Human Rights Protocol Considerations Research Group N. ten Oever (editor)
Internet-Draft ARTICLE 19
Intended status: Informational November 12, 2017
Expires: May 16, 2018

Guidelines for Human Rights Protocol Considerations
draft-tenoever-hrpc-guidelines-01

Abstract

This document proposes guidelines for human rights considerations, 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].

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on May 16, 2018.

Table of Contents

1. Introduction

2. Vocabulary used

3. Model for developing human rights protocol considerations

This section 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 their human rights impact. It should however 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 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 impact human rights and how these technical concept 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.

3.1. Human rights threats

Human rights threats on the Internet come in a myriad of forms. Protocols and standards can harm or enable the right to freedom of expression, right to non-discrimination, right to equal protection, right to participate in cultural life, arts and science, 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 the malpractices of a government or other authority. A person whose communications are monitored may be prevented 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 or extrajudicial killing or detention on the basis of information gathered by state agencies through information leakage in protocols.

This section details several ’common’ threats to human rights, indicating how each of these can lead to human rights violations/ harms and present 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 the 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 the security considerations [BCP72], but privacy guidelines [RFC6973] or reviews, let alone human rights impact assessments of protocols are not standardized or implemented.

Many threats, enablers and risks are linked to different rights. This is not unsurprising 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 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 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.

3.2. 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, but rather contribute to them and look at the design process from a human rights perspective.

Protocols and Internet Standard might benefit from a documented discussion of potential human rights risks arising from potential misapplications of the protocol or technology described in the 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 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.

3.2.1. Connectivity

Question(s): 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 other 'services' 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 and 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

3.2.2. Privacy

Question(s): Did you have a look at the Guidelines in the Privacy Considerations for Internet Protocols [RFC6973] section 7? 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 logged data 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 in 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

3.2.3. Content agnosticism

Question(s): If your protocol impacts packet handling, does it use user data (packet data that is not included in the header)? Is it making 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 exception where it comes to effective traffic handling, for instance where it comes to 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 on the basis of 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

3.2.4. 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 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: 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

3.2.5. Internationalization

Question(s): 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 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 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 CCS 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 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 are both content and protocol elements, such as the identifiers.) If the Internet wants to be a global network of networks, the protocols should work with other languages than English and other character sets 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 localization

Impacts:

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

3.2.6. Censorship resistance

Question(s): Does this protocol introduce new identifiers or reuse existing identifiers (e.g. MAC addresses) that might be associated with persons or content? Does your protocol make it apparent or transparent when access to a resource it restricted? Can your

protocol contribute to filtering in a way it could be implemented to censor data or services? Could this be designed to ensure this doesn't happen?

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

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. [RFC4941] This is why Privacy Extensions for Stateless Address Autoconfiguration in IPv6 have been introduced. [RFC4941]

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 - which allows 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

3.2.7. 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 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 it possible be done without)? Are you aware of any patents that would prevent your standard from being fully implemented [RFC3979] [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 '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 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 speficiations produced by others 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 reasonable protection. Reasonable patent protection should includes but is not limited to cryptographic primitives.

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 to DPI, this document describes a system that does not rely upon DPI, and is instead based in open IETF standards and open source applications.

Impacts:

- Right to freedom of expression

- Right to participate in cultural life, arts and science

3.2.8. 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? Does your protocol allow there to be well-defined extension points? Do these extension points allow for open innovation?

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 Internet service providers, 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].

Example: Heterogeneity is inevitable and needs be supported by design. Multiple types of hardware must be allowed for, e.g. transmission speeds differing by at least 7 orders of magnitude, various computer word lengths, and hosts ranging from memory-starved microprocessors up to massively parallel supercomputers. 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 [RFC1958].

Impacts:

- Right to freedom of expression

- Right to political participtation

3.2.9. Pseudonymity

Question(s): Have you considered the Privacy Considerations for Internet Protocols [RFC6973], especially section 6.1.2 ? Does the protocol collect personally derived data? Does the protocol generate or process anything that can be, or 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 their device or address? Does this protocol generate personally derived data, and if so how will that data be handled?

