

Internet Research Task Force (IRTF)
Request for Comments: 8280
Category: Informational
ISSN: 2070-1721

N. ten Oever
ARTICLE 19
C. Cath
Oxford Internet Institute
October 2017

Research into Human Rights Protocol Considerations

Abstract

This document aims to propose guidelines for human rights considerations, similar to the work done on the guidelines for privacy considerations (RFC 6973). The other parts of this document explain the background of the guidelines and how they were developed.

This document is the first milestone in a longer-term research effort. It has been reviewed by the Human Rights Protocol Considerations (HRPC) Research Group and also by individuals from outside the research group.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Research Task Force (IRTF). The IRTF publishes the results of Internet-related research and development activities. These results might not be suitable for deployment. This RFC represents the consensus of the Human Rights Protocol Considerations Research Group of the Internet Research Task Force (IRTF). Documents approved for publication by the IRSG are not a candidate for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8280>.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

1. Introduction	4
2. Vocabulary Used	6
3. Research Questions	12
4. Literature and Discussion Review	12
5. Methodology	15
5.1. Data Sources	17
5.1.1. Discourse Analysis of RFCs	17
5.1.2. Interviews with Members of the IETF Community	17
5.1.3. Participant Observation in Working Groups	17
5.2. Data Analysis Strategies	18
5.2.1. Identifying Qualities of Technical Concepts That Relate to Human Rights	18
5.2.2. Relating Human Rights to Technical Concepts	20
5.2.3. Mapping Cases of Protocols, Implementations, and Networking Paradigms That Adversely Impact Human Rights or Are Enablers Thereof	21
6. Model for Developing Human Rights Protocol Considerations	40
6.1. Human Rights Threats	40
6.2. Guidelines for Human Rights Considerations	42
6.2.1. Connectivity	43
6.2.2. Privacy	43
6.2.3. Content Agnosticism	44
6.2.4. Security	45
6.2.5. Internationalization	46
6.2.6. Censorship Resistance	47
6.2.7. Open Standards	48
6.2.8. Heterogeneity Support	50
6.2.9. Anonymity	51
6.2.10. Pseudonymity	51
6.2.11. Accessibility	53
6.2.12. Localization	53
6.2.13. Decentralization	54
6.2.14. Reliability	55
6.2.15. Confidentiality	56
6.2.16. Integrity	58
6.2.17. Authenticity	59
6.2.18. Adaptability	60
6.2.19. Outcome Transparency	61
7. Security Considerations	61
8. IANA Considerations	61
9. Research Group Information	62
10. Informative References	62
Acknowledgements	80
Authors' Addresses	81

1. Introduction

"There's a freedom about the Internet: As long as we accept the rules of sending packets around, we can send packets containing anything to anywhere." [Berners-Lee]

"The Internet isn't value-neutral, and neither is the IETF."
[RFC3935]

The ever-growing interconnectedness of the Internet and society increases the impact of the Internet on the lives of individuals. Because of this, the design and development of the Internet infrastructure also have a growing impact on society. This has led to a broad recognition that human rights [UDHR] [ICCPR] [ICESCR] have a role in the development and management of the Internet [UNGA2013] [NETmundial]. It has also been argued that the Internet should be strengthened as an enabling environment for human rights [Brown].

This document aims to (1) expose the relationship between protocols and human rights, (2) propose possible guidelines to protect the Internet as an enabling environment for human rights in future protocol development, in a manner similar to the work done for privacy considerations [RFC6973], and (3) increase the awareness, in both the human rights community and the technical community, of the importance of the technical workings of the Internet and its impact on human rights.

Document authors who want to apply this work to their own can go directly to Section 6 of this document.

Open, secure, and reliable connectivity is necessary (although not sufficient) to exercise human rights such as freedom of expression and freedom of association [FOC], as defined in the Universal Declaration of Human Rights [UDHR]. The purpose of the Internet is to be a global network of networks that provides unfettered connectivity to all users, and for any content [RFC1958]. This objective of stimulating global connectivity contributes to the Internet's role as an enabler of human rights. The Internet has given people a platform to exchange opinions and gather information; it has enabled people of different backgrounds and genders to participate in the public debate; it has also allowed people to congregate and organize. Next to that, the strong commitment to security [RFC1984] [RFC3365] and privacy [RFC6973] [RFC7258] in the Internet's architectural design contributes to the strengthening of the Internet as an enabling environment for human rights. One could even argue that the Internet is not only an enabler of human rights but that human rights lie at the base of, and are ingrained in, the architecture of the networks that make up the Internet. Internet

connectivity increases the capacity for individuals to exercise their rights; the core of the Internet -- its architectural design -- is therefore closely intertwined with the human rights framework [CathFloridi]. The quintessential link between the Internet's infrastructure and human rights has been argued by many. [Bless1], for instance, argues that "to a certain extent, the Internet and its protocols have already facilitated the realization of human rights, e.g., the freedom of assembly and expression. In contrast, measures of censorship and pervasive surveillance violate fundamental human rights." [DeNardis15] argues that "Since the first hints of Internet commercialization and internationalization, the IETF has supported strong security in protocol design and has sometimes served as a force resisting protocol-enabled surveillance features." By doing so, the IETF enabled the manifestation of the right to privacy, through the Internet's infrastructure. Additionally, access to freely available information gives people access to knowledge that enables them to help satisfy other human rights; as such, the Internet increasingly becomes a precondition for human rights rather than a supplement.

Human rights can be in conflict with each other, such as the right to freedom of expression and the right to privacy. In such cases, the different affected rights need to be balanced. To do this, it is crucial that the impacts on rights are clearly documented in order to mitigate potential harm. This research aims to ultimately contribute to making that process tangible and practical for protocol developers. Technology can never be fully equated with a human right. Whereas a specific technology might be a strong enabler of a specific human right, it might have an adverse impact on another human right. In this case, decisions on design and deployment need to take this into account.

The open nature of the initial technical design and its open standards, as well as developments like open source, fostered freedom of communication. What emerged was a network of networks that could enable everyone to connect and to exchange data, information, and code. For many, enabling such connections became a core value. However, as the scale and the commercialization of the Internet grew, topics like access, rights, and connectivity have been forced to compete with other values. Therefore, important characteristics of the Internet that enable human rights might be degraded if they're not properly defined, described, and protected as such. Conversely, not protecting characteristics that enable human rights could also result in (partial) loss of functionality and connectivity, along with other inherent parts of the Internet's architecture of networks. New protocols, particularly those that upgrade the core infrastructure of the network, should be designed to continue to enable fundamental human rights.

The IETF has produced guidelines and procedures to ensure and galvanize the privacy of individuals and security of the network in protocol development. This document aims to explore the possibility of developing similar procedures for guidelines for human rights considerations to ensure that protocols developed in the IETF do not have an adverse impact on the realization of human rights on the Internet. By carefully considering the answers to the questions posed in Section 6 of this document, document authors should be (1) able to produce a comprehensive analysis that can serve as the basis for discussion on whether the protocol adequately protects against specific human rights threats and (2) potentially stimulated to think about alternative design choices.

This document was developed within the framework of the Human Rights Protocol Considerations (HRPC) Research Group, based on discussions on the HRPC mailing list (Section 9); this document was also extensively discussed during HRPC sessions. This document has received eleven in-depth reviews on the mailing list, and it received many comments from inside and outside the IRTF and IETF communities.

2. Vocabulary Used

In the discussion of human rights and Internet architecture, concepts developed in computer science, networking, law, policy-making, and advocacy are coming together [Dutton] [Kaye] [Franklin] [RFC1958]. The same concepts might have a very different meaning and implications in other areas of expertise. In order to foster a constructive interdisciplinary debate and minimize differences in interpretation, the following glossary is provided. It builds as much as possible on existing definitions; when definitions were not available in IETF documents, definitions were taken from other Standards Development Organizations (SDOs) or academic literature.

Accessibility: "Full Internet Connectivity", as described in [RFC4084], to provide unfettered access to the Internet.

The design of protocols, services, or implementations that provide an enabling environment for people with disabilities.

The ability to receive information available on the Internet.

Anonymity: The condition of an identity being unknown or concealed [RFC4949].

Anonymous: A state of an individual in which an observer or attacker cannot identify the individual within a set of other individuals (the anonymity set) [RFC6973].

Authenticity: The property of being genuine and able to be verified and be trusted [RFC4949].

Blocking: The practice of preventing access to resources in the aggregate [RFC7754]. Both blocking and filtering can be implemented at the level of "services" (web hosting or video streaming, for example) or at the level of particular "content" [RFC7754].

Censorship: Technical mechanisms, including both blocking and filtering, that certain political or private actors around the world use to block or degrade Internet traffic. For further details on the various elements of Internet censorship, see [Hall].

Censorship resistance: Methods and measures to mitigate Internet censorship.

Confidentiality: The property that data is not disclosed to system entities unless they have been authorized to know the data [RFC4949].

Connectivity: The extent to which a device or network is able to reach other devices or networks to exchange data. The Internet is the tool for providing global connectivity [RFC1958]. Different types of connectivity are further specified in [RFC4084].

The end-to-end principle, interoperability, distributed architecture, resilience, reliability, and robustness in combination constitute the enabling factors that result in connectivity to, and on, the Internet.

Content agnosticism: Treating network traffic identically regardless of content.

Decentralized: Implementation or deployment of standards, protocols, or systems without one single point of control.

End-to-end principle: The principle that application-specific functions should not be embedded into the network and thus stay at the endpoints. In many cases, especially when dealing with failures, the right decisions can only be made with the corresponding application-specific knowledge, which is available at endpoints not in the network.

The end-to-end principle is one of the key architectural guidelines of the Internet. The argument in favor of the end-to-end approach to system design is laid out in the

fundamental papers by Saltzer, Reed, and Clark [Saltzer] [Clark]. In these papers, the authors argue in favor of radical simplification: system designers should only build the essential and shared functions into the network, as most functions can only be implemented at network endpoints. Building features into the network for the benefit of certain applications will come at the expense of others. As such, in general system designers should attempt to steer clear of building anything into the network that is not a bare necessity for its functioning. Following the end-to-end principle is crucial for innovation, as it makes innovation at the edges possible without having to make changes to the network, and it protects the robustness of the network. [RFC2775] further elaborates on various aspects of end-to-end connectivity.

Federation: The possibility of connecting autonomous and possibly centralized systems into a single system without a central authority.

Filtering: The practice of preventing access to specific resources within an aggregate [RFC7754].

Heterogeneity: "The Internet is characterized by heterogeneity on many levels: devices and nodes, router scheduling algorithms and queue management mechanisms, routing protocols, levels of multiplexing, protocol versions and implementations, underlying link layers (e.g., point-to-point, multi-access links, wireless, FDDI, etc.), in the traffic mix and in the levels of congestion at different times and places. Moreover, as the Internet is composed of autonomous organizations and internet service providers, each with their own separate policy concerns, there is a large heterogeneity of administrative domains and pricing structures." [FIArch]

As a result, per [FIArch], the heterogeneity principle proposed in [RFC1958] needs to be supported by design.

Human rights: Principles and norms that are indivisible, interrelated, unalienable, universal, and mutually reinforcing. Human rights have been codified in national and international bodies of law. The Universal Declaration of Human Rights [UDHR] is the most well-known document in the history of human rights. The aspirations from [UDHR] were later codified into treaties such as the International Covenant on Civil and Political Rights [ICCPR] and the International Covenant on Economic, Social and Cultural Rights [ICESCR], after which signatory countries were

obliged to reflect them in their national bodies of law. There is also a broad recognition that not only states have obligations vis-a-vis human rights, but non-state actors do as well.

Integrity: The property that data has not been changed, destroyed, or lost in an unauthorized or accidental manner [RFC4949].

Internationalization (i18n): The practice of making protocols, standards, and implementations usable in different languages and scripts (see Section 6.2.12 ("Localization")).

"In the IETF, 'internationalization' means to add or improve the handling of non-ASCII text in a protocol" [RFC6365].

A different perspective, more appropriate to protocols that are designed for global use from the beginning, is the definition used by the World Wide Web Consortium (W3C) [W3Ci18nDef]:

"Internationalization is the design and development of a product, application or document content that enables easy localization for target audiences that vary in culture, region, or language."

Many protocols that handle text only handle one charset (US-ASCII), or they leave the question of encoding up to local guesswork (which leads, of course, to interoperability problems) [RFC3536]. If multiple charsets are permitted, they must be explicitly identified [RFC2277]. Adding non-ASCII text to a protocol allows the protocol to handle more scripts, hopefully all scripts in use in the world. In today's world, that is normally best accomplished by allowing Unicode encoded in UTF-8 only, thereby shifting conversion issues away from ad hoc choices.

Interoperable: A property of a documented standard or protocol that allows different independent implementations to work with each other without any restriction on functionality.

Localization (l10n): The practice of translating an implementation to make it functional in a specific language or for users in a specific locale (see Section 6.2.5 ("Internationalization")).

(cf. [RFC6365]): The process of adapting an internationalized application platform or application to a specific cultural environment. In localization, the same semantics are preserved while the syntax may be changed [FRAMEWORK].

Localization is the act of tailoring an application for a different language, script, or culture. Some internationalized applications can handle a wide variety of languages. Typical users only understand a small number of languages, so the program

must be tailored to interact with users in just the languages they know. The major work of localization is translating the user interface and documentation. Localization involves not only changing the language interaction but also other relevant changes, such as display of numbers, dates, currency, and so on. The better internationalized an application is, the easier it is to localize it for a particular language and character-encoding scheme.

Open standards: Conform with [RFC2026], which states the following: "Various national and international standards bodies, such as ANSI, ISO, IEEE, and ITU-T, develop a variety of protocol and service specifications that are similar to Technical Specifications defined here. National and international groups also publish 'implementors' agreements' that are analogous to Applicability Statements, capturing a body of implementation-specific detail concerned with the practical application of their standards. All of these are considered to be 'open external standards' for the purposes of the Internet Standards Process."

Openness: Absence of centralized points of control -- "a feature that is assumed to make it easy for new users to join and new uses to unfold" [Brown].

Permissionless innovation: The freedom and ability to freely create and deploy new protocols on top of the communications constructs that currently exist.

Privacy: The right of an entity (normally a person), acting on its own behalf, to determine the degree to which it will interact with its environment, including the degree to which the entity is willing to share its personal information with others [RFC4949].

The right of individuals to control or influence what information related to them may be collected and stored, and by whom and to whom that information may be disclosed.

Privacy is a broad concept relating to the protection of individual or group autonomy and the relationship between an individual or group and society, including government, companies, and private individuals. It is often summarized as "the right to be left alone", but it encompasses a wide range of rights, including protections from intrusions into family and home life, control of sexual and reproductive rights, and communications secrecy. It is commonly recognized as a core right that underpins human dignity and other values such as freedom of association and freedom of speech.

The right to privacy is also recognized in nearly every national constitution and in most international human rights treaties. It has been adjudicated upon by both international and regional bodies. The right to privacy is also legally protected at the national level through provisions in civil and/or criminal codes.

Reliability: Ensures that a protocol will execute its function consistently as described and function without unexpected results. A system that is reliable degenerates gracefully and will have a documented way to announce degradation. It also has mechanisms to recover from failure gracefully and, if applicable, allow for partial healing [dict].

Resilience: The maintaining of dependability and performance in the face of unanticipated changes and circumstances [Meyer].

Robustness: The resistance of protocols and their implementations to errors, and resistance to involuntary, legal, or malicious attempts to disrupt their modes of operation [RFC760] [RFC791] [RFC793] [RFC1122]. Or, framed more positively, a system can provide functionality consistently and without errors despite involuntary, legal, or malicious attempts to disrupt its mode of operation.

Scalability: The ability to handle increased or decreased system parameters (number of end systems, users, data flows, routing entries, etc.) predictably within defined expectations. There should be a clear definition of its scope and applicability. The limits of a system's scalability should be defined. Growth or shrinkage of these parameters is typically considered by orders of magnitude.

Strong encryption / cryptography: Used to describe a cryptographic algorithm that would require a large amount of computational power to defeat it [RFC4949]. In the modern usage of the definition of "strong encryption", this refers to an amount of computing power currently not available, not even to major state-level actors.

Transparency: In this context, linked to the comprehensibility of a protocol in relation to the choices it makes for users, protocol developers, and implementers, and to its outcome.

Outcome transparency is linked to the comprehensibility of the effects of a protocol in relation to the choices it makes for users, protocol developers, and implementers, including the comprehensibility of possible unintended consequences of protocol choices (e.g., lack of authenticity may lead to lack of integrity and negative externalities).

