

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 5, 2010

JM. Valin
Octasic Inc.
S. Borilin
SPIRIT DSP
K. Vos
Skype
C. Montgomery
Xiph.Org Foundation
R. Chen
Broadcom Corporation
March 4, 2010

Guidelines for the Codec Development Within the IETF
draft-valin-codec-guidelines-03

Abstract

This document provides general guidelines for work on developing and specifying a codec within the IETF. These guidelines cover the development process, evaluation, requirements conformance, and intellectual property issues.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on September 5, 2010.

Copyright Notice

Copyright (c) 2010 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 BSD License.

Table of Contents

- [1.](#) Introduction [3](#)
- [2.](#) Development Process [4](#)
- [3.](#) Evaluation, Testing, and Characterization [7](#)
- [4.](#) Requirements Conformance [8](#)
- [5.](#) Intellectual Property [10](#)
- [6.](#) Relationship with Other SDOs [13](#)
- [7.](#) Security Considerations [15](#)
- [8.](#) IANA Considerations [16](#)
- [9.](#) Acknowledgments [17](#)
- [10.](#) References [18](#)
 - [10.1.](#) Normative References [18](#)
 - [10.2.](#) Informative References [18](#)
- Authors' Addresses [20](#)

1. Introduction

This document describes a suggested process for work at the IETF on standardization of a codec that is optimized for use in interactive Internet applications and that can be widely implemented and easily distributed among application developers, service operators, and end users. Nothing in this document shall be taken to override official IETF policies and procedures as codified in [BCP 9](#), [BCP 78](#), [BCP 79](#), or any other approved document. In case the suggestions in this document conflict with official policies and procedures, the official policies and procedures shall rule.

The authors welcome discussion and comments related to the topics presented in this document. The preferred venue is the <codec@ietf.org> mailing list, for which archives and subscription information are available at <https://www.ietf.org/mailman/listinfo/codec>.

2. Development Process

The process outlined here is intended to maximize the transparency of work on a codec within the IETF. Such work might involve development of a completely new codec, adaptation of an existing codec to meet the requirements, or integration between two or more existing codecs that results in an improved codec combining the best aspects of each codec. To enable such procedural transparency, the contributor of an existing codec must be willing to cede change control to the IETF and should have sufficient knowledge of the codec to assist in the work of adapting it or applying some of its technology to the development or improvement of other codecs. Furthermore, contributors need to be aware that any codec that results from work within the IETF is likely to be different from any existing codec that was contributed to the Internet Standards Process.

Work on codec development is expected to proceed as follows:

1. IETF participants will identify the requirements to be met by an Internet codec, in the form of an Internet-Draft.
2. Interested parties will actively solicit the contribution of existing or proposed new codecs to the Internet Standards Process. Ideally these contributions will occur in the form of Internet-Drafts to enable the widest community review, although any contributions made in accordance with [BCP 78](#) and [BCP 79](#) will be acceptable (e.g., via mailing list or a presentation during a WG session). If there is any IPR related to a contributed codec, the contribution should be accompanied by an appropriate IPR disclosure as specified in [BCP79](#) (IPR issues are described in more detail under [Section 5](#)).
3. As contributions are received and discussed within the working group, the group will gain a clearer understanding of what is achievable within the design space. As a result, the authors of the requirements document should iteratively clarify and improve their document to reflect the emerging working group consensus. This will likely involve collaboration with IETF working groups in other areas, such as collaboration with working groups in the Transport area to identify important aspects of packet transmission over the Internet, with working groups in the RAI area to ensure that information about and negotiation of the codec can be easily represented at the signalling layer, and with working groups in the Transport area to understand the degree of rate adaptation desirable. In parallel with this work, interested parties should evaluate the contributions at a higher level to see which requirements might be met by each codec.

4. Once a sufficient number of proposals has been received, the interested parties will identify the strengths, weaknesses, and innovative aspects of the contributed codecs. This step will consider not only the codecs as a whole, but also key features of the individual algorithms (predictors, quantizers, transforms, etc.).
5. It is expected that none of the contributed codecs will meet all of the defined requirements. Therefore, it is expected that IETF participants will choose a `_starting point_` for the reference implementation to facilitate the development process. This starting point codec will meet as many of the requirements as possible, but probably will need to be adjusted through an iterative development process in order to meet all of the requirements (or as many requirements as possible). The starting point codec might be one of the contributed codecs (especially if it is the only codec that meets most of the requirements), a combination of two or more of the contributed codecs, or an entirely new codec. None of the decisions taken at this step will be definitive. In particular, IETF participants will not provide a "rubber stamp" for any contributed codec.
6. IETF participants will attempt to iteratively improve each component of the starting point reference implementation, where by "component" we mean individual algorithms such as predictors, transforms, quantizers, and entropy coders. The participants will proceed by trying new designs, applying ideas from the contributed codecs, evaluating "proof of concept" ideas, and using their expertise in codec development to improve the starting point codec. Any aspect of the starting point codec might be changed (even the fundamental principles of the codec) or the participants might start over entirely by scrapping the starting point codec and designing a completely new one. The overriding goal shall be to design a codec that will meet the requirements defined in the requirements document. Given the IETF's open standards process, any interested party will be able to contribute to this work, whether or not they submitted an Internet-Draft for one of the contributed codecs. The codec itself will be normatively specified with code in an Internet-Draft.
7. In parallel with work on the codec reference implementation, developers and other interested parties should perform evaluation of the codec as described under [Section 3](#), IETF participants should define (within the AVT Working Group) the codec's payload format for use with the Real-time Transport Protocol [[RTP](#)], and application developers should start testing the codec by implementing it in code and deploying it in actual Internet

