

3

Tools and Technology of Internet Filtering

Steven J. Murdoch and Ross Anderson

Internet Background

TCP/IP is the unifying set of conventions that allows different computers to communicate over the Internet. The basic unit of information transferred over the Internet is the Internet protocol (IP) *packet*. All Internet communication—whether downloading Web pages, sending e-mail, or transferring files—is achieved by connecting to another computer, splitting the data into packets, and sending them on their way to the intended destination.

Specialized computers known as *routers* are responsible for directing packets appropriately. Each router is connected to several communication *links*, which may be cables (fiber-optic or electrical), short-range wireless, or even satellite. On receiving a packet, the router makes a decision of which outgoing link is most appropriate for getting that packet to its ultimate destination. The approach of encapsulating all communication in a common format (IP) is one of the major factors for the Internet's success. It allows different networks, with disparate underlying structures, to communicate by hiding this nonuniformity from application developers.

Routers identify computers (hosts) on the Internet by their IP address, which might look like 192.0.2.166. Since such numbers are hard to remember, the domain name system (DNS) allows mnemonic names (domain names) to be associated with IP addresses. A host wishing to make a connection first looks up the IP address for a given name, then sends packets to this IP address. For example, the Uniform Resource Locator (URL) `www.example.com/page.html` contains the domain name “`www.example.com.`” The computer that performs the domain-name-to-IP-address lookup is known as a DNS resolver, and is commonly operated by the Internet service provider (ISP)—the company providing the user with Internet access.

During connection establishment, there are several different ways in which the process can be interrupted in order to perform censorship or some other filtering function. The next section describes how a number of the most relevant filtering mechanisms operate. Each mechanism has its own strengths and weaknesses and these are discussed later. Many of the blocking mechanisms are effective for a range of different Internet applications, but in this chapter we concentrate on access to the Web, as this is the current focus of Internet filtering efforts.

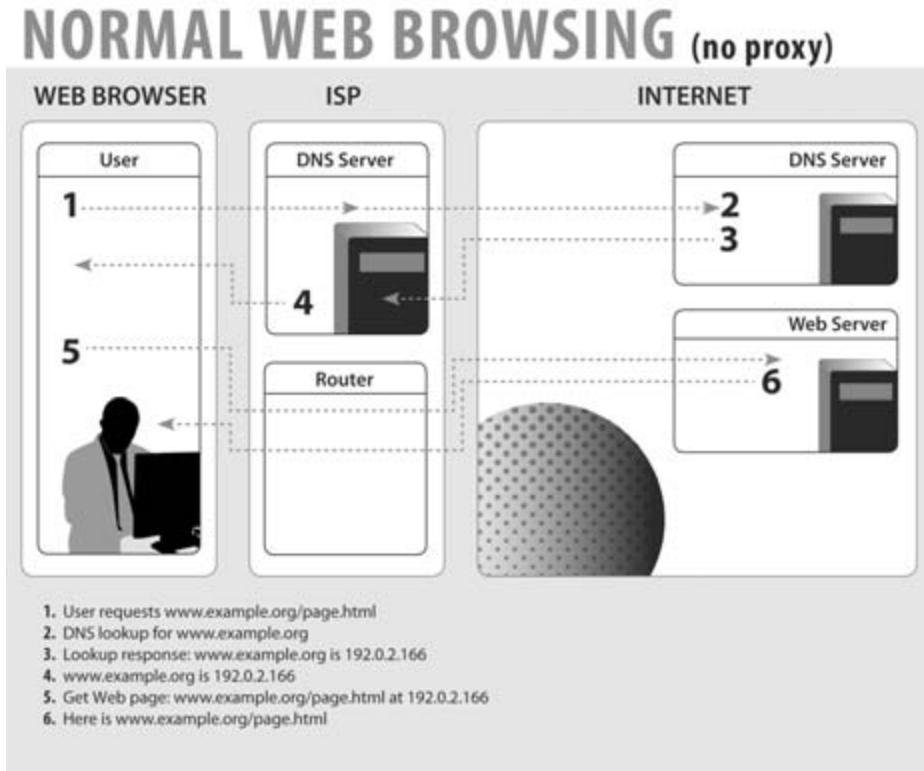


Figure 3.1

Steps in accessing a Web page via normal Web browsing without a proxy.

Figure 3.1 shows an overview of how a Web page (`http://www.example.com/page.html`) is downloaded. The first stage is the DNS lookup (steps 1–4), as mentioned above, where the user first connects to their ISP’s DNS resolver, which then connects to the Web site’s DNS server to find the IP address of the requested domain name—“`www.example.com`.” Once the IP address is determined, a connection is made to the Web server and the desired page—“`page.html`”—is requested (steps 5–6).

