



Document Number:	IMTC1015	Date:	11 March 2016
Working Group:	SIP Parity Activity Group	Status (draft, approved, obsolete):	Approved
Title:	SIP Video Profile Best Practices		
Purpose:	Implementation Guideline		

Summary

The document defines the SIP Video Profile, covering bandwidth, flow control, and intra-frame request use cases and proposed best practices.

Document history

Revision	Date	Description
1	7 December 2015	Initial version with only boilerplate changes from IMTC1013
2	8 December 2015	Added clarifications for rtcp-fb based on findings from SuperOp 2015
3	11 March 2016	Approved as IMTC1015

Copyright statement and disclaimer

Copyright 2016 International Multimedia Telecommunications Consortium, Inc. (“IMTC”). All rights reserved, except as expressly delineated herein. May include contributions for which Copyright is owned by IMTC contributors. See IMTC Copyright Policy [<http://www.imtc.org/documents/policy-documents/>] for additional information.

Unless authorized by the IMTC in writing, this document (including translations) may only be copied, distributed, or used for the purpose of implementing the specification in the document or preparing suggested revisions for consideration by the IMTC. Except as provided in the preceding sentence, the document may not be modified. If you reproduce or translate this document, this copyright notice and the disclaimer notice provided below must be included on the first page of any copy or translation. If you translate this document, you must translate the notice and the disclaimer notice provided below in the language used in the rest of the translation.

DISCLAIMER OF WARRANTIES: The specification in this document is provided “AS IS”. IMTC and every contributor to the specification contained in this document hereby disclaim to the greatest extent possible under applicable law all warranties, express or implied, including but not limited to any warranty of merchantability, fitness for a particular purpose, or non-infringement. You are advised that Implementation of the specification may require licenses to patents or other intellectual property rights owned by IMTC, by contributors to the IMTC, or by third parties. IMTC disclaims ANY representation that any such intellectual property rights will be available to you.

Table of Contents

1. Introduction	3
2. Asymmetric Negotiation	3
2.1 Asymmetric Bandwidth to Home	3
2.2 Video Encode and Decode Computational Complexity	3
3. Bandwidth Indication	4
4. RTP/AVPF Profile	9
5. Flow Control	10
6. Intra Frame Request	10
7. H.264 Video Specifics	11
7.1 RFCs to be Taken Into Account	11
7.2 H.264 Capabilities Declaration	11
8. References	13

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. In scope are video telephony services, point-to-point as well as multi-party. They are conversational services with delay constraints and typically involving one video and one audio stream.

2. Asymmetric Negotiation

In video telephony services 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 384 kbps but transmitting only 128 kbps.

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 an 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, subject to the constraint in RFC 6184 that the level part is the only part of the profile-level-id that changes. 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 use `max-recv-level` for expressing ability to receive higher level than expressed in the profile-level field per RFC 6184. 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 that 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) bandwidth specifier indicates the maximum supported bandwidth excluding IP/UDP/RTP overhead. The AS (RFC 4566) bandwidth specifier indicates the maximum supported bandwidth including IP/UDP/RTP overhead. These specifiers should be used as shown in the following examples. Use of the TIAS bandwidth specifier at the session level and the video m line level is mandatory. It is recommended that the AS bandwidth specifier be used at the session level for backward compatibility and also at media level, both for audio and video, for compatibility with multi-media telephony services over IMS [10].

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

It is highly recommended to use the AS bandwidth specifier for audio stream at media level. The use of application layer redundancy to handle packet losses, i.e. sending some or all speech frames multiple times in different RTP packets, in scenarios where the delay constraints do not allow for re-transmissions, increases the bandwidth requirements rapidly over what could be implied by the codec used. The use of the AS specifier for audio streams is in line with IMS which requires that the bandwidth is declared in the media scope for all audio and video streams [11]. However, if no bandwidth parameter is specified for the audio stream in the received SDP, implementations may imply it from the codec used. In this case, it is recommended that implementations use the guidelines contained in 3GPP [10] to calculate the IP/UDP/RTP overhead for the audio stream.

