

Network Security and Measurement

- DNS Measurements -

Prof. Dr. Thomas Schmidt

<http://inet.haw-hamburg.de> | t.schmidt@haw-hamburg.de

Agenda

How can we measure the DNS?

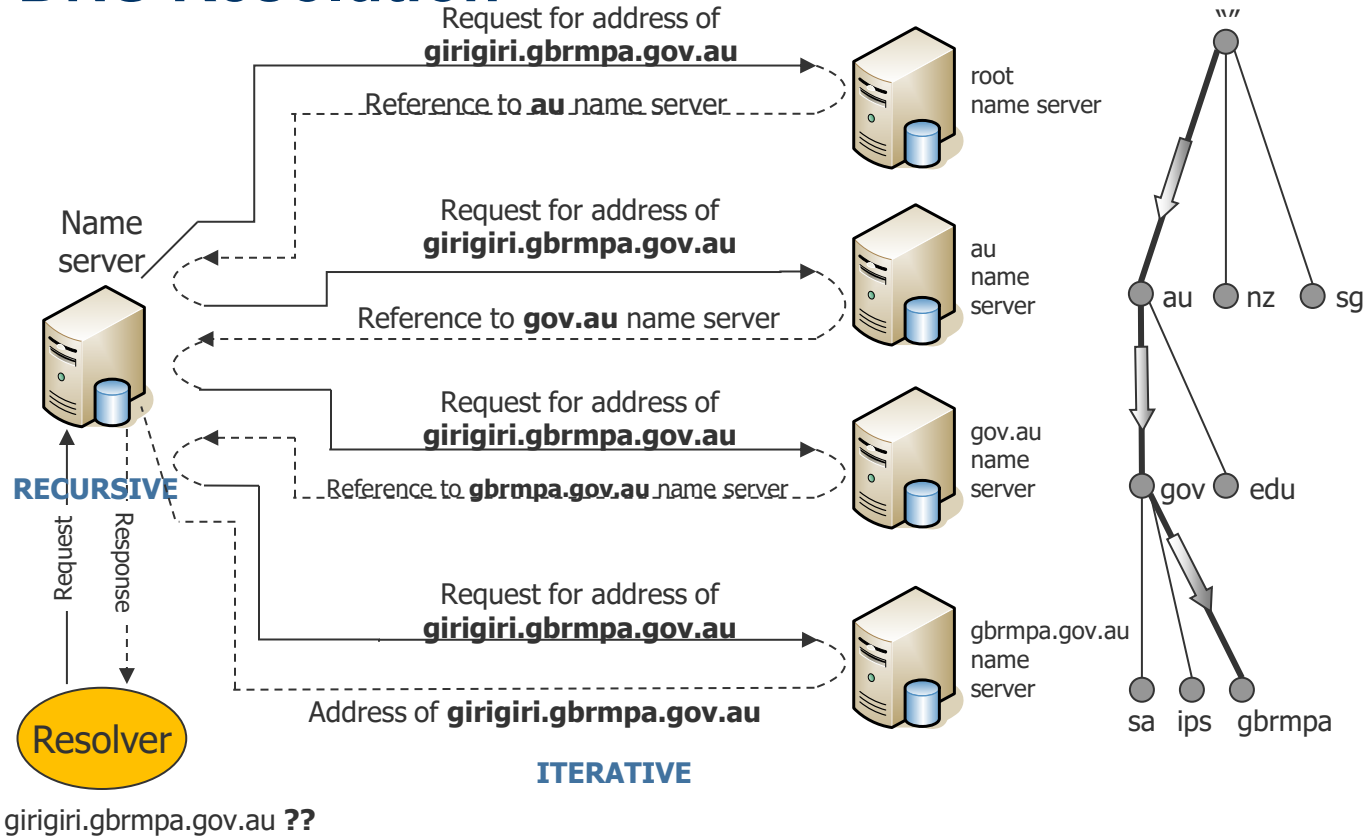
How should we design an active DNS measurement infrastructure?

How can you measure DNS impact?
Hijacking Internet resources from
expired DNS domains

