Guidelines for Human Rights Protocol and Architecture Considerations

Abstract

This document sets guidelines for human rights considerations for developers working on network protocols and architectures, similar to the work done on the guidelines for privacy considerations [RFC6973]. This is an updated version of the guidelines for human rights considerations in [RFC8280].

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

This informational document has consensus for publication from the Internet Research Task Force (IRTF) Human Right Protocol Considerations Research (HRPC) Group. It has been reviewed, tried, and tested by both by the research group as well as by researchers and practitioners from outside the research group. The research group acknowledges that the understanding of the impact of Internet protocols and architecture on society is a developing practice and is a body of research that is still in development.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

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

The questions are based on the research performed by the Human
Rights Protocol Considerations (HRPC) research group which has been
documented before these considerations. The research establishes
that human rights relate to standards and protocols, and offers a
common vocabulary of technical concepts that influence human rights
and how these technical concepts can be combined to ensure that the
Internet remains an enabling environment for human rights. With
this, the contours of a model for developing human rights protocol
considerations has taken shape.

This document is an iteration of the guidelines that can be found in
[RFC8280]. The methods for conducting human rights reviews (Section
3.2), and guidelines for human rights considerations (Section 3.3)
in this document are being tested for relevance, accuracy, and
validity. [HR-RT] The understanding of what human rights are is
based on the Universal Declaration of Human Rights [UDHR] and
subsequent treaties that jointly form the body of international
human rights law [UNHR].

This document does not provide a detailed taxonomy of the nature of
(potential) human rights violations, whether direct or indirect,
long-term or short-term, certain protocol choices might present. In
part because this is highly context-dependent, and in part, because
this document aims to provide a practical set of guidelines.
However, further research in this field would definitely benefit
developers and implementers.

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

This informational document has consensus for publication from the Internet Research Task Force (IRTF) Human Right Protocol Considerations Research Group. It has been reviewed, tried, and tested by both by the research group as well as by researchers and practitioners from outside the research group. The HRPC research group acknowledges that the understanding of the impact of Internet protocols and architecture on society is a developing practice and is a body of research that is still in development.

2. **Human rights threats**

Threats to the exercise of human rights on the Internet come in many forms. Protocols and standards may harm or enable the right to freedom of expression, right to freedom of information, right to non-discrimination, right to equal protection, right to participate in cultural life, arts and science, right to freedom of assembly and association, right to privacy, and the right to security. An end-user who is denied access to certain services or content may be unable to disclose vital information about the malpractices of a government or other authority. A person whose communications are monitored may be prevented or dissuaded from exercising their right to freedom of association or participate 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 individuals perceived as threats to the state are subjected to torture, extra-judicial killing or detention on the basis of information gathered by state agencies through the monitoring of network traffic.

This document presents several examples of how 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 a work in progress and by no means a perfect solution for assessing human rights risks in Internet protocols and systems. Certain specific human rights threats are indirectly considered in Internet protocols as part of the security considerations [BCP72], but privacy considerations [RFC6973] or reviews, let alone human rights impact assessments of protocols, are neither standardized nor 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 information and communication technologies (ICTs) in general and protocols and standards in particular [Bless]: "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 UN Human Rights Council Resolution 20/8 was adopted in 2012, affirming that "the same rights that people have offline must also be protected online." [UNHRC2016] 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 the moral and material interests resulting from any scientific, literary or artistic production of which [they are] the author (Art. 27), and privacy (Art. 12)." A partial catalog of human rights related to Information and Communications Technologies, 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.

3. Conducting human rights reviews

Ideally, protocol developers and collaborators should incorporate human rights considerations into the design process itself (see Guidelines for human rights considerations). This section provides guidance on how to conduct a human rights review, i.e., gauge the impact or potential impact of a protocol or standard on human rights.

Human rights reviews can be done by any participant, and can take place at different stages of the development process of an Internet-Draft. Generally speaking, it is easier to influence the development of a technology at earlier stages than at later stages. This does not mean that reviews at last-call are not relevant, but they are less likely to result in significant changes in the reviewed document.

Human rights review can be done by document authors, document shepherds, members of review teams, advocates, or impacted communities to influence the standard development process. IETF documents can benefit from people with different knowledges, perspectives, and backgrounds, especially since their implementation can impact many different communities as well.

Methods for analyzing technology for specific human rights impacts are still quite nascent. Currently, five methods have been explored by the human rights review team, often in conjunction with each other:

3.1. Analyzing drafts based on guidelines for human rights considerations model

This analysis of Internet-Drafts uses the model as described in section 4. The outlined categories and questions can be used to review an Internet-Draft. The advantage of this is that it provides a known overview, and document authors can go back to this document as well as [RFC8280] to understand the background and the context.

3.2. Analyzing drafts based on their perceived or speculated impact

When reviewing an Internet-Draft, specific human rights impacts can become apparent by doing a close reading of the draft and seeking to understand how it might affect networks or society. While less structured than the straight use of the human rights considerations model, this analysis may lead to new speculative understandings of links between human rights and protocols.

