

Internet Architecture Board (IAB)
Request for Comments: 7754
Category: Informational
ISSN: 2070-1721

R. Barnes
A. Cooper
O. Kolkman
D. Thaler
E. Nordmark
March 2016

Technical Considerations for Internet Service Blocking and Filtering

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. When it is possible to do so, 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. We observe that certain filtering and blocking approaches can cause unintended consequences to third parties, and we discuss the limits of efficacy of various approaches.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Architecture Board (IAB) and represents information that the IAB has deemed valuable to provide for permanent record. It represents the consensus of the Internet Architecture Board (IAB). Documents approved for publication by the IAB are not a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7754>.

Copyright Notice

Copyright (c) 2016 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.

Table of Contents

| | |
|--|----|
| 1. Introduction | 4 |
| 2. Filtering Examples | 5 |
| 3. Characteristics of Blocking Systems | 7 |
| 3.1. The Party Who Sets Blocking Policies | 8 |
| 3.2. Purposes of Blocking | 8 |
| 3.2.1. Blacklist vs. Whitelist Model | 9 |
| 3.3. Intended Targets of Blocking | 9 |
| 3.4. Components Used for Blocking | 10 |
| 4. Evaluation of Blocking Design Patterns | 11 |
| 4.1. Criteria for Evaluation | 11 |
| 4.1.1. Scope: What set of hosts and users are affected? . . | 12 |
| 4.1.2. Granularity: How specific is the blocking? Will | |
| blocking one service also block others? | 12 |
| 4.1.3. Efficacy: How easy is it for a resource or service to | |
| avoid being blocked? | 13 |
| 4.1.4. Security: How does the blocking impact existing trust | |
| infrastructures? | 14 |
| 4.2. Network-Based Blocking | 15 |
| 4.2.1. Scope | 16 |
| 4.2.2. Granularity | 17 |
| 4.2.3. Efficacy and Security | 17 |
| 4.2.4. Summary | 20 |
| 4.3. Rendezvous-Based Blocking | 20 |
| 4.3.1. Scope | 21 |
| 4.3.2. Granularity | 21 |
| 4.3.3. Efficacy | 21 |
| 4.3.4. Security and Other Implications | 22 |
| 4.3.5. Examples | 22 |
| 4.3.6. Summary | 23 |
| 4.4. Endpoint-Based Blocking | 24 |
| 4.4.1. Scope | 24 |
| 4.4.2. Granularity | 24 |
| 4.4.3. Efficacy | 25 |
| 4.4.4. Security | 25 |
| 4.4.5. Server Endpoints | 25 |
| 4.4.6. Summary | 26 |
| 5. Security Considerations | 26 |
| 6. Conclusion | 27 |
| 7. Informative References | 28 |
| IAB Members at the Time of Approval | 32 |
| Acknowledgments | 33 |
| Authors' Addresses | 33 |

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 perform filtering to restrict 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. However, governments, service providers, and enterprises are increasingly seeking to block or filter access to certain content, traffic, or services without the knowledge or agreement of affected users. Where these organizations do not directly control networks themselves, they commonly aim to make use of intermediary systems to implement the blocking or filtering.

While blocking and filtering remain highly contentious in many cases, the desire to restrict communications or 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 implemented 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. The scope of the analyzed blocking is limited to intentional blocking, not accidental blocking due to misconfiguration or as an unintentional side effect of something else.

Filtering may be considered legal, illegal, ethical, or unethical in different places, at different times, and by different parties. This document is intended for those who are conducting filtering or are considering conducting filtering and 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.

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 parties and for different purposes.

Firewalls: Firewalls of various sorts are very commonly employed at many points in today's Internet [RFC2979]. They can be deployed either on end hosts (under user or administrator control) or in the network, typically at network boundaries. While the Internet Security Glossary [RFC4949] contains an extended definition of a firewall, informally, most people would tend to think of a firewall as simply "something that blocks unwanted traffic" (see [RFC4948] for a discussion on many types of unwanted traffic). While there are many sorts of firewalls, there are several specific types of firewall functionality worth noting.

