

## Chapter 12

# MEASURING NAME SYSTEM HEALTH

Emiliano Casalicchio, Marco Caselli, Alessio Coletta, Salvatore Di Blasi  
and Igor Nai Fovino

**Abstract** Modern critical infrastructure assets are exposed to security threats arising from their use of IP networks and the Domain Name System (DNS). This paper focuses on the health of DNS. Indeed, due to the increased reliance on the Internet, the degradation of DNS could have significant consequences for the critical infrastructure. This paper describes the Measuring Naming System (MeNSa), a framework designed to provide a formal methodology, metrics and tools for evaluating DNS health. Additionally, it proposes a process for aggregating health and security metrics to provide potential threat indicators. Results from a scenario-based experiment demonstrate the utility of the framework and aggregation metrics.

**Keywords:** Domain Name System, security, aggregation metrics

### 1. Introduction

Critical infrastructure assets such as electric power grids, gas pipelines, and telecommunications and banking systems are increasingly reliant on information and communications technologies. Information and communication technologies provide opportunities to enhance and optimize services and efficiently manage remote installations. As consequence, however, the information and communications infrastructures that enable these services have become vital to the proper operation of critical infrastructure assets.

This paper focuses on the Domain Name System (DNS) infrastructure. DNS is a hierarchical naming system that “maps” Internet domain names to corresponding IP addresses. Often viewed as a phone book, the operation of DNS is essential to the proper functioning of the Internet. Without DNS, it would be practically impossible for users to navigate the Internet or use web service applications. Due to the growing interconnectivity of critical infrastructure assets, a DNS fault under certain conditions could have serious national and international implications [14].

DNS security concerns and the potential impact were discussed during the Internet Corporation for Assigned Names and Numbers (ICANN) symposia in 2009 and 2010 [17, 18]. From these two symposia emerged the concept of DNS health as a means for expressing the status of DNS.

This paper presents results from the Measuring Naming System (MeNSa) effort [12]. The primary goal of the project is to design a formal methodology, metrics and tools for evaluating DNS health. The paper presents the architecture of the framework [6], sample metrics [7] and the operation schema [5]. Additionally, it describes a process for aggregating health and security metrics via structured indices.

## 2. Domain Name System

This section provides an overview of DNS. Additionally, it discusses vulnerabilities and the associated impact on information and communications technology infrastructures.

### 2.1 DNS Overview

The DNS infrastructure is composed of geographical and logical entities that are organized in a hierarchical fashion. The topmost level of the hierarchy is the root domain, while the next subordinate level consists of top-level domains (TLDs). Each TLD, in turn, can have many sub-domains, called second-level or enterprise-level domains. Entities associated with the root domain are called root operators. Registries are the organizations that manage name servers related to a TLD. To facilitate the administration process, DNS defines the concept of a zone – a portion of the domain name space for which administrative responsibility is delegated.

A DNS query to resolve an Internet domain name originates from a client component to either an authoritative name server or a caching name server. Note that this process can be iterative or recursive. A response is generated that provides the IP address corresponding to the Internet domain name. A zone transfer represents an operation where a secondary name server refreshes its records with the primary name servers. This process enables a secondary name server to maintain synchronization with the primary name server. DNS dynamic services provide the ability to dynamically add and/or delete a subset of the resource records for an existing domain, to delete an entire domain, or to create a new domain. DNS administrative services also include tasks performed by the responsible entity to provide an appropriate level of service and to ensure security.

### 2.2 DNS Threats

DNS was designed in the 1980s with little concern for security. Because DNS functionality has, for the most part, remained unchanged, several intrinsic

vulnerabilities exist. DNS threats can be broadly classified into three main categories [23]: (i) data corruption; (ii) denial of service (DoS); and (iii) privacy.

Data corruption is defined as the unauthorized modification of DNS data and includes repository corruption and system corruption. Repository corruption is the debasement of databases containing authoritative data necessary for DNS operations (e.g., resource records and zone files). System corruption is the alteration of the authenticity of DNS responses. Note that weaknesses in the design of the DNS protocol are often exploited in data corruption attacks. Examples include cache poisoning, route injection and man-in-the-middle attacks. The well-known Kaminsky attack [19] is a concrete example of this class of attacks.

