

Internet Censorship in China: Where Does the Filtering Occur?

Xueyang Xu, Z. Morley Mao, and J. Alex Halderman

Department of Computer Science and Engineering, University of Michigan,
2260 Hayward Street, Ann Arbor, MI 48109
{xueyang,zmao,jhalderm}@umich.edu
<http://www.cse.umich.edu>

Abstract. China filters Internet traffic in and out of the country. In order to circumvent the firewall, it is helpful to know where the filtering occurs. In this work, we explore the AS-level topology of China's network, and probe the firewall to find the locations of filtering devices. We find that even though most filtering occurs in border ASes, choke points also exist in many provincial networks. The result suggests that two major ISPs in China have different approaches placing filtering devices.

Keywords: Censorship, topology, network measurement.

1 Introduction

In this work, we explore where Intrusion Detection System (IDS) devices of the Great Firewall of China (GFC) are placed for keyword filtering at AS and router level. Knowing where IDSes are attached helps us better understand the infrastructure of the firewall, gain more knowledge about its behavior and find vantage point for future circumvention techniques.

China has the world's most complex Internet censorship system, featuring IP blocking, keyword filtering, DNS hijacking and so on [1]. IP blocking is the earliest filtering mechanism. It is easy to circumvent, because webmasters can always change their IP and DNS record. Besides, censors are very prudent to do DNS hijacking nowadays due to the risk of affecting the network in other countries [2]. In this paper, we focus on the most effective filtering mechanism of GFC, keyword filtering.

According to [4], the filtering occurs more at AS-level rather than strictly along the border routers. This paper answers the question whether all censorship occurs at border AS, and how filtering occurs inside those ASes. We first explore the AS-level topology of China's network. In this part, we explore which Chinese ASes are directly peered with foreign ones and which are internal ones. We call those peered with foreign network *border AS*, and the others *internal AS*. The resulting AS-level topology shows that the best vantage point to place filtering device is in the border ASes.

To find where IDS devices are attached at router level, we select a set of web servers in China and probe with HTTP GET packets that contain known

keywords. In order to find more filtering devices, we manually select web servers to ensure their geographical diversity, as opposed to previous work that uses top websites in search result. This diversity is desirable, because it helps us to find more routing paths across China, and with more paths, we can discover more filtering devices.

The result shows that most filtering devices are in the border ASes, but a small portion is not. It is possible that there is a trend of placing filtering devices outside of border ASes. The number of router interfaces that have filtering devices attached for CHINANET is stable since 2007, while the second largest filtering force CNCGROUP has increasing number of filtering interfaces. Moreover, CHINANET's filtering is decentralized, while CNCGROUP has their IDS devices mostly in the backbone. A decentralized placement of filtering devices can facilitate censor to monitor domestic traffic.

The rest of the paper is organized as follows. Section 2 introduces the related work on measurement of the China's network censorship. Section 3 presents our result on AS topology of China's network. We locate filtering devices at router level in Section 4 to find how they are related to AS-level topology and the device placement strategies of different ISPs. Section 5 concludes the paper.

2 Related Work

An early work in the censorship measurement field is [3]. This paper analyzes the keyword filtering mechanism of GFC, and is a good source of background knowledge. They claim that the mechanism is based on an out-of-band intrusion detection system at border routers. The system emits forged reset packet to both destination and source, but packets themselves go through the router unhindered. Therefore, both source and destination ignoring forged reset packet makes the system entirely ineffective. They also claim that the firewall does not maintain a state.

An influential paper in this field is [4]. In the measurement study part, the most significant discovery is that unlike commonly believed, the censorship system in China is like a panopticon, where filtering does not occur strictly at border routers, but rather more centralized at AS level. They find that some filtering occurs 13 hops past border. In our work, we provide a more fine-grained analysis of where those filtering devices are located, answering whether all filtering occurs at border AS, and where IDS devices are attached at router level. They also discover that the firewall is stateful, namely a GET packet with keyword itself will not trigger the firewall. Rather, a complete TCP handshake is required. This contradicts with [3]. The paper also demonstrates that the RST packets sent by IDS devices are more complicated than before. The TTL of RST is now crafted, so we cannot identify the location of IDS devices by simply looking at the TTL values. Therefore, we identify the location of filtering devices by sending probe packets with increasing TTL values, and see when we receive RST packets from censors, as proposed in this work.

The most recent work in this field is [5]. This paper reports the discontinuation of keyword filtering in HTTP response on most routes, while that in HTTP GET

request is still prevalent. They investigate whether the firewall has a state and yield a result that the firewall is stateful only in part of the country. All 3 works have conflicting views of whether the firewall is stateful. Their latest tests have done in August 2009.

