



**Implementation Agreement for SDP Usage &
Codec Negotiation for MSF Release 3**

MSF-IA-SDP.002-FINAL

MultiService Forum Implementation Agreement

Contribution Number: msf2006.046.04

Document Filename: MSF-IA-SDP.002-FINAL

Working Group: Protocol and Control

Title: Implementation Agreement for SDP Usage & Codec Negotiation for MSF Release 3

Editor: Tom Taylor, Nortel
+1 613 763 1496
taylor@nortel.com

Working Group Chairperson: Chris Gallon, Fujitsu

Date: 24 May 2007

Abstract: This contribution has a two-fold purpose. At the signalling level, it profiles the usage of SDP for media sessions in SIP, H.248, and MGCP. At the media level, it profiles the usage of RTP and RTCP and individual payload types within MSF networks.

The MultiService Forum (MSF) is responsible for developing Implementation Agreements or Architectural Frameworks which can be used by developers and network operators to ensure interoperability between components from different vendors. MSF Implementation Agreements are formally ratified via a Straw Ballot and then a Principal Member Ballot.

Draft MSF Implementation Agreements or Architectural Framework may be published before formal ratification via Straw or Principal Member Ballot. In order for this to take place, the MSF Technical Committee must formally agree that a draft Implementation Agreement or Architectural Framework should be progressed through the balloting process. A Draft MSF Implementation Agreement or Architectural Framework is given a document number in the same manner as an Implementation Agreement.

Draft Implementation Agreements may be revised before or during the full balloting process. The revised document is allocated a new major or minor number and is published. The original Draft Implementation Agreement or Architectural Framework remains published until the Technical Committee votes to withdraw it.

After being ratified by a Principal Member Ballot, the Draft Implementation Agreement or Architectural Framework becomes final. Earlier Draft Implementation Agreements or Architectural Frameworks remain published until the Technical Committee votes to withdraw them.

The use of capitalization of the key words "MUST", "SHALL", "REQUIRED", "MUST NOT", "SHOULD NOT", "SHOULD", "RECOMMENDED", "NOT RECOMMENDED", "MAY" or "OPTIONAL" is as described in section V-B of the MSF Technical Committee Operating Procedures.

The goal of the MSF is to promote multi-vendor interoperability as part of a drive to accelerate the deployment of next generation networks. To this end the MSF looks to adopt pragmatic solutions in order to maximize the chances for early deployment in real world networks.

To date the MSF has defined a number of detailed Implementation Agreements and detailed Test Plans for the signaling protocols between network components and is developing additional Implementation Agreements and Test Plans addressing some of the other technical issues such as QoS and Security to assist vendors and operators in deploying interoperable solutions.

The MSF welcomes feedback and comment and would encourage interested parties to get involved in this work program. Information about the MSF and membership options can be found on the MSF website <http://www.msforum.org/>

Note: Attention is called to the possibility that use or implementation of this MSF Implementation Agreement may require use of subject matter covered by intellectual property rights owned by parties who have not authorized such use. By publication of this Implementation Agreement, no position is taken by MSF as its Members with respect to the existence or validity of any intellectual property rights in connection therewith, nor does any warranty, express or implied, arise by reason of the publication by MSF of this Implementation Agreement. Moreover, the MSF shall not have any responsibility whatsoever for determining the existence of IPR for which a license may be required for the use or implementation of an MSF Implementation Agreement, or for conducting inquiries into the legal validity or scope of such IPR that is brought to its attention.

DISCLAIMER

The information in this publication is believed to be accurate as of its publication date. Such information is subject to change without notice and the MultiService Forum is not responsible for any errors or omissions. The MultiService Forum does not assume any responsibility to update or correct any information in this publication. Notwithstanding anything to the contrary, neither the MultiService Forum nor the publisher make any representation or warranty, expressed or implied, concerning the completeness, accuracy, or applicability of any information contained in this publication. No liability of any kind whether based on theories of tort, contract, strict liability or otherwise, shall be assumed or incurred by the MultiService Forum, its member companies, or the publisher as a result of reliance or use by any party upon any information contained in this publication. All liability for any implied or express warranty of merchantability or fitness for a particular purpose is hereby disclaimed.

The receipt or any use of this document or its contents does not in any way create by implication or otherwise:

Any express or implied license or right to or under any MultiService Forum member company's patent, copyright, trademark or trade secret rights which are or may be associated with the ideas, techniques, concepts or expressions contained herein; nor

Any warranty or representation that any MultiService Forum member companies will announce any product(s) and/or service(s) related thereto, or if such announcements are made, that such announced product(s) and/or service(s) embody any or all of the ideas, technologies, or concepts contained herein; nor

Any commitment by a MultiService Forum company to purchase or otherwise procure any product(s) and/or service(s) that embody any or all of the ideas, technologies, or concepts contained herein; nor

Any form of relationship between any MultiService Forum member companies and the recipient or user of this document.

