Economic Incentives on DNSSEC Deployment: Time to Move from Quantity to Quality

Tho Le*†, Roland van Rijswijk-Deij†‡, Luca Allodi* and Nicola Zannone*
*Eindhoven University of Technology
†SURFnet
‡University of Twente
tho.le3839@gmail.com, r.m.vanrijswijk@utwente.nl, l.allodi@tue.nl, n.zannone@tue.nl

***Abstract*—The security extensions to the DNS (DNSSEC) currently cover approximately 3% of all domains worldwide. In response to the low deployment of DNSSEC, a few top-level domains started offering 'per-domain' economic incentives to encourage adoption of the protocol by offering a yearly discount on each signed domain. However, it remains unclear whether these incentives are well-balanced and foster the overall security of the infrastructure as well as its deployment at scale.**

In this paper we argue that, in the presence of fixed costs of deployment, misaligned 'per-domain' incentives may have the collateral effect of encouraging large operators to *massively* deploy *unsecure* implementations of DNSSEC, whereas smaller operators, for which the effect of the economic incentive is negligible, may not significantly benefit from it. To investigate this, we study the security of DNSSEC deployment at scale, particularly in TLDs that offer economic incentives. We find that the security of DNSSEC implementations in the wild poorly reflects standard recommendations, particularly for tasks that cannot be solved by triggering a flag in the DNS software service (e.g. key rollover). Further, we find that, on average, large operators deploy weak DNSSEC security more frequently than small DNSSEC operators, suggesting that current incentives are ineffective in promoting a secure adoption and in deterring insecure implementations. We conclude the paper with actionable recommendations for TLD registry operators to improve the alignment of economic incentives with secure DNSSEC requirements.

***Index Terms*—DNS; DNSSEC; measurement; network security; economic incentives**

I. Introduction

The Domain Name System (DNS) maps human-readable domain names to addresses that can be conveniently interpreted by a machine [1]. Unfortunately, the lack of security mechanisms for authentication and integrity verification of DNS responses makes DNS vulnerable to the so-called cache poisoning attacks [2], whereby an attacker disguises its own network resources under the name of a legitimate domain.

The DNS Security Extensions (DNSSEC) [3] have been introduced to address this problem via digital signatures. At the time of writing, more than 90% of Top Level Domains (TLDs) are DNSSEC-enabled [4]. Unfortunately, an opposite observation applies to DNSSEC adoption by second-level domains, with only approximately 3% of all domains worldwide using DNSSEC [5]. In order to encourage DNSSEC deployment, several policy measures have been taken, for instance by ICANN, which has been organising day-long workshops dedicated to DNSSEC at every public ICANN meeting since 2005 [6]. Similarly, a number of country-code TLD (ccTLD) registries offer economic incentives to operators, calculated on a 'per-domain' basis [7], [8]. For example, `.nl` and `.se` offer a discount to registrars for every signed zone: `.nl` offers €0.28 off the registration cost of €3.40 per year [7] and `.se` offers a 6 SEK discount on the registration cost of 90 SEK per year [9]. These initiatives seem to have led to some success in terms of amount of signed domains [8]–[10]; for example, the ccTLD `.nl` has become the largest DNSSEC zone in a short time following the introduction of the incentives [10], and the `.se` registry observed a rapid increase of more than 160,000 DNSSEC adoptions overnight once the registry doubled their incentive program to a 5% discount [11].

On the other hand, it remains unexplored whether the increased *volume* of DNSSEC deployments is reflected in a *secure* implementation, or if sloppy security practices cripple the effectiveness of DNSSEC adoption at its core. To evaluate this, we perform a longitudinal study on DNSSEC adoption that investigates the quality of its deployment in terms of adherence to standard security best practices defined by the NIST [12], [13]. Based on observations on existing 'per-domain' economic incentives for `.se` and `.nl` domains, we further speculate that this type of incentive may be misaligned in promoting a complete and correct deployment of DNSSEC: large operators (that manage hundreds of thousands of domains) may be encouraged to deploy DNSSEC at scale, focusing on *quantity* as opposed to the *security* of the deployment; on the other hand, smaller operators (that manage only a few thousand domains) may not gain substantial benefits from incentive schemes, and may therefore deploy DNSSEC at smaller scales but with higher average security (e.g. out of considerations for market competitiveness or differentiation). To evaluate this effect, we rely on a unique, large-scale dataset from the OpenINTEL project that systematically crawls DNS records on a daily basis for all second-level domains under the most common TLDs [14], [15]. Our ultimate goal is to provide additional insights on devising effective policies that encourage *secure* deployments of DNSSEC on top of the mere adoption of the protocol, and provide several suggestions in this direction as a conclusion to this work.

Our contribution is threefold: (1) we perform the first large-scale study of security aspects of DNSSEC in the presence of economic incentives for the deployment of this technology;

978-1-5386-3416-5/18/$31.00 © 2018 IEEE

(2) we evaluate the difference in compliance to security best practices between *large* and *small* operators, and find that the former generally perform worse than the latter for high-complexity tasks such as key rollover; (3) we provide actionable recommendations for registries to improve the impact of economic incentives on the quality of DNSSEC deployments.

The remainder of the paper is organized as follows. Section II presents background information on DNSSEC. Section III introduces the dataset and research methodology. Section IV presents the results of our study and Section V discusses our findings and provides recommendations to registries. Section VI discusses related work. Finally, Section VII concludes the paper and suggests directions for future work.

II. Background

This section first presents an overview of DNSSEC from an operator's perspective. Next, the actors involved in the administration of the DNS are introduced. Finally, we discuss deployment costs and the issues of existing incentive programs.