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

Example: Designing a standard that exposes personal data, it is important to consider ways to mitigate the obvious impacts. While pseudonyms cannot be simply reverse engineered - some early approaches simply took approaches such as simple hashing of IP addreses, these could then be simply 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 pseudoyms, for instance to: hide their gender, protect themselves against harassment, protect their families’ privacy, frankly discuss sexuality, or develop a artistic or journalistic persona without retribution from an employer, (potential) customers, or social surrounding. [geekfeminism] The difference between anonymity and pseudonymity is that a pseudonym often is 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

3.2.10. Accessibility

Question(s): 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 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 ability [W3CAccessibility]. Sometimes in the design of protocols, websites, web technologies, or web tools, barriers are created that exclude people from using the Web.

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.

Impacts:

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

3.2.11. Localization

Question(s): Does your protocol uphold the standards of internationalization? Have 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]. 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 Internationalization).

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 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 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]. This protocol labels the information content with an identifier for the language in which it is written. And this allows 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

3.2.12. Decentralization

Question(s): 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? How can the use of your protocol be used to implicate users? Does your protocol create additional centralized points of control?

Explanation: Decentralization is one of the central technical concepts of the architecture of the networks, and 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 1990’s there has been a slow move to move away from decentralization, to the detriment of the technical benefits of having a decentralized Internet.

Example: The bits traveling the Internet are increasingly susceptible to monitoring and censorship, from both governments and Internet service providers, 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 in to. 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

3.2.13. Reliability

Question(s): Is your protocol fault tolerant? Does it 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 and error resistant as described, and function without unexpected result. 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. 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 is good. From a security perspective, this can be very bad. As with confidentiality, the growth of the Internet and fostering innovation in services depends on users having confidence and trust [RFC3724] in the network. 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], 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

3.2.14. Confidentiality

Question(s): Does this protocol expose information related to identifiers or data? If so, does it do so to each other 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 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 clear text 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 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].

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 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

3.2.15. 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, like 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 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 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

3.2.16. 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, 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.

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

Impacts:

- Right to privacy

- Right to freedom of expression

- Right to security

3.2.17. Adaptability

Question(s): Is your protocol written in such a way that is 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 '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. 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.

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 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

- Freedom of expression

- Freedom of assembly and association

3.2.18. Outcome Transparency

Question(s): Are the effects of your protocol fully and easily comprehensible, including with respect to unintended consequences of protocol choices?

Explanation: certain technical choice may have unintended consequences.

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.

Impacts: - Freedom of expression - Privacy - Freedom of assembly and association - Access to information

3.2.19. Anonymity

Example: Often protocols expose personal data, it is important to consider ways to mitigate the obvious privacy impacts. 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 virtual private network) can be helpful to protect from third parties likely to eavesdrop packets exchanged between the tunnel endpoints.

Question(s): Does you protocol make use of persistent identifiers? Can it be done without them? If your protocol collects data and distributes it (see [RFC6235]), you should anonymize the data, but keep in mind that "anonymizing" data is notoriously hard. Do not think that just dropping the last byte of an IP address "anonymizes" data. If your protocol allows for identity management, there should be a clear barrier between the identities to ensure that they cannot (easily) be associated with each other. Did you have a look at the Privacy Considerations for Internet Protocols [RFC6973], especially section 6.1.1 ?

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. Anonymity is an inherent part of the right to freedom of opinion and expression and the right to privacy. Avoid adding identifiers, options or configurations that create or might lead to patterns or regularities that are not explicitely required by the protocol.

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

Impacts:

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

4. Document Status

5. Acknowledgements

6. Security Considerations

As this document concerns a research document, there are no security considerations.

7. IANA Considerations

This document has no actions for IANA.

8. 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 [1]. 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

9. References

9.1. Informative References

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

[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.

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

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

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

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

[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>.

[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", STD 7, 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>.

[RFC3979] Bradner, S., Ed., "Intellectual Property Rights in IETF Technology", RFC 3979, DOI 10.17487/RFC3979, March 2005, <https://www.rfc-editor.org/info/rfc3979>.

[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>.

[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>.

[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>.

[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>.

[Saltzer] Saltzer, J., Reed, D., and 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/>.

[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>.

9.2. URIs

[1] mailto:hrpc@ietf.org

Author’s Address

Niels ten Oever
ARTICLE 19

EMail: mail@nielstenoever.net