

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: 2 January 2021

B.E. Carpenter
Univ. of Auckland
S. Farrell
Trinity College Dublin
1 July 2020

Additional Criteria for Nominating Committee Eligibility
draft-carpenter-eligibility-expand-03

Abstract

This document defines a process experiment under RFC 3933 that temporarily updates the criteria for qualifying volunteers to participate in the IETF Nominating Committee. It therefore also updates the criteria for qualifying signatories to a community recall petition. The purpose is to make the criteria more flexible in view of increasing remote participation in the IETF and a probable decline in face-to-face meetings. This document temporarily varies the rules in RFC 8713.

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 <https://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 2 January 2021.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

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

1. Introduction	2
2. Term of the Experiment	3
3. Goals	3
4. Criteria	3
5. Open Questions	4
6. Available data	4
7. Possible Future Work	5
8. IANA Considerations	5
9. Security Considerations	5
10. Acknowledgements	5
11. Normative References	5
Appendix A. Change Log	6
A.1. Draft-02 to -03	6
A.2. Draft-01 to -02	6
A.3. Draft-00 to -01	6
Authors' Addresses	6

1. Introduction

According to [RFC8713], the IETF Nominating Committee is populated from a pool of volunteers with a specified record of attendance at IETF plenary face-to-face meetings. In view of the unexpected cancellation of the IETF 107 meeting, the risk of future cancellations, the probability of less frequent meetings in future in support of sustainability, and a general increase in remote participation, this document defines a process experiment [RFC3933] of fixed duration to use additional criteria to qualify volunteers.

Also according to [RFC8713], the qualification for signing a community petition for the recall of certain IETF office-holders is that same as for the Nominating Committee. This document does not change that, but see Section 7.

The source for this is at <https://github.com/sftcd/elig/> and PRs are welcome there. Discussion on the eligibility-discuss@ietf.org list is also welcome.

2. Term of the Experiment

The cancellation of the in-person IETF 107 meeting, and the risk of IETF 108 also being cancelled, mean that the current criteria are in any case seriously perturbed for the next two years. The experiment therefore needs to start as soon as possible. However, the experiment does not apply to the selection of the 2020-2021 Nominating Committee.

The experiment will cover the two IETF Nominating Committee cycles starting in 2021 and 2022. As soon as the 2022-2023 Nominating Committee is seated, the IESG must consult the Nominating Committee chairs involved and publish a report on the results of the experiment. The IESG must then also begin a community discussion of whether to amend [RFC8713] in time for the 2023 Nominating Committee cycle.

3. Goals

The goals of the additional criteria are as follows:

- * Mitigate the issue of active remote (or rarely in-person) participants being disenfranchised in the NomCom and recall processes.
- * Prepare for an era in which face-to-face plenary meetings are less frequent (thus extending the issue to many, perhaps a majority, of participants).
- * Ensure that those eligible are true "participants" with enough current understanding of IETF practice and people to make informed decisions.
- * The criteria must be algorithmic so that the Secretariat can check them mechanically.

4. Criteria

There will be several alternative paths to qualification, replacing the single criterion in section 4.14 of [RFC8713]. Any one of the paths is sufficient, unless the person is otherwise disqualified under section 4.15 of [RFC8713]:

- * Path 1: As per [RFC8713], the person has attended 3 out of the last 5 in-person IETF meetings.

- This criterion has been retained for backward compatibility. It has not been extended to include remote participation in meetings, since the additional criteria below are intended to capture significant participation without regard to meeting attendance.
- * Path 2: Has been a WG Chair or Secretary, or a BOF Chair, within the last 3 years.
- * Path 3: (Feedback on path 3 from draft-01 was uniformly negative so ignore this, we'll leave placeholder text here for now just to avoid renumbering.)
- * Path 4: Has served in the IESG or IAB, or has been appointed to a formal role by the IESG or IAB, within the last 5 years.
- * Path 5: Has been a listed author of at least 2 IETF stream RFCs and/or WG-adopted drafts within the last 5 years.

5. Open Questions

- * Should we consider how many nomcom voting members qualify via which paths? For example, would it be ok if all 10 nomcom voting members qualified via path 4 in one year?
- * Certain criteria were rejected as not truly indicating effective IETF participation. These included authorship of individual Internet-Drafts, sending email to IETF lists, reviewing drafts, etc. Since the criteria must be objectively and mechanically measurable, no qualitative evaluation of an individual's contributions is considered.

6. Available data

An analysis of how some of the above criteria would affect the number of NomCom-qualified participants if applied in June 2020 has been performed. The results are presented below in Venn diagrams as Figure 1 to Figure 4. Note that the numbers shown may differ slightly from manual counts due to database mismatches.

See <https://github.com/sftcd/elig/blob/master/venn-2018-2019.pdf>

Figure 1: 2019 - 5 years

See <https://github.com/sftcd/elig/blob/master/venn-2018-2019.pdf>

Figure 2: 2019 - 3 years

See <https://github.com/sftcd/elig/blob/master/venn-2018-2019.pdf>

Figure 3: 2018 - 5 years

See <https://github.com/sftcd/elig/blob/master/venn-2018-2019.pdf>

Figure 4: 2018 - 3 years

7. Possible Future Work

- * Combined paths (e.g., a person who partly satisfies Path 2 and Path 5); otherwise known as a "points system". That seems to involve work/complexity either for the secretariat or for the volunteer.
- * Tweaking the "time decay" in each of the path definitions that ensures recent participation is more highly valued.
- * Separating the NomCom volunteer criteria from the recall petitioner criteria.

8. IANA Considerations

This document makes no request of IANA.

9. Security Considerations

This document should not affect the security of the Internet.

10. Acknowledgements

Useful comments were received from John Klensin, Warren Kumari, Michael Richardson, Martin Thomson, (to be completed)

The data analysis was mainly done by Robert Sparks.

11. Normative References

- [RFC3933] Klensin, J. and S. Dawkins, "A Model for IETF Process Experiments", BCP 93, RFC 3933, DOI 10.17487/RFC3933, November 2004, <<https://www.rfc-editor.org/info/rfc3933>>.

[RFC8713] Kucherawy, M., Ed., Hinden, R., Ed., and J. Livingood, Ed., "IAB, IESG, IETF Trust, and IETF LLC Selection, Confirmation, and Recall Process: Operation of the IETF Nominating and Recall Committees", BCP 10, RFC 8713, DOI 10.17487/RFC8713, February 2020, <<https://www.rfc-editor.org/info/rfc8713>>.

Appendix A. Change Log

A.1. Draft-02 to -03

- * Adjusted criteria according to comments received
- * Added data

A.2. Draft-01 to -02

- * Made this an RFC 3933 process experiment
- * Eliminated path based on directorate reviews, used to be: "Has submitted at least 6 reviews as a member of an official IETF review team within the last 3 years."
- * Other comments from IETF107 virtual gendispatch meeting handled

A.3. Draft-00 to -01

- * Added author

Authors' Addresses

Brian E. Carpenter
The University of Auckland
School of Computer Science
PB 92019
Auckland 1142
New Zealand

Email: brian.e.carpenter@gmail.com

Stephen Farrell
Trinity College Dublin
College Green
Dublin
Ireland

Email: stephen.farrell@cs.tcd.ie

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: 11 July 2021

B.E. Carpenter
Univ. of Auckland
S. Farrell
Trinity College Dublin
7 January 2021

Additional Criteria for Nominating Committee Eligibility
draft-carpenter-eligibility-expand-10

Abstract

This document defines a process experiment under RFC 3933 that temporarily updates the criteria for qualifying volunteers to participate in the IETF Nominating Committee. It therefore also updates the criteria for qualifying signatories to a community recall petition. The purpose is to make the criteria more flexible in view of increasing remote participation in the IETF and a reduction in face-to-face meetings. The experiment is of fixed duration and will apply to one, or at most two, consecutive Nominating Committee cycles, starting in 2021. This document temporarily varies the rules in RFC 8713.

Discussion Venues

This note is to be removed before publishing as an RFC.

Discussion of this document takes place on the ad hoc mailing list (eligibility-discuss@ietf.org), which is archived at <https://mailarchive.ietf.org/arch/browse/eligibility-discuss/> (<https://mailarchive.ietf.org/arch/browse/eligibility-discuss/>).

Source for this draft can be found at <https://github.com/sftcd/elig> (<https://github.com/sftcd/elig>).

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 <https://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 11 July 2021.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

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

- 1. Introduction 3
- 2. Term and Evaluation of the Experiment 3
- 3. Goals 4
- 4. Criteria 4
 - 4.1. Clarifying Detail 5
- 5. Omitted Criteria 6
- 6. IANA Considerations 6
- 7. Security Considerations 6
- 8. Acknowledgements 6
- 9. Normative References 6
- Appendix A. Available data 7
- Appendix B. Change Log 11
 - B.1. Draft-09 to -10 11
 - B.2. Draft-08 to -09 11
 - B.3. Draft-07 to -08 11
 - B.4. Draft-06 to -07 11
 - B.5. Draft-05 to -06 12
 - B.6. Draft-04 to -05 12
 - B.7. Draft-03 to -04 12
 - B.8. Draft-02 to -03 12
 - B.9. Draft-01 to -02 12
 - B.10. Draft-00 to -01 13
- Authors' Addresses 13

1. Introduction

According to [RFC8713], the IETF Nominating Committee (NomCom) is populated from a pool of volunteers with a specified record of attendance at IETF plenary meetings, which were assumed to be face-to-face meetings when that document was approved. In view of the cancellation of the IETF 107, 108, 109 and 110 face-to-face meetings, the risk of future cancellations, the probability of less frequent face-to-face meetings in future in support of sustainability, and a general increase in remote participation, this document defines a process experiment [RFC3933] of fixed duration (described in Section 2) to use modified and additional criteria to qualify volunteers.