A. DNSSEC deployment

DNSSEC adds *authenticity* and *integrity* to the DNS using digital signatures. An operator that wants to deploy DNSSEC needs to perform a number of additional tasks compared to managing a domain that just uses regular DNS. First, the operator needs to generate cryptographic keys to use for signing. In the most common DNSSEC setup, this involves generating two keys, a so-called Key Signing Key (KSK) (this key is only used to sign the keyset for a domain) and a so-called Zone Signing Key (ZSK) (this key is used to sign the individual DNS records in a domain). An alternative, but less frequently used configuration, combines the roles of these two keys in a single key, called a Combined Signing Key (CSK). The operator needs to pick a secure algorithm for the key, and should select an appropriate, cryptographically strong key size. Second, the operator needs to sign all records in the domain. Since signatures in DNSSEC have a limited validity, this is a task that needs to be repeated at regular intervals. Third, the operator has to upload its KSK to the parent zone – using a dedicated DNS record called Delegation Signer (DS) – to create a so-called *chain of trust* (Fig. 1 shows an example). Finally, the operator needs to replace signing keys at regular intervals. As we will see in Section III-B, where we discuss best practices, this interval depends on the cryptographic strength of keys. In general, weaker keys have to be replaced more frequently. Key rollover is one of the more complex aspects of DNSSEC. Replacement of the KSK requires an interaction with the parent zone, and more generally, due to the dynamics of the DNS (e.g. caching of records), replacing a key requires the introduction of new keys prior to use, waiting for new keys to propagate across the Internet, and retaining old keys for some period after they have been retired to allow for records signed with old keys persisting for some time in caches. For a more detailed discussion of the intricacies of key management, we refer to RFC 6781 [16].

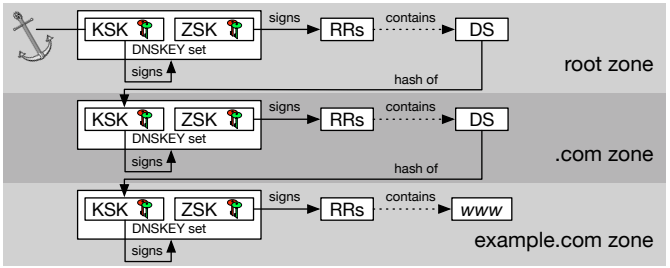

Fig. 1. An overview of DNSSEC

B. Registry, registrar, registrant and DNS operator

The root of the DNS is managed by the Internet Corporation for Assigned Names and Numbers (ICANN). ICANN delegates the responsibility of maintaining TLDs to organizations called *registries*. Most registries allow third-parties, known as *registrars*, to sell these domain names to the public, and buyers or owners of domain names are called *registrants*. This Registry-Registrar-Registrant channel forms the core procedure for the registration and administration of domain names. Another important role in DNS is the one of *DNS operator* which is the entity hosting and managing the actual DNS records. Although a registrar can also be a DNS operator, these two roles are distinct and should not be confused.

C. Deployment costs and misaligned 'per-domain' incentives

The European Network and Information Security Agency (ENISA) performed a survey on the cost of DNSSEC deployment [17], and concluded that the cost is not clearly correlated with the number of managed domains or handled queries: whereas it seems unlikely that there is no infrastructural overhead scaling with the size of the DNSSEC deployment, its associated costs may 'collapse' orders of magnitude below ordinary deployment and maintenance costs [17]. At the same time TLD registries, in an attempt to encourage the adoption of DNSSEC, adopted a 'per-domain' incentive mechanism that rewards the operator for each new signed domain. Hence, the total reward scales linearly with the number of signed domains.

We observe that, in the presence of fixed costs of deployment [17], 'per-domain' incentives may create a perverse mechanism whereby signing DNS records 'en-masse' may be encouraged, without specific considerations on the security of the deployment, which is not a condition for eligibility in these incentive mechanisms. On the contrary, such mechanisms may simply provide an economic advantage for those operators that deploy DNSSEC *at scale* (i.e. those that have enough critical mass for the cumulative incentive to represent a significant fraction of the fixed deployment costs), while providing virtually no cost relief to smaller operators for which the discount represents only a negligible fraction of the overall costs. This effect would have severe negative consequences on the DNSSEC market by unfairly favouring large operators over small operators, while at the same time not incentivizing a secure deployment of DNSSEC and therefore undermining the very infrastructure they are intendend to promote.

III. Approach

The goal of this study is to evaluate the effect of 'per-domain' economic incentives on secure DNSSEC deployments. In particular, we speculate that the attribution of these economic benefits, in the absence of specific security eligibility requirements, increases net cost advantages for large operators, which cumulatively cash in incentive payouts for the mass of signed domains they manage, without favoring a *correct* and *secure* deployment of DNSSEC; on the other hand, small DNS operators may deploy DNSSEC with different motivations than cost benefits (e.g. competitive market advantage), as the small volume of domains they manage does not provide enough critical mass for the incentive to cover a significant fraction of overall deployment costs. If this holds, we would then expect the average *small* operator to deploy DNSSEC (1) less frequently and (2) with higher average security than *large* operators. We study this by investigating the difference in the overall compliance to NIST security guidelines between operators that massively benefit from the incentive (large operators), and those that do not (small operators). Our running hypothesis is:

> *Hypothesis:* Despite the presence of 'per-domain' economic incentives in `.nl` and `.se`, large DNS operators deploy DNSSEC with lower compliance to security guidelines than small DNS operators.

A. Data

Our analysis relies on datasets collected by a large-scale active measurement platform called OpenINTEL [14], [15]. OpenINTEL crawls DNS records daily for all second-level domains under various TLDs together comprising the majority of the global DNS namespace [14]. In this work, we use a subset of the datasets collected by OpenINTEL. Specifically, we use those datasets that cover the full `.com`, `.net` and `.org` TLDs, because these are some of the largest TLDs,[1] and those that cover the `.nl` and `.se` ccTLDs, because these two ccTLDs have economic incentive programmes to incentivize DNSSEC deployment. Tab. I provides an overview of the specific OpenINTEL datasets used in this study. We collect the following record types: (1) Start of Authority (SOA) records and their Resource Record Signature (RRSIG), (2) DS records and (3) DNS Public Key (DNSKEY) records with associated RRSIGs. These records allow us to assess the quality of the DNSSEC deployment for a domain in terms of NIST's recommended best practices as described in Tab. II.

a) Identifying DNS Operators for Each Zone: To meaningfully distinguish between operators we first need to identify which DNS operator administers a domain. We utilize the `MNAME` field in an SOA resource record. The `MNAME` field identifies the master name server for a domain [19]. Specifically, we determine the DNS operator of a domain by considering the root of the master name server's domain name. For example, if the `MNAME` field of a domain is *ns0.transip.nl*, we identify that this domain is operated by *transip.nl*.

