

Measuring the Global Domain Name System

E. Casalicchio, Univ. of Rome Tor Vergata

M. Caselli and A. Coletta, Global Cyber Security Center (GCSEC)

Abstract

The Internet is a worldwide distributed critical infrastructure, and it is composed of many vital components. While IP routing is the most important service, today the Domain Name System can be classified as the second most important, and has been defined as a critical infrastructure as well. DNS enables naming services used by every networked application and therefore by every networked critical infrastructure. Without DNS all services used in daily life activities (e.g., commerce, finance, industrial process control, logistics, transportation, health care) become unavailable. A big challenge is to guarantee the proper level of DNS health. Providing DNS health requires monitoring the system, analyzing its behavior, and planning and actuating corrective actions. There are several initiatives in this field, all claiming to be able to measure the DNS health from a local perspective. The reality is a bit different and many challenges are still open: no standard metric exist (only a shared list of five health indicators); no common rules to compute health indicators are agreed; no common concept of regular DNS behavior is defined. The Measuring the Naming System (MeNSa) project proposes a formal and structured methodology and a set of metrics for the evaluation of the DNS health and security levels. This article discusses the problem of measuring the DNS health level and introduces the main concepts of the MeNSa project. Finally, using a real case study, the problem of metrics aggregation is discussed.

ritical infrastructures are increasingly incorporating in a massive way networked components. This trend obviously allowed the services provided to be enhanced and optimized, distributed self-orchestration mechanisms to be implemented, and remote installations to be managed efficiently. On the other hand, as a consequence, the Internet infrastructure used to realize these services must be considered now as part of the critical infrastructure themselves. For example, in [1] the authors proposed an architecture to secure the Domain Name System (DNS) with the final goal of protecting critical infrastructures. A recent study [2] showed how an attack on the DNS might impact several levels of the operation of energy smart grids.

The DNS, as part of the Internet infrastructure, constitutes the backbone of the modern cyber world. In 2007, the Internet Engineering Task Force's (IETF's) DNS Extensions Working Group (DNSEXT) identified the DNS as "a critical Internet infrastructure" because it resolves billions of queries per day in support of global communications and commerce. In 2011, Steve Gibbard, at the spring DNS-OARC¹ workshop in San Francisco, California, stressed that the DNS is a critical infrastructure. However, the DNS being a completely distributed infrastructure and completely seamless to end users has con-

tributed to making it one of the less considered infrastructures of the Internet when speaking of cyber security. This is no longer true; Kaminsky's exploit, and attacks that in the last years have taken advantage of the weaknesses of the DNS in order to damage cyber-infrastructure, have posed serious questions about the security, safety, and stability of this system.

The community has widely discussed the problem of the security of DNS and its impact on the cyber society. In 2010 [3] the concept of DNS health came out as a way of defining when the DNS system is well functioning, taking as an example human body health. The concept of DNS health is developed around five main indicators: availability, coherency, integrity, resiliency, and speed. We believe that such a list needs to be extended with the concepts of stability and security, since a system cannot be described as "healthy" if it is not stable and can become quickly unhealthy if it is not secure. Stability is intended as the desired DNS feature to function reliably and predictably day by day (e.g., protocols and standards). Stability facilitates universal acceptance and usage. Hereafter, when we talk about DNS health we include the concepts of security, stability, and resiliency.

The concepts expressed in [3] remained at the abstract level, without suggesting how to assess such a complex property. What is indeed unclear in this definition is the way in which it would be possible, from real, measurable observation, to obtain "numbers" or, better, indices, that can be used to quantify these global properties and, at last, to summarize the "level of health" of the portion of DNS under analysis. Shortly speaking, no standard frameworks for DNS measurement

¹ DNS-OARC, the Domain Name System Operation Analysis and Research Center, is a non-profit membership organization that seeks to improve the security, stability, and understanding of the Internet's DNS infrastructure. The main DNS operators are members of DNS-OARC.

exist, and no standard metrics have been defined. Moreover, how to define a common concept of *regular* DNS behavior and how to develop a standard framework for data/information sharing are still open issues [4].