3.3. Expert interviews

Interviews with document authors, active members of the Working Group, or experts in the field can help explore the characteristics of the protocol and its effects. There are two main advantages to this approach: one the one hand, it allows the reviewer to gain a deeper understanding of the (intended) workings of the protocol; on the other hand, it also allows for the reviewer to start a discussion with experts or even document authors, which can help the review gain traction when it is published.

3.4. Interviews with impacted persons and communities

Protocols impact users of the Internet. Interviews can help the reviewer understand how protocols affect the people that use the protocols. Since human rights are best understood from the perspective of the rights-holder, this approach will improve the understanding of the real world effects of the technology. At the same time, it can be hard to attribute specific changes to a particular protocol, this is of course even harder when a protocol has not been (widely) deployed.

3.5. Tracing impacts of implementations

The reality of deployed protocols can be at odds with the expectations during the protocol design and development phase [RFC8980]. When a specification already has associated running code,

the code can be analyzed either in an experimental setting or on the Internet where its impact can be observed. In contrast to reviewing the draft text, this approach can allow the reviewer to understand how the specifications works in practice, and potentially what unknown or unexpected effects the technology has.

4. **Guidelines for human rights considerations**

This section provides guidance for document authors in the form of a questionnaire about protocols and how technical decisions can shape the exercise of human rights. The questionnaire may be useful at any point in the design process, particularly after the document authors have developed a high-level protocol model as described in [RFC4101]. These guidelines do not seek to replace any existing referenced specifications, but rather 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 Request For Comments (RFC). 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 considerations section (following the example set in [RFC6973]).

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.

Also note that while the section uses the word, 'protocol', the principles identified in these questions may be applicable to other types of solutions (extensions to existing protocols, architecture for solutions to specific problems, etc.).

4.1. Intermediaries

Question(s): Does your protocol depend on or allow for protocol-specific functions at intermediary nodes?

Explanation: The end-to-end principle [Saltzer] holds that certain functions can and should be performed at 'ends' of the network. [RFC1958] states "that in very general terms, the community believes that the goal is connectivity […] and the intelligence is end to end rather than hidden in the network." When a protocol exchange includes both endpoints and an intermediary, there are new opportunities for failure, especially when the intermediary is not under control of either endpoint, or even largely invisible to it, as, for instance, in intercepting HTTPS proxies [https-interception]. This pattern also contributes to ossification, because the intermediaries may impose protocol restrictions – sometimes in violation of the specification – that prevent the endpoints from using more modern protocols, as described in Section 9.3 of [RFC8446].

Note that intermediaries are distinct from services: in the former case the third party element is part of the protocol exchange, whereas in the latter the endpoints communicate explicitly with the service. The client/server pattern provides clearer separation of responsibilities between elements than having an intermediary. However, even in client/server systems, it is often good practice to provide for end-to-end encryption between endpoints for protocol elements which are outside of the scope of the service, as in the design of MLS [I-D.ietf-mls-protocol].

Example: Encryption between the endpoints can be used to protect the protocol from interference by intermediaries. The encryption of transport layer information in QUIC [RFC9000] and of the TLS Server Name Indication field [I-D.ietf-tls-esni] are examples of this practice. One consequence of this is to limit the extent to which network operators can inspect traffic, requiring them to have control of the endpoints in order to monitor their behavior.

Impacts:

  *Right to freedom of expression

  *Right to freedom of assembly and association

4.2. Connectivity

Questions(s): Is your protocol optimized for low bandwidth and high latency connections? Could your protocol also be developed in a stateless manner?

Also considering the fact that network quality and conditions vary across geography and time, it is also important to design protocols such that they are reliable even on low bandwidth and high latency connections.

Impacts:

  *Right to freedom of expression

  *Right to freedom of assembly and association

4.3. Reliability

Question(s): Is your protocol fault tolerant? Does it downgrade gracefully, i.e., with mechanisms for fallback and/or notice? 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 and resiliency ensures that a protocol will execute its function consistently and error resistant as described, and function without unexpected result. Measures for reliability in protocols assure users that their intended communication was successfully executed.

A system that is reliable degrades gracefully and will have a documented way to announce degradation. It will also have mechanisms to recover from failure gracefully, and if applicable, will allow for partial healing.

It is important here to draw a distinction between random degradation and malicious degradation. Some attacks against previous versions of TLS, for example, exploited TLS' ability to gracefully downgrade to non-secure cipher suites [FREAK][Logjam]– from a functional perspective, this is useful; from a security perspective, this can be disastrous.

For reliability, it is necessary that services notify the users if a delivery fails. In the case of real-time systems, in addition to the reliable delivery, the protocol needs to safeguard timeliness.

Example: In the modern IP stack structure, a reliable transport layer requires an indication that transport processing has successfully completed, such as given by TCP's ACK message [RFC0793]. Similarly, an application layer protocol may require an application-specific acknowledgment 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

4.4. Content signals

Question(s): Does your protocol include explicit or implicit
plaintext elements, either in the payload or headers, that can be
used for differential treatment? Is there a way minimise leaking of
such data to network intermediaries? If not, is there a way for
deployments of the protocol to make the differential treatment
(including prioritisation of certain traffic), if any, auditable for
negative impacts on net neutrality?