A DoS attack renders a service unavailable to legitimate users. These attacks usually impact a specific service (e.g., targeting assets that rely on the proper functioning of the DNS) or create wide-ranging outages (e.g., degrading general Internet functionality).

A privacy threat relates to the loss or theft of personal information. One example is reading a DNS cache to discern an individual's browsing activities. The consideration of privacy threats is beyond the scope of this paper. However, we intend to consider privacy issues in our future work related to DNS health.

### 2.3 DNS Incidents

The first security flaws in the DNS protocol were identified in the early 1990s when Bellovin [4] and Vixie [25] discovered how to spoof name-based authentication systems using cache contamination attacks. The security extension DNSSEC was proposed in 1997 to address the identified vulnerabilities [10, 11]. Further cache poisoning vulnerabilities discovered by Kaminsky led to the development of additional specifications, namely RFC 4033 [1], RFC 4034 [2] and RFC 4035 [3].

Two major attacks have been reported on DNS root servers. The first attack, which occurred in 2002 and lasted approximately one hour, simultaneously targeted all thirteen DNS root servers [26]. The performance and availability of nine servers were degraded during the attack; in response, the Anycast protocol was implemented in eleven root servers. The second global attack occurred in 2007 [15]. This attack was larger in scale, however, only the two root servers that had not adopted the Anycast solution were impacted.

Root servers are not the only DNS components that are vulnerable. Several DNS hijacking attacks that targeted domain name registrars have been reported. In June 2008, for example, the ICANN website was the victim of a defacement attack resulting from the compromise of its name registrar [16]. Another attack compromised a large e-bill payment site that redirected visitors to an alternate website and installed malicious code on their machines [20]. In 2009, the New Zealand version of Microsoft's MSN website was compromised after attackers penetrated the country's primary domain name registrar [9]. Similarly, in 2009, a domain name registrar in Puerto Rico was compromised, resulting in the redirection of local websites for major companies such as

Google, Microsoft and Yahoo [22]. Also in 2009, malicious entities used cache poisoning to redirect the login page for a major Brazilian bank to a fraudulent website that stole user credentials [13].

A recent study by the Global Cyber Security Center detailed how a DNS attack could impact operations in a smart grid [14]. Indeed, the increasing dependency of critical infrastructure assets on information and communications technologies warrants security solutions that ensure that DNS is adequately protected.

## 2.4 DNS Health and Security

The security, stability and resiliency of DNS have received significant attention over the past few years. Following the 2009 and 2010 DNS symposia [17, 18], ICANN specified the following indicators for DNS health:

- **Availability:** The ability of DNS to be operational and accessible when required.
- **Coherency:** The ability of DNS to accurately resolve name queries; this is one of the core principles of DNS. For example, if the IP address 192.0.2.1 is resolved to `www.foo.example.com`, then the coherency principle implies that the name `www.foo.example.com` should resolve to the IP address 192.0.2.1.
- **Integrity:** The ability of DNS to guard against improper data modification or destruction; this includes ensuring information non-repudiation and authenticity.
- **Resiliency:** The ability of DNS to effectively respond and recover to a known, desired and safe state in the event of a disturbance.
- **Security:** The ability of DNS to limit or protect itself from malicious activities (e.g., unauthorized system access, fraudulent representation of identity and interception of communications).
- **Speed:** The performance of DNS with respect to response time and throughput. Note that, in addition to queries, speed applies to maintenance, administration and management operations.
- **Stability:** The ability of DNS to function in a reliable and predictable manner (e.g., protocols and standards). Stability is important because it facilitates universal acceptance and usage.
- **Vulnerability:** The likelihood that a DNS weakness can be exploited by one or more threats.

Several studies have examined DNS traffic measurement techniques and performance metrics [8, 21, 24]. However, hardly any research has examined DNS

health in relation to the prescribed indicators. This paper focuses on the security, resiliency and vulnerability indicators for deriving DNS health metrics associated with our MeNSa framework.

### 3. MeNSa Framework

The 2009 and 2010 ICANN DNS symposia that introduced the concept of DNS health, also identified the following requirements:

- The need for viable indicators of DNS health for different DNS actors (i.e., root server operators, non-root authoritative name server operators, recursive caches, open DNS resolvers and end users).
- The need to understand and refine proper methods and techniques for measuring DNS health indicators.
- The need to refine and improve existing metrics for availability, coherency, integrity, resiliency, security, speed, stability and vulnerability.
- The need for metric threshold levels that identify when DNS health has degraded below acceptable standards.

