Journal of Information Technology Impact
Vol. 11, No. 2, pp. 145-152, 2011

Effectiveness of Internet Content Filtering

Davis Gossett [1] Jack D. Shorter [2]
Texas A&M University – Kingsville
USA

Abstract

This paper describes some recent implementations of large-scale Internet Content Filtering on the ISP or national level. It discusses the reasoning and purpose for establishing these systems and the people responsible for them. Cultural, ideological, and legal problems with content filtering programs will be explored. Technical reasons for the ineffectiveness of these implementations will then be discussed. Some recent examples of content filtering appear to be quite prevalent in all levels of education. According to many educators and most administrators there is a feeling for the need of content filtering on the Internet. The problem is when that filtering goes too far. This paper hopes to describe the myriad of fallacies involved with any filtering implementation by pointing to numerous current examples and their technical shortcomings.

Keywords: ISPs, firewalls, blocked websites.

Introduction

The Internet has revolutionized the way that people disseminate information around the globe. It offers instantaneous access to nearly any kind of digital resource. Unfortunately, this applies to both legal and illegal content equally. The Internet was built from the ground up as a decentralized network with the ability to reroute itself around physical damage or censorship. This poses a significant challenge to those organizations which are responsible for policing the trafficking of illegal materials.

The primary proponents of Internet Content Filtering are government and law enforcement officials, though a few ISPs (most notably in the UK) have voluntarily implemented content filtering programs of their own volition. The vast majority of countries have laws against crimes such as child pornography, defamation, and fraud. Selected governments who actively attempt to shut down this type of illegal operation run into significant problems when attempting to enforce laws on the Internet. The Internet has no legal boundaries or jurisdictions, so it is impossible for law enforcement agencies to take physical action against the vast majority of websites which break their laws. This is the case, even if a similar law applies wherever the site is physically hosted. Law enforcement agencies are left with only the option of discouraging (and potentially prosecuting) the potential customers of these operations on their own soil through some kind of monitoring or filtering mechanism, if they wish to fight these crimes at all. Content filters can also be used to prevent potential victims from reaching sites run by foreign scamming operations and identity thieves.

Some Examples of Content Filtering

When you think of International content filtering, you must always include China when giving examples of egregious blocking and filtering with government support. Currently China is the home to the world's largest population of Internet users, numbering around 446 million by the end of 2010. Unfortunately, the country's Internet environment remains extremely restrictive. It is characterized by a sophisticated, multi-layered control apparatus. Recently the system has been further enhanced, institutionalized, and decentralized. The blocks on social network sites such as Facebook and Twitter are now permanent. The Chinese Internet appears to be more like an Intranet at this time. China has also emerged as a key global source of cyber-attacks, with targets ranging from groups reporting on Chinese human rights abuses to international financial, defense, and technology companies. The above restrictions were offset somewhat by the Internet's continued growth as a primary source of news, a forum for discussion, and a mobilization channel for many Chinese (Kelly & Cook, 2911).

Some recent examples of content filtering appear to be quite prevalent in all levels of education. According to many educators and most administrators there is a feeling for the need of content filtering on the Internet. The problem is when that filtering goes too far. Universities are starting to embrace social networking (most notably - Facebook and Twitter) as a way to keep their students engaged and aware of what is happing at both the college and university level. Many of our colleagues are looking for ways to utilize this technology in their classroom. At the very same time, university administrators (due to economic short falls in most states) are restricting access to these very same sites due to perceived cost savings (at least with wireless hookups). It appears that the decision to filter or not to filter is causing quite a conundrum. It is almost a sure bet that most Higher Education students will want to call the shots and demand open access. After all, educational institutes are like any other business, they have to keep the customer happy (Lota, 2011).

