

# SIP: Session Initiation Protocol

## Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026.

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>

To view the list Internet-Draft Shadow Directories, see <http://www.ietf.org/shadow.html>.

## Copyright Notice

Copyright (c) The Internet Society (2002). All Rights Reserved.

### Abstract

The Session Initiation Protocol (SIP) is an application-layer control (signaling) protocol for creating, modifying and terminating sessions with one or more participants. These sessions include Internet telephone calls, multimedia distribution and multimedia conferences.

SIP invitations used to create sessions carry session descriptions which allow participants to agree on a set of compatible media types. SIP makes use of elements called proxy servers to help route requests to the users current location, assist in firewall traversal, and provide features to users. SIP also provides a registration function that allows them to upload their current location for use by proxy servers. SIP runs ontop of several different transport protocols.

## Contents

<b>1</b>	<b>Introduction</b>	<b>7</b>
<b>2</b>	<b>Overview of SIP Functionality</b>	<b>8</b>
<b>3</b>	<b>Terminology</b>	<b>8</b>
<b>4</b>	<b>Overview of Operation</b>	<b>9</b>
<b>5</b>	<b>Structure of the Protocol</b>	<b>14</b>
<b>6</b>	<b>Definitions</b>	<b>15</b>
<b>7</b>	<b>SIP Messages</b>	<b>18</b>
7.1	Requests . . . . .	19
7.2	Responses . . . . .	19
7.3	Header Fields . . . . .	20

34	7.3.1	Header Field Format . . . . .	20
35	7.3.2	Header Field Classification . . . . .	22
36	7.3.3	Compact Form . . . . .	22
37	7.4	Bodies . . . . .	22
38	7.4.1	Message Body Type . . . . .	22
39	7.4.2	Message Body Length . . . . .	23
40	7.5	Framing SIP messages . . . . .	23
41	<b>8</b>	<b>General User Agent Behavior</b>	<b>23</b>
42	8.1	UAC Behavior . . . . .	24
43	8.1.1	Generating the Request . . . . .	24
44	8.1.2	Sending the Request . . . . .	26
45	8.1.3	Processing Responses . . . . .	26
46	8.2	UAS Behavior . . . . .	28
47	8.2.1	Authentication/Authorization . . . . .	28
48	8.2.2	Method Inspection . . . . .	28
49	8.2.3	Header Inspection . . . . .	28
50	8.2.4	Content Processing . . . . .	29
51	8.2.5	Applying Extensions . . . . .	29
52	8.2.6	Processing the Request . . . . .	30
53	8.2.7	Generating the Response . . . . .	30
54	8.2.8	Stateless UAS Behavior . . . . .	30
55	8.3	Redirect Servers . . . . .	31
56	<b>9</b>	<b>Canceling a Request</b>	<b>32</b>
57	9.1	Client Behavior . . . . .	32
58	9.2	Server Behavior . . . . .	33
59	<b>10</b>	<b>Registrations</b>	<b>33</b>
60	10.1	Overview of Usage . . . . .	33
61	10.2	Construction of the REGISTER request . . . . .	35
62	10.2.1	Adding Bindings with REGISTER . . . . .	35
63	10.2.2	Removing Bindings with REGISTER . . . . .	36
64	10.2.3	Fetching Bindings with REGISTER . . . . .	37
65	10.2.4	Refreshing Registrations . . . . .	37
66	10.2.5	Discovering a Registrar . . . . .	37
67	10.3	Processing of REGISTER at the Registrar . . . . .	37
68	<b>11</b>	<b>Querying for Capabilities</b>	<b>39</b>
69	11.1	Construction of OPTIONS Request . . . . .	40
70	11.2	Processing of OPTIONS Request . . . . .	40
71	<b>12</b>	<b>Dialogs</b>	<b>41</b>
72	12.1	Creation of a Dialog . . . . .	42
73	12.1.1	UAS behavior . . . . .	42
74	12.1.2	UAC behavior . . . . .	43

