The Man in the Middlebox

Violations of End-to-End Encryption

JASMINE PELED, BENDERT ZEVENBERGEN, AND NICK FEAMSTER

Jasmine Peled currently works on computer network analysis at the Department of Defense. She recently graduated from Princeton University, where she studied computer science and philosophy. Her work at Princeton focused on how undergraduate computer science courses can better incorporate material about ethics in order to encourage students to consider the ethical and societal implications of the technologies they develop. Jasmine's senior thesis, "Towards a Pedagogy of Principles: Teaching Ethics in Computer Science," received Princeton's Outstanding Senior Thesis Award. jasminepeled21@gmail.com

Ben Zevenbergen is a visiting professional specialist at the Center for Information Technology Policy at Princeton. His work mostly consists of multidisciplinary investigations in the ethical, social, and legal impacts of Internet technologies, and vice versa. At CITP Ben is working on the engineering ethics and political theory impacts of artificial intelligence. Ben is currently finishing a PhD at the Oxford Internet Institute about the research ethics for technical projects that involve unsuspecting Internet users as data subjects. benzevenberger@princeton.edu

We consider the ethical issues of the paper "Multi-Context TLS (mcTLS): Enabling Secure In-Network Functionality in TLS" [8], which presents a method to extend the Transport Layer Security (TLS) protocol to allow it to support middleboxes. Specifically, to what extent should third parties be able to decrypt traffic between two Internet endpoints for various purposes, ranging from performance to security? This is the first column in a series about ethics that we hope will encourage ongoing discussion and debate in the research community about ethical considerations that may arise in the course of networking, security, and systems research.

Ongoing research in the computer science communities of security, privacy, and networking investigates and develops network applications and appliances that may improve Internet performance and security, often by modifying traffic en route between two Internet endpoints. Middleboxes constitute one such example of this capability; middleboxes are defined as "any intermediary box performing functions apart from normal, standard functions of an IP router on the data path between a source host and destination host" [1]. Middlebox functionality includes transcoding videostreams to different bit rates or detecting attacks, often through inspection of the contents of a packet's payload.

Because some of this functionality can require inspecting the contents of network traffic, these middleboxes may need to break end-to-end encryption, decrypting traffic midstream to facilitate operating on packet contents. mcTLS describes mechanisms for breaking the end-to-end encryption of TLS specifically to enable middleboxes to view and edit data and metadata.

Middleboxes and End-to-End Encryption

The rise of end-to-end encryption is generally heralded as a positive development, as it protects both the integrity and confidentiality of communications between Internet endpoints, thus protecting sensitive transactions and preserving user privacy.

On the other hand, if traffic is encrypted, conventional middleboxes have difficulty performing any operation that depends on seeing packet contents. In response, researchers have grappled with this problem in various ways [6]. One approach involves developing techniques that can still operate on encrypted traffic, including techniques that can perform operations on packet headers alone [5] or limited types of operations on encrypted messages [11]. Yet, certain types of operations that require deep packet inspection may be either inefficient or ineffective when payloads are encrypted; thus, another approach involves developing a "backdoor" of sorts that allows an Internet service provider (ISP) to decrypt encrypted communications in flight.

ISPs have developed an increased interest in deploying middleboxes that perform operations on traffic that is en route between source and destination. For example, ISPs often deploy middleboxes that perform intrusion detection and detect a range of different types of attacks; these middleboxes may also perform certain performance optimizations, such as transcoding a videostream to a lower bit rate or performing other types of optimizations (e.g., WAN acceleration, load balancing). These operations may depend on at least inspecting traffic contents; in some cases, the traffic contents may even be modified.

Nick Feamster is a Professor in the Computer Science Department at Princeton University and the Deputy Director of the Princeton University Center for Information Technology Policy (CITP). He was formerly a Professor at Georgia Tech, and received his MEng and PhD degrees from MIT. He has won many awards for his networking research, at ACM SIGCOMM, IMC, and USENIX NSDI. Nick is also an avid distance runner, having completed nearly 20 marathons and the Comrades ultra-marathon in South Africa.
feamster@cs.princton.edu

Multi-context TLS (mcTLS) is one such technology; it permits ISPs to decrypt secure, end-to-end sessions of TLS Internet traffic by third parties, allowing them to control, read, and write the data in the communications. The authors of the paper [8] outline several technical advantages to mcTLS:

- In-network functions may be more effective at scale, in contrast to relying on endpoint-based functionality alone.
- Middleboxes may be useful for both users and service operators in terms of speed and data storage.
- Middleboxes may help protect personal information by acting as a watchdog over applications that may leak data unwittingly.

mcTLS is based on the premise that, just like end-to-end encryption, middleboxes are a "useful part of the Internet and are here to stay." More generally, the question of whether (and how) middleboxes should have access to encrypted communications is under active discussion in industry standards organizations, such as the Internet Engineering Task Force (IETF) [7].

