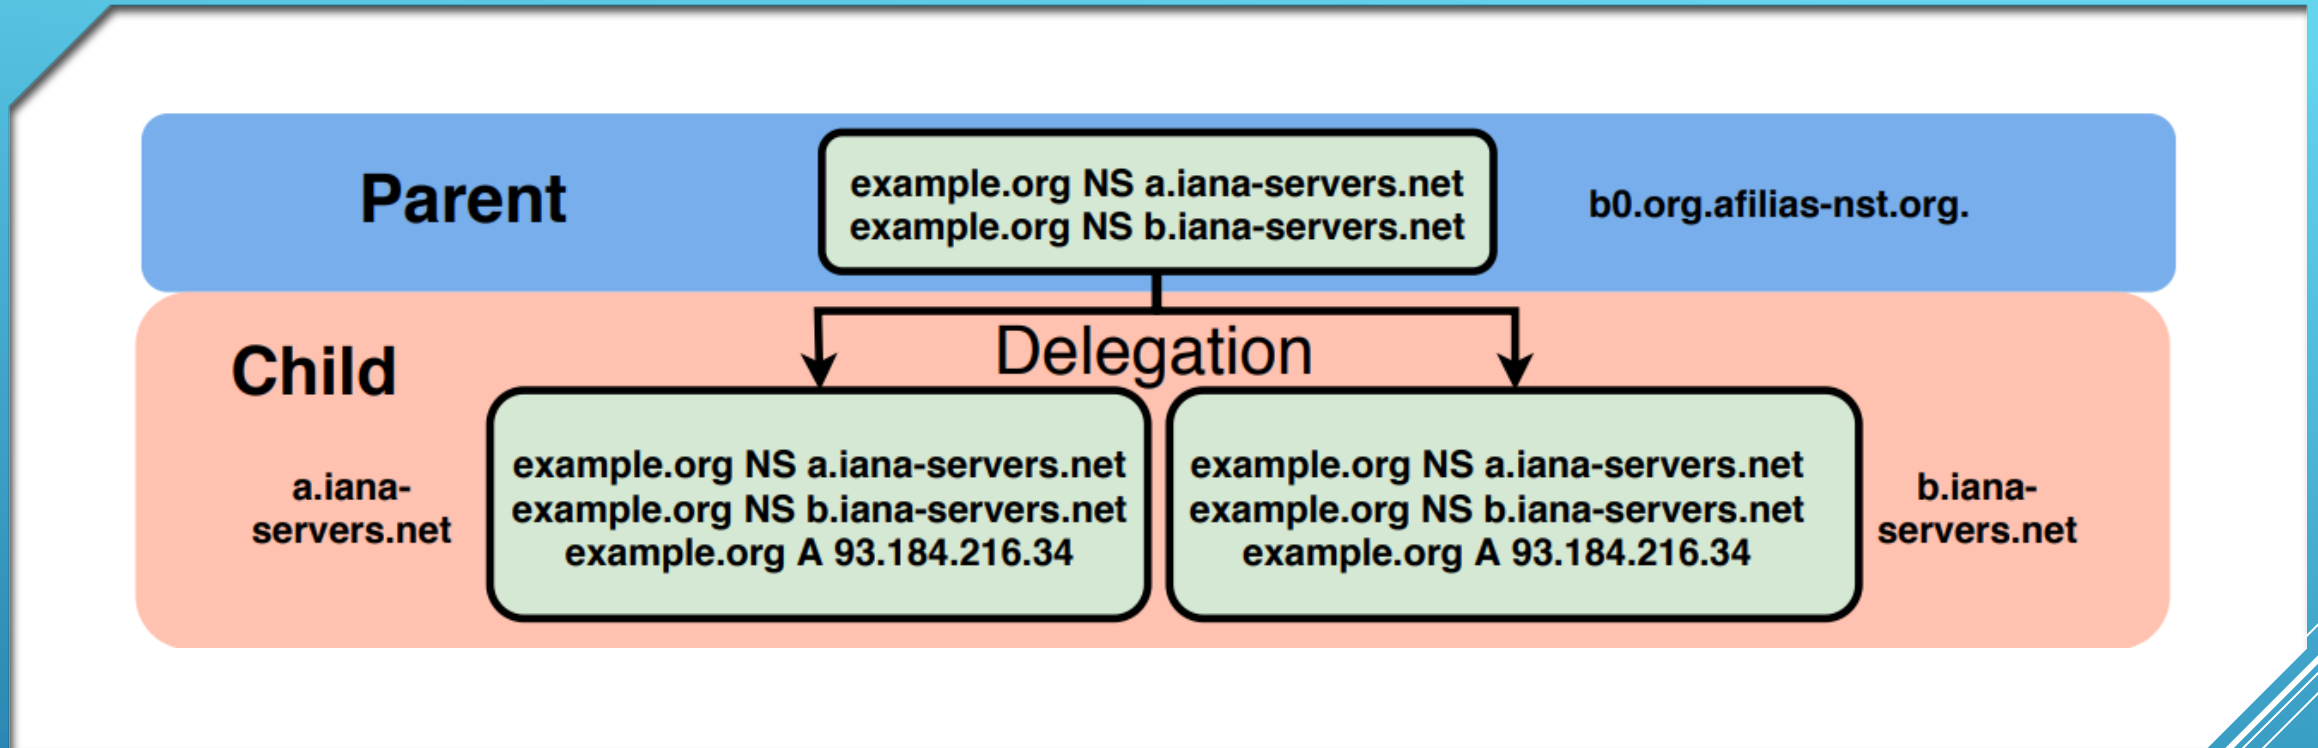


Provide a broad characterization  
of inconsistencies in DNS  
delegations



Investigate the practical  
consequences of these  
inconsistencies.