Implementation or use of specific MultiService Forum Implementation Agreements, Architectural Frameworks or recommendations and MultiService Forum specifications will be voluntary, and no company shall agree or be obliged to implement them by virtue of participation in the MultiService Forum.

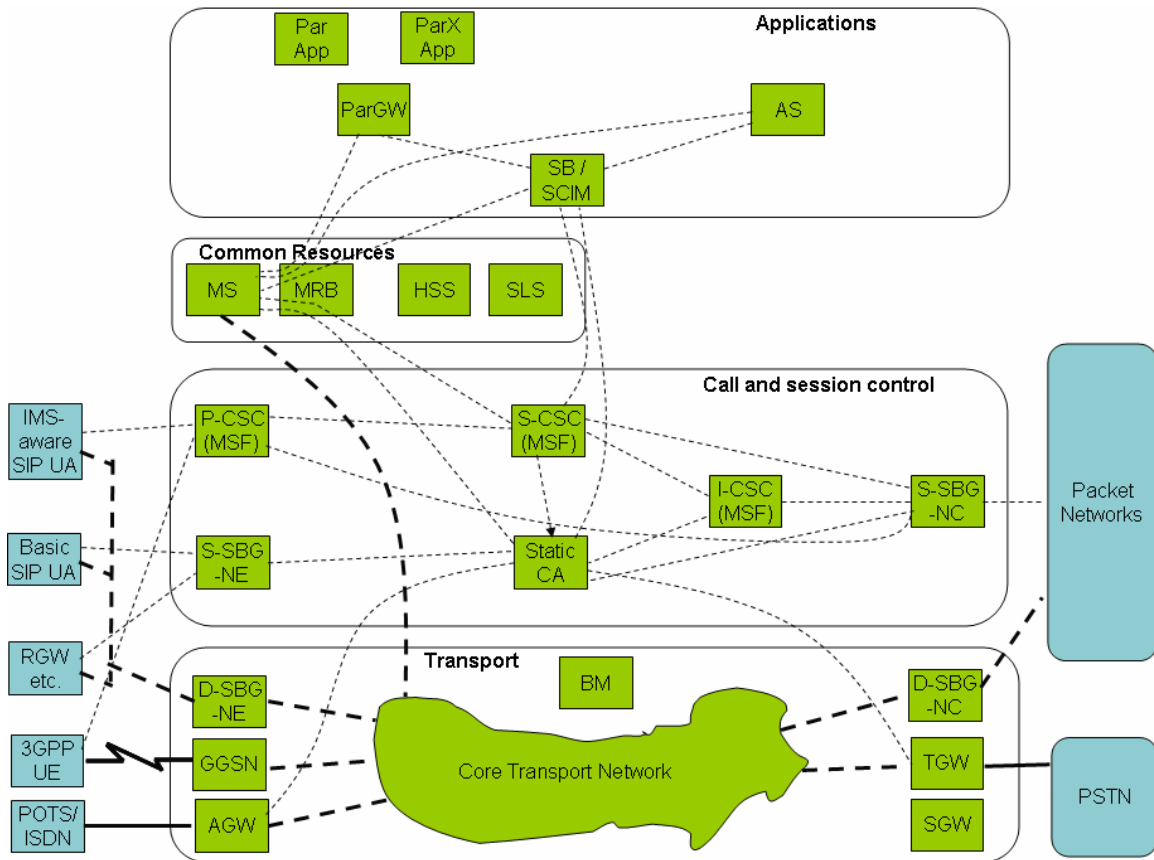
For addition information contact:

MultiService Forum
48377 Fremont Blvd., Suite 117
Fremont, CA 94538 USA
Phone: +1 510 492-4050
Fax: +1 510 492-4001
info@msforum.org
<http://www.msforum.org>

1 Introduction

As shown in Figure 1, the MSF R3 network supports a variety of components. Most of these components either process signalling relating to the setup of media sessions, process the resulting media flows, or both. An agreement on the minimum support these components will provide for features of the Session Description Protocol (SDP) at the signalling level and for RTP, RTCP, and individual media formats at the level of media flows is desirable for successful operation of the MSF R3 network.

The signalling interfaces affected by this Implementation Agreement are shown as lighter dashed lines in Figure 1. The media flows affected are shown as heavier dashed lines.



Taken from Figure 1 of MSF-ARCH-003.00-FINAL. Interfaces not carrying SDP have been deleted.

Figure 1 – MSF R3 architecture with interfaces affected by this Implementation Agreement

This Implementation Agreement is written with three objectives in mind:

- to provide a smooth evolution from the MSF R2 network;
- to match as well as possible the ETSI TISPAN specifications for the Next Generation Network (NGN);

- to provide compatibility with the specific requirements of 3GPP User Equipment (3GPP UE).

2 References

This section provides a set of references which appear to be relevant for MSF consideration.

The following references pertain to media formats and the use of RTP and RTCP.

