SIP Video Profile

Bandwidth, Flow Control and Intra-frame Request Use Cases & Proposed Best Practices

Version: 1.15

Date: 3 May 2010

Contacts:

Ashish Goyal (agoyal@lifesize.com), Charles Eckel (eckelcu@cisco.com),

Matt Collier (mcollier@lifesize.com), Patrick Monfort (patrick.monfort@orange-ftgroup.com), Roni Even (ron.even.tlv@gmail.com)

Table of Contents

1. Introduction

The aim of this document is to list the use cases that are associated with the bandwidth, flow control, intra-frame request functionalities, and to provide recommendations on best current practice in this field. This document should serve as a reference for video services based on SIP/SDP, from the point of view of a user-agent with video capacity.

2. Asymmetric Negotiation

Video is fundamentally different from audio in the sense that there are use cases where asymmetric media flows are desirable and hence a mechanism is needed to allow implementations to have asymmetric media flows. Here are some typical example scenarios.

2.1 Asymmetric Bandwidth to Home

A large number of broadband home users have an asymmetric bandwidth service for which bitrates available for download are significantly greater than upload bitrates. Video is well suited to take advantage of this property by allowing receive video quality/bitrate to be much greater than transmit video quality/bitrate.

2.2 Video Encode and Decode Computational Complexity

The video encode operations are computationally significantly more expensive than the decode operations. Most video UA implementations can decode a much better resolution than they can encode.

The above scenario require the video components of SDP specifications to be expressed in a declarative fashion, i.e. the offer as well as the answer contain the maximum bitrate/profile-level the UA can support receiving rather than restricting it to being negotiated as a symmetric offer answer parameter.

The bandwidth specified using TIAS / AS – should be considered as receive bandwidth capability and not as negotiated call bandwidth. As an illustration if a UA receives an offer with the bandwidth modifier b=TIAS:128000, it would be legal for it to respond with a different capability, e.g. b=TIAS:384000 in answer. For this example offer/answer exchange, the UA may end up receiving 384kbps but transmitting only 128kbps.

The capabilities expressed in video codec parameters - e.g., profile-level / max-br/max-mbps etc. – should be considered as receive capability and not negotiated capability. As an illustration, if a UA receives offer with H.264 SDP a=fmtp:96 profile-level-id=42801d, it would be legal for it to respond with a higher capability a=fmtp:96 profile-level-id=42801f in the answer. The bandwidth specified in an SDP answer can be different from the bandwidth appearing in the associated SDP offer. In such a case, the call may end up as the above UA receiving higher resolution (say HD) but transmitting only CIF.

New implementations are recommended to start honouring max-recv-level for expressing ability to receive higher level than expressed in the profile-level field per RFC 3984-bis. Use of level-asymmetry-allowed parameter is also recommended to negotiate whether level asymmetry is allowed for the call.

However implementations are recommended to be lenient towards endpoints which do not advertise level-asymmetry-allowed and should still support asymmetric negotiation with them.

3. Bandwidth Indication

There is a need to signal the bandwidth corresponding to each video stream in the SDP.

The TIAS (RFC 3890) and AS (RFC 4566) bandwidth specifiers should be used as shown in the following example. Use of TIAS bandwidth specifier at the session level and the video m line level is mandatory. It is recommended that AS bandwidth specifier also be used at the session level for backward compatibility.

Notes:

> The maximum receive bandwidth specified in an SDP answer can be different from the bandwidth specified in the corresponding SDP offer.

For an audio stream (or when the bandwidth is implied by the codec used), bandwidth signalling at the corresponding media level is not necessary and is therefore not recommended. If no bandwidth parameter (b=) is present at the session level, then a bandwidth parameter must be present at every media level where needed.

Sample SDP specification including bandwidth parameters with TIAS and AS specifiers:

```
v=0
o=anonymous 1240218157 1240218157 IN IP4 10.193.128.35
s=-
i=myUserAgent
c=IN IP4 10.193.128.35
b=TIAS:256000
b=AS:256
t=0 0
m=audio 6000 RTP/AVP 9 8 0 18 116
a=sendrecv
a=rtpmap:9 G722/8000
a=rtpmap:8 PCMA/8000
a=rtpmap:0 PCMU/8000
a=rtpmap:18 G729/8000
a=fmtp:18 annexb=no
a=rtpmap:116 telephone-event/8000
a=fmtp:116 0-15
m=video 6002 RTP/AVP 96 97 34 31
b=TIAS:256000
a=rtpmap:96 H264/90000
a=fmtp:96 profile-level-id=428014
a=sendrecv
a=rtpmap:97 H263-1998/90000
a=fmtp:97 CIF=1;QCIF=1;I=1;J=1;T=1;N=4;K=1
```

```
a=rtpmap:34 H263/90000
a=fmtp:34 CIF=1;QCIF=1
a=rtpmap:31 H261/90000
a=fmtp:31CIF=1;QCIF=1
```

Important:

In this example, b=TIAS:256000/b=AS:256 at the session level means that the maximum capability on bandwidth for all streams is 256 Kbps.