75	12.2 Requests within a Dialog . . . . .	43
76	12.2.1 UAC Behavior . . . . .	44
77	12.2.2 UAS behavior . . . . .	45
78	12.3 Termination of a Dialog . . . . .	46
79	<b>13 Initiating a Session</b>	<b>46</b>
80	13.1 Overview . . . . .	46
81	13.2 Caller Processing . . . . .	47
82	13.2.1 Creating the Initial INVITE . . . . .	47
83	13.2.2 Processing INVITE Responses . . . . .	48
84	13.3 Callee Processing . . . . .	49
85	13.3.1 Processing of the INVITE . . . . .	49
86	<b>14 Modifying an Existing Session</b>	<b>51</b>
87	14.1 UAC Behavior . . . . .	51
88	14.2 UAS Behavior . . . . .	52
89	<b>15 Terminating a Session</b>	<b>52</b>
90	15.1 Terminating a Dialog with a BYE . . . . .	53
91	15.1.1 UAC Behavior . . . . .	53
92	15.1.2 UAS Behavior . . . . .	53
93	<b>16 Proxy Behavior</b>	<b>54</b>
94	16.1 Overview . . . . .	54
95	16.2 Stateful Proxy . . . . .	54
96	16.3 Request Validation . . . . .	55
97	16.4 Making a Routing Decision . . . . .	57
98	16.5 Request Processing . . . . .	58
99	16.6 Response Processing . . . . .	62
100	16.7 Handling transport errors . . . . .	66
101	16.8 CANCEL Processing . . . . .	66
102	16.9 Stateless proxy . . . . .	66
103	<b>17 Transactions</b>	<b>68</b>
104	17.1 Client transaction . . . . .	69
105	17.1.1 INVITE Client Transaction . . . . .	69
106	17.1.2 non-INVITE Client Transaction . . . . .	73
107	17.1.3 Matching Responses to Client Transactions . . . . .	75
108	17.1.4 Handling Transport Errors . . . . .	75
109	17.2 Server Transaction . . . . .	75
110	17.2.1 INVITE Server Transaction . . . . .	75
111	17.2.2 non-INVITE Server Transaction . . . . .	77
112	17.2.3 Matching Requests to Server Transactions . . . . .	79
113	17.3 RTT Estimation . . . . .	79
114	<b>18 Reliability of Provisional Responses</b>	<b>80</b>

115	18.1 UAS Behavior . . . . .	80
116	18.2 UAC Behavior . . . . .	82
117	<b>19 Transport</b>	<b>82</b>
118	19.1 Clients . . . . .	82
119	19.1.1 Sending Requests . . . . .	82
120	19.1.2 Receiving Responses . . . . .	83
121	19.2 Servers . . . . .	84
122	19.2.1 Receiving Requests . . . . .	84
123	19.2.2 Sending Responses . . . . .	84
124	19.3 Framing . . . . .	85
125	19.4 Error Handling . . . . .	85
126	<b>20 Security Considerations</b>	<b>85</b>
127	20.1 Transport and Network Layer Security . . . . .	86
128	20.2 SIP Authentication . . . . .	87
129	20.2.1 Framework . . . . .	87
130	20.2.2 User to User Authentication . . . . .	88
131	20.2.3 Proxy to User Authentication . . . . .	89
132	20.2.4 Authentication Schemes . . . . .	91
133	20.3 SIP Encryption . . . . .	92
134	20.4 Denial of Service . . . . .	92
135	<b>21 Common Message Components</b>	<b>93</b>
136	21.1 SIP Uniform Resource Indicators . . . . .	94
137	21.1.1 SIP URI components . . . . .	94
138	21.1.2 Character escaping requirements . . . . .	96
139	21.1.3 Example SIP URIs . . . . .	97
140	21.1.4 SIP URI Comparison . . . . .	97
141	21.1.5 Forming Requests from a SIP URI . . . . .	99
142	21.2 Option Tags . . . . .	99
143	21.3 Tags . . . . .	99
144	<b>22 Header Fields</b>	<b>100</b>
145	22.1 Accept . . . . .	101
146	22.2 Accept-Encoding . . . . .	101
147	22.3 Accept-Language . . . . .	101
148	22.4 Alert-Info . . . . .	103
149	22.5 Allow . . . . .	104
150	22.6 Authentication-Info . . . . .	104
151	22.7 Authorization . . . . .	104
152	22.8 Call-ID . . . . .	104
153	22.9 Call-Info . . . . .	105
154	22.10 Contact . . . . .	105
155	22.11 Content-Disposition . . . . .	106
156	22.12 Content-Encoding . . . . .	106