Example: When network intermediaries are able to determine the type
of content that a packet is carrying then they can use that
information to discriminate in favor of one type of content and
against another. This impacts users ability to send and receive the
content of their choice.

As recommended in [RFC8558] protocol designers should avoid the
construction of implicit signals of their content. In general,
protocol designers should avoid adding explicit signals for
intermediaries. In certain cases, it may be necessary to add such
explicit signals, but designers should only do so when they provide
clear benefit to end users (see [RFC8890] for more on the priority
of constituencies). In these cases, the implications of those signal
for human rights should be documented.

Note that many protocols provide signals that are intended for
endpoints that can be used as implicit signals by intermediaries for
traffic discrimination, either based on content (e.g., TCP port
numbers) or sender/receiver (IP addresses). Where possible, these
should be protected from intermediaries by encryption. In many cases
– e.g., IP address – these signals are difficult to remove, but in
other cases, such as TLS Application Layer Protocol Negotiation
[RFC7301], there are active efforts to protect this data
[I-D.ietf-tls-esni].

   *Right to freedom of expression

   *Right to non-discrimination

   *Right to equal protection

4.5. Internationalization

Question(s): Does your protocol or specification define text string elements, in the payload or headers, that have to be understood or entered by humans? Does your specification 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 specification mandate a 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 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):

  "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." {{W3Ci18nDef}}

Many protocols that handle text only handle one charset (US-ASCII), or leave the question of what coded character set and encoding are used up to local guesswork (which leads, of course, to interoperability problems). 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 representing users across the world. In today's world, that is normally best accomplished by allowing Unicode encoded in UTF-8 only.

In current IETF practice [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 are both content and protocol elements, such as identifiers.) Although this is reasonable practice for non-user visible elements, given the IETF's mission to make the Internet a global network of networks, [RFC3935] developers should provide full and equal support for all scripts and character sets in the user-facing features of protocols and for any content they carry.

Example: See localization

Impacts:

  *Right to freedom of expression

*Right to political participation

      *Right to participate in cultural life, arts and science

4.6. Localization

   Question(s): Does your protocol uphold the standards of
   internationalization? Have you made any concrete steps towards
   localizing your protocol for relevant audiences?

   Explanation: 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)
   [W3Ci18nDef]. For our purposes, it can be described as the practice
   of translating an implementation to make it functional in a specific
   language or for users in a specific locale (see
   Internationalization). Internationalization is related to
   localization, but they are not the same. Internationalization is a
   necessary precondition for localization.

   Example: The Internet is a global medium, but many of its protocols
   and products are developed with a certain audience in mind, that
   often share particular characteristics like knowing how to read and
   write in American Standard Code for Information Interchange (ASCII)
   and knowing English. This limits the ability of a large part of the
   world's online population from using the Internet in a way that is
   culturally and linguistically accessible. An example of a standard
   that has taken into account the view that individuals like to have
   access to data in their native language can be found in [RFC5646].
   The document describes a way to label information with an identifier
   for the language in which it is written. And this allows information
   to be presented and accessed in more than one language.

   Impacts:

      *Right to non-discrimination

      *Right to participate in cultural life, arts and science

      *Right to freedom of expression

4.7. Open Standards

   Question(s): Is your protocol fully documented in 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 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 do without it)? Are you aware of any patents that would prevent your standard from being fully implemented [RFC8179] [RFC6701]?

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: "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 an "open process", i.e., "any interested person can participate in the work, know what is being decided, and make [their] voice heard on the issue."

Open standards (and open source software) allow users to glean information about how the tools they are using work, including the tools' security and privacy properties. They additionally allow for permissionless innovation, which is important to maintain 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 with reasonable protection for patent infringement claims, so it can also be implemented in open source or free software. Patents have often held back open standardization or been used against those deploying open standards, particularly in the domain of cryptography [newegg]. An exemption of this is sometimes made when a protocol is standardized that normatively relies on specifications produced by others standards development organizations (SDOs) that are not freely available. 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 fair, reasonable and non-discriminatory terms.

Example: [RFC6108] describes a system for providing critical end-user notifications to web browsers, which has been deployed by Comcast, an Internet Service Provider (ISP). Such a notification system is being used to provide near-immediate notifications to customers, such as to warn 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, that document 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

4.8. Heterogeneity Support

Question(s): 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 it remain usable and open if the context changes?

Explanation: 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 ISPs, each with their own separate policy concerns, there is a large heterogeneity of administrative domains and pricing structures. As a result, the heterogeneity principle proposed in [RFC1958] needs to be supported by design [FIArch].

Heterogeneity support in protocols can thus enable a wide range of devices and (by extension) users to participate on the network.

