

Internet Architecture Board
Internet-Draft
Intended status: Informational
Expires: January 16, 2014

R. Barnes
BBN Technologies
A. Cooper
CDT
O. Kolkman
NLnet Labs
July 15, 2013

Technical Considerations for Internet Service Blocking and Filtering
draft-iab-filtering-considerations-03.txt

Abstract

The Internet is structured to be an open communications medium. This openness is one of the key underpinnings of Internet innovation, but it can also allow communications that may be viewed as undesirable by certain parties. Thus, as the Internet has grown, so have mechanisms to limit the extent and impact of abusive or objectionable communications. Recently, there has been an increasing emphasis on "blocking" and "filtering," the active prevention of such communications. This document examines several technical approaches to Internet blocking and filtering in terms of their alignment with the overall Internet architecture. In general, the approach to blocking and filtering that is most coherent with the Internet architecture is to inform endpoints about potentially undesirable services, so that the communicants can avoid engaging in abusive or objectionable communications.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 16, 2014.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	Filtering Examples	4
3.	Characteristics of Blocking Systems	6
3.1.	Entities that set blocking policies	6
3.2.	Purposes of blocking	6
3.3.	Intended targets of blocking	7
3.4.	Components used for blocking	7
4.	Evaluation of Blocking Design Patterns	9
4.1.	Criteria for evaluation	9
4.1.1.	Scope: What content or services can be blocked?	9
4.1.2.	Granularity: How specific is the blocking? Will blocking one service also block others?	10
4.1.3.	Efficacy: How easy is it for a resource or service to avoid being blocked?	11
4.1.4.	Security: How does the blocking impact existing trust infrastructures?	11
4.2.	Network-Based Blocking	12
4.2.1.	Scope	12
4.2.2.	Granularity	13
4.2.3.	Efficacy and security	14
4.2.4.	Summary	15
4.3.	Rendezvous-Based Blocking	16
4.3.1.	Scope	16
4.3.2.	Granularity	17
4.3.3.	Efficacy	17
4.3.4.	Security and other implications	17
4.3.5.	Examples	17
4.3.6.	Summary	19
4.4.	Endpoint-Based Blocking	19
4.4.1.	Scope and granularity	20
4.4.2.	Efficacy	20
4.4.3.	Security	21
4.4.4.	Summary	21

5.	Security Considerations	21
6.	Conclusion	22
7.	Informative References	23
	Authors' Addresses	26

[1.](#) Introduction

The original design goal of the Internet was to enable communications between hosts. As this goal was met and people started using the Internet to communicate, however, it became apparent that some hosts were engaging in communications that were viewed as undesirable by certain parties. The most famous early example of undesirable communications was the Morris worm [[Morris](#)], which used the Internet to infect many hosts in 1988. As the Internet has evolved into a rich communications medium, so too have mechanisms to restrict communications viewed as undesirable, ranging from acceptable use policies enforced through informal channels to technical blocking mechanisms.

Efforts to restrict or deny access to Internet resources and services have evolved over time. As noted in [[RFC4084](#)], some Internet service providers impose restrictions on which applications their customers may use and which traffic they allow on their networks. These restrictions are often imposed with customer consent, where customers may be enterprises or individuals. Increasingly, however, both governmental and private sector entities are seeking to block or filter access to certain content, traffic, or services without the knowledge or agreement of affected users. Where these entities do not directly control networks themselves, they commonly aim to make use of intermediary systems to effectuate the blocking or filtering.

While blocking and filtering remain highly contentious in many cases, the desire to restrict access to content will likely continue to exist.

The difference between "blocking" and "filtering" is a matter of scale and perspective. "Blocking" often refers to preventing access to resources in the aggregate, while "filtering" refers to preventing access to specific resources within an aggregate. Both blocking and filtering can be effectuated at the level of "services" (web hosting or video streaming, for example) or at the level of particular "content." For the analysis presented in this document, the distinction between blocking and filtering does not create meaningfully different conclusions. Hence, in the remainder of this document, we will treat the terms as being generally equivalent and applicable to restrictions on both content and services.

This document aims to clarify the technical implications and trade-offs of various blocking strategies and to identify the potential for different strategies to potentially cause harmful side effects ("collateral damage") for Internet users and the overall Internet architecture. This analysis is limited to technical blocking mechanisms. Enforcement of blocking via contractual terms or legal action is out of scope, although at times these actions ultimately result in the application of technical mechanisms.

Filtering may be considered legal, illegal, ethical, or unethical in different places, at different times, and by different parties. This document is intended for an audience of entities that are conducting filtering or are considering conducting filtering and who want to understand the implications of their decisions with respect to the Internet architecture and the trade-offs that come with each type of filtering strategy. This document does not present formulas on how to make those trade-offs; it is likely that filtering decisions require knowledge of context-specific details. Whether particular forms of filtering are lawful in particular jurisdictions raises complicated legal questions that are outside the scope of this document. For similar reasons questions about the ethics of particular forms of filtering are also out of scope.

In [[SAC-056](#)], ICANN's Security and Stability Advisory Committee (SSAC) assessed the aspects of blocking using the DNS. This document attempts to take a broader perspective on blocking and filtering and generalizes from some of SSAC's findings.

2. Filtering Examples

Blocking systems have evolved alongside the Internet technologies they seek to restrict. Looking back at the history of the Internet, there have been several such systems deployed by different entities and for different purposes.

Firewalls: Firewalls are a very common tool used for service blocking, employed at many points in today's Internet. Typically, firewalls block according to content-neutral rules, e.g., blocking all inbound connections or outbound connections on certain ports, protocols and network layer addresses. More advanced configurations perform deep packet inspection or traffic flow analysis and filter or block based on rich (content-specific) rules and policies. Many firewalls include web filtering capabilities (see below). Firewalls can be deployed either on end hosts (under user or administrator control), or at network boundaries.

