International Multimedia Telecommunications Consortium

Rich Media - Anywhere, Anytime

Document Number:	**IMTC1011**	**Date:**	June 14, 2012
Working Group:	Unified Communications	**Status** (draft, approved, obsolete)**:**	Approved
Title:	**Telepresence Interoperability Protocol (TIP) Version 8.0**		
Purpose:	Technical Specification		

Summary

This document specifies the Telepresence Interoperability Protocol (TIP) Version 8.0. The TIP protocol specification describe how to multiplex multiple screens, multiple audio streams, as well as an auxiliary-data screen into a single Real-Time Transport Protocol (RTP) flows, one each for video and audio. It enables point to point and multipoint sessions as well as a mix of multi-screen and single-screen endpoints.

Previous versions of the TIP protocol were not published as IMTC numbered documents.

Document history

Revision	Date	Description
1	June 14, 2012	Initial version – approved for release

REPRESENTATIONS, WARRANTIES, CONDITIONS OR OTHER TERMS WHICH MAY BE IMPLIED (BY COMMON LAW, STATUTE OR OTHERWISE) WITH RESPECT TO THE TIP SPECIFICATION AND/OR THE LICENSE TO THE TIP SPECIFICATION AND SUB-LICENSE TO THE ESSENTIAL CLAIMS GRANTED UNDER THIS AGREEMENT, IN WHOLE OR IN PART, ALL OF WHICH ARE EXCLUDED, INCLUDING, WITHOUT LIMITATION, ANY IMPLIED TERM OF FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, TITLE OR NON-INFRINGEMENT. NONE OF THE IMTC, CISCO, OR ANY AFFILIATE OF EITHER OF THEM, SHALL HAVE ANY DUTY TO INDEMNIFY, HOLD HARMLESS OR DEFEND YOU, ANY OF YOUR AFFILIATES, OR ANY LICENSEE OR OTHER THIRD PARTY FROM AND AGAINST ANY LOSS, DAMAGE, LAWSUITS, PROCEEDINGS, CLAIMS OR SIMILAR ACTIONS THAT ARISE OR RESULT FROM ANY USE OR IMPLEMENTATION OF THE TIP SPECIFICATION, IN WHOLE OR IN PART, INCLUDING, BUT NOT LIMITED TO, ANY USE, MANUFACTURE, OFFER TO SELL, SALE, IMPORTATION, OR DISTRIBUTION OF ANY COMPLIANT IMPLEMENTATION.

Telepresence Interoperability Protocol ("TIP")
Implementation License

Published July 2010

This is a legal agreement ("Agreement") between the International Multimedia Telecommunications Consortium, Inc. ("IMTC") and the individual or entity that accesses, downloads, uses, or reviews the Telepresence Interoperability Protocol Specification ("You"). IMTC and You will be referred to collectively in this Agreement as "the Parties". Cisco Systems Inc. ("Cisco") is an intended third-party beneficiary of this Agreement.

If You want a license from IMTC to implement the TIP Specification, You must agree to these terms. This is an offer to be accepted only on the terms set forth in this Agreement. You accept these terms by accessing, downloading, using, or reviewing the TIP Specification. Once accepted, IMTC grants You the license and sub-license below, provided that You fully comply with these terms. If You do not agree to any of these terms, then You are not licensed under this Agreement. If You attempt to alter or modify any of these terms, then the offer is no longer valid and is revoked.

TIP is a protocol for interoperability between videoconferencing products, including streaming of audio, video, and data to and from videoconferencing products. IMTC offers this license to encourage interoperability between competitive videoconferencing products that exist as of the date of this license or that may come to exist in the future.

Neither the TIP Specification nor any portion of the TIP Specification may be printed out, reproduced, distributed, resold, or transmitted for any purpose other than the use of TIP to create Compliant Implementations (as defined below), including, without limitation, commercial purposes, such as selling copies of this documentation or providing paid-for support services.

Definitions

An **"Affiliate"** of any entity shall mean any entity that owns or controls (directly or indirectly), is owned or controlled by (directly or indirectly), or is under direct or indirect common ownership or control with, such entity. For purposes of this definition of "Affiliate," the phrase "owns or controls" shall mean direct or indirect control of more than 50 percent of the voting power to elect directors of a corporation in the case of a corporate entity, or for any other entity, the power to direct the management of such entity.

An **"Assertion of an Essential Patent Claim"** means any assertion, threat to assert, or participation in the assertion of a lawsuit, proceeding, claim, or similar action, by You or one of your Affiliates, or by a Transferee (as defined below) to which you have directly or indirectly transferred an Essential Patent Claim, against any person or entity, including Cisco, the IMTC, or any Affiliate of either of them, claiming that the making, having made, using, selling, offering for sale, importation, or other disposition of a Compliant Implementation infringes an Essential Patent Claim You or Your Affiliate owns, now or in the future. For the avoidance of doubt, the term "Assertion of an Essential Patent Claim" does not include the assertion, threat to assert, or

Bishop Ranch 6 • 2400 Camino Ramon, Suite 375• San Ramon • CA USA• 94583 • phone +1.925.275.6600 • fax +1.925.275.6691

participation in the assertion of: (a) any claim of any patent that is not an Essential Patent Claim; or (b) any Essential Patent Claim that is essential to the creation of a Compliant Implementation, to the extent that such Essential Patent Claim is infringed by any making, having made, using, selling, offering for sale, importation, or other disposition other than the making, having made, using, selling, offering for sale, importation, or other disposition of a Compliant Implementation.

"Compliant Implementation" means only the portions of Your product or Your Affiliate's product that are fully compliant with the TIP Specification.

"Essential Patent Claim" means a claim of a patent, whether or not issued as of the date of this Agreement, that is necessarily infringed by implementing the required portions of the TIP Specification. A claim is necessarily infringed only when there is no technically or commercially reasonable way to avoid infringement of that claim when following the requirements of the TIP Specification to make a Compliant Implementation. Notwithstanding the foregoing, the term Essential Claim does not include any claim: (a) that, if licensed, would require the payment of royalties by the licensor to a person or entity that is not affiliated with the licensor; or (b) that is essential only to the implementation of a standard that is referenced in the TIP Specification.

"TIP Specification" means the TIP specification found at http://www.imtc.org /[insert]. Only the TIP Specification published by IMTC (and versions previously published by Cisco or its Affiliates at www.cisco.com/go/tip) will be considered the TIP Specification for purposes of this Agreement. Any specification not published by IMTC or Cisco, or one that incorporates the TIP Specification in part, in whole, or by reference shall not be considered the TIP Specification for purposes of this Agreement.

A "**Transferee"** means any direct or indirect transferee, assignee, or purchaser of an Essential Patent Claim (or a patent containing an Essential Patent Claim, or the right to assert an Essential Patent Claim or a patent containing an Essential Patent Claim) You or Your Affiliate owns, now or in the future.

Copyright License
Upon Your acceptance of these terms, IMTC grants You and each of Your Affiliates a non-exclusive, royalty-free, non-transferable, non-sublicensable, personal, and worldwide license to use and reproduce the attached TIP Specification, solely in connection with Your creation of Compliant Implementations. The right to reproduce is limited to the creation of such limited number of copies of the Specification as are reasonably required for your creation of Compliant Implementations. No rights are granted to distribute, display, perform, modify, or create derivative works of the Specification, provided that You may propose modifications to the Specification solely in connection with Your participation in the IMTC Telepresence Interoperability Protocol Activity Group.

Patent Sub-license
Upon Your acceptance of these terms, IMTC grants You and each of Your Affiliates a non-exclusive, royalty-free, non-transferable, non-sub-licensable, personal, worldwide sub-license under all Essential Patent Claims owned, by Cisco or its Affiliates, now or in the future, solely to make, have made, use, sell, offer to sell, import, and distribute Compliant Implementations. For the avoidance of doubt, nothing in the previous sentence grants You a license or sub-license under: (a) any claim of any patent licensable by Cisco that is not an Essential Patent Claim; or (b) any Essential Patent Claim that is infringed by any making, having made, using selling, offering for sale, importation, or other disposition other than the making, having made, using, selling, offering for sale, importation, or other disposition of a Compliant Implementation.

Bishop Ranch 6 • 2400 Camino Ramon, Suite 375• San Ramon • CA USA• 94583 • phone +1.925.275.6600 • fax +1.925.275.6691

Prohibited Uses
The licenses and sub-licenses granted to You and Your Affiliates by IMTC in the TIP Specification, including those granted in the Patent Sub-license, are conditioned upon Your agreement to use the TIP Specification only for the purpose of interoperating between videoconferencing products, and for no other purpose. No right or license to any intellectual property owned by IMTC or Cisco is granted under this agreement for any other use.

Defensive Suspension
Effective upon the Assertion of an Essential Patent Claim by You or Your Affiliate against Cisco or its Affiliate, the IMTC or its Affiliate, or any other licensee or sub-licensee (or Affiliate of such licensee or sub-licensee) of an Essential Patent Claim under a license or sub-license provided to such licensee or sub-licensee of Essential Patent Claims, all sub-licenses granted to You or Your Affiliate under this Agreement to any Essential Patent Claims shall terminate immediately, and retroactively.

Surrender of IMTC License Prior to Assertion.
In the event that You or Your Affiliate shall determine that it wishes to engage in an Assertion of an Essential Patent Claim, You or Your Affiliate will first provide written notice to the IMTC stating that You and Your Affiliate are surrendering the sub-license to Essential Patent Claims granted by this Agreement. You acknowledge that the IMTC will promptly forward the notice You or Your Affiliate provides to all IMTC members as of the date on which the notice is received. Neither You nor Your Affiliate may engage in any Assertion of an Essential Patent Claim until one year following the date of its notice to the IMTC.

Effect of Surrender.
For the avoidance of doubt, the license to Essential Patent Claims You and Your Affiliates are granted under this Agreement shall continue in force during the one-year period following the date of the notice You or Your Affiliate provides. However, if You or Your Affiliate makes an Assertion of an Essential Patent Claim after the end of the one-year waiting period, any licenses or sub-licenses to Essential Patent Claims that You or Your Affiliate enjoyed under this Agreement are retroactively terminated, with the same effect as if such licenses had never existed. Further, You and Your Affiliate acknowledge and agree that when You Assert Essential Patent Claims, You may only seek reasonable and non-discriminatory (RAND) damages and/or RAND license fees, but may not seek an injunction or order of exclusion. Likewise, the entityagainst which You or Your Affiliate Asserts Essential Patent Claims may Assert any Essential Patent Claims it owns against You or Your Affiliate, but may only seek RAND damages and/or RAND license fees.

Application of Defensive Suspension and Surrender Provisions to Transferees.
Unless You (or Your Affiliate) and each direct and indirect transferor of an Essential Patent Claim formerly owned by You or Your Affiliate would be able to engage in the Assertion of an Essential Patent Claim consistent with the Defensive Suspension and Surrender provisions in the three previous paragraphs of this Agreement, no Transferee of an Essential Patent Claim formerly owned by You or Your Affiliate will have the right to engage in such an Assertion. If the transfer of an Essential Patent Claim occurs before the end of the one-year notice period referenced in the paragraph titled “Effect of Surrender”, no Transferee of an Essential Patent Claim formerly owned by You or Your Affiliate shall engage in the Assertion of an Essential Patent Claim. In other words, any such Transferee receives the patent as if it is perpetually licensed to IMTC and IMTC’s past, present, and future sub-licensees of Essential Patent Claims. If the transfer occurs after the end of the one-year waiting period there are no such limitations on the Transferee or the transferred patent, except that the limitation to seek only RAND license fees, including the prohibition on seeking injunctive relief or an order of exclusion, shall run with each transferred Essential Patent Claim to bind any Transferee of such Essential Patent Claim. You and each of Your Affiliates agree to use all reasonable efforts to provide

Bishop Ranch 6 • 2400 Camino Ramon, Suite 375• San Ramon • CA USA• 94583 • phone +1.925.275.6600 • fax +1.925.275.6691

written notice to any Transferee of an Essential Patent Claim formerly owned by You or Your Affiliate of the limitation on the ability of such Transferee to seek any remedy other than RAND license fees when such Transferee engages in the Assertion of an Essential Patent Claim.

Reservation of Rights

Copyright to the attached TIP Specification is the property of IMTC. All rights not expressly granted herein are reserved.

Updates

IMTC reserves the right to issue updates to any of the TIP Documents. Updates will be provided on IMTC's website, at http://www.imtc.org/tip

Feedback and Improvements

IMTC welcomes feedback to and suggested improvements to, the TIP Documents. Feedback and suggested improvements may be sent by e-mail to: tip_info@imtc.org

Please note that by providing such feedback or suggested improvements to IMTC, You acknowledge and agree that IMTC, any IMTC Affiliate, any IMTC member entity (and any Affiliate of any of them) and any other implementer of TIP may use such feedback and/or suggested improvements in connection with the TIP Specification without obligation of any kind to You, including by incorporating Your feedback or suggested improvements into future versions of the TIP Specification. Your submission of feedback constitutes Your agreement that IMTC may incorporate any feedback You provide into the TIP Specification (including any updates), and that You are licensing IMTC and its Affiliates, Cisco and its Affiliates, and all other implementers of the TIP Documents on the same terms on which IMTC is making the TIP Specification available to You and other actual and prospective implementers.

