Network Working Group
Internet-Draft
Intended status: Informational
Expires: August 26, 2021

J. Arkko
Ericsson
S. Farrell
Trinity College Dublin

February 22, 2021

Internet Threat Model Evolution: Background and Principles
draft-arkko-farrell-arch-model-t-redux-01

Abstract

Communications security has been at the center of many security improvements in the Internet. The goal has been to ensure that communications are protected against outside observers and attackers.

This memo suggests that the existing RFC 3552 threat model, while important and still valid, is no longer alone sufficient to cater for the pressing security and privacy issues seen on the Internet today. For instance, it is often also necessary to protect against endpoints that are compromised, malicious, or whose interests simply do not align with the interests of users. While such protection is difficult, there are some measures that can be taken and we argue that investigation of these issues is warranted.

It is particularly important to ensure that as we continue to develop Internet technology, non-communications security related threats, and privacy issues, are properly understood.

Table of Contents

1. Introduction

Communications security has been at the center of many security improvements in the Internet. The goal has been to ensure that communications are protected against outside observers and attackers. At the IETF, this approach has been formalized in BCP 72 [RFC3552], which defined the Internet threat model in 2003.

The purpose of a threat model is to outline what threats exist in order to assist the protocol designer. But RFC 3552 also ruled some threats to be in scope and of primary interest, and some threats out of scope [RFC3552]:

> The Internet environment has a fairly well understood threat model. In general, we assume that the end-systems engaging in a protocol exchange have not themselves been compromised. Protecting against an attack when one of the end-systems has been compromised is extraordinarily difficult. It is, however, possible to design protocols which minimize the extent of the damage done under these circumstances.
>
> By contrast, we assume that the attacker has nearly complete control of the communications channel over which the end-systems communicate. This means that the attacker can read any PDU (Protocol Data Unit) on the network and undetectably remove, change, or inject forged packets onto the wire.

However, the communications-security -only threat model is becoming outdated. Some of the causes for this are:

- Success! Advances in protecting most of our communications with strong cryptographic means. This has resulted in much improved communications security, but also highlights the need for addressing other, remaining issues. This is not to say that communications security is not important, it still is: improvements are still needed. Not all communications have been protected, and even out of the already protected communications, not all of their aspects have been fully protected. Fortunately, there are ongoing projects working on improvements.

- Adversaries have increased their pressure against other avenues of attack, from supply-channel attacks, to compromising devices to legal coercion of centralized endpoints in conversations.

- New adversaries and risks have arisen, e.g., due to creation of large centralized information sources.

- While communications-security does seem to be required to protect privacy, more is needed, especially if endpoints choose to act against the interests of their peers or users.

In short, attacks are migrating towards the currently easier targets, which no longer necessarily include direct attacks on traffic flows. In addition, trading information about users and ability to influence them has become a common practice for many Internet services, often without users understanding those practices.

This memo suggests that the existing threat model, while important and still valid, is no longer alone sufficient to cater for the pressing security and privacy issues on the Internet. For instance, while it continues to be very important to protect Internet communications against outsiders, it is also necessary to protect systems against endpoints that are compromised, malicious, or whose interests simply do not align with the interests of the users.

Of course, there are many trade-offs in the Internet on who one chooses to interact with and why or how. It is not the role of this memo to dictate those choices. But it is important that we understand the implications of different practices. It is also important that when it comes to basic Internet infrastructure, our chosen technologies lead to minimal exposure with respect to the non-communications threats.

It is particularly important to ensure that non-communications security related threats are properly understood for any new Internet technology. While the consideration of these issues is relatively new in the IETF, this memo provides some initial ideas about potential broader threat models to consider when designing protocols for the Internet or when trying to defend against pervasive monitoring. Further down the road, updated threat models could result in changes in BCP 72 [RFC3552] (guidelines for writing security considerations) and BCP 188 [RFC7258] (pervasive monitoring), to include proper consideration of non-communications security threats.

It may also be necessary to have dedicated guidance on how systems design and architecture affect security. The sole consideration of communications security aspects in designing Internet protocols may lead to accidental or increased impact of security issues elsewhere. For instance, allowing a participant to unnecessarily collect or receive information may lead to a similar effect as described in [RFC8546] for protocols: over time, unnecessary information will get used with all the associated downsides, regardless of what deployment expectations there were during protocol design.

This memo does not stand alone. To begin with, it is a continuation of earlier work by the two authors [I-D.farrell-etm] [I-D.arkko-arch-internet-threat-model] [I-D.arkko-farrell-arch-model-t]. There are also other documents discussing this overall space, e.g. [I-D.lazanski-smart-users-internet] [I-D.arkko-arch-dedr-report].

The rest of this memo is organized as follows. Section 2 makes some observations about the situation, with respect to communications

security and beyond. The section also provides a number of real-world examples.