3 China's AS-Level Topology

Crandall et al. [4] claims that the firewall is better described as a panopticon, where filtering does not strictly occur at the border and suggests that the filtering is more at AS-level. Inspired by their work, we want to explore whether only border ASes are involved in the filtering and how filtering occurs inside those ASes. This knowledge is important, because if internal routers also have filtering devices attached, censors would have the capabilities to monitor and filter domestic traffic, which is considered not true before. It is believed that Chinese censors do not filter domestic websites technically, presumably because of the heavy domestic traffic flow; rather, the domestic Internet censorship is about social control, human surveillance, peer pressure, and self-censorship. [15]

In this section, we provide a more comprehensive view of China's AS-level topology that lays the foundation for Section 4.

3.1 Methodology

The first step is to find the mapping between AS numbers that belongs to China and their corresponding IP prefixes. Finding the mapping between IP prefix and AS number is known to be hard. We propose a methodology that yields a coarse-grained result. We first get the list of ASes that headquarter in China from APNIC [6]. An estimated mapping between IP prefix and AS number is extracted from the archival file obtained from Routeview [7] and RIPE [8] collectors. For each prefix entry in the archival files, we claim that its corresponding AS is the last AS in AS_PATH.

We acknowledge that the methodology is an estimation, because 1) we do not address the inaccuracies introduced by router interfaces that have addresses belonging to neighboring AS, and 2) the list of ASes is incomplete as the assignment record from APNIC does not capture all traffic originated from China.

We take the archival data collected between May and June 2010 from all collectors of Routeview and RIPE. We parse more than 300 MRT files, and this effort yields 408,688 AS-prefix mappings. Among them, 11,824 are in China's address space. In 136 AS numbers assigned to China, we find 76 corresponding prefixes of them.

In order to get as many peerings between China and other countries as possible, we traceroute from PlanetLab [14] nodes all over the world. A script is written to traceroute from each PlanetLab node outside of China to each of 76 Chinese ASes that we have their corresponding IP prefixes. We take the first IP in a prefix as the sample IP to which we traceroute.

For each hop in the traceroute result, we attempt to map them back into AS number using our estimated mappings. For those that we fail to map back, the whois server of Team Cymru [9] that returns IP to ASN mapping is consulted.

From the traceroute result, we construct an estimated AS-level topology of China’s network. Once the first hop inside China’s address space is noted, we add its corresponding AS number to a graph and denote it as a border AS. The corresponding AS numbers of all following hops are also added and are denoted as internal ASes. In addition, we also include the immediate AS that precedes each border AS, annotated as external AS.

CIDR report [16] analyzes the BGP table within AS2.0 and generates an aggregation report for each individual AS. For each AS, the report contains a list of its adjacent ASes and its announced prefix. To include the result of CIDR report into our topology, we crawl its website. For the report of each AS, we download the list of its adjacent ASes. We use the largest AS in China (AS4134) to bootstrap, and do a breadth-first search over its adjacent AS list. The search terminates whenever we encounter an AS not belonging to China.

In the resulting topology graph, the names of ASes are obtained from [9]. We use the name to imply the ISP that an AS belongs to.

3.2 Results

We find 138 internal, 24 border and 92 external ASes. Our result shows 133 unique peerings with external ASes. Among them, 62 belong to CHINANET and 23 belong to CNCGROUP. These two ISPs possess 63.9% of China’s total peerings with other countries. Table 1 shows the breakdown of ISPs in China that have the most number of unique peerings with foreign ASes according to our experiment. The resulting topology serves as the foundation of the experiment in the second part of the paper, while the following are some interesting observations that are worth further investigation.

Table 1. Chinese ISP with most number of unique peerings to foreign AS

ISP	AS Numbers	Peerings
CHINANET	4134, 4809, 4812, 23724, 17638	62 (46.6%)
CNCGROUP	4837, 9929, 17621, 4808	23 (17.3%)
TEIN	24489, 24490	8 (6.0%)
CNNIC	37958, 24151, 45096	8 (6.0%)
CERNET	4538, 4789	9 (6.8%)
Other	9808, 9394, 4847, 7497, 9298, 23911	23 (17.3%)

It is observed that some border ASes do not peer with any internal AS at all. These include 37958, 24151, 45096, 24489 and 24490. The first three belong to CNNIC, the national Internet registry of China. Even though it is possible

that the lack of internal peering is due to our experimental error, we speculate that the CNNIC ASes are used for special purposes. A future work could be exploring whether these ASes have different filtering rules. Another owner of this kind of AS is Trans-Eurasia Information Network, the traffic through which should be transit traffic, which means that both the source and destination are not in China's address space. We do not expect to see filters being installed in Trans-Eurasia ASes.