Example: Heterogeneity significantly contributed to the success of the internet architecture [Zittrain]. Niels Bohr famously said: "Prediction is very difficult, especially if it's about the future", this also holds true for future uses of the internet architecture and infrastructure. Therefore, as a rule of thumb it is important to - as far as possible - design your protocol for different devices and uses, especially at lower layers of the stack. However, if you choose not to do this, it could be relevant to document the reasoning for that.

Impacts:

   *Right to freedom of expression

   *Right to political participation

4.9. Adaptability

Question(s): Question: Is your protocol written in a modular fashion
and does it facilitate or hamper extensibility? In this sense, does
your protocol impact permissionless innovation? (See Open Standards)

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. 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 impact of protocols on maintaining or
reducing permissionless innovation to ensure the Internet can
continue to develop.

Adaptability and permissionless innovation can be used to shape
information networks as preferenced by groups of users. Furthermore,
a precondition of adaptability is the ability of the people who can
adapt the network to be able to know and understand the network.
This is why adaptability and permissionless innovation are
inherently connected to the right to education and the right to
science as well as the right to freedom of assembly and association
as well as the right to freedom of expression. Since it allows the
users of the network to determine how to assemble, collaborate, and
express themselves.

Example: WebRTC generates audio and/or video data. WebRTC can be
used in different locations by different parties; WebRTC's standard
application programming interfaces (APIs) are 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 need to be
adaptable, and allow for permissionless innovation.

Impacts:

   *Right to education

   *Right to science

   *Right to freedom of expression

   *Right to freedom of assembly and association

4.10. Integrity

Question(s): 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 for 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 it has not been (intentionally or unintentionally) altered.

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

Alice wants to communicate with Bob. Alice sends a message to Bob, which Corinne intercepts and modifies. 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

4.11. Authenticity

Question(s): Do you have sufficient measures to confirm the truth of an attribute of a single piece of data or entity? Can the attributes get garbled along the way (see security)? If relevant, have you implemented IPsec, DNS Security (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 to prevent certain attacks or unauthorized access and use of data.

At the same time, authentication should not be used as a way to prevent heterogeneity support, as is often done for vendor lock-in or digital rights management.

Example: Authentication of data is important to prevent vulnerabilities, and attacks from on-path attackers. 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
potentially alters) the message to Bob. Bob cannot see that the data
did not come from Alice but from Corinne.

With 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 is unable to verify whether that the data
came from Alice.

Impacts:

   *Right to privacy

   *Right to freedom of expression

   *Right to security

4.12. Confidentiality

Question(s): Does the protocol expose the transmitted data over the
wire? Does the protocol expose information related to identifiers or
data? If so, what does it reveal to each protocol entity (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 the sharing
or express individuals' preferences to recipients or intermediaries
with regard to the collection, use, or disclosure of their personal
data? 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 clear text operation?

Explanation: Confidentiality refers to keeping your data secret from unintended listeners [BCP72]. The growth of the Internet depends on users having confidence that the network protects their personal data [RFC1984]. The possibility of pervasive monitoring and surveillance undermines users' trust, and can be mitigated by ensuring confidentiality, i.e., passive attackers should gain little or no information from observation or inference of protocol activity. [RFC7258][RFC7624].

Example: Protocols that do not encrypt their payload make the entire content of the communication available to the idealized attacker along their path. 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 the use of more recent standards like DNS over TLS [RFC7858] or DNS over HTTPS [RFC8484], 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

4.13. Security

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

Explanation: Security is not a single monolithic property of a protocol or system, but rather 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, security

goals obviously interlock, but they can also be independently
provided. [BCP72].

Typically, any protocol operating on the Internet can be the target
of passive attacks (when the attacker can access and read packets on
the network); active attacks (when an attacker is capable of writing
information to the network packets). [BCP72]

Example: See [BCP72].

Impacts:

   *Right to freedom of expression

   *Right to freedom of assembly and association

   *Right to non-discrimination

   *Right to security

4.14. Privacy

Question(s): Did you have a look at the Guidelines in the Privacy
Considerations for Internet Protocols [RFC6973] section 7? Does your
protocol maintain the confidentiality of metadata? Could your
protocol counter traffic analysis? Does your protocol adhere to data
minimization principles? Does your document identify potentially
sensitive data logged by your protocol and/or for how long that
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 privacy

   *Right to non-discrimination

4.15. Anonymity and Pseudonymity

Question(s): Does your protocol make use of identifiers? Are these identifiers persistent? Are they used across multiple contexts? Is it possible for the user to reset or rotate them without negatively impacting the operation fo the protocol? Are they visible to others besides the protocol endpoints? Are they tied to real-world identities? Have you considered the Privacy Considerations for Internet Protocols [RFC6973], especially section 6.1.2?

Explanation: Most protocols depend on the use of some kind of identifier in order to correlate activity over time and space. For instance:

  *IP addresses are used as an identity for the source and
   destination for IP datagrams.

  *QUIC connection identifiers are used to correlate packets
   belonging to the same connection.

  *HTTP uses cookies to correlate multiple HTTP requests from the
   same client.