It appears that China is not the only country in the eastern Asian area that is working hard to suppress freedom of expression on the Internet. Thailand has seen its share of Internet censorship since the military coup of 2006. While obstacles to access are profound in Thailand, the fact that you have a military regime has got to be impeding growth and freedom of use of the Internet in this country. However, one of the premises of this paper is that in the long run content filtering will not keep the general population from eventually seeking and obtaining informational freedom. This eventually turns into a longing for the eradication of a repressive social structure in many countries today. Thailand appears to be no exception to this premise. Once again it appears that advanced web applications such as the video-sharing site YouTube, the social networking site Facebook, the Twitter microblogging platform, and international blog-hosting services like Blogger are freely available in Thailand. Therefore, freedom of thought and expression are alive and well in Thailand (Kelly & Cook, 2011).

Public opinion concerning filtering schemes is generally quite negative. In 2009, backlash against the Australian government's proposed filtering scheme was a good example of this sentiment. According to a national telephone poll, only 5 percent of the Australian population wanted ISPs to be responsible for protecting children from on-line predators, and only 4 percent wanted the responsibility to fall on the shoulders of the government (Moses, 2011). Over 60 percent of the people interviewed by Netspace, an Australian ISP, about the subject were strongly opposed to mandatory filtering, with only 6.3 percent of the 10,000 people interviewed agreeing strongly with the proposed policy. To many, these numbers are astonishing.

One of the primary reasons for public discontent regarding Internet Content Filtering is the general lack of transparency involved in these operations. Making available a list of addresses which are blocked as part of the list can have the opposite effect of helping those who wish to find the blocked content to circumvent the restrictions. The Australian Communications and Media Authority, which is the governmental body responsible for managing the list of blocked foreign child pornography sites in the Australian trials, has refused to disclose exactly which sites are on the list (Foo, 2009). However, they did release some disturbing statistics regarding the makeup of the sites that were blocked. Of the 1370 websites blocked under their program, only 864 of the blocked sites were classified as child pornography. Over 400 of the sites which were blocked by the program were still pornographic, but were classified as X18+. No reason was given as to why these perfectly legal sites had been blocked, but the lack of judicial review over the list does not sit well with most of the affected citizens.

The Australian filtering trials are not alone in their lack of oversight regarding the list of blocked websites. In late 2008, six British ISPs simultaneously added Wikipedia.org to their child pornography block lists (Metz, 2008). The reason for this action was due to an article containing a controversial album cover by a 70's band which contains a picture of a nude girl. One of the side effects of implementing this block, however, was that the transparent proxies used to implement the ban forced almost all of the British traffic going to Wikipedia to be routed through a small handful of IP addresses. This caused serious issues for the administrators of the website. They could not individually ban people who break their rules any longer, meaning that they must either allow all defamation from these addresses to continue unhindered, or ban all users of those ISPs from editing Wikipedia. Either of these outcomes has far-reaching negative consequences for all users of one of the most important and influential websites on the Internet. However, as everyone knows Wikipedia has caused a maelstrom due to the fact that they keep releasing top secret documents obtained in illegal ways under the guise of freedom of information and the fact that they feel the United States is the perpetrator of great evil in the world. If the site could have been shut down or filtered it would have been. It was not.

The real reason that most detractors of Internet Content Filtering are wary of governmental control over the Internet is because of the possibility of using these systems to serve propagandists or some groups political goals. This fear is well-founded, as nearly every large-scale filtering scheme has indeed gone beyond blocking only illegal web pages and added legitimate sites to their filter lists, either mistakenly or intentionally. One of the most startling examples of this occurred in early 2008, when Nikki (2008) began to protest against the web filter put in place in his country of Finland. He began performing research on the nature of the filter and compiled a list of the banned sites, which had been kept secret, on his own web page. A few weeks later, his site was added to the block list and became inaccessible to any Finnish residents. Not only did this overstep the bounds of the content filtering system's purpose and trample over his right to free speech, but it was diametrically opposed to the entire purpose of content filtering in the first place. Since the site was run by a Finnish owner on Finnish soil, any Finnish laws broken by the site could be easily enforced by the local law.