In the literature there are many studies related to DNS traffic measurement and performance metrics (e.g., [5–7]), but very few report on the measurement and quantification of security, stability, and resiliency, or, in general, DNS health. Moreover, to the best of our knowledge, a complete framework dealing with DNS health evaluation has not yet been developed.

In this article, we present our “answer” to this challenge. After a brief description of the MeNSa project [8], we show how it is possible to aggregate several different metrics related to the DNS system, after identifying a well defined measurement point of view, in order to obtain aggregate indicators of its health. Nevertheless, in the article we consider the end-user perspective, and the solution we propose can be applied to any observation point and for any set of metrics.

The article is organized as follows. We briefly describe DNS vulnerabilities and show examples of real incidents. We present the MeNSa project. We discuss the metrics aggregation problem and describe the proposed solution. We introduce a case study to show how metric aggregation can be computed and used to evaluate a real case. We then conclude the article.

Vulnerabilities and Incidents

The DNS threats can be broadly classified into three main categories: *data corruption*, *denial of service*, and *privacy*.

Data corruption clusters all types of incidents related to the unauthorized modification of DNS data. These incidents can happen in every part of the DNS propagation chain and can be related to corruption of repositories (e.g., databases containing resource records, zone files, DNSSEC keys, and so on); cache consistency [9]; alteration of the authenticity of DNS responses; and protocol issues, which deal with design flaws of the DNS protocol. Examples of protocol issues are cache poisoning (i.e., the well known Kaminsky’s attack), route injections, and man-in-the-middle threats. The first security flaws in the protocol were discovered in the early 1990s: in [10, 11] the authors pointed out how it would be possible to fool name-based authentication systems by means of cache contamination attacks. To solve this problem a security extension, DNSSEC, was proposed (IETF RFCs 2065 and 2535). Recent vulnerabilities discovered by Dan Kaminsky via cache poisoning attacks finally led to the development of the latest generation specifications, IETF RFC 4033, 4034, and 4035.

Remaining in the context of data corruption, several DNS hijacking attacks have been reported since 2008, where the attacked domain name registrars were the target. For example, in 2008, a large e-bill payment site was compromised (targeting its domain name registrar) by redirecting its visitors to a crafted web address, later attempting to install malicious code on the visitors’ machines. A major case of a successful cache poisoning attack against the DNS infrastructure was reported in Brazil in 2009, against one of the major Brazilian banks: the login page redirection toward a fraudulent site caused the theft of users’ access credentials. In 2012, the *.ke domains had a good share of data corruption attacks; for example, 103 of the Government of Kenya’s websites (.go.ke) were hacked in one night [12].

Denial of service attacks are aimed at impacting the DNS infrastructure composed of DNS servers and the network con-

nections. There have been two major reported distributed denial of service (DDoS) attacks on the root servers, in 2002 and 2007. The first attack covered a timeline of around one hour and targeted simultaneously all of the 13 root DNS servers, affecting overall performance and in particular degrading the availability level for some of them. In light of this attack, the Anycast protocol was implemented in several root servers, mitigating the second wave of global coordinated DDoS attacks, which occurred in 2007.

A successful Denial of Service attack on a regional name server or on a name server in a higher position in the hierarchy can completely make inaccessible distributed applications (from web site to control systems) in a entire DNS zone or geographical region.

Finally, *privacy* threats are related, for example, to snooping of DNS caches. Privacy, despite its relevance, is out of the scope of our investigation.

The MeNSa Project

The scope of the MeNSa project [8] is to define a methodology and a set of metrics to quantify the global health level of the DNS. The DNS community agrees on the fact that while it is a common practice to individually monitor the DNS subsystems to observe if the traffic parameters deviate from the average, it is a challenge to extract knowledge on the more global DNS behavior and its “normality” [4].

The key actions we propose to face this challenge are:

- To refine and improve existing metrics for DNS health indicators
- To define a metric aggregation model to merge measured metrics into a few indicators
- To identify metric threshold levels that allow the DNS community to trigger when the behavior is normal or abnormal

While in the long term the MeNSa project would provide a solution to all the above items, in this article, we concentrate our attention on *metric aggregation*.

The most relevant concepts behind the MeNSa methodology are summarized in the following.