157	22.13 Content-Language . . . . .	106
158	22.14 Content-Length . . . . .	107
159	22.15 Content-Type . . . . .	107
160	22.16 CSeq . . . . .	107
161	22.17 Date . . . . .	107
162	22.18 Error-Info . . . . .	108
163	22.19 Expires . . . . .	108
164	22.20 From . . . . .	108
165	22.21 In-Reply-To . . . . .	109
166	22.22 Max-Forwards . . . . .	109
167	22.23 MIME-Version . . . . .	109
168	22.24 Organization . . . . .	109
169	22.25 Priority . . . . .	110
170	22.26 Proxy-Authenticate . . . . .	110
171	22.27 Proxy-Authorization . . . . .	110
172	22.28 Proxy-Require . . . . .	111
173	22.29 RAck . . . . .	111
174	22.30 Record-Route . . . . .	111
175	22.31 Require . . . . .	111
176	22.32 Retry-After . . . . .	111
177	22.33 Route . . . . .	112
178	22.34 RSeq . . . . .	112
179	22.35 Server . . . . .	112
180	22.36 Subject . . . . .	112
181	22.37 Supported . . . . .	113
182	22.38 Timestamp . . . . .	113
183	22.39 To . . . . .	113
184	22.40 Unsupported . . . . .	113
185	22.41 User-Agent . . . . .	114
186	22.42 Via . . . . .	114
187	22.43 Warning . . . . .	114
188	22.44 WWW-Authenticate . . . . .	115
189	<b>23 Response Codes</b> . . . . .	<b>116</b>
190	23.1 Provisional 1xx . . . . .	116
191	23.1.1 100 Trying . . . . .	116
192	23.1.2 180 Ringing . . . . .	116
193	23.1.3 181 Call Is Being Forwarded . . . . .	116
194	23.1.4 182 Queued . . . . .	116
195	23.1.5 183 Session Progress . . . . .	117
196	23.2 Successful 2xx . . . . .	117
197	23.2.1 200 OK . . . . .	117
198	23.3 Redirection 3xx . . . . .	117
199	23.3.1 300 Multiple Choices . . . . .	117
200	23.3.2 301 Moved Permanently . . . . .	117

201	23.3.3 302 Moved Temporarily . . . . .	117
202	23.3.4 305 Use Proxy . . . . .	118
203	23.3.5 380 Alternative Service . . . . .	118
204	23.4 Request Failure 4xx . . . . .	118
205	23.4.1 400 Bad Request . . . . .	118
206	23.4.2 401 Unauthorized . . . . .	118
207	23.4.3 402 Payment Required . . . . .	118
208	23.4.4 403 Forbidden . . . . .	118
209	23.4.5 404 Not Found . . . . .	118
210	23.4.6 405 Method Not Allowed . . . . .	119
211	23.4.7 406 Not Acceptable . . . . .	119
212	23.4.8 407 Proxy Authentication Required . . . . .	119
213	23.4.9 408 Request Timeout . . . . .	119
214	23.4.10 410 Gone . . . . .	119
215	23.4.11 413 Request Entity Too Large . . . . .	119
216	23.4.12 414 Request-URI Too Long . . . . .	119
217	23.4.13 415 Unsupported Media Type . . . . .	119
218	23.4.14 420 Bad Extension . . . . .	120
219	23.4.15 421 Extension Required . . . . .	120
220	23.4.16 480 Temporarily Unavailable . . . . .	120
221	23.4.17 481 Call/Transaction Does Not Exist . . . . .	120
222	23.4.18 482 Loop Detected . . . . .	120
223	23.4.19 483 Too Many Hops . . . . .	120
224	23.4.20 484 Address Incomplete . . . . .	120
225	23.4.21 485 Ambiguous . . . . .	121
226	23.4.22 486 Busy Here . . . . .	121
227	23.4.23 487 Request Terminated . . . . .	121
228	23.4.24 488 Not Acceptable Here . . . . .	121
229	23.5 Server Failure 5xx . . . . .	121
230	23.5.1 500 Server Internal Error . . . . .	121
231	23.5.2 501 Not Implemented . . . . .	122
232	23.5.3 502 Bad Gateway . . . . .	122
233	23.5.4 503 Service Unavailable . . . . .	122
234	23.5.5 504 Server Time-out . . . . .	122
235	23.5.6 505 Version Not Supported . . . . .	122
236	23.5.7 513 Message Too Large . . . . .	122
237	23.6 Global Failures 6xx . . . . .	122
238	23.6.1 600 Busy Everywhere . . . . .	123
239	23.6.2 603 Decline . . . . .	123
240	23.6.3 604 Does Not Exist Anywhere . . . . .	123
241	23.6.4 606 Not Acceptable . . . . .	123
242	<b>24 Examples</b>	<b>123</b>
243	24.1 Registration . . . . .	123
244	24.2 Session Setup . . . . .	124