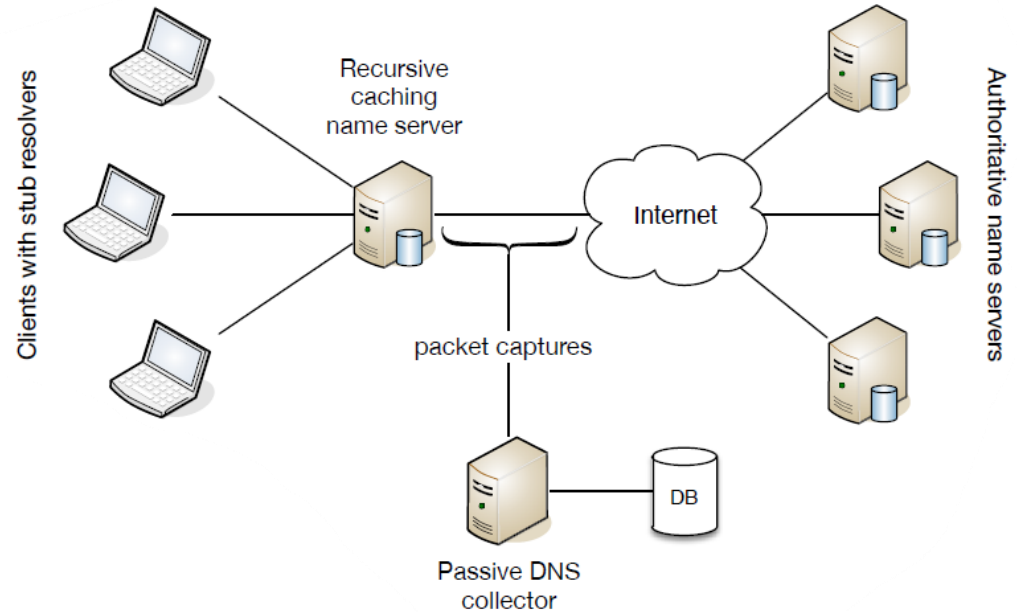
Technical Challenge

MEASURING THE DNS

Recap DNS Resolution



Passive DNS measurements: Typical setup




Passive DNS measurements: Typical setup

Examples: dnsdb.info and pDNS

Two key downsides

One sees what clients asked
(bias)

No control over query time
(unsuitable for time series)



Active DNS measurements

Actively query the DNS from a pre-fetched name list

- Toplist of Webservers (e.g., Alexa)
- Public sub-TLD lists

Purposefully define queries w.r.t.

- Resolvers
- Query types

Big Data Challenge

AN INFRASTRUCTURE FOR MEASURING THE DNS

"Can we measure (large parts of) the
global DNS on a daily basis?"

[Roland van Rijswijk-Deij et al.]

OpenINTEL: <https://www.openintel.nl>

Performs active measurements, sending a fixed set of queries for all covered domains **once every 24 hours**

gTLDs:

.com, .net, .org, .info, .mobi, .aero, .asia, .name, .biz, .gov

+ almost 1200 "new" gTLDs (.xxx, .xyz, .amsterdam, .berlin, ...)

ccTLDs:

.nl, .se, .nu, .ca, .fi, .at, .dk, .ru, .ppp, .us,

Big data in context

One **human genome** is about **$3 \cdot 10^9$ DNA base pairs**

OpenINTEL collects **over $2.3 \cdot 10^9$ DNS records each day** (about 3/4 of a human)

Since **February 2015** they collected **over $4.5 \cdot 10^{18}$ results (4.5 trillion)**
or: **over one billion (10^9) human genomes**

Goals

- G1 Measure every single domain in a top-level domain (TLD)
- G2 Be able to measure even the largest TLD (.com)
- G3 Measure a fixed set of relevant resource records for each domain
- G4 Measure each domain once per day
- G5 Store at least one year's worth of data
- G6 Analyze data efficiently
- G7 Scalability

Challenges

C1 (relates to G3)

Query volume
(.com 123M names in 2015 * x queries)

C2 (relates to C1)

Query pacing
Don't overload authoritative servers

C3
(relates to G5 and G6)

Storage
Assuming each query returns 10,7B, 240GB/day
for .com

C4

Robustness

C5

Ease of operation

System design: Software


Bare metal

- + fast
- High risks of bugs

Off-the-shelf
DNS software

- + long-term experiences
- slower

System design: Scalability



Stage 1: Input data collection

Zone files of top-level domains (TLDs)

Only some TLD (.se, .nu) zone files are public

Dedicated agreements w/ registries

Each database has two tables

Active domains

All domains since start of measurement,
including timestamps when domain was first
seen, last removed, reappeared

Stage 2: Measurements

Cluster manager organizes chunk (a set of domains that were last measured), added to a pool of worker

Worker nodes reports back to manager when work finished, enriches data by meta-data (IP2AS, Geo mapping), submit results to storage

LDNS and Unbound to handle DNS requests

UNBOUND is a DNS resolver

It provides caching

Why is this important?

Distributes queries evenly over
authoritative name servers


Responsible measurements

```
inet6num:      xxxx:xxx:xxxx::/48
netname:       UTwente-OpenINTEL
descr:         University of Twente
descr:         Faculty EEMCS/DACS
descr:         OpenINTEL Active DNS Measurements
descr:         See http://www.openintel.nl/
                for more information
country:       NL
admin-c:       RVR180-RIPE
tech-c:        RVR180-RIPE
status:        ALLOCATED-BY-LIR
mnt-by:        SN-LIR-MNT
mnt-irt:        irt-SURFcert
created:       2018-06-26T08:53:10Z
last-modified: 2018-06-26T08:53:10Z
source:        RIPE
```

Clearly marked the address space from which OpenINTEL measures (including **reverse DNS** and **RIPE DB**)

Very few complaints received

Top-Lists: WWW vs. non-WWW domain names



Wählisch et al., ACM HotNets, 2015

Stage 3: Storage and analysis

Two-tiered approach

- (1) Store in Apache Avro file format
Structured, self-describing data
serialization format + compression; flat
schema, single DNS record is one row
- (2) Convert to Parquet (Hadoop), columnar
format stores all data in single column
sequentially (makes aggregation across
single or few columns + compression
efficient)

Input zone characteristics & worker time

TLD	Registry	#domains	(% of DNS)	Stage I time (Mar-Dec 2015)	
				mean	σ
.com	Verisign	123.1M	(41.2%)	4h 17 min.	1h 15 min.
.net	Verisign	15.6M	(5.2%)	45 min.	31 min.
.org	PIR	10.9M	(3.6%)	19 min.	6 min.
<i>total</i>		149.6M	(50.0%)	5h 20 min.	1h 20 min.

