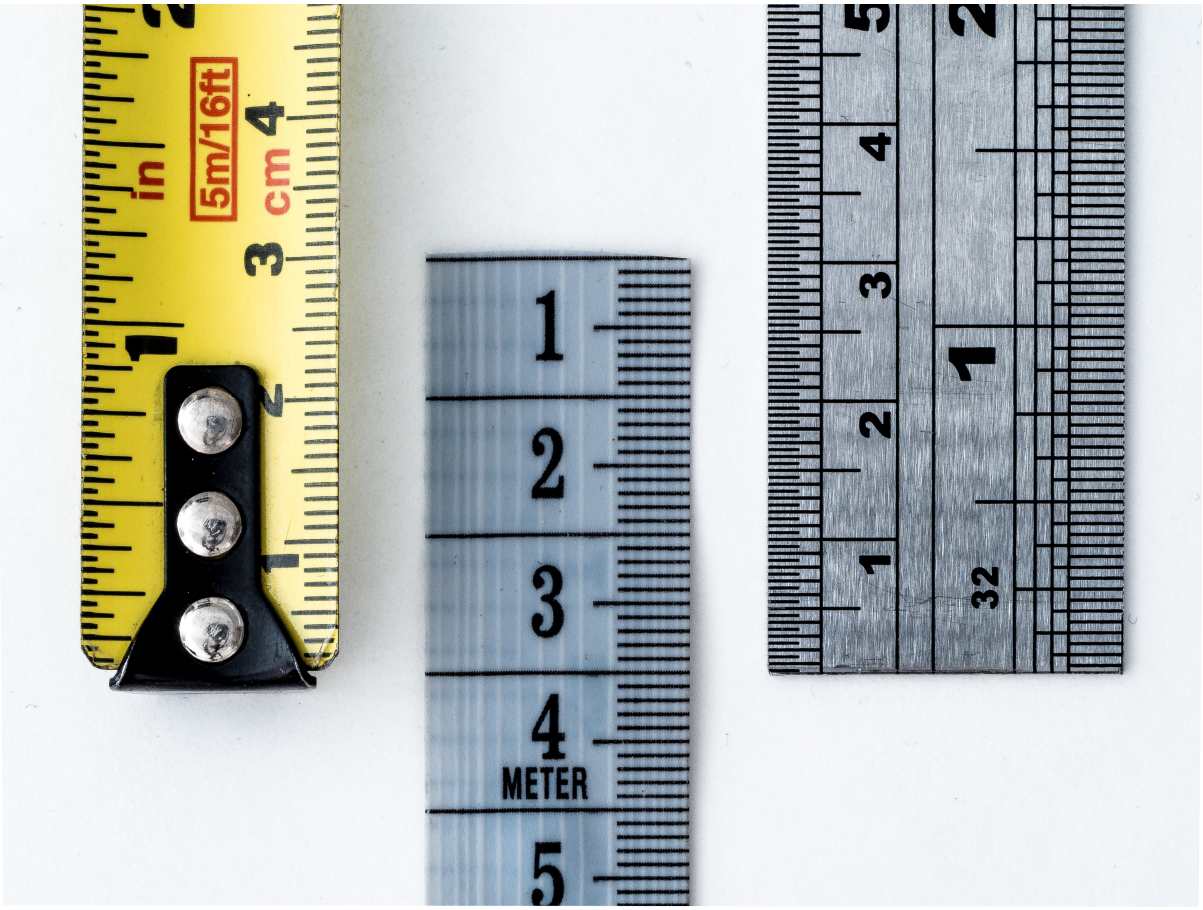

DNSSEC Deployment Metrics Research

2022/08/08

MORITZ MÜLLER, SIDN Labs
JELTE JANSEN, SIDN Labs
MARCO DAVIDS, SIDN Labs
WILLEM TOOROP, NLnet Labs

Executive Summary

The DNS Security Extensions (DNSSEC) add integrity and origin authenticity to the DNS, *the* Internet's naming system. Without DNSSEC, information in the DNS can be manipulated, rendering Internet users vulnerable to a plethora of different threats.

Measuring the deployment of DNSSEC is crucial for gaining insights into deployment drivers, barriers, and its overall state. These insights enable the ICANN community and other stakeholders to plan and influence the future of DNSSEC. Since the publication of DNSSEC, its deployment has been measured by both the DNS community and the academic community. These efforts have resulted in a large number of measurement techniques which focus on different aspects of DNSSEC and the DNS ecosystem.

This is the report on the DNSSEC Deployment Metrics Research, conducted on behalf of ICANN, with three deliverables: (i) to carry out a survey of academic and industry literature related to the deployment of DNSSEC, (ii) to inventory different techniques and metrics used to measure aspects of DNSSEC deployment across the Internet, and (iii) to recommend ICANN with relevant DNSSEC deployment metrics.

This report lists 64 metrics identified in literature, for measuring aspects of DNSSEC deployment at recursive resolvers, domain names, end-users, DNS software, and the wider DNS ecosystem. The identified metrics cover a wide range of aspects and include both metrics directly related to DNSSEC deployment, e.g. whether a domain name is signed, or a recursive resolver is validating, to metrics indirectly related to DNSSEC deployment, e.g. the reliability of the underlying transport. In order to identify the metrics, we searched not only scientific publications, but also publications and presentations by the DNS industry. A comprehensive bibliography can be found at the end of this report.

Most identified metrics can be measured using multiple measurement techniques. For example, we can measure whether a recursive resolver is validating by querying specific domain names and observing the answers, or by observing queries by the resolver at authoritative name servers. Each measurement technique has different advantages and disadvantages. In this report, we describe an assessment framework that allows us and the community to assess

measurement techniques based on three different attributes: the coverage that can be achieved with the measurement technique, whether there are barriers to achieving the coverage, and whether the measurement techniques create reproducible results. We apply that framework to the measurement techniques identified in this study, but the framework can also be applied to new measurement techniques that emerge in the future. For each of the identified metrics, we recommend one measurement technique that we think is most suitable.

Finally, this report makes recommendations to ICANN and the community regarding the DNSSEC metrics that should be considered for further use. In order to make those recommendations, we define three goals that ICANN and the wider community might want to achieve in the future. First: *protecting DNS transactions with DNSSEC*. Instead of focussing on deployment numbers only, the community might want to focus on how many DNS transactions are actually protected. Second: *increasing DNSSEC deployment quality*. In recent years, the community has put a lot of effort into getting DNSSEC deployed as widely as possible. In the future, instead of focusing on quantity, the community might want to focus on the quality of DNSSEC deployment. Third: *making DNSSEC deployments future-proof*. The DNS and security landscape is ever changing. The community might therefore want to focus on preparing DNSSEC deployments for these future changes. For all three goals, we commend metrics that provide the insights needed for goal realisation.

Contents

Abstract 1
Contents 4
1 Introduction 5
2 Background 7
2.1 DNSSEC signing 7
2.2 DNSSEC validation 8
2.3 Clients 8
3 Approach 9
3.1 Assessment framework 9
3.2 Literature study 9
3.3 Recommendations 10
3.4 Limitations 10
4 Assessment framework 12
4.1 Requirement 1: Coverage 13
4.2 Requirement 2: Feasibility 13
4.3 Requirement 3: Reproducibility 14
5 Measurement techniques 17
5.1 Active measurements 17
5.2 Passive measurements 18
5.3 Manual collection of data 19
6 Metrics 20
6.1 Domain name metrics 20
6.2 Resolver metrics 30
6.3 Other metrics 35
7 Recommendations 40
7.1 Goal 1: Protecting DNS transactions with DNSSEC 40
7.2 Goal 2: Increasing DNSSEC deployment quality 41
7.3 Goal 3: Making DNSSEC deployments future-proof 42
References 44
A Overview of metrics and related studies 50
B Acronyms 53

1 INTRODUCTION

DNS Security Extensions (DNSSEC) add authenticity and integrity to the DNS. That allows for the protection of DNS data itself, but also provides additional building blocks to add or improve security in other protocols. However, 17 years after the publication of the current DNSSEC standards, and 12 years after the root zone was signed, the deployment of DNSSEC remains patchy.

This report contains the results of DNSSEC Deployment Metrics Research, conducted on behalf of ICANN.[1] This project has three *deliverables*:

(I) A survey of academic and industry literature related to the deployment of DNSSEC
(II) An inventory of the different techniques and metrics used to measure aspects of DNSSEC deployment across the Internet
(III) A comprehensive report detailing the inventory and advising ICANN on relevant DNSSEC deployment metrics

This document constitutes *deliverable III*.

This report does not attempt to explain why DNSSEC deployment numbers are the way they are, or why DNSSEC deployment is higher in some areas than in others. Nor does this report directly suggest initiatives or incentives to improve deployment numbers. That said, to evaluate the effectiveness of any initiative, one requires a comprehensive way of measuring deployment.

In this study, we have identified 64 different metrics, but not all are necessarily relevant to ICANN and the broader Domain Name System (DNS) community. This report helps ICANN to identify the metrics that we deem most relevant. We do so by by defining three *deployment goals* that ICANN and the DNS community might want to achieve. For each *deployment goal*, we recommend a selection of metrics that provide ICANN and the community with the insights needed for goal realisation. The goals are as follows:

(1) Protecting DNS transactions with DNSSEC
(2) Increasing DNSSEC deployment quality
(3) Making DNSSEC deployments future-proof

The first goal, *protecting DNS transactions with DNSSEC* is derived from two, high-level, deployment metrics: (i) the number of DNSSEC-signed zones that are published and (ii) the number of DNS resolvers that validate DNSSEC signatures. Because of the way the DNS works and the DNS is used, with the hierarchical distribution of authoritative name servers, the ability to cache data in resolvers, and the uneven usage of Internet resources, those metrics are increasingly difficult to measure accurately on a global scale. We therefore define a deployment goal combining the two metrics above, namely the number of client queries that are protected.

[1] https://www.icann.org/en/announcements/details/request-for-proposal-researching-dnssec-deployment-metrics-17-5-2021-en

That goal is both a metric in its own right and a function of the first two; the more zones are signed, and the more resolvers validate the signatures, the more client queries are protected.

The second goal, *increasing DNSSEC deployment quality*, goes beyond deployment numbers, and aims for high quality deployments. By measuring how many operational problems occur with, or because of, DNSSEC deployments, the causes of those problems can be identified and addressed. Furthermore, if the prevalence of problems is low, the use of measured data to demonstrate that fact can reduce fears of complexity, and increase trust in the deployability of DNSSEC.

The third goal, *making DNSSEC deployments future-proof*, is related to the second goal, but focusses on the future of DNSSEC instead, of current DNSSEC deployments. We argue that certain requirements need to be fulfilled to make sure that DNSSEC continues to serve the DNS and the Internet community in the future, and we propose metrics for measuring the extent to which that is the case.

The remainder of this report is structured as follows. In Section 2, we briefly provide background information required for the reader of this report. Then, in Section 3, we describe the approach we took to achieve *deliverable I* and *deliverable II*. In Section 4, we introduce our assessment framework, used to evaluate measurement techniques for metric collection. That is followed by Section 5, in which we give an overview of the different measurement techniques and their attributes. In Section 6, we list all the metrics identified in the course of the project, provide a short description and recommend suitable measurement techniques. that section addresses *deliverable I*. We then recommend relevant metrics to ICANN and the DNS community in Section 7. For a comprehensive list of metrics and the related publications, see Appendix A (*deliverable II*).

2 BACKGROUND

DNSSEC adds integrity and origin authenticity to the DNS. It does so by enabling domain name operators to cryptographically sign their zones, and allowing everyone who queries information in the DNS to validate the signatures. In this section, we introduce the most relevant aspects for this study: DNSSEC, DNS, the underlying transport protocol and the ecosystem in which DNSSEC is being deployed. This introduction also defines the scope for deployment metrics, discussed later in this report.

For more detailed information about the different components and aspects see the tutorial paper by Van der Toorn et al. [van der Toorn et al. 2022] or the related standard documents.

2.1 DNSSEC signing

This section is structured to follow the main steps domain name operators need to take to roll out DNSSEC and maintain DNSSEC for their domain names. In most cases, we describe the steps from the perspective of a second-level domain name operator, but the relevant aspects mostly apply to other levels of the DNS hierarchy as well. We assume that the operator has already registered their domain name with a registrar, who has communicated the registration to the registry for the relevant Top Level Domain (TLD). Note that domain name registrants may handle DNS infrastructure operations themselves or, rely on third party DNS operators.

Signing. When deploying DNSSEC for a domain name, an operator needs to make a number of choices. First, a zone can be signed using different cryptographic signing algorithms. Each signing algorithm has its own attributes (e.g. key and signature size, or validation and signing speed) and not every algorithm is equally widely supported by DNS signing software and hardware, at validating clients or in the DNS ecosystem. Also, an operator needs to decide how to prove that a give record is *not* part of the zone. Here, operators have the choice between NSEC and NSEC3 [Laurie et al. 2008], each of which has its own set of configurable parameters. Also, the operator has the choice of splitting their keys into a Zone Signing Key (ZSK) and a Key Signing Key (KSK), or relying on one Combined Signing Key (CSK) only.

The operator needs to make sure that the underlying DNS software or the Hardware Security Module (HSM), if used, supports the chosen parameters.

After signing the zone, the operator should communicate their public KSK to their parent. Usually, that means relying on their registrar to relay the key information to the relevant registry. It is therefore important that the registrar and the registry support DNSSEC and the chosen parameters as well. In some cases, it is possible to relay the key to the parent directly, using CDS/CDNSKEY [Gudmundsson and Wouters 2017]. Regardless of whether the key material is communicated directly or indirectly, the parent needs to have its zone signed as well, so that DNSSEC creates a chain of trust up to the root. The root's public key acts as the trust anchor for signed domain names.

Operations. A DNSSEC signature has an expiration date. The lifetime of a signature is decided by the operator. A new signature needs to be created before the old one expire. Also, it is common practice to replace signing keys on a regular basis. With both operations, the operator needs to be careful to follow the right timing schedule and the correct procedures. Otherwise, recursive resolvers may be unable to validate the domain name's signatures. When a operator replaces their KSK, they need to communicate the change to their parent.

A replacement key may be based on the same cryptographic algorithm as the one it replaces, or on another algorithm. A switch may be required if, for example, the "old" algorithm is no longer considered insecure. For the most parts, an algorithm switch is similar to replacing the KSK.

There are other operational choices relevant to the deployment of DNSSEC as well, even though they might not be directly related to the DNSSEC protocol itself. For example, signed DNSSEC resource records necessitate larger responses. Operators therefore need to make sure that sufficiently large messages can be transmitted reliably. Moreover, the larger size of a signed zone might make it attractive for attackers to be misused in a Distributed Denial of Service (DDoS) reflection attack. An operator may therefore want to take certain steps to make their zone less attractive to attackers.

