Avoiding Unintended Harm to Internet Infrastructure

The Internet Architecture Board (IAB) provides long-range technical direction for Internet development, ensuring the Internet continues to grow and evolve as a platform for global communication and innovation. It also provides oversight of several administrative activities and relationships on behalf of the Internet Engineering Task Force (IETF). The IAB is chartered both as a committee of the IETF and an advisory body of the Internet Society. RFC 2850 documents further details about the IAB.

The IETF is a global organization whose goal is to make the Internet work better, with individual participation from implementers and users of technology, governments, and civil society organizations. The IETF is responsible for the key Internet technology standards, including IP, TCP, DNS, BGP, TLS, and HTTP, to name but a few. For further information about the purpose and mission of the IETF, see RFC 3935.

Many legislative and regulatory efforts around the world are currently focusing on how to preserve the Internet's benefits to society while curtailing undesirable uses.

While members of the technical community that designs, builds, and operates the Internet have a variety of positions regarding these efforts, the Internet Architecture Board (IAB) has a responsibility to inform legislators and regulators when proposals might have impacts on the viability of the Internet as a whole, considering all of its uses and users.

This statement discusses possible unintended effects policy and regulatory proposals may have on the Internet'. It develops our recent submission regarding Australia's Assistance and Access Bill 2018 into a more general overview of the potential impacts of such proposals, along with our recommendations for avoiding adverse outcomes.

It focuses on trust as a core mechanism of a secure and functional Internet, and how applying legal instruments intended for consumer- or business-facing services and products to Internet infrastructure services can have serious and undesirable impact upon the Internet, with an ultimate result of fragmenting the Internet itself.

Encryption as a Foundation for Trust

Over the past few years, Internet applications have increasingly adopted encryption (for example, using Transport Layer Security (TLS) or end-to-end messaging security), both as part of a long-term effort to improve Internet security and as a response to pervasive monitoring and other attacks.

Some jurisdictions have responded to these developments by proposing that authorities gain access to encrypted data through legal requirements upon vendors and operators of encryption products and services.

After vigorous representation of the risks of doing so by a broad cross-section of the community (perhaps best represented by the Keys Under Doormats paper), it now appears that weakening cryptography with legal instruments is no longer a focus in many jurisdictions.

The IAB supports this emerging norm; encryption is a core building block of a secure, trustworthy Internet. Proposals that allow third-party access to encrypted communications inevitably weaken security for users and the communication systems they use – not only directly on the Internet, but also throughout their daily lives (for example, in healthcare and banking), thanks to the ubiquity of Internet protocols.

Keeping the mechanism of encryption free from tampering is one necessary tool to ensure trust, but is not sufficient alone to ensure a secure, functional Internet.

Trust as a Building Block of the Internet

More recent discussions have focused on the endpoints of communication. They are of interest because many applications on the Internet are hosted by someone who is not participating in their operation directly, but still has access to the content of communication, or metadata related to it (also known as intercept-related information). Thus, they can provide access without “breaking” encryption.

For example, if a suspect has an account on a social network, law enforcement may believe that that account’s contents could aid an investigation. Some jurisdictions have sought to codify legal processes by which law enforcement can obtain that content from the social network provider.

The designers of such processes need to consider many factors, and we take no position on whether or not it is advisable to enact them for services *on* the Internet. However, when legal instruments designed for consumer-facing services and products (e.g., shopping, file storage, chat) apply to Internet infrastructure (such as routing, Web security, and domain names), there can be significantly harmful effects to the Internet itself. Doing so risks the many social and economic benefits that the Internet provides.

These efforts put the Internet at risk because many of the underlying protocols that make the Internet work rely on relationships as well as technology to function correctly. As a result, they are critical to the healthy function of the Internet and uniquely impacted by such proposals.

We give an incomplete list of examples below.

Trust and Internet Routing

The Internet consists of tens of thousands of networks operated by independent network operators. The Border Gateway Protocol (BGP) is the mechanism that glues them together, routing traffic towards its destination.

Network operators publish BGP “announcements” to inform each other about the parts of the Internet to which they can send traffic. Currently, there are limited technical means for network operators to verify the authenticity of BGP announcements. Thus, proper routing of Internet traffic relies largely upon trust between network operators.

If authorities could require a network operator to issue a BGP announcement directing traffic to a network under law enforcement control, the resulting misrouting would create collateral damage, since doing so would unavoidably reroute a considerable amount of unrelated network traffic.

However, more fundamentally, doing so would damage that operator’s reputation with other networks and thereby limit the operator’s ability to continue as a viable network, due to the unwillingness of other networks to trust their BGP relationships with the operator. Customer trust in these networks would likewise be eroded, and together, these effects would encourage consolidation of power into the hands of a few, very large networks -- ultimately raising costs, shifting market dynamics with a detrimental impact on the economy.

Trust and the Web

Another example of the vitality of trust is the Web Public Key Infrastructure (PKI), upon which the security of the Web (along with several other applications) relies.