- o Stateless Packet Filtering: Stateless packet filters block according to content-neutral rules, e.g., blocking all inbound connections or outbound connections on certain ports, protocols, or network-layer addresses. For example, blocking outbound connections to port 25.
- o Stateful Packet Filtering: More advanced configurations require keeping state used to enforce flow-based policies, e.g., blocking inbound traffic for flows that have not been established.

- o Deep Packet Inspection: Yet more advanced configurations perform deep packet inspection and filter or block based on the content carried. Many firewalls include web filtering capabilities (see below).

Web Filtering: HTTP and HTTPS are common targets for blocking and filtering, typically targeted at specific URIs. 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 URI 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 on network- and transport-layer headers (IP address and/or port), 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" presenting their own, dynamically generated HTTPS certificates. (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 Blacklists use the reverse DNS to flag whether an IP address is a known spam source [RFC5782]. Spam filters can be installed on user devices (e.g., in a mail client), operated by a mail domain on behalf of users, or outsourced to a third party that acts as an intermediate MX proxy.

Domain Name Seizure: A number of approaches are used to block or modify resolution of a domain name. One approach is to make use of ICANN's Uniform Dispute Resolution Policy (URDP) for the purposes of dealing with fraudulent use of a name. Other authorities may require that domains be blocked within their jurisdictions. Substantial research has been performed on the value and efficacy of such seizures [Takedown08] [BlackLists14].

The precise method of how domain names are seized will vary from place to place. One approach in use is for queries to be redirected to resolve to IP addresses of the authority that hosts information about the seizure. The effectiveness of domain seizures will similarly vary based on the method. In some cases, the person whose name was seized will simply use a new name. In other cases, the block may only be effective within a region or when specific name service infrastructure is used.

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 URI, current versions of Google Chrome and Firefox use the Google Safe Browsing service to determine whether or not a given URI 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.

Ingress filtering: Network service providers use ingress filtering [RFC2827] [RFC3704] as a means to prevent source address spoofing which is used as a part of other attacks.

Data loss prevention (DLP): Enterprise and other networks are concerned with potential leaking of confidential information, whether accidental or intentional. Some of the tools used for this are similar to the main subject of this document of blocking and filtering. In particular, enterprise proxies might be part of a DLP solution.

3. Characteristics of Blocking Systems

At a generic level, blocking systems can be characterized by four attributes: the party who 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. The Party Who Sets Blocking Policies

Parties that institute blocking policies include governments, courts, enterprises, network operators, reputation trackers, application providers, and individual end users. A government might create laws based on cultural norms and/or their elected mandate. Enterprises might use cultural, industry, or legal norms to guide their policies.

There can be several steps of translation and transformation from the original intended purpose -- first to laws, then to (government) regulation, followed by high-level policies in, e.g., network operators, and from those policies to filtering architecture and implementation. Each of those steps is a potential source of unintended consequences as discussed in this document.

In some cases, the policy setting entity is the same as the entity that enforces the policy. 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, the policy setting entity is different from the entity that enforces the policy. Such policy might be imposed upon the enforcing entity, such as in the case of blocking initiated by governments, or the enforcing entity might explicitly choose to use policy set by others, such as in the case of a reputation system used by a spam filter or safe browsing service. Because a policy might be enforced by others, it is best if it can be expressed in a form that is independent of the enforcing technology.

3.2. Purposes of Blocking

There are a variety of motivations to filter:

- 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. 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).

3.2.1. Blacklist vs. Whitelist Model

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

Blocking systems are instituted 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 URIs whose content is deemed by the enterprise to be inappropriate for the workplace. 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, when domain name seizures described above are intended to address specific web pages associated with illegal activity, 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 that may be unrelated to the illegal activity. Depending on where the block is imposed within the DNS hierarchy, entirely unrelated organizations may be impacted.

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 typically an application-layer protocol between two or more 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. An example of a rendezvous service is the domain name system. Distributed Hash Tables (DHTs) have also been used as rendezvous services.

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 endpoint then responds to the HTTP request.