Web Filtering: HTTP and HTTPS are common targets for blocking and filtering, typically targeted at specific URLs. Some enterprises use

HTTP blocking to block non-work-appropriate web sites, and several nations require HTTP and HTTPS filtering by their ISPs in order to block content deemed illegal. HTTPS is a challenge for these systems, because the URL in an HTTPS request is carried inside the encrypted channel. To block access to content made accessible via HTTPS, filtering systems thus must either block based only on IP address, or else obtain a trust anchor certificate that is trusted by endpoints (and thus act as a man in the middle). These filtering systems often take the form of "portals" or "enterprise proxies." These portals present their own HTTPS certificates that are invalid for any given domain according to normal validation rules, but may still be trusted if the user installs a security exception. (See further discussion in [Section 5](#).)

Spam Filtering: Spam filtering is one of the oldest forms of content filtering. Spam filters evaluate messages based on a variety of criteria and information sources to decide whether a given message is spam. For example, DNS Reverse Black Lists use the reverse DNS to flag whether an IP address is a known spam source [[RFC5782](#)]. Spam filters are typically either installed on user devices (e.g., in a mail client) or operated by a mail domain on behalf of users.

Domain name seizure: In recent years, US law enforcement authorities have been issuing legal orders to domain name registries to seize domain names associated with the distribution of counterfeit goods and other alleged illegal activity [[US-ICE](#)]. When domain names are seized, DNS queries for the seized names are typically redirected to resolve to U.S. government IP addresses that host information about the seizure. The effectiveness of domain seizures is limited by application mobility -- applications using the seized name can switch to using another name. Seizures can also have overbroad effects, since access to content is blocked not only within the jurisdiction of the seizure, but globally, even when it may be affirmatively legal elsewhere [[RojaDirecta](#)]. When domain redirection is effected via redirections at intermediate resolvers rather than at authoritative servers, it directly contradicts end-to-end assumptions in the DNS security architecture [[RFC4033](#)], potentially causing validation failures by validating end-nodes.

Safe Browsing: Modern web browsers provide some measures to prevent users from accessing malicious web sites. For instance, before loading a URL, current versions of Google Chrome and Firefox use the Google Safe Browsing service to determine whether or not a given URL is safe to load [[SafeBrowsing](#)]. The DNS can also be used to store third party information that mark domains as safe or unsafe [[RFC5782](#)].

Manipulation of routing and addressing data: Governments have recently intervened in the management of IP addressing and routing information in order to maintain control over a specific set of DNS servers. As part of an internationally coordinated response to the DNSChanger malware, a Dutch court ordered the RIPE NCC to freeze the accounts of several resource holders as a means to limit the resource holders' ability to use certain address blocks [[GhostClickRIPE](#)] (also see [Section 4.3](#)). These actions have led to concerns that the number resource certification system and related secure routing technologies developed by the IETF's SIDR working group might be subject to government manipulation as well [[RFC6480](#)], potentially for the purpose of denying targeted networks access to the Internet.

3. Characteristics of Blocking Systems

At a generic level, blocking systems can be characterized by four attributes: the entity that sets the blocking policy, the purpose of the blocking, the intended target of the blocking, and the Internet component(s) used as the basis of the blocking system.

[3.1.](#) Entities that set blocking policies

Parties that institute blocking policies include governments, enterprises, network operators, application providers, and individual end users. In some cases, these parties use their own technical assets to conduct blocking; for example, a network operator might install a firewall in its own networking equipment, or a web application provider might block responses between its web server and certain clients. In other cases, particularly in the case of blocking initiated by governments, the entity that institutes the blocking policy works with other entities to effectuate blocking using technical assets that it does not control.

[3.2.](#) Purposes of blocking

Entities may be motivated to filter for a variety of purposes:

- o Preventing or responding to security threats. Network operators, enterprises, application providers, and end users often block communications that are believed to be associated with security threats or network attacks.
- o Restricting objectionable content or services. Certain communications may be viewed as undesirable, harmful, or illegal by particular governments, enterprises, or users (e.g., parents). Governments may seek to block communications that are deemed to be defamation, hate speech, obscenity, intellectual property infringement, or otherwise objectionable. Enterprises may seek to

restrict employees from accessing content that is not deemed to be work appropriate. Parents may restrict their children from accessing content or services targeted for adults.

- o Restricting access based on business arrangements. Some networks are designed so as to only provide access to certain content or services ("walled gardens"), or to only provide limited access until end users pay for full Internet services (captive portals provided by hotspot operators, for example).

Note that the purpose for which blocking occurs often dictates whether the blocking system operates on a blacklist model, where communications are allowed by default but a subset are blocked, or a whitelist model, where communications are blocked by default with only a subset allowed. Captive portals, walled gardens, and sandboxes used for security or network endpoint assessment usually require a whitelist model since the scope of communications allowed is narrow. Blocking for other purposes often uses a blacklist model since only individual content or traffic is intended to be blocked.

3.3. Intended targets of blocking

Entities institute blocking systems so as to target particular content, services, endpoints, or some combination of these. For example, a "content" filtering system used by an enterprise might block access to specific URLs whose content is deemed by the enterprise to be inappropriate for the work place. This is distinct from a "service" filtering system that blocks all web traffic (perhaps as part of a parental control system on an end user device), and also distinct from an "endpoint" filtering system in which a web application blocks traffic from specific endpoints that are suspected of malicious activity.