3. Research Questions

The Human Rights Protocol Considerations (HRPC) Research Group in the Internet Research Task Force (IRTF) embarked on its mission to answer the following two questions, which are also the main two questions that this document seeks to answer:

1. How can Internet protocols and standards impact human rights, by either enabling them or creating a restrictive environment?
2. Can guidelines be developed to improve informed and transparent decision-making about the potential impact of protocols on human rights?

4. Literature and Discussion Review

Protocols and standards are regularly seen as merely performing technical functions. However, these protocols and standards do not exist outside of their technical context, nor do they exist outside of their political, historical, economic, legal, or cultural context. This is best exemplified by the way in which some Internet processes and protocols have become part and parcel of political processes and public policies: one only has to look at the IANA transition, [RFC7258] ("Pervasive Monitoring Is an Attack"), or global innovation policy, for concrete examples [DeNardis15]. According to [Abbate], "protocols are politics by other means." This statement would probably not garner IETF consensus, but it nonetheless reveals that protocols are based on decision-making, most often by humans. In this process, the values and ideas about the role that a particular technology should perform in society are embedded into the design. Often, these design decisions are partly "purely technical" and partly inspired by a certain world view of how technology should function that is inspired by personal, corporate, and political views. Within the community of IETF participants, there is a strong desire to solve technical problems and to minimize engagement with political processes and non-protocol-related political issues.

Since the late 1990s, a burgeoning group of academics and practitioners researched questions surrounding the societal impact of protocols, as well as the politics of protocols. These studies vary in focus and scope: some focus on specific standards [Davidson-etal] [Musiani]; others look into the political, legal, commercial, or social impact of protocols [BrownMarsden] [Lessig] [Mueller]; and yet others look at how the engineers' personal set of values get translated into technology [Abbate] [CathFloridi] [DeNardis15] [WynsbergheMoura].

Commercial and political influences on the management of the Internet's infrastructure are well documented in the academic literature and will thus not be discussed here; see [Benkler], [Brown-etal], [DeNardis15], [Lessig], [Mueller], and [Zittrain]. It is sufficient to say that the IETF community consistently tries to push back against the standardization of surveillance and certain other issues that negatively influence an end user's experience of, and trust in, the Internet [DeNardis14]. The role that human rights play in engineering, infrastructure maintenance, and protocol design is much less clear.

It is very important to understand how protocols and standards impact human rights, in particular because SDOs are increasingly becoming venues where social values (like human rights) are discussed, although often from a technological point of view. These SDOs are becoming a new focal point for discussions about "values by design" and the role of technical engineers in protecting or enabling human rights [Brown-etal] [Clark-etal] [DeNardis14] [CathFloridi] [Lessig] [Rachovitsa].

In the academic literature, five clear positions can be discerned in relation to the role of human rights in protocol design and how to account for these human rights in protocol development: Clark et al. [Clark-etal] argue that there is a need to design "for variation in outcome -- so that the outcome can be different in different places, and the tussle takes place within the design (...)" [as] "Rigid designs will be broken; designs that permit variation will flex under pressure and survive." They hold that human rights should not be hard-coded into protocols for three reasons: First, the rights in the UDHR are not absolute. Second, technology is not the only tool in the tussle over human rights. And last but not least, it is dangerous to make promises that can't be kept. The open nature of the Internet will never, they argue, be enough to fully protect individuals' human rights.

Conversely, Brown et al. [Brown-etal] state that "some key, universal values -- of which the UDHR is the most legitimate expression -- should be baked into the architecture at design time." They argue that design choices have offline consequences and are able to shape the power positions of groups or individuals in society. As such, the individuals making these technical decisions have a moral obligation to take into account the impact of their decisions on society and, by extension, human rights. Brown et al. recognize that values and the implementation of human rights vary across the globe. Yet they argue that all members of the United Nations have found "common agreement on the values proclaimed in the Universal Declaration of Human Rights. In looking for the most legitimate set

of global values to embed in the future Internet architectures, the UDHR has the democratic assent of a significant fraction of the planet's population, through their elected representatives."

The main disagreement between these two academic positions lies mostly in the question of whether (1) a particular value system should be embedded into the Internet's architectures or (2) the architectures need to account for a varying set of values.

A third position, which is similar to that of Brown et al., is taken by [Broeders], in which Broeders argues that "we must find ways to continue guaranteeing the overall integrity and functionality of the public core of the Internet." He argues that the best way to do this is by declaring the backbone of the Internet -- which includes the TCP/IP protocol suite, numerous standards, the Domain Name System (DNS), and routing protocols -- a common public good. This is a different approach than those of [Clark-et al] and [Brown-et al] because Broeders does not suggest that social values should (or should not) be explicitly coded into the Internet, but rather that the existing infrastructure should be seen as an entity of public value.

Bless and Orwat [Bless2] represent a fourth position. They argue that it is too early to make any definitive claims but that there is a need for more careful analysis of the impact of protocol design choices on human rights. They also argue that it is important to search for solutions that "create awareness in the technical community about impact of design choices on social values" and "work towards a methodology for co-design of technical and institutional systems."

Berners-Lee and Halpin [BernersLeeHalpin] represent a fifth position. They argue that the Internet could lead to even newer capacities, and these capacities may over time be viewed as new kinds of rights. For example, Internet access may be viewed as a human right in and of itself if it is taken to be a precondition for other rights, even if it could not have been predicted at the time that the UDHR was written (after the end of World War II).

It is important to contextualize the technical discussion with the academic discussions on this issue. The academic discussions are also important to document, as they inform the position of the authors of this document. The research group's position is that hard-coding human rights into protocols is complicated and changes with the context. At this point, it is difficult to say whether or not hard-coding human rights into protocols is wise or feasible. Additionally, there are many human rights, but not all are relevant for information and communications technologies (ICTs). A partial

catalog (with references to sources) of human rights related to ICTs can be found in [Hill2014]. It is, however, important to make conscious and explicit design decisions that take into account the human rights protocol considerations guidelines developed below. This will contribute to the understanding of the impact that protocols can have on human rights, for both developers and users. In addition, it contributes to (1) the careful consideration of the impact that a specific protocol might have on human rights and (2) the dissemination of the practice of documenting protocol design decisions related to human rights.

Pursuant to the principle of constant change, because the function and scope of the Internet evolve, so does the role of the IETF in developing standards. Internet Standards are adopted based on a series of criteria, including high technical quality, support by community consensus, and their overall benefit to the Internet. The latter calls for an assessment of the interests of all affected parties and the specifications' impact on the Internet's users. In this respect, the effective exercise of the human rights of the Internet users is a relevant consideration that needs to be appreciated in the standardization process insofar as it is directly linked to the reliability and core values of the Internet [RFC1958] [RFC2775] [RFC3439] [RFC3724].

This document details the steps taken in the research into human rights protocol considerations by the HRPC Research Group to clarify the relationship between technical concepts used in the IETF and human rights. This document sets out some preliminary steps and considerations for engineers to take into account when developing standards and protocols.

5. Methodology

Mapping the relationship between human rights, protocols, and architectures is a new research challenge that requires a good amount of interdisciplinary and cross-organizational cooperation to develop a consistent methodology.

The methodological choices made in this document are based on the political-science-based method of discourse analysis and ethnographic research methods [Cath]. This work departs from the assumption that language reflects the understanding of concepts. Or, as [Jabri] holds, policy documents are "social relations represented in texts where the language contained within these texts is used to construct meaning and representation." This process happens in society [Denzin] and manifests itself in institutions and organizations [King], exposed using the ethnographic methods of semi-structured interviews and participant observation. Or, in non-academic

language, the way the language in IETF/IRTF documents describes and approaches the issues they are trying to address is an indication of the underlying social assumptions and relationships of the engineers to their engineering. By reading and analyzing these documents, as well as interviewing engineers and participating in the IETF/IRTF working groups, it is possible to distill the relationship between human rights, protocols, and the Internet's infrastructure as it pertains to the work of the IETF.

The discourse analysis was operationalized using qualitative and quantitative means. The first step taken by the authors and contributors was reading RFCs and other official IETF documents. The second step was the use of a Python-based analyzer, using the "Bigbang" tool, adapted by Nick Doty [Doty], to scan for the concepts that were identified as important architectural principles (distilled on the initial reading and supplemented by the interviews and participant observation). Such a quantitative method is very precise and speeds up the research process [Ritchie]. But this tool is unable to understand "latent meaning" [Denzin]. In order to mitigate these issues of automated word-frequency-based approaches and to get a sense of the "thick meaning" [Geertz] of the data, a second qualitative analysis of the data set was performed. These various rounds of discourse analysis were used to inform the interviews and further data analysis. As such, the initial rounds of quantitative discourse analysis were used to inform the second rounds of qualitative analysis. The results from the qualitative interviews were again used to feed new concepts into the quantitative discourse analysis. As such, the two methods continued to support and enrich each other.

The ethnographic methods of the data collection and processing allowed the research group to acquire the data necessary to "provide a holistic understanding of research participants' views and actions" [Denzin] that highlighted ongoing issues and case studies where protocols impact human rights. The interview participants were selected through purposive sampling [Babbie], as the research group was interested in getting a wide variety of opinions on the role of human rights in guiding protocol development. This sampling method also ensured that individuals with extensive experience working at the IETF in various roles were targeted. The interviewees included individuals in leadership positions (Working Group (WG) chairs, Area Directors (ADs)), "regular participants", and individuals working for specific entities (corporate, civil society, political, academic) and represented various backgrounds, nationalities, and genders.

5.1. Data Sources

In order to map the potential relationship between human rights and protocols, the HRPC Research Group gathered data from three specific sources:

5.1.1. Discourse Analysis of RFCs

To start addressing the issue, a mapping exercise analyzing Internet infrastructure and protocol features vis-a-vis their possible impact on human rights was undertaken. Therefore, research on (1) the language used in current and historic RFCs and (2) information gathered from mailing-list discussions was undertaken to expose core architectural principles, language, and deliberations on the human rights of those affected by the network.

5.1.2. Interviews with Members of the IETF Community

Over 30 interviews with the current and past members of the Internet Architecture Board (IAB), current and past members of the Internet Engineering Steering Group (IESG), chairs of selected working groups, and RFC authors were done at the IETF 92 meeting in Dallas in March 2015 to get an insider's understanding of how they view the relationship (if any) between human rights and protocols, and how this relationship plays out in their work. Several of the participants opted to remain anonymous. If you are interested in this data set, please contact the authors of this document.

5.1.3. Participant Observation in Working Groups

By participating in various working groups, in person at IETF meetings, and on mailing lists, information about the IETF's day-to-day workings was gathered, from which general themes, technical concepts, and use cases about human rights and protocols were extracted. This process started at the IETF 91 meeting in Honolulu and continues today.

5.2. Data Analysis Strategies

The data above was processed using three consecutive strategies: mapping protocols related to human rights, extracting concepts from these protocols, and creation of a common glossary (detailed under Section 2). Before going over these strategies, some elaboration on the process of identifying technical concepts as they relate to human rights is needed:

5.2.1. Identifying Qualities of Technical Concepts That Relate to Human Rights

5.2.1.1. Mapping Protocols and Standards to Human Rights

By combining data from the three data sources named above, an extensive list of protocols and standards that potentially enable the Internet as a tool for freedom of expression and association was created. In order to determine the enabling (or inhibiting) features, we relied on direct references in the RFCs as related to such impacts, as well as input from the community. Based on this analysis, a list of RFCs that describe standards and protocols that are potentially closely related to human rights was compiled.

5.2.1.2. Extracting Concepts from Selected RFCs

The first step was to identify the protocols and standards that are related to human rights and to create an environment that enables human rights. For that, we needed to focus on specific technical concepts that underlie these protocols and standards. Based on this list, a number of technical concepts that appeared frequently were extracted and used to create a second list of technical terms that, when combined and applied in different circumstances, create an enabling environment for exercising human rights on the Internet.

5.2.1.3. Building a Common Vocabulary of Technical Concepts That Impact Human Rights

While interviewing experts, investigating RFCs, and compiling technical definitions, several concepts of convergence and divergence were identified. To ensure that the discussion was based on a common understanding of terms and vocabulary, a list of definitions was created. The definitions are based on the wording found in various IETF documents; if the definitions were not available therein, definitions were taken from other SDOs or academic literature, as indicated in Section 2.

5.2.2. Relating Human Rights to Technical Concepts

The technical concepts listed in the steps above have been grouped according to their impact on specific rights, as mentioned in the interviews done at IETF 92 as well as the study of literature (see Section 4 ("Literature and Discussion Review") above).

This analysis aims to assist protocol developers in better understanding the roles that specific technical concepts have with regard to their contribution to an enabling environment for people to exercise their human rights.

This analysis does not claim to be a complete or exhaustive mapping of all possible ways in which protocols could potentially impact human rights, but it presents a mapping of initial concepts based on interviews and on discussion and review of the literature.

Technical Concepts	Rights Potentially Impacted
Connectivity Privacy Security Content agnosticism Internationalization Censorship resistance Open standards Heterogeneity support	Right to freedom of expression
Anonymity Privacy Pseudonymity Accessibility	Right to non-discrimination
Content agnosticism Security	Right to equal protection
Accessibility Internationalization Censorship resistance Connectivity	Right to political participation
Open standards Localization Internationalization Censorship resistance Accessibility	Right to participate in cultural life, arts, and science, and Right to education

Connectivity Decentralization Censorship resistance Pseudonymity Anonymity Security	Right to freedom of assembly and association
Reliability Confidentiality Integrity Authenticity Anonymity	Right to security

Figure 2: Relationship between Specific Technical Concepts with Regard to Their Contribution to an Enabling Environment for People to Exercise Their Human Rights

5.2.3. Mapping Cases of Protocols, Implementations, and Networking Paradigms That Adversely Impact Human Rights or Are Enablers Thereof

Given the information above, the following list of cases of protocols, implementations, and networking paradigms that either adversely impact or enable human rights was formed.

It is important to note that the assessment here is not a general judgment on these protocols, nor is it an exhaustive listing of all the potential negative or positive impacts on human rights that these protocols might have. When these protocols were conceived, there were many criteria to take into account. For instance, relying on a centralized service can be bad for freedom of speech (it creates one more control point, where censorship could be applied), but it may be a necessity if the endpoints are not connected and reachable permanently. So, when we say "protocol X has feature Y, which may endanger freedom of speech," it does not mean that protocol X is bad, much less that its authors were evil. The goal here is to show, with actual examples, that the design of protocols has practical consequences for some human rights and that these consequences have to be considered in the design phase.

5.2.3.1. IPv4

The Internet Protocol version 4 (IPv4), also known as "Layer 3" of the Internet and specified with a common encapsulation and protocol header, is defined in [RFC791]. The evolution of Internet communications led to continued development in this area, "encapsulated" in the development of version 6 (IPv6) of the protocol [RFC8200]. In spite of this updated protocol, we find that 23 years after the specification of IPv6 the older IPv4 standard continues to account for a sizable majority of Internet traffic. Most of the issues discussed here (Network Address Translators (NATs) are a major exception; see Section 5.2.3.1.2 ("Address Translation and Mobility")) are valid for IPv4 as well as IPv6.

The Internet was designed as a platform for free and open communication, most notably encoded in the end-to-end principle, and that philosophy is also present in the technical implementation of IP [RFC3724]. While the protocol was designed to exist in an environment where intelligence is at the end hosts, it has proven to provide sufficient information that a more intelligent network core can make policy decisions and enforce policy-based traffic shaping, thereby restricting the communications of end hosts. These capabilities for network control and for limitations on freedom of expression by end hosts can be traced back to the design of IPv4, helping us to understand which technical protocol decisions have led to harm to this human right. A feature that can harm freedom of expression as well as the right to privacy through misuse of IP is the exploitation of the public visibility of the host pairs for all communications and the corresponding ability to differentiate and block traffic as a result of that metadata.

5.2.3.1.1. Network Visibility of Source and Destination

The IPv4 protocol header contains fixed location fields for both the source IP address and destination IP address [RFC791]. These addresses identify both the host sending and the host receiving each message; they also allow the core network to understand who is talking to whom and to practically limit communication selectively between pairs of hosts. Blocking of communication based on the pair of source and destination is one of the most common limitations on the ability for people to communicate today [CAIDA] and can be seen as a restriction of the ability for people to assemble or to consensually express themselves.

Inclusion of an Internet-wide identified source in the IP header is not the only possible design, especially since the protocol is most commonly implemented over Ethernet networks exposing only link-local identifiers [RFC894].