False Positives

Another issue concerning the lack of oversight for Internet filter lists is the possibility of accidental false positives. While it is understandable for some websites to slips through the cracks, there have been a number of cases where the blocked site is both very well-known and clearly has no ties to any illegal, immoral, or politically-charged content. The Finnish content filtering system erroneously blocked the W3C website in a startling display of ineptitude (Lehto, 2008). The W3C is one of the primary standards bodies which maintain important universal formats such as HTML and XML for the Internet. The blockage of this page would seem to point towards one of two things. Either there is a total lack of competent review for sites which have been submitted to the body responsible for the filtering program, or some kind of automated program is responsible for analyzing and adding sites to the blacklist and it does not function correctly. Since it is unlikely that the W3C website was intentionally submitted into the system, and because the W3C URL is transparently embedded in the source code for most properly-written HTML pages on the Internet, the latter possibility seems far more likely.

While countries with freedom of speech are forced to add shady entries to their block lists behind closed doors, many of those without similar protections have no qualms about doing it in the open. By far, the most extensive Internet Content Filtering system in place today is deployed by the PRC government. Nicknamed the “Great Firewall of China” by most westerners, the filtering system goes beyond simple website blocks and actively filters other Internet services such as Instant Messaging and E-mail. The program has gone so far as to make a concerted effort to block all “lewd” content on the Internet by employing thousands of workers to report offenders (Powell, 2009). While this is obviously a monstrous undertaking, the filtering is done on such a massive scale that they may actually succeed in this endeavor. The Chinese filtering system is not limited to illegal or “immoral” content, however. Most western news sources, such as the *New York Times* web site, are inaccessible from behind the Chinese firewall (Bradsher, 2008). One of the primary goals of the filter is to prevent the flow of information which is “politically harmful”. They have even coerced outside companies such as Google to remove all references to events such as the Tiananmen Square massacre, the Dalai Lama, and Falun Gong. During the 2008 Olympics, the IOC forced the PRC to temporarily lift the filtering restrictions for the duration of the event (Taylor, 2008). During this period, the English-language version of Wikipedia was accessible to users in China, but not the Chinese-language version. These changes were only temporary, however, and the firewall has since returned to its previous state.

Costs

Although Internet Content Filtering programs are both unpopular and some will argue ineffective, there are nevertheless a large number of them in place around the world. The costs of implementing these programs are often passed straight down to the consumer, driving up the cost of Internet access for all users. The Australian filtering program was budgeted to cost as much as

$128 million (Clarke, 2009). However, this doesn't take into consideration the indirect costs incurred by ISPs such as extra help desk staff hired due to increased calls from customers who cannot access websites. The costs of wider-scale deployments such as the one in China can only be speculated about, but it is most likely an order of magnitude higher, at the very least.

Aside from direct monetary costs, one of the biggest costs incurred by content filters is decreased network performance. In the trials conducted for the proposed Australian filtering system, some figures showed that regular Internet browsing could be slowed down by as much as 87 percent by the cheaper solutions that were being explored first (Pillion, 2009). Although other filtering solutions are not nearly as detrimental to network performance as these, the improved performance comes with a price tag. Another important cost to consider when implementing a content filtering system is the potential for false positives. Even the best of the content filters being considered in the Australian operation still had a 1% false positive rate (Meloni, 2008). The potential economic damage caused by arbitrarily blocking the websites of legitimate businesses can only be speculated about, to say nothing of the unquantifiable value of other sites which are accidentally blocked by these filters.

Technical Limitations

While there are many disputes over whether Internet Content Filtering should be performed, there is one point that can be almost universally agreed upon by anyone with experience in the field of Information Technology. Internet Content Filtering simply does not work. Trying to prevent the flow of information on the Internet is a costly arms-race at best, and a complete waste of time and money with no chance of success at worst. No matter what kind of filtering approach is used, those who are determined to access the blocked content can almost always skirt around the restrictions with very little effort. Unfortunately the bureaucrats in charge of said content filtering can claim all is well and the children are safe, whether they really are or not.

The most basic and obvious method of content filtering is through manipulation of DNS records. Most home users receive their DNS server assignments through the DHCP protocol from their ISP. The ISP can configure its DNS servers to resolve certain addresses to a “block-page”, or simply to nothing at all. Some free DNS services such as OpenDNS allow for private homes and businesses to easily manage this kind of filter with a few short clicks (Diehl, 2009). However, there is nothing stopping users from simply changing their computer to point towards a different, non-filtered DNS server. Another issue with DNS filtering is its tendency towards “scatter-shot” blocking. It is required to block entire domain names at once, leading to large amounts of collateral damage in certain hosting arrangements. In essence, this approach is equivalent to erasing records out of a local phone book. The numbers are still there for you to call if you know what they are, and nothing is stopping you from picking up a different phone book and looking it up yourself.