As discussed in [Section 4](#), the design of a blocking system may affect content, services, or endpoints other than those that are the intended targets. For example, the domain name seizures described above target particular web pages associated with illegal activity, but by removing the domains from use, they affect all services made available by the hosts associated with those names, including mail services and web services unrelated to the illegal activity.

3.4. Components used for blocking

Broadly speaking, the process of delivering an Internet service involves three different components:

1. Endpoints: The actual content of the service is an application layer protocol between two Internet hosts. In many protocols, there are two endpoints, a client and a server.
2. Network services: The endpoints communicate by way of a collection of IP networks that use routing protocols to determine how to deliver packets between the endpoints.
3. Rendezvous services: Service endpoints are typically identified by identifiers that are more "human-friendly" than IP addresses. Rendezvous services allow one endpoint to figure out how to contact another endpoint based on an identifier.

Consider, for example, an HTTP transaction fetching the content of the URI `<http://example.com/index.html>`. The client endpoint is an end host running a browser. The client uses the DNS as a rendezvous service when it performs a AAAA query to obtain the IP address for the server name "example.com". The client then establishes a connection to the server, and sends the actual HTTP request. The server then responds to the HTTP request.

As another example, in the SIP protocol, the client and server are IP phones, and the rendezvous service is provided by an application-layer SIP proxy as well as the DNS.

Blocking access to Internet content, services, or endpoints is done by controlling one or more of the components involved in the provision of the communications involved in accessing the content, services or endpoints. In the HTTP example above, the successful completion of the HTTP request could have been prevented in several ways:

- o [Endpoint] Preventing the client from making the request
- o [Endpoint] Preventing the server from responding to the request
- o [Network] Preventing the request from reaching the server
- o [Network] Preventing the response from reaching the client
- o [Network] Preventing the client from reaching the DNS server
- o [Rendezvous] Preventing the DNS server from providing the client the correct IP address of the server

Most entities that desire to block communications will have access to only one or two components, and therefore their choices for how to effectuate blocking will be limited. End users and application

providers can usually only control their own software and hardware, which means that they are limited to endpoint-based filtering. Some network operators offer filtering services that their customers can activate individually, in which case end users might have network-based filtering systems available to them. Network operators can control their own networks and the rendezvous services for which they provide infrastructure support (e.g., DNS resolvers) or to which they may have access (e.g., SIP proxies), but not usually endpoints. Enterprises usually have access to their own networks and endpoints for filtering purposes. Governments might make arrangements with the operators or owners of any of the three components that exist within their jurisdictions to effectuate filtering.

In the next section, blocking systems designed according to each of the three patterns -- network services, rendezvous services, and endpoints -- are evaluated for their technical and architectural implications. The analysis is as agnostic as possible as to which kind of entity sets the blocking policy (government, end user, network operator, application provider, or enterprise), but in some cases the way in which a particular blocking design pattern is used might differ depending on the entity that desires to block. For example, a network-based firewall provided by an ISP that parents can elect to use for parental control purposes will likely function differently from one that all ISPs in a particular jurisdiction are required to use by the local government, even though in both cases the same component (network) forms the basis of the blocking system.

4. Evaluation of Blocking Design Patterns

4.1. Criteria for evaluation

To evaluate the technical implications of each of the blocking design patterns, we compare them based on four criteria: scope, granularity, efficacy, and security.

4.1.1. Scope: What content or services can be blocked?

The Internet is comprised of many distinct autonomous networks and applications, which means that the impact of a blocking system will only be within a defined scope. For example, blocking within an access network will only affect a relatively small, well-defined set of users (namely, those connected to the access network), but can affect all applications for those users. Blocking effectuated by an application provider can affect users across the entire Internet, but only for that specific application. Thus the scope of the impact might be narrow in one dimension (set of users or set of applications affected) but broad in another. In some cases, applications and rendezvous services are so intertwined with each other that filtering

a single service or in a single network location can have broad effects in multiple directions. Blocking systems are generally viewed as less objectionable if the scope of their impact is as narrow as possible while still being effective.

4.1.2. Granularity: How specific is the blocking? Will blocking one service also block others?

Internet applications are built out of a collection of loosely-coupled components or "layers." Different layers serve different purposes, and rely on or offer different functions such as routing, transport, and naming (see [[RFC1122](#)], especially [Section 1.1.3](#)). The functions at these layers are developed autonomously and almost always operated by different entities. For example, in many networks, physical and link-layer connectivity is provided by an "access provider", IP routing is performed by an "Internet service provider," and application-layer services are provided by completely separate entities (e.g., web servers). Upper-layer protocols and applications rely on combinations of lower-layer functions in order to work. Functionality at higher layers tends to be more specialized, so that many different specialized applications can make use of the same generic underlying network functions.

As a result of this structure, actions taken at one layer can affect functionality or applications at higher layers. For example, manipulating routing or naming functions to restrict access to a narrow set of resources via specific applications will likely affect all applications that depend on those functions. As with the scope criteria, blocking systems are generally viewed as less objectionable when they are highly granular and do not cause collateral damage to content or services unrelated to the target of the blocking.

Even within the application layer, the granularity of blocking can vary depending on how targeted the blocking system is designed to be. Blocking all traffic associated with a particular application protocol is less granular than blocking only traffic associated with a subset of application instances that make use of that protocol. Sophisticated heuristics that make use of information about the application protocol, lower-layer protocols, payload signatures, source and destination addresses, inter-packet timing, packet sizes, and other characteristics are sometimes used to narrow the subset of traffic to be blocked.

Design flaws in blocking systems may also cause the effects of blocking to be overbroad. For example, web filtering systems in India and China have been shown to cause "collateral damage" by unwittingly blocking users in Oman and the US from accessing web sites in Germany and Korea [[IN-OM-filtering](#)][CCS-GFC-collateral-damage].