Section 3 discusses some high-level principles that relate to these changes, and could be used to tackle some of the emerging issues.

Comments are solicited on these and other aspects of this document. The best place for discussion is on the model-t list. (https://www.ietf.org/mailman/listinfo/model-t)

2. Attack Landscape

This section discusses the evolving landscape of security vulnerabilities, threats, and attacks.

2.1. Communications Security Improvements

Being able to ask about threat model improvements is due to progress already made: The fraction of Internet traffic that is cryptographically protected has grown tremendously in the last few years. Several factors have contributed to this change, from Snowden revelations to business reasons and to better available technology such as HTTP/2 [RFC7540], TLS 1.3 [RFC8446], QUIC [I-D.ietf-quic-transport].

In many networks, the majority of traffic has flipped from being cleartext to being encrypted. Reaching the level of (almost) all traffic being encrypted is no longer something unthinkable but rather a likely outcome in a few years.

At the same time, technology developments and policy choices have driven the scope of cryptographic protection from protecting only the pure payload to protecting much of the rest as well, including far more header and meta-data information than was protected before. For instance, efforts are ongoing in the IETF to assist encrypting transport headers [I-D.ietf-quic-transport], server domain name information in TLS [I-D.ietf-tls-esni], and domain name queries [RFC8484].

There have also been improvements to ensure that the security protocols that are in use actually have suitable credentials and that those credentials have not been compromised, see, for instance, Let's Encrypt [RFC8555], HSTS [RFC6797], HPKP [RFC7469], and Expect-CT [I-D.ietf-httpbis-expect-ct].

This is not to say that all problems in communications security have been resolved - far from it. But the situation is definitely different from what it was a few years ago. Remaining issues will be

and are worked on; the fight between defense and attack will also continue. Communications security will stay at the top of the agenda in any Internet technology development.

2.2. Beyond Communications Security

There are, however, significant issues beyond communications security in the Internet.

To begin with, it is not necessarily clear that one can trust all the endpoints in any protocol interaction, including the user's own devices. Managed or closed ecosystems with multiple layers of hardware and software have made it harder to understand or influence what your devices do.

The situation is different, but not necessarily better on the side of servers. Even for applications that are for user-to-user communication, a typical pattern of communications is almost always via an intermediary that has at least as much information as the other parties have. For instance, these intermediaries are typically endpoints for any transport layer security connections, and able to see much communications or other messaging in cleartext. There are some exceptions, of course, e.g., messaging applications with end-to-end confidentiality protection.

For instance, while e-mail transport security [RFC7817] has become much more widely deployed in recent years, progress in securing e-mail messages between users has been much slower. This has lead to a situation where e-mail content is considered a critical resource by some mail service providers who use the content for machine learning, advertisement targeting, and other purposes unrelated to message delivery. Equally however, it is unclear how some useful anti-spam techniques could be deployed in an end-to-end encrypted mail universe (with today's end-to-end mail security protocols) and there are many significant challenges should one desire to deploy end-to-end email security at scale.

Services that are not about user-to-user to communication often collect information about the user.

Even services that are part of the infrastructure may have security issues. For instance, despite progress in protecting DNS query protocols with encryption (see, e.g., [RFC7858] and [RFC8484]), DNS resolver services themselves may be targets for attack or sources for leaks. For instance, the services may collect information or be vulnerable targets of denial-of-service attacks, attacks to steal user browsing history information, or be the target of surveillance activities and government information requests. Infrastructure

services with information about a large number of users is collected in small number of services are particularly attractive targets for these attacks.

With the growth of trading users' information by many of these parties, it becomes necessary to take precautions against endpoints that are compromised, malicious, or whose interests simply do not align with the interests of the users.

In general, many recent attacks relate more to information than communications. For instance, personal information leaks typically happen via information stored on a compromised server rather than capturing communications. There is little hope that such attacks can be prevented entirely. Again, the best course of action seems to be avoid the disclosure of information in the first place, or at least to not perform that in a manner that makes it possible that others can readily use the information.

2.3. Types of Attacks

This section discusses a few classes of attacks that are relevant for our consideration.

2.3.1. Misuse of Accidental Vulnerabilities

Not all adversarial behaviour starts as deliberate, some is initiated due to various levels of carelessness and/or due to erroneous assumptions about the environments in which those applications currently run at. Nevertheless, a leak or vulnerability can be exploited by others that misuse the data for their own purposes.

Some attacks in this category include:

- Virtualisation exposing secrets, for example, Meltdown and Spectre [MeltdownAndSpectre] [Kocher2019] [Lipp2018] and other similar side-channel attacks.

- Compromised badly-maintained web sites or services, e.g., [Passwords] or Amazon S3 leaks.

