

Network Working Group
Internet-Draft
Expires: August 11, 2007

C. Boulton
Ubiquity Software Corporation
T. Melanchuk
BlankSpace
S. McGlashan
Hewlett-Packard
A. Shiratzky
Radvision
February 7, 2007

**A Control Framework for the Session Initiation Protocol (SIP)
draft-boulton-sip-control-framework-05**

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on August 11, 2007.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

This document describes a Framework and protocol for application deployment where the application logic and processing are distributed. The framework uses the Session Initiation Protocol

(SIP) to establish an application-level control mechanism between application servers and associated external servers such as media servers.

The motivation for the creation of this Framework is to provide an interface suitable to meet the requirements of a distributed, centralized conference system, as defined by the XCON work group of the IETF. It is not, however, limited to this scope and it is envisioned that this generic Framework will be used for a wide variety of de-coupled control architectures between network entities.

Table of Contents

1.	Introduction	4
2.	Conventions and Terminology	4
3.	Overview	6
4.	Locating External Server Resources	10
5.	Control Client SIP UAC Behavior - Control Channel Setup . . .	10
5.1.	Control Client SIP UAC Behavior - Media Dialogs	12
6.	Control Server SIP UAS Behavior - Control Channel Setup . . .	12
7.	Control Channel Keep-Alive	14
7.1.	Timeout Negotiation	14
7.2.	Generating Keep-Alive Messages	15
8.	Control Framework Interactions	16
8.1.	Constructing Requests	17
8.1.1.	Sending CONTROL	18
8.1.2.	Sending REPORT	18
8.2.	Constructing Responses	20
9.	Response Code Descriptions	21
9.1.	200 Response Code	21
9.2.	202 Response Code	21
9.3.	400 Response Code	21
9.4.	403 Response Code	21
9.5.	481 Response Code	21
9.6.	500 Response Code	22
10.	Control Packages	22
10.1.	Control Package Name	22
10.2.	Framework Message Usage	22
10.3.	Common XML Support	23
10.4.	CONTROL Message Bodies	23
10.5.	REPORT Message Bodies	23
10.5.1.	Events	23
10.6.	Examples	24
11.	Network Address Translation (NAT)	24
12.	Formal Syntax	24
12.1.	SIP Formal Syntax	24
12.2.	Control Framework Formal Syntax	24

13.	Examples	27
14.	Security Considerations	32
15.	IANA Considerations	32
15.1.	IANA Registration of the 'escs' Option Tag	32
15.2.	Control Package Registration Information	32
15.2.1.	Control Package Registration Template	32
15.3.	SDP Transport Protocol	32
15.3.1.	TCP/ESCS	32
15.3.2.	TCP/TLS/ESCS	32
15.4.	SDP Attribute Names	32
15.5.	SIP Response Codes	32
16.	Acknowledgments	32
17.	Appendix A	32
17.1.	Common Dialog/Multiparty Reference Schema	32
18.	References	34
18.1.	Normative References	34
18.2.	Informative References	34
	Authors' Addresses	36
	Intellectual Property and Copyright Statements	37

1. Introduction

Applications are often developed using an architecture where the application logic and processing activities are distributed. Commonly, the application logic runs on "application servers" whilst the processing runs on external servers, such as "media servers". This document focuses on the framework and protocol between the application server and external processing server. The motivation for this framework comes from a set of requirements for Media Server Control, which can be found in the 'Media Control Protocol Framework' document[8]. While the Framework is not media server control specific, it is the primary driver and use case for this work. It is intended that the framework contained in this document will be used for a plethora of appropriate device control scenarios.

This document does not define a SIP based extension that can be used directly for the control of external components. The framework mechanism must be extended by other documents that are known as "Control Packages". A comprehensive set of guidelines for creating "Control Packages" is described in [Section 10](#).

Current IETF transport device control protocols, such as megaco [7], while excellent for controlling media gateways that bridge separate networks, are troublesome for supporting media-rich applications in SIP networks, because they duplicate many of the functions inherent in SIP. Rather than relying on single protocol session establishment, application developers need to translate between two separate mechanisms.

Application servers traditionally use SIP third party call control [RFC 3725](#) [12] to establish media sessions from SIP user agents to a media server. SIP, as defined in [RFC 3261](#) [2], also provides the ideal rendezvous mechanism for establishing and maintaining control connections to external server components. The control connections can then be used to exchange explicit command/response interactions that allow for media control and associated command response results.

2. Conventions and Terminology

In this document, [BCP 14](#)/RFC 2119 [1] defines the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL". In addition, [BCP 15](#) indicates requirement levels for compliant implementations.

The following additional terms are defined for use in this document:

B2BUA: A B2BUA is a Back-to-Back SIP User Agent.

Control Server: A Control Server is an entity that performs a service, such as media processing, on behalf of a Control Client. For example, a media server offers mixing, announcement, tone detection and generation, and play and record services. The Control Server in this case, has a direct RTP [[15](#)] relationship with the source or sink of the media flow. In this document, we often refer to the Control Server simply as "the Server".

Control Client: A Control Client is an entity that requests processing from a Control Server. Note that the Control Client may not have any processing capabilities whatsoever. For example, the Control Client may be an Application Server (B2BUA) or other endpoint requesting manipulation of a third-party's media stream, that terminates on a media server acting in the role of a Control Server. In this document, we often refer to the Control Client simply as "the Client".

Control Channel: A Control Channel is a reliable connection between a Client and Server that is used to exchange Framework messages. The term "Connection" is used synonymously within this document.

Framework Message: A Framework Message is a message on a Control Channel that has a type corresponding to one of the Methods defined in this document. A Framework message is often referred to by its method, such as a "CONTROL message".

Method: A Method is the type of a framework message. Three Methods are defined in this document: SYNCH, CONTROL, REPORT, and K-ALIVE.

Control Command: A Control Command is an application level request from a Client to a Server. Control Commands are carried in the body of CONTROL messages. Control Commands are defined in separate specifications known as "Control Packages".

framework transaction: A framework transaction is defined as a sequence composed of a control framework message originated by either a Control Client or Control Server and responded to with a control Framework response code message. Note that the control framework has no "provisional" responses. A control framework transaction MUST complete within Transaction-Timeout time.

extended framework transaction: An extended framework transaction is used to extend the lifetime of a CONTROL method transaction when the Control Command it carries cannot be completed within Command-Timeout milliseconds. A Server extends the lifetime of a CONTROL method transaction by sending a 202 response code followed by one or more REPORT transactions as specified in [Section 8.1.2](#). Extended framework transactions allow command failures to be discovered at the transaction layer.

Transaction-Timeout: the maximum allowed time between a control Client or Server issuing a framework message and receiving a corresponding response. The value for the timeout should be based on a multiple of the network RTT plus Command-Timeout milliseconds to allow for message parsing and processing.

3. Overview

This document details mechanisms for establishing, using, and terminating a reliable channel using SIP for the purpose of controlling an external server. The following text provides a non-normative overview of the mechanisms used. Detailed, normative guidelines are provided later in the document.

Control channels are negotiated using standard SIP mechanisms that would be used in a similar manner to creating a SIP voice session. Figure 1 illustrates a simplified view of the proposed mechanism. It highlights a separation of the SIP signaling traffic and the associated control channel that is established as a result of the SIP interactions.

The use of SIP for the specified mechanism provides many inherent capabilities which include:-

- o Service location - Use SIP Proxies or Back-to-Back User Agents for discovering Control Servers.
- o Security mechanisms - Leverage established security mechanisms such as Transport Layer Security (TLS) and Client Authentication.
- o Connection maintenance - The ability to re-negotiate a connection, ensure it is active, audit parameters, and so forth.
- o Agnostic - Generic protocol allows for easy extension.