4.1.3. Efficacy: How easy is it for a resource or service to avoid being blocked?

For blacklist-style blocking, the distributed and mobile nature of Internet resources limits the effectiveness of blocking actions. A service that is blocked in one jurisdiction can often be moved or re-instantiated in another jurisdiction (see, for example, [[Malicious-Resolution](#)]). Likewise, services that rely on blocked resources can often be rapidly re-configured to use non-blocked resources. If a web site is prevented from using a domain name or set of IP addresses, the web site can simply move to another domain name or network. In a process known as "snowshoe spamming," a spam originator uses addresses in many different networks as sources for spam. This technique is already widely used to spread spam generation across a variety of resources and jurisdictions to prevent spam blocking from being effective.

In the presence of either blacklist or whitelist systems, users may choose to use different sets of protocols or otherwise alter their traffic characteristics to circumvent filters. As discussed in [[I-D.blanchet-iab-internetoverport443](#)], many applications shift their traffic to port 80 or 443 when other ports are blocked. If voice communication based on SIP [[RFC3261](#)] is blocked, users are likely to use proprietary protocols that allow them to talk to each other. Thus distribution and mobility can hamper efforts to block communications in a number of ways.

4.1.4. Security: How does the blocking impact existing trust infrastructures?

Modern security mechanisms rely on trusted hosts communicating via a secure channel without intermediary interference. Protocols such as TLS and IPsec [[RFC5246](#)][RFC4301] are designed to ensure that each endpoint of the communication knows the identity of the other endpoint, and that only the endpoints of the communication can access the secured contents of the communication. For example, when a user connects to a bank's web site, TLS ensures that the user's banking information is securely communicated to the bank and nobody else, ensuring the data remains confidential while in transit.

Some blocking strategies require intermediaries to insert themselves within the end-to-end communications path, potentially breaking security properties of Internet protocols. In these cases it can be difficult or impossible for endpoints to distinguish between attackers and "authorized" entities conducting blocking.

4.2. Network-Based Blocking

Being able to block access to resources without the consent or cooperation of either endpoint to a communication is viewed as a desirable feature by some entities that deploy blocking systems. Systems that have this property are often implemented using intermediary devices in the network, such as firewalls or filtering systems. These systems inspect traffic as it passes through the network, decide based on the characteristics or content of a given communication whether it should be blocked, and then block or allow the communication as desired. For example, web filtering devices usually inspect HTTP requests to determine the URL being requested, compare that URL to a list of black-listed or white-listed URLs, and allow the request to proceed only if it is permitted by policy (or at least not forbidden). Firewalls perform a similar function for other classes of traffic in addition to HTTP. Some blocking systems focus on specific application-layer traffic, while others, such as router ACLs, filter traffic based on lower layer criteria (transport protocol and source or destination addresses or ports).

Intermediary systems used for blocking are often not far from the edge of the network. For example, many enterprise networks operate firewalls that block certain web sites, as do some residential ISPs. In some cases, this filtering is done with the consent or cooperation of the affected endpoints. PCs within an enterprise, for example, might be configured to trust an enterprise proxy, a residential ISP might offer a "safe browsing" service, or mail clients might authorize mail servers on the local network to filter spam on their behalf. These cases share some of the properties of the "Endpoint-Based Blocking" scenarios discussed in [Section 4.4](#) below, since the endpoint has made an informed decision to authorize the intermediary to block on its behalf and is therefore unlikely to attempt to circumvent the blocking. From an architectural perspective, however, they may create many of the same problems as network-based filtering conducted without consent.

4.2.1. Scope

Network-based approaches to blocking run into several technical issues that limit their viability in practice. In particular, many issues arise from the fact that an intermediary needs to have access to a sufficient amount of traffic to make its blocking determinations.

For residential or consumer networks with many egress points, the first challenge to obtaining this traffic is simply gaining access to the constituent packets. The Internet is designed to deliver packets hop-by-hop from source to destination -- not to any particular point along the way. In practice, inter-network routing is often asymmetric, and for sufficiently complex local networks, intra-network traffic flows can be asymmetric as well [[asymmetry](#)].

This asymmetry means that an intermediary in a network with many egress points will often see only one half of a given communication (if it sees any of it at all), which may limit the scope of the communications that it can filter. For example, a filter aimed at requests destined for particular URLs cannot make accurate blocking decisions if it is only in the data path for HTTP responses and not requests. Routing can sometimes be forced to be symmetric within a given network using routing configuration, NAT, or layer-2 mechanisms (e.g., MPLS), but these mechanisms are frequently brittle, complex, and costly -- and can sometimes result in reduced network performance relative to asymmetric routing. Enterprise networks may also be less susceptible to these problems if they route all traffic through a small number of egress points.

4.2.2. Granularity

Once an intermediary in a network has access to traffic, it must identify which packets must be filtered. This decision is usually based on some combination of information at the network layer (e.g., IP addresses), transport layer (ports), or application layer (URLs or other content). Blocking based on application-layer attributes can be potentially more granular and less likely to cause collateral damage than blocking all traffic associated with a particular address, which can impact unrelated occupants of the same address. However, more narrowly focused targeting may be more complex, less efficient, or easier to circumvent than filtering that sweeps more broadly, and entities that seek to block may balance these attributes against each other when choosing a blocking system.

4.2.3. Efficacy and security

Regardless of the layer at which blocking occurs, it may be open to circumvention, particularly in cases where network endpoints have not authorized the blocking. The communicating endpoints can deny the intermediary access to attributes at any layer by using encryption (see below). IP addresses must be visible, even if packets are protected with IPsec, but blocking based on IP addresses can be trivial to circumvent. A filtered site may be able to quickly change its IP address using only a few simple steps: changing a single DNS record and provisioning the new address on its server or moving its services to the new address. Indeed, in the face of IP-based blocking in some networks, services such as The Pirate Bay are now using cloud hosting services so that their IP addresses are difficult for intermediaries to predict [[BT-TPB](#)][TPB-cloud].