Filtering Mechanisms

The goals of deploying a filtering mechanism vary depending on the motivations of the organization deploying them. They may be to make a particular Web site (or individual Web page) inaccessible to those who wish to view it, to make it unreliable, or to deter users from even attempting to access it in the first place. The choice of mechanism will also depend upon the

capability of the organization that requests the filtering—where they have access to, the people against whom they can enforce their wishes, and how much they are willing to spend. Other considerations include the number of acceptable errors, whether the filtering should be overt or covert, and how reliable it is (both against ordinary users and those who wish to bypass it). The next section discusses these trade-offs, but first we describe a range of mechanisms available to implement a filtering regime.

Here, we discuss only how access is blocked once the list of resources to be blocked is established. Building this list is a considerable challenge and a common weakness in deployed systems. Not only does the huge number of Web sites make building a comprehensive list of prohibited content difficult, but as content moves and Web sites change their IP addresses, keeping this list up-to-date requires a lot of effort. Moreover, if the operator of the site wishes to interfere with the blocking, the site could be moved more rapidly than it would be otherwise.

TCP/IP Header Filtering

An IP packet consists of a *header* followed by the data the packet carries (the *payload*). Routers must inspect the packet header, as this is where the destination IP address is located. To prevent targeted hosts being accessed, routers can be configured to drop packets destined for IP addresses on a blacklist. However, each host may provide multiple services, such as hosting both Web sites and e-mail servers. Blocking based solely on IP addresses will make all services on each blacklisted host inaccessible.

Slightly more precise blocking can be achieved by additionally blacklisting the *port number*, which is also in the TCP/IP header. Common applications on the Internet have characteristic port numbers, allowing routers to make a crude guess as to the service being accessed. Thus, to block just the Web traffic to a site, a censor might block only packets destined for port 80 (the normal port for Web servers).

Figure 3.2 shows where this type of blocking may be applied. Note that when the blocking is performed, only the IP address is inspected, which is why multiple domain names that share the same IP address will be blocked, even if only one is prohibited.

TCP/IP Content Filtering

TCP/IP header filtering can only block communication on the basis of where packets are going to or coming from, not what they contain. This can be a problem if it is impossible to establish the full list of IP addresses containing prohibited content, or if some IP address contains enough noninfringing content to make it unjustifiable to totally block all communication with it. There is a finer-grained control possible: the content of packets can be inspected for banned keywords.

As routers do not normally examine packet content but just packet headers, extra equipment may be needed. Typical hardware may be unable to react fast enough to block the infringing packets, so other means to block the information must be used instead. As packets

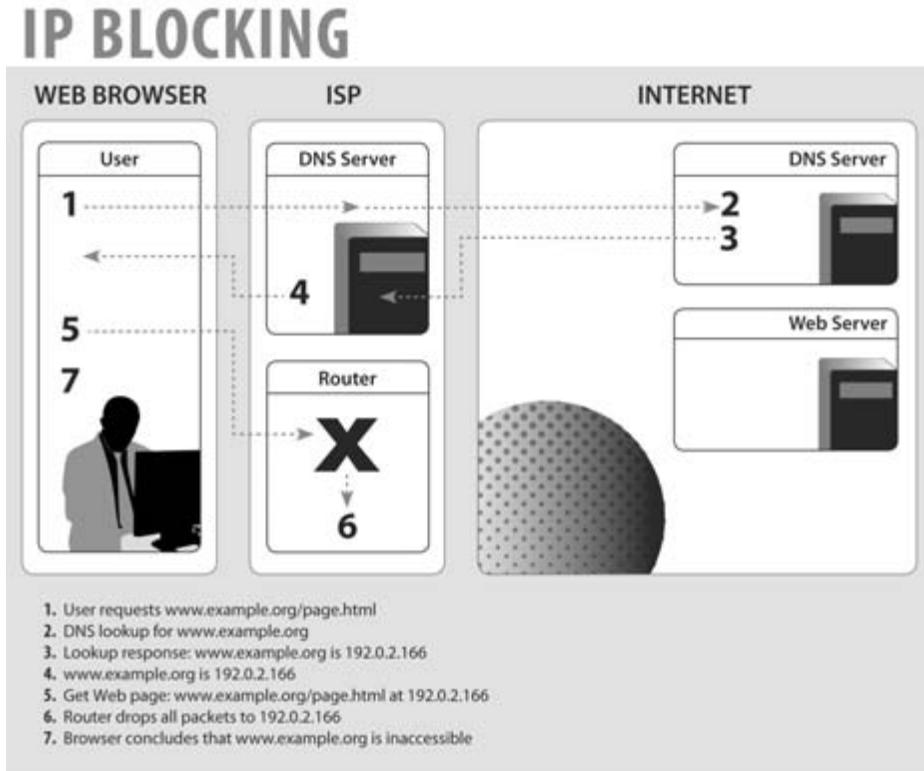


Figure 3.2
IP blocking.

have a maximum size, the full content of the communication will likely be split over multiple packets. Thus while the offending packet will get through, the communication can be disrupted by blocking subsequent packets. This may be achieved by blocking the packets directly or by sending a message to both of the communicating parties requesting they terminate the conversation.¹

Another effect of the maximum packet size is that keywords may be split over packet boundaries. Devices that inspect each packet individually may then fail to identify infringing keywords. For packet inspection to be fully effective, the stream must be reassembled, which adds additional complexity. Alternatively, an HTTP proxy filter can be used, as described later.