Despite the specification of these requirements, the realization of DNS health metrics is still at a primitive stage. This section describes the main components of the MeNSa framework for deriving DNS health metrics. Interested readers are referred to [5–7] for additional details about MeNSa.

#### 3.1 Framework Components

Figure 1 shows the primary components of the MeNSa framework along with their functional relationships. The DNS reference model specifies the attributes that must be measured in order to discern DNS health levels. Note that the point of view (PoV) is an inherent part of the DNS reference model that specifies DNS health from a local perspective for components and actors. A set of use cases provide detailed scenarios that outline the functional interactions between DNS components and actors. Measurement techniques and tools specify methods for obtaining the information necessary to compute the metrics. Metrics are derived that quantify DNS health based on inputs from the other primary components.

#### 3.2 Reference Architecture

Figure 2 presents a graphical display of the reference DNS architecture. The user application (e.g., Internet browser) is the actor that generates DNS queries. The application service provider is the actor that provides distributed services and applications, primarily via web service technologies. The name server resolves queries for a specific zone and can function as a master or slave. The resolver is a name server, often owned and managed by an Internet service provider (ISP), that receives DNS queries and either resolves the queries or

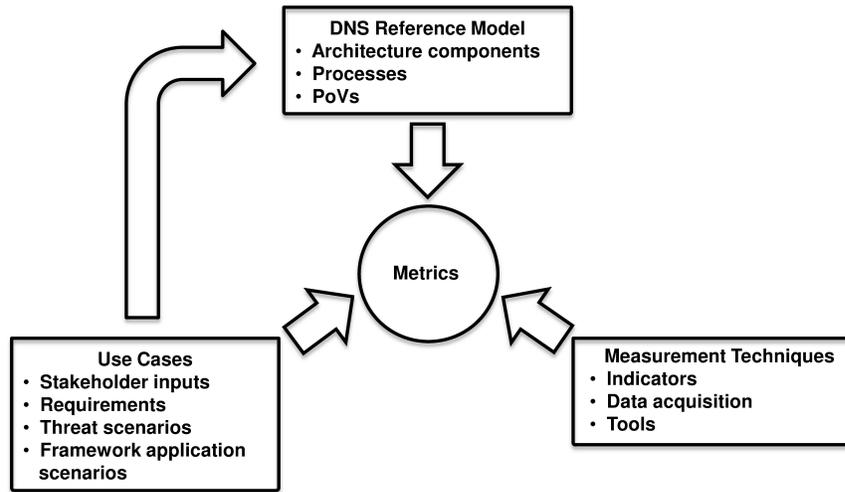


Figure 1. MeNSA framework.

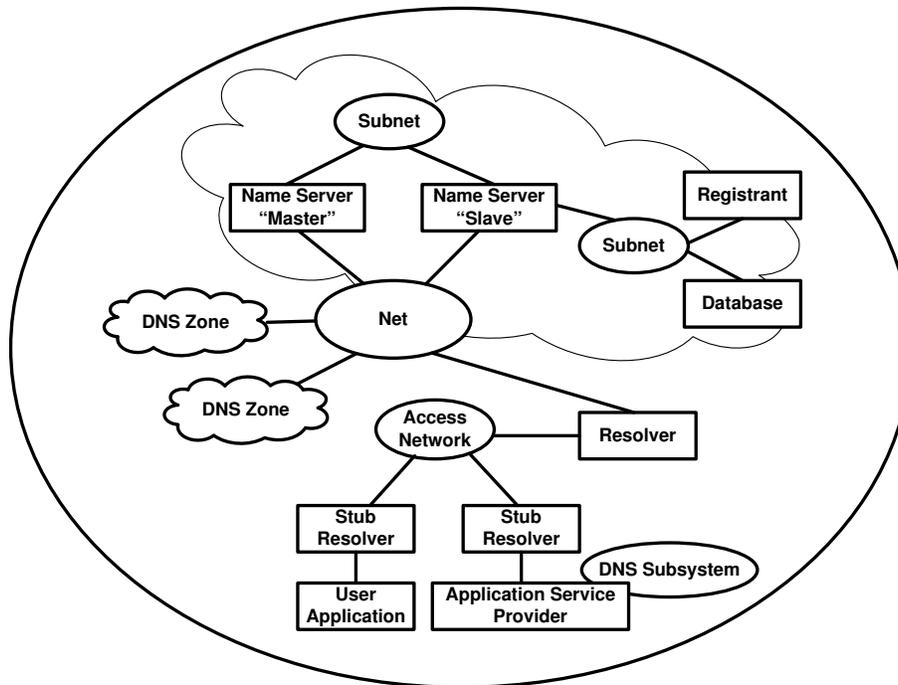


Figure 2. Reference architecture.