Third-Party Information

The TIP Specification may contain links to third-party websites and standards and specifications (for example, H.264 and the IETF RTP specification) that are not under the control of IMTC, or of Cisco, or of any Affiliate of either of them. Neither IMTC nor Cisco nor any Affiliate of either of them is responsible for the content on any linked site or referenced standard or specification. If You access a third-party website mentioned in the TIP Specification, then You do so at Your own risk. IMTC and Cisco provide these links and references only as a convenience, and the inclusion of the link or reference does not imply that IMTC, Cisco, or any Affiliate of either endorses or accepts any responsibility for the content on those third-party sites or of the referenced standards and specifications. Nothing in this Agreement grants any right, license or interest in or to any third party technology referenced in the TIP Specification. You acknowledge that the owners of intellectual property rights asserted to be necessary to implement any referenced standard may or may not make those intellectual property rights available to implementers.

No Warranty

THE TIP SPECIFICATION AND ALL LICENSES AND SUB-LICENSES GRANTED UNDER THIS AGREEMENT ARE PROVIDED "AS IS". NEITHER IMTC NOR CISCO, NOR ANY AFFILIATE OF EITHER OF THEM, MAKES ANY REPRESENTATION, WARRANTY, CONDITION OR OTHER TERM, EXPRESS OR IMPLIED (BY COMMON LAW, STATUTE OR OTHERWISE), AND EACH DISCLAIMS ANY AND ALL REPRESENTATIONS, WARRANTIES, CONDITIONS OR OTHER TERMS WHICH MAY BE IMPLIED (BY COMMON LAW, STATUTE OR OTHERWISE) WITH RESPECT TO THE TIP SPECIFICATION AND/OR THE LICENSE TO THE TIP SPECIFICATION AND SUB-LICENSE TO THE ESSENTIAL CLAIMS GRANTED UNDER THIS AGREEMENT, IN WHOLE OR IN PART, ALL OF WHICH ARE EXCLUDED,

Bishop Ranch 6 • 2400 Camino Ramon, Suite 375• San Ramon • CA USA• 94583 • phone +1.925.275.6600 • fax +1.925.275.6691

INCLUDING, WITHOUT LIMITATION, ANY IMPLIED TERM OF FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, TITLE OR NON-INFRINGEMENT. NONE OF THE IMTC, CISCO, OR ANY AFFILIATE OF EITHER OF THEM, SHALL HAVE ANY DUTY TO INDEMNIFY, HOLD HARMLESS OR DEFEND YOU, ANY OF YOUR AFFILIATES, OR ANY LICENSEE OR OTHER THIRD PARTY FROM AND AGAINST ANY LOSS, DAMAGE, LAWSUITS, PROCEEDINGS, CLAIMS OR SIMILAR ACTIONS THAT ARISE OR RESULT FROM ANY USE OR IMPLEMENTATION OF THE TIP SPECIFICATION, IN WHOLE OR IN PART, INCLUDING, BUT NOT LIMITED TO, ANY USE, MANUFACTURE, OFFER TO SELL, SALE, IMPORTATION, OR DISTRIBUTION OF ANY COMPLIANT IMPLEMENTATION.

Limitation Of Liability
NEITHER IMTC, NOR CISCO, NOR ANY AFFILIATE OF EITHER OF THEM, SHALL BE LIABLE UNDER OR IN CONNECTION WITH THIS AGREEMENT (WHETHER FOR BREACH OF WARRANTY, IN CONTRACT, TORT, NEGLIGENCE, STRICT LIABILITY, OR OTHERWISE):

(A) FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, PUNITIVE OR INCIDENTAL DAMAGES; OR

(B) FOR ANY DIRECT LOSSES, INCLUDING WITHOUT LIMITATION COSTS OF PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, LOSS OF PROFITS, LOSS OF REVENUE, LOSS OF USE, LOSS OF OPPORTUNITY, INTERRUPTION OF BUSINESS; OR

(C) FOR ANY LOSSES, DAMAGES, COSTS, CLAIMS OR OTHER LIABILITIES ARISING FROM OR RELATED TO THIS AGREEMENT WHATSOEVER; OR

(D) FOR ANY CLAIM ARISING FROM OR RELATED TO IMTC'S DISTRIBUTION OR PUBLICATION OF NOTICES YOU PROVIDE TO IMTC UNDER THIS AGREEMENT;

IN EACH CASE HOWEVER CAUSED OR ARISING, PROVIDED THAT NOTHING IN THIS AGREEMENT SHALL EXCLUDE OR LIMIT THE LIABILITY OF THE IMTC, CISCO, OR ANY AFFILIATE OF EITHER OF THEM:

(1) FOR DEATH OR PERSONAL INJURY RESULTING FROM ITS NEGLIGENCE; OR

(2) FOR FRAUD; OR

(3) TO THE EXTENT SUCH LIABILITY CANNOT BE EXCLUDED OR LIMITED BY APPLICABLE LAW.

YOU ACKNOWLEDGE THAT IMTC WOULD NOT HAVE ENTERED INTO THIS AGREEMENT ON THE TERMS HEREOF WITHOUT THE LIMITATIONS OF LIABILITY SET FORTH HEREIN, AND THAT CISCO WOULD NOT HAVE PERMITTED IMTC TO GRANT THE SUB-LICENSES IN THIS AGREEMENT WITHOUT THE LIMITATIONS OF LIABILITY SET FORTH HEREIN.

Governing Law
This Agreement shall be construed in accordance with, and all disputes hereunder shall be governed by, the laws of England and Wales. The Courts of England and Wales shall have exclusive jurisdiction over any disputes relating to the Agreement.

No Requirement to Implement

Bishop Ranch 6 • 2400 Camino Ramon, Suite 375• San Ramon • CA USA• 94583 • phone +1.925.275.6600 • fax +1.925.275.6691

Nothing in this Agreement will be construed as requiring You to use or implement any of the TIP Specifications, or limit You from developing any products that compete with products developed by the IMTC or Cisco provided that You do not violate any intellectual property right owned by the IMTC or Cisco in the course of such development.

Bishop Ranch 6 • 2400 Camino Ramon, Suite 375• San Ramon • CA USA• 94583 • phone +1.925.275.6600 • fax +1.925.275.6691

Telepresence Interoperability Protocol (TIP)

Version 8.0

This document can only be distributed under the TIP License Agreement.

Table of Contents

1 Introduction 3
2 Background Information 3
3 Media Channel Establishment 8
3.1 NAT/Firewall Traversal 8
3.2 TIP Negotiation 8
3.3 Secure Channel 9
4 TIP Multiplex 11
4.1 Positional Multiplex 11
4.2 TIP RTCP APP Extensions 16
4.2.1 Multiplex Control 17
4.2.2 Network Path Measurement 20
4.2.3 Media Flow Control 22
4.2.4 Video Refresh Request 24
4.2.5 Media Options 26
4.2.6 SPI Map 38
4.2.7 Request to Send 40
4.2.8 Notify 43
4.2.9 Acknowledgement 44
4.2.10 RTP Feedback 46
4.3 Positional Rules 49
4.3.1 Control Position 49
4.3.2 Center, Left, Right Positions 50
4.3.3 Auxiliary (AUX) Position 50
4.3.4 Legacy Positions 51
4.4 RTCP Rules 52
5 Summary of Changes from TIP 7 to TIP v8 52
6 Backward Compatibility 52
6.1 TIP Version Compatibility 53
6.2 TIP Restricted Media Compatibility 54
6.3 TIP Unrestricted Media Constraints Compatibility 55
7 Call Flow Examples 56
7.1 Basic TIP MUX Setup 56
7.2 Secure TIP MUX Setup with DTLS 57
7.3 Secure TIP MUX Setup with DTLS and EKT 57
8 References 58

1 Introduction

Generally, TIP systems are high-end, high-definition video conferencing devices capable of handling multiple audio and video streams. TIP devices negotiate the number of media streams they will exchange in the TIP messaging, but to external infrastructure entities these appear as traditional Voice over IP (VoIP) devices with audio and video capabilities.

TIP devices can be endpoints, including single and multi-screen systems participating in point-to-point and multipoint sessions. In the case of a multipoint session, endpoints will exchange TIP messaging with a multipoint control unit (MCU) that implements TIP. For purposes of this document, an MCU refers to a multipoint device that may or may not terminate, decode or transcode any of the video media before forwarding on to the rest of the endpoints in a multipoint session. The TIP protocol allows an MCU to indicate if it will transcode the media or remain focused on only providing multipoint services.

The TIP protocol provides a variety of ways to indicate or negotiate endpoint or MCU number of streams, the configuration as well as preferences or restrictions. A set of accompanying documents, called “TIP Implementation Profile” guides or similar, define exactly what protocol and media options in TIP are used to achieve tested interoperability with a particular TIP implementer’s specific product(s). These accompanying documents may be updated by the TIP implementer from time to time as product capabilities evolve, potentially asynchronous from revisions to the TIP protocol.

This document specifies TIP version 8 (TIP v8), which is the third version of the TIP protocol to be published. TIP v8 extends TIP v7 [10] by providing mechanisms to enable per stream video profile indicationlegacy stream maximum bitrate indication, a new constrained video media indication and the ability to devices to indicate a preference to use BFCP for shared content stream(s). Please reference section 5 for a more detailed summary. A backwards-compatibility strategy is also defined in section 6 to allow convergence between TIP v6, TIP v7 and TIP v8 devices trying to interoperate in the same session.

2 Background Information

The TIP protocol is designed around the Internet Engineering Task Force (IETF) standards for VoIP and Video Conferencing. A TIP device would use a standard, such as SIP [1] for call signaling and management, and the Real-Time Transport Protocol (RTP) [4] or SRTP [11] for media transmission. In the call scenarios discussed in this document, the media transmission uses only IP unicast. No IP multicast is required.

TIP devices will offer both audio and video streams during the call setup phase. TIP devices do not express their multiple stream capability during call setup.

The following supported media formats are carried in their respective IETF RTP payloads:

- Audio
  - AAC-LD
    - Bitrate: 64 kbps/channel
    - RTP Payload: IETF RFC 3640, AAC-hbr mode
    - Default Dynamic Payload Number: 96
  - G.711 (u-law)
    - RTP Payload: IETF RFC 3351
    - Static Payload Number: 0
  - G.722
    - RTP Payload: IETF RFC 3351
    - Static Payload Number: 9
  - DTMF
    - RTP Payload: IETF RFC 2833
    - Default Dynamic Payload Number: 101
- Video
  - H.264 Baseline Profile
    - Image sizes: 1080p, 720p, 1024x768, 352x288
    - Bitrates: 4 Mbps to 300 kbps
    - RTP Payload: IETF RFC 3984, packetization mode 1 and mode 0
    - Default Dynamic Payload Number: 112

Implementing this protocol requires a detailed knowledge of the RTP protocol [4] and [5], SRTP protocol [11] and Encrypted Key Transport (EKT) [12]. As a reference, the following sections provide you with the packet diagrams for RTP, Real-Time Control Protocol (RTCP), SRTP, Secure RTCP (SRTCP) and EKT packet diagrams are repeated below from [4], [11] and [12]:

RTP Header

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|V=2|P|X|  CC   |M|     PT      |       sequence number         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           timestamp                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|           synchronization source (SSRC) identifier            |
+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
|            contributing source (CSRC) identifiers             |
|                             ....                              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

RTCP Header

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|V=2|P| subtype |       PT      |             length            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              SSRC                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

SRTP Packet with optional EKT

```
       0                   1                   2                   3
     0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+<+
    |V=2|P|X|  CC   |M|     PT      |       sequence number         | |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ |
    |                           timestamp                           | |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ |
    |           synchronization source (SSRC) identifier            | |
    +=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+ |
    |            contributing source (CSRC) identifiers             | |
    |                               ....                            | |
  +>+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ |
  | |                           payload  ...                        | |
  | |                               +-------------------------------+ |
  | |                               | RTP padding   | RTP pad count | |
  +>+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ |
  | ~                      SRTP MKI (OPTIONAL)                      ~ |
  | +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ |
  | ~                       Authentication Tag                      ~ |
  | +=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+ |
  | ~                      Encrypted Master Key                     ~ |
  | +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ |
  | |                        Rollover Counter                       | |
  | +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ |
  | |    Initial Sequence Number    |   Security Parameter Index    | |
  | +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+<+
  |                                                                   |
  +- Encrypted Portion*                        Authenticated Portion ---+
```

SRTCP Packet with optional EKT

```
     0                   1                   2                   3
     0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+<+
    |V=2|P|    RC   |   PT=SR or RR   |             length          | |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ |
    |                         SSRC of sender                        | |
  +>+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+ |
  | ~                          sender info                          ~ |
  | +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ |
  | ~                         report block 1                        ~ |
  | +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ |
  | ~                         report block 2                        ~ |
  | +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ |
  | ~                              ...                              ~ |
  | +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ |
  | |V=2|P|    SC   |  PT=SDES=202  |             length            | |
  | +=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+ |
  | |                          SSRC/CSRC_1                          | |
  | +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ |
  | ~                           SDES items                          ~ |
  | +=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+ |
  | ~                              ...                              ~ |
  +>+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+ |
  | |E|                         SRTCP index                         | |
  | +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ |
  | ~                     SRTCP MKI (OPTIONAL)                      ~ |
  | +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ |
  | ~                       Authentication Tag                      ~ |
  | +=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+ |
  | ~                      Encrypted Master Key                     ~ |
  | +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ |
  | |                        Rollover Counter                       | |
  | +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ |
  | |    Initial Sequence Number    |   Security Parameter Index    | |
  | +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+<+
  |                                                                   |
  +-- Encrypted Portion                    Authenticated Portion -----+
```

The following types of RTCP packets may be used by TIP devices:

Description	Abbreviation	Type
Sender Report	SR	200
Receiver Report	RR	201
Session Description	SDES	202
Goodbye	BYE	203
Application	APP	204
Feedback	FB	205
Extended Report	XR	207

Table 1: RTCP Packet Types used by TIP Devices

3 Media Channel Establishment

For purposes of this document, we assume that call signaling has been established and the potential network addresses to be used for the media channel between two entities in the media path have been determined. This section details how the media channel and the rules for its operation are established.