Input zone characteristics & worker time

TLD	Registry	#domains	(% of DNS)	Stage I time (Mar-Dec 2015)	
				mean	σ
.com	Verisign	123.1M	(41.2%)	4h 17 min.	1h 15 min.
.net	Verisign	15.6M	(5.2%)	45 min.	31 min.
.org	PIR	10.9M	(3.6%)	19 min.	6 min.
<i>total</i>		149.6M	(50.0%)	5h 20 min.	1h 20 min.


TLD	#worker VMs	averages over Mar-Dec 2015			
		time (batch)		time (total)	
		mean	σ	mean	σ
.com	80	54 min.	6 min.	17h 10 min.	2h 23 min.
.net	10	52 min.	8 min.	14h 29 min.	2h 15 min.
.org	10	37 min.	4 min.	7h 19 min.	57 min.

Query results

TLD	results for December 31, 2015				averages over Mar-Dec 2015			
	#results	#domains	size (uncompressed)		results/domain		failed domains	
					mean	σ	mean	σ
.com	1419M	122.3M	28.8GB	(211.6GB)	11.75	0.07	0.83%	0.17%
.net	166M	15.5M	3.4GB	(24.3GB)	11.05	0.15	1.21%	0.19%
.org	125M	10.7M	2.5GB	(18.4GB)	11.77	0.09	1.60%	0.22%
<i>total</i>	1709M	148.5M	34.8GB	(254.3GB)	11.68	0.08	0.92%	0.17%

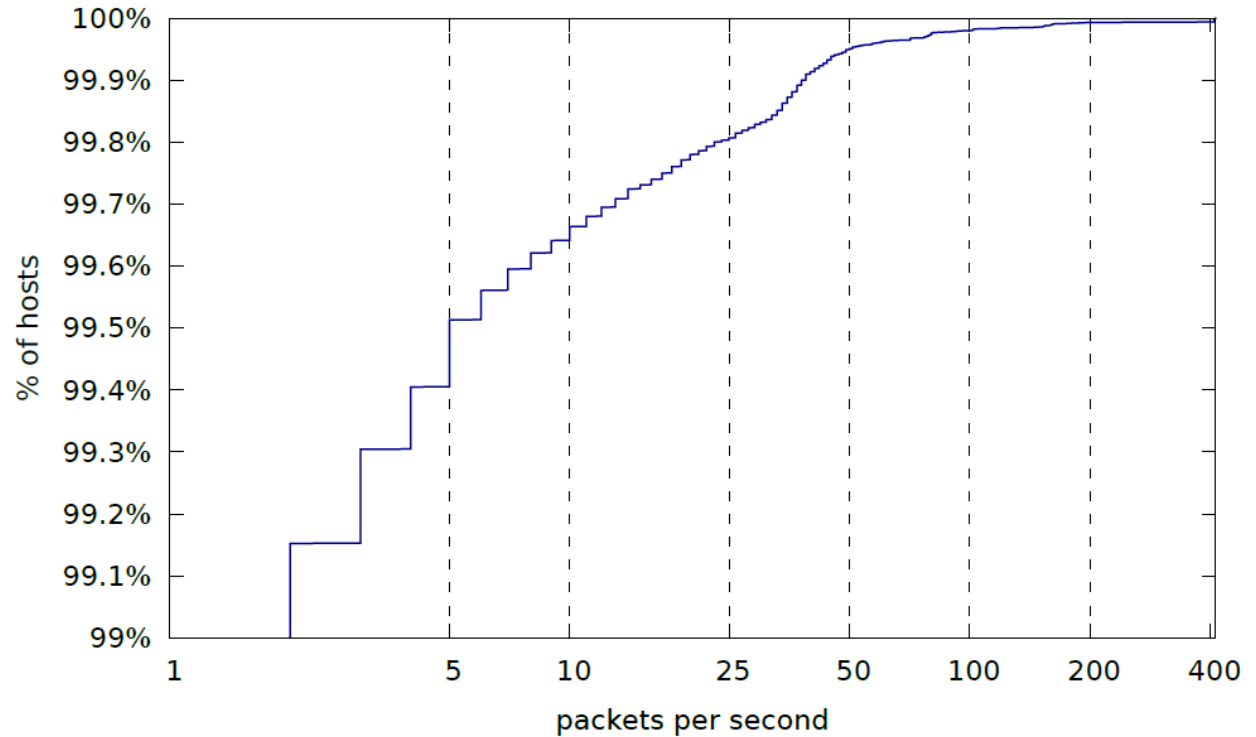
Measurement overhead

Put to context:
 Passive measurements would sample flow data at SURFnet (180 institutes, 1 million users)



How much traffic do individual IP addresses receive?

Analyze outgoing flows for 24 hours, ordered by average number of packets per second



APPLICATIONS OF OPENINTEL

Growing use of email service providers

March – December 2015

Which email provider handles most emails of the .com domain?