- Supply-chain attacks, for example, the [TargetAttack] or malware within pre-installed applications on Android phones [Bloatware].

- Breaches of major service providers, that many of us might have assumed would be sufficiently capable to be the best large-scale "Identity providers", for example, Yahoo (https://www.wired.com/story/yahoo-breach-three-billion-accounts/), Facebook (https://www.pcmag.com/news/367319/facebook-

stored-up-to-600m-user-passwords-in-plain-text and https://www.cnet.com/news/facebook-breach-affected-50-million-people/), Telcos (https://www.zdnet.com/article/us-telcos-caught-selling-your-location-data-again-senator-demands-new-laws/ and https://www.zdnet.com/article/millions-verizon-customer-records-israeli-data/), Google (https://www.wsj.com/articles/google-exposed-user-data-feared-repercussions-of-disclosing-to-public-1539017194), and Microsoft (https://motherboard.vice.com/en_us/article/ywyz3x/hackers-could-read-your-hotmail-msn-outlook-microsoft-customer-support).

2.3.2. Misbehaving Applications

There are many examples of application developers doing their best to protect the security and privacy of their users or customers. That's just the same as the case today where we need to consider in-network actors as potential adversaries despite the many examples of network operators who both act in the best interests of their users and succeed in defending against attacks from others.

In short, there are applications that do not act in the best interests of their users.

This can also happen indirectly. Despite the best efforts of curators, so-called App-Stores frequently distribute malware of many kinds and one recent study [Curated] claims that simple obfuscation enables malware to avoid detection by even sophisticated operators. Given the scale of these deployments, distribution of even a small percentage of malware-infected applications can affect a large number of people. The end result is an application that

Applications may also mislead users. Many web sites today provide some form of privacy policy and terms of service, that are known to be mostly unread [Unread]. This implies that, legal fiction aside, users of those sites have not in reality agreed to the specific terms published and so users are therefore highly exposed to being exploited by web sites, for example [Cambridge] is a recent well-publicised case where a service provider abused the data of 87 million users via a partnership. While many web site operators claim that they care deeply about privacy, it seems prudent to assume that some do not in fact care about user privacy in ways with which many of their users would agree.

2.3.3. Network Infrastructure Attacks

The network infrastructure may also work in an inappropriate manner. For instance, a Virtual Private Network (VPN) may misrepresent how it carries the users' traffic, for example misrepresenting the countries

in which they provide vantage points [Vpns]. A user's home network equipment may also be malicous or compromised. For example, one study [Home] reports on a 2011 attack that affected 4.5 million DSL modems in Brazil. The absence of software update [RFC8240] has been a major cause of these issues and rises to the level that considering this as intentional behaviour by device vendors who have chosen this path is warranted.

2.3.4. Untrustworthy Devices

Traditionally, there's been an implied trust in various parts of the system - such as the user's own device, nodes inside a particular network perimeter, or nodes under a single administrative control.

Client endpoint implementations were never fully trusted, but the environments in which those endpoints exist are changing. Users may not have as much control over their own devices as they used to, due to manufacturer-controlled operating system installations and locked device ecosystems. And within those ecosystems, even the applications that are available tend to have privileges that users by themselves might not desire those applications be granted, such as excessive rights to media, location, and peripherals. There are also designated efforts by various authorities to hack end-user devices as a means of intercepting data about the user.

Examples of these issues are too many to list, for instance, so-called "smart" televisions spying on their owners and one survey of user attitudes [SmartTV].

There are similar issues with larger, networked systems. As these systems evolve over time, they get used and connected in different ways, run in virtual environments, and expanded for new functions. Old assumptions and security mechanisms may no longer be applicable in these new environments, leading to security vulnerabilities.

Even in a closed network with carefully managed components there may be compromised components, as evidenced in the most extreme way by the Stuxnet worm that operated in an airgapped network.

Every system runs large amount of software, and it is often not practical or even possible to prevent compromised code even in a high-security setting, let alone in commercial or private networks. Installation media, physical ports, both open source and proprietary programs, firmware, or even innocent-looking components on a circuit board can be suspect [TinyChip]. In addition, complex underlying computing platforms, such as modern CPUs with underlying security and management tools are prone to problems. Analysis for the security of many interesting real-world systems now commonly needs to include

cross-component attacks, e.g., the use of car radios and other externally communicating devices as part of attacks launched against the control components such as brakes in a car [Savage].

Untrustworthy systems can also cause damage to other parties. Examples of this range from attacks of badly constructed IoT devices [DynDDoS] to large Internet services that become single points of failure [I-D.arkko-arch-infrastructure-centralisation].

2.3.5. Tracking

One of the biggest threats to user privacy on the Web is ubiquitous tracking. This is often done to support advertising based business models.

