

**The Internet is for End Users  
draft-nottingham-for-the-users-02**

Abstract

Internet standards serve and are used by a variety of communities. This document contains guidelines for explicitly identifying them, serving them, and determining how to resolve conflicts between their interests, when necessary.

It also motivates considering end users as the highest priority concern for Internet standards.

Note to Readers

The issues list for this draft can be found at <https://github.com/mnot/I-D/labels/for-the-users> .

The most recent (often, unpublished) draft is at <https://mnot.github.io/I-D/for-the-users/> .

Recent changes are listed at <https://github.com/mnot/I-D/commits/gh-pages/for-the-users> .

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 April 21, 2016.

## Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

<a href="#">1.</a>	<a href="#">Introduction</a>	<a href="#">2</a>
<a href="#">1.1.</a>	<a href="#">Notational Conventions</a>	<a href="#">4</a>
<a href="#">2.</a>	<a href="#">The Internet is for End Users</a>	<a href="#">4</a>
<a href="#">3.</a>	<a href="#">Identifying Relevant Parties</a>	<a href="#">5</a>
<a href="#">3.1.</a>	<a href="#">Handling Change in Relevant Parties</a>	<a href="#">6</a>
<a href="#">3.2.</a>	<a href="#">Avoiding Unnecessary Parties</a>	<a href="#">6</a>
<a href="#">4.</a>	<a href="#">IANA Considerations</a>	<a href="#">7</a>
<a href="#">5.</a>	<a href="#">Security Considerations</a>	<a href="#">7</a>
<a href="#">6.</a>	<a href="#">References</a>	<a href="#">7</a>
<a href="#">6.1.</a>	<a href="#">Normative References</a>	<a href="#">7</a>
<a href="#">6.2.</a>	<a href="#">Informative References</a>	<a href="#">7</a>
<a href="#">Appendix A.</a>	<a href="#">Acknowledgements</a>	<a href="#">8</a>
	<a href="#">Author's Address</a>	<a href="#">9</a>

## [1.](#) Introduction

As the Internet has become prevalent in many societies, it has also unavoidably become a profoundly political thing; it has helped overthrow governments, revolutionize social orders, control populations and reveal people's secrets. It has created wealth for some individuals and companies, while destroying others'.

The IETF, while focused on technical matters, is not neutral about the purpose of its work [[RFC3935](#)]:

The IETF community wants the Internet to succeed because we believe that the existence of the Internet, and its influence on economics, communication, and education, will help us to build a better human society.



However, the IETF is most comfortable making purely technical decisions; our process is defined to favor technical merit, through our well-known bias towards "rough consensus and running code".

Nevertheless, the running code that results from our process (when things work well) inevitably has an impact beyond technical considerations, because the underlying decisions afford some uses, while discouraging others. Or, in the words of Lawrence Lessig [[CODELAW](#)]:

Ours is the age of cyberspace. It, too, has a regulator... This regulator is code -- the software and hardware that make cyberspace as it is. This code, or architecture, sets the terms on which life in cyberspace is experienced. It determines how easy it is to protect privacy, or how easy it is to censor speech. It determines whether access to information is general or whether information is zoned. It affects who sees what, or what is monitored. In a host of ways that one cannot begin to see unless one begins to understand the nature of this code, the code of cyberspace regulates.

All of this raises the question: Who do we go through the pain of rough consensus and write the running code for?

There are a variety of identifiable parties in the larger Internet community that standards can provide benefit to, such as (but not limited to) end users, network operators, schools, equipment vendors, specification authors, specification implementers, content owners, governments, non-governmental organisations, social movements, employers, and parents.

Successful specifications will provide some benefit to all of the relevant parties, because standards do not represent a zero-sum game. However, there are often situations where we need to balance the benefits of a decision between two (or more) parties.

We regularly decide to take up work against those who attempt to use the Internet for goals that we do not believe are beneficial; for example, those who attempt to disrupt Internet access (denial-of-service attackers) and those who seek to obtain data or control over a system that is not authorised by its administrator.

Additionally, efforts are sometimes brought to the IETF that represent the needs of some parties but at the expense of others. When presented with such a proposal, we need to decide how to handle it.