A variety of alternative designs do exist, such as the Accountable and Private Internet Protocol [APIP] and High-speed Onion Routing at the Network Layer (HORNET) [HORNET] as well as source routing. The latter would allow the sender to choose a predefined (safe) route and spoofing of the source IP address, which are technically supported by IPv4, but neither are considered good practice on the Internet [Farrow]. While projects like [TorProject] provide an alternative implementation of anonymity in connections, they have been developed in spite of the IPv4 protocol design.

5.2.3.1.2. Address Translation and Mobility

A major structural shift in the Internet that undermined the protocol design of IPv4, and significantly reduced the freedom of end users to communicate and assemble, was the introduction of network address translation [RFC3022]. Network address translation is a process whereby organizations and autonomous systems connect two networks by translating the IPv4 source and destination addresses between them. This process puts the router performing the translation in a privileged position, where it is predetermined which subset of communications will be translated.

This process of translation has widespread adoption despite promoting a process that goes against the stated end-to-end process of the underlying protocol [NATusage]. In contrast, the proposed mechanism to provide support for mobility and forwarding to clients that may move -- encoded instead as an option in IP [RFC5944] -- has failed to gain traction. In this situation, the compromise made in the design of the protocol resulted in a technology that is not coherent with the end-to-end principles and thus creates an extra possible hurdle for freedom of expression in its design, even though a viable alternative exists. There is a particular problem surrounding NATs and Virtual Private Networks (VPNs) (as well as other connections used for privacy purposes), as NATs sometimes cause VPNs not to work.

5.2.3.2. DNS

The Domain Name System (DNS) [RFC1035] provides service discovery capabilities and provides a mechanism to associate human-readable names with services. The DNS is organized around a set of independently operated "root servers" run by organizations that function in line with ICANN's policy by answering queries for which organizations have been delegated to manage registration under each Top-Level Domain (TLD). The DNS is organized as a rooted tree, and this brings up political and social concerns over control. TLDs are maintained and determined by ICANN. These namespaces encompass several classes of services. The initial namespaces, including ".com" and ".net", provide common spaces for expression of ideas,

though their policies are enacted through US-based companies. Other namespaces are delegated to specific nationalities and may impose limits designed to focus speech in those forums, to both (1) promote speech from that nationality and (2) comply with local limits on expression and social norms. Finally, the system has recently been expanded with additional generic and sponsored namespaces -- for instance, ".travel" and ".ninja" -- that are operated by a range of organizations that may independently determine their registration policies. This new development has both positive and negative implications in terms of enabling human rights. Some individuals argue that it undermines the right to freedom of expression because some of these new generic TLDs have restricted policies on registration and particular rules on hate speech content. Others argue that precisely these properties are positive because they enable certain (mostly minority) communities to build safer spaces for association, thereby enabling their right to freedom of association. An often-mentioned example is an application like .gay [CoE].

As discussed in [RFC7626], DNS has significant privacy issues. Most notable is the lack of encryption to limit the visibility of requests for domain resolution from intermediary parties, and a limited deployment of DNSSEC to provide authentication, allowing the client to know that they received a correct, "authoritative" answer to a query. In response to the privacy issues, the IETF DNS Private Exchange (DPRIVE) Working Group is developing mechanisms to provide confidentiality to DNS transactions, to address concerns surrounding pervasive monitoring [RFC7258].

Authentication through DNSSEC creates a validation path for records. This authentication protects against forged or manipulated DNS data. As such, DNSSEC protects directory lookups and makes it harder to hijack a session. This is important because interference with the operation of the DNS is currently becoming one of the central mechanisms used to block access to websites. This interference limits both the freedom of expression of the publisher to offer their content and the freedom of assembly for clients to congregate in a shared virtual space. Even though DNSSEC doesn't prevent censorship, it makes it clear that the returned information is not the information that was requested; this contributes to the right to security and increases trust in the network. It is, however, important to note that DNSSEC is currently not widely supported or deployed by domain name registrars, making it difficult to authenticate and use correctly.

5.2.3.2.1. Removal of Records

There have been a number of cases where the records for a domain are removed from the name system due to political events. Examples of this removal include the "seizure" of wikileaks [BBC-wikileaks] and the names of illegally operating gambling operations by the United States Immigration and Customs Enforcement (ICE) unit. In the first case, a US court ordered the registrar to take down the domain. In the second, ICE compelled the US-based registry in charge of the .com TLD to hand ownership of those domains over to the US government. The same technique has been used in Libya to remove sites in violation of "our Country's Law and Morality (which) do not allow any kind of pornography or its promotion." [techyum]

At a protocol level, there is no technical auditing for name ownership, as in alternate systems like Namecoin [Namecoin]. As a result, there is no ability for users to differentiate seizure from the legitimate transfer of name ownership, which is purely a policy decision made by registrars. While DNSSEC addresses the network distortion events described below, it does not tackle this problem.

(Although we mention alternative techniques, this is not a comparison of DNS with Namecoin: the latter has its own problems and limitations. The idea here is to show that there are several possible choices, and they have consequences for human rights.)

5.2.3.2.2. Distortion of Records

The most common mechanism by which the DNS is abused to limit freedom of expression is through manipulation of protocol messages by the network. One form occurs at an organizational level, where client computers are instructed to use a local DNS resolver controlled by the organization. The DNS resolver will then selectively distort responses rather than request the authoritative lookup from the upstream system. The second form occurs through the use of Deep Packet Inspection (DPI), where all DNS protocol messages are inspected by the network and objectionable content is distorted, as can be observed in Chinese networks.

A notable instance of distortion occurred in Greece [Ververis], where a study found evidence of both (1) DPI to distort DNS replies and (2) more excessive blocking of content than was legally required or requested (also known as "overblocking"). Internet Service Providers (ISPs), obeying a governmental order, prevented clients from resolving the names of domains, thereby prompting this particular blocking of systems there.

At a protocol level, the effectiveness of these attacks is made possible by a lack of authentication in the DNS protocol. DNSSEC provides the ability to determine the authenticity of responses when used, but it is not regularly checked by resolvers. DNSSEC is not effective when the local resolver for a network is complicit in the distortion -- for instance, when the resolver assigned for use by an ISP is the source of injection. Selective distortion of records is also made possible by the predictable structure of DNS messages, which makes it computationally easy for a network device to watch all passing messages even at high speeds, and the lack of encryption, which allows the network to distort only an objectionable subset of protocol messages. Specific distortion mechanisms are discussed further in [Hall].

Users can switch to another resolver -- for instance, a public resolver. The distorter can then try to block or hijack the connection to this resolver. This may start an arms race, with the user switching to secured connections to this alternative resolver [RFC7858] and the distorter then trying to find more sophisticated ways to block or hijack the connection. In some cases, this search for an alternative, non-disrupting resolver may lead to more centralization because many people are switching to a few big commercial public resolvers.

5.2.3.2.3. Injection of Records

Responding incorrectly to requests for name lookups is the most common mechanism that in-network devices use to limit the ability of end users to discover services. A deviation that accomplishes a similar objective and may be seen as different from a "freedom of expression" perspective is the injection of incorrect responses to queries. The most prominent example of this behavior occurs in China, where requests for lookups of sites deemed inappropriate will trigger the network to return a false response, causing the client to ignore the real response when it subsequently arrives [greatfirewall]. Unlike the other network paradigms discussed above, injection does not stifle the ability of a server to announce its name; it instead provides another voice that answers sooner. This is effective because without DNSSEC, the protocol will respond to whichever answer is received first, without listening for subsequent answers.

5.2.3.3. HTTP

The Hypertext Transfer Protocol (HTTP) version 1.1 [RFC7230] [RFC7231] [RFC7232] [RFC7233] [RFC7234] [RFC7235] [RFC7236] [RFC7237] is a request-response application protocol developed throughout the 1990s. HTTP factually contributed to the exponential growth of the

Internet and the interconnection of populations around the world. Its simple design strongly contributed to the fact that HTTP has become the foundation of most modern Internet platforms and communication systems, from websites to chat systems and computer-to-computer applications. In its manifestation in the World Wide Web, HTTP radically revolutionized the course of technological development and the ways people interact with online content and with each other.

However, HTTP is also a fundamentally insecure protocol that doesn't natively provide encryption properties. While the definition of the Secure Sockets Layer (SSL) [RFC6101], and later of Transport Layer Security (TLS) [RFC5246], also happened during the 1990s, the fact that HTTP doesn't mandate the use of such encryption layers by developers and service providers was one of the reasons for a very late adoption of encryption. Only in the middle of the 2000s did we observe big ISPs, such as Google, starting to provide encrypted access to their web services.

The lack of sensitivity and understanding of the critical importance of securing web traffic incentivized certain (offensive) actors to develop, deploy, and utilize interception systems at large and to later launch active injection attacks, in order to swipe large amounts of data and compromise Internet-enabled devices. The commercial availability of systems and tools to perform these types of attacks also led to a number of human rights abuses that have been discovered and reported over the years.

Generally, we can identify traffic interception (Section 5.2.3.3.1) and traffic manipulation (Section 5.2.3.3.2) as the two most problematic attacks that can be performed against applications employing a cleartext HTTP transport layer. That being said, the IETF is taking steady steps to move to the encrypted version of HTTP, HTTP Secure (HTTPS).

While this is commendable, we must not lose track of the fact that different protocols, implementations, configurations, and networking paradigms can intersect such that they (can be used to) adversely impact human rights. For instance, to facilitate surveillance, certain countries will throttle HTTPS connections, forcing users to switch to (unthrottled) HTTP [Aryan-et al].

5.2.3.3.1. Traffic Interception

While we are seeing an increasing trend in the last couple of years to employ SSL/TLS as a secure traffic layer for HTTP-based applications, we are still far from seeing a ubiquitous use of encryption on the World Wide Web. It is important to consider that the adoption of SSL/TLS is also a relatively recent phenomenon.

Email providers such as riseup.net were the first to enable SSL by default. Google did not introduce an option for its Gmail users to navigate with SSL until 2008 [Rideout] and turned TLS on by default later, in 2010 [Schillace]. It took an increasing amount of security breaches and revelations on global surveillance from Edward Snowden before other mail service providers followed suit. For example, Yahoo did not enable SSL/TLS by default on its webmail services until early 2014 [Peterson].

TLS itself has been subject to many attacks and bugs; this situation can be attributed to some fundamental design weaknesses, such as lack of a state machine (which opens a vulnerability for triple handshake attacks) and flaws caused by early US government restrictions on cryptography, leading to cipher-suite downgrade attacks (Logjam attacks). These vulnerabilities are being corrected in TLS 1.3 [Bhargavan] [Adrian].

HTTP upgrading to HTTPS is also vulnerable to having an attacker remove the "s" in any links to HTTPS URIs from a web page transferred in cleartext over HTTP -- an attack called "SSL Stripping" [sslstrip]. Thus, for high-security use of HTTPS, IETF standards such as HTTP Strict Transport Security (HSTS) [RFC6797], certificate pinning [RFC7469], and/or DNS-Based Authentication of Named Entities (DANE) [RFC6698] should be used.

As we learned through Snowden's revelations, intelligence agencies have been intercepting and collecting unencrypted traffic at large for many years. There are documented examples of such mass-surveillance programs with the Government Communications Headquarters's (GCHQ's) Tempora [WP-Tempora] and the National Security Agency's (NSA's) XKeyscore [Greenwald]. Through these programs, the NSA and the GCHQ have been able to swipe large amounts of data, including email and instant messaging communications that have been transported in the clear for years by providers unsuspecting of the pervasiveness and scale of governments' efforts and investment in global mass-surveillance capabilities.

However, similar mass interception of unencrypted HTTP communications is also often employed at the national level by some democratic countries, by exercising control over state-owned ISPs and through the use of commercially available monitoring, collection, and censorship equipment. Over the last few years, a lot of information has come to public attention on the role and scale of a surveillance industry dedicated to developing different types of interception gear, making use of known and unknown weaknesses in existing protocols [RFC7258]. We have several records of such equipment being sold and utilized by some regimes in order to monitor entire segments of a population, especially at times of social and political

distress, uncovering massive human rights abuses. For example, in 2013, the group Telecomix revealed that the Syrian regime was making use of Blue Coat products in order to intercept cleartext traffic as well as to enforce censorship of unwanted content [RSF]. Similarly, in 2011, it was found that the French technology firm Amesys provided the Gadhafi government with equipment able to intercept emails, Facebook traffic, and chat messages at a country-wide level [WSJ]. The use of such systems, especially in the context of the Arab Spring and of civil uprisings against the dictatorships, has caused serious concerns regarding significant human rights abuses in Libya.

5.2.3.3.2. Traffic Manipulation

The lack of a secure transport layer under HTTP connections not only exposes users to interception of the content of their communications but is more and more commonly abused as a vehicle for actively compromising computers and mobile devices. If an HTTP session travels in the clear over the network, any node positioned at any point in the network is able to perform man-in-the-middle attacks; the node can observe, manipulate, and hijack the session and can modify the content of the communication in order to trigger unexpected behavior by the application generating the traffic. For example, in the case of a browser, the attacker would be able to inject malicious code in order to exploit vulnerabilities in the browser or any of its plugins. Similarly, the attacker would be able to intercept, add malware to, and repackage binary software updates that are very commonly downloaded in the clear by applications such as word processors and media players. If the HTTP session were encrypted, the tampering of the content would not be possible, and these network injection attacks would not be successful.

While traffic manipulation attacks have long been known, documented, and prototyped, especially in the context of Wi-Fi and LAN networks, in the last few years we have observed an increasing investment in the production and sale of network injection equipment that is both commercially available and deployed at scale by intelligence agencies.

For example, we learned from some of the documents provided by Edward Snowden to the press that the NSA has constructed a global network injection infrastructure, called "QUANTUM", able to leverage mass surveillance in order to identify targets of interest and subsequently task man-on-the-side attacks to ultimately compromise a selected device. Among other attacks, the NSA makes use of an attack called "QUANTUMINSERT" [Haagsma], which intercepts and hijacks an unencrypted HTTP communication and forces the requesting browser to redirect to a host controlled by the NSA instead of the intended website. Normally, the new destination would be an exploitation

service, referred to in Snowden documents as "FOXACID", which would attempt to execute malicious code in the context of the target's browser. The Guardian reported in 2013 that the NSA has, for example, been using these techniques to target users of the popular anonymity service Tor [Schneier]. The German Norddeutscher Rundfunk (NDR) reported in 2014 that the NSA has also been using its mass-surveillance capabilities to identify Tor users at large [Appelbaum].

Recently, similar capabilities used by Chinese authorities have been reported as well in what has been informally called the "Great Cannon" [Marcak], which raised numerous concerns on the potential curb on human rights and freedom of speech due to the increasingly tighter control of Chinese Internet communications and access to information.

Network injection attacks are also made widely available to state actors around the world through the commercialization of similar, smaller-scale equipment that can be easily acquired and deployed at a country-wide level. Certain companies are known to have network injection gear within their products portfolio [Marquis-Boire]. The technology devised and produced by some of them to perform network traffic manipulation attacks on HTTP communications is even the subject of a patent application in the United States [Googlepatent]. Access to offensive technologies available on the commercial lawful interception market has led to human rights abuses and illegitimate surveillance of journalists, human rights defenders, and political activists in many countries around the world [Collins]. While network injection attacks haven't been the subject of much attention, they do enable even unskilled attackers to perform silent and very resilient compromises, and unencrypted HTTP remains one of the main vehicles.

There is a new version of HTTP, called "HTTP/2" [RFC7540], which aims to be largely backwards compatible while also offering new options such as data compression of HTTP headers, pipelining of requests, and multiplexing multiple requests over a single TCP connection. In addition to decreasing latency to improve page-loading speeds, it also facilitates more efficient use of connectivity in low-bandwidth environments, which in turn enables freedom of expression; the right to assembly; the right to political participation; and the right to participate in cultural life, arts, and science. [RFC7540] does not mandate TLS or any other form of encryption, nor does it support opportunistic encryption even though opportunistic encryption is now addressed in [RFC8164].

5.2.3.4. XMPP

The Extensible Messaging and Presence Protocol (XMPP), specified in [RFC6120], provides a standard for interactive chat messaging and has evolved to encompass interoperable text, voice, and video chat. The protocol is structured as a federated network of servers, similar to email, where users register with a local server that acts on their behalf to cache and relay messages. This protocol design has many advantages, allowing servers to shield clients from denial of service and other forms of retribution for their expression; it is also designed to avoid central entities that could control the ability to communicate or assemble using the protocol.

Nonetheless, there are plenty of aspects of the protocol design of XMPP that shape the ability for users to communicate freely and to assemble via the protocol.