3.1 NAT/Firewall Traversal

The first challenge to the establishment of the media channel is the traversal of any NAT/Firewalls that may exist between the two media processing entities (endpoints or MCUs).

To help with the NAT/Firewall traversal TIP requires the use of symmetric RTP/RTCP ports [6]. A TIP device MUST transmit all packets in a symmetric fashion, i.e., that the IP address and UDP port of a transmitted packet exactly match the IP address and UDP port for reception of the corresponding media stream. However, when receiving packets from the remote peer, implementations SHOULD accept packets with a UDP source port that differs from the destination UDP port that is being used when transmitting.

3.2 TIP Negotiation

This section describes how an implementation detects that its remote peer is capable of understanding TIP. If the remote peer does not implement TIP, an implementation SHOULD operate in a strictly standards compliant mode where only a single audio or video stream is sent in the media channel, and no TIP specific control extensions are used.

TIP uses APP packet, which is an extension of RTCP and allows control information to be transmitted. TIP uses the RTCP private application name “xcts”, eXtended Control for Telepresence Systems, for all of its control packets.

TIP defines an RTCP APP MUXCTRL packet which informs the remote peer about media multiplexing capabilities. This includes the number of media streams that a device is willing to transmit and willing to receive for a given media channel of a particular type (audio or video). It also allows the receiver to specify which media positions it is willing to transmit and receive. The reception of an RTCP APP MUXCTRL packet or an RTCP APP MUXCTRL ACK packet specifying the TIP private application name, “xcts”, is used as acknowledgement that TIP is understood by the remote peer.

TIP requires that the local system’s MUXCTRL be acknowledged by the remote peer and that the MUXCTRL of the remote peer be received before communication under the rules of TIP can begin. The RTCP APP packets and processing rules are described in section 4.2.

The TIP negotiation MUST begin by sending the RTCP APP MUXCTRL on the RTCP UDP channel(s) indicated by the media channel establishment.

For a short interval (recommendation is 15 seconds), the MUXCTRL packet should be periodically retransmitted, until an acknowledgement is received or the interval elapses without a response.

The reception of either a MUXCTRL or MUXCTRL ACK packet on the RTCP UDP channel SHOULD be interpreted as support of TIP.

An implementation SHOULD acknowledge all MUXCTRL packets it receives. The lack of any response MUST be interpreted as a lack of support for TIP, and hence the standard RTCP traffic should be sent on the RTCP UDP channel as per the standard mode of operation for [4], unless the peer has signaled otherwise in SDP [9].

If media security has been negotiated, the RTCP packets described in this section would instead be SRTCP packets.

For a message flow of a basic TIP control channel setup, refer to section 7.1.

3.3 Secure Channel

Optionally, the media channel may be secured (encrypted and authenticated) via SRTP/SRTCP. The channel is considered to be secure when the two communicating devices have negotiated secure media and the security association parameters. How the security associations are accomplished is outside the scope of this document; however the following are generally relevant considerations to a TIP device implementation:

First, the security establishment process occurs before any media related activity occurs on the UDP channel.

Second, with a secure media channel, Datagram Transport Layer Security (DTLS) [13, 14] packets may also occur within the media channel. These packets can be de-multiplexed from RTP traffic by examination of the first byte of the packet. The details are described in section 5.1.2 of [14].

Third, once security association is established, secure TIP negotiation can start.

For a message flow of a secure TIP control channel setup using DTLS, refer to section 7.2.

The following are specific TIP protocol items used to facilitate multipoint encryption:

TIP devices MAY use Encrypted Key Transport (EKT) extension to SRTP [12] to coordinate SRTP contexts between transmitters and receivers in a multipoint session where the MCU does not decrypt and re-encrypt packets.

Note that a TIP endpoint MUST NOT encrypt the Audio Activity Metric (see section 4.2.5.2) or the Video Refresh Flag (see section 4.2.5.4), so that entities such as MCU do not need to perform crypto operations to read the relevant information.

A TIP MCU device that sets the TIP multipoint focus parameter (see section 4.2.1) has negotiated EKT capability if it is communicating with an endpoint that has indicated in TIP its ability to receive EKT security parameters (see section 4.2.5).

A TIP endpoint has negotiated EKT capability if it is communicating with a TIP MCU that has indicated in TIP its ability to transmit EKT security parameters (see section 4.2.5).

A TIP MCU that has negotiated EKT capability MUST send an SPIMAP (see section 4.2.6) packet every 250 msec until one or more of the following conditions are met:

a- MCU receives an ACK packet for its SPIMAP packet.

b- MCU receives an SRTP or SRTCP packet with an EKT extension that has the SPI value that was specified in the SPIMAP packet

c- After recommended interval of 5 seconds have passed

In the above (a) & (b) conditions, the MCU should consider EKT negotiated. In condition (c), the MCU MUST consider EKT negotiation to have failed.

An endpoint that has negotiated EKT capability with a MCU MUST not send any media packets (SRTP/SRTCP) until it has received an SPIMAP packet from the MCU and it has sent back an ACK packet. If no SPIMAP packet is received within the specified time frame, the endpoint MUST consider the EKT negotiation to have failed.

Once EKT has been negotiated for a TIP session, EKT authentication tags MUST be appended to both SRTP and SRTCP packets. It is RECOMMENDED that all SRTCP packets carry complete EKT tags (tags with final bit set to 1 that carries the full SRTP context). The one exception is for the SRTCP feedback packets, which might carry abbreviated EKT tags with a final bit set to zero that do not carry the SRTP context.

It is STRONGLY RECOMMENDED that the first video packet of a video refresh frame carries the full EKT authentication tag. It is also STRONGLY RECOMMENDED that periodically a packet with full EKT tag is transmitted for all video and audio streams.

For a message flow of a secure TIP control channel setup using DTLS and EKT, see section 7.3.

4 TIP Multiplex

4.1 Positional Multiplex

The key function of TIP is to enable multiple media streams, which belong to the same media type, to be transported on the same UDP channel within the boundaries of standards compliance. To accomplish this there needs to be a multiplex point for the multiple streams based on a label, here after referred to as the “position.”

All SSRC fields MUST be random 32 bit values that are unique across the RTP session as per the RTP and SRTP standards [4]. Each unique entity that is a RTP or RTCP transmitter should have a unique SSRC value.

The TIP positional multiplex value MUST be added to all RTP packets as the first CSRC value. (Note that multiple CSRC values are still allowed, but the first one must contain the TIP positional values.) The RTP CC field should be set appropriately to indicate the addition of the CSRC value. This document will henceforth refer to this value as the MUX-CSRC.

The following packet diagram illustrates these rules. For more information about rules for positional multiplex, see section 4.2.

TIP RTP Packet

```
        0                   1                   2                   3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+<+
     |V=2|P|X| CC=1  |M|     PT      |       sequence number         | |
     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ |
     |                           timestamp                           | |
     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ |
     |           synchronization source (SSRC) identifier            | |
     +=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+ |
     |    contributing “positional” source (MUX-CSRC) identifier     | |
   +>+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ |
   | |                          payload  ...                         | |
   | |                               +-------------------------------+ |
   | |                               | RTP padding   | RTP pad count | |
   +>+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+<+
```

TIP multiplexes multiple RTP media streams of the same type (audio or video) on one UDP channel.

For the video stream multiplex, typically all the streams use the same encoding format, H.264. Therefore, the RTP payload number is the same for all streams. Each video stream identifies its particular media configuration options via use of in-band H.264 parameter sets (SPS/PPS).

For the audio stream multiplex, multiple encoding formats such as AAC-LD, and/or G.711, and G.722 audio streams may be used depending on the device capabilities and the call topology. Any of the audio streams may include DTMF sent via RFC 2833 encoding. Multiple RTP payload numbers may be used within the audio multiplex.

The RTP media streams sent on the same UDP channel are multiplexed by using different RTP MUX-CSRC values for each stream. Each RTP stream has its own independent RTP sequence number and RTP timestamp space and a receiver should be prepared for a change in the value of the SSRC to be accompanied by a change of sequence number and timestamp. Note that multiplexing different, yet related, RTP streams of the same media is allowed in RTP as long as certain requirements are met [4].

Additionally, the bits of the MUX-CSRC field are subdivided and assigned specific semantics:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|           Sampling Clock ID           | OutPos|XmitPos|RcvPos |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Bits (Big-Endian)	Mask (Big-Endian)	Description
31-12	0xFFFFF000	Sampling Clock ID
11-8	0x00000F00	Output Position
7-4	0x000000F0	Transmitter Position
3-0	0x0000000F	Receiver Position

Table 2: TIP RTP header CSRC field

The Sampling Clock ID is a random number that is chosen by the transmitting endpoint to designate the media sampling clock used for this stream. If media streams share the same clock source then they should use the same value of the Sampling Clock ID. This allows TIP receivers to detect when two media streams should be strictly synchronized. This is consistent with [4]. In situations where there is no literal sampling clock, e.g. for control packets, a unique random number for the transmitter should be selected and used for this subfield.

The stream positions denote a spatial relationship or functional use of the stream. Audio and video streams with the same positions are intended to accompany each other -- i.e., there is a physical correspondence between the left audio position and the left video position.

The transmitter position denotes a both an RTP transmitter and an input channel as these are tightly bound.

The receiver position denotes an RTP receiver and a default output channel.

The output position is an option that can be used to direct a receiver to output media to a non-default channel. If the option is not used, the value of the output position SHOULD be zero. The Audio Dynamic Output option described later in this document uses this output position.

The definition of the video stream position values are listed in the following table:

Value	Position	Description
0	Control (see section 4.3.1)	Multiplex control
1	Center (see section 4.3.2)	Physical center
2	Left (see section 4.3.2)	Physical left
3	Right (see section 4.3.2)	Physical right
4	Aux 5fps (see section 4.3.3)	Auxiliary running at 1 or 5fps e.g. presentation, doc camera
5	Aux 30fps (see section 4.3.3)	Auxiliary running at 30 fps
6	<Not Used>	NA
7-8	<Not Used>	NA
9	Legacy Center (see section 4.3.4)	Legacy interoperability version of center
10	Legacy Left (see section 4.3.4)	Legacy interoperability version of left
11	Legacy Right (see section 4.3.4)	Legacy interoperability version of right
12	<Not Used>	NA
13-15	<Not Used>	NA

Table 3: Video stream position values

The definition of the audio stream position values are listed in the following table:

Value	Position	Description
0	Control (see section 4.3.1)	Multiplex control
1	Center (see section 4.3.2)	Physical center
2	Left (see section 4.3.2)	Physical left
3	Right (see section 4.3.2)	Physical right
4	Aux (see section 4.3.3)	Auxiliary audio
5	<Not Used>	NA
6	<Not Used>	NA
7-8	<Not Used>	NA
9	<Not Used>	NA
10	<Not Used>	NA
11	<Not Used>	NA
12	Legacy Mix (see section 4.3.4)	Legacy interoperability mix of all audio positions
13-15	<Not Used>	NA

Table 4: Audio stream position values

The transmitting source is allowed to switch the media streams (i.e. MUX-CSRCs) that are sent at any time. However, the total number of streams sent MUST NOT exceed the value negotiated with the receiver. For example, a transmitter may be capable and allowed to send any one of its camera streams to a receiver, but it may only send one at any given time. (Note that for video streams, the transmitting source should make provisions to ensure that a newly switched video stream can be decoded by the receiver, e.g. generation of an IDR frame for H.264.)