# OUR CONTRIBUTION



A WELL CONFIGURED DELEGATION

- ▶ We study delegation consistency between parent (TLD) and child (SLD) zones for all active second-level domain names of .com, .net, and .org.
- ▶ We analyse more than 166M domain names (50% of the DNS namespace)
- ▶ 80% of these domain names exhibit consistency.
- ▶ 8% (13 million domains) DO NOT!

ARE THE DOMAINS IN THE DNS  
WELL CONFIGURED?

- ▶ We study delegation consistency between parent (TLD) and child (SLD) zones for all active second-level domain names of .com, .net, and .org.
- ▶ We analyse more than 166M domain names (50% of the DNS namespace)
- ▶ 80% of these domain names exhibit consistency.
- ▶ 8% (13 million domains) DO NOT!

ARE THE DOMAINS IN THE DNS  
WELL CONFIGURED?



01

Parent and children have a disjoint NSSet

02

Parent NSSet is a subset of children NSSet

03

Parent NSSet is a superset of children NSSet

04

Parent and children NSSet have some common elements and some different elements.

WHICH KIND OF INCONSISTENCY WE FOUND?

- ▶ In 55% of domains with delegation inconsistency, parents and children has a disjoint NSSet.
- ▶ Half of these domains are consistent at IP level
- ▶ Half are NOT!
- ▶ 16 TLDs present this inconsistency in the root zone, but all are consistent at IP level.


b0.org.afilias-nst.org (.org Auth NS) - Parent

example.org.	86400	IN	NS	a.iana-servers.net.
example.org.	86400	IN	NS	b.iana-servers.net.

a.iana-servers.net. (example.org Auth NS) - Child

example.org.	86400	IN	NS	c.iana-servers.net.
example.org.	86400	IN	NS	d.iana-servers.net.

# PARENT AND CHILDREN HAVE A DISJOINT NSSET

- 
- ▶ Different servers, which could be lame delegation.
  - ▶ Even if IP level is coherent, keep A records in sync makes misconfiguration easy.
  - ▶ Behaviour of resolver is not predictable!

## DISJOINT NSSET CONSEQUENCES

India's .in registry had ns[1–6].neustar.in as NS records at the parent (Root), and [ns1-ns6].registry.in at the child.

Both NSSets pointed to the same A/AAAA records.

On 2019-10-30 we notified them and on 2019-11-02 they fixed the inconsistency.

15 other internationalized ccTLDs run by India had the same issue with their NSSet, and were also fixed

## INDIA'S .IN REGISTRY

- ▶ In 30% of domains with delegation inconsistency, parent NS-Set is a subset of children NS-Set.
- ▶ 18 TLDs present this inconsistency in the root zone.