- [1] RFC 3550, "RTP: A Transport Protocol for Real-Time Applications", July 2003.
- [2] RFC 3551, "RTP Profile for Audio and Video Conferences with Minimal Control", July 2003.
- [3] draft-ietf-avt-rfc3555bis-part2-00.txt, "Media Type Registration of Payload Formats in the RTP Profile for Audio and Video Conferences", February 2006.
- [4] RFC 2429, "RTP Payload Format for the 1998 Version of ITU-T Rec. H.263 Video (H.263+)", October 1998.
- [5] RFC 3267, "RTP payload format and file storage format for the Adoptive Multi-Rate (AMR) and Adaptive Multi-Rate Wideband (AMR-WB) audio codecs", June 2002.
- [6] RFC 3389, "RTP Payload for Comfort Noise", September 2002.
- [7] RFC 3556, "Session Description Protocol (SDP) Bandwidth Modifiers for RTP Control Protocol (RTCP) Bandwidth", July 2003.
- [8] RFC 3611, "RTP Control Protocol Extended Reports (RTCP XR)", November 2003.
- [9] RFC 3984, "RTP payload Format for H.264 Video", February 2005.
- [10] RFC 4040, "RTP payload format for a 64 kbit/s transparent call", April 2005.
- [11] RFC 4103, "RTP Payload for Text Conversation", June 2005.
- [12] RFC 4145, "TCP-Based Media Transport in the Session Description Protocol (SDP)", September 2005.
- [13] RFC 4351, "Real-Time Transport Protocol (RTP) Payload for Text Conversation Interleaved in an Audio Stream", January 2006.
- [14] RFC 4733, "RTP Payload for DTMF Digits, Telephony Tones and Telephony Signals", November 2006.
- [15] RFC 4734, "Definition of Events For Modem, FAX, and Text Telephony Signals", November 2006.
- [16] ITU-T Recommendation T.38 (09/2005), "Procedures for real-time Group 3 facsimile communication over IP networks".
- [17] ITU-T Recommendation V.152 (01/2005), "Procedures for supporting voice-band data over IP networks".
- [18] ITU-T Recommendation V.152 Corrigendum 1 (09/2005).
- [19] 3GPP TS 26.235 V7.0.0 (2005-12), "Technical Specification Group Services and System Aspects; Packet switched conversational multimedia applications; Default codecs (Release 7)".

- [20] ISO/IEC 14496-2, "Information technology Coding of audio-visual objects Part 2: Visual", 12/01/2001.

The following references pertain to the Session Description Protocol (SDP) and extensions that may be used in the MSF R3 network.

- [21] RFC4566, "SDP: Session Description Protocol", July 2006.
- [22] RFC 3605, "Real Time Control Protocol (RTCP) attribute in Session Description Protocol (SDP)", October 2003.
- [23] RFC 3312, "Integration of Resource Management and Session Initiation Protocol (SIP)", October 2002.
- [24] RFC 3524, "Mapping of Media Streams to Resource Reservation Flows", April 2003.
- [25] RFC 3388, "Grouping of Media Lines in the Session Description Protocol (SDP)", December 2002.
- [26] RFC 4032, "Update to the Session Initiation Protocol (SIP) Preconditions Framework", March 2005.
- [27] RFC 4091, "The Alternative Network Address Types (ANAT) Semantics for the Session Description Protocol (SDP) Grouping Framework", June 2005.

The following references define the signalling protocols within which SDP is used:

- [28] RFC 3261, "SIP: Session Initiation Protocol", June 2002.
- [29] RFC 3264, "An Offer/Answer Model with the Session Description Protocol (SDP)", June 2002.
- [30] RFC 4317, "Session Description Protocol (SDP) Offer/Answer Examples", December 2005.
- [31] RFC 3435, "Media Gateway Control Protocol (MGCP), Version 1.0", January 2003.
- [32] ITU-T Recommendation H.248.1 (05/02), "Gateway control protocol: Version 2", May 2002.

The following references specify the 3GPP/ETSI TISPAN profile for SDP and codec usage in the IMS portion of the network. As such, they are the starting point for this Implementation Agreement.

- [33] Draft ETSI ES 283 003 V<0.7.13> (2006-2), "Telecommunications and Internet Converged Services and Protocols for Advanced Networking (TISPAN), IP Multimedia Call Control Protocol based on Session Initiation Protocol (SIP) and Session Description Protocol (SDP) Stage 3 Protocol specification 3GPP TS 24.229 Release 7, modified", also known as DES 03019v0713.

- [34] 3GPP TS 24.229 V7.4.0 (2006-06), "IP multimedia call control protocol based on Session Initiation Protocol (SIP) and Session Description Protocol (SDP); Stage 3 (Release 7)".
- [35] 3GPP TS 26.235 V7.0.0 (2005-12): "Packet switched conversational multimedia applications; Default codecs (Release 7)".

The following reference is the document which the present document replaces for the purposes of GMI 2006.