As mentioned in the previous list, one of the main benefits of using SIP as the session control protocol is the "Service Location" facilities provided. This applies at both a routing level, where [RFC 3263](#) [4] provides the physical location of devices, and at the Service level, using Caller Preferences[13] and Callee Capabilities[14]. The ability to select a Control Server based on Service level capabilities is extremely powerful when considering a distributed, clustered architecture containing varying services (for example Voice, Video, IM). More detail on locating Control Server resources using these techniques is outlined in [Section 5](#) of this document.

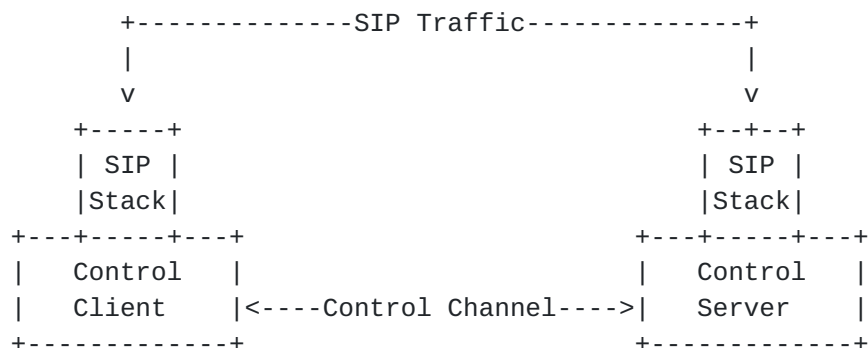


Figure 1: Basic Architecture

The example from Figure 1 conveys a 1:1 connection between the Control Client and the Control Server. It is possible, if required, for multiple control channels using separate SIP dialogs to be established between the Control Client and the Control Server entities. Any of the connections created between the two entities can then be used for Server control interactions. The control connections are agnostic to any media sessions. Specific media session information can be incorporated in control interaction commands (which themselves are defined in external packages) using the XML schema defined in [Section 17](#). The ability to have multiple control channels allows for stronger redundancy and the ability to manage high volumes of traffic in busy systems.

[Editors Note: Still under discussion. How does an app server know, when there are multiple external servers, which specific server has any given media session? Next version of the draft will discuss the correlation procedures. The App server needs a control channel with the media server and needs to know which channel to use once the media session has been established. Sounds like a GRUU usage?]

Consider the following simple example for session establishment between a Client and a Server (Note: Some lines in the examples are removed for clarity and brevity). Note that the roles discussed are logical and can change during a session, if the Control Package allows.

The Client constructs and sends a SIP INVITE request to the external Server. The request contains the SIP option tag "escs" in a SIP "Require" header for the purpose of forcing the use of the mechanism described in this document. The SDP payload includes the required information for control channel negotiation. The COMEDIA [\[6\]](#) specification for setting up and maintaining reliable connections is used (more detail available in later sections).

The client MUST include details of control packages that are supported and, more specifically, that will be used within the control channel created. This is achieved through the inclusion of a SIP "Control-Packages" header. The "Control-Packages" header is defined and described later in this document.

Client Sends to External Server:


```
INVITE sip:External-Server@example.com SIP/2.0
To: <sip:External-Server@example.com>
From: <sip:Client@example.com>;tag=64823746
Require: escs
Control-Packages: <example-package>
Call-ID: 7823987HJHG6
Content-Type: application/sdp
```

```
v=0
o=originator 2890844526 2890842808 IN IP4 controller.example.com
s=-
c=IN IP4 controller.example.com
m=application 7575 TCP/ESCS
a=setup:active
a=connection:new
```

On receiving the INVITE request, the external Server supporting this mechanism generates a 200 OK response containing appropriate SDP.

External Server Sends to Client:

```
SIP/2.0 200 OK
To: <sip:External-Server@example.com>;tag=28943879
From: <sip:Client@example.com>;tag=64823746
Call-ID: 7823987HJHG6
Content-Type: application/sdp
```

```
v=0
o=originator 2890844526 2890842808 IN IP4 controller.example.com
s=-
c=IN IP4 mserver.example.com
m=application 7563 TCP/ESCS
a=setup:passive
a=connection:new
```

The Control Client receives the SIP 200 OK response and extracts the relevant information (also sending a SIP ACK). It creates an outgoing (as specified by the SDP 'setup:' attribute) TCP connection to the Control Server. The connection address (taken from 'c=') and port (taken from 'm=') are used to identify the remote part in the new connection.

Once established, the newly created connection can be used to exchange control language requests and responses. If required, after the control channel has been setup, media sessions can be established

using standard SIP third party call control.

[Editors Note: See previous note: this is where we may need to mention how an App Server knows which external Server is responsible for any given media session.]

Figure 4 provides a simplified example where the proposed framework is used to control a User Agent's RTP session. (1) in brackets represents the SIP dialog and dedicated control channel previously described in this overview section.

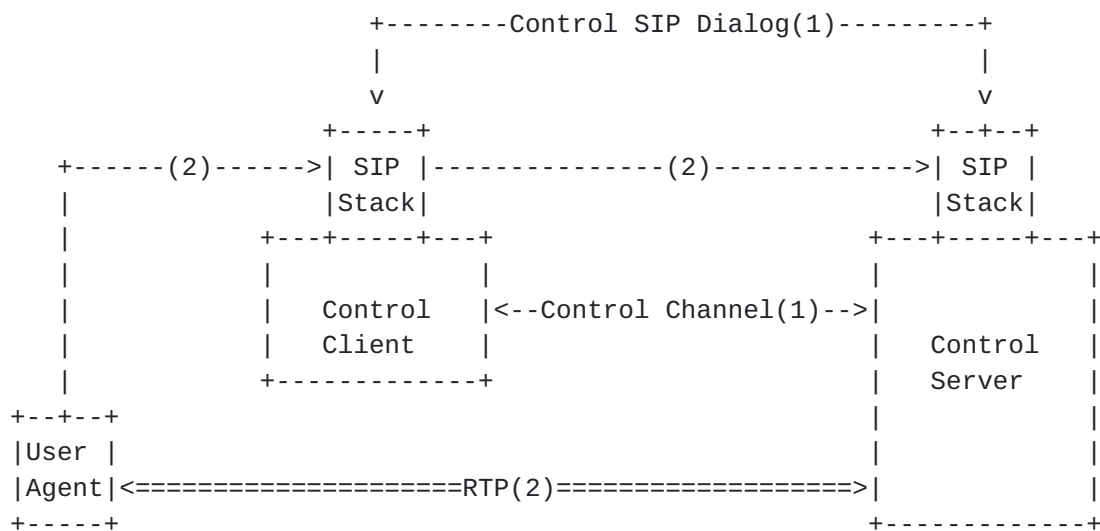


Figure 4: Participant Architecture

(2) from Figure 4 represents the User Agent SIP dialog interactions and associated media flow. A User Agent would create a SIP dialog with the Control Client entity. The Control Client entity will also create a related dialog to the Control Server (B2BUA type functionality). Using the interaction illustrated by (2), the User Agent is able to negotiate media capabilities with the Control Server using standard SIP mechanisms as defined in [RFC 3261](#) [2] and [RFC 3264](#) [5].

If not present in the SDP received by the Control Client from the User Agent(2), a media label SDP attribute, which is defined in [10], should be inserted for every media description (identified as m= line as defined in [9]). This provides flexibility for the Control Client, because it can generate control messages that specify a particular Media stream (between User Agent and Control Server) within a SIP media dialog. If a Media label is not included in the control message, it applies to all media associated with the dialog.

