

Poisoning the Well – Exploring the Great Firewall’s Poisoned DNS Responses

Oliver Farnan
University of Oxford
Wolfson Building
Parks Road
Oxford, UK
oliver.farnan@cs.ox.ac.uk

Alexander Darer
University of Oxford
Wolfson Building
Parks Road
Oxford, UK
alexander.darer@cs.ox.ac.uk

Joss Wright
Oxford Internet Institute
1 St Giles
Oxford, UK
joss.wright@oii.ox.ac.uk

Keywords

Censorship; DNS Poisoning; Great Firewall; China; Internet Security

CCS Concepts

•**Security and privacy** → **Privacy protections**; Web protocol security; Social aspects of security and privacy; Spoofing attacks;

Abstract

One of the primary filtering methods that the Great Firewall of China (GFW) relies on is poisoning DNS responses for certain domains. When a DNS request is poisoned by the GFW, multiple DNS responses are received - both legitimate and poisoned responses. While most prior research into the GFW focuses on the poisoned responses, ours also considers the legitimate responses from the DNS servers themselves. We find that even when we ignored the immediate poisoned responses, the cache from the DNS servers themselves are also poisoned. We also find and discuss the IP addresses within the DNS responses we get; in particular 9 IP addresses that are returned as a result for many different poisoned domains. We present the argument that this type of attack may not be primarily targeted directly at users, but at the underlying DNS infrastructure within China.

1. BACKGROUND

Much of the internet in China is censored. The Chinese state employs a multi-pronged approach for censorship, that includes both technical and non-technical means. This is colloquially referred to as the Great Firewall (GFW).

One of the key technical methods used by the GFW is DNS poisoning. When the GFW observes DNS queries to certain domains, it responds by sending a poisoned DNS response to the requesting DNS resolver. Due to its position in the network, this typically reaches the requesting DNS resolver before the response from the DNS server. This results

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WPES’16, October 24 2016, Vienna, Austria

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-4569-9/16/10...\$15.00

DOI: <http://dx.doi.org/10.1145/2994620.2994636>

in the requesting DNS resolver caching the poisoned DNS response, and ignoring the response from the DNS server itself. Throughout this paper we describe the DNS response that comes from the DNS server as the ‘legitimate’ response, and the response that comes from the GFW as the ‘poisoned’ response.

Most previous work on Chinese DNS censorship has focused on poisoned responses from GFW infrastructure, and not legitimate responses from DNS servers in China. We believe that the legitimate response deserves further investigation. Our analysis looks at both legitimate and poisoned DNS responses. We look at whether DNS servers themselves respond with the correct IP addresses, or if they too are poisoned. We hope that this will allow us to follow up with an analysis into geographical or logical correlation between results, and between any incorrect IP addresses that we receive in responses.

We find that most DNS servers within China are themselves poisoned, often with the same IP addresses that are returned from the GFW itself. We believe that DNS servers within China are themselves being poisoned by the same process known to affect users’ DNS queries. We describe the incorrect IP addresses that come from these queries, and their frequency.

We highlight 9 IP addresses that are frequently and independently given in responses to DNS queries to servers under the influence of the GFW. These IP addresses are returned in both legitimate and poisoned results. There appears to be no pattern or relationship between these IP addresses, either in terms of logical address space, or geographical registration. There does not appear to be active hosts located at these IP addresses.

2. RELATED RESEARCH

The seminal paper in this area is written by Lowe et al. from 2007[12]. Lowe introduces DNS poisoning in China, and provides an analysis of its implementation. It presents a technique to start mapping this censorship (via editing the TTL field of the DNS query) that has proved useful in subsequent analyses of the GFW. This work built upon earlier work by Clayton et al.[6], which was focused not on DNS based censorship, but on HTTP traffic being filtered by the insertion of TCP RST packets into the TCP stream, causing connections to drop.

More recently, a thorough analysis of the GFW’s DNS poisoning was presented in a pair of Anonymous papers produced in 2013 and 2014. In 2013 Anonymous (

ail.com) began to look into the collateral damage of DNS filtering[3]. They looked at how DNS queries in other countries were affected by the GFW's DNS interception. Their paper deals primarily with DNS servers performing recursive lookups to servers within China, resulting in pollution of their results. They cite examples of lookups that were made where neither the requester nor resolving server resided within China, yet the results were being poisoned because the resolving server tried to answer by recursively requesting the answer from a DNS server within China. They found that 26% of open recursive resolvers worldwide were vulnerable to result pollution in this manner.