- [36] MSF-IA-SDP.001-FINAL, " Implementation Agreement for SDP Usage & Codec Negotiation for GMI 2004", MultiService Forum, February 2004.

The following reference defines the components to which this agreement applies.

- [37] MSF-ARCH-003.00-FINAL (msf2005.179.04), "MSF Release 3 Architecture", Multi-Service Forum, June 2006.

3 Terminology

The following definitions are used in this Implementation Agreement:

application server (*note small initial letters*)

Any one of the Application Server, Parlay/Parlay X Gateway, or Service Logic Gateway as defined in [37].

SDP sender

A component of the MSF R3 network that originates an SDP body in a SIP message. Specifically, this includes the basic SIP UA, the IMS-aware SIP UA (including the 3GPP UE), an application server, the Call Agent, and the Media Server.

Bandwidth requestor

A component of the MSF R3 network that sends requests to the Bandwidth Manager and/or to the D-SBG-NE or D-SBG-NC relating to the allocation of bandwidth to individual sessions. This includes the S-SBG-NE, the S-SBG-NC, the P-CSC (MSF), and the Call Agent.

4 SDP Profile

The primary reference for the Session Description Protocol (SDP) shall be RFC 4566 [21].

The IMS usage of SDP is profiled in 3GPP TS 24.229 [34] section 6 and Annex A.3, with additional information on procedures for modification of SDP in the cases of NAT at the customer premises or NA(P)T and NA(P)T-PT in Annexes F and G respectively. The TISPAN endorsement of TS 24.229 [33] modifies section 6 but leaves Annex A.3 untouched. The requirements for SDP usage listed in these two documents shall in principle apply to the components of the MSF R3 Architecture. Section 4.1 of this Implementation Agreement adds details on the application of [33] and [34] section 6 and Annexes F and G. Section 4.2 of this Implementation Agreement adds details on the applicability of [34] Annex A.3.

4.1 Application of SDP procedures

Section 6 of 3GPP TS 24.229 [34] and its TISPAN endorsement [33] spell out the procedures for use of SDP, for specified functional elements. These procedures apply only to the use of SDP in SIP. This section summarizes each requirement in section 6 and specifies its applicability within the MSF R3 network for GMI 2006.

Encryption of SDP body

Section 6 states that the UE **MUST NOT** encrypt the SDP body, and that the P-CSCF and S-CSCF **MAY** reject a request containing an encrypted SDP body.

This requirement is applicable without change to the MSF R3 network. In specific terms, the basic and IMS-aware SIP UAs **MUST NOT** encrypt SDP bodies. The S-SBG-NE, P-CSC (MSF), Call Agent and S-CSC (MSF) **MAY** reject requests containing encrypted SDP bodies.

Presence of SDP in SIP messages

Section 6 states that an SDP offer **MUST** be present in an INVITE generated by a UE or MGCF. In the MSF R3 network, this requirement applies to the basic and IMS-aware SIP UAs and the Call Agent.

Section 6 states that SDP may be present in SIP messages only when it is desired to modify the current session description or when required by the rules of RFC 3261 [28]. This requirement applies without modification to SDP senders in the MSF R3 network.

Required usage of "b=" at the media level

Section 6 states that for "video" and "audio" media types the UE **MUST** include the "b=" line at the media level, using the "AS" bandwidth modifier as specified in [21]. When RTP transport is specified on the "m=" line, the UE **MUST** include two additional "b=" lines at the media level, one containing the "RS" bandwidth modifier, the other containing the "RR" bandwidth modifier, as specified in RFC 3556 [7]. The P-CSCF and S-CSCF use the information on the "b=" lines when deciding whether the session description conforms to local policy. In the TISPAN network this policy verification includes resource authorization over the Gq' interface. The peer to the SDP sender uses the "RS" and "RR" values when deciding on the rate of sending of RTCP sender and receiver reports.

In the MSF R3 network, all SDP senders **MUST** provide the "b=" line with the "AS" bandwidth modifier. By definition, IMS-aware SIP UAs **MUST** provide the "b=" lines with the "RS" and "RR" values. Bandwidth requestors **MUST** take account of the values associated with the "RS" and "RR" bandwidth modifiers when requesting bandwidth, and in their absence **MUST** assume an increment of 5% above the value given for the "AS" bandwidth modifier, as recommended in RFC 3550 [1].

Grouping of media lines

The P-CSCF may use the grouping capability specified in RFC 3388 [25] to group media streams that should be carried over the same access bearer. The UE is required to inspect incoming SDP for evidence of such grouping, and allocate media flows to access bearers accordingly.

Note that the usage of bearer grouping is access-dependent, as explained in Annexes B, D, E and H of [34]. For GMI 2006, it shall apply to 3GPP UEs, other types of IMS-aware SIP UAs and to the P-CSC (MSF).