As another example, in the SIP protocol, the two endpoints communicating 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 [Endpoint] Preventing the client from making the DNS request needed to resolve example.com
- 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 servers
- o [Network] Preventing the DNS responses from reaching the client
- o [Rendezvous] Preventing the DNS servers from providing the client the correct IP address of the server

Those who desire to block communications will typically have access to only one or two components; therefore their choices for how to perform 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 perform 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 who 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 who desires a 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 set of hosts and users are affected?

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 topological scope. For example, blocking within an access network will only affect a well-defined set of users (namely, those connected to the access network). Blocking performed by an application provider can affect users across the entire Internet.

Blocking systems are generally viewed as less objectionable if the scope of their impact is as narrow as possible while still being effective, and as long as the impact of the blocking is within the administrative realm of the policy setting entity. As mentioned previously, enterprise blocking systems are commonly deployed, and will generally have impact on enterprise users. However, design flaws in blocking systems may cause the effects of blocking to be overbroad. For example, at least one service provider blocking content in accordance with a regulation has ended up blocking content for downstream service providers because it filtered routes to particular systems and did not distribute the original information to downstream service providers in other jurisdictions [IN-OM-filtering]. Other service providers have accidentally leaked such black hole routes beyond the jurisdiction [NW08]. A substantial amount of work has gone into BGP security to avoid such attacks, but deployment of such systems lags.

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 parties. 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 other 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 [RFC4924].

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.

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

Although blocking a resource or service might have some immediate effect, efficacy must be evaluated in terms of whether it is easy to circumvent. Simply doing a one-time policy is often unlikely to have lasting efficacy (e.g., see [CleanFeed] and [BlackLists14]). Experience has shown that, in general, blacklisting requires continual maintenance of the blacklist itself, both to add new entries for unwanted traffic and deleting entries when offending content is removed. Experience also shows that, depending on the nature of the block, it may be difficult to determine when to unblock. For instance, if a host is blocked because it has been compromised and used as a source of attack, it may not be plainly evident when that site has been fixed.

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 reconfigured to use non-blocked resources. If a web site is prevented from using a domain name or set of IP addresses, the content can simply be moved to another domain name or network, or use alternate syntaxes to express the same resource name (see the discussion of false negatives in [RFC6943]).

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, there are several ways in which a user or application can try to circumvent the filters.

The users may choose to use different sets of protocols or otherwise alter their traffic characteristics to circumvent the filters. In some cases, applications may shift their traffic to port 80 or 443 when other ports are blocked. Or, services may be tunneled within other services, proxied by a collaborating external host (e.g., an anonymous redirector), or simply run over an alternate port (e.g., port 8080 vs port 80 for HTTP). Another means of circumvention is alteration of the service behavior to use a dynamic port negotiation phase, in order to avoid use of a constant port address.

One of the primary motivations for arguing that HTTP/2 should be encrypted by default was that unencrypted HTTP 1.1 traffic was sometimes blocked or improperly processed. Users or applications shifting their traffic to encrypted HTTP has the effect of circumventing filters that depend on the HTTP plaintext payload.

If voice communication based on SIP [RFC3261] is blocked, users are likely to use applications which use proprietary protocols that allow them to talk to each other.

Some filtering systems are only capable of identifying IPv4 traffic and therefore, by shifting to IPv6, users may be able to evade filtering. Using IPv6 with header options, using multiple layers of tunnels, or using encrypted tunnels can also make it more challenging for blocking systems to find transport ports within packets, making port-based blocking more difficult. 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 Transport Layer Security (TLS) [RFC5246] and IPsec [RFC4301] are designed to ensure that each endpoint of the communication knows the identity of the other endpoint(s) 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 [RFC4924]. In these cases, it can be difficult or impossible for endpoints to distinguish between attackers and "authorized" parties conducting blocking. For example, an enterprise firewall administrator could gain access to users' personal bank accounts when users on the enterprise network connect to bank web sites.

Finally, one needs to evaluate whether a blocking mechanism can be used by an end user to efficiently locate blocked resources that can then be accessed via other mechanisms that circumvent the blocking mechanism. For example, Clayton [CleanFeed] showed how special treatment in one blocking system could be detected by end users in order to efficiently locate illegal web sites, which was thus counterproductive to the policy objective of the blocking mechanism.