Currently, these kinds of decisions occur in an ad hoc fashion, often without explicitly being discussed. This approach works reasonably well in many cases; even if a party is not directly represented in the process, there are often advocates for their interests, and ultimately protocols that disadvantage a particular party tend to be either rejected by it or eventually replaced.

However, we do sometimes expend a considerable amount of energy mitigating potential harm to under-represented members of the Internet community, and often such harm is not so onerous or obvious as to dissuade them from using something (e.g., [[RFC6265](#)]).

In other words - because our decisions have ethical implications, we should consider their impact and determine whether it is within our core values, and do so in a well-defined, open fashion.

To facilitate that, [Section 3](#) outlines a set of guidelines for identifying the relevant parties to an Internet standard. The aim of doing so is to both clarify the decision-making process, and to aid external parties when engaging with and judging the results of the standards process.

In doing so, it becomes clear that Internet standards that give the highest priority to end users have the best chance of success, and of helping the IETF to succeed in its mission. As a result, [Section 2](#) mandates that other parties cannot have a higher priority.

### **[1.1.](#) Notational Conventions**

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

## **[2.](#) The Internet is for End Users**

Internet standards MUST NOT consider any other party to have higher priority over end users.

While networks need to be managed, employers and equipment vendors need to meet business goals, etc., the IETF's mission is to "build a better human society" [[RFC3935](#)] and - on the Internet - society is composed of what we call "end users."

Furthermore, the success of the Internet to date is arguably due largely to its bias towards end user concerns; without a firm preference for their benefit, trust in the Internet will erode, and its value - for everyone - will be greatly diminished.



This does not mean that end users have ultimate priority; there may be cases where genuine technical need of another party requires that end user requirements compromise. However, such tradeoffs need to be carefully examined, and avoided when there are alternate means of achieving the desired goals. If they cannot be, these choices and reasoning SHOULD be carefully documented.

For example, IPv6 [[RFC2460](#)] identifies each client with a unique address - even though this provides a way to track end user activity and helps identify them - because it is technically necessary to provide networking (and despite this, there are mechanisms like [[RFC4941](#)] to mitigate this effect, for those users who desire it).

This also does not mean that the IETF community has any specific insight into what is "good for end users"; as before, we will need to interact with the greater Internet community and apply our process to help us make decisions, deploy our protocols, and ultimately determine their success or failure.

### **3. Identifying Relevant Parties**

The relevant parties to an Internet standard MUST be documented, along with their interrelationships.

For example, HTML does so using the "priority of constituencies" in the HTML Design Principles [[PRIORITY](#)]:

In case of conflict, consider users over authors over implementors over specifiers over theoretical purity. In other words costs or difficulties to the user should be given more weight than costs to authors; which in turn should be given more weight than costs to implementors; which should be given more weight than costs to authors of the spec itself, which should be given more weight than those proposing changes for theoretical reasons alone. Of course, it is preferred to make things better for multiple parties at once.

Note how the relative priority is explicit; this is intentional and encouraged. However, it need not be a strict ranking in all cases; in some areas, it can be more useful to give equal weight to parties, so as to encourage the tussle [[TUSSLE](#)].

Likewise, the responsibilities of, or expectations upon, different parties to a standard can vary greatly. For example, end users of Web browsers cannot be reasonably expected to make informed decisions about security, and therefore design decisions there are biased towards default security. When applicable, the expectations upon a party SHOULD be documented.





Extensions to existing standards **MUST** consider how they interact with the extended standard's relevant parties. If they are not yet documented, this **SHOULD** be done in coordination with that standard's community and the IESG.

The burden of this documentation need not be high; if HTML can do it in a paragraph, so can most other standards. While it might be appropriate in a separate document (e.g., a requirements or use cases draft) or the specification itself, documenting relevant parties in the WG charter has considerable benefits, since it clarifies their relationships up-front.

Inevitably, documenting and interpreting these roles will become controversial; this is to be expected, and is still preferable to avoiding the discussion. The point is to make it explicit, so that the affected parties can be made aware of the discussion, and judge the outcome.

### **3.1. Handling Change in Relevant Parties**

Changes in the use, deployment patterns, legal context, or other factors of a standard can bring pressure to re-balance the priorities of existing parties, or insert new ones (usually, when a standard is either extended or evolved).

