Automated detection of network intervention policies based on anomalous Domain Name System server responses

The Chokepoint Project
info@chokepointproject.net

February 10, 2015

Abstract

It is an established fact that manipulation of the Domain Name System (DNS) is used to implement censorship policies on the internet. This paper describes the implementation of an automated system to monitor the responses of DNS servers and classify these responses according to the extent the response is likely to have been manipulated in a way that is consistent with the intent to obstruct access to specific websites and services.

1 Background

The internet provides unprecedented communications capabilities. The social and political ramifications have in recent years manifested themselves in a dramatic fashion. In fact, A number of authors have commented on the instrumental role of the internet in social movements such as the Tahrir protests [1], the Tunesian revolution and the Occupy movement [2].

Conversely, the internet provides unprecedented capabilities for surveillance. Government-operated systems that allow for large scale interception of internet traffic have been in operation as early as 1997 and likely earlier [3] [4].

Since then, these systems have evolved to provide sophisticated intervention and disruption capabilities in addition to surveillance capabilities [5] [6].

Given that these systems are typically managed by government institutions which operate under a high level of secrecy, details about their functioning are scarce. This raises several questions, such as the extent to which their operation is in accordance with national and international law, in particular the right to privacy and the right to freedom of expression [7].

Detecting the presence of these systems and reporting on the mechanisms by which they function is therefore important not just from an academic or technical perspective, but also as a prerequisite for evidence-based policy and human rights analysis.

2 How manipulation of the Domain Name System can be used to obstruct access to information

One form of network intervention is based on manipulation of the Domain Name System (DNS).

The DNS is a decentralized directory system. Its primary function is to translate human readable internet domain names to numerical internet addresses ('IP addresses') that correspond to networked devices. The Domain Name System is organized as a hierarchical tree of DNS servers. Typically, a user connecting to the internet will be assigned a DNS server that is in close proximity[1] to that user. This DNS server then translates internet domain names to IP addresses for said user.

[1] Close proximity from a network topological view.

In most cases, end users do not interact with the Domain Name System directly. However, behind the scenes, the DNS is instrumental in determining the route taken to reach a website, whether that website is also the website a user *intended* to visit, and whether the website can be reached at all.

As such, DNS manipulation can be used to implement or support internet censorship and surveillance policies. The decentralized nature of the DNS allows for highly localized implementations of DNS based filtering policies. The visibility and effects of DNS based manipulations can be limited to customers of a specific Internet Service Provider (ISP) or employees of a single company. They can also be applied to an entire country or specific regions in that country.

DNShonest Joss Wright at the Oxford Internet Institute has researched the application of DNS manipulation techniques in China [8].

As part of this research, Wright developed a set of software tools ('DNShonest') that allow large scale, semi-automated scanning, to aid in the identification of manipulated DNS responses.

DNShonest identifies manipulation of DNS requests by asking a DNS server for the IP addresses of a number of well-known domains. The responses are then compared to the responses from a "trusted" DNS server. When the responses are substantially different, this is an indication that some form of DNS manipulation may be occurring.

3 Implementation of a system to detect manipulation of the Domain Name System

Wright's DNShonest is not (and was not intended) to be a monitoring system: it is a research tool meant to sporadically obtain a large set of data, which can be further mined for evidence of DNS manipulations using software tools and manual analysis. A monitoring system, by contrast, must run continuously and have the ability to identify DNS manipulations unattended without relying on manual analysis.

The Chokepoint Project, as part of its mission to monitor and report on network interventions in near real-time, has implemented a continuous DNS monitoring system which runs unattended, based on an augmented version of Wright's DNShonest tool.

The remainder of this paper describes the functioning of this monitoring system, the key differences between this monitoring system and Wright's DNShonest, and discusses some limitations of monitoring DNS manipulation in general.

The Chokepoint Project's DNS monitoring system The detection of DNS manipulation relies on the detection of anomalous responses from DNS servers. Anomalous responses can be detected by comparing them against a "known good" response.

This process is complicated by a number of factors, such as there being no single "good" response. For example, different DNS servers may legitimately return different answers for the same domain. Several of these limitations are discussed in more detail in the section 5.

Overview The DNS monitoring system implemented by The Chokepoint Project works by periodically[2] performing a "run". A run consists of the following steps:

- Selection of DNS servers
- Selection of domain names
- DNS probing
  - Query each DNS server for the IP address of each domain name
  - Classify the responses of each DNS server as valid, maybe valid, invalid, or ambiguous
- HTTP probing
  - For responses classified as ambiguous during DNS probing, perform HTTP probing
  - HTTP probing tests whether the returned IP address points to a webserver that hosts the website associated with the requested

[2] At the time of writing, The Chokepoint Project performs two runs per day.

    domain name, by performing a HTTP request on the IP address returned by the DNS server
  - Classify the results of the HTTP probing as valid, maybe valid, or invalid

- Enrichment
  - GeoIP lookups
  - Autonomous System information
- Reporting
  - Summarize the results by domain, DNS server, and autonomous system

The following sections describe these steps in detail.