Identify top MX records


Group by second-level domain

Manual classification

Clouds providers, top three the usual suspects:
Google (4.09M domain), MS Office 365 (948k domains), Yahoo (609k domains)


In general, most dominant mail handler is GoDaddy (27M domains)

Growing use of **cloud** email providers



Growing use of **cloud** email providers


Side note:
 Middle of May 2015, sharp decline for some top MX SLDs, which belonged to a service that specialized in domain parking



Example 2: DNS resilience


The **attack on Dyn in 2016** shows the risk of sharing DNS infrastructure

Data from OpenINTEL shows that many key customers switched to using two DNS providers





DNS resilience: Topological AS diversity

- **Topological diversity** is important to **protect against denial-of-service**
- Vast **majority of .com** domains has **name servers** located **in a single AS**
- For **.nl** almost **half of domains** have **name servers** in at least **two AS-es**



DNS resilience: Topological prefix diversity

- Majority of .com and .nl have name servers in multiple prefixes, yet 15% only have name servers in a single prefix (IPv4)**



Stupidest thing you can put in a TXT record

In TXT they found

HTML snippets

JavaScript

Windows Powershell code

Other scripting languages (bash, python, ...)

PEM-encoded X.509 certificates

Snippets of DNS zone files

The winner is ...

The winner is ...

```
-----BEGIN RSA PRIVATE KEY-----  
MIICXwIBAAKBgQC36kRNc50wG3uDlRy00xU+9X5LYlhdj0D+ax6BiC27W7iweVwf  
wupxsMvLBhhgegptc5tqb1puXPkCxA6aHwhToFtKSEy4fIWTjWoRthy07SSLsFAC  
koXP++JxZ7bIakqdj5wAyIJ53zSJU7wKImH1Eha7+Myip9LG8HPfsZtY3wIDAQAB  
... ← I left this part out...  
-----END RSA PRIVATE KEY-----
```

The winner is ...

```
-----BEGIN RSA PRIVATE KEY-----
MIICXwIBAAKBgQC36kRNc50wG3uDlRy00xU+9X5LYlhdj0D+ax6BiC27W7iweVwf
wupxsMvLBhhgegptc5tqb1puXPkCxA6aHwhToFtKSEy4fIWTjWoRthy07SSLsFAC
koXP++JxZ7bIakqdj5wAyIJ53zSJU7wKImH1Eha7+Myip9LG8HPfsZtY3wIDAQAB
... ← I left this part out...
-----END RSA PRIVATE KEY-----
```

- Why, oh why, oh why... **oh wait, someone's trying to configure DKIM --- D'oh!**

```
<redacteddomain.tld> IN TXT "v=DKIM1; k=rsa;
p=MIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQC36kRNc50wG3uDlRy00xU+9X5LYlhdj
0D+ax6BiC27W7iweVwfwupxsMvLBhhgegptc5tqb1puXPkCxA6aHwhToFtKSEy4fIWTjWoR
thy07SSLsFACkoXP+JxZ7bIakqdj5wAyIJ53zSJU7wKImH1Eha7+Myip9LG8HPfsZtY3wID
AQAB"
```

MATCH!!!



Discussion

OpenINTEL provides useful data but only for DNS that is homogenous across multiple vantage points, which conflicts with CDNs

Content delivery networks are location-sensitive and reply to DNS queries differently, dependent on the origin of the querier

Literature

R. van Rijswijk-Deij, M. Jonker, A. Sperotto and A. Pras, "[A High-Performance, Scalable Infrastructure for Large-Scale Active DNS Measurements](#)," in *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 6, pp. 1877-1888, June 2016.
doi: [10.1109/JSAC.2016.2558918](#)

Talk by R. van Rijswijk-Deij at RIFE 78

This is the author's version of an article that has been published in this journal. Changes were made to this version by the publisher prior to publication.
The final version of record is available at <https://doi.org/10.1109/JSAC.2016.2558918>

A High-Performance, Scalable Infrastructure for Large-Scale Active DNS Measurements

Roland van Rijswijk-Deij, Mattijs Jonker, Anna Sperotto, and Aiko Pras

Abstract—The Domain Name System (DNS) is a core component of the Internet. It performs the vital task of mapping human readable names into machine readable data (such as IP addresses, which hosts handle e-mail, etc.). The content of the DNS reveals a lot about the technical operations of a domain. Thus, studying the state of large parts of the DNS over time reveals valuable information about the evolution of the Internet.

We collect a unique long-term dataset with daily DNS measurements for all domains under the main top-level domains on the Internet (including .com, .net and .org, comprising 50% of the global DNS name space). This paper discusses the challenges of performing such a large-scale active measurement. These challenges include scaling the daily measurement to collect data for the largest TLD (.com, with 123M names) and ensuring that a measurement of this scale does not impose an unacceptable burden on the global DNS infrastructure. The paper discusses the design choices we have made to meet these challenges and documents the design of the measurement system we implemented based on these choices. Two case studies related to cloud e-mail services illustrate the value of measuring the DNS at this scale. The data this system collects is valuable to the network research community. Therefore, we end the paper by discussing how we make the data accessible to other researchers.

Index Terms—DNS; active measurements; cloud; Internet evolution

I. INTRODUCTION

THE Domain Name System (DNS), plays a crucial role in the day-to-day operation of the Internet. It performs the vital task of translating human readable names – such as `www.example.com` – into machine readable information. Almost all networked services depend on the DNS to store information about the service. Often this information is about what IP address to contact, but also whether or not e-mail received from another host is legitimate or should be treated as spam. Thus, measuring the DNS provides a wealth of data about the Internet, ranging from operational practices, to the stability of the infrastructure, to security. Consider, for example, e-mail handling. In the DNS, the MX record type specifies which hosts handle e-mail for a domain. Thus, examining which MX records are present can tell us, for example, if e-mail handling for that domain is outsourced to a cloud provider such as Google, Microsoft or Yahoo. Another example is the monitoring of protocol adoption such as IPv6 and DNSSEC. The analysis of AAAA or DNSKEY resource

R. van Rijswijk-Deij, M. Jonker, A. Sperotto and A. Pras are with the Design and Analysis of Communications (DMACS) group at the faculty for Electrical Engineering, Mathematics and Computer Science of the University of Twente, Enschede, the Netherlands.