  *Email uses email addresses of the form example@example.com to
   identify senders and receivers.

In general, these identifiers serve a necessary function for protocol operations, by allowing them to maintain continuity. However, they can also create privacy risks. There are two major ways in which those risks manifest:

  *The identifier may itself reveal the user's identity in some way
   or be tied to an identifier which does, as is the case when E.164
   (telephone) numbers are used as identifiers for instant messaging
   systems.

  *While the identifier may not reveal the user's identity, it may
   make it possible to link enough of a user's behavior to threaten
   their privacy, as is the case with HTTP cookies.

Because identifiers are necessary for protocol operation, true anonymity is very difficult to achieve, but there are practices which promote user privacy even when identifiers are used.

Impacts:

  *Right to non-discrimination

  *Right to freedom of expression

*Right to political participation

   *Right to freedom of assembly and association

4.15.1. Pseudonymity

   In general, user privacy is better preserved when identifiers are
   pseudonymous (not tied to a user's real-world identity).

   Example: In the development of the IPv6 protocol, it was discussed
   to embed a Media Access Control (MAC) address into unique IP
   addresses. This would make it possible for eavesdroppers and other
   information collectors to identify when different addresses used in
   different transactions actually correspond to the same node. This is
   why standardization efforts like Privacy Extensions for Stateless
   Address Autoconfiguration in IPv6 [RFC4941] and MAC address
   randomization [draft-zuniga-mac-address-randomization] have been
   pursued.

   Note that it is often attractive to try to create a pseudonym from a
   persistent identifier. This can be very difficult to do correctly in
   a way that does not allow for recovering the persistent identifiers.

   Example: A common practice in Web tracking is to "encrypt" email
   addresses by hashing them, thus allegedly making them "non-
   personally identifying". However, because hash functions are public
   operations, it is possible to dictionary search candidate email
   addresses and recover the original address [email-hashing].

4.15.2. Unlinkability

   Even true pseudonymous identifiers can present a privacy risk if
   they are used across a wide enough scope. User privacy is better
   preserved if identifiers have limited scope both in time and space.

   Example: An example is Dynamic Host Configuration Protocol (DHCP)
   where sending a persistent identifier as the client name was not
   mandatory but, in practice, done by many implementations, before
   [RFC7844].

   Example: Third party cookies in HTTP allow trackers to correlate
   HTTP traffic across sites. This is the foundation of a whole
   ecosystem of Web tracking. Increasingly, Web browsers are
   restricting the use of third party cookies in order to protect user
   privacy.

4.16. Censorship resistance

   Question(s): Does your protocol architecture facilitate censorship?
   Does it include "choke points" which are easy to use for censorship?

Does it expose identifiers which can be used to selectively block certain kinds of trafic? Could it be designed to be more censorship resistant? Does your protocol make it apparent or transparent when access to a resource is restricted and the reasons why it is restricted?

Explanation: Governments and service providers block or filter content or traffic, often without the knowledge of end-users. [RFC7754] See [draft-irtf-pearg-censorship] for a survey of censorship techniques employed across the world, which lays out protocol properties that have been exploited to censor access to information. Censorship resistance refers to the methods and measures to prevent Internet censorship.

Example: The current design of the Web has a number of architectural choke points where it is possible for censors to intervene. These include obtaining the control of the domain name itself, DNS blocking at either the protocol layer or at the resolver, IP address blocking, and blocking at the Web server. There has been extensive work on content distribution systems which are intended to be more censorship resistant, and some, such as BitTorrent, are in wide use, but these systems may have inferior reliability and performance compared to the Web (e.g., they do not support active content on the server).

Example: Identifiers of content exposed within a protocol might be used to facilitate censorship by allowing the censor to determine which traffic to block. DNS queries, the "host" request header in an HTTP request, the Server Name Indication (SNI) in a Transport Layer Security (TLS) ClientHello are all examples of protocol elements that can travel in plaintext and be used by censors to identify what content a user is trying to access. [draft-irtf-pearg-censorship]. Protocol mechanisms such as Encrypted Client Hello [I-D.ietf-tls-esni] or DNS over HTTPS [RFC8484] that encrypt metadata provide some level of resistance to this type of protocol inspection. Full traffic encryption systems such as Tor [https://torproject.org] can also be used by people access otherwise censored resources.

Example: As noted above, one way to censor Web traffic is to require the server to block it or require internet service providers to block requests to the server. In HTTP, denial or restriction of access can be made apparent by the use of status code 451, which allows server operators and intermediaries to operate with greater transparency in circumstances where issues of law or public policy affect their operation [RFC7725]. If a protocol potentially enables censorship, protocol designers should strive towards creating error codes that capture different scenarios (blocked due to

administrative policy, unavailable because of legal requirements,
etc.) to minimize ambiguity for end-users.

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

4.17. Outcome Transparency

Question(s): Are the intended and forseen effects of your protocol
documented and easily comprehensible?