It is possible for bandwidth at the media level to be set at 256 Kbps also, but the sender should not use total bandwidth greater than the total allowed at the session level. In this case the audio + video are limited to 256 Kbps.

Clarify: max for the session is 256 Kbps, including audio, this means the max of video is equal to 256 Kbps minus the audio bitrate used.

Note H.264

max-br: this SDP parameter is used in H.264 per RFC 3984-bis. The negotiation of a H.264 video stream must comply with RFC 3984-bis.

4. RTP/AVPF Profile

The video implementations should support RTP/AVPF profile per RFC 4585. Supporting RTP/AVPF allows implementations to use advanced RTCP mechanisms, like requesting intra frame and temporary bitrate change indication, which are essential for video streams.

Video endpoints that support RTP/AVPF profile may signal m lines with RTP/AVPF attributes yet specify the profile as RTP/AVP for backward compatibility with earlier implementations that do not support the RTP/AVFP profile. Receivers of such signalling should be lenient in accepting signalling. Any new implementations should also be able to handle m lines signalled as RTP/AVPF.

Here is a sample SDP of advertising AVPF attributes within an m line with profile specified as RTP/AVP.

```
v=0
o=anonymous 1240218157 1240218157 IN IP4 10.193.128.35
s=-
i=myUserAgent
c=IN IP4 10.193.128.35
b=TIAS:256000
b=AS:256
t=0 0
m=audio 6000 RTP/AVP 9 8 0 18 116
a=sendrecv
a=rtpmap:9 G722/8000
a=rtpmap:8 PCMA/8000
a=rtpmap:0 PCMU/8000
a=rtpmap:18 G729/8000
a=fmtp:18 annexb=no
a=rtpmap:116 telephone-event/8000
a=fmtp:116 0-15
m=video 6002 RTP/AVP 96 97 34 31
b=TIAS:256000
a=rtpmap:96 H264/90000
a=fmtp:96 profile-level-id=428014
a=sendrecv
a=rtpmap:97 H263-1998/90000
a=fmtp:97 CIF=1;QCIF=1;I=1;J=1;T=1;N=4;K=1
a=rtpmap:34 H263/90000
a=fmtp:34 CIF=1;QCIF=1
a=rtpmap:31 H261/90000
a=fmtp:31CIF=1;QCIF=1
a=rtcp-fb:* nack pli
```

The leniency in the signalling of RTP/AVPF attributes within RTP/AVP m lines is applicable to codec control messages defined via RFC 5104 as well.

5. Flow Control

RFC 5104 codec control messages should be supported by video implementations.

The recommended mechanism to signal temporary bitrate changes is by using tmmbr (RFC 5104 codec control messages). However, tmmbr cannot be used to signal higher bitrate than negotiated for the session using TIAS or AS.

Re-INVITE should be used for permanent session bandwidth modification. Here is a sample SDP that describes how to advertise support for RTCP feedback tmmbr capability

```
m=video 6002 RTP/AVPF 96 97 34 31
```

```
b=TIAS:256000
a=rtpmap:96 H264/90000
a=fmtp:96 profile-level-id=428014
a=sendrecv
a=rtpmap:97 H263-1998/90000
a=fmtp:97 CIF=1;QCIF=1;I=1;J=1;T=1;N=4;K=1
a=rtpmap:34 H263/90000
a=fmtp:34 CIF=1;QCIF=1
a=rtpmap:31 H261/90000
a=fmtp:31 CIF=1;QCIF=1
a=rtcp-fb:* ccm tmmbr
```

6. Intra Frame Request

RFC 5104 codec control messages should be supported by video implementations.

Intra Frame Request: the recommended way of supporting intra-frame-requests is to support RTCP feedback. For backward compatibility reasons, the SIP INFO (RFC 5168) method should also be supported. In the event that neither RTCP feedback nor the SIP INFO method is supported, the implementation should have a mechanism to periodically send an intra-frame.