3.1 Selection of DNS servers to monitor

The monitoring system maintains a database of DNS servers. Each record is associated with a country code, describing the country in which this DNS server operates and two timestamps: a creation date that specfies when the record was added to the database, and an expiration date that specifies when the record is no longer relevant.

The timestamps allow the system to accommodate the fact that DNS servers are constantly added and removed from the network.

Based on this information, the monitoring system selects the currently active DNS servers for a country.

3.2 Selection of domain names to monitor

The monitoring system maintains a database of domain names in a way that is identical to the database of DNS servers.

It is important to mention that the selection of domain names (and to a lesser extent, the selection of DNS servers) crucially determines the results of the monitoring system. The monitoring system can only detect DNS manipulation of the selected domain names: additionally (and more importantly), the monitoring system is blind to DNS manipulation of domain names that are not part of the selection. This limitation is discussed in more detail in section 5.

3.3 DNS probing

After the DNS servers and domain names have been selected, the monitoring system contacts every selected DNS server and requests the IP address of each selected domain (this process is known as a "DNS lookup").

The monitoring system then performs a DNS lookup on a trusted DNS server, to compare against the responses of all the tested DNS servers.

For every query, the monitoring system stores the following results:

`error_code`
    An error code: 0 means no error

`reverse_lookup`
    The reverse lookup for the IP address returned by the DNS server

`report_country_code`
    The country code of the run

`report_datetime`
    Date and time of the run

`dns_response`
    The response from the DNS server

`local_result`
    The response from the trusted DNS server

`tested_domain`
    The domain that was tested

`requested_DNS server`
    The IP address of the DNS server

Classification The results of the DNS probes are classified in the following way:

`valid`
    The response is valid. This is the case when:

1. the first two octets of the IP address returned by the remote DNS server are identical to the first two octets of the IP address returned by the trusted DNS server, or

2. the reverse lookup of the IP address matches at least 2 top-level components from the tested domain name.

`probably_lie`
The response is probably a lie. This is the case when the remote DNS server did not return any IN A records.

`lie`
The response is certainly a lie. This is the case when

1. the remote DNS server returns ’127.0.0.1’, or
2. the remote DNS server returns error_code of 1, indicating that the domain does not exist (i.e. NXDOMAIN).

`unknown`
This is the case when none of the other cases apply. These items are subjected to additional tests by way of HTTP probing, explained below.

3.4 HTTP probing

The results classified as `unknown` during DNS probing are subjected to additional testing using HTTP probing. HTTP probing involves making a HTTP request to the IP addresses returned by the remote DNS server, then analyzing the HTTP responses.

First, the IP addresses returned by the remote DNS server (the ‘remotely resolved IPs’) are added to a probing list. The IP address returned by the trusted DNS server (the ‘locally resolved IP’) is also added to this list.

Then, a HTTP request is made to all IP addresses, with the `Host:` HTTP header set to the domain name that is expected to be associated with these IP addresses. The response headers and response body for each request are stored.

Classification Based on the stored response headers and response body, the monitoring system performs a further classification as follows:

`valid`
The DNS server response is valid. This is the case when:

1. any HTTP request to a remotely resolved IP yields a correct redirect, i.e. returns a `Location:` HTTP header that contains a domain name that looks like the tested domain name, or
2. the responses from the locally resolved IP are similar to the response of the last remotely resolved IP. This means that both the HTTP response status code (2xx, 3xx, ...) and response body need to be similar. Similarity of response body is determined by matching the `<title>` tag in the responses (in case of a 2xx response) or by matching the `Location:` header redirect (in case of a 3xx response).

`maybe_lie`
The response may be a lie. This is the case when:

1. any of the HTTP statuses from remotely resolved IPs are 4xx/5xx and there are no 2xx/3xx statuses, or
2. the HTTP statuses from remotely resolved IPs differ, or
3. the result of the HTTP probe was inconclusive, i.e. none of the other cases apply.

