March 2026

Mandated DNS Blocking

A Guide for Legal Professionals

Carl Gahnberg

Internet Society

Table of Contents

Executive Summary

The Domain Name System (DNS) provides a critical function for the Internet. It enables people to use familiar words, like "example.com", instead of long numerical addresses to reach websites, send emails, and use online services. Without this naming layer, the Internet would not be the practical, global communications system the world depends on today.

Because the DNS offers this function, it is sometimes treated as a convenient point of enforcing public policy, especially since voluntary or parental DNS filtering is a common practice. Governments and courts increasingly turn to mandated DNS blocking to stop access to certain content, whether to combat child exploitation, copyright infringement, restrict online gambling, or address politically sensitive or unlawful material.

The message of this report is simple: Mandated DNS blocking may look like a straightforward technical fix to enforce public policy, but in practice, it is blunt, costly, and even counterproductive. Understanding why requires viewing the DNS as a shared global infrastructure and recognizing the impact when it is tasked with a job it was not designed for.

Why DNS Blocking Is a Blunt Tool

DNS resolution only translates a domain name to the IP address of a web server, not the full path of a resource. For example, when a user enters "https://example.com/page1" in their web browser, the DNS only resolves the "example.com" part to an IP address. The rest of the path, i.e., the "/page1" part (or any specific image, video, or file), is handled once a connection has been established with the web server. This means DNS blocking can only apply to entire domain names, not to individual pages or files.

DNS Blocking Breaks Things

If a single page under a domain is unlawful, blocking at the DNS level prevents access to all the other, lawful material hosted under the same domain. The impact can be particularly significant for shared platforms such as social networks, blogging services, or cloud services, where millions of distinct users might rely on the same domain name. A single block can therefore disrupt vast amounts of unrelated content and services. This includes important implications for security, as DNS blocking orders hinder the widespread adoption of Domain Name System Security Extensions (DNSSEC).

DNS Blocking Is Easy to Circumvent and Ineffective

Determined users have many easily accessible (and even automated) circumvention techniques to avoid DNS blocking. This drastically reduces the long-term effectiveness of DNS blocking.

Furthermore, DNS blocking only prevents a resolver from translating a blocked domain name into its corresponding IP address. It does not remove the underlying content from the Internet. In practice, the material usually remains available and can be accessed again once it is tied to a new domain name.

DNS Blocking Has Effects Across Geographies and Jurisdictions

Policies for blocking content can affect users far beyond the intended jurisdiction unless the operator can reliably and very quickly (within milliseconds) distinguish users' locations. However, this is complicated by the limitations of IP-based geolocation, which can result in lawful users outside the intended jurisdiction being denied access, while targeted users may still be able to access it.

Furthermore, because the DNS is a global system, blocking measures often collide with other national legal frameworks. What may be prohibited in one jurisdiction may be legal in another, and a blocking order issued domestically cannot easily be reconciled with such differences. The operator could then face a legal dilemma: comply with the order and risk violating the rights of users or laws of another jurisdiction, or refuse and face penalties in the jurisdiction imposing the block.

DNS Blocking Comes at a Cost

Implementing DNS blocking can be costly for both operators and users. Implementing the block, particularly when applied selectively by jurisdiction, requires not only technical changes to resolver infrastructure but also ongoing operational investment. The more fragmented the rules become, the greater the operational burden on DNS operators and the greater chance of error in over- or under-blocking. From a user's perspective, blocking also risks affecting service quality as integrating geolocation and blocklist checks into DNS resolution increases latency. Although delays of just a few milliseconds may seem negligible, at the scale of the Internet, they can have a significant impact on user experience as complex pages load much slower as each component is reached via DNS and undergoes a blocklist check. This can also cause users to switch to alternative resolvers, which undermines both compliance and business goals.

For all these reasons, mandated DNS blocking is the wrong tool for public policy enforcement. If online harms need to be addressed, interventions should focus on the content itself, the actors responsible, and measures grounded in due process and international cooperation. The DNS exists to make the Internet usable, not to serve as a mechanism of control. Preserving its universality, reliability, and security is essential to maintaining an open, resilient, and global Internet.

Introduction

The Domain Name System (DNS) provides a critical function to the Internet. It provides the naming function that lets people use familiar words, like “example.com”, instead of long numerical addresses to reach websites, send emails, and use online services. Without this naming layer, the Internet would not be the practical, global communications system the world depends on today.

Because the DNS offers this function, it is sometimes treated as a convenient point of enforcing public policy, especially since voluntary or parental DNS filtering is a common practice. Governments and courts increasingly turn to mandated DNS blocking to stop access to certain content, whether to combat child exploitation, address copyright infringement, restrict online gambling, or address politically sensitive or unlawful material. At first glance, this approach appears straightforward: if a name cannot be translated to its corresponding IP address, the content cannot be reached.

However, this apparent simplicity is misleading. The DNS was designed for universality, consistency, and usability, not as a public policy enforcement tool. Repurposing it to this end introduces technical and operational side effects, creates legal and jurisdictional tensions, and undermines important security measures that keep users safe. Importantly, these interventions are inherently blunt: they cannot target individual pages or files, are prone to collateral damage, and are easily circumvented because content is obscured but not actually removed. Furthermore, the global, decentralized nature of the DNS makes location-based blocking hard to confine to a single jurisdiction.

This report is specifically focused on this form of mandated DNS blocking. It begins with a technical primer on the Internet, how its addressing and naming functions work together, and how DNS resolution happens in practice. It then examines the real-world effects of DNS blocking, including the risks of over-blocking and fragmentation, the incompatibility with security mechanisms, and the ease with which blocked content remains accessible. It also considers geography and jurisdiction, how a global naming system intersects with national laws, and the practical costs of implementing blocking orders.

The message of this report is simple: DNS blocking may appear to be a straightforward technical fix to enforce public policy, but in practice, it is blunt, costly, and even counterproductive. Understanding why requires seeing the DNS as a shared global infrastructure and recognizing the impact when put to do a job it was not designed for.

Key Legal Considerations at a Glance

Any measure requiring DNS blocking should be evaluated against the following legal principles: the scope of the order should be precise and limited to what is necessary ("narrow tailoring"); what the order requires should be proportionate and not excessive ("proportionality"); the order respects the jurisdiction and laws of other countries ("comity"); and due process has been applied ("due process'). The questions below are intended as a quick reference guide to help legal professionals apply these principles when assessing DNS blocking measures.

Narrow Tailoring

Context: The DNS can only block at the level of domain names, not individual files, URLs, or specific pieces of content.

1. Would the block prevent access to lawful material hosted under the same domain?
2. Would the block affect users or services that are not implicated in the underlying claim?

Proportionality

Context: DNS blocking is technically blunt and can be easily circumvented (e.g. by changing resolvers, use of VPNs, etc.), and more direct and targeted measures may exist.

1. Could the same objective be achieved through more precise actions aimed at the content provider, or actors closer to the violation (e.g., the hosting provider)?
2. Is the block likely to be ineffective and easily circumvented?
3. What impact could the order have on innocent users and service providers?

Comity

Context: The DNS is geographically distributed across borders, and many DNS resolvers can serve users and services located in multiple jurisdictions. A DNS blocking order could create cross-border effects.

1. Does the DNS resolver serve users in multiple jurisdictions?
2. Would the block affect users and services in other jurisdictions and potentially create conflicts of laws and rights?

Due Process

Context: DNS resolver operators may be forced to block a large number of domains in ways they cannot safely or accurately control. This increases the risk of errors, over-blocking, and the global propagation of mistakes.

1. Is there a reliable process for correcting mistakes, updating blocking orders, and addressing operators' questions so that errors or over-blocking can be corrected?
2. Has the operator been given timely notice of the block, including its legal basis, and is there appropriate oversight for how the blocklist might expand?

Understanding the Domain Name System (DNS) - A Technical Primer

The Internet: A Unique Networking Model

The Internet is commonly referred to as a "network of networks," reflecting its core purpose of enabling communications across different computer networks. While this is an accurate description, it also downplays the Internet's broader significance. In fact, the idea of interconnecting computer networks predates the Internet, but earlier systems were typically closed, limited to specific organizations, or restricted to particular technologies[1]. By contrast, the Internet represents a fundamentally different approach, of a distributed and collaborative model based on open standards, which has enabled it to grow to a global scale. From this view, the Internet is not merely a generic data network or a collection of technical protocols. It is a distinct *networking model*, aimed at voluntarily interconnecting diverse and independent networks into a single global communications system[2].

The Internet can be understood as an infrastructure providing a simple function: the ability to send data packets between computers located in different networks operated by different organizations. The connected computers then rely on shared standards to build applications on top of this infrastructure. For example, computers connected to the Internet created the World Wide Web by adopting the Hypertext Transfer Protocol (HTTP), which web browsers and servers use to exchange web pages. This design, in which the network simply delivers packets while the connected computers decide how to use them, explains why the Internet is often described as a general-purpose technology. The same underlying system supports email, the Web, instant messaging, video streaming, online banking, and countless other applications. The choice of what to build and how to use connectivity rests with the connected computers, not with the network itself.

The success of this networking model, in terms of its global reach, capacity for innovation, and continuous evolution, lies as much in its governance as in its technology. The Internet is not centrally controlled but operates through the voluntary adoption of open standards developed in multistakeholder processes[3]. Each network retains autonomy while participating in a larger whole, which creates resilience and adaptability.

The ability for the independent networks to function as a single system is, in turn, dependent on two key functions. First is "addressing", which provides each connected device (such as a server, laptop, or

---

[1] Barry M. Leiner et al., *A Brief History of the Internet* (Internet Society, 2009), https://www.internetsociety.org/internet/history-internet/brief-history-internet/.

[2] Internet Society, *The Internet Way of Networking (IWN): Defining the Critical Properties of the Internet* (2019), https://www.internetsociety.org/iwn/.

[3] 'Why the Multistakeholder Approach Works', *Internet Society*, 26 April 2016, https://www.internetsociety.org/resources/doc/2016/internet-governance-why-the-multistakeholder-approach-works/.

phone) with a numerical identifier so that data packets find their way across the networks and reach the correct destination. Second is "naming", which translates human-readable names into numerical identifiers, allowing people to locate resources in a usable way. Together, addressing and naming make the Internet both globally coherent and practically accessible, turning it into a global communications medium.