[1] At the end of Q3 2017, `.com`, `.net` and `.org` together make up 47.2% of the global DNS namespace [18].

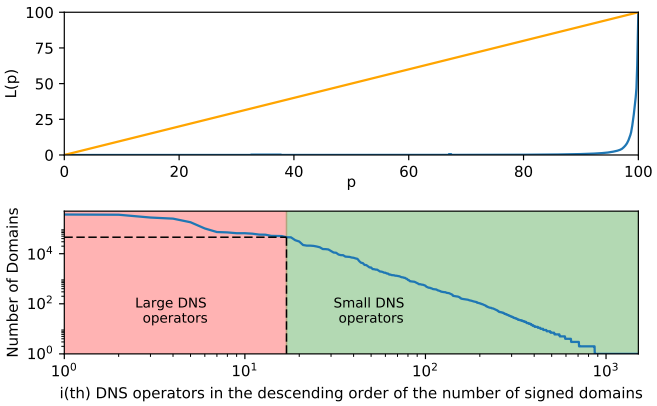

Fig. 2. Classification of large and small DNS operators for `.nl`

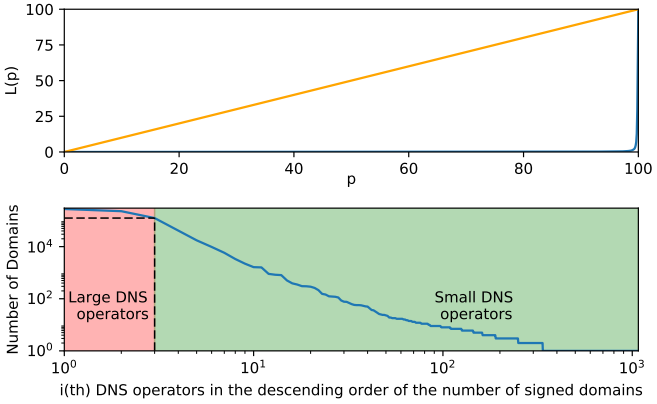

Fig. 3. Classification of large and small DNS operators for `.se`

b) Distinguishing between Small and Large DNS Operators: Since a clear-cut definition of large and small DNS operators is not available, we define a criterion to classify DNS operators. Specifically, we empirically observe that the domain distribution by operators roughly follows the Pareto principle (also known as the 80/20 rule); that is, a small number of DNS operators account for the vast majority of domains, whereas the remaining mass of operators only manages a fraction of the overall volume of domains. This distinction is reported pictorially in Figs. 2 and 3 for the `.nl` and `.se` TLDs, the focus of this paper. The top row of each figure presents a Lorenz curve [20] describing the p percent of DNS operators that manage the $L(p)$ percent of DNSSEC-signed domains. The diagonal line presents the equality state in which each DNS operator is responsible for the same amount of signed domains. The large gap between the two Lorenz curves and equality lines indicates a significant inequality in the distribution of domains per DNS operator. In that light, only a few DNS operators account for most of the signed domains. The second row of each figure illustrates our clear-cut classification between large and small DNS operators. Following the trend reported in Fig. 3, we label as large those DNS operators responsible for more than 80% of signed domains in `.nl` and `.se`; the others are classified as small DNS operators. The full list of large DNS operators is presented in Tab. VI.

TABLE I
OVERVIEW OF THE DATA USED FOR THIS STUDY.

TLDs	Measurement Period	#Domains
.com	2015-02-28 - 2017-07-31	116,814,548
.net	2015-02-28 - 2017-07-31	13,011,428
.org	2015-02-28 - 2017-07-31	9,373,214
.nl	2016-02-09 - 2017-07-31	5,440,975
.se	2016-06-07 - 2017-07-31	1,440,244

TABLE II
NIST DNSSEC BEST PRACTICES

Aspects	NIST recommendation
Key size	- ECDSA keys. - RSA: KSKs >= 2048 bits and ZSKs >= 1024 bits.
Key algorithm	- Recommended: Algorithms 8 and 10. - Highly recommended: Algorithms 13 and 14.
Key rollover	**KSKs/CSKs:** - ECDSA keys and and RSA keys (with key size >=2048 bits): rollover within 24 months. **ZSKs:** - 1024-bit RSA keys: rollover within 90 days. - RSA keys' size between 1024 - 2048 bits: rollover within 12 months. - ECDSA keys and RSA keys (with key size >= 2048 bits): rollovers within 24 months.

B. Evaluation of DNSSEC deployment security

Although there is no universal agreement on criteria for secure DNSSEC deployment, several works propose guidelines for DNSSEC deployment: RFC 6781 [16], the Good Practices Guide for Deploying DNSSEC by ENISA [21] and two guides by National Institute of Standards and Technology (NIST), namely the Secure Domain Name System (DNS) Deployment Guide [12] and Recommendations for Key Management (part 3) [13]. However, the ENISA guide is outdated (2010), and RFC 6781 only provides generic recommendations. The NIST guides, on the other hand, offer more recent and detailed recommendations; hence, we use these as the basis for best practices for DNSSEC deployment. Tab. II presents an overview of NIST best practices. In this study, the "quality" of a DNSSEC deployment refers to its adherence to these recommendations.

It is worth mentioning that we have left out some recommendations from the NIST guides. In particular, we did not consider recommendations on the key rollover approach (e.g., pre-publish for ZSK and double signature for KSK) and key algorithm rollover, since these do not directly affect the quality of a signed zone. Moreover, we did not consider the recommendation on the validity period of signatures over DNSKEY records. This recommendation is controversial, as short validity periods limit an operator's ability to perform maintenance and troubleshooting in case of problems.