A non 2xx class SIP response received for the INVITE request indicates that no SIP dialog has been created and is treated as specified [RFC 3261](#) [2]. One exception to this is the "496" (TODO: need to pick an appropriate response code) response code whose operation is defined in [Section 6](#).

4. Locating External Server Resources

Section will describe mechanisms for locating an external Server.

5. Control Client SIP UAC Behavior - Control Channel Setup

On creating a new SIP INVITE request, a UAC can insist on using the mechanisms defined in this document. This is achieved by inserting a SIP Require header containing the option tag 'escs'. A SIP Require header with the value 'escs' MUST NOT be present in any other SIP request type.

If on creating a new SIP INVITE request, a UAC does not want to insist on the usage of the mechanisms defined in this document but merely that it supports them, a SIP Supported header MUST be included in the request with the option tag 'escs'.

The SIP INVITE MUST include a SIP "Control-Packages" header which MUST contain at least one valid entry and can contain multiple control packages if required.

If a reliable response is received (as defined [RFC 3261](#) [2] and [RFC 3262](#) [3]) that contains a SIP Require header containing the option tag 'escs', the mechanisms defined in this document are applicable to the newly created dialog.

Before the UAC can send a request, it MUST include a valid session description using the Session Description Protocol defined in [9]. The following information defines the composition of some specific elements of the SDP payload that MUST be adhered to for compliancy to this specification.

The Connection Data line in the SDP payload is constructed as specified in [9]:

```
c=<nettype> <addrtype> <connection-address>
```

The first sub-field, <nettype>, MUST equal the value "IN". The second sub-field, <addrtype>, MUST equal either "IP4" or "IP6". The third sub-field for Connection Data is <connection-address>. This

supplies a representation of the SDP originators address, for example dns/IP representation. The address will be the network address used for connections in this specification.

Example:

```
c=IN IP4 controller.example.com
```

The SDP MUST contain a corresponding Media Description entry for compliance to this specification:

```
m=<media> <port> <proto>
```

The first "sub-field" <media> MUST equal the value "application". The second sub-field, <port>, MUST represent a port on which the constructing client can receive an incoming connection if required. The port is used in combination with the address specified in the 'Connection Data line defined previously to supply connection details. If the constructing client can't receive incoming connections it MUST still enter a valid port range entry. The use of the port value '0' has the same meaning as defined in the SDP specification[9]. The third sub-field, <proto>, MUST equal the value "TCP/ESCS" as defined in [Section 15.3.2](#) of this document.

[Editors note: Need to cover other protocols so not TCP specific]

The SDP MUST also contain a number of SDP media attributes(a=) that are specifically defined in the COMEDIA specification. The attributes provide connection negotiation and maintenance parameters. A client conforming to this specification SHOULD support all the possible values defined for media attributes from the COMEDIA [6] specification but MAY choose not to support values if it can definitely determine they will never be used (for example will only ever initiate outgoing connections). It is RECOMMENDED that a Controlling UAC initiate a connection to an external Server but that an external Server MAY negotiate and initiate a connection using COMEDIA, if network topology prohibits initiating connections in a certain direction. An example of the attributes is:

```
a=setup:active  
a=connection:new
```

This example demonstrates a new connection that will be initiated from the owner of the SDP payload. The connection details are contained in the SDP answer received from the UAS. A full example of an SDP payload compliant to this specification can be viewed in

[Section 3](#). Once the SDP has been constructed along with the remainder of the SIP INVITE request (as defined in [RFC 3261](#) [2]), it can be sent to the appropriate location. The SIP dialog and appropriate control connection is then established.

5.1. Control Client SIP UAC Behavior - Media Dialogs

It is intended that the Control framework will be used within a variety of architectures for a wide range of functions. One of the primary functions will be the use of the control channel to apply specific Control package commands to co-existing SIP dialogs that have been established with the same remote server, for example the manipulation of audio dialogs connected to a media server.

Such co-existing dialogs will pass through the Control Client (see Figure 4) entity and may contain more than one Media Description (as defined by "m=" in the SDP). The Control Client SHOULD include a media label attribute (B2BUA functionality), as defined in [10], for each "m=" definition. A Control Client constructing the SDP MAY choose not to include the media label SDP attribute if it does not require direct control on a per media stream basis.

This framework identifies the common re-use of referencing media dialogs and has specified a connection reference attribute that can optionally be imported into any Control Package. It is intended that this will reduce repetitive specifying of dialog reference language. The schema can be found in Section 17.1 in [Appendix A](#).

Similarly, the ability to identify and apply commands to a group of media dialogs is also identified as a common structure that could be defined and re-used (for example playing a prompt to all participants in a Conference). The schema for such operations can also be found in Section 17.1 in [Appendix A](#).

Support for both the common attributes described here is specified as part of each Control Package definition, as detailed in [Section 10](#).

6. Control Server SIP UAS Behavior - Control Channel Setup

On receiving a SIP INVITE request, an external Server(UAS) inspects the message for indications of support for the mechanisms defined in this specification. This is achieved through the presence of the SIP Supported and Require headers containing the option tag 'escs'. If the external Server wishes to construct a reliable response that conveys support for the extension, it should follow the mechanisms defined in [RFC 3261](#) [2] for responding to SIP supported and Require headers. If support is conveyed in a reliable SIP provisional

response, the mechanisms in [RFC 3262](#) [3] MUST also be used.

When constructing a SIP success response, the SDP payload MUST be constructed using the semantics(Connection, Media and attribute) defined in [Section 5](#) using valid local settings and also with full compliance to the COMEDIA[6] specification. For example, the SDP attributes included in the answer constructed for the example offer provided in [Section 5](#) would look as illustrated below:

```
a=setup:passive
a=connection:new
```

Once the SIP success response has been constructed, it is sent using standard SIP mechanisms. Depending on the contents of the SDP payloads that were negotiated using the Offer/Answer exchange, a reliable connection will be established between the Controlling UAC and external Server UAS entities. The connection is now available to exchange commands, as defined in "Control Packages" and described in [Section 10](#). The state of the SIP Dialog and the associated Control channel are now explicitly linked. If either party wishes to terminate a Control channel is simply issues a SIP termination request (SIP BYE request). The Control Channel therefore lives for the duration of the SIP dialog.

If the UAS does not support the extension contained in a SIP Supported or Require header it MUST respond as detailed in [RFC 3261](#) [2]. If the UAS does support the SIP extension contained in a SIP Require or Supported header but does not support one or more of the Control packages, as represented in the "Control-Packages" SIP header, it MUST respond with a SIP "496 Unknown Control Package" response code. The error response MUST conform to [RFC 3261](#) [2] and MUST also include a "Control-Packages" SIP header which lists the control packages from the request that the UAS does not support. This provides the Controlling UAC with an explicit reason for failure and allows for re-submission of the request without the un-supported control package.

A SIP entity receiving a SIP OPTIONS request MUST respond appropriately as defined in [RFC 3261](#) [2]. This involves providing information relating to supported SIP extensions in the 'Supported' message header. For this extension a value of 'escs' MUST be included. Additionally, a SIP entity MUST include all the additional control packages that are associated with the Control channel. This is achieved by including a 'Control-Packages' SIP message header listing all relevant supported Control package tokens.