DNS Tampering

Most Internet communication uses domain names rather than IP addresses, particularly for Web browsing. Thus, if the domain name resolution stage can be filtered, access to infringing

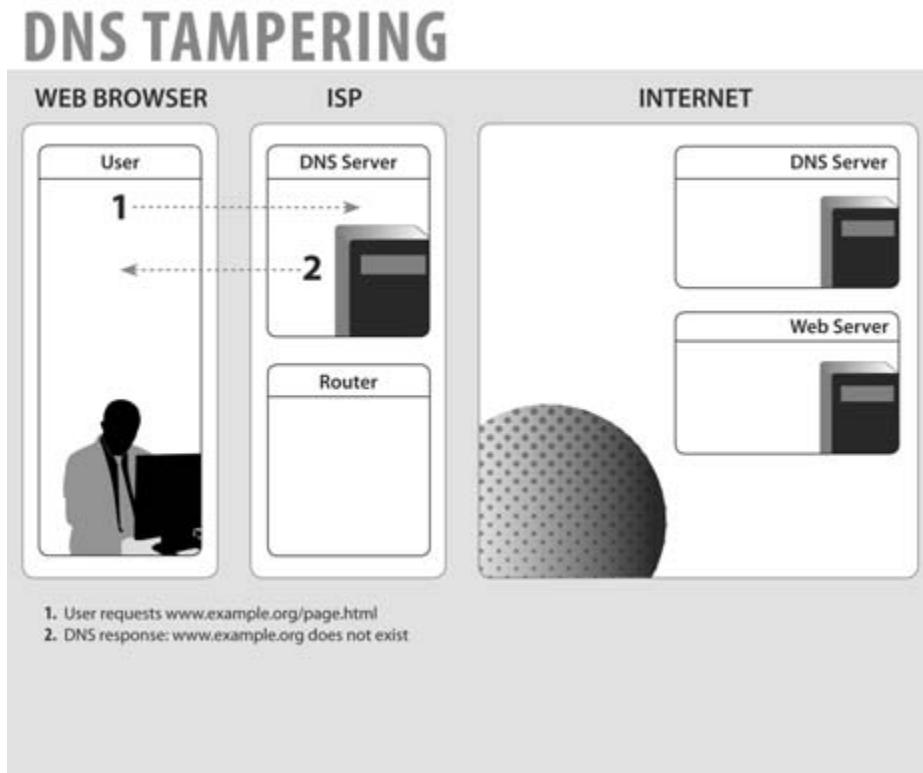


Figure 3.3
DNS tampering via filtering mechanism.

sites can be effectively blocked. With this strategy, the DNS server accessed by users is given a list of banned domain names. When a computer requests the corresponding IP address for one of these domain names, an erroneous (or no) answer is given. Without the IP address, the requesting computer cannot continue and will display an error message.²

Figure 3.3 shows this mechanism in practice. Note that at the stage the blocking is performed, the user has not yet requested a page, which is why all pages under a domain name will be blocked.

HTTP Proxy Filtering

An alternative way of configuring a network is to not allow users to connect directly to Web sites but force (or just encourage) all users to access Web sites via a *proxy server*. In addition to relaying requests, the proxy server may temporarily store the Web page in a *cache*. The advantage of this approach is that if a second user of the same ISP requests the same

