

Enhancing DNS Security using Dynamic Firewalling with Network Sensors

Joao Afonso

Foundation for National Scientific Computing
Lisbon, Portugal
e-mail: joao.afonso@fccn.pt

Pedro Veiga

Department of Informatics
University of Lisbon
Lisbon, Portugal
e-mail: pedro.veiga@di.fc.ul.pt

Abstract— Security problems that plague network services today are increasing at a dramatic pace especially with the continuous improvement of network transmission rates and the total amount of data exchanged. This translates not only into more incidents but also to new types of attacks with network incidents becoming more and more frequent. A significant part of the attacks occur at Top Level Domains (TLD) who have the task of ensuring the correct functioning of Domain Name System (DNS) zones. In this article we discuss a solution developed and tested at FCCN (Foundation for National Scientific Computing), the TLD manager for the .PT domain. The system consists of a series of network sensors that monitor the network in real-time and can dynamically detect, prevent, or limit the scope of the attempted intrusions or other types of attacks to the DNS service, thus improving its global availability.

Keywords—DNS; security; intrusion detection system; real-time; monitoring.

I. INTRODUCTION

DNS is a critical application for the reliable and trustworthy operation of the Internet. DNS servers assume a pivotal role in the normal functioning of Internet Protocol (IP) networks today and any disturbance to their normal operation can have a dramatic impact on the service they provide and on the global Internet. Although based on a small set of basic rules, stored in files, and distributed hierarchically, the DNS service has evolved into a very complex and, at the same time, very vulnerable system [1].

According to recent studies [2], there are nearly 11.7 million public DNS servers on the Internet. It is estimated that nearly 52% of them, due to improper configuration, allow arbitrary queries (thus allowing denial of service attacks or “poisoning” of the cache). About 31.1% of the servers also allow for the transfer of their DNS zones.

There are still nearly 33% of situations where the authoritative nameservers of an area are on the same network, which facilitates Denial of Service (DOS), a frequent attack to the DNS.

Furthermore, the types of attacks targeting the DNS are becoming more sophisticated, making them more difficult to detect and control in real-time. Examples are the attacks by

Fast Flux (ability to quickly move the DNS information about the domain to delay or evade detection) and its recent evolution to Double Flux.

One of these attacks is the conficker [3] worm, first appeared on October 2008, but also known as Code Red, Blaster, Sasser and SQL Slammer.

Every type of computer, using a Microsoft Operating System can potentially be infected. Attempts to estimate the populations of conficker have lead to different figures but all these estimates exceed millions of personal computers. Conficker made use of domain names instead of IP address in order to make its attack networks resilient against detection and takedown.

The ICANN (Internet Corporation for Assigned Names and Numbers) created a list containing the domains that could be used in each TLD in such attacks to simplify the work of identifying attacked domains.

A central aspect of the security system that we propose and have implemented is the ability to statistically collect useful data about network traffic for a DNS resolver and use it to identify classes of harmful traffic to the normal operation of the DNS infrastructure.

In addition to collecting data, the system can take protective actions by detecting trends and patterns in the traffic data that might suggest a new type of attack or simply to record important parameters to help improve the performance of the overall DNS system.

The fact that the DNS is based on an autonomous database, distributed by hierarchy, means that whatever solution we use to monitor, it must respect this topology.

In this paper, we propose a distributed system using a network of sensors, which operate in conjunction with the DNS servers of one or more TLDs, monitoring in real-time the data that passes through them and taking actions when considered adequate.

The ability to perform real-time analysis is crucial in the DNS area since it may be necessary to immediately act in case of abuse or attack, by blocking a particular access and notifying other cooperating sensors on the origin of the problem, since several types of attacks may be directed to other DNS components.

The use of a Firewall solution, whose triggering rules are dynamically generated by the network sensors, is a fundamental component of the system. This way we can filter attacking systems efficiently and return to the initial status when the potential threat has ceased to exist.

With this approach we aim to guarantee an autonomous functioning of the platform without the need of human intervention.

The use of network alarms can also help in monitoring the correct functioning of the whole solution. Special care has been taken to minimize the detection of false positives and false negatives.