7. Control Channel Keep-Alive

It is reasonable to expect this document to be used in various network architectures. This will include a wide range of deployments where the clients could be co-located in a secured, private domain or spread across disparate domains that require traversal of devices such as Network Address Translators (NAT) and Firewalls (which is covered in [Section 11](#)). It is important, therefore, that this document provides a 'keep-alive' mechanism that enables the control channel being created to firstly be kept active during times of inactivity (most Firewalls have a timeout period after which connections are closed) and also provide the ability for application level failure detection. It should be noted at this point that the following procedures apply explicitly to the control channel being created and for details relating to a SIP keep-alive mechanism implementers should seek guidance from SIP Outbound [[11](#)]. The following 'keep-alive' procedures SHOULD be implemented by all entities but MAY NOT be implemented if it can be guaranteed that deployments will only occur with entities in a co-located domain. It should be noted that choosing to not implement the 'keep-alive' mechanism in this section, even when in a co-located architecture, will reduce the ability to detect application level errors - especially during long periods of in-activity.

7.1. Timeout Negotiation

During the SIP dialog negotiation the clients will also negotiate a timeout period for the control channel 'keep-alive' mechanism. The following rules MUST be obeyed in conjunction with COMEDIA[6]:

- o If the Client initiating the SIP INVITE request has a COMEDIA 'setup' attribute equal to 'active', the 'k-alive' 'Control-Packages' SIP header parameter MUST be included. The value of the 'k-alive' header parameter SHOULD be in the range of 95 and 120 seconds (this is consistent with SIP Outbound[11]). The client generating the subsequent answer ('passive' client) MUST copy the 'k-alive' 'Control-Packages' header parameter into the response with the same value.
- o If the Client initiating the SIP INVITE request has a COMEDIA 'setup' attribute equal to 'passive', the 'k-alive' Control-Packages SIP header parameter MUST NOT be included. The client generating the subsequent answer ('active' client) MUST include the 'k-alive' header parameter. The value of the 'k-alive' header parameter SHOULD be in the range of 95 and 120 seconds.
- o If the Client initiating the SIP INVITE request has a COMEDIA 'setup' attribute equal to 'actpass', the 'k-alive' 'Control-Packages' header parameter MUST be included. The value of the 'k-alive' header parameter SHOULD be in the range of 95 to 120 seconds. If the client generating the subsequent answer places a

value of 'active' in the COMEDIA 'setup' attribute, It MUST include a 'k-alive' header parameter and MAY choose a value that is different from that received in the request. The value SHOULD be in the range 95 to 120 seconds. If the client generating the subsequent answer places a value of 'passive' in the COMEDIA 'setup' attribute, it MUST copy the 'k-alive' header and value from the request into the 'k-alive' 'Control-Packages' SIP header parameter.

- o The 'Control-Packages' 'k-alive' header parameter MUST not be included when the COMEDIA 'setup' attribute is equal to 'holdconn'.
- o Following the previous steps ensures that the entity initiating the control channel connection is always the one specifying the keep-alive timeout period. It will always be the initiator of the connection who generates the 'K-ALIVE' Control Framework level messages. The following section describes in more detail how to generate the Control Framework 'K-ALIVE' message.

7.2. Generating Keep-Alive Messages

Once the SIP dialog has been established using the offer answer mechanism and the control channel has been established (including the initial identity handshake using SYNCH as discussed in [Section 8](#)), both the 'active' and 'passive' (as defined in COMEDIA[6]) clients MUST start a keep-alive timer equal to the value negotiated during the offer/answer exchange (from the 'k-alive' 'Control-Packages' SIP header parameter).

When acting as an 'active' client, a 'K-ALIVE' Control Framework message MUST be generated before the local 'keep-alive' timer fires. An active client is free to send the K-ALIVE Control Framework message when ever it chooses. A guideline of 80% of the local 'keep-alive' timer is suggested. The 'passive' client MUST generate a 200 OK Control Framework response to the K-ALIVE message and reset the local 'keep-alive' timer. No other Control Framework response is valid. On receiving the 200 OK Control Framework message, the 'active' client MUST reset the local 'keep-alive' timer. If no 200 OK response is received to the K-ALIVE Control Framework message, before the local 'keep-alive' timer fires, the 'active' client SHOULD tear down the SIP dialog and recover the associated control channel resources. The 'active' client MAY choose to try and recover the connection by renegotiation using COMEDIA.

When acting as a 'passive' client, a 'K-ALIVE' Control Framework message MUST be received before the local 'keep-alive' timer fires. The 'passive' client MUST generate a 200 OK control framework response to the K-ALIVE Control Framework message. On sending the 200 OK response, the 'passive' client MUST reset the local 'keep-

alive' timer. If no K-ALIVE message is received before the local 'keep-alive' timer fires, the 'passive' client SHOULD tear down the SIP dialog and recover the associated control channel resources. The 'active' client MAY try to and recover the connection by renegotiating using COMEDIA.

8. Control Framework Interactions

The use of the COMEDIA specification in this document allows for a Control Channel to be set up in either direction as a result of the SIP INVITE transaction. While providing a flexible negotiation mechanism, it does provide certain correlation problems between the channel and the overlying SIP dialog. Remember that the two are implicitly linked and so need a robust correlation mechanism. A Control Client receiving an incoming connection (whether it be acting in the role of UAC or UAS) has no way of identifying the associated SIP dialog as it could be simply listening for all incoming connections on a specific port. As a consequence, some rules are applied to allow a connecting (defined as 'active' role in COMEDIA) client to identify the associated SIP dialog that triggered the connection. The following steps provide an identification mechanism that MUST be carried out before any other signaling is carried out on the newly created Control channel.

- o Once connected, the client initiating the connection (as determined by COMEDIA) MUST immediately send a Control Framework SYNCH request. The SYNCH request will be constructed as defined in [Section 12.2](#) and MUST only contain one message header, 'dialog-id', which contains the SIP dialog information.
- o The 'dialog-id' message header is constructed by concatenating the Local-tag, Call-ID and Remote-tag (as defined in [Section 12.2](#)) from the SIP dialog and separating with a '~'. See syntax defined in Section 17.1 in [Appendix A](#) and examples in [Section 10.6](#). For example, if the SIP dialog had values of 'Local-tag=HKJDH', 'Remote-tag=JJSUSHJ' and 'Call-ID=8shKUHSUKHW@example.com' - the 'dialog-id' header would look like this:
'dialog-id=HKJDH-8shKUHSUKHW@example.com~JJSUSHJ'.
- o The client who initiated the connection MUST then send the SYNCH request. It should then wait for a period of 5 seconds to receive a response. It MAY choose a longer time to wait but it should not be shorter than 5 seconds.
- o If no response is received for the SYNCH control message, a timeout occurs and the control channel is terminated along with the associated SIP dialog (issue a BYE request).
- o If the client who initiated a connection receives a 481 response, this implies that the SYNCH request was received but no associated SIP dialog exists. This also results in the control channel being terminated along with the associated SIP dialog (issue a BYE

- request).
- o All other error responses received for the SYNCH request are treated as detailed in this specification and also result in the termination of the control channel and the associated SIP dialog (issue a BYE request).
 - o The receipt of a 200 response to a SYNCH message implies that the SIP dialog and control connection have been successfully correlated. The control channel can now be used for further interactions.

Once a successful control channel has been established, as defined in [Section 5](#) and [Section 6](#) (and the connection has been correlated, as described in previous paragraph), the two entities are now in a position to exchange relevant control framework messages. The remainder of this section provides details of the core set of methods and responses that MUST be supported for the core control framework. Future extensions to this document MAY define new methods and responses.