2.2 DNSSEC validation

Operators that enable DNSSEC validation on recursive resolvers influence which signing algorithms are supported through their choice of resolver software and underlying cryptographic libraries. Also, they need to make sure that they have configured the public key of the root zone as a trust anchor. Instead of validating signatures and enforcing validation failures, resolver operators can decide to validate signatures but to not act on validation errors, or can even decide to only request DNSSEC records and not perform validation at all.

Recursive resolver operators need to make sure that they have reliable transport to the authoritative name servers, and to their clients. Of course, they need to do that even if DNSSEC is not deployed, but the demands are higher with DNSSEC because of the increased message sizes. Resolver operators may additionally decide to protect the communication channels between client and authoritative name server against eavesdropping, by enabling support for encrypted DNS transport protocols. The other end of the communication channel needs to support the transport protocol as well. As a side-effect, the transport protocols lower the barrier for transmitting large DNS messages.

2.3 Clients

At present, clients that would like to request information from the DNS usually rely on their recursive resolvers to perform DNSSEC validation. In the interest of redundancy, many clients have two or more recursive resolvers at their disposal. A client is only fully protected by DNSSEC when all recursive resolvers have enabled validation.

3 APPROACH

In order to identify, assess, and recommend DNSSEC metrics, we split the work into three stages. First, we develop a framework to assess techniques used to collect DNSSEC metrics. A metric can be collected using different measurement techniques, but not necessarily all are suitable for collecting the metric over the long term or on a large scale. Second, we perform a literature study, identifying relevant metrics for this project. Third, we formulate recommendations, regarding the metrics ICANN should consider.

In this section, we describe the approaches to each stage in more detail.

3.1 Assessment framework

The assessment framework, which is described Section 4, is intended to help us and future researchers to find the most suitable measurement technique for collecting a given metric. We identify three requirements that a measurement technique should fulfil and develop a rating system to assess the performance of a technique for each requirement.

We identify the requirements in close contact with ICANN, drawing on our own expertise in DNS measurements (e.g. [Müller et al. 2019b, 2020; Toorop 2019]), and on feedback from the community. We have also presented the requirements to the DNS and measurement communities at ICANN, IRTF and CENTR meetings [Müller 2022a,b,c].

3.2 Literature study

The goal of the literature study presented in Section 6 is to identify relevant metrics for measuring DNSSEC deployment.

Relevant metrics. We consider a metric to be relevant if it falls into one of the following categories. *(i)* The metric measures aspects directly related to the DNSSEC protocol, such as whether a domain name is signed or whether a resolver performs DNSSEC validation. *(ii)* The metric measures aspects related to the DNS protocol that can have an effect on DNSSEC deployments, such as support for EDNS(0), or whether DNS response rate limiting is enabled on the authoritative name server. *(iii)* The metric measures aspects related to the underlying transport protocols that have an influence on transporting DNSSEC records, e.g. support for TCP fallback. *(iv)* The metric measures aspects related to DNS and Internet security that are orthogonal to DNSSEC deployment or secure other aspects of the DNS, such as deployment of RPKI or DNS-over-HTTPS (DoH).

Relevant venues and publications. In order to identify metrics, we focus on high-tier security and Internet measurement conferences. We manually assess each publication presented at such a conference, identifying whether the publication discusses metrics that fall into one of the above categories. If a publication discusses a relevant metric, then we also go through the publication's references. That allows us to identify interesting publications that were not

Table 1. Academic and other venues studied for relevant metrics.

Conference/Venue	Years
ACM Internet Measurement Conference (IMC)	2010 – 2021
USENIX Security	2004 – 2021
USENIX Summit on Hot Topics in Security (HotSec)	2006 – 2021
USENIX Workshop on Offensive Technologies (WOOT)	2011 – 2020
USENIX Workshop on Free and Open Communications on the Internet (FOCI)	2014 – 2020
USENIX Symposium on Usable Privacy and Security (SOUP)	2018 – 2020
RIPE Meetings	2006 – 2022
DNS-OARC Workshops	2015 – 2022
CENTR Research & Development Workshops	2012 – 2022

published at higher-tier and more mainstream venues. We search for relevant metrics in the referenced publications as well.

We also assess publications and presentations made at other venues, namely at meetings of the RIPE community, DNS Operations, Analysis, and Research Center (DNS-OARC) workshops, and Council of European National Top-Level Domain Registries (CENTR) meetings. We apply the same criteria to metrics presented at those venues. If a metric was presented by the same author at both, an academic venue and an industry venue, we generally include only the reference to the academic venue.

See Table 1 for a list of assessed venues.

3.3 Recommendations

Finally, in Section 7, we recommend the metrics that we believe should be taken into account in order to measure DNSSEC deployment as comprehensive as possible, and explain how they can be used. A measurement should not be performed for its own the sake, but only if it supports the realisation of a greater goal. Depending on the goal it might differ which *aspects* of DNSSEC deployment are relevant, and therefore which metrics should be taken into account.

For that reason, we formulate three goals that affect DNSSEC deployment and that we believe are relevant for ICANN and the DNS community. The goals relate to both current and future DNSSEC deployments. For each goal, we recommend metrics and associated measurement techniques. The recommended metrics and techniques should help ICANN and the community to gain the insights necessary to make strategic decisions that influence DNSSEC deployment, and to achieve their greater goals.

3.4 Limitations

We discuss two sets of limitations: limitations relating to our own work and limitations of the collected metrics.

Study. We carry out an extensive literature study, but we cannot rule out the possibility that we miss some relevant metrics and measurement studies. First, because we focus on well-known venues and journals, we potentially miss studies published at local and small-scale workshops. Second, because we collect only metrics described in the English language. Nevertheless, we are confident that we cover the vast majority of relevant metrics due to our own expertise in DNS, DNSSEC, and Internet measurements.

Metrics. As with all measurements on the Internet, no metric is able to reflect the reality with 100% accuracy [Paxson 2004]. In the DNS, for example, anycast allows name servers to be distributed across different locations, but it also makes it harder to measure all the authoritative name servers for a domain name. Also, a recursive resolver might only be reachable from within a certain network, making it harder or even in some cases impossible to measure. Finally, the information in the DNS name space is public, but not necessarily known to everyone. In other words, anyone can request any information in the public DNS name space, but only if they know that the information exists. That too is a challenge in relation to DNS metric collection, and can limit the overall coverage. Some zones (e.g. the root and some TLDs) publish their zone files, enabling researchers to measure all the child names of a parent, but 100% coverage of all domain names is not feasible.

We discuss the limitations of the various measurement techniques in more detail in Section 5.

Table 2. Assessment framework for measurement techniques

	Coverage	Feasibility	Reproducibility
Rating			
++	Covers every instance of a component.	Coverage can be achieved independently and with public data only. Also, costs have no impact on the coverage.	The technique is extensively documented and well understood.
+	Covers a representative share of a component.	Coverage can be achieved independently and with public data only. Costs can have some impact on the coverage.	The technique is documented and important parameters of the technique are known.
=	Covers large parts of a component, but with some unknown bias.	Coverage depends partially on third parties. Costs can have some impact on the coverage.	The documentation lacks some of the important parameters.
-	Covers only small parts of a component with a known bias.	Coverage depends largely on third parties or, costs have large impact on the coverage.	The technique lacks the majority of the most important parameters.
- -	Covers only small parts of a component with an unknown bias.	Coverage depends completely or almost completely on third parties.	The documentation lacks completely or is not accessible.

4 ASSESSMENT FRAMEWORK

Metrics can be collected using various measurement techniques. This framework is intended to help us and other researchers to select the most suitable measurement techniques. Every measurement technique, used to collect DNSSEC deployment metrics comes with advantages and disadvantages. The framework allows us to assess the measurement techniques identified in this project in a semi-structured way.

First, we identify which requirements the measurement techniques should fulfil. Fulfilling a requirement contributes to *deliverable I* to gain *"the most comprehensive insight into DNSSEC deployment"*. We use input by ICANN Office of the CTO (OCTO), the DNS and domain-name community, and our own judgment to define the requirements.

Many measurement techniques cannot fully fulfil a requirement. We therefore introduce a rating between - - (does not fulfil the requirement) and *++* (fully fulfils the requirement) for each requirement.

In the next section, we apply the framework to the identified measurement techniques.

4.1 Requirement 1: Coverage

We define three basic requirements. The first is *Coverage*, defined as the extent to which a measurement technique covers a DNS component, as described in Section 2. The perfect measurement technique can cover every instance of a component and would thus receive the rating *++*. As an example, a measurement technique able to collect data on the name servers of every existing domain name would achieve 100% coverage.

In many cases, although it is not be possible to cover 100% of a component, it is possible to cover a representative proportion. In such cases, the technique receives the rating *+*. So, for example, that rating is given to a technique that is able to measure the recursive resolvers that answer 80% of the queries of the internet population, where it is clear which part of the population the technique cannot cover.

A measurement technique that can cover large parts of the component, but without clarity as to which parts of the component are missed, receives the rating =. An example could be a measurement technique that can observe queries to two of the three authoritative name servers of a TLD. In that case, the technique has access to roughly two thirds of the overall queries, but it is not clear which queries are missing, or where the missing queries originated.

Subpar measurement techniques receive the rating -. That rating is given to a technique that can cover only small parts of a component and has a significant, but known, bias. Examples could include active measurements of recursive resolvers with a limited set of vantage points located in only in the western hemisphere.

Measurement techniques that cover only small parts of a component and have unknown biases receive the rating - -. These could include a novel technique that collects telemetry information from resolvers that have already been upgraded to share the telemetry.

4.2 Requirement 2: Feasibility

In many cases, a measurement technique can achieve full coverage only *theoretically*. In practice, the coverage of a DNS measurement technique depends largely on the number of available vantage points or the known domain names. Also, some measurement techniques require larger investments when scaled up or run for an extended, continuous period. A measurement technique that achieves its coverage only by relying on public data and has no significant costs would receive the rating *++*. An example could be a measurement technique that embeds code on a website under the control of the researcher, triggering HTTP requests from which the validation status of the client's resolver can be deduced. Even though the total

coverage might be limited, the coverage can be achieved by relying solely on data collected by the researcher at low costs.

We give measurement techniques that can run without relying on third-party data, but which require some investment, a + rating. As an example, a measurement technique that relies on authoritative name server traffic under the control of the researcher would receive this rating. Here, the researcher has full access to the data, but storing and processing the data requires some investment, due to the volume of the data involved.

The rating = applies to measurement techniques that rely on third-party data or vantage points to achieve their full coverage. In such cases, the cost again has an impact on the coverage. This rating could, for example, apply to a technique that measures domain names. Complete coverage would require the collection of zone files, which might not always be publicly available.

Measurement techniques that rely on third parties to a large extent, or whose coverage is substantially influenced by cost, receive the rating -. For example, a technique that would require a significant investment to achieve its full coverage would receive such a rating.

Finally, the lowest rating - - applies to measurement techniques that rely fully on third parties, either for their vantage points or for their measurement targets. The rating - - is not applied to measurement techniques that rely on third party platforms if the cost does not negatively impact the coverage. Examples included measurements on authoritative name servers that are not under the control of the researcher.

4.3 Requirement 3: Reproducibility

Metrics are useful to a broader community only if they are collected in a transparent and reproducible way. If someone is to make strategic decisions based on the collected metrics, it must be clear how the results have been collected and processed. It is therefore also important that all the main parameters of the underlying measurement technique are known. Parameters include, for example, the locations of the vantage points, how many measurements have been carried out, whether and how the data is cleaned before it is being processed, and how it is processed.

Measurement techniques whose main parameters are known receive the rating ++. Active measurement that relies on open-source software, running on vantage points known to the community, and with an extensive public documentation would receive the highest rating.

A measurement technique receives the rating + if some steps are undocumented, but that does not affect the measurement results when the measurements are repeated by another researcher. As an example, a measurement technique that lacks documentation on some post-processing steps that do not affect its accuracy would receive this rating.

If any relevant parameters are missing, a measurement technique cannot receive a rating higher than =. This could, for example, apply to measurement techniques that rely on third-party measurement platforms. Here, the researcher might not always have full control over

the selected vantage points, but still be able to demonstrate the exact result processing steps followed.

If most of the main parameters are missing, a measurement technique receives the rating -. That is the case with, for example, a measurement technique that relies on a proxy network, which provides only limited documentation.

Measurement techniques whose main parameters are largely undocumented, or whose main parameters are proprietary and therefore unavailable to the community, receive the lowest rating - -. For example, this could include measurements run on a commercial measurement platform whose post-processing steps are not disclosed.

Table 3. Measurement techniques used to collect DNS metrics and their performance in terms of coverage, feasibility, and reproducibility.

Technique	Description	Coverage	Feasibility	Reproducibility
Active measurements				
Active DNS probes	Active DNS queries from vantage points (VPs) under the control of the researcher. Example: Python script running on one or multiple local machines.	+	=	++
Advertising network	Code distributed with advertising network to trigger DNS queries or embedded on website statically. Example: Google Ads [Google 2022]	=	=	+
Proxy network	Proxy network, used for example provide access to geo-blocked content, that allows sending packets from the participating users. Example: Bright Data [Bright Data 2022]	-	-	-
Measurement Platform	Platform, implemented to send active probes from one or multiple vantage points. Example: RIPE Atlas [RIPE NCC Staff 2015]	=	=	+
Response manipulation at Name Server (NS)	Manipulating DNS responses in order to observe client behaviour. Example: Detect if resolvers validate by responding with bogus signatures and observing retries.	+	+	++
Active resolver telemetry	Diagnosis information, voluntarily exposed by DNS component through active probing. Example: RFC 8509 [Huston et al. 2018]	+	=	++
Passive measurements				
Traffic traces at NS	Raw traffic collected at one or multiple authoritative name servers, controlled by the researcher. Example: ENTRADA [Wullink et al. 2016]	+	+	++
Aggregated traffic traces at NS	Aggregated traffic collected at one or multiple authoritative name servers or recursive resolvers, controlled by a third party. Example: ITHITOOLS [Durand and Huitema 2022]	+	+	++
Passive resolver telemetry	Diagnosis information, voluntarily exposed by DNS component through passive measurements. Example: RFC 8145 [Wessels et al. 2017]	+	=	++

5 MEASUREMENT TECHNIQUES

In this section, we introduce the measurement techniques identified in the literature. We then apply the assessment framework defined in the previous section to assess their coverage, feasibility, and reproducibility. In some cases, a given measurement techniques can be used to collect multiple metrics. In such a case, to provide an overview and avoid repetition, we provide only a basic description of the measurement technique here. In Section 6, we state which technique can be used for each metric, and recommend the one that is most suitable.