With the RTCP APP extensions, there are packets can be generated and/or consumed by a control entity that does not have a definitive position. In these situations the reserved value of “Control” (0) SHOULD be used for the position in the MUX-CSRC.

4.2 TIP RTCP APP Extensions

A number of control extensions are used in TIP. This section describes the TIP specific extensions carried in RTCP APP packets. Note that each of the TIP specific RTCP APP messages MUST be sent within compound RTCP packets as per [4]. The way to accomplish this is to prepend an “empty” receiver report (RR) and session description (SDES with CNAME) before the APP packet.

RTCP Packet with Empty RR, SDES CNAME, and TIP APP extension

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|V=2|P|  RC=0   |   PT=RR=201   |          length=1             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              SSRC                             |
+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
|V=2|P|  SC=1   |  PT=SDES=202  |          length               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              SSRC                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|    CNAME=1    |     length    |   user and domain name      ...
+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
|V=2|P| subtype |   PT=APP=204  |          length               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              SSRC                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           name=xcts                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
~                              ...                              ~
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

By default the TIP APP extensions are hop-by-hop, and are not required to be forwarded by MCUs. However some of the extensions do need to be forwarded. These extensions either specify a rule by which forwarding should be performed, or they include a “target” field that specifies the MUX-CSRC value of the intended recipient (from the perspective of the RTCP packet transmitter).

Note that currently TIP does not attempt to abide by the RTCP bandwidth limitations discussed in [4] with respect to the RTCP APP packets. The packets are small in size, limited in number and used only over unicast connections, and should minimally affect the RTCP bandwidth allocation.

4.2.1 Multiplex Control

This packet is used to negotiate and establish the parameters of the TIP multiplex. It MUST be sent after the media channel is established but before any media is transmitted. It MAY be sent later in the call to update some of the multiplex parameters.

RTCP APP MUXCTRL packet format

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|V=2|P|subtype=1|   PT=APP=204  |          length               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              SSRC                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           name=xcts                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| MV=8  | profile |  options    |numXmitStreams | numRcvStreams |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|             NTP timestamp, most significant word              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|             NTP timestamp, least significant word             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Conference ID, most significant word                         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Conference ID, least significant word                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         xmitPositions         |       rcvPositions            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| numShared     |   reserved    |         sharedPositions       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

RTCP APP Packet field definitions should follow [4] and:

`subtype`: 5 bits

MUST be set to ‘1’ to designate the MUXCTRL packet type.

`SSRC`: 32 bits

MUX SSRC field which is a random number

`name`: 4 bytes

MUST be set to the value ‘xcts’.

MUXCTRL Packet field definitions:

`mux version (MV)`: 4 bits

The version of the TIP multiplex. Implementations of this version MUST specify the value ‘8’.

`profile`: 4 bits TIP employs RTP profiles beyond the standard AVP profile described in [5]. These profiles specifically support for a unique Feedback Profile that is used to carry packet ACK/NACK information in a manner aligned with [8]. The MUXCTRL packet can be used to negotiate use of extended RTP profiles if it has not already been accomplished by call signaling. The allowed values assign numeric values to the IANA defined profiles. The values are in the following table:

Value	Name	Specification
0	RTP/AVP	RFC 3351
1	RTP/SAVP	RFC 3711
2	RTP/AVPF	TIP flavor of RFC 4585
3	RTP/SAVPF	TIP flavor of RFC 5124
4-15	Unused	NA

Table 5: MUXCTRL profile values

`options`: 8 bits

The following option bits are currently defined in the following table:

Bits	Value	Description
0	0x01	Sender is a multipoint focus device
1	0x02	Sender is a transcoding device
2-7	NA	Unused, should be zero.

Table 6: MUXCTRL sender –type option values

`num xmit streams`: 8 bits

The number of simultaneous media streams that can be transmitted within this multiplex by the system sending the MUXCTRL information.

`num rcv streams`: 8 bits

The number of simultaneous media streams that can be received within this multiplex by the system sending the MUXCTRL information. A transmitter MUST never send more than the number of streams the receiver is willing to accept.

`ntp timestamp`: 64 bits

The NTP time associated with the creation of the MUXCTRL information. Note that retransmissions with the same NTP timestamp are permitted.

`conference identifier`: 64 bits

A conference identifier used to associate with a particular multipoint call. The value of zero is reserved to indicate that the field is unused and invalid, e.g. in a point-to-point call.

`xmit positions:` 16 bits

A bitmask of the available transmit positions. Transmit position i is available if the i'th bit is set to 1, otherwise it is not allowed. See Table 3 for video position values and Table 4 Table 4 for audio position values.

`rcv positions:` 16 bits

A bitmask of the available receive positions. Receive position i is available if the i'th bit is set to 1, otherwise it is not allowed. See Table 3 for video position values and Table 4 for audio position values.

`num shared streams`: 8 bits

The number of streams that are characterized as shared. A shared stream has the following properties:

- there can only be one TIP device transmitting on a shared position at a time
- a TIP device can either transmit on a shared position or receive on a shared position but not both
- default state for both TX and RX of a shared stream is OFF

Control of the shared streams should be negotiated using the REQTOSEND packet as specified in section 4.2.7. Media should not be transmitted until the REQTOSEND negotiation has been completed.

`reserved`: 8 bits

Should be set to 0.

`Shared positions`: 16 bits

A bitmask of the shared positions. Shared transmission is enabled for position i if the i'th bit is set to 1, otherwise it is not a shared position. See Table 3 for video position values and Table 4 for audio position values.

It is RECOMMENDED to send MUXCTRL packets every 250 ms, for the first 15 seconds of the call (for a maximum of 60 transmissions). The NTP timestamp should be the same value for all retransmissions of the MUXCTRL packet. Once a MUXCTRL ACK packet is received from the peer with a received NTP time that matches the transmitted MUXCTRL packet, no more transmissions should be made.

In version 6 of the TIP protocol, a sender was allowed to changes its transmit and receive positions dynamically during a TIP session. It is STRONGLY RECOMMENDED that a sender NOT change its transmit and receive positions.

These requests MUST be explicitly acknowledged by the receiver using the RTCP APP ACK scheme described in section 4.2.9.

4.2.2 Network Path Measurement

To enable accurate measurements of the network path between peers, TIP specifies the RTCP APP packet extension "ECHO". This packet is similar in function to an ICMP ECHO, but has some significant advantages. By performing the echo function in-band, the same network path and QoS are applied to the measurement packets and the media packets. ICMP ECHO packets are frequently filtered by intermediate systems such as firewalls and the in-band RTCP messages are less likely to be filtered. .

TIP devices can use these RTCP APP ECHO packets to measure network path latency. Indirectly this also provides a keep-alive test for the network and the remote peer. The current best practice is to transmit an ECHO packet once per second during a call on each UDP channel, and measure averages over 10 second periods.

RTCP APP ECHO packet format

```
0                   1                   2                   3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|V=2|P|subtype=4|   PT=APP=204  |          length               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              SSRC                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           name=xcts                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Transmit NTP timestamp, most significant word         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Transmit NTP timestamp, least significant word        |
```

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Receive NTP timestamp, most significant word          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Receive NTP timestamp, least significant word         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

RTCP APP Packet field definitions should follow [4] and:

`subtype`: 5 bits

MUST be set to ‘4’ to designate the ECHO packet type.

`SSRC`: 32 bits

MUX SSRC field which is a random number

`name`: 4 bytes

MUST be set to the value ‘xcts’.

ECHO Packet field definitions:

`transmit ntp timestamp`: 64 bits

The NTP time of the transmission of the ECHO request. An ECHO response should contain the same transmission time as the request which generated the response.

`receive ntp timestamp`: 64 bits

The NTP time of the reception of the ECHO request. If this field is zero, then the packet is an ECHO request. If this field is non-zero, then the packet is an ECHO response.

There is no guarantee that the transmitter and receiver share a common NTP clock. This behavior is consistent with the IETF RTP model. Therefore, the receive timestamp may not fall between the transmit timestamp and the NTP time of the ECHO requestor’s reception of the ECHO response. The ECHO requestor can measure round trip time, and estimate one-way latency if symmetric network paths are assumed. The requestor can also detect if a common NTP clock is shared with the ECHO responder based on the expected value of the receive timestamp and the value actually returned.

These requests MUST NOT be explicitly acknowledged by the receiver using the RTCP APP ACK scheme described in section 4.2.9. These requests use their own acknowledgement mechanism.

4.2.3 Media Flow Control

This flow control applies to the media codec and not the RTP transmitter or receiver. The RTP entity should remain active awaiting either a request to continue media processing or call termination via the SIP control plane.

There is one packet format defined for FLOWCTRL, but two subtypes are used to distinguish between messages destined for transmitters and those destined for receivers.

RTCP APP FLOWCTRL packet format

```
0                   1                   2                   3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|V=2|P| subtype |   PT=APP=204  |          length               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              SSRC                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           name=xcts                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Transmit NTP timestamp, most significant word         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Transmit NTP timestamp, least significant word        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       Flow Control Opcode                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              Target                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|              .. Optional Opcode Specific Data ..              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

RTCP APP Packet field definitions should follow [4] and:

`subtype`: 5 bits

MUST be set to ‘5’ to designate the transmit flow control packet type or ‘6’ to designate the receive flow control packet type.

`name`: 4 bytes

MUST be set to the value ‘xcts’.

FLOWCTRL Packet field definitions:

`transmit ntp timestamp`: 64 bits

The NTP time of the transmission of the FLOWCTRL request. This can be used by the receiver to detect out of order or duplicate requests.

`flow control opcode`: 4 bytes

The flow control operation to perform, based on the tables below.

`target`: 4 bytes

The MUX-CSRC of the target that is being requested to act on the flow control state. Typically this value would be provided from the RTP media packets being received from the transmitter.

These requests MUST be explicitly acknowledged by the receiver using the RTCP APP ACK scheme described in section 4.2.9.

Value	Description
0	Start transmission notification
1	Stop transmission notification
2	H.264 control
3 – 65535	Reserved

Table 7: TXFLOWCTRL opcode values

Value	Description
0	Start reception notification
1	Stop reception notification
2 – 65535	Reserved

Table 8: RXFLOWCTRL opcode values

For the TXFLOWCTRL H.264 control message the following opcode specific data MUST be appended to the end of the base TXFLOWCTRL message.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                            Bitrate                            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|    H.264 Level Integer        |      H.264 Level Decimal      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        H.264 max-mbps                         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        H.264 max-fs                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        H.264 max-fps                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

TXFLOWCTRL H.264 control packet field definitions:

`Bitrate:  32 bits`

The bitrate to be transmitted in units of kilobits per second.

`H.264 Level Integer`: 16 bits

The integer portion of the H.264 level.

`H.264 Level Decimal`: 16 bits

The decimal portion of the H.264 level.

`H.264 max-mbps`: 32 bits

The maximum number of macroblocks per second allowed in units of macroblocks per second.

`H.264 max-fs`: 32 bits

The maximum frame size in units of macroblocks.

`H.264 max-fps`: 32 bits

The maximum number of frames per second, in units of hundredths of frames per second.

4.2.4 Video Refresh Request

When using an MCU with a multipoint focus, which is switching video streams, there is a need to have a mechanism to request that video encoders generate a new I Frame (IDR Frame in H.264). The following RTCP APP packet is the mechanism for that request.

This mechanism is not intended as the mechanism for repair of packet loss. The feedback mechanism detailed below should be used for that purpose.

RTCP APP REFRESH packet format

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|V=2|P|subtype=8|   PT=APP=204  |          length               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              SSRC                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           name=xcts                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Transmit NTP timestamp, most significant word         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Transmit NTP timestamp, least significant word        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              Target                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                               Flags                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

RTCP APP Packet field definitions should follow [4] and:

`subtype`: 5 bits

MUST be set to ‘8’ to designate the REFRESH packet type.

`name`: 4 bytes

MUST be set to the value ‘xcts’.

REFRESH Packet field definitions:

`transmit ntp timestamp`: 64 bits

The NTP time of the transmission of the REFRESH request. This can be used by the receiver to detect out of order or duplicate requests.

`target`: 4 bytes

The MUX-CSRC of the target encoder that is being requested to refresh. Typically this value would be provided from the RTP media packets being received from the transmitter.

`flags`: 4 bytes

Flags is a bit field that specifies the characteristics of the requested refresh frame. Currently only bit 0, the least significant bit, and 1 are used. Bit 0 specifies the receiver’s preference as to the type of the refresh frame (IDR versus GDR). Bit 1 specifies whether the transmitter SHOULD send the refresh frame type preferred by the receiver or whether the transmitter MAY send the refresh frame preferred by the receiver. All other bits MUST be set to zero.

Bit 0	Bit 1	Description
0	0	Receiver prefers and requires an IDR
1	0	Receiver prefers and requires a GDR (when negotiated)
0	1	Receiver prefers an IDR but does not require it
1	1	Receiver prefers a GDR but does not require it

Table 9: REFRESH Packet field values

These requests MUST be explicitly acknowledged by the receiver using the RTCP APP ACK scheme described in section 4.2.9.

4.2.5 Media Options

TIP provides several media encoding options, and extensions of the standard RTP media payloads to allow TIP devices to fully utilize advanced capabilities. By default, these options are disabled. They can be enabled via use of the RTCP APP MEDIAOPTS packet. The details of each media option are described in the following sections.

Some media options are specific to the media type, e.g audio-specific, while others are independent of the media type. Audio-specific and video-specific media options can, and do, share the same identifiers as they are transmitted using two different TIP control sessions.

For the same media type, some media options are specific to a media format, e.g. H.264-specific, while others are media format independent. In TIP V7 options that are specific to a media format such as H.264 MUST NOT share the same identifiers with media options specific to another media format such as H.263. Such sharing would either require restricting the negotiation of more than one media format in call signaling which is not compatible with call signaling protocols such as SIP/SDP. Future TIP releases can extend the Media Option packet format to bind media options to RTP payload types in which case such sharing of identifiers can be allowed.

RTCP APP MEDIAOPTS V3 packet format

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|V=2|P|subtype=7|   PT=APP=204  |          length               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              SSRC                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           name=xcts                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Transmit NTP timestamp, most significant word         |
```