The audio and video codecs shown in all samples in this document are simply chosen for the purpose of illustrating the use of SDP parameters. No attempt is made therein to mandate specific codecs in the profile.

Sample SDP specification including bandwidth parameters with TIAS and AS specifiers at session level and at media level for all media streams:

```
v=0
o=anonymous 1240218157 1240218157 IN IP4 10.193.128.35
s=-
i=myUserAgent
c=IN IP4 10.193.128.35
b=TIAS:320000
b=AS:350
t=0 0
m=audio 6000 RTP/AVP 9 8 0 18 116
b =TIAS :64000
b=AS:80
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
b=TIAS:256000
b=AS:270
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
```

Important:

In this example, `b=TIAS:320000/b=AS:350` at the session level means that the maximum capability on bandwidth for all streams is 320 Kbps for the RTP payload and 350 Kbps including IP/UDP/RTP overhead. From that bandwidth a maximum of 256 Kbps (270 including IP/UDP/RTP overhead) can be used for the video stream and a maximum of 64 Kbps (80 with overhead) can be used for the audio stream.

3GPP [10] contains guidelines how to calculate the IP/UDP/RTP overhead for the audio stream. For a 64 Kbps audio stream with 20 ms packetization time, the overhead is 16 Kbps. The same specification contains several examples of IP/UDP/RTP overhead for video stream. This is estimated to be around 5% of the video bitstream.

Note: max-br: this SDP parameter is used in H.264 per RFC6184 to indicate a receive capability higher than the one derived from the profile level. The negotiation of a H.264 video stream must comply with RFC6184. This value should not be higher than the one indicated in the TIAS specifier.

Sample SDP specification including bandwidth parameters with TIAS and AS specifiers at session level and at media level for video streams:

```
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:270
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
b=TIAS:256000
b=AS:270
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
```

Important:

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

Clarification: 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, even if 256 Kbps is specified for the video media level.

SDP sample containing only TIAS bandwidth specifier:

```
v=0
o=anonymous 1240218157 1240218157 IN IP4 10.193.128.35
s=-
i=myUserAgent
c=IN IP4 10.193.128.35
b=TIAS:320000
t=0 0
m=audio 6000 RTP/AVP 9 8 0 18 116
b =TIAS :64000
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
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
```

Important:

An implementation that is not able to calculate the IP/UDP/RTP overhead for a particular media may just indicate the TIAS parameter in the SDP. A receiver of such SDP may send application data, i.e. payload on top of RTP, using up to the indicated bandwidth.

Not including AS specifier however jeopardizes interoperability with 3GPP MMTel clients, many of which don't support the TIAS specifier.

SDP sample with TIAS and AS set to the same value:

```
v=0
o=anonymous 1240218157 1240218157 IN IP4 10.193.128.35
s=-
i=myUserAgent
c=IN IP4 10.193.128.35
b=TIAS:320000
b=AS:320
t=0 0
m=audio 6000 RTP/AVP 9 8 0 18 116
b=TIAS :64000
b=AS:64
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
b=TIAS:256000
b=AS:256
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
```

Important:

An implementation which is not able to calculate the IP/UDP/RTP overhead for a particular media may include the AS parameter set to the same value as the TIAS parameter in the SDP. A well-behaved

implementation receiver of such SDP should calculate the IP/UDP/RTP overhead for the used codec and subtract it from the received AS value. Thereafter it may send application data, i.e. payload on top of RTP, using up to the calculated bandwidth, which will obviously be smaller than the received TIAS value.

If the receiver however is also not able to calculate the IP/UDP/RTP overhead and sends application data using up to the bandwidth indicated in the TIAS value, there is a risk to exceed the total bandwidth allocated by intermediary nodes according to the AS value and thus of loss of data.