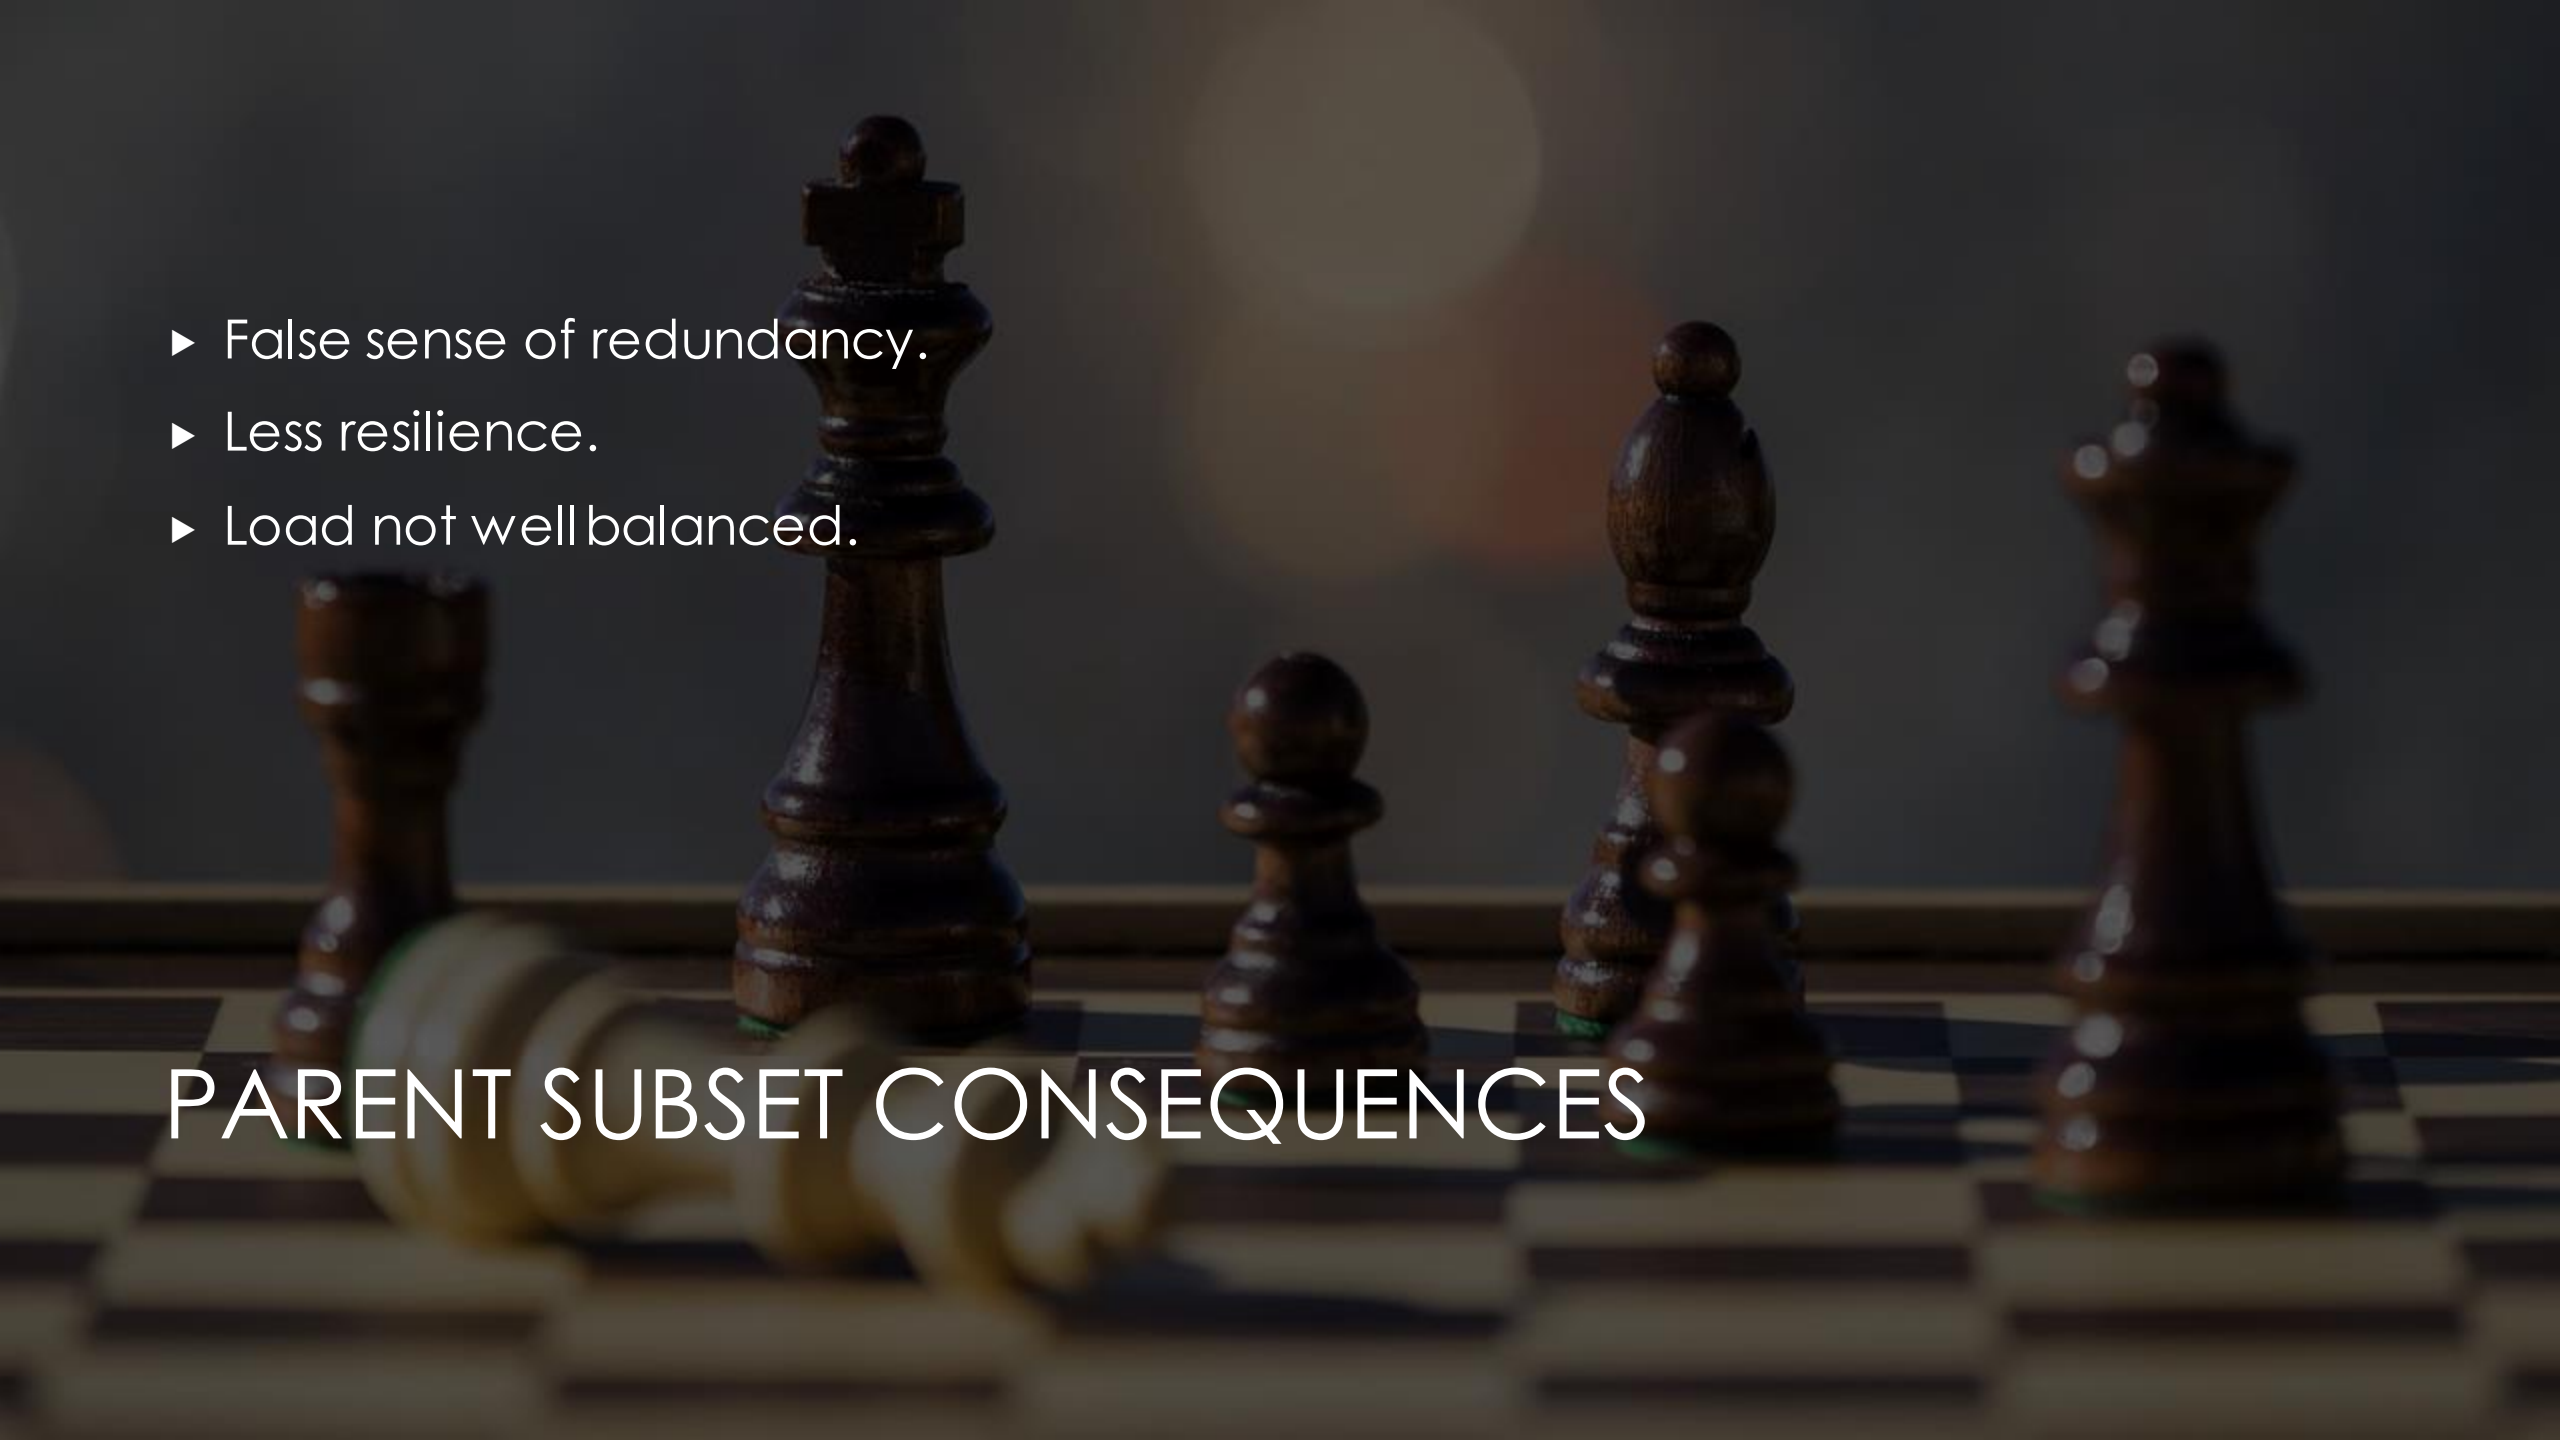
b0.org.afiliast.org (.org Auth NS) - Parent

