

International Journal of Computer Science and Mobile Computing

A Monthly Journal of Computer Science and Information Technology

ISSN 2320-088X

IJCSMC, Vol. 3, Issue. 9, September 2014, pg.442 – 449

RESEARCH ARTICLE



An Approach for Determining the Health of the DNS

TEJASWINI YADAV C.Y¹, BALAJI RAJENDRAN², RAJANI P³

¹CSE Department, Sree Vidyanikethan Engineering College, INDIA

²CSE Department, Sree Vidyanikethan Engineering College, INDIA

³Center for Development of Advanced Computing, INDIA

¹teju.cy@gmail.com, ²balaji@cdac.in, ³prajanireddy@gmail.com

Abstract— Domain Name system (DNS) is a global and decentralized system comprising of several types of nodes across geographies that are critical in resolving billions of translation queries – from IP addresses to domain names and vice versa – in the Internet at any given moment. A stable and healthy DNS is therefore important for the smooth functioning of the internet which in turn is dependent on the millions of nodes that comprise it. However determining the health of the global DNS at any given moment remains infeasible, as it will require conducting millions of probes at every level in the DNS hierarchy. In this paper, we propose a simple approach that will be able to approximate the health of the DNS, by determining the critical nodes in the DNS hierarchy, which is passively and periodically monitored. The proposed approach is evaluated with an emulated setup and the initial results are encouraging.

Keywords— Domain Name System, DNS Health, Critical Nodes, Response time

I. INTRODUCTION

Domain Name System (DNS) is a global and decentralized system comprising of several types of nodes across geographies that are used for resolving IP addresses to domain names and vice-versa. The ever-growing Internet has now led to the classification of Domain Name system (DNS) as the second most critical component of the Internet, and has been defined as a critical infrastructure as well [1]. A big challenge is to ensure and guarantee the proper level of DNS health for a resilient and robust Internet. A healthy DNS is important for the smooth functioning of the internet. As DNS is a global system, determining the health of the overall DNS, is an enormous challenge as it involves millions of nodes, and may require millions of probes installed and configured, which is not feasible owing to several practical constraints including threats to privacy. We here in propose that the health of such a distributed and global system can be determined by determining the health of the few critical nodes and links i.e., those critical nodes that have the potential to stabilize or destabilize the entire DNS system thereby affecting the overall health of the DNS.

As DNS is decentralized there will be several critical nodes that will affect the DNS system. The critical nodes can be defined as the node that will de-stabilize the DNS system thereby affecting the overall health of the DNS. The criticality of the DNS node can be determined by factors such as the number of Internet nodes accessible through it, number of users accessing it

at any given point of time. The health of such critical nodes can be determined using factors such as the responsiveness of the node, queries responded over time, round trip time taken etc...

The network of critical nodes and their respective health go a long way in influencing the health of the overall DNS system. The critical nodes could be found anywhere in the DNS network - root servers, top-level domain name server, authoritative name server, recursive resolver, public DNS service mapper etc.

II. RELATED WORK

Domain Name System (DNS) is a highly distributed, hierarchical naming system for computers, services and resources connected to the internet, which associates information such as network addresses with a human readable name. DNS is used to map a hostname in the application layer to an IP address in the network layer. DNS therefore provides a simple service for lookup and translation of URL into IP address and vice versa. Initially, when the Internet was small, a file called "hosts.txt" was used to serve as a lookup table [2]. However, as the size increased, it became difficult to access and maintain. In order to resolve the complexity, DNS was introduced. It was developed in November 1983 by Paul Mockapetris. As DNS is decentralized there is no need to maintain data in a big and single file. The DNS is organized as an inverted tree structure with the root at the top. Internet Corporation for Assigned Names and Numbers (ICANN) – an organization is in charge of maintaining the root zone. Under root node there will be top level domains which includes country code top level domains (ccTLDs), unsponsored top level domains, generic top level domains (gTLDs) etc. In the DNS hierarchy there are 13 root name servers named from A to M that contains copies of the root zone. These 13 root servers are replicated across the world. DNS security is very important for the Internet, because all the users use the DNS services indirectly. Any compromise to the DNS could cause a heavy loss to the user. Some attacks on DNS are cache poisoning, DOS attacks, DDOS attacks and botnet attacks. Mainly the health of the DNS depends on Security, Stability and Resiliency (SSR) [3]-[4]. There are five vital signs of DNS which are used as a part of health assessment of the system. They are: Coherency, Integrity, speed, availability and resiliency.