To test our hypothesis, we compare the quality of DNSSEC deployment between large and small DNS operators based on three aspects: key algorithm, key size and key rollover. For key algorithm and key size, we do this by inspecting DNSKEY records from the input datasets we obtain from OpenINTEL for each signed domain (as single operators may not perform uniformly over all managed domains). Tracking key rollover requires more complex considerations: as changes happen over time, we need to track the set of DNSKEY records for each signed domain. Furthermore, we have to track the signature (RRSIG record) for the SOA record, in order to establish when a new key is first used, and when an old key is retired. This needs to be done on a day-to-day basis over the entire duration of our datasets, and requires processing of millions of records for each calendar day in the dataset. For the comparison, each of the three aspects is evaluated against the best practices as shown in Tab. II. As our dataset comprises measurements on the full population of .nl and .se domains, we compare observation frequencies for the hypothesis testing.

TABLE III
OVERVIEW OF DNSSEC DEPLOYMENT (JULY 31ST, 2017)

TLD	#Signed domains			%Signed domains
	KSK/ZSK	CSK	Total	
.com	932,334	4,079	936,413	0.80%
.net	140,322	765	141,087	1.08%
.org	104,942	566	105,508	1.13%
.nl	2,709,503	119,681	2,829,184	52.00%
.se	721,090	16,236	737,326	51.19%

IV. RESULTS

This section presents our results. First, we present demographics for DNSSEC deployment from our datasets. We then test the security of DNSSEC deployment for operators that heavily benefit from the incentive (large operators), and for operators that do not substantially benefit from it (small operators).

A. Overview of TLDs

Earlier work by Chung et al. [22], which studied the state of DNSSEC deployment in .com, .net and .org, showed that DNSSEC adoption in these TLDs is low, and that there are serious security issues. They find use of weak keys, weak signing algorithms and a large number of domains that fail to deploy DNSSEC completely. That is: they find several domains that are signed, but for which a corresponding secure delegation with a DS record in the parent zone is missing. In follow-up work, Chung et al. [23] also shed light on the role that domain name registrars play in the deployment of DNSSEC, especially their vital role in creating full deployments including secure delegations. In this work, we focused on a different aspect of DNSSEC deployment, specifically on the deployment of DNSSEC in the presence of economic incentives to deploy DNSSEC.

Tab. III shows an overview of DNSSEC adoption in the five analyzed TLDs, as of July 31, 2017. We observe that .nl and .se are the TLDs that achieve the highest levels of DNSSEC adoption. This may be a direct consequence of the incentive programs promoted by the registries responsible for these TLDs. Notably, .nl is the largest DNSSEC zone whereas we observed a very low percentage of DNSSEC deployment in three popular TLDs, namely .com, .net and .org. In absolute terms, the number of signed domains in this TLD is more than double that of the domains in .com, .net

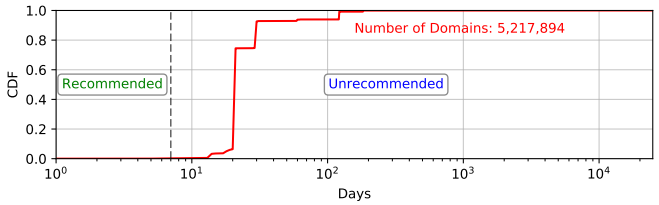

Fig. 4. Cumulative Distribution Function (CDF) of validity periods for signatures covering DNSKEY RRSets.

TABLE IV
DEPLOYMENT DIFFERENCES BETWEEN LARGE AND SMALL OPERATORS (JULY 31ST, 2017)

	Large operators			Small operators		
TLD	**#Domains**	**#Signed**	**%**	**#Domains**	**#Signed**	**%**
.com	93,464,626	712,162	0.76%	23,349,922	224,251	0.96%
.net	10,412,605	114,687	1.10%	2,598,823	26,400	1.02%
.org	7,501,310	85,166	1.14%	1,871,904	20,342	1.09%
.nl	4,353,518	2,736,393	62.85%	1,087,457	92,791	8.53%
.se	1,153,129	723,532	62.75%	287,115	13,794	4.80%

and .org combined. Interestingly, as Chung et al. [23] also observe, a large fraction (more than 34%) of DNSSEC-signed domains in .com, .net and .org can be attributed to operators that also manage a large number of signed domains in the .nl and .se TLDs. We speculate that this may actually be a side-effect of the economic incentives: the intuition behind this is that operators aiming to qualify for the economic incentives in .nl or .se use a single strategy for all domains they manage, and thus sign everything they manage, including domains in other TLDs. Another hint that this intuition may hold is the fact that we observe a large number of partial DNSSEC deployments for these operators in .com, .net and .org. We consider a deployment partial if the domain is signed (i.e. there are keys present, in the form of DNSKEY records, and there are signatures present in the form of RRSIG records), but there is no secure delegation (using a DS record) in the parent zone. Such a partial deployment is effectively useless, as it means all the effort is expended to sign a domain, whereas no one will be able to validate the signatures along the chain of trust. Approximately half of DNSSEC deployments in .com, .net and .org, that can be attributed to operators that have a large presence in .nl and .se, are partial deployments.[2] It seems likely that these operators may simply not have bothered to create secure delegations in .com, .net and .org as there is no incentive (in economic terms) for them to do so.

With respect to signing schemes, we observe that the KSK/ZSK scheme is significantly more common than the CSK one. The most likely explanation for this is that most DNSSEC software implementations use the KSK/ZSK scheme by default. This is despite the fact that CSK is generally preferable to the KSK/ZSK scheme, as it leads to smaller DNS responses for the DNSKEY record type, and is therefore an important tool to reduce the risk of packet fragmentation (which can lead to availability issues [24]), and to reduce the potential for DNS amplification attacks that abuse a signed domain [25]. Further, recall from Section III-B that we explicitly chose not to consider NIST's recommendation of a maximum signature validity period of 7 days for signatures over DNSKEY records due to its controversial nature; an analysis of current re-signing practices (Fig. 4) confirms this, and indicates that virtually no operator follows NIST's recommendation of a maximum validity period of 7 days. Quite differently, the most common DNSKEY re-signing periods are 21, 30 or 122 days.