forwards them to the next server in the DNS hierarchy. The stub resolver is the operating system component that receives DNS requests from applications and sends them to the resolver. The net and subnet components represent the various network interconnections (e.g., LAN and Internet communication channels). The registrant represents the administrator of a zone. Databases store DNS information for each respective zone. The DNS zone is a specific DNS domain that is managed by a single administrative entity. Finally, the DNS subsystem represents an autonomous naming system that is isolated from the global DNS.

### 3.3 Point of View

The MeNSa framework is intended to provide DNS health awareness for end users, application service providers and operators (e.g., resolvers, name servers and registrars). Depending on their role and access to various components, each actor may have one or more views of DNS. Note that the perception of health is limited in scope by processes and components that each actor can observe and control.

The PoV concept helps categorize components that a specific DNS actor can observe and measure. Additionally, PoVs help identify the information that is required from other actors to properly assess DNS health. The six PoVs incorporated in the MeNSa framework are: (i) end user PoV; (ii) application service provider PoV; (iii) resolver PoV; (iv) name server PoV; (v) zone PoV; and (vi) global PoV.

Of particular interest in this work is the end user PoV, which represents the perspective from which each user can evaluate DNS. The components associated with the end user PoV are the user application, stub resolver and net. The specific operation of interest is the DNS lookup process.

### 3.4 Metrics

The proposed metrics are intended to evaluate DNS health based on vulnerability, security and resiliency. Table 1 provides example categories and metrics associated with the MeNSa framework. The vulnerability metrics are associated with repository corruption, system corruption and denial of service. Indicators for repository corruption include data staleness, zone drift/zone thrash and data coherence. System corruption indicators include zone transfer failure, DNS spoofing and cache poisoning. Denial of service indicators include DNS request variation, bandwidth consumption and traffic variation. Metrics for resiliency include indicators for mean time to discovery, mean time between failures and operational availability. Finally, security metrics are associated with indicators for attack surface, attack depth, attack escalation speed and annual loss expectancy. Interested readers are referred to [7] for a comprehensive list of derived metrics used in the MeNSa effort.

Table 1. DNS health and security metrics.

Indicator	Metric
Data Staleness	Percent of resource records that differ across authoritative servers
Zone Drift/Zone Thrash	Probability of incurring zone drift and zone thrash
Data Coherence	Percent of responses that differ between queries to the parent zone and authoritative server
Zone Transfer Failure	Number of failed zone transfer operations
DNS Spoofing	Probability of being spoofed
Cache Poisoning	Percent of content that differs between cache and authoritative data
DNS Request Variation	Variance of the number of requests per second
Bandwidth Consumption	Percent of available bandwidth
Traffic Variation	Variance of the incoming DNS traffic rate
Mean Time to Discovery	Average response time
Mean Time between Failures	Average time between invalid responses
Operational Availability	Percent of time executing at the expected service level
Attack Surface	Percent of nodes vulnerable to a certain type of attack
Attack Depth	Percent of nodes impacted by an attack
Attack Escalation Speed	Time required to affect a specified number of nodes
Annual Loss Expectancy	Financial loss as a result of incidents in one year

### 3.5 Framework Application

This section describes the main phases of the MeNSa framework and how the framework can be used in an operational environment. The application of the framework is organized into three macro phases: (i) preliminary diagnosis; (ii) service level objectives (SLOs) and scenario definition; and (iii) detailed diagnosis and measurement.

In the preliminary diagnosis phase, an initial evaluation of DNS health is conducted based on a subset of the metrics associated with the respective PoV. In the SLOs and scenario definition phase, one or more threat scenarios are derived given the PoV and representative indices. The detailed diagnosis and measurement phase assesses the perceived health level, achievable SLOs, causes of SLO violations and improvement actions.