applications to identify any potential problems.

8. Once IETF participants agree that the codec being developed meets the requirements (e.g., via a working group last call), IETF participants can begin the task of characterizing the codec. The characterization process is described under [Section 3](#).

3. Evaluation, Testing, and Characterization

Lab evaluation of the codec being developed should happen throughout the development process because it will help ensure that progress is being made toward fulfillment of the requirements. There are many ways in which continuous evaluation can be performed. For minor, uncontroversial changes to the codec it should usually be sufficient to use objective measurements (e.g., PESQ, PEAQ, and SegSNR) validated by informal subjective evaluation. For more complex changes (e.g., when psychoacoustic aspects are involved) or for controversial issues, internal testing should be performed. An example of internal testing would be to have individual participants rate the decoded samples using one of the established testing methodologies, such as ITU-R BS.1534 (MUSHRA).

Throughout the process, it will be important to make use of the Internet community at large for real-world distributed testing. This will enable many different people with different equipment and use cases to test the codec and report any problems they experience. In the same way, third-party software developers will be encouraged to integrate the codec (with a warning about the bit-stream not being final) and provide feedback on its performance in real-world use cases.

Characterization of the final codec must be based on the reference implementation only (and not on any "private implementation"). This can be performed by independent testing labs or, if this is not possible, using the testing labs of the organizations that contribute to the Internet Standards Process. Packet loss robustness should be evaluated using actual loss patterns collected from use over the Internet, rather than theoretical models. The goals of the characterization phase are to:

- o ensure that the requirements have been fulfilled
- o guide the IESG in its evaluation of the resulting work
- o assist application developers in understanding whether the codec is suitable for a particular application

4. Requirements Conformance

It is the responsibility of the working group to define criteria for evaluating conformance, including but not limited to comparison tools and test vectors. The following text provides suggestions for consideration by the working group.

First, any codec specified by the IETF must include source code for a normative reference implementation, specified in an Internet-Draft. This will aid in interoperability testing and requirements conformance.

Second, it is best to use the minimum number of normative requirements that will ensure complete interoperability between implementations. In practice this generally means that only the decoder needs to be normative, so that the encoder can improve over time. This also enables different tradeoffs between quality and complexity.

Third, to reduce the risk of bias towards certain CPU/DSP architectures, ideally the decoder specification should not require "bit-exact" conformance with the reference implementation. The output of a decoder implementation should only be "close enough" to the output of the reference decoder. A comparison tool should be provided along with the codec to verify objectively that the output of a decoder is likely to be perceptually indistinguishable from that of the reference decoder. However, an implementation may still wish to produce an output that is bit-exact with the reference implementation to simplify the testing procedure. The packet loss concealment (PLC) algorithm need not be normative. However, if any part of the PLC requires specific information transmitted by the encoder in the bit-stream, then the corresponding aspect of the reference encoder must be normative.

Fourth, an encoder implementation should not be required to make use of all the "features" (tools) in the bit-stream definition. However, the codec specification may require that an encoder implementation be able to generate any possible bit-rate. Unless a particular "profile" is defined in the specification, the decoder must be able to decode all features of the bit-stream. The decoder must also be able to handle any combination of bits, even combinations that cannot be generated by the reference encoder. It is recommended that the decoder specification shall define exactly how the decoder should react to "impossible" packets. However, an encoder must never generate such packets that do not conform to the bit-stream definition.

Compressed test vectors should be provided as a means to verify

conformance with the decoder specification. These test vectors should exercise all paths in the decoder (100% code coverage).

While the exact encoder will not be specified, it is recommended to specify objective measurement targets for an encoder, below which use of a particular encoder implementation is not recommended. For example, one such specification could be: "the use of an encoder whose PESQ MOS is less than 0.1 below the reference encoder in the following conditions is not recommended".