Following on from this in 2014, Anonymous (zion.vlab2@gmail.com) provided a more thorough analysis of GFW DNS poisoning[4]. They built upon the earlier work from Lowe et al.[12] and combined it with the King method[11] to map out where DNS poisoning was occurring. They found that the majority of locations where DNS poisoning was taking place were within the border ASs of China's internet, and was primarily targeted at requests going into or out-of China. As well as this, they performed a large scale evaluation of the domains that are filtered by the GFW. They attempted to resolve all listed Alexa domains (130 million individual domains) and found that of these; around 35 thousand were censored by the GFW. Through subsequent analyses they were able to identify exact terms that were filtered by the GFW. They then offered a method for attempting to estimate the amount of requests a node deals with, and provided this analysis for a single node.

Simultaneously, Wright presented work indicating that it is wrong to view the GFW as a homogeneous filter across the entire country and provided evidence for regional variation in censorship[17]. Again, this evidence focused on DNS poisoning performed by the GFW, and how queries changed depending on where on the network they were intercepted. Wright observed that the responses from the GFW were different depending on where in the country they were intercepted, and found evidence for a decentralisation of filtering based on centrally coordinated policy. This is supported by a preceding piece by Xu et al.[18], suggesting that different ASs and ISPs within China performed filtering differently.

There has been other analytical work performed on the type of content the GFW filters. In 2007 Crandall et al. presented a technique for determining which[7] characters and keywords the GFW attempted to filter. They found that the GFW did not attempt to block all communication that could be considered harmful, but instead filtered enough to encourage self-censorship. They argued that in this sense, it was closer to a panopticon than a firewall. In a similar discovery, King et al.[11] found that agents of censorship in China did not attempt to filter all communication, but instead focused on that which could have real world consequences. Specifically, they realised the primary focus for censorship was often discussion that encouraged collective expression. Their analysis was not based on direct filtering by the GFW like earlier analyses, but instead on content which had been posted but then later removed.

3. METHODOLOGY

Most previous GFW analyses have focused on the first DNS responses received to queries. This is typically a poisoned response, coming from systems belonging to the GFW, rather than from the DNS servers[4]. For our analysis we

cover not just this first response but also data from any subsequent DNS responses. This ensures that legitimate responses from DNS servers are captured.

We obtained a list of DNS servers from Public DNS Server list[14]. Of these 1949 were registered in China. We cross referenced this list against the MaxMind[13] GeoIP database, and removed 78 servers listed as outside of China. This left 1871 public DNS servers in China, according to both Public DNS Server List and MaxMind.

We choose 15 domains to act as our test domains. This consisted of 5 control domains, and 10 domains that were previously known to have been poisoned by the GFW. The choice of these domains was based on the Alexa most popular domains site[2], taken on 18/04/2016.

We used the Alexa top domains list and the ViewDNS.info Chinese Firewall Testing Tool[15] to check each domain listed, sequentially from the most popular domain downwards. This tool makes a DNS request to a variety of servers in China, and flags the domain if any of the requests return an incorrect response. Any domains which showed signs of filtering were then manually checked with DIG to confirm the result.

The most popular domain found to return incorrect IP addresses was google.com. This process was then repeated on sequentially less popular domains, until 10 domains returning incorrect IP addresses were found. These popular filtered domains were:

- google.com
- youtube.com
- facebook.com
- wikipedia.org
- twitter.com
- instagram.com
- blogspot.com
- imgur.com
- github.com
- blogger.com

The control domains were the 5 domains listed as the most popular within China, and consisted of:

- baidu.com
- qq.com
- taobao.com
- sina.com.cn
- weibo.com

We sent a DNS request to each of the 1871 public DNS servers for each of the 15 domains, resulting in 28065 total DNS requests. Each DNS request had a unique and incremental DNS transaction ID value, which was used to determine which response matched to which request. This was needed as many requests had multiple responses due to the GFW. All traffic during experimentation was recorded directly from network traffic, rather than from the software DNS resolver making the request. This was then stored in packet capture format. This process of recording the network traffic itself ensured that all DNS responses were captured, even in cases where multiple DNS answers came for the same query – for example when the GFW had poisoned the request – and that each response could be correctly attributed to the request that prompted it.

4. FINDINGS

Most requests that were made to DNS servers under the influence of the GFW came back with two DNS responses: the legitimate response from the server, and the poisoned response from the GFW. Our initial expectations were that for DNS requests made to servers under the influence of the GFW, we would receive one correct DNS response from

the server, and one incorrect DNS response from the GFW infrastructure.

This was not the case. On most occasions both the legitimate and the poisoned DNS responses were incorrect.

4.1 Multiple Poisoned Responses

Some DNS requests returned more than two responses. In many cases, none of the responses returned the correct IP address for the domain. For example, sending a DNS request for “blogger.com” to 140.206.217.2 elicited three responses, one from the legitimate DNS server and two from GFW infrastructure, none of which were correct.