245	<b>25 Augmented BNF for the SIP Protocol</b>	<b>129</b>
246	25.1 Basic Rules . . . . .	130
247	<b>26 IANA Considerations</b>	<b>144</b>
248	26.1 Option Tags . . . . .	144
249	26.1.1 Registration of 100rel . . . . .	144
250	26.2 Warn-Codes . . . . .	145
251	26.3 Header Field Names . . . . .	145
252	26.4 Method and Response Codes . . . . .	146
253	<b>27 Changes Made in Version 00</b>	<b>146</b>
254	<b>28 Changes Made in Version 01</b>	<b>151</b>
255	<b>29 Changes Made in Version 02</b>	<b>151</b>
256	<b>30 Changes Made in Version 03</b>	<b>153</b>
257	<b>31 Changes Made in Version 04</b>	<b>155</b>
258	<b>32 Changes Made in Version 05</b>	<b>156</b>
259	<b>33 Changes Made in Version 06</b>	<b>159</b>
260	<b>34 Acknowledgments</b>	<b>161</b>
261	<b>35 Authors' Addresses</b>	<b>162</b>

## 262 1 Introduction

263 There are many applications of the Internet that require the creation and management of a session, where  
264 a session is considered an exchange of data between an association of participants. The implementation  
265 of these services is complicated by the practices of participants; users may move between endpoints, they  
266 may be addressable by multiple names, and they may communicate in several different media - sometimes  
267 simultaneously. Numerous protocols have been authored that carry various forms of real-time multimedia  
268 session data such as voice, video, or text messages. SIP works in concert with these protocols by enabling  
269 Internet endpoints (called “user agents”) to discover one another and to agree on a characterization of a  
270 session they would like to share. For locating prospective session participants, SIP relies on an infrastructure  
271 of network hosts (called “proxy servers”) to which user agents can send registrations, invitations to sessions  
272 and other requests. SIP is an agile, general-purpose tool for creating, modifying and terminating sessions  
273 that works independently of underlying transport protocols and without dependency on the type of session  
274 that is being established.

## 2 Overview of SIP Functionality

The Session Initiation Protocol (SIP) is an application-layer control protocol that can establish, modify, and terminate multimedia sessions (conferences) such as Internet telephony calls. SIP can also invite participants to already existing sessions. A SIP entity issuing an invitation for an already existing session does not necessarily have to be a member of the session to which it is inviting. Media can be added to (and removed from) an existing session. SIP transparently supports name mapping and redirection services, which supports *personal mobility* [1, p. 44] - users can maintain a single externally visible identifier (SIP URI) regardless of their network location.

SIP supports five facets of establishing and terminating multimedia communications:

**User location:** determination of the end system to be used for communication;

**User availability:** determination of the willingness of the called party to engage in communications;

**User capabilities:** determination of the media and media parameters to be used;

**Session setup:** “ringing”, establishment of session parameters at both called and calling party;

**Session handling:** including transfer and termination of sessions, modifying session parameters, and invoking services.

SIP is not a vertically integrated communications system. SIP is rather a component of the overall IETF multimedia data and control architecture that incorporates protocols such as RSVP (RFC 2205 [2]) for reserving network resources, the real-time transport protocol (RTP) (RFC 1889 [3]) for transporting real-time data and providing QOS feedback, the real-time streaming protocol (RTSP) (RFC 2326 [4]) for controlling delivery of streaming media, the session announcement protocol (SAP) [5] for advertising multimedia sessions via multicast and the session description protocol (SDP) (RFC 2327 [6]) for describing multimedia sessions. Therefore, SIP should be used in conjunction with other protocols in order to provide complete services to the users. However, the basic functionality and operation of SIP does not depend on any of these protocols.

SIP does not provide services. SIP rather provides primitives that can be used to implement different services. For example, SIP can locate a user and deliver an opaque object to his current location. If this primitive is used to deliver a session description written in SDP, for instance, the parameters of a session can be agreed between endpoints. If the same primitive is used to deliver a photo of the caller as well as the session description, a “caller ID” service can be easily implemented. As this example shows, a single primitive is typically used to provide several different services. Consequently, generality is more important than efficiency when designing SIP primitives.

SIP does not offer conference control services such as floor control or voting and does not prescribe how a conference is to be managed. SIP can be used to initiate a session that uses some other conference control protocol. SIP does not allocate multicast addresses and does not reserve network resources.