example.org.	86400	IN	NS	a.iana-servers.net.
example.org.	86400	IN	NS	b.iana-servers.net.

a.iana-servers.net. (example.org Auth NS) - Child

example.org.	86400	IN	NS	a.iana-servers.net.
example.org.	86400	IN	NS	b.iana-servers.net.
example.org.	86400	IN	NS	c.iana-servers.net.
example.org.	86400	IN	NS	d.iana-servers.net.

# PARENT NSSET IS A SUBSET OF THE CHILDREN NSSET

- 
- ▶ False sense of redundancy.
  - ▶ Less resilience.
  - ▶ Load not well balanced.

## PARENT SUBSET CONSEQUENCES

AT&T's main domain att.com had a parent NSSet containing [ns1...ns3].attdns.com, whereas the child had [ns1...ns4].attdns.com.

We notified AT&T of this misconfiguration.

On 24/10/2019 the issue was resolved and the fourth name server (ns4.attdns.com) was also added to the parent

## AT&T CASE

- ▶ In 8% of domains with delegation inconsistency, parent NS-Set is a superset of children NS-Set.
- ▶ 10 TLDs present this inconsistency in the root zone.

b0.org.afilias-nst.org (.org Auth NS) – Parent

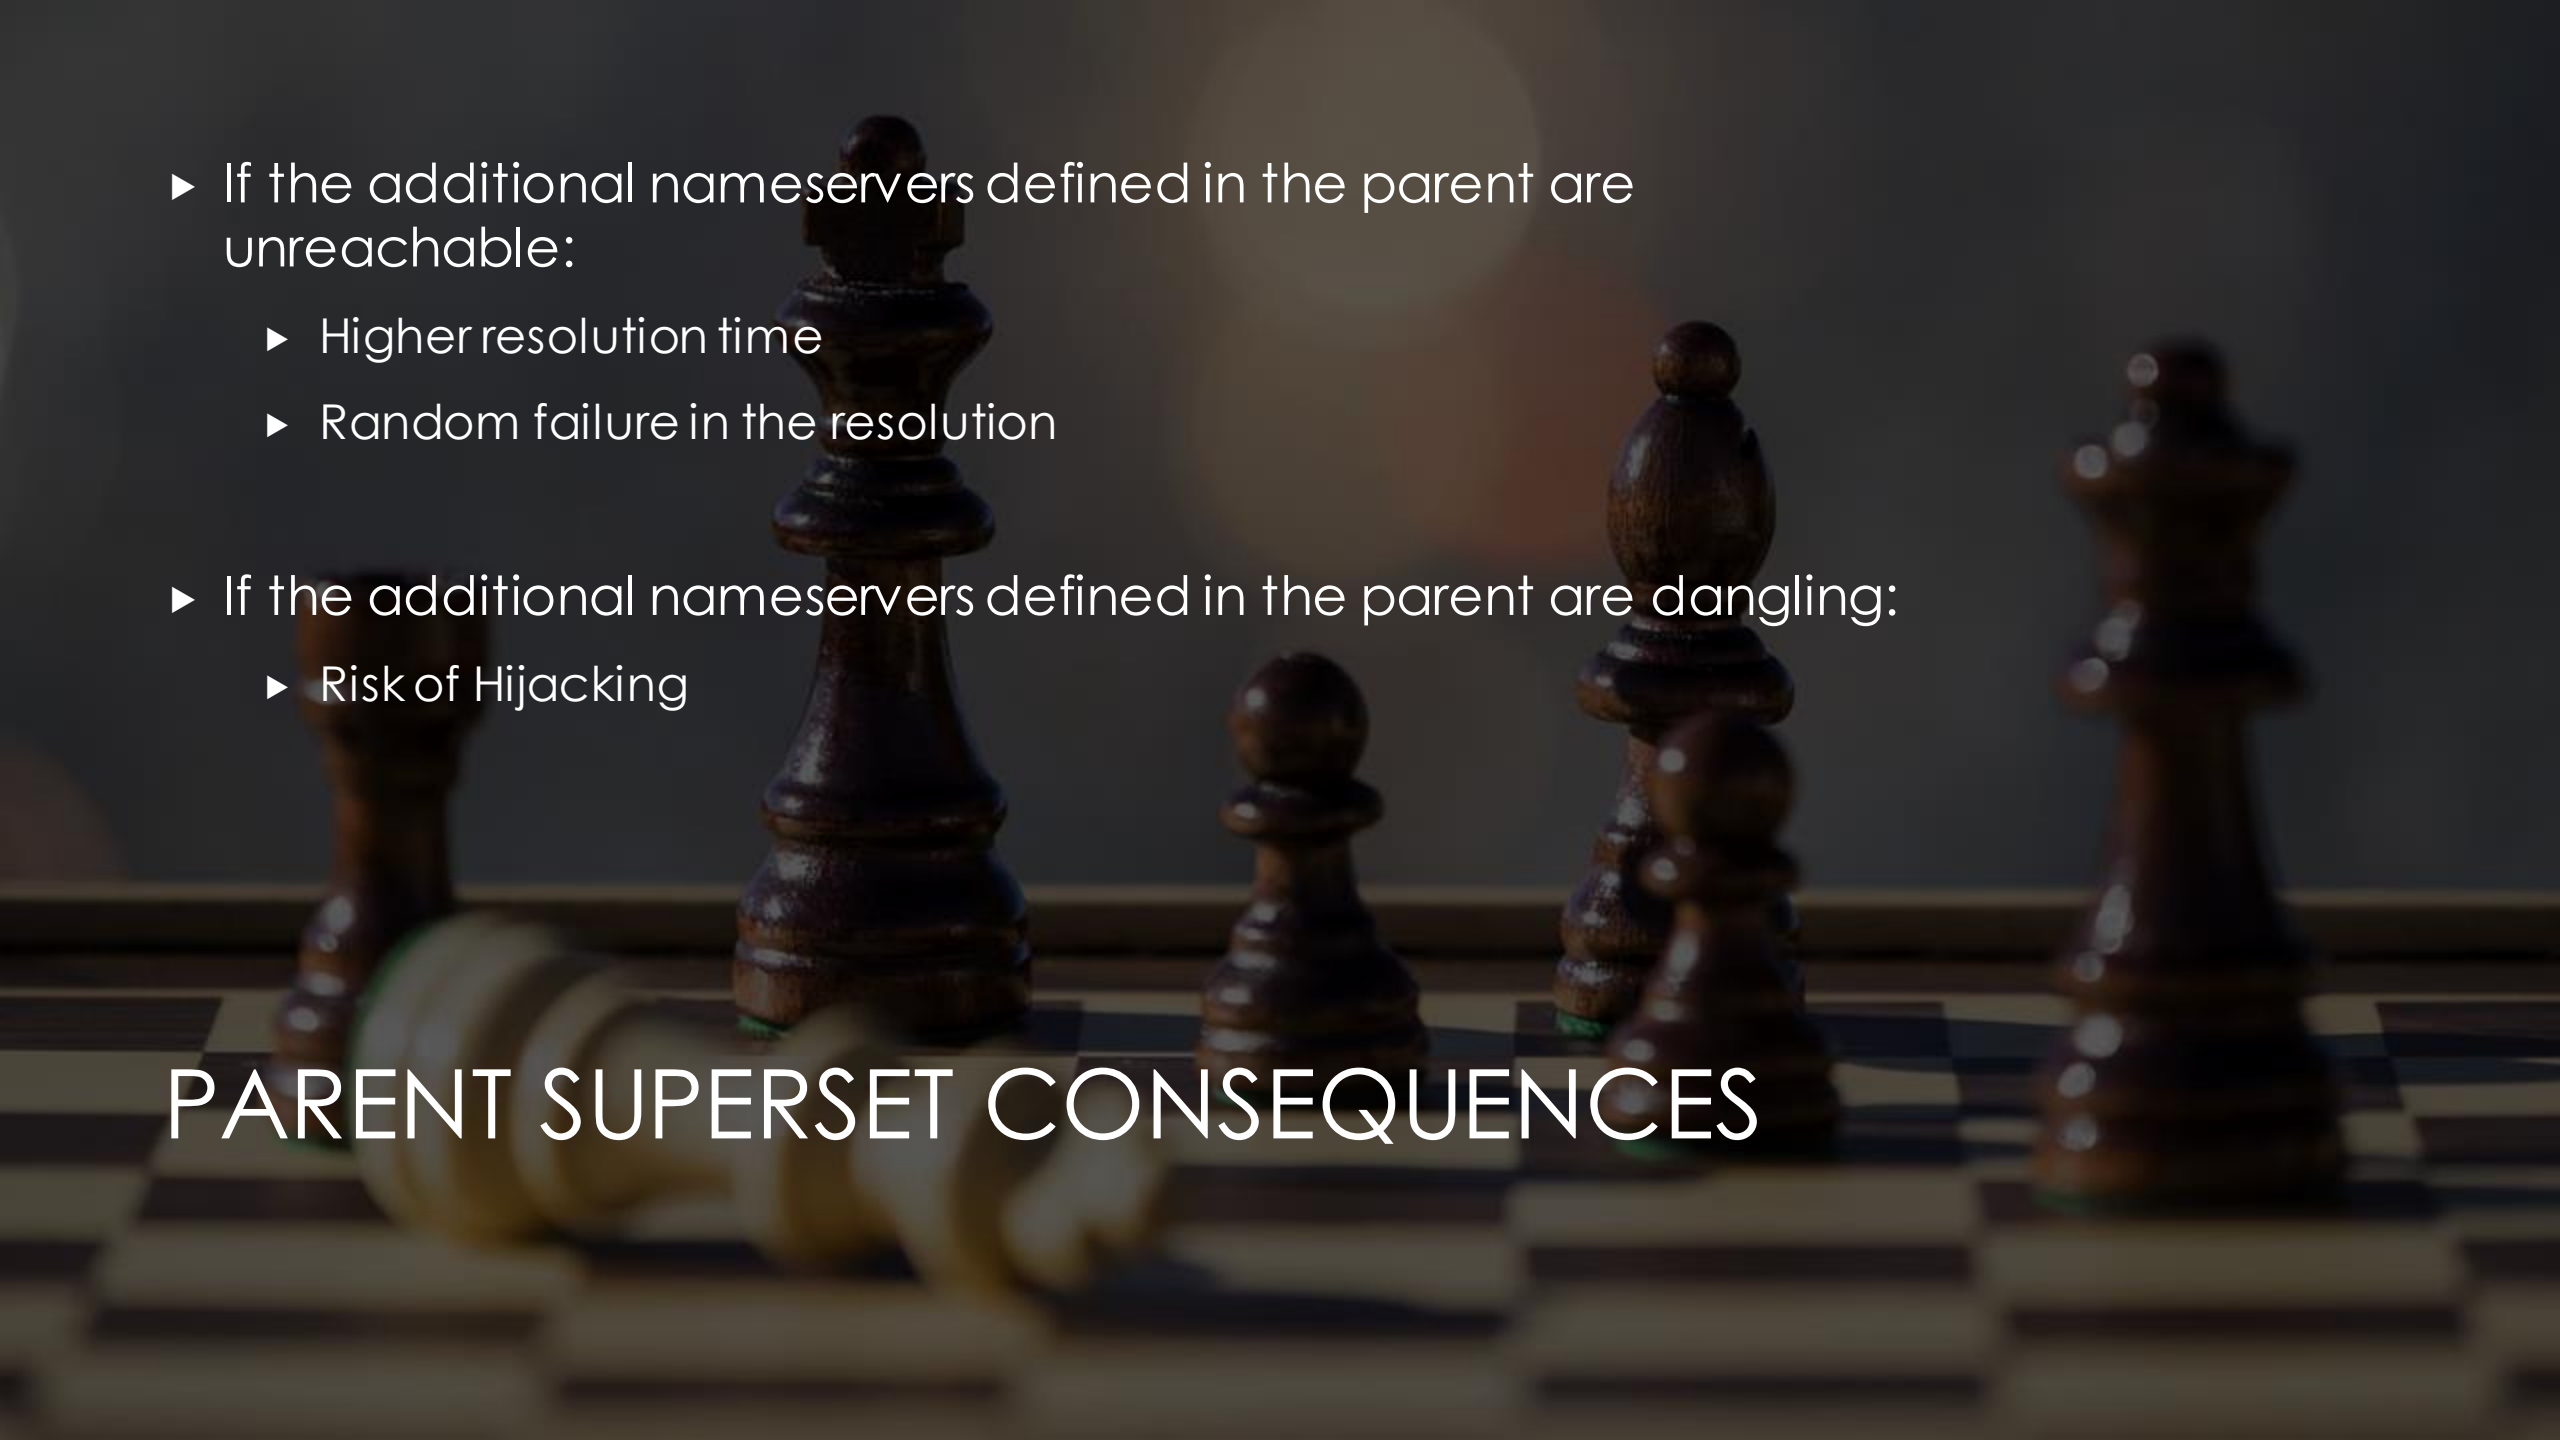
example.org.	86400	IN	NS	a.iana-servers.net.
example.org.	86400	IN	NS	b.iana-servers.net.
example.org.	86400	IN	NS	c.iana-servers.net.
example.org.	86400	IN	NS	d.iana-servers.net.