While some people may be sanguine about this kind of tracking, others consider this behaviour unwelcome, when or if they are informed that it happens, [Attitude] though the evidence here seems somewhat harder to interpret and many studies (that we have found to date) involve small numbers of users. Historically, browsers have not made this kind of tracking visible and have enabled it by default, though some recent browser versions are starting to enable visibility and blocking of some kinds of tracking. Browsers are also increasingly imposing more stringent requirements on plug-ins for varied security reasons.

Third party tracking

One form of tracking is by third parties. HTTP header fields (such as cookies, [RFC6265]) are commonly used for such tracking, as are structures within the content of HTTP responses such as links to 1x1 pixel images and (ab)use of Javascript APIs offered by browsers [Tracking].

Whenever a resource is loaded from a server, that server can include a cookie which will be sent back to the server on future loads. This includes situations where the resource is loaded as a resource on a page, such as an image or a JavaScript module. When loading a resource, the server is aware of the top-level page that the resource is used on, through the use of the Referer HTTP header [RFC7231]. those loads include a Referer header which contains the top-level page from which that subresource is being loaded.

The combination of these features makes it possible to track a user across the Web. The tracker convinces a number of content sites ("first parties") to include a resource from the tracker site. This resource can perform some function such as displaying an advertisement or providing analytics to the first party site. But

the resource may also be simply a tracker. When the user visits one of the content sites, the tracker receives both a Referer header and the cookie. For an individual user with a particular browser, the cookie is the same regardless of which site the tracker is on. This allows the tracker to observe what pages within the set of content sites the user visits. The resulting information is commonly used for targeting advertisements, but it can also be used for other purposes.

This capability itself constitutes a major threat to user privacy. Additional techniques such as cookie syncing, identifier correlation, and fingerprinting make the problem even worse [I-D.wood-pearg-website-fingerprinting].

As a given tracker will not be on all sites, that tracker has incomplete coverage. However, trackers often collude (a practice called "cookie syncing") to combine the information from different tracking cookies.

Sometimes trackers will be embedded on a site which collects a user identifier, such as social media identity or an e-mail address. If the site can inform the tracker of the identifier, that allows the tracker to tie the identifier to the cookie.

While a browser may block cookies, fingerprinting browsers often allows tracking the users. For instance, features such as User-Agent string, plugin and font support, screen resolution, and timezone can yield a fingerprint that is sometimes unique to a single user [AmIUnique] and which persists beyond cookie scope and lifetime. Even in cases where this fingerprint is not unique, the anonymity set may be sufficiently small that, coupled with other data, this yields a unique, per-user identifier. Fingerprinting of this type is more prevalent on systems and platforms where data-set features are flexible, such as desktops, where plugins are more commonly in use. Fingerprinting prevention is an active research area; see [Boix2018] for more information.

Other types of tracking linked to web tracking

Third party web tracking is not the only concern. An obvious tracking danger exists also in popular ecosystems - such as social media networks - that house a large part of many users' online existence. There is no need for a third party to track the user's browsing as all actions are performed within a single site, where most messaging, viewing, and sharing activities happen.

Browsers themselves or services used by the browser can also become a potential source of tracking users. For instance, the URL/search bar

service may leak information about the user's actions to a search provider via an "autocomplete" feature. [Leith2020]

Tracking through users' IP addresses or DNS queries is also a danger. This may happen by directly observing the cleartext IP or DNS traffic, though DNS tracking may be preventable via DNS protocols that are secured end-to-end. But the DNS queries are also (by definition) seen by the used DNS recursive resolver service, which may accidentally or otherwise track the users' activities. This is particularly problematic if a large number of users employ either a commonly used ISP service or an Internet-based resolver service [I-D.arkko-arch-infrastructure-centralisation]. In contrast, use of a DNS recursive that sees little traffic could equally be used for tracking. Similarly, other applications, such an mail or instant messaging protocols, that can carry HTML content can be integrated with web tracking. For instance, intentional tracking are also seen many times per day by email users - in one study [Mailbug] the authors estimated that 62% of leakage to third parties was intentional, for example if leaked data included a hash of the recipient email address.

Tracking happens through other systems besides the web, of course. For instance, some mail user agents (MUAs) render HTML content by default (with a subset not allowing that to be turned off, perhaps particularly on mobile devices) and thus enable the same kind of adversarial tracking seen on the web. Attempts at such intentional tracking are also seen many times per day by email users - in one study [Mailbug] the authors estimated that 62% of leakage to third parties was intentional, for example if leaked data included a hash of the recipient email address.

3. Principles

Based on the above issues, it is necessary to pay attention to the following aspects:

- Security of devices, including the user's own devices.
- Security of data at rest, in various parts of the system.
- Tracking and identification of users and their devices.
- Role of intermediaries, and in particular information passing through them.