If application content is encrypted with a security protocol such as IPsec or TLS, then the intermediary will require the ability to decrypt the packets to examine application content. Since security protocols are designed to provide end-to-end security (i.e., to prevent intermediaries from examining content), the intermediary would need to masquerade as one of the endpoints, breaking the authentication in the security protocol, reducing the security of the users and services affected, and interfering with legitimate private communication. Besides, various techniques that use public databases with whitelisted keys (e.g. DANE) enable users to detect these sort of intermediaries. Those users are then likely to act as if the service is blocked.

If the intermediary is unable to decrypt the security protocol, then its blocking determinations for secure sessions can only be based on unprotected attributes, such as IP addresses, protocol IDs and port numbers. Some blocking systems today still attempt to block based on these attributes, for example by blocking TLS traffic to known proxies that could be used to tunnel through the blocking system.

However, as the Telex project recently demonstrated, if an endpoint cooperates with a relay in the network (e.g., a Telex station), it can create a TLS tunnel that is indistinguishable from legitimate traffic [[Telex](#)]. For example, if an ISP used by a banking website were to operate a Telex station at one of its routers, then a blocking system would be unable to distinguish legitimate encrypted banking traffic from Telex-tunneled traffic (potentially carrying content that would have been filtered).

Thus, in principle in a blacklist system it is impossible to block tunneled traffic through an intermediary device without blocking all secure traffic. (The only limitation in practice is the requirement

for special software on the client.) In most cases, blocking all secure traffic is an unacceptable consequence of blocking, since security is often required for services such as online commerce, enterprise VPNs, and management of critical infrastructure. If governments or network operators were to force these services to use insecure protocols so as to effectuate blocking, they would expose their users to the various attacks that the security protocols were put in place to prevent.

Some operators may assume that only blocking access to resources available via unsecure channels is sufficient for their purposes -- i.e., that the size of the user base that will be willing to use secure tunnels and/or special software to circumvent the blocking is low enough to make blocking via intermediaries worthwhile. Under that assumption, one might decide that there is no need to control secure traffic, and thus that network-based blocking is an attractive option.

However, the longer such blocking systems are in place, the more likely it is that efficient and easy-to-use tunneling tools will become available. The proliferation of the Tor network, for example, and its increasingly sophisticated blocking-avoidance techniques demonstrate that there is energy behind this trend [[Tor](#)]. Thus, network-based blocking becomes less effective over time.

Network-based blocking is a key contributor to the arms race that has led to the development of these kinds of tools, the result of which is to create unnecessary layers of complexity in the Internet. Before content-based blocking became common, the next best option for network operators was port blocking, the widespread use of which has driven more applications and services to use ports (80 and 443 most commonly) that are unlikely to be blocked. In turn, network operators shifted to finer-grained content blocking over port 80, content providers shifted to encrypted channels, and operators began seeking to identify those channels (although doing so can be resource-prohibitive, especially if tunnel endpoints begin to change frequently). Because the premise of network-based blocking is that endpoints have incentives to circumvent it, this cat-and-mouse game is an inevitable by-product of this form of blocking.

4.2.4. Summary

In sum, network-based blocking is only effective in a fairly constrained set of circumstances. First, the traffic needs to flow through the network in such a way that the intermediary device has access to any communications it intends to block. Second, the blocking system needs an out-of-band mechanism to mitigate the risk of secure protocols being used to avoid blocking (e.g., human

analysts identifying IP addresses of tunnel endpoints). If the network is sufficiently complex, or the risk of tunneling too high, then network-based blocking is unlikely to be effective, and in any case this type of blocking drives the development of increasingly complex layers of circumvention. Network-based blocking can be done without the cooperation of either endpoint to a communication, but it has the serious drawback of breaking end-to-end security assurances in some cases.

4.3. Rendezvous-Based Blocking

Internet applications often require or rely on support from common, global rendezvous services, including the DNS, certificate authorities, WHOIS databases, and Internet Route Registries. These services control or register the structure and availability of Internet applications by providing data elements that are used by application code. Some applications also have their own specialized rendezvous services. For example, to establish an end-to-end SIP call the end-nodes (terminals) will rely on presence and session information supplied by SIP servers.

Global rendezvous services are comprised of generic technical databases intended to record certain facts about the network. The DNS, for example, stores information about which servers provide services for a given name; the RPKI about which entities have been allocated IP addresses. To offer specialized Internet services and applications, different entities rely on these generic records in different ways. Thus the effects of changes to the databases can be much more difficult to predict than, for example, the effect of shutting down a web server (which fulfills the specific purpose of serving web content).

Although rendezvous services are discussed as a single category, the precise characteristics and implications of blocking each kind of rendezvous service are slightly different. This section provides examples to highlight these differences.

4.3.1. Scope

In the case of government-initiated blocking, the servers that are used to provide rendezvous services exist within specific jurisdictions, and their operators are thus subject to jurisdictional laws. It is thus possible for laws to be structured to effectuate blocking by imposing obligations on the operators of rendezvous services within a jurisdiction, either via direct government action or by allowing private actors to demand blocking (e.g., through lawsuits).

The scope of blocking conducted by other entities will depend on which servers those entities can access. For example, network operators and enterprises may be capable of conducting blocking using their own DNS resolvers or application proxies within their networks, but not authoritative servers controlled by others.