Finally, we compared the extent to which DNSSEC is deployed by large and small operators respectively. We do this in order to test our intuition that economic incentives on a 'per-domain' basis are more favorable for large operators. If this is the case, we expect to see higher deployment rates for large operators. Tab. IV shows the result of this analysis. As the table shows, there is a clear difference in deployment rates, but this difference *only occurs in TLDs with economic incentives*. This strongly suggests that our intuition holds; in both .nl and .se the fraction of domains for which DNSSEC is deployed is an order of magnitude higher for large operators.

B. DNSSEC Security in .nl and .se

In this section, we evaluate the difference in security levels between DNSSEC deployments from large operators and DNSSEC deployments from small operators. In particular, we focus on the .nl and .se top-level domains where incentives are provided to evaluate whether these also encourage a secure deployment. First, we provide an overview of the average compliance to NIST best practices by large and small operators. Then, we analyze the compliance of single operators to evaluate whether the observed effect can be explained by one or few 'outlier' operators. The results are used to test our running hypothesis defined in Section III.

1) Overview of large and small operators: We first provide a birds-eye view of the data by providing figures on average compliance per domain for large and small operators.

a) Comparison of key algorithms: The first block of Tab. V illustrates the comparison of key algorithms chosen by large and small DNS operators for domains under .nl and .se, based on the latest snapshot of data on July 31, 2017. In .nl small DNS operators perform better, albeit by a small margin, than large DNS operators. In .se the figure seems to be inverted. This effect may be attributed to large operators in other TLDs, that happen to have only a small presence in .se. Still, the overall difference between large and small operators appears to be relatively close for the algorithm criterion.

b) Comparison of key size: We compare the RSA key sizes used by large and small operators based on the latest snapshot on July 31, 2017. The second and third blocks of Tab. V show the comparison of KSK and ZSK key size between the two groups of DNS operators. Similarly to the previous case, large operators perform slightly worse than small operators, with a more marked difference for .se domains. As we will see further down, like with poor algorithm choices, this can be attributed to a single large operator that does not comply with best practices. With regards

[2]Partial deployments in .com: 45.4%; .net: 46.5%; .org: 51.0%.

TABLE V
COMPARISON OF KEY ALGORITHM, RSA KEY SIZE CHOICE AND ZSK ROLLOVER IMPLEMENTATION BETWEEN LARGE AND SMALL DNS OPERATORS IN .NL AND .SE.

TLD	Operator type	%Recommended	%Unrecommended
	(a) *Key algorithm*		
.nl	Large	76.60%	23.40%
	Small	87.36%	12.64%
.se	Large	100.00%	0.00%
	Small	94.26%	5.74%
	(b) *KSK RSA key size*		
.nl	Large	97.55%	2.45%
	Small	99.17%	0.83%
.se	Large	83.70%	16.30%
	Small	96.25%	3.75%
	(b) *ZSK RSA key size*		
.nl	Large	100.00%	0.00%
	Small	99.92%	0.08%
.se	Large	100.00%	0.00%
	Small	98.26%	1.74%
	(c) *ZSK rollover*		
.nl	Large	8.19%	91.81%
	Small	39.36%	60.64%
.se	Large	6.21%	93.79%
	Small	43.00%	57.00%

TABLE VI
LARGE DNS OPERATORS IN TLDS .NL AND .SE

DNS operator	Master NS†	#Signed	Algorithm	KSK size	ZSK size	ZSK Rollover
	TLD .nl					
TransIP	*.transip.net.	265,341	✗	✓	⚠+	✗
	*.transip.nl.	206,254	✗	✓	⚠+	✗
	*.sonexo.eu.	75,256	✓	✓	⚠+	✗
	ns0.nl.	50,273	✗	✓	⚠+	✗
Metaregistrar BV	*.metaregistrar.nl.	386,913	✓	✓	⚠+	✗
Hostnet BV Network	*.hostnet.nl.	359,793	✓	✓	⚠+	✗
Cyso Hosting	*.firstfind.nl.	246,385	✓	✓	⚠+	✗
Argeweb BV	*.argewebhosting.eu.	101,993	✓	✓	⚠+	✗
Openprovider	*.openprovider.nl.	79,367	✓	✓	⚠+	✗
Village Media BV	*.webhostingserver.nl.	67,150	✓	✓	⚠+	✗
Hosting2GO	*.hosting2go.nl.	64,568	✓	✓	⚠+	✗
Flexwebhosting BV	*.flexwebhosting.nl.	60,753	✓	✓	⚠+	✗
Internedservices	*.is.nl.	57,033	✓	✓	⚠+	✗
Neostrada	*.neostrada.nl.	56,295	✓	✓	⚠+	✗
One.com	*.one.com.	55,397	✓	✗	✓	?
PCextreme	*.pcextreme.nl.	50,102	✓	✓	⚠+	✗
AXC B.V.	*.axc.nl.	47,861	✓	✓	⚠+	✗
	TLD .se					
Loopia AB	*.loopia.se.	282,604	✓	✓	⚠+	✗
One.com	*.one.com.	221,372	✓	⚠⋆	⚠+	✗
Binero AB	*.binero.se.	123,131	✓	✓	⚠+	✗

Legend: ✓: meets recommendation; ✗: does not meet recommendation; ⚠: only partially meets recommendation; ?: unknown.
†The master name server from the SOA records is used to identify the operator, as described in Section III-A.
⋆About half of One.com *.se* domains use unrecommended KSK sizes.
+These operators have 1024-bit ZSKs that require regular key rollovers according to the best practice (Tab. II); as the rollover column shows, however, they do not perform key rollover for ZSK.

to the ZSK key size, all DNS operators sign with a key of suitable length. This is not the whole story however, as the most commonly chosen key length (1024 bits) requires regular key rollovers to be performed. As we will see, this is an area where almost all of our large operators perform poorly.

c) Comparison of key rollover: We also compared the ZSK key rollover implementations of large and small DNS operators. We did not evaluate the rollover for KSKs and CSKs due to the lack of sufficient data (the period over which we have data is too short, given that NIST recommends to rollover KSK and CSKs up to every 2 years). The last block of Tab. V shows the percentage of domains in the *recommended* and *unrecommended* categories. As can be seen, small DNS operators perform significantly better than large DNS operators in both .nl and .se. If we look at the two main causes for the high percentages in the unrecommended category, we observe that while a small fraction (less than 1%) can be explained by late key rollovers, the majority of domains in this category (over 90%) have never had their ZSKs replaced over the entire duration of our datasets. This shows that not performing key rollovers is the biggest problem in DNSSEC. Since the key rollover process is quite complex and is not required in order to be eligible for economic incentives, large DNS operators may choose to avoid the extra effort and risk of performing regular key rollovers. Furthermore, we observe this behavior consistently over all large DNS operators in .nl and .se.

