A Distributed DNS Traffic Monitoring System

Luca Deri, Lorenzo Luconi Trombacchi, Maurizio Martinelli, Daniele Vannozzi
IIT/CNR
Via Giuseppe Moruzzi 1
56124 Pisa, Italy.
{ luca.deri, lorenzo.luconi, maurizio.martinelli, daniele.vannozzi}@iit.cnr.it

Abstract— **The domain name system (DNS) is a complex distributed database on which most Internet services rely on. As its monitoring is critical, it is necessary to continuously monitor DNS traffic for identifying anomalies, measuring performance, and generating usage statistics.**

This paper describes the design and implementation of a distributed realtime DNS monitoring system, that is able to monitor the authoritative name servers of the .it country code Top Level Domain (ccTLD). In addition to the production of usage records, it is able to understand trends, characterize economical relationships, and also track suspicious activities.

Keywords-component; Domain name system, traffic measurement.

I. Introduction and Motivation

The domain name system (DNS) is a distributed database system that allows numeric IP addresses used in the Internet protocol suite to be associated with human-readable names. The DNS structure is organized as an inverted tree with the root at the top. Each node in the tree has a text label which identifies the node relative to its parent. Each node (or domain) can be further divided into additional partitions, originating in this case a new subtree (or subdomain). A Top Level Domain (TLD) is a "first level" domain, so it is a child of the root. Management of TLDs is delegated by the Internet Corporation for Assigned Names and Numbers (ICANN) [38], which is also in charge of maintaining the root zone. The top-level domain space is mainly organized in country code Top Level Domains (ccTLDs), un/sponsored TLDs (e.g. .com, .net, .travel), and generic Top Level Domains (gTLDs). National domains are conventionally specified using the two-letter ISO 3166-1 country code, and are known as ccTLDs [1]. The DNS protocol [2] is based on the client/server paradigm. A DNS server stores DNS records for a set of domains for which it is authoritative (i.e. responsible), and answers to database queries that have been performed using the DNS protocol. Every zone has a configured set of DNS authoritative servers. The client-side of the DNS is called resolver, and it is responsible for translating a domain name into an IP address or vice-versa. The address resolution mechanism is a sequence of queries used to resolve an address starting with the top level domain label. Using a file that contains the list of known root servers, the resolver first contacts a root name server, in order to obtain the address of one of the DNS servers authoritative for the TLD. Then it queries the obtained TLD server in order to obtain the address of the server authoritative for the second-level domain. This sequence is repeated until the address is resolved. This means that in order to resolve www.xxx.it, unless a similar query was issued and cached previously, the resolver needs to query one of the name servers authoritative for the ccTLD in order to know the name server authoritative for the xxx.it domain. The consequence is that ccTLD servers will be involved in all address resolutions, and thus observe all queries for such country code.

A. Related Work

As the DNS is a complex distributed database on which several Internet services rely on, its monitoring is a crucial activity that has attracted the interest of the research community since long time [14]. Tool such as dnstop, dsc [4] and TreeTop [7] can be used to analyze DNS traffic, and also create reports based on the observed traffic. Others tools such as Nagios [17] and SmokePing [18] can be effectively used to detect name server failures, as well as monitor DNS response time and jitter. Security [8], performance [9] and traffic visualization [5] are other areas where research on DNS is currently focusing. Since the rise of the Internet [12], DNS is also appealing for companies that are using it for various reasons not immediately related to its governance [11], including traffic redirection for non existing domains, Internet user profiling [6] and bad ISP practices that use the naming service for increasing their profits and perhaps resell information about DNS queries performed by users [19].

B. Motivation

The authors of this paper are working at the Institute of Informatics and Telematics of the Italian National Research Council of Pisa (IIT-CNR) which the .it ccTLD. In the past few years IIT-CNR has started a research project [15] [16] focusing on the design and implementation of a passive DNS monitoring system, whose aim is to analyze DNS traffic in order to understand Internet users trends and interests, and also track anomalous traffic pattern behaviors (e.g. DoS attempts and DNS attacks). Analysis of DNS traffic is a widespread activity, as this is one of the core protocols on which the Internet is relying. Nevertheless an area on which both the research community and TLD's seems not to have focused yet, is the analysis of DNS traffic for understanding the evolution and trends of Internet users, similar to web search and traffic reports such as Google Zeitgeist [3] and Akamai State of the Internet [10]. Furthermore monitoring DNS activities allowed use to analyze relatively little traffic (the .it DNSs serve in total about 7 million requests/hour)

978-1-4577-1379-8/12/$26.00 ©2012 IEEE

when compared to complex application-protocols probes that instead need to decode a much larger traffic volume (not to mention that they are unable to analyze encrypted traffic) that needs to be diverted to probes by using network taps or span ports.