Resource reservation and preconditions

In the 3GPP model, the originating UE makes the initial request for QoS resources through the radio network and eventually receives a response. In support of this procedure, the UE may use the QoS preconditions procedure (RFC 3312 [23] and RFC 4032 [26]) if it supports it, to indicate when it has not yet been allocated the resources to handle a session description it is sending. It also must set the affected media streams to inactive using the "a:inactive" attribute until the necessary resources are available, and then make the streams active in an updated session description. The MGCF originating a call uses a slight modification of these procedures. Similarly the terminating UE and MGCF may respond to preconditions if supported, and must set streams inactive until they have the resources to send and/or receive them.

SDP senders in the MSF R3 network **MUST** support the "a:inactive" procedure for disabling of flows before resources are available to carry them. SDP senders **SHOULD** support the preconditions procedure. SDP senders **MUST** follow the section 6 procedures for use of precondition signalling if they support it. Of the set of SDP senders, the basic and IMS-aware SIP UAs **MUST** follow the procedures applicable to the UE in section 6. The remaining components **MUST** follow the procedures applicable to the MGCF.

488 responses and further offers

Section 6 states that when rejecting a session description offer because of a local policy violation, the P-CSCF or S-CSCF must return a 488 Not Acceptable Here response containing the set of media types, codecs and other parameters that the P-CSCF or S-CSCF will allow. The UE in generating a subsequent offer is required to choose a subset from the values returned in the 488 response, retaining the order of preference indicated in that response.

In the MSF R3 network, the P-CSC (MSF), S-CSC (MSF), S-SBG-NE, and the Call Agent **MUST** generate 488 responses as specified by section 6 if they reject a session description offer based on local policy. By definition, the IMS-aware SIP UA **MUST** follow the section 6 procedure regarding use of the content of 488 responses.

The basic SIP UA may or may not use the content of 488 responses to shape a subsequent offer. If it does not, it may be unable to establish a call.

Address types

Section 6 provides that a UE that does not support an address type presented on a "c=" line of an SDP offer in a SIP INVITE must reject the INVITE with a 488 response containing a 301 Warning header indicating "incompatible network address format".

For the MSF R3 network, the following requirements apply:

- By definition, an IMS-aware SIP UA **MUST** conform to the section 6 requirement.
- A basic SIP UA **SHOULD** conform to the section 6 requirement.
- The S-SBG-NE, S-SBG-NC, and P-CSC (MSF) may need to modify addresses in session descriptions when NA(P)T or NA(P)T-PT is being applied.

Sending of inapplicable SDP line types

Section 6 specifies that the MGCF must not send SDP session descriptions containing the "i=", "u=", "e=", "p=", "r=", or "z=" lines in either offers or responses, and must ignore them when received.

In the MSF R3 network, this restriction applies to the Call Agent in its role as MGC (i.e. for calls to and from gateway-served lines and trunks). In its role as call server for basic SIP UAs, the Call Agent shall behave in accordance with the Proxy role as specified in the next section.

MGCF interaction with MGW

Section 6 states that the MGCF must offer only the codecs supported by the applicable media gateway, adding the telephone-event payload type to support DTMF.

In the MSF R3 network, this requirement applies to the Call Agent in its role as MGC.

Support of TCP media streams

Section 6 requires the UE to apply the procedures of RFC 4145 [12] if TCP media streams are required and must be negotiated using SDP.

In the MSF R3 network, this requirement applies to all SDP senders.

Modification of SDP when customer premises ("hosted") NAT is detected

Annex F describes procedures applicable at the P-CSCF to detect and support sessions involving a UE behind a NAT. Annex F.3 specifically deals with the modification of SDP addresses and ports as required.

In the MSF R3 network, this requirement applies to the P-CSC (MSF) and the S-SBG-NE.

Modification of SDP in support of NA(P)T and NA(P)T-PT

Annex G describes procedures applicable at the P-CSCF when that entity controls the application of NA(P)T and NA(P)T-PT. Section 5.10.6.3 extends those procedures to the IBCF.

In the MSF R3 network, these procedures apply to the P-CSC (MSF), the S-SBG-NE and the S-SBG-NC.

4.2 Requirements set by the Annex A.3 profile

Annex A.3 of [34] defines separate SDP profiles for the UA role and for the Proxy role. The functional elements performing the UA role are the UE, the MGCF, the AS, and the IMS-ALG.

In the MSF R3 network, the list of components performing the UA role are the 3GPP UE, the basic and IMS-aware SIP UAs, the Call Agent in its role of MGC, application servers, the Media Server, and the following when acting as B2BUAs: the S-SBG-NE, the S-SBG-NC, and the P-CSC (MSF).

The profile applicable to the UA role is specified as follows:

- All devices performing the UA role must support QoS preconditions (RFC 3312 [23] and RFC 4032 [26]).

In the MSF R3 network this provision applies to all of the components listed above except the basic SIP UA and the Media Server.

- The UA must support grouping of media lines (RFC 3388 [25]) and mapping of media streams to resource reservation flows (RFC 3524 [24]).

In the MSF R3 network this applies only to IMS-aware SIP UAs and the P-CSC (MSF).