During this experiment, the eligibility criteria for signing recall petitions - which [RFC8713] defines to be the same as those for NomCom eligibility - are consequently also modified as described in this document. This experiment has no other effect on the recall process.

2. Term and Evaluation of the Experiment

The cancellation of the in-person IETF 107 through 110 meetings means that the current criteria are in any case seriously perturbed for at least two years. The experiment therefore needs to start as soon as possible. However, the experiment did not apply to the selection of the 2020-2021 Nominating Committee, which was performed according to [RFC8788].

The experiment will initially cover the IETF Nominating Committee cycle that begins in 2021. As soon as the entire 2021-2022 Nominating Committee is seated, the IESG must consult the 2021-2022 Nominating Committee chair and the 2020-2021 Nominating Committee chair (who will maintain NomCom confidentiality) and publish a report on the results of the experiment. Points to be considered are whether the experiment has produced a sufficiently large and diverse pool of individuals, whether enough of those individuals have volunteered to produce a representative Nominating Committee with good knowledge of the IETF, and whether all the goals in Section 3 have been met. If possible, a comparison with results from the previous procedure (i.e., RFC 8713) should be made.

The IESG must then also begin a community discussion of whether to:

1. Amend [RFC8713] in time for the 2022-2023 Nominating Committee cycle; or

2. Prolong the current experiment for a second and final year with additional clarifications specific to the 2022-2023 cycle; or
3. Run a different experiment for the next nominating cycle; or
4. Revert to [RFC8713].

The IESG will announce the results of the consensus determination of this discussion in good time for the 2022-2023 Nominating Committee cycle to commence.

In the event of prolongation of this experiment for a second year, the IESG will repeat the consultation, report and community discussion process accordingly, but this document lapses at the end of the 2022-2023 Nominating Committee cycle.

3. Goals

The goals of the modified and additional criteria are as follows:

- * Mitigate the issue of active remote (or rarely in-person) participants being disenfranchised in the NomCom and recall processes.
- * Enable the selection of a 2021-2022 NomCom, and possibly a 2022-2023 NomCom, when it is impossible for anyone to have attended three out of the last five IETF meetings in person.
- * Prepare for an era in which face-to-face plenary meetings are less frequent (thus extending the issue to many, perhaps a majority, of participants).
- * Ensure that those eligible have enough current understanding of IETF practices and people to make informed decisions.
- * Provide algorithmic criteria, so that the Secretariat can check them mechanically against available data.

4. Criteria

This experiment specifies several alternative paths to qualification, replacing the single criterion in section 4.14 of [RFC8713]. Any one of the paths is sufficient, unless the person is otherwise disqualified under section 4.15 of [RFC8713]:

- * Path 1: The person has registered for and attended 3 out of the last 5 IETF meetings. For meetings held entirely online, online registration and attendance counts as attendance. For the

2021-2022 Nominating Committee, the meetings concerned will be IETF 106, 107, 108, 109, and 110. Attendance is as determined by the record keeping of the secretariat for in-person meetings, and based on being a registered person who logged in for at least one session of an online IETF meeting.

- * Path 2: Has been a Working Group Chair or Secretary within the 3 years prior to the day the call for NomCom volunteers is sent to the community.
- * Path 3: Has been a listed author or editor (on the front page) of at least 2 IETF stream RFCs within the last 5 years prior to the day the call for NomCom volunteers is sent to the community. An Internet-Draft that has been approved by the IESG and is in the RFC Editor queue counts the same as a published RFC, with the relevant date being the date the draft was added to the RFC Editor queue. For avoidance of doubt, the 5 year timer extends back to the date 5 years before the date when the call for NomCom volunteers is sent to the community.

Notes:

- * Path 1 corresponds approximately to [RFC8713], modified as per [RFC8788].
- * Path 3 includes approved drafts, since some documents spend a long time in the RFC Editor's queue.
- * Path 3 extends to 5 years because it commonly takes 3 or 4 years for new documents to be approved in the IETF stream, so 3 years would be too short a sampling period.
- * All the required data are available to the IETF Secretariat from meeting attendance records or the IETF data tracker.

4.1. Clarifying Detail

Path 1 does not qualify people who register and attend face-to-face meetings remotely. That is, it does not qualify remote attendees at IETF 106, because that meeting took place prior to any question of cancelling meetings.

If the IESG prolongs this experiment for a second year, as allowed by Section 2, the IESG must also clarify how Path 1 applies to IETF 111, 112 and 113.

5. Omitted Criteria

During community discussions of this document, certain criteria were rejected as not truly indicating effective IETF participation, or as being unlikely to significantly expand the volunteer pool. These included authorship of individual or Working-Group-adopted Internet-Drafts, sending email to IETF lists, reviewing drafts, acting as a BOF Chair, and acting in an external role for the IETF (liaisons etc.).

One path, service in the IESG or IAB within the last 5 years, was found to have no benefit since historical data show that such people always appear to be qualified by another path.

Since the criteria must be measurable by the Secretariat, no qualitative evaluation of an individual's contributions is considered.

6. IANA Considerations

This document makes no request of IANA.

7. Security Considerations

This document should not affect the security of the Internet.

8. Acknowledgements

Useful comments were received from Abdussalam Baryun, Alissa Cooper, Lars Eggert, Adrian Farrel, Bron Gondwana, Russ Housley, Chrsitian Huitema, Ben Kaduk, John Klensin, Victor Kuarsingh, Warren Kumari, Barry Leiba, Eric Rescorla, Michael Richardson, Rich Salz, Ines Robles, Martin Thomson and Magnus Westerlund.

The data analysis was mainly done by Robert Sparks. Carsten Bormann showed how to represent Venn diagrams in ASCII art.

9. Normative References

- [RFC3933] Klensin, J. and S. Dawkins, "A Model for IETF Process Experiments", BCP 93, RFC 3933, DOI 10.17487/RFC3933, November 2004, <<https://www.rfc-editor.org/info/rfc3933>>.

- [RFC8713] Kucherawy, M., Ed., Hinden, R., Ed., and J. Livingood, Ed., "IAB, IESG, IETF Trust, and IETF LLC Selection, Confirmation, and Recall Process: Operation of the IETF Nominating and Recall Committees", BCP 10, RFC 8713, DOI 10.17487/RFC8713, February 2020, <<https://www.rfc-editor.org/info/rfc8713>>.
- [RFC8788] Leiba, B., "Eligibility for the 2020-2021 Nominating Committee", BCP 10, RFC 8788, DOI 10.17487/RFC8788, May 2020, <<https://www.rfc-editor.org/info/rfc8788>>.

Appendix A. Available data

An analysis of how some of the above criteria would affect the number of NomCom-qualified participants if applied in August 2020 has been performed. The results are presented below in Venn diagrams as Figure 1 to Figure 4. Note that the numbers shown differ slightly from manual counts due to database mismatches, and the results were not derived at the normal time of the year for NomCom formation. The remote attendee lists for IETF 107 and 108 were used, although not yet available on the IETF web site.

A specific difficulty is that the databases involved inevitably contain a few inconsistencies such as duplicate entries, differing versions of a person's name, and impersonal authors. (For example, "IAB" qualifies under Path 3, and one actual volunteer artificially appears not to qualify.) This underlines that automatically generated lists of eligible and qualified people will always require manual checking.

The first two diagrams illustrate how the new paths (2 and 3) affect eligibility numbers compared to the meeting participation path (1). Figure 1 gives the raw numbers, and Figure 2 removes those disqualified according to RFC 8713. The actual 2020 volunteer pool is shown too.