4.2. Network-Based Blocking

Being able to block access to resources without the consent or cooperation of either endpoint is viewed as a desirable feature by some 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 URI being requested, compare that URI to a list of blacklisted or whitelisted URIs, and allow the request to proceed only if it is permitted by policy. 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 Access Control Lists (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

In the case of government-initiated blocking, network operators subject to a specific jurisdiction may be required to block or filter. Thus, it is possible for laws to be structured to result in blocking by imposing obligations on the operators of networks within a jurisdiction, either via direct government action or by allowing private actors to demand blocking (e.g., through lawsuits).

Regardless of who is responsible for a blocking policy, enforcement can be done using Stateless Packet Filtering, Stateful Packet Filtering, or Deep Packet Inspection as defined in Section 2. While network-based Stateless Packet Filtering has granularity issues discussed in Section 4.2.2, network-based Stateful Packet Filtering and Deep Packet Inspection approaches often run into several technical issues that limit their viability in practice. For example, 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 step to obtaining this traffic is simply gaining access to the constituent packets. The Internet is designed to deliver packets independently from source to destination -- not to any particular point along the way. Thus, the sequence of packets from the sender can only be reliably reconstructed at the intended receiver. In addition, inter-network routing is often asymmetric, and for sufficiently complex local networks, intra-network traffic flows can be asymmetric as well [asymmetry]. Thus, packets in the reverse direction use a different set of paths than the forward direction.

This asymmetry means that an intermediary in a network with many egress points may, depending on topology and configuration, see only one half of a given communication, which may limit the scope of the communications that it can filter. For example, a filter aimed at requests destined for particular URIs cannot make accurate blocking decisions based on the URI if it is only in the data path for HTTP responses and not requests, since the URI is not included in the responses. Asymmetry may be surmountable given a filtering system

with enough distributed, interconnected filtering nodes that can coordinate information about flows belonging to the same communication or transaction, but depending on the size of the network this may imply significant complexity in the filtering system. 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 (URIs or other content). Deep Packet Inspection type 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 those who seek to block must 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 [BT-TPB].

Indeed, Poort, et al. [Poort] found that "any behavioural change in response to blocking access to The Pirate Bay has had no lasting net impact on the overall number of downloaders from illegal sources, as new consumers have started downloading from illegal sources and people learn to circumvent the blocking while new illegal sources may be launched, causing file sharing to increase again", and that these

results "are in line with a tendency found in the literature that any effects of legal action against file sharing often fade out after a period of typically six months."

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, or resort to statistical methods to guess what the content is. Since security protocols are generally 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 [RFC6698]) 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, port numbers, packet sizes, and packet timing. 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 [Telex] 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. For example, if an ISP used by a banking web site 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 from that system. (The only limitation in practice is the requirement for special software on the client.) Those who require that secure traffic be blocked from such sites risk blocking content that would be valuable to their users, perhaps impeding substantial economic activity. Conversely, those who are hosting a myriad of content have an incentive to see that law abiding content does not end up being blocked.

Governments and network operators should, however, take care not to encourage the use of insecure communications in the naming of security, as doing so will invariably 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 such 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.

One reason above all stands as an enormous challenge to network-based blocking: the Internet was designed with the premise that people will want to connect and communicate. IP will run on anything up to and including carrier pigeons [RFC1149]. It often runs atop TLS and has been made to run on other protocols that themselves run atop IP. Because of this fundamental layering approach, nearly any authorized avenue of communication can be used as a transport. This same "problem" permits communications to succeed in the most challenging of environments.

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. The fact that network-based blocking is premised on this lack of cooperation results in arms races that increase the complexity of both application design and network design.

4.3. Rendezvous-Based Blocking

Internet applications often require or rely on support from common, global rendezvous services, including the DNS, certificate authorities, search engines, 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) 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, and the Resource Public Key Infrastructure (RPKI) stores information about which organizations have been allocated IP addresses. To offer specialized Internet services and applications, different people 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 operators of servers used to provide rendezvous service that are subject to a specific jurisdiction may be required to block or filter. Thus, it is possible for laws to be structured to result in 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 others will depend on which servers they 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.