R. van Rijswijk-Deij is also with SURFnet bv, the National Research and Education Network in Utrecht, the Netherlands.
Manuscript received September 9, 2015; revised March 3, 2016.

records can provide ground truth about the adoption of, and operational practices for these protocols over time. Finally, DNS data can also play a vital role in security research, for instance for studying botnets, phishing and malware.

The DNS has been the focus of, or used in, past measurement studies. These studies, however, had a limited scope, in time, coverage of DNS records or number of domains measured. It remains highly challenging to measure the DNS in a comprehensive, large-scale, and long-term manner. Nonetheless, because this type of measurement can provide such valuable information about the evolution of the Internet, we challenged ourselves to do precisely this. Our research goal is to perform daily active measurements of all domains in the main top-level domains (TLDs) on the Internet (including .com, .net and .org, together comprising 50% of the global DNS name space) and to collect this data over long periods of time potentially spanning multiple years.

This paper focuses on the challenges of achieving this goal by answering the following main research question: *"How can one perform a daily active DNS measurement of a significant proportion of all domains on the Internet?"*. The main contributions of the paper are that we show how to:

- Scale such a measurement to cope with the largest TLD (.com with 123M names).
- Ensure that the traffic such a measurement generates does not adversely affect the global DNS infrastructure.
- Efficiently store and analyse the collected data.

Our measurements create a novel large-scale dataset of great value to the research community as well as in other contexts (e.g. for security and forensic purposes). Our ultimate goal therefore is to make the data accessible to others. How we will do this is discussed at the end of the paper.

Finally, in order to validate our system in practice and to illustrate potential uses of the data it collects, we performed two case studies. Given the growing research interest in cloud services, the case studies focus on the use of cloud e-mail services. Based on ten months of data collected by the measurement system between March 2015 and January 2016, we studied the following questions:

- Is Google the most popular cloud mail service provider, or are others, such as Microsoft or Yahoo, more popular?
- Which of these three providers sees the fastest growth?
- Do domains that use these cloud mail services use the Sender Policy Framework (SPF) [1] to combat e-mail forgery, especially since most providers support SPF?

Structure of this paper – Section II introduces our long-term research goals and the challenges that achieving these

Long-term Study


KEY TRANSITIONS IN DNSSEC

Longitudinal Measurement Study: 2005 - 2020


- Measurement tool: Secspider - <https://secspider.net/>
- Crawling DNSSEC from Root/TLDs downward, using zone files, NSEC walking and hitlists
- > 9.5 million DNSSEC zones

Observed during 15 years:

- 35,882,395 distinct DNSSEC keys
- 58,193,197 points in time when keys were added or removed
- Total of ≈ 19 million key transitions