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/AVPF 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.

```
m=video 6002 RTP/AVP 96 97 34  
b=TIAS:256000  
b=AS:270  
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=rtcp-fb:* nack pli  
a=rtcp-fb:* ccm tmmbr  
a=rtcp-fb:* ccm fir
```

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. Note that the example above illustrates correct conformance to the requirements in RFC 5104, section 7.1, that one and only one `rtcp-fb-ccm-param` may appear within any `rtcp-fb` line.

5. Flow Control

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

The recommended mechanism to signal temporary bitrate change is using TMMBR (RFC5104 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
b=TIAS:256000
b=AS:270
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=rtcp-fb:* ccm tmmbr
```

6. Intra Frame Request

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

Full Intra Request (fir): 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. SIP INFO should be used only in cases in which the preferred RTCP feedback mechanism is not successfully negotiated. 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
b=TIAS:256000
b=AS:270
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=rtcp-fb:* nack pli
a=rtcp-fb:* ccm fir
```

Picture Loss Indication (PLI) is the recommended method for the receiver to react to picture losses. FIR should be used when the decoder cannot recover without a decoder refresh point [5].

7. H.264 Video Specifics

7.1 RFCs to be Taken Into Account

RFC 6184 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 6184 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 6184) make it possible to specify HD video format and to declare H.264 level 3.0 or lower.

7.2 H.264 Capabilities Declaration

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 lowest common denominator for vendors to implement to improve interoperability.

profile-level-id:

While specified as optional (as are all parameters) in RFC 6184, 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. If not included, the default value is 420010, as specified in RFC 6184.

level-asymmetry-allowed:

This parameter may be used to explicitly signal support, or lack of support, of level asymmetry, as described in section 2.2. Implementations are encouraged to use this parameter as described in RFC 6184, with the exception that when the parameter is not present, the value be inferred to be equal to 1. This is to aid with backward compatibility with implementations that existed prior to addition of this parameter in RFC 6184.

max-recv-level:

This parameter may be used to declare the highest level supported when receiving, as described in section 2.2. Implementations are encouraged to use this parameter in conjunction with level-asymmetry-allowed as described in RFC 6184.

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 6184. 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. However, implementations should be lenient to SDP offers that contain this parameter, i.e. should not reject the SDP offer, in order to facilitate interoperability with MMTel (IMS Multi Media Telephony) terminals. If no in-band parameters are received, a FIR should be sent.

Note: MMTel recommends signalling the information out-of-band but also specifies the support of in-band signalling for parameter modifications [10].

in-band-parameter-sets:

Implementations are encouraged to include the in-band-parameter-sets parameter, as described in RFC 6184, to indicate whether or not out-of-band parameter sets in sprop-parameter-sets and sprop-level-parameter-sets are discarded.

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.

If omitted, the default packetization-mode 0 is implied. When using packetization-mode 1, it must be included explicitly in the SDP.

RFC 6184 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 interoperability 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 samples indicate a capability to receive HD resolution video.

8. References

- [1] IETF RFC 3984, "RTP Payload Format for H.264 Video" (obsoleted by RFC 6184)
- [2] ITU-T Rec. H.264 | ISO/IEC 14496-10 AVC, "Advanced video coding for generic audiovisual services"
- [3] IETF RFC6184, "RTP Payload Format for H.264 Video"
- [4] IETF RFC4585, "Extended RTP Profile for Real-time Transport Control"
- [5] IETF RFC5104, "Codec Control Messages in the RTP Audio-Visual Profile with Feedback (AVPF)"
- [6] IETF RFC5168, "XML Schema for Media Control"
- [7] IETF RFC3890, "A Transport Independent Bandwidth Modifier for the Session Description Protocol (SDP)"
- [8] IETF RFC4566, "SDP: Session Description Protocol"
- [9] ITU-T Rec. H.241, "Extended video procedures and control signals for H.300-series terminals"
- [10] 3GPP TS 26.114 "IP Multimedia System (IMS); Multimedia Telephony; media handling and interaction"
- [11] 3GPP TS.24.229 "IP multimedia call control protocol based on SIP and SDP"