5.2.3.4.1. User Identification

The XMPP specification [RFC6120] dictates that clients are identified with a resource (<node@domain/home> / <node@domain/work>) to distinguish the conversations to specific devices. While the protocol does not specify that the resource must be exposed by the client's server to remote users, in practice this has become the default behavior. In doing so, users can be tracked by remote friends and their servers, who are able to monitor the presence of not just the user but of each individual device the user logs in with. This has proven to be misleading to many users [Pidgin], since many clients only expose user-level rather than device-level presence. Likewise, user invisibility so that communication can occur while users don't notify all buddies and other servers of their availability is not part of the formal protocol and has only been added as an extension within the XML stream rather than enforced by the protocol.

5.2.3.4.2. Surveillance of Communication

XMPP specifies the standard by which communications channels may be encrypted, but it does not provide visibility to clients regarding whether their communications are encrypted on each link. In particular, even when both clients ensure that they have an encrypted connection to their XMPP server to ensure that their local network is unable to read or disrupt the messages they send, the protocol does not provide visibility into the encryption status between the two servers. As such, clients may be subject to selective disruption of communications by an intermediate network that disrupts communications based on keywords found through DPI. While many operators have committed to only establishing encrypted links from

their servers in recognition of this vulnerability, it remains impossible for users to audit this behavior, and encrypted connections are not required by the protocol itself [XMPP-Manifesto].

In particular, Section 13.14 of the XMPP specification [RFC6120] explicitly acknowledges the existence of a downgrade attack where an adversary controlling an intermediate network can force the inter-domain federation between servers to revert to a non-encrypted protocol where selective messages can then be disrupted.

5.2.3.4.3. Group Chat Limitations

Group chat in XMPP is defined as an extension within the XML specification of XMPP (<https://xmpp.org/extensions/xep-0045.html>). However, it is not encoded or required at a protocol level and is not uniformly implemented by clients.

The design of multi-user chat in XMPP suffers from extending a protocol that was not designed with assembly of many users in mind. In particular, in the federated protocol provided by XMPP, multi-user communities are implemented with a distinguished "owner" who is granted control over the participants and structure of the conversation.

Multi-user chat rooms are identified by a name specified on a specific server, so that while the overall protocol may be federated, the ability for users to assemble in a given community is moderated by a single server. That server may block the room and prevent assembly unilaterally, even between two users, neither of whom trust or use that server directly.

5.2.3.5. Peer-to-Peer

Peer-to-Peer (P2P) is a distributed network architecture [RFC5694] in which all the participant nodes can be responsible for the storage and dissemination of information from any other node (see [RFC7574], an IETF standard that discusses a P2P architecture called the "Peer-to-Peer Streaming Peer Protocol" (PPSPP)). A P2P network is a logical overlay that lives on top of the physical network and allows nodes (or "peers") participating in it to establish contact and exchange information directly with each other. The implementation of a P2P network may vary widely: it may be structured or unstructured, and it may implement stronger or weaker cryptographic and anonymity properties. While its most common application has traditionally been file-sharing (and other types of content delivery systems), P2P is a popular architecture for networks and applications that require (or encourage) decentralization. Prime examples include Bitcoin and other proprietary multimedia applications.

In a time of heavily centralized online services, P2P is regularly described as an alternative, more democratic, and resistant option that displaces structures of control over data and communications and delegates all peers to be equally responsible for the functioning, integrity, and security of the data. While in principle P2P remains important to the design and development of future content distribution, messaging, and publishing systems, it poses numerous security and privacy challenges that are mostly delegated to individual developers to recognize, analyze, and solve in each implementation of a given P2P network.

5.2.3.5.1. Network Poisoning

Since content, and sometimes peer lists, are safeguarded and distributed by their members, P2P networks are prone to what are generally defined as "poisoning attacks". Poisoning attacks might be aimed directly at the data that is being distributed, for example, (1) by intentionally corrupting the data, (2) at the index tables used to instruct the peers where to fetch the data, or (3) at routing tables, with an attempt to provide connecting peers with lists of rogue or nonexistent peers, with the intention to effectively cause a denial of service on the network.

5.2.3.5.2. Throttling

P2P traffic (and BitTorrent in particular) represents a significant percentage of global Internet traffic [Sandvine], and it has become increasingly popular for ISPs to perform throttling of customers' lines in order to limit bandwidth usage [torrentfreak1] and, sometimes, probably as an effect of the ongoing conflict between copyright holders and file-sharing communities [wikileaks]. Such throttling undermines the end-to-end principle.

Throttling the P2P traffic makes some uses of P2P networks ineffective; this throttling might be coupled with stricter inspection of users' Internet traffic through DPI techniques, possibly posing additional security and privacy risks.

5.2.3.5.3. Tracking and Identification

One of the fundamental and most problematic issues with traditional P2P networks is a complete lack of anonymization of their users. For example, in the case of BitTorrent, all peers' IP addresses are openly available to the other peers. This has led to ever-increasing tracking of P2P and file-sharing users [ars]. As the geographical location of the user is directly exposed, as could also be his identity, the user might become a target of additional harassment and attacks of a physical or legal nature. For example, it is known that

in Germany law firms have made extensive use of P2P and file-sharing tracking systems in order to identify downloaders and initiate legal actions looking for compensations [torrentfreak2].

It is worth noting that there are some varieties of P2P networks that implement cryptographic practices and that introduce anonymization of their users. Such implementations may be proved to be successful in resisting censorship of content and tracking of network peers. A prime example is Freenet [freenet1], a free software application that is (1) designed to make it significantly more difficult to identify users and content and (2) dedicated to fostering freedom of speech online [freenet2].

5.2.3.5.4. Sybil Attacks

In open-membership P2P networks, a single attacker can pretend to be many participants, typically by creating multiple fake identities of whatever kind the P2P network uses [Douceur]. Attackers can use Sybil attacks to bias choices that the P2P network makes collectively to the attacker's advantage, e.g., by making it more likely that a particular data item (or some threshold of the replicas or shares of a data item) is assigned to attacker-controlled participants. If the P2P network implements any voting, moderation, or peer-review-like functionality, Sybil attacks may be used to "stuff the ballots" to benefit the attacker. Companies and governments can use Sybil attacks on discussion-oriented P2P systems for "astroturfing" or creating the appearance of mass grassroots support for some position where in reality there is none. It is important to know that there are no known complete, environmentally sustainable, and fully distributed solutions to Sybil attacks, and routing via "friends" allows users to be de-anonymized via their social graph. It is important to note that Sybil attacks in this context (e.g., astroturfing) are relevant to more than P2P protocols; they are also common on web-based systems, and they are exploited by governments and commercial entities.

Encrypted P2P and anonymous P2P networks have already emerged. They provide viable platforms for sharing material [Tribler], publishing content anonymously, and communicating securely [Bitmessage]. These platforms are not perfect, and more research needs to be done. If adopted at large, well-designed and resistant P2P networks might represent a critical component of a future secure and distributed Internet, enabling freedom of speech and freedom of information at scale.

5.2.3.6. Virtual Private Networks

The VPNs discussed here are point-to-point connections that enable two computers to communicate over an encrypted tunnel. There are multiple implementations and protocols used in the deployment of VPNs, and they generally diversify by encryption protocol or particular requirements, most commonly in proprietary and enterprise solutions. VPNs are commonly used to (1) enable some devices to communicate through peculiar network configurations, (2) use some privacy and security properties in order to protect the traffic generated by the end user, or both. VPNs have also become a very popular technology among human rights defenders, dissidents, and journalists worldwide to avoid local monitoring and eventually also to circumvent censorship. VPNs are often debated among human rights defenders as a potential alternative to Tor or other anonymous networks. Such comparisons are misleading, as some of the privacy and security properties of VPNs are often misunderstood by less tech-savvy users and could ultimately lead to unintended problems.

As VPNs have increased in popularity, commercial VPN providers have started growing as businesses and are very commonly picked by human rights defenders and people at risk, as they are normally provided with an easy-to-use service and, sometimes, even custom applications to establish the VPN tunnel. Not being able to control the configuration of the network, let alone the security of the application, assessing the general privacy and security state of common VPNs is very hard. Such services have often been discovered to be leaking information, and their custom applications have been found to be flawed. While Tor and similar networks receive a lot of scrutiny from the public and the academic community, commercial or non-commercial VPNs are far less analyzed and understood [Insinuator] [Alshalan-et al], and it might be valuable to establish some standards to guarantee a minimal level of privacy and security to those who need them the most.

5.2.3.6.1. No Anonymity against VPN Providers

One of the common misconceptions among users of VPNs is the level of anonymity that VPNs can provide. This sense of anonymity can be betrayed by a number of attacks or misconfigurations of the VPN provider. It is important to remember that, in contrast to Tor and similar systems, VPNs were not designed to provide anonymity properties. From a technical point of view, a VPN might leak identifiable information or might be the subject of correlation attacks that could expose the originating address of a connecting user. Most importantly, it is vital to understand that commercial and non-commercial VPN providers are bound by the law of the jurisdiction in which they reside or in which their infrastructure is

located, and they might be legally forced to turn over data of specific users if legal investigations or intelligence requirements dictate so. In such cases, if the VPN providers retain logs, it is possible that a user's information could be provided to the user's adversary and lead to his or her identification.

5.2.3.6.2. Logging

Because VPNs are point-to-point connections, the service providers are in fact able to observe the original location of connecting users, and they are able to track at what time they started their session and, eventually, also to which destinations they're trying to connect. If the VPN providers retain logs for a long enough time, they might be forced to turn over the relevant data or they might be otherwise compromised, leading to the same data getting exposed. A clear log-retention policy could be enforced, but considering that countries enforce different levels of data-retention policies, VPN providers should at least be transparent regarding what information they store and for how long it is being kept.

5.2.3.6.3. Third-Party Hosting

VPN providers very commonly rely on third parties to provision the infrastructure that is later going to be used to run VPN endpoints. For example, they might rely on external dedicated server providers or on uplink providers. In those cases, even if the VPN provider itself isn't retaining any significant logs, the information on connecting users might be retained by those third parties instead, introducing an additional collection point for the adversary.

5.2.3.6.4. IPv6 Leakage

Some studies proved that several commercial VPN providers and applications suffer from critical leakage of information through IPv6 due to improper support and configuration [PETS2015VPN]. This is generally caused by a lack of proper configuration of the client's IPv6 routing tables. Considering that most popular browsers and similar applications have been supporting IPv6 by default, if the host is provided with a functional IPv6 configuration, the traffic that is generated might be leaked if the VPN application isn't designed to manipulate such traffic properly.

5.2.3.6.5. DNS Leakage

Similarly, VPN services that aren't handling DNS requests and aren't running DNS servers of their own might be prone to DNS leaking that might not only expose sensitive information on the activity of a user but could also potentially lead to DNS hijacking attacks and subsequent compromises.

5.2.3.6.6. Traffic Correlation

Some VPN implementations appear to be particularly vulnerable to identification and collection of key exchanges that, some Snowden documents revealed, are systematically collected and stored for future reference. The ability of an adversary to monitor network connections at many different points over the Internet can allow them to perform traffic correlation attacks and identify the origin of certain VPN traffic by cross-referencing the connection time of the user to the endpoint and the connection time of the endpoint to the final destination. These types of attacks, although very expensive and normally only performed by very resourceful adversaries, have been documented [SPIEGEL] to be already in practice, and they could completely nullify the use of a VPN and ultimately expose the activity and the identity of a user at risk.

5.2.3.7. HTTP Status Code 451

"Every Internet user has run into the '404 Not Found' Hypertext Transfer Protocol (HTTP) status code when trying, and failing, to access a particular website" [Cath]. It is a response status that the server sends to the browser when the server cannot locate the URL. "403 Forbidden" is another example of this class of code signals that gives users information about what is going on. In the "403" case, the server can be reached but is blocking the request because the user is trying to access content forbidden to them, typically because some content is only for identified users, based on a payment or on special status in the organization. Most of the time, 403 is sent by the origin server, not by an intermediary. If a firewall prevents a government employee from accessing pornography on a work computer, it does not use 403.

As surveillance and censorship of the Internet are becoming more commonplace, voices were raised at the IETF to introduce a new status code that indicates when something is not available for "legal reasons" (like censorship):

The 451 status code would allow server operators to operate with greater transparency in circumstances where issues of law or public policy affect their operation. This transparency may be beneficial to both (1) these operators and (2) end users [RFC7725].

The status code is named "451" in reference to both Bradbury's famous novel "Fahrenheit 451" and to 451 degrees Fahrenheit (the temperature at which some claim book paper autoignites).

During the IETF 92 meeting in Dallas, there was discussion about the usefulness of 451. The main tension revolved around the lack of an apparent machine-readable technical use of the information. The extent to which 451 is just "political theatre" or whether it has a concrete technical use was heatedly debated. Some argued that "the 451 status code is just a status code with a response body"; others said it was problematic because "it brings law into the picture." Still others argued that it would be useful for individuals or for organizations like the "Chilling Effects" project that are crawling the Web to get an indication of censorship (IETF discussion on 451 -- author's field notes, March 2015). There was no outright objection during the Dallas meeting against moving forward on status code 451, and on December 18, 2015, the IESG approved "An HTTP Status Code to Report Legal Obstacles" (now [RFC7725]) for publication. HTTP status code 451 is now an IETF-approved HTTP status code that signals when resource access is denied as a consequence of legal demands.

What is interesting about this particular case is that not only technical arguments but also the status code's outright potential political use for civil society played a substantial role in shaping the discussion and the decision to move forward with this technology.

It is nonetheless important to note that HTTP status code 451 is not a solution to detect all occasions of censorship. A large swath of Internet filtering occurs in the network, at a lower level than HTTP, rather than at the server itself. For these forms of censorship, 451 plays a limited role, as typical censoring intermediaries won't generate it. Besides technical reasons, such filtering regimes are unlikely to voluntarily inject a 451 status code. The use of 451 is most likely to apply in the case of cooperative, legal versions of content removal resulting from requests to providers. One can think of content that is removed or blocked for legal reasons, like copyright infringement, gambling laws, child abuse, etc. Large

Internet companies and search engines are constantly asked to censor content in various jurisdictions. 451 allows this to be easily discovered -- for instance, by initiatives like the Lumen Database.

Overall, the strength of 451 lies in its ability to provide transparency by giving the reason for blocking and giving the end user the ability to file a complaint. It allows organizations to easily measure censorship in an automated way and prompts the user to access the content via another path (e.g., Tor, VPNs) when (s)he encounters the 451 status code.

Status code 451 impacts human rights by making censorship more transparent and measurable. It increases transparency by signaling the existence of censorship (instead of a much broader HTTP error message such as HTTP status code 404) as well as providing details of the legal restriction, which legal authority is imposing it, and to what class of resources it applies. This empowers the user to seek redress.

5.2.3.8. DDoS Attacks

Many individuals, including IETF engineers, have argued that DDoS attacks are fundamentally against freedom of expression. Technically, DDoS attacks are attacks where one host or multiple hosts overload the bandwidth or resources of another host by flooding it with traffic or making resource-intensive requests, causing it to temporarily stop being available to users. One can roughly differentiate three types of DDoS attacks:

1. volume-based attacks (which aim to make the host unreachable by using up all its bandwidth; often-used techniques are UDP floods and ICMP floods)
2. protocol attacks (which aim to use up actual server resources; often-used techniques are SYN floods, fragmented packet attacks, and "ping of death" [RFC4949])
3. application-layer attacks (which aim to bring down a server, such as a web server)

DDoS attacks can thus stifle freedom of expression and complicate the ability of independent media and human rights organizations to exercise their right to (online) freedom of association, while facilitating the ability of governments to censor dissent. When it comes to comparing DDoS attacks to protests in offline life, it is important to remember that only a limited number of DDoS attacks solely involved willing participants. In the overwhelming majority of cases, the clients are hacked hosts of unrelated parties that

have not consented to being part of a DDoS (for exceptions, see Operation Ababil [Ababil] or the Iranian Green Movement's DDoS campaign at election time [GreenMovement]). In addition, DDoS attacks are increasingly used as an extortion tactic.

All of these issues seem to suggest that the IETF should try to ensure that their protocols cannot be used for DDoS attacks; this is consistent with the long-standing IETF consensus that DDoS is an attack that protocols should mitigate to the extent they can [BCP72]. Decreasing the number of vulnerabilities in protocols and (outside of the IETF) the number of bugs in the network stacks of routers or computers could address this issue. The IETF can clearly play a role in bringing about some of these changes, but the IETF cannot be expected to take a positive stance on (specific) DDoS attacks or to create protocols that enable some attacks and inhibit others. What the IETF can do is critically reflect on its role in the development of the Internet and how this impacts the ability of people to exercise their human rights, such as freedom of expression.