The DNS reference model defining the boundaries of the system we want to measure. Figure 1 shows the simplified architecture we consider. The *end user application* (e.g., browser, Apps, thin/fat clients) generates DNS queries, and can have advanced features such as prefetching and internal caching. Name servers work at a different level of the hierarchy, from root zones to local caches. Also of great importance are the Anycast resolvers.

The set of metrics to quantify the health and security level of the DNS. The metrics we propose are intended to evaluate the health of the DNS by measuring the DNS along three dimensions: vulnerabilities, security, and resiliency. Examples of metrics, clustered by threat category, are reported in Table 1. A comprehensive description can be found in [8].

The set of measurement techniques and tools put in place to gather information needed to compute metrics. How measurements are implemented depends on two main factors:

- What can be measured from which point
- The time horizon of data collection (e.g., seconds, hours, days, or months)

Measurement techniques and data collection issues are out of the scope of the project (and of this article).

The concept of point of view (PoV). A PoV is intended as the perspective of a DNS actor/component in observing, using, operating, and influencing the global DNS. Potential users of the MeNSa methodology fall into one of the following categories: *end users*, who are mostly unaware of the DNS function and operation; *service providers*, such as the Internet




Figure 1. DNS reference architecture considered in the MeNSa project.

and application service providers; and *operators* (e.g., resolvers, name servers, registrars). The definition of different points of view is intended to categorize which components can be observed and measured by a specific DNS actor and what information is needed from other DNS actors to properly assess the level of DNS health perceived. This categorization will allow, for each PoV, defining a set of health indicators and a set of measurable metrics needed to evaluate the indicator of interest. The six points of view we defined are end-user PoV, application service provider PoV, resolver PoV, name server PoV, zone PoV, and global PoV (Fig. 1).

From each PoV it is possible to measure the perceived health level mixing two sources of information (Fig. 2): direct observations of the global DNS behavior and the Internet traffic (e.g., through active and passive measurement); and data shared by other PoVs (e.g., in the form of aggregated measures or anonymized data). The strength of this distributed approach is that measurement is kept local, and only aggregated information is shared. Local data collection is usually performed by DNS operators that daily collect tons of data to manage their infrastructure; therefore, there is no overhead. Sharing aggregated information (e.g., health indices) introduces negligible network overhead. The only additional tasks for a DNS operator are data processing to compute health and security metrics and data sharing. Both tasks imply the agreement of providers on a set of metrics and sharing rules. As reported in [4], the community is moving in this direction.

Using this approach, it should be possible to build a global picture of the health state of the DNS.

Methodology in Action

The methodology operation is organized in three macro phases:

- *Preliminary diagnosis* that, from the chosen PoV, performs a first evaluation of the health level perceived conducting simple measurements and assessments.
- *The definition of the service level objectives (SLOs) and scenario phase*, given that the PoV is allowed to select one or more threat scenarios and the measurable and representative indices.
- *The detailed diagnosis and measurement phase* assessing the

DNS health level perceived; the achievable SLOs; the causes of SLO violation and improvement actions. The detailed diagnosis and measurement phase is organized in three stages: selection of metrics, measurement, and aggregation.

At the *aggregation* stage, all the measures collected are combined to provide aggregated indices summarizing the DNS health level perceived by the PoV, what the achievable SLOs are, and finally, what the cause of health degradation could be and possible solutions. Measures are performed using network and DNS measurement tools typically used by the community.

Metrics Aggregation

Internet measurement theories and practices suggest that to understand the system behavior, it is better to

look at a set of metrics rather than a single indicator. Trade-offs exist between different metrics [13], and a single value can be misleading. This is true from a technical or scientific perspective, while non-technically skilled users or decision makers are typically affected by “mononumerosis,” an undue focus on a single measured value (as defined in the 1990s by Cindy Bickerstaff of the IETF IPPM metrics working group).

The concept of DNS health is multi-faceted and, as above mentioned, can hardly be captured considering the measurements on a single metric. However, non-technically skilled users do not appreciate the details of a multitude of metrics and prefer one or very few simple indicators. To overcome this problem we propose to compute a limited set of indicators that have three fundamental features:

- They provide aggregate technical level metrics.
- They provide a general understandable (by any user) system state description.
- They give a measure of the security level intended not only as prevention against unauthorized access but also for performance, stability, and resiliency.