This has been the motivation for this work. Namely, we wanted to create a simple yet effective country-wide distributed monitoring platform able to passively monitor DNS traffic for the purpose of understanding trends and interests of a country, geo-locate Internet users, areas of digital-divide, as well suspicious activities [25, 26]. As we maintain the .it ccTLD, measurement results as well detected anomalies are used to both improve the DNS infrastructure and inform domain registrars of the detected issues.

The used methodology required both the analysis of DNS packet payload and displacement of various software monitoring probes at the .it name servers. This has been done in order to have a comprehensive view of all queries performed for the .it domain. Results have been matched against the information records stored in the .it domain database. The outcome of this project, is a novel approach at DNS traffic analysis that is not limiting its scope to traffic volume and query type as most tool do, but also tries to obtain more detailed information about the evolution of a country, interesting domains and its trends (a.k.a. Zeitgeist). Although this work has been validated and deployed on .it name servers, it is general enough to be applied to other contexts not limited to ccTLDs, but also to ISPs and large companies.

The rest of the paper is organized as follows. Section 2 describes the used measurement methodology and monitoring architecture. Section 3 presents the main results obtained while monitoring traffic. Section 4 highlights some open issues, future work items and extensions for the measurement architecture described on this paper.

II. DNS Monitoring Architecture

The DNS is a distributed architecture. In order to distribute the load across name servers, the DNS protocol specifies that clients have to select one of them from the list of authoritative name servers. Modern DNS implementations base their choice on various metrics (e.g. response and round-trip time) and no longer pick a random name server, but rather use the one that they believe is best. Moreover, the use of DNS anycast servers [20, 28] makes impossible to obtain trustworthy results when monitoring only one DNS server for a given domain. Thus, as we planned to monitor a whole ccTLD, we decided that monitoring all ccTLD name servers was the only way to validate these claims. At this stage, we decided to monitor both .it unicast and anycast servers, located at some major Italian Internet eXchange Points (IXPs), in addition to the master server placed in the .it Registry premises. In order not to put constraints on specific DNS implementations and create a replicable architecture, we decided to extend nProbe [21], an open-source NetFlow/IPFIX [13] passive monitoring probe, by developing a plugin for passive DNS traffic analysis. Monitoring systems have been deployed at the main national IXP (Internet Exchange Point) where .it DNS servers are located. This has allowed us to analyze all DNS traffic directed towards unicast servers, and also monitor two anycast servers located in northern (Milan) and center (Rome) Italy.

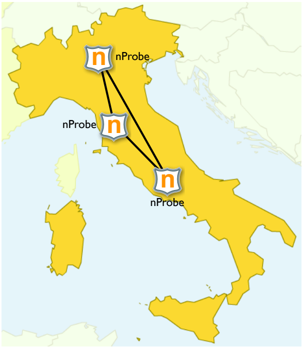

Figure 1. .it DNS Monitoring Infrastructure

nProbe can both export DNS analysis information to a remote NetFlow/IPFIX collector, and dump it to log files on the host where nProbe is running. A typical log entryproduced by nProbe, that can also be exported via NetFlow/IPFIX, has the following format:

```
1302192056.149|A.B.C.D|XXXXX|US|Sunnyvale|194.119.192.34|itgeo.nic.it|0|NOERROR|0|1|A|22165||r.dns.it;dns.nic.it;ns2.nic.it;c.dns.it;nameserver.cnr.it;itgeo.mix-it.net
```

and it contains (sensitive data has been obfuscated on this example to preserve privacy):

- IP address (A.B.C.D), autonomous system (XXXXX), location (US, Sunnyvale) of the DNS client.
- IP address (194.119.192.34) of the name server.
- Record request (itgeo.nic.it), type (A), response type /code (NOERROR, 0) and transaction Id (22165).
- List of DNS authoritative name servers for this request.