5.1 Active measurements

We differentiate between measurements that require sending active probes and measurements that rely on collecting data from one or more VP. The first four techniques have in common that they all involve actively sending queries towards authoritative name servers or recursive resolvers, but they differ in terms of the trigger events. For example, researchers can send out queries directly from a machine under their control, or can rely on an advertisement network that facilitates advertising impressions shown to end users. In all cases, the queries are initiated on demand, whereas the queries observed using passive measurement techniques are not triggered by a researcher.

Active DNS probes. Sends queries from VPs under the control of the researchers to targets of their choice. Targets include recursive resolvers and authoritative name servers. High *coverage* can be achieved if the list of targets is complete. Anycast or forwarding resolvers can limit the coverage achieved (-). *Feasibility* is mediocre (=). A complete list of targets might be hard to collect (e.g. for all second-level domain names or all resolvers in the IPv6 address space). The technique is highly *reproducible* if standard DNS software is used, the measurement targets are public, and the data processing is well documented (*++*).

Advertising network. Uses online advertisements to distribute dynamic or static code to end-user devices. When loaded, the code issues requests (usually HTTP) to resources under the control of the researcher. These requests are preceded by DNS queries, relying on the recursive resolvers of the client. Researchers can derive information about the recursive resolvers used by analysing the DNS queries, and the ad network's status. *Coverage* of end users can be high, but insights into server systems is lacking (=). The more money is invested in the measurements, the more end users will be served with the advertisement. Running large scale measurements for an extended period is *feasible* (=). *Reproducibility* is dependent on the information and settings provided by the ad network. We assume that the ad network's clients receive sufficient information, e.g. on where the advertisement is shown (*+*).

Proxy network. Uses proxy network participants as VPs. Commercial proxy networks can enable third parties to send out queries through their participating clients. Similar to the advertising network, this can help researchers to derive information about the recursive resolvers used by the client base. These networks cover mostly end users and *coverage* is dependent

on proxy network coverage. We expect it to be smaller than that of advertisement networks (-). Again, coverage is dependent on financial resources. If sufficient funds are available, continuous measurements are *feasible* (+). *Reproducibility* is dependent on the information and settings made available by the proxy provider. We expect limited, opaque possibilities for tuning measurements (-).

Measurement platform. Relies on VPs tailored for running common active measurements, and enables researchers to trigger DNS queries directly. With RIPE Atlas, for example, a researcher can send queries and modify a limited set of parameters. In many cases, the platform vantage points are run by volunteers, which introduces a bias towards a technical community and limits *coverage* (=). Costs are usually low, but sometimes require credits that can be earned through hosting a measurement VP. Coverage depends on third parties, which makes large-scale measurements less *feasible* (=). On the other hand, measurements carried out with measurement platforms are usually very easy to *reproduce* (+). Measurement platforms are well documented and raw measurement results are often public by default.

Response manipulation on NS. Involves manipulation of DNS responses on the authoritative name server to derive information about the behaviour of recursive resolvers. For example, deliberately responding with the `TC`-flag set should cause the resolver to retry the query via TCP. From that, researchers can deduce whether a resolver supports TCP fall-back. *Coverage* depends greatly on the name server whose responses are manipulated but, if that is the root, for example, it can be high (+). Because the technique can be implemented on the researchers' infrastructure, its *feasibility* is good (+). The measurement technique is also readily *reproducible*, since the main parameters of the measurement are known (++).

Active resolver telemetry. Involves sharing information about resolvers attributes directly with the researcher. Recursive resolvers and authoritative name servers can make information about their operational status available to active queriers. That is the case with root sentinel queries (RFC 8509 [Pauly and Wouters 2019]), for example, and with CHAOS queries. The *coverage* of the measurement technique can be high, if implemented in an official standard (+). In practice, however, high coverage is less *feasible*, since it depends on the resolvers supporting this telemetry protocol (=). That might be low if the telemetry protocol has only been standardised only recently. Also, coverage depends on the number of vantage points (see also "active DNS probes"). On the other hand, measurements can be highly *reproducible*, if the telemetry protocol is an official standard (++).

5.2 Passive measurements

Traffic traces on NS or recursive resolvers. Involves collecting unfiltered traffic traces on DNS components themselves. Traffic collected on one or multiple authoritative name servers can reveal information about the querying resolvers. Similarly, traffic traces collected on recursive resolvers can reveal information about their clients and the queried authoritative name servers.

As with “response manipulation”, *coverage* depends on the server whose traffic is collected and can be high if the client base is large (as with the root) (+). Collecting traffic on a large scale, however, might become less *feasible* if traffic is collected for a longer period, or if researchers depend on a third party (+). Collecting data is straightforward, and, if the post-processing steps are well documented, the technique is easy to reproduce (++).

Aggregated traffic traces on NS or recursive resolvers. Involves collecting processed DNS traffic traces on authoritative name servers or recursive resolvers. With this technique, data is aggregated, usually by a third party. Once again, *coverage* depends on the vantage point where the data is collected, but can be high if the VP is, for example, the root (+). The fact that data is collected in an aggregated form can make it easier for third parties to collect the data. Aggregating data might address the security and privacy concerns of third parties and requires less storage. That makes larger-scale measurements more *feasible* (++). If data is aggregated in a standardised and transparent way, measurements are *reproducible* (++).

Passive resolver telemetry. Involves sharing information on the attributes of resolvers directly with the researcher. In contrast to “active resolver telemetry", resolvers can also share telemetry information with others on an unsolicited basis. This is the case with RFC 8145 [Wessels et al. 2017], where resolvers share trust anchor configurations with the root servers, usually once per day. As with active telemetry protocols, *coverage* can be high (+), but is heavily dependent on the deployment (feasibility =). Again, measurements are highly *reproducible* when the telemetry protocol is an official standard (++).

5.3 Manual collection of data

Some aspects cannot be measured automatically or do not relate to technical aspects of DNSSEC deployment. Examples are whether a registrar supports DNSSEC or whether operators rely on an HSMs to secure their keys and create their signatures. More examples are listed in Section 6.3.

Collecting these metrics could require the manual inspection of software, the manual registration of domain names by individual registrars, the distribution of questionnaires to operators, and other approaches. Such techniques are necessary in some cases, and we also mention them in our recommendations (Section 7). However, we do not describe them in more detail, because the range of possible approaches is very wide.

6 METRICS

In this section, we discuss metrics that are described in literature or that were presented at conferences. First, we describe each metric in more detail. We then describe how the metric can be measured, using the technique that we have identified as most suitable. In some cases, we also recommend techniques whose use for measurement of the metric is not described in the literature, but which we nevertheless believe to be the most suitable. Finally, we make reference to a publication that describes the measurement of the metrics using the recommended technique. Tables 4, 5, and 6 give an overview of the identified metrics. Table 7 in Appendix A lists every study that discusses each metric.

We first discuss metrics concerning the attributes of domain names and their authoritative name servers, followed by metrics concerning the attributes of recursive resolvers, and by metrics concerning the attributes of stub resolvers and their clients. Finally, we discuss metrics relating to the DNS ecosystem, including registries, registrars and DNS operators. Each metric is identified by its category, e.g. MRx for a resolver metrics, or MDy for a domain name metric.

6.1 Domain name metrics

Table 4. Domain name and authoritative name server metrics and their suitable measurement techniques. Techniques marked with ❍ can measure the metric, but techniques marked with ☆ are recommended.

	Metric	Description	Active DNS probes	Advertising network	Proxy network	Measurement platform	Traffic traces at NS	Aggregated traffic traces at NS	Aggregated traffic traces at resolver	Response manipulation at NS	Resolver telemetry report
Domain Name and Name Server Metrics											
MD1	DNSSEC publication	Are DNSSEC records available for a domain name	☆								
MD2	DNSSEC validity	Is the domain name *secure*	☆						❍		
MD3	DNSSEC validation errors	Can the RRs, partially, not be validated	☆								
MD4	Key attribute: Key length	What is the bit-length of the used keys	☆								

Table 4. Domain name and authoritative name server metrics and their suitable measurement techniques. Techniques marked with ❍ can measure the metric, but techniques marked with ☆ are recommended.

	Metric	Description	Active DNS probes	Advertising network	Proxy network	Measurement platform	Traffic traces at NS	Aggregated traffic traces at NS	Aggregated traffic traces at resolver	Response manipulation at NS	Resolver telemetry report
MD5	Key attribute: Time To Live (TTL)	What is the TTL of the key records	☆								
MD6	Key attribute: shared keys	How many zones share the same key	☆								
MD7	Key attribute: Keys in the RR set	How many public keys are in the `DNSKEY` RR set	☆								
MD8	Key attribute: algorithm	Which cryptographic algorithm is used to sign a domain name	☆								
MD9	Key attributes: `DS` digest type	Which hash algorithm is used to calculate the `DS` record	☆								
MD10	Key attribute: RSA modulus	Which prime number is used when creating the public-private RSA key pair	☆								
MD11	Key attribute: ZSK/KSK or CSK	Is a domain name signed using a CSK	☆								
MD12	Key attribute: Key-tag exists multiple times	Does the key-tag exist more than once in the same zone or in other zones	☆								
MD13	Key attribute: Response size	How large are the responses for `DNSKEY` queries	☆			❍	❍				
MD14	NSEC/NSEC3 usage	Which protocol is used to proof denial of existence	☆			❍	❍				
MD15	Transport: Path MTU	What is the largest supported MTU by the name servers	❍						☆		
MD16	Transport: TCP support	Does the name servers support TCP	☆								

Table 4. Domain name and authoritative name server metrics and their suitable measurement techniques. Techniques marked with ❍ can measure the metric, but techniques marked with ☆ are recommended.

	Metric	Description	Active DNS probes	Advertising network	Proxy network	Measurement platform	Traffic traces at NS	Aggregated traffic traces at NS	Aggregated traffic traces at resolver	Response manipulation at NS	Resolver telemetry report
MD17	EDNS(0) Support	Does the authoritative name server fully support Extension Mechanisms for DNS (EDNS0) [Damas et al. 2013]	☆								
MD18	Transport: DNS-over-TLS (DoT)	Is the authoritative name servers reachable via DoT	☆								
MD19	Transport: DNS-over-QUIC (DoQ)	Is the authoritative name servers reachable via DoQ	☆								
MD20	Operations: Signature lifetime	How long is the lifetime of the published RRSIGs	☆								
MD21	Operations: Rollover frequency ZSK	How often is the ZSK rolled	☆								
MD22	Operations: Rollover frequency KSK	How often is the KSK rolled	☆								
MD23	Operations: Rollover correctness	Are keys and algorithms rolled without going insecure or bogus	☆								
MD24	Operations: Rollover type	Which process is followed to roll ZSK, KSK or algorithm	☆								
MD25	Operations: Algorithm rollover	Has a domain name ever rolled its algorithm	☆								
MD26	Operations: Name Server operator	Which organisation/instance operates the authoritative name server	☆								
MD27	Operations: Operational complexity	How many operators are responsible for managing the authoritative name servers	☆								
MD28	Operations: Effects on query load	Does DNSSEC increases the traffic at authoritative name servers					❍	☆			

Table 4. Domain name and authoritative name server metrics and their suitable measurement techniques. Techniques marked with ❍ can measure the metric, but techniques marked with ☆ are recommended.

	Metric	Description	Active DNS probes	Advertising network	Proxy network	Measurement platform	Traffic traces at NS	Aggregated traffic traces at NS	Aggregated traffic traces at resolver	Response manipulation at NS	Resolver telemetry report
MD29	Operations: Effects on TCP load	Does DNSSEC cause increase in TCP traffic at authoritative name servers	☆				❍				
MD30	Operations: CDS/CDNSKEY publication	Does the zone publish CDS/CDNSKEY RRs	☆								
MD31	Name Server: DNS Cookie support	Does the authoritative name servers support DNS Server Cookies [Sury et al. 2021]	☆								
MD32	Name Server: Response rate limiting	Does the authoritative name servers employ response rate limiting	☆								
MD33	Name Server: Minimal responses to query type ANY	Does the name server provides minimal responses when receiving a query for type ANY [Abley et al. 2019]	☆								
MD34	RPKI: Served from RPKI signed resources	Is the authoritative name servers located in a network which is covered by a valid ROA	☆								

MD1: DNSSEC publication. Measures whether DNSSEC records are published for DNS zones. *Recommended measurement technique:* Active DNS probes – by querying for signed records, e.g. SOA. Coverage of the metric depends on the available list of domain names. See [Chung et al. 2017a].

MD2: DNSSEC validity. Measures whether DNSSEC records are valid. The number of DNSSEC-specific errors that are caused by operational mistakes is relevant for DNSSEC adoption, as this hinders day-to-day DNS use. Perhaps more importantly, these problems do not occur on networks that do not use validating resolvers, giving the impression that such networks have

better connectivity than networks that do offer DNSSEC protection. That may cause operators to stop validating DNSSEC signatures. Therefore, it is important to track operational DNSSEC errors, and reduce them as much as possible. *Recommended measurement technique:* Active DNS probes – by querying and validating DNSSEC records. Coverage of the metric depends on the available list of domain names. See [Osterweil et al. 2008].

MD3: DNSSEC validation errors. Measures which DNSSEC errors occur. Whereas *MD2* merely reflects whether a DNSSEC record is valid or not, this metric states the actual error types. Collecting statistics on DNSSEC error types can be useful in efforts to reduce problems, as this can direct policy and tooling to help operators avoid the most prevalent issues. *Recommended measurement technique:* Active DNS probes – by querying and validating DNSSEC records. Coverage of the metric depends on the available list of domain names. See [Osterweil et al. 2008].

MD4: Key attribute: key length. Measures the length in bits of the keys used (in relation to the algorithm used). This is relevant in relation to RFC 8624 [Wouters and Sury 2019]. This should be measured in combination with the algorithm used and is not relevant for, for example, the ECDSA algorithms because of their fixed key length. *Recommended measurement technique:* Active DNS probes – by requesting keys and signatures of signed domain names. Coverage of the metric depends on the available list of domain names. See [Chung et al. 2017a].

MD5: Key attribute: Time To Live. The TTLs of DNSSEC resource records have an influence on key replacement and on the impact of outages caused by DNSSEC misconfigurations. *Recommended measurement technique:* Active DNS probes – by requesting keys and signatures of signed domain names and collecting the Round-Trip Times (RTTs). Coverage of the metric depends on the available list of domain names. See [Chung et al. 2017a].

MD6: Key attribute: shared keys. Measures the number of zones that share the same key. A key that is shared by numerous domain names is more valuable to attackers. *Recommended measurement technique:* Active DNS probes – by requesting keys of signed domain names. Coverage of the metric depends on the available list of domain names. See [Chung et al. 2017a].