8.1. Constructing Requests

An entity acting as a Control Client is now able to construct and send new requests on a control channel and MUST adhere to the syntax defined in [Section 12](#). Control Commands MUST also adhere to the syntax defined by the Control Packages negotiated in [Section 5](#) and [Section 6](#) of this document. A Control Client MUST create a unique transaction and associated identifier per request transaction. The transaction identifier is then included in the first line of a control framework message along with the method type (as defined in the ABNF in [Section 12](#)). The first line starts with the SCFW token for the purpose of easily extracting the transaction identifier. The transaction identifier MUST be globally unique over space and time. All required mandatory and optional control framework headers are then inserted into the control message with appropriate values (see relevant individual header information for explicit detail). A "Control-Package" header MUST also be inserted with the value indicating the Control Package to which this specific request applies (Multiple packages can be negotiated per control channel).

Any framework message that contains an associated payload MUST also include a 'Content-Length' message header which represents the size of the message body in decimal number of octets. If no associated payload is to be added to the message, a 'Content-Length' header with a value of '0' MUST be included.

When all of the headers have been included in the framework message, it is sent down the control channel established in [Section 5](#).

It is a requirement that a Server receiving such a request respond quickly with an appropriate response (as defined in [Section 8.2](#)). A Control Client entity needs to wait for Transaction-Time time for a response before considering the transaction a failure.

[8.1.1.1.](#) Sending CONTROL

A 'CONTROL' message is used by Control Client to invoke control commands on a Control Server. The message is constructed in the same way as any standard Control Framework message, as discussed previously in [Section 8.1](#) and defined in [Section 12](#). A CONTROL message MAY contain a message body. The explicit control command(s) of the message payload contained in a CONTROL message are specified in separate Control Package specifications. These specifications MUST conform to the format defined in [Section 10.4](#).

[8.1.1.2.](#) Sending REPORT

A 'REPORT' message is used by a Control Server in two situations. The first situation occurs when processing of a Control Command extends beyond a Command-Timeout. In this case a 202 response is returned. Status updates and the final results of the command are then returned in subsequent REPORT messages. The second situation allows REPORT to be used as an event notification mechanism where events are correlated with the original CONTROL message. In this case, REPORT messages may be sent after the original transaction or extended transaction has completed.

All REPORT messages MUST contain the same transaction ID in the request start line that was present in the original CONTROL transaction. This allows both extended transactions and event notifications to be correlated with the original CONTROL transaction.

[8.1.2.1.](#) Reporting the Status of Extended Transactions

On receiving a CONTROL message, a Control Server MUST respond within Command-Timeout with a status code for the request, as specified in [Section 8.2](#). If the command completed within that time, a 200 response code would have been sent. If the command did not complete within that time, the response code 202 would have been sent indicating that the requested command is still being processed and the CONTROL transaction is being extended. The REPORT method is used to update the status of the extended transaction.

A Control Server issuing a 202 response MUST immediately issue a REPORT message. The initial REPORT message MUST contain a 'Seq' (Sequence) message header with a value equal to '1' (It should be noted that the 'Seq' numbers at both Control Client Control Server

for framework messages are independent). The initial REPORT message MUST also contain a 'Status' message header with a value of 'pending'. This initial REPORT message MUST NOT contain a message body and is primarily used to establish a subsequent message transaction based on the initial CONTROL message.

All REPORT messages for an extended CONTROL transaction MUST contain a 'Timeout' message header. This header will contain a value in delta seconds that represents the amount of time the recipient of the REPORT message must wait before assuming that there has been a problem and terminating the extended transaction and associated state. On receiving a REPORT message with a 'Status' header of 'pending' or 'update', the Control Client MUST reset the counter for the associated extended CONTROL transaction to the indicated timeout period. If the timeout period approaches with no intended REPORT messages being generated, the entity acting as a Control Framework UAS for the interaction MUST generate a REPORT message containing, as defined in this paragraph, a 'Status' header of 'pending'. Such a message acts as a timeout refresh and in no way impacts the extended transaction, because no message body or semantics are permitted. It is RECOMMENDED that a minimum value of 10 and a maximum of ?? is used for the value of the 'Timeout' message header. It is also RECOMMENDED that a Control Server refresh the timeout period of the CONTROL transaction at an interval that is not too close to the expiry time. A value of 80% of the timeout period could be used, for example a timeout period of 10 seconds would be refreshed after 8 seconds.

Subsequent REPORT messages that provide additional information relating to the extended CONTROL transaction MUST also include and increment by 1 the 'Seq' header value. They MUST also include a 'Status' header with a value of 'update'. These REPORT messages sent to update the extended CONTROL transaction status MAY contain a message body, as defined by individual Control Packages and specified in [Section 9.5](#). A REPORT message sent updating the extended transaction also acts as a timeout refresh, as described earlier in this section. This will result in a transaction timeout period at the initiator of the request being reset to the interval contained in the 'Timeout' message header.

When all processing for an extended CONTROL transaction has taken place, the entity acting as a Control Server MUST send a terminating REPORT message. The terminating REPORT message MUST increment the value in the 'Seq' message header by the value of '1' from the previous REPORT message. It MUST also include a 'Status' header with a value of 'terminate' and MAY contain a message body. A Control Framework UAC can then clean up any pending state associated with the original control transaction.

8.1.2.2. Reporting Asynchronous Events

Commands that are carried in CONTROL messages can request that the Server notify the Client about events that occur sometime in the future. It is not desirable to use extended Control transactions for these types of commands for two reasons. First, an event never occurring is often correct behavior. Second, events may occur long after the original request for their notification.

REPORT messages can be used to notify events. REPORT messages that notify events MUST contain a 'Status' header of 'Notify'. They MUST NOT contain either a 'Timeout' or 'Seq' header and any such headers MUST be ignored when the REPORT message has a 'Status' of 'notify'. The REPORT message MAY contain a message body.

8.2. Constructing Responses

A Control Client or Server, on receiving a request, MUST generate a response within Command-Time time. The response MUST conform to the ABNF defined in [Section 12](#). The first line of the response MUST contain the transaction identifier used in first line of the request, as defined in [Section 8.1](#). Responses MUST NOT include the 'Status' or 'Timeout' message headers - if they are included they have no meaning or semantics.

A Control Client or Server MUST then include a status code in the first line of the constructed response. A CONTROL request that has been understood, and either the relevant actions for the control command have completed or a control command error is detected, uses the 200 status code as defined in [Section 9.1](#). A 200 response MAY include message bodies. If a 200 response does contain a payload it MUST include a Content-Length header. A 200 is the only response defined in this specification that allows a message body to be included. A client receiving a 200 class response then considers the control command completed. A CONTROL request that is received and understood but requires processing that extends beyond Command-Time time will return a 202 status code in the response. This will be followed immediately by an initial REPORT message as defined in [Section 8.1.2](#). A Control Package SHOULD explicitly define the circumstances under which either 200 or 202 with subsequent processing takes place.

If a Control Client or Server encounters problems with either a REPORT or CONTROL request, an appropriate error code should be used in the response, as listed in [Section 9](#). The generation of a non 2xx class response code to either a CONTROL or REPORT message will result in failure of the transaction, and all associated state and resources should be terminated. The response code may provide an explicit

indication of why the transaction failed, which might result in a re-submission of the request.

[timm]: how can an error response provide an explicit indication of the reason for the transaction failure when only a 200 response allows message bodies?

9. Response Code Descriptions

The following response codes are defined for transaction responses to methods defined in [Section 8.1](#). All response codes in this section MUST be supported and can be used in response to both CONTROL and REPORT messages except that a 202 MUST NOT be generated in response to a REPORT message.

Note that these response codes apply to framework transactions only. Success or error indications for control commands MUST be treated as the result of a control command and returned in either a 200 response or REPORT message.

