Mandated DNS Blocking

Critical Considerations

7 November 2025

Summary

Mandated DNS blocking, often presented as a straightforward policy solution, is ineffective, harmful, and impractical. It is **ineffective** because it is easily circumvented by users and fails to remove targeted content, which simply reappears under new domain names. It is **harmful** because this blunt instrument cannot distinguish between lawful and unlawful material, leading to overblocking, fragmentation of the Internet's global naming system, and the failure of interconnected services, including critical security protocols. Finally, it is **impractical** because the DNS is not bound by geography, meaning national blocking orders against global resolvers create unintended and widespread extraterritorial effects. Mandated blocking is the wrong tool for a role the DNS was never designed to play. To address online harms, interventions must focus on the content itself and the actors responsible—not on compromising the universality, reliability, and security of the Internet's core infrastructure.

Introduction

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.

This approach seems straightforward: if a name cannot be translated to its IP address, the content cannot be reached. However, this 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. Furthermore, the global, decentralized nature of the DNS makes location-based blocking hard to confine to a single jurisdiction.

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.

Understanding the DNS (A Very Brief Primer)

The DNS provides the naming function that allows the Internet to work in practice. 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 offer translation between human-friendly names, such as "example.com", into the numerical identifiers that the 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 delivering data packets.

The translation process between names and addresses can be understood through two key components of the DNS. First is the *recursive resolver*, which is a server that receives the user's request to resolve a given domain name (e.g., "example.com") and is responsible for finding its corresponding IP address. To do so, the resolver sequentially queries a series of *authoritative name servers*, which are the servers that hold the official records for specific portions, or "zones", of the DNS. The process begins at the top of the hierarchy, the so-called "root" zone, which tells the resolver where to find the servers for the next level (such as, ".com"). Those servers, in turn, point the resolver to the authoritative servers for the requested domain (such as "example.com"), which finally provide the IP address.

This highlights two important characteristics of the DNS:

- **DNS resolution is consistent, but not persistent:** 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 details of what is being accessed, such as a specific web page path[1], image, or video file, are handled later

[1] A web page's 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/DNSblocking.html", the domain name is "example.com", while "/articles/2025/DNS blocking.html" is the page path. DNS resolution is ignorant of anything but the domain.

by the web server and browser (or some other application). As a result, the DNS operates only at the level of domain names, not at the level of individual pages or other resources.

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[2]. Instead, most devices rely on a separate recursive resolver, usually operated by their Internet service provider (ISP) or another service.

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, called 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 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.

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 based on improved speed, advanced security features, or enhanced privacy practices. They may be adopted by individuals who reconfigure their devices, app developers who need consistent DNS performance, or by entire networks that choose to outsource DNS resolution rather than operate their own resolvers[3].

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.

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

[3] Farzaneh Badiei and Sebastian Castro, *Resolving the Future – The DNS Layer and the Power to Navigate the Internet* (Digital Medusa, 2025), https://digitalmedusa.org/wp-content/uploads/2025/04/DNS-Resolvers-2025-Final.pdf.

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 recursive resolver returns the corresponding IP address so the browser (or other application) 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.

Policymakers sometimes point to the 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 it can modify or disable at any time. Users may, for instance, enable parental-control filters to block pornography, or security lists to protect against malware and phishing sites[4]. 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[5]. 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 differ in their technical and operational consequences. Voluntary filtering is typically narrow, transparent, and locally administered. For example, if overblocking occurs it can be quickly detected and resolved, with limited impact on the wider Internet. In contrast, mandated

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

[5] Internet Society, *Policy Brief: Perspectives on Internet Content Blocking* (2025), https://www.internetsociety.org/resources/policybriefs/2025/perspectives-on-internet-content-blocking/.

blocking risks introducing significant collateral effects by imposing inconsistent name resolution across jurisdictions.

Why DNS Blocking Is a Blunt Tool

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 through other resolvers or via direct connection whenever the IP address is available.

DNS blocking is easy to circumvent

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.

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 DNS blocks by local ISPs 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[6].

Secondly, the user can use a virtual private network (VPN) or the Tor network[7]. 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.