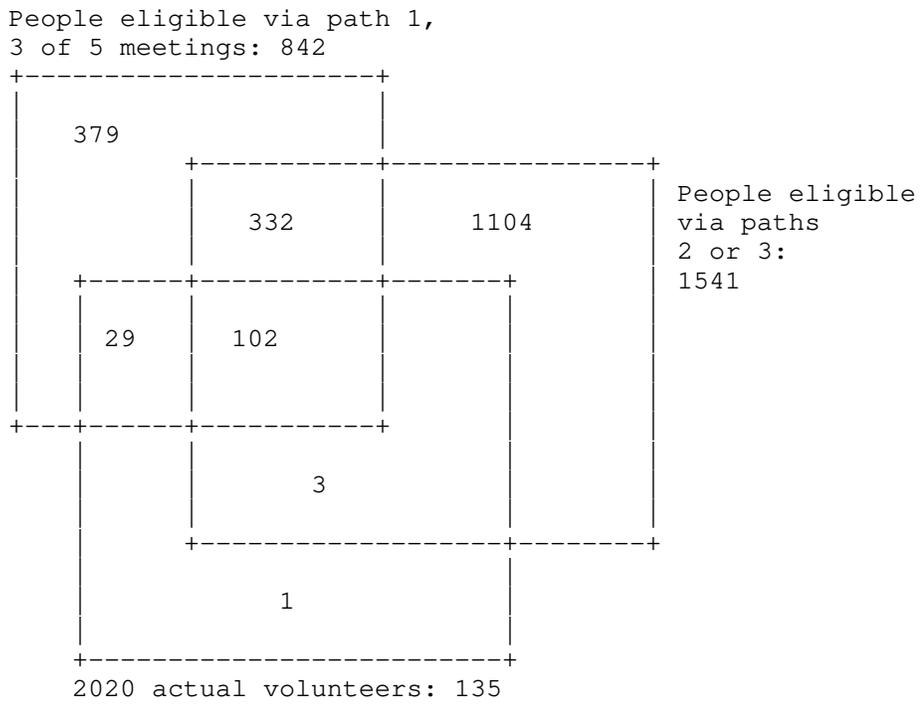


Figure 1: All paths, before disqualification

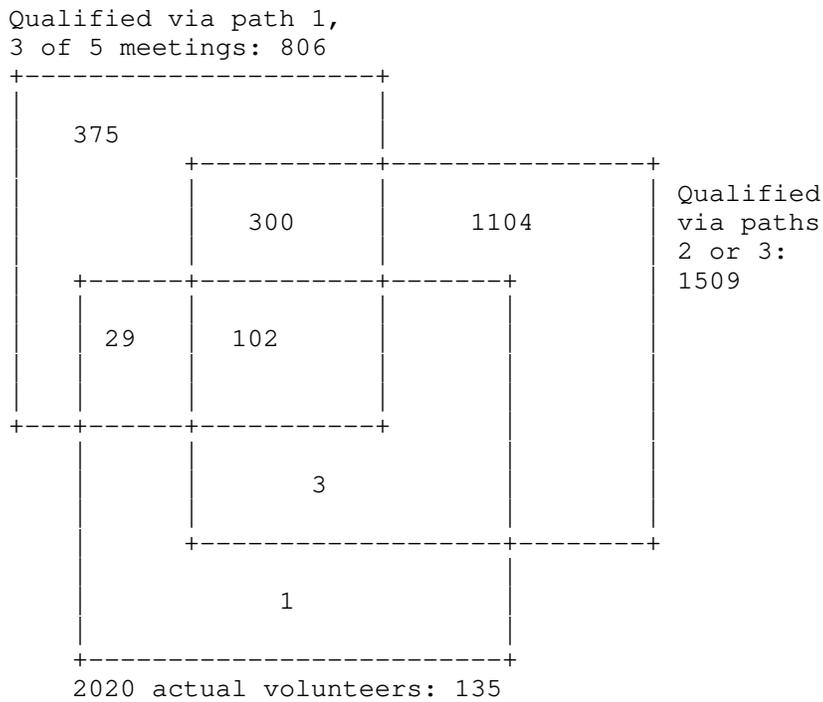


Figure 2: All paths, after disqualification

Figure 3 and Figure 4 illustrate how the new paths (2 and 3) interact with each other, also before and after disqualifications. The discarded path via IESG and IAB service (Section 5) is also shown, as path "I". The data clearly show that path "I" has no practical value.

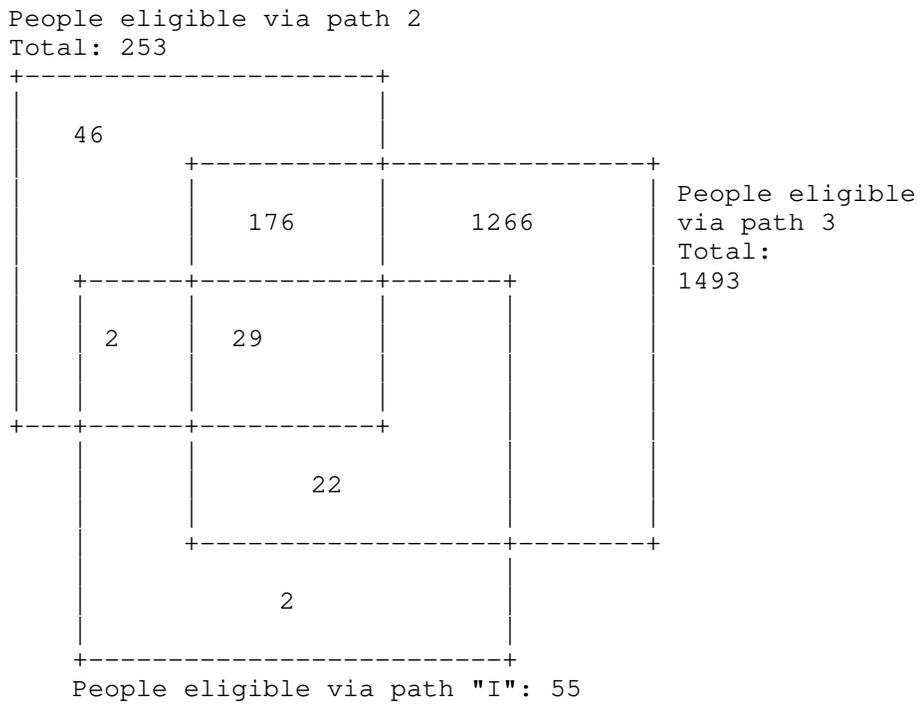


Figure 3: New paths, before disqualification

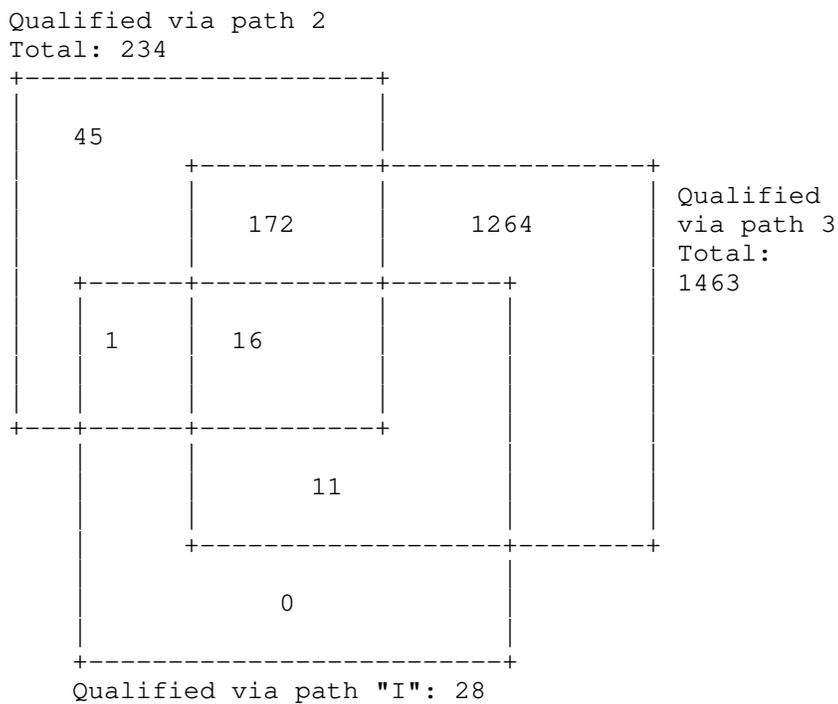


Figure 4: New paths, after disqualification

Appendix B. Change Log

This section is to be removed before publishing as an RFC.

B.1. Draft-09 to -10

- * IESG comments

B.2. Draft-08 to -09

- * IETF Last Call comments

B.3. Draft-07 to -08

- * AD review comments

B.4. Draft-06 to -07

- * Clarifications following reviews by Lars Eggert, Victor Kuarsingh and Barry Leiba

- * Added ASCII art versions of Venn diagrams
- B.5. Draft-05 to -06
- * Allowed for IETF 110 decision
 - * Resolved open issue
 - * Removed "future work" section
 - * Editorial improvements
- B.6. Draft-04 to -05
- * Adjusted criteria according to comments received
 - * Removed previous path 3
 - * Renumbered paths
 - * Updated diagrams
 - * Editorial improvements
- B.7. Draft-03 to -04
- * Adjusted criteria according to comments received
 - * Shortened period to one year (initially)
 - * Renumbered paths
 - * Updated diagrams
 - * Editorial improvements
- B.8. Draft-02 to -03
- * Adjusted criteria according to comments received
 - * Added data
- B.9. Draft-01 to -02
- * Made this an RFC 3933 process experiment

- * Eliminated path based on directorate reviews, used to be: "Has submitted at least 6 reviews as a member of an official IETF review team within the last 3 years."
- * Other comments from IETF107 virtual gendispatch meeting handled

B.10. Draft-00 to -01

- * Added author

Authors' Addresses

Brian E. Carpenter
The University of Auckland
School of Computer Science
PB 92019
Auckland 1142
New Zealand

Email: brian.e.carpenter@gmail.com

Stephen Farrell
Trinity College Dublin
College Green
Dublin
Ireland

Email: stephen.farrell@cs.tcd.ie

Network Working Group
Internet-Draft
Intended status: Best Current Practice
Expires: December 2, 2020

B. Carpenter
Univ. of Auckland
F. Gont
SI6 Networks
M. Richardson
Sandelman Software Works
May 31, 2020

Process for Working Group Adoption of Drafts
draft-carpenter-gendispatch-draft-adoption-00

Abstract

IETF working groups often formally adopt drafts. This document specifies minimum requirements for this process, thereby extending RFC 2418. It also describes how an adopted draft may be withdrawn from the working group process.