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Transmit NTP timestamp, least significant word         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|           Version              |           Reserved             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           target SSRC                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                         Transmit options                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          Receive options                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   Option tag  |           Option value                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
~                             ...                                 ~
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|Last Option tag|           Option value                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           target SSRC                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                         Transmit options                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          Receive options                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   Option tag  |           Option value                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
~                             ...                                 ~
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|Last Option tag|           Option value                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
~              Additional Target SSRCs and options                ~
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|Last Option tag|           Option value                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```


RTCP APP Packet field definitions should follow [4] and:

`subtype`: 5 bits

MUST be set to '7' to designate the MEDIAOPTS packet type.

`name`: 4 bytes

MUST be set to the value 'xcts'.

MEDIAOPTS Packet field definitions:

`transmit ntp timestamp`: 64 bits

The NTP time of the transmission of the request. This can be used by the receiver to detect out of order or duplicate requests.

`version`: 2 bytes

The version of the MEDIAOPTS packet. Implementations compliant with this specification MUST set the value to ’3’.

`reserved`: 2 bytes

MUST be set to 0.

The V3 MEDIAOPTS packet supports sending differing media options on a per SSRC basis. Each SSRC section begins with a ‘target SSRC’ field and continues until the ‘Last Option Tag’ field is seen. The overall packet length can be used to identify when the last SSRC specific section has been processed.

`target ssrc`: 4 bytes

This field identifies the SSRC with which the options should be associated. If the SSRC is set to 0x00000000, the options in the associated section will be used as default values for SSRCs that do not have a more specific set of options.

`transmit options`: 4 bytes

The media options that the sender of the MEDIAOPTS packet is willing to transmit. See the following table for media option value assignments.

`receive options`: 4 bytes

The media options that the sender of the MEDIAOPTS packet is willing to receive. See the following table for media option value assignments.

Each option tag/value is 32 bits, and is subdivided as follows:

`Option tag`: 1 byte

A unique identifier for the option. (See below for assignments).

`Option value`: 3 bytes

A 24 bit value whose interpretation is determined by the option tag (See below for definitions).

These MEDIAOPTS requests MUST be explicitly acknowledged by the receiver using the RTCP APP ACK scheme described in section 4.2.9.

Audio-specific[1] media options field bits for an audio multiplex are assigned the following:

Bits	Value	Description
0	0x00000001	Audio Activity Metric
1	0x00000002	Audio Dynamic Output Channels
2	0x00000004	Capable of G.722 for Legacy Audio
3	0x00000008	Using G.722 for Legacy Audio
4-23	N/A	Unused set to 0

Table 10: Audio-specific media options values

Video-specific[2] media options field bits for a video multiplex are assigned the following:

Bits	Value	Description
0	0x00000001	Refresh Flag
1	0x00000002	Inband Parameter Sets (SPS/PPS) – H.264-specific
2	0x00000004	Arithmetic Coding (CABAC) – H.264-specific
3	0x00000008	Long-Term Reference Pictures (LTRP) – H.264-specific
4	N/A	Unused set to zero
5	0x00000020	Aux Video FPS Bit 0
6	N/A	Unused set to zero
7	N/A	Unused set to zero
8	0x00000100	Gradual Decoder Refresh (GDR) – H.264-specific
9	0x00000200	Aux Video FPS Bit 1
10	0x00000400	8x8 Transforms (High Profile) – H.264-specific
11	0x00000800	XGA 5 or 1 fps unrestricted media
12	0x00001000	720p unrestricted media
13	0x00002000	1080p unrestricted media
14	0x00004000	XGA 30 fps unrestricted media
15	0x00008000	Unrestricted Media Constraints
16	0x00010000	Prefer BFCP for auxiliary control
17-23	N/A	Unused set to zero

[1] Options that are specific to a media format will be indicated as such in the table
[2] Options that are specific to a media format will be indicated as such in the table

Table 11: Video Specific media options values

The media options field bits shared by both audio and video multiplexes are assigned the following values:

Bits	Value	Description
24	N/A	Unused set to zero
25	0x02000000	EKT
26-31	N/A	Unused

Table 12: Media options values used by both Audio and Video

An option is enabled if it has been offered by the local system and the complementary option has been offered by the remote peer. For example, if the local system has offered to transmit an option, and the remote peer has offered to receive the option, the option is enabled. Otherwise it is not. In pseudo-code:

```
bool transmitOptionEnabled = localTransmitOffered & remoteReceivedOffered;
bool receiveOptionEnabled = localReceivedOffered & remoteTransmitOffered;
```

An implementation can independently offer to transmit or receive a media option. For instance, a TIP endpoint might offer to transmit the Audio Activity Metric, but not offer to receive this option.

As required, new MEDIAOPTS packets can be sent to renegotiate the media options in use. This occurrence can happen during a multipoint conference due to the entry or departure of an endpoint with more limited capabilities than the other members of the conference.

Any intermediate entity that has enabled media options and is forwarding the media packets to other entities must take responsibility for reversing these options if the downstream receiver has not enabled them. The burden of interoperability with standards compliant implementations falls on the option enabler.

4.2.5.1 Audio Dynamic Output Channels

Typically, audio RTP streams are statically mapped to an output channel (speaker) configuration. In the highly dynamic environment of a multi-channel, multipoint conference that static mapping is suboptimal. TIP supports the capability of multiple audio streams with dynamically changing output channel mappings. This enables TIP devices to mix multiple audio streams when more than one stream is mapped to a particular output channel.

To allow for highly dynamic switching of audio streams to different output channels, as with an MCU with a multipoint focus, this media option enables the interpretation of the output position of the RTP MUX-CSRC field. The bitmask 0x00000F00 is used to select the bits of the MUX-CSRC that represent the desired output channel. The position values defined in Table 4 apply for this position.

4.2.5.2 Audio Activity Metric

TIP supports devices capable of generating and receiving a real-time voice activity confidence metric for audio input channels. This media option transmits that 1 byte metric at the end of the standard RTP audio payload for the codec in use. If more than one audio frame is sent per packet, the metric is the average value of the included audio frames.

If used with SRTP, then the appended metric byte MUST NOT be encrypted (as part of the RTP payload it typically would be encrypted under SRTP). It MUST however still consume a byte of any keystream used by the encryption transform. Note this rule is functionally relevant only if additional encrypted media options are later defined and appended to the RTP payload. Also the (unencrypted) value MUST be included in the SRTP packet authentication code (HMAC). These rules enable intermediate entities, such as MCU with a multipoint focus, to do stream switching without the need to decrypt and then potentially re-encrypt every audio packet.

Implementations that allow for transmission of this media option should ensure that the additional byte does not induce IP layer fragmentation due to MTU limitations.

4.2.5.3 G.722 Legacy Audio

TIP supports devices capable of sending or receiving an extra audio stream encoded with a different encoding format to that negotiated in SIP. This is achieved within the multiplex via use of the Legacy Mix position. By default, the encoding of this audio stream is G.711. This media option allows signaling of the capability to use G.722 as an alternate encoding. Support for this alternate encoding may be only for transmit, only for receive, or both depending on the capabilities of the system.

If a given transmitter/receiver pair between peers is capable of supporting the G.722 encoding for the legacy audio, then an additional media option bit is defined to signal the actual desire to use the capability.

Note that these two audio encodings use different, static RTP payload numbers (G.711 uses 0 and G.722 uses 9). Also no encoding specific parameters are needed. Hence these two media option bits are sufficient to accomplish the equivalent of an SDP offer/answer [2] negotiation between these two encodings. Receivers can also easily distinguish the encodings to enable graceful handling of mid-call media transitions between the two audio encodings.

4.2.5.4 Video Refresh Flag

For multipoint sessions, it is desirable for receivers to be able to easily distinguish video refresh points in the RTP packet stream. This can be done with video codec and RTP payload specific knowledge, but it requires fairly complicated code. To ease this task, TIP provides an option where such video refresh packets can be marked with an explicit refresh flag. This flag is carried in a single byte appended to the end of the standard RTP video payload for the codec in use. The extra byte is sent for every video packet, with the byte set as follows:

Byte	Description
0	Not a refresh packet
1	First packet of an IDR frame
2	First packet of a GDR frame
3	First packet of an LTRP 0 candidate frame
4	First packet of an LTRP 1 candidate frame
5	First packet of a repair frame associated with LTRP 0
6	First packet of a repair frame associated with LTRP 1
7	First packet of a repair frame not using LTRP

Table 13: Video Refresh Flag field values

Note this refresh flag is analogous to the carriage of the audio activity metric described in the previous section, and hence the same rules apply for use with SRTP. The appended flag byte MUST NOT be encrypted. It MUST however still consume a byte of any keystream used by the encryption transform. Also the (unencrypted) value MUST be included in the SRTP packet authentication code (HMAC).

Implementations that allow for transmission of this media option should ensure that the additional byte does not induce IP layer fragmentation due to MTU limitations.

4.2.5.5 Video Parameter Sets

To support highly dynamic video switching among a set of compatible video streams, as with an MCU with a multipoint focus, there is a need to have video configuration information travel in-band with the video data. In H.264 this video configuration information is contained in "parameter sets", specifically the sequence parameter set (SPS) and the picture parameter set (PPS). Since video switching can only usefully occur at video refresh points (IDR frames), the video parameter sets should accompany these frames. When enabled, this media option instructs the video transmitter to include the parameter sets with the refresh frames.

This in-band option for parameter sets is allowed by [7] but only when a means of providing reliable transport exists. For TIP devices, this reliable transport is accomplished via the RTP feedback mechanism described later in this document.

4.2.5.6 Video Coding Mode

TIP allows devices which operate in H.264 baseline profile, which is defined by the ITU to use context adaptive variable length coding (CAVLC), to use context adaptive binary arithmetic coding (CABAC), that yields increased performance, i.e., a lower bit rate. Although CABAC is part of the H.264 standard, it is not allowed under a strict (i.e., interoperable) definition of the baseline profile. When enabled this media option allows the video codec to use CABAC in conjunction with the baseline profile.

4.2.5.7 Video Long Term Reference Pictures

This media option enables use of this encoding feature if both peers are capable of using it.

4.2.5.8 Aux Video FPS

TIP supports three discrete frame rate values for auxiliary video streams: 1 frame per second, 5 frames per second, or 30 frames per second. A TIP device will advertise the highest number of frames per second it can support, lower levels are supported implicitly. Note that frame rates 1 and 5 always use position 4, while frame rate 30 always uses position 5. An aux transmitter MUST always advertise (via its REQTOSEND packet described later) at least one position that satisfies the remote endpoint’s supported frame rate. An aux transmitter MAY advertise additional positions if it is capable of generating them.

This advertisement is done using the two Aux Video FPS bit as shown in the following table:

Bit 0	Bit 1	Required Frame Rate	Position 4 Frame Rate
0	0	5 FPS	5 FPS
0	1	1 FPS	1 FPS
1	0	30 FPS	5 FPS
1	1	N/A	N/A

Table 14: AUX Video Frames Per Second indication

4.2.5.9 Gradual Decoder Refresh

This bit indicates that the sender supports sending (TX options) or receiving (RX options) gradual decoder refresh frames. When enabled on the transmitter, the video encoder will generate GDR frames in place of IDR frames.

4.2.5.10 High Profile

This bit indicates that the sender supports sending (TX options) or receiving (RX options) high profile H.264 video media. When enabled on the transmitter, the video encoder (high definition and 30FPS XGA only) will generate media using 8x8 Transforms.