The remaining of the paper is structured as follows: Section II provides information regarding related work. Section III introduces our proposed methodology. In section IV, we describe the solution. Section V presents a case study used to validate the solution. In Section VI, the results gathered in the case study are analyzed. Section VII describes the process of evaluation and treatment of false positives and negatives. Finally, Section VIII presents some conclusions and directions for further work.

II. RELATED WORK

One of the first attempts in this area was a tool called sqldjbdns developed by Guenter and Kolar [4]. Their proposal uses a modified version of the traditional BIND [5] working together with a Structured Query Language (SQL) version inside a Relational database management system (RDBMS). For DNS clients, this solution is transparent and there is no difference from classic BIND.

Zdrnja presented a system for Security Monitoring of DNS traffic [6], using network sensors without interfering with the DNS servers to be monitored. This is a transparent solution that does not compromise the high availability needed for the DNS service.

Vixie proposed a DNS traffic capture utility called, DNSCap [7]. This tool is able to produce binary data using pcap format, either on standard output or in successive dump files. The application is similar to tcpdump [8] – command line tool for monitoring network traffic, and has finer grained packet recognition tailored for DNS transactions and protocol options, allowing for instance to see the full DNS message when tcpdump only shows a one-line summary.

Another tool available is DSC - DNS Statistics Collector [9]. DSC is an application for collecting and analyzing statistics from busy DNS servers. Major features include the ability to parse, summarize and search inside DNS queries detail. All data is stored in an SQL database. This tool, can work inside a DNS server or in another server that "captures" bi-directional traffic for a DNS node.

Kristoff also proposed an automated incident response system using BIND query logs [10]. This particular system, besides the common statistical analysis, also provides information regarding the kind of consultations operated. All information is available through the Web based portal. Each security incident can result in port deactivation.

III. METHODOLOGY

A. Architecture

The architecture of the system that we developed aims to improve the security, performance and efficiency of the DNS protocol, removing all unwanted traffic and reinforcing the resilience of a Top Level Domain. We propose an architecture comprising an integrated protection of multiple DNS servers, working together with several network sensors that apply live rules to a dedicated firewall, acting as a traffic shaping element.

Sensors carefully located in the network monitor all the traffic going to the DNS infrastructure, identify potentially harmful traffic using an algorithm that we have developed and tested and use this information to isolate traffic that has been identified as a security threat.

Several networks sensor monitor different parts of the infrastructure and exchange information related to security attacks. In this way, as shown in Fig. 1, it should also be possible to exchange critical security information between the sensors. In addition to an increase in performance, this operation should prevent an attack on a server from a source, identified by another sensor as malicious. This scenario is relevant since some kinds of attacks are directed to several components of the DNS infrastructure.

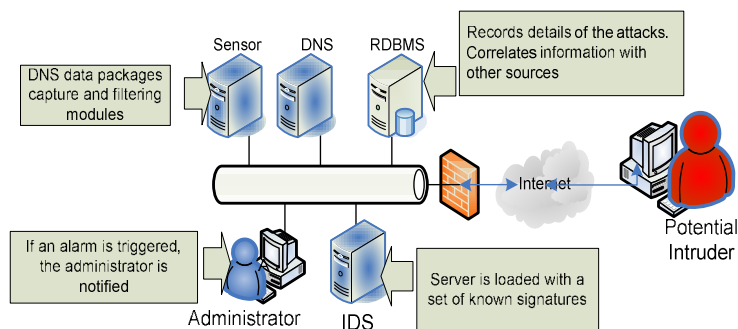


Figure1. Diagram of the desired solution

B. Heuristic

One of the crucial parts of our work is the algorithm to identify traffic potentially harmful to the DNS. In order to implement the stated hypothesis in the architecture and keep the DNS protocol as efficient as possible, it is necessary to apply a heuristic, which in real time, evaluates all the information collected from different sources and applies convenient weights to each component and act accordingly.

The components that we have chosen to have impact in the security incidents of DNS are: the number of occurrences, analysis of type of queries been made, the amount of time between occurrences, the number of probes affected and information reported from intrusion detection systems.

