

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 10, 2010

JM. Valin
Octasic Inc.
P. Saint-Andre
Cisco
S. Borilin
SPIRIT DSP
K. Vos
Skype
C. Montgomery
Xiph.Org Foundation
R. Chen
Broadcom Corporation
October 7, 2009

Guidelines for Proposed CODEC Working Group
draft-valin-codec-guidelines-01

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/lid-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 April 10, 2010.

Copyright Notice

Copyright (c) 2009 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 in effect on the date of

Internet-Draft

Codec Guidelines

October 2009

publication of this document (<http://trustee.ietf.org/license-info>).
Please review these documents carefully, as they describe your rights
and restrictions with respect to this document.

Internet-Draft

Codec Guidelines

October 2009

Abstract

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

Table of Contents

1.	Introduction	4
2.	Development Process	6
3.	Evaluation, Testing, and Characterization	9
4.	Requirements Conformance	10
5.	Intellectual Property	12
6.	Relationship with Other SDOs	14
7.	Security Considerations	16
8.	IANA Considerations	17
9.	Acknowledgments	18
10.	References	19
10.1.	Normative References	19
10.2.	Informative References	19
	Authors' Addresses	21

1. Introduction

This document describes a suggested process for the specification of easily distributable audio codecs within the Internet Standards Process [[PROCESS](#)]. Although several easily distributable audio codecs have been developed "in the wild" outside the IETF or any other standards development organization (SDO), there are several reasons for pursuing such work within the IETF:

- o The IETF has produced a large "stack" of easily distributable protocols and formats for use in a wide variety of Internet applications. However, feedback received from application developers and service operators indicates that the lack of standardized and easily distributable audio codecs has inhibited the development and deployment of interoperable Internet audio applications.
- o Specification of such codecs within the Internet Standards Process would ensure more transparent, stringent, and trusted change control than is currently available "in the wild".
- o In order to develop audio codecs that are optimized for use over the Internet, it is important to gain input from the entire Internet community and to incorporate cross-area review from the full range of technical experts who work within the IETF, including areas such as transport, routing, operations, architecture, applications, and security.
- o The Internet Standards Process provides a transparent process for

the development of Internet technologies, in which any interested member of the Internet community can participate.

- o The IETF's policies toward Intellectual Property Rights (IPR), as codified in [BCP 79](#) [[IPR](#)], facilitate the development of Internet audio codecs that will be easily distributable to all parties that might want to implement or deploy such codecs.

This document describes suggested processes for work on development of easily distributable Internet audio codecs within the Internet Standards Process, in particular as guidance to the proposed Codec Working Group. 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.

Although the authors of this document consider formation of a working group to be the most productive path forward because of the

consensus-building mechanisms inherent to working groups, work could also be pursued through independent submissions, to which the procedural guidelines suggested here would apply.

In addition, although this document sometimes speaks of "codec" in the singular in order to simplify the text, IETF participants could develop more than one codec; any decisions about the number of codecs to work on would be subject to normal consensus processes within the IETF.

The authors welcome discussion and comments related to the topics presented in this document. The preferred venue is the [<codec@ietf.org>](mailto: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 codecs within the IETF. Such work might involve development of one or more completely new codecs, adaptation of an existing codec to meet the requirements specified in the accompanying requirements document [[REQS](#)], 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 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. One or more IETF participants will identify the requirements to be met by Internet codecs, 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 (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 [[REQS](#)] should iteratively clarify and improve their document to reflect the emerging working group consensus. In parallel, 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 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 audio applications to identify any potential problems.
8. Once IETF participants agree that the codec being developed meets the requirements (e.g., through consensus calls on the relevant discussion list or via a working group last call if a working group is formed), IETF participants can begin the task of characterizing the codec. The characterization process is

3. Evaluation, Testing, and Characterization

Lab evaluation of the codecs 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 listening to a few samples. 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 audio 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 would not have a "bit-exact" definition. The output of a decoder implementation should only be "close enough" to the output of the reference decoder. A comparison tool should be used to verify objectively that the output of a decoder is likely to be perceptually indistinguishable from that of the reference decoder. The packet loss concealment (PLC) algorithm need not be normative. However, if any part of the PLC requires specific information (in the bit-stream) transmitted by the encoder, then the corresponding aspect of the reference encoder must be normative.

Fourth, an encoder implementation should not be required to be able to generate all the "features" in the bit-stream definition. However, it may be required (by the codec specification) to be able to generate any possible bit-rates. Unless any "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 if that combination cannot be generated by the reference encoder. It is recommended that the decoder

specification defines exactly how the decoder should react to "impossible" packets.