a.iana-servers.net. (example.org Auth NS) - Child

example.org.	86400	IN	NS	a.iana-servers.net.
example.org.	86400	IN	NS	b.iana-servers.net.

# PARENT NSSET IS A SUPERSET OF THE CHILDREN NSSET



- 
- ▶ If the additional nameservers defined in the parent are unreachable:
    - ▶ Higher resolution time
    - ▶ Random failure in the resolution
  - ▶ If the additional nameservers defined in the parent are dangling:
    - ▶ Risk of Hijacking

## PARENT SUPERSSET CONSEQUENCES

- ▶ In 7% of domains with delegation inconsistency, Parent and children NSSet have some common elements and some different elements.
- ▶ 8 TLDs present this inconsistency in the root zone.
- ▶ All risk and consequences mentioned before are applicable.

b0.org.afilias-nst.org (.org Auth NS) - Parent

example.org.	86400	IN	NS	a.iana-servers.net.
example.org.	86400	IN	NS	b.iana-servers.net.
example.org.	86400	IN	NS	c.iana-servers.net.

a.iana-servers.net. (example.org Auth NS) - Child

example.org.	86400	IN	NS	a.iana-servers.net.
example.org.	86400	IN	NS	b.iana-servers.net.
example.org.	86400	IN	NS	d.iana-servers.net.

# REST CATEGORY



We investigate the consequences of such inconsistencies, by emulating the four categories of NSSet mismatches.



We use RIPE Atlas, measuring each unique resolver as seen from their probes physically distributed around the world (3.3k ASes).



Our goal is to study these consequences in terms of query load distribution in a controlled environment, where the authoritative name servers are in the same network

# IMPLICATIONS OF NSSET INCONSISTENCY IN THE WILD



**RIPE NCC**  
RIPE Atlas

```
:: QUESTION SECTION:
;example.org.          IN      A

:: ANSWER SECTION:
example.org.          16807  IN      A      93.184.216.34

:: AUTHORITY SECTION:
iana-servers.net.    1800  IN      NS      a.iana-servers.net.
iana-servers.net.    1800  IN      NS      b.iana-servers.net.
iana-servers.net.    1800  IN      NS      c.iana-servers.net.
iana-servers.net.    1800  IN      NS      ns.icann.org.


:: ADDITIONAL SECTION:
a.iana-servers.net.  1800  IN      A      199.43.135.53
a.iana-servers.net.  1800  IN      AAAA   2001:500:8f::53
b.iana-servers.net.  1800  IN      A      199.43.133.53
b.iana-servers.net.  1800  IN      AAAA   2001:500:8d::53
c.iana-servers.net.  1800  IN      A      199.43.134.53
c.iana-servers.net.  1800  IN      AAAA   2001:500:8e::53
```

```
:: QUESTION SECTION:
;example.org.          IN      A

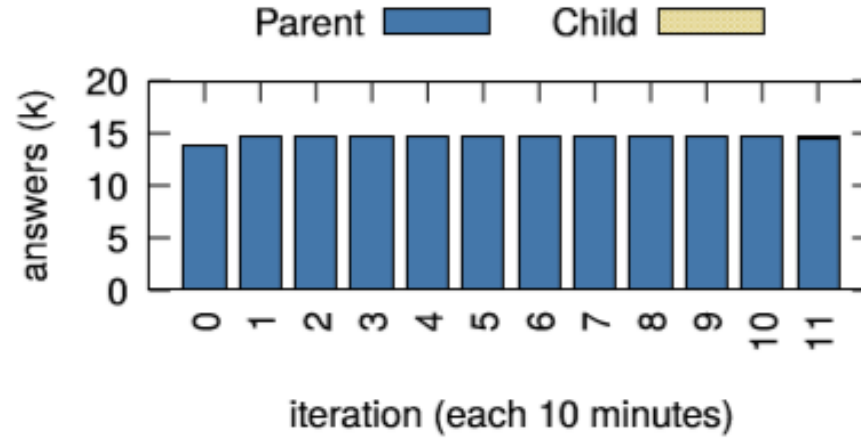
:: ANSWER SECTION:
example.org.          16807  IN      A      93.184.216.34

:: Query time: 31 msec
:: SERVER: 8.8.4.4#53(8.8.4.4)
:: WHEN: Mon Mar 23 16:07:23 CET 2020
:: MSG SIZE rcvd: 56
```