4.2.5.11 Restricted Media

For each of the resolutions, a receiver that can receive either restricted or unrestricted video media at the resolution will set the corresponding RX options unrestricted bit to 1. If the receiver can receive only restricted media at that resolution then the bit will be set to 0.

For each of the resolutions, a transmitter that can transmit only unrestricted video media at the resolution will set the corresponding TX options bit to 1. If the transmitter can transmit either restricted or unrestricted media at that resolution then the bit will be set to 0.

The following table shows for a given resolution what the transmitter and receiver behavior should be. Additional compatibility information can be found in section 6.2.

RX	TX	Description
0	0	Only restricted media will be sent at this resolution
0	1	This resolution is not supported
1	0	Either restricted or unrestricted media can be sent at this resolution
1	1	Only unrestricted media will be sent at this resolution

Table 15: Restricted Video indication

Media at a given resolution qualifies as restricted if it meets the following requirements.

The following Restricted Media requirements are in place for 720p and 1080p, as well as XGA when transmitted at 30 fps.

1. MUST support H.264 Main Profile or High Profile (if negotiated) compatible bit stream.
2. MUST support CAVLC and optionally CABAC.
3. MUST support fixed frame rate: 30 fps or 29.97 fps.
4. MUST generate one macro block-row per slice.
    a. 1280x720: 45 slices per frame, each slice must be 80 macro blocks in length
    b. 1920x1072: 67 slices per frame, each slice must be 120 macro blocks in length
    c. 1024x768: 48 slices per frame, each slice must be 64 macro blocks in length
5. MUST support deblocking_filter_control_present = 1 & disable_deblocking_filter_idc = 1

6. MUST only use interprediction blocks of size 16x16.
7. MUST only use one reference picture.
8. Each NAL MUST be less than 3320 bytes in size.
9. POC MUST increment by 2 every frame.
10. When high profile is negotiated, MUST support seq_scaling_matrix_present_flag=0 and pic_scaling_matrix_present_flag=0, MUST NOT use monochrome pictures, MUST NOT use quantization matrices.
11. Long term reference pictures (LTRPs) MAY be used for error concealment. LTRPs require a four frame picture buffer. LTRPs use H.264 standard picture buffer management. MUST support long_term_frame_idx < 2, MUST use memory_management_control_operation = 6. LTRPs are optional and can be disabled.
12. Gradual decoder refresh pictures (GDRs) MAY be used at the start of a sequence instead of an instantaneous decoder refresh picture (IDR). A recovery point SEI precedes the GDR. The SEI always has the broken_link_flag set to 0. GDRs are optional and can be disabled.
13. Must not use SPS IDs in the range 10-20, inclusive.

The following requirements are in place for XGA when transmitted at either 1 or 5FPS.

1. MUST support H.264 Baseline Profile.
2. MUST support 1024x768 resolution.
3. MUST support Fixed frame rate: 1 or 5 fps
4. MUST support the use of one reference picture from previous frame.
5. MUST support one macro block-row per slice, 48 slices per frame, each slice must be 64 macro blocks in length.
6. MUST support deblocking_filter_control_present = 1 & disable_deblocking_filter_idc = 1
7. MUST only support inter-prediction blocks of size 16x16.
8. Each NAL MUST be less than 3320 bytes in size.
9. POC MUST increment by 2 every frame.

4.2.5.12 Unrestricted Media Constraints

When operating at an unrestricted resolution a device may indicate that it cannot receive certain features of unrestricted H.264 video by setting the RX bit to 1. A device that has no such constraints will set the RX bit to 0.

A device capable of transmitted unrestricted video that satisfies the constraints will set the TX bit to 1. A device not capable of transmitting unrestricted video with constraints will set the TX bit to 0.

The constraints required by this option are:

1. Each NAL MUST be less than 3320 bytes in size
2. CABAC encoding MUST NOT be used
3. High Profile encoding MUST NOT be used

The following table shows what the transmitter and receiver behavior should be.

RX	TX	Description
0	0	No constraints required
0	1	No constraints required
1	0	Incompatible
1	1	Unrestricted Media with Constraints required

Table 16: Unrestricted Media Constraints indication

When operating in an Unrestricted Media with Constraints mode, the individual CABAC and High Profile MEDIAOPTS bit indications are superseded by this mode which dictates these are not to be used. Thus, transmitters must not send encoded video media using those features or profiles regardless of what a receiver indicated in those MEDIAOPTS bits when the unrestricted media constraints mode is negotiated.

When operating in an unrestricted media without constraints mode, the individual CABAC and High Profile bits must be examined to determine the state of those features. When operating in the default restricted media mode, the individual CABAC and High Profile bits must be examined to determine the state of those features.

4.2.5.13 Prefer BFCP for auxiliary control

This bit MUST have the same value set for TX and RX option fields.

When set to 1, this bit indicates that the sender prefers to use a separate media line and the BFCP protocol [16] for auxiliary video control instead of using the TIP mechanism outlined in 4.2.7.

This bit MUST be set to 0 if a BFCP controlled media line was not negotiated in the SIP/SDP offer/answer. This bit MAY be set to 1 if a BFCP controlled media line was negotiated in the SIP/SDP offer/answer.

When both TIP peers set this to 1, then the following TIP negotiation fields MUST be ignored

- MUXCTRL numShared
- MUXCTRL sharedPositions
- MEDIAOPTS Aux Video FPS Bit 0
- MEDIAOPTS Pre-Allocated fixed presentation bandwidth
- MEDIAOPTS On-Demand presentation bandwidth
- MEDIAOPTS Aux Video FPS Bit 1
- MEDIAOPTS XGA 5 or 1FPS unrestricted media
- MEDIAOPTS XGA 30FPS unrestricted media

and, the following TIP messages MUST NOT be transmitted and MUST be ignored if received.

- REQTOSEND
- REQTOSEND-ACK
- Video TXFLOWCTRL with target CSRC source position of 4 or 5
- Video RXFLOWCTRL with target CSRC source position of 4 or 5
- Video REFRESH with target CSRC source position of 4 or 5
- Video RTP packets with a CSRC source or destination position of 4 or 5

4.2.5.14 EKT

This signals the ability to send and receive group EKT Security parameters and support of EKT tag transmission in both RTP and RTCP as specified in [12].

If no initial security associations have been exchanged between the endpoint and the MCU prior to the start of the TIP session, then no EKT support SHOULD be indicated on either side.

Per section 3.2, transmission of EKT tags MUST NOT start till after the negotiation of the EKT capability in MediaOptions and the successful exchange of the SPIMAP packet (see section 4.2.6).

4.2.5.15 Options Tags and Values

The media options tags are assigned as follows:

Tag	Description
0	Reserved
1	Transmitter Profile (see below for values)
2	Unused

3	Unused
4	Unused
5	Legacy audio mix control. Bitfield indicating the audio positions which should be used as part of the legacy audio mix output of a TIP device.
6	Transmitter’s optimal receive video resolution, expressed as a number of rows (i.e. 1080 for 1080p).
7	Legacy video bitrate in kilobits per second. The maximum bitrate (per stream) to be reserved for the legacy video streams. If not present, the default value of 704Kbps should be used. Note that this value is only relevant when the legacy video positions are active (as determined by MUXCTRL position negotiation).
8-254	Unused
255	Last option tag/value for this section. The next word in the packet, if any, starts a new section.

Table 17: Media Options Tags

The transmitter profile values are assigned as follows:

Value	Description
0	Reserved
1	Satellite deployment
2	Public Internet

Table 18: Transmitter Deployment Profiles

Note: The purpose of the current options is to enable TIP devices to adapt their network error condition algorithms to values that are appropriate for varying deployment scenarios that can occur in different calls/conferences.

4.2.6 SPI Map

TIP supports the use of EKT to communicate and update SRTP contexts between transmitters and receivers in a multipoint session. EKT depends on the sharing of a set of security parameters between the transmitters and receivers of an SRTP session. Each security parameter set is uniquely identified by an index (SPI). The SPI is transmitted as part of each EKT tag to allow receivers to use the same security parameters used by the transmitter. TIP uses the SPIMAP packet to communicate the security parameter sets and indexes used in a multipoint session to all participants in the session.

RTCP APP SPIMAP packet format

```
0                   1                   2                   3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|V=2|P|subtype=12|   PT=APP=204  |          length               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              SSRC                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           name=xcts                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Transmit NTP timestamp, most significant word           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Transmit NTP timestamp, least significant word          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|             SPI               |  SRTP PProfile  | EKT PProfile  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          SRTP Salt ...                          |
|                               +---------------------------------+
|                               | KEK ...                         |
+-------------------------------+                                 +
|                                                                 |
|                               +---------------------------------+
|                               | Padding                         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

RTCP APP Packet field definitions should follow [5] and:

`subtype`: 5 bits

MUST be set to ‘12’ to designate the SPIMAP packet type.

`name`: 4 bytes

MUST be set to the value ‘xcts’.

SPIMAP Packet field definitions:

`transmit ntp timestamp`: 64 bits

The NTP time of the transmission of the request. This can be used by the receiver to detect out of order or duplicate requests.

`SPI`: 2 bytes

The session protection index. Note the least significant bit MUST be set to zero.

`SRTP PProfile`: 1 byte

This represents the SRTP Protection Profile as specified in [15]. Note that in [15] two bytes are used to represent the SRTP Protection Profile while TIP only uses the second byte. For example, the default SRTP crypto suite will be represented as 0x01 instead of {0x00, 0x01}.

For this specification the default protection profile MUST be supported. The SRTP Protection Profile parameters for this profile (0x01) are:

SRTP crypto suite: AES_128_CM
SRTP master key length: 128 bits
SRTP salt key length: 112 bits
SRTP/SRCTP authentication function: HMAC-SHA1
SRTP/SRTCP authentication key length: 160 bits
SRTP/SRTCP authentication tag length: 80 bits

`EKT PProfile`: 1 byte

This represents the EKT cipher suite used for the encryption of the SRTP master key before transmitting it in the EKT extension. It also represents the length of the KEK key.

For this specification, the default profile (value 0x01) MUST be supported and will have the following parameters:

EKT crypto suite: AES_128_ECB
KEK length: 128 bits

`SRTP Salt`: variable number of bytes

The SRTP salt value. The length is determined by the SRTP Protection Profile value. For the default SRTP Protection Profile, the length is 14 bytes (112 bits).

`KEK`: variable number of bytes

The Key Encrypting Key. The length is determined by the EKT Protection Profile value. For the default EKT Protection Profile, the length is 16 bytes (128 bits).

`padding`: variable number of bytes

RTCP packets must be a multiple of 32 bits in length. As many zeroed bytes as necessary should be added after the KEK to meet this requirement.

These requests SHOULD be explicitly acknowledged by the receiver using the RTCP APP ACK scheme described in section 4.2.9. These requests MUST be encrypted and authenticated using the security association negotiated prior to the start of the TIP session.

4.2.7 Request to Send

TIP supports the notion of dynamic media sources that come and go during the course of a call. An example of this is a presentation sharing feature.

When a new media source needs to be added to an existing session the TIP device MUST send a REQTOSEND message to its remote peer. The remote peer can ACK or NACK the request.

The local TIP device MUST only begin transmission of the new media source after it has received an ACK from the remote peer.

When a media source needs to be removed from an existing session the local TIP device MUST send a REQTOSEND message to the remote peer indicating that the media source has stopped. The local system MAY stop transmitting media for this position before sending the REQTOSEND.

The REQTOSEND packet MUST be sent over only one of the media control channels, it is RECOMMENDED to send this packet over the video channel when possible.

RTCP APP REQTOSEND packet format

```
0                   1                   2                   3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|V=2|P|subtype=15|  PT=APP=204  |           length              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              SSRC                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           name=xcts                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Transmit NTP timestamp, most significant word         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Transmit NTP timestamp, least significant word        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                             flags                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|       VideoPositions          |         AudioPositions        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

RTCP APP Packet field definitions should follow [5] and:

`subtype`: 5 bits

MUST be set to ‘15’ to designate the REQTOSEND packet type.

`name`: 4 bytes

MUST be set to the value ‘xcts’.

REQTOSEND Packet field definitions:

`transmit ntp timestamp`: 64 bits

The NTP time of the transmission of the request. This can be used by the receiver to detect out of order or duplicate requests.

`flags`: 4 bytes

Set of flags specifying the interpretation of the positions fields. Values are:

Value	**Description**
0	Stop transmission notification
1	Request to start transmission
2-0xFFFFFFFF	Unused

Table 19: REQTOSEND Flags values

`VideoPositions`: 2 bytes

A bitmask of the video transmit positions being modified.

`AudioPositions`: 2 bytes

A bitmask of the audio transmit positions being modified.

These requests SHOULD NOT be explicitly acknowledged by the receiver using the RTCP APP ACK scheme described in section 4.2.9. Instead the receiver should use the above packet format with the following field definitions.