## 3 Terminology

In this document, the key words “MUST”, “MUST NOT”, “REQUIRED”, “SHALL”, “SHALL NOT”, “SHOULD”, “SHOULD NOT”, “RECOMMENDED”, “MAY”, and “OPTIONAL” are to be interpreted as described in RFC 2119 [7] and indicate requirement levels for compliant SIP implementations.



## 313 4 Overview of Operation

314 This section introduces the basic operations of SIP using simple examples. Note that this section is tutorial  
315 in nature and does not contain any normative statements.

316 The first example shows the basic functions of SIP: location of an end point, signal of a desire to com-  
317 municate, negotiation of session parameters to establish the session, and teardown of the session once es-  
318 tablished.

319 Figure 1 shows a typical example of a SIP message exchange between two users, Alice and Bob. (Each  
320 message is labeled with the letter “F” and a number for reference by the text.) In this example, Alice uses a  
321 SIP application on her PC (referred to as a softphone) to call Bob on his SIP phone over the Internet. Also  
322 shown are two SIP proxy servers that act on behalf of Alice and Bob to facilitate the session establishment.  
323 This typical arrangement is often referred to as the “SIP trapezoid” as shown by the geometric shape of the  
324 dashed lines in Figure 1.

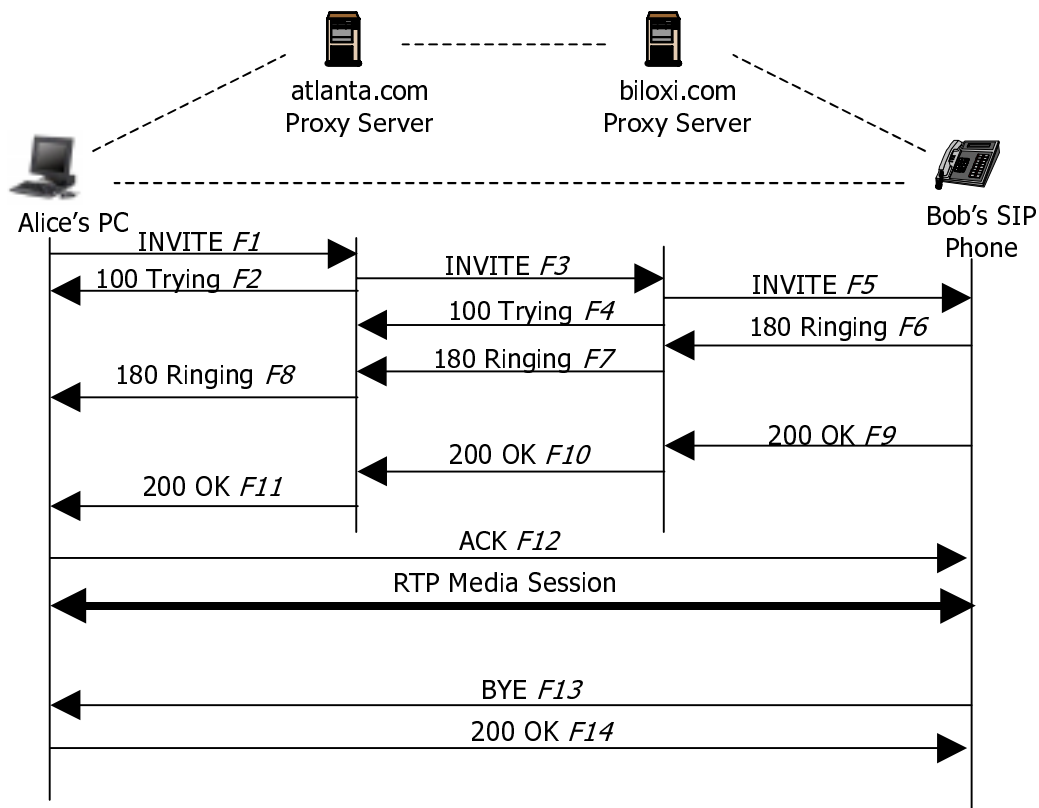


Figure 1: SIP session setup example with SIP trapezoid

325 Alice “calls” Bob using his SIP identity, a type of Uniform Resource Identifier (URI) called a SIP URI  
326 and defined in Section 21.1. It has a similar form to an email address, typically containing a username and  
327 a host name. In this case, it is sip:bob@biloxi.com, where biloxi.com is the domain of Bob’s SIP service  
328 provider (which can be an enterprise, retail provider, etc). Alice also has a SIP URI of sip:alice@atlanta.com.  
329 Alice might have typed in Bob’s URI or perhaps clicked on a hyperlink or an entry in an address book.