4.3.2. Granularity

Blocking based on global rendezvous services tends to be overbroad, causing collateral damage to legitimate uses of a resource. For example, a given address or domain name might host both legitimate services and services that governments desire to block. A service hosted under a domain name and operated in a jurisdiction where it is considered undesirable might be considered legitimate in another jurisdiction; a blocking action in the host jurisdiction would deny legitimate services in the other.

4.3.3. Efficacy

The distributed nature of the Internet limits the efficacy of blocking based on rendezvous services. If the Internet community realizes that a blocking decision has been made and wishes to counter it, then local networks can "patch" the authoritative data that the rendezvous service provides to avoid the blocking (although the development of DNSSEC and the RPKI are causing this to change by requiring updates to be authorized). In the DNS case, registrants whose names get blocked can relocate their resources to different names.

4.3.4. Security and other implications

Blocking of global rendezvous services also has a variety of other implications that may reduce the stability, accessibility, and usability of the global Internet. Infrastructure-based blocking may erode the trust in the general Internet and encourage the development of parallel or "underground" infrastructures causing forms of Internet balkanisation, for example. This risk may become more acute as the introduction of security infrastructures and mechanisms such as DNSSEC and RPKI "hardens" the authoritative data -- including blocked names or routes -- that the existing infrastructure services provide. Those seeking to circumvent the blocks may opt to use less-secure but unblocked parallel services. As applied to the DNS, these considerations are further discussed in ISOC's whitepaper on DNS filtering [[ISOCFiltering](#)], but they also apply to other global Internet resources.

4.3.5. Examples

Below we provide a few specific examples for routing, DNS, and WHOIS services. These examples demonstrate that for these types of rendezvous services (services that are often considered a global commons), jurisdiction-specific legal and ethical motivations for blocking can both have collateral effects in other jurisdictions and be circumvented because of the distributed nature of the Internet.

In 2008, Pakistan Telecom attempted to deny access to YouTube within Pakistan by announcing bogus routes for YouTube address space to peers in Pakistan. YouTube was temporarily denied service on a global basis as a result of a route leak beyond the Pakistan ISP's scope, but service was restored in approximately two hours because network operators around the world re-configured their routers to ignore the bogus routes [[RenesysPK](#)]. In the context of SIDR and secure routing, a similar re-configuration could theoretically be done if a resource certificate were to be revoked in order to block routing to a given network.

In the DNS realm, one of the recent cases of US law enforcement seizing domain names involved RojaDirecta, a Spanish web site. Even though several of the affected domain names belonged to Spanish entities, they were subject to blocking by the US government because certain servers were operated in the US. Government officials required the operators of the parent zones of a target name (e.g., "com" for "example.com") to direct queries for that name to a set of US-government-operated name servers. Users of other services under a target name (e.g. e-mail) would thus be unable to locate the servers providing services for that name, denying them the ability to access these services.

Similar workarounds as those that were used in the Pakistan Telecom case are also available in the DNS case. If a domain name is blocked by changing authoritative records, network operators can restore service simply by extending TTLs on cached pre-blocking records in recursive resolvers, or by statically configuring resolvers to return un-blocked results for the affected name. However, depending on availability of valid signature data, these types of workarounds will not work with DNSSEC-signed data.

The action of the Dutch authorities against the RIPE NCC, where RIPE was ordered to freeze the accounts Internet resource holders, is of a similar character. By controlling the account holders' WHOIS information, this type of action limited the ability of the ISPs in question to manage their Internet resources. This example is slightly different from the others because it does not immediately impact the ability of ISPs to provide connectivity. While ISPs use (and trust) the WHOIS databases to build route filters or use the databases for trouble-shooting information, the use of the WHOIS

databases for those purposes is voluntary. Thus, seizure of this sort may not have any immediate effect on network connectivity, but it may impact overall trust in the common infrastructure. It is similar to the other examples in that action in one jurisdiction can have broader effects, and in that the global system may encourage networks to develop their own autonomous solutions.

4.3.6. Summary

In summary, rendezvous-based blocking can sometimes be used to immediately block a target service by removing some of the resources it depends on. However, such blocking actions can have harmful side effects due to the global nature of Internet resources and the fact that many different application-layer services rely on generic, global databases for rendezvous purposes. The fact that Internet resources can quickly shift between network locations, names, and addresses, together with the autonomy of the networks that comprise the Internet, can mean that the effects of rendezvous-based blocking can be negated on short order in some cases. For some applications, rendezvous services are optional to use, not mandatory. Hence they are only effective when the endpoint or the endpoint's network chooses to use them; they can be routed around by choosing not to use the rendezvous service or migrating to an alternative one. To adapt a quote by John Gilmore, "The Internet treats blocking as damage and routes around it".

4.4. Endpoint-Based Blocking

[TO DO: Should this section only address "user"-based blocking or should it also discuss server-based blocking (e.g., a web server that prevents abusive users from accessing content?)]

Internet users and their devices constantly make decisions as to whether to engage in particular Internet communications. Users decide whether to click on links in suspect email messages; browsers advise users on sites that have suspicious characteristics; spam filters evaluate the validity of senders and messages. If the hardware and software making these decisions can be instructed not to engage in certain communications, then the communications are effectively blocked because they never happen.

There are several systems in place today that advise user systems about which communications they should engage in. As discussed above, several modern browsers consult with "Safe Browsing" services before loading a web site in order to determine whether the site could potentially be harmful. Spam filtering is one of the oldest types of filtering in the Internet; modern filtering systems typically make use of one or more "reputation" or "blacklist"

databases in order to make decisions about whether a given message or sender should be blocked. These systems typically have the property that many filtering systems (browsers, MTAs) share a single reputation service. Even the absence of a provisioned PTR records for an IP address may result in email messages not being accepted.