5. Intellectual Property

Producing an unencumbered codec is desirable for the following reasons:

- o It is the experience of a wide variety of application developers and service providers that encumbrances such as licensing and royalties make it difficult to implement, deploy, and distribute audio applications for use by the Internet community.
- o It is beneficial to have low-cost options whenever possible because standalone voice services are being commoditized and small, innovative development teams often cannot afford to pay per-channel licensing fees and royalties.
- o Many market segments are moving away from selling hard-coded hardware devices and toward freely distributing end-user software; this is true of numerous large application providers and even telcos themselves.
- o Compatibility with the licensing of typical open source applications implies the need to avoid encumbrances, including even the requirement to obtain a license for implementation, deployment, or use (even if the license does not require the payment of a fee).

Therefore, a codec that can be widely implemented and easily distributed among application developers, service operators, and end users is preferred. Many existing codecs that might fulfill some or most of the technical attributes listed above are encumbered in various ways. For example, patent holders might require that those wishing to implement the codec in software, deploy the codec in a service, or distribute the codec in software or hardware need to request a license, enter into a business agreement, pay licensing fees or royalties, or adhere to other special conditions or restrictions. Because such encumbrances have made it difficult to widely implement and easily distribute high-quality audio codecs across the entire Internet community, the working group prefers unencumbered technologies in a way that is consistent with [BCP 78](#) and [BCP 79](#). In particular, the working group shall heed the preference stated in [BCP 79](#): "In general, IETF working groups prefer technologies with no known IPR claims or, for technologies with claims against them, an offer of royalty-free licensing." Although this preference cannot guarantee that the working group will produce an unencumbered codec, the working group shall follow [BCP 79](#), and adhere to the spirit of [BCP 79](#). The working group cannot explicitly rule out the possibility of adopting encumbered technologies; however, the working group will try to avoid encumbered technologies

that require royalties or other encumbrances that would prevent such technologies from being easy to redistribute and use.

The following guidelines will help to maximize the odds that the codec will be unencumbered:

1. In accordance with [BCP 79](#) [[IPR](#)], contributed codecs should preferably use technologies with no known IPR claims or technologies with an offer of royalty-free (RF) licensing.
2. Given two, nearly equivalent technologies, the working group should prefer technologies where any intellectual property rights have expired. In most jurisdictions, that means technologies where the first publication was at least twenty years ago. This protects the work from both current IPR claims as well as potential IPR the group is not aware of.
3. Whenever possible, the working group should use technologies that are perceived by the participants to be safer with regard to IPR issues.
4. Contributors must disclose IPR as specified in [BCP 79](#).
5. If a technology under consideration is known to be covered by a patent, the patent holder should be contacted and asked to license the patent under acceptable RF terms (if the patent holder is also a contributor, the license may have already been specified in the IPR disclosure).
6. In cases where no RF license can be obtained regarding a patent, the group should consider alternative algorithms or methods, even if they result in lower quality, higher complexity, or otherwise less desirable characteristics (in most cases, the degradation will likely be small once the best alternative has been identified).
7. In accordance with [BCP 78](#) [[TRUST](#)], the source code for the reference implementation should be made available under a BSD-style license (or whatever license is defined as acceptable by the IETF Trust when the Internet-Draft defining the reference implementation is published).

IETF participants should be aware that, given the way patents work in most countries, the resulting codec can never be guaranteed to be free of patent claims because some patents may not be known to the contributors, some patent applications may not be disclosed at the time the codec is developed, and only courts of law can determine the validity and breadth of patent claims. However, these observations

are no different within the Internet Standards Process than they are for standardization of codecs within other SDOs (or development of codecs outside the context of any SDO), and furthermore are no different for codecs than for other technologies worked on within the IETF. In all these cases, the best approach is to minimize the risk of unknowingly incurring encumbrance on existing patents. Despite these precautions, participants need to understand that, practically speaking, it is nearly impossible to guarantee that implementors will not incur encumbrance on existing patents.

6. Relationship with Other SDOs

It is understood that other SDOs are also involved in the codec development and standardization, including but not necessarily limited to:

- o The Telecommunication Standardization Sector (ITU-T) of the International Telecommunication Union (ITU), in particular Study Group 16
- o The Moving Picture Experts Group (MPEG)
- o The European Telecommunications Standards Institute (ETSI)
- o The 3rd Generation Partnership Project (3GPP)
- o The 3rd Generation Partnership Project 2 (3GPP2)

It is important to ensure that such work does not constitute uncoordinated protocol development, of the kind described in [\[UNCOORD\]](#) in the following principle:

[T]he IAB considers an essential principle of the protocol development process that only one SDO maintains design authority for a given protocol, with that SDO having ultimate authority over the allocation of protocol parameter code-points; defining the intended semantics, interpretation, and actions associated with those code-points.