The detailed diagnosis and measurement phase is further organized into three stages: (i) metric selection; (ii) measurement; and (iii) aggregation. The selection of metrics is an off-line process. The MeNSa framework enables users to predefine a set of validated metrics for each perceived threat scenario and PoV. The measurement stage involves data collection and the computation of the selected metrics. Note that we use a “bottom up” measurement model [5, 6] that first acquires information from other PoVs. Certain indices (e.g., network reachability and traffic load) help discern if a measurement can be effected by critical states of the infrastructure.

The aggregation stage combines the results from the measurement stage to provide aggregated indices that summarize DNS health as perceived by the

PoV. The indicators determine achievable SLOs, causes of health degradation and possible solutions. In the MeNSa framework, data aggregation is accomplished according to the following definitions:

- $M = \{m_1, \dots, m_M\}$  is the set of metrics used to evaluate DNS health and security.
- $D_i$  is the domain of the  $i$ -th metric.
- $v_{i,j} \in D_i$  defines values for the metric  $m_i$ . Note that  $j = 1, \dots, n$  where  $n$  is the number of computed values.
- $q_i: D_i \rightarrow [0, 1]$  is a “quality mapping” for metrics  $m_i$  with  $q_i$  transforming the measured values  $v_{i,j}$  into a dimensionless quality value  $q_{i,j} = q_i(v_{i,j})$ . Note that  $q_{i,j} = 1$  is the highest quality value and  $q_{i,j} = 0$  is the lowest quality value.
- $\{w_k\}$  is a set of “weight vectors,” where  $w_k = (w_{k,1} \dots w_{k,M})$  is a vector of weights such that  $w_{k,i} \in [0, 1]$  and  $\sum_{i=1}^M w_{k,i} = 1$ . Each vector  $w_k$  defines the aggregation of the  $M$  metrics corresponding to the  $k$ -th result.

Given the above definitions, the aggregation process can be specified as:

1. Choose a set of metrics to be aggregated and calculate  $n$   $v_{i,j}$  values.
2. Define a quality mapping  $q_i$  for each metric and transform the measured values into quality values  $q_{i,j} = q_i(v_{i,j})$ .
3. Aggregate the quality metrics by averaging the quality values using a weights vector  $v_k$ .

## 4. Experimental Evaluation

This section evaluates the utility of the MeNSa framework and the application of the associated metrics. A scenario-based experiment is used to demonstrate how a subset of defined metrics can be computed and aggregated for the end user PoV.

### 4.1 Measurements and Metrics

The experimental testbed consisted of a Windows machine running Firefox 8.0 and connected to the Internet through the Italian ISP Fastweb (7 Mbps nominal). DNS queries were sent to Fastweb’s recursive resolvers.

Data was collected during ten web browsing sessions ranging in duration from 10 minutes to 15 minutes and lasting a total of two hours. Data from each session was collected for aggregation, yielding  $n = 10$  values for each metric. The following metrics were computed and aggregated:

- **Incoming Bandwidth Consumption (IBC):** This is computed by dividing the total amount of incoming data by the duration of the measurement session. The domain of this metric is  $[0, IBC_M]$  and the metric

is measured in Mbps, where the value  $IBC_M$  is the nominal maximum bandwidth declared by the ISP.

- **Incoming Traffic Variation (ITV):** This measures the bandwidth variance for sessions. For a session  $i$ , ITV is given by:

$$ITV = \frac{IBC_i - IBC_{i-1}}{length_i}$$

where  $IBC_i$  is the incoming bandwidth consumption measured in the  $i$ -th session and  $length_i$  is the duration of the session. The domain of this metric is  $[-ITV_M, ITV_M]$  and the metric is measured in Mbps<sup>2</sup> where

$$ITV_M = \max_i \frac{IBC_M}{length_i}.$$

- **Traffic Tolerance (TT):** This specifies the round trip time (RTT) of an IP packet traveling between the end user's node and the ISP's recursive resolver. The domain of the metric is  $[0, +\infty]$  and the metric is measured in seconds.
- **Stub Resolver Cache Poisoning (CP):** This specifies the percentage of poisoned entries in the cache. The domain is  $[0, 100]$  with every entry in the cache being verified against a set of trusted recursive resolvers.
- **DNS Requests per Second (DNSR):** This is the total number of DNS queries in a session. The domain is  $[0, +\infty]$ .
- **Rate of Repeated Queries (RRQ):** This is the number of repeated DNS queries in a session. Under normal conditions, a name is resolved only once due to DNS caching. Many DNS queries for the same name during the same session could be an indicator of malicious activity. The domain is  $[0, +\infty]$ .