6. Model for Developing Human Rights Protocol Considerations

This section outlines a set of human rights protocol considerations for protocol developers. It provides questions that engineers should ask themselves when developing or improving protocols if they want to understand their impact on human rights. It should, however, be noted that the impact of a protocol cannot be solely deduced from its design; its usage and implementation should also be studied to form a full assessment of the impact of the protocol on human rights.

The questions are based on the research performed by the HRPC Research Group. This research was documented prior to the writing of these considerations. The research establishes that human rights relate to standards and protocols; it also offers a common vocabulary of technical concepts that impact human rights and how these technical concepts can be combined to ensure that the Internet remains an enabling environment for human rights. With this, a model for developing human rights protocol considerations has taken shape.

6.1. Human Rights Threats

Human rights threats on the Internet come in a myriad of forms. Protocols and standards can either harm or enable the right to freedom of expression; the right to non-discrimination; the right to equal protection; the right to participate in cultural life, arts, and science; the right to freedom of assembly and association; and the right to security. An end user who is denied access to certain services, data, or websites may be unable to disclose vital information about malpractice on the part of a government or other

authority. A person whose communications are monitored may be prevented from exercising their right to freedom of association or participation in political processes [Penney]. In a worst-case scenario, protocols that leak information can lead to physical danger. A realistic example to consider is when, based on information gathered by state agencies through information leakage in protocols, individuals perceived as threats to the state are subjected to torture, extrajudicial killings, or detention.

This section details several "common" threats to human rights, indicating how each of these can lead to harm to, or violations of, human rights. It also presents several examples of how these threats to human rights materialize on the Internet. This threat modeling is inspired by [RFC6973] ("Privacy Considerations for Internet Protocols"), which is based on security threat analysis. This method is by no means a perfect solution for assessing human rights risks in Internet protocols and systems; it is, however, the best approach currently available. Certain specific human rights threats are indirectly considered in Internet protocols as part of their security considerations [BCP72], but privacy guidelines [RFC6973] or reviews, let alone the assessments of the impact of protocols on human rights, are not standardized or implemented.

Many threats, enablers, and risks are linked to different rights. This is not surprising if one takes into account that human rights are interrelated, interdependent, and indivisible. Here, however, we're not discussing all human rights, because not all human rights are relevant to ICTs in general and to protocols and standards in particular [Bless1]:

The main source of the values of human rights is the International Bill of Human Rights that is composed of the Universal Declaration of Human Rights [UDHR] along with the International Covenant on Civil and Political Rights [ICCPR] and the International Covenant on Economic, Social and Cultural Rights [ICESCR]. In the light of several cases of Internet censorship, the Human Rights Council Resolution 20/8 was adopted in 2012 [UNHRC2016], affirming "... that the same rights that people have offline must also be protected online ...". In 2015, the Charter of Human Rights and Principles for the Internet [IRP] was developed and released. According to these documents, some examples of human rights relevant for ICT systems are human dignity (Art. 1 UDHR), non-discrimination (Art. 2), rights to life, liberty and security (Art. 3), freedom of opinion and expression (Art. 19), freedom of assembly and association (Art. 20), rights to equal protection, legal remedy, fair trial, due process, presumed innocent (Art. 7-11), appropriate social and international order (Art. 28),

participation in public affairs (Art. 21), participation in cultural life, protection of intellectual property (Art. 27), and privacy (Art. 12).

A partial catalog of human rights related to ICTs, including economic rights, can be found in [Hill2014].

This is by no means an attempt to exclude specific rights or prioritize some rights over others. If other rights seem relevant, please contact the authors of this document.

6.2. Guidelines for Human Rights Considerations

This section provides guidance for document authors in the form of a questionnaire about protocols and their (potential) impact. The questionnaire may be useful at any point in the design process, particularly after document authors have developed a high-level protocol model as described in [RFC4101]. These guidelines do not seek to replace any existing referenced specifications; rather, they contribute to them and look at the design process from a human rights perspective.

Protocols and Internet Standards might benefit from a documented discussion of potential human rights risks arising from potential misapplications of the protocol or technology described in the RFC in question. This might be coupled with an Applicability Statement for that RFC.

Note that the guidance provided in this section does not recommend specific practices. The range of protocols developed in the IETF is too broad to make recommendations about particular uses of data or how human rights might be balanced against other design goals. However, by carefully considering the answers to the following questions, document authors should be able to produce a comprehensive analysis that can serve as the basis for discussion on whether the protocol adequately takes specific human rights threats into account. This guidance is meant to help the thought process of a human rights analysis; it does not provide specific directions for how to write a human rights protocol considerations section (following the example set in [RFC6973]), and the addition of a human rights protocol considerations section has also not yet been proposed. In considering these questions, authors will need to be aware of the potential of technical advances or the passage of time to undermine protections. In general, considerations of rights are likely to be more effective if they are considered given a purpose and specific use cases, rather than as abstract absolute goals.

6.2.1. Connectivity

Questions:

- Does your protocol add application-specific functions to intermediary nodes?
- Could this functionality be added to end nodes instead of intermediary nodes?
- Is your protocol optimized for low bandwidth and high-latency connections?
- Could your protocol also be developed in a stateless manner?

Explanation: The end-to-end principle [Saltzer] holds that "the intelligence is end to end rather than hidden in the network" [RFC1958]. The end-to-end principle is important for the robustness of the network and innovation. Such robustness of the network is crucial to enabling human rights like freedom of expression.

Example: Middleboxes (which can be content delivery networks, firewalls, NATs, or other intermediary nodes that provide "services" other than routing) serve many legitimate purposes. But the protocols guiding them can influence individuals' ability to communicate online freely and privately. The potential for abuse, intentional and unintentional censoring, and limiting permissionless innovation -- and thus, ultimately, the impact of middleboxes on the Internet as a place of unfiltered, unmonitored freedom of speech -- is real.

Impacts:

- Right to freedom of expression
- Right to freedom of assembly and association

6.2.2. Privacy

Questions:

- Did you have a look at the guidelines in Section 7 of [RFC6973] ("Privacy Considerations for Internet Protocols")?
- Could your protocol in any way impact the confidentiality of protocol metadata?

- Could your protocol counter traffic analysis?
- Could your protocol improve data minimization?
- Does your document identify potentially sensitive data logged by your protocol and/or for how long that data needs to be retained for technical reasons?

Explanation: "Privacy" refers to the right of an entity (normally a person), acting on its own behalf, to determine the degree to which it will interact with its environment, including the degree to which the entity is willing to share its personal information with others [RFC4949]. If a protocol provides insufficient privacy protection, it may have a negative impact on freedom of expression as users self-censor for fear of surveillance or find themselves unable to express themselves freely.

Example: See [RFC6973].

Impacts:

- Right to freedom of expression
- Right to non-discrimination

6.2.3. Content Agnosticism

Questions:

- If your protocol impacts packet handling, does it use user data (packet data that is not included in the header)?
- Does your protocol make decisions based on the payload of the packet?
- Does your protocol prioritize certain content or services over others in the routing process?
- Is the protocol transparent about the prioritization that is made (if any)?

Explanation: "Content agnosticism" refers to the notion that network traffic is treated identically regardless of payload, with some exceptions when it comes to effective traffic handling -- for instance, delay-tolerant or delay-sensitive packets based on the header.

Example: Content agnosticism prevents payload-based discrimination against packets. This is important because changes to this principle can lead to a two-tiered Internet, where certain packets are prioritized over others based on their content. Effectively, this would mean that although all users are entitled to receive their packets at a certain speed, some users become more equal than others.

Impacts:

- Right to freedom of expression
- Right to non-discrimination
- Right to equal protection

6.2.4. Security

Questions:

- Did you have a look at [BCP72] ("Guidelines for Writing RFC Text on Security Considerations")?
- Have you found any attacks that are somewhat related to your protocol yet considered out of scope for your document?
- Would these attacks be pertinent to the features of the Internet that enable human rights (as described throughout this document)?

Explanation: Most people speak of security as if it were a single monolithic property of a protocol or system; however, upon reflection one realizes that it is clearly not true. Rather, security is a series of related but somewhat independent properties. Not all of these properties are required for every application. Since communications are carried out by systems and access to systems is through communications channels, these goals obviously interlock, but they can also be independently provided [BCP72].

Example: See [BCP72].

Impacts:

- Right to freedom of expression
- Right to freedom of assembly and association
- Right to non-discrimination
- Right to security

6.2.5. Internationalization

Questions:

- Does your protocol have text strings that have to be understood or entered by humans?
- Does your protocol allow Unicode? If so, do you accept texts in one charset (which must be UTF-8) or several (which is dangerous for interoperability)?
- If character sets or encodings other than UTF-8 are allowed, does your protocol mandate proper tagging of the charset?
- Did you have a look at [RFC6365]?

Explanation: "Internationalization" refers to the practice of making protocols, standards, and implementations usable in different languages and scripts (see Section 6.2.12 ("Localization")). "In the IETF, 'internationalization' means to add or improve the handling of non-ASCII text in a protocol" [RFC6365].

A different perspective, more appropriate to protocols that are designed for global use from the beginning, is the definition used by the W3C [W3C118nDef]: "Internationalization is the design and development of a product, application or document content that enables easy localization for target audiences that vary in culture, region, or language."

Many protocols that handle text only handle one charset (US-ASCII), or they leave the question of what coded character set (CCS) and encoding are used up to local guesswork (which leads, of course, to interoperability problems) [RFC3536]. If multiple charsets are permitted, they must be explicitly identified [RFC2277]. Adding non-ASCII text to a protocol allows the protocol to handle more scripts, hopefully all scripts in use in the world. In today's world, that is normally best accomplished by allowing Unicode encoded in UTF-8 only.

In the current IETF policy [RFC2277], internationalization is aimed at user-facing strings, not protocol elements, such as the verbs used by some text-based protocols. (Do note that some strings, such as identifiers, are both content and protocol elements.) If the Internet wants to be a global network of networks, the protocols should work with languages other than English and character sets other than Latin characters. It is therefore crucial that at least the content carried by the protocol can be in any script and that all scripts are treated equally.

Example: See Section 6.2.12 ("Localization").

Impacts:

- Right to freedom of expression
- Right to political participation
- Right to participate in cultural life, arts, and science

6.2.6. Censorship Resistance

Questions:

- Does this protocol introduce new identifiers or reuse existing identifiers (e.g., Media Access Control (MAC) addresses) that might be associated with persons or content?
- Does your protocol make it apparent or transparent when access to a resource is restricted?
- Can your protocol contribute to filtering in such a way that it could be implemented to censor data or services? If so, could your protocol be designed to ensure that this doesn't happen?

Explanation: "Censorship resistance" refers to the methods and measures to prevent Internet censorship.

Example: When IPv6 was developed, embedding a MAC address into unique IP addresses was discussed. This makes it possible, per [RFC4941], for "eavesdroppers and other information collectors to identify when different addresses used in different transactions actually correspond to the same node." This is why privacy extensions for stateless address autoconfiguration in IPv6 [RFC4941] have been introduced.

Identifiers of content exposed within a protocol might be used to facilitate censorship, as in the case of application-layer-based censorship, which affects protocols like HTTP. Denial or restriction of access can be made apparent by the use of status code 451, thereby allowing server operators to operate with greater transparency in circumstances where issues of law or public policy affect their operation [RFC7725].

Impacts:

- Right to freedom of expression
- Right to political participation
- Right to participate in cultural life, arts, and science
- Right to freedom of assembly and association

6.2.7. Open Standards

Questions:

- Is your protocol fully documented in such a way that it could be easily implemented, improved, built upon, and/or further developed?
- Do you depend on proprietary code for the implementation, running, or further development of your protocol?
- Does your protocol favor a particular proprietary specification over technically equivalent and competing specification(s) -- for instance, by making any incorporated vendor specification "required" or "recommended" [RFC2026]?
- Do you normatively reference another standard that is not available without cost (and could you possibly do without it)?
- Are you aware of any patents that would prevent your standard from being fully implemented [RFC6701] [RFC8179]?

Explanation: The Internet was able to be developed into the global network of networks because of the existence of open, non-proprietary standards [Zittrain]. They are crucial for enabling interoperability. Yet, open standards are not explicitly defined within the IETF. On the subject, [RFC2026] states the following: "Various national and international standards bodies, such as ANSI, ISO, IEEE, and ITU-T, develop a variety of protocol and service specifications that are similar to Technical

Specifications defined" at the IETF. "National and international groups also publish 'implementors' agreements' that are analogous to Applicability Statements, capturing a body of implementation-specific detail concerned with the practical application of their standards. All of these are considered to be 'open external standards' for the purposes of the Internet Standards Process." Similarly, [RFC3935] does not define open standards but does emphasize the importance of "open process": any interested person can participate in the work, know what is being decided, and make his or her voice heard on the issue. Part of this principle is the IETF's commitment to making its documents, WG mailing lists, attendance lists, and meeting minutes publicly available on the Internet.

Open standards are important, as they allow for permissionless innovation, which in turn is important for maintaining the freedom and ability to freely create and deploy new protocols on top of the communications constructs that currently exist. It is at the heart of the Internet as we know it, and to maintain its fundamentally open nature, we need to be mindful of the need for developing open standards.

All standards that need to be normatively implemented should be freely available and should provide reasonable protection against patent infringement claims, so that it can also be implemented in open-source or free software. Patents have often held back open standardization or have been used against those deploying open standards, particularly in the domain of cryptography [Newegg]. An exemption is sometimes made when a protocol that normatively relies on specifications produced by other SDOs that are not freely available is standardized. Patents in open standards or in normative references to other standards should have a patent disclosure [notewell], royalty-free licensing [patentpolicy], or some other form of reasonable protection. Reasonable patent protection should include, but is not limited to, cryptographic primitives.

Example: [RFC6108] describes a system deployed by Comcast, an ISP, for providing critical end-user notifications to web browsers. Such a notification system is being used to provide almost-immediate notifications to customers, such as warning them that their traffic exhibits patterns that are indicative of malware or virus infection. There are other proprietary systems that can perform such notifications, but those systems utilize Deep Packet Inspection (DPI) technology. In contrast to DPI, [RFC6108] describes a system that does not rely upon DPI and is instead based on open IETF standards and open-source applications.

Impacts:

- Right to freedom of expression
- Right to participate in cultural life, arts, and science

6.2.8. Heterogeneity Support

Questions:

- Does your protocol support heterogeneity by design?
- Does your protocol allow for multiple types of hardware?
- Does your protocol allow for multiple types of application protocols?
- Is your protocol liberal in what it receives and handles?
- Will your protocol remain usable and open if the context changes?
- Does your protocol allow well-defined extension points? If so, do these extension points allow for open innovation?

Explanation: [FIArch] notes the following: "The Internet is characterized by heterogeneity on many levels: devices and nodes, router scheduling algorithms and queue management mechanisms, routing protocols, levels of multiplexing, protocol versions and implementations, underlying link layers (e.g., point-to-point, multi-access links, wireless, FDDI, etc.), in the traffic mix and in the levels of congestion at different times and places. Moreover, as the Internet is composed of autonomous organizations and internet service providers, each with their own separate policy concerns, there is a large heterogeneity of administrative domains and pricing structures." As a result, as also noted in [FIArch], the heterogeneity principle proposed in [RFC1958] needs to be supported by design.

Example: Heterogeneity is inevitable and needs to be supported by design. For example, multiple types of hardware must be allowed for transmission speeds differing by at least seven orders of magnitude, various computer word lengths, and hosts ranging from memory-starved microprocessors up to massively parallel supercomputers. As noted in [RFC1958], "Multiple types of application protocol must be allowed for, ranging from the simplest such as remote login up to the most complex such as distributed databases."

Impacts:

- Right to freedom of expression
- Right to political participation

6.2.9. Anonymity

Question:

- Did you have a look at [RFC6973] ("Privacy Considerations for Internet Protocols"), especially Section 6.1.1 of that document?

Explanation: "Anonymity" refers to the condition of an identity being unknown or concealed [RFC4949]. Even though full anonymity is hard to achieve, it is a non-binary concept. Making pervasive monitoring and tracking harder is important for many users as well as for the IETF [RFC7258]. Achieving a higher level of anonymity is an important feature for many end users, as it allows them different degrees of privacy online.

Example: Protocols often expose personal data; it is therefore important to consider ways to mitigate the obvious impacts on privacy. A protocol that uses data that could help identify a sender (items of interest) should be protected from third parties. For instance, if one wants to hide the source/destination IP addresses of a packet, the use of IPsec in tunneling mode (e.g., inside a VPN) can help protect against third parties likely to eavesdrop packets exchanged between the tunnel endpoints.