Our system uses the following formula to evaluate a parameter that measures the likelihood of the occurrence of a security incident:

$$f(x) = O \cdot 0,2 + C \cdot 0,2 + G \cdot 0,15 + N \cdot 0,25 + I \cdot 0,20$$

The following factors are considered:

- Occurrences (O) - Represents the number of times (instances) that a given source was blocked weighted according to the factors indicated in Table I.

TABLE I – CONTRIBUTION OF THE NUMBER OF OCCURRENCES OF A SOURCE IN MALICIOUS HEURISTIC

Occurrences	Weight
1	25%
2	50%
3	75%
4 or more	100%

- Analysis (C) - Real-time evaluation of the deviation of the values observed relatively to the average recorded, based on the criteria and weights identified in Table II.

TABLE II – CONTRIBUTION OF EVENTS TYPIFIED AS POTENTIALLY MALICIOUS ACCORDING TO THE HEURISTIC

Event	Weight
Entire zone transfer attempt (AXFR)	100%
Partial transfer zone attempt (IXFR)	50%
Incorrect query volume, 50 to 75% on average per source	75%
Incorrect query volume exceeding 75%	100%
Query volume, up 50%, the average number of access by origin	50%

Note that the estimates apply the moving average, for the determination of reference values, given the continuous collection of data.

- Time between occurrences (G) - time since last occurrence of a given event, distributed with the weights indicated below.

TABLE III – WEIGHT OF DIFFERENT TIME BETWEEN EACH OCCURRENCE

Time	Weight
Less than 1 Minute	100%
Less than 1 Hour	75%
Less than 1 Day	50%
Less than 1 Week	25%

- Incidence (N) - Number of probes that report blocks in the same source.
For the calculation, we use the expression:

$$N = \frac{1}{\#Total_Sensors - \#Sensors_Attacked}$$

- Intrusion Detection Systems (I) - We considered the use of the Snort platform, a free tool that recognizes a large number of signatures of security incidents related to the DNS service.

TABLE IV – INTERCONNECTION WITH TEMPORAL DATA GATHERED FROM INTRUSION DETECTION SYSTEMS

Metric: Common Vulnerability Scoring System (CVSS)	Weight
Low level	34%
Middle level	67%
High level	100%

An activation of a rule in the Firewall will require:

1. The formula above has a value equal to or greater than 0.25;
2. The combination of two or more criteria of the formula.

Exception: when receiving information from all the other sensors, in which case a single criterium is sufficient;

3. That the source is not in the White List, that contains privileged sources that should never be blocked.

In this way we avoid compromising the Internet service, considering the key role played by DNS, the White List protects key addresses from being blocked in case of false positives events.

This list is created from a record of trusted sources, allowing all addresses listed here to be protected from being added to the Firewall rules.

One example is the list of internal addresses, and the DNS servers of ISPs.

On the opposite side, the removal of a rule in the firewall will require the following assumptions to occur simultaneously:

1. Exceeded the quarantine period, based on the parameters in use;
2. The expression of activation (heuristic) no longer checks the referenced source.

IV. PROPOSED SOLUTION

A. Diagram

As shown in Fig. 2, this solution is based on a network of sensor engines that analyze all traffic flowing into the DNS server in the form of valid or invalid queries, process the information received from other probes and issue restrictions for specific network addresses. In case an abnormal behavior is detected or there is suspicious behavior from a certain network address, it will be blocked in the firewall and the other probes notified so they can act accordingly. The system can also calculate the response time for each operation to evaluate the performance of the server.