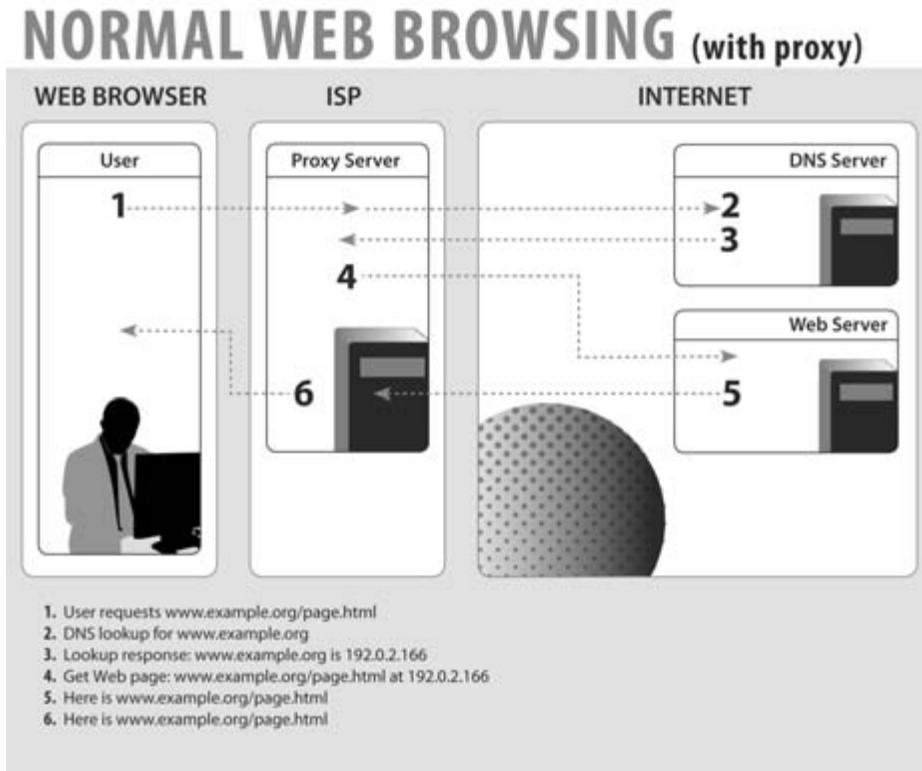


Figure 3.4

Normal Web browsing with a proxy.

page, it will be returned directly from the cache, rather than connecting to the actual Web server a second time. From the user's perspective this is better since the Web page will appear faster, as they never have to connect outside their own ISP. It is also better for the ISP, as connecting to the Web server will consume (expensive) bandwidth, and rather than having to transfer pages from a popular site hundreds of times, they need only do this once. Figure 3.4 shows how the use of a proxy differs from the normal case.

However, as well as improving performance, an HTTP proxy can also block Web sites. The proxy decides whether requests for Web pages should be permitted, and if so, sends the request to the Web server hosting the requested content. Since the full content of the request is available, individual Web pages can be filtered, not just entire Web servers or domains.

An HTTP proxy may be nontransparent, requiring that users configure their Web browsers to send requests via it, but its use can be forced by deploying TCP/IP header filtering to block normal Web traffic. Alternatively, a transparent HTTP proxy may intercept outgoing Web

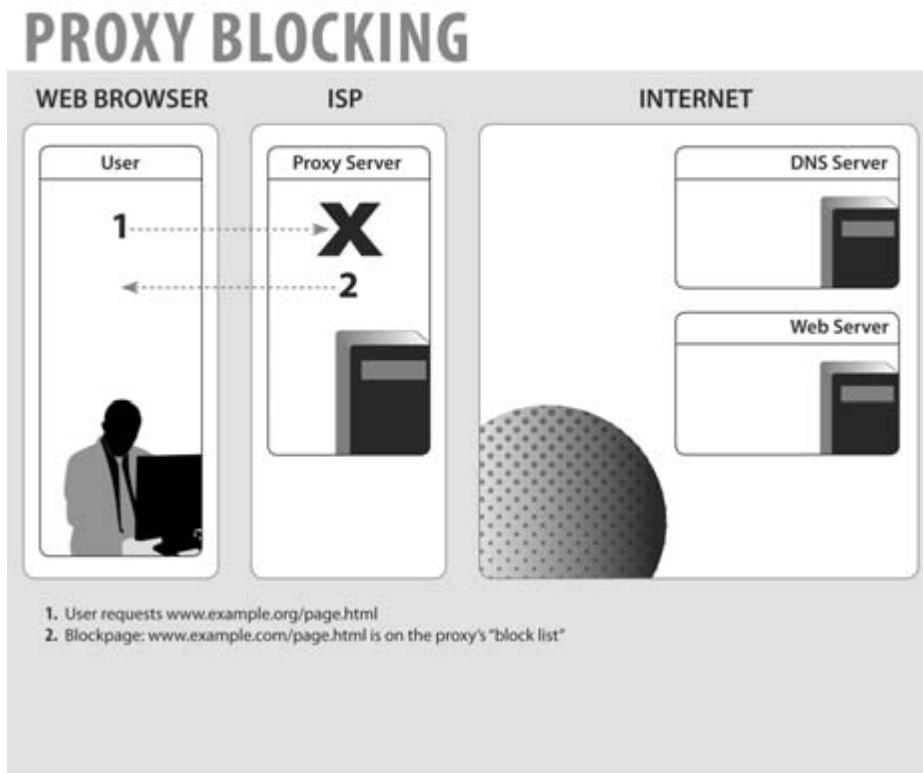


Figure 3.5
HTTP proxy blocking.

requests and send them to a proxy server. While being more complex to set up, this option avoids any configuration changes on the user's computer.

Figure 3.5 shows how HTTP proxy filtering is applied. The ISP structure is different from figure 3.1 because the proxy server must intercept all requests. This gives it the opportunity of seeing both the Web site domain name and which page is requested, allowing more precise blocking when compared to TCP/IP header or DNS filtering.