Here is a sample SDP that describes how to advertise support for RTCP feedback fir capability.

```
m=video 6002 RTP/AVPF 96 97 34 31
b=TIAS:256000
a=rtpmap:96 H264/90000
a=fmtp:96 profile-level-id=428014
a=sendrecv
a=rtpmap:97 H263-1998/90000
a=fmtp:97 CIF=1;QCIF=1;I=1;J=1;T=1;N=4;K=1
a=rtpmap:34 H263/90000
a=fmtp:34 CIF=1;QCIF=1
a=rtpmap:31 H261/90000
a=fmtp:31 CIF=1;QCIF=1
a=rtcp-fb:* ccm fir
```

7. H.264 Video Specifics

7.1 RFCs to be Taken Into Account

RFC 3984-bis (http://tools.ietf.org/id/draft-ietf-avt-rtp-RFC3984bis-09.txt) defines the SDP parameters that must be used to declare a video stream using this codec.

Note: section 8.2.2 of RFC 3984 indicates some rules regarding the value of the "profile-level-id", "packetization-mode' and 'sprop-deint-buf-req' (check [1-sec. 8.2.2]). According to RFC 3984, the "level" value of the profile-level-id parameter must be symmetric.

RFC 3984-bis updates these rules. In particular, it allows an SDP answer to change the level value in a corresponding SDP offer.

Additional parameters, as described per RFC 3984 (and RFC 3984-bis) make it - possible to specify HD video format and to declare H.264 level 3.0 or lower.

7.2 H.264 Capabilities Declaration

RFC 3984 defines the RTP Payload Format for H.264, an extract of which describes H.264 parameters as follows:

“However, the variability and flexibility of the H.264 codec leads to a wide array of optional parameters. Some of these parameters are implemented by many endpoints while others are rarely implemented in the mainstream. The purpose of this document is to establish a baseline, lowest-common-denominator for vendors to implement to improve interoperability.

profile-level-id:

While specified as optional (as are all parameters) in RFC 3984, the 'profile-level-id' parameter is fundamental to the setup of the codec, and is also required for any further parameters to be specified. Hence all implementations should include this parameter in their SDPs, and should interpret it when receiving it.

max-mbps, max-fs, max-cpb, max-dpb, and max-br:

These parameters allow the implementation to specify that they can support certain features of H.264 at higher rates and values than those signalled by their level (set with profile-level-id). Implementations need not include these parameters in their SDP, but should interpret them when receiving them, allowing them to send the highest quality of video possible.

max-smbps: Implementers may be interested in MaxStaticMBPS defined in RFC 3894-bis. At this stage, implementations should at the least ensure they do not behave undesirably (e.g. by crashing) when receiving this parameter (or other, unknown parameters) and may wish to honour it.

sprop-parameter-sets:

H.264 allows sequence and picture information to be sent both in-band, and out-of-band. SIP video implementations should signal this information in-band, conforming to the model prevalent in H.323 and in the overwhelming majority of existing SIP video implementations, and hence this parameter should not be included.

packetization-mode:

The codec can be broken up into smaller packets in a number of different ways. While these smaller fragments may be necessary in the future to cover cases such as high-quality video over mobile phone, current implementations shall support packetization-mode of 0 (no additional packetization).

Most of the further parameters are only needed if packetization-mode is not 0: these and other parameters are not required to be included in the SDP. The additional parameters if included should not cause the answerer to crash.

It is recommended that the implementations start supporting packetization mode 1.

RFC 3984-bis should be followed as far as all specifics are concerned.

SDP Sample to Declare H.264 / HD Video Format

Sample H.264 – HD (720p30) SDP parameters

Here follows some SDP as advertised with no interop issues.

```
m=video 60002 RTP/AVP 96
a=rtpmap:96 H264/90000
a=fmtp:96 profile-level-id=42801f
```

Alternatively one could advertise:

```
m=video 60002 RTP/AVP 96
a=rtpmap:96 H264/90000
a=fmtp:96 profile-level-id=428014; max-fs=3600; max-mbps=108000; max-br=14000
```

Both indicate a capability to receive HD resolution video.

8. References

[1] S. Wenger, M.M. Hannuksela, T. Stockhammer, M. Westerlund, D. Singer, "RTP Payload Format for H.264 Video", RFC 3984, February 2005

[2] ITU-T Rec. H.264 | ISO/IEC 14496-10 AVC

[3] RFC 3984-bis "RTP Payload Format for H.264 Video”

[4] RFC4585 – “Extended RTP Profile for Real-time Transport Control"

[5] RFC5104 “Codec Control Messages in the RTP Audio-Visual Profile with Feedback (AVPF)”

[6] RFC 5168 “XML Schema for Media Control”

[7] RFC3890 “A Transport Independent Bandwidth Modifier for the Session Description Protocol (SDP)”

[8] RFC4566 “SDP: Session Description Protocol”

[9] H.241 “Extended video procedures and control signals for H.300-series terminals”