MD7: Key attributes: number of keys in RR set. A large number of unnecessarily long-lasting DNSKEYs in an RRset may cause larger than needed responses, reducing the retrievability of the RRset and making it more prone to exploitation in amplification denial-of-service attacks. *Recommended measurement technique:* Active DNS probes – by querying for `DNSKEY` records. Coverage of the metric depends on the available list of domain names. See [van Rijswijk-Deij et al. 2014]

MD8: Key attributes: used algorithm. Measure the algorithms used for signing zones. *Recommended measurement technique:* Active DNS probes – by querying for `DNSKEY` records. Coverage of the metric depends on the available list of domain names. One study that focused

on the usage of various DNSKEY algorithms for signing over time is [Müller et al. 2020]. The study leverages data collected by the OpenINTEL platform [van Rijswijk-Deij et al. 2016b].

MD9: Key attributes: DS digest type. Measure the various digest types in use by Delegation Signer resource records. *Recommended measurement technique:* Active DNS probes – by querying for `DS` records. Coverage of the metric depends on the available list of domain names. In [Müller et al. 2020], the usage of DS digest types over time is also inventoried based on data from the OpenINTEL platform [van Rijswijk-Deij et al. 2016b].

MD10: Key attributes: vulnerable RSA keys due to even moduli. If an RSA key is not generated carefully, it can have vulnerabilities which may compromise the encryption algorithm. Sometimes this can be determined from the public key alone. *Recommended measurement technique:* Active DNS probes – by querying for `DNSKEY` records. Coverage of the metric depends on the available list of domain names. An inventory of misconfigurations in popular domains can be found in [Dai et al. 2016].

MD11: ZSK/KSK or CSK. Measures if a domain name is relying on a split between ZSK and KSK or whether it relies on a CSK. This has implications for key management and response size. *Recommended measurement technique:* Active DNS probes – by querying for `DNSKEY` records. Coverage of the metric depends on the available list of domain names. See [Le et al. 2018].

MD12: Key-tag exists multiple times. Measures if the key tag exists multiple times, either in the same zone or in multiple zones. Duplicated key tags can increase the workload for validating resolvers seeking to identify the correct key for validating a signature. *Recommended measurement technique:* Active DNS probes – by querying for `DNSKEY` records. Coverage of the metric depends on the available list of domain names. See [van Rijswijk 2019].

MD13: Response size. Measures the response size associated with domain name DNSKEY queries. Large responses could be abused in DDoS attacks or could lead to message fragmentation. *Recommended measurement technique:* Active DNS probes – by querying for `DNSKEY` records. Coverage of the metric depends on the available list of domain names. See [Müller et al. 2019b].

MD14: NSEC/NSEC3 Usage. Measures which approaches domain names choose for authenticated denial of existence of records. Relevant, for example, to measure whether hash iterations for NSEC3 records are sufficiently low. *Recommended measurement technique:* Active DNS probes – by querying for likely non-existent records. See [Wander 2017].

MD15: Transport: Path MTU. Measures the maximum MTU between the authoritative servers for a zone and the resolver(s) querying those servers. This is relevant for DNSSEC, as fragmentation issues may result in operational problems. *Recommended measurement technique:* Traffic traces from (validating) DNS resolvers, which can be searched for fragmented IP packets. See [Van Den Broek et al. 2014].

MD16: Transport: TCP support. Measures if a domain name resolves via TCP. This is relevant for DNSSEC, because DNSSEC-signed responses are bigger than unsigned ones, increasing the probability of truncated replies. *Recommended measurement technique:* Active DNS probes – by requesting resource records via TCP from all authoritative name servers. Coverage of the metric depends on the available list of domain names. See [Osterweil et al. 2008].

MD17: EDNS(0) Support. Measures whether name servers support the Extension Mechanisms for DNS (EDNS(0)) [Damas et al. 2013] and whether they comply with the standard. Supporting EDNS(0) is a prerequisite for supporting DNSSEC and lays the foundation for other DNS extensions, like DNS Cookies [Sury et al. 2021]. *Recommended measurement technique:* Active DNS probes, sending a number of queries using different EDNS(0) parameters. See the "DNS-Compliance-Testing" tool, for examples [Risk 2015].

MD18: Transport: DoT. Measures whether name servers support DNS-over-TLS (DoT), which protects DNS traffic between a client and a server against eavesdropping and interference, by encrypting the traffic. The relation to DNSSEC is that relying on TLS for transport would remove the risk of message fragmentation, thereby allowing for larger messages. Larger messages might become necessary if new signing algorithms are adopted in the future [Müller et al. 2020]. *Recommended measurement technique:* Active DNS probes, testing DoT capabilities of known name servers. See [Kosek et al. 2022].

MD19: Transport: DoQ. Measures whether authoritative name servers support DoQ, which protects DNS traffic between a client and a server against eavesdropping and interference, by encrypting the traffic with TLS and using on QUIC as a transport protocol. As with DoT, DoQ would make the transmission of larger DNS messages more reliable. *Recommended measurement technique:* Active DNS probes, testing DoQ capabilities of known name servers.

MD20: Signature lifetime. Measures the lifetime of `RRSIG` records. A shorter lifetime requires operators to re-sign a record more frequently but decreases the time window in which a compromised signature is trusted by a recursive resolver. Usually, the TTL of an `RRSIG` record is shorter than the signature lifetime. *Recommended measurement technique:* Active DNS probes – by querying for signed records, e.g. `SOA`. Coverage of the metric depends on the available list of domain names. See [Chung 2021].

MD21: Rollover frequency ZSK. Measures how often ZSKs are replaced (rolled over). Higher rollover frequency reduces the opportunity for a key to be compromised. *Recommended measurement technique:* Active DNS probes – by querying for `DNSKEY` records frequently, e.g. once per day. Coverage of the metric depends on the available list of domain names. See [Deccio 2011].

MD22: Rollover frequency KSK. Measures how often KSKs are rolled. Higher rollover frequency reduces opportunity for a key to be compromised. In contrast to *MD21*, rolling the KSK requires the replacement to be communicated to the parent. That also increases the risk of

misconfiguration and validation failures. *Recommended measurement technique:* Active DNS probes – by querying for DNSKEY records, e.g. once per day. Coverage of the metric depends on the available list of domain names. See [Deccio 2011].

MD23: Rollover correctness. Measures whether a key rollover is performed in accordance with best common practices [Kolkman et al. 2012], while keeping the zone in a *secure* state at every stage of the rollover. Applies to key and algorithm rollovers. *Recommended measurement technique:* Active DNS probes – by querying for DNSKEY records and DS records (for KSK and algorithm rollovers). e.g. once per hour. Coverage of the metric depends on the available list of domain names. See [Osterweil et al. 2021].

MD24: Rollover type. The process followed to roll the ZSK, KSK, CSK or algorithm. *Recommended measurement technique:* Active DNS probes – by querying for DNSKEY records and DS records (for KSK and algorithm rollovers). e.g. once per hour. Coverage of the metric depends on the available list of domain names, which is obviously simpler for the root zone than for TLDs. Platforms such as OpenINTEL can be used. The root zone forms a special case, for which resolver telemetry (RFC 8145) might also be useful for monitoring a key rollover. See [Osterweil et al. 2021].

MD25: Algorithm rollover. Whether a domain name has ever rolled its algorithm. *Recommended measurement technique:* Active DNS probes – by querying for DNSKEY records and DS records (for KSK and algorithm rollovers). e.g. once per hour. Coverage of the metric depends on the available list of domain names. Platforms such as OpenINTEL can be used. See [Müller et al. 2020].

MD26: Operational complexity. E.g. how many operators are responsible for managing the authoritative name servers. Measuring this can be a challenge in certain (anycasted) situations. For example, the .nl TLD previously used its own autonomous system number, while the operation was performed by a third party, on hardware provided by yet another party. *Recommended measurement technique:* No recommended technique. Measured in [Deccio et al. 2011] using an active measurement technique involving the fetching of name server names and information in the MNAME field of the SOA record. However, we do not expect that technique to be very precise.

MD27: Effects on query load. DNSSEC can cause an increase in TCP traffic on authoritative name servers, because there is an increased risk of truncated replies. *Recommended measurement technique:* Aggregated traffic traces at NS.

MD28: Operations: Query load effects of DNSSEC. Effects on query load and DNSSEC-related increase in traffic on authoritative name servers *Recommended measurement technique:* Traffic traces on authoritative name servers See [Minda 2011]

MD29: Operations: TCP query load effects of DNSSEC. Effects on TCP load and DNSSEC-related increase in TCP traffic on authoritative name servers *Recommended measurement technique:* Traffic traces on authoritative name servers See [Minda 2011]

MD30: Operations: CDS/CDNSKEY publication. Measure support for the intention to provide key material from child to parent. *Recommended measurement technique:* Active DNS probes – by querying for `CDS` and `CDNSKEY` records, preferably from a platform with access to a large number of zone files, such as OpenINTEL [van Rijswijk-Deij et al. 2016b].

MD31: Name server attributes: DNS cookie support. Whether authoritative name servers support DNS server cookies. *Recommended measurement technique:* Active DNS probes – by including client cookies in the queries to the authoritative name servers found in zone files. A large number of authoritative name servers can be found from a platform with access to a large number of zone files, such as OpenINTEL [van Rijswijk-Deij et al. 2016b]. The measurements should preferably be done from active DNS probes from multiple vantage points in order to reach all potential anycasted instances of an authoritative server.

MD32: Name Server: Response Rate Limiting. Measures whether authoritative name servers employ Response Rate Limiting (RRL). DNSSEC increases the impact of DNS amplification attacks, as DNSSEC responses are much larger than plain DNS responses. RRL is a technique for mitigating such attacks. *Recommended measurement technique:* Active DNS probes. Note that it requires bursts of traffic to trigger RRL, which may cause operational issues for operators, especially those that do not support rate limiting. Care needs to be taken not to overload name servers when probing for RRL support. See [Deccio et al. 2019].

MD33: Name Server: Minimal responses to query type ANY. Measures whether name servers implement [Abley et al. 2019]. This is not necessarily DNSSEC-related, and does not measure DNSSEC deployment itself. It can, however, be a useful metric for general DNS operations since DNSSEC signatures can make large ANY responses even larger. *Recommended measurement technique:* Active DNS probes. Coverage of the metric depends on the available list of domain names. See [van der Toorn et al. 2021].

MD34: RPKI: Served from RPKI signed resources. Measures whether name servers for domain names use RPKI to protect against route hijacks. This is not directly related to DNSSEC; it is another level of network protection against attacks and misconfigurations by third parties. *Recommended measurement technique:* Combination of name server address data collected by probing with information provided by a route collector. See [Müller 2021].

Reverse tree. The discussed metrics also apply to the reverse DNS tree, even though it is not discussed in the referenced literature. Full coverage of the reverse IPv4 name space is easier to achieve than full coverage of the IPv6 space. While in both cases the name space is limited, the sheer number of possible IPv6 addresses can pose a challenge. If access to Reginal Internet Registry (RIR) databases is available, the search space can be limited.

Table 5. Metrics and their suitable measurement techniques. Techniques marked with ❍ can measure the metric, but techniques marked with ☆ are recommended.

	Metric	Description	Active DNS probes	Advertising network	Proxy network	Measurement platform	Traffic traces at NS	Aggregated traffic traces at NS	Aggregated traffic traces at resolver	Response manipulation at NS	Resolver telemetry report
Resolver Metrics											
MR1	EDNS(0) Support	Does the resolver fully support EDNS0 [Damas et al. 2013].	❍	☆	❍	❍	❍				
MR2	DNSSEC capable	Does a resolver request DNSSEC records	❍	☆	❍	❍	❍	❍			
MR3	DNSSEC validation	Does a resolver validate any DNSSEC record	❍	☆	❍	❍	❍	❍			
MR4	DNSSEC validation and enforcement	Does a resolver validate DNSSEC records and does it return SERVFAIL on a validation error	❍	❍	❍	☆					
MR5	Treatment of specific DNSSEC validation failures	How does a resolver treat specific bogus domain names, e.g. wrong label count or single RRSIG missing	❍	☆	❍	❍	❍	❍			
MR6	Algorithm support	Does a resolver validate signatures of a certain algorithm	❍	☆	❍	❍					
MR7	DNSSEC trust anchor support	Which DNSSEC trust anchors are configured	❍	❍	❍	❍	❍	❍			☆
MR8	Transport: Path MTU	Which DNS message size can a resolver support before a message gets fragmented	❍				☆				
MR9	Transport: TCP support	Does a resolver fall back on fragmented DNS responses	❍	❍			☆				
MR10	DNS protocol: DoH support	Does a resolver support DoH between stub and recursive resolver	❍			☆					
MR11	DNS protocol: DoT support	Does a resolver support DoT	❍			☆					
MR12	DNS protocol: DoQ support	Does a resolver support DoQ	❍			☆					
MR13	DNS Cookie support	Does a resolver send queries with the `COOKIE` option set	❍	☆	❍	❍	❍				

Table 5. Metrics and their suitable measurement techniques. Techniques marked with ❍ can measure the metric, but techniques marked with ✯ are recommended.

	Metric	Description	Active DNS probes	Advertising network	Proxy network	Measurement platform	Traffic traces at NS	Aggregated traffic traces at NS	Aggregated traffic traces at resolver	Response manipulation at NS	Resolver telemetry report
MR14	Telemetry: RFC 8145	Does a resolver share trust anchor information with the root through RFC 8145					✯	❍			❍
MR15	Telemetry: RFC 8509	Does a resolver share trust anchor information with clients through root sentinel queries (RFC 8509)	✯	❍	❍		❍	❍			❍
MR16	RPKI: Served from RPKI signed resources	Is the resolver located in a network which is covered by a valid ROA		✯	❍	❍	❍	❍			
MR17	RPKI: Route origin validation	Does the network of the resolver perform route origin validation		✯	❍	❍					

6.2 Resolver metrics

MR1: EDNS(0) Support at resolvers. Measures a resolvers' support for the Extension Mechanisms for EDNS0 [Damas et al. 2013] and its compliance with the standard. Supporting EDNS0 is a prerequisite for supporting DNSSEC and lays the foundation for other DNS extensions, like DNS Cookies [Sury et al. 2021]. *Recommended measurement technique:* Using an advertisement network to trigger DNS queries to a client's resolvers for a domain name under the researcher's control and observing whether resolvers indicate EDNS support.

MR2: DNSSEC Capable. Measures whether resolvers are capable of fetching DNSSEC records. *Recommended measurement technique:* Using traffic traces on third-party NSs to observe whether resolvers query with the `DO` flag set, and whether they request `DNSKEY` or `DS` records occasionally. Low-effort measurement with wide coverage if performed on popular name servers, such as the root, but sheds no light on forwarding resolvers. See [Gudmundsson and Crocker 2011].