These topics are discussed below. There are obviously many detailed technical questions and approaches to tackling them. However, in

this memo we wish to focus on higher level architectural principles that might guide us in thinking about about the topics.

3.1. Trusting Devices

In general, this means that one cannot entirely trust even a closed system where you picked all the components yourself, let alone typical commercial, networked and Internet-connected systems.

PRINCIPLE: Consider all system components as potentially untrustworthy, and consider the implications of their compromise.

There may also be ways to mitigate damages, should a compromise occur.

3.2. Protecting Information

Data leaks have become the primary concern. Even trusted, well-managed parties can be problematic, such as when large data stores attract attempts to use that data in a manner that is not consistent with the users' interests.

Mere encryption of communications is not sufficient to protect information.

PRINCIPLE: Consider information passed to another party as a publication to at least some number of entities.

This principle applies even if the communications that carry that information are encrypted.

3.3. Tracking

Information leakage is particularly harmful in situations where the information can be traced to an individual, such as is the case with any information that users would consider private, be it about messages to another users, browsing history, or even user's medical information.

PRINCIPLE: Assume that every interaction with another party can result in fingerprinting or identification of the user in question.

In many cases there are readily available user identifiers in data that is leaked, such as was the case with a recent medical information leak in Finland [Vastaamo]. But even when such identifiers are not present, in many communication methods, there is ample opportunity for narrowing down which entity is connecting,

through geolocation, fingerprinting, and correlation with other information.

3.4. Role of End-to-End

[RFC1958] notes that "end-to-end functions can best be realised by end-to-end protocols":

> The basic argument is that, as a first principle, certain required end-to-end functions can only be performed correctly by the end-systems themselves. A specific case is that any network, however carefully designed, will be subject to failures of transmission at some statistically determined rate. The best way to cope with this is to accept it, and give responsibility for the integrity of communication to the end systems. Another specific case is end-to-end security.

The "end-to-end argument" was originally described by Saltzer et al [Saltzer]. They said:

> The function in question can completely and correctly be implemented only with the knowledge and help of the application standing at the endpoints of the communication system. Therefore, providing that questioned function as a feature of the communication system itself is not possible.

These functional arguments align with other, practical arguments about the evolution of the Internet under the end-to-end model. The endpoints evolve quickly, often with simply having one party change the necessary software on both ends. Whereas waiting for network upgrades would involve potentially a large number of parties from application owners to multiple network operators. The end-to-end model supports permissionless innovation where new innovation can flourish in the Internet without excessive wait for other parties to act.

But the details matter. What is considered an endpoint? What characteristics of Internet are we trying to optimize?

There is a significant difference between actual endpoints from a user's interaction perspective (e.g., another user) and from a system perspective (e.g., a third party relaying a message). Such intermediaries can provide a useful service. As [I-D.thomson-tmi] points out, networks themselves would not exist without intermediaries that can forward communications to others.

PRINCIPLE: Limit the use of intermediaries, and what they can do.

PRINCIPLE: Pass information only between the "real ends" of a conversation, unless the information is necessary for a useful function in an intermediary.

For instance, a transport connection between two components of a system is not an end-to-end connection even if it encompasses all the protocol layers up to the application layer. It is not end-to-end, if the information or control function it carries actually extends beyond those components. For instance, just because an e-mail server can read the contents of an e-mail message does not make it a legitimate recipient of the e-mail.

This memo also proposes to focus on the "need to know" aspect in systems. Information should not be disclosed, stored, or routed in cleartext through parties that do not absolutely need to have that information. This relates to the discussion in [I-D.thomson-tmi], in that the valuable functions provided by intermediaries need to be balanced against the information that they need to perform their function. And, in a lot of cases unnecessary information is provided to intermediaries, which leads to privacy and other problems.

4. Security Considerations

The entire memo covers the security considerations.

5. IANA Considerations

There are no IANA considerations in this work.

6. Informative References

[AbuseCases]
McDermott, J. and C. Fox, "Using abuse case models for security requirements analysis", IEEE Annual Computer Security Applications Conference (ACSAC'99), https://www.acsac.org/1999/papers/wed-b-1030-john.pdf , 1999.

[AmIUnique]
INRIA, ., "Am I Unique?", https://amiunique.org , 2020.

[Attitude]
"User Perceptions of Sharing, Advertising, and Tracking", Symposium on Usable Privacy and Security (SOUPS), https://www.usenix.org/conference/soups2015/proceedings/presentation/chanchary , 2015.

[avleak] Cox, J., "Leaked Documents Expose the Secretive Market for Your Web Browsing Data", https://www.vice.com/en_us/article/qjdkq7/avast-antivirus-sells-user-browsing-data-investigation , 2020.

[Bloatware] Gamba, G., Rashed, M., Razaghpanah, A., Tapiado, J., and N. Vallina, "An Analysis of Pre-installed Android Software", arXiv preprint arXiv:1905.02713 (2019) , 2019.

