Network Working Group                                    JM. Valin
Internet-Draft                                             Mozilla
Intended status: Informational                          S. Borilin
Expires: July 12, 2012                                 SPIRIT DSP
                                                            K. Vos

                                                    C. Montgomery
                                              Xiph.Org Foundation
                                                         R. Chen
                                            Broadcom Corporation
                                                 January 9, 2012

### Guidelines for Development of an Audio Within the IETF
### draft-ietf-codec-guidelines-07

Abstract

   This document provides general guidelines for work on developing and
   specifying an interactive audio 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 in full conformance with the
   provisions of BCP 78 and BCP 79.

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

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

   This Internet-Draft will expire on July 12, 2012.

Table of Contents

## 1.  Introduction

This document describes a process for work in the IETF codec WG on
standardization of an audio 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.

2.  **Development Process**

   The process outlined here is intended to make the work on a codec
   within the IETF transparent, predictable, and well organized, in a
   way that is consistent with [PROCESS].  Such work might involve
   development of a completely new codec, adaptation of an existing
   codec to meet the requirements of the working group, or integration
   between two or more existing codecs that results in an improved codec
   combining the best aspects of each.  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 improvemnet 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 an interactive audio 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 are encouraged to make contributions proposing
       existing or new codecs, or elements thereof, to the codec WG as
       long as these contributions are within the scope of the WG.
       Ideally, these contributions should be in the form of Internet
       Drafts, although other forms of contributions are also possible,
       as discussed in [PROCESS].

   3.  Given the importance of IPR to the activities of the working
       group, any IPR disclosures must be made in a timely way.
       Contributors are required, as described in [IPR], to disclose any
       known IPR, both first and third party.  Timely disclosures are
       particularly important, since those disclosures may be material
       to the decision process of the working group.

   4.  As contributions are received and discussed within the working
       group, the group should 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 is likely to 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 and to understand the degree of
       rate adaptation desirable, and with working groups in the RAI

area to ensure that information about and negotiation of the
codec can be easily represented at the signaling layer.  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.

5.  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.).

6.  Interested parties are encouraged to collaborate together and
    combine the best ideas from the various codec contributions into
    a consolidated codec definition, representing the merging of some
    of the contributions.  Through this iterative process, the number
    of proposals will reduce and consensus will generally form around
    one of them.  At that point, the working group should adopt that
    document as a working group item, forming a baseline codec.

7.  IETF participants should then attempt to iteratively add to or
    improve each component of the baseline codec reference
    implementation, where by "component" we mean individual
    algorithms such as predictors, transforms, quantizers, and
    entropy coders.  The participants should 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 baseline codec.  Any aspect of the
    baseline codec might be changed (even the fundamental principles
    of the codec) or the participants might start over entirely by
    scrapping the baseline 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
    [codec-requirements].  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 should be normatively
    specified with code in an Internet-Draft.

8.  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 PAYLOAD Working Group) the codec's
    payload format for use with the Real-time Transport Protocol
    [RTP].  Ideally, application developers should test the codec by
    implementing it in code and deploying it in actual Internet
    applications.  Unfortunately, developers will frequently wait

until RFC or until a stable bitstream is guaranteed before
deployment.  As such, this is a nice-to-have and not a
requirement for this process.  Lab implementations are certainly
encouraged.

9.  The group will produce a testing results document.  The document
will be a living document that captures testing done before the
codec stabilized, after it has stabilized, and after the codec
specification is issued as an RFC.  The document serves the
purpose of helping the group determine whether the codec meets
the requirements.  Any testing done after the codec RFC is issued
helps implementors understand the final performance of the codec.
The process of testing is described in 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 [ITU-T P.862], PEAQ [ITU-R
   BS.1387-1], and segmental signal-to-noise ratio) 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 developers will be encouraged to integrate
   the codec into their software (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

   The exact methodology for the characterization phase can be
   determined the working group.  Because the IETF does not have testing
   resources of its own, it has to rely on the resources of its
   participants.  For this reason, even if the group agrees that a
   particular test is important, if no one volunteers to do it, or if
   volunteers do not complete it in a timely fashion, then that test
   should be discarded.  This ensures that only important tests be done,

and in particular those tests which are important to participants.

4.  Specifying the Codec

   Specifying a codec requires careful consideration around what is
   required vs. what is left to the implementation.  The following text
   provides guidelines for consideration by the working group:

   1.  Any audio codec specified by the codec Working Group must include
       source code for a normative software implementation, documented
       in an Internet Draft destined for standards track RFC.  This
       implementation will be used to verify conformance of an
       implementation.  Although a text description of the algorithm
       should be provided, its use should be limited to helping the
       reader in understanding the source code.  Should the description
       contradict the source code, the latter shall take precedence.
       For convenience, the source code may be provided in compressed
       form, with base64 [BASE64] encoding.

   2.  Because of the size and complexity of most codecs, it is possible
       that even after publishing the RFC, bugs will be found in the
       reference implementation, or differences between the
       implementation and the text description.  As usual, an errata of
       the RFC should be maintained.  Although a public software
       repository containing the current reference implementation is
       desirable, the normative implementation would still be the RFC.

   3.  It is the intention of the group to allow the greatest possible
       choice of freedom in implementing the specification.
       Accordingly, the number of binding RFC2119 keywords will be the
       minimum that still allows interoperable 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 trade-offs between quality and complexity.

   4.  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.  In that case, the
       output of a decoder implementation should only be "close enough"
       to the output of the reference decoder and 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.  An
       implementation may still wish to produce an output that is bit-
       exact with the reference implementation to simplify the testing
       procedure.