- Addressing: Every device connected to the Internet is assigned by the network it connects to a unique numerical identifier called an Internet Protocol (IP) address. Like a postal address, it ensures that data sent from one device can reach the correct destination. The IP addresses are allocated in blocks to network operators by the Regional Internet Registries (RIRs), which manage IP address allocations on a regional basis.
- Naming: While IP addresses are essential for computers, they are difficult for people to remember or use. To make the Internet more usable, the DNS translates human-readable names (such as "example.com") into IP addresses (e.g., 203.0.113.42). Much like IP addresses, domain names are managed through an allocation system. They are organized in a hierarchical and decentralized manner, with responsibility for each "zone" delegated to a zone operator that manages a set of authoritative servers. At the top levels of the hierarchy, this role is carried out by registries (for example, those managing *.com* or *.org*), often working with registrars who resell names directly to end-users. This naming layer, and the infrastructure that supports it, enables people to interact with the Internet in a practical and intuitive way.

The distinction between addressing and naming is important for understanding how the Internet works and the impact of potential policy interventions. Whereas IP addresses are the numbers used by computers to deliver data to the correct destination, domain names are a layer of abstraction that makes it easier for people to use the Internet. Measures that alter one do not necessarily affect the other, and technical controls applied at either layer can have very different operational and policy consequences.

The DNS Lookup – A Practical Guide

As described in the previous section, DNS provides the naming function that enables the Internet to function. While computers use IP addresses to identify one another, these are difficult for people to remember or use. Thus, the role of the DNS is to translate human-friendly names, such as "example.com", into the numerical identifiers that networks require to deliver data. In doing so, it allows people to interact with the Internet in a simple and intuitive way, while computers continue to rely on the numerical identifiers needed for packet delivery.

The translation process between names and addresses involves several components working together. A simplified model of these components is visualized in the following diagram[4]:

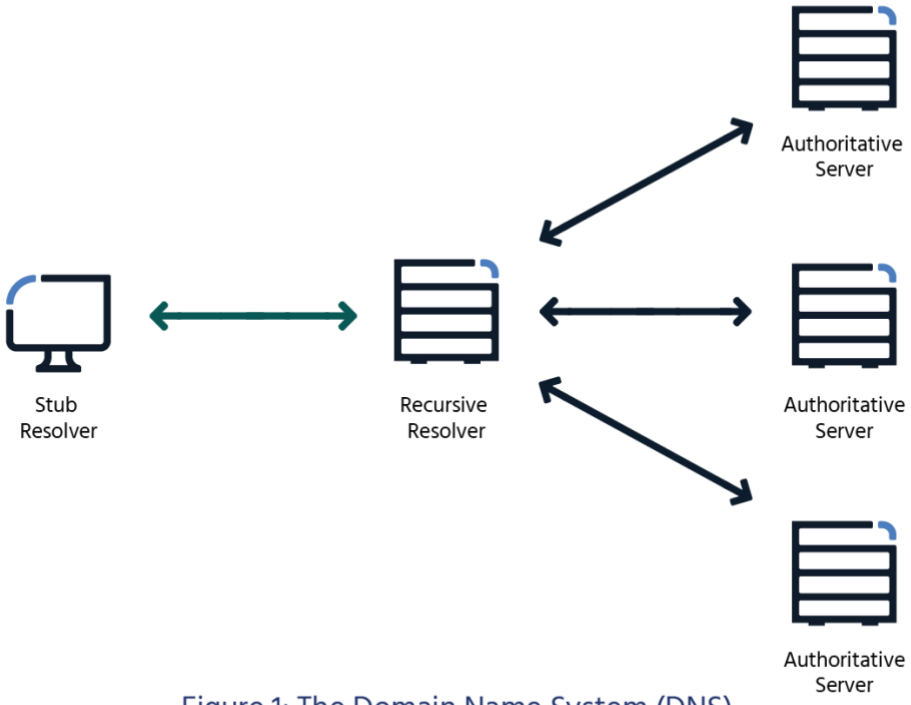

Figure 1: The Domain Name System (DNS)

There are three different components of the DNS in this diagram:

1. Stub resolver: lightweight software on the user's device that forwards a domain name query from an application, like a web browser, to a recursive resolver
2. Recursive resolver: a server that receives the query from the stub resolver and finds the correct answer by in turn querying a set of authoritative name servers
3. Authoritative name servers: servers that hold the authoritative records for a zone (e.g., ".com" or "example.com"). They tell resolvers which IP addresses match the domain names in their section of the DNS hierarchy.

Typically, DNS resolution involves querying authoritative name servers successively, starting from the "root zone" of the DNS, and on through the DNS hierarchy until an authoritative name server can provide a response to the query. For example, the following diagram illustrates one possible scenario to resolve the domain name "www.example.com" to an IP address:

[4] Internet Society, Introduction to DNS Privacy (2018), https://www.internetsociety.org/resources/doc/2018/introduction-to-dns-privacy/.

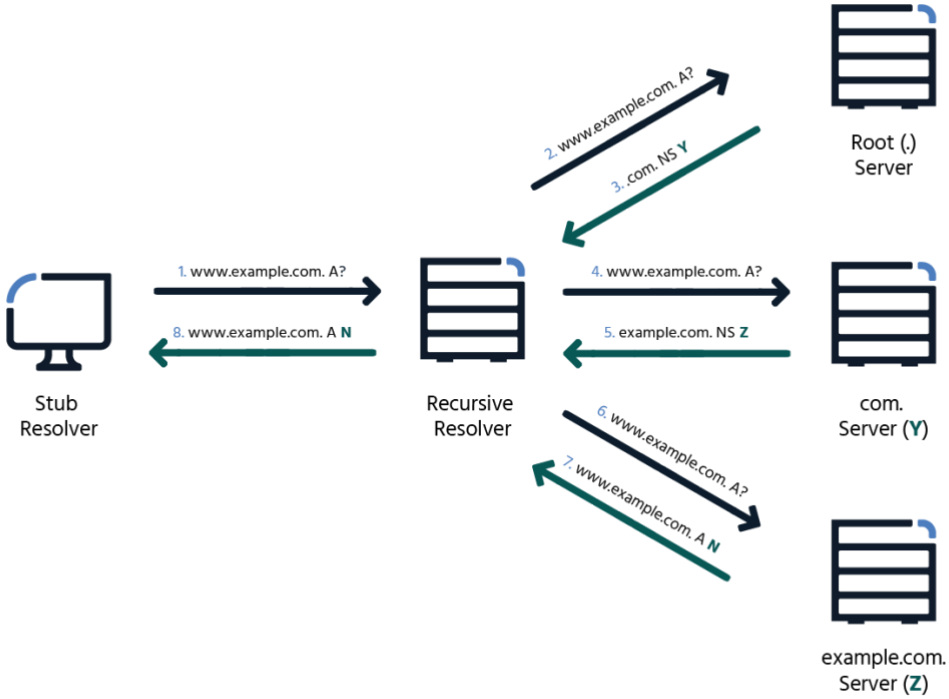

Figure 2: DNS Resolution Example

The first transaction involves the query (step 1) from the stub resolver to the recursive resolver, which will only respond (in step 8) once it has obtained the answer to the query as a result of multiple queries to authoritative DNS servers.

The recursive resolver will typically resend the query (step 2) to one of the authoritative DNS servers for the root zone (“.”), which will respond (step 3) with the list of authoritative name servers for the “.com” zone.

The recursive resolver will then resend the query (step 4) to one of the authoritative DNS servers for the “.com” zone, which will respond (step 5) with a list of authoritative DNS servers for the “example.com” zone.

Finally, the recursive resolver will send the query (step 6) to one of the authoritative DNS servers for the domain. That server replies (step 7) with the correct IP address for www.example.com, thereby completing the resolution process by providing the answer to the original query.

This simplified walk-through of DNS resolution highlights two characteristics of the DNS that are important to keep in mind:

- **DNS resolution is consistent, but not persistent:** While DNS answers are consistent at the time of the query, they are not necessarily permanent. The operator of the authoritative server can update the records so that a domain points to different IP addresses over time.

In fact, many domains are even configured to return multiple IP addresses in a single response, allowing users to be directed to servers more closely located to them. This flexibility of the DNS supports redundancy, traffic distribution, and fast recovery from outages.

- **DNS resolution occurs before any specific web page or file is requested:** The DNS only translates the domain name into an IP address. The details of what is being accessed, such as a specific page path[5], image, or video file, are handled later by the web server and browser (or some other application). As a result, the DNS operates only at the domain name level, not at the level of individual pages or other resources. This has significant implications for the effectiveness and proportionality of DNS-based interventions.

Private vs. Public Recursive DNS Resolvers

While it is technically possible to place the recursive resolver on the user's device to perform the full process of name resolution on its own, this is uncommon and largely impractical. Instead, most devices rely on a separate recursive resolver, usually operated by their Internet service provider (ISP) or another service. This delegation has important advantages[6]:

- First, it reduces complexity on the device itself. Rather than carrying out all the steps of contacting multiple authoritative servers, the device can simply delegate the query to the recursive resolver and receive the final answer.
- Second, it improves performance. Because recursive resolvers serve many users, they can *cache* (i.e., temporarily save) recent answers. If the same information is needed again, whether by the same user or another user of the same resolver, the cached response can be provided immediately. This avoids the delay of repeating the full lookup.
- Third, it can offer additional privacy protections. Instead of every device directly revealing its queries to multiple authoritative servers around the world, those queries are passed through the recursive resolver. This reduces the exposure of individual user activity across the global DNS infrastructure.

For these reasons, the provisioning of the recursive resolver has traditionally been supplied by the user's ISP or, in the case of an enterprise network, by the network administrator. In this model, known as a *private recursive resolver,* the recursive resolver is operated within the access network and is primarily available to that network's subscribers or members. Under this arrangement, the recursive

[5] A page path is the portion of a URL that comes after the domain name and identifies a specific resource on the web server. For example, in "https://example.com/articles/2025/DNS blocking.html", the domain name is "example.com", while "/articles/2025/DNS blocking.html" is the page path.