Let us now discuss how DNS health metrics can be aggregated. Formally, a PoV is associated with:

- A set of M metrics $\{m_1, \dots, m_M\}$. Let D_i be the *domain* of the metric m_i (i.e., the measured values v_{i1}, v_{i2}, \dots) of m_i belongs to D_i .
- A set of M quality mappings $q_i : D_i \rightarrow [0,1]$, one for each metric m_i . The mapping q_i transforms the measured value v_{ij} into a dimensionless quality value $q_{ij} = q(v_{ij})$, where 0 indicates the lowest quality and 1 indicates the highest one.
- A set of *aggregated indicators*. Each indicator I_k is fully defined by its *vector of weights* $w_k = (w_{k1} \dots w_{kM})$ such that $\sum_{i=1}^M w_{ki} = 1$.

Different techniques can be used to aggregate metrics. These techniques do not depend on the specific PoV and should satisfy two properties. First, the aggregation process should not depend on the metrics to be aggregated and the aggregated index. Second, the timescale of the phenomena observed should not influence the aggregation, that is, the aggregation process should be capable of handling variable timescales and sampling frequencies.

For example, given a metric m_i and an observation time period T , and a chosen sampling time interval, it is possible to

Category	Measure	Metric
Repository Corruption	Data Staleness	Percentage of differing SOA serial numbers across all auth. servers numbers over a time period.
	Zone drift/Zone trash	Probability of incurring in zone drift and zone thrash status
	NS Parent/Child Data Coherence	Percentage of differences between the responses to NS queries to the parent zone with the responses to NS queries among all authoritative servers for the zone within one serial number
System Corruption	Cache Poisoning	Percentage of differences between the contents of caches vs. authoritative data
	Zone Transfer failure	Number of failed zone transfer operations
	DNS spoofing	Probability of being spoofed and probability of being spoofed over a time period
Denial of Service	Variation of DNS Request per Second	Variation of the requests number per second
	Incoming bandwidth consumption	Percentage of the available bandwidth
	Incoming traffic variation	Variation of incoming DNS traffic
Resiliency	Mean Time to Incident Discovery	Average value over a long observation period
	Operational Mean Time Between Failures	Average value over a long observation period
	Operational Availability	Percentage of the mean time an ICT system is running at the normal service level over the observation time period
Security	Attack Surface	Percentage of nodes of a target system that is susceptibility to a certain type of attack.
	Attack Deepness	Percentage of impacted nodes of a system as consequence of an attack
	Attack Escalation Speed	Attacks in a time unit variations
	Annualized Loss Expectancy	Dollars loss as consequence of incidents per years

Table 1. Examples of DNS health and security metrics.

identify S_i sessions $\{s_{ij}\}$ with $j = 1 \dots S_i$, and to compute S_i values of the metric.

Having computed the v_{ij} values of m_i , it is possible to evaluate the quality values q_{ij} through the quality function q_i . Then the mean value \bar{q}_i and the standard deviation Δq_i over the S_i sessions are computed as the quality value of the metric m_i and the corresponding uncertainty level, respectively.

The aggregated indicators are computed as weighted averages using their vectors of weights.

An estimation of the uncertainty can be expressed by the squared weighted average, as is standard in error theory. Formally, the k th aggregator and the error estimations are computed as

$$I_K = \sum_{i=1}^M w_{ik} q_i \quad \Delta I_k = \sqrt{\sum_{i=1}^M w_{ik}^2 \Delta q_i^2}$$

Figure 3 shows an example of the aggregation process. In the example six metrics are considered, which are IBC, ITV, TT, CP DNSR and RRQ (see the next section for the descriptions of their meanings). Metrics are measured and transformed into quality values at step 1. At step 2 the average value and the error for each metric are computed. Finally, at

step 3, the values representing each single metric are aggregated using the weighted average. The choice of weight w_{ik} is arbitrary and depends on the relevance of metric m_i for the health index I_k . An example is provided later.

A Case Study

As an example to clarify the use of the proposed methodology we consider the end-user PoV, which represents the perspective from which each user can evaluate the naming service. From the end-user PoV, the components involved in the resolution process are the end-user application, the stub resolver, and the network, while the only operation of interest is the DNS lookup process.