9.1. 200 Response Code

The 200 code indicates the completion of a successful transaction.

9.2. 202 Response Code

The 202 response code indicates the completion of a successful transaction with additional information to be provided at a later time through the REPORT mechanism defined in [Section 8.1.2](#).

9.3. 400 Response Code

The 400 response indicates that the request was syntactically incorrect.

9.4. 403 Response Code

The 400 response indicates that the requested operation is illegal.

9.5. 481 Response Code

The 481 response indicates that the intended target of the request does not exist.

9.6. 500 Response Code

The 500 response indicates that the recipient does not understand the request

10. Control Packages

"Control Packages" are intended to specify behavior that extends the the capability defined in this document. "Control Packages" are not allowed to weaken "MUST" and "SHOULD" strength statements that are detailed in this document. A "Control Package" may strengthen "SHOULD" to "MUST" if justified by the specific usage of the framework.

In addition to normal sections expected in a standards-track RFC and SIP extension documents, authors of "Control Packages" need to address each of the issues detailed in the following subsections. The following sections MUST be used as a template and included appropriately in all Control-Packages.

10.1. Control Package Name

This section MUST be present in all extensions to this document and provides a token name for the Control Package. The section MUST include information that appears in the IANA registration of the token. Information on registering control package event tokens is contained in [Section 15](#). The package name MUST also register a version number for the package. This enables updates to the package to be registered where appropriate. An initial version of a package MUST start with the value '1.0'. Subsequent versions MUST increment this number if the same package name is to be used. The exact increment is left to the discretion of the package author.

10.2. Framework Message Usage

The Control Framework defines a number of message primitives that can be used to exchange commands and information. There are no limitations restricting the directionality of messages passed down a control channel. This section of a Control package document should explicitly detail the control messages that can be used as well as provide an indication of directionality between entities. This will include which role type is allowed to initiate a request type.

[Editors Note: Need to examine text.]

10.3. Common XML Support

This optional section is only included in a Control Package if the attributes for media dialog or Conference reference are required. The Control Package will make strong statements (MUST strength) if the XML schema defined in Section 17.1 in [Appendix A](#) is to be supported. If only part of the schema is required (for example just 'connection-id' or just conf-id), the Control Package will make equally strong (MUST strength) statements.

10.4. CONTROL Message Bodies

This mandatory section of a Control Package defines the control body that can be contained within a CONTROL command request, as defined in [Section 8](#) (or that no control package body is required). This section should indicate the location of detailed syntax definitions and semantics for the appropriate body types.

10.5. REPORT Message Bodies

This mandatory section of a Control Package defines the REPORT body that can be contained within a REPORT command request, as defined in [Section 8](#) (or that no report package body is required). This section should indicate the location of detailed syntax definitions and semantics for the appropriate body types. It should be noted that the Control Framework specification does allow for payloads to exist in 200 responses to CONTROL messages (as defined in this document). An entity that is prepared to receive a payload type in a REPORT message MUST also be prepared to receive the same payload in a 200 response to a CONTROL message.

10.5.1. Events

A Control Package can optionally include one or more subscriptions that allow a controlling client to receive specific event updates in REPORT message bodies. The mechanisms that installs/un-installs subscriptions is not specified in document and is considered out of scope. Event notifications are always carried in REPORT messages MUST conform to the rules detailed in [Section 8.1.2.2](#). This section of a Control Package definition MUST specify details of the payload expected to be received from subscriptions that have been installed.

[Editors Note: Ongoing discussions relating to a generic subscription/event mechanism across all packages.]

10.6. Examples

It is strongly RECOMMENDED that Control Packages provide a range of message flows that represent common flows using the package and this framework document.

11. Network Address Translation (NAT)

[Editors Note: This section will look at geographically distributed systems where NAT traversal might be an issue. It will look at both the SIP media dialog traversal and the control channel traversal.]

12. Formal Syntax

12.1. SIP Formal Syntax

The ABNF for the "Control-Packages" SIP header is as follows:

```
Control-Packages = "Control-Packages" HCOLON control-package-value
                  *(COMMA control-package-value) *( SEMI package-params )
control-package-value = control-package-name "/" control-package-version
control-package-name = token
control-package-version = 1*DIGIT "." 1*DIGIT
package-params = k-alive-param *(SEMI extension-param)
k-alive-param = "keep-alive" EQUAL delta-seconds
delta-seconds = 1*DIGIT
extension-param = generic-param
```

12.2. Control Framework Formal Syntax

The Control Framework interactions use the UTF-8 transformation format as defined in [RFC3629](#) [[16](#)]. The syntax in this section uses the Augmented Backus-Naur Form (ABNF) as defined in [RFC2234](#) [[17](#)].

```
control-req-or-resp = control-request / control-response
control-request = control-req-start headers [control-content]
control-response = control-resp-start headers
control-req-start = pSCFW SP transact-id SP method CRLF
control-resp-start = pSCFW SP transact-id SP status-code [SP comment] CRLF
comment = utf8text
```

```
pSCFW = %x53.43.46.57; SCFW in caps
transact-id = alpha-num-token
method = mCONTROL / mREPORT / mSYNCH / mK-ALIVE / other-method
```


mCONTROL = %x43.4F.4E.54.52.4F.4C; CONTROL in caps
mREPORT = %x52.45.50.4F.52.54; REPORT in caps
mSYNCH = %x53.59.4E.43.48; SYNCH in caps
mK-ALIVE = %x4B.2D.41.4C.49.56.45;K-ALIVE in caps

other-method = 1*UPALPHA
status-code = 3DIGIT ; any code defined in this and other documents

headers = Content-Length
 /Control-Package
 /Status
 /Seq
 /Timeout
 /Dialog-id
 /ext-header

Content-Length = "Content-Length:" SP 1*DIGIT
Control-Package = "Control-Package:" SP 1*alpha-num-token
Status = "Status:" SP ("pending" / "update" / "terminate")
Timeout = "Timeout:" SP 1*DIGIT
Seq = "Seq:" SP 1*DIGIT
Dialog-id = "Dialog-id:" SP dialog-id-string

dialog-id-string = alpha-num-token "~" alpha-num-token ["~" alpha-num-token]

alpha-num-token = alphanum 3*31alpha-num-token-char
alpha-num-token-char = alphanum / "." / "-" / "+" / "%" / "="

control-content = Content-Type 2CRLF data CRLF

Content-Type = "Content-Type:" SP media-type
media-type = type "/" subtype *(";" gen-param)
type = token
subtype = token

gen-param = pname ["=" pval]
pname = token
pval = token / quoted-string

token = 1*(%x21 / %x23-27 / %x2A-2B / %x2D-2E
 / %x30-39 / %x41-5A / %x5E-7E)
 ; token is compared case-insensitive

quoted-string = DQUOTE *(qdtext / qd-esc) DQUOTE
qdtext = SP / HTAB / %x21 / %x23-5B / %x5D-7E
 / UTF8-NONASCII
qd-esc = (BACKSLASH BACKSLASH) / (BACKSLASH DQUOTE)
BACKSLASH = "\"

UPALPHA = %x41-5A

ALPHANUM = ALPHA / DIGIT

data = *OCTET

ext-header = hname ":" SP hval CRLF

hname = ALPHA *token

hval = utf8text

utf8text = *(HTAB / %x20-7E / UTF8-NONASCII)

UTF8-NONASCII = %xC0-DF 1UTF8-CONT

/ %xE0-EF 2UTF8-CONT

/ %xF0-F7 3UTF8-CONT

/ %xF8-Fb 4UTF8-CONT

/ %xFC-FD 5UTF8-CONT

UTF8-CONT = %x80-BF

The following table details a summary of the headers that can be contained in Control Framework interactions. The "where" columns details where headers can be used:

R: header field may only appear in requests;

r: header field may only appear in responses;

2xx, 4xx, etc.: A numerical value or range indicates response codes with which the header field can be used;

An empty entry in the "where" column indicates that the header field may be present in all requests and responses.