[6] Internet Society, *Introduction to DNS Privacy.*

resolver is typically configured automatically through network settings. As a result, most users are unaware of the specific resolver they are using since it is provided automatically as part of their Internet access service.

However, in recent years, there has also been a growth in the use of so-called *public recursive resolvers.* These are operated by third-party organizations, such as Google (8.8.8.8), Cloudflare (1.1.1.1), and Quad9 (9.9.9.9), that make their resolvers openly accessible to anyone on the Internet. Public resolvers are often promoted on the basis of improved speed, advanced security features, or enhanced privacy practices. They may be adopted by individuals who reconfigure their devices, or by entire networks that choose to outsource DNS resolution rather than operate their own resolvers.

The operational and structural differences between private and public recursive resolvers are significant in several ways. Notably, private resolvers are typically tied, both technically and contractually, to the access network in which they operate, making them part of the user's direct service relationship. Public resolvers, by contrast, are offered as independent services that operate globally and may be located in different jurisdictions from their users. Public resolvers typically have no process of signing up for the service and operate at no cost to the user, with few terms or conditions. Consequently, they are also unaware of the geographic, policy, or personal information related to any specific user.

Security and Resiliency of the DNS

The DNS was originally designed to provide a simple translation between human-readable names and numerical IP addresses, but it was also built with the capacity to scale to a global system. Over time, the Internet community has also introduced additional mechanisms to strengthen its resilience and protect it against failures and malicious tampering.

- Redundancy in the DNS: Redundancy has been a fundamental design feature of the DNS since the beginning. At the top of the hierarchy, queries start with the root zone, which is served not by a single operator but by 13 root name servers, each operated by independent organizations. In practice, these servers are not a single machine—they are distributed across hundreds of servers worldwide using a technique called anycast[7]. From a recursive resolver's perspective, this means a query to a root server will be answered by the closest available server. If one server or even an entire region becomes unreachable, the query can be answered by another server.

The same principle of redundancy is often applied throughout the DNS hierarchy. Authoritative name servers for individual domains are frequently deployed in multiple locations and networks. If one authoritative server fails, recursive resolvers can switch to another listed for that zone. By ensuring that

[7] 'Root Servers', accessed 2 September 2025, https://www.iana.org/domains/root/servers; 'DNS Anycast: Concepts and Use Cases', accessed 2 September 2025, https://www.catchpoint.com/dns-monitoring/dns-anycast.

no single machine or network outage can prevent a domain from being resolved, redundancy has been a key factor in the DNS's continuous operation at a global scale[8].

- **Caching helps improve efficiency and resiliency**: Recursive resolvers also improve performance and reduce strain on the system by caching the answers they receive. If users request the same domain name, the resolver can provide the cached answer immediately, without repeating the full resolution process. This makes DNS lookups faster for users, while also reducing the number of queries that need to traverse the hierarchy of authoritative name servers.

Caching also contributes to the DNS's resiliency during times of heavy demand or network disruption. If authoritative name servers become temporarily unreachable, cached records can continue to provide correct answers until their validity expires (this is controlled by the record's "time-to-live" (TTL) values that are set by the domain operator[9]).

- **DNSSEC provides authenticity and integrity**: While the DNS was originally designed to make the Internet usable and scalable, it also relied on trust. For example, the system assumes that the answer a resolver receives is accurate, which creates opportunities for abuse. Attackers can insert false information into the system or trick users into being redirected to fraudulent websites that appear to be legitimate. To strengthen the DNS against this kind of manipulation, the Internet community developed DNS Security Extensions (DNSSEC). DNSSEC enables the cryptographic signing of information in DNS records, allowing resolvers to verify that the answers they receive are authentic and have not been altered.[10] The goal is to give users confidence that when they enter a domain name in their browser, whether to reach a bank, an e-commerce site, or a government service, they are taken to the legitimate destination.

Because the DNS operates as a decentralized system, it is not possible to simply "switch on" DNSSEC for everyone at once. Instead, each part of the system, whether registries, registrars, domain operators, or recursive resolvers, must adopt it independently. As a result, deployment has been gradual and remains uneven, leaving some domains protected while others are not.

---

[8] Michael A. Patton et al., *Selection and Operation of Secondary DNS Servers*, Request for Comments RFC 2182 (Internet Engineering Task Force, 1997), https://doi.org/10.17487/RFC2182.

[9] 'How DNS Caching Works', accessed 25 August 2025, https://dnsprivacy.org.uk/docs/how-it-works/dns-caching.html.

[10] 'DNSSEC – What Is It and Why Is It Important? - ICANN', accessed 2 September 2025, https://www.icann.org/resources/pages/dnssec-what-is-it-why-important-2019-03-05-en; 'DNSSEC News, Guides and Tutorials', *Internet Society*, 11 June 2020, https://www.internetsociety.org/deploy360/dnssec/.

DNS Blocking

DNS blocking (also referred to as DNS-based filtering) alters the normal operation of the DNS to prevent users from accessing specific domain names. As described above, under standard conditions, when a user types a domain name (e.g., "example.com") into a web browser, their resolver returns the corresponding IP address so the browser can connect to the correct server.

However, with DNS blocking, the resolver is configured to check requested names against a block list before completing the lookup. If the queried name is on the list, the resolver will return a modified or false response instead of the actual IP address. This can take several forms[11]:

- **Redirect**: returning the IP address of a different server, often one that displays a notice indicating that access to the requested site is blocked
- **NXDOMAIN**: returning a response indicating that the domain name does not exist
- **Null or error response**: in some cases, the resolver simply returns an error or a non-functional address (e.g., 0.0.0.0, which is a non-existent destination)

Regardless of the method, the effect is the same: the user cannot reach the intended destination through DNS resolution. Redirect also has the potential to introduce security and privacy issues for users, for example, if a different server is not properly secured or uses tracking technologies to observe users who attempt to access the redirected site.

DNS blocking is often promoted as a tool for enforcing broader public policy goals, such as preventing copyright infringement, restricting access to online gambling, or controlling politically sensitive or unlawful content. However, as the following chapters will show, DNS blocking is a blunt and largely ineffective instrument for such purposes. Because it operates only at the level of entire domains, it cannot distinguish between lawful and unlawful material hosted under the same domain. This makes it prone to over-blocking, where legitimate services are swept up in the block, as well as under-blocking, since the same prohibited content can quickly reappear under a different domain[12]. The result is a measure that is easy to bypass for determined users, yet still disruptive for lawful services and ordinary users.

However, it's important to note that this does not mean that all uses of filtering out DNS responses are unwarranted. Operators sometimes deploy DNS filtering in ways that are limited, transparent, and designed to protect users. For example, resolvers may block access to domains known to distribute malware, prevent phishing sites from collecting credentials, or enforce internal policies within

[11] ICANN SSAC, *DNS Blocking Revisited*, SAC127 (2025), https://itp.cdn.icann.org/en/files/security-and-stability-advisory-committee-ssac-reports/sac127-dns-blocking-revisited-16-05-2025-en.pdf.

[12] ICANN SSAC, *DNS Blocking Revisited*; Vinton Cerf et al., 'Concerns over DNS Blocking', *Medium*, 24 June 2023, https://medium.com/@vgcerf/concerns-over-dns-blocking-988ef546a100.

organizations such as schools or enterprises. In these situations, the filtering is voluntary, narrowly scoped, and generally understood and specifically desired by the affected users as part of operating in a managed environment[13].

Policymakers sometimes point to this use of voluntary filtering as evidence that mandated DNS blocking is merely an extension of existing practices. However, this comparison is misleading. While both approaches involve modifying DNS responses, they differ fundamentally in purpose, implementation, and technical effect.

First, voluntary filtering is driven by the individual user, organization, or Internet Service Provider (ISP) that selects a blocklist that can be modified or disabled at any time. Users may, for instance, enable parental-control filters to block pornography, or security lists to protect against malware and phishing sites[14]. Their filtering does not affect any other users. In contrast, mandated DNS blocking is an expression of public policy, issued by government agencies or courts, whose decisions are binding on all operators within their jurisdiction. The authority to decide what may or may not be resolved thus shifts from the network edge to a centralized authority (the state), where decisions are imposed universally and without user input.

Secondly, in voluntary systems, the DNS filtering is typically implemented close to the user, e.g., on home routers, corporate firewalls, or ISP-level resolvers, which the user can replace or reconfigure. The filtering, therefore, operates within a confined and transparent relationship[15]. Mandated blocking, by contrast, requires system-level intervention where resolvers across an entire jurisdiction (and potentially beyond) must be reconfigured to comply with an order. Such measures remove the practical ability for users or networks to choose alternative resolvers, transforming a local management choice into a nation-level policy.

Finally, the two approaches also differ in their technical and operational consequences. Voluntary filtering is typically narrow, transparent, and locally administered. For example, if over-blocking occurs, it can be quickly detected and resolved, with limited impact on the wider Internet. In contrast, mandated blocking risks introducing significant collateral effects by imposing inconsistent name resolution across jurisdictions.

The distinction between these voluntary uses of DNS filtering and mandated DNS blocking is important. The former is focused on safeguarding users from clearly harmful activity, typically with consent and within defined and transparent boundaries. The latter seeks to enforce public policy objectives by altering the core naming function of the Internet, introducing risks to its reliability and universality. It is

[13] ICANN SSAC, *DNS Blocking Revisited.*

[14] Richard Barnes et al., *RFC 7754 Technical Considerations for Internet Service Blocking and Filtering* (2016), 77, https://datatracker.ietf.org/doc/rfc7754/.

[15] Internet Society, *Perspectives on Internet Content Blocking: An Overview* (2017), https://www.internetsociety.org/resources/doc/2017/perspectives-on-internet-content-blocking/.

this second category, of mandated DNS blocking, that raises the most serious concerns and that the remainder of this report examines in detail.