This case study and the set of experiments that follows are designed with the following objectives in mind: to show examples of aggregated indicators, to show how a set of the metrics can be computed and how these metrics can be aggregated, and to explain how the values assumed by the aggregated indexes should be interpreted.

The aggregate indicators we consider are the following:

- Total evaluation (TE) index. It gives a global assessment of the PoV aggregating all the considered measurable metrics




Figure 2. Concept of PoV. The PoV creates its own knowledge of the health state of the DNS mixing local and remote information. Locally collected information is shared with other PoVs.

- Protocol issues (PI) index. It estimates possible DNS protocol problems (e.g., cache poisoning)
- Denial of service (DoS) index. It evaluates how improbable a DoS is in a given scenario.
- Network (NET) index. It estimates the performance of the network components.
- Stub resolver (SR) index. It evaluates the performance of the stub resolver (i.e., the operating system libraries that implement the DNS queries).

We chose these indices because:

- They are common in PoVs.
- They suit well the metrics we have chosen and those that are used by the DNS community. Such health indices are versatile, and indeed could be bound with different metrics depending on the specific PoV and available data. This flexibility also allows us to cope with evolving security threats.
- They give a measure of the security level intended not only as prevention against unauthorized access, but also for performance, stability, and resiliency.

Moreover, these indices are pragmatic (the DNS community does not like theoretical solutions very much) and are simple enough to be understood by non-technical people.

Measurements and Metrics

We set up a testbed where two client machines running Windows OS and Firefox 8.0 query the DNS from two different Internet service providers (ISPs). To be specific, a client was connected to the Internet through the Italian ISP Fastweb using as its access point the GCSEC laboratory (located in South East Rome, Italy); and the other client was connected to the Internet through the GARR network using as its access point the University of Rome “Tor Vergata” (UTV). DNS resolutions are demanded by Fastweb and UTV resolvers, respectively.

During the tests we collected traffic from 12 web browsing sessions each. Every session lasted from 10 min to a total of 2 h. Collected traffic is analyzed to get a measure of the metrics.

In the MeNSa project we identified a large set of metrics

useful to assess DNS health and security. As explained in the project deliverables [8], we started by considering all the most important threat scenarios for the DNS. A metric is interesting if it is capable to track system changes and deviation from normal behavior. In the following experiments we select only the metrics computable in the end-user PoV that are able to represent the system dynamic in a timespan of 2 h. These metrics are:

- Incoming bandwidth consumption (IBC), the ratio between the total amount of incoming bits during a session and the duration of the session.
- Incoming traffic variation (ITV), defined, for each session i , as the variation of IBC measured in session i in respect to the value of IBC measured in section $i - 1$.
- Traffic tolerance (TT), measuring the round-trip time (RTT) of an IP packet flowing between the end-user node and the ISP’s recursive resolver.
- Stub resolver cache poisoning

(CP), measuring the percentage of poisoned entries of the cache. Every entry of the cache is checked against a set of known recursive resolvers.

- DNS requests per seconds (DNSR), giving the total number of DNS queries in the session.
- Rate of repeated queries (RRQ) returns the number of repeated DNS queries in a session. During normal behavior, in a short time period, a name should be resolved only once because of DNS caching. If there are many DNS queries for the same name in the same session, this could be a marker of misbehavior.

IBC and ITV are measured using NetAnalyzer; TT is measured using ping. DNSR and RRQ are measured monitoring the session with WireShark and analyzing the resulting PCAP file. Finally, CP is measured by dumping the cache and parsing its content vs. authoritative DNS servers. The comparison is done immediately after the section to avoid resolver caches expiration.

DNS Health Evaluation

Figure 3 shows an example of quality values computed for each session and the results of the metric-based aggregation explained earlier. The figure also shows the weight values. The TE index gives an overall evaluation; thus, it must aggregate all the available metrics with the same weight. The PI index in our case only refers to the cache poisoning problems, because in our test we decided to measure only this protocol issue. Thus, the corresponding vector of weights consists only of the cache poisoning metric. The DoS index aggregates all the metrics but the cache poisoning one because it focuses on network traffic. The NET index focuses only on network related metrics, equally considered. The SR index focuses only on stub resolver measures, giving more importance (≈ 75 percent) to DNSR and RRQ.