330 SIP is based on an HTTP-like request/response transaction model. Each transaction consists of a request  
331 that invokes a particular “Method”, or function, on the server, and at least one response. In this example, the  
332 transaction begins with Alice’s softphone sending an INVITE request addressed to Bob’s SIP URI. INVITE  
333 is an example of a SIP method which specifies the action that the requestor (Alice) wants the server (Bob)  
334 to take. The INVITE request contains a number of header fields. Header fields are named attributes that  
335 provide additional information about a message. The ones present in an INVITE include a unique identifier  
336 for the call, the destination address, Alice’s address, and information about the type of session that Alice  
337 wishes to establish with Bob. The INVITE (message F1 in Figure 1) might look like this:

```
338 INVITE sip:bob@biloxi.com SIP/2.0
339 Via: SIP/2.0/UDP 10.1.3.3:5060
340 To: Bob <sip:bob@biloxi.com>
341 From: Alice <sip:alice@atlanta.com>;tag=1928301774
342 Call-ID: a84b4c76e66710@10.1.3.3
343 CSeq: 314159 INVITE
344 Contact: <sip:alice@10.1.3.3>
345 Content-Type: application/sdp
346 Content-Length: 142
347
348 (Alice’s SDP not shown)
```

349 The first line of the text-encoded message contains the method name (INVITE). The lines that follow  
350 are a list of header fields. This example contains a minimum required set. The headers are briefly described  
351 below:

352 **Via** contains the IP address (10.1.3.3), port number (5060), and transport protocol (UDP) on which Alice  
353 is expecting to receive responses to this request.

354 **To** contains a display name (Bob) and a SIP URI (sip:bob@biloxi.com) towards which the request was  
355 originally directed.

356 **From** also contains a display name (Alice) and a SIP URI (sip:alice@atlanta.com) that indicate the  
357 originator of the request. This header field also has a **tag** parameter containing a pseudorandom string  
358 (1928301774) that was added to the URI by the softphone. It is used for identification purposes.

359 **Call-ID** contains a globally unique identifier for this call, generated by the combination of a pseudoran-  
360 dom string and the softphone’s IP address. The combination of the **To**, **From**, and **Call-ID** completely define  
361 a peer-to-peer SIP relationship between Alice and Bob, and is referred to as a “dialog”.

362 **CSeq** or Command Sequence contains an integer and a method name. The **CSeq** number is incremented  
363 for each new request, and is a traditional sequence number.

364 **Contact** contains a SIP URI that represents a direct route to reach or contact Alice, usually composed  
365 of a username at an IP address. While the **Via** header field tells other elements where to send the response,  
366 the **Contact** header field tells other elements where to send future requests for this dialog.

367 **Content-Type** contains a description of the message body (not shown).

368 **Content-Length** contains an octet (byte) count of the message body.

369 The complete set of SIP header fields is defined in Section 22.

370 The details of the session, type of media, codec, sampling rate, etc. are not described using SIP. Rather,  
371 the body of a SIP message contains a description of the session, encoded in some other protocol format.  
372 One such format is Session Description Protocol (SDP) [6]. This SDP message (not shown in the example)

373 is carried by the SIP message in a way that is analogous to a document attachment being carried by an email  
374 message, or a web page being carried in an HTTP message.

375 Since the softphone does not know the location of Bob or the SIP server in the biloxi.com domain, the  
376 softphone sends the INVITE to the SIP server that serves Alice's domain, atlanta.com. The IP address of the  
377 atlanta.com SIP server could have been configured in Alice's softphone, or it could have been discovered by  
378 DHCP, for example.

379 The atlanta.com SIP server is a type of SIP server known as a proxy server. A proxy server receives SIP  
380 requests and forwards them on behalf of the requestor. In this example, the proxy server receives the INVITE  
381 request and sends a 100 Trying response back to Alice's softphone. The 100 Trying response indicates  
382 that the INVITE has been received and that the proxy is working on her behalf to route the INVITE to the  
383 destination. Responses in SIP use a three-digit code followed by a descriptive phrase. This response contains  
384 the same To, From, Call-ID, and CSeq as the INVITE, which allows Alice's softphone to correlate this  
385 response to the sent INVITE. The atlanta.com proxy server locates the proxy server at biloxi.com, possibly  
386 by performing a DNS (Domain Name Service) lookup to find the SIP server that serves the biloxi.com  
387 domain. This is described in [8]. As a result, it obtains the IP address of the biloxi.com proxy server and  
388 forwards, or proxies, the INVITE request there. Before forwarding the request, the atlanta.com proxy server  
389 adds an additional Via header field that contains its own IP address (the INVITE already contains Alice's IP  
390 address in the first Via). The biloxi.com proxy server receives the INVITE and responds with a 100 Trying  
391 response back to the Atlanta.com proxy server to indicate that it has received the INVITE and is processing  
392 the request. The proxy server consults a database, generically called a location service, that contains the  
393 current IP address of Bob. (We shall see in the next section how this database can be populated.) The  
394 biloxi.com proxy server adds another Via header with its own IP address to the INVITE and proxies it to  
395 Bob's SIP phone.