MR3: DNSSEC validation. Measures whether resolvers are validating DNSSEC records. *Recommended measurement technique:* Using an advertising network to measure end-user validation by embedding code in advertisements that trigger HTTP requests to web resources. Some of the web sources have “bogus" signatures, which enables researchers to deduce whether a client relies on validating resolvers. Advertising networks can cover resolvers used by end-users and be deployed on a large scale. However, some resolvers, such as those used by email servers, however, cannot be covered. To increase transparency, raw measurement results should be published. See [Lian et al. 2013].

MR4: DNSSEC validation and enforcement. Measures whether resolvers return SERVFAIL on validation errors. *Recommended measurement technique:* Measurement platform sending queries to resolvers for specific preconfigured domains. Such domains could also contain multiple types of validation failures, such as separate domain names that return bogus or expired signatures. Such domains are most suitable for measurements where there is full control over the queries that are sent to the resolver. See [Müller 2016].

MR5: Treatment of specific DNSSEC validation failures. Similar to *MR3*, but enables more fine-grained distinctions to be made where recursive resolvers do perform DNSSEC validation. Distinctions reported in the literature include whether a validating resolver considers a signature *bogus* when the label count in the `RRSIG` record does not match with the number of labels in the signed `RR` or when a record has no signature at all. *Recommended measurement technique:* As with *MR3*, the researcher can add appropriate test cases to a zone under their control to cover additional failure scenarios. See [Lian et al. 2013].

MR6: Algorithm support. Which signing algorithms resolvers support. *Recommended measurement technique:* Advertising network with embedded Java script in advertisements, which in turn, trigger requests to domain names signed with various algorithms, using valid or bogus signatures. Depending on the success rate, we can derive whether a resolver is validating a certain algorithm, but not another. Advertising networks can cover resolvers used by end-users and can be deployed on a large scale. However, some resolvers, such as those used for example by email servers, cannot be covered. See [Huston 2021b].

MR7: Trust Anchor. Measure which root DNSSEC trust anchors resolvers have available for validation. This metric is relevant mainly during or in the run-up to a root KSK rollover, when a new root KSK for future use is pre-published using [StJohns 2007] and on the iana website, as described in [Abley et al. 2016]. Metric *MR3* already reveals whether the *current* root trust anchor is available, but does not indicate what other root trust anchors the resolver has available, if they are not currently in active use. That information is obtainable only from resolver software itself by means of telemetry signalling. *Recommended measurement technique:* Combining results from active and passive resolver telemetry, as for metrics *MR14* and *MR15*, as described in [Müller et al. 2019b]. However, coverage and feasibility could be substantially improved using additional as of yet non-existent resolver telemetry (or otherwise

determined metrics) which would enable measurement of the query volume for which a resolver is responsible and the position of a resolvers in relation to DNS forwarders, caches and other middle-boxes.

MR8: Transport: Path MTU. Measures the maximum transmission unit (MTU) on the path between the client and the server. When the path MTU is exceeded there may be connectivity problems due to fragmentation issues. Path MTU becomes more important as DNSSEC is more widely deployed, since DNSSEC packets are larger than plain DNS packets. *Recommended measurement technique:* To measure the number of operational issues fragmentation causes, the best approach would be traffic traces from name servers, as described in[Van Den Broek et al. 2014], for instance. However, to measure the potential impact on a more global scale, we recommend an approach as described in [Koolhaas and Slokker 2020].

MR9: Transport: TCP support. Measures whether resolvers support TCP transport. Fallback to TCP when a response from an authoritative server over UDP has the 'truncated' flag set is mandatory, but not always implemented, and sometimes blocked by a firewall. Issues related to this are more prevalent with DNSSEC, since packets are larger than plain DNS packets. *Recommended measurement technique:* Traffic traces on NS. See [Moura 2021]

MR10: DNS Protocol: DoH support. Measures whether resolvers support DNS-over-HTTPS. DoH [Hoffman and McManus 2018], together with DoT [Hu et al. 2016] and DoQ [Huitema et al. 2022], is indirectly related to DNSSEC, since it enables the transport of larger DNS messages. More reliable transport of larger DNS messages gives DNSSEC deployments greater flexibility in their choice of signing algorithms and keys. *Recommended measurement technique:* DoH cannot be measured with RIPE Atlas probes, but RIPE Atlas anchors are an option, among other things such as the NLNOG ring.

MR11: DNS Protocol: DoT support. Measures whether resolvers support DoT. Like DoH and DoQ, DoT provides message confidentiality between client and resolver, and resolver between and authoritative name server and enables larger message sizes. *Recommended measurement technique:* RIPE Atlas can be used to measure DNS-over-TLS support on resolvers, but at the time of writing that option can be selected only by using the API or via tools that interact directly with the API directly such as the command line tools.

MR12: DNS Protocol: DoQ support. Measures whether the resolver supports DoQ. Like DoH and DoT, DoQ provides message encryption and optional authentication, but not origin authentication. DoQ is also relevant to DNSSEC deployment, since it enables larger message sizes (see *MR10*). *Recommended measurement technique*: Measured from measurement platforms like RIPE Atlas. At the point of writing the report, RIPE Atlas does not support DoQ. [Kosek et al. 2022] describes the measurement of DoQ deployment by actively scanning the whole IPv4 address space. That technique therefore measures only open DNS resolvers.

MR13: DNS Cookie support. Measuring whether recursive resolvers send queries with the COOKIE option set. DNS cookies provides limited protection to DNS servers and clients against a variety of increasingly common denial-of-service and amplification forgery attacks [Eastlake 3rd and Andrews 2016]. That is relevant to DNSSEC, since DNSSEC-signed messages are often misused in amplification attacks. *Recommended measurement technique:* Using an advertisement network to trigger DNS queries to a client's resolvers for a domain name the researcher's control and observing queries with the COOKIE option set. See [Lian et al. 2013] for measurements with ad networks and [Davis and Deccio 2021] for measuring DNS cookie support.

MR14: Telemetry: RFC 8145. A metric and a measurement technique. Passive measurement of the DNSSEC trust anchors that a resolver has available for validation for a specific origin (e.g. the root). The resolver signals to the origin of the trust-anchors which trust anchors it has available for DNSSEC validation using a Key Tag Query, as described in Section 5 of [Wessels et al. 2017]. *Recommended measurement technique:* The data is collected on the authoritative name servers that serve the origin of the trust anchor, as described in [Müller et al. 2019b]. This specific telemetry has the following disadvantages: (i) it is impossible to query resolvers with problems for further state information, (ii) there is no telemetry on the query volume a resolver processes, making it hard to judge how relevant or risky a resolver with problems is, and (iii) the Key Tag Query may propagate through upstream systems (NATs, DNS forwarders, caches and other middle-boxes), leading to distorted signals and hiding systems with actual problems. We recommend combining this metric with measurements that would resolve or alleviate those disadvantages.

MR15: Telemetry: RFC 8509. Active measurement of the root DNSSEC trust anchors that a resolver has available for validation. A querier asks the resolver performing DNSSEC validation for the availability of the trust anchor by way of a special name. The resolver provides the availability with a non-standard return code (SERVFAIL) for that name. *Recommended measurement technique:* Ideally using advertising networks, but we are not aware of a study describing the use of that technique for the purpose indicated. For a study that uses the RIPE Atlas measurement platform, see[Müller et al. 2019b]. This telemetry shares the last two disadvantages of *MR14*. We recommend combining this metric with measurements that would resolve or alleviate those disadvantages.

MR16: RPKI: Served from RPKI signed resources. Are resolvers located in networks with valid Route Origin Authorization (ROA). Such recursive resolvers are better protected against routing configuration mistakes and hijacks, provided the networks on the path to the client are validating ROAs (see also *MR17*). *Recommended measurement technique:* Identifying recursive resolvers via an advertisement network, then using RPKI relying party software such as Routinator to verify whether the prefixes of the recursive resolver locations publish valid ROAs. Note, that recursive resolvers can forward their queries to other forwarding resolvers. With the proposed measurement techniques, we can identify only the upstream resolvers. We

cannot determine whether forwarding resolvers "downstream" are served from RPKI-signed resources. See [Lian et al. 2013] for using an advertisement for measurements and [NLnet Labs 2022] for a description of a RPKI relying party software.

MR17: RPKI route origin validation. Does the network of recursive resolvers perform RPKI route origin validation. If a network does, the recursive resolver is less vulnerable to sending queries to authoritative name servers in hijacked networks. *Recommended measurement technique:* Using advertising network and embedding code in online advertisements to trigger queries to a domain name under the researcher's control. One authoritative name server needs to be located in a network that has published an invalid ROA. A downside of that measurement technique is the fact that we cannot identify which network on the path from the recursive resolver to the authoritative name server performs the validation. See [Toorop 2020] for a description of the measurement setup and [Lian et al. 2013] for using an advertisement for measurements.

6.3 Other metrics

Table 6. Client, software, and ecosystem metrics and their suitable measurement techniques. Techniques marked with ❍ can measure the metric, but techniques marked with ☆ are recommended. If no technique is marked then we do not recommend any technique or the metric cannot be measured automatically.

	Metric	Description	Active DNS probes	Advertising network	Proxy network	Measurement platform	Traffic traces at NS	Aggregated traffic traces at NS	Aggregated traffic traces at resolver	Response manipulation at NS	Resolver telemetry report
Client Metrics											
MC1	DNSSEC validation	Does a client rely on validating resolvers		☆	❍	❍					
MC2	Transport: DoH	Does a client rely on DoH as a transport channel to its recursive resolver							☆		
MC3	Upstream resolver type	Does a client rely on internal, forwarding or external recursive resolvers		☆	❍	❍					
DNS Software Metrics											
MS1	RFC compliance	Is a software compliant with a certain RFC									
MS2	Default settings	Which are the default parameters of DNS software									
MS3	Cryptographic library	What is the default library used for cryptographic operations									
MS4	Software deployment	Which software is deployed in the wild	☆								
DNS Ecosystem Metrics											
ME1	DNSSEC support	Does the component of the ecosystem support DNSSEC									
ME2	DNSSEC default	Is DNSSEC enabled by default									
ME3	DNSSEC settings	Can DNSSEC be configured through a web inerface									

Table 6. Client, software, and ecosystem metrics and their suitable measurement techniques. Techniques marked with ❍ can measure the metric, but techniques marked with ✰ are recommended. If no technique is marked then we do not recommend any technique or the metric cannot be measured automatically.

	Metric	Description	Active DNS probes	Advertising network	Proxy network	Measurement platform	Traffic traces at NS	Aggregated traffic traces at NS	Aggregated traffic traces at resolver	Response manipulation at NS	Resolver telemetry report
ME4	Algorithm support	Which algorithms are supported									
ME5	CDS/CDNSKEY support	Does the component scan for CDS/CDNSKEY records									
ME6	DNSSEC fees	Is DNSSEC free of charge									

In addition to metrics for domain names and recursive resolvers, we have identified a number of other metrics that we believe are relevant for measuring DNSSEC deployment.

We have assigned the metrics to three categories. The first category described below consists of metrics relating to aspects of the DNS *clients*, e.g. end users or services that rely on recursive resolvers to look up names in the DNS. The second is metrics concerning DNSSEC-related capabilities and features in DNS *software*. Those include, among others, recursive resolver software, name server software and signing software. The third is metrics concerning aspects of DNSSEC deployment in the *DNS ecosystem*. Registries and registrars are part of this ecosystem, as well as DNS providers.

6.3.1 Clients. Clients rely on recursive resolvers to query records in the DNS. For redundancy, clients can have multiple recursive resolvers configured. Also, they can rely on their own, local, recursive resolvers.

MC1: DNSSEC validation. Measures whether clients rely on a validating resolvers. All of the resolvers at a client's disposal may perform DNSSEC validation, or only a subset of them may do so. *Recommended measurement technique:* Using an advertising network and triggering DNS queries to a domain name under the researcher's control by embedding advertisements

on web sites. Each query must include a unique identifier to attribute multiple queries to one client. See [Osterweil et al. 2008].

MC2: Transport: DoH. Measures whether clients rely on DoH to communicate with their recursive resolvers. An encrypted transport channel between a client and a validating resolver decreases the risk of the client to becoming a victim of DNS manipulation without performing validation itself. A requirement is that the client trusts its recursive resolver. *Recommended measurement technique:* Browser telemetry. Browsers are among the most common implementers of DoH. Browser vendors could provide ICANN with anonymised information on DoH usage. That would enable ICANN to collect information across multiple recursive resolvers. Browser telemetry including information about DoH usage is not publicly available. Alternatively, researchers can rely on statistics published by resolver operators. Resolvers that are chosen by client software by default are of interest (e.g. Cloudflare in the case of the Firefox browser).

MC3: Upstream resolver type. Measures clients rely on resolvers located in their own networks (*internal*), or on resolvers in their own networks that forwards queries to another recursive resolver (*forwarding*), or on resolvers in other network (*external*). Such information is relevant for DNSSEC because it reveals information about resolver centralisation which, in turn, can be a barrier to, or an enabler of DNSSEC deployment. *Recommended measurement technique:* using a measurement platform and triggering DNS queries to a domain name under the researcher's control and observing queries at the authoritative name server. The network of the IP address is then compared with the IP address of the measurement vantage point and the configured recursive resolver. See [Müller et al. 2020].

6.3.2 Software. DNS software can have an influence on DNSSEC deployment. For example, if a feature is not supported by mainstream DNS software, then it is also less likely that the feature gains wider deployment. Also, operators might not compile the recent software versions manually but only rely on the versions shipped with their operating system, with the result that they are not running the latest versions with the latest features. Furthermore, the default settings chosen by software developers can have an influence as well. For example, a mainstream software vendor that stops supporting a certain DNSSEC signing algorithm can have a major impact on overall support of this algorithm.

The following metrics have in common that it is not feasible to measure them automatically and on a larger scale. In most cases, they require the manual analysis of changelogs and software documentation. Nevertheless, we believe that these metrics have an influence on DNSSEC deployment and are therefore worth consideration.

MS1: RFC compliance. What RFCs software supports. Usually, it takes time before DNS software implements changes or new features defined in an RFC. Wider deployment of changes proposed in an RFC cannot be expected until most major DNS software vendors have updated their software accordingly. If possible, compliance can be determined by establishing whether

software or an organisation follows the mandatory (MUST [Bradner 1997]) specifications in an RFC. *Recommended measurement technique:* Manual study of changelogs and manuals.

MS2: Default settings. The default settings in DNS software can have an influence on DNS deployment. For example, if the vendor of a popular DNS software who decides to enable validation by default, that can increase overall DNSSEC validation as soon as the new version is widely deployed. *Recommended measurement technique:* Manual study of changelogs and manuals.

MS3: Cryptographic library. DNS software relies on libraries to perform cryptographic functions. If, for example, a cryptographic library does not support a certain signing algorithm, the recursive resolver software that relies on that library cannot validate signatures created with the algorithm in question. That can hinder wider deployment of new cryptographic algorithms. *Recommended measurement technique:* Manual study of changelogs and manuals. See [Müller et al. 2020] for examples.