Such changes **MUST NOT** diminish the priority of existing relevant parties without informed consent. Note that this may preclude the change completely, as it is often impossible to gain the informed consent of a large or diffuse group (e.g., end users).

For example, there has been increasing pressure to change HTTP [[RFC7230](#)] to make it more amenable to optimization, filtering, and interposition of other value-added services, especially in the face of wider use of encryption (through HTTPS URIs). However, since HTTPS is already defined as a two-party protocol with end-to-end encryption, inserting a third party in any fashion would violate the expectations of two existing parties; end users and content publishers. Therefore, the HTTP Working Group has refused to consider such changes.

### **3.2. Avoiding Unnecessary Parties**

In protocol design, intermediation is often thought of as "those parties on the direct path between two people attempting to communicate"; e.g., middleboxes, proxies and so on.

When discussing the parties relevant to an Internet standard, this definition can be expanded to include those parties that have the



ability to prevent or control communication between two parties. This naturally includes middleboxes, but can also include third parties not directly on-path.

For example, HTTP has on-path intermediaries (proxies, gateways, etc.), but also off-path intermediaries, in the form of the DNS registrar, the DNS server, and also the Certificate Authority if TLS is in use. Certificate Transparency [RFC6962] potentially adds yet another intermediary to this protocol suite.

While there might be a good technical reason to interpose such an intermediary, it also introduces a new party, and thus needs to be done with due consideration of the impact on other parties.

Therefore, unnecessary parties SHOULD be avoided when possible in Internet standards.

#### **4. IANA Considerations**

This document does not require action by IANA.

#### **5. Security Considerations**

This document does not have direct security impact; however, applying its guidelines (or failing to) might affect security positively or negatively.

#### **6. References**

##### **6.1. Normative References**

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.

##### **6.2. Informative References**

[CODELAW] Lessig, L., "Code Is Law: On Liberty in Cyberspace", 2000, <<http://harvardmagazine.com/2000/01/code-is-law-html>>.

[PRIORITY] van Kesteren, A. and M. Stachowiak, "HTML Design Principles", November 2007, <<http://www.w3.org/TR/html-design-principles/#priority-of-constituencies>>.

- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), DOI 10.17487/RFC2460, December 1998, <<http://www.rfc-editor.org/info/rfc2460>>.
- [RFC3935] Alvestrand, H., "A Mission Statement for the IETF", [BCP 95](#), [RFC 3935](#), DOI 10.17487/RFC3935, October 2004, <<http://www.rfc-editor.org/info/rfc3935>>.
- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", [RFC 4941](#), DOI 10.17487/RFC4941, September 2007, <<http://www.rfc-editor.org/info/rfc4941>>.
- [RFC6265] Barth, A., "HTTP State Management Mechanism", [RFC 6265](#), DOI 10.17487/RFC6265, April 2011, <<http://www.rfc-editor.org/info/rfc6265>>.
- [RFC6962] Laurie, B., Langley, A., and E. Kasper, "Certificate Transparency", [RFC 6962](#), DOI 10.17487/RFC6962, June 2013, <<http://www.rfc-editor.org/info/rfc6962>>.
- [RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", [RFC 7230](#), DOI 10.17487/RFC7230, June 2014, <<http://www.rfc-editor.org/info/rfc7230>>.
- [TUSSE] Clark, D., Sollins, K., Wroclawski, J., and R. Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet", 2002, <<http://groups.csail.mit.edu/ana/Publications/PubPDFs/Tussle2002.pdf>>.

## **Appendix A. Acknowledgements**

Thanks to Jacob Appelbaum for making the suggestion that led to this document.

Thanks also to the WHATWG for blazing the trail.

Thanks to Edward Snowden for his comments regarding the priority of end users at IETF93.

Thanks to Harald Alvestrand for his substantial feedback and Stephen Farrell, Joe Hildebrand, Russ Housley, Niels ten Oever, and Martin Thomson for their suggestions.



Author's Address

Mark Nottingham

Email: [mnot@mnot.net](mailto:mnot@mnot.net)

URI: <http://www.mnot.net/>