Our result indicates that border ASes in this country are peered with at least 20 foreign countries. Among them, U.S. is the largest one that has a peering count of 52. Hong Kong and Japan follows U.S., and have 21 and 11 peerings respectively. This information is useful in future work to find whether GFC defines different policies for different countries.

3.3 Discussion

We then organize the resulting topology hierarchically. In order to do that, we select border ASes as roots and grow trees under them with internal ASes as children. The depth of the tree is only 2, meaning that to get to any AS we discovered in China, we only need to traverse at most 2 other ASes. In fact, only 18 out of 138 internal nodes are at level 2.

Most of the internal ASes (87.0%) are within direct reach of border ASes. The names of border ASes suggest that most of them belong to backbone, and there are just 24 of them. This implies that the best vantage points for efficient content filtering are in the border/backbone ASes since they can easily serve as choke points, given that IDS devices have enough power and the censors do not intend to monitor domestic traffic.

4 Locating Filtering Devices

As the key step of this study, we make efforts to find as many filtering devices as we can to see their relationship with AS topology. Before we get started, here we provide some brief background of the firewall. As suggested by [3], IDS devices are attached externally to routers and thus out-of-band. The IDS terminates TCP connection by sending multiple spoofed RST packets to both ends of the communication. Within a period after that, all traffic between these two parties is blocked by RST packets, no matter whether a keyword is included in the packets.

For detailed description of the behavior of GFC, please refer to [3]. In this section, we discover to router interfaces at which locations IDS devices are attached.

4.1 Statefulness of the Firewall

A firewall being stateful means that we need to establish a TCP connection with a legal handshake to trigger the firewall [3]. If we directly send a TCP packet

that contains an HTTP GET with a known keyword but without a handshake, a stateful firewall would not send any RST packet. On the other hand, if the firewall is not stateful, any TCP packet with keyword, regardless of the existence of TCP connection, would trigger it.

Previous works [3] and [4] have contradicting result of whether the firewall is stateful, and [5] claims that part of the firewall is stateful. After sending a single packet with known keyword to the first IP of 11,824 Chinese prefixes, we observe no firewall activity at all. Assuming that the firewall behaves the same for all IPs in a prefix, the result indicates that the firewall is now totally stateful.

The firewall being stateful is meaningful. It can at least make probing in this kind of studies difficult. With a stateful firewall, we need to find servers in China that accept TCP connection to determine the position of filtering devices, rather than just probe the first IP of all prefixes with a packet that contains a keyword. With a stateless firewall, we can easily get a comprehensive set of filtering devices by probing all prefixes. On the other hand, in a stateful firewall, it is time-consuming to find an active server in each prefix, because it requires port scanning. Therefore, a stateful firewall makes probing more difficult and reduces the completeness of this kind of measurement.

4.2 Websites Probed

Since we are unable to probe each prefix to get a complete list of filtering devices, it is necessary to select a set of websites that are in different part of the country to achieve better completeness. Most previous work selects websites from the top result from search engine. This is biased, because top websites are likely to be clustered in some big cities in China. A CNNIC report [13] states that 51.2% websites are in 5 provinces, and there are 32 provinces in China. The least represented 17 provinces only have 10.8% of total number of websites in China. Furthermore, 13 provinces have less than 1% representation. Therefore, we cannot achieve our goal of getting as many filtering device as possible by employing their methodology.

Consequently, we carefully select web servers geographically across the entire country to probe. Our list of website covers all provinces and three major ISPs in this country, CHINANET, CNCGROUP and CERNET. To cover all provinces, we gather a list that contains the websites of all provincial governments. This list is obtained from the website of the central government [10]. The list of websites of provincial branches of CHINANET and CNCGROUP is also collected from Google search. Moreover, from a Chinese web resource guide [11], we collect a number of popular local websites. Taking CERNET, which is not a public network but mainly serves academic institutes, into account, we include websites of many universities in different parts of the country into our list.

Our final list contains 1594 websites. To show that they are geographically diverse, we query the most popular IP geolocation database in China [17]. The result is shown in the Appendix.

4.3 Algorithm

We probe our list of Chinese websites described in 4.1.2 to find the location of filtering devices at router level. Our methodology is similar to the one used by [4] and [5]. In short, the algorithm sends probing packets that contain known keywords with increasing TTLs.

For each IP of websites in our list, we first determine if it is online and whether it accepts TCP connection by establishing a TCP connection and sending an innocuous HTTP GET request. If we receive RST packets or the connection timed out, we skip it and proceed to the next website. Otherwise, we establish another connection with it, but this time, we send an HTTP GET with a known keyword `falun` that triggers the firewall.