Explanation: Certain technical choices may have unintended
consequences. Have you described the central use case(s) for your
protocol with a clear description of expected behavior and how it
may, or may not, impact other protocols, implementations, user
expectations, or behavior? Have you reviewed other protocols that
solve similar problems, or make use of similar mechanisms, to see if
there are lessons that can be learnt from their use and misuse?

Example: Lack of authenticity may lead to lack of integrity and
negative externalities, of which spam is an example. Lack of data
that could be used for billing and accounting can lead to so-called
"free" arrangements which obscure the actual costs and distribution
of the costs, for example the barter arrangements that are commonly
used for Internet interconnection; and 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. Unexpected outcomes might not be
technical, but rather architectural, social or economic. Therefore
it is of importance to document the intended outcomes and other
possible outcomes that have been considered.

Impacts:

  *Right to freedom of expression

  *Right to privacy

  *Right to freedom of assembly and association

  *Right to access to information

4.18. Accessibility

Question(s): Is your protocol designed to provide an enabling environment for all? Have you looked at the W3C Web Accessibility Initiative for examples and guidance?

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

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

Another example is the work done in the AVT and AVTCORE working groups in the IETF that enables text conversation in multimedia, text telephony, wireless multimedia and video communications for sign language and lip-reading (i.e., [RFC9071]).

Impacts:

  *Right to non-discrimination

  *Right to freedom of assembly and association

  *Right to education

  *Right to political participation

4.19. Decentralization

Question(s): Can your protocol be implemented without a single point of control? If applicable, can your protocol be deployed in a federated manner? Does your protocol create additional centralized points of control?

Explanation: Decentralization is one of the central technical concepts of the architecture of the Internet, 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 even if one or several nodes are disabled. With the commercialization of the

Internet in the early 1990s, there has been a slow move away from decentralization, to the detriment of the technical benefits of having a decentralized Internet. For a more detailed discussion of this topic, please see [arkkoetal].

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 the increased centralization of the network that creates central infrastructure points that can be tapped into. The creation of peer-to-peer networks and the development of voice-over-IP protocols using peer-to-peer technology in combination with 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

4.20. Remedy

Question(s): Can your protocol facilitate a negatively impacted party's right to remedy without disproportionately impacting other parties' human rights, especially their right to privacy?

Explanation: Providing access to remedy by states and corporations is a part of the UN Guiding Principles on Business and Human Rights [UNGP]. Access to remedy may help victims of human rights violations in seeking justice, or allow law enforcement agencies to identify a possible violator. However, current mechanisms in protocols that try to enable 'attribution' to individuals impede the exercise of the right to privacy. The former UN Special Rapporteur for Freedom of Expression has also argued that anonymity is an inherent part of freedom of expression [Kaye]. Considering the potential adverse impact of attribution on the right to privacy and freedom of expression, enabling attribution on an individual level is most likely not consistent with human rights.

Example: Adding personally identifiable information to data streams as a means to enable the human right to remedy might help in identifying a violator of human rights and provide access to remedy, but this would disproportionally affect all users right to privacy, anonymous expression, and association. Furthermore, there are some recent advances in enabling abuse detection in end-to-end encrypted messaging systems, which also carry some risk to users' privacy [messenger-franking][hecate].

Impacts:

 *Right to remedy

 *Right to security

 *Right to privacy

4.21. Misc. considerations

Question(s): Have you considered potential negative consequences
(individual or societal) that your protocol or document might have?

Explanation: Publication of a particular RFC under a certain status
has consequences. Publication as an Internet Standard as part of the
Standards Track may signal to implementers that the specification
has a certain level of maturity, operational experience, and
consensus. Similarly, publication of a specification an experimental
document as part of the non-standards track would signal to the
community that the document "may be intended for eventual
standardization but [may] not yet [be] ready" for wide deployment.
The extent of the deployment, and consequently its overall impact on
end-users, may depend on the document status presented in the RFC.
See [BCP9] and updates to it for a fuller explanation.

5. Document Status

This RG document lays out best practices and guidelines for human
rights reviews of network protocols, architectures and other
Internet-Drafts and RFCs.

6. Acknowledgements

Thanks to:

 *Corinne Cath-Speth for work on [RFC8280].

 *Reese Enghardt, Joe Hall, Avri Doria, Joey Salazar, Corinne Cath-
  Speth, Farzaneh Badii, Sandra Braman, Colin Perkins, John Curran,
  Eliot Lear, Mallory Knodel, Brian Trammell, Jane Coffin, Eric
  Rescorla, Sofía Celi and the hrpc list for reviews and
  suggestions.

 *Individuals who conducted human rights reviews for their work and
  feedback: Amelia Andersdotter, Shane Kerr, Beatrice Martini,
  Karan Saini, and Shivan Kaul Sahib.

7. Security Considerations

Article three of the Universal Declaration of Human Rights reads: "Everyone has the right to life, liberty and security of person.". This article underlines the importance of security and its interrelation with human life and liberty, but since human rights are indivisible, interrelated and interdependent, security is also closely linked to other human rights and freedoms. This document seeks to strengthen human rights, freedoms, and security by relating and translating these concepts to concepts and practices as they are used in Internet protocol and architecture development. The aim of this is to secure human rights and thereby improve the sustainability, usability, and effectiveness of the network. The document seeks to achieve this by providing guidelines as done in section three of this document.