Summary

- **The Internet is a general-purpose infrastructure.** It simply moves data packets between computers, while applications such as the Web or email are implemented at the edges by the connected devices. This design allows the Internet to support an open-ended and constantly evolving range of uses.
- **Addressing and naming are essential functions.** IP addresses identify devices with numerical identifiers, while the DNS translates those numbers into human-readable names that people can use. Because these functions are foundational, interventions at either can have wide and often unintended consequences.
- **The DNS operates across borders by design.** Neither names nor addresses map neatly onto geography, and both are allocated on functional rather than jurisdictional lines. This global character is a source of strength.
- **The separation of naming and addressing is deliberate.** Names make the Internet easier for people, while addresses allow machines to communicate. This separation allows flexibility.
- **DNS resolution is distributed and hierarchical.** Users' devices pass queries to recursive resolvers, which in turn consult authoritative name servers arranged from the root zone downward. This design ensures resilience but also means the DNS depends on many actors across different jurisdictions.
- **Recursive resolvers are central to how queries are handled.** Most users rely on resolvers provided by their ISP (private resolvers) or on global services such as Google, Cloudflare, or Quad9 (public resolvers). Private resolvers are typically tied to local networks, while public resolvers operate globally, complicating jurisdictional enforcement of blocking orders.
- **DNS answers are consistent but not persistent.** A domain may change its IP address over time or be configured to return multiple addresses for redundancy and traffic management. This flexibility supports reliability, but also means DNS data cannot be fixed in place by a single intervention.
- **DNS resolution happens before content is served.** The DNS only translates the domain name into an IP address, while the actual content request (a page, image, or video) happens later, once a connection has been established between the user's computer and the web server.

- **The DNS was designed for usability, not public policy enforcement.** It solves the problem of making network addresses memorable and reliable for human use. Repurposing it as a public policy enforcement tool is incompatible with its intended function.

DNS Blocking: Unintended Consequences

Over-blocking

As described in the previous chapter, DNS resolution only translates a domain name to the IP address of a web server, not the full path of a resource. For example, when a user enters https://example.com/page1 in their web browser, the DNS only resolves the "example.com" part to an IP address. The rest of the path, i.e., the "/page1" part (or any specific image, video, or file), is handled once a connection has been established with the web server. This means DNS blocking can only apply to entire domain names, not to individual pages or files[16].

This limitation makes DNS blocking a blunt tool. If a single page under a domain is unlawful, blocking at the DNS level prevents access to all the other, lawful material hosted under the same domain. The impact can be particularly significant for shared platforms such as social networks, blogging services, or cloud services, where millions of distinct users might rely on the same domain name. A single block can therefore disrupt vast amounts of unrelated content.

An illustrative case occurred in 2011, when US authorities seized the domain name "mooo.com" as part of an enforcement initiative targeting child exploitation material. The specific target was a set of subdomains allegedly being used to host the illegal content. However, thousands of individuals and small organizations also relied on the domain to host their own websites under subdomains of "mooo.com". At the time of the seizure, the domain was supporting roughly 84,000 distinct subdomains, ranging from personal blogs to small businesses and community projects. When authorities redirected "mooo.com" to government-controlled name servers, all of those subdomains immediately broke. Visitors attempting to access perfectly legitimate sites under "mooo.com" could no longer reach their intended destinations. Instead, they were greeted with a standard US Department of Homeland Security (DHS) seizure banner claiming the domain had been taken due to child pornography. This meant that tens of thousands of completely innocent site operators were not only cut off from their users but also had their reputations damaged by being falsely associated with child exploitation.[17]

[16] ICANN SSAC, *SSAC Advisory on Impacts of Content Blocking via the DNS*, SAC056 (2012), https://itp.cdn.icann.org/en/files/security-and-stability-advisory-committee-ssac-reports/sac-056-en.pdf.
[17] Jerry Brito, '"Operation Protect Our Children" Accidentally Shutters 84,000 Sites', *Time*, 17 February 2011, https://techland.time.com/2011/02/17/operation-protect-our-children-accidentally-shutters-84000-sites/.

Fragmentation and Collateral Damage

As described above, the DNS was designed to provide universal naming for Internet resources, ensuring that a given domain name resolves consistently, regardless of where or by whom the query is made. This universality has been one of the Internet's greatest strengths, enabling seamless communication, interoperability, and innovation across borders[18].

When naming stops being universal, because some resolvers return altered or incomplete answers, the consistency of the user experience is lost. Anyone who has attempted to reach a familiar site from a network where DNS filtering is in place, only to find an unexpected page or an error, understands the usability cost of such fragmentation. The more DNS answers differ depending on where or how a lookup occurs, the greater the erosion in the Internet's economic and social value. For example, domain names are embedded in hyperlinks, email messages, shared on social media, and distributed across networks and jurisdictions. A URL in an email, for instance, is expected to lead the recipient to the same resource the sender intended, irrespective of ISP or country. This feature has been fundamental to global commerce, free expression, and the development of the Internet as a unified network of networks.

However, the universality of the DNS is not only a matter of convenience for users. From a technical perspective, many Internet applications depend on it. Online services rarely exist as a single system hosted on one server. Instead, what a user experiences as a unified website or application is typically a collection of components distributed across multiple servers. For example, the basic structure of a web page may be delivered from one server, while images or video are embedded from another, the login function may rely on an external authentication service, and additional features can be drawn from other third-party providers. These elements, each located on different servers, are assembled into a seamless user experience, and their integration typically relies on domain names to locate and connect the different parts[19].

This reliance on names rather than numbers is not just about the convenience of using easily remembered names. The DNS also allows service operators to change the underlying infrastructure without disrupting the applications that depend on it. The domain name remains stable, while the operator can update the IP address it points to in order to add servers, balance traffic, or other upgrades[20]. Consider a website that relies on a third party for authentication. Its code may call a URL, such as "example.com/api/auth/login", to access an application programming interface (API) that is located on the authentication provider's servers. If the authentication provider needs to scale or improve its service, it can add servers or redistribute traffic simply by updating the DNS records. The

[18] Internet Society, *The Internet Way of Networking (IWN): Defining the Critical Properties of the Internet*.

[19] 'Populating the Page: How Browsers Work - Performance | MDN', MDN Web Docs, 11 August 2025, https://developer.mozilla.org/en-US/docs/Web/Performance/Guides/How_browsers_work.

[20] 'What Is DNS-Based Load Balancing? | DNS Load Balancing', accessed 1 September 2025, https://www.cloudflare.com/learning/performance/what-is-dns-load-balancing/.

website owner does not need to change anything (the code can keep pointing to the same URL) or even be aware of the upgrade. This flexibility has been essential for the growth of online services, enabling them to scale globally and adapt infrastructure without requiring constant changes by users or developers.

These dependencies mean that the blocking of a domain can lead to failures across unrelated sites and applications. A notable case occurred in 2016 when Turkish authorities ordered ISPs to block GitHub in response to a leak of politically sensitive emails that had been uploaded on the site. Because developer tools such as Font Awesome and the Homebrew package manager relied on GitHub, the block triggered a nationwide outage that lasted for nearly 18 hours[21]. Similarly, in 2015, Chinese filtering of "connect.facebook.net", a domain used to load Facebook's "Like" button and other social plugins, disabled that functionality on millions of websites that had embedded the feature, even though the sites themselves were unrelated to Facebook[22].

From this view, domain names are not just typed into a browser but form a connective tissue for online applications. They are embedded in applications' code, which means that every call to these external servers, whether to authenticate a user, display a map, image, or stream a video, ultimately depends on resolving a domain name into the correct IP address. If a resolver provides an altered or inconsistent response, the primary service may fail altogether, even if the user never notices which part of the system failed[23].

Security Implications: DNSSEC and Incompatibility with Blocking

As previously described, the Internet community developed DNS Security Extensions (DNSSEC) to address vulnerabilities in the DNS. DNSSEC allows DNS records to be cryptographically signed so that recursive resolvers can verify that the answers they receive are authentic and unaltered. When properly deployed, DNSSEC prevents "machine-in-the-middle" attacks in which false DNS responses are injected to redirect users to malicious sites.

The strength of DNSSEC lies in its ability to ensure the integrity and accuracy of the information received from DNS. By design, only the legitimate operator of a domain's authoritative server can create the cryptographic signatures that validate its DNS records. A resolver that performs DNSSEC validation will therefore reject responses that do not carry a valid signature.

---

[21] Efe Kerem Sozeri, 'How Hacktivist Group RedHack Gamed Turkey's Censorship Regime', *The Daily Dot*, 12 October 2016, https://www.dailydot.com/debug/redhack-gamed-turkey-censorship/; Editorial, 'Dropbox, Google Drive and Microsoft OneDrive Cloud Services Blocked in Turkey Following Leaks', *Turkey Blocks*, 8 October 2016, https://turkeyblocks.org/2016/10/08/google-drive-dropbox-blocked-in-turkey/; 'Turkey Blocks Web Drives after Email Leak', Technology, *BBC News*, 10 October 2016, https://www.bbc.com/news/technology-37608553.

[22] *China Censors Facebook.Net, Blocks Sites With "Like" Buttons – Krebs on Security*, 28 April 2015, https://krebsonsecurity.com/2015/04/china-censors-facebook-net-blocks-sites-with-like-buttons/.

[23] ICANN SSAC, *DNS Blocking Revisited.*

This creates a direct conflict with DNS blocking since any attempt to redirect a query to something other than the requested domain violates the way DNSSEC is designed to work. If a resolver returns a forged IP address to comply with a blocking order, e.g., by sending the user to a page stating that the site is blocked, that response cannot be signed by the domain's legitimate operator and would be rejected by DNSSEC validation.

In fact, even blocking methods that avoid redirection undermine DNSSEC. Returning a false "NXDOMAIN" (indicating that a domain does not exist) for a signed domain also breaks the chain of trust. From the user's perspective, a DNSSEC failure triggered by a blocking order is indistinguishable from one caused by an attacker tampering with the system. Without a way to tell the difference, applications that rely on secure name resolution may fall back to less secure modes or fail altogether. As a result, this weakens the very protections DNSSEC was designed to provide[24].