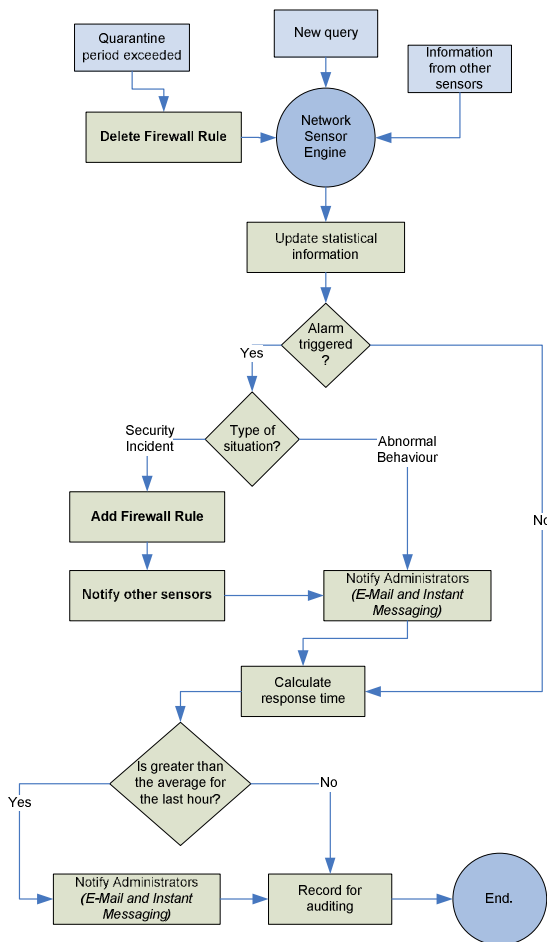


Figure 2. Block Diagram of proposed solution

For each rule inserted in the sensor firewall, there will be a period of quarantine and, at the end of this time, the sensor will evaluate the behavior of that source, to decide the needed to remove or keep that rule enabled, as shown in Fig. 2.

B. Network data flow

According to our design, all data that flows through the probe heading for the DNS server is treated according to a standard set of global firewall rules, followed by specific local rules regarding to the addresses that are being blocked in real time. The queries are then delivered to the parser to be analyzed and stored in the RDBMS. At the top is the system of alarms and the Web portal.

All information collected is stored in a database implemented in MySQL [11]. Taking into consideration the need to optimize the performance of the queries and to reduce the volume of information stored, the data is divided into a number of different tables.

The conversion of the IP address of source and destination (DNS server) into an integer format, has allowed for much more efficient data storage, and a significant improvement in the overall performance of the solution.

The information regarding all queries made, is stored daily into a log, and kept available during the next 30 days.

Two tables containing the set of rules that are dynamically applied – add or removed, based on situations that have been triggered - control the correct operation of the firewall. For auditing purposes every action is registered.

The information required for auditing and statistical tasks never expires.

C. Statistical analysis and performance evaluation

The statistical information collected and stored in the database has a significant amount of detail. It is possible, for example, to calculate, for each sensor, the evolution of queries per unit of time (hour, day, etc) badly formatted requests, DNS queries of rare types and determine the sources that produce the larger number of consultations. It is also possible to see the standard deviation of a given measure so we can relate it to that is seen with the other hits [14].

The performance of the DNS protocol responses is permanently measured, regarding the response time per request. Data is constantly registered and an alarm is raised in case normal response times are exceeded.

V. CASE STUDY

Our proposal have been under development since September 2006 at FCCN – who has the responsibility to manage, register and maintain the domains under the .PT TLD.

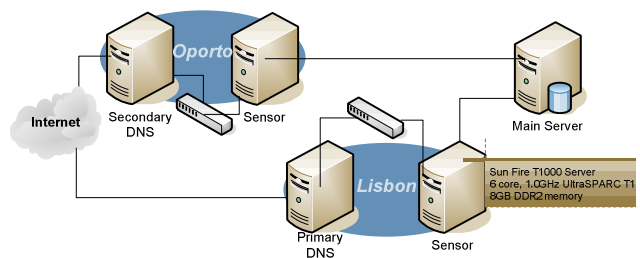


Figure 3. Working Sensors coupled with DNS servers

At present time, there are two sensors running, one at the primary DNS and another working together with a secondary DNS server.

The network analyzer is tshark [15], and the firewall used is IPFilter [12]. The real time parser was programmed in Java, collecting the information received from the tshark. The Web server is running Apache with PHP.

Regarding the Xmpp server [13], we choose the Jive messenger platform.

All modules are integrated together.

The entire sensor solution, as described above, as well as the web platform we developed is on-line from the 1st of January 2007, and the data from the various agents is being collected from the 10th of May 2008.