4.4.1. Scope and granularity

Endpoint-based blocking lacks some of the limitations of rendezvous-based blocking: while rendezvous-based blocking can only see and affect the rendezvous service at hand (e.g., DNS name resolution), endpoint-based blocking (depending on how it is designed) can have visibility into the entire application, across all layers and transactions. This visibility can provide endpoint-based blocking systems with a much richer set of information for making narrow blocking decisions.

In an endpoint-based blocking system, blocking actions are performed autonomously, by individual endpoints or their delegates. The effects of blocking are thus usually local in scope, minimizing the effects on other users or other, legitimate services.

4.4.2. Efficacy

Endpoint-based blocking deals well with mobile adversaries. If a blocked service relocates resources or uses different resources, a rendezvous- or network-based blocking approach may not be able to affect the new resources (at least not immediately). A network-based blocking system may not even be able to tell whether the new resources are being used, if the previously blocked service uses secure protocols. By contrast, endpoint-based blocking systems can detect when a blocked service's resources have changed (because of their full visibility into transactions) and adjust blocking as quickly as new blocking data can be sent out through a reputation system.

The primary challenge to endpoint-based blocking is that it requires the cooperation of endpoints. Where this cooperation is willing, this is a fairly low barrier, requiring only reconfiguration or software update. Where cooperation is unwilling, it can be challenging to enforce cooperation for large numbers of endpoints. That challenge is exacerbated when the endpoints are a diverse set of static, mobile or visiting endpoints. If cooperation can be achieved, endpoint-based blocking can be much more effective than other approaches because it is so coherent with the Internet's architectural principles.

4.4.3. Security

Endpoint-based blocking is performed at one end of an Internet communication, and thus avoids the problems related to end-to-end security mechanisms that network-based blocking runs into and the challenges to global trust infrastructures that rendezvous-based blocking creates.

4.4.4. Summary

Out of the three design patterns, endpoint-based blocking is the least likely to cause collateral damage to Internet services or the overall Internet architecture. Endpoint-based blocking systems have full visibility into all layers involved in a communication, allowing blocking to be narrowly targeted. Adversary mobility can be accounted for as soon as reputation systems are updated with new adversary information. The one potential drawback of endpoint-based blocking is that it requires the endpoint's cooperation; effectuating blocking at an endpoint when it is not in the endpoint's interest is therefore difficult to accomplish.

5. Security Considerations

The primary security concern related to Internet service blocking is the effect that it has on the end-to-end security model of many Internet security protocols. When blocking is enforced by an intermediary with respect to a given communication, the blocking system may need to obtain access to confidentiality-protected data to make blocking decisions. Mechanisms for obtaining such access often require the blocking system to defeat the authentication mechanisms built into security protocols.

For example, some enterprise firewalls will dynamically create TLS certificates under a trust anchor recognized by endpoints subject to blocking. These certificates allow the firewall to authenticate as any website, so that it can act as a man-in-the-middle on TLS connections passing through the firewall. This is not unlike an external attacker using compromised certificates to intercept TLS connections.

Modifications such as these obviously make the firewall itself an attack surface. If an attacker can gain control of the firewall or compromise the key pair used by the firewall to sign certificates, he will have access to the unencrypted data of all current and recorded TLS sessions for all users behind that firewall, in a way that is undetectable to users. Besides, if the compromised key-pairs can be extracted from the firewall, all users, not only those behind the firewall, that rely on that public key are vulnerable.

When blocking systems are unable to inspect and surgically block secure protocols, it is tempting to completely block those protocols. For example, a web blocking system that is unable to inspect HTTPS connections might simply block any attempted HTTPS connection. However, since Internet security protocols are commonly used for critical services such as online commerce and banking, blocking these protocols would block access to these services as well, or worse, force them to be conducted over insecure communication.

Security protocols can, of course, also be used as mechanisms for blocking services. For example, if a blocking system can insert invalid credentials for one party in an authentication protocol, then the other end will typically terminate the connection based on the authentication failure. However, it is typically much simpler to simply block secure protocols than to exploit those protocols for service blocking.

6. Conclusion

Because it is least likely to create technical or architectural problems, endpoint-based blocking is the form of Internet service blocking that is least harmful to the Internet. From a technical perspective, it is the most preferred option because it maintains transparency of the network, vests functionality at the endpoints, can be applied granularly so as to avoid collateral damage, and accommodates mobile adversaries. Entities seeking to filter and for whom endpoint-based blocking is a potential choice should view its technical benefits as distinct advantages compared to the other approaches.

In reality, the various approaches discussed above are all applied for different reasons, and particular entities may not consider endpoint-based filtering to be viable. Often, the choice of a filtering solution is constrained by practical limitations on which parts of the network are under the control of the entity implementing filtering, and which parts of the network are trusted to cooperate. For example, an ISP that is subject to filtering requirements might implement a network-based filtering approach because it cannot be sure that endpoints will cooperate in filtering. As discussed above, government agencies tasked with disabling certain foreign web sites have done so by manipulating infrastructure servers that are within their own jurisdictions, based on legal claims to obtain access to those servers. An enterprise with filtering requirements might require employees to install a certain filtering software package on enterprise-owned PCs.

It is therefore realistic to expect that certain entities will continue to attempt to conduct network- or rendezvous-based filtering

since they may not have control over the endpoints they wish to affect or because the endpoints do not have incentives to consent to the filtering. In some cases, an approach that combines one of these with endpoint-based filtering can help strike a better balance. For example, a filtering system might make it possible for some endpoints to cooperate or "opt in" to additional endpoint-based filtering, rather than deploying a purely network-based solution.