In summary, DNSSEC and DNS blocking are fundamentally incompatible. Where DNSSEC seeks to guarantee authenticity through cryptographic validation, DNS blocking introduces deliberate inauthenticity. The result is a clash where widespread adoption of DNSSEC is undermined by DNS blocking orders.

Summary

- **DNS blocking operates only at the domain name level.** Because the DNS translates only the domain name (e.g., *example.com*) into an IP address, it cannot block individual pages, files, or resources. This limitation makes DNS blocking a blunt tool at risk of over-blocking, preventing access to large amounts of lawful material hosted alongside unlawful content.
- **Fragmentation undermines universality.** The DNS was designed to provide consistent answers worldwide. When resolvers return altered results, users experience a form of fragmentation where the same domain may resolve differently depending on location or provider. This erodes trust, reduces usability, and diminishes the Internet's economic and social value.
- **Online services depend on DNS integration.** Websites and applications are rarely single systems. They are built from many components, such as third-party authentication, content delivery, or embedded tools, that rely on domain names to function together. Altering DNS responses can therefore cause cascading failures in services far removed from the primary target.

DNS blocking conflicts with DNSSEC. DNS Security Extensions (DNSSEC) provide cryptographic assurance that DNS responses are authentic and unaltered. Blocking measures that redirect queries or

[24] Steve Crocker et al., *Security and Other Technical Concerns Raised by the DNS Filtering Requirements in the PROTECT IP Bill* (Authors (Affiliations provided for identification only), 2011), https://www.circleid.com/pdf/PROTECT-IP-Technical-Whitepaper-Final.pdf.

return false answers break this chain of trust by deliberately introducing inauthenticity, making the two approaches fundamentally incompatible.

DNS Blocking: Efficacy of Measures

Circumvention Techniques

DNS blocking does not remove content from the Internet. It only prevents a particular resolver from resolving the IP addresses for the server where the content is hosted. The content remains accessible via other resolvers or through direct connection if the IP address is known. Determined users have many circumvention techniques to avoid such blocks that are easily accessible (and even automated), which reduces the long-term effectiveness of DNS blocking.

First, the user can simply switch resolvers. Since DNS blocking is typically implemented on specific resolvers, such as those operated by an Internet Service Provider (ISP), a user can often bypass the block by directing queries to a different resolver. Most users start with a resolver chosen by their ISP that has been automatically configured through network settings during setup. However, users can override this setting and instead point to another resolver (e.g., a public resolver), or even run a resolver on their own device. For example, during the 2014 Twitter ban in Turkey, the use of Google's public resolver to circumvent local ISP blocks became so widespread that the resolver's IP address (8.8.8.8.) was spray-painted on walls as information for how to bypass the censorship[25]. Furthermore, and in addition to well-known public resolvers, researchers have identified more than a million lesser-known recursive resolvers that are widely distributed across geographic regions and network topology, openly reachable, and that could be used in a similar way[26].

Secondly, the user can use a virtual private network (VPN) or the Tor network[27]. These tools encrypt and redirect all Internet traffic (including DNS queries) through servers located in another network (potentially outside the affected jurisdiction). From the perspective of the ISP's resolver, no queries are being made since DNS resolution is done via the external servers. VPNs and Tor are increasingly mainstream tools, making circumvention easier for non-technical users.

Finally, there is also the option to bypass the DNS altogether by connecting directly to the server's IP address. A related method is to bypass the DNS altogether by using the computer's "hosts file". This is a simple text file that allows a user to manually specify which IP address a particular domain name should point to. By adding an entry in the hosts file, effectively hardwiring an IP address to a specific domain name, the computer will connect directly to that IP address without needing to contact the DNS

[25] Aaron Souppouris, 'Turkish Citizens Use Google to Fight Twitter Ban', The Verge, 21 March 2014, https://www.theverge.com/2014/3/21/5532522/turkey-twitter-ban-google-dns-workaround.

[26] Meng Luo et al., 'Who Are Querying for Me? Measuring the Dependency and Centralization in Recursive Resolution', *Computers & Security* 156 (September 2025): 104501, https://doi.org/10.1016/j.cose.2025.104501.

[27] ICANN SSAC, *DNS Blocking Revisited.*

resolver[28]. For non-technical users, this can certainly be challenging, but tools exist that automate this type of change. For example, the "MafiaaFire" browser plugin, developed in response to the domain seizures in the US, automatically redirected users from blocked domains to replacement domains without requiring manual changes. In effect, it kept the browsing experience the same while bypassing the DNS[29].

These examples show that DNS blocking is ineffective for people with basic technical skills, with widely available tools that can be used to evade blocks. As a result, DNS blocking functions less as a comprehensive enforcement tool and more as a barrier for those unwilling or unable to seek alternatives.

Blocked Content Staying Up

DNS blocking only prevents a resolver from translating a blocked domain name into its corresponding IP address. It does not remove the underlying content from the Internet. In practice, the material usually remains available and can be accessed again once the user finds an alternative resolution path or if it becomes associated with a new domain name.

This dynamic is especially visible in fast-moving contexts such as phishing, extremist propaganda, or copyright infringement. In such cases, targeted operators can register new domains and point them to the same servers almost immediately. As a result, a block placed on one domain may only temporarily affect access until the site resurfaces under another[30].

More advanced techniques amplify this challenge. Domain generation algorithms (DGAs) can automatically create thousands of possible domain names each day, of which only a handful need to be registered. While the operator of the domain knows which names will be valid at a given time, for those attempting to block the content, the task becomes challenging as they may need to block thousands of potential domains.[31]

Additionally, mirroring, where volunteers or sympathetic groups replicate the same material on different servers under new domains, further undermines the effectiveness of DNS blocking. Once shared across social networks and messaging platforms, these new URLs quickly spread, ensuring that the content remains accessible to users regardless of the initial block[32].

---

[28] 'Understanding Hosts Files - What You Need to Know | Lenovo IE', accessed 3 September 2025, https://www.lenovo.com/ie/en/glossary/hosts-file/.

[29] Ian Paul, 'Mozilla Refuses to Help Censor the Internet', *PCWorld*, 6 May 2011, https://www.pcworld.com/article/491182/mozilla_refuses_to_help_censor_the_internet.html.

[30] Barnes et al., *RFC 7754 Technical Considerations for Internet Service Blocking and Filtering*.

[31] 'Dynamic Resolution: Domain Generation Algorithms, Sub-Technique T1568.002 - Enterprise | MITRE ATT&CK®', accessed 2 September 2025, https://attack.mitre.org/techniques/T1568/002/.

[32] 'Collateral Freedom | RSF', 28 April 2022, https://rsf.org/en/collateral-freedom.

Taken together, these strategies mean that the effectiveness of DNS blocking is short-lived. Initial disruptions may be measurable in the first hours or days, but they tend to erode quickly as new domains proliferate and mirrors sustain availability. Maintaining complete and up-to-date block lists becomes increasingly resource-intensive, and attempts to broaden their scope risk sweeping up lawful services and creating widespread collateral damage.

Summary

- **DNS blocking does not remove content from the Internet.** It only prevents a specific resolver from answering queries for a given domain. The underlying material remains online and can still be reached through other resolvers, direct IP access, or alternative techniques.
- **Circumvention is easy and widespread.** Users can bypass DNS blocks by switching to a different resolver, using VPNs, or Tor to route traffic outside the blocking jurisdiction. Those determined to access content, with even modest technical knowledge, can quickly evade DNS-based restrictions. This makes DNS blocking a weak enforcement tool that functions more as a barrier to the unaware than a meaningful deterrent.
- **Blocked material typically stays up.** Operators of targeted sites can reappear under new domain names within hours, leaving block lists one step behind. Automated techniques and mirroring further amplify this challenge.

DNS Blocking: Issues of Geography and Jurisdiction

The DNS and Geography

The Internet's architecture does not align with geographical or jurisdictional boundaries. While data transmission inevitably occurs over physical infrastructure located in specific places, the addressing and naming systems that enable Internet communication operate independently of geographical borders. For example, IP addresses are allocated by the five Regional Internet Registries based on an organization's operational needs, such as the requirements of running a network and serving customers, rather than its nationality. As a result, an address block assigned to an organization in one jurisdiction may be used to serve users in many others[33].

A similar logic applies to the DNS. Responsibility for each "zone" (e.g., ".com" or "example.com") rests with its authoritative operator, which may be located in a jurisdiction entirely different from the domain's registrant or the users who rely on it. For example, a domain might be registered through a

[33] *Resource Ranges by RIR | APNIC*, n.d., accessed 2 September 2025, https://www.apnic.net/manage-ip/manage-resources/address-status/by-rir/.

registrar in one country, its authoritative name servers operated from another, and the actual web servers hosting the content are in yet another location.

Recursive resolvers add even further complexity. For example, a public resolver located in Country A can serve queries from millions of users in Countries B, C, and D. Policies for blocking content applied to that resolver may thus affect users far beyond the intended jurisdiction unless the operator can reliably distinguish user location, which is complicated by the limitations of IP-based geolocation (further discussed in the next section).

For DNS blocking, this illuminates two important considerations. First, the location of a resolver does not necessarily correspond to the location of its users, making jurisdiction-specific enforcement challenging. Second, the role of the DNS is only to convey the IP address of the server hosting the content, and that server may be located outside the jurisdiction imposing the block. As cases like *The Pirate Bay* illustrate, even repeated efforts to block access through DNS resolvers or seize domain names have not made the underlying content disappear. Instead, the site has reappeared under new names or on servers in other jurisdictions, demonstrating that resolver-based blocking or domain seizures in one country cannot ensure global removal[34].

In sum, the DNS and IP addressing are not organized around national borders. Geographical blocking measures are difficult to enforce, often require cooperation across jurisdictions, and carry a high risk of unintended cross-border effects.

DNS Blocking and Public Recursive Resolvers

Public recursive resolvers, such as Google (8.8.8.8), Cloudflare (1.1.1.1), and Quad9 (9.9.9.9), offer their service to users across many jurisdictions. Unlike private resolvers operated by ISPs, which are typically tied to a national or regional user base consisting of users who pay for service and are specifically known to the service provider, public resolvers are designed to be global. Thus, a blocking order imposed on such a service risks being applied indiscriminately to all of its users, regardless of their location. This means that, unless the operator introduces mechanisms for location-based filtering, the block is effectively extraterritorial.