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 <https://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 December 2, 2020.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

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

1. Introduction	2
2. Consequences of WG Adoption of an Internet-Draft	2
3. Rules for Adoption of an Internet-Draft	3
4. Withdrawal of an Adopted Internet-Draft	4
5. IANA Considerations	5
6. Security Considerations	5
7. Acknowledgements	5
8. References	6
8.1. Normative References	6
8.2. Informative References	6
Appendix A. Change Log	6
A.1. Draft-00	6
Authors' Addresses	6

1. Introduction

According to [RFC2418], the Internet-Drafts (I-D) mechanism is a "resource for posting and disseminating in-process copies of working group documents." However, most I-Ds start as individual contributions and only become working group documents by a WG decision generally referred to as "adoption." As noted in [RFC7221], this process was not previously documented as a formal step in the IETF WG process. This has sometimes led to confusion about the significance of such adoption and about how it fits into the IETF standards process. The present document is intended to define a few formal rules about adoption to reduce such confusion.

2. Consequences of WG Adoption of an Internet-Draft

After a draft has been formally adopted by a WG, its original authors no longer have formal change control of the text. In addition to the normal consequence of posting a draft, i.e., that it becomes an IETF Contribution under [RFC5378], all future substantive changes to the draft require WG consensus and are no longer at the authors' sole discretion.

As a practical matter, the original authors usually continue to edit the document and make routine editorial decisions, but substantive changes must be referred to the WG and require WG rough consensus, consistently with [RFC2418]. It is also possible that new authors or editors will join the draft, or that previous authors may withdraw.

Adoption represents a commitment that the WG will spend time and effort on the draft, but it does not guarantee that the draft will reach WG consensus and be submitted to the IESG for publication as an RFC.

3. Rules for Adoption of an Internet-Draft

A WG Adoption Call of an I-D is not a required step of the IETF standards process. The WG chairs decide what documents belong in the WG, and can create new documents by fiat. A simple situation would be if a WG decides that an existing document should be split into two pieces: There is no reason to adopt each piece, that is needless bureaucracy. A WG that decides to create a design team to solve a problem has implicitly agreed to adopt the result. To not adopt the result is to say that the results of the WG mandated design team does not deserve first class agenda time. Such a design team would have been created, for instance, when a WG can not decide between two competing individual drafts and decides to merge them.

It is legitimate for a draft to be submitted to the IESG as described in [RFC2026] without a formal adoption by a WG.

If WG Chairs choose to consult the WG about adopting a document, this is the recommended process. The WG Chairs should also consider the additional guidelines in [RFC7221].

- o Any participant may request the adoption of a draft, after there has been a period of technical discussion of the draft in the relevant WG.
- o WG Chairs have discretion about when to issue an WG call for adoption, but they should do so regardless of their own opinions, when the WG discussion shows that there is clear interest in the draft in question.
- o A WG Chair or WG Secretary must send a formal WG call for adoption of a draft to the WG mailing list with at least two weeks time to respond.
- o This proposal should remind all participants, not just the authors, of their obligation to disclose relevant intellectual property rights (IPR) under [RFC8179].
- o Participants should consider the following aspects when responding to the WG call for adoption:
 - * The draft must fit within the current WG charter, or would do so with a simple modification to the charter.

- * The purpose of the draft should be clear.
 - * The proposal should be useful.
 - * The quality of writing should be sufficient for document to serve as the basis further work.
 - * There should be no strong technical objections.
 - * Any IPR disclosures should be acceptable.
 - * The work should not be in conflict with work elsewhere in the IETF.
- o An informal summary of these criteria is: Is this a problem the WG wants to solve in a way approximately as described in the draft?
 - o After this period, a WG Chair must, in a timely fashion, consider the comments and discussion in order to judge whether there is rough consensus to adopt the draft, and whether there is enough interest in the work that its completion is likely.
 - o If there is such consensus, this WG Chair will announce the result and, if positive, will request the authors to post a new version using an appropriate naming convention.
 - o This whole process is subject to the appeals process of [RFC2026].
4. Withdrawal of an Adopted Internet-Draft

It sometimes happens that an adopted draft does not reach WG consensus to be submitted to the IESG for publication as an RFC due to lack of interest, lack of effort, or lack of consensus. In such a case, it may be desirable for the WG to formally withdraw the WG draft, such that it is explicitly removed from the WG's agenda and returned to the authors' control.

The withdrawal of WG document should be the result of an explicit decision by the relevant WG, and should follow the following recommendations.

- o Upon evidence that progress on a WG draft has been stalled for a considerable period of time, a WG chair should evaluate the reasons of the apparent lack of progress. Such reasons may include lack of interest, lack of effort, or lack of consensus.
- o When progress on a document has been stalled for a considerable period of time, a WG chair, in consultation with the WG draft

authors and editors, should attempt to resume progress by taking appropriate actions that will normally depend on the nature of the lack of progress. For example, a WG draft that has been stalled due to apparent lack of interest may benefit from a call for a number of volunteers to produce detailed reviews of the WG draft. Similarly, a WG draft that has been stalled due to lack of effort by its authors/editors may benefit from the incorporation of new WG draft editors or the replacement of some of the existing ones.

- o If after successive failed attempts to make progress on a WG draft its completion remains unlikely, the WG Chairs may, at their own discretion, conclude that it is time for the WG to consider the formal withdrawal of the WG draft.
- o In such case, a WG Chair or WG Secretary must send a formal WG consensus call for withdrawal of the WG draft to the WG mailing list with at least two weeks time to respond, explaining the events that have triggered the aforementioned consensus call.
- o After this period, a WG Chair must, in a timely fashion, consider the comments and discussion in order to judge whether there is any concrete evidence that completion of the work may now be feasible, or whether completion of the work remains unlikely.
- o If further progress on the document remains unlikely, the WG Chair will announce the result of the consensus call and the formal withdrawal of the WG document. This will result in the document being removed from the WG's agenda and returned to the authors' control.
- o This whole process is subject to the appeals process of [RFC2026].

5. IANA Considerations

This document makes no request of IANA.

6. Security Considerations

This document should not affect the security of the Internet.

7. Acknowledgements

Useful comments were received from [TBD] ...

8. References

8.1. Normative References

- [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>>.
- [RFC2418] Bradner, S., "IETF Working Group Guidelines and Procedures", BCP 25, RFC 2418, DOI 10.17487/RFC2418, September 1998, <<https://www.rfc-editor.org/info/rfc2418>>.
- [RFC5378] Bradner, S., Ed. and J. Contreras, Ed., "Rights Contributors Provide to the IETF Trust", BCP 78, RFC 5378, DOI 10.17487/RFC5378, November 2008, <<https://www.rfc-editor.org/info/rfc5378>>.
- [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>>.

8.2. Informative References

- [RFC7221] Farrel, A. and D. Crocker, Ed., "Handling of Internet-Drafts by IETF Working Groups", RFC 7221, DOI 10.17487/RFC7221, April 2014, <<https://www.rfc-editor.org/info/rfc7221>>.

Appendix A. Change Log

A.1. Draft-00

- o Original version

Authors' Addresses

Brian E. Carpenter
The University of Auckland
School of Computer Science
PB 92019
Auckland 1142
New Zealand

Email: brian.e.carpenter@gmail.com

Fernando Gont
SI6 Networks
Evaristo Carriego 2644
Haedo, Provincia de Buenos Aires 1706
Argentina

Email: fgont@si6networks.com
URI: <https://www.si6networks.com>

Michael Richardson
Sandelman Software Works

Email: mcr+ietf@sandelman.ca
URI: <https://www.sandelman.ca/mcr/>

Network Working Group
Internet-Draft
Intended status: Best Current Practice
Expires: January 9, 2021

M. Knodel
Center for Democracy & Technology
N. ten Oever
Texas A&M University
July 08, 2020

Terminology, Power, and Inclusive Language in Internet-Drafts and RFCs
draft-knodel-terminology-03

Abstract

This document argues for moving away from certain specific language conventions sometimes used by RFC authors and the RFC Production Centre in order to encourage the use of more inclusive terminology in Internet-Drafts that are work in progress, and in new RFCs that may be published in any of the RFC series. The document also provides examples of inclusive terminology as precise alternatives for these conventions.

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 <https://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 January 9, 2021.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

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

1. Introduction

According to [RFC7322], "The ultimate goal of the RFC publication process is to produce documents that are readable, clear, consistent, and reasonably uniform," and one function of the RFC Editor is to "[c]orrect larger content/clarity issues; flag any unclear passages for author review." Documents that are published as RFCs are first worked on as Internet-Drafts.

Given the importance of communication between people developing RFCs, Internet-Drafts (I-D's), and related documents, it is worth considering the effects of terminology that has been identified as exclusionary. This document argues that certain obviously exclusionary terms should be avoided and replaced with alternatives. We propose nothing more than additional care in the choice of language just as care is taken in defining standards and protocols themselves.