8. IANA Considerations

This document has no actions for IANA.

9. Research Group Information

The discussion list for the IRTF Human Rights Protocol Considerations Research Group is located at the e-mail address hrpc@ietf.org. Information on the group and information on how to subscribe to the list is 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

[arkkoetal]
          Arkko, J., Trammell, B., Notthingham, M., Huitema, C., Thomson, M., Tantsure, J., and N. ten Oever, "Considerations on Internet Consolidation and the Internet Architecture", 2019, <https://datatracker.ietf.org/doc/html/draft-arkko-iab-internet-consolidation-02>.

[BCP72]   IETF, "Guidelines for Writing RFC Text on Security Considerations", 2003, <https://datatracker.ietf.org/doc/bcp72/>.

[BCP9]   Bradner, S. and IETF, "The Internet Standards Process -- Revision 3", 1996, <https://datatracker.ietf.org/doc/rfc2026/>.

[Bless]   Bless, R. and C. Orwat, "Values and Networks", 2015.

[Brown]
Brown, I. and M. Ziewitz, "A Prehistory of Internet Governance", Research Handbook on Governance of the Internet. Cheltenham, Edward Elgar , 2013.

[draft-ietf-ohai-ohttp] Thomson, M. and C.A. Wood, "Oblivious DNS Over HTTPS", 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-ohai-ohttp>.

[draft-irtf-pearg-censorship] Hall, J., Aaron, M., Adams, S., Jones, B., and N. Feamster, "A Survey of Worldwide Censorship Techniques", 2020, <https://tools.ietf.org/html/draft-irtf-pearg-censorship>.

[draft-zuniga-mac-address-randomization] Zuniga, J.C., Bernardos, C.J., and A. Andersdotter, "MAC address randomization", 2020, <https://tools.ietf.org/html/draft-ietf-madinas-mac-address-randomization>.

[email-hashing] Acar, G., Englehardt, S., and A. Narayanan, "Four cents to deanonymize: Companies reverse hashed email addresses", n.d., <https://freedom-to-tinker.com/

2018/04/09/four-cents-to-deanonymize-companies-reverse-hashed-email-addresses/>.

[FIArch]   "Future Internet Design Principles", January 2012, <http://www.future-internet.eu/uploads/media/FIArch_Design_Principles_V1.0.pdf>.

[FREAK]   "Tracking the FREAK Attack", 2015, <https://web.archive.org/web/20150304002021/https://freakattack.com/>.

[geekfeminism] Geek Feminism Wiki, "Pseudonymity", 2015, <http://geekfeminism.wikia.com/wiki/Pseudonymity>.

[hecate]   Issa, R., Alhaddad, N., and M. Varia, "Hecate, Abuse Reporting in Secure Messengers with Sealed Sender", 2022, <https://eprint.iacr.org/2021/1686>.

[Hill2014]   Hill, R., "Partial Catalog of Human Rights Related to ICT Activities", 2014, <http://www.apig.ch/UNIGE%20Catalog.pdf>.

[HR-RT]   "Human Rights Reviews", 2022, <https://github.com/IRTF-HRPC/reviews>.

[HTML5]   W3C, "HTML5", 2014, <https://www.w3.org/TR/html5/>.

[https-interception]
             Durumeric, Z., Ma, Z., Springall, D., Barnes, R., Sullivan, N., Bursztein, E., Bailey, M., Halderman, J., and V. Paxson, "The Security Impact of HTTPS Interception", 2017.

[I-D.ietf-mls-protocol]
             Barnes, R., Beurdouche, B., Robert, R., Millican, J., Omara, E., and K. Cohn-Gordon, "The Messaging Layer Security (MLS) Protocol", 2023, <https://datatracker.ietf.org/doc/draft-ietf-mls-protocol/>.

[I-D.ietf-tls-esni] Rescorla, E., Oku, K., Sullivan, N., and C. A. Wood, "TLS Encrypted Client Hello", Work in Progress, Internet-Draft, draft-ietf-tls-esni-17, 9 October 2023,

<https://datatracker.ietf.org/doc/html/draft-ietf-tls-esni-17>.

[ICCPR]   United Nations General Assembly, "International Covenant on Civil and Political Rights", 1976, <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>.

[IRP]   Internet Rights and Principles Dynamic Coalition, "10 Internet Rights & Principles", 2014, <http://internetrightsandprinciples.org/site/wp-content/uploads/2014/06/IRPC_10RightsandPrinciples_28May2014-11.pdf>.

[Kaye]   Kaye, D., "The use of encryption and anonymity in digital communications", 2015, <https://www.ohchr.org/EN/HRbodies/HRC/RegularSessions/Session29/Documents/A.HRC.29.32_AEV.doc>.