The next step in the filtering war is to block unwanted websites by their IP address instead of their FQDN. This requires significantly more investment and overhead, as this means that you must analyze the headers of all packets that pass through your network and compare them to a list of bad addresses. According to an interview of a policy chief from a peering cooperative by The Register, many smaller ISPs in Britain refuse to implement the “Cleanfeed” filtering system deployed by many of the larger ISPs because "it would mean spending a lot of money on something that simply does not work." (Williams, 2009). One of the ISPs which has agreed to deploy the system has stated that it has only done so in order to prove that it cannot work to the politicians pushing for its adoption (Strover, 2009). The Cleanfeed system first checks against a list of "suspect" IP addresses as they pass through a gateway router. If there is a match, it passes the packets on to a proxy server which then does more granular page-level checking and blocks only those pages which contain the blacklisted material. However, this method of filtering is bypassed just as easily as regular DNS filtering. Since there is no deep-packet inspection going on, websites can simply be hosted on ports other than 80 to avoid the blocks by disguising themselves as something other than web traffic. Users can also connect to websites through proxy servers outside of the filtering system, who will bounce the requests in their name to avoid the blacklist.

Some Conclusions

The next logical step in the filtering game, then, is to perform deep-packet inspection in order to stop web traffic to blacklisted sites which do not occur on the standard ports. Inspecting the entire packet in order to determine the type of traffic is a relatively simple, but expensive task. It requires orders of magnitude more processing time than simple header inspection since the payload contained in a packet is longer than the header. This process vastly increases hardware costs and decreases network performance significantly. Even after investing in the equipment required to filter content in this manner, however, it is yet again bypassed incredibly easily. This time, all that is required to beat this mechanism is standard encryption, either via https or a VPN/SSH tunnel. Blocking or proxying all encrypted content breaks a significant portion of the Internet, while also opening users up to potential security risks.

Another issue to address is the ease of access to proxy servers. There are so many of them on the Internet, that adding all of them to the block list is essentially an impossible task. Although proxy servers are important for Internet anonymity, we can assume that by this point in the arms race, all of your customers are being treated like criminals anyway so this is unlikely to be a consideration. Although employing the number of people required to track down and block every proxy server on the Internet is not feasible for most ISPs, the Great Firewall of China has worked hard at accomplishing this task. However, there is nothing preventing people from running personal or private non-advertised proxy servers outside the boundary of the firewall. It is essentially impossible to prevent private proxy servers from being used to circumvent nearly any firewall scheme.

However, personal proxy servers are probably beyond the means of your average overly-curious Internet user who might run afoul of Internet Content Filters. Since packet headers have to remain unencrypted in order to remain usable by the ISP routers, and we now have a complete list of banned IP addresses for both websites and proxy servers. The battle seems essentially won against all but the most dedicated users. Well, not quite. Of course, the Internet has invented a number of ways to circumvent even a firewall as extensive as this one. Programs such as Freenet have created a decentralized, encrypted peer-to-peer network which sits on top of the regular Internet infrastructure. Other programs such as Tor can be easily integrated into web browsers in the form of plugins, which give one-click anonymity through a distributed encrypted network. In order to block users from using these meta-networks to bypass content filtering, the filtering operation would need access to each and every computer running the software. This is essentially QED (for now) when it comes to the filtering arms-race.