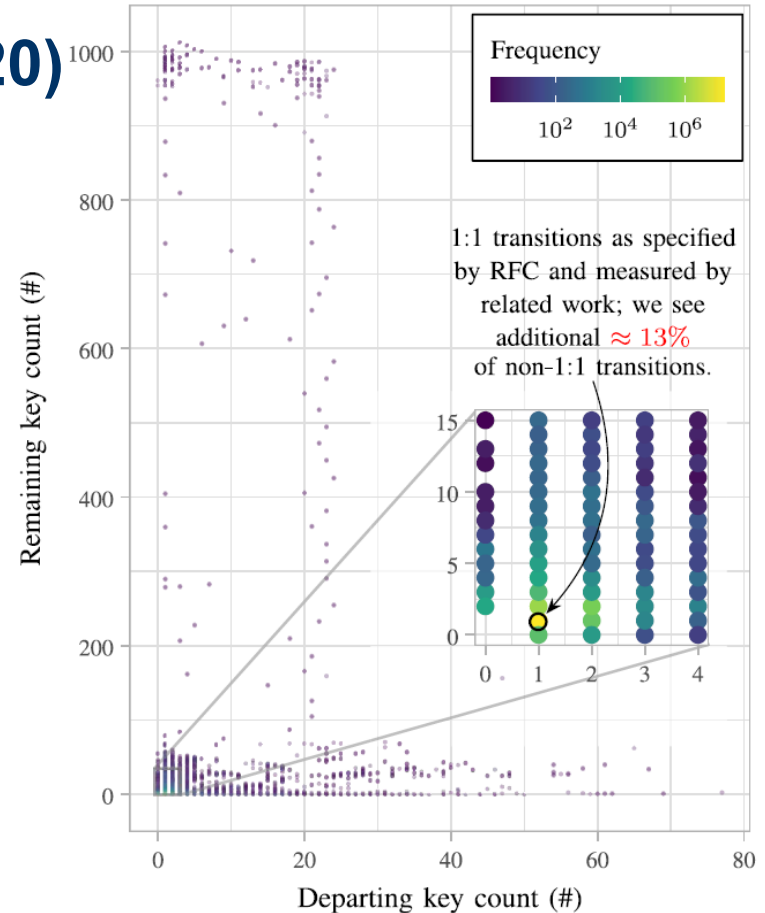


Anatomy of a 1:1 Key Transition

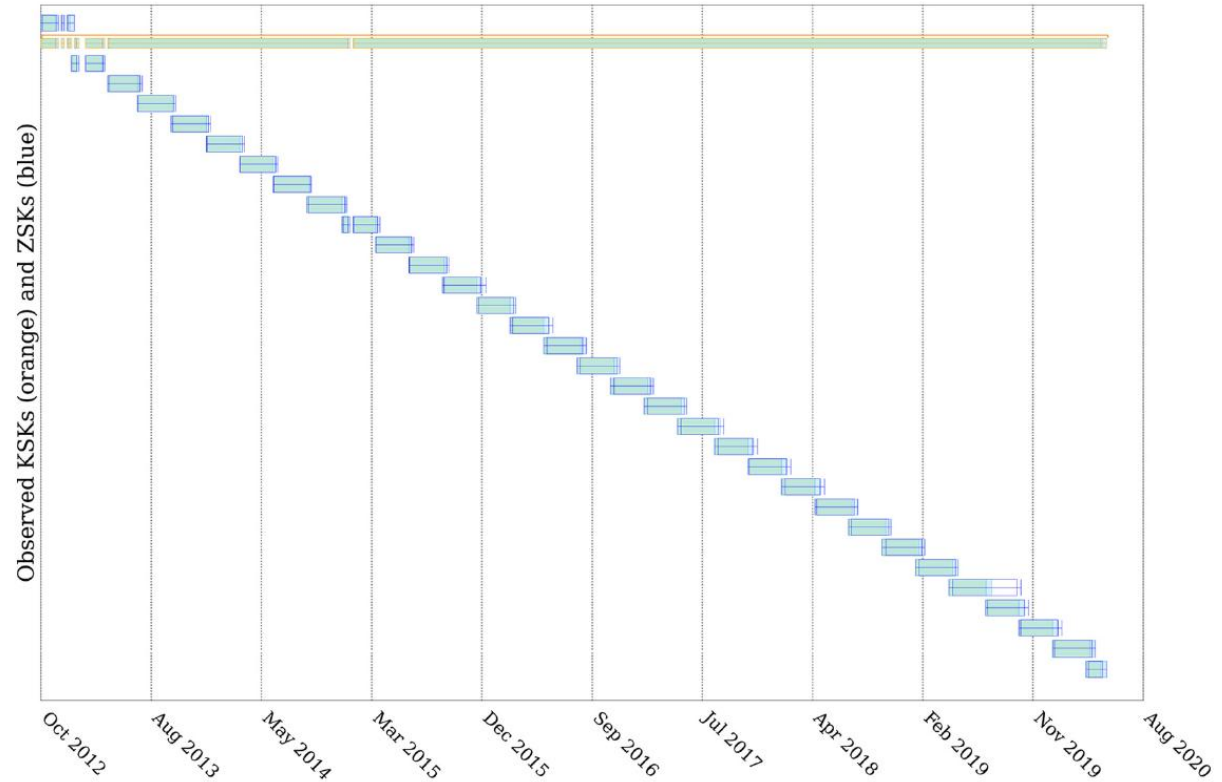


Observed Key Transitions (‘05-‘20)


- Surprisingly many variants of key transitions
- 13% non-1:1
- Some transitions largely increase the number of DNSKEYs of a domain



Regular Example: .com Zone



Irregular Example: Up to 55 Simultaneously Active Keys



Lessons Learned

- DNSSEC key management still challenging
 - Threats of broken trust chains
 - Threats of high amplification by keys
- Intricate temporal interplay of
 - DNS record TTLs and caching
 - Signature lifetimes
 - Key lifetimes
- Still incomplete automation and tooling
 - Particular problem:
DS records that refer to externals

Literature

E. Osterweil, P. F. Tehrani, TC. Schmidt, M. Wählisch: **From the Beginning: Key Transitions in the First 15 Years of DNSSEC**, *IEEE Transactions on Network and Service Management (TNSM)*, Vol. 19, No. 4, p. 5265–5283, December 2022.

Doi: [10.1109/TNSM.2022.3195406](https://doi.org/10.1109/TNSM.2022.3195406)

From the Beginning: Key Transitions in the First 15 Years of DNSSEC

Eric Osterweil¹, Pouyan Fotouhi Tehrani¹, Thomas C. Schmidt¹, *Member, IEEE*, and Matthias Wählisch¹, *Member, IEEE*

Abstract—When the global rollout of the DNS Security Extensions (DNSSEC) began in 2005, a first-of-its-kind trial started: The complexity of a core Internet protocol was magnified in favor of better security for the overall Internet. Thereby, the scale of the loosely-federated delegation in DNS became an unprecedented cryptographic key management challenge. Though fundamental for current and future operational success, our community lacks a clear notion of how to empirically evaluate the process of securely transitioning keys. In this paper, we propose two building blocks to formally characterize and assess key transitions. First, the *anatomy of key transitions*, i.e., measurable and well-defined properties of key changes; and second, a novel *classification model* based on this anatomy for describing key transition practices in abstract terms. This abstraction allows for classifying operational behavior. We apply our proposed transition anatomy and transition classes to describe the global DNSSEC deployment. Specifically, we use measurements from




Fig. 1. Notable DNS(SEC) deployment events (blue) and security incidents (red) during the measurement periods of related work (black) and this work (orange).