This document presents arguments for why exclusionary terms should be avoided in Internet-Drafts and RFCs, describes the problems introduced by some specific terms, and proposes alternative language. The terms discussed in this document include "master-slave" and "whitelist-blacklist". There is a final section on additional considerations and general action points to address future RFCs and I-D's. Lastly, a summary of recommendations is presented.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119][RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Terminology and Power in Internet-Drafts and RFCs

According to the work of scholar Heather Brodie Graves from 1993, "one goal of the application of rhetorical theory in the technical communication classroom is to assess the appropriateness of particular terms and to evaluate whether these terms will facilitate or hinder the readers' understanding of the technical material" [BrodieGravesGraves]. This implies that in order to effectively communicate the content of I-Ds and RFCs to all readers, it is important for Authors to consider the kinds of terms or language

conventions that may inadvertently get in the way of effective communication. She continues, "complex and subtle configurations of sexist, racist, or ethnocentric language use in technical documents can derail or interfere with readers' ability and desire to comprehend and follow important information."

Indeed, problems of language are problems of everyday speech. Racist and sexist language is rampant and similarly counter-productive in other sectors, notably social work [Burgest]. The terms "master-slave," treated in detail below are present in other realms of technology, notably "automotive clutch and brake systems, clocks, flip-flop circuits, computer drives, and radio transmitters" [Eglash].

However as noted in the research by Ron Eglash, this seemingly entrenched technical terminology is relatively recent. It is not too late for these terms to be replaced with alternative metaphors that are more accurate, clearer, less distracting, and that do not offend their readers. Language matters and metaphors matter. Indeed, metaphors can be incredibly useful devices to make more human the complex technical concepts presented in RFCs. Metaphors should not be avoided, but rather taken seriously. Renowned linguist George Lakoff argued in 1980 that the ubiquitous use of metaphors in our everyday speech indicates a fundamental instinct to "structure our most basic understandings of experience" [Lakoff]. Metaphors structure relationships, and they frame possibilities and impossibilities [Wyatt].

Like Graves, this document recognises the monumental challenge of addressing linguistics and power, and attempts to "promote awareness that may lead to eventual wide-spread change" [BrodieGravesGraves] and suggests first steps for actions that may remedy the inadvertent use of undesirable terms'. To that end, the list below is a tersely written set of IETF-specific arguments as to why the RFC Editor should be encouraged to correct other content and clarity issues with respect to excluding language and metaphors:

- o The RFC series is intended to remain online in perpetuity. Societal attitudes to offensive and excluding language shift over time in the direction of more empathy, not less.
- o That excluding terms in RFCs are largely hidden from the larger public, or read only by engineers, is no excuse to ignore social-level reactions to the terms. If the terms would be a poor choice for user-facing application features, the terms should be avoided in technical documentation and specifications, too.

- o At the time of this drafting, the digital technology community has a problem with monoculture. And because the diversity of the technical community is already a problem, a key strategy to breaking monoculture is to ensure that technical documentation is addressed to a wide audience and multiplicity of readers.
- o And yet the technical community already includes members who take offense to these terms. Eradicating the use of excluding terminology in official RFCs recognises the presence of and acknowledges the requests from black and brown engineers and from women and gender-non-conforming engineers to avoid the use of exclusive terminology.

This document does not try to prescribe terminology shifts for any and all language that could be deemed exclusive. Instead what follow are specific alternative suggestions to "master-slave" and "white-blacklist" and the rationale for the use of the alternatives. Additional considerations are presented in a subsequent section.

3.1. Master-Slave

Master-slave is an offensive and exclusive metaphor that will and should never become fully detached from history. Aside from being unprofessional and exclusive it stifled the participation of students whom Eglash interviewed for his research. He asks: "If the master-slave metaphor affected these tough-minded engineers who had the gumption to make it through a technical career back in the days when they may have been the only black persons in their classes, what impact might it have on black students who are debating whether or not to enter science and technology careers at all?" [Eglash]

Aside from the arguably most important reason outlined above, these terms are becoming less used and therefore increasingly less compatible as more communities move away from its use (eg [NIST], [Python], [Drupal], [Github] and [Django]). The usage of 'master' and 'slave' in hardware and software has been halted by the Los Angeles County Office of Affirmative Action, the Django community, the Python community and several other programming languages. This was done because the language is offensive and hurts people in the community [Django2]. Root operator Internet Systems Consortium stopped using the terms because they were asked to [ISC].

In addition to being inappropriate and arcane, the master-slave metaphor is both technically and historically inaccurate. For instance, in DNS the 'slave' is able to refuse zone transfers on the ground that it is malformed. The metaphor is incorrect historically given the most recent centuries during which "the role of the master was to abdicate and the role of the slave was to revolt"

[McClelland]. Yet in another sense slavery is also not 'just an historic term', whereas freedom from slavery is a human-rights issue [UDHR], it continues to exist in the present [Wikipedia]. Furthermore, this term set wasn't revived until recently, after WWII, and after many of the technologies that adopted it were already in use with different terminology [Eglash].

Lastly, we present not an additional rationale against their use, but an indicator of actual racism in the community that has been surfaced as a result of this larger debate among technologists, "I don't believe in PC (political correctness), mostly because the minorities constantly use it to get away with anything" [Jansens]. This illustrates the need to, as Graves is cited above as saying, continue to raise awareness within our community for eventual, lasting change on the continued front of struggle against the racists amongst us.

3.1.1. Suggested Alternatives

There are also many other relationships that can be used as metaphors, Eglash's research calls into question the accuracy of the master-slave metaphor. Fortunately, there are ample alternatives for the master-slave relationship. Several options are suggested here and should be chosen based on the pairing that is most clear in context:

- o Primary-secondary
- o Primary-replica
- o Active-standby
- o Writer-reader

Since the use of master-slave is becoming less common in other technical communities, it is best to simply duplicate the metaphor being used by comparable or interoperable technologies. The IETF can show positive leadership in the technical community by setting standards without using offensive and exclusive metaphors.

For the DNS, RFC 8499 defines the current best practise for DNS terminology and uses the term pair 'primary' and 'secondary' [RFC8499].

3.2. Blacklist-Whitelist

The metaphorical use of white-black to connote good-evil is exclusive. While master-slave might seem like a more egregious example of racism, white-black is arguably worse because it is more

pervasive and therefore more insidious. While recent headlines have decried the technical community's use of master-slave, there is far less discussion about white-black despite its importance. There is even a name for this pervasive language pitfall: the association of white with good and black with evil is known as the "bad is black effect" [Grewal].

Indeed, there is an entire book on the subject, written by renowned authority on race, Frantz Fanon. In his book "Black Skin, White Masks," Fanon makes several persuasive arguments that standard language encodes subconscious in-group, out-group preferences [Fanon].

In the case of blacklist-whitelist in the technical documentation of I-Ds and RFCs, it is entirely a term of art and an arbitrary metaphorical construct with no technical merit. There are scientific uses of black that are related to light- black holes are black because light cannot escape them. Blacklist-whitelist is not a metaphor for lightness or darkness, it is a good-evil metaphor and therefore this trope has significant impact on how people are seen and treated. As we've seen with metaphors, its use is pervasive and, though not necessarily conscious, perceptions do get promulgated through culture and repetition.

As with master-slave, we save our technical argument for last, referencing and presenting first the reasons for the use of non-offensive, alternative terminology for the sake of our humanity. Indeed, our technical argument is incredibly succinct: Why use a metaphor when a direct description is both succinct and clear? There can be absolutely no ambiguity if one uses the terms, as suggested below, allow-block rather than white-black.

3.2.1. Suggested Alternatives

There are alternatives to this terminology set that vastly improve clarity because they are not even metaphors without adding a single additional character. The alternatives proposed here say exactly what they mean. Examples of specifications that use these alternative terms are also provided for illustration purposes.

- o Accept-list and Drop-list (see for example [RFC8612], [RFC8782], and [RFC8783])
- o Blocklist-allowlist
- o Deny-allow
- o Droplist-accesslist

- o Drop-permit
- o Block-permit

3.3. Other Considerations

As described in the preceding sections, the language used in technical documentation, like all written text, creates and reinforces expectations and stereotypes. We propose nothing more than additional care in the choice of language just as care is taken in defining standards and protocols themselves. The two examples provided above are not the only cases of offensive language to be avoided, and many more can be collected. However, these two examples are particularly significant and require immediate action. We use this section to broaden the context of other offensive and excluding terminologies to encompass additional concerns.

There are many other metaphors present in technical documentation that are "terms of art" but that have no technical basis whatsoever.

If any of these metaphors is offensive there is no excuse for its continued use. A term like "man-in-the-middle" is not technically useful. It is not a standard term, not as clear as its alternative "on-path attacker", and should therefore be avoided. When presented with the opportunity to employ the use of metaphors or to unthinkingly repeat terms of art that connote gender or race, Authors should simply find a better way to explain themselves. A fun read on the politics of colloquial speech by George Orwell should dissuade any clever Author from using tired explanatory metaphors [Orwell].