2) Detailed analysis of single operators: To evaluate whether the results presented above may be the result of a specific skew in the data (e.g. 'bad' operators that manage a significant fraction of records), we now evaluate the specific deployments at the operator level. Tab. VI shows the detailed analysis for large operators. In general, we find that the descriptive results reported in the previous section hold in the detailed analysis as well. Overall, we observe that large operators perform well for security configurations that can be addressed by a one-time setting in the service configuration. For example, the key size comes by default with the server configuration and requires an effortless change at installation time to be set up correctly. The triangle for ZSKs indicates that, whereas DNS operators satisfy, at large, the NIST requirement on the key size, an appropriate setting for ZSKs can only be identified by the *combination* of key size with rollover frequency (see Tab. II). For example, a ZSK of 1024 bits, while in principle acceptable, needs to be rolled at least every 90 days. In this respect, the systematic lack of compliance to the key rollover mechanism for DNSSEC deployments in large operators leads to a general inadequacy of ZSK key sizes. This reflects the well documented complexity of managing key rollovers in all its phases, including announcement, publication, and retiring of old keys (see, e.g., RFC 6781, Section 4.1).

On the other hand, we observe a few cases where the behavior of single, individual operators may explain some of the divergences observed (both in a positive and a negative direction). The operator *TransIP* has non-uniform algorithm

choices over the four `.nl` master name servers it manages. In particular for three of its servers it uses DNSSEC algorithm 7 (RSA-SHA1), which is considered insecure. This operator by itself explains the apparent misaligned performance of large operators for choice of algorithm reported in Tab. V, whereas all other operators perform adequately. Similarly, *One.com* is the only operator that fails to meet the requirement for KSK size. At the same time, this same operator is the only one to choose a ZSK size greater than 1024 bits (*One.com* uses 1280-bit ZSKs for a large fraction of domains). We report this in Tab. VI as a question mark for key rollover, as this larger ZSK requires less frequent rollovers (similar to KSKs, so every 1-2 years). As our dataset spans less than two years, we cannot yet make claims about rollover behavior in this case. Overall, the result that is certainly confirmed for large operators is the lack of compliance on key rollover. For conciseness, we continue our analysis on small operators focusing on this aspect only.

As shown in Tab. V, the average small operator performs significantly better than large operators where key rollover is concerned. Given the high number of operators classified as *small*, to simplify the breakdown analysis by operator we randomly sampled 100 small operators from the 4,432 small operators for `.nl` and 1,232 small operators for `.se`.[3] We classify these as behaving correctly or badly in terms of operating ZSK rollover by employing a standard 'majority voting' mechanism, whereby we mark an operator as behaving correctly if it implements ZSK rollover following NIST guidelines for 50%+1 of the domains it is responsible for. Fig. 5 shows the distribution of this analysis. Overall, we observe that the incidence of 'bad' operators do not show any clear dependence with the number of signed domains. This indicates that for this category the averages reported in the previous sections can be considered good representations of the performance of small operators. We note that repeating the analysis over multiple samples reveals that the reported distribution is stable.

3) Summary: Overall, we observe that large DNS operators do not appropriately implement the more complex aspects of DNSSEC – such as key rollovers – and seem driven to only fulfill the minimum requirements to sign DNSSEC records. This ineffective security deployment happens *in spite* of the presence of economic incentives that are to the advantage of large operators: the hypothesized 'perverse' effect whereby lack of controls on deployment security favors deployment en-masse but do not provide a significant incentive for secure deployments finds support in the data.

V. Recommendations to Registry Operators

Economic incentives have successfully stimulated the *quantity* of DNSSEC deployment. Our study shows, however, that economic incentives do not provide a similar effect on the *quality* of DNSSEC deployment. Therefore, we advocate changing the focus of economic incentives from quantity toward quality. In this spirit, we have shared our results with the registries for `.nl` and `.se` to solicit their feedback on the work and encourage them – and other registries that have, or are considering, similar incentive schemes – to incorporate quality as a specific metric in future incentive programs.

[3]A plot of the full distribution, not reported here for clarity of representation, leads to the same observations.

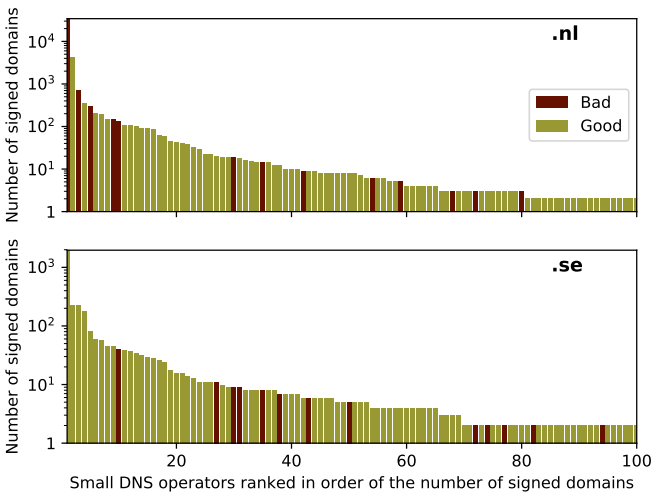

Fig. 5. Compliance of small DNS operators to NIST rollover guidelines

Our findings strongly point in the direction of the ineffectiveness of incentivizing DNSSEC deployments without a clear metric for the *security* of the adopted measures. This demands the establishment of sound and reproducible metrics for DNSSEC security that scale well with the infrastructure and that can be used as a basis for more balanced economic incentives that promote a *secure* deployment of the standard as opposed to solely a *massive* one. These aspects have also emerged from a private discussion with the Swedish registry *Internetstiftelsen i Sverige* (IIS), that confirmed that whereas the initial intention of their incentive was to promote the quantity of the deployment, the complexity of a correct enforcement of security guidelines may hinder its quality, as our study shows. Following this discussion, IIS stated that incentivizing quality is a long term goal and that IIS may use its incentive programs as a tool to spur on quality. In particular, they may launch two DNSSEC incentive programs: one program similar to the current one and another with higher rewards but stricter requirements on the quality of DNSSEC deployment. In a similar vein, SIDN – the registry operator for `.nl` – has indicated that they will use the results of our study as input for discussions on their economic incentive program for DNSSEC deployment when this program is up for renewal in 2018.