A natural question concerns whether the increased in-network capabilities that result from breaking end-to-end encryption offer benefits that outweigh the risks of harm to stakeholders. A related question concerns whether the development and deployment of such research should focus on technologies that weaken end-to-end encryption in favor of potentially improved security and performance, versus technologies that can operate on traffic with encrypted payloads, potentially with reduced effectiveness.

The Appropriate Ethical Lens

Ethical analysis can take many forms, which are best understood on a spectrum. On one end of the spectrum is *normative ethics*—as practiced in academic philosophy—which studies reasoning methods such as utilitarianism, deontology, and virtue ethics. *Ethics compliance frameworks* such as research ethics or medical ethics—which consist of more formal procedures for specific professions—are on the other end of the spectrum. In between these two approaches to ethics are several other, more applied types of ethics sub-disciplines, such as information ethics, technology ethics, computer ethics, data ethics, bioethics, animal ethics, among many others. Compliance-ethics frameworks typically consist of "check-box exercises" that may be rooted in law; applied ethics have some generally agreed upon methodologies for reasoning about sectors of society; and normative ethics studies the reasoning methods themselves. For this article, it is relevant to establish whether man-in-the-middle technologies such as mcTLS should be analyzed through the lens of research ethics or through a different approach.

The framework of research ethics is typically an appropriate lens for an academic paper. This framework is commonly applied to a study or experimentation when (1) it presents research in the formal sense, and (2) when the research is conducted with human subjects. In the United States, research in the formal sense is defined in the US Code of Federal Regulations on the Protection of Human Subjects as a "systematic investigation, including research development, testing and evaluation, designed to develop or contribute to generalizable knowledge" [10].

Once it has been established that a given paper constitutes research, the next question is whether the authors conduct research on human subjects. Formal regulations on the protection of human subjects in research apply to persons who conduct research (e.g., the

Common Rule [2]). Although security and networking researchers typically see themselves as conducting research on technical systems, the Internet is more properly understood as a socio-technical system in which humans and technology interact. Humans will often be implicated in data collection.

The mcTLS technology aims to intercept the Internet traffic of humans, though the paper discussed in this column merely proposes a novel functionality but does not actually present data from experimentation on live Internet traffic. Instead, the paper presents the research and development of a new technology. Therefore, the formal framework of research ethics (such as the Common Rule) need not be applied. However, even when formal requirements do not demand research conform to a research ethics checklist, researchers should still assess the broader ethical impact of their work. After all, research that does not constitute "human-subject research" may still affect people, and this series of columns seeks to bring to mind some questions that researchers should be asking themselves.

Research into computers and networked systems have traditionally challenged the principles laid out in existing research ethics procedures, such as the Belmont Report. In response, several computer science communities embraced the Menlo Report [3], which interprets the principles of the Belmont Report [4] and applies them to computer and information security and measurement research specifically. Additional networked systems ethics guidelines were developed through lengthy processes of reflection and iteration in workshops by scholars from many different disciplines [9]. Because the Menlo Report is more applicable to experimentation with human subjects on the Internet, the analysis in this article will lean on the concepts presented in Networked Systems Ethics Guidelines [9].

Technology Ethics Analysis

The Networked Systems Ethics Guidelines suggest that researchers aim to understand a technology within the social context where it operates. This social context includes an analysis of the stakeholders, the aims, benefits, risks of harm, meaning of collected data in context, shifts in power, and an understanding of the affected values. The guidelines then suggest analyzing the impact of the values on stakeholders and the socio-technical environment, the values themselves, and any foreseeable unintended consequences. It is useful to link these analyses to the technical sources of the original design. When the impact of technical alternatives have been considered in minimizing risks of harm, the guidelines suggest managing the residual risks through information governance methods, also known as responsible data stewardship. We will preface each section with a question from the guidelines.

Aims and Benefits

What are the aims and benefits of the project? How will the research benefit society and specific stakeholders?

The technology presented in the mcTLS paper [8] realizes a technology to intercept, analyze, and possibly manipulate Internet traffic that has been encrypted on an end-to-end basis. The proposed tool would replace previous "hacks," which ostensibly decrease security in the existing all-or-nothing security model. The authors state the aims of the mcTLS project concretely as follows: (1) to optimize network resource usage, (2) to improve user experience, and (3) to protect clients and servers from security threats. This tool would only be applied with the consent of all the parties involved in the connection.