RTCP APP Packet field definitions should follow [5] and:

`subtype`: 5 bits

MUST be set to '31' to designate the REQTOSENDACK packet type.

`name`: 4 bytes

MUST be set to the value 'xcts'.

REQTOSENDACK Packet field definitions:

`transmit ntp timestamp`: 64 bits

The NTP time of the transmission of the request being ack'ed, ie copied from the REQTOSEND packet.

`flags`: 4 bytes

The flags field should match the flags field in the REQTOSEND packet.

`VideoPositions`: 2 bytes

A bitmask of the dynamic video positions acknowledged as changing state.

`AudioPositions`: 2 bytes

A bitmask of the dynamic audio positions acknowledged as changing state.

4.2.8 Notify

The NOTIFY message is a generic notification message used by TIP devices to convey real time information related to changes in the media state. Where possible, out of band methods of conveying this information should be preferred and where available specific TIP messages (e.g. REQTOSEND) should be preferred over using a NOTIFY.

RTCP APP NOTIFY packet format

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|V=2|P|subtype=1 |  PT=APP=204  |           length              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              SSRC                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           name=xcts                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Transmit NTP timestamp, most significant word         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Transmit NTP timestamp, least significant word        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|      Tag      |  ValueLength  |            Value...           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+                               +
|                                                               |
|                               +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                               |      Tag      |  ValueLength  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Value...                                                     |
|                               +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                               | Padding                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

RTCP APP Packet field definitions should follow [5] and:

`subtype`: 5 bits

MUST be set to '3' to designate the NOTIFY packet type.

`name`: 4 bytes

MUST be set to the value 'xcts'.

NOTIFY Packet field definitions:

`transmit ntp timestamp`: 64 bits

The NTP time of the transmission of the request. This can be used by the receiver to detect out of order or duplicate requests.

One or more tag / length / value (TLV) combinations, bounded by overall packet length.

Each TLV consists of:

tag: 1 byte, a unique identifier

length: 1 byte, the number of bytes of the value that immediately follows. Zero is allowed.

value: variable number [0-255] of bytes

`padding`: variable number of bytes

RTCP packets must be a multiple of 32 bits in length. As many zero'ed bytes as necessary should be added after the last TLV to meet this requirement

The TLV tags are assigned as follows:

Tag	Description	Length	Value
0	Reserved	NA	NA
1	Unused	NA	NA
2	Unused	NA	NA
3	Security State irrespective of whether SRTP or RTP is being used for media transport	1	0 to indicate non-secure, 1 to indicate secure
4-255	Unused	NA	NA

Table 20: NOTIFY Tags

4.2.9 Acknowledgement

Most of the RTCP APP extensions described in this document benefit from having an explicit acknowledgement from the receiver. In this section, a generic ACK packet is defined for use with these extensions and any future extensions that may require it.

The most significant bit of the RTCP APP subtype is defined to represent an ACK bit. The least significant nibble (4 bits) of the subtype is used to carry the extension identifier of the packet being acknowledged. This results in the following allocation of the subtype values for TIP extensions.

Subtypes within the 'xcts' namespace

Subtype	Extension

0	Reserved
1	MUXCTRL
2	Unused
3	NOTIFY
4	ECHO
5	TXFLOWCTRL
6	RXFLOWCTRL
7	MEDIAOPTS
8	REFRESH
9	Unused
10	Unused
11	Unused
12	SPIMAP
13	Unused
14	Unused
15	REQTOSEND
16	Reserved
17	MUXCTRL ACK
18	Unused
19	NOTIFY ACK
20	Reserved for ECHO ACK
21	TXFLOWCTRL ACK
22	RXFLOWCTRL ACK
23	MEDIAOPTS ACK
24	REFRESH ACK
25	Unused
26	Unused
27	Unused
28	SPIMAP ACK
29	Unused
30	Unused
31	REQTOSEND ACK

Table 21: ACK Packet Subtype values

RTCP APP ACK packet format

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|V=2|P| subtype |   PT=APP=204  |          length               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              SSRC                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           name=xcts                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Transmit NTP timestamp, most significant word         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Transmit NTP timestamp, least significant word        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

RTCP APP Packet field definitions, should follow [4] and:

`subtype`: 5 bits

MUST be set to the subtype field of the RTCP APP packet being acknowledged plus 16 (bitwise-or of 0x10)

`name`: 4 bytes

MUST be set to the value ‘xcts’.

ACK Packet field definitions:

`transmit ntp timestamp`: 64 bits

The NTP time contained in the RTCP APP extension packet being acknowledged.

Reception of RTCP APP packets with NTP timestamps that are older than the last ACK transmitted SHOULD NOT be acknowledged.

Reception of RTCP APP packets with NTP timestamps equal to the last ACK transmitted SHOULD be re-acknowledged (as the ACK may have been lost).

4.2.10 RTP Feedback

Depending on the capabilities of the encoder to adjust to packet loss, TIP devices may use an RTP profile such as [8] that allows for timely receiver feedback on specific packet reception or loss. Use of this profile is indicated in the RTCP APP MUXCTRL packet discussed in section 4.2.1 by setting the profile field value to ‘2’ signaling the RTP/AVPF profile.

When used, TIP devices will generate transport layer feedback messages (RTCP packet type RTPFB). The FMT parameter of these messages types is set to private value 30, and a custom feedback control information (FCI) block is used.

RTCP RTPFB packet format with custom FCI

```
0                   1                   2                   3
   0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |V=2|P|  FMT=30 |  PT=RTPFB=205 |          length               |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                  SSRC of packet sender                        |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                  MUX-CSRC of media source                     |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

   |             PID               |  PPA byte 0   |   PPA byte 1  |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |   PPA byte 2  |   PPA byte 3  |  PPA byte 4   |   PPA byte 5  |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |   PPA byte 6  |   PPA byte 7  |  PPA byte 8   |   PPA byte 9  |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |   PPA byte 10 |   PPA byte 11 |  PPA byte 12  |   PPA byte 13 |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |           Reserved            | PPAm byte 0   |  PPAm byte 1  |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |  PPAm byte 2  |  PPAm byte 3  | PPAm byte 4   |  PPAm byte 5  |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |  PPAm byte 6  |  PPAm byte 7  | PPAm byte 8   |  PPAm byte 9  |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |  PPAm byte 10 |  PPAm byte 11 | PPAm byte 12  |  PPAm byte 13 |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

RTCP RTPFB fields:

See [8].

FCI fields:

`packet ID (PID)`: 16 bits

The RTP sequence number of the media SSRC that is being acknowledged as received.

`previous packet acknowledgements (PPA)`: 112 bits

A bitmask indicating whether the previous 112 RTP packets relative to the packet identified by the PID field have been received or not. The position of the packet with the RTP sequence number X, where 0 < PID – X < 113 can be found using the following formulas:

PPA byte position = (112 + X – PID) / 8
PPA bit position = (112 + X – PID) % 8

All variables should be unsigned 16 bit numbers to allow for proper sequence number wrap-around handling. Note that the byte and bit positions are 0 based. A bit value of 0 indicates no packet reception, a bit value of 1 indicates reception.

The length of this bit field is designed to provide adequate redundancy against loss of the feedback packets themselves up to the expected maximum bitrate of a typical session.

`previous packet acknowledgements mask (PPAm)`: 112 bits

A bitmask indicating which of the 112 bits in the PPA field are valid. The RTP sequence number to bit mapping is the same as the PPA field. A bit value of 0 indicates the corresponding PPA bit is invalid, a bit value of 1 indicated it is valid. Invalid PPA bit values should be ignored by the receiver as they do not indicate an ACK or a NACK.

The PPAm field is optional; its presence should be detected by the length of the RTPFB packet.

The primary purpose of the PPAm field is to allow RTPFB packets to be sent out of order without implicitly indicating reception or no reception for the previous packets in the PPA bitfield. A PPAm bit set to zero indicates that the corresponding PPA bit CAN NOT be used to establish whether the associated packet has been received or not.

TIP only supports the feedback profile for video streams. The TIP RTCP RTPFB packets should be generated by the receiver at the time that a video frame is consumed. For a typical implementation of feedback, this is when the packets corresponding to a video frame are dequeued from the receiver's jitter buffer. The number of RTPFB packets generated should have a one-to-one correspondence with the received video frame rate. Note however that the PPA field will correspond to both the current video frame, and a number of previous video frames. This is intentional, and provides for retransmission of the ACK/NAK data to protect against loss of the RTPFB packets themselves.

Note that each of the TIP specific RTCP RTPFB messages SHOULD be sent within compound RTCP packets as per [4] and [8]. The simplest way to accomplish this is to prepend an “empty” receiver report (RR) and session description (SDES with CNAME) before the APP packet.

RTCP Packet with empty RR, SDES CNAME and TIP RTPFB

```
0                   1                   2                   3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|V=2|P|  RC=0   |   PT=RR=201   |          length=1             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              SSRC                             |
+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
|V=2|P|  SC=1   |  PT=SDES=202  |          length               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              SSRC                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|    CNAME=1    |     length    |   user and domain name      ...
+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
|V=2|P|  FMT=30 | PT=RTPFB=205  |          length               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  SSRC of packet sender                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  MUX-CSRC of media source                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|             PID               |  PPA byte 0   |   PPA byte 1  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   PPA byte 2  |  PPA byte 3   |  PPA byte 4   |   PPA byte 5  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   PPA byte 6  |  PPA byte 7   |  PPA byte 8   |   PPA byte 9  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   PPA byte 10 |  PPA byte 11  |  PPA byte 12  |   PPA byte 13 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```


4.3 Positional Rules

4.3.1 Control Position

The control position (value 0) of the multiplex is reserved for use with RTCP packets that carry TIP APP extension packets that pertain to the entire multiplex and not a specific media stream within the multiplex, e.g., MUXCTRL.

RTP packets MUST NOT use the control position for either the transmitter or receiver positions. They MAY use it as an “undefined” value for the output position when that option is not used.

4.3.2 Center, Left, Right Positions

The center (value 1), left (value 2), and right (value 3) positions correspond to the physical segments of a typical “three screens” TIP device. A typical “single screen” TIP device will only have the center position available.

4.3.3 Auxiliary (AUX) Position

The auxiliary positions are typically used for two purposes: local 3-way audio conferencing and audio and video presentation (aka auto-collaborate). These functions are not always used in a call, and hence the availability of these positions dynamically changes within the TIP multiplex. Note that if a TIP device is performing both functions, the auxiliary audio stream is a mix of the presentation audio and the conference sources.

Unlike the physical positions, the auxiliary positions have a fixed mapping from auxiliary transmitter to auxiliary receiver. Additionally MCUs MUST NOT remap the positions of auxiliary video streams as these streams have special properties (such as frame rate) that differentiate this video stream from the other physical position video streams.

A TIP device will always advertise the availability of the auxiliary audio stream. If there is not a conference or presentation source for this audio, no media will be sent. If a source becomes available the TIP device MAY immediately start sending media without any notification.

Changes in the availability of video auxiliary streams will result in the transmission of a REQTOSEND packet informing the remote peer of the change. When a REQTOSEND packet is ACK’ed it indicates that the remote peer is ready to receive the auxiliary vide. If a REQTOSEND packet indicating the desire to transmit is received by an endpoint that is already transmitting, the endpoint SHOULD stop transmission, configure itself to be an auxiliary receiver, and ACK the REQTOSEND. Subsequently should the endpoint receive a REQTOSEND indicating the remote peer no longer wishes to transmit, the endpoint SHOULD send a new REQTOSEND to reassert control of the auxiliary channels. At any time should an endpoint no longer desire to transmit it MUST send a REQTOSEND off packet, this should occur whether the endpoint currently owns the channel or not.

This policy does impose an additional burden on MCUs that support the auxiliary video stream. The MCU needs to mediate the transmission of the auxiliary video stream among the conference participants. Only one requestor's REQTOSEND request will be ACK'ed at any time, all other endpoints will be receivers. When the current presenting endpoint stops presenting using REQTOSEND, the MCU MAY select another endpoint to gain presentation control by sending that endpoint a REQTOSEND off message. The endpoint SHOULD then reassert control of the auxiliary channel. If that does not happen then the MCU MAY select another endpoint.

The AUX FPS bits in the MEDIAOPTS packet control the frame rate of the auxiliary video streams. An endpoint requesting to take control of presentation must send a REQTOSEND packet that provides at least one stream that satisfies the requested auxiliary frame rate. For example, if the MEDIAOPTS frame rate is 5fps, then a REQTOSEND packet only advertising position '5' (which is always 30fps) would be invalid. However, if an endpoint is capable of generating multiple auxiliary frame rates, it may advertise more streams than are necessary to satisfy the MEDIAOPTS frame rate. For example, if the MEDIAOPTS frame is 5fps, then a REQTOSEND packet advertising position '4' and position '5' is valid. In this case the REQTOSEND-ACK packet will inform the endpoint of which streams should be transmitted (one or both).