From a technical perspective, operators can implement filters that approximate a user's geographical location based on their IP address. However, this approach is far from precise[35]. This type of location-based filtering is imperfect and can misclassify users, especially in regions where IP allocations are fluid,

[34] 'The Pirate Bay Returns to Its .Org Domain', IT Pro, 25 May 2016, https://www.itpro.com/intellectual-property/26615/the-pirate-bay-returns-to-its-org-domain.

[35] Geoff Huston, 'Where Are You? A Look at GeoIP', Tech Matters, *APNIC Blog*, 16 December 2025, https://blog.apnic.net/2025/12/16/where-are-you-a-look-at-geoip/; 'Geolocation Accuracy', MaxMind, 12 March 2025, https://support.maxmind.com/hc/en-us/articles/4407630607131-Geolocation-Accuracy; Jon Worley, 'IP Geolocation: The Good, The Bad, & The Frustrating', 11 June 2018, https://www.arin.net/vault/blog/2018/06/11/ip-geolocation-the-good-the-bad-the-frustrating/.

shared across borders, or reassigned dynamically. This can result in lawful users outside the intended jurisdiction being denied access, while targeted users may still pass through.

Furthermore, adding location-based filtering also risks conflicting with the privacy and neutrality commitments of many public resolver operators. Services such as Cloudflare's or Quad9's public resolvers explicitly market themselves as global, consistent, and privacy-preserving alternatives to ISP-operated resolvers[36]. This includes policies such as not logging, profiling, or altering users' DNS queries depending on who or where the user is. In contrast, location-based filtering would require resolvers to collect or infer this type of information for every query, undermining both neutrality and privacy. Imposing blocking obligations on public resolvers is therefore not only a technical adjustment but also requires operators to compromise on important principles that underpin their service.

Recent cases illustrate the reality of these challenges. In 2022, the Court of Milan ordered Cloudflare to block access to three torrent websites through its public recursive resolver. Cloudflare argued that its resolver is a global service and does not distinguish between queries from Italian users and those from millions of users elsewhere. Therefore, enforcing the order would apply Italian law worldwide. Despite these concerns, the court upheld the injunction, effectively turning a national block into a global restriction.

Courts in Germany have taken a different approach. In *Universal v. Cloudflare*, the Higher Regional Court of Cologne rejected a request to impose blocking on Cloudflare's public resolver, explaining that a resolver only translates names into IP addresses and should not be used as an enforcement point.[37] [38] This reasoning is also in line with a case against Quad9, where the Higher Regional Court in Dresden overturned an earlier order requiring the Swiss-based nonprofit operator to block piracy domains globally, recognizing that recursive resolvers are neutral infrastructure and should not be responsible for deciding what content is accessible[39].

Enforcement Challenges and Conflicts of Law

Because the DNS is a global system, blocking measures often conflict with the diverse array of national legal frameworks. What may be prohibited in one jurisdiction may be legal in another, and a blocking order issued domestically cannot easily be reconciled with such differences. This means that when the resolver in question is operated by a foreign organization, such as a public resolver operator, the order may be unenforceable or even unlawful under the laws of the jurisdiction in which the operator is based.

---

[36] 'Quad9 | A Public and Free DNS Service for a Better Security and Privacy', Quad9, accessed 2 September 2025, https://quad9.net/; '1.1.1.1 Public DNS Resolver', Cloudflare Docs, 13 August 2024, https://developers.cloudflare.com/1.1.1.1/privacy/public-dns-resolver/.

[37] Notably, the Regional High Court of Cologne also argued that DNS services are protected by the EU's Digital Services Act (DSA).

[38] 'Latest Copyright Decision in Germany Rejects Blocking through Global DNS Resolvers', *The Cloudflare Blog*, 5 December 2023, https://blog.cloudflare.com/latest-copyright-decision-in-germany-rejects-blocking-through-global-dns-resolvers/.

[39] 'Quad9 Turns the Sony Case Around in Dresden', *Quad9*, 6 December 2023, https://quad9.net/.

This problem is particularly acute for public resolvers, which serve users globally and across borders. When a national regulator or court directs a public resolver operator to implement blocking, they are effectively asking that operator to apply domestic laws to users abroad. The operator could then face a legal dilemma: comply with the order and risk violating the rights of users or laws of another jurisdiction, or refuse and face penalties in the jurisdiction imposing the block. Furthermore, if one jurisdiction insists that a global operator enforce its laws abroad, others may follow suit, leading to a patchwork of overlapping and inconsistent blocking requirements. The result would be a highly fragmented DNS, in which the accessibility of domain names depends on which jurisdiction's rules dominate the resolver being used.

Furthermore, blocking obligations can also set precedents that other jurisdictions imitate or expand. Once mandated DNS blocking becomes a normalized policy tool for specific content categories, such as copyright infringement, there are no technical barriers for adding new categories. For example, policymakers might point to existing measures to justify extending blocking practices to new domains, including political criticism, religious commentary, or other content that is lawful or protected in many jurisdictions. In this way, a narrowly targeted blocking mechanism can gradually evolve into a broad tool for restricting access to information, placing operators in increasingly untenable legal positions[40].

Jurisdictional conflicts are a recurring challenge whenever national authorities seek to enforce domestic content rules vis-à-vis globally distributed Internet services. When the targeted service is legally established elsewhere, the enforcing state's authority to require infrastructure operators to act is often uncertain. This tension of authority can emerge in any governance domain, but is especially visible where multiple legal orders coexist or place limits on one another's territorial reach.

For example, in the EU, online services are regulated based on the principle of country of origin, meaning that the regulatory authority resides with the member state in which the provider is established, and restrictions by other states are only permitted under narrow conditions. A recent case in Germany illustrates the tension that can emerge, where German authorities were initially allowed to issue restrictions directly to a Cyprus-based content provider. However, when authorities sought to force DNS resolvers to block the associated domain, such measures were deemed incompatible with the EU's internal market. In other words, while the order against the content provider remained valid, the effort to enforce it through DNS blocking did not.[41]

Importantly, the case illustrates a broader point where, even if a state can lawfully act against a content provider, that does not mean it can compel an infrastructure operator (like a DNS resolver), to enforce those measures. Where national orders conflict with foreign or higher-level legal rules, it is not always clear if blocking measures are permitted at all. For resolver operators serving users in multiple

[40] Quad9 Foundation, *The Public Risk of Governments Controlling DNS Providers* (2025), https://www.quad9.net.

[41] 'Verwaltungsgericht Düsseldorf: Access-Provider Müssen Pornografische Internetangebote Vorerst Nicht Sperren', accessed 1 December 2025, https://www.vg-duesseldorf.nrw.de/behoerde/presse/pressemitteilungen/2025/25_11_19/index.php.

countries, this creates legal uncertainty, as they must navigate overlapping and sometimes contradictory obligations.

Alternatively, operators may simply choose to withdraw from markets rather than implement complex compliance systems that could expose them to liability in other jurisdictions. For example, Cisco's OpenDNS service chose to end operations in France (and later in Belgium) after being ordered to block sites that were allegedly housing content that was infringing on copyright[42]. The compliance challenge is especially acute for smaller operators, who lack the resources to manage contradictory orders from multiple jurisdictions.

The result of these legal disputes is a growing body of decisions that point in different directions, with no common framework for reconciling conflicts. Absent international coordination, attempts to enforce DNS blocking will either prove ineffective or create harmful side effects, introducing new risks for the stability and universality of the DNS itself.

Intermediary Liability Protections and Enforcement Pressures

Mandated DNS blocking raises important questions about how far governments may require infrastructure operators to intervene against online content. As described above, DNS resolvers do not create, host, or even see the content they help users reach. They merely translate domain names into IP addresses so that users can reach the server where the content is hosted. Many legal systems recognize this distinction through intermediary liability protections, which exempt infrastructure services from responsibility for user-generated content. Such protections reflect a widely shared principle that liability should lie with those who create or control the unlawful content, not with the intermediaries that facilitate the communication. Legislation such as Brazil's Marco Civil da Internet, the United States' Section 230, the EU's Digital Service Act (DSA), and comparable protections in other countries exemplify this approach.[43]

However, despite such safeguards, DNS resolvers sometimes face enforcement pressure when authorities rely on administrative, consumer protection, copyright, online safety, or risk-prevention powers to mandate blocking. This can place operators in untenable positions to police content they cannot access, and assume responsibilities more appropriately directed at hosting or content providers.

Recent case law in Germany provides a helpful illustration of how courts can address this tension through proportionality and a graduated approach to enforcement. In a ruling by the Administrative Court of Frankfurt am Main[44], an order requiring a major Internet Service Provider (ISP) to impose DNS

[42] 'DNS Piracy Blocking Orders: Google, Cloudflare, and OpenDNS Respond Differently * TorrentFreak', accessed 2 September 2025, https://torrentfreak.com/dns-piracy-blocking-orders-google-cloudflare-and-opendns-respond-differently-250511/.

[43] *A Policy Framework for Internet Intermediaries and Content* (Internet Society, 2024), https://www.internetsociety.org/resources/doc/2024/a-policy-framework-for-internet-intermediaries-and-content/.

[44] Judgment of October 23, 2024 – 7 K 800/22.F (MMR 2025, 849) (Administrative Court of Frankfurt am Main), accessed 14 January 2026, https://www.rv.hessenrecht.hessen.de/bshe/document/LARE240001433.

blocking against an unlicensed financial-service website was ruled unlawful and disproportionate since the regulator had not first attempted more direct and effective measures. The judgment made clear that DNS blocking should be a measure of last resort, and only after authorities have exhausted actions against the content provider and, secondly, the hosting provider (as they would be closer to the legal violation) — including through international legal assistance where necessary.

As legal frameworks evolve, these examples underscore a broader principle: Internet intermediaries, including DNS resolvers, should not be used for enforcement against content they cannot control, nor should they be the first point of intervention when more direct avenues exist. Safeguarding this principle is essential to prevent overreach, avoid fragmentation, and preserve the resilience and global reach of the Internet.