396 Bob's SIP phone receives the INVITE and alerts Bob to the incoming call from Alice so that Bob can  
397 decide whether or not to answer the call - i.e. Bob's phone rings. Bob's SIP phone sends an indication of  
398 this in a 180 Ringing response, which is routed back through the two proxies in the reverse direction. Each  
399 proxy uses the Via header to determine where to send the response and removes its own address from the  
400 top. As a result, although DNS and location service lookups were required to route the initial INVITE, the  
401 180 Ringing response can be returned to the caller without lookups or without state being maintained in the  
402 proxies. This also has the desirable property that each proxy that sees the INVITE will also see all responses  
403 to the INVITE.

404 When Alice's softphone receives the 180 Ringing response, it passes this information to Alice, perhaps  
405 using an audio ringback tone or by displaying a message on Alice's screen.

406 In this example, Bob decides to answer the call. When he picks up the handset, his SIP phone sends a  
407 200 OK response to indicate that the call has been answered. The 200 OK contains a message body with the  
408 SDP media description of the type of session that Bob is willing to establish with Alice. As a result, there  
409 is a two-phase exchange of SDP messages; Alice sent one to Bob, and Bob sent one back to Alice. This  
410 two-phase exchange provides basic negotiation capabilities and is based on a simple offer/answer model. If  
411 Bob did not wish to answer the call or was busy on another call, an error response would have been sent  
412 instead of the 200 OK, which would have resulted in no media session being established. The complete list  
413 of SIP response codes is in Section 23. The 200 OK (message F9 in Figure 1) might look like this:

```
414 SIP/2.0 200 OK
415 Via: SIP/2.0/UDP 10.2.1.1:5060;branch=4b43c2ff8.1
416 Via: SIP/2.0/UDP 10.1.1.1:5060;branch=77ef4c2312983.1
```

```
417 Via: SIP/2.0/UDP 10.1.3.3:5060
418 To: Bob <sip:bob@biloxi.com>;tag=a6c85cf
419 From: Alice <sip:alice@atlanta.com>;tag=1928301774
420 Call-ID: a84b4c76e66710@10.1.3.3
421 CSeq: 314159 INVITE
422 Contact: <sip:bob@10.4.1.4>
423 Content-Type: application/sdp
424 Content-Length: 131
425
426 (Bob's SDP not shown)
```

427 The first line of the response contains the response code (200) and the reason phrase (OK). The remain-  
428 ing lines contain header fields. The *Via* header fields, *To*, *From*, *Call-ID*, and *CSeq* are all copied from  
429 the *INVITE* request. (Note that there are three *Via* headers - one added by Alice's SIP phone, one added by  
430 the atlanta.com proxy, and one added by the biloxi.com proxy.) Also note that Bob's SIP phone has added a  
431 *tag* parameter to the *To* header field. This tag will be incorporated by both User Agents into the dialog and  
432 will be included in all future requests and responses in this call. The *Contact* header field contains a URI at  
433 which Bob can be directly reached at his SIP phone. The *Content-Type* and *Content-Length* refer to the  
434 message body (not shown) that contains Bob's SDP media information.

435 In addition to DNS and location service lookups shown in this example, proxy servers can make arbitrar-  
436 ily complex "routing decisions" to decide where to send a request. For example, if Bob's SIP phone returned  
437 a 486 Busy Here response, the biloxi.com proxy server could proxy the *INVITE* to Bob's voicemail server.  
438 A proxy server can also send an *INVITE* to a number of locations at the same time. This type of parallel  
439 search is known as "forking".