MS4: Software deployment. As discussed in relation to metrics *S1* to *S4*, the deployed software and software version can have an impact on DNSSEC deployment. It is therefore relevant to know which DNS software is deployed in the wild. *Recommended measurement technique:* Active DNS probes using custom software. By sending crafted, not standard, queries to resolvers or name servers, one can deduce which software is being deployed with some certainty. That might be more efficient than relying on DNS CHAOS queries only. Coverage depends on the domain name target list and the reachable recursive resolvers. FPDNS [DNS OARC 2020] can send such queries.

MS5: HSM support. In some DNSSEC deployments, such as at the root and some TLDs, HSMs are used to store signing keys in a secure manner and to create signatures. If an HSM does not support a certain algorithm, organisations relying on that HSM cannot roll to the algorithm in question. *Recommended measurement technique:* Manually inventory of HSM types and versions used at the root and TLDs. Studying product information to get an overview of supported algorithms.

6.3.3 DNS ecosystem. The DNS ecosystem is comprised of parties involved in registering domain names and publishing domain names in the DNS. Such parties include registries, managing TLDs, registrars that manage the registration of a domain name, and DNS providers, responsible for publishing the zone of a domain name. A registrar can also have the role of a DNS providers. Registrants are not obliged to rely on DNS providers, but they nevertheless often do so for convenience, reliability, or performance. Not all metrics discussed in this section apply to all components of the ecosystem.

As with metrics relating to DNS software, it is challenging to measure these metrics automatically and on a larger scale. Given its direct relationship with registrars and registrants,

we suggest that ICANN could ask those entities to regularly report metrics. That would allow ICANN to measure the metrics automatically to some extent.

ME1: DNSSEC support. Do registries, registrars or DNS providers support DNSSEC. Depending on the entity, support might include having the TLD signed and publishing the DS records of child domain names, communicating the DS records or DNSKEY records to the parents, or signing the zone. *Recommended measurement technique:* Study of the functionalities supported by DNS providers and registrars. DNSSEC signing of TLDs can be deduced using *MD2*. See [Chung et al. 2017b].

ME2: DNSSEC default. Do DNS providers sign zones by default. *Recommended measurement technique:* Manual study of the features of DNS providers. See [Chung et al. 2017b].

ME3: DNSSEC settings. How can a user configure DNSSEC. In some cases, it might only be possible to enable DNSSEC only by contacting the DNS provider or registrar manually, which might deter DNSSEC deployment. *Recommended measurement technique:* Manual study of the features of DNS providers and registrars. See [Chung et al. 2017b].

ME4: Algorithm support. Which signing algorithms are supported by registries, registrars or DNS providers. Registries and registrars are not technically required to support the uploading of keys or DS records based on all algorithms, and some may not do so. If a DNS provider is responsible for signing a domain name, then their back-end limits the supported signing algorithms. *Recommended measurement technique:* Manual study of the features of DNS providers, registrars, and registries. See [Müller et al. 2020].

ME5: CDS/CDNSKEY support. Registrars and registries can scan for CDS/CDNSKEY records [Gudmundsson and Wouters 2017] and update the DNSSEC settings of their clients or their children domain names automatically. *Recommended measurement technique:* Manual study of the features of registrars or registries. Measurements could be automated to some extent, e.g. by registering a domain name in the relevant TLD or via the relevant registrar, publishing CDS/CDNSKEY records and observing possible changes in the parent zone using active probes. See [Caletka 2022] for a manually maintained list of CDS/CDNSKEY support in the DNS ecosystem and in software.

ME6: DNSSEC fees. Does the registrar or DNS provider charges for enabling DNSSEC. Fees could deter DNSSEC deployment. *Recommended measurement technique:* Manual study of the features of DNS providers or registrars. See [Chung et al. 2017b].

7 RECOMMENDATIONS

In the previous section, we listed 64 direct or indirect metrics of DNSSEC deployment. In this section, we make recommendations about which metrics ICANN could measure in order to gain a comprehensive view of DNSSEC deployment. Measuring DNSSEC deployment is not an end in itself, but rather a means of supporting strategic decision-making in order to achieve a higher goal.

We have therefore formulated three goals related to DNSSEC deployment that ICANN and the DNS community might want to achieve in the medium and long term. For each goal, we list which metrics should be collected in order to make informed decisions and, thus, to achieve the goal. These goals are not exhaustive; other goals could be formulated by ICANN and the community. We acknowledge that collecting the metrics for the first goal might not be easy, but we think that it is still a goal worth pursuing. We believe that the metrics of the second and third goals are feasible.

7.1 Goal 1: Protecting DNS transactions with DNSSEC

The total number of signed domain names or the number of validating resolvers might give a distorted view if widely-used domain names are not protected and popular resolvers are not validating. We therefore propose a combination of metrics focussed on the number of transactions protected with DNSSEC. Unless end-user queries are measured directly, these measurements should take the "popularity" and "importance" of a domain name into account as well as the number of clients relying on a recursive resolver.[2]

Collecting metrics directly. Ideally, we recommend collecting this metric directly on recursive resolvers. One approach would be to extend ICANN's own Identifier Technology Health Indicators (ITHI) programme [Durand and Huitema 2022]. The goal of ITHI is to monitor the health of the DNS ecosystem, through a set of metrics. It is an open-source project[3], already collecting metrics on the root servers and a number of recursive resolvers. ITHI derives metrics from DNS traffic. Current metrics related to DNSSEC include the share of resolvers querying the root with the DNSSEC OK (DO) flag set and whether TLDs are signed.

We recommend extending the ITHI programme with an additional metric: the number of incoming queries that trigger a validated response.

That is a combination of the following two metrics, identified in the literature:

- *MR4 – DNSSEC validation and enforcement*: Resolvers running the ITHI tools report whether they are validating DNSSEC and enforcing validation errors. We cannot reliably deduce whether a recursive resolver is validating from DNS traffic only. Operators running

[2]This collection of metrics was inspired by discussion with the community, namely Alex Mayerhofer (NIC.AT) as well as Sebastian Castro (.IE).

[3]https://github.com/private-octopus/ithitools

the ITHI tool therefore need to indicate whether they are validating when they join the ITHI programme.
- *MD2 – DNSSEC validity*: Resolvers measures whether queried domain names are signed with DNSSEC and whether the signature are valid.

Additionally, resolvers should report rough estimates of their client populations (in unique source IP addresses). For privacy reasons, we recommend not reporting exact numbers, but reporting population in orders of magnitude (e.g. to the power of ten). ITHI tools are open source and, as indicated in Section 5, create reproducible results, and can achieve high levels of coverage if widely deployed.

Collecting metrics indirectly. Getting ITHI tools more widely deployed on resolvers might be challenging. We therefore propose a second set of metrics and measurement techniques to gain similar insights.

First, we recommend to using a measurement platform or an advertisement network to determine whether recursive resolvers have enabled validation (*MR4*) and to get an estimate about their client bases. See, for example, the work by [Huston 2021b,c]. That yields an estimate of the share of the internet population relying on validating resolvers.

Second, we recommend using on "top" lists like the Tranco list [Le Pochat et al. 2019] to identify popular domain names. Active measurements can be used to test whether the domain names are signed (*MD2*).

Assuming a long-tail distribution of queries, where the most popular domain names receive 90% of queries or more [Federrath et al. 2011], we can combine the results of *MR4* and *MD2* to obtain an estimate of how many *transactions* on the Internet are protected with DNSSEC.

We acknowledge that neither the direct approach nor the indirect approach we recommend provides insights into DNSSEC coverage amongst low-profile but potentially relevant domain names. We do, however, believe that our proposed approaches are the only way to develop an objective metric that is not influenced by subjective criteria such as the "importance" of a domain name.

7.2 Goal 2: Increasing DNSSEC deployment quality

DNSSEC in its current form has been around for 17 years. Whereas in the beginning, the goal was to maximise deployment, the goal now could be to improve current DNSSEC deployment so as to maximise conformance to best security and DNS/DNSSEC practices. Quality initiatives are in place at, for example, the registries of the country code Top Level Domains (ccTLDs) for the Netherlands [SIDN 2019] and Sweden, which involve monitoring domain names and incentivising state-of-the-art DNSSEC deployments.

In order to achieve this goal, we recommend measuring the following domain name metrics:

- *MD3 – validation errors* to understand the types of misconfigurations.

- *MD8 – key algorithm* to measure whether recent and secure algorithms are being used. In combination with *MD4* (key-length) to identify inappropriately short RSA-based keys.
- *MD14 – NSEC/NSEC3 usage* to measure whether guidelines on setting NSEC3 parameters are followed [Hardaker and Dukhovni 2022].
- *MD23 – Rollover correctness* to measure if key and algorithm rollovers are carried out, following the guidelines in RFC 6781 [Kolkman et al. 2012].
- *MD15 and MD16 – Path MTU and TCP support* to understand whether authoritative name servers provide reliable transport.
- *MD7 – Keys in RRs* in combination with *MD8*, *MD31* (DNS cookie support) and *MD32* (Response rate limiting) to identify zones that are attractive to attackers for misuse in DDoS reflection attacks.

On the recursive resolver side, we recommend measuring:

- *MR6 – algorithm support*, to identify resolvers that do not support modern signing algorithms or that still support insecure algorithms
- *MR8 – Path MTU* in combination with *MR9* (TCP support) to understand whether recursive resolvers provide reliable transport.
- *MR13 – DNS Cookie support* to understand whether recursive resolvers provide mechanisms to prevent DDoS attacks.

See Section 6 for recommended measurement techniques. We believe that these metrics can give ICANN an overview of the current state of DNSSEC deployment. For background information see [Dukhovni and Hardaker 2022], where researchers publish some of the discussed metrics for domain names and [Toorop 2018], where some of the discussed resolver metrics are published.

Note that the list of metrics needs to be kept under review and extended as the need arises. For example, when new extensions to the DNS are being developed or when best current practices change.

7.3 Goal 3: Making DNSSEC deployments future-proof

An operator cannot simply deploy DNSSEC once and then forget about it. DNSSEC needs maintenance, e.g. to upgrade to more secure and efficient signing algorithms, or to roll ZSKs and KSKs. In the long term, it may be necessary to adopt signing algorithms that require more invasive changes to the DNS protocol. In order to make sure that current DNSSEC deployments can transition to new signing algorithms as easily as possible, a number of prerequisites must be satisfied. The following metrics cover relevant aspects to get a better understanding of the degree to which current deployments are indeed future-proof.

- *MD15 and MD16 – Path MTU and TCP support* to understand whether authoritative name servers provide reliable transport, and are also able to transfer larger messages in the future. That might become necessary, for example, with the adoption of new quantum-safe algorithms.

- *MD21 and MD22 – ZSK and KSK rollover frequency* to see whether domain name operators practice the replacement of their keys, and are therefore ready to move to other algorithms if necessary in the future.
- *MD25 – Algorithm rollover* to see whether domain name operators have rolled their algorithm before and are therefore ready for future algorithm rollovers.

On the recursive resolver side, we recommend measuring:

- *MR6 – algorithm support*, to identify resolvers that do not support modern signing algorithms or that still supporting insecure algorithms. This measurement can also indicate whether recursive resolvers are being maintained.
- *MR8 – Path MTU* in combination with *MR9* (TCP support) to understand whether recursive resolvers provide reliable transport, and are also able to transfer larger messages in the future.

Additionally, we recommend to manually inspecting algorithm support (*ME4*) and CDS/CDNSKEY support (*ME5*) in the DNS ecosystem regularly, as well as algorithm support in relevant DNS software (*MS1* and *MS3*).

By measuring those metrics, ICANN and the DNS community will get a better picture of the current DNSSEC ecosystem's readiness for future transitions to other signing algorithms.

REFERENCES

J. Abley, O. Gudmundsson, M. Majkowski, and E. Hunt. 2019. Providing Minimal-Sized Responses to DNS Queries That Have QTYPE=ANY. RFC 8482 (Proposed Standard). https://doi.org/10.17487/RFC8482

J. Abley, J. Schlyter, G. Bailey, and P. Hoffman. 2016. DNSSEC Trust Anchor Publication for the Root Zone. RFC 7958 (Informational). https://doi.org/10.17487/RFC7958

Timm Böttger, Felix Cuadrado, Gianni Antichi, Eder Leão Fernandes, Gareth Tyson, Ignacio Castro, and Steve Uhlig. 2019. An Empirical Study of the Cost of DNS-over-HTTPS. In *Proceedings of the Internet Measurement Conference (IMC '19)*. Association for Computing Machinery, New York, NY, USA, 15–21. https://doi.org/10.1145/3355369.3355575

S. Bradner. 1997. Key words for use in RFCs to Indicate Requirement Levels. RFC 2119 (Best Current Practice). https://doi.org/10.17487/RFC2119 Updated by RFC 8174.

Bright Data. 2022. Bright Data – The World's #1 Web Data Platform. https://brightdata.com/.

Ondřej Caletka. 2021. Deployment of CDS: Automating DNSSEC maintenance. https://ripe82.ripe.net/wp-content/uploads/presentations/62-Deployment_of_CDS.pdf. In *RIPE 82*.

Ondřej Caletka. 2022. Support for CDS/CDNSKEY/CSYNC updates. https://github.com/oskar456/cds-updates.

Nicolas Canceill. 2014. Measure validation by triggering DNS lookups on RIPE atlas probes. https://ripe68.ripe.net/presentations/232-slides.pdf. In *RIPE 68, DNS-WG*.

Tijay Chung. 2021. Understanding DNSSEC debugging patterns using DNSVIZ. https://indico.dns-oarc.net/event/40/contributions/891/. In *DNS-OARC 36*.

Taejoong Chung, Roland van Rijswijk-Deij, Balakrishnan Chandrasekaran, David Choffnes, Dave Levin, Bruce M. Maggs, Alan Mislove, and Christo Wilson. 2017a. A Longitudinal, {End-to-End} View of the {DNSSEC} Ecosystem. 1307–1322. https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/chung

Taejoong Chung, Roland van Rijswijk-Deij, David Choffnes, Dave Levin, Bruce M. Maggs, Alan Mislove, and Christo Wilson. 2017b. Understanding the role of registrars in DNSSEC deployment. In *Proceedings of the 2017 Internet Measurement Conference (IMC '17)*. Association for Computing Machinery, 369–383. https://doi.org/10.1145/3131365.3131373

Tianxiang Dai, Haya Shulman, and Michael Waidner. 2016. DNSSEC Misconfigurations in Popular Domains. In *Cryptology and Network Security (Lecture Notes in Computer Science)*, Sara Foresti and Giuseppe Persiano (Eds.). Springer International Publishing, 651–660. https://doi.org/10.1007/978-3-319-48965-0_43

J. Damas, M. Graff, and P. Vixie. 2013. Extension Mechanisms for DNS (EDNS(0)). RFC 6891 (Internet Standard). https://doi.org/10.17487/RFC6891

Jacob Davis and Casey Deccio. 2021. A Peek into the DNS Cookie Jar. In *Passive and Active Measurement*, Oliver Hohlfeld, Andra Lutu, and Dave Levin (Eds.). Springer International Publishing, Cham, 302–316.