`probably_lie`
The response is probably a lie. This is the case when:

1. all remotely resolved IPs return 4xx/5xx responses, or
2. the HTTP status from the last remotely resolved IP[3] does not match the HTTP status for the locally resolved IP, or
3. there is a mismatch between the response from the locally resolved IP and the last remotely resolved IP[4], i.e. there is a mismatch in either the `<title>` tag (in case of

[3] To simplify the implementation, we assume that the response from the last remotely resolved IP is representative of the responses by all the remotely resolved IPs. This assumption is clearly not valid in general. The effects of this assumption are relatively small, because most queries yield only a single IP address, and those queries which yield multiple IP addresses generally involve sites hosted on highly homogenous content distribution networks, where all servers return identical responses.

[4] As above.

a 2xx response) or a mismatch in the Location: header redirect (in case of a 3xx response).

`lie`
: The response is certainly a lie. This is the case when all HTTP probes have failed or returned 4xx/5xx codes.

3.5 Enrichment

After collecting the results, the data is augmented with two additional sources of information:

1. GeoIP lookups
2. Autonomous System information

GeoIP lookups The Chokepoint Project maintains archives of the freely available MaxMind GeoIP database. The latest version of the database is periodically downloaded from MaxMind.

This database is used to map remote DNS server IP addresses to a longitude and latitude, as well as a country and city names.

The accuracy of the free MaxMind database is limited. It cannot always resolve an IP address to a city name and the longitude/latitude information is often approximate at best. The accuracy varies significantly with country/region. As a rule of thumb, the most highly developed regions offer the most accurate information.

Autonomous System information The Chokepoint Project maintains two databases that can be used to map an IP address to an Autonomous System name. An Autonomous System (AS) is set of routers under a single technical administration[5]. The monitoring system uses this information to provide an indication of the administrative entity responsible for the DNS server which returns anomalous responses.

The two databases used in determining the Autonomous System information are:

- The Routing Information Base
- The APNIC Routing Report

At regular intervals, The Chokepoint Project downloads an updated Routing Information Base (RIB) file[9]. This database is used to to look up the Autonomous System Number (ASN) for an IP address.

The Chokepoint Project also regularly fetches the latest Routing Report from the Asia-Pacific Network Network Information Center (APNIC) [10], processes the information, and stores it in a database. This database is used to look up the organization name that corresponds to an Autonous System Number.

3.6 Reporting

The augmented results data is summarized in the following reports:

1. Summarized by report ID
2. Summarized and aggregated by domain name
3. Summarized and aggregated by AS
4. Summarized and aggregated by DNS server

3.6.1 Summarized by report ID

This report includes the following information per run, where each run is identified by a report ID:

`domains_tested`
: The total number of domains tested.

`nameservers_queried`
: The total number of DNS servers queried.

`queries_total`
: The total number of queries performed for this run (nominally, this is `domains_tested` multiplied by `DNS servers_queried`).

`queries_replied`
: The number of queries where the DNS server responded.

`queries_errored_out`
: The number of queries that failed, i.e. the DNS server did not respond at all, or responded with an error code other than NXDOMAIN.

`queries_lie`
: The number of queries classified as `lie`.

[5] See RFC 1930, `http://tools.ietf.org/html/rfc1930`

`queries_probably_lie`
: The number of queries classified as `probably_lie`.

`queries_maybe_lie`
: The number of queries classified as maybe_lie.

3.6.2 Summarized and aggregated by domain name

This report is similar to the summarized report, but grouped by domain name, where different counts (such as `queries_total`) refer to the number of queries performed for that particular domain name.

3.6.3 Summarized and aggregated by AS

This report is similar to the summarized report, but grouped by ASN, where different counts (such as `queries_total`) refer to the number of queries performed for that particular AS. In addition, the report contains the following fields:

`requested_nameserver_asn`
: The Autonomous System Number that this DNS server belongs to, as reported by the RIB [9].

`requested_nameserver_asn`
: The name of the organization responsible for the Autonomous System that this DNS server belongs to, as reported by the APNIC Routing Report [10].

In addition, the report contains corresponding fields for the actual DNS server that was queried, in case the query was redirected by the requested DNS server.

3.6.4 Summarized and aggregated by DNS server

This report is similar to the summarized report, but grouped by DNS server, so that the different counts (such as `queries_total`) refer to the number of queries performed for that particular DNS server. In addition, the report includes the following fields:

`requested_nameserver_city`
: The city of the DNS server, as reported by the MaxMind GeoIP database.

`requested_nameserver_lat`
: The latitude of the DNS server, as reported by the MaxMind GeoIP database.

`requested_nameserver_lon`
: The longitude of the DNS server, as reported by the MaxMind GeoIP database.

`requested_nameserver_asn`
: The Autonomous System Number that this DNS server belongs to, as reported by the RIB data [9].