The remaining columns list the specified methods and the presence of a specific header:

m: The header field is mandatory.

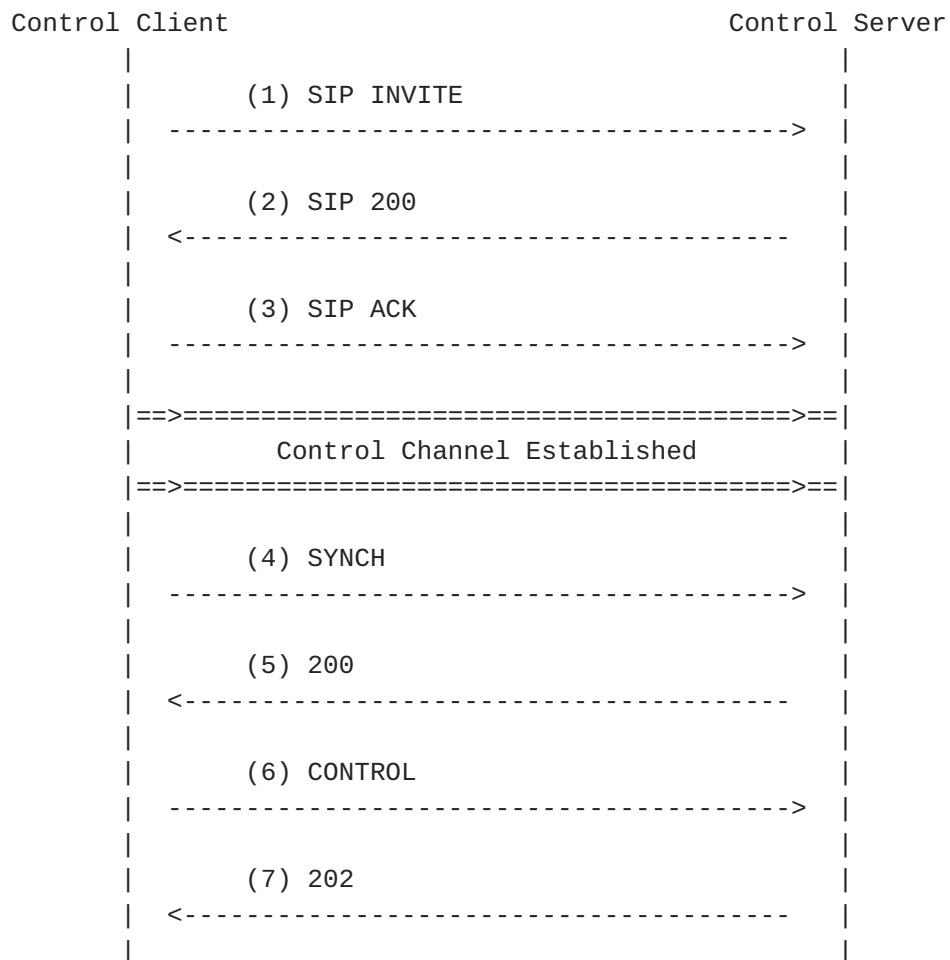
o: The header field is optional.


Header field	Where	CONTROL	REPORT	SYNCH
Content-Length		o	o	-
Control-Package	R	m	-	-
Seq		-	m	-
Status	R	-	m	-
Timeout	R	-	m	-
Dialog-id	R	-	-	m

Figure 11: Table 1

13. Examples

The following examples provide an abstracted flow of Control Channel establishment and Control Framework message exchange. The SIP signaling is prefixed with the token 'SIP'. All other messages are Control Framework interactions defined in this document.





1. Control Client->Control Server (SIP): INVITE
sip:control-server@example.com


```
INVITE sip:control-server@example.com SIP/2.0
To: <sip:control-server@example.com>
From: <sip:control-client@example.com>;tag=8937498
Via: SIP/2.0/UDP control-client.example.com;branch=z9hG412345678
CSeq: 1 INVITE
Require: escs
Control-Packages: <example-package>
Call-ID: 893jhoeihjr8392@example.com
Contact: <sip:control-client@pc1.example.com>
Content-Type: application/sdp
```

```
v=0
o=originator 2890844526 2890842808 IN IP4 controller.example.com
s=-
c=IN IP4 control-client.example.com
m=application 7575 TCP/ESCS
a=setup:active
a=connection:new
```

2. Control Server->Control Client (SIP): 200 OK

```
SIP/2.0 200 OK
To: <sip:control-server@example.com>;tag=023983774
From: <sip:control-client@example.com>;tag=8937498
Via: SIP/2.0/UDP control-client.example.com;branch=z9hG412345678
CSeq: 1 INVITE
Require: escs
Control-Packages: <example-package>
Call-ID: 893jhoeihjr8392@example.com
Contact: <sip:control-client@pc2.example.com>
Content-Type: application/sdp
```

```
v=0
o=originator 2890844526 2890842808 IN IP4 controller.example.com
s=-
c=IN IP4 control-server.example.com
m=application 7575 TCP/ESCS
a=setup:passive
a=connection:new
```

3. Control Client->Control Server (SIP): ACK

4. Control Client opens a TCP connection to the Control Server.
The connection can now be used to exchange control framework
messages. Control Client->Control Server (Control Framework
Message): SYNCH.

SCFW 8djae7khauj SYNCH

Dialog-id: 8937498~893jhoeihjr8392@example.com~023983774

5. Control Server-->Control Client (Control Framework Message):
200.

SCFW 8djae7khauj 200

6. Control Client opens a TCP connection to the Control Server.
The connection can now be used to exchange control framework
messages. Control Client-->Control Server (Control Framework
Message): CONTROL.

SCFW i387yeiqyiq CONTROL
Control-Package: <package-name>
Content-Length: 11

<XML BLOB/>

7. Control Server-->Control Client (Control Framework Message):
202.

SCFW i387yeiqyiq 202

8. Control Server-->Control Client (Control Framework Message):
REPORT.

SCFW i387yeiqyiq REPORT
Seq: 1
Status: pending
Timeout: 10

9. Control Client-->Control Server (Control Framework Message):
200.

SCFW i387yeiqyiq 200
Seq: 1

10. Control Server-->Control Client (Control Framework Message):
REPORT.

SCFW i387yeiqyiq REPORT
Seq: 2
Status: update
Timeout: 10

<XML BLOB/>

11. Control Client-->Control Server (Control Framework Message):
200.

SCFW i387yeiqyiq 200
Seq: 2

12. Control Server-->Control Client (Control Framework Message):
REPORT.

SCFW i387yeiqyiq REPORT
Seq: 3
Status: terminate
Timeout: 10

<XML BLOB/>

13. Control Client-->Control Server (Control Framework Message):
200.

SCFW i387yeiqyiq 200
Seq: 3

14. Control Client->Control Server (SIP): BYE

BYE sip:control-client@pc2.example.com SIP/2.0
To: <sip:control-server@example.com>
From: <sip:control-client@example.com>;tag=8937498
Via: SIP/2.0/UDP control-client.example.com;branch=z9hG423456789
CSeq: 2 BYE
Require: escs
Control-Packages: <example-package>
Call-ID: 893jhoeihjr8392@example.com

15. Control Server->Control Client (SIP): 200 OK

SIP/2.0 200 OK
To: <sip:control-server@example.com>;tag=023983774
From: <sip:control-client@example.com>;tag=8937498
Via: SIP/2.0/UDP control-client.example.com;branch=z9hG423456789
CSeq: 2 BYE
Require: escs
Control-Packages: <example-package>
Call-ID: 893jhoeihjr8392@example.com

14. Security Considerations

Security Considerations to be included in later versions of this document.