   5.  To ensure freedom of implementation, decoder-side only error
       concealment does not need to be specified, although the reference
       implementation should include the same packet loss concealment

(PLC) algorithm as used in the testing phase.  Is it up to the
working group to decide whether minimum requirements on PLC
quality will be required for compliance with the specification.
Obviously, any information signaled in the bitstream intended to
aid PLC needs to be specified.

6.  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 how the decoder should react to "impossible" packets (e.g.
reject, consider as valid).  However, an encoder must never
generate such packets that do not conform to the bit-stream
definition.

7.  Compressed test vectors should be provided as a means to verify
conformance with the decoder specification.  These test vectors
should be designed to exercise as much of the decoder code as
possible.

8.  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 better 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
       multimedia applications for use by the Internet community.

    o  It is beneficial to have low-cost options whenever possible,
       because innovation - the hallmark of the Internet - is hampered
       when small development teams cannot deploy an application because
       of usage-based 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 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.

When considering license terms for technologies with IPR claims against them, some members of the working group have expressed their preference for license terms which:

o  are available to all, worldwide, whether or not they are working group participants

o  extend to all essential claims owned or controlled by the licensor

o  do not require payment of royalties, fees or other consideration

o  do not require licensees to adhere to restrictions on usage (though, licenses which apply only to implementation of the standard are acceptable)

o  do not otherwise impede the ability of the codec to be implemented in open-source software projects

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 licensing.

2.  As described in BCP 79, the working group should use technologies that are perceived by the participants to be safer with regard to IPR issues.

3.  Contributors must disclose IPR as specified in BCP 79.

4.  In cases where no royalty-free license can be obtained regarding a patent, BCP 79 suggests that the working group consider alternative algorithms or methods, even if they result in lower quality, higher complexity, or otherwise less desirable characteristics.

5.  In accordance with BCP 78 [TRUST], the source code for the reference implementation must 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).

Many IPR licenses specify that a license is granted only for technologies which are adopted by the IETF as a standard.  While

reasonable, this has the unintended side-effect of discouraging
implementation prior to RFC status.  Real-world implementation is
beneficial for evaluation of the codec.  As such, entities making IPR
license statements are encouraged to use wording which permits early
implementation and deployment.

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 implementers
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) of the International
      Organization for Standardization and International
      Electrotechnical Commission (ISO/IEC)

   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, since the intention of
   this process is that two possible outcomes might occur:

   1.  The IETF adopts an existing audio codec, and specifies that it is
       the "anointed" IETF Internet codec.  In such a case, codec
       ownership lies entirely with the SDO which produced the codec,
       and not with the IETF, OR

   2.  The IETF produces a new codec.  Even if this codec uses concepts,
       algorithms, or even source code from a codec produced by another
       SDO, the IETF codec is a specification unto itself and under
       complete control of the IETF.  Any changes or enhancements made
       by the original SDO to the codecs whose components the IETF used
       are not applicable to the IETF codec.  Such changes would be
       incorporated as a consequence of a revision or extension of the
       IETF RFC.  In no case should the new codec reuse a name or code

      point from another SDO.

   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, as outlined in [SECGUIDE] and in the security considerations of [codec-requirements].  More specifically, the resulting codec must avoid being subject to denial of service [DOS] and buffer overflows, and should take into consideration the impact of VBR [srtp-vbr].

## 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 B. 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.

[RTP]       Schulzrinne, H., Casner, S., Frederick, R., and V.
            Jacobson, "RTP: A Transport Protocol for Real-Time
            Applications", STD 64, RFC 3550, July 2003.

[BASE64]    Josefsson, S., "The Base16, Base32, and Base64 Data
            Encodings", RFC 4648, October 2006.

[SECGUIDE]
            Rescorla, E. and B. Korver, "Guidelines for Writing RFC
            Text on Security Considerations", BCP 72, RFC 3552,
            July 2003.

[codec-requirements]
            Valin, JM. and K. Vos, "Requirements for an Internet Audio
            Codec", RFC 6366, August 2011.

[srtp-vbr]
            Perkins, C. and JM. Valin, "Guidelines for the use of
            Variable Bit Rate Audio with Secure RTP",
            draft-ietf-avtcore-srtp-vbr (work in progress), July 2011.

[UNCOORD]   Bryant, S. and M. Morrow, "Uncoordinated Protocol
            Development Considered Harmful", RFC 5704, November 2009.

[ITU-R BS.1534]

              "Method for the subjective assessment of intermediate
              quality levels of coding systems", ITU-R
              Recommendation BS.1534, January 2003.

   [ITU-T P.862]
              "Perceptual evaluation of speech quality (PESQ): An
              objective method for end-to-end speech quality assessment
              of narrow-band telephone networks and speech codecs",
              ITU-T Recommendation P.862, October 2007.

   [ITU-R BS.1387-1]
              "Method for objective measurements of perceived audio
              quality", ITU-R Recommendation BS.1387-1, November 2001.

Authors' Addresses

     Jean-Marc Valin
     Mozilla
     650 Castro Street
     Mountain View, CA  94041
     USA


     Email: jmvalin@jmvalin.ca


     Slava Borilin
     SPIRIT DSP


     Email: borilin@spiritdsp.net


     Koen Vos


     Email: koen.vos@skype.net


     Christopher Montgomery
     Xiph.Org Foundation


     Email: xiphmont@xiph.org


     Raymond (Juin-Hwey) Chen
     Broadcom Corporation


     Email: rchen@broadcom.com