The health of DNS can be determined with the help of critical nodes on the network. Determining DNS health requires monitoring the system, analyzing its behavior, planning and initiating corrective actions. 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. Analyzing the DNS health is one of the key challenges, due to following reasons: There don't exist any common standards or metrics; no common indicators; no common rules exist for computing the values of generally accepted indicators; no agreed terms for the normal DNS behavior; no common sharing or pooling availability [5]. So, here we propose a mechanism for analyzing the DNS health by using necessary DNS tools and we identify the critical nodes on the DNS system and then conducting passive probes on it.

A. Detection And Classification of Critical Nodes

For determining the Health of the DNS we have to detect the critical nodes and classify according to the patterns observed. We are considering the critical nodes as all the root nameservers, top 10 gTLDs (.com, .net, .org, .info, .biz, .mobi, .asia, .name, .tel, .pro), and top 15 ccTLDs (.tk, .de, .uk, .cn, .nl, .ru, .eu, .br, .ar, .au, .fr, .it, .pl, .ca, .us, .ch, .in, .es, .co, .be) and some nodes that had prone to attack earlier which are listed in table 1. We considered them as critical nodes based on number of sub domains and number of users accessed. The classifications of problems affecting critical nodes are configuration problems, DNS server and Hardware problem (Router), Heavy traffic load, Malware/DDOS/DOS attack and botnet attack. The propagation impact will depend on the number of nodes that were affected. For example, consider a G-root and L-root servers, both are the root servers that will maintain several ccTLDs and gTLDs. The propagation impact will be high if there was any attack on these root servers. On Feb 6th 2007 there was a DDOS attack on the G root and L root servers. The attack happened on these root servers because these servers haven't installed Anycast. Propagation impact was high. At least six root servers were attacked but only two are badly affected (G and L). The g-root was maintained by U.S Department of Defense and L-root by ICANN [6].

Similar DDOS attack was happened on the root servers on Oct 21st 2002. It was happened because the attacker sends many ICMP pings using a botnet to each of the servers. Propagation impact was low. The attack targeted all the 13 root servers for an hour. The routers were protected by packet filters which were configured to block all ICMP pings; they didn't sustain much damage.

Another attack was happened on March 18th 2014 on Google public DNS server of Venezuela and Brazil. The server was hijacked for nearly 22 minutes and the attack was based on BGP which is very hard to detect. The propagation impact was high, and the government agencies, financial institution and enterprises of that region were inaccessible.

B. Analysis of Critical Nodes

After we zeroed on these critical nodes, we conducted probes over a continuous period of time. We probed the identified critical nodes to find the response time (the time it took to get the response from the node). The results were accumulated over a period of time, to identify thresholds that can be used as a cut-off to flag the critical state of the critical node. After collecting the results over a period of time, the average response time was calculated. For continuous probing and storing the data we have used the POSTGRES database, which easily lends for executing scripts and conducting analysis.

TABLE 1
DETECTION AND CLASSIFICATION OF CRITICAL NODES

Node	Occurred in	Incident	Propagation impact	Effect	Problem
G root and L root (root servers)	Feb 6 th 2007	It happened on those root servers because they haven't installed Anycast.	Propagation impact was high	At least six root servers were attacked but only two of them are badly affected. The g-root which was run by the U.S department of defense and L-root by ICANN	DDOS attack
Root Servers	Oct 21 st 2002	It happened because the attacker sends many ICMP pings using a botnet to each of the servers	Propagation impact was low.	The attack targeted all the 13 root name servers for an hour. The routers were protected by packet filters which were configured to block all ICMP pings, they didn't sustain much damage.	DDOS attack

Venezuela and Brazil	March 18 th 2014	The attack happened on Google DNS server of Venezuela and Brazil. The attack is based on BGP, and it is very hard to detect. The attackers performed man-in-middle attack such kind of attacks is ideal for cyber espionage operations.	Propagation impact was high.	The Google DNS server 8.8.8.8/32 was hijacked for 22 minutes. The impact was mostly on government agencies, financial institution and enterprises.	Hijacking of DNS server (using BGP)
----------------------	-----------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------	-------------------------------------

III. IMPLEMENTATION AND ANALYSIS

In this section, we analyze the global DNS health by identifying few critical nodes with large number of sub nodes or large number of users dependent on it. After identifying a select set of critical nodes, we monitored those critical nodes on a periodic basis and calculated the average response time. The command-line utility DIG [7] came in handy in conducting these probes which was wrapped over a script that executed in the background. For instance, for probing a root name server, the following command was used.

```
#dig @server_ip_address name-server
```

If we want to probe an a.root-name server, the following command was used.

```
#dig @198.41.0.4 a.root-servers.net
```

The output of the above command is:

```
;; Query time: 157msec
;; SERVER: 198.41.0.4#53(198.41.0.4)
;; When: Sun Aug 31 12:00:01 IST 2014
MSG SIZE rcvd: 769
```

As shown above, similar probing was carried out on the select critical nodes as shown in table and the data was stored in postgresql.