15. IANA Considerations

15.1. IANA Registration of the 'escs' Option Tag

15.2. Control Package Registration Information

15.2.1. Control Package Registration Template

15.3. SDP Transport Protocol

15.3.1. TCP/ESCS

15.3.2. TCP/TLS/ESCS

15.4. SDP Attribute Names

15.5. SIP Response Codes

16. Acknowledgments

The authors would like to thank Ian Evans and Michael Bardzinski of Ubiquity Software, Adnan Saleem of Convedia, and Dave Morgan for useful review and input to this work. Eric Burger contributed to the early phases of this work.

17. [Appendix A](#)

During the creation of the Control Framework it has become clear that there are number of components that are common across multiple packages. It has become apparent that it would be useful to collect such re-usable components in a central location. In the short term this appendix provides the place holder for the utilities and it is the intention that this section will eventually form the basis of an initial 'Utilities Document' that can be used by Control Packages.

17.1. Common Dialog/Multiparty Reference Schema

The following schema provides some common attributes for allowing Control Packages to apply specific commands to a particular SIP media dialog (also referred to as Connection) or conference. If used

within a Control Package the Connection and multiparty attributes will be imported and used appropriately to specifically identify either a SIP dialog or a conference instance. If used within a package, the value contained in the 'connection-id' attribute MUST be constructed by concatenating the 'Local' and 'Remote' SIP dialog identifier tags as defined in [RFC3261](#) [2]. They MUST then be separated using the '~' character. So the format would be:

'Local Dialog tag' + '~' + 'Remote Dialog tag'

As an example, for an entity that has a SIP Local dialog identifier of '7HDY839' and a Remote dialog identifier of 'HJKSkyHS', the 'connection-id' attribute for a Control Framework command would be:

7HDY839~HJKSkyHS

If a session description has more than one media description (as identified by 'm=' in [9]) it is possible to explicitly reference them individually. When constructing the 'connection-id' attribute for a command that applies to a specific media ('m=') in an SDP description, an optional third component can be concatenated to the Connection reference key. It is again separated using the '~' character and uses the 'label' attribute as specified in [10]. So the format would be:

'Local Dialog tag' + '~' + 'Remote Dialog tag' + '~' + 'Label Attribute'

As an example, for an entity that has a SIP Local dialog identifier of '7HDY839', a Remote dialog identifier of 'HJKSkyHS' and an SDP label attribute of 'HUwkuh7ns', the 'connection-id' attribute for a Control Framework command would be:

7HDY839~HJKSkyHS~HUwkuh7ns

It should be noted that Control Framework requests initiated in conjunction with a SIP dialog will produce a different 'connection-id' value depending on the directionality of the request, for example Local and Remote tags are locally identifiable.

As with the Connection attribute previously defined, it is also useful to have the ability to apply specific control framework commands to a number of related dialogs, such as a multiparty call. This typically consists of a number of media dialogs that are logically bound by a single identifier. The following schema allows for control framework commands to explicitly reference such a grouping through a 'conf' XML container. If used by a Control Package, any control XML referenced by the attribute applies to all related media dialogs. Unlike the dialog attribute, the 'conf-id'

attribute does not need to be constructed based on the overlying SIP dialog. The 'conf-id' attribute value is system specific and should be selected with relevant context and uniqueness.

The full schema follows:

```
<?xml version="1.0" encoding="UTF-8"?>
  <xsd:schema targetNamespace="urn:ietf:params:xml:ns:control:framework-
attributes"
  xmlns:xs="http://www.w3.org/2001/XMLSchema"
  xmlns="urn:ietf:params:xml:ns::control:framework-attributes"
  elementFormDefault="qualified" attributeFormDefault="unqualified">
    <!-- xs:include schemaLocation="common-schema.xsd"/ -->

    <xsd:attributeGroup name="framework-attributes">
      <xsd:annotation>
        <xsd:documentation>SIP Connection and Conf Identifiers</
xsd:documentation>
      </xsd:annotation>

      <xsd:attribute name="connection-id" type="xsd:string"/>
      <xsd:attribute name="conf-id" type="xsd:string"/>

    </xsd:attributeGroup>
  </xs:schema>
```

18. References

18.1. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

18.2. Informative References

- [2] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", [RFC 3261](#), June 2002.
- [3] Rosenberg, J. and H. Schulzrinne, "Reliability of Provisional Responses in Session Initiation Protocol (SIP)", [RFC 3262](#), June 2002.
- [4] Rosenberg, J. and H. Schulzrinne, "Session Initiation Protocol (SIP): Locating SIP Servers", [RFC 3263](#), June 2002.

- [5] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", [RFC 3264](#), June 2002.
- [6] Yon, D. and G. Camarillo, "TCP-Based Media Transport in the Session Description Protocol (SDP)", [RFC 4145](#), September 2005.
- [7] Groves, C., Pantaleo, M., Anderson, T., and T. Taylor, "Gateway Control Protocol Version 1", [RFC 3525](#), June 2003.
- [8] Dolly, M., "Media Control Protocol Requirements", [draft-dolly-xcon-mediactrlframe-02](#) (work in progress), September 2006.
- [9] Handley, M., "SDP: Session Description Protocol", [draft-ietf-mmusic-sdp-new-26](#) (work in progress), January 2006.
- [10] Levin, O. and G. Camarillo, "The SDP (Session Description Protocol) Label Attribute", [draft-ietf-mmusic-sdp-media-label-01](#) (work in progress), January 2005.
- [11] Jennings, C. and R. Mahy, "Managing Client Initiated Connections in the Session Initiation Protocol (SIP)", [draft-ietf-sip-outbound-07](#) (work in progress), January 2007.
- [12] Rosenberg, J., Peterson, J., Schulzrinne, H., and G. Camarillo, "Best Current Practices for Third Party Call Control (3pcc) in the Session Initiation Protocol (SIP)", [BCP 85](#), [RFC 3725](#), April 2004.
- [13] Rosenberg, J., Schulzrinne, H., and P. Kyzivat, "Indicating User Agent Capabilities in the Session Initiation Protocol (SIP)", [RFC 3840](#), August 2004.
- [14] Rosenberg, J., Schulzrinne, H., and P. Kyzivat, "Caller Preferences for the Session Initiation Protocol (SIP)", [RFC 3841](#), August 2004.
- [15] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, [RFC 3550](#), July 2003.
- [16] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, [RFC 3629](#), November 2003.
- [17] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", [RFC 2234](#), November 1997.

Authors' Addresses

Chris Boulton
Ubiquity Software Corporation
Building 3
Wern Fawr Lane
St Mellons
Cardiff, South Wales CF3 5EA

Email: cboulton@ubiquitysoftware.com

Tim Melanchuk
BlankSpace

Email: tim.melanchuk@gmail.com

Scott McGlashan
Hewlett-Packard
Gustav III:s boulevard 36
SE-16985 Stockholm, Sweden

Email: scott.mcglashan@hp.com

Asher Shiratzky
Radvision
24 Raoul Wallenberg st
Tel-Aviv, Israel

Email: ashers@radvision.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