- All devices performing the UA role must support bandwidth modifiers for RTCP bandwidth (RFC 3556 [7]) if they describe audio or video streams running over RTP transport.

In the MSF R3 network this applies to all of the listed components in their UA roles except the basic SIP UA.

- The sender must include the following lines in any session description: "v=", "o=", "s=", "t=". The sender may in addition include any of the following lines: "i="; "c="; "b=", "k=", and "a=" at the session level; "m=", "i=", "c=", "b=", "k=", and "a=" at the media level. The "b=" line is required at both the session and media levels when audio or video over RTP is described. The requirement for the presence of "c=" lines at either level is as specified in [21] section 5.7.

In the MSF R3 network this applies to all listed components in their UA roles, with the additional restrictions on the Call Agent acting as MGC that were specified in the previous section.

- The receiver must be able to process the "v=", "o=", "s=", "i=", "c=", "b=", "t=", and "a=" lines at the session level. At the media level, it must be able to process the "m=", "c=", and "a=" lines. It may support "k=" lines at either level, or "i=" lines at the media level.

In the MSF R3 network this applies to listed components in their UA roles, except that the Call Agent acting as MGC MUST ignore the "i=" line as specified in the previous section.

- Support of specific attributes is unspecified except as the use of the "a=curr", "a=des", and "a=conf" attributes is implied by [23] and [26] and the "a=mid" and "a=group" attributes is implied by [25].

In the MSF R3 network, all of the listed components in their UA roles MUST understand the following attributes: "a=ptime", "a=maxptime", "a=rtpmap", "a=recvonly", "a=sendrecv", "a=sendonly", "a=inactive", and "a=fmtp". All components supporting preconditions MUST support "a=curr", "a=des", and "a=conf". The P-CSC (MSF) and the IMS-aware SIP UA MUST support "a=mid" and "a=group". Finally, the P-CSC (MSF), the S-SBG-NE, and the S-SBG-NC SHOULD support the ANAT (alternative address) semantics of the grouping capability as defined in RFC 4091 [27] when acting as B2BUAs.

In the MSF R3 network, the list of components performing a Proxy role are the P-CSC (MSF), S-SBG-NE, and S-SBG-NC when not acting as B2BUAs, the S-CSC (MSF), the Call Agent when serving basic SIP UAs, and the Media Resource Broker. The profile applicable to the Proxy role is specified as follows:

- The P-CSCF only must support grouping of media lines [25], mapping of media lines to resource reservations [24], and bandwidth modifiers for RTCP bandwidth [7].

In the MSF network, the first two requirements apply to the P-CSC (MSF). The P-CSC (MSF), the S-SBG-NE, and the S-SBG-NC MUST understand and process bandwidth modifiers for RTCP bandwidth so that they can request the required total session bandwidth from the Bandwidth Manager for each flow they are admitting.

- Entities fulfilling the Proxy role must pass on all lines received down to the media-level "m=" line. As receivers, these entities must process the "v=" and each "m=" line, but ignore all the other session-level lines. No information is provided in Annex A.3 about the sending or receiving media-level lines other than "m=".

In the MSF R3 network, components fulfilling the Proxy role MUST pass on all lines of the SDP.

- Entities fulfilling the Proxy role must pass on "a=curr", "a=des", and "a=conf" attributes received. The P-CSCF may support insertion of "a=mid" and must support insertion of "a=group". On the receiving side, these entities must process "a=curr", "a=des", and "a=conf" if they support resource preconditions (not a capability applicable to the Proxy role) or otherwise ignore them. The P-CSCF must process "a=mid" if supported, and must process "a=group".

In the MSF R3 network, the P-CSCF-specific requirements apply to the P-CSC (MSF). The remaining requirements apply to all components acting in a Proxy role.

5 SDP to SDP Interworking

The primary restrictions and discrepancies of each of SIP [28], MGCP [31], and H.248 [32] are listed below:

Protocol	Restriction
MGCP	optional use of o=, s=, and t= lines (original RFC), use of sendonly / recvonly attributes is optional - indeed not normally used.
H.248	multiple session descriptions permitted per SDP block, optional use of o=, s=, and t= lines for MGC/CA to MG direction, multiple m= lines prohibited in any one stream.
SIP	SDP conveys view of session instead of receive capabilities, offer / answer protocol required, use of sendonly / recvonly / inactive / sendrecv attributes, Answer must contain the same m= lines as the offer, in the same order, except that lines disabled by setting port to 0 may be replaced by new ones. Answer may contain additional m= lines.

5.1 Recommendations for SDP Handling at the Call Agent

The Call Agent must perform SDP-SDP interworking to reconcile the different usages of SDP within the assorted control protocols. In order to perform this role, the Call Agent must decide:

- i) which parts of SDP are required to be sent on an interface, and
- ii) which parts of SDP must be received on an interface.