[Logjam]   Adrian, D., Bhargavan, K., and . et al, "Imperfect Forward Secrecy, How Diffie-Hellman Fails in Practice", 2015, <https://weakdh.org/imperfect-forward-secrecy-ccs15.pdf>.

[messenger-franking] Grubbs, P., Lu, J., and T. Ristenpart, "Message Franking via Committing Authenticated Encryption", 2017, <https://eprint.iacr.org/2017/664>.

[newegg]   Mullin, J., "Newegg on trial: Mystery company TQP rewrites the history of encryption", 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] W3C, "W3C Patent Policy", 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>.

[Pouwelse] Pouwelse, Ed, J., "Media without censorship", 2012, <https://tools.ietf.org/html/draft-pouwelse-censorfree-scenarios>.

[RFC0793]   Postel, J., "Transmission Control Protocol", RFC 793, DOI
            10.17487/RFC0793, September 1981, <https://www.rfc-
            editor.org/info/rfc793>.

[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>.

[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>.

[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>.

[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>.

[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>.

[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>.

[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>.

[RFC6235]   Boschi, E. and B. Trammell, "IP Flow Anonymization
             Support", RFC 6235, DOI 10.17487/RFC6235, May 2011,
             <https://www.rfc-editor.org/info/rfc6235>.

[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>.

[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>.

[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>.

[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>.

[RFC7301]   Friedl, S., Popov, A., Langley, A., and E. Stephan,
             "Transport Layer Security (TLS) Application-Layer

Protocol Negotiation Extension", RFC 7301, DOI 10.17487/
RFC7301, July 2014, <https://www.rfc-editor.org/info/
rfc7301>.

[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>.

[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>.

[RFC7844]   Huitema, C., Mrugalski, T., and S. Krishnan, "Anonymity
Profiles for DHCP Clients", RFC 7844, DOI 10.17487/
RFC7844, May 2016, <https://www.rfc-editor.org/info/
rfc7844>.

[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>.

[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>.

[RFC8280]   ten Oever, N. and C. Cath, "Research into Human Rights
Protocol Considerations", RFC 8280, DOI 10.17487/RFC8280,
October 2017, <https://www.rfc-editor.org/info/rfc8280>.

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

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

[RFC8558]   Hardie, T., Ed., "Transport Protocol Path Signals", RFC
8558, DOI 10.17487/RFC8558, April 2019, <https://www.rfc-
editor.org/info/rfc8558>.

[RFC8890]   Nottingham, M., "The Internet is for End Users", RFC
              8890, DOI 10.17487/RFC8890, August 2020, <https://
              www.rfc-editor.org/info/rfc8890>.

[RFC8980]   Arkko, J. and T. Hardie, "Report from the IAB Workshop on
              Design Expectations vs. Deployment Reality in Protocol
              Development", RFC 8980, DOI 10.17487/RFC8980, February
              2021, <https://www.rfc-editor.org/info/rfc8980>.

[RFC9000]   Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based
              Multiplexed and Secure Transport", RFC 9000, DOI
              10.17487/RFC9000, May 2021, <https://www.rfc-editor.org/
              info/rfc9000>.

[RFC9071]   Hellström, G., "RTP-Mixer Formatting of Multiparty Real-
              Time Text", RFC 9071, DOI 10.17487/RFC9071, July 2021,
              <https://www.rfc-editor.org/info/rfc9071>.

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

[UDHR]   United Nations General Assembly, "The Universal
              Declaration of Human Rights", 1948, <http://www.un.org/
              en/documents/udhr/>.

[UNGP]   United Nations, "United Nations Guiding Principles on
              Business and Human Rights", 2011, <https://www.ohchr.org/
              documents/publications/
              guidingprinciplesbusinesshr_en.pdf>.

[UNHR]   United Nations, "The Core International Human Rights
              Instruments and their monitoring bodies", 2011, <https://
              www.ohchr.org/en/professionalinterest/pages/
              coreinstruments.aspx>.

[UNHRC2016]   United Nations Human Rights Council, "UN Human Rights
              Council Resolution "The promotion, protection and
              enjoyment of human rights on the Internet" (A/HRC/32/L.

20)", 2016, <https://documents-dds-ny.un.org/doc/UNDOC/
LTD/G16/131/89/PDF/G1613189.pdf?OpenElement>.

[W3CAccessibility] W3C, "Accessibility", 2015, <https://www.w3.org/
standards/webdesign/accessibility>.

[W3Ci18nDef] W3C, "Localization vs. Internationalization", 2010,
<http://www.w3.org/International/questions/qa-i18n.en>.

[Zittrain] Zittrain, J., "The Future of the Internet - And How to
Stop It", Yale University Press , 2008, <https://
dash.harvard.edu/bitstream/handle/1/4455262/
Zittrain_Future%20of%20the%20Internet.pdf?sequence=1>.

Authors' Addresses

Gurshabad Grover

Email: gurshabad@cis-india.org

Niels ten Oever
University of Amsterdam

Email: mail@nielstenoever.net