Security and privacy in the public interest: Singing to the choir

A talk on how to confront the discordant effects of security and privacy features by placing people at the center of technology design.

Introduction

Security and privacy standardisation introduces tensions that impact the public interest. This invited talk discusses how protocol designers and public interest advocates can together confront the discordant effects of security and privacy features by placing people at the center of technology design. For the sake of balance, there is a work in progress framework based on a paper that outlines the public interest detriments of more private DNS lookups: Internet measurements become opaque, service provision consolidates, abuse mitigation is harder, accessibility features break, and even risks of internet shutdowns and censorship become greater. And yet it is in the public interest and the interest of standards communities to properly research and mitigate these tensions so as to remove barriers to the ubiquitous adoption of strong privacy and security techniques.

It is my aim to help document the most common public interest tensions, or human rights considerations, for the standardisation of security and privacy protocols. This documentation doesn't necessarily alleviate or avoid those tensions, but it does hope to ease them.

At the heart of privacy concerns is the wanton dissipation of user data, including (or especially) metadata. Attempts to reduce or encrypt data, especially on the network, leads to a whole set of problems with current architectures.

At the heart of security concerns is the identification and mitigation of abuse and risk. Attempts to classify network data and learn about traffic patterns is directly at odds with security (as we know), but can also negatively impact user-centric mitigations. It all depends on your threat model, as they say.

I will talk more at length about the former because while it is rather dull to pit privacy against security, it is far more novel to pit privacy against other domains relevant to the public interest. However, by the

end of the talk I hope that I have sparked enough ideas that we can begin a dialogue about how to adapt this framework for security researchers as well as privacy engineers.

So, to that end, I will get into the 4-5 areas of tension with the public interest when addressing privacy and security problems through standards: measurement and research, consolidation and verticalisation, abuse and threat mitigation, usability and accessibility, and shutdowns and censorship.

Measurement and research

Measurement of the network, adoption of protocols and behaviours, as well as security research and cybersecurity mitigation are critical for the public interest.

Internet research through measurement is important to the public interest because knowledge of network performance and operation are critical for empowering users as consumers, in sustaining a healthy internet environment, and also in the monitoring of behaviours that might impact human rights, such as freedom of expression and censorship. It can be argued that web privacy measurement studies have played an important role in highlighting privacy abuses on the internet.

There are three main tensions that arise at this intersection:

1) Measurements done for research cost user privacy, even though such measurements play a fundamentally democratic role in performance monitoring. Especially in large scale measurements, gathering consent of the user can prove to be a hard issue, but researchers highlight various ways with which consent could (and should) be gathered (informed, proxy, or implied) so as to conduct internet research in an ethical manner. Furthermore, they urge researchers to employ various safety considerations in measurement, alongside risk analysis and security considerations.
2) Internet measurement becomes more difficult when user data is more private. For example, DNS is critical network intelligence for censorship measurement and recently lookups are more private. However DoH/DoT provision might actually be an improvement given how probes have used DoH as a trusted, unimpeachable source with which to compare DNS poisoning techniques used for censorship.
3) Securitisation of one network may be detrimental to securitisation of the internet at large. Threat intelligence must find a balance between disclosing vulnerability and information sharing such that threats can be monitored and mitigated more broadly. Established networks of trust might include more public interest participation.

Future areas of work: privacy preserving measurement, use of private and secure protocols by probes themselves, and folding in public interest institutions into the security incident response community.

Consolidation and verticalisation

Many applications and services that run on the internet are increasingly consolidated. Proxy-based solutions to security and privacy shifts protocol preferences into applications and to the detriment of provider and traffic diversity. There is a special risk to privacy of users by centralising data with fewer providers. And it creates an environment conducive to monopolies, such as consolidation in DNS resolvers, and more directly through vertical integration, which is itself a consolidated architecture that weakens market competition.

The centralization of user data raises the following concerns: data mining, law enforcement and intelligence agencies gaining access to information, conflicts in jurisdictional privacy laws, and creating single points of failure and targets.

A tension emerges between decentralisation, and privacy and security protocols that use intermediary architectures and proxies. Yet the technologies themselves need not be delivered in this manner and public interest advocates have argued for developers and providers to combat the tendencies of centralisation and point towards user preference in choosing whether and how to use them.

The latter point raises the issue of user competence and knowledge in internet protocols, opening suggestions for user empowerment, and in immediate practice, brings the urgency to think about the power of the default setting and user agen*ts* over user agenc*y*.