As the reader can see from this log excerpt, nProbe does not focus on reporting information about the query (e.g. the numeric IP address of itgeo.nic.it), but rather on other information such as the client address, query response, and authoritative name servers. The latter information is important because being nProbe deployed at a ccTLD, the name server is unable to resolve the address as it is not authoritative for most domains, but rather it can point the client to the authoritative name servers for the requested domain. Additional information such as response time is also measured by nProbe, and it can be used for monitoring name server performance.

It is worth to remark that the implementation of the DNS plugin in nProbe has requested some architectural changes in the probe. In fact, the flow key that is usually defined as a tuple (<VLAN>, <IP src/dst>, <Port src/dst>, <Protocol>, <TOS>) has been extended with a new field that is the DNS transaction Id. This extension has been required as multiple record requests can share the same original flow key. In addition we have defined new flow elements for NetFlow/IPFIX for exporting information about DNS queries/responses. They include the domain to resolve, the query id and type, the return code and number of answer, as well the list of authoritative name servers for the searched domain. To the best of author's knowledge, advanced DNS support in a NetFlow/IPFIX probe is a novel contribution of this research work, not present on other probes.

Although NetFlow/IPFIX probes are often deployed on a (semi-)centralized architecture where a central collector receives flows from various probes, we decided to use nProbe-generated logs instead of flows. Beside privacy motivations that could have been solved using a secure channel for delivering flow to a central collector, the main motivation behind this choice is the probe placement inside an IXP peering network. In fact, IXPs charge peers using various criteria, including the volume of intra-peer traffic. This means that DNS queries are not part of this traffic, whereas flows delivered towards the core of the ccTLD network are, thus increasing the yearly peering contract costs. For this reason, we compute statistics on the node where the probe is active, while delivering the results (and not the raw data flows) to the central collector in order to produce an aggregated view of the ccTLD DNS traffic. As real time monitoring is not a requirement, we aggregate results with one hour granularity by parsing log files produced by the probe on the past hour.

The result of this aggregation is dumped with a great level of detail on an hourly database, and also with less detail on a daily database. The difference between these two databases is that the hourly database also contains the raw data that can be used to drill-down whenever necessary, whereas the daily database only contains aggregated data.

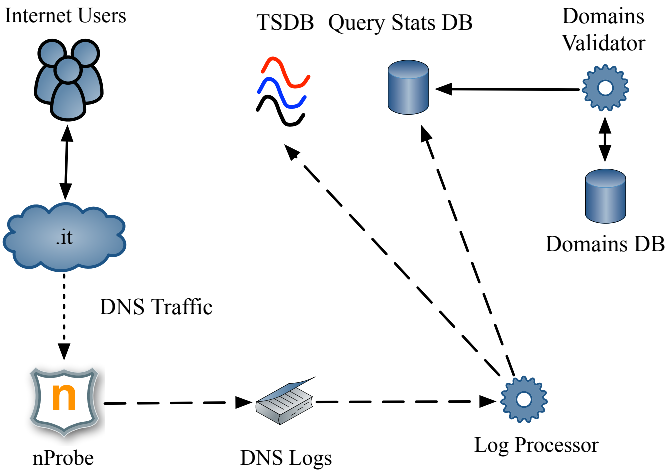

Figure 2. DNS Traffic Processing Architecture

Currently the log processing application is written in python and it dumps data into a MySQL database. We decided not to instruct nProbe to dump data directly into the database as we do not want to risk loosing records due to database slowdown in case of traffic spikes. Please note that statistics of existing domains are limited by the number of registered .it domains (at the time of writing over 2.2 million). Instead, statistics on non-existing domains can have a much higher cardinality, as it is not unusual to detect clients which scan the .it namespace (probably) aiming at creating a map of the registered domains.
For each Internet domain we produce a set of time series for tracking the number of queries according to the origin AS (Autonomous System), and well time series of non-existent domains (NXDOMAIN) grouped per source AS. In total the number of time series we need to maintain is about 25 millions. In early system prototypes, we have used the popular RRD database [23] for storing time series but we have faced major performance issues when updating million time series. This is because RRD is a file-based database, where each file can store just a few time series. The consequence is that at each update interval we had to open/update/close several million RRD files causing a significant load on the server. Replacing the RRD database with a No-SQL key-value database such as Redis [23] improved the performance significantly, but it was not a scalable solution as the database speed derived from the fact that all data was kept in memory. Using it would have required us to use a server with several tenth of GB making this solution impractical.