The following set of quality mapping functions for the metrics are employed:

- **Incoming Bandwidth Consumption (IBC):** The quality mapping  $q: [0, IBC_M] \rightarrow [0, 1]$  for the IBC metric is defined as:

$$q(x) = 1 - \frac{x}{IBC_M}$$

where  $IBC_M$  is the maximum bandwidth value provided by the ISP.

- **Incoming Traffic Variation (ITV):** The quality mapping  $q: [-ITV_M, ITV_M] \rightarrow [0, 1]$  for the ITV metric is defined as:

$$q(x) = 1 - \frac{|x|}{ITV_M}.$$

Table 2. Measurements and quality ratings for Sessions 1, 2 and 3.

	IBC		ITV		TV		CP		DNSR		RRQ	
	Mbps	q	Mbps <sup>2</sup>	q	s	q	%	q	#	q	#	q
$S_1$	11.8	0.998	0	1	0.80	0.80	9.96	0.90	0.87	1	0.84	0.84
$S_2$	11.9	0.997	0.0054	0.999	0.74	0.74	6.67	0.93	0.33	0	0.89	0.79
$S_3$	13.9	0.997	0.0002	0.999	0.78	0.78	10.40	0.89	0.24	0	0.74	0.74

- **Traffic Tolerance (TT):** The quality mapping  $q: [0, +\infty] \rightarrow [0, 1]$  for the TT metric is defined as:

$$q(x) = \begin{cases} 1 & x \leq RTT_{Av} \\ -\frac{x}{RTT_{Av}} + 2 & RTT_{Av} < x \leq 2RTT_{Av} \\ 0 & x > 2RTT_{Av} \end{cases}$$

where  $RTT_{Av}$  is the average RTT value during the session.

- **Cache Poisoning in the Stub Resolver (CP-SR):** The quality mapping  $q: [0, 100] \rightarrow [0, 1]$  for the CP-SR metric is defined as:

$$q(x) = 1 - \frac{x}{100}.$$

- **DNS Requests per Second (DNSR):** The quality mapping compares the current DNS behavior against a previous reference. The quality mapping  $q$  for the DNSR metric is defined as:

$$q(x) = \begin{cases} 1 - \frac{x}{2 \cdot DNSR_{Av}} & 0 \leq x \leq 2DNSR_{Av} \\ 0 & x > 2DNSR_{Av} \end{cases}$$

where  $DNSR_{Av}$  is the average number of the DNS requests per second during the session.

- **Rate of Repeated Queries (RRQ):** The quality mapping  $q$  for the RRQ metric is defined as:

$$q(x) = 1 - \frac{x}{R_M}$$

where  $R_M$  is the maximum number of DNS requests in the current session. Note that  $R_M$  changes for different sessions.

## 4.2 Aggregation and Experimental Results

Table 2 shows the measurement values and related quality ratings for the experiment. For brevity, data for Sessions 4 through 10 are not presented.

Table 3. Session 1 aggregate results.

	IBC q = 0.998	ITV q = 1	TV q = 0.801	CP q = 0.9	DNSR q = 1	RRQ q = 0.842	Aggregate Result
TE	0.19	0.19	0.19	0.05	0.19	0.19	0.927
PI	0.00	0.00	0.00	1.00	0.00	0.00	0.900
DoS	0.20	0.20	0.20	0.00	0.20	0.20	0.928
Net	0.33	0.33	0.33	0.00	0.00	0.00	0.932
SR	0.00	0.00	0.00	0.12	0.44	0.44	0.918

For every session, the quality ratings of the metrics are aggregated for the end user PoV indices. Table 3 presents the following aggregate results for the first session:

- **Total Evaluation Index (TE):** This provides a global assessment of the PoV aggregated over all considered metrics.
- **Protocol Issues Index (PI):** This estimates possible DNS protocol problems. The index is related to the cache poisoning metric.
- **Denial of Service Index (DoS):** This evaluates how improbable DoS is in a given scenario. The DoS index aggregates all the metrics except for cache poisoning.
- **Net Index (Net):** This estimates the performance of the network component. The Net index aggregates incoming bandwidth consumption, incoming traffic variation and traffic tolerance.
- **Stub Resolver Index (SR):** This evaluates stub resolver performance. The SR index aggregates cache poisoning, DNS requests variation per second and rate of repeated queries.