Up until recently, strict English grammarians like Orwell decried the use of the neutral pronoun "they". Without a neutral singular pronoun, "he" is assumed as the default singular pronoun when the gender of the person is unknown or ambiguous. However, that has changed, and it is now widely accepted that "they" can be used as a neutral singular pronoun. Since it is unlikely that all implementers and infrastructure operators are of any particular gender, "he" should never be used to refer to a person in I-Ds and RFCs. An Author who uses male examples sets male-ness as a standard.

Militarised metaphors are also a pervasive problem in language, perhaps even more so in technical communities because of the historical and actual relationship between technology and war. We welcome additional examples of terminology that might be avoided through more awareness and thoughtfulness.

4. Summary of Recommendations

To summarise, we have bulleted some very concrete action points that can be taken by Editors, reviewers and Authors, both present and future as they develop and publish Internet-Drafts and new RFCs.

Authors SHOULD: * Replace the excluding term "master-slave" with more accurate alternatives, for instance from the list of Section 3.1. * Replace the excluding term "blacklist-whitelist" with more accurate alternative, for instance from the list of suggested alternatives at Section 3.2. * Reflect on their use of metaphors generally * Use the neutral "they" as the singular pronoun, and * Consider changing existing exclusive language in current (reference) implementations [socketwench] * Consult the style sheet maintained by the RFC editor.

RFC Editor MUST: * Offer alternatives for excluding terminology as an important act of correcting larger editorial issues and clarifying technical concepts and * Maintain a style sheet that collects all terms that have been considered and indicate whether they are deemed acceptable, and if not what terms Authors should consider instead * Suggest to Authors that even when referencing other specifications that have not replaced offensive terminology, the Authors could use another term in their document and include a note to say that they have used the new term as a replacement for the term used in the referenced document.

5. Further reading

'Anyone can edit', not everyone does: Wikipedia and the gender gap' by Ford, Heather and Wajcman, Judy (2017) Social Studies of Science. ISSN 0306-3127

Grant, Barbara M. "Master--slave dialogues in humanities supervision...<https://doi.org/10.1177/1474022207084880>

Miller, Carolyn. "A Humanistic Rationale for Technical Writing"

6. Security Considerations

Security is dependent on a wide range of actors that are implementing technical documentation. Therefore it is crucial that language is clear, and understood by all that need to implement this documentation. Correct and inclusive language is therefore conducive for secure implementations of technical documentation.

7. IANA Considerations

This document has no actions for IANA.

8. References

8.1. 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, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

8.2. Informative References

- [BrodieGravesGraves] Heather Brodie Graves, . and . Roger Graves, "Masters, slaves, and infant mortality: Language challenges for technical editing", *Technical Communication Quarterly*, 7:4, 389-414 , 1998, <<https://doi.org/10.1080/10572259809364639>>.
- [Burgest] Burgest, David., "'Racism in Everyday Speech and Social Work Jargon.'", *Social Work*, vol. 18, no. 4, 1973, pp. 20-25 , 1973, <www.jstor.org/stable/23711113>.
- [Django] fcurella, ., "#22667 replaced occurrences of master-slave terminology with leader/follower #2692", 2014, <<https://github.com/django/django/pull/2692>>.
- [Django2] lynncyryn, ., "comment on #22667 replaced occurrences of master-slave terminology with leader/follower #2692", 2014, <<https://github.com/django/django/pull/2692#issuecomment-44221563>>.
- [Drupal] Xano, ., "Replace 'master-slave' terminology with 'primary/replica'", 2014, <<https://www.drupal.org/project/drupal/issues/2275877>>.
- [Eglash] Ron Eglash, ., "Broken Metaphor: The Master-Slave Analogy in Technical Literature.", *Technology and Culture*, vol. 48 no. 2, 2007, pp. 360-369. , 2007, <<https://doi.org/10.1353/tech.2007.0066>>.

- [Fanon] Fanon, F., "Black skin, white masks", 1952.
- [Github] Kevin Truong, . and VICE, "Github to Remove 'Master/Slave' Terminology From its Platform", June 2020, <https://www.vice.com/en_us/article/k7qbyv/github-to-remove-masterslave-terminology-from-its-platform>.
- [Grewal] Grewal, D., "The 'Bad Is Black' Effect", 2017, <<https://www.scientificamerican.com/article/the-bad-is-black-effect/>>.
- [ISC] Internet Systems Consortium, ., "@ISCdotORG reply tweet", 2017, <<https://twitter.com/ISCdotORG/status/943152507211071489>>.
- [Jansens] Bart Jansens, ., "I don't believe in PC", 2008, <https://www.drupal.org/project/project_issue_file_review/issues/343414#comment-1164514>.
- [Lakoff] George Lakoff, . and . Mark Johnson, "Metaphors We Live By", U of Chicago P, 1980. , n.d..
- [McClelland] McClelland, J., "We need better metaphors", 2011, <<https://current.workingdirectory.net/posts/2011/master-slave>>.
- [NIST] Eric Geller, . and Politico, "Agency to end use of technology terms such as 'master' and 'slave' over racist associations", June 2020, <<https://www.politico.com/news/2020/06/25/agency-ends-use-technology-terms-racist-associations-339880>>.
- [Orwell] George Orwell, ., "Politics and the English Language", 1946.
- [Python] Daniel Oberhaus, ., "'master-slave' Terminology Was Removed from Python Programming Language", 2018, <https://motherboard.vice.com/en_us/article/8x7akv/masterslave-terminology-was-removed-from-python-programming-language>.
- [RFC7322] Flanagan, H. and S. Ginoza, "RFC Style Guide", RFC 7322, DOI 10.17487/RFC7322, September 2014, <<https://www.rfc-editor.org/info/rfc7322>>.

- [RFC8499] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", BCP 219, RFC 8499, DOI 10.17487/RFC8499, January 2019, <<https://www.rfc-editor.org/info/rfc8499>>.
- [RFC8612] Mortensen, A., Reddy, T., and R. Moskowitz, "DDoS Open Threat Signaling (DOTS) Requirements", RFC 8612, DOI 10.17487/RFC8612, May 2019, <<https://www.rfc-editor.org/info/rfc8612>>.
- [RFC8782] Reddy, K. T., Ed., Boucadair, M., Ed., Patil, P., Mortensen, A., and N. Teague, "Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal Channel Specification", RFC 8782, DOI 10.17487/RFC8782, May 2020, <<https://www.rfc-editor.org/info/rfc8782>>.
- [RFC8783] Boucadair, M., Ed. and T. Reddy, K., Ed., "Distributed Denial-of-Service Open Threat Signaling (DOTS) Data Channel Specification", RFC 8783, DOI 10.17487/RFC8783, May 2020, <<https://www.rfc-editor.org/info/rfc8783>>.
- [socketwench] socketwench, ., "Even in tech, words matter", 2018, <<https://deninet.com/blog/2018/09/09/even-tech-words-matter>>.
- [UDHR] United Nations General Assembly, "The Universal Declaration of Human Rights", 1948, <<http://www.un.org/en/documents/udhr/>>.
- [Wikipedia] Wikipedia, "Slavery in the 21st century", 2018, <https://en.wikipedia.org/wiki/Talk:Slavery_in_the_21st_century>.
- [Wyatt] Sally Wyatt, ., "Danger! Metaphors at Work in Economics, Geophysiology, and the Internet", Science, Technology, and Human Values, Volume: 29 issue: 2, page(s): 242-261 , 2004.

Authors' Addresses

Mallory Knodel
Center for Democracy & Technology

Email: mknodel@cdt.org

Niels ten Oever
Texas A&M University

Email: mail@nielstenoever.net

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 25 February 2024

M. Knodel
Center for Democracy & Technology
N. ten Oever
University of Amsterdam
24 August 2023

Terminology, Power, and Inclusive Language in Internet-Drafts and RFCs
draft-knodel-terminology-14

Abstract

There has been extensive discussion in and around the IETF community about the use of technical terminology, which could be interpreted as exclusionary. The document below is published as an artefact of the discussion because it sparked many debates and inspired several actions in the IETF community. This, however, does not say anything about whether the opinions it holds are correct or incorrect. Since the debate about technology, language, and its implications will probably never be finished, we offer this document for reference in future discussions about the topic.

It is important to note that this is not standard, it does not represent IETF consensus, and should not be misconstrued as anything other than the authors views.

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 <https://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 25 February 2024.

Copyright Notice

Copyright (c) 2023 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1. Introduction	2
2. Terminology and Power in Internet-Drafts and RFCs	3
2.1. Master-Slave	4
2.2. Blacklist-Whitelist	6
2.3. Other Considerations	7
3. Summary of Recommendations	8
4. Epilogue	9
5. Further reading	9
6. Security Considerations	9
7. IANA Considerations	9
8. References	9
8.1. Normative References	9
8.2. Informative References	10
Authors' Addresses	12

1. Introduction

According to [RFC7322], The ultimate goal of the RFC publication process is to produce documents that are readable, clear, consistent, and reasonably uniform, and one function of the RFC Editor is to [c]orrect larger content/clarity issues; flag any unclear passages for author review. Documents that are published as RFCs are first worked on as Internet-Drafts.

Given the importance of communication between people developing RFCs, Internet-Drafts (I-Ds), and related documents, it is worth considering the effects of terminology that has been identified as exclusionary. This document argues that certain obviously exclusionary terms should be avoided and replaced with alternatives. We propose nothing more than additional care in the choice of language just as care is taken in defining standards and protocols themselves.