One final point is that maybe consolidation is good, actually. There are roles for consolidators when centralisation provides useful functions, that we term **tactical centralization**. One such function is easily deploying security and privacy enhancements via software to as many end users as possible in all corners of the globe, such as walled-garden e2ee messaging or DoH provision in major browsers. Another function is end-user resilience made possible through services that are "too big to block" (TBTT), such as domain fronting, which works (worked!) beautifully for both privacy and censorship. But if consolidators are a persistent reality, they must at least be responsible consolidators.

Abuse and threat mitigation

Abuse and threat mitigation are the essence of security. But while network data is leveraged to detect and mitigate abusive behaviour on networks, some of the techniques for which may be impossible or simply more difficult with privacy enhancements.

Privacy enhancing technologies can have the unintended consequence of making abuse mitigation harder. For example, deploying DoH and DoT would render invisible malicious domains, making it hard to detect and block them. Losing the ability to mitigate abuse on a network is a loss in the public interest. This has knock-on effects through other layers as well, such as when dealing with moderation of abusive content, not just abusive network behaviour. Controversial opinion: the refusal of platforms to mitigate illegal or harmful content has put undue pressure on lower layers that has traded off more widespread security and privacy crises.

But while the heart of security concerns rests upon the identification and mitigation of abuse and risk, network operators and systems administrators recognise that there are larger internet- and sector-wide mechanisms in place that facilitate information sharing. Those should be privacy conscious. They should also include public interest advocates and at-risk community representatives.

Future work should scope privacy-compatible abuse mitigation and network management, such as communication systems to share threat intelligence, privacy as a requirement for security functions, and monitoring functions in cooperation with user agent endpoints.

Usability and accessibility

Restricting signals on the network– be it for privacy or security reasons– impacts end users and user agents from performing desirable intermediary functions that may be tailored to meet specific needs.

Aside from the reality that security and privacy enhancing end-user tools tend to be harder to use, there are technical incompatibilities with architectures that obfuscate information about users and systems. One example is an end-to-end encrypted messaging app that does not interoperate with a voice assistant, often for very good privacy and security reasons but to the detriment of a visually impaired user.

Future work for security and privacy standards must take into consideration usability for end users or adoption speed will be diminished.

Shutdowns and censorship

Internet shutdowns violate the rights to free expression, access to information, and assembly. Internet shutdowns, either temporary or longer-term, can negatively impact the economy and other social and cultural life. Shutdowns, or temporal censorship at-scale, are a worst-case scenario for all people, not just users as they often come in moments of crises and impact public services as well as protest and dissent.

The reverberating effects of these extreme measures are at the heart of “the dictator’s dilemma.” Yet for some digital authoritarians, a fully controlled but flawed information flow is better than free-flowing information. Privacy and security protocols and applications designed to circumvent controls, checkpoints and surveillance can trigger a reaction that is more extreme and detrimental.

The question any security and privacy standards developer is: is there the possibility of authoritarian governance models’ to resort to blanket shutdown based on privacy respecting protocols, wholesale, whether or not end-users have even chosen to use them?

To the security experts in the room

What other issues are inherently at odds with security, especially, but also privacy? How do you go about creating harmony for *your* users, eg those with a similar threat model? When we consider all end users, including dissidents and underconnected populations, what other tensions arise when security protocols become ubiquitous?

Conclusions

To conclude there are resonant areas of new work where public interest issues interplay.

The vast public interest benefits of measurement and research of the internet should incentivise future areas of work for standardising security and privacy:

- privacy preserving measurement,
- use of private and secure protocols by probes themselves, and
- folding in public interest institutions into the security incident response community.

Mitigating the consolidated and verticalised network effects requires more work for security and privacy standards in the areas of:

- meaningful user choice and empowerment,
- default settings and user agents, and
- the responsible leveraging of popularity and TBTB.

While security standards are ultimately about abuse and threat mitigation, this may not always align with the threat model of the public interest or specific at-risk user communities. Where privacy and security are at clear tension here, there are additional considerations including:

- privacy-compatible abuse mitigation and network management,
- secure/private communication systems to share threat intelligence,
- privacy as a requirement for security functions, and
- monitoring functions in cooperation with user agent endpoints.

Usability and accessibility are particular challenges for security and privacy standards therefore future work must take them into consideration or users are left behind and speed of adoption will be diminished.

Some privacy and security standards are aimed at mitigating human rights risks for communities under authoritarian regimes. Retribution must be considered:

- Is there the possibility an authoritarian will resort to a blanket shutdown based on end-users or services employing this standard?
- Does this standard take that risk with explicit user consent?
- What is the likely scale of a provoked reaction?

With each iteration of this talk I learn something new. Usually it is additional complexity, rather than simplification of the problem space. Nonetheless it is a valuable exercise to begin a dialogue about how to build a mental model for security researchers and privacy engineers that can more systematically address the interests of the public, which includes end users and at-risk populations.