Targeted content is not removed

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. 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

[6] 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.

[7] 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.

servers almost immediately. As a result, a block placed on one domain may only temporarily affect access until the site resurfaces under another[8].

DNS blocking breaks things

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[9].

This 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.

Furthermore, 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[10]. These dependencies also mean that the blocking of a domain can lead to failures across unrelated sites and applications. 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[11].

The practice of DNS blocking also has important implications for security. For instance, the Internet community developed the 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 "man-in-the-middle" attacks in which false DNS responses are injected to redirect users to malicious sites. 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

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

[9] 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.

[10] '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.

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

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. Where DNSSEC seeks to guarantee authenticity through cryptographic validation, DNS blocking introduces deliberate inauthenticity. The result is a clash where DNS blocking orders undermine widespread adoption of DNSSEC.[12]

DNS blocking is impractical

The Internet's architecture does not align with geography or jurisdictions. 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, a domain might be registered through a registrar in one country, its authoritative name servers operated from another, and the actual web servers hosting the content are in yet another location. To the network, all these scattered locations might seem very "close" in terms of milliseconds rather than miles or kilometers, even though they are geographically dispersed.

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 and very quickly (within milliseconds), distinguish user location, which is complicated by the limitations of IP-based geolocation.

From a technical perspective, operators can implement filters that approximate a user's geographical location based on the user's IP address. However, this approach is far from precise[13]. This type of location-based filtering is imperfect and can misclassify users, especially in regions where IP allocations are fluid, 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[14]. 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.

[12] 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.

[13] '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/.

[14] '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/.

Because the DNS is a global system, blocking measures often collide with many 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 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[15].

Finally, 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. A simple block list applied globally is relatively easy to configure, but targeted blocking based on user location or legal jurisdiction increases complexity and cost significantly. The more fragmented the rules become, the greater the operational burden on DNS operators and the greater chance of error in over- or underblocking. While the incremental cost of supporting an additional blocking rule may be modest, the costs multiplies as more jurisdictions impose different requirements. The result may be market consolidation, where only the largest players can operate at scale, reducing competition and diversity in DNS service provision.[16]

From a user's perspective, blocking also risks affecting service quality as integrating geolocation checks into DNS resolution can increase latency. Although delays of just a few milliseconds may seem insignificant, at the scale of the Internet they impact user experience because complex pages load much slower as each component is reached via DNS and undergo e.g. a geolocation check. This can also cause users to switch to alternative resolvers, which undermines both compliance and business goals.

Conclusion

As a public policy tool, DNS blocking is ineffective, harmful, and impractical.

It doesn't work because it can be easily circumvented and fails to remove the targeted material that is still accessible, including under a new domain name. It breaks things by overblocking lawful content, fragmenting global name resolution, and causing collateral failures across interconnected services, including security. It's impractical since the DNS is not bound by geography, and global resolvers mean that national blocking orders produce extraterritorial effects.

[15] Ernesto van der Sar, 'French Piracy Blocking Order Goes Global, DNS Service Quad9 Vows to Fight', *TorrentFreak*, 12 December 2024, https://torrentfreak.com/french-piracy-blocking-order-goes-global-dns-service-quad9-vows-to-fight-241212/ .

[16] David Abecassis et al., The Economic Cost of Network Blocking (Analysys Mason, 2025), https://www.analysysmason.com/consulting/reports/network-blocking-economic-impact-jul25/.

For all these reasons, mandated DNS blocking is the wrong tool for public policy enforcement. A role it was never designed to play.

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.

References

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/

Badiei, Farzaneh, and Sebastian Castro. Resolving the Future – The DNS Layer and the Power to Navigate the Internet. Digital Medusa, 2025. https://digitalmedusa.org/wp-content/uploads/2025/04/DNS-Resolvers-2025-Final.pdf

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/

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

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

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

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/

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

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

Sar, Ernesto van der. 'French Piracy Blocking Order Goes Global, DNS Service Quad9 Vows to Fight'. TorrentFreak, 12 December 2024. https://torrentfreak.com/french-piracy-blocking-order-goes-global-dns-service-quad9-vows-to-fight-241212/

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

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