Hybrid TCP/IP and HTTP Proxy

As the requests intercepted by an HTTP proxy must be reassembled from the original packets, decoded, and then retransmitted, the hardware required to keep up with a fast Internet connection is very expensive. So systems like the BT Cleanfeed project³ were created, which give the versatility of HTTP proxy filtering at a lower cost. It operates by building a list of the IP addresses of sites hosting prohibited content, but rather than blocking data flowing to these

servers, the traffic is redirected to a transparent HTTP proxy. There, the full Web address is inspected and if it refers to banned content, it is blocked; otherwise the request is passed on as normal.

Denial of Service

Where the organization deploying the filtering does not have the authority (or access to the network infrastructure) to add conventional blocking mechanisms, Web sites can be made inaccessible by overloading the server or network connection. This technique, known as a Denial-of-Service (DoS) attack, could be mounted by one computer with a very fast network connection; more commonly, a large number of computers are taken over and used to mount a distributed DoS (DDoS).

Domain Deregistration

As mentioned earlier, the first stage of a Web request is to contact the local DNS server to find the IP address of the desired location. Storing all domain names in existence would be infeasible, so instead so-called recursive resolvers store pointers to other DNS servers that are more likely to know the answer. These servers will direct the recursive resolver to further DNS servers until one, the “authoritative” server, can return the answer.

The domain name system is organized hierarchically, with country domains such as “.uk” and “.de” at the top, along with the nongeographic top-level domains such as “.org” and “.com.” The servers responsible for these domains delegate responsibility for subdomains, such as example.com, to other DNS servers, directing requests for these domains there. Thus, if the DNS server for a top-level domain deregisters a domain name, recursive resolvers will be unable to discover the IP address and so make the site inaccessible.

Country-specific top-level domains are usually operated by the government of the country in question, or by an organization appointed by it. So if a site is registered under the domain of a country that prohibits the hosted content, it runs the risk of being deregistered.

Server Takedown

Servers hosting content must be physically located somewhere, as must the administrators who operate them. If these locations are under the legal or extra-legal control of someone who objects to the content hosted, the server can be disconnected or the operators can be required to disable it.

Surveillance

The above mechanisms inhibit the access to banned material, but are both crude and possible to circumvent. Another approach, which may be applied in parallel to filtering, is to monitor which Web sites are being visited. If prohibited content is accessed (or attempted to be accessed) then legal (or extra-legal) measures could be deployed as punishment.

If this fact is widely publicized, it will discourage others from attempting to access banned content, even if the technical measures for preventing it are inadequate. This type of publicity has been seen in China with Jingjing and Chacha,⁴ two cartoon police officers who inform Internet users that they are being monitored and encourage them to report suspected rule-breakers.

Social Techniques

Social mechanisms are often used to discourage users from accessing inappropriate content. For example, families may place the PC in the living room where the screen is visible to all present, rather than somewhere more private, as a low-key way of discouraging children from accessing unsuitable sites. A library may well situate PCs so that their screens are all visible from the librarian's desk. An Internet café may have a CCTV surveillance camera. There might be a local law requiring such cameras, and also requiring that users register with government-issue photo ID. There is a spectrum of available control, ranging from what many would find sensible to what many would find objectionable.

Comparison of Mechanisms

Each mechanism has different properties of who can deploy systems based around them, what the cost will be, and how effective the filtering is. In this section we compare these properties.

Positioning of System and Scope of Blocking

No single entity has absolute control of the entire Internet, so those who wish to deploy filtering systems are limited in where they can deploy the required hardware or software. Likewise a particular mechanism will block access only to the desired Web site by a particular group of Internet users.

In-line filtering mechanisms (HTTP proxies, TCP/IP header/content filtering, and hybrid approaches) may be placed at any point between the user and the Web server, but to be reliable they must be at a *choke point*—a location that all communication must go through. This could be near the server to block access to it from all over the world, but this requires access to the ISP hosting the server (and they could simply disconnect it completely).

More realistically, these mechanisms are deployed near or in the user's ISP, thereby blocking content from users of its network. For countries with tightly controlled Internet connectivity, these measures can also be placed at the international gateway(s), which makes circumvention more difficult and avoids ISPs being required to take any special action. The positioning of surveillance mechanisms share the same requirements.

DNS tampering is more limited, in that it must be placed at the recursive resolver used by users and is normally within their ISP. The actual list of blocked sites could, however, be

managed on a per-country basis by mandating that all ISPs look up domain names through the government-run DNS server.

Server takedown must be done by the ISP hosting the server and domain deregistration by the registry maintaining the domain use by the Web site. This will usually be a country top-level domain and so be controlled by a government. The physical location of the server need not correspond to the country code used.

Denial-of-Service attacks are the most versatile in terms of location, in that the attacker may be anywhere and an effective attack will prevent access from anywhere.