Impacts:

- Right to non-discrimination
- Right to political participation
- Right to freedom of assembly and association
- Right to security

6.2.10. Pseudonymity

Questions:

- Have you considered [RFC6973] ("Privacy Considerations for Internet Protocols"), especially Section 6.1.2 of that document?
- Does the protocol collect personally derived data?

- Does the protocol generate or process anything that can be, or that can be tightly correlated with, personally identifiable information?
- Does the protocol utilize data that is personally derived, i.e., derived from the interaction of a single person or from their device or address?
- Does this protocol generate personally derived data? If so, how will that data be handled?

Explanation: Pseudonymity -- the ability to use a persistent identifier that is not immediately linked to one's offline identity -- is an important feature for many end users, as it allows them different degrees of disguised identity and privacy online.

Example: When designing a standard that exposes personal data, it is important to consider ways to mitigate the obvious impacts. While pseudonyms cannot easily be reverse-engineered -- for example, some early approaches used such techniques as simple hashing of IP addresses that could in turn be easily reversed by generating a hash for each potential IP address and comparing it to the pseudonym -- limiting the exposure of personal data remains important.

"Pseudonymity" means using a pseudonym instead of one's "real" name. There are many reasons for users to use pseudonyms -- for instance, to hide their gender; protect themselves against harassment; protect their families' privacy; frankly discuss sexuality; or develop an artistic or journalistic persona without retribution from an employer, (potential) customers, or social surroundings [geekfeminism]. The difference between anonymity and pseudonymity is that a pseudonym is often persistent.

"Pseudonymity is strengthened when less personal data can be linked to the pseudonym; when the same pseudonym is used less often and across fewer contexts; and when independently chosen pseudonyms are more frequently used for new actions (making them, from an observer's or attacker's perspective, unlinkable)."
[RFC6973]

Impacts:

- Right to non-discrimination
- Right to freedom of assembly and association

6.2.11. Accessibility

Questions:

- Is your protocol designed to provide an enabling environment for people who are not able-bodied?
- Have you looked at the W3C Web Accessibility Initiative [W3CAccessibility] for examples and guidance?

Explanation: The Internet is fundamentally designed to work for all people, whatever their hardware, software, language, culture, location, or physical or mental ability. When the Internet meets this goal, it is accessible to people with a diverse range of hearing, movement, sight, and cognitive abilities [W3CAccessibility]. Sometimes, in the design of protocols, websites, web technologies, or web tools, barriers that exclude people from using the Web are created.

Example: The HTML protocol as defined in [HTML5] specifically requires that (with a few exceptions) every image must have an "alt" attribute to ensure that images are accessible for people that cannot themselves decipher non-text content in web pages.

Impacts:

- Right to non-discrimination
- Right to freedom of assembly and association
- Right to education
- Right to political participation

6.2.12. Localization

Questions:

- Does your protocol uphold the standards of internationalization?
- Have you taken any concrete steps towards localizing your protocol for relevant audiences?

Explanation: Per [W3Ci18nDef], "Localization refers to the adaptation of a product, application or document content to meet the language, cultural and other requirements of a specific target market (a 'locale')." It is also described as the practice of

translating an implementation to make it functional in a specific language or for users in a specific locale (see Section 6.2.5 ("Internationalization")).

Example: The Internet is a global medium, but many of its protocols and products are developed with a certain audience in mind; this audience often shares particular characteristics like knowing how to read and write in ASCII and knowing English. This limits the ability of a large part of the world's online population to use the Internet in a way that is culturally and linguistically accessible. An example of a protocol that has taken into account the view that individuals like to have access to data in their native language can be found in [RFC5646]; such a protocol would label the information content with an identifier for the language in which it is written and would allow information to be presented in more than one language.

Impacts:

- Right to non-discrimination
- Right to participate in cultural life, arts, and science
- Right to freedom of expression

6.2.13. Decentralization

Questions:

- Can your protocol be implemented without one single point of control?
- If applicable, can your protocol be deployed in a federated manner?
- What is the potential for discrimination against users of your protocol?
- Can your protocol be used to negatively implicate users (e.g., incrimination, accusation)?
- Does your protocol create additional centralized points of control?

Explanation: Decentralization is one of the central technical concepts of the architecture of networks and is embraced as such by the IETF [RFC3935]. It refers to the absence or minimization of centralized points of control -- "a feature that is assumed to

make it easy for new users to join and new uses to unfold" [Brown]. It also reduces issues surrounding single points of failure and distributes the network such that it continues to function if one or several nodes are disabled. With the commercialization of the Internet in the early 1990s, there has been a slow trend toward moving away from decentralization, to the detriment of any technical benefits that having a decentralized Internet otherwise provides.

Example: The bits traveling the Internet are increasingly susceptible to monitoring and censorship, from both governments and ISPs, as well as third (malicious) parties. The ability to monitor and censor is further enabled by increased centralization of the network, creating central infrastructure points that can be tapped into. The creation of P2P networks and the development of voice-over-IP protocols using P2P technology in combination with a distributed hash table (DHT) for scalability are examples of how protocols can preserve decentralization [Pouwelse].

Impacts:

- Right to freedom of expression
- Right to freedom of assembly and association

6.2.14. Reliability

Questions:

- Is your protocol fault tolerant?
- Does your protocol degrade gracefully?
- Can your protocol resist malicious degradation attempts?
- Do you have a documented way to announce degradation?
- Do you have measures in place for recovery or partial healing from failure?
- Can your protocol maintain dependability and performance in the face of unanticipated changes or circumstances?

Explanation: Reliability ensures that a protocol will execute its function consistently, be error resistant as described, and function without unexpected results. A system that is reliable degenerates gracefully and will have a documented way to announce degradation. It also has mechanisms to recover from failure

gracefully and, if applicable, to allow for partial healing. It is important here to draw a distinction between random degradation and malicious degradation. Many current attacks against TLS, for example, exploit TLS's ability to gracefully degrade to older cipher suites; from a functional perspective, this ability is good, but from a security perspective, it can be very bad. As with confidentiality, the growth of the Internet and fostering innovation in services depend on users having confidence and trust [RFC3724] in the network. For reliability, it is necessary that services notify users if packet delivery fails. In the case of real-time systems, the protocol needs to safeguard timeliness in addition to providing reliable delivery.

Example: In the modern IP stack structure, a reliable transport layer requires an indication that transport processing has successfully completed, such as the indication given by TCP's ACK message [RFC793] and not simply an indication from the IP layer that the packet arrived. Similarly, an application-layer protocol may require an application-specific acknowledgement that contains, among other things, a status code indicating the disposition of the request (see [RFC3724]).

Impacts:

- Right to freedom of expression
- Right to security

6.2.15. Confidentiality

Questions:

- Does this protocol expose information related to identifiers or data? If so, does it do so to each of the other protocol entities (i.e., recipients, intermediaries, and enablers) [RFC6973]?
- What options exist for protocol implementers to choose to limit the information shared with each entity?
- What operational controls are available to limit the information shared with each entity?
- What controls or consent mechanisms does the protocol define or require before personal data or identifiers are shared or exposed via the protocol? If no such mechanisms or controls are specified, is it expected that control and consent will be handled outside of the protocol?

- Does the protocol provide ways for initiators to share different pieces of information with different recipients? If not, are there mechanisms that exist outside of the protocol to provide initiators with such control?
- Does the protocol provide ways for initiators to limit which information is shared with intermediaries? If not, are there mechanisms that exist outside of the protocol to provide users with such control?
- Is it expected that users will have relationships that govern the use of the information (contractual or otherwise) with those who operate these intermediaries?
- Does the protocol prefer encryption over cleartext operation?
- Does the protocol provide ways for initiators to express individuals' preferences to recipients or intermediaries with regard to the collection, use, or disclosure of their personal data?

Explanation: "Confidentiality" refers to keeping a user's data secret from unintended listeners [BCP72]. The growth of the Internet depends on users having confidence that the network protects their personal data [RFC1984].

Example: Protocols that do not encrypt their payload make the entire content of the communication available to the idealized attacker along their path [RFC7624]. Following the advice in [RFC3365], most such protocols have a secure variant that encrypts the payload for confidentiality, and these secure variants are seeing ever-wider deployment. A noteworthy exception is DNS [RFC1035], as DNSSEC [RFC4033] does not have confidentiality as a requirement. This implies that, in the absence of changes to the protocol as presently under development in the IETF's DNS Private Exchange (DPRIVE) Working Group, all DNS queries and answers generated by the activities of any protocol are available to the attacker. When store-and-forward protocols are used (e.g., SMTP [RFC5321]), intermediaries leave this data subject to observation by an attacker that has compromised these intermediaries, unless the data is encrypted end to end by the application-layer protocol or the implementation uses an encrypted store for this data [RFC7624].

Impacts:

- Right to privacy
- Right to security

6.2.16. Integrity

Questions:

- Does your protocol maintain, assure, and/or verify the accuracy of payload data?
- Does your protocol maintain and assure the consistency of data?
- Does your protocol in any way allow the data to be (intentionally or unintentionally) altered?

Explanation: "Integrity" refers to the maintenance and assurance of the accuracy and consistency of data to ensure that it has not been (intentionally or unintentionally) altered.

Example: Integrity verification of data is important for preventing vulnerabilities and attacks such as man-in-the-middle attacks. These attacks happen when a third party (often for malicious reasons) intercepts a communication between two parties, inserting themselves in the middle and changing the content of the data. In practice, this looks as follows:

Alice wants to communicate with Bob.
Corinne forges and sends a message to Bob, impersonating Alice.
Bob cannot see that the data from Alice was altered by Corinne.
Corinne intercepts and alters the communication as it is sent between Alice and Bob.
Corinne is able to control the communication content.

Impacts:

- Right to freedom of expression
- Right to security

6.2.17. Authenticity

Questions:

- Do you have sufficient measures in place to confirm the truth of an attribute of an entity or of a single piece of data?
- Can attributes get garbled along the way (see Section 6.2.4 ("Security"))?
- If relevant, have you implemented IPsec, DNSSEC, HTTPS, and other standard security best practices?

Explanation: Authenticity ensures that data does indeed come from the source it claims to come from. This is important for preventing (1) certain attacks or (2) unauthorized access to, and use of, data.

Example: Authentication of data is important for preventing vulnerabilities and attacks such as man-in-the-middle attacks. These attacks happen when a third party (often for malicious reasons) intercepts a communication between two parties, inserting themselves in the middle and posing as both parties. In practice, this looks as follows:

Alice wants to communicate with Bob.
Alice sends data to Bob.
Corinne intercepts the data sent to Bob.
Corinne reads and alters the message to Bob.
Bob cannot see that the data did not come from Alice but instead came from Corinne.

When there is proper authentication, the scenario would be as follows:

Alice wants to communicate with Bob.
Alice sends data to Bob.
Corinne intercepts the data sent to Bob.
Corinne reads and alters the message to Bob.
Bob can see that the data did not come from Alice but instead came from Corinne.

Impacts:

- Right to privacy
- Right to freedom of expression
- Right to security

6.2.18. Adaptability

Questions:

- Is your protocol written in such a way that it would be easy for other protocols to be developed on top of it or to interact with it?
- Does your protocol impact permissionless innovation (see Section 6.2.1 ("Connectivity") above)?

Explanation: Adaptability is closely interrelated with permissionless innovation; both maintain the freedom and ability to freely create and deploy new protocols on top of the communications constructs that currently exist. Permissionless innovation is at the heart of the Internet as we know it. To maintain the Internet's fundamentally open nature and ensure that it can continue to develop, we need to be mindful of the impact of protocols on maintaining or reducing permissionless innovation.

Example: WebRTC generates audio and/or video data. In order to ensure that WebRTC can be used in different locations by different parties, it is important that standard JavaScript APIs be developed to support applications from different voice service providers. Multiple parties will have similar capabilities; in order to ensure that all parties can build upon existing standards, these standards need to be adaptable and allow for permissionless innovation.

Impacts:

- Right to education
- Right to freedom of expression
- Right to freedom of assembly and association

6.2.19. Outcome Transparency

Question:

- Are the effects of your protocol fully and easily comprehensible, including with respect to unintended consequences of protocol choices?

Explanation: Certain technical choices may have unintended consequences.

Example: Lack of authenticity may lead to lack of integrity and negative externalities; spam is an example. Lack of data that could be used for billing and accounting can lead to so-called "free" arrangements that obscure the actual costs and distribution of the costs -- for example, (1) the barter arrangements that are commonly used for Internet interconnection and (2) the commercial exploitation of personal data for targeted advertising, which is the most common funding model for the so-called "free" services such as search engines and social networks.

Impacts:

- Right to freedom of expression
- Right to privacy
- Right to freedom of assembly and association
- Right to access to information

7. Security Considerations

As this document discusses research, there are no security considerations.

8. IANA Considerations

This document does not require any IANA actions.

9. Research Group Information

The discussion list for the IRTF Human Rights Protocol Considerations Research Group is located at the email address <hrpc@ietf.org>. Information on the group and information on how to subscribe to the list are provided at <<https://www.irtf.org/mailman/listinfo/hrpc>>.

Archives of the list can be found at
<<https://www.irtf.org/mail-archive/web/hrpc/current/index.html>>.

10. Informative References

- [Ababil] Danchev, D., "Dissecting 'Operation Ababil' - an OSINT Analysis", September 2012, <<http://ddanchev.blogspot.be/2012/09/dissecting-operation-ababil-osint.html>>.
- [Abbate] Abbate, J., "Inventing the Internet", MIT Press, 2000, <<https://mitpress.mit.edu/books/inventing-internet>>.
- [Adrian] Adrian, D., Bhargavan, K., Durumeric, Z., Gaudry, P., Green, M., Halderman, J., Heninger, N., Springall, D., Thome, E., Valenta, L., VanderSloot, B., Wustrow, E., Zanella-Beguelin, S., and P. Zimmermann, "Imperfect Forward Secrecy: How Diffie-Hellman Fails in Practice", Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, pp. 5-17, DOI 10.1145/2810103.2813707, October 2015.
- [Alshalan-et al] Alshalan, A., Pisharody, S., and D. Huang, "A Survey of Mobile VPN Technologies", IEEE Communications Surveys & Tutorials, Volume 18, Issue 2, pp. 1177-1196, DOI 10.1109/COMST.2015.2496624, 2016, <<http://ieeexplore.ieee.org/document/7314859/?arnumber=7314859>>.
- [APIP] Naylor, D., Mukerjee, M., and P. Steenkiste, "Balancing accountability and privacy in the network", SIGCOMM '14, Proceedings of the 2014 ACM Conference on SIGCOMM, pp. 75-86, DOI 10.1145/2740070.2626306, October 2014, <<https://dl.acm.org/citation.cfm?id=2626306>>.
- [Appelbaum] Appelbaum, J., Gibson, A., Goetz, J., Kabisch, V., Kampf, L., and L. Ryge, "NSA targets the privacy-conscious", 2014, <http://daserste.ndr.de/panorama/aktuell/nsa230_page-1.html>.

- [ars] Anderson, N., "P2P researchers: use a blocklist or you will be tracked... 100% of the time", October 2007, <<http://arstechnica.com/uncategorized/2007/10/p2p-researchers-use-a-blocklist-or-you-will-be-tracked-100-of-the-time/>>.
- [Aryan-etal] Aryan, S., Aryan, H., and J. Alex Halderman, "Internet Censorship in Iran: A First Look", 2013, <<https://jhalderm.com/pub/papers/iran-focil3.pdf>>.
- [Babbie] Babbie, E., "The Basics of Social Research", Cengage, Belmont, CA, 2017.
- [BBC-wikileaks] BBC, "Whistle-blower site taken offline", February 2008, <<http://news.bbc.co.uk/2/hi/technology/7250916.stm>>.
- [BCP72] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, July 2003, <<https://www.rfc-editor.org/info/bcp72>>.
- [Benkler] Benkler, Y., "The Wealth of Networks - How Social Production Transforms Markets and Freedom", Yale University Press, New Haven and London, 2006, <<http://is.gd/rxUpTQ>>.
- [Berners-Lee] Berners-Lee, T. and M. Fischetti, "Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web", HarperCollins, p. 208, 1999.
- [BernersLeeHalpin] Berners-Lee, T. and H. Halpin, "Internet Access is a Human Right", 2012, <<http://www.ibiblio.org/hhalpin/homepage/publications/def-timbl-halpin.pdf>>.
- [Bhargavan] Bhargavan, K., Delignat-Lavaud, A., Fournet, C., Pironti, A., and P. Strub, "Triple Handshakes and Cookie Cutters: Breaking and Fixing Authentication over TLS", 2014 IEEE Symposium on Security and Privacy, pp. 98-113, DOI 10.1109/SP.2014.14, May 2014.
- [Bitmessage] Bitmessage, "Bitmessage Wiki", March 2017, <https://bitmessage.org/wiki/Main_Page>.