Summary

- **The DNS is not organized around national borders.** Addressing and naming are allocated on operational rather than territorial lines, meaning that a resolver or server in one country may serve millions of users elsewhere. This makes jurisdiction-specific enforcement difficult and often ineffective.
- **Public recursive resolvers are designed to be global.** Services like Google, Cloudflare, and Quad9 handle queries from users worldwide. A blocking order against one of them risks being applied indiscriminately across all users, effectively turning a local policy into a global rule.
- **Geolocation filtering is imprecise and risks conflicts.** Operators can try to infer user location via IP-based tools, but these techniques are imprecise, and errors are common. Beyond the technical flaws, such filtering also typically conflicts with the neutrality and privacy commitments that many public resolver operators make to their users.
- **Courts have diverged sharply in their treatment of resolvers.** Some courts have upheld blocking orders against public resolvers, while others have rejected them, highlighting a lack of a coherent global framework.
- **Conflicts of law create fragmentation risks.** A blocking order that is valid in one jurisdiction may be unlawful in another, placing operators in untenable positions that threaten the universality and stability of the DNS.
- **DNS resolvers are Internet intermediaries.** Many regulatory frameworks recognize that they do not create, host, or control content, which makes them ill-suited as points of intervention compared to more direct approaches.

DNS Blocking: Implementation Costs

Infrastructure and Monitoring Costs

Implementing DNS-based blocking, particularly when applied selectively by jurisdiction, requires not only technical changes to resolver infrastructure but also ongoing operational investment. A simple global block list is relatively easy to configure, but targeted blocking based on user location or legal jurisdiction significantly increases complexity and cost. The more fragmented the rules become, the greater the operational burden on DNS operators.

To apply blocking rules selectively, this typically requires:

- **Geolocation databases.** Before deciding whether to block a domain, the provider first needs to know where each user is located. Geolocation databases provide this information by matching IP addresses to approximate physical locations. However, more accurate databases are typically commercial products, and depending on query volume and the level of detail required, they can cost thousands of dollars per month.[45]
- **Specialized software.** Once the user location is known, providers need to add code that must run for each DNS request to check them against location-specific rules. This is a significant change to DNS resolvers, which are normally designed for speed and simplicity (just answer queries as fast as possible).[46]
- **Additional hardware.** Normally, a resolver performs one task: look up the answer to a query. With location-based blocking, it must now perform two or three tasks: look up the answer, determine the user's location, and check whether that answer is permitted in that location. Each additional check requires CPU processing time and memory. In practice, this can significantly increase the workload on DNS servers, thereby increasing costs for hardware, electricity, and maintenance.

While the incremental cost of supporting an additional blocking rule may be modest if there is no distinction given to jurisdictional application, the expense can multiply as more jurisdictions impose different requirements.

[45] 'How Much Is an IP? Understanding IP Data Pricing in 2025', accessed 22 January 2026, https://ipinfo.io/blog/ip-geolocation-data-pricing.
[46] David Abecassis et al., *The Economic Cost of Network Blocking* (Analysys Mason, 2025), https://www.analysysmason.com/consulting/reports/network-blocking-economic-impact-jul25/.

Personnel and Operational Overhead

Implementing DNS blocking at scale requires staff to monitor the technical implementation, handle user complaints, and ensure ongoing compliance with legal orders. This includes maintaining audit logs of actions and responding to inquiries from regulators and affected domain owners. As blocking obligations expand, staff must also track updates to block lists, verify new entries, and assess whether the requests are legally valid. Taken together, such tasks come to require extensive resources and ongoing operational costs for operators.

Furthermore, to the extent mandated DNS blocking becomes an ongoing requirement in a given jurisdiction, the scope and frequency of requests also tend to expand. Targeted measures that may have begun narrowly can quickly grow into bulk submissions covering a large number of domains, and blocklists may even be updated on a near-weekly basis. With such volume and pace, it becomes extremely challenging for operators to verify the legal basis of new orders, including their accuracy and proportionality[47]. And if the processes for corrections to remedy mistakes or over-blocking are unclear (or even absent), this further amplifies the challenge for operators (and affected domain owners).

Where multiple jurisdictions are involved, staff must also track changes in relevant laws and court rulings, coordinate with legal teams, and, in some cases, appear in cross-border compliance proceedings. Providers that are legally or contractually required to demonstrate compliance would have to validate that demands for blocking are, in fact, valid, which implies additional costs for each site to be researched and claims confirmed before inclusion in a blocklist. Costs for these validation and objection processes can be extensive, unless the provider chooses to accept automated demands for blocking without validation, thereby risking enforcement of invalid, mistaken, or abusive demands. This burden falls particularly heavily on smaller operators that may lack dedicated compliance departments. Large, well-resourced operators may be able to absorb the compliance costs, but smaller and non-profit operators risk being excluded from the market if they cannot afford the staffing and compliance infrastructure demanded by jurisdiction-specific blocking orders. The result may be market consolidation, where only the largest players can operate at scale, reducing competition and diversity in DNS service provisioning[48].

Furthermore, when professionally operated (yet increasingly constrained) resolvers lose users' trust, end-users may migrate to services that are compromised in a different way. This can fuel the rise of "below-the-radar" resolvers that have no official standing in any jurisdiction and operate with no transparency or for purposes that conflict with user safety and data integrity. An example of similar

[47] Steven Hawley, *Italy: Piracy Shield Cited for 'Indiscriminate' Blocking, Evasion by Pirates Using IPv6 - Piracy Monitor*, 1 October 2025, https://piracymonitor.org/italy-piracy-shield-found-committing-indiscriminate-blocking-losing-to-evolving-evasion-techniques/.
[48] Abecassis et al., *The Economic Cost of Network Blocking*.

patterns has been observed in research on free VPN services that harvest data and expose users to the risk of having their credentials stolen[49].

Performance Degradation and Reputation Risks for Operators

Integrating geolocation and blocklist checks into DNS resolution increases latency. While delays of only a few milliseconds may sound trivial, at Internet scale, they affect user experience and can drive users toward alternative resolvers, undermining both compliance and business objectives.

Services that apply location-based analysis may become measurably slower, which will result in lower user adoption. Sometimes this impact on user adoption is the result of testing and word of mouth. However, systems configured with multiple resolvers will typically switch to alternatives when the primary is slow or unresponsive, and some automated tools actively shift traffic based on performance[50]. As a result, adding geographic-based analysis can reduce traffic, simply because software configured with multiple resolvers will choose the fastest option.

Beyond performance, operators may also face reputational risks. DNS resolvers are, in general, expected to return accurate answers, but for many public resolver operators, this is also a core part of their mission and business proposition. This is why users choose their service over alternatives. When blocking orders force a resolver to provide answers that are incomplete, misleading, or otherwise inconsistent with the global DNS system, this can cause significant harm to the operator's reputation.

Summary

- **Mandated DNS blocking requires substantial investments in infrastructure and monitoring.** Location-based blocking often demands new resolver logic, geolocation systems, monitoring tools, and sustained operational capacity. As legal requirements fragment across jurisdictions, the burden on operators increases accordingly.
- **Jurisdiction-based blocking is complex and resource-intensive.** While blocking a domain globally is technically straightforward, implementing differentiated blocking at scale is resource-intensive. Costs include ongoing staffing, verification, error management, and compliance functions that many smaller operators cannot sustain, risking reduced diversity in the DNS ecosystem.

[49] Muhammad Ikram et al., 'An Analysis of the Privacy and Security Risks of Android VPN Permission-Enabled Apps', *Proceedings of the 2016 Internet Measurement Conference* (New York, NY, USA), IMC '16, 2016, 349–64, https://doi.org/10.1145/2987443.2987471.

[50] Alicia Sykes, 'DNS Query Performance', Networking Toolbox, accessed 26 January 2026, https://networkingtoolbox.net/diagnostics/dns/query-performance.

- **Location-based blocking impacts performance.** Integrating location checks into DNS resolution adds latency. Even small delays, when multiplied across billions of queries, degrade user experience and can drive users toward alternative resolvers.

Degraded performance and reliability create reputational risks. If a resolver becomes slower or forced to return manipulated answers, users and network operators may perceive it as less trustworthy. This harms reputation, reduces adoption, and can incentivize users to switch operators.

Conclusion

DNS provides a critical function in making the Internet usable and coherent. It ensures users can interact with a wide range of online services while allowing those services to grow and adapt. These qualities, of universality, consistency, and usability, have been central to the Internet's success. They are also the reasons why using the DNS for policy enforcement creates significant problems.

DNS blocking may seem straightforward: if the name does not resolve, the content cannot be reached. However, as described in this report, such a simplified understanding is deceptive. Blocking at the domain level cannot distinguish between lawful and unlawful content, which means that legitimate services are often caught in the process. Moreover, DNS blocking fragments the universality of the naming system (so the same name may no longer yield the same result everywhere) while undermining security mechanisms such as DNSSEC, which rely on authentic, untampered responses. The approach is also ineffective because it can be easily circumvented via alternative resolvers, VPNs, or by the content appearing under a new domain name.

DNS blocking is not simply a technical inconvenience but creates wider risks for users, businesses, and the Internet as a whole. Over-blocking can interfere with lawful expression and commerce. Jurisdictional conflicts can interfere with operations across borders, imposing one country's rules on users everywhere. Compliance obligations can also be costly, particularly for smaller operators, pushing the market toward consolidation around a handful of large providers.

In practice, DNS blocking is at best a temporary barrier for casual users, and at worst a measure that erodes trust in the DNS itself. The more it is used as an enforcement tool, the more it threatens the universality and resilience of one of the Internet's critical functions.

It is important for policy discussions to recognize these limits. If online harms need to be addressed, interventions should focus on the content itself, the actors responsible, and on approaches grounded in due process and international cooperation. For courts applying narrow tailoring, proportionality tests, due process, and principles of comity, the technical realities illustrated in this report show that DNS blocking is difficult to reconcile with such standards.

The DNS was designed to make the Internet usable, not to act as a tool of censorship or control. Preserving its universality, reliability, and security is essential if the Internet is to remain an open, resilient, and global network.