Finally, social influence is most effectively applied by the country that can impose legal sanctions on the people who are infringing the restrictions, be that people accessing banned Web sites or people publishing banned content.

Error Rate

All the mechanisms suffer from the possibility of errors that may be of two kinds: “false positives”—where sites that were not intended to be blocked are inaccessible, and “false negatives”—where sites are accessible despite the intention that they be blocked. There is commonly a trade-off between these two properties, which are also known as overblocking and underblocking. The trade-off between false positives and false negatives is a pervasive issue in security engineering, appearing in applications from biometric authentication to electronic warfare. The *Receiver Operating Characteristic* (ROC) is the term given to the curve that maps the trade-off between false negative and false positive. Tweaking a parameter typically moves the operating point of the system along the curve; for example, one may obtain fewer false negatives but at the cost of more false positives. In general, the way to improve this trade-off is to devise more precise ways of discriminating between desired and undesired results. This will, in general, shift the ROC curve, so that false negatives and false positives may be reduced at the same time.

TCP/IP header filtering is comparatively crude and must block an entire IP address or address range, which may host multiple Web sites and other services. Taking into account the port number makes the discrimination more precise in that it might limit the blocking to only Web traffic, but this still will often include several hundred Web sites.⁵ Server takedown makes the discrimination less precise, in that it will also make all content on the server inaccessible (including content not served over the Web at all).

DNS tampering and domain deregistration will allow individual Web sites to be blocked but, with the exception of e-mail, which may be handled differently at the DNS level, all services on that domain will be made inaccessible. Both may be more precise than packet header filtering, as multiple servers may be hosted on one machine, and blacklisting that machine may take down many Web sites other than the target site.

TCP/IP content filtering allows particular keywords to be filtered, allowing individual Web pages to be blocked. It does run the risk of missing keywords that are split over multiple packets, but this would be unusual for standard Web browsers.

HTTP proxy and hybrid approaches give the greatest flexibility, allowing blocking both by full Web page URL and by Web page content.

Denial-of-Service attacks are the most crude of the options discussed. Since they normally make sites inaccessible by saturating the network infrastructure, rather than the server itself, many servers could be blocked unintentionally, and perhaps the entire ISP hosting the prohibited content.

Surveillance and the threat of legal measures can be effective, as the human element allows much greater subtlety. Even if the authorities have not discovered a site that should be blocked, self-censorship will still discourage users from attempting to access it. However, such measures are also likely to result in overblocking by creating a climate of fear.

Detectability

Given adequate access to computers that are being blocked from accessing certain Web sites, it is possible to reliably detect most of the mechanisms already discussed. Mechanisms at the server side are more difficult. For example, although the server being blocked can detect Denial of Service, it may be difficult to differentiate from a legitimate "flash crowd." Similarly, a server that has been taken down, or whose domain name has been deregistered for reasons of blocking, appears the same as one that has suffered a hardware failure or DNS misconfiguration.

Surveillance is extremely difficult to detect technically if it has been competently implemented. However, the results of surveillance (arrests or warnings) are often made visible in order to deter future infringement of the rules. So it may be possible to infer the existence of surveillance, but law enforcement agencies may choose to hide precisely how they obtained the information used for targeting.

Circumventability

Although the mechanisms discussed will block access to prohibited resources to users who have configured their computers in a normal way, the protections may be circumvented. However, the effort and skills required vary.

DNS filtering is comparatively easy to bypass by the user selecting an alternative recursive resolver. This type of circumvention may be made more difficult by blocking access to external DNS servers, but doing so would be disruptive to normal activities and could also be bypassed.

TCP/IP header filtering, HTTP proxies, and hybrid proxies may all be fooled by redirecting traffic through an open proxy server. Such servers may be set up accidentally by computer users who misconfigure their own computers. Alternatively, a proxy could be specifically designed for circumventing Internet filtering. Here, the main challenge is to discover an open proxy as many are shut down rapidly due to spammers abusing them, or blocked by organizations that realize they are being used for circumvention.

TCP/IP content filtering will not be resisted by a normal HTTP proxy as the keywords will still be present when communicating with the proxy server. However, encrypted proxy servers may be used to hide what is being accessed through them.

Server takedown, Denial of Service, and domain deregistration are more difficult to resist and require effort on the part of the service operator rather than those who access the Web site. Moving the service to a different location is comparatively easy, as is changing the domain name—particularly if the service has planned for this possibility. More difficult is to notify their users of the new address before the attack is repeated.

Reliability

Even where users are not attempting to circumvent the system, they may still be able to access the prohibited resource. Provided they are implemented correctly and the hardware is capable of handling the required processing, all except Denial of Service and social techniques will block all accesses. The problem with Denial-of-Service attacks is that when systems are overloaded, they will drop some requests at random. This results in some connections, which the censor intended to block, getting through. With social techniques, if someone is simply unaware of the risks they may visit the banned site regardless.