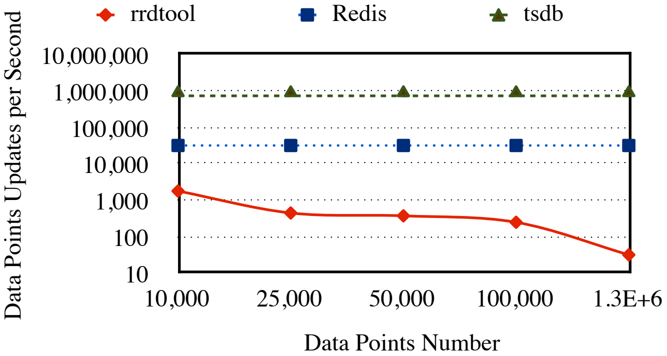

Figure 3. Time series update processing speed (logarithmic scale).

For this reason we have developed a new type of database named TSDB (Time Series DataBase) [24] able to outperform of two order of magnitudes the Redis database. In TSDB time series are maintained on disk on compressed format and decompressed on the fly when necessary. The use of memory allows the tool to perform almost one million updates/sec enabling us to perform all needed updates at each time interval within a minute, requirement for near-realtime traffic monitoring not achievable with the other tools.

As shown in figure 2, in addition to live traffic processing, we have created another tool that for each registered .it Internet domain, performs some periodic checks including:

- Consistency check of registered DNS records (e.g. DNS zone check).
- Check if a domain is operational or if it is just a stub/placeholder.
- Identify common DNS configuration errors and non-optimal configurations.

Beside gathering statistical data, one of the goal of our project being these checks run by the .it ccTLD, is to report back to domain registrar all these issues so that they can improve/fix them on DNS records. This is because we believe that addressing these issues can help improving the quality of the .it DNS as well reduce problems due to poor configuration that might lead to domain unreachability.
In order to evaluate the DNS experience, we have also

mapped DNS requests to AS path (i.e. the sequence of AS that have been traversed by a given DNS client) in order to understand what are the top ASs that are traversed by DNS clients. In order to do that the nProbe BGP plugin has been used: such plugin acts as a BGP daemon to which our border gateway connects by delivering BGP updates.

The identification of top traversed ASs is very important as it enable us to identify locations where future DNS servers for the .it should be placed. This is because placing DNSs where most traffic is flowing allows us to reduce the amount of requests received by national DNS servers as well reduce the DNS response time as they are placed closed to the source of requests. Measurement reports are accessible via a web 2.0 interface that can display both information about a specific observation point, and aggregated for all monitored name servers.

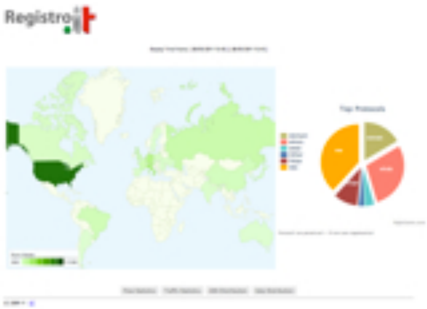

Figure 4. Web Monitoring Console

All monitoring data have been geolocated, thus it is possible to know where requests have been originated. This is important as mapping domain names to a typology (e.g. news, sport, culture, food) it is possible to create a map of interests. For instance we know what kind of information is mostly accessed by US internet users, information that can help the national government to promote specific products on selected countries.

III. MEASUREMENT RESULTS

Before describing some measurement results, it is worth to note that the local law forbids techniques that divert unsolicited traffic towards specific monitoring systems. Fortunately this restriction does not apply to DNS traffic as resolvers contact our servers where the monitoring system is deployed.