`requested_nameserver_asn`
: The name of the organization responsible for the Autonomous System that this DNS server belongs to, as reported by the APNIC data [10].

In addition, the report contains corresponding fields for the actual DNS server that was queried, in case the query was redirected by the requested DNS server.

4 Results

The DNShonest research paper by Joss Wright[8] was based on a survey of 86 domains and 208 DNS servers in China.

For purposes of comparison, we have applied our DNShonest implementation to these same set of domains and DNS servers.

The 86 domains include popular domains like `www.google.com` and `dropbox.com`, as well as websites that provide access to circumvention technology, such as `www.torproject.org`, and websites hosting politically sensitive material, such as `wikileaks.org`.

Since each domain is tested against each DNS server, this yields a total of 17680 queries.

Of these 17680 queries, the original implementation [6] identifies 4081 lies (with 742 queries whitelisted). 5506 queries yield a "no valid DNS servers" response, 924 queries yield a "no answer" response and 4031 queries yield a "timeout". The monitoring system takes about twenty minutes to complete a single run of this magnitude.

[6] Specifically, the find_all_lies.py script that is part of the dnshonest source distribution.

classification	original	modified
lies	4081	3558
whitelisted	742	0
no valid DNS servers	5506	5368
no answer	924	917
timeout	4031	4246

Table 1: Comparison of results between DNShonest tool and the DNS monitoring system implemented by The Chokepoint Project

Our implementation identifies 3558 lies, 86 responses that are probably lies, and 197 responses that may be lies. 5368 queries yield a "no valid DNS servers" response, 917 queries yield a "no answer" and 4246 queries yield a timeout.

5 Limitations

The DNS monitoring system is not perfect: it does not comprehensively monitor all possible DNS servers and domain names, and it yields both false positives (classifying valid responses as anomalous) and false negatives (classifying anomalous responses as valid).

While the heuristics employed by the monitoring system were designed to err on the side of caution, i.e. to avoid classifying a response as anomalous in case of doubt, the monitoring system can still classify responses incorrectly due to the wide variety of possible responses, the highly localized nature of DNS manipulation, and the variation in network conditions.

Some of these incorrect classifications can be addressed by incorporating more sophisticated heuristics[7]

The selection of DNS servers One of the fundamental limitations is the impossibility to obtain an authoritative list of all DNS servers in a country or region. Even if such a list were readily available, the monitoring system would have to be expanded to exhaustively test all of them on a continuous basis.

For this reason, ideally, the test would be limited to a *representative selection* of DNS servers in a country or region.

However, determining which criteria would allow the (semi) automatic construction of a representative list of DNS servers and how this list would remain up-to-date over time, is still an open question.

The selection of domain names An equally fundamental, and probably more acute limitation, is the selection of domain names to test. As with DNS servers, it is not possible to obtain an authoritative list of all domain names to test in a country or region. The selection of domain names is further complicated by two factors of a socio-political nature.

First, virtually all types of DNS manipulation serve some social or political goal: access to services or websites may be manipulated or censored for legal, moral or political reasons. The selection of domain names to monitor for manipulation therefore incorporates an inherent value judgment as to which types of potential manipulation are more or less significant than others.

Second, the significance of manipulation of any particular domain name is highly dependent on regional and cultural context, which requires local knowledge to appraise accurately.

As with the selection of DNS servers, a methodology yielding a *representative selection* of domain names is ideal. For the reasons mentioned above, it seems inevitable that any selection will be based on subjective criteria, limiting the ability to automate the selection process.

Detection prevention An actor manipulating the DNS may want to obscure the fact that this manipulation is taking place. This can be done by either limiting the manipulation to a subset of the network that is outside the range of the monitoring system (e.g. by geographically limiting the manipulation) or by presenting unmanipulated results to the monitor-

[7] For example, the monitoring system used to incorrectly classify responses for domain names hosted on the Akamai Content Distribution Network because the HTTP probing mechanism did not include a valid HTTP User-Agent header. This issue was fixed on February 5 2015.

ing system (e.g. by only manipulating queries if they do not originate at the monitoring system).

Mitigating this effect would require the development of technological tools that help obfuscate the monitoring process and a broad geographical and network-topological distribution of the monitoring infrastructure.

Load-balancing and content delivery networks A domain name can resolve to more than a single IP address. DNS-based load balancing is commonly used to distribute network requests over multiple servers by returning a different address on each DNS request.

In addition, websites are increasingly hosted on large content distribution networks, such as Akamai's or Google's. The domain names for these websites can resolve to hundreds of IP addresses, which means that any of those hundreds of IP addresses constitute a valid response.