Organizations implementing technical filtering systems must also build a list of sites and pages to block. This is a considerable undertaking if the content to be blocked is a type of content, such as pornography, rather than a specific site, such as an opposing political party. There are commercial filtering products that contain a regularly updated list of material commonly objected to, but even this is likely to miss significant content. Keyword filtering (whether at TCP/IP packet level or by HTTP proxy) mitigates this partially, as only the prohibited keywords need to be listed, rather than enumerating all sites that contain them, but sites aware of this technique can simply not use the offending keyword and select an equivalent term.

Cost and Speed

The cost of deploying a filtering mechanism depends on the complexity of the hardware required to implement it. Also, due to the limited market, specialized Internet filtering equipment is comparatively expensive, so if general purpose facilities can be used to implement filtering, the cost will be lower.

Both of these factors result in TCP/IP header filtering being the cheapest option available. Routers already implement logic for redirecting packets based on destination IP address and adding so-called null routing entries, which discard packets to banned sites, is fairly easy. However, routers can only handle up to a maximum number of rules at a time, so this could become a problem in routers working near their limit. Adding port numbers to these rules requires some additional facilities within the router, but as only the header needs to be inspected, the speed penalty of enabling this is small.

TCP/IP content filtering requires inspecting the payload of the IP packet, which is not ordinarily done by routers. Additional hardware may be required, which, for the data rates found on high-speed Internet links, would be expensive. A cheaper option, which reduces reliability but would considerably decrease cost, is for the filter to examine IP packets as they pass, rather than stopping them for the duration of the examination. Now the filtering equipment is not a bottleneck and may be slower, at the cost of missing some packets. When an infringement of policy is detected, the filtering hardware could send a message to both ends of the connection, requesting that they terminate.

DNS tampering is also very inexpensive as recursive resolvers need not respond particularly rapidly and existing configuration options in DNS servers can be used to implement filtering.

HTTP proxies require connections to be built by reassembling the constituent packets—which requires substantial resources, thereby making this option expensive. Hybrid HTTP proxies are more complex to set up, but once this is done, they are only slightly more expensive than IP filtering despite their much higher versatility. This is because the expensive stage—the HTTP proxy—receives only a small proportion of the traffic, and so need not be particularly powerful.

The cost of Denial-of-Service attacks is difficult to quantify as the scale required depends on how capable the target server is and how fast its Internet connection is. Also, it will likely be illegal to mount this attack, at least on the territory of another country. Legality also affects surveillance, domain deregistration, and server takedown; while easy to do, these mechanisms require adequate legal or extra-legal provisions before ISPs will perform them.

Insertion of False Information

If access to a prohibited Web site is blocked, depending on the mechanism, the user experience will vary. For TCP/IP header and content filtering and Denial of Service it will appear as if there has been an error, which may be desirable if the filtering is intended to be covert. The other options, DNS tampering, proxy and hybrid proxy, domain deregistration, and server takedown all give the option of displaying replacement content. This could be a notification that the site is blocked, to be open about the filtering regime, or it could be a spoofed error message, to be covert. Also, it could be false information, pretending to be from the authors of the content, but actually from somewhere else.

Strategic and Tactical Considerations

It can be useful to compare filtering for censorship with filtering for other purposes. Wiretapping systems, firewalls, and intrusion detection systems share many of the same attributes and problems. In general, such systems may be strategic or tactical. A country may collect strategic communications intelligence by intercepting all traffic with a hostile country regardless of its type, source, or destination using a mechanism such as a tap into a cable. It

may also collect tactical communications intelligence in the context of a criminal investigation by wiretapping the phones of particular suspects or by instructing their ISPs to copy IP traffic to an analysis facility.

Similarly, censorship can be strategic or tactical. Strategic censorship may include permanent blocking of porn sites, or of news sites such as the BBC and CNN; this may be done at the DNS level or by blocking a range of IP addresses. An example of tactical censorship might be interference during an election with the Web servers of an opposition group; this might be done by a service-denial attack or some other relatively deniable technique.

Censorship systems interact in various ways with other types of filtering. Where communications are decentralized, for example, through many blogs and bulletin boards, the censor may use classic communications-intelligence techniques such as traffic analysis and snowball sampling in order to trace sites that are candidates for suppression. (*Snowball sampling* refers to tracking a suspect's contacts and then their contacts recursively, adding suspects as a snowball adds snow when rolling downhill.) Countersurveillance techniques may therefore become part of many censorship resistance strategies.

The interaction between censorship and surveillance is not new. During the early 1980s, the resistance in Poland used radios that operated in bands also used by the BBC and Voice of America; the idea was that the Russians would have to turn off their jammers in order to use radio-direction finding to locate the dissidents. Today, many news sites have blogs or other facilities that third parties can use to communicate with each other; so if a censor is reluctant to jam *The Guardian* newspaper, then its dissidents could use blog posts on *The Guardian* site to talk to each other, using pseudonyms. But many of the novel and interesting interactions have to do with applications.