At this moment, we wait for 5 seconds for the connection to completely die down. This allows the real and spoofed RST exchanges among source, destination and the firewall to complete.

Since the firewall is already triggered and now all further traffic between two endpoints, no matter considered harmful or not, is blocked by the firewall for a period, a simple ACK packet would trigger the firewall. Therefore, we send ACK packets with increasing TTL, and stop whenever we receive RST from a filtering device and record the IP address revealed by the ICMP packet that the router interface to which the filtering device is attached sends. We skip and record the website if the keyword does not trigger the reset in case of whitelisted websites.

4.4 Results

We found 495 router interfaces that have filtering devices attached to in our experiment, 106 more than in [4]. The proportions of filtering interfaces that each ISP has are as follows: CHINANET has 79.4%; CNCGROUP possess 17.4%, and the rest 3.2% belong to other ISPs. We get largely identical proportion for CHINANET as in [4], but for CNCGROUP, our percentage is three times of their result. Our result suggests that the placement of IDS device of CHINANET is stable since 2007, and the filtering power of CNCGROUP is growing and now counts for almost one fifth of all filtering interfaces in the country. We can derive that the filtering capability of CHINANET is mature, as the increase in traffic has not made it too overloaded to force them adding more filtering interfaces for several years.

Table 2 shows what ASes the filtering devices belong to. We consult the whois server of Team Cymru [9] for IP to ASN mapping.

Not surprisingly, most of the filtering devices belong to the border ASes. However, we find that some of them are in internal ASes. The proportion is small (2.9%), so it is prone to errors introduced by inaccurate IP to AS number mapping. However, it is still worth noting. We will continue to monitor this number, to see if there exists a trend that censors deploy more and more filtering devices to internal ASes.

All except for two internal filtering interfaces belong to CHINANET, and none belongs to CNCGROUP. Since this is particularly questionable, we examine our

Table 2. ASes that contain filtering devices

AS Number	AS Name	Number of Filtering Interfaces
Border ASes		481
4134	CHINANET-BACKBONE	374
4812	CHINANET-SH-AP	9
4837	CHINA169-BACKBONE CNCGROUP	82
9929	CNCNET-CN	4
4538	ERX-CERNET-BKB	4
9808	CMNET-GD	5
9394	CRNET	3
Non-border ASes		14
23650	CHINANET-JS-AS-AP	4
17785	CHINATELECOM-HA-AS-AP	4
37943	CNNIC-GIANT	3
38356	TIMENET	1
17633	CHINATELECOM-SD-AS-AP	1
4813	BACKBONE-GUANGDONG-AP	1

traceroute log more carefully. As a result, we find that if the first router interface in China’s address space belongs to CHINANET, it rarely conducts any filtering. Pursuing this further, we find that many filtering router interfaces do not seem to belong to the same prefix as that of the first few router interfaces into the country, so we whois CHINANET’s filtering interfaces to find more.

The result is interesting. Despite the name of AS4134 suggests, only 49 of 374 filtering interfaces actually belong to the backbone of CHINANET, and the rest of them are actually belong to provincial branch companies of CHINANET. In AS4134, we find that 16 provinces have their own filtering devices. Counting Shanghai that is not represented in AS4134 but has its own AS number, 80% of 21 provinces that CHINANET serves [12] do their own filtering. The provinces that are observed not having their own filtering are Shaanxi, Gansu, Qinghai and Ningxia. According to a CNNIC report [13], the number of IP addresses in these 4 provinces only counts 2.5% of the nation’s total number of IPs. Table 3 shows where the filtering devices are located in AS4134. We only list the provinces that are in the service area of CHINANET.

This implies that CHINANET, instead of filtering strictly along the border, offloads the burden to its provincial network. On the other hand, CNCGROUP has most of its filtering devices in the backbone rather than provincial network, and all its filtering is done within very few hops into China’s address space. We also whois the IP address of those filtering devices, and find that 74 out of 82, or 90% of filtering devices belongs to the backbone of CNCGROUP. This indicates that two major ISPs in China have different approaches placing their filtering devices.