This document presents arguments for why exclusionary terms should be avoided in Internet-Drafts and RFCs and as an exercise describes the problems introduced by some specific terms and why their proposed alternatives improve technical documentation. The example terms

discussed in this document include master-slave and whitelist-blacklist. There is a final section on additional considerations and general action points to address future RFCs and I-Ds. Lastly, a summary of recommendations is presented.

2. Terminology and Power in Internet-Drafts and RFCs

According to the work of scholar Heather Brodie Graves from 1993, one goal of the application of rhetorical theory in the technical communication classroom is to assess the appropriateness of particular terms and to evaluate whether these terms will facilitate or hinder the readers understanding of the technical material [BrodieGravesGraves]. This implies that in order to effectively communicate the content of I-Ds and RFCs to all readers, it is important for authors to consider the kinds of terms or language conventions that may inadvertently get in the way of effective communication. She continues, complex and subtle configurations of sexist, racist, or ethnocentric language use in technical documents can derail or interfere with readers ability and desire to comprehend and follow important information.

Indeed, problems of language are problems of everyday speech. Racist and sexist language is rampant and similarly counter-productive in other sectors, notably social work [Burgest]. The terms master-slave, treated in detail below are present in other realms of technology, notably automotive clutch and brake systems, clocks, flip-flop circuits, computer drives, and radio transmitters [Eglash].

However as noted in the research by Ron Eglash, this seemingly entrenched technical terminology is relatively recent. It is not too late for these terms to be replaced with alternative metaphors that are more accurate, clearer, less distracting, and that do not offend their readers. Language matters and metaphors matter. Indeed, metaphors can be incredibly useful devices to make more human the complex technical concepts presented in RFCs. Metaphors should not be avoided, but rather taken seriously. Renowned linguist George Lakoff argued in 1980 that the ubiquitous use of metaphors in our everyday speech indicates a fundamental instinct to structure our most basic understandings of experience [Lakoff]. Metaphors structure relationships, and they frame possibilities and impossibilities [Wyatt].

Like Graves, this document recognises the monumental challenge of addressing linguistics and power, and attempts to promote awareness that may lead to eventual wide-spread change [BrodieGravesGraves] and suggests first steps for actions that may remedy the inadvertent use of undesirable terms. To that end, the list below is a tersely

written set of IETF-specific arguments as to why the RFC Editor should be encouraged to correct other content and clarity issues with respect to exclusionary language and metaphors:

1. The RFC series is intended to remain online in perpetuity. Societal attitudes to offensive and exclusionary language shift over time in the direction of more empathy, not less.
2. That exclusionary terms in RFCs are largely hidden from the larger public, or read only by engineers, is no excuse to ignore social-level reactions to the terms. If the terms would be a poor choice for user-facing application features, the terms should be avoided in technical documentation and specifications, too.
3. At the time of this drafting, the digital technology community has a problem with monoculture. And because the diversity of the technical community is already a problem, a key strategy to breaking monoculture is to ensure that technical documentation is addressed to a wider audience and greater multiplicity of readers.
4. And yet the technical community already includes members who take offense to these terms. Eradicating the use of exclusionary terminology in official RFCs recognises the presence of and acknowledges the requests from black and brown engineers and from women and gender-non-conforming engineers to avoid the use of exclusionary terminology.

This document does not try to prescribe terminology shifts for any and all language that could be deemed exclusionary. Instead what follow are two examples of specific alternative suggestions to master-slave and white-blacklist and the rationale for the use of the alternatives. Suggested actions for handling additional considerations are presented in a subsequent section.

2.1. Master-Slave

Master-slave is an offensive and exclusionary metaphor that will and should never become fully detached from history. Aside from being unprofessional and exclusionary it stifled the participation of students whom Eglash interviewed for his research. He asks: If the master-slave metaphor affected these tough-minded engineers who had the gumption to make it through a technical career back in the days when they may have been the only black persons in their classes, what impact might it have on black students who are debating whether or not to enter science and technology careers at all? [Eglash]

Aside from the arguably most important reason outlined above, these terms are becoming less used and therefore increasingly less compatible as more communities move away from its use (eg [NIST], [Python], [Drupal], [Github] and [Django]). The usage of master and slave in hardware and software has been halted by the Los Angeles County Office of Affirmative Action, the Django community, the Python community and several other programming languages. This was done because the language is offensive and hurts people in the community [Django2]. Root operator Internet Systems Consortium also stopped using the terms [ISC].

In addition to being inappropriate and arcane, the master-slave metaphor is both technically and historically inaccurate. For instance, in DNS the slave is able to refuse zone transfers on the ground that it is malformed. The metaphor is incorrect historically given the most recent centuries during which the role of the master was to abdicate and the role of the slave was to revolt [McClelland]. Yet in another sense slavery is also not just an historic term, whereas freedom from slavery is a human-rights issue [UDHR], it continues to exist in the present [Wikipedia]. Furthermore, this term set wasnt revived until recently, after WWII, and after many of the technologies that adopted it were already in use with different terminology [Eglash].

Ultimately master-slave is a poor choice since it is 1) being used less frequently already 2) in a variety of applications 3) to correct perceived exclusionary effects 4) at the request of concerned members of the technical community.

To find alternatives to master-slave, one can look to myriad existing implementations. There are also many other relationships that can be used as metaphors, Eglashs research calls into question the accuracy of the master-slave metaphor. An alternative should be chosen based on the pairing that is most clear in context:

- * Primary-secondary based on authority. See for example [RFC8499].
- * Primary-replica based originality.
- * Active-standby based on state.
- * Writer-reader based on function.

2.2. Blacklist-Whitelist

The metaphorical use of white-black to connote good-evil is exclusive. While master-slave might seem like a more egregious example of racism, white-black is arguably worse because it is more pervasive and therefore more insidious. While recent headlines have decried the technical community's use of master-slave, there is far less discussion about white-black despite its importance. There is even a name for this pervasive language pitfall: the association of white with good and black with evil is known as the bad is black effect [Grewal].

Indeed, there is an entire book on the subject, written by renowned authority on race, Frantz Fanon. In his book *Black Skin, White Masks*, Fanon makes several persuasive arguments that standard language encodes subconscious in-group, out-group preferences [Fanon].

In the case of blacklist-whitelist in the technical documentation of I-Ds and RFCs, it is entirely a term of art and an arbitrary metaphorical construct with no technical merit. There are scientific uses of black that are related to light black holes are black because light cannot escape them. The dark web is not counter-related to light; it is a metaphor for evil. Blacklist-whitelist is not a metaphor for lightness or darkness, it is a good-evil metaphor and therefore this trope has significant impact on how people are seen and treated. As we've seen with metaphors, its use is pervasive and, though not necessarily conscious, perceptions get promulgated through culture and repetition.

As with master-slave, we save our technical argument for last, referencing and presenting first the reasons for the use of non-offensive, alternative terminology for the sake of our humanity. Indeed, our technical argument is incredibly succinct: Why use a metaphor when a direct description is both succinct and clear? There can be absolutely no ambiguity if one uses the terms, as suggested below, allow-block rather than white-black.

There are alternatives to this terminology set that vastly improve clarity because they are not even metaphors, they're descriptions. The alternatives proposed here say exactly what they mean.

* Accept-list and drop-list for threat signaling. See for example [RFC8612], [RFC8782], and [RFC8783]).

* Blocklist-allowlist, deny-allow, exempt-allowlist or block-permit for permissions.

2.3. Other Considerations

As described in the preceding sections, the language used in technical documentation, like all written text, creates and reinforces expectations and stereotypes. We propose nothing more than additional care in the choice of language just as care is taken in defining standards and protocols themselves. The two examples provided above are not the only cases of exclusionary language to be avoided, and many more can be collected. We use this section to broaden the context of other offensive and exclusionary terminologies to encompass additional concerns, why spotting and eradicating problematic terminologies is a valid endeavour for authors and editors of technical documentation and how this might be systematised.

There are many other metaphors present in technical documentation that are terms of art but that have no technical basis whatsoever. If any of these metaphors is offensive there is no excuse for its continued use. A term like man-in-the-middle is not technically useful. It is not a standard term, not as clear as its alternative on-path attacker, and should therefore be avoided. When presented with the opportunity to employ the use of metaphors or to unthinkingly repeat terms of art that connote gender or race, authors should simply find a better way to explain themselves. A fun read on the politics of colloquial speech by George Orwell should dissuade any clever author from using tired explanatory metaphors [Orwell].

Gendered pronouns and sexism are common place but easy to spot and replace. Up until recently, strict English grammaticists like Orwell decried the use of the neutral pronoun they. Without a neutral singular pronoun, he is assumed as the default singular pronoun when the gender of the person is unknown or ambiguous. However, that has changed, and it is now widely accepted that they can be used as a neutral singular pronoun. Since it is unlikely that all implementers and infrastructure operators are of any particular gender, he should never be used to refer to a person in I-Ds and RFCs. An author who uses male examples sets male-ness as a standard.