Recommendations: Since incentive programs have shown their effectiveness in encouraging the growth of DNSSEC adoption, we suggest an actionable plan to use incentive programs to improve the quality of DNSSEC deployment. Specifically, key rollover is the most critical problem as discussed in the previous section; hence, addressing this issue will significantly enhance the security of DNSSEC deployment. As we alluded to in Section III-B, however, it is challenging for registries to measure key rollover for domains under their management. On the other hand, it is easier for registries to

incorporate the key size and key algorithm requirements in their incentive programs; therefore, we suggest (1) and (2) for immediate actions, and (3) for a long-term plan:

1) Registries should at least include guidelines for key algorithms to be used as a mandatory part of their incentives. In this respect, elliptic curve algorithms are emerging as a viable alternative to RSA since they offer equal or higher security than RSA alongside other benefits [25]. Therefore, we suggest that registries strongly encourage a move towards elliptic curve algorithms.
2) Registries should explicitly define required key sizes in their incentive programs. Since elliptic curve algorithms are introduced with fixed key sizes, registries can omit this suggestion. If RSA is the preferred algorithm, guidelines for each key type should be clearly stated (as in Tab. II). As a result, keys with better key algorithms and key sizes can be used safely for a longer time, which will essentially address the current key rollover issue.
3) For a long-term plan, checking for regular key rollovers should be a part of incentive programs. Given that this is decidedly non-trivial to implement, we suggest the use of platforms like OpenINTEL, that provide established methodologies to perform this measurement. In this context, we note that, where possible, the OpenINTEL platform releases its datasets as open data[4].

VI. Related Work

This section reviews related DNSSEC studies and highlights the differences with this study.

We first consider studies of the challenges of DNSSEC deployment. Yang et al. [26] discuss the challenges that DNSSEC deployment has faced and draw lessons from it for the design of cryptographic protocols tailored to large-scale deployment. Osterweil et al. [27] examined DNSSEC deployment via three critical metrics: availability, verifiability and validity, and highlighted a number of unexpected challenges.

Since the interest in DNSSEC has been growing recently, there are more efforts in measuring and quantifying DNSSEC deployment. The Internet Society published an updated state of DNSSEC deployment [28] in which they highlight that 89% of TLDs and 47% of ccTLDs are signed. Adrichem et al. [29] examined the misconfiguration of DNSSEC deployment and found that over 4% of evaluated domains contain misconfigurations. Similarly, Deccio et al. [30] analyzed a six-month dataset to examine DNSSEC deployment issues and found out that nearly 20% of zones experienced invalid signatures. Notably, Wander [31] studied server-side DNSSEC adoption, in which he analyzed 22 months of data and provided insights into key management, NSEC/NSEC3 and signature validation in TLDs and 3.4 million second-level domains. Chung et al. [22] performed a longitudinal measurement study of the DNSSEC ecosystem and found critical problems in signature validity, key size, key rollover and DS record maintenance. Recently, Chung et al. [23] investigated why DNSSEC adoption remains relatively low by focusing on the role of registrars.

The main difference between previous studies and this work is that we focus on the quality of DNSSEC deployment in the presence of economic incentives for DNSSEC deployment. To this end, we study two ccTLDs, `.nl` and `.se`, that have such economic incentive programs and have significant numbers of DNSSEC-signed domains (in fact, `.nl` currently has the largest number of signed domains in absolute numbers of any TLD).

VII. Conclusions and Future Work

DNSSEC is a vital tool to improve the security of the DNS and the security of the Internet. Some TLD operators have stimulated deployment of DNSSEC through economic incentives. There is evidence that this has led to an increase in the *quantity* of DNSSEC-signed domains [10], [11]. In this paper we investigated whether these incentives have also had a positive influence on the *quality* (in terms of security) of these deployments. Our results clearly show that, unfortunately, economic incentives may be ineffective in promoting a secure deployment of DNSSEC. This appears to suggest that in the absence of clear quality criteria that need to be met to qualify for economic incentives, operators will simply perform the bare minimum in terms of DNSSEC deployment to reap the benefits of the incentive programs.

Given the security benefits of DNSSEC, it is of paramount importance that the technology is deployed *securely*, according to well-established best practices. To this end, we advocate that economic incentives for deploying DNSSEC must include quality criteria. Given the apparent success in boosting the *quantity* of DNSSEC deployment, we have good hopes that amending incentives to include quality requirements will equally help boost the *quality* of DNSSEC deployment. To aid registry operators in implementing such quality requirements, we provide a set of actionable recommendations in this paper, that can help design new incentive schemes.

Future Work – As we discussed in our recommendations in Section V, it is difficult to measure some aspects of DNSSEC best practices, such as, e.g., regular key rollovers. We are therefore actively working with registry operators to perform this task in a reliable and systematic way, based on data from the OpenINTEL platform. In addition to this, we see a need for qualitative studies on how operators deploy DNSSEC. Existing studies such as the one from ENISA [21] are outdated and do not consider recent developments in the area of DNSSEC deployment, especially for key rollover. The study of current practices would provide insights into the actual costs and efforts sustained by operators for DNSSEC deployment. The results of such a study can help in defining future schemes for economic incentives stimulating an effective deployment of DNSSEC and in improving DNSSEC best practices.

Acknowledgements – This work was funded by SURF, the Netherlands collaborative organisation for ICT in higher education and research. The research leading to the results presented in this paper was made possible by OpenINTEL [15], a joint project of SURFnet, the University of Twente and SIDN.