References

A Policy Framework for Internet Intermediaries and Content. Internet Society, 2024. https://www.internetsociety.org/resources/doc/2024/a-policy-framework-for-internet-intermediaries-and-content/.

Abecassis, David, Andrew Daly, and Dalya Glickman. *The Economic Cost of Network Blocking*. Analysys Mason, 2025. https://www.analysysmason.com/consulting/reports/network-blocking-economic-impact-jul25/.

Barnes, Richard, Alissa Cooper, Olaf Kolkman, Dave Thaler, and Erik Nordmark. *RFC 7754 Technical Considerations for Internet Service Blocking and Filtering*. 2016. https://datatracker.ietf.org/doc/rfc7754/.

BBC News. 'Turkey Blocks Web Drives after Email Leak'. Technology. 10 October 2016. https://www.bbc.com/news/technology-37608553.

Brito, Jerry. '"Operation Protect Our Children" Accidentally Shutters 84,000 Sites'. *Time*, 17 February 2011. https://techland.time.com/2011/02/17/operation-protect-our-children-accidentally-shutters-84000-sites/.

Cerf, Vinton, Stephen D. Crocker, Mirja Kühlewind, et al. 'Concerns over DNS Blocking'. *Medium*, 24 June 2023. https://medium.com/@vgcerf/concerns-over-dns-blocking-988ef546a100.

China Censors Facebook.Net, Blocks Sites With "Like" Buttons – Krebs on Security. 28 April 2015. https://krebsonsecurity.com/2015/04/china-censors-facebook-net-blocks-sites-with-like-buttons/.

Cloudflare Docs. '1.1.1.1 Public DNS Resolver'. 13 August 2024. https://developers.cloudflare.com/1.1.1.1/privacy/public-dns-resolver/.

'Collateral Freedom | RSF'. 28 April 2022. https://rsf.org/en/collateral-freedom.

Crocker, Steve, David Dagon, Dan Kaminsky, Danny McPherson, and Paul Vixie. *Security and Other Technical Concerns Raised by the DNS Filtering Requirements in the PROTECT IP Bill*. Authors (Affiliations provided for identification only), 2011. https://www.circleid.com/pdf/PROTECT-IP-Technical-Whitepaper-Final.pdf.

'DNS Anycast: Concepts and Use Cases'. Accessed 2 September 2025. https://www.catchpoint.com/dns-monitoring/dns-anycast.

'DNS Piracy Blocking Orders: Google, Cloudflare, and OpenDNS Respond Differently * TorrentFreak'. Accessed 2 September 2025. https://torrentfreak.com/dns-piracy-blocking-orders-google-cloudflare-and-opendns-respond-differently-250511/.

'DNSSEC – What Is It and Why Is It Important? - ICANN'. Accessed 2 September 2025. https://www.icann.org/resources/pages/dnssec-what-is-it-why-important-2019-03-05-en.

'DNSSEC News, Guides and Tutorials'. *Internet Society*, 11 June 2020. https://www.internetsociety.org/deploy360/dnssec/.

'Dynamic Resolution: Domain Generation Algorithms, Sub-Technique T1568.002 - Enterprise | MITRE ATT&CK®'. Accessed 2 September 2025. https://attack.mitre.org/techniques/T1568/002/.

Editorial. 'Dropbox, Google Drive and Microsoft OneDrive Cloud Services Blocked in Turkey Following Leaks'. *Turkey Blocks*, 8 October 2016. https://turkeyblocks.org/2016/10/08/google-drive-dropbox-blocked-in-turkey/.

Hawley, Steven. *Italy: Piracy Shield Cited for 'Indiscriminate' Blocking, Evasion by Pirates Using IPv6 - Piracy Monitor*. 1 October 2025. https://piracymonitor.org/italy-piracy-shield-found-committing-indiscriminate-blocking-losing-to-evolving-evasion-techniques/.

'How DNS Caching Works'. Accessed 25 August 2025. https://dnsprivacy.org.uk/docs/how-it-works/dns-caching.html.

'How Much Is an IP? Understanding IP Data Pricing in 2025'. Accessed 22 January 2026. https://ipinfo.io/blog/ip-geolocation-data-pricing.

Huston, Geoff. 'Where Are You? A Look at GeoIP'. Tech Matters. *APNIC Blog*, 16 December 2025. https://blog.apnic.net/2025/12/16/where-are-you-a-look-at-geoip/.

ICANN SSAC. *DNS Blocking Revisited*. SAC127. 2025. https://itp.cdn.icann.org/en/files/security-and-stability-advisory-committee-ssac-reports/sac127-dns-blocking-revisited-16-05-2025-en.pdf.

ICANN SSAC. *SSAC Advisory on Impacts of Content Blocking via the DNS*. SAC056. 2012. https://itp.cdn.icann.org/en/files/security-and-stability-advisory-committee-ssac-reports/sac-056-en.pdf.

Ikram, Muhammad, Narseo Vallina-Rodriguez, Suranga Seneviratne, Mohamed Ali Kaafar, and Vern Paxson. 'An Analysis of the Privacy and Security Risks of Android VPN Permission-Enabled Apps'. *Proceedings of the 2016 Internet Measurement Conference* (New York, NY, USA), IMC '16, 2016, 349–64. https://doi.org/10.1145/2987443.2987471.

Internet Society. *Introduction to DNS Privacy*. 2018. https://www.internetsociety.org/resources/doc/2018/introduction-to-dns-privacy/.

Internet Society. *Perspectives on Internet Content Blocking: An Overview*. 2017. https://www.internetsociety.org/resources/doc/2017/perspectives-on-internet-content-blocking/.

Internet Society. *The Internet Way of Networking (IWN): Defining the Critical Properties of the Internet*. 2019. https://www.internetsociety.org/iwn/.

IT Pro. 'The Pirate Bay Returns to Its .Org Domain'. 25 May 2016. https://www.itpro.com/intellectual-property/26615/the-pirate-bay-returns-to-its-org-domain.

Judgment of October 23, 2024 – 7 K 800/22.F (MMR 2025, 849) (Administrative Court of Frankfurt am Main). Accessed 14 January 2026. https://www.rv.hessenrecht.hessen.de/bshe/document/LARE240001433.

'Latest Copyright Decision in Germany Rejects Blocking through Global DNS Resolvers'. *The Cloudflare Blog*, 5 December 2023. https://blog.cloudflare.com/latest-copyright-decision-in-germany-rejects-blocking-through-global-dns-resolvers/.

Leiner, Barry M., Vinton G. Cerf, David D. Clark, et al. *A Brief History of the Internet*. Internet Society, 2009. https://www.internetsociety.org/internet/history-internet/brief-history-internet/.

Luo, Meng, Qiuyun Wang, Jianrong Zhang, et al. 'Who Are Querying for Me? Measuring the Dependency and Centralization in Recursive Resolution'. *Computers & Security* 156 (September 2025): 104501. https://doi.org/10.1016/j.cose.2025.104501.

MaxMind. 'Geolocation Accuracy'. 12 March 2025. https://support.maxmind.com/hc/en-us/articles/4407630607131-Geolocation-Accuracy.

MDN Web Docs. 'Populating the Page: How Browsers Work - Performance | MDN'. 11 August 2025. https://developer.mozilla.org/en-US/docs/Web/Performance/Guides/How_browsers_work.

Patton, Michael A., Scott O. Bradner, Robert Elz, and Randy Bush. *Selection and Operation of Secondary DNS Servers*. Request for Comments RFC 2182. Internet Engineering Task Force, 1997. https://doi.org/10.17487/RFC2182.

Paul, Ian. 'Mozilla Refuses to Help Censor the Internet'. *PCWorld*, 6 May 2011. https://www.pcworld.com/article/491182/mozilla_refuses_to_help_censor_the_internet.html.

Quad9. 'Quad9 | A Public and Free DNS Service for a Better Security and Privacy'. Accessed 2 September 2025. https://quad9.net/.

Quad9 Foundation. *The Public Risk of Governments Controlling DNS Providers*. 2025. https://www.quad9.net.

'Quad9 Turns the Sony Case Around in Dresden'. *Quad9*, 6 December 2023. https://quad9.net/.

Resource Ranges by RIR | APNIC. n.d. Accessed 2 September 2025. https://www.apnic.net/manage-ip/manage-resources/address-status/by-rir/.

'Root Servers'. Accessed 2 September 2025. https://www.iana.org/domains/root/servers.

Souppouris, Aaron. 'Turkish Citizens Use Google to Fight Twitter Ban'. The Verge, 21 March 2014. https://www.theverge.com/2014/3/21/5532522/turkey-twitter-ban-google-dns-workaround.

Sozeri, Efe Kerem. 'How Hacktivist Group RedHack Gamed Turkey's Censorship Regime'. *The Daily Dot*, 12 October 2016. https://www.dailydot.com/debug/redhack-gamed-turkey-censorship/.

Sykes, Alicia. 'DNS Query Performance'. Networking Toolbox. Accessed 26 January 2026. https://networkingtoolbox.net/diagnostics/dns/query-performance.

'Understanding Hosts Files - What You Need to Know | Lenovo IE'. Accessed 3 September 2025. https://www.lenovo.com/ie/en/glossary/hosts-file/.

'Verwaltungsgericht Düsseldorf: Access-Provider Müssen Pornografische Internetangebote Vorerst Nicht Sperren'. Accessed 1 December 2025. https://www.vg-duesseldorf.nrw.de/behoerde/presse/pressemitteilungen/2025/25_11_19/index.php.

'What Is DNS-Based Load Balancing? | DNS Load Balancing'. Accessed 1 September 2025. https://www.cloudflare.com/learning/performance/what-is-dns-load-balancing/.

'Why the Multistakeholder Approach Works'. *Internet Society*, 26 April 2016. https://www.internetsociety.org/resources/doc/2016/internet-governance-why-the-multistakeholder-approach-works/.

Worley, Jon. 'IP Geolocation: The Good, The Bad, & The Frustrating'. 11 June 2018. https://www.arin.net/vault/blog/2018/06/11/ip-geolocation-the-good-the-bad-the-frustrating/.