The first, perhaps obvious to the reader, result that we validated is that well-behaved DNS client really perform queries in round-robin mode as we basically see the same amount of queries from all the three monitoring systems. This does not apply to DNS scanners (i.e. hosts that issue a large number of requests likely for creating a map of registered Internet domains). In fact thanks to our distributed monitoring platform, we have realized that very often those scanners use a specific DNS server to which issue requests. This means that in order to detect scanners, we need to monitor all DNS servers and not just a few; this as scanners do not follow the DNS round-robin.

In order to distinguish between a domain scanner and domain that is invalid/misspelled we have created a tool that:

- Filters out invalid requests (e.g. _ldap._tcp $e52fced9-8106-48a0-9c86-69c5d32d8c92.domains._msdcs.abcd.it) or human mistakes (e.g.we have found many DNS requests containing email addresses, probably because people type email addresses in web browsers).
- Analyzes NXDOMAIN responses and tries to match a request with a registered domain using the Levenshtein distance [35] that allows the tool to check if a NXDOMAIN is likely a misspelled name rather than a completely non existing domain. In our records we have verified that about 4% of non-misspelled NXDOMAIN are then registered within few weeks.

DNS scanners can be easily identified with our system, as they do not often use sophisticated algorithms for hiding them beside a slow scan approach, just to avoid them to be detected using traffic-based tools:

- The number of NXDOMAIN per searched domain is always one (i.e. they do not repeat the query for a given domain in case of NXDOMAIN) whereas statistically even a DNS server of a large ISP which issues many NXDOMAIN request, some invalid queries are repeated.
- Even if NXDOMAIN queries come from various hosts for reducing the number of queries per host, grouping these hosts per AS ease the scanner detection process.
- A DNS scanner when grouped per AS, typically issues more NXDOMAIN queries than the sum of all other DNS clients.

DNS scanners are not too frequent, as usually we identify just a few scanners per month but none of them has been able to create any problem to the DNS infrastructure also thanks to the anycast servers. The reason why we are tracking and blocking them is that beside putting load on DNS servers, they often create maps of the registered domains sold on the Internet that might be used by spammers to send unsolicited emails. Note that monitoring DNS traffic at ccTLD servers is a unique position as it allows us to monitor NXDOMAIN responses. This is because when a domain is resolved, the DNS resolver first contacts the .it ccTLD for obtaining the DNS authoritative servers for the searched domain, thus making our monitoring positions privileged. On the other hand we can observe attacks only directed to the .it domain servers, as our systems have no ability to monitor individual .it domain DNS server that in our knowledge are those mostly subject to attacks.

In addition to detection and mitigation of these type of issues, we also analyzed how the DNS infrastructure is robust and immune to attacks. The methodology we used is the analysis of DNS servers for each domain registration. In particular we have created a quick traceroute tool (i.e. able to perform a traceroute in a limited amount of time) derived from the tracepath tool, we used to create the list of hops necessary for reaching the configured DNS servers from the .it network. We define as overlapping DNS servers, two or more DNS servers that are connected to the same router as last hop. For instance if we analyze domain X with configured DNSa and DNSb, these DNS are overlapping if the traceroute for these servers has as last hop the same router IP. Using this methodology we have realized that 54% of .it domains have overlapping servers, and out of them, 95% of domains have critical overlapping i.e. at least half of the DNS servers configured for a specific domain overlap. We reported this information to domain registrars in order to improve the quality of the .it DNS and thus make it more resilient to attacks and network disruption.

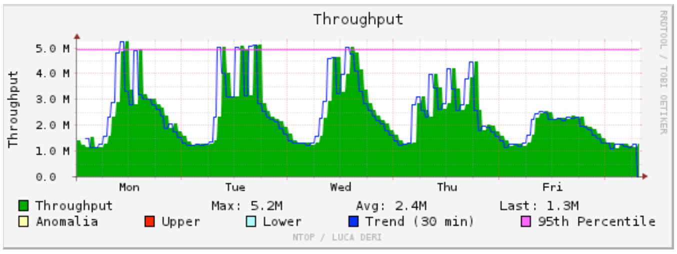

Figure 5. Typical Daily Traffic per Authoritative .it DNS Server

In average each authoritative DNS server for .it receives between 1 and 6 Mbps of DNS traffic that boils down to 3 to 7 M queries/hour. NXDOMAIN queries account for 5% of requests, whereas most queries are for A (~70%) and MX (~15%) records. The static DNS validation tool reports that:

- 2 % of registered .it domains are placeholders.
- 98 % of registered domains have a www.domainname.it A record with a web site configured.
- 7% of registered domains have no MX record whereas most of these domains have a web server configured. This is because most of these sites (e.g. apple.it) have www.domain.it with email configured at domainname.com.