# MINIMAL RESPONSES




(a) Results for normal responses




(b) Results with minimal responses

# DISJOINT NSSET EXPERIMENTS




(a) Results for normal responses

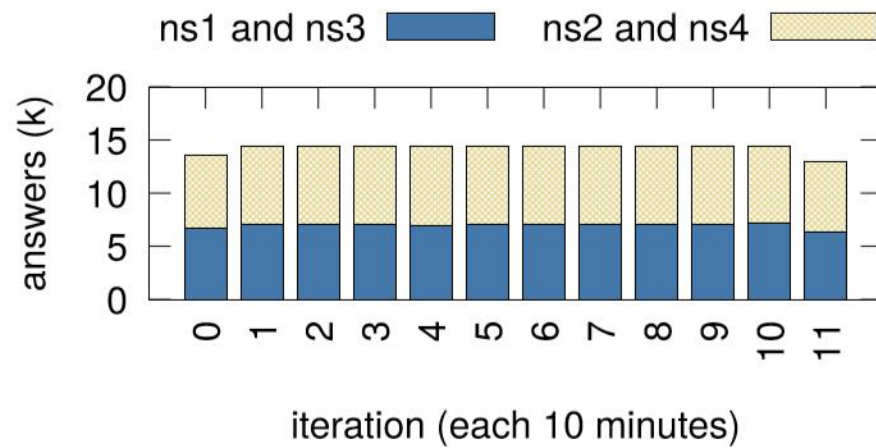


(b) Results with minimal responses

# SUBSET NS SETS EXPERIMENTS

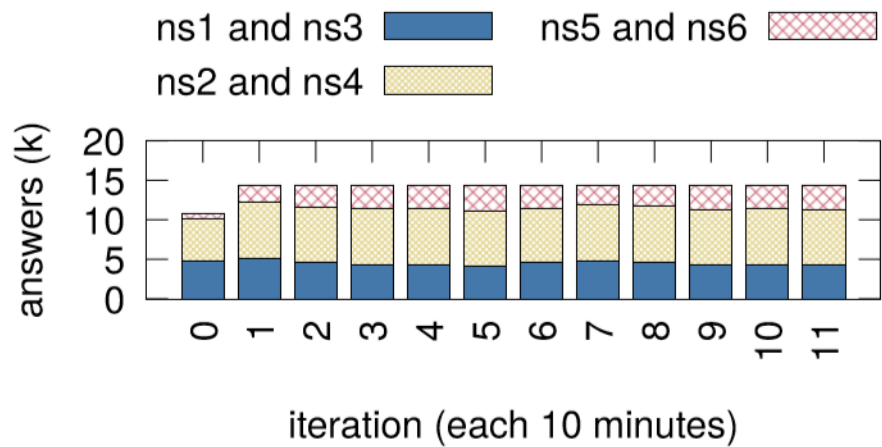


(a) Results for normal responses

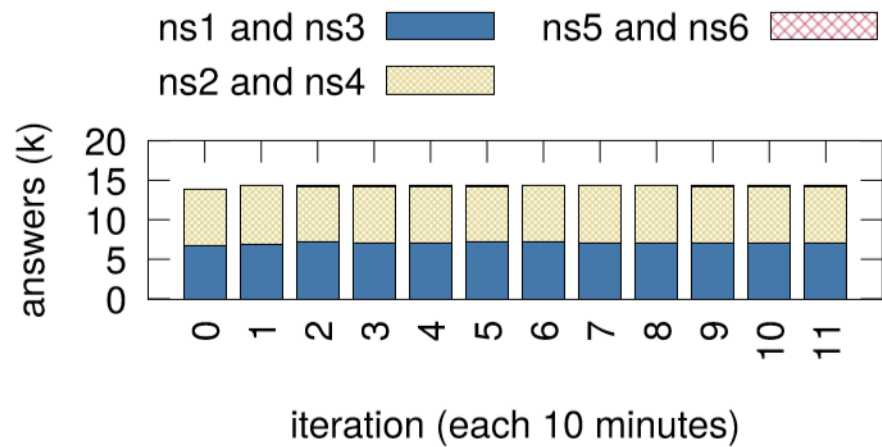


(b) Results with minimal responses

# SUPERSET NS SETS EXPERIMENTS



(a) Results for normal responses



(b) Results with minimal responses

# REST NS SETS EXPERIMENTS



- ▶ Having inconsistent NSSets in parent and child authoritative servers impacts how queries are distributed among name servers.
- ▶ For all evaluated cases, queries will be unevenly distributed among authoritative servers.
- ▶ The servers listed at the parent zone will receive more queries than the ones specified in the child.
- ▶ Minimal responses has an impact on resolver behaviour.

## CONSEQUENCES

- ▶ We focus on evaluating specific DNS resolver software to understand how they behave in case of NS-Set Inconsistency.
- ▶ We pay attention as to whether resolvers follow RFC2181, which specifies how resolvers should rank data in case of inconsistency.
- ▶ The RFC states that child authoritative data should be preferred.
- ▶ We evaluate four popular DNS resolver implementations: BIND, Unbound, Knot, PowerDNS and Windows.

# RESOLVER SOFTWARE EVALUATION

- i. We ask the resolver for an A record of a subdomain in our test zone
- ii. We ask for the NS record of the zone
- iii. We se ask first an A query followed by an NS query, to understand if resolvers use non-authoritative cached NS information to answer to the following query violating §5.4.1 of RFC2181
- iv. We invert this order to understand if authoritative record are overwritten by non-authoritative ones in the cache.

## FOUR TESTS

Knot and Unbound comply with RFC2181 ranking specification.

In (i) BIND packaged for Ubuntu did not: it caches only information from the parent and does not override it with information from the child.

In (i) and (iii), BIND from source sends the parent an explicit NS query before performing the A query.

In (iii) PowerDNS packaged for CentOS 6 and Ubuntu Xenial, and Windows (all) use the cached non-authoritative information to answer the NS query in the test, not conforming to RFC2181.

## RESULTS

- ▶ The problem of Parent-Child consistency is addressed in RFC7477.
- ▶ RFC7477 introduces a method to automatically keep records in the parent in sync
- ▶ The sync is performed through a periodical polling of the child using SOA records and a new type of record (CSYNC).
- ▶ Unfortunately, RFC7477 lacks deployment.

## RFC 7477 CHILD-TO-PARENT SYNCHRONIZATION IN DNS

- ▶ RFC1034 states that the NS records at both parent and child should be “consistent and remain so”
- ▶ We discover a significant part of the namespace is misconfigured and this has consequences for the resolution process.
- ▶ We strongly advise operators to verify their zones and follow RFC1034 and to consider supporting CSYNC DNS records.
- ▶ We also recommend that resolver vendors conform to the authoritative information ranking in RFC2181.

## CONCLUSION



QUESTIONS?