The final result of each aggregated index for the end user PoV is computed as the average of the results over all ten sessions. The variances are computed to provide estimates of the uncertainty of the results. Figure 3 shows the final values.

### 4.3 Discussion

The total evaluation index is the primary consideration for the end user PoV – it reflects the overall DNS health using components that can be measured by end users. In the investigated scenario, minor disruptions are deemed to be acceptable (e.g., temporary DNS failures that require the reloading of web pages). For this reason, total evaluation index values less than one are acceptable in a properly functioning system. With the MeNSa framework, it is possible to quantify service levels and to verify if SLOs are violated. As an example, in our experiment, the total evaluation value was computed to be 0.833 with an uncertainty value  $\pm 0.134$ . Such a value quantifies DNS health as perceived by the

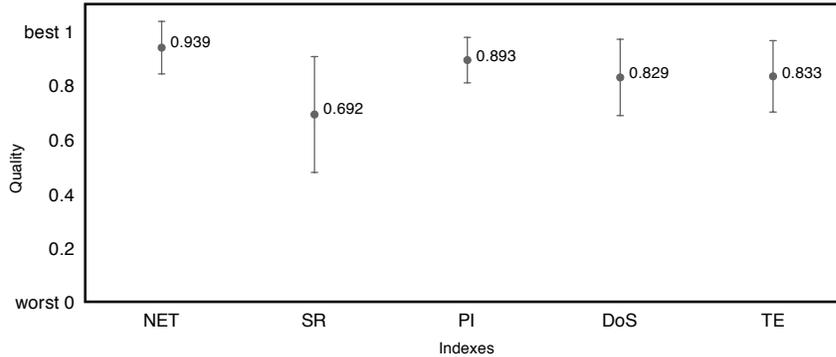


Figure 3. Aggregated results and index uncertainty over all sessions.

end user. The other aggregated results provide insight into the performance of different aspects of the system. This information is valuable because it can help identify components that require further scrutiny in the event of malfunctions.

Our calculations show that the stub resolver is the component that has the highest likelihood to have problems; this is because the stub resolver index value of 0.692 is far from one. In contrast, the Net component evaluation is 0.939 and has a low degree of uncertainty. It is important to note that further analysis is possible if the aggregated results of the recursive resolver PoV are available as an input metric for the end user PoV. In other words, the outputs of a PoV can be used as input metrics to another PoV to provide results with finer granularity. Aggregating PoV values with available local metrics increases the accuracy of an overall assessment and refines the evaluation of single components.

Our investigation also focused on threat scenarios that could affect a targeted infrastructure. Indeed, some of the results provide insights into the likelihood of certain threats or attacks. For example, the high values of the protocol issues and denial of service indices (0.893 and 0.829, respectively) indicate, with a high degree of certainty, that the system was not affected by protocol issues or denial-of-service attacks during the measurement period.

It is important to note that the results presented above cannot be generalized and must be validated using larger sets of experiments. Nevertheless, the study demonstrates the ability to measure and aggregate DNS health metrics.

## 5. Conclusions

DNS is a critical component of the Internet. Indeed, without DNS services the majority of Internet applications would not function properly. The increasing use of information and communications technologies in critical infrastructure assets makes it vital to protect DNS – targeted attacks that degrade DNS could cause serious consequences to modern society.

The MeNSa framework provides a formal methodology for evaluating DNS health based on requirements identified by the DNS community. The experi-

mental results demonstrate that end user metrics can be aggregated to verify the level of service and identify potential threats. Our future research will continue our efforts at validating the MeNSa framework using larger data sets and also expand the framework to consider other points of view.