For a given call, some SDP parameters may pass end-to-end whilst others may pass only between the Call Agent and the media gateways it controls. The Call Agent should respect the internet "rule" about being liberal with what is received and conservative with what is sent in order to facilitate inter-operability.

Handling of the different SDP components is discussed below:

- **v=** The SDP version line is required to separate multiple session descriptions in a block. It could also be useful in future to detect different versions of SDP (although this is not likely to happen in fact, since the updating of RFC 2327 to RFC 4566 [21] did not cause the version to be incremented). The currently expected version is v=0. The first occurrence of this line in a SDP block should be transited. An H.248 interface can support subsequent 'v=' lines and their related session descriptions.¹ In order to deal with this mismatch between H.248 and the other control protocols, it is recommended that, for this IA, both the Call Agent & MG shall use an over-specified m= line (i.e. containing multiple codecs) within a single SDP session rather than multiple SDP sessions each of which contain a separate m= line. This results in H248 becoming aligned with the other protocols and aids inter-working.
- **o=** The origin line is intended to create a unique identifier for the session, and to contain a sequence number so that updates to the session can be correctly ordered. For this IA, the Call Agent has the option of either passing through an o= line or not. If not, the Call Agent must itself generate an o= line. Note that the SIP offer/answer model requires that subsequent offer/answer exchanges within the same dialogue keep the o= line unchanged except for increments to the version number.
- **s=** This is the session name. For this IA, the Call Agent has the option of either passing through an s= line or not. In the latter case, the Call Agent must itself generate an s= line. The line 's=-' may be used as a default session name. The SIP offer/answer model requires that this line remain unchanged for subsequent offer/answer exchanges within the same dialogue.
- **i=** The session information is an optional component of SDP, and is not required by the Call Agent. As indicated in the previous section, the Call Agent **MUST NOT** include the "i=" line in any session description it sends in SIP signalling, and must ignore any "i=" line received. The Call Agent **MAY** pass media-level "i=" lines through to an IAD/RGW/ATA, but **SHOULD NOT** pass them to an AGW or TGW. (Acting as a SIP proxy to basic SIP UA end points, the Call Agent passes the received SDP without change.)

¹ H.248 uses the multiple session description mechanism to negotiate a choice between alternative sets of capabilities. The IETF is currently defining new capability negotiation mechanisms for SDP. A future version of this IA could describe the interworking between these new mechanisms and those used by H.248.

The 'i=' line is intended for human consumption. The originator may use it within a media description to name the media stream. In complex multimedia connections that involve multiple streams of similar type, this can be used to unambiguously distinguish the streams. It is possible, though unlikely, that an IAD/RGW/ATA will display received "i=" lines to an user.

- **u=**
e=
p=
z=
r= As indicated in the previous section, the Call Agent in its MGC role **MUST NOT** include these lines in SDP contained in SIP signalling that it sends, and **MUST** ignore them in SDP it receives. This Implementation Agreement specifies that the Call Agent **MUST** delete these lines in SDP passed on to media gateways.
- **c=** Connection information may be provided either at the session level or within each media description. This field should be transited by the Call Agent. Note that where NAT is involved, there will be a need to alter the transport address in the c=line. However, this function is performed by the S-SBG-NE and does not have to be considered by the Call Agent.
- **b=** Call Agent support of the "b=" line in SIP signalling is specified in the previous sections. The Call Agent **MUST** include this line in SDP sent to IADs/RGWs/ATAs, since the S-SBG-NE requires this information in order to request bandwidth from the Bandwidth Manager. However, it is not necessary for the Call Agent to send the "b=" line in SDP sent to AGWs or RGWs, since the Call Agent itself is responsible for reserving bandwidth for media flows originated by them.
- **t=** This is a mandatory parameter. Session times are not controlled via SDP. For this IA, the Call Agent has the option of either transiting a t= line or not. In the latter case, the Call Agent must itself generate a t= line and the default setting of 't=0 0' should be used. The SIP offer/answer model requires that this line remain unchanged for subsequent offer/answer exchanges within the same dialogue.
- **k=** Encryption keys are an optional component of SDP. Until the details of Call Agent security (and its dependence on SDP) have been determined, the Call Agent need not send a 'k=' line to their external interfaces, and **MAY** discard any 'k=' lines received from external units. MGs should not generate k=lines and should ignore the parameter when received from the Call Agent.
- **a=** Attribute lines may appear at the session level, or within media descriptions. The only current use for session level attributes is to control the BICC request/accept protocol. Since this is only relevant on a per-hop basis, session level attribute lines should not be transited. The Call Agent must extract and insert session level attribute lines in their external BICC protocol as necessary. For other external interfaces, the Call Agent need not send session level attribute lines and **MAY** discard any session level attribute lines received.

Media level attribute lines are optional and are used for a number of purposes (e.g. to convey codec related information, to qualify the related SDP or to exchange QOS related parameteres etc.). For this IA, it is proposed that all media attribute lines are transited by the Call Agent and sent on an external interface. It is recognised that recipient end points MAY ignore some or all of these attributes.