4.2 Nine repeated IP addresses

Many poisoned and incorrect responses from the GFW returned an IP address from a small set of incorrect IP addresses. Not only would the same incorrect IP address be returned for multiple requests to the same domain, but often the same incorrect IP addresses were observed as answers for different filtered domains. For example, requests to both facebook.com and blogger.com returned results for 37.61.54.158.

We also observed these same IP addresses as responses within both legitimate and poisoned DNS responses. In many cases we observed both DNS responses coming from this small set of IP addresses.

In total we found 9 IP addresses that were repeated for incorrect results for requests made into the GFW. These IP addresses were not all observed with the same frequency, with the most common address appearing over 11 times more than the least frequent address. These IP addresses are registered with, and geolocate to, different locations and AS around the world, with no obvious pattern. The occurrences figures below came from running the experiment four times over four different days.

These IP addresses were:

- 37.61.54.158 - 54076 occurrences
- 93.46.8.89 - 14642 occurrences
- 59.24.3.173 - 4848 occurrences
- 78.16.49.15 - 4841 occurrences
- 203.98.7.65 - 4832 occurrences
- 243.185.187.39 - 4755 occurrences
- 159.106.121.75 - 4754 occurrences
- 46.82.174.68 - 4711 occurrences
- 8.7.198.45 - 4683 occurrences

4.3 History of Null 9 IP Addresses

There is evidence of the use of these addresses going back several years. There has been some limited observation of these addresses in the academic community[17], and informal observation from others[5][19]. Lowe’s work from 2007 identified 8 different IP addresses being returned as results from similar requests[12].

We used a passive DNS replication database to check for historic use of these IP addresses within DNS. Passive DNS replication is a technique to replicate domain information by passively collecting historic DNS queries and their responses[16]. We searched the passive DNS database DNSDB[1] for instances of these IP addresses.

DNSDB records show that before 2010 these IP addresses were not associated with these filtered domains, and had few, if any, domains associated with them. Some of them had not historically been observed occurring in any DNS

records at all. For example, for 37.61.54.158, which is the most frequently observed IP address in the set of 9 repeat incorrect IP addresses, up until 2010 there were no domains which returned this IP address. Starting in 2010 there were 1494 domains, then in 2011 there were 6111 domains, in 2012 there were 2008 domains, etc. This process begun on June 30th 2010, starting with the domain wdxxx.com. Throughout July 2010 this increased to 55 domains, and by the end of the year DNS servers in China were returning this IP address for at least 1494 domains.

- 2009 - 0 domains
- 2010 - 1494 domains
- 2011 - 6111 domains
- 2012 - 2008 domains
- 2013 - 1556 domains
- 2014 - 1000000+ domains
- 2015 - 3501 domains
- 2016 - 754664 domains (as of July)

Note that these values include subdomains. 2014 and 2016’s high figures are caused by domains using highly variable strings as subdomains. In 2014 the domain jpteacn had over 1,000,000 listed subdomains, ranging from a.jpteacn to 999999999999.jpteacn. In 2016 the domain jjj.com has 744492 listed subdomains, mostly made from subdomains such as zzz96706.jjj.com, zzz85965.jjj.com, etc. As the GFW blocks each subdomain individually, these appear to be attempts to avoid its poisoning.

We scanned each of these IP addresses, using both ICMP echo requests and TCP SYN requests to all possible ports. There was no response to any packet we sent. This either means that there is no host located at these IP addresses, or that if there is a host there the responses are filtered, for example either at the network interface or by an outbound firewall.

5. DISCUSSION

When we started this work we expected to find that the legitimate response from DNS servers within the GFW contained the correct address for filtered sites, or at least a different incorrect IP address than the one set by the GFW. Instead what we found indicates that the DNS servers themselves have had their results poisoned. Whilst this finding may be obvious to some, we could not find this discovery stated outright in the existing literature. Indeed, some of the papers working in this area appear to make the assumption that the underlying DNS infrastructure itself could be trusted.

Most of the previous analyses of the GFW’s DNS poisoning have discussed the effects of this poisoning on users within China, rather than on the DNS servers themselves. While there has been discussion of collateral DNS poisoning with recursive queries, this has been from the perspective of users outside of China[3][5][8].

We do not believe we are alone in our ignorance of this phenomenon. Indeed, several past studies of the GFW have proposed methods of avoiding poisoning that assumes the servers themselves can be trusted, without specifying the need to configure for the use of alternative DNS servers[12]. Suggested methods include:

- Using TCP for DNS queries
- Using UDP on a non-standard port
- Ignore the first received DNS response

- Identify and ignore poisoned responses

As public DNS servers within China appear to be poisoned themselves, none of these methods will work on their own as the infrastructure itself is poisoned. Instead users must also configure their local DNS resolver to point to an unpoisoned DNS server outside of the influence of the GFW. This includes not just those that are physically within China, but also those that could be affected by collateral censorship[3].