[Boix2018] Gomez-Boix, A., Laperdrix, P., and B. Baudry, "Hiding in the crowd: an analysis of the effectiveness of browser fingerprinting at large scale", Proceedings of the 2018 world wide web conference , 2018.

[Cambridge] Isaak, J. and M. Hanna, "User Data Privacy: Facebook, Cambridge Analytica, and Privacy Protection", Computer 51.8 (2018): 56-59, https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=8436400 , 2018.

[Curated] Hammad, M., Garcia, J., and S. MaleK, "A large-scale empirical study on the effects of code obfuscations on Android apps and anti-malware products", ACM International Conference on Software Engineering 2018, https://www.ics.uci.edu/~seal/publications/2018ICSE_Hammad.pdf , 2018.

[DynDDoS] York, K., "Dyn's Statement on the 10/21/2016 DNS DDoS Attack", Company statement: https://dyn.com/blog/dyn-statement-on-10212016-ddos-attack/ , 2016.

[GDPRAccess] EU, ., "Right of access by the data subject", Article 15, GDPR, https://gdpr-info.eu/art-15-gdpr/ , n.d..

[Home] Nthala, N. and I. Flechais, "Rethinking home network security", European Workshop on Usable Security (EuroUSEC), https://ora.ox.ac.uk/objects/uuid:e2460f50-579b-451b-b14e-b7be2decc3e1/download_file?safe_filename=bare_conf_EuroUSEC2018.pdf&file_format=application%2Fpdf&type_of_work=Conference+item , 2018.

[I-D.arkko-arch-dedr-report]
Arkko, J. and T. Hardie, "Report from the IAB workshop on Design Expectations vs. Deployment Reality in Protocol Development", draft-arkko-arch-dedr-report-00 (work in progress), November 2019.

[I-D.arkko-arch-infrastructure-centralisation]
Arkko, J., "Centralised Architectures in Internet Infrastructure", draft-arkko-arch-infrastructure-centralisation-00 (work in progress), November 2019.

[I-D.arkko-arch-internet-threat-model]
Arkko, J., "Changes in the Internet Threat Model", draft-arkko-arch-internet-threat-model-01 (work in progress), July 2019.

[I-D.arkko-farrell-arch-model-t]
Arkko, J. and S. Farrell, "Challenges and Changes in the Internet Threat Model", draft-arkko-farrell-arch-model-t-04 (work in progress), July 2020.

[I-D.farrell-etm]
Farrell, S., "We're gonna need a bigger threat model", draft-farrell-etm-03 (work in progress), July 2019.

[I-D.ietf-httpbis-expect-ct]
estark@google.com, e., "Expect-CT Extension for HTTP", draft-ietf-httpbis-expect-ct-08 (work in progress), December 2018.

[I-D.ietf-quic-transport]
Iyengar, J. and M. Thomson, "QUIC: A UDP-Based Multiplexed and Secure Transport", draft-ietf-quic-transport-34 (work in progress), January 2021.

[I-D.ietf-tls-esni]
Rescorla, E., Oku, K., Sullivan, N., and C. Wood, "TLS Encrypted Client Hello", draft-ietf-tls-esni-09 (work in progress), December 2020.

[I-D.lazanski-smart-users-internet]
Lazanski, D., "An Internet for Users Again", draft-lazanski-smart-users-internet-00 (work in progress), July 2019.

[I-D.thomson-tmi]
Thomson, M., "Principles for the Involvement of Intermediaries in Internet Protocols", draft-thomson-tmi-01 (work in progress), January 2021.

[I-D.wood-pearg-website-fingerprinting]
Goldberg, I., Wang, T., and C. Wood, "Network-Based Website Fingerprinting", draft-wood-pearg-website-fingerprinting-00 (work in progress), November 2019.

[Kocher2019]
Kocher, P., Horn, J., Fogh, A., Genkin, D., Gruss, D., Haas, W., Hamburg, M., Lipp, M., Mangard, S., Prescher, T., Schwarz, M., and Y. Yarom, "Spectre Attacks: Exploiting Speculative Execution", 40th IEEE Symposium on Security and Privacy (S&P'19) , 2019.

[Leith2020]
Leith, D., "Web Browser Privacy: What Do Browsers Say When They Phone Home?", In submission, https://www.scss.tcd.ie/Doug.Leith/pubs/browser_privacy.pdf , March 2020.

[Lipp2018]
Lipp, M., Schwarz, M., Gruss, D., Prescher, T., Haas, W., Fogh, A., Horn, J., Mangard, S., Kocher, P., Genkin, D., Yarom, Y., and M. Hamburg, "Meltdown: Reading Kernel Memory from User Space", 27th USENIX Security Symposium (USENIX Security 18) , 2018.