We have collected from public web sites the list of public DNS servers. Looking at query records, almost no queries come from such public servers. Whereas we see that aggregating DNS queries per AS, DNS queries coming from Google and OpenDNS ASs (i.e. two popular public DNS servers) rank 2nd on this list; note that we do not what portion of these queries come from public DNS servers. Despite the fact that public reports from TLD registries are usually based on the number of registered domains, we have decided to investigate what percentage of these domains are active (i.e. we observe queries for the domain). In average, every day we observe queries for about half of the registered domains. Our system is running since about two months, thus we plan to continuously monitor this metric in order to understand whether it changes over time, in particular for domains associated with seasonal events.

Another metric we are measuring is the (potential) economical value of domain registrars. In particular based on domain responses and not on domain registration records, we aggregate the number of domains that are resolved by specific DNS. This information is obtained from DNS responses, where the authoritative DNSs for the queried domains are returned. The HHI (Herfindahl-Hirschman Index) index [27], a widespread indicator of competition among companies, has been used to evaluate the concentration ratio of registrar. The result has confirmed that even if in Italy the first 10 registrars do register more than 60% of Italian domains (HHI national index of 1389 making this a not very competitive market), most of those registrars do not host the active domains, thus making the hosting market much more competitive than the registry market (the HHI index based on name servers is less than 200).

In order to classify DNS requests not just using widespread criteria such as AS/network/client, we have classified Internet sites according to the information they offer. We have used various techniques for tagging the content of .it internet domain. At the moment we have divided web sites in several macro-categories and for each category we daily report about domain queries. For privacy reasons, we are not allowed to publish reports about specific domains but it is possible to disclose a report about the number of queries grouped by categories as shown in the above figure. We are currently studying how specific trends (e.g. blogging vs. newspapers, country house vs. hotels) change overtime (e.g. during summer). One of the results we would like to achieve in the long run, is find out whether trends on specific topics we observe on the Internet also match the statistics which are published every year by official statistical companies and institutions. The following table shows the current query distribution divided per macro categories.

TABLE I
DNS QUERIES GROUPER PER MACRO-CATEGORIES

Category	Query Percentage
News	48.0%
Mobile (Phones)	29.5%
Computing	14.3%
Italian News	11%
Bet	10.4%
Tourism	10.4%
Sport (no Soccer)	5.20%
Soccer	4.8%
Internet Games	4.71%
Internet (Generic)	4.42%
Sex	2.9%
Politics	1.5%

Part of the effort of every ccTLD organization, is to support IPv6 in addition to the IPv4 protocol. We have created a tool that analyzes the domain registrations and realized that only as little as 5% of domains support IPv6 queries. Looking at DNS traffic, we observed only 2% of queries for IPv6 records (AAAA). Unfortunately we cannot compare this data with other ccTLDs in order to understand if the diffusion of IPv6 we observed for .it is similar to other ccTLDs.

IV. OPEN ISSUES AND FUTURE WORK

As we are monitoring only .it domain names, our work does not take into account all italian domains that have been registered under another TLD such as .com and .net. Although we acknowledge that this can be a limitation of our methodology, the local law does not allow us to place probes across the country to collect DNS traffic statistics on requests that are not sent to our servers.

We are extending our platform to collect comprehensive information about the use of anycast in DNS and produce best practices for servers deployment.

We are planning to refine the mechanism used to tag a site with respect to its content. In particular we are considering to extend our system with a web crawler that could attempt to visit www.<registered domain>.it (note that not all registered domains have a web site) and based on the content of the site, tag it automatically according to the categories we defined. Doing this, we will greatly complement the information currently gathered by our system.