Case Study

HIJACKING INTERNET RESOURCES WHEN DOMAIN NAMES EXPIRE

Motivation 1: Long-term abuse of IP prefixes



home | my eBay | site map | sign in/out

ebay **Browse** Sell Services Search Help Community
item view

[See this item](#) in eBay's new look for this page.

/16 CLASS B - 65534 IP's GRANDFATHERED !!!!
Item # 3029809556

[Electronics & Computers Networking & Telecom Other](#)
[Electronics & Computers Wholesale Lots Networking & Telecom](#)

	Current bid	US \$6,800.00 <small>(reserve not yet met)</small>	Starting bid	US \$0.01
	Quantity	1	# of bids	29 Bid history
	Time left	8 days, 0 hours +	Location	Houston
	Started	Jun-09-03 22:34:11 PDT	Country/Region	United States /Houston
	Ends	Jun-19-03 22:34:11 PDT	Mail this auction to a friend	
			Watch this item	


Featured Auction

[csutter2002](#) (170 ★)

Seller (rating)
Feedback rating: 170 with 100% positive feedback reviews ([Read all reviews](#))
 Member since: Jun-23-02. Registered in United States
[View seller's other items](#) | [Ask seller a question](#) | [Safe Trading Tips](#)


Regular **prefix** hijacking

The Abandoned Side of the Internet



Regular **prefix** hijacking

The Abandoned Side of the Internet



Motivation 2: The LINKTEL INCIDENT

A new hijacking attack

SOS to NANOG from a Russian ISP under attack
Unnoticed for 6 months due to business struggles
Forensic analysis of the incident one year later


Complex attack plan with a hand-picked target

The victim's DNS domain had expired, which enabled administrative take-over of its Internet resources

No BGP activity for the victim's IP prefixes, which enabled stealthy hijack of the prefixes and the AS

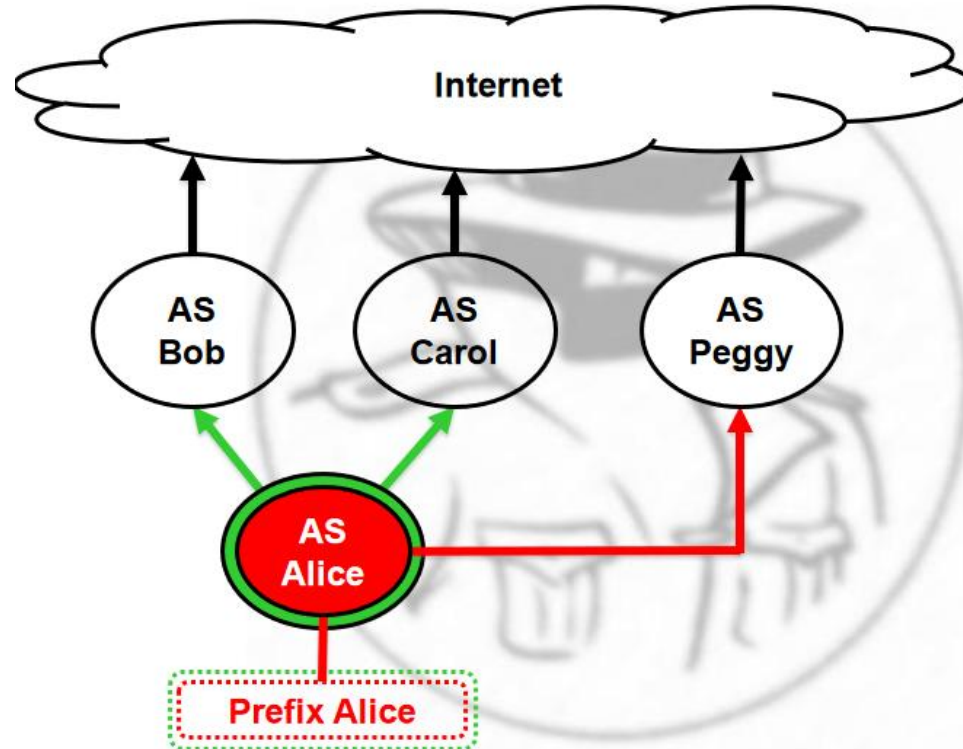
AS hijacking

The Abandoned Side of the Internet



AS hijacking

The Abandoned Side of the Internet



Precondition for successful attacks

Today, origin validation is based on

- ISP info in Internet Routing Registries (IRR)
- Social exchange (email conversation)
- IRR, RPKI entries binding an AS to a prefix

Imagine a company going (temporarily) out of business. Eventually, without cash flow...

- Its DNS domain is going to expire
- Its BGP activity terminates
- Its IRR entries remain

What are we looking for

Given this knowledge, an attacker can easily impersonate a hand-picked victim by

- Re-registration of the DNS domain
- Claiming ownership and misleading any upstream ISP

Our approach is similar

- Find resource groups under same administration
- Identify groups that reference expired domains only
- Cross-check time of last IRR update
- Take into account BGP history
- Evaluate gain (e.g. number of abandoned prefixes)

Recap: RIPE database

RIPE maintains an IRR database for the European service region

- Daily snapshots are available (mostly anonymized)
- We analyzed 2.5 years of archived snapshots (Feb 23, 2012 –July 9, 2014)

```
inetnum: 194.28.196.0 - 194.28.199.255
netname: UA-VELES
descr: LLC „Unlimited Telecom“
descr: Kyiv
notify: internet@veles-isp.com.ua
mnt-by: VELES-MNT
```

```
aut-num: AS51016
as-name: VALES
descr: LLC „Unlimited Telecom“
notify: internet@veles-isp.com.ua
mnt-by: VELES-MNT
```

Grouping objects by maintainer

Maintainer groups


- Group by unique mnt-by references of all objects
- Yields 48,802 disjoint groups

We disregard groups...