VI. RESULTS

We present here the results of 12 months of data collection (between 1st of May 2010 and 1st May 2011). The Average number of requests to the primary DNS server is up to 14,459,356 per day (167 per sec.).

The performance of the data analysis program is above 1240 requests processed per sec. (filtered, validated and inserted in the database).

Using the data collected by the sensors, during this time period, we were able to:

- Collect useful statistical information. E.g., daily statistics by type of DNS protocol registers accessed (Fig. 4).

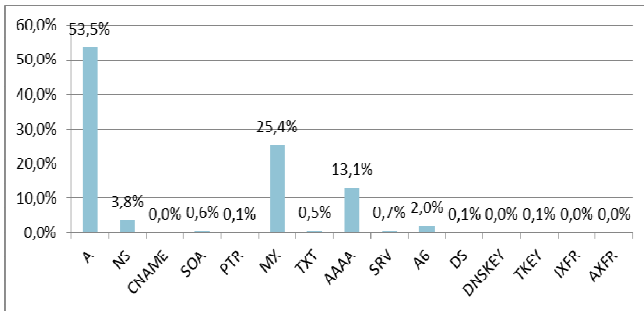


Figure 4. Statistical analysis by type of records accessed

- Detect examples of abnormal use (that are not security incidents). For example we were able to detect that a given IP was using the primary .PT DNS server as location resolver. The number of queries made was excessive when compared with the average value per source, reaching values close to some Internet Service Providers that operate under the .PT domain.
- Detect situations of abuse, including denial of service attacks, with the execution of massive queries. In last

12 months of analysis there are 17 DOS attacks triggered.

They were instantly blocked, and addresses placed in quarantine (Table V).

TABLE V. EXAMPLES WHEN THE SENSOR DETECTED SITUATIONS THAT REQUIRED THE FIREWALL RULES TO CHANGE.

Source Address	Date / Time	Operation	Sensor
xx.xx.200.35	2011-07-12 01:03:12	Add rule	xx.xx.21.62
xx.xx.17.212	2011-07-12 03:25:19	Remove rule	xx.xx.32.63
xx.xx.117.51	2011-07-13 01:23:17	Add rule	xx.xx.21.62
xx.xx.94.139	2011-07-13 02:27:11	Add rule	xx.xx.31.63
xx.xx.13.231	2011-07-14 03:42:52	Remove rule	xx.xx.21.62

- Improve DNS protocol performance repairing situations of inefficient parameterization of the DNS server.

On the DNS server side, considering the capacity of the probe to determine the processing time for each consultation, it is possible to detect cases of excessive delay, which was later confirmed to coincide with of moments of zone update.

Considering the daily progress of DNS queries, before and after applying shaping heuristic to the protocol we obtain an improvement between values of 5.3% (minimum) and 19.4% (maximum).

VII. REDUCING FALSE POSITIVES AND FALSE NEGATIVES

For the treatment of false positives and negatives, it is important to evaluate the results obtained with each of the components in separate – from intrusion detection solution and the solution proposed in the hypothesis presented here, and after evaluating the results that accrue from the application of heuristic indicated above.

It is inevitable given the occurrence of false positives - Type I error (when events are identified as security incidents that do not correspond to real situations) as well as false negatives – Type II error (when there are security incidents that are not detected by the solution in use).

And such an occurrence is the case both in traditional IDS solution - in this case the SNORT, as in the methodology implemented here in the form of prototype with the DNS server at . PT (Country-code top-level domain designed for Portugal).

The way to mitigate such situations is to combine the values obtained by both mechanisms, thus seeking an end result, as close to reality as possible, as identified in Table 6.