- [Bless1] Orwat, C. and R. Bless, "Values and Networks - Steps Toward Exploring their Relationships", ACM SIGCOMM Computer Communication Review, Volume 46, Number 2, pp. 25-31, DOI 10.1145/2935634.2935640, April 2016, <<http://www.sigcomm.org/sites/default/files/ccr/papers/2016/April/0000000-0000003.pdf>>.
- [Bless2] Bless, R. and C. Orwat, "Values and Networks", July 2015, <<https://www.ietf.org/proceedings/93/slides/slides-93-hrpc-2.pdf>>.
- [Broeders] Broeders, D., "The public core of the Internet. An international agenda for Internet governance", The Netherlands Scientific Council for Government Policy (WRR) Report No. 94 (under "Reports to the government"), 2015, <<https://english.wrr.nl/publications/reports/2015/10/01/the-public-core-of-the-internet>>
- [Brown] Ziewitz, M. and I. Brown, Ed., "A Prehistory of Internet Governance", Research Handbook on Governance of the Internet, Part 1, Chapter 1 (pp. 3-26), Edward Elgar Publishing Ltd, Cheltenham, DOI 10.4337/9781849805049, 2013.
- [Brown-etal] Brown, I., Clark, D., and D. Trossen, "Should Specific Values Be Embedded In The Internet Architecture?", ReARCH '10, Proceedings of the Re-Architecting the Internet Workshop, Article No. 10, DOI 10.1145/1921233.1921246, November 2010, <http://conferences.sigcomm.org/co-next/2010/Workshops/REARCH/ReArch_papers/10-Brown.pdf>.
- [BrownMarsden] Brown, I. and C. Marsden, "Regulating Code: Good Governance and Better Regulation in the Information Age", MIT Press, 2013, <<https://mitpress.mit.edu/books/regulating-code>>.
- [CAIDA] Dainotti, A., Squarcella, C., Aben, E., Claffy, K., Chiesa, M., Russo, M., and A. Pescapé, "Analysis of Country-wide Internet Outages Caused by Censorship", DOI 10.1109/TNET.2013.2291244, December 2013, <http://www.caida.org/publications/papers/2014/outages_censorship/outages_censorship.pdf>.

- [Cath] Cath, C., "A Case Study of Coding Rights: Should Freedom of Speech Be Instantiated in the Protocols and Standards Designed by the Internet Engineering Task Force?", August 2015, <<https://www.ietf.org/mail-archive/web/hrpc/current/pdf36GrmRM84S.pdf>>.
- [CathFloridi] Cath, C. and L. Floridi, "The Design of the Internet's Architecture by the Internet Engineering Task Force (IETF) and Human Rights", April 2017.
- [Clark] Clark, D., "The Design Philosophy of the DARPA Internet Protocols", SIGCOMM '88, Proceedings of the ACM CCR, Volume 18, Number 4, pp. 106-114, DOI 10.1145/52324.52336, August 1988.
- [Clark-etal] Clark, D., Wroclawski, J., Sollins, K., and R. Braden, "Tussle in cyberspace: defining tomorrow's Internet", IEEE/ACM Transactions on Networking (TON) archive, Volume 13, Issue 3, pp. 462-475, DOI 10.1109/TNET.2005.850224, June 2005, <<https://dl.acm.org/citation.cfm?id=1074049>>.
- [CoE] Council of Europe, "Applications to ICANN for Community-based New Generic Top Level Domains (gTLDs): Opportunities and challenges from a human rights perspective", 2016, <<https://rm.coe.int/CoERMPublicCommonSearchServices/DisplayDCTMContent?documentId=09000016806b5a14>>.
- [Collins] Collins, K., "Hacking Team's oppressive regimes customer list revealed in hack", July 2015, <<http://www.wired.co.uk/news/archive/2015-07/06/hacking-team-spyware-company-hacked>>.
- [Davidson-etal] Davidson, A., Morris, J., and R. Courtney, "Strangers in a Strange Land: Public Interest Advocacy and Internet Standards", Telecommunications Policy Research Conference, Alexandria, VA, September 2002, <<https://www.cdt.org/files/publications/piais.pdf>>.
- [DeNardis14] DeNardis, L., "The Global War for Internet Governance", Yale University Press, 2014, <<https://www.jstor.org/stable/j.ctt5vkz4n>>.

- [DeNardis15] DeNardis, L., "The Internet Design Tension between Surveillance and Security", IEEE Annals of the History of Computing, Volume 37, Issue 2, DOI 10.1109/MAHC.2015.29, 2015, <<http://is.gd/7GANFy>>.
- [Denzin] Denzin, N., Ed., and Y. Lincoln, Ed., "The SAGE Handbook of Qualitative Research", SAGE Handbooks, Thousand Oaks, CA, 2011, <<http://www.amazon.com/SAGE-Handbook-Qualitative-Research-Handbooks/dp/1412974178>>.
- [dict] BusinessDictionary.com, "Reliability (dictionary entry)", WebFinance, Inc., 2017, <<http://www.businessdictionary.com/definition/reliability.html>>.
- [Doty] Doty, N., "Automated text analysis of Requests for Comment (RFCs)", 2014, <<https://github.com/npdoty/rfc-analysis>>.
- [Douceur] Douceur, J., "The Sybil Attack", 2002, <<https://www.microsoft.com/en-us/research/wp-content/uploads/2002/01/IPTPS2002.pdf>>.
- [Dutton] Dutton, W., Dopatka, A., Law, G., and V. Nash, "Freedom of Connection, Freedom of Expression: The Changing Legal and Regulatory Ecology Shaping the Internet", 2011, <<http://www.unesco.org/new/en/communication-and-information/resources/publications-and-communication-materials/publications/full-list/freedom-of-connection-freedom-of-expression-the-changing-legal-and-regulatory-ecology-shaping-the-internet/>>.
- [Farrow] Farrow, R., "Source Address Spoofing", 2016, <<https://technet.microsoft.com/library/cc723706.aspx>>.
- [FIArch] "Future Internet Design Principles", January 2012, <http://www.future-internet.eu/uploads/media/FIArch_Design_Principles_V1.0.pdf>.
- [FOC] Ministers of the Freedom Online Coalition, "The Tallinn Agenda - Recommendations for Freedom Online", 2014, <<https://www.freedomonlinecoalition.com/wp-content/uploads/2014/04/FOC-recommendations-consensus.pdf>>.

- [FRAMEWORK] ISO/IEC, "Information technology - Framework for internationalization", prepared by ISO/IEC JTC 1/SC 22/WG 20 ISO/IEC TR 11017, 1998.
- [Franklin] Franklin, U., "The Real World of Technology", June 1999, <<http://houseofanansi.com/products/the-real-world-of-technology-digital>>.
- [freenet1] Freenet, "What is Freenet?", n.d., <<https://freenetproject.org/whatis.html>>.
- [freenet2] Clarke, I., "The Philosophy behind Freenet", n.d., <<https://freenetproject.org/pages/about.html>>.
- [geekfeminism] Geek Feminism Wiki, "Pseudonymity", 2015, <<http://geekfeminism.wikia.com/wiki/Pseudonymity>>.
- [Geertz] Geertz, H. and C. Geertz, "Kinship in Bali", University of Chicago Press, Chicago, 1975, <<http://press.uchicago.edu/ucp/books/book/chicago/K/bo25832222.html>>.
- [Googlepatent] Google, "Method and device for network traffic manipulation", 2012, <<https://www.google.com/patents/EP2601774A1?cl=en>>.
- [greatfirewall] Anonymous, "Towards a Comprehensive Picture of the Great Firewall's DNS Censorship", 4th USENIX Workshop on Free and Open Communications on the Internet (FOCI) '14, August 2014, <<https://www.usenix.org/system/files/conference/focil4/focil4-anonymous.pdf>>.
- [GreenMovement] Villeneuve, N., "Iran DDoS", 2009, <<https://www.nartv.org/2009/06/16/iran-ddos/>>.
- [Greenwald] Greenwald, G., "XKeyscore: NSA tool collects 'nearly everything a user does on the internet'", July 2013, <<https://www.theguardian.com/world/2013/jul/31/nsa-top-secret-program-online-data>>.

- [Haagsma] Haagsma, L., "Deep dive into QUANTUM INSERT", April 2015, <<http://blog.fox-it.com/2015/04/20/deep-dive-into-quantum-insert/>>.
- [Hall] Hall, J., Aaron, M., Jones, B., and N. Feamster, "A Survey of Worldwide Censorship Techniques", Work in Progress, draft-hall-censorship-tech-04, July 2016.
- [Hill2014] Hill, R., "Partial Catalog of Human Rights Related to ICT Activities", May 2014, <<http://www.apig.ch/UNIGE%20Catalog.pdf>>.
- [HORNET] Chen, C., Asoni, D., Barrera, D., Danezis, G., and A. Perrig, "HORNET: High-speed Onion Routing at the Network Layer", CCS '15, Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, pp. 1441-1454, DOI 10.1145/2810103.2813628, October 2015, <<https://dl.acm.org/citation.cfm?id=2813628>>.
- [HTML5] Hickson, I., Ed., Berjon, R., Ed., Faulkner, S., Ed., Leithead, T., Ed., Navara, E., Ed., O'Connor, E., Ed., and S. Pfeiffer, Ed., "HTML5", W3C Recommendation, October 2014, <<https://www.w3.org/TR/html5/>>.
- [ICCPR] United Nations General Assembly, "International Covenant on Civil and Political Rights", 1966, <<http://www.ohchr.org/EN/ProfessionalInterest/Pages/CCPR.aspx>>.
- [ICESCR] United Nations General Assembly, "International Covenant on Economic, Social and Cultural Rights", 1966, <<http://www.ohchr.org/EN/ProfessionalInterest/Pages/CESCR.aspx>>.
- [Insinuator] Schiess, N., "Vulnerabilities & attack vectors of VPNs (Pt 1)", August 2013, <<https://www.insinuator.net/2013/08/vulnerabilities-attack-vectors-of-vpns-pt-1/>>.
- [IRP] Internet Rights and Principles Dynamic Coalition, "10 Internet Rights & Principles", 2017, <<http://internetrightsandprinciples.org/site/campaign/>>.
- [Jabri] Jabri, V., "Discourses on violence: conflict analysis reconsidered", Manchester University Press, 1996.

- [Kaye] Kaye, D., "Freedom of expression and the private sector in the digital age", 2016, <<http://www.ohchr.org/EN/Issues/FreedomOpinion/Pages/PrivateSectorintheDigitalAge.aspx>>.
- [King] King, C., "Power, Social Violence and Civil Wars", Chapter 8 of "Leashing the Dogs of War: Conflict Management in a Divided World", United States Institute of Peace Press, Washington, D.C., 2007.
- [Lessig] Lessig, L., "Code and Other Laws of Cyberspace, Version 2.0 ('Codev2')", Basic Books, New York, 2006, <<http://codev2.cc/>>.
- [Marcak] Marcak, B., Weaver, N., Dalek, J., Ensafi, R., Fifield, D., McKune, S., Rey, A., Scott-Railton, J., Deibert, R., and V. Paxson, "China's Great Cannon", April 2015, <<https://citizenlab.org/2015/04/chinas-great-cannon/>>.
- [Marquis-Boire] Marquis-Boire, M., "Schrodinger's Cat Video and the Death of Clear-Text", August 2014, <<https://citizenlab.org/2014/08/cat-video-and-the-death-of-clear-text/>>.
- [Meyer] Meyer, J., "Defining and Evaluating Resilience: A Performability Perspective", presentation at International Workshop on Performability Modeling of Computer and Communication Systems, September 2009.
- [Mueller] Mueller, M., "Networks and States: The Global Politics of Internet Governance", MIT Press, DOI 10.7551/mitpress/9780262014595.001.0001, 2010, <<https://mitpress.mit.edu/books/networks-and-states>>.
- [Musiani] Musiani, F., "Giants, Dwarfs and Decentralized Alternatives to Internet-based Services: An Issue of Internet Governance", Westminster Papers in Communication and Culture, 10(1), pp. 81-94, DOI 10.16997/wpcc.214, 2015, <<https://www.westminsterpapers.org/articles/10.16997/wpcc.214/>>.
- [Namecoin] Namecoin, "Namecoin", 2015, <<https://namecoin.info/>>.

- [NATusage] Maier, G., Schneider, F., and A. Feldmann, "NAT usage in Residential Broadband networks", PAM: International Conference on Passive and Active Network Measurement Lecture Notes in Computer Science, Volume 6579, Springer, Berlin and Heidelberg, DOI 10.1007/978-3-642-19260-9_4, 2011, <<http://www.icsi.berkeley.edu/pubs/networking/NATusage11.pdf>>.
- [NETmundial] NETmundial, "NETmundial Multistakeholder Statement", April 2014, <<http://netmundial.br/wp-content/uploads/2014/04/NETmundial-Multistakeholder-Document.pdf>>.
- [Newegg] Mullin, J., "Newegg on trial: Mystery company TQP rewrites the history of encryption", November 2013, <<http://arstechnica.com/tech-policy/2013/11/newegg-on-trial-mystery-company-tqp-re-writes-the-history-of-encryption/>>.
- [notewell] IETF, "Note Well", 2015, <<https://www.ietf.org/about/note-well.html>>.
- [patentpolicy] Weitzner, D., Ed., "W3C Patent Policy", World Wide Web Consortium, February 2004, <<https://www.w3.org/Consortium/Patent-Policy-20040205/>>.
- [Penney] Penney, J., "Chilling Effects: Online Surveillance and Wikipedia Use", 2016, <http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2769645>.
- [Peterson] Peterson, A., Gellman, B., and A. Soltani, "Yahoo to make SSL encryption the default for Webmail users. Finally.", October 2013, <https://www.washingtonpost.com/news/the-switch/wp/2013/10/14/yahoo-to-make-ssl-encryption-the-default-for-webmail-users-finally/?utm_term=.a17eca45ddfe>.
- [PETS2015VPN] Perta, V., Barbera, M., Tyson, G., Haddadi, H., and A. Mei, "A Glance through the VPN Looking Glass: IPv6 Leakage and DNS Hijacking in Commercial VPN clients", DOI 10.1515/popets-2015-0006, 2015, <<http://www.eecs.qmul.ac.uk/~hamed/papers/PETS2015VPN.pdf>>.

- [Pidgin] js and Pidgin Developers, "[XMPP] Invisible mode violating standard", 2007, <<https://developer.pidgin.im/ticket/4322>>.
- [Pouwelse] Pouwelse, J., Ed., "Media without censorship (CensorFree) scenarios", Work in Progress, draft-pouwelse-censorfree-scenarios-02, October 2012.
- [Rachovitsa] Rachovitsa, A., "Engineering and lawyering privacy by design: understanding online privacy both as a technical and an international human rights issue", International Journal of Law and Information Technology, Volume 24, Issue 4, pp. 374-399, DOI 10.1093/ijlit/eaw012, December 2016, <<https://academic.oup.com/ijlit/article/24/4/374/2566975/Engineering-and-lawyering-privacy-by-design>>.
- [RFC760] Postel, J., "DoD standard Internet Protocol", RFC 760, DOI 10.17487/RFC0760, January 1980, <<https://www.rfc-editor.org/info/rfc760>>.
- [RFC791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <<https://www.rfc-editor.org/info/rfc791>>.
- [RFC793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, DOI 10.17487/RFC0793, September 1981, <<https://www.rfc-editor.org/info/rfc793>>.
- [RFC894] Hornig, C., "A Standard for the Transmission of IP Datagrams over Ethernet Networks", STD 41, RFC 894, DOI 10.17487/RFC0894, April 1984, <<https://www.rfc-editor.org/info/rfc894>>.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <<https://www.rfc-editor.org/info/rfc1035>>.
- [RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, DOI 10.17487/RFC1122, October 1989, <<https://www.rfc-editor.org/info/rfc1122>>.
- [RFC1958] Carpenter, B., Ed., "Architectural Principles of the Internet", RFC 1958, DOI 10.17487/RFC1958, June 1996, <<https://www.rfc-editor.org/info/rfc1958>>.