Using our testbed we set up two experiments: one reproduces normal working conditions and the other a cache incoherency scenario. The first experiment consists of two different sets of measurements that can be compared: one

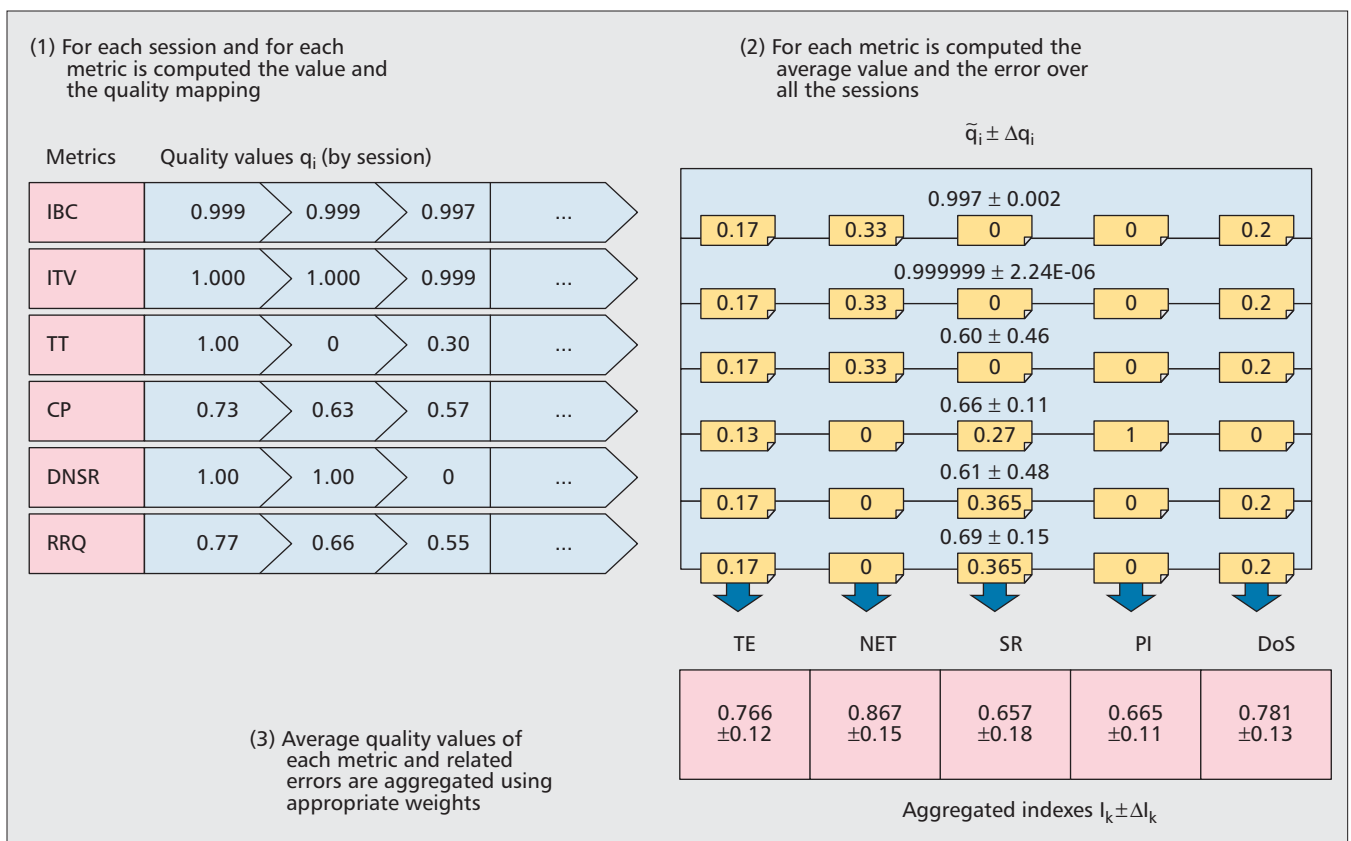


Figure 3. An example of quality values and results of the metric-based aggregation. Steps 1–3 are executed in sequence. The weights that define the aggregation (represented in yellow boxes) are tuned in the validation phase of the methodology.