Naylor et al. [8] state some further benefits that could be considered as secondary goals. For example, the authors mention that the in-network services may increase competition, innovation, and choice for end-users. Another stated benefit is that the use of middleboxes may reduce energy consumption by all stakeholders on the Internet.

The aims and benefits appear to be presented from the point of view of an ISP or network operator. The interests of end-users on the Internet are scarcely considered. The second-order benefits to society are difficult—if not impossible—to prove or support with evidence, and the paper does not consider some of the unintended social harms that may result from this tool, particularly the fact that breaking end-to-end encryption in this way will give the network operator complete power to read users' Internet traffic.

Privacy

Which definitions or explanations will be used to assess a value? Is the risk of harm high, medium, or low?

The interception and possible processing and dissemination of end-users' Internet traffic data may be considered a violation of their privacy. The concept of privacy is vague and illusive, however, and has thus been difficult to define precisely. Privacy may be best understood as an umbrella term referring to a group of related concepts, issues, and values that protect the individual's private life from intrusions by others. The use of mcTLS on end users' encrypted traffic violates the sub-category of *information privacy*, especially if their data is further processed or disseminated to third parties.

Privacy violations can be harmful in immaterial ways, though they may also reveal information about persons that can lead to physical, financial, reputational, or other types of harm, depending on the actor who receives the information and decides to act upon it. Different types of information have different types

of impact on persons when revealed, depending on the context. Given the mediating role of the Internet to support modern life, encrypted Internet traffic intercepted by mcTLS will likely contain a large variety of information types, concerning a large and diverse set of persons.

To assess the risk of harm, one must consider the type of attacker who may be interested in the information that mcTLS may expose, the level of technical sophistication they have, what actions could be taken based on the new knowledge, and what the consequences would be for an Internet user. Given the large amount of Internet traffic generated by a variety of end-users that mcTLS could intercept, all types of attackers—from individual hackers to well-resourced government surveillance actors—should be taken into account. Further, mcTLS creates a point of failure for a variety of actors to gain access to Internet traffic through both security vulnerabilities and traditional legal procedures.

Further, mcTLS poses threats to privacy by altering the context in which certain information is processed and handled. Information that may be acceptable for both endpoints of communication to view should not necessarily be shared with third parties. For example, a user may choose to enter Personally Identifiable Information (PII) into a healthcare site in order to receive personalized care, but sharing this information in one context does not constitute approval for their ISP to share it with other companies. This could violate the Health Insurance Portability and Accountability Act (HIPAA), as well as the trust that users place in their ISP to keep communications and data private.

Due to the large variety of users, stakeholders, and their purpose for using the Internet, it is nearly impossible to generalize the risk of harm and define it precisely and meaningfully. This makes it especially challenging to assess the ethical tradeoffs presented by an emerging technology. Further, what may be considered harmless today may become a much larger threat in future. For example, the creation of new data sets may allow identification of Internet users in ways that cannot be foreseen today.

Violations of end-user privacy may be justified to some extent by gaining their consent or when serving the greater good. However, an informed consent notice or other justifications should be based on factual information and informed assessments rather than self-serving arguments of increased efficiency. The complex and international nature of the Internet complicates such an analysis, because risks of privacy harm should first be defined and identified for all affected Internet users in their contexts. This is, of course, a near impossible task.

Autonomy, Consent, and Choice

Do you need to rely on informed consent from participants and stakeholders? Which stakeholders carry the burdens of the study?

The Belmont Report gives guidance regarding the respect for autonomy, balancing the value of autonomy with the interests of others:

"To respect autonomy is to give weight to autonomous persons' considered opinions and choices while refraining from obstructing their actions unless they are clearly detrimental to others" [4].

To achieve the aims and deliver the benefits identified in the paper, the existing security that users currently enjoy due to end-to-end encryption will be violated. Of course, most Internet users may not have a full understanding of the security mechanisms currently in place or even awareness of the existence of end-to-end encryption in the first place. This situation raises the question of whether taking away a good that users enjoy unwittingly as a means to achieve another end—the relative benefit of which is itself debatable—is a valid justification.

Informed consent is widely considered to be a mechanism that operationalizes the concept of autonomy of Internet users. Indeed, the authors state that both endpoints of a connection within which an mcTLS is deployed must consent to its use. However, similarly to the realm of healthcare, a key aspect of informed consent is being informed of reasonable alternatives to the proposed action. In the context of mcTLS, respect for autonomy may be understood as the obligation to fully inform an Internet user of the benefits and risks of harm in their particular context. The rejection of these benefits and risks of harm should not lead to a suspension of their Internet connection but possibly to access an alternative network within which the mcTLS tool is not operational.