References

Bradsher, K. (2008). China Blocks Access to the Times's Website. *The New York Times.* Retrieved October 6, 2009, from http://www.nytimes.com/2008/12/20/world/asia/20china.html?_r=2&ref=todayspaper

Clarke, T. (2009). Xenophon speaks out against Internet content filtering. *Computerworld.* Retrieved November 8, 2009, from http://www.computerworld.com.au/article/278285/xenophon_speaks_against_internet_content_filtering?fp=4194304&fpid=1

Diehl, M. (2009). Web Content Filtering with OpenDNS. *Linux Journal*. Retrieved November 8, 2009, from http://www.linuxjournal.com/content/web-content-filtering-opendns

Foo, F. (2009). Row Over Web Blacklist. *Australian IT.* Retrieved November 7, 2009, from http://www.australianit.news.com.au/story/0,24897,25096792-15306,00.html

Kelly, S., & Cook, S. (Eds.). (2011). Freedom on the Net 2011 – A Global Assessment of Internet and Digital Media. *Ifap*. Retrieved May 22, 2011, from http://www.ifap.ru/library/book497.pdf

Kelly, S., & Cook, S. (Eds.). (2011). Freedom on the Net 2011 – A Global Assessment of Internet and Digital Media. *Scribd*. Retrieved May 23, 2011, from http://www.scribd.com/doc/53256748/Freedom-on-the-Net-2011/

Lehto, T. (2008). The error was corrected quickly today. *Tietokone*. Retrieved March 7, 2009, from http://translate.google.com/translate?u=http%3A%2F%2Fwww.tietokone.fi%2Fuutta%2Fuutinen.asp%3Fnews_id%3D35075%26tyyppi%3D1&hl=en&ie=UTF-8&sl=fi&tl=en

Lota, S. (2011). Internet Content Filtering in Education. *Sukhwant's Blog*. Retrieved May 23, 2011, from sukhwantlota.blogspot.com/2011/02/internet-content-filtering-in-Eduction.html

Meloni, M. (2008). The high price of internet filtering. *ABC News.* Retrieved Sept. 8, 2009, from http://www.abc.net.au/news/stories/2008/10/24/2399876.htm

Metz, C. (2008). Brith ISPs censor Wikipedia over 'child porn' album cover. *The Register*. Retrieved August 6, 2009, from http://www.theregister.co.uk/2008/12/07/brit_isps_ censor_wikipedia/

Moses, A. (2009). Web Censorship Plan Heads Towards a Dead End. *The Sydney Morning Herald*. Retrieved November 6, 2009, from http://www.smh.com.au/articles/2009/02/26/1235237810486.html

Nikki, M. (2008). Lapsiporno.info and the Finnish Internet censorship. *Lapsiporno.info*. Retrieved July 5, 2009, from http://lapsiporno.info/english-2008-02-15.html

Pillion, A. (2009). Mandatory filtering won't slow net access. *Australian IT*. Retrieved October 7, 2009, from http://www.australianit.news.com.au/story/0,24897,25040381-5013038,00.html

Powell, G. (2009). China makes arrests to stop 'vulgar' content. *Tech.Blorge*. Retrieved October 6, 2009, from http://tech.blorge.com/Structure:%20/2009/01/16/china-makes-arrests-to-stop-internet-porn/

Strover, R. (2009). Clean Feed. *PC Update*. Retrieved November 7, 2009, from http:// www.melbpc.org.au/ pcupdate/2902/2902article7.htm

Taylor, S. (2008). China allows access to English Wikipedia. *Reuters*. Retrieved March 8, 2009, from http://in.reuters.com/article/technologyNews/idINIndia-32865420080405

Williams, C. (2009). Small ISPs reject call to filter out child abuse sites. *The Register*. Retrieved November 22, 2009, from http://www.theregister.co.uk/2009/02/25/iwf_small_ isps/

[1] Davis Gossett was an Undergraduate Research Assistant at the College of Business Administration, Texas A&M University – Kingsville. He can be reached at: 1115 University Blvd., CBA Room 119, Kingsville, Texas 78363, U.S.A. Email: jorrix@gmail.com.

[2] Dr. Jack D. Shorter is a Professor of Information Systems at the College of Business Administration, Texas A&M University – Kingsville. He can be reached at: 1115 University Blvd., CBA Room 119, Kingsville, Texas 78363, U.S.A. Email: jack.shorter@tamuk.edu; Phone: + (361) 593-2130; Fax: + (361) 593-3708.