TABLE 2

PROBING AND STORING THE DATA IN POSTGRESQL

Id	Ip address	Response time	Date	Time
1	198.41.0.4	175msec	2014-08-31	12:00:01
2	192.228.79.201	235msec	2014-08-31	12:00:01
3	192.33.4.12	241msec	2014-08-31	12:00:02
4	199.7.91.13	61msec	2014-08-31	12:00:02
5	192.203.230.10	60msec	2014-08-31	12:00:02

Let us take an example for Google public DNS server (8.8.8.8). We probed this node continuously for a period of time and stored the query time for finding the average response time as given in the following table.

TABLE 3

RESPONSE TIMES OF GOOGLE PUBLIC DNS SERVER AT REGULAR INTERVALS OF TIME

Response time msec	Date	Time
109	2014-08-31	11:07:05
131	2014-08-31	11:25:27
32	2014-08-31	11:30:32
33	2014-08-31	11:35:41
28	2014-08-31	11:40:15

The response time for Google Public DNS Server is as below:

TABLE 4
HOURLY AVERAGE RESPONSE TIME FOR GOOGLE PUBLIC DNS SERVER

Response time (msec)	Time	Date
57.16	11:00	2014-08-31
39.16	12:00	2014-08-31
32	01:00	2014-08-31
30.28	02:00	2014-08-31
29.17	03:00	2014-08-31

The average response times were then plotted as depicted in the following figure and sent to the DNS Visualizer module, which will then display to the DNS administrator.




Figure 1: Average Response Time of Google Public DNS Server

IV. EXPERIMENTS AND RESULTS

We simulated a top-level domain (.xy) with thousands of sub-domains under it as illustrated in Figure 2. We then constructed a DNS Query prober that probed the top-level domain with queries to several of its sub-domains, at regular intervals of time, and the average response times were calculated on hourly basis as given in Figure 3.




Figure 2: .xy node hierarchy

We are going to probe that domain continuously for a period of time.

The response time of .xy is as below:




Figure 3: Average Response Time Graph for .xy node in normal condition

An attack was simulated on the .xy node that made the domain struggle to respond to the queries. The propagation impact will depend on the number of sub-domains under it; though measuring it is not in the scope of this work. During attack, the response times conducted by the prober indicated considerable difference as given in Figure 4.




Figure 4: Average Response Time Graph for .xy node after attack

V. CONCLUSION

We have proposed a mechanism to analyze the global DNS health by identifying critical nodes that were defined as either having a large number of sub-nodes under it or having a number of users dependent on it. We then propose to monitor those identified critical nodes on a periodic basis and determine their average response times for the legitimate queries by using passive mechanisms without any intrusive probes or affecting the privacy of the users. The average response times will then indicate abnormal behaviors in case of attacks such as Denial-of-Service being carried out on the segments of the critical nodes.

The proposed mechanism is evaluated with an emulated setup of a top-level domain and sub-domains under it, and the response times were analyzed which differed in case of attacks indicating abnormal behaviors. This proves the proposed concept which can be further refined in future to conduct detailed analysis

REFERENCES

- [1] Casalicchio .E, Caselli .M, Coletta .A “Measuring the Global Domain Name system” Network, IEEE, volume: 27 Issue:1
- [2] P. Mockapetris, Domain Names - Implementation and Specification, RFC 1035, 1987.
- [3] ICANN, “Measuring the Health of the Domain Name System,” *Report of the 2nd Annual Symp. DNS Security, Stability, & Resiliency*, Kyoto, Japan, 2010.
- [4] Global Cyber Security Center (GCSEC), “Measuring the Naming System (MeNSa) Project,” <http://www.gcsec.org/activity/research/dns-security-and-stability>.
- [5] ICANN, GCSEC, DNS-OARC, “Final Report of the 3rd Global DNS Stability, Security and Resiliency Symposium,” 2011, Rome, Italy.
- [6] <https://www.icann.org/en/system/files/files/factsheet-dns-attack-08mar07-en.pdf>
- [7] <http://manpages.ubuntu.com/manpages/lucid/man1/dig.1.html>