The Web PKI enables users to authenticate the identity of the Web sites they visit, preventing “Man-in-the-Middle” and many other attacks. It relies on “certificates” issued by Certificate Authorities (CAs); browsers can cryptographically verify that a CA issued a given certificate. There are many CAs; they are independent entities that Web browsers, Web sites, and Web users all must trust for the Web to be secure.

If a legal instrument could be interpreted as allowing authorities to require a CA to issue a certificate improperly, it would damage that CA’s ability to function. Any forced issuance of an improper certificate would necessarily become apparent through Certificate Transparency (CT) logs, which are public repositories that document certificate issuances.

Thus, an intentional mis-issuance would jeopardize a CA’s continued operation, which relies solely upon the trust granted to it by browsers, Web sites, and ultimately the end-users of the Internet.

Even if a CA is not required to mis-issue a certificate, the possibility that compulsion is possible harms it, because its reputation is the basis of its function. If CAs were commonly subject to such instruments in multiple jurisdictions, it would make the Web less reliable and trustworthy overall, with corresponding economic and social impact.

This harm could occur because the CA's potential customers (Web sites) and users (people browsing) are distributed across the globe, not just in the affected jurisdiction, and those stakeholders do not typically have the same protections available to local entities. Even if they did, most would prefer to avoid interacting with a potentially compromised entity in a foreign jurisdiction.

Trust and Domain Names

The Domain Name System (DNS) is another place where trust – or the lack thereof – can cause irreparable damage to businesses and users.

The DNS provides a way to look up the Internet Protocol (IP) address associated with a domain name. Many Internet users rely on a DNS resolver provided by another entity, such as their ISP, to provide DNS answers[1], leaving them vulnerable to a variety of attacks. In particular, returning the wrong (or no) address breaks the trust that users have in this infrastructure.

Most Internet users, therefore, rely heavily on trust in their DNS resolver and - unless they use encrypted protocols like DNS over TLS - any network between them and the resolver.

Interference in these and similar parts of the Internet infrastructure harms not only the specific entities involved but also the Internet overall. If a legal instrument can be perceived as even *potentially* compromising such trusted Internet infrastructure, it has the potential to reduce trust in the Internet itself, and change how people use it -- with corresponding social and economic impact.

Impact on Internet Evolution

Beyond the impact of violating trust in the infrastructure, it is also essential to consider how legal constraints can unintentionally harm the future of the Internet.

[1] Though Domain Name System Security Extensions (DNSSEC) offers a means of ensuring that results are accurate, the protocol is not thoroughly implemented and deployed. Even where DNSSEC is deployed, many Internet users rely on their DNS resolver to perform validation.

“The Internet” is not one monolithic application or service, but rather a dynamic substrate for a continually changing set of applications and uses. It provides a basis for competition and interoperation, and a platform for the delivery of applications.

Applications built on top of the Internet (e.g., “apps” and Web sites), on the other hand, are diverse in their size, business models (or lack thereof), their policies, their audiences, and their security properties.

Much of the value of the Internet comes from the ability of new applications and services to be provided without hindrance; so-called “permissionless innovation.” Legal and regulatory interventions that hinder such innovation will, over time, encourage consolidation of power into the hands of a few large actors, creating so-called “walled gardens” and impossible barriers to market entry.

For example, a legal instrument designed to curb the worst aspects of social networks can have severe effects on smaller applications if applied too broadly. This, in turn, can cement the dominance of a small number of very large social networks, preventing the emergence of new and valuable applications.

In other words, the Internet is not a static system; it continually evolves, and interventions that appear reasonable today may not only be unworkable in the future, they may also prevent future evolution at a considerable social and economic cost.

Recommendations

We encourage efforts to develop legal and regulatory instruments that apply to the Internet to consider the following:

- Continue to support the widespread deployment and use of strong encryption. Laws and policies should refrain from imposing requirements that could have the effect -- intentionally or unintentionally -- of weakening encryption or deterring its use in any manner.

- When creating requirements to allow law enforcement access or control, explicitly exempt Internet infrastructure services; in particular, communication between network operators (e.g., using BGP), DNS server operators (authority and otherwise), and PKI vendors (e.g., CAs and their resellers). These exemptions should be clearly and unambiguously stated, to avoid even the perception that they could be compromised.

- Endpoint-focused interventions and capabilities should be as targeted as possible; for example, if large social networks are a primary concern, they should be explicitly listed, rather than targeting all Internet services.

- Involve all affected communities when considering legal or regulatory interventions. For example, when considering regulation or law with technical implications for

Internet services, it is not sufficient to vet proposals with “big tech” companies, as they will affect other parties who rely on those services. Restricting the consultation process to them only emphasizes the control and power they have on the Internet, leading to more consolidation of power in their hands in the long term. Limiting involvement to reviewing a proposed regulation is also problematic; deeper engagement by all stakeholders is necessary.