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

The codec should be easily distributable, with as few restrictions as possible. This can be achieved by handling IPR issues in accordance with [BCP 78](#) [[TRUST](#)] and [BCP 79](#) [[IPR](#)]. IPR is a very important issue for the following reasons:

- o The IETF already provides a comprehensive "stack" of easily distributable protocols for communication over the Internet. However, there is a conspicuous gap in this stack for media codecs.
- o Many market segments are moving from selling hard-coded hardware devices to freely distributing end-user software, including large application providers and even telcos themselves.
- 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 royalties.
- o Most applications provided in the Internet space will benefit from easily distributable codecs.

In accordance with [BCP 78](#) [[TRUST](#)], the source code for the reference

implementation should be available under the BSD license or whatever license is defined as acceptable by the IETF Trust when the relevant Internet-Drafts are published. Also, in accordance with [BCP 79 \[IPR\]](#), the codecs should preferably use technologies with no known IPR claims or technologies with an offer of royalty-free (RF) licensing. Compatibility with the licensing of typical open source applications requires the avoidance of restrictive IPR claims.

There are several ways to maximize the odds that the codec will be royalty-free. First, when 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. In cases where no RF license can be obtained regarding a patent, then alternative algorithms or methods should be considered, 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.

IETF participants should be aware that, given the way patents work in most countries, the resulting codecs can never be guaranteed to be free of patent claims because:

- o Some patents may not be known to the contributors.
- o Some patents applications may not be disclosed at the time the codec is developed.
- o Only the courts 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 are no different 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.

To help minimize the risk of unknowingly incurring encumbrance on existing patents, the WG should prefer technologies that were published more than twenty years in the past over more recent

technologies, at least when the two are mostly equivalent. Similarly, whenever possible the WG should use technologies that are perceived to be safer by the participants. Despite these precautions, the 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 development and standardization of media codecs, 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)

By specifying guidelines for working on media codecs within the Internet Standards Process, IETF participants are not attempting to duplicate work being performed by other SDOs. Clearly, there is no "natural monopoly" over codec work, since SDOs such as ITU-T, MPEG, ETSI, 3GPP, and 3GPP2 have defined codecs in parallel.

With regard to Study Group 16 of the ITU-T, the work envisioned within the Internet Standards Process differs as follows:

- o The codec(s) produced within the Internet Standards Process will specifically address problems related to transmission on the Internet:
 - * The codec(s) will be optimized for transport using the Real-time Transport Protocol [[RTP](#)], including secure transport as described in [[SRTP](#)]
 - * The codec(s) will take into account end-to-end parameter negotiation using Internet signalling technologies such as Session Initiation Protocol [[SIP](#)], Session Description Protocol [[SDP](#)], and the Extensible Messaging and Presence Protocol [[XMPP](#)] extensions for media negotiation as specified in [[Jingle](#)]
 - * Robustness to packet loss will be a very important aspect to consider
 - * Purely wireless applications will not be considered

- * Issues with circuit-switched networks will not be considered
- o As discussed under [Section 5](#), IPR issues will be handled in accordance with IETF IPR policies, with the goal of developing

easily distributable codecs

With regard to MPEG, the work envisioned within the Internet Standards Process differs as follows:

- o The focus will be on interactive audio transmission rather than storage
- o As discussed under [Section 5](#), IPR issues will be handled in accordance with IETF IPR policies, with the goal of developing easily distributable codecs
- o The emphasis will be on the quality of the reference implementations (encoder and decoder), which will be the basis for characterization

Although there is already sufficient expertise available among participants who are active within the IETF, additional contributions are welcome. Because IETF work happens through the efforts of individuals, individuals who contribute to the work of other SDOs, as well as the SDOs themselves, are welcome to contribute to the work within the Internet Standards Process, in the in the following ways:

- o Contributors to other SDOs can participate in the Internet Standards Process, as can any other interested parties
- o The ITU-T can contribute its existing characterization and evaluation techniques, methods, experience and expertise, and thus facilitate the testing process
- o Any other SDO can provide input to the process through liaison statements.

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

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, and Henry Sinnreich. We also like to thank Cullen Jennings and Gregory Lebovitz for their advice.

[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

- [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.
- [REQS] Valin, J., Borilin, S., Vos, K., Montgomery, C., and J. Chen, "Codec Requirements", [draft-valin-codec-requirements-00](#) (work in progress), September 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.
- [XMPP] Saint-Andre, P., Ed., "Extensible Messaging and Presence

Valin, et al.

Expires April 10, 2010

[Page 20]

Internet-Draft

Codec Guidelines

October 2009

Authors' Addresses

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

Email: jean-marc.valin@octasic.com

Peter Saint-Andre

Cisco

Email: Peter.SaintAndre@WebEx.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