However, if a service is hosted and operated within a jurisdiction where it is considered legitimate, then blocking access at a global rendezvous service (e.g., one within a jurisdiction where it is considered illegitimate) might deny services in jurisdictions where they are considered legitimate. This type of collateral damage is lessened when blocking is done at a local rendezvous server that only has local impact, rather than at a global rendezvous server with global impact.

4.3.2. Granularity

Blocking at a global rendezvous service can be overbroad if the resources blocked support multiple services, since blocking service can cause collateral damage to legitimate uses of other services. For example, a given address or domain name might host both legitimate services as well as services that some would desire to block.

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 a global 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.

Endpoints can also choose not to use a particular rendezvous service. They might switch to a competitor or use an alternate mechanism (for example, IP literals in URIs to circumvent DNS filtering).

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 fragmentation, 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 RFC 2826 [RFC2826], in the advisory [SAC-056] from ICANN's Security and Stability Advisory Committee (SSAC), and in the Internet Society'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 Pakistani ISP's scope, but service was restored in approximately two hours because network operators around the world reconfigured their routers to ignore the bogus routes [RenesysPK]. In the context of SIDR and secure routing, a similar reconfiguration 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 U.S. law enforcement seizing domain names involved RojaDirecta, a Spanish web site. Even though several of the affected domain names belonged to Spanish organizations, they were subject to blocking by the U.S. government because certain servers were operated in the United States.

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 U.S.-government-operated name servers. Users of other services (e.g., email) under a target name 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 unblocked results for the affected name. However, depending on the 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 of 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

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, Mail Transfer Agents (MTAs)) share a single reputation service. Even the absence of provisioned PTR records for an IP address may result in email messages not being accepted.

4.4.1. Scope

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. Granularity

Endpoint-based blocking avoids 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 can potentially see 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. Support for narrow granularity depends on how the application protocol client and server are designed, however. A typical endpoint-based firewall application

may have less ability to make fine-grained decisions than an application that does its own blocking (see [RFC7288] for further discussion).

4.4.3. 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.

4.4.4. 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.

4.4.5. Server Endpoints

In this discussion of endpoint-based blocking, the focus has been on the consuming side of the end-to-end communication, mostly the client side of a client-server type connection. However, similar considerations apply to the content-producing side of end-to-end communications, regardless of whether that endpoint is a server in a client-server connection or a peer in a peer-to-peer type of connection.

For instance, for blocking of web content, narrow targeting can be achieved through whitelisting methods like password authentication, whereby passwords are available only to authorized clients. For example, a web site might only make adult content available to users who provide credit card information, which is assumed to be a proxy for age.

The fact that content-producing endpoints often do not take it upon themselves to block particular forms of content in response to requests from governments or other parties can sometimes motivate those latter parties to engage in blocking elsewhere within the Internet.

If a service is to be blocked, the best way of doing that is to disable the service at the server endpoint.

4.4.6. 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 can potentially see into all layers involved in a communication, allowing blocking to be narrowly targeted and can minimize unintended consequences. Adversary mobility can be accounted for as soon as reputation systems are updated with new adversary information. One potential drawback of endpoint-based blocking is that it requires the endpoint's cooperation; implementing blocking at an endpoint when it is not in the endpoint's interest is therefore difficult to accomplish because the endpoint's user can disable the blocking or switch to a different endpoint.

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 web site, 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, the attacker 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.

We must also consider the possibility that a legitimate administrator of such a firewall could gain access to privacy-sensitive information, such as the bank accounts or health records of users who access such secure sites through the firewall. These privacy considerations motivate legitimate use of secure end-to-end protocols that often make it difficult to enforce granular blocking policies.

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

Filtering will continue to occur on the Internet. We conclude that, whenever possible, filtering should be done on the endpoint. Cooperative endpoints are most likely to have sufficient contextual knowledge to effectively target blocking; hence, such blocking minimizes unintended consequences. It is realistic to expect that at times filtering will not be done on the endpoints. In these cases, promptly informing the endpoint that blocking has occurred provides necessary transparency to redress any errors, particularly as they relate to any collateral damage introduced by errant filters.