Unfortunately there is no reliable, automated way to obtain all the possible addresses that constitute a valid response. This makes it difficult to detect, based on the DNS server response alone, whether that response is valid or anomalous.

To address this issue, the monitoring system implements a HTTP probing mechanism, which heuristically compares the contents of a web page to the expected contents of the web page.

Confidence level Not every anomalous response from a DNS server implies that DNS manipulation is taking place. The DNS server or website may be misconfigured or temporarily unavailable, or there may be subtle differences between the expected response and the actual response that make it hard to categorically classify a result as valid or invalid.

For this reason, the monitoring system implements a classification system to classify probe results in different categories: `lie`, `probable_lie` and `maybe_lie`, to distinguish between responses that are obviously false and responses that are just unexpected.

6 Conclusions

Limiting access to information by manipulating the Domain Name System is easy and cheap, since it relies on network infrastructure that is already in place and requires no additional knowledge beyond what is needed to set up and manage the network in the first place. Because it can be deployed without conspicuous changes to the network infrastructure, and because of the decentralized nature of the Domain Name System, the impact of DNS manipulation is mostly invisible (except to the impacted users), and may easily evade detection.

While knowledgeable users can circumvent DNS manipulation, it seems plausible that DNS manipulation can effectively staunch the dissemination of information to a broad audience, and it is an established fact that DNS manipulation has been used to implement censorship policies[11][12].

The DNS monitoring system implemented by The Chokepoint Project is capable of performing hundreds of thousands of queries per hour. This corresponds to testing hundreds of domain names against hundreds of DNS servers per hour. While these numbers are impressive when considering that the DNS monitoring system runs on modest hardware and a limited network connection, it is also apparent that DNS monitoring for truly representative lists of domain names and DNS servers on a global scale will require scaling up the system by several orders of magnitude.

The HTTP probing mechanism allows the DNS monitoring system to accurately test and classify many of the domain names that are hosted on Akamai or Google content delivery networks[8]. Augmenting the results with data from APNIC [10] and RIB data [9] identifies the organizations suspected of performing DNS manipulation.

By capturing and storing data over time, the DNS monitoring system allows for comparisons over time. Since the results of the system may be influenced by changes to the detection heuristics, the system also includes an annotation system to record and be able

[8] This is an improvement over Wright's DNShonest tool, which always considers IP addresses belonging to these CDN's to be valid.

to account for these changes when making historical comparisons.

The fundamental problem of how to select representative DNS servers and domain names remains an open question. We hope that our reporting on DNS manipulation as it takes place helps spur efforts to begin tackling this difficult problem.

7 Acknowledgments

We would like to thank Joss Wright at the Oxford Internet Institute, who graciously provided us with the DNShonest source code, valuable advice, and great pub food.

References

[1] Tufekci, Z. and Wilson, C. (2012), *Social Media and the Decision to Participate in Political Protest: Observations From Tahrir Square.* Journal of Communication, 62: 363-379. doi: 10.1111/j.1460-2466.2012.01629.x

[2] Castells, M. (2012) *Networks of Outrage and Hope: Social Movements in the Internet Age*, Polity Press

[3] Schmid, G. et al. (2001) *Report on the existence of a global system for the interception of private and commercial communications (ECHELON interception system)* (2001/2098(INI)), European Parliament

[4] Smith, S. P. et al. (2000) *Independent Review of the Carnivore System*, IIT Research Institute

[5] Appelbaum, J. (2013) *Shopping for Spy Gear: Catalog Advertises NSA Toolbox*, Spiegel Online

[6] Conaway, S. n.d. (2009) *The Great Firewall: How China Polices Internet Traffic*, Certification Magazine

[7] Articles 12 and 19 of the UN Universal Declaration of Human Rights

[8] Wright, J. (2014) *Regional Variation in Chinese Internet Filtering. Information*, Communication & Society 17 (1) 121-141.

[9] RIB files are downloaded from `http://archive.routeviews.org/bgpdata`

[10] APNIC files are downloaded from `http://thyme.apnic.net/ap-data`

[11] Anonymous (2014) *Towards a Comprehensive Picture of the Great Firewall's DNS Censorship*, 4th USENIX Workshop on Free and Open Communications on the Internet (FOCI 14)

[12] Dornseif, M. (2004) *Government mandated blocking of foreign Web content*, Jan von Knop, Wilhelm Haverkamp, Eike Jessen (Editors) Security, E-Learning, E-Services: Proceedings of the 17. DFN-Arbeitstagung ueber Kommunikationsnetze, Duesseldorf 2003, ISBN 3-88579-373-3; Series: Lecture Notes in Informatics ISSN 1617-5468; Pages 617-648