[4]See https://openintel.nl/, under "Data Access".

REFERENCES

[1] P. Mockapetris, "Domain Names - Concepts and Facilities," RFC 1034, 1987.

[2] S. M. Bellovin, "Using the Domain Name System for System Break-ins," in *Proceedings of 5th USENIX Security Symposium*, 1995.

[3] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose, "DNS Security Introduction and Requirements," RFC 4033, 2005.

[4] ICANN Research. TLD DNSSEC Report. http://stats.research.icann.org/dns/tld_report/.

[5] R. van Rijswijk-Deij, "Improving DNS Security, A Measurement-Based Approach," Ph.D. dissertation, University of Twente, 2017.

[6] ICANN. Celebrating 10 Years of DNSSEC Workshops. https://singapore52.icann.org/en/schedule/wed-dnssec/presentation-dnssec-workshop-topics-11feb15-en.pdf.

[7] (2014) ISDN Verlengt DNSSEC Kortingsregeling tot 1 Juli 2018. https://www.ispam.nl/archives/38957/sidn-verlengt-dnssec-kortingsregeling-tot-1-juli-2018/.

[8] (2014) DNSSEC Deployment in Sweden. https://london50.icann.org/en/schedule/wed-dnssec/presentation-dnssec-deployment-sweden-25jun14-en.pdf.

[9] J. Brännlund, Internetstiftelsen i Sverige. Personal Communication.

[10] M. Davids. (2016) DNSSEC in .nl. https://www.sidnlabs.nl/downloads/presentations/SIDN-Labs-InternetNL-20160316.pdf.

[11] R. Mohan. (2012) Slowly Cracking the DNSSEC Code at ICANN 43. https://afilias.info/blogs/ram-mohan/slowly-cracking-dnssec-code-icann-43.

[12] R. Chandramouli and S. Rose, "Secure Domain Name System (DNS) Deployment Guide," National Institute of Standards and Technology, NIST SP 800-81-2, 2013.

[13] E. B. Barker and Q. H. Dang, "Recommendation for Key Management Part 3: Application-Specific Key Management Guidance," National Institute of Standards and Technology, NIST SP 800-57 Pt3 Rev 1, 2015.

[14] R. van Rijswijk-Deij, M. Jonker, A. Sperotto, and A. Pras, "A High-Performance, Scalable Infrastructure for Large-Scale Active DNS Measurements," *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 7, 2016.

[15] OpenINTEL Active DNS Measurement Project. https://www.openintel.nl/.

[16] O. Kolkman, W. Mekking, and R. Gieben, "DNSSEC Operational Practices, Version 2," RFC 6781, 2012.

[17] Security Tools and Architectures Section of ENISA and Deloitte Enterprise Risk Services, "The Costs oF DNSSEC Deployment," European Network and Information Security Agency, ENISA Reports, 2009.

[18] Verisign, "The Domain Name Industry Brief Q3 2017," Reston, VA, USA, Tech. Rep. 4, 2017. [Online]. Available: https://www.verisign.com/assets/domain-name-report-Q32017.pdf

[19] P. Mockapetris, "Domain Names - Implementation and Specification," RFC 1035, 1987.

[20] M. O. Lorenz, "Methods of Measuring the Concentration of Wealth," *Publications of the American Statistical Association*, vol. 9, no. 70, pp. 209–219, 1905.

[21] P. Saragiotis, "Good Practices Guide for Deploying DNSSEC," European Union Agency for Network and Information Security, ENISA Reports, 2010.

[22] T. Chung, R. van Rijswijk-Deij, B. Chandrasekaran, D. Choffnes, D. Levin, B. M. Maggs, A. Mislove, and C. Wilson, "A Longitudinal, End-to-End View of the DNSSEC Ecosystem," in *Proceedings of the 26th USENIX Security Symposium*. USENIX Association, 2017, pp. 1307–1322.

[23] T. Chung, R. van Rijswijk-Deij, D. Choffnes, D. Levin, B. M. Maggs, A. Mislove, and C. Wilson, "Understanding the Role of Registrars in DNSSEC Deployment," in *Proceedings of the 17th ACM SIGCOMM Conference on Internet Measurement*, 2017.

[24] G. van den Broek, R. van Rijswijk-Deij, A. Sperotto, and A. Pras, "DNSSEC Meets Real World: Dealing with Unreachability Caused by Fragmentation," *IEEE Communications Magazine*, vol. 52, no. April, pp. 154–160, 2014.

[25] R. van Rijswijk-Deij, A. Sperotto, and A. Pras, "Making the Case for Elliptic Curves in DNSSEC," *ACM Computer Communication Review*, vol. 45, no. 5, 2015.

[26] H. Yang, E. Osterweil, D. Massey, S. Lu, and L. Zhang, "Deploying Cryptography in Internet-Scale Systems: A Case Study on DNSSEC," *IEEE Transactions on Dependable and Secure Computing*, vol. 8, no. 5, pp. 656–669, 2011.

[27] E. Osterweil, M. Ryan, D. Massey, and L. Zhang, "Quantifying the Operational Status of the DNSSEC Deployment," in *Proceedings of the 8th ACM SIGCOMM Conference on Internet Measurement*. ACM, 2008, pp. 231–242.

[28] "State of DNSSEC Deployment," Internet Society, Reports, 2016.

[29] N. L. M. van Adrichem, N. Blenn, A. R. Lua, X. Wang, M. Wasif, F. Fatturrahman, and F. A. Kuipers, "A Measurement Study of DNSSEC Misconfigurations," *Security Informatics*, vol. 4, no. 1, p. 8, 2015.

[30] C. Deccio, J. Sedayao, K. Kant, and P. Mohapatra, "A Case for Comprehensive DNSSEC Monitoring and Analysis Tools," in *Proceedings of Securing and Trusting Internet Names Workshop*. National Physical Laboratory, 2011.

[31] M. Wander, "Measurement Survey of Server-Side DNSSEC Adoption," in *Proceedings of Network Traffic Measurement and Analysis Conference*. IEEE, 2017.