- [RFC1984] IAB and IESG, "IAB and IESG Statement on Cryptographic Technology and the Internet", BCP 200, RFC 1984, DOI 10.17487/RFC1984, August 1996, <<https://www.rfc-editor.org/info/rfc1984>>.
- [RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, DOI 10.17487/RFC2026, October 1996, <<https://www.rfc-editor.org/info/rfc2026>>.
- [RFC2277] Alvestrand, H., "IETF Policy on Character Sets and Languages", BCP 18, RFC 2277, DOI 10.17487/RFC2277, January 1998, <<https://www.rfc-editor.org/info/rfc2277>>.
- [RFC2775] Carpenter, B., "Internet Transparency", RFC 2775, DOI 10.17487/RFC2775, February 2000, <<https://www.rfc-editor.org/info/rfc2775>>.
- [RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", RFC 3022, DOI 10.17487/RFC3022, January 2001, <<https://www.rfc-editor.org/info/rfc3022>>.
- [RFC3365] Schiller, J., "Strong Security Requirements for Internet Engineering Task Force Standard Protocols", BCP 61, RFC 3365, DOI 10.17487/RFC3365, August 2002, <<https://www.rfc-editor.org/info/rfc3365>>.
- [RFC3439] Bush, R. and D. Meyer, "Some Internet Architectural Guidelines and Philosophy", RFC 3439, DOI 10.17487/RFC3439, December 2002, <<https://www.rfc-editor.org/info/rfc3439>>.
- [RFC3536] Hoffman, P., "Terminology Used in Internationalization in the IETF", RFC 3536, DOI 10.17487/RFC3536, May 2003, <<https://www.rfc-editor.org/info/rfc3536>>.
- [RFC3724] Kempf, J., Ed., Austein, R., Ed., and IAB, "The Rise of the Middle and the Future of End-to-End: Reflections on the Evolution of the Internet Architecture", RFC 3724, DOI 10.17487/RFC3724, March 2004, <<https://www.rfc-editor.org/info/rfc3724>>.
- [RFC3935] Alvestrand, H., "A Mission Statement for the IETF", BCP 95, RFC 3935, DOI 10.17487/RFC3935, October 2004, <<https://www.rfc-editor.org/info/rfc3935>>.

- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, DOI 10.17487/RFC4033, March 2005, <<https://www.rfc-editor.org/info/rfc4033>>.
- [RFC4084] Klensin, J., "Terminology for Describing Internet Connectivity", BCP 104, RFC 4084, DOI 10.17487/RFC4084, May 2005, <<https://www.rfc-editor.org/info/rfc4084>>.
- [RFC4101] Rescorla, E. and IAB, "Writing Protocol Models", RFC 4101, DOI 10.17487/RFC4101, June 2005, <<https://www.rfc-editor.org/info/rfc4101>>.
- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, DOI 10.17487/RFC4941, September 2007, <<https://www.rfc-editor.org/info/rfc4941>>.
- [RFC4949] Shirey, R., "Internet Security Glossary, Version 2", FYI 36, RFC 4949, DOI 10.17487/RFC4949, August 2007, <<https://www.rfc-editor.org/info/rfc4949>>.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <<https://www.rfc-editor.org/info/rfc5246>>.
- [RFC5321] Klensin, J., "Simple Mail Transfer Protocol", RFC 5321, DOI 10.17487/RFC5321, October 2008, <<https://www.rfc-editor.org/info/rfc5321>>.
- [RFC5646] Phillips, A., Ed., and M. Davis, Ed., "Tags for Identifying Languages", BCP 47, RFC 5646, DOI 10.17487/RFC5646, September 2009, <<https://www.rfc-editor.org/info/rfc5646>>.
- [RFC5694] Camarillo, G., Ed., and IAB, "Peer-to-Peer (P2P) Architecture: Definition, Taxonomies, Examples, and Applicability", RFC 5694, DOI 10.17487/RFC5694, November 2009, <<https://www.rfc-editor.org/info/rfc5694>>.
- [RFC5944] Perkins, C., Ed., "IP Mobility Support for IPv4, Revised", RFC 5944, DOI 10.17487/RFC5944, November 2010, <<https://www.rfc-editor.org/info/rfc5944>>.

- [RFC6101] Freier, A., Karlton, P., and P. Kocher, "The Secure Sockets Layer (SSL) Protocol Version 3.0", RFC 6101, DOI 10.17487/RFC6101, August 2011, <<https://www.rfc-editor.org/info/rfc6101>>.
- [RFC6108] Chung, C., Kasyanov, A., Livingood, J., Mody, N., and B. Van Lieu, "Comcast's Web Notification System Design", RFC 6108, DOI 10.17487/RFC6108, February 2011, <<https://www.rfc-editor.org/info/rfc6108>>.
- [RFC6120] Saint-Andre, P., "Extensible Messaging and Presence Protocol (XMPP): Core", RFC 6120, DOI 10.17487/RFC6120, March 2011, <<https://www.rfc-editor.org/info/rfc6120>>.
- [RFC6365] Hoffman, P. and J. Klensin, "Terminology Used in Internationalization in the IETF", BCP 166, RFC 6365, DOI 10.17487/RFC6365, September 2011, <<https://www.rfc-editor.org/info/rfc6365>>.
- [RFC6698] Hoffman, P. and J. Schlyter, "The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA", RFC 6698, DOI 10.17487/RFC6698, August 2012, <<https://www.rfc-editor.org/info/rfc6698>>.
- [RFC6701] Farrel, A. and P. Resnick, "Sanctions Available for Application to Violators of IETF IPR Policy", RFC 6701, DOI 10.17487/RFC6701, August 2012, <<https://www.rfc-editor.org/info/rfc6701>>.
- [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>>.
- [RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, DOI 10.17487/RFC6973, July 2013, <<https://www.rfc-editor.org/info/rfc6973>>.
- [RFC7230] Fielding, R., Ed., and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, DOI 10.17487/RFC7230, June 2014, <<https://www.rfc-editor.org/info/rfc7230>>.

- [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>>.
- [RFC7232] Fielding, R., Ed., and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Conditional Requests", RFC 7232, DOI 10.17487/RFC7232, June 2014, <<https://www.rfc-editor.org/info/rfc7232>>.
- [RFC7233] Fielding, R., Ed., Lafon, Y., Ed., and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Range Requests", RFC 7233, DOI 10.17487/RFC7233, June 2014, <<https://www.rfc-editor.org/info/rfc7233>>.
- [RFC7234] Fielding, R., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Caching", RFC 7234, DOI 10.17487/RFC7234, June 2014, <<https://www.rfc-editor.org/info/rfc7234>>.
- [RFC7235] Fielding, R., Ed., and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Authentication", RFC 7235, DOI 10.17487/RFC7235, June 2014, <<https://www.rfc-editor.org/info/rfc7235>>.
- [RFC7236] Reschke, J., "Initial Hypertext Transfer Protocol (HTTP) Authentication Scheme Registrations", RFC 7236, DOI 10.17487/RFC7236, June 2014, <<https://www.rfc-editor.org/info/rfc7236>>.
- [RFC7237] Reschke, J., "Initial Hypertext Transfer Protocol (HTTP) Method Registrations", RFC 7237, DOI 10.17487/RFC7237, June 2014, <<https://www.rfc-editor.org/info/rfc7237>>.
- [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. Slevvi, "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>>.

- [RFC7574] Bakker, A., Petrocco, R., and V. Grishchenko, "Peer-to-Peer Streaming Peer Protocol (PPSPP)", RFC 7574, DOI 10.17487/RFC7574, July 2015, <<https://www.rfc-editor.org/info/rfc7574>>.
- [RFC7624] Barnes, R., Schneier, B., Jennings, C., Hardie, T., Trammell, B., Huitema, C., and D. Borkmann, "Confidentiality in the Face of Pervasive Surveillance: A Threat Model and Problem Statement", RFC 7624, DOI 10.17487/RFC7624, August 2015, <<https://www.rfc-editor.org/info/rfc7624>>.
- [RFC7626] Bortzmeyer, S., "DNS Privacy Considerations", RFC 7626, DOI 10.17487/RFC7626, August 2015, <<https://www.rfc-editor.org/info/rfc7626>>.
- [RFC7725] Bray, T., "An HTTP Status Code to Report Legal Obstacles", RFC 7725, DOI 10.17487/RFC7725, February 2016, <<https://www.rfc-editor.org/info/rfc7725>>.
- [RFC7754] Barnes, R., Cooper, A., Kolkman, O., Thaler, D., and E. Nordmark, "Technical Considerations for Internet Service Blocking and Filtering", RFC 7754, DOI 10.17487/RFC7754, March 2016, <<https://www.rfc-editor.org/info/rfc7754>>.
- [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>>.
- [RFC8164] Nottingham, M. and M. Thomson, "Opportunistic Security for HTTP/2", RFC 8164, DOI 10.17487/RFC8164, May 2017, <<https://www.rfc-editor.org/info/rfc8164>>.
- [RFC8179] Bradner, S. and J. Contreras, "Intellectual Property Rights in IETF Technology", BCP 79, RFC 8179, DOI 10.17487/RFC8179, May 2017, <<https://www.rfc-editor.org/info/rfc8179>>.
- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.
- [Rideout] Rideout, A., "Making security easier", July 2008, <<http://gmailblog.blogspot.de/2008/07/making-security-easier.html>>.

- [Ritchie] Ritchie, J. and J. Lewis, "Qualitative Research Practice: A Guide for Social Science Students and Researchers", SAGE Publishing, London, 2003, <<http://www.amazon.co.uk/Qualitative-Research-Practice-Students-Researchers/dp/0761971106>>.
- [RSF] Reporters Without Borders (RSF), "Syria using 34 Blue Coat servers to spy on Internet users", January 2016, <<https://rsf.org/en/news/syria-using-34-blue-coat-servers-spy-internet-users>>.
- [Saltzer] Saltzer, J., Reed, D., and D. Clark, "End-to-End Arguments in System Design", ACM Transactions on Computer Systems (TOCS), Volume 2, Number 4, pp. 277-288, DOI 10.1145/357401.357402, November 1984.
- [Sandvine] Sandvine, "Sandvine: Over 70% Of North American Traffic Is Now Streaming Video And Audio", December 2015, <<https://www.sandvine.com/pr/2015/12/7/sandvine-over-70-of-north-american-traffic-is-now-streaming-video-and-audio.html>>.
- [Schillace] Schillace, S., "Default https access for Gmail", January 2010, <<http://gmailblog.blogspot.de/2010/01/default-https-access-for-gmail.html>>.
- [Schneier] Schneier, B., "Attacking Tor: how the NSA targets users' online anonymity", October 2013, <<http://www.theguardian.com/world/2013/oct/04/tor-attacks-nsa-users-online-anonymity>>.
- [SPIEGEL] SPIEGEL, "Prying Eyes - Inside the NSA's War on Internet Security", December 2014, <<http://www.spiegel.de/international/germany/inside-the-nsa-s-war-on-internet-security-a-1010361.html>>.
- [sslstrip] Marlinspike, M., "Software >> sslstrip", 2011, <<https://moxie.org/software/sslstrip/>>.
- [techyum] Violet, "Official - vb.ly Link Shortener Seized by Libyan Government", October 2010, <<http://techyum.com/2010/10/official-vb-ly-link-shortener-seized-by-libyan-government/>>.
- [TorProject] The Tor Project, "Anonymity Online", 2006, <<https://www.torproject.org/>>.

- [torrentfreak1] Van der Sar, E., "Is Your ISP Messing With BitTorrent Traffic? Find Out", January 2014, <<https://torrentfreak.com/is-your-isp-messing-with-bittorrent-traffic-find-out-140123/>>.
- [torrentfreak2] Andy, "Lawyers Sent 109,000 Piracy Threats in Germany During 2013", March 2014, <<https://torrentfreak.com/lawyers-sent-109000-piracy-threats-in-germany-during-2013-140304/>>.
- [Tribler] Delft University of Technology, Department EWI/PDS/Tribler, "About Tribler", 2013, <<https://www.tribler.org/about.html>>.
- [UDHR] United Nations General Assembly, "The Universal Declaration of Human Rights", 1948, <<http://www.un.org/en/universal-declaration-human-rights/index.html>>.
- [UNGA2013] United Nations General Assembly, "UN General Assembly Resolution "The right to privacy in the digital age" (A/C.3/68/L.45)", 2013, <<https://documents-dds-ny.un.org/doc/UNDOC/LTD/N13/576/77/PDF/N1357677.pdf?OpenElement>>.
- [UNHRC2016] United Nations Human Rights Council, "The promotion, protection and enjoyment of human rights on the Internet", Resolution A/HRC/32/L.20, 2016, <http://ap.ohchr.org/documents/alldocs.aspx?doc_id=20340>.
- [Ververis] Ververis, V., Kargiotakis, G., Filasto, A., Fabian, B., and A. Alexandros, "Understanding Internet Censorship Policy: The Case of Greece", 5th USENIX Workshop on Free and Open Communications on the Internet (FOCI) '15, August 2015, <<https://www.usenix.org/system/files/conference/focil15/focil15-paper-ververis-update.pdf>>.
- [W3CAccessibility] World Wide Web Consortium, "Accessibility", 2016, <<https://www.w3.org/standards/webdesign/accessibility>>.
- [W3Ci18nDef] Ishida, R. and S. Miller, "Localization vs. Internationalization", World Wide Web Consortium, April 2015, <<http://www.w3.org/International/questions/qa-il18n.en>>.

- [wikileaks] Sladek, T. and E. Broese, "Market Survey: Detection & Filtering Solutions to Identify File Transfer of Copyright Protected Content for Warner Bros. and movielabs", 2011, <<https://wikileaks.org/sony/docs/05/docs/Anti-Piracy/CDSA/EANTC-Survey-1.5-unsecured.pdf>>.
- [WP-Tempora] Wikipedia, "Tempora", September 2017, <<https://en.wikipedia.org/wiki/Tempora>>.
- [WSJ] Sonne, P. and M. Coker, "Firms Aided Libyan Spies", The Wall Street Journal, August 2011, <<http://www.wsj.com/articles/SB10001424053111904199404576538721260166388>>.
- [WynsbergheMoura] Nguyen, B., Ed., van Wynsberghe, A., van Wynsberghe, A., and G. Moreira Moura, "The concept of embedded values and the example of internet security", June 2013, <<http://doc.utwente.nl/87095/>>.
- [XMPP-Manifesto] Saint-Andre, P. and XMPP Operators, "A Public Statement Regarding Ubiquitous Encryption on the XMPP Network", March 2014, <<https://raw.githubusercontent.com/stpeter/manifesto/master/manifesto.txt>>.
- [Zittrain] Zittrain, J., "The Future of the Internet - And How to Stop It", Yale University Press & Penguin UK, 2008, <https://dash.harvard.edu/bitstream/handle/1/4455262/Zittrain_Future%20of%20the%20Internet.pdf?sequence=1>.

Acknowledgements

A special thanks to all members of the HRPC Research Group who contributed to this document. The following deserve a special mention:

- Joana Varon for helping draft the first iteration of the methodology and previous drafts, and for directing the film "Net of Rights" and working on the interviews at IETF 92 in Dallas.
- Daniel Kahn Gillmor (dkg) for helping with the first iteration of the glossary (Section 2) as well as a lot of technical guidance, support, and language suggestions.
- Claudio Guarnieri for writing the first iterations of the case studies on VPNs, HTTP, and P2P.
- Will Scott for writing the first iterations of the case studies on DNS, IP, and XMPP.
- Avri Doria for proposing writing a glossary in the first place, help with writing the initial proposals and Internet-Drafts, her reviews, and her contributions to the glossary.

Thanks also to Stephane Bortzmeyer, John Curran, Barry Shein, Joe Hall, Joss Wright, Harry Halpin, and Tim Sammut, who made a lot of excellent suggestions, many of which found their way directly into the text. We want to thank Amelia Andersdotter, Stephen Farrell, Stephane Bortzmeyer, Shane Kerr, Giovane Moura, James Gannon, Alissa Cooper, Andrew Sullivan, S. Moonesamy, Roland Bless, and Scott Craig for their reviews and for testing the HRPC guidelines in the wild. We would also like to thank Molly Sauter, Arturo Filasto, Nathalie Marechal, Eleanor Saitta, Richard Hill, and all others who provided input on this document or the conceptualization of the idea. Thanks to Edward Snowden for his comments at IETF 93 in Prague regarding the impact of protocols on the rights of users.

Authors' Addresses

Niels ten Oever
ARTICLE 19

Email: mail@nielstenoever.net

Corinne Cath
Oxford Internet Institute

Email: corinnecath@gmail.com