Alternatively, an ISP or network operator could choose to base the legitimacy of the increase in power on a more paternalistic approach, whereby they interpret their duty of care to justify the use of mcTLS, along with its benefits and risks of harm. This constitutes a use of power over Internet users that may require some balancing through accountability mechanisms (see the Accountability section, below). For example, the ISP or network operator may choose to publish their considered justification for the use of mcTLS in their network, along with a technical description that allows some auditing of their system, as well as an information governance (or data stewardship) statement to which it can be held accountable by end-users. It is critical, though, that these explanations of benefits and potential harms posed by mcTLS do not simply use technical jargon to scare off the average user from understanding the full implications of middlebox technologies, so that supposed informed consent is, in fact, informed.

Many of these ethical concerns regarding privacy, autonomy, and choice could be resolved through agreements between ISPs and users about whether mcTLS will be implemented and how user data will be used. However, the next two sections present ethical challenges to the deployment of mcTLS which do not have such clear solutions.

Stakeholders and Power Shifts

Are particular stakeholders empowered or disempowered as a result of this project? Which values will the project conceivably impact?

ISPs and network operators will be the actors that implement and have access to mcTLS; these actors ultimately make the decision to implement and deploy such systems. These actors already have significant power over information flows, as the de facto gatekeepers to the Internet with the ability to control, manipulate, and, in some cases, observe data flows between their subscribers and other sites on the Internet. mcTLS further amplifies their power over Internet users, giving them the ability to observe the contents of network communications that might otherwise have been encrypted.

Internet users, on the other hand, will be disempowered over the collection and use of their data. Once a user has given consent to the use of mcTLS on their traffic, it will be difficult to control how their Internet traffic is collected, processed, and further disseminated, which may result in a violation of privacy. An informed consent notice referring to end-to-end encryption and the functionality of mcTLS is unlikely to be meaningful to most Internet users. First, an informed consent notice is unlikely to give the end-user meaningful information regarding the creation of a single-point-of-failure within their Internet traffic and the possible attackers or interested parties that may subsequently gain access to their data. Further, a rejection of the mcTLS tool on their Internet traffic may lead the ISP or network operator to suspend Internet access of the end-users, thereby offering users a choiceless choice (or Hobson's choice) whereby the user is asked to agree with a technically complex violation of their encrypted end-to-end connection. This may constitute a violation of their autonomy.

The mcTLS paper does not differentiate between Internet users in its analysis of benefits and harms. It is important to note that the benefits to some users can result in vastly increasing risks of harm for other users. For example, the use of middleboxes on the Internet traffic of oppressed peoples or whistleblowers in countries where the rule of law is not as effective as the authors' home country should be considered.

Unintended Consequences

Does the project potentially set a precedent for unethical methodologies that could be misused by others in the future?

Although developers of new technologies may not be directly responsible for misuses of their products under the law or under typical "checklist" research ethics restrictions, developers should still take care to mitigate potential unintended negative consequences. It is therefore important that researchers engage actively with the possibility that their methods and technologies may be misused, and design ways to mitigate those identified risks and harms. The most common ways projects influence future malevolent technology uses is through function creep and precedent setting. The following questions can help address the future concerns of creating a technology that enables a so-called "back door" into end-to-end encrypted Internet traffic.

Function creep occurs when functionality of a technology is used for other purposes than for which it was originally intended. Researchers and developers may want to consider for which other—more malevolent—aims the mcTLS technology can be used. It is relevant to consider a wide array of threat actors that would have an interest in using mcTLS for their own aims. When even companies such as Experian and Equifax are unable to keep their data secure, it is important to consider whether users can truly expect ISPs to protect their information and how adding a third party complicates this. How could the developers mitigate these foreseeable malevolent uses through their technical design?

Precedent-setting occurs when other researchers or developers can point at the use of mcTLS's technology or functionality to justify the development and use of new technologies. Technology is typically a double-edged sword that can be used for both good and bad purposes. It is therefore important to interrogate the use of precedents critically. Developers should consider how other future malevolent developers can utilize the existence and use of mcTLS to justify the development and use of technologies that cause more harms. For example, does the interception of end-to-end encrypted traffic by ISPs for efficiency in finding malware justify the interception of encrypted traffic to create profiles of Internet users for law enforcement?