The work envisioned by this guidelines document is not "uncoordinated" in the sense described in the foregoing quote, for the following reasons:

- o Internet signalling technologies are designed to enable the negotiation of any codecs that are supported in a particular application (such signalling technologies include the Session Initiation Protocol [\[SIP\]](#), Session Description Protocol [\[SDP\]](#), and the Extensible Messaging and Presence Protocol [\[XMPP\]](#) extensions for media negotiation as specified in [\[Jingle\]](#)).
- o Internet transport technologies such as the Real-time Transport Protocol [\[RTP\]](#) (including secure transport as described in [\[SRTP\]](#)) are designed to support any codec for which RTP packetization rules have been defined.
- o To the extent that codec work at the IETF differs from work at other SDOs, it will focus on issues that are specific to the Internet, including robustness to packet loss and other aspects of

packet transmission over the Internet. Issues that are specific to non-Internet transports (e.g., radio communication and circuit-switched networks) are specifically out of scope.

Although there is already sufficient codec expertise available among IETF participants to complete the envisioned work, additional contributions are welcome within the framework of the Internet Standards Process, in the following ways:

- o Individuals who are technical contributors to codec work within other SDOs can participate directly in codec work within the IETF.
- o Other SDOs can contribute their expertise (e.g., codec characterization and evaluation techniques) and thus facilitate the testing of a codec produced by the IETF.
- o Any SDO can provide input to IETF work through liaison statements.

However, it is important to note that final responsibility for the development process and the resulting codec will remain with the IETF as governed by [BCP 9](#) [[PROCESS](#)].

Finally, there is precedent for the contribution of codecs developed elsewhere to the ITU-T (e.g., AMR Wideband was standardized originally within 3GPP). This is a model to explore as the IETF coordinates further with the ITU-T in accordance with the collaboration guidelines defined in [[COLLAB](#)].

[7.](#) Security Considerations

The procedural guidelines for codec development do not have security considerations. However, the resulting codec needs to take appropriate security considerations into account, for example as outlined in [[DOS](#)] and [[SECGUIDE](#)].

[8.](#) IANA Considerations

This document has no actions for IANA.

[9.](#) Acknowledgments

We would like to thank all the other people who contributed directly or indirectly to this document, including Jason Fischl, Gregory Maxwell, Alan Duric, Jonathan Christensen, Julian Spittka, Michael Knappe, Timothy Terriberry, Christian Hoene, Stephan Wenger and Henry Sinnreich. We also like to thank Cullen Jennings and Gregory Lebovitz for their advice. Special thanks to Peter Saint-Andre, who originally co-authored this document.

[10.](#) References

[10.1.](#) Normative References

- [IPR] Bradner, S., "Intellectual Property Rights in IETF Technology", [BCP 79](#), [RFC 3979](#), March 2005.
- [PROCESS] Bradner, S., "The Internet Standards Process -- Revision 3", [BCP 9](#), [RFC 2026](#), October 1996.
- [TRUST] Bradner, S. and J. Contreras, "Rights Contributors Provide to the IETF Trust", [BCP 78](#), [RFC 5378](#), November 2008.

[10.2.](#) Informative References

- [COLLAB] Fishman, G. and S. Bradner, "Internet Engineering Task Force and International Telecommunication Union - Telecommunications Standardization Sector Collaboration Guidelines", [RFC 3356](#), August 2002.
- [DOS] Handley, M., Rescorla, E., and IAB, "Internet Denial-of-Service Considerations", [RFC 4732](#), December 2006.
- [Jingle] Ludwig, S., Saint-Andre, P., Egan, S., McQueen, R., and D. Cionoiu, "Jingle RTP Sessions", XSF XEP 0167, June 2009.
- [RTP] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, [RFC 3550](#), July 2003.
- [SDP] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", [RFC 4566](#), July 2006.
- [SECGUIDE] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", [BCP 72](#), [RFC 3552](#), July 2003.
- [SIP] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", [RFC 3261](#), June 2002.
- [SRTP] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", [RFC 3711](#), March 2004.
- [UNCOORD] Bryant, S. and M. Morrow, "Uncoordinated Protocol

Development Considered Harmful", [RFC 5704](#), November 2009.

[XMPP] Saint-Andre, P., Ed., "Extensible Messaging and Presence Protocol (XMPP): Core", [RFC 3920](#), October 2004.

Authors' Addresses

Jean-Marc Valin
Octasic Inc.
4101, Molson Street
Montreal, Quebec
Canada

Email: jean-marc.valin@octasic.com

Slava Borilin
SPIRIT DSP

Email: borilin@spiritdsp.net

Koen Vos
Skype

Email: koen.vos@skype.net

Christopher Montgomery
Xiph.Org Foundation

Email: xiphmont@xiph.org

Raymond (Juin-Hwey) Chen
Broadcom Corporation

Email: rchen@broadcom.com