collected at GCSEC and another one at UTV (hereafter referred to as the GCSEC laboratory and Uniroma2 tests). Figure 4 shows the results of the first experiment. The TE index values computed at Uniroma2 and GCSEC are 0.79 and 0.76, respectively, showing that the overall performance is good enough in both cases. A further analysis of the other indices may be useful to increase the performance. The PI and DoS indexes give insights on the possible *issues* of the system, while the NET and SR indexes focus on the performance of the *system components*. The DoS index shows an equally good result in both settings (about 0.8). Indeed, there was no DoS issue, and no further investigation is needed. On the other hand, in the GCSEC experiment the PI index (about 0.67) is lower than the corresponding index in the Uniroma2 experiment (about 0.85). Thus, it emerges that the GCSEC access network should improve DNS security (e.g., changing its Internet access provider or changing contract or managing its own cache). The NET index shows that the network component worked properly in both tests (about 0.86). Instead, the SR index also shows good results, but the value 0.65 measured in the laboratory test suggests some possible improvement in the performance of the stub resolver used.

In the second experiment we simulated some cache poisoning in order to validate our methodology. We manually corrupted 10 percent of the DNS cache entries in the GCSEC laboratory. The TE decreases to 0.7. The NET index is still evaluated around 0.8, but the SR assessment goes down to 0.6. These results entail the presence of problems in the DNS libraries of the operating system as expected. Going further, through the measurement process we discovered that we clearly suffer from some protocol issues since $PI = 0.38$. The DoS indicator, however, remains above 0.75. Figure 5 contains the results of this experiment,




Figure 4. Comparison of the health and security level perceived by end users using the DNS from two different ISP.

where the normal behavior data were measured in the laboratory.

The results can lead to practical actions. Comparing the SR and PI indicators enables spotting the cache poisoning problem. Indeed, the result of the analysis should suggest refreshing the DNS cache. Repeating the same evaluation afterward would further validate this suggested action.

The results we obtained cannot be generalized and must be validated with a larger set of experiments. Our goal was to show that measurement is possible, and how metrics can be used and aggregated to investigate DNS health.

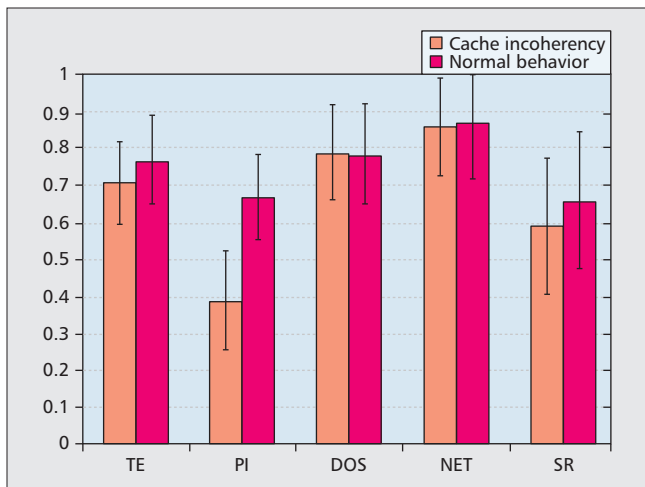


Figure 5. Comparison of the health and security level perceived by end users when a cache is poisoned. Poisoning has been artificially introduced in our collected data set.

Concluding Remarks

The Domain Name System constitutes the hidden backbone of the Internet. Without its services almost all the applications making use of the public network would not be able to operate in an efficient manner. The massive use of information and communications technology systems in critical infrastructures puts the DNS under the lights as a new potential source of disservices.

The DNS community in the last few years has started to reflect on the need of methodologies for assessing the health of the global DNS system. In this article, after describing at a high level the MeNSa project, designed to fulfill this need, we provided the results of the first tests on field, showing how, from the end-user PoV, metrics can be aggregated and used as a tool to verify the level of service perceived and the presence or absence of threats. The aggregation method proposed is general enough to be applied to any PoV. Of course, the implementation, that is, the choice of aggregated indices, of the set of metrics and the definition of the quality mapping functions are strictly related to the PoV and the goal of the analysis.