Although there have previously been disagreements about whether the GFW poisons results centrally or along border nodes[7][18][10][9], Anonymous provided strong evidence for the border theory in 2014[4]. They also found that only a small number of filtered requests within China are actively poisoned: 4% – 16%.

We postulate that these findings indicate that the primary use of the GFW's DNS poisoning isn't to poison DNS requests of users, but to corrupt the cache of the DNS servers. The majority of Internet users in China make DNS requests to regional servers, and these requests are unlikely to pass border ASs and receive poisoned responses. While much research focuses on the direct DNS poisoning and how to avoid it, direct poisoning of user queries appears to be a supplementary effect.

While we find the repeated use of 9 IP addresses in DNS responses interesting, we are unable to find any connection between them, or hosts located at these addresses. There is evidence that they have been in use since 2010, and that other IP addresses were used for the same purpose previously[12].

6. CONCLUSION

We performed an analysis of DNS responses from public DNS servers under the influence of the GFW. We focused not just on the poisoned responses, but also looked at the legitimate responses from DNS servers. We found that in many cases the legitimate responses were pointing to the same IP addresses as the poisoned responses, suggesting that the servers themselves may have been poisoned.

We also observed 9 incorrect IP addresses that are repeated from both legitimate DNS servers as well as the GFW infrastructure itself. We have yet to explore why the GFW is responding to DNS requests with these specific IP addresses, and have not found any evidence that there are hosts listening at these addresses.

Our findings indicate that even if the GFW does not poison a particular DNS request, if for example a request does not pass any poisoning nodes, the results are still unreliable. They indicate that several proposed methods for avoiding the GFW may not be sufficient alone and that additional steps must be taken, including the use of trusted servers outside of the control of the GFW.

We postulate that this indicates that the GFW's DNS poisoning technique is aimed less at users, and more at the DNS infrastructure itself. This is supported by evidence from Anonymous' 2014 review of the GFW[4]. We believe there is need for further investigation into the propagation of domain information on DNS servers within and around the GFW. This information is often not trustworthy, even when dealing with some of the most popular domains in the world such as google.com, facebook.com and wikipedia.org.

7. REFERENCES

- [1] Farsight Security - DNS Database. <https://www.dnsdb.info/>. Accessed: 2015-08-28.
- [2] Alexa - Actionable Analytics for the Web. <http://www.alexa.com/>, 2016.
- [3] Anonymous. The collateral damage of internet censorship by DNS injection. *ACM SIGCOMM CCR*, 42(3), 2012.
- [4] Anonymous. Towards a Comprehensive Picture of the Great Firewall's DNS Censorship. In *4th USENIX Workshop on Free and Open Communications on the Internet (FOCI 14)*, San Diego, CA, Aug. 2014. USENIX Association.
- [5] M. A. Brown, D. Madory, A. Popescu, and E. Zmijewski. Dns tampering and root servers. *Presentation, Renesys Corporation*, 2010.
- [6] R. Clayton, S. J. Murdoch, and R. N. Watson. Ignoring the great firewall of China. In *Privacy Enhancing Technologies*, pages 20–35. Springer, 2006.
- [7] J. R. Crandall, D. Zinn, M. Byrd, E. T. Barr, and R. East. ConceptDoppler: a weather tracker for internet censorship. In *ACM Conference on Computer and Communications Security*, pages 352–365, 2007.
- [8] M. V. Ereche. Odd Behaviour on One Node in I root-server, 2010. <https://lists.dns-oarc.net/pipermail/dns-operations/2010-March/005260.html>.
- [9] GFW Technology Review. <http://gfwrev.blogspot.com/>, 2010.
- [10] Online Censorship In China – GreatFire.org. <https://en.greatfire.org/>, 2016.
- [11] G. King, J. Pan, and M. E. Roberts. How censorship in China allows government criticism but silences collective expression. *American Political Science Review*, 107(02):326–343, 2013.
- [12] G. Lowe, P. Winters, and M. L. Marcus. The great DNS wall of China. *MS, New York University*. Accessed December, 21, 2007.
- [13] MaxMind: IP Geolocation and Online Fraud Prevention. <https://www.maxmind.com/>, 2016.
- [14] Public DNS Server List. <http://public-dns.info/>, 2016.
- [15] View DNS Info. <http://http://viewdns.info/>, 2016.
- [16] F. Weimer. Passive dns replication. In *FIRST conference on computer security incident*, page 98, 2005.
- [17] J. Wright. Regional variation in Chinese internet filtering. *Information, Communication & Society*, 17(1):121–141, 2014.
- [18] X. Xu, Z. M. Mao, and J. A. Halderman. Internet censorship in China: Where does the filtering occur? In *Passive and Active Measurement*, pages 133–142. Springer, 2011.
- [19] E. Zmijewski. Accidentally importing censorship. *Renesys Blog*, March, 30, 2010.