When the risks of harm to stakeholders and potential unintended consequences have been identified, the researchers may pinpoint the technological causes of harms. For example, the main cause of harms is the creation of a back door and concentrated point of access for encrypted Internet traffic. Researchers should consider ways to address these issues and justify why alternative designs (or not acting at all) may be most beneficial.

Accountability

Which measures are taken to allow affected stakeholders to address concerns effectively?

Accountability is the concept that allows actors to be held liable or answerable for their actions. When an actor gains power over other stakeholders from the introduction of a technology, and the new actions may violate particular values, this increase in power should be accompanied by an increase in accountability. Accountability thus serves as a rebalancing mechanism.

Several governance mechanisms exist to allow for the exercise of accountability. For example, data governance policies can include codes of practice for employees and organizations within a sector to limit the extent to which technologies may be (mis) used. Other mechanisms include a statement of data collection policies, data retention periods for collected data, mitigation strategies for unforeseen risks, and limits on the further use or dissemination of collected data. Technical measures include information security strategies, de-identification of collected data, and further encryption of retained data. Meaningful accountability can be achieved when an organization is transparent about these policies and technical choices, as it allows third parties to audit and limit the exercise of power.

Conclusion

The introduction of technology in an environment will inevitably empower some actors over others. This is also true for mcTLS, a tool that breaks the end-to-end encryption of Internet traffic to achieve some beneficial ends, such as increased efficiency in identifying and solving security issues. However, the means by which these ends are achieved may conceivably cause harms to individual Internet users due to the shift in power over Internet traffic. End-users' autonomy and privacy are likely violated, which have further social consequences. The developers may explore options to remedy these violations through technical means. However, not all problems are solvable through technology. Therefore, the actors who employ a technology such as mcTLS should consider rebalancing their newly gained power over Internet users with accountability mechanisms, allowing for transparency (and audibility) of the systems and clear information governance policies to which affected parties can hold the operators to account.

References

[1] B. Carpenter and S. Brim, "Middleboxes: Taxonomy and Issues," 2002: https://tools.ietf.org/html/rfc3234 .

[2] "Federal Policy for the Protection of Human Subjects ('Common Rule')," 1991: https://www.hhs.gov/ohrp/regulations-and-policy/regulations/common-rule/index.html.

[3] D. Dittrich and E. Kenneally, "The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research," US Department of Homeland Security, 2012: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2445102.

[4] National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, "The Belmont Report: Ethical Principles and Guidelines for the Protection of Human Subjects of Research," Department of Health, Education, and Welfare, 1979: https://www.hhs.gov/ohrp/regulations-and-policy/belmont-report/read-the-belmont-report/index.html.

[5] G. Gu, R. Perdisci, J. Zhang, W. Lee, "BotMiner: Clustering Analysis of Network Traffic for Protocol-and Structure-Independent Botnet Detection," in *Proceedings of the 17th USENIX Security Symposium* (USENIX Security '08), pp. 139–154: http://static.usenix.org/events/sec08/tech/full_papers/gu/gu_html/.

[6] K. Moriarty, "TLS Security and Data Center Monitoring: Searching for a Path Forward," August 2017: https://www.rsa.com/en-us/blog/2017-08/tls-security-and-data-center-monitoring-searching-for-a-path-forward.

[7] K. Moriarty and A. Morton, "Effects of Pervasive Encryption on Operators," 2018: https://tools.ietf.org/html/draft-mm-wg-effect-encrypt-14.

[8] D. Naylor, K. Schomp, M. Varvello, I. Leontiadis, J. Blackburn, D. R. López, K. Papagiannaki, P. R. Rodriguez, and P. Steenkiste, "Multi-Context TLS (mcTLS): Enabling Secure In-Network Functionality in TLS," in *ACM SIGCOMM Computer Communication Review*, vol. 45 (August 2015), pp. 199–212.

[9] "Networked Systems Ethics—Guidelines," last modified on July 10, 2017: http://networkedsystemsethics.net/index.php?title=Networked_Systems_Ethics_-_Guidelines.

[10] "Code of Federal Regulations, Title 45, Public Welfare, and Part 46, Protection of Human Subjects," Department of Health and Human Services, 2009: https://www.hhs.gov/ohrp/sites/default/files/ohrp/policy/ohrpregulations.pdf.

[11] J. Sherry, C. Lan, R. A. Popa, and S. Ratnasamy, "Blindbox: Deep Packet Inspection over Encrypted Traffic," in *ACM SIGCOMM Computer Communication Review*, vol. 45 (August 2015), pp. 213–226.