References

- [1] Y. Huang, D. Arsenault, and A. Sood, "SCIT-DNS: Critical Infrastructure Protection Through Secure DNS Server Dynamic Updates," *J. High Speed Networks*, vol. 15, no. 1/2006, IOS Press, pp. 5–19.
- [2] I. Nai Fovino, S. Di Blasi, and A. Rigoni, "The Role of the DNS in the Secure and Resilient Operation of CIs: The Energy System Example,"

- CRITIS, Lucerne, Switzerland, Sept. 2011.
- [3] ICANN, "Measuring the Health of the Domain Name System," Report of the 2nd Annual Symp. DNS Security, Stability, & Resiliency, Kyoto, Japan, 2010.
- [4] ICANN, GCSEC, DNS-OARC, "Final Report of the 3rd Global DNS Stability, Security and Resiliency Symposium," 2011, Rome, Italy.
- [6] R. Liston, S. Srinivasan, and E. Zegura, "Diversity in DNS Performance Measures," *Proc. 2nd ACM SIGCOMM Wksp. Internet Measurement*, 2002 ACM, New York, NY, USA, pp. 19–31.
- [5] S. Castro et al., "A Day at the Root of the Internet," *ACM SIGCOMM Comp. Commun.*, 2008, Rev. 38, 5, pp. 41–46.
- [7] B. Ager et al., "Comparing DNS Resolvers in the Wild," *Proc. 10th ACM SIGCOMM Conf. Internet Measurement*, pp. 15–21.
- [8] Global Cyber Security Center (GCSEC), "Measuring the Naming System (MeNSa) Project," <http://www.gcsec.org/activity/research/dns-security-and-stability>.
- [9] X. Chen et al., "Maintaining Strong Cache Consistency for the Domain Name System," *IEEE Trans. Knowledge and Data Engineering*, vol. 19, no. 8, Aug. 2007.
- [10] S. M. Bellovin, "Using the Domain Name System for System Break-ins," *Proc. 5th USENIX UNIX Security Symp.*, Salt Lake City, UT, June 1995.
- [11] P. Vixie, "DNS and BIND Security Issues," *Proc. 5th USENIX UNIX Security Symp.*, Salt Lake City, UT, June 1995.
- [12] J. G. Kagwe and M. Masinde, "Survey on DNS Configurations, Interdependencies, Resilience and Security for *.ke Domains," *Proc. 2nd ACM Symp. Computing for Development*, 2012.
- [13] N. Brownlee, C. Loosley, "Fundamentals of Internet Measurement: A Tutorial," *CMG J. Computer Resource Management*, vol. 102, 2001.

Biographies

EMILIANO CASALICCHIO (emiliano.casalicchio@uniroma2.it), Ph.D., is a researcher in the Department of Civil Engineering and Computer Science of the University of Roma "Tor Vergata." Since 1998, his research has mainly focused on large-scale distributed systems, with a specific emphasis on performance oriented design of algorithms and mechanisms for resource allocation and management. Domains of interest have been distributed web servers, grid, service oriented architectures, cloud systems, as well as critical infrastructure protection. He is the of about 70 international publications, and his research is and has been funded by the Italian Ministry of Research, CNR, ENEA, the European Community, and private companies.

MARCO CASELLI (marco.caselli@gcsec.org) received his B.S. degree from the University of Palermo and his M.S.E. degree from Sapienza University, Rome. He is currently working toward his Ph.D. degree at the University of Twente. His research interests include critical infrastructure protection as well as industrial control systems' security. He has published scientific papers on DNS security in international conferences. Before starting his Ph.D. he worked for GCSEC.

ALESSIO COLETTA (alessio.coletta@gcsec.org) graduated in computer science at Scuola Normale Superiore di Pisa and possesses a background in computer science, specifically related to ICT security from both theoretical and practical points of view. He had research experience at the University of Pisa, Italy. He worked as a scientific officer at the Joint Research Centre of the European Commission, and he currently works at GCSEC in Rome, on research activities about ICT and industrial security, security policies, malware, critical infrastructure protection, digital identity, and incident response teams.