If a presentation controller receives a MEDIAOPTS update that changes the requested auxiliary frame rate and the current presentation controller offer (i.e. original REQTOSEND) does not satisfy the new requested auxiliary frame rate, then the presentation controller MUST send a new REQTOSEND which satisfies the requested frame rate. For example, assume that the initial frame rate was 30fps and the initial REQTOSEND contained only position '5', and a new MEDIAOPTS is received which changes the frame rate to 5fps. The presentation controller MUST generate a new REQTOSEND that contains position '4' (either in addition to or in lieu of position '5').

4.3.4 Legacy Positions

The legacy video positions (values 9-11) provide TIP devices the ability to transmit lower image resolution versions of the HD video streams. If enabled, they allow an MCU to simultaneously provide both high definition video streams to Telepresence sites and a standard definition streams to devices that do not support HD video.

Only if both peers offer the legacy positions will the streams be transmitted from the TIP endpoint to the MCU. The positional MUX-CSRC of the video legacy stream will have the legacy value for both the transmitter and receiver positions.

The HD video decoders should be capable of handling SD video.

An MCU can transmit a legacy video stream to the center, left, or right position video receivers. The inband SPS/PPS should be inserted immediately before the IDR frame that begins the first frame of the new image size video. Hence, the MCU is able to map the legacy video stream to any available display according to the user experience policies of the product. In this case the positional CSRC of the video stream transmitted from the MCU to the TIP endpoint should have the legacy value for the transmitter position, and the selected physical position (center, left, or right) for the receiver position.

4.4 RTCP Rules

TIP devices operating under TIP will generate RTCP sender and receiver reports (SRs and RRs) in alignment with [4]. The receipt of the sender reports is used by TIP receivers to achieve correct synchronization of the media streams (though a less precise fallback mechanism should be used if no SRs are received).

5 Summary of Changes from TIP 7 to TIP v8

The following list summarizes the changes between version 7 and version 8 of this specification:

- Media Flow Control: Add a new TXFLOWCTRL sub-message allowing control over the H.264 profile configuration of a video encoder per stream. Reference section 4.2.3.
- Media Options: Added a new media option tag allowing a legacy video receiver to set the maximum bitrate to be reserved for the legacy streams. Reference section 4.2.5.15.
- Media Options: Added a new media option allowing the negotiation of constrained video media, “Unrestricted Media Constraints,” used to when a decoder has a couple of prescribed constraints when supporting “unrestricted” video media. Reference section 4.2.5.12.
- Media Options: Added a new media option allowing a sender to indicate that it prefers to use a separate media line and the BFCP protocol for auxiliary video control instead of the multiplexed streams for auxiliary video available in TIP. Devices could then proceed to use a separate BFCP media line if such was successfully negotiated during SIP call establishment. Reference section 4.2.5.13.

6 Backward Compatibility

For implementations that support multiple versions of TIP, the following strategy is recommended. For this backward compatibility to work, it is RECOMMENDED that devices implementing TIP v8 also implement at least the TIP v7 specification.

It is RECOMMENDED that both peers initiate transmission of the MUXCTRL packet at the start of the association of the UDP channel using the highest TIP version supported. If a MUXCTRL packet is received from the remote peer with an earlier version number that is supported by the receiving system, then the receiving device should fallback to the matching version of TIP and restart the TIP negotiation phase. Using this rule, an implementation MUST ignore any MUXCTRL packets with a version number higher than the maximum one supported by the implementation. It is the responsibility of the entity with the more recent implementation to make provisions for interoperability with older implementations. All further use of the protocol within the same session should follow the older version of the protocol.

An alternate strategy is to wait for the remote peer to send a MUXCTRL packet. This strategy is NOT RECOMMENDED as it will result in failure to negotiate if followed by both peers.

6.1 TIP Version Compatibility

The following table outlines the behavior of a TIP system when encountering a TIP device of a lower or higher version. The version is indicated by four bits in the RTCP APP MUXCTRL packet as defined in section 4.2.1, Multiplex Control of the TIP specification [15].

Remote Version	Local Version	Local Behavior
6	6	Proceed with negotiation
6	7	Downgrade to v6 or end TIP negotiation
6	8	Downgrade to v6 or end TIP negotiation
7	6	Wait for remote downgrade
7	7	Proceed with negotiation
7	8	Downgrade to v7 or end call
8	6	Wait for remote downgrade
8	7	Wait for remote downgrade
8	8	Proceed with negotiation

Table 22: TIP Version Compatibility

The decision to proceed, downgrade or end a call is always placed upon the higher TIP version device. A higher TIP version device that can support a lower version indicated by a remote TIP device, should downgrade to that version and proceed with the negotiation. A higher TIP version device that cannot support the lower version of a remote device should end the call or may attempt to complete the call without TIP by proceeding to send only media compatible with what was negotiated in SDP as suggested in section 3.2 of the TIP v7 specification [15].

A lower version TIP device that never receives a compatible version in the negotiation from the remote device should continue to wait for the TIP negotiation timeout, and then drop the call.

6.2 TIP Restricted Media Compatibility

The ability to negotiate restricted vs unrestricted media, such as either for 720p or 1080p resolutions for main video positions, was introduced in TIP version 7 [15].

The table below shows the transmitter and receiver “mode” that can be derived by examining the restricted bits, 12 and 13, in the MEDIAOPTS packet in TIP v7, section 4.2.5. Examining the local MEDIAOPTS TX bit and the remote MEDIAOPTS RX bit derives the transmitter mode. Examining the local RX bit and the remote TX bit derives the receiver mode. This process can be repeated for each resolution. Note that all TIP v6 devices are required to send and receive restricted media. Therefore a TIP v6 device is equivalent to a TIP v7 or 8 device advertising RX=0 and TX=0 for all the restricted resolutions.

A transmitter in unrestricted mode should be sending unrestricted media, but since transmitting restricted is compatible with receiver indicating unrestricted (* foot note), it may optionally send restricted media. A transmitter in restricted mode must send restricted media (i.e. transmitting unrestricted would not be compatible with a restricted receiver). A receiver in unrestricted mode must be able to decode either restricted or unrestricted media. A receiver in restricted mode must be able to decode restricted media.

A device that wishes to be compatible with the broadest possible range of TIP version 7 and 8 deployments, would need to indicate TX=0, RX=1 for all restricted resolutions in MEDIAOPTS.

Transmitter Mode	**TX**	**RX**	**Receiver Mode**
Restricted	0	0	Restricted
Unrestricted ([3])	0	1	Unrestricted
Invalid	1	0	Invalid
Unrestricted ([4])	1	1	Unrestricted

Table 23: Unrestricted Media Negotiation

[3] Transmitter is assumed by its peers to be transmitting Unrestricted media, but it is acceptable for such a transmitter to send Restricted media to a receiver in an Unrestricted mode.

[4] Transmitter is assumed by its peers to be transmitting Unrestricted media, but it is acceptable for such a transmitter to send Restricted media to a receiver in an Unrestricted mode.

6.3 TIP Unrestricted Media Constraints Compatibility

The ability to negotiate constrained media was (proposed) added in TIP version 8.

Unrestricted Media Constraints negotiation is only applied when operating in an unrestricted mode (per section 3 above). The table below shows the transmitter and receiver "mode" that can be derived by examining the constrained bits in the MEDIAOPTS packet. Examining the local MEDIAOPTS TX bit and the remote MEDIAOPTS RX bit derives the transmitter mode. Examining the local RX bit and the remote TX bit derives the receiver mode.

A transmitter in unrestricted mode may send unrestricted, constrained, or restricted media (** foot note). A transmitter in constrained mode may send either constrained or restricted media. A receiver in unrestricted mode must be able to receive unrestricted, constrained, or restricted media. A receiver in constrained mode must be able to receive constrained, unrestricted or restricted media.

A device that wishes to be compatible with all possible TIP version 8 deployments, would advertise TX=1, RX=0 for the constrained bit in MEDIAOPTS.

Transmitter Mode	TX	RX	Receiver Mode
Unconstrained([5])	0	0	Unconstrained
Invalid	0	1	Invalid
Unconstrained([6])	1	0	Unconstrained
Constrained	1	1	Constrained

Table 24: Unrestricted Media Constraints Compatibility

A device that wishes to be compatible with all possible TIP version 7 and 8 deployments, MUST assume that all TIP version 7 devices can only receive Constrained Unrestricted media when they have indicated unrestricted media. Thus, a TIP version 7 device is equivalent to a TIP version 8 device advertising RX=1 and TX=1 for all the Constrained resolutions in MEDIAOPTS per above.

[5] Transmitter is assumed by its peers to be transmitting Unconstrained, Unrestricted media, but it is acceptable for such a transmitter to send Constrained, Unrestricted media or Restricted media to a receiver in an Unconstrained mode.
[6] Transmitter is assumed by its peers to be transmitting Unconstrained, Unrestricted media, but it is acceptable for such a transmitter to send Constrained, Unrestricted media or Restricted media to a receiver in an Unconstrained mode.

7 Call Flow Examples

7.1 Basic TIP MUX Setup

Note that the sequence of MUXCTRL and MEDIAOPTS exchanges is just an example; several other possibilities exist as each peer's exchange is independent. It is required that a successful MUXCTRL occur before any other TIP MUX extension, including MEDIAOPTS, is sent.

```
Peer A                    Peer B
  |                         |
  |        MUXCTRL          |
  |------------------------>|
  |      MUXCTRL ACK        |
  |<------------------------|
  |        MUXCTRL          |
  |<------------------------|
  |      MUXCTRL ACK        |
  |------------------------>|
  |                         |
  |       MEDIAOPTS         |
  |------------------------>|
  |     MEDIAOPTS ACK       |
  |<------------------------|
  |       MEDIAOPTS         |
  |<------------------------|
  |     MEDIAOPTS ACK       |
  |------------------------>|
  |                         |
  |        RTP/RTCP         |
  |<=======================>|
  |                         |
```

7.2 Secure TIP MUX Setup with DTLS

```
Peer A                  Peer B
 |                          |
 |          DTLS            |
 |<========================>|
 | SRTCP: MUXCTRL/MEDIAOPTS |
 |<========================>|
 |        SRTP/SRTCP        |
 |<========================>|
```

7.3 Secure TIP MUX Setup with DTLS and EKT

The following diagram illustrates the differences when a secure channel is used. Note that the MUXCTRL and MEDIAOPTS are carried in SRTCP. Note also that to start transmission of EKT tags as part of SRTCP and SRTP packets, EKT is negotiated in MediaOptions and the SPIMAP packet has successfully been sent by the MCU and received by the endpoint. Note that the SPIMAP packet is transmitted without an EKT tag (including retransmissions) whiles the SPIMAP ACK is transmitted with an EKT tag.

```
Endpoint                         MCU
  |                               |
  |          DTLS                 |
  |<=============================>|
  | SRTCP: MUXCTRL/MEDIAOPTS/EKT  |
  |<=============================>|
  | SRTCP: SPIMAP Packet          |
  |<==============================|
  | SRTCP+EKT tag: SPIMAK ACK     |
  |==============================>|
  |   SRTP+EKT tag/SRTCP+EKT tag  |
  |<=============================>|
```

8 References

The following documents provide the context for the issues discussed in this document:

[1] IETF RFC 3261 “SIP: Session Initiation Protocol”
[2] IETF RFC 3264 “An Offer/Answer Model with the Session Description Protocol (SDP)”
[3] IETF RFC 2327 “SDP: Session Description Protocol”
[4] IETF RFC 3550 “RTP: A Transport Protocol for Real-Time Applications”
[5] IETF RFC 3551 “RTP Profile for Audio and Video Conferences with Minimal Control”
[6] IETF RFC 4961 “Symmetric RTP/RTP Control Protocol (RTCP)”
[7] IETF RFC 3984 “RTP Payload Format for H.264 Video”
[8] IETF RFC 4585 “Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)”
[9] IETF RFC 5761 “Multiplexing RTP Data and Control Packets on a Single Port”
[10] Telepresence Interoperability Protocol (TIP) Version 6
[11] IETF RFC 3711 “Secure Real-time Transport Protocol (SRTP)”, includes SAVP
[12] [EKT] Internet-Draft “Encrypted Key Transport for Secure RTP” version 03 http://tools.ietf.org/html/draft-ietf-avt-srtp-ekt-03
[13] IETF RFC 4347 “Datagram Transport Layer Security”
[14] IETF RFC 5764 “Datagram Transport Layer Security (DTLS) Extension to Establish Keys for the Secure Real-time Transport Protocol (SRTP)”
[15] Telepresence Interoperability Protocol (TIP) Version 7
[16] IETF Internet-Draft “Revision of the Binary Floor Control Protocol (BFCP)” http://tools.ietf.org/html/draft-ietf-bfcpbis-rfc4582bis

End of Document