[Mailbug] Englehardt, S., Han, J., and A. Narayanan, "I never signed up for this! Privacy implications of email tracking", Proceedings on Privacy Enhancing Technologies 2018.1 (2018): 109-126, https://www.degruyter.com/downloadpdf/j/popets.2018.2018.issue-1/popets-2018-0006/popets-2018-0006.pdf , 2018.

[MeltdownAndSpectre]
CISA, ., "Meltdown and Spectre Side-Channel Vulnerability Guidance", Alert (TA18-004A), https://www.us-cert.gov/ncas/alerts/TA18-004A , 2018.

[Passwords]
com, haveibeenpwned., "Pwned Passwords", Website https://haveibeenpwned.com/Passwords , 2019.

[RFC1958] Carpenter, B., Ed., "Architectural Principles of the Internet", RFC 1958, DOI 10.17487/RFC1958, June 1996, <https://www.rfc-editor.org/info/rfc1958>.

[RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, DOI 10.17487/RFC3552, July 2003, <https://www.rfc-editor.org/info/rfc3552>.

[RFC6265] Barth, A., "HTTP State Management Mechanism", RFC 6265, DOI 10.17487/RFC6265, April 2011, <https://www.rfc-editor.org/info/rfc6265>.

[RFC6797] Hodges, J., Jackson, C., and A. Barth, "HTTP Strict Transport Security (HSTS)", RFC 6797, DOI 10.17487/RFC6797, November 2012, <https://www.rfc-editor.org/info/rfc6797>.

[RFC7231] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", RFC 7231, DOI 10.17487/RFC7231, June 2014, <https://www.rfc-editor.org/info/rfc7231>.

[RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, DOI 10.17487/RFC7258, May 2014, <https://www.rfc-editor.org/info/rfc7258>.

[RFC7469] Evans, C., Palmer, C., and R. Sleevi, "Public Key Pinning Extension for HTTP", RFC 7469, DOI 10.17487/RFC7469, April 2015, <https://www.rfc-editor.org/info/rfc7469>.

[RFC7540] Belshe, M., Peon, R., and M. Thomson, Ed., "Hypertext Transfer Protocol Version 2 (HTTP/2)", RFC 7540, DOI 10.17487/RFC7540, May 2015, <https://www.rfc-editor.org/info/rfc7540>.

[RFC7817] Melnikov, A., "Updated Transport Layer Security (TLS) Server Identity Check Procedure for Email-Related Protocols", RFC 7817, DOI 10.17487/RFC7817, March 2016, <https://www.rfc-editor.org/info/rfc7817>.

[RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <https://www.rfc-editor.org/info/rfc7858>.

[RFC8240] Tschofenig, H. and S. Farrell, "Report from the Internet of Things Software Update (IoTSU) Workshop 2016", RFC 8240, DOI 10.17487/RFC8240, September 2017, <https://www.rfc-editor.org/info/rfc8240>.

[RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <https://www.rfc-editor.org/info/rfc8446>.

[RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <https://www.rfc-editor.org/info/rfc8484>.

[RFC8546] Trammell, B. and M. Kuehlewind, "The Wire Image of a Network Protocol", RFC 8546, DOI 10.17487/RFC8546, April 2019, <https://www.rfc-editor.org/info/rfc8546>.

[RFC8555] Barnes, R., Hoffman-Andrews, J., McCarney, D., and J. Kasten, "Automatic Certificate Management Environment (ACME)", RFC 8555, DOI 10.17487/RFC8555, March 2019, <https://www.rfc-editor.org/info/rfc8555>.

[Saltzer] Saltzer, J., Reed, D., and D. Clark, "End-To-End Arguments in System Design", ACM TOCS, Vol 2, Number 4, pp 277-288 , November 1984.

[Savage] Savage, S., "Modern Automotive Vulnerabilities: Causes, Disclosures, and Outcomes", USENIX , 2016.

[SmartTV] Malkin, N., Bernd, J., Johnson, M., and S. Egelman, "What Can't Data Be Used For? Privacy Expectations about Smart TVs in the U.S.", European Workshop on Usable Security (Euro USEC), https://www.ndss-symposium.org/wp-content/uploads/2018/06/ eurousec2018_16_Malkin_paper.pdf" , 2018.

[Sybil] Viswanath, B., Post, A., Gummadi, K., and A. Mislove, "An analysis of social network-based sybil defenses", ACM SIGCOMM Computer Communication Review 41(4), 363-374, https://conferences.sigcomm.org/sigcomm/2010/papers/sigcomm/p363.pdf , 2011.

[TargetAttack] Osborne, C., "How hackers stole millions of credit card records from Target", ZDNET, https://www.zdnet.com/article/how-hackers-stole-millions-of-credit-card-records-from-target/ , 2014.