Discussion

Communication is now a part of more and more applications. Some of these are designed for communication, such as Skype, but bring new capabilities; in Skype's case, it provides encrypted communications and is also widely used. Previously, users of cryptography would be likely to draw attention to themselves, especially in authoritarian countries. Today, Skype and other voice-over IP (VoIP) products are used to save money on telephone bills, and provide voice privacy as a side effect.

Another example is given by Google Docs & Spreadsheets. Google purchased an online word-processor product (Writely) and now makes it available to Internet users in many countries as Google Docs & Spreadsheets. People keep their private documents on Google's servers and edit them online via a Web-based interface. Such a document can be shared instantly with other users; this provides a convenient channel for communications. In this case, the communications are the side effect; the reason people use Google Docs & Spreadsheets is to avoid spending money on Microsoft Office.

The general picture is that censors—and wiretappers—perpetually lag behind the wave of innovation. In the 1960s, computer companies fought with telephone companies for less restricted access to the network, and the telephone companies called government agencies to their aid so as to protect their business models (which involved owning all network-attached devices). By the mid-1980s only a few authoritarian states banned the private ownership of modems, and the security agencies of developed countries had acquired the capability to intercept data communications. The explosion in popularity of fax machines in the mid-1980s put the agencies on the back foot again; a handwritten fax still gives reasonable protection against automated surveillance. When e-mail and the Web took off in the mid-1990s, the agencies scrambled to catch up, with proposals for laws restricting cryptography, which turned out to be irrelevant to the real problems that emerged, and more recent proposals for the retention of communications data. Modern Google users may be largely unaffected by all this—their searches, e-mails, word processing, and group communications may all be cohosted.

It will be interesting, to say the least, to see how states deal with the move to edge-based computing. Developed countries tend to observe a distinction between wiretapping that gives access to content, and traffic analysis that gives access merely to traffic data. Most countries require a higher level of warrantry for access to the former. However, the move to the edge blurs the distinction between traffic and content, and there must eventually be a question as to whether this might undermine the existing controls on state interference with communications. Other countries may be less limited by legal scruples than by technical capability and by access. Application service providers such as Google and Yahoo! have to cooperate with the authorities in countries like China where they maintain a physical presence, but may not make all applications available. Small authoritarian states, that enjoy neither the physical presence of the main service providers nor the technological capability, may find their ability to exert technical control over information flows seriously compromised, and may have to rely largely on legal and social mechanisms.

To sum up, the Internet has borders—just like meatspace—and the quality of the borders depends on the situation of the country that erects them.

Conclusion

Ten years ago, Internet utopians like John Perry Barlow held out the prospect that state-sponsored barriers to communication would be swept away, leading to a significant improvement in the human condition. Some less-developed countries denounced this as “U.S. Information Imperialism.”

The Internet is more complex than previous mechanisms (such as the postal system and telephones). Control is not impossible, but it requires more sophistication, and the censors are continually playing catch-up as technological innovation changes the game. The migration of

communications into the application domain will increase complexity further, and raise interesting policy questions—but there may be different questions in developed and less-developed countries.

The utopians are sometimes seen as having lost; the Internet does have borders now. However, information is much more free than it was ten years ago, and the real question is whether one sees the glass as being half empty or half full. There is no doubt that modern communications technologies—including the mobile phone as well as the Internet—have greatly facilitated the dissemination of news, cultural exchanges, and political activism. Even in developed countries, new technologies from blogs to videophones have increased the potential for surveillance, but have also helped people hold officials to account.

Notes

1. Richard Clayton, Steven J. Murdoch, and Robert N. M. Watson, "Ignoring the Great Firewall of China" in 6th Workshop on Privacy Enhancing Technologies (Cambridge, England, 28–30 June 2006).
2. Maximilian Dornseif, "Government Mandated Blocking of Foreign Web Content," in *Security, E-Learning, E-Services: Proceedings of the 17*, ed. J. von Knop, W. Haverkamp, and E. Jessen (DFN-Arbeitstagung über Kommunikationsnetze, 2003).
3. Richard Clayton, "Failures in a Hybrid Content Blocking System," in Fifth Workshop on Privacy Enhancing Technologies (Dubrovnik Cavtat, Croatia, 30 May–1 June 2005).
4. Xiao Qiang, "Image of Internet Police: Jingjing and Chacha Online," *China Digital Times*, http://chinadigitaltimes.net/2006/01/image_of_internet_police_jingjing_and_chacha_online_hon.php (accessed February 19, 2007).
5. Ben Edelman, "Web Sites Sharing IP Addresses: Prevalence and Significance," Berkman Center for Internet and Society (September 2003).