Casey Deccio. 2011. Maintenance, Mishap, and Mending in DNSSEC Deployment.

Casey Deccio, Derek Argueta, and Jonathan Demke. 2019. A Quantitative Study of the Deployment of DNS Rate Limiting. In *2019 International Conference on Computing, Networking and Communications (ICNC)*. 442–447. https://doi.org/10.1109/ICCNC.2019.8685601

Casey Deccio, Jeff Sedayao, Krishna Kant, and Prasant Mohapatra. 2011. Quantifying and Improving DNSSEC Availability. In *2011 Proceedings of 20th International Conference on Computer Communications and Networks (ICCCN)*. 1–7. https://doi.org/10.1109/ICCCN.2011.6005908 ISSN: 1095-2055.

Sander Degen. 2011. DNSSEC Client Behaviour. https://ripe62.ripe.net/presentations/145-DNS_Client_analysis_presentation_RIPE.pdf. In *RIPE 62*.

DNS OARC. 2020. Net::DNS::Fingerprint. https://github.com/kirei/fpdns.

Trinh Viet Doan, Irina Tsareva, and Vaibhav Bajpai. 2021. Measuring DNS over TLS from the Edge: Adoption, Reliability, and Response Times. In *Passive and Active Measurement*, Oliver Hohlfeld, Andra Lutu, and Dave

Levin (Eds.). Springer International Publishing, Cham, 192–209.
Viktor Dukhovni and Wes Hardaker. 2022. DDNSSEC and DANE Deployment Statistics. https://stats.dnssec-tools.org/.
Alain Durand and Christian Huitema. 2022. ICANN Identifier Technology Health Indicators (ITHI). https://ithi.research.icann.org/.
D. Eastlake 3rd and M. Andrews. 2016. Domain Name System (DNS) Cookies. RFC 7873 (Proposed Standard). https://doi.org/10.17487/RFC7873 Updated by RFC 9018.
Hannes Federrath, Karl-Peter Fuchs, Dominik Herrmann, and Christopher Piosecny. 2011. Privacy-preserving DNS: analysis of broadcast, range queries and mix-based protection methods. In *European Symposium on Research in Computer Security*. Springer, 665–683.
Pawel Foremski, Oliver Gasser, and Giovane C. M. Moura. 2019. DNS Observatory: The Big Picture of the DNS. In *Proceedings of the Internet Measurement Conference (IMC '19)*. Association for Computing Machinery, New York, NY, USA, 87–100. https://doi.org/10.1145/3355369.3355566
Kensuke Fukuda, Shinta Sato, and Takeshi Mitamura. 2013. A technique for counting DNSSEC validators. In *2013 Proceedings IEEE INFOCOM*. 80–84. https://doi.org/10.1109/INFCOM.2013.6566739 ISSN: 0743-166X.
Google. 2022. Google Ads – Get More Customers & Generate Leads with Online Ads. https://ads.google.com.
O. Gudmundsson and P. Wouters. 2017. Managing DS Records from the Parent via CDS/CDNSKEY. RFC 8078 (Proposed Standard). https://doi.org/10.17487/RFC8078
Ólafur Gudmundsson and Stephen D Crocker. 2011. Observing DNSSEC validation in the wild. https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.362.5934&rep=rep1&type=pdf
Wes Hardaker and Viktor Dukhovni. 2022. *Guidance for NSEC3 parameter settings*. Internet-Draft draft-ietf-dnsop-nsec3-guidance-10. Internet Engineering Task Force. https://datatracker.ietf.org/doc/html/draft-ietf-dnsop-nsec3-guidance-10 Work in Progress.
P. Hoffman and P. McManus. 2018. DNS Queries over HTTPS (DoH). RFC 8484 (Proposed Standard). https://doi.org/10.17487/RFC8484
Austin Hounsel, Paul Schmitt, Kevin Borgolte, and Nick Feamster. 2021. Can Encrypted DNS Be Fast?. In *Passive and Active Measurement*, Oliver Hohlfeld, Andra Lutu, and Dave Levin (Eds.). Springer International Publishing, Cham, 444–459.
Z. Hu, L. Zhu, J. Heidemann, A. Mankin, D. Wessels, and P. Hoffman. 2016. Specification for DNS over Transport Layer Security (TLS). RFC 7858 (Proposed Standard). https://doi.org/10.17487/RFC7858 Updated by RFC 8310.
Christian Huitema, Sara Dickinson, and Allison Mankin. 2022. DNS over Dedicated QUIC Connections. RFC 9250. https://doi.org/10.17487/RFC9250
Geoff Huston. 2012. DNSSEC Measurement - A slightly closer look. https://ripe65.ripe.net/presentations/177-2012-09-26-dnssec.pdf. In *RIPE 65*.
Geoff Huston. 2021a. Measurement of DNSSEC Validation with Edwards Curve Cryptography. https://indico.dns-oarc.net/event/39/contributions/868/. In *DNS-OARC 35a*.
Geoff Huston. 2021b. Measurement of DNSSEC Validation with RSA-4096. https://indico.dns-oarc.net/event/40/contributions/888/. In *DNS-OARC 36*.
Geoff Huston. 2021c. Measuring DNS Flag Day 2020. https://indico.dns-oarc.net/event/37/contributions/806/. In *DNS-OARC 34*.
G. Huston, J. Damas, and W. Kumari. 2018. A Root Key Trust Anchor Sentinel for DNSSEC. RFC 8509 (Proposed Standard). https://doi.org/10.17487/RFC8509
Neda Kianpour and Taylor Shaw. 2019. Challenges and Successes of DNSSEC Signing an F5 BigIP DNS Hosted Zone. https://ripe79.ripe.net/wp-content/uploads/presentations/34-Challenges-and-Successes-of-DNSSEC-Signing-an-F5-BigIP-DNS-Hosted-Zone.pdf. In *RIPE 79, DNS-WG*.

O. Kolkman, W. Mekking, and R. Gieben. 2012. DNSSEC Operational Practices, Version 2. RFC 6781 (Informational). https://doi.org/10.17487/RFC6781

Axel Koolhaas and Tjerd Slokker. 2020. Defragmenting DNS - Determining the optimal maximum UDP response size for DNS. https://indico.dns-oarc.net/event/36/contributions/776/. In *DNS-OARC 32b*.

Mike Kosek, Trinh Viet Doan, Malte Granderath, and Vaibhav Bajpai. 2022. One to Rule Them All? A First Look at DNS over QUIC. In *Passive and Active Measurement*, Oliver Hohlfeld, Giovane Moura, and Cristel Pelsser (Eds.). Springer International Publishing, Cham, 537–551.

Christian Kreibich, Nicholas Weaver, Boris Nechaev, and Vern Paxson. 2010. Netalyzr: illuminating the edge network. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement (IMC '10)*. Association for Computing Machinery, 246–259. https://doi.org/10.1145/1879141.1879173

B. Laurie, G. Sisson, R. Arends, and D. Blacka. 2008. DNS Security (DNSSEC) Hashed Authenticated Denial of Existence. RFC 5155 (Proposed Standard). https://doi.org/10.17487/RFC5155 Updated by RFCs 6840, 6944, 9077, 9157.

Tho Le, Roland van Rijswijk-Deij, Luca Allodi, and Nicola Zannone. 2018. Economic incentives on DNSSEC deployment: Time to move from quantity to quality. In *NOMS 2018 - 2018 IEEE/IFIP Network Operations and Management Symposium*. 1–9. https://doi.org/10.1109/NOMS.2018.8406223 ISSN: 2374-9709.

Victor Le Pochat, Tom Van Goethem, Samaneh Tajalizadehkhoob, Maciej Korczyński, and Wouter Joosen. 2019. Tranco: A Research-Oriented Top Sites Ranking Hardened Against Manipulation. In *Proceedings of the 26th Annual Network and Distributed System Security Symposium (NDSS 2019)*. https://doi.org/10.14722/ndss.2019.23386

Edward Lewis. 2012. Looking at TLD DNSSEC Practices. https://ripe64.ripe.net/presentations/46-RIPE64PlenaryTLDDNSSEC.pdf. In *RIPE 64*.

Wilson Lian, Eric Rescorla, Hovav Shacham, and Stefan Savage. 2013. Measuring the Practical Impact of {DNSSEC} Deployment. 573–588. https://www.usenix.org/conference/usenixsecurity13/technical-sessions/paper/lian

Chaoyi Lu, Baojun Liu, Zhou Li, Shuang Hao, Haixin Duan, Mingming Zhang, Chunying Leng, Ying Liu, Zaifeng Zhang, and Jianping Wu. 2019. An End-to-End, Large-Scale Measurement of DNS-over-Encryption: How Far Have We Come?. In *Proceedings of the Internet Measurement Conference (IMC '19)*. Association for Computing Machinery, New York, NY, USA, 22–35. https://doi.org/10.1145/3355369.3355580

Bill Manning. 2006. CADR: DNSSEC Authenticated DNS Registry. https://meetings.ripe.net/ripe-53/presentations/caadr_update.pdf. In *RIPE 53*.

Jiarun Mao, Michael Rabinovich, and Kyle Schomp. 2022. Assessing Support for DNS-over-TCP in the Wild. In *Passive and Active Measurement*, Oliver Hohlfeld, Giovane Moura, and Cristel Pelsser (Eds.). Springer International Publishing, Cham, 487–517.

Masato Minda. 2011. Changes to JP DNS traffic by DNSSEC. https://ripe62.ripe.net/presentations/150-dnssec-in-jp-10.pdf. In *RIPE 62, DNS-WG*.

Robert Mortimer. 2021. The state of DNS security records. https://indico.dns-oarc.net/event/37/contributions/817/. In *DNS-OARC 34*.

Giovane Moura. 2021. Fragmentation, truncation, and timeouts: are large DNS messages falling to bits?. https://indico.dns-oarc.net/event/38/contributions/851/. In *DNS-OARC 35*.

Moritz Müller. 2022a. DNSSEC Deployment Metrics Research. In *DNSSEC and Security Workshop ICANN73*.

Moritz Müller. 2022b. DNSSEC Deployment Metrics Research. In *MAPRG, IRTF-113*.

Moritz Müller. 2022c. DNSSEC Deployment Metrics Research. In *CENTR 20th RnD Workshop*.

Moritz Müller, Jins de Jong, Maran van Heesch, Benno Overeinder, and Roland van Rijswijk-Deij. 2020. Retrofitting Post-Quantum Cryptography in Internet Protocols: A Case Study of DNSSEC. *ACM SIGCOMM Computer Communication Review* 50, 4 (2020).

Moritz Müller. 2016. Measuring DNSSEC Configuration of Upstream Resolvers with RIPE Atlas. https://ripe72.ripe.net/wp-content/uploads/presentations/103-RIPE72_Measure_DNSSEC_with_Atlas.pdf. In *RIPE 72*.

Moritz Müller. 2021. RPKI and the DNS: role of big players is crucial. https://www.sidnlabs.nl/en/news-and-blogs/rpki-and-the-dns-role-of-big-players-is-crucial.

Moritz Müller, Taejoong Chung, Alan Mislove, and Roland van Rijswijk-Deij. 2019a. Rolling With Confidence: Managing the Complexity of DNSSEC Operations. *IEEE Transactions on Network and Service Management* 16, 3 (2019), 1199–1211. https://doi.org/10.1109/TNSM.2019.2916176

Moritz Müller, Matthew Thomas, Duane Wessels, Wes Hardaker, Taejoong Chung, Willem Toorop, and Roland van Rijswijk-Deij. 2019b. Roll, Roll, Roll your Root: A Comprehensive Analysis of the First Ever DNSSEC Root KSK Rollover. In *Proceedings of the Internet Measurement Conference (IMC '19)*. Association for Computing Machinery, 1–14. https://doi.org/10.1145/3355369.3355570

Moritz Müller, Willem Toorop, Taejoong Chung, Jelte Jansen, and Roland van Rijswijk-Deij. 2020. The Reality of Algorithm Agility: Studying the DNSSEC Algorithm Life-Cycle. In *Proceedings of the ACM Internet Measurement Conference (IMC '20)*. Association for Computing Machinery, 295–308. https://doi.org/10.1145/3419394.3423638

Wolfgang Nagele. 2011. DNSSEC and the RIPE NCC. https://ripe62.ripe.net/presentations/27-RIPE62_WolfgangNagele_DNSSEC_and_the_RIPE_NCC.pdf. In *RIPE 62*.

NLnet Labs. 2022. Routinator 3000. https://nlnetlabs.nl/projects/rpki/routinator/.

Eric Osterweil, Michael Ryan, Dan Massey, and Lixia Zhang. 2008. Quantifying the operational status of the DNSSEC deployment. In *Proceedings of the 8th ACM SIGCOMM conference on Internet measurement conference - IMC '08*. ACM Press, 231. https://doi.org/10.1145/1452520.1452548

Eric Osterweil, Pouyan Fotouhi Tehrani, Thomas C. Schmidt, and Matthias Wählisch. 2021. From the Beginning: Key Transitions in the First 15 Years of DNSSEC. *CoRR* abs/2109.08783 (2021). arXiv:2109.08783 https://arxiv.org/abs/2109.08783

T. Pauly and P. Wouters. 2019. Split DNS Configuration for the Internet Key Exchange Protocol Version 2 (IKEv2). RFC 8598 (Proposed Standard). https://doi.org/10.17487/RFC8598

Vern Paxson. 2004. Strategies for Sound Internet Measurement. In *Proceedings of the 4th ACM SIGCOMM Conference on Internet Measurement*. 263–271.

James Richard. 2021. An analysis of DNSSEC signed domains appearing within a large web crawl dataset. https://indico.dns-oarc.net/event/40/contributions/886/. In *DNS-OARC 36*.

RIPE NCC Staff. 2015. RIPE Atlas: A Global Internet Measurement Network. *Internet Protocol Journal (IPJ)* 18, 3 (Sep 2015).

Victoria Risk. 2015. Impact of unknown EDNS options on the DNS. https://indico.dns-oarc.net/event/24/contributions/379/attachments/334/591/Winstead-EDNS.pdf. In *DNS-OARC Fall Workshop*.

Craig A. Shue and Andrew J. Kalafut. 2013. Resolvers Revealed: Characterizing DNS Resolvers and their Clients. 12, 4 (2013), 14:1–14:17. https://doi.org/10.1145/2499926.2499928

SIDN. 2019. Registrar Scorecard yields great results. https://www.sidn.nl/en/news-and-blogs/registrar-scorecard-yields-great-results.