TABLE VI – ORIGINAL DECISION VS MID-TERM REVIEW

<i>Original decision (intrusion detection system)</i>	<i>Mid-term review</i>
True positive	Correct
False Positive	Type I error
True Negative	Correct
False Negative	Type II error

TABLE VII – DECISION WITH VS WITHOUT THE METHODOLOGY

<i>Original decision</i>	<i>Decision without the methodology</i>	<i>Decision with the methodology</i>
True positive	Accept	Correct
	Reject	Type II Error
False Positive	Accept	Type I Error
	Reject	Correct
True Negative	Accept	Correct
	Reject	Type I Error
False Negative	Accept	Type II Error
	Reject	Correct

The analysis of the data obtained for a period of 12 months, demonstrated a reduction in 21% of false positives identified by the IDS solution when operated in isolation. The identification of type II errors is now possible in 8% of cases, which previously went unnoticed.

Given the mutual dependence on heuristic established between the two data sources - IDS and the proposed platform, we cannot verify the occurrence of false positives as a result of implementation of the method proposed here. It is, however, possible that false negatives that arise due to situations considered in the IDS solution, can be canceled later in the final assessment.

In any case, considering the relevance of the DNS service on the proper operation of the Internet, the disruption caused by a blockage of a given origin and the fact that there is no absolute certainty that this is a clear case of a security incident, makes it preferable, in this particular case, to allow the occurrence of a controlled number of false negatives

VIII. CONCLUSION AND FUTURE WORK

The solution presented here builds upon the existing solutions that collect statistical information regarding DNS services, by adding the ability to detect and control security incidents in real time. It also adds the advantage of operating in a distributed way, allowing the exchange of information between cooperating probes, and the reinforcement of its own security, even before it is threatened.

In order to improve the ability to detect abnormal behaviors patterns, the solution can be evaluated using a data mining approach.

Currently, the solution presented does not allow the processing of addresses in the IPv6 format. The technical aspects that led to this situation are linked to the need to optimize the performance of the data recorder application making it possible to store the data from all consultations. Nevertheless, all queries made to IPv6 addresses are contained in this solution (AAAA types).

We are also working on extending the data correlation capabilities of the system by adding information collected from other sources (intrusion detection systems for instance). We anticipate that this could be a valuable approach to reduce considerably the number of false positives and negatives [16].

REFERENCES

- [1] P. Vixie, "DNS Complexity", ACM Queue vol. 5, no. 3, April 2007.
- [2] D. Wessels, "A Recent DNS Survey", DNS-OARC, November 2007.
- [3] Dave Piscitello, "Conficker Summary and Review", ICANN, May 2010.
- [4] SQLDNS website, [http://home.tiscali.cz:8080/~cz210552/sqldns.html]. Last accessed on 27 July 2011.
- [5] BIND website, [http://www.isc.org/products/BIND]. Last accessed on 27 July 2011.
- [6] Bojan Zdrnja, "Security Monitoring of DNS traffic", May 2006.
- [7] Paul Vixie, D. Wessels, "DNSCAP – DNS traffic capture utility", CAIDA Workshop, July 2007.
- [8] Duane Wessels, "Whats New with DSC", DNS-OARC, November 2007.
- [9] Lawrence Berkeley National Laboratory. Tcpcdump website http://www.tcpdump.org.
- [10] John Kristoff, "An Automated Incident Response System Using BIND Query Logs", June 2006.
- [11] MySQL website – (Open Source Database), [http://www.mysql.com]. Last accessed on 27 July 2011.
- [12] IP FILTER – TCP/IP Firewall/NAT Software, [http://coombs.anu.edu.au/~avalon]. Last accessed on 27 July 2011.
- [13] P. Saint-Andre, Ed., Extensible Messaging and Presence Protocol (XMPP): Core, RFC 3920, 2004.
- [14] Joao Afonso and Edmundo Monteiro, "Development of an Integrated Solution for Intrusion Detection: A Model Based on Data Correlation", in Proc. of the IEEE ICNS'06, International Conference on Networking and Services - ICNS'06, Silicon Valley, USA, July 2006.
- [15] Tshark website – The Wireshark Network Analyzer, [http://www.wireshark.org]. Last accessed on 27 July 2011.
- [16] Joao Afonso and Pedro Veiga, "Protecting the DNS Infrastructure of a Top Level Domain: Real-Time monitoring with Network Sensors", WSNS 2008, 4th IEEE – International Workshop on Wireless and Sensor Networks Security, Atlanta, USA, 29 September – 2 October 2008.