The total bandwidth of CHINANET’s international connection is 616703Mbps, and that of CNCGROUP is 330599Mbps [13]. Moreover, the number of peerings with foreign AS of CHINANET is 3 times of that of CNCGROUP. CHINANET, as a larger operator that has international bandwidth 2 times of

Table 3. Locations of filtering devices in AS4134

Province	# Devices	Percentage
Backbone	49	13.10%
Guangdong	84	22.46%
Fujian	29	7.75%
Hunan	28	7.49%
Hubei	24	6.42%
Sichuan	22	5.88%
Yunnan	21	5.61%
Guangxi	19	5.08%
Jiangsu	19	5.08%
Zhejiang	15	4.01%
Guizhou	14	3.74%
Jiangxi	14	3.74%
Hainan	11	2.94%
Chongqing	10	2.67%
Anhui	6	1.60%
Unidentified	6	1.60%
Xinjiang	2	0.53%
Tibet	1	0.27%

CNCGROUP, needs to filter more network traffic. Placing all filtering devices in backbone might have created a bottleneck for CHINANET, and allowed some unwanted traffic to go through. This might partly explain why they have different IDS placement strategies.

Another implication is that the filtering devices being in the provincial network allows censor to inspect inter-province traffic. Even though there is no evidence that they are doing this right now, this arrangement makes the future deployment of stricter firewall that censors domestic traffic easier.

5 Conclusion

Chinese censors impose strict restrictions on international Internet traffic. In order to understand the national-scale intrusion detection system better, this is the first study dedicated to explore both AS and router-level structures of China's censored network.

As a preparation, the first part of the paper presents our approximate result of China's AS-level Internet topology. We manage to collect the peering among 265 China-related ASes. In the second part of our work, we probe the firewall in an attempt to gather as many filtering interfaces as we can and to relate AS topology to the location of those filtering devices. We find that most filtering occurs in border ASes, but two major filtering ISP's have different approaches placing their filtering devices. CHINANET does not place most of its filtering devices in its backbone, but rather distribute the work to provincial networks. This makes censoring domestic traffic easier.

References

1. The Great Firewall Revealed, <http://www.internetfreedom.org/files/WhitePaper/ChinaGreatFirewallRevealed.pdf>
2. China's Great Firewall spreads overseas, <http://www.networkworld.com/news/2010/032510-chinas-great-firewall-spreads.html>
3. Clayton, R., Murdoch, S., Watson, R.: Ignoring the Great Firewall of China. In: Danezis, G., Golle, P. (eds.) PET 2006. LNCS, vol. 4258, pp. 20–35. Springer, Heidelberg (2006)
4. Crandall, J., Barr, E.: ConceptDoppler: A Weather Tracker for Internet Censorship. In: 14th ACM Conference on Computer and Communications Security (2007)
5. Park, J., Crandall, J.: Empirical Study of a National-Scale Distributed Intrusion Detection System: Backbone-Level Filtering of HTML Responses in China. In: The Proceedings of the 30th International Conference on Distributed Computing Systems (2010)
6. APNIC delegated internet number resource, <http://ftp.apnic.net/stats/apnic/delegated-apnic-latest>
7. University of Oregon Route Views Archive Project, <http://archive.routeviews.org/>
8. RIPE NCC Projects, <http://www.ripe.net/projects/ris/rawdata.html>
9. Team Cymru, <http://www.team-cymru.org/Services/ip-to-asn.html>
10. The Central People's Government of the People's Republic of China, http://www.gov.cn/zwgk/2008-04/23/content_952239.htm
11. Dao Hang Wang, <http://www.daohang.com/>
12. China Comservice, <http://www.chinaccs.com.hk/eng/about/history.htm>
13. The 26th Statistical Reports on the Internet Development in China, <http://www.cnnic.cn/uploadfiles/pdf/2010/7/15/100708.pdf>
14. PlanetLab, <http://www.planet-lab.org>
15. The Connection Has Been Reset, http://msl1.mit.edu/furdlog/docs/atlantic/2008-03_atlantic_fallows_chinese_firewall.pdf
16. CIDR Report, <http://www.cidr-report.org/as2.0/>
17. Chunzhen IP geolocation database, <http://www.cz88.net>

Appendix: Geographical Locations of Probed Websites

This is a list of provinces represented in our probed websites. It is a result after querying the database of [17] dated December 30, 2010. The number in bracket is the number of probed websites in that province.

Shanghai (24), Yunnan (36), Inner Mongolia (29), Beijing (94), Hubei (48), Guangdong (115), Fujian (59), Jilin (25), Sichuan (65), Liaolin (63), Tianjin (16), Ningxia (15), Anhui (43), Shandong (73), Shanxi (31), Guangxi (39), Xinjiang (28), Jiangsu (82), Jiangxi (45), Hebei (48), Henan (63), Zhejiang (69), Hainan (21), Hunan (44), Gansu (35), Shaanxi(36), Tibet (4), Guizhou (28), Chongqing (15), Qinghai (5), Heilongjiang (28), Other (268)