Blacklist approaches are often a game of "cat and mouse", where those with the content move it around to avoid blocking. Or, the content may even be naturally mirrored or cached at other legitimate sites such as the Internet Archive Wayback Machine [Wayback]. At the same time, whitelists provide similar risks because sites that had "acceptable" content may become targets for "unacceptable content", and similarly, access to perfectly inoffensive and perhaps useful or productive content is unnecessarily blocked.

From a technical perspective, there are no perfect or even good solutions -- there is only least bad. On that front, we posit that a hybrid approach that combines endpoint-based filtering with network filtering may prove least damaging. An endpoint may choose to participate in a filtering regime in exchange for the network providing broader unfiltered access.

Finally, we note that where filtering is occurring to address content that is generally agreed to be inappropriate or illegal, strong cooperation among service providers and governments may provide additional means to identify both the victims and the perpetrators through non-filtering mechanisms, such as partnerships with the finance industry to identify and limit illegal transactions.

7. Informative References

[asymmetry]

John, W., Dusi, M., and K. Claffy, "Estimating routing symmetry on single links by passive flow measurements", Proceedings of the 6th International Wireless Communications and Mobile Computing Conference, IWCMC '10, DOI 10.1145/1815396.1815506, 2010, <http://www.caida.org/publications/papers/2010/estimating_routing_symmetry/estimating_routing_symmetry.pdf>.

[BlackLists14]

Chachra, N., McCoy, D., Savage, S., and G. Voelker, "Empirically Characterizing Domain Abuse and the Revenue Impact of Blacklisting", Workshop on the Economics of Information Security 2014, <<http://www.econinfosec.org/archive/weis2014/papers/Chachra-WEIS2014.pdf>>.

[BT-TPB]

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

[CleanFeed]

Clayton, R., "Failures in a Hybrid Content Blocking System", Fifth Privacy Enhancing Technologies Workshop, PET 2005, DOI 10.1007/11767831_6, 2005, <<http://www.cl.cam.ac.uk/~rnc1/cleanfeed.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>>.

[IN-OM-filtering]

Citizen Lab, "Routing Gone Wild: Documenting upstream filtering in Oman via India", 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>>.

[NW08]

Marsan, C., "YouTube/Pakistan incident: Could something similar whack your site?", Network World, March 2008, <<http://www.networkworld.com/article/2284273/software/youtube-pakistan-incident--could-something-similar-whack-your-site-.html>>.

[Poort]

Poort, J., Leenheer, J., van der Ham, J., and C. Dumitru, "Baywatch: Two approaches to measure the effects of blocking access to The Pirate Bay", Telecommunications Policy 38:383-392, DOI 10.1016/j.telpol.2013.12.008, 2014, <<http://staff.science.uva.nl/~vdham/research/publications/1401-Baywatch.pdf>>.

- [RenesysPK] Brown, M., "Pakistan hijacks YouTube", February 2008, <<http://research.dyn.com/2008/02/pakistan-hijacks-youtube-1/>>.
- [RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, DOI 10.17487/RFC1122, October 1989, <<http://www.rfc-editor.org/info/rfc1122>>.
- [RFC1149] Waitzman, D., "Standard for the transmission of IP datagrams on avian carriers", RFC 1149, DOI 10.17487/RFC1149, April 1990, <<http://www.rfc-editor.org/info/rfc1149>>.
- [RFC2826] Internet Architecture Board, "IAB Technical Comment on the Unique DNS Root", RFC 2826, DOI 10.17487/RFC2826, May 2000, <<http://www.rfc-editor.org/info/rfc2826>>.
- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827, May 2000, <<http://www.rfc-editor.org/info/rfc2827>>.
- [RFC2979] Freed, N., "Behavior of and Requirements for Internet Firewalls", RFC 2979, DOI 10.17487/RFC2979, October 2000, <<http://www.rfc-editor.org/info/rfc2979>>.
- [RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, DOI 10.17487/RFC3261, June 2002, <<http://www.rfc-editor.org/info/rfc3261>>.
- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, DOI 10.17487/RFC3704, March 2004, <<http://www.rfc-editor.org/info/rfc3704>>.
- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, DOI 10.17487/RFC4033, March 2005, <<http://www.rfc-editor.org/info/rfc4033>>.
- [RFC4084] Klensin, J., "Terminology for Describing Internet Connectivity", BCP 104, RFC 4084, DOI 10.17487/RFC4084, May 2005, <<http://www.rfc-editor.org/info/rfc4084>>.

- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <<http://www.rfc-editor.org/info/rfc4301>>.
- [RFC4924] Aboba, B., Ed. and E. Davies, "Reflections on Internet Transparency", RFC 4924, DOI 10.17487/RFC4924, July 2007, <<http://www.rfc-editor.org/info/rfc4924>>.
- [RFC4948] Andersson, L., Davies, E., and L. Zhang, "Report from the IAB workshop on Unwanted Traffic March 9-10, 2006", RFC 4948, DOI 10.17487/RFC4948, August 2007, <<http://www.rfc-editor.org/info/rfc4948>>.
- [RFC4949] Shirey, R., "Internet Security Glossary, Version 2", FYI 36, RFC 4949, DOI 10.17487/RFC4949, August 2007, <<http://www.rfc-editor.org/info/rfc4949>>.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <<http://www.rfc-editor.org/info/rfc5246>>.
- [RFC5782] Levine, J., "DNS Blacklists and Whitelists", RFC 5782, DOI 10.17487/RFC5782, February 2010, <<http://www.rfc-editor.org/info/rfc5782>>.
- [RFC6480] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", RFC 6480, DOI 10.17487/RFC6480, February 2012, <<http://www.rfc-editor.org/info/rfc6480>>.
- [RFC6698] Hoffman, P. and J. Schlyter, "The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA", RFC 6698, DOI 10.17487/RFC6698, August 2012, <<http://www.rfc-editor.org/info/rfc6698>>.
- [RFC6943] Thaler, D., Ed., "Issues in Identifier Comparison for Security Purposes", RFC 6943, DOI 10.17487/RFC6943, May 2013, <<http://www.rfc-editor.org/info/rfc6943>>.
- [RFC7288] Thaler, D., "Reflections on Host Firewalls", RFC 7288, DOI 10.17487/RFC7288, June 2014, <<http://www.rfc-editor.org/info/rfc7288>>.

- [RojaDirecta] Masnick, 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] ICANN SSAC, "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/>>.
- [Takedown08] Moore, T. and R. Clayton, "The Impact of Incentives on Notice and Take-down", Workshop on the Economics of Information Security 2008, <<http://www.econinfosec.org/archive/weis2008/papers/MooreImpact.pdf>>.
- [Telex] Wustrow, E., Wolchok, S., Goldberg, I., and J. Halderman, "Telex: Anticensorship in the Network Infrastructure", <<https://telex.cc/>>.
- [Tor] "Tor Project: Anonymity Online", <<https://www.torproject.org/>>.
- [Wayback] "Internet Archive: Wayback Machine", <<http://archive.org/web/>>.

IAB Members at the Time of Approval

Jari Arkko
Mary Barnes
Marc Blanchet
Ralph Droms
Ted Hardie
Joe Hildebrand
Russ Housley
Erik Nordmark
Robert Sparks
Andrew Sullivan
Dave Thaler
Brian Trammell
Suzanne Woolf

Acknowledgments

Thanks to the many reviewers who provided helpful comments, especially Bill Herrin, Eliot Lear, Patrik Faltstrom, Pekka Savola, and Russ White. NLnet Labs is also acknowledged as Olaf Kolkman's employer during most of this document's development.

Authors' Addresses

Richard Barnes
Mozilla
Suite 300
650 Castro Street
Mountain View, CA 94041
United States

Email: rlb@ipv.sx

Alissa Cooper
Cisco
707 Tasman Drive
Milpitas, CA 95035
United States

Email: alcoop@cisco.com

Olaf Kolkman
Internet Society

Email: kolkman@isoc.org

Dave Thaler
Microsoft
One Microsoft Way
Redmond, WA 98052
United States

Email: dthaler@microsoft.com

Erik Nordmark
Arista

Email: nordmark@arista.com