While this document has focused on technical mechanisms used to filter Internet content, a variety of non-technical mechanisms may also be available depending on the particular context and goals of the public or private entity seeking to restrict access to content. For example, purveyors of illegal online content can be pursued through international cooperation, by using the criminal justice system, and by targeting the funding that supports their activities through collaboration with financial services companies [click-trajectories]. Thus even in cases where endpoint-based filtering is not viewed as a viable means of restricting access to content, entities seeking to filter may find other strategies for achieving their goals that do not involve the operation of the Internet.

Those with a desire to filter should take into account the limitations discussed in this document and holistically assess the space of technical and non-technical solutions at their disposal and the likely effectiveness of each combination of approaches.

7. Informative References

[BT-TPB] Meyer, D., "BT blocks The Pirate Bay", June 2012, <<http://www.zdnet.com/bt-blocks-the-pirate-bay-4010026434/>>.

[CCS-GFC-collateral-damage]
, "The Collateral Damage of Internet Censorship by DNS Injection", July 2012, <<http://conferences.sigcomm.org/sigcomm/2012/paper/ccr-paper266.pdf>>.

[GhostClickRIPE]
RIPE NCC, "RIPE NCC Blocks Registration in RIPE Registry Following Order from Dutch Police", 2012, <<http://www.ripe.net/internet-coordination/news/about-ripe-ncc-and-ripe/ripe-ncc-blocks-registration-in-ripe-registry-following-order-from-dutch-police>>.

[I-D.blanchet-iab-internetoverport443]
Blanchet, M., "Implications of Blocking Outgoing Ports Except Ports 80 and 443", [draft-blanchet-iab-internetoverport443-01](#) (work in progress), October 2012.

[IN-OM-filtering]

Citizen Lab, ., "Routing Gone Wild", July 2012, <<https://citizenlab.org/2012/07/routing-gone-wild/>>.

[ISOCFiltering]

Internet Society, "DNS: Finding Solutions to Illegal On-line Activities", 2012, <<http://www.internetsociety.org/what-we-do/issues/dns/finding-solutions-illegal-line-activities>>.

[Malicious-Resolution]

Dagon, D., Provos, N., Lee, C., and W. Lee, "Corrupted DNS Resolution Paths: The Rise of a Malicious Resolution Authority", 2008, <http://www.citi.umich.edu/u/provos/papers/ndss08_dns.pdf>.

[Morris] Kehoe, B., "The Robert Morris Internet Worm", 1992, <<http://groups.csail.mit.edu/mac/classes/6.805/articles/morris-worm.html>>.

[RFC1122] Braden, R., "Requirements for Internet Hosts - Communication Layers", STD 3, [RFC 1122](#), October 1989.

[RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", [RFC 3261](#), June 2002.

[RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", [RFC 4033](#), March 2005.

[RFC4084] Klensin, J., "Terminology for Describing Internet Connectivity", [BCP 104](#), [RFC 4084](#), May 2005.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", [RFC 4301](#), December 2005.

[RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", [RFC 5246](#), August 2008.

[RFC5782] Levine, J., "DNS Blacklists and Whitelists", [RFC 5782](#), February 2010.

[RFC6480] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", [RFC 6480](#), February 2012.

[RenesysPK]

Brown, M., "Pakistan hijacks YouTube", February 2008, <http://www.renesys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml>.

[RojaDirecta]

Masnack, M., "Homeland Security Seizes Spanish Domain Name That Had Already Been Declared Legal", 2011, <<http://www.techdirt.com/articles/20110201/10252412910/homeland-security-seizes-spanish-domain-name-that-had-already-been-declared-legal.shtml>>.

[SAC-056] , "SSAC Advisory on Impacts of Content Blocking via the Domain Name System", October 2012, <<http://www.icann.org/en/groups/ssac/documents/sac-056-en.pdf>>.

[SafeBrowsing]

Google, "Safe Browsing API", 2012, <<https://developers.google.com/safe-browsing/>>.

[TPB-cloud]

, "The Pirate Cloud", October 2012, <<http://thepiratebay.se/blog/224>>.

[Telex] Wustrow, E., Wolchok, S., Goldberg, I., and J. Halderman, "Telex: Anticensorship in the Network Infrastructure", August 2011, <<https://telex.cc/>>.

[Tor] , "Tor Project: Anonymity Online", 2012, <<https://www.torproject.org/>>.

[US-ICE] U.S. Immigration and Customs Enforcement, "Operation in Our Sites", 2011, <<http://www.ice.gov/doclib/news/library/factsheets/pdf/operation-in-our-sites.pdf>>.

[asymmetry]

John, W., Dusi, M., and K. Claffy, "Estimating routing symmetry on single links by passive flow measurements", 2010, <<http://dl.acm.org/citation.cfm?id=1815506>>.

[click-trajectories]

Levchenko, K., Pitsillidis, A., Chacra, N., Enright, B., Felegyhazi, M., Grier, C., Halvorson, T., Kreibich, C., Liu, H., McCoy, D., Weaver, N., Paxson, V., Voelker, G., and S. Savage, "Click Trajectories: End-to-End Analysis of the Spam Value Chain", 2011, <<http://cseweb.ucsd.edu/~savage/papers/Oakland11.pdf>>.

Authors' Addresses

Richard Barnes
BBN Technologies
1300 N. 17th St
Arlington, VA 22209
USA

Phone: +1 703 284 1340
Email: rbarnes@bbn.com

Alissa Cooper
CDT
1634 Eye St. NW, Suite 1100
Washington, DC 20006
USA

Email: acooper@cdt.org

Olaf Kolkman
NLnet Labs
Science Park 400
Amsterdam 1422 JB
Netherlands

Email: olaf@nlnetlabs.nl