Geoffrey Sisson. 2010. DNS SURVEY: OCTOBER 2010. http://dns.measurement-factory.com/surveys/201010/.

M. StJohns. 2007. Automated Updates of DNS Security (DNSSEC) Trust Anchors. RFC 5011 (Internet Standard). https://doi.org/10.17487/RFC5011

O. Sury, W. Toorop, D. Eastlake 3rd, and M. Andrews. 2021. Interoperable Domain Name System (DNS) Server Cookies. RFC 9018 (Proposed Standard). https://doi.org/10.17487/RFC9018

Willem Toorop. 2018. DNSThought - Everything you ever wanted to know about caching resolvers but were afraid to ask. https://indico.dns-oarc.net/event/29/contributions/654/. In *DNS-OARC 29*.

Willem Toorop. 2019. Seeing the effects of DNS Flag Day in action. https://indico.dns-oarc.net/event/31/contributions/681/. In *DNS-OARC 30*.

Willem Toorop. 2020. The Current State of DNS Resolvers and RPKI Protection. https://indico.dns-oarc.net/event/36/contributions/775/. In *DNS-OARC 32b*.

Niels LM van Adrichem, Norbert Blenn, Antonio Reyes Lúa, Xin Wang, Muhammad Wasif, Ficky Fatturrahman, and Fernando A Kuipers. 2015. A measurement study of DNSSEC misconfigurations. *Security Informatics* 4, 1 (2015), 1–14.

Gijs Van Den Broek, Roland Van Rijswijk-Deij, Anna Sperotto, and Aiko Pras. 2014. DNSSEC meets real world: dealing with unreachability caused by fragmentation. 52, 4 (2014), 154–160. https://doi.org/10.1109/MCOM.2014.6828880

Olivier van der Toorn, Johannes Krupp, Mattijs Jonker, Roland van Rijswijk-Deij, Christian Rossow, and Anna Sperotto. 2021. ANYway: Measuring the Amplification DDoS Potential of Domains. In *2021 17th International Conference on Network and Service Management (CNSM)*. 500–508. https://doi.org/10.23919/CNSM52442.2021.9615596

Olivier van der Toorn, Moritz Müller, Sara Dickinson, Cristian Hesselman, Anna Sperotto, and Roland van Rijswijk-Deij. 2022. Addressing the challenges of modern DNS: A comprehensive tutorial. *Computer Science Review* 45 (2022), 100469. https://doi.org/10.1016/j.cosrev.2022.100469

Roland van Rijswijk. 2019. Tag You're It: Revisiting the Reality of DNSSEC Keytags. https://ripe78.ripe.net/wp-content/uploads/presentations/5-20190520-RIPE-78-DNS-wg-Keytags.pdf. In *RIPE 78, DNS-WG*.

Roland van Rijswijk-Deij, Mattijs Jonker, and Anna Sperotto. 2016a. On the adoption of the elliptic curve digital signature algorithm (ECDSA) in DNSSEC. In *2016 12th International Conference on Network and Service Management (CNSM)*. IEEE, 258–262.

Roland van Rijswijk-Deij, Mattijs Jonker, Anna Sperotto, and Aiko Pras. 2016b. A high-performance, scalable infrastructure for large-scale active DNS measurements. *IEEE journal on selected areas in communications* 34, 6 (2016), 1877–1888.

Roland van Rijswijk-Deij, Anna Sperotto, and Aiko Pras. 2014. DNSSEC and its potential for DDoS attacks: a comprehensive measurement study. In *Proceedings of the 2014 Conference on Internet Measurement Conference (IMC '14)*. Association for Computing Machinery, 449–460. https://doi.org/10.1145/2663716.2663731

Roland van Rijswijk-Deij, Anna Sperotto, and Aiko Pras. 2015. Making the Case for Elliptic Curves in DNSSEC. *SIGCOMM Comput. Commun. Rev.* 45, 5 (sep 2015), 13–19. https://doi.org/10.1145/2831347.2831350

Patrik Wallström. 2012. Quality of DNS and DNSSEC in the .se Zone. https://ripe64.ripe.net/presentations/87-Patrik_Wallstrom_RIPE64_2012-04-18_DNSSEC_in_.SE.pdf. In *RIPE 64*.

Matthäus Wander. 2017. Measurement survey of server-side DNSSEC adoption. In *2017 Network Traffic Measurement and Analysis Conference (TMA)*. 1–9. https://doi.org/10.23919/TMA.2017.8002913

Matthäus Wander and Torben Weis. 2013. Measuring Occurrence of DNSSEC Validation. In *Passive and Active Measurement (Lecture Notes in Computer Science)*, Matthew Roughan and Rocky Chang (Eds.). Springer, 125–134. https://doi.org/10.1007/978-3-642-36516-4_13

Nicolas Weaver, Christian Kreibich, Boris Nechaev, and Vern Paxson. 2011. Implications of Netalyzr's DNS measurements. In *Proceedings of the First Workshop on Securing and Trusting Internet Names (SATIN)*.

D. Wessels, W. Kumari, and P. Hoffman. 2017. Signaling Trust Anchor Knowledge in DNS Security Extensions (DNSSEC). RFC 8145 (Proposed Standard). https://doi.org/10.17487/RFC8145 Updated by RFC 8553.

P. Wouters and O. Sury. 2019. Algorithm Implementation Requirements and Usage Guidance for DNSSEC. RFC 8624 (Proposed Standard). https://doi.org/10.17487/RFC8624 Updated by RFC 9157.

Maarten Wullink, Giovane CM Moura, Moritz Müller, and Cristian Hesselman. 2016. ENTRADA: A high-performance network traffic data streaming warehouse. In *NOMS 2016-2016 IEEE/IFIP Network Operations and Management Symposium*. IEEE, 913–918.

Hao Yang, Eric Osterweil, Dan Massey, Songwu Lu, and Lixia Zhang. 2011. Deploying Cryptography in Internet-Scale Systems: A Case Study on DNSSEC. 8, 5 (2011), 656–669. https://doi.org/10.1109/TDSC.2010.10

Yingdi Yu, Duane Wessels, Matt Larson, and Lixia Zhang. 2013. Check-Repeat: A new method of measuring DNSSEC validating resolvers. In *2013 Proceedings IEEE INFOCOM*. 3147–3152. https://doi.org/10.1109/INFCOM.2013.6567129 ISSN: 0743-166X.

A OVERVIEW OF METRICS AND RELATED STUDIES

Table 7. DNSSEC metrics and related studies

	Metric	References
Domain Name and Name Server Metrics		
MD1	DNSSEC publication	[Chung 2021; Chung et al. 2017a,b; Dai et al. 2016; Deccio et al. 2011; Le et al. 2018; Lewis 2012; Manning 2006; Mortimer 2021; Müller et al. 2019a; Nagele 2011; Richard 2021; Sisson 2010; van Rijswijk-Deij et al. 2016a, 2014; Wallström 2012; Yang et al. 2011]
MD2	DNSSEC validity	[Chung 2021; Chung et al. 2017a; Dai et al. 2016; Deccio 2011; Foremski et al. 2019; Lewis 2012; Mortimer 2021; Nagele 2011; Osterweil et al. 2008; Sisson 2010; van Adrichem et al. 2015; Wallström 2012]
MD3	DNSSEC validation errors	[Chung 2021; Deccio 2011; Gudmundsson and Crocker 2011; Nagele 2011; Sisson 2010; van Adrichem et al. 2015; Wallström 2012; Wander 2017]
MD4	Key attribute: Key length	[Chung et al. 2017a; Dai et al. 2016; Le et al. 2018; Lewis 2012; van Rijswijk 2019; Wallström 2012; Wander 2017]
MD5	Key attribute: TTL	[Chung et al. 2017a; Lewis 2012; Wallström 2012]
MD6	Key attribute: shared keys	[Chung et al. 2017a]
MD7	Key attribute: Keys in the RR set	[Lewis 2012; Manning 2006; van Rijswijk-Deij et al. 2014; Wallström 2012]
MD8	Key attribute: algorithm	[Chung 2021; Dai et al. 2016; Le et al. 2018; Lewis 2012; Manning 2006; Müller et al. 2020; Sisson 2010; van Rijswijk 2019; van Rijswijk-Deij et al. 2016a; Wallström 2012; Wander 2017]
MD9	Key attributes: DS digest type	[Dai et al. 2016; Müller et al. 2020]
MD10	Key attribute: RSA modulus	[Dai et al. 2016; van Rijswijk 2019]
MD11	Key attribute: ZSK/KSK or CSK	[Le et al. 2018; Lewis 2012; Wallström 2012; Wander 2017]
MD12	Key attribute: Key-tag exists multiple times	[van Rijswijk 2019]
MD13	Key attribute: Response size	[Kianpour and Shaw 2019; Müller et al. 2019b; van Rijswijk-Deij et al. 2015]
MD14	NSEC/NSEC3 usage	[Sisson 2010; Wander 2017]
MD15	Transport: Path MTU	[Osterweil et al. 2008]
MD16	Transport: TCP support	[Mao et al. 2022; Osterweil et al. 2008]
MD17	EDNS(0) Support	[Risk 2015; Sisson 2010; Weaver et al. 2011]
MD18	Transport: DoT	[Doan et al. 2021; Hounsel et al. 2021]
MD19	Transport: DoQ	–
MD20	Operations: Signature lifetime	[Chung 2021; Sisson 2010; van Adrichem et al. 2015; van Rijswijk 2019]
MD21	Operations: Rollover frequency ZSK	[Chung et al. 2017a; Deccio 2011; Le et al. 2018; Yang et al. 2011]

Table 7. DNSSEC metrics and related studies

	Metric	References
MD22	Operations: Rollover frequency KSK	[Chung et al. 2017a; Deccio 2011; Le et al. 2018; Yang et al. 2011]
MD23	Operations: Rollover correctness	[Chung et al. 2017a; Müller et al. 2019a, 2020]
MD24	Operations: Rollover type	[Chung et al. 2017a; Mao et al. 2022; Müller et al. 2020; van Adrichem et al. 2015]
MD25	Operations: Algorithm rollover	[Müller et al. 2020]
MD26	Operations: Name Server operator	[Le et al. 2018; Müller et al. 2020; van Rijswijk-Deij et al. 2016a]
MD27	Operations: Operational complexity	[Gudmundsson and Crocker 2011]
MD28	Operations: Effects on query load	[Minda 2011]
MD29	Operations: Effects on TCP load	[Minda 2011]
MD30	Operations: CDS/CDNSKEY publication	[Caletka 2021]
MD31	Name Server: DNS Cookie support	[Davis and Deccio 2021]
MD32	Name Server: Response rate limiting	[Deccio et al. 2019]
MD33	Name Server: Minimal responses to query type `ANY`	[van der Toorn et al. 2021]
MD34	RPKI: Served from RPKI signed resources	[Müller 2021]
Resolver Metrics		
MR1	EDNS(0) Support	[Kreibich et al. 2010; Minda 2011; Van Den Broek et al. 2014]
MR2	DNSSEC capable	[Canceill 2014; Chung et al. 2017a; Fukuda et al. 2013; Gudmundsson and Crocker 2011; Huston 2012, 2021b; Kreibich et al. 2010; Lian et al. 2013; Müller 2016; Müller et al. 2019b; Shue and Kalafut 2013; Weaver et al. 2011]
MR3	DNSSEC validation	[Canceill 2014; Chung et al. 2017a; Gudmundsson and Crocker 2011; Huston 2012; Lian et al. 2013; Müller 2016; Müller et al. 2019b; Yu et al. 2013]
MR4	DNSSEC validation and enforcement	[Canceill 2014; Huston 2012; Müller 2016; Müller et al. 2019a; Wander and Weis 2013]
MR5	Treatment of specific DNSSEC validation failures	[Canceill 2014; Lian et al. 2013; Müller 2016]
MR6	Algorithm support	[Huston 2021a; Lian et al. 2013; Müller et al. 2020]
MR7	DNSSEC trust anchor support	[Müller et al. 2019b]
MR8	Transport: Path MTU	[Van Den Broek et al. 2014; Weaver et al. 2011]
MR9	Transport: TCP support	[Huston 2021c; Lian et al. 2013; Moura 2021; Van Den Broek et al. 2014; Weaver et al. 2011]
MR10	DNS protocol: DoH support	[Böttger et al. 2019; Hounsel et al. 2021; Lu et al. 2019]
MR11	DNS protocol: DoT support	[Doan et al. 2021; Hounsel et al. 2021; Lu et al. 2019]
MR12	DNS protocol: DoQ support	[Kosek et al. 2022]
MR13	DNS Cookie support	[Davis and Deccio 2021]
MR14	Telemetry: RFC 8145	[Müller et al. 2019b]

Table 7. DNSSEC metrics and related studies

	Metric	References
MR15	Telemetry: RFC 8509	[Müller et al. 2019b]
MR16	RPKI: Served from RPKI signed resources	–
MR17	RPKI: Route origin validation	[Toorop 2020]
Client Metrics		
MC1	DNSSEC validation	[Osterweil et al. 2008]
MC2	Transport: DoH	–
MC3	Upstream resolver type	[Degen 2011; Müller et al. 2020]
DNS Software Metrics		
MS1	RFC compliance	–
MS2	Default settings	–
MS3	Cryptographic library	[Müller et al. 2020]
MS4	Software deployment	[Sisson 2010; **?**]
DNS Ecosystem Metrics		
ME1	DNSSEC support	[Chung et al. 2017b]
ME2	DNSSEC default	[Chung et al. 2017b]
ME3	DNSSEC settings	[Chung et al. 2017b]
ME4	Algorithm support	[Müller et al. 2020]
ME5	CDS/CDNSKEY support	[Caletka 2022]
ME6	DNSSEC fees	[Chung et al. 2017b]

B ACRONYMS

ccTLD country code Top Level Domain.
CENTR Council of European National Top-Level Domain Registries.
CSK Combined Signing Key.

DDoS Distributed Denial of Service.
DNS Domain Name System.
DNS-OARC DNS Operations, Analysis, and Research Center.
DNSSEC DNS Security Extensions.
DoH DNS-over-HTTPS.
DoQ DNS-over-QUIC.
DoT DNS-over-TLS.

EDNS0 Extension Mechanisms for DNS.

HSM Hardware Security Module.

ITHI Identifier Technology Health Indicators.

KSK Key Signing Key.

NS Name Server.

OCTO ICANN Office of the CTO.

RIR Reginal Internet Registry.
ROA Route Origin Authorization.
RTT Round-Trip Time.

TLD Top Level Domain.
TTL Time To Live.

VP vantage point.

ZSK Zone Signing Key.

Notes:
Cover photo by William Warby on Unsplash