V. Conclusion

This paper has presented a distributed platform for law-compliant, permanent DNS traffic monitoring. In addition to monitoring network related indicators such as number of queries and their distribution across ASs, we also wanted to understand the trends and interests of a country by analyzing queries to ccTLD domain servers. Novel contributions of this paper are manyfold and include:

- The creation of a time-series compressed database named TSDB, that allows efficient data update and retrieval.
- Definition of NetFlow/IPFIX extensions for DNS traffic monitoring that have been implemented in nProbe, our open-source NetFlow probe.
- Creation of a permanent, scalable, distributed monitoring platform for near-realtime analysis of DNS traffic able to monitor over 100 M queries/day per node.

Acknowledgment

Our thanks to Italian Internet domain Registry (nic.it) that has greatly supported this research work.

References

[1] J. Postel, Domain Name System Structure and Delegation, RFC 1591, 1994.

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

[3] Google Inc., Google Zeitgeist 2010, http://www.google.com/zeitgeist, 2010

[4] The Measurement Factory, dnstop and dsc tools, http://dns.measurement-factory.com/tools/, 2006.

[5] P. Ren et al., Visualizing DNS Traffic, Proc. of VizSEC'06, 2006.

[6] C. L. Simon, Launching the DNS war: dot-com Privatization and the Rise of Global Internet Governance, PhD Thesis, University of Miami, 2006.

[7] D. Plonka and P. Barford, Context-aware Clustering of DNS Query Traffic, Proc. of IMC'08 Conference, 2008.

[8] B. Zdrnja, Security Monitoring of DNS Traffic, University of Auckland, May 2006.

[9] K. Sato et al., DNS Query Sent by Heavy Users and DNS Prefetch Effect, Nanog 51, 2009.

[10] Akamai Technologies, State of the Internet: Q4 2010 Report, http://www.akamai.com/stateoftheinternet/, 2010.

[11] P. Vixie, What DNS is Not, ACM Queue, Issue 52, November 2009.

[12] B. Wellman and C. Haythornthwaite, The Internet in Everyday Life, Wiley-Blackwell, 2002.

[13] B. Claise, Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information, RFC 5101, 2008.

[14] D. Wessels and M. Fomenkov, Wow, That's a Lot of Packets, Technical Report, 2003.

[15] A. Bonaccorsi, M.Martinelli, C.Rossi, I. Serrecchia, and D. Vannozzi, Internet Diffusion and Internet DomainsL Looking for a New Metric: the Case of Registrations by Italian Individuals, Proc. of ICWI Workshop, 2002.

[16] M. Martinelli, C. Rossi, I. Serrecchia, and A.D. Soldato, Measuring Internet Diffusion in Italy, Proc. of IFIP TC6/ WG6.4 Workshop on Internet Technologies, 2002.

[17] D Josephsen, Building a Monitoring Infrastructure with Nagios, Prentice Hall, ISBN 0132236931, 2007.

[18] T. Oetiker, Monitoring your IT Gear: the MRTG story, IT Professional, Vol. 3, Issue 6, 2001.

[19] M. Michlick, Verisign to Profit from Rootserver Data?, Domain Name News Magazine, October 2007.

[20] T. Hardie, Distributing Authoritative Name Servers via Shared Unicast Addresses, RFC 3258, 2002.

[21] L. Deri, nProbe: an Open Source NetFlow Probe for Gigabit Networks, Proc. of Terena TNC Conference, 2003.

[22] T. Oetiker, RRDtool: Round Robin Database Tool, http://oss.oetiker.ch/rrdtool/.

[23] Jeremy Zawodny, Redis: Lightweight key/value Store That Goes the Extra Mile, Linux Magazine, August 31, 2009.

[24] L. Deri, S. Mainardi and F. Fusco, tsdb: A Compressed Database For Time Series, Proceedings of TMA Workshop, March 2012.

[25] M. Antonakakis, D. Dagon, X. Luo, R. Perdisci, W. Lee, and J. Bellmor, A Centralized Monitoring Infrastructure for Improving DNS Security, Proc. 13th RAID Conference, September 2010.

[26] L. Bilge, E. Kirda, C. Kruegel, and M. Balduzzi, EXPOSURE: Finding Malicious Domains Using Passive DNS Analysis. Proc. of 18th NDSS Symposium, February 2011.

[27] A. Hirschman, The Paternity of an Index, The American Economic Review Vol. 54, No. 5, 1964.

[28] S. Sarat, V. Pappas, and A. Terzis, On the Use of Anycast in DNS, Proc. of ICCCN Conference, 2006.