[TinyChip] Robertson, J. and M. Riley, "The Big Hack: How China Used a Tiny Chip to Infiltrate U.S. Companies", https://www.bloomberg.com/news/features/2018-10-04/the-big-hack-how-china-used-a-tiny-chip-to-infiltrate-america-s-top-companies , October 2018.

[Toys] Chu, G., Apthorpe, N., and N. Feamster, "Security and Privacy Analyses of Internet of Things Childrens' Toys", IEEE Internet of Things Journal 6.1 (2019): 978-985, https://arxiv.org/pdf/1805.02751.pdf , 2019.

[Tracking] Ermakova, T., Fabian, B., Bender, B., and K. Klimek, "Web Tracking-A Literature Review on the State of Research", Proceedings of the 51st Hawaii International Conference on System Sciences, https://scholarspace.manoa.hawaii.edu/bitstream/10125/50485/paper0598.pdf , 2018.

[Troll] Stewart, L., Arif, A., and K. Starbird, "Examining trolls and polarization with a retweet network", ACM Workshop on Misinformation and Misbehavior Mining on the Web, https://faculty.washington.edu/kstarbi/examining-trolls-polarization.pdf , 2018.

[Unread] Obar, J. and A. Oeldorf, "The biggest lie on the internet{:} Ignoring the privacy policies and terms of service policies of social networking services", Information, Communication and Society (2018): 1-20 , 2018.

[Vastaamo] Redcross Finland, ., "Read this if your personal data was leaked in the Vastaamo data system break-in", https://www.redcross.fi/news/20201029/read-if-your-personal-data-was-leaked-vastaamo-data-system-break , October 2020.

[Vpns] Khan, M., DeBlasio, J., Voelker, G., Snoeren, A., Kanich, C., and N. Vallina, "An empirical analysis of the commercial VPN ecosystem", ACM Internet Measurement Conference 2018 (pp. 443-456), https://eprints.networks.imdea.org/1886/1/imc18-final198.pdf , 2018.

Appendix A. Contributors

Eric Rescorla and Chris Wood provided much of the text in Section 2.3.5. Martin Thomson's excellent document [I-D.thomson-tmi] also inspired some of the work in Section 3.

Appendix B. Acknowledgements

The authors would like to thank the IAB:

Alissa Cooper, Wes Hardaker, Ted Hardie, Christian Huitema, Zhenbin Li, Erik Nordmark, Mark Nottingham, Melinda Shore, Jeff Tantsura, Martin Thomson, Brian Trammel, Mirja Kuhlewind, and Colin Perkins.

The authors would also like to thank the participants of the IETF SAAG meeting where this topic was discussed:

Harald Alvestrand, Roman Danyliw, Daniel Kahn Gilmore, Wes Hardaker, Bret Jordan, Ben Kaduk, Dominique Lazanski, Eliot Lear, Lawrence Lundblade, Kathleen Moriarty, Kirsty Paine, Eric Rescorla, Ali Rezaki, Mohit Sethi, Ben Schwartz, Dave Thaler, Paul Turner, David Waltemire, and Jeffrey Yaskin.

The authors would also like to thank the participants of the IAB 2019 DEDR workshop:

Tuomas Aura, Vittorio Bertola, Carsten Bormann, Stephane Bortzmeyer, Alissa Cooper, Hannu Flinck, Carl Gahnberg, Phillip Hallam-Baker, Ted Hardie, Paul Hoffman, Christian Huitema, Geoff Huston, Konstantinos Komaitis, Mirja Kuhlewind, Dirk Kutscher, Zhenbin Li, Julien Maisonneuve, John Mattson, Moritz Muller, Joerg Ott, Lucas Pardue, Jim Reid, Jan-Frederik Rieckers, Mohit Sethi, Melinda Shore, Jonne Soininen, Andrew Sullivan, and Brian Trammell.

The authors would also like to thank the participants of the November 2016 meeting at the IETF:

Carsten Bormann, Randy Bush, Tommy C, Roman Danyliw, Ted Hardie, Christian Huitema, Ben Kaduk, Dirk Kutscher, Dominique Lazanski, Eric Rescorla, Ali Rezaki, Mohit Sethi, Melinda Shore, Martin Thomson, and Robin Wilton ... (missing many people... did we have minutes other than the list of actions?) ...

Thanks for specific comments on this text to: Ronald van der Pol.

Finally, the authors would like to thank numerous other people for insightful comments and discussions in this space.

Authors' Addresses

Jari Arkko
Ericsson
Valitie 1B
Kauniainen
Finland

Email: jari.arkko@piuha.net

Stephen Farrell
Trinity College Dublin
College Green
Dublin
Ireland

Email: stephen.farrell@cs.tcd.ie