- **m=** A media description is introduced by an ‘m=’ line. Media descriptions should be transited unchanged by the Call Agent. Note that if NAT is involved, there will be a need to alter the port address information in this line. However, this is the task of the S-SBG-NE and does not concern the Call Agent.

H.248 only permits a single media description block in a session description. This m= line may include multiple codecs. H248 requires a separate session description if multiple m= lines are required to be supported. Other protocols such as SIP and MGCP support multiple m= lines within a single session descriptor. Thus, the Call Agent must reconcile these different mappings/views. For voice services, in order to maximise interoperability, the MGC shall apply some rules to limit the flexibility of Megaco and provide greater alignment with the other protocols. Thus, this IA recommends the following:

- The Call Agent specifies only a single session in the Local Descriptor to its MG,
- The related m=line may include a number of codecs,
- The Call Agent always sets ReserveGroup=False. Due to the fact that there is only a single session block, this is somewhat redundant.
- The Call Agent always sets ReserveValue=True for the initial GW connection. The GW returns an SDP block containing (in general) a list of codecs.
- The Call Agent has the option of either setting ReserveValue=True or ReserveValue=False for the second GW connection. In either case, the second GW is provided with a Remote Descriptor containing multiple codecs. The terminating GW selects a list of codecs which are the common set of its capabilities and the codec list in the Remote Descriptor. Based on the setting of ReserveValue, the GW returns either a single codec (a list of 1) or a list of codecs in its Local Descriptor. It is assumed that the order of the list reflects that of the Remote Descriptor in order to realise a symmetric connection. In turn, the codec list from the terminating GW is passed to the first GW as its Remote Descriptor. The result of this is that symmetric codec lists are provided in each direction. The same codec will be at the head of the list and would be used to provide a symmetric connection.

6 Codecs and RTP/RTCP usage

6.1 Baseline codecs

The set of default codecs for 3GPP terminals is defined in [35]. This set consists of the following.

If the terminal supports audio:

- AMR narrowband as a minimum requirement;
- AMR wideband if the terminal supports wideband speech at 16 KHz sampling rate;
- telephone-event is recommended to support DTMF.

If the terminal supports video:

- H.263 Profile 0 Level 45 as the minimum requirement;
- H.263 Profile 3 Level 45 preferred;
- ISO/IEC 14496-2 [20] (MPEG-4 Visual) Simple Profile at Level 0b and H.264 (AVC) Baseline Profile at Level 1b as "should support". Additional requirements are provided when using H.264 (AVC).

If the terminal supports real-time text:

- RFC 4103 [11].

Non-3GPP equipment has a requirement for the following:

- G.711 A-law and μ -law;
- G.729;
- G.723.1;
- telephony-event and DTMF events as defined in [14].

The MSF R3 requirements are as follows:

1. The P-CSC (MSF), S-CSC (MSF), application servers, the Media Server, and the Call Agent MUST support negotiation of any of the codecs just listed and the additional ones listed below if proposed by an end point.
2. The Media Server MUST support both G.711 and AMR narrowband, to provide a transcoder of last resort between the two codecs.

The S-CSC (MSF) or Call Agent MAY call upon an Application Server to insert transcoding and reoriginate the request to or from a dependent end point if a codec mismatch condition is detected (through return of a 488 Not Acceptable Here response with a 304 warning code).

6.2 Additional codecs

Components supporting V.152 voice-band data [17], [18] MUST also support the telephony-event payload [14] and the event codes for ANS, /ANS, ANSam, and /ANSam [15].

6.3 Packetization intervals

It is expected that in a real VoIP network, there will be a policy of enforcing packetization periods. This enforcement could be limited to a range (e.g. limit G.711 to 10-40 ms) or could apply an absolute value (e.g. 20 ms for G.711).

Packetization at a media gateway can either be influenced by the Call Agent (e.g. MGCP LCO) or else set by the media gateway MIB. For this IA, the following assumptions are made :-

- The Call Agent is permitted to mandate the packetization period via MGCP LCO or H248 Local Descriptor (i.e. using the media attributeptime).
- The MG shall obey any mandated packetization period from the Call Agent.
- If no packetization period is specified by the Call Agent, then the MG shall assume a default value. This default value is dependent on the chosen codec (i.e. different codecs have different defaults).
- Different packetization periods for transmit and receive directions are possible for a given connection. For this IA, it is assumed that a 20 ms packetization period shall be used where possible for the given codec. All of the previously cited audio codecs support a 20 ms packetization period with the exception of G.723.1, which uses a 30 ms packetization. This does mean that the media gateway should interpret any Call Agent supplied packetization period as a preference and be able to override that preference if it is incompatible with the selected codec.

6.4 Usage of payload types

Components MUST use "a.rtpmap" mapping attribute for every format listed on the m= line. This MUST be done even when the payload type number is the one reserved as a static payload type for that codec.