## References

- [1] R. Arends, R. Austein, M. Larson, D. Massey and S. Rose, DNS Security Introduction and Requirements, RFC 4033, 2005.
- [2] R. Arends, R. Austein, M. Larson, D. Massey and S. Rose, Resource Records for the DNS Security Extensions, RFC 4034, 2005.
- [3] R. Arends, R. Austein, M. Larson, D. Massey and S. Rose, Protocol Modifications for the DNS Security Extensions, RFC 4035, 2005.
- [4] S. Bellovin, Using the domain name system for system break-ins, *Proceedings of the Fifth USENIX UNIX Security Symposium*, p. 18, 1995.
- [5] E. Casalicchio, M. Caselli, D. Conrad, J. Damas and I. Nai Fovino, Framework Operation, the Web User PoV, Technical Report, Version 1.1, Global Cyber Security Center, Rome, Italy, 2011.
- [6] E. Casalicchio, M. Caselli, D. Conrad, J. Damas and I. Nai Fovino, Reference Architecture, Models and Metrics, Technical Report, Version 1.5, Global Cyber Security Center, Rome, Italy, 2011.
- [7] E. Casalicchio, D. Conrad, J. Damas, S. Di Blasi and I. Nai Fovino, DNS Metric Use Cases, Technical Report, Version 1.0, Global Cyber Security Center, Rome, Italy, 2011.
- [8] S. Castro, D. Wessels, M. Fomenkov and K. Claffy, A day at the root of the Internet, *ACM SIGCOMM Computer Communication Review*, vol. 38(5), pp. 41–46, 2008.
- [9] D. Danchev, Hackers hijack DNS records of high profile New Zealand sites, *ZDNet*, San Francisco, California, April 21, 2009.
- [10] D. Eastlake, Domain Name System Security Extensions, RFC 2535, 1999.
- [11] D. Eastlake and C. Kaufman, Domain Name System Security Extensions, RFC 2065, 1997.
- [12] Global Cyber Security Center, The Measuring Naming System Project, Rome, Italy ([www.gcsec.org/activity/research/dns-security-and-stability](http://www.gcsec.org/activity/research/dns-security-and-stability)), 2012.
- [13] D. Goodin, Cache-poisoning attack snares top Brazilian bank, *The Register*, April 22, 2009.
- [14] I. Nai Fovino, S. Di Blasi and A. Rigoni, The role of the DNS in the secure and resilient operation of critical infrastructures: The energy system example, presented at the *Sixth International Conference on Critical Information Infrastructure Security*, 2011.

- [15] Internet Corporation for Assigned Names and Numbers, Root Server Attack on 6 February 2007, DNS Attack Factsheet 1.1, Los Angeles, California, 2007.
- [16] Internet Corporation for Assigned Names and Numbers, Response to Recent Security Threats, ICANN Technical Report, Los Angeles, California, 2008.
- [17] Internet Corporation for Assigned Names and Numbers, Security, Stability and Resiliency of the Domain Name System, ICANN Technical Report, Los Angeles, California, 2009.
- [18] Internet Corporation for Assigned Names and Numbers, Measuring the Health of the Domain Name System, Report of the Second Annual Symposium on DNS Security, Stability and Resiliency (Kyoto, Japan), Los Angeles, California, 2010.
- [19] D. Kaminsky, It's the end of the cache as we know it, presented at *Black Hat USA*, 2008.
- [20] B. Krebs, Hackers hijacked large e-bill payment site, *Washington Post*, December 3, 2008.
- [21] R. Liston, S. Srinivasan and E. Zegura, Diversity in DNS performance measures, *Proceedings of the Second ACM SIGCOMM Workshop on Internet Measurement*, pp. 19–31, 2002.
- [22] E. Mills, Puerto Rico sites redirected in DNS attack, *CNET*, San Francisco, California, April 27, 2009.
- [23] M. Santcroos and O. Kolkman, DNS Threat Analysis, Technical Document 2006-SE-01 version 1.0, NLnet Labs, Amsterdam, The Netherlands, 2007.
- [24] Y. Sekiya, K. Cho, A. Kato and J. Murai, Research of method for DNS performance measurement and evaluation based on benchmark DNS servers, *Electronics and Communications in Japan (Part I: Communications)*, vol. 89(10), pp. 66–75, 2006.
- [25] P. Vixie, DNS and BIND security issues, *Proceedings of the Fifth USENIX UNIX Security Symposium*, p. 19, 1995.
- [26] P. Vixie, G. Sneeringer and M. Schleifer, 21 Oct 2002 Root Server Denial of Service Attack – Report, Technical Report, Internet Systems Consortium, Redwood City, California, 2002.