- Of zero-size (unreferenced maintainers)
- With multiple or without any DNS names
- Without inet-num or aut-num objects

We merge groups by identical DNS names, leading to a total of 7,907 remaining groups


Size of maintainer groups



RIPE database objects

Object type	Frequency	DNS references	
inetnum	3,876,883	1,350,537	(34.84%)
domain	658,689	97,557	(14.81%)
route	237,370	50,300	(21.19%)
inet6num	231,355	8,717	(3.77%)
organisation	82,512	0	(0.00%)
mntner	48,802	0	(0.00%)
aut-num	27,683	6,838	(24.70%)
role	20,684	14,430	(69.76%)
as-set	13,655	2,500	(18.31%)
route6	9,660	723	(7.48%)
irt	321	162	(50.47%)
Total	5,239,201	1,531,764	(29.24%)

Lifetime of domain names



Extracted domain names

More than 1.5 M references to DNS names, of which 21,061 are distinct

Whois queries yield 214 expired DNS names

65 of 7,907 groups reference expired DNS names

Top5 TLDs

.com	27.9%
.ru	21.5%
.net	13.0%
.se	4.8%
.co.uk	3.5%

Top5 TLDs (expired)

.ru	20.1%
.it	16.4%
.com	9.8%
.dk	9.8%
.net	7.0%

Refinement by active measures

The RIPE db could be simply outdated


Time since last database update

- Top-10% of valid groups changed within 2 months
- Top-10% of expired groups changed within 6 months
- DNS expiry and update behavior correlate


Time since last BGP update

- Search for prefixes and ASes of the maintainer groups
- Analysis of 2.5 years of archived BGP routing tables
- Key findings: 90% of valid resources are active in BGP, in contrast to 75% of expired resources

Time since last DB update



Time since last BGP activity



Abandoned Resources

Expired DNS names

- 65 disjoint resource groups reference expired domains
- These groups hold 773 /24 networks and 54 ASes

BGP activity for these resources

- 75% are still in use (but impersonation is possible, i.e. a hijack would disrupt operational use)
- 13 groups show no activity for more than 6 months

Summary

- Correlation of archived RIPE databases, BGP tables and DNS registration data over a period of 30 months
- We found that in total, more than a /18 network is abandoned, waiting to be stealthily hijacked!

We need better ownership validation to secure unused resources!

Literature

Johann Schlamp, Josef Gustafsson, Matthias Wählisch,
Thomas C. Schmidt, Georg Carle,

[The Abandoned Side of the Internet: Hijacking Internet Resources When Domain Names Expire](#),

In: *Proc. of 7th International Workshop on Traffic Monitoring and Analysis (TMA)*, (Moritz Steiner, Pere Barlet-Ros, Olivier Bonaventure: Ed.), ser. LNCS, Vol. **9053**, pp. 188--201, Heidelberg: Springer-Verlag, 2015.
DOI: https://doi.org/10.1007/978-3-319-17172-2_13

The Abandoned Side of the Internet: Hijacking Internet Resources When Domain Names Expire

Johann Schlamp¹(✉), Josef Gustafsson¹, Matthias Wählisch²,
Thomas C. Schmidt³, and Georg Carle¹

¹ Technische Universität München, München, Germany

{schlamp,gustafss,carle}@net.in.tum.de

² Freie Universität Berlin, Berlin, Germany

m.waehlich@fu-berlin.de

³ HAW Hamburg, Hamburg, Germany

schmidt@informatik.haw-hamburg.de

Abstract. The vulnerability of the Internet has been demonstrated by prominent IP prefix hijacking events. Major outages such as the China Telecom incident in 2010 stimulate speculations about malicious intentions behind such anomalies. Surprisingly, almost all discussions in the current literature assume that hijacking incidents are enabled by the lack of security mechanisms in the inter-domain routing protocol BGP.

In this paper, we discuss an attacker model that accounts for the hijacking of network ownership information stored in Regional Internet Registry (RIR) databases. We show that such threats emerge from abandoned Internet resources (e.g., IP address blocks, AS numbers). When DNS names expire, attackers gain the opportunity to take resource ownership by re-registering domain names that are referenced by corresponding RIR database objects. We argue that this kind of attack is more attractive than conventional hijacking, since the attacker can act in full anonymity on behalf of a victim. Despite corresponding incidents have been observed in the past, current detection techniques are not qualified to deal with these attacks. We show that they are feasible with very little effort, and analyze the risk potential of abandoned Internet resources for the European service region: our findings reveal that currently 73 /24 IP prefixes and 7 ASes are vulnerable to be stealthily abused. We discuss countermeasures and outline research directions towards preventive solutions.

1 Introduction

Internet resources today are assigned by five Regional Internet Registrars (RIRs). These non-profit organisations are responsible for resources such as blocks of IP addresses or numbers for autonomous systems (ASes). Information about the status of such resources is maintained in publicly accessible RIR databases, which are frequently used by upstream providers to verify ownership for customer networks. In general, networks are vulnerable to be hijacked by attackers due to