Besides race and gender, our world is full of metaphors rooted in oppression, ableism, and colonialism. Militarised metaphors are also a pervasive problem in language, perhaps even more so in technical communities because of the historical and actual relationship between technology and war.

While it is not our intention to be exhaustive we hope to have made a persuasive case for authors and editors to pay attention to the finer details of metaphor, and the ways power is replicated in technical documentation unless detailed attention is paid. The example terms

above master-slave and blacklist-whitelist are already less common. If the IETF community has learned anything from the debate over the use of these terms, and this document, it is that language matters to us deeply as members of society and as engineers. And because language, and society, change over time, we must approach future concerns with some degree of dispassion when the arguments presented in the first section can be clearly applied.

There is harm in protracted discussion that weighs the validity IETF participants experiences with exclusionary terminology. The IETFs own discussions surfaced expressions and defense of racist views that pushed away participants and observers. This illustrates the need to, as Graves is cited above as saying, continue to raise awareness within our community for eventual, lasting change on the continued front of struggle against the racists amongst us. Yet we recommend a living stylesheet, rather than repeated RFCs, be used as a mechanism for monitoring exclusionary language in IETF documents [inclusiveterminology].

It is there that we welcome additional examples of terminology that might be avoided through more awareness and thoughtfulness.

3. Summary of Recommendations

To summarise, we have outlined some very concrete action points that can be taken by editors, reviewers and authors, both present and future as they develop and publish Internet-Drafts and new RFCs.

Authors can consider to: * Replace the exclusionary terms master-slave and blacklist-whitelist with more accurate alternatives. * Read and reflect upon the repository of exclusionary terminology [inclusiveterminology]. * Reflect on their use of metaphors generally. * Consider changing existing exclusionary language in current (reference) implementations [socketwrench]. * Consult the RFC style sheet maintained by the RFC editor and the community that can be found at <https://github.com/ietf/terminology> .

During the publication process, publishers (such as the RFC Editor) are advised to: * Offer alternatives for exclusionary terminology as an important act of correcting larger editorial issues and clarifying technical concepts and * Maintain the IETF repository that collects all terms that have been considered and indicate whether they are deemed acceptable, and if not what terms authors should consider instead. * Suggest to authors that even when referencing other specifications that have not replaced offensive terminology, the authors could use another term in their document and include a note to say that they have used the new term as a replacement for the term used in the referenced document.

4. Epilogue

This document built a compendium of scholarship, activist campaigns, and the will of technologists who had pointed out general and specific issues with technical terms. This sparked a significant discussion in the IETF. Concretely the documents writing resulted in a statement by the IESG [Statement] on on Inclusive Language and its mainstreaming with the [in-solidarity-bot]. The authors chose to seek publication of this document as a historical marker of that discussion and as a contribution to social and restorative justice.

5. Further reading

Ford, Heather., Wajcman, Judy. 2017. Anyone can edit, not everyone does: Wikipedia and the gender gap Social Studies of Science. ISSN 0306-3127

Grant, Barbara M. 2008. Masterslave dialogues in humanities supervision Arts and Humanities in Higher Education, Volume: 7 issue: 1, page(s): 9-27 <https://doi.org/10.1177/1474022207084880>

Miller, Carolyn, R. 1979. A Humanistic Rationale for Technical Writing College English, Vol. 40, No. 6, pp. 610-617

6. Security Considerations

Security is dependent on a wide range of actors that are implementing technical documentation. Therefore it is crucial that language is clear, and understood by all that need to implement this documentation. Correct and inclusive language is therefore conducive for secure implementations of technical documentation.

7. IANA Considerations

This document has no actions for IANA.

8. References

8.1. 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, <<https://www.rfc-editor.org/info/rfc2119>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

8.2. Informative References

- [RFC7322] Flanagan, H. and S. Ginoza, "RFC Style Guide", RFC 7322, DOI 10.17487/RFC7322, September 2014, <<https://www.rfc-editor.org/info/rfc7322>>.
- [RFC8499] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", BCP 219, RFC 8499, DOI 10.17487/RFC8499, January 2019, <<https://www.rfc-editor.org/info/rfc8499>>.
- [RFC8782] Reddy.K, T., Ed., Boucadair, M., Ed., Patil, P., Mortensen, A., and N. Teague, "Distributed Denial-of-Service Open Threat Signaling (DOTS) Signal Channel Specification", RFC 8782, DOI 10.17487/RFC8782, May 2020, <<https://www.rfc-editor.org/info/rfc8782>>.
- [RFC8783] Boucadair, M., Ed. and T. Reddy.K, Ed., "Distributed Denial-of-Service Open Threat Signaling (DOTS) Data Channel Specification", RFC 8783, DOI 10.17487/RFC8783, May 2020, <<https://www.rfc-editor.org/info/rfc8783>>.
- [RFC8612] Mortensen, A., Reddy, T., and R. Moskowitz, "DDoS Open Threat Signaling (DOTS) Requirements", RFC 8612, DOI 10.17487/RFC8612, May 2019, <<https://www.rfc-editor.org/info/rfc8612>>.
- [Burgest] Burgest, David R., "Racism in Everyday Speech and Social Work Jargon.", *Social Work*, vol. 18, no. 4, 1973, pp. 2025 , 1973, <www.jstor.org/stable/23711113>.
- [Eglash] Ron Eglash, ., "Broken Metaphor: The Master-Slave Analogy in Technical Literature.", *Technology and Culture*, vol. 48 no. 2, 2007, pp. 360-369. , 2007, <<https://doi.org/10.1353/tech.2007.0066>>.
- [BrodieGravesGraves] Heather Brodie Graves, . and . Roger Graves, "Masters, slaves, and infant mortality: Language challenges for technical editing", *Technical Communication Quarterly*, 7:4, 389-414 , 1998, <<https://doi.org/10.1080/10572259809364639>>.
- [Wyatt] Sally Wyatt, ., "Danger! Metaphors at Work in Economics, Geophysiology, and the Internet", *Science, Technology, and Human Values*, Volume: 29 issue: 2, page(s): 242-261 , 2004.

- [Lakoff] George Lakoff, . and . Mark Johnson, "Metaphors We Live By", U of Chicago P, 1980. , n.d..
- [Orwell] George Orwell, ., "Politics and the English Language", 1946.
- [McClelland] McClelland, J., "We need better metaphors", 2011, <<https://current.workingdirectory.net/posts/2011/master-slave>>.
- [UDHR] United Nations General Assembly, "The Universal Declaration of Human Rights", 1948, <<http://www.un.org/en/documents/udhr/>>.
- [Fanon] Fanon, F., "Black skin, white masks", 1952.
- [Jansens] Bart Jansens, ., "I don't believe in PC", 2008, <https://www.drupal.org/project/project_issue_file_review/issues/343414#comment-1164514>.
- [Python] Daniel Oberhaus, ., "'master-slave' Terminology Was Removed from Python Programming Language", 2018, <https://motherboard.vice.com/en_us/article/8x7akv/masterslave-terminology-was-removed-from-python-programming-language>.
- [Django] fcurella, ., "#22667 replaced occurrences of master-slave terminology with leader/follower #2692", 2014, <<https://github.com/django/django/pull/2692>>.
- [Django2] lynncyryn, ., "comment on #22667 replaced occurrences of master-slave terminology with leader/follower #2692", 2014, <<https://github.com/django/django/pull/2692#issuecomment-44221563>>.
- [Wikipedia] Wikipedia, "Slavery in the 21st century", 2018, <https://en.wikipedia.org/wiki/Talk:Slavery_in_the_21st_century>.
- [Drupal] Xano, ., "Replace 'master-slave' terminology with 'primary/replica'", 2014, <<https://www.drupal.org/project/drupal/issues/2275877>>.
- [Grewal] Grewal, D., "The 'Bad Is Black' Effect", 2017, <<https://www.scientificamerican.com/article/the-bad-is-black-effect/>>.

- [socketwench] socketwench, ., "Even in tech, words matter", 2018, <<https://deninet.com/blog/2018/09/09/even-tech-words-matter>>.
- [ISC] Internet Systems Consortium, ., "@ISCdotORG reply tweet", 2017, <<https://twitter.com/ISCdotORG/status/943152507211071489>>.
- [Github] Kevin Truong, . and VICE, "Github to Remove 'Master/Slave' Terminology From its Platform", June 2020, <https://www.vice.com/en_us/article/k7qbyv/github-to-remove-masterslave-terminology-from-its-platform>.
- [NIST] Eric Geller, . and Politico, "Agency to end use of technology terms such as 'master' and 'slave' over racist associations", June 2020, <<https://www.politico.com/news/2020/06/25/agency-ends-use-technology-terms-racist-associations-339880>>.
- [inclusiveterminology] IETF, "Inclusive terminology in IETF Documents", August 2020, <<https://github.com/ietf/terminology>>.
- [Statement] IESG, "IESG Statement on on Inclusive Language", May 2021, <<https://www.ietf.org/about/groups/iesg/statements/on-inclusive-language/>>.
- [in-solidarity-bot] IETF, "Inclusive terminology in IETF Documents", August 2020, <<https://github.com/ietf/terminology>>.

Authors' Addresses

Mallory Knodel
Center for Democracy & Technology
Email: mknodel@cdt.org

Niels ten Oever
University of Amsterdam
Email: mail@nielstenoever.net