440 In this case, the 200 OK is routed back through the two proxies and is received by Alice's softphone  
441 which then stops the ringback tone and indicates that the call has been answered. Finally, an acknowledge-  
442 ment message, *ACK*, is sent by Alice to Bob to confirm the reception of the final response (200 OK). Note  
443 that in this example, the *ACK* is sent directly from Alice to Bob, bypassing the two proxies. This is be-  
444 cause, through the *INVITE*/200 OK exchange, the two SIP user agents have learned each other's IP address  
445 through the *Contact* header fields, which was not known when the initial *INVITE* was sent. The lookups  
446 performed by the two proxies are no longer needed, so they drop out of the call flow. This completes the  
447 *INVITE*/200/*ACK* three-way handshake used to establish SIP sessions and is the end of the transaction. Full  
448 details on session setup are in Section 13.

449 Alice and Bob's media session has now begun, and they send media packets using the format agreed to  
450 in the exchange of SDP. In general, the end-to-end media packets take a different path from the SIP signaling  
451 messages.

452 During the session, either Alice or Bob may decide to change the characteristics of the media session.  
453 This is accomplished by sending a re-*INVITE* containing a new media description. If the change is accepted  
454 by the other party, a 200 OK is sent, which is itself responded to with an *ACK*. This re-*INVITE* references  
455 the existing dialog so the other party knows that it is to modify an existing session instead of establishing a  
456 new session. If the change is not accepted, an error response, such as a 406 Not Acceptable, is sent, which  
457 also receives an *ACK*. However, the failure of the re-*INVITE* does not cause the existing call to fail - the  
458 session continues using the previously negotiated characteristics. Full details on session modification is in  
459 Section 14.

460 At the end of the call, Bob disconnects (hangs up) first, and generates a *BYE* message. This *BYE* is

461 routed directly to Alice's softphone, again bypassing the proxies. Alice confirms receipt of the BYE with  
462 a 200 OK response, which terminates the session and the BYE transaction. Note that no ACK is sent - an  
463 ACK is only sent in response to a response to an INVITE request. The reasons for this special handling for  
464 INVITE will be discussed later, but relate to the reliability mechanisms in SIP, the length of time it can take  
465 for a ringing phone to be answered, and forking. For this reason, request handling in SIP is often classified  
466 as either INVITE or non- INVITE, referring to all other methods besides INVITE. Full details on session  
467 termination is in Section 15.

468 Full details of all the messages shown in the example of Figure 1 are shown in Section 24.2.

469 In some cases, it may be useful for proxies in the SIP signaling path to see all the messaging between  
470 the endpoints for the duration of the session. For example, if the biloxi.com proxy server wished to remain  
471 in the SIP messaging path beyond the initial INVITE, it would add to the INVITE a required routing header  
472 field known as **Record-Route** that contained a URI resolving to the proxy. This information would be  
473 received by both Bob's SIP phone and (due to the **Record-Route** header field being passed back in the 200  
474 OK) Alice's softphone and stored for the duration of the dialog. The biloxi.com proxy server would then  
475 receive and proxy the ACK, BYE, and 200 OK to the BYE. Each proxy can independently decide to receive  
476 subsequent messaging, and that messaging will go through all proxies that elect to receive it. Common uses  
477 of this capability are firewall traversal and mid-call feature implementation.

478 Registration is another common operation in SIP. Registration is one way that the biloxi.com server  
479 can learn the current location of Bob. Upon initialization, and at periodic intervals, Bob's SIP phone sends  
480 REGISTER messages to a server in the biloxi.com domain known as a SIP registrar. The REGISTER  
481 messages associate Bob's SIP URL (sip:bob@biloxi.com) with the machine he is currently logged in at  
482 (conveyed as a SIP URL in the **Contact** header). The registrar writes this association, also called a binding,  
483 to a database, called the *location service*, where it can be used by the proxy in the biloxi.com domain. Often,  
484 a registrar server for a domain is co-located with the proxy for that domain. It is an important concept that  
485 the distinction between types of SIP servers is logical, not physical.

486 Bob is not limited to registering from a single device. For example, both his SIP phone at home and  
487 the one in the office could send registrations. This information is stored together in the location service and  
488 allows a proxy to perform various types of searches to locate Bob. Similarly, more than one user can be  
489 registered on a single device at the same time.

490 The location service is just an abstract concept. It generally contains information that allows a proxy  
491 to input a URI and get back a translated URI that tells the proxy where to send the request. Registrations  
492 are one way to create this information, but not the only way. Arbitrarily complex mapping functions can be  
493 programmed, at the discretion of the administrator.

494 Finally, it is important to note that in SIP, registration is used for routing incoming SIP requests and  
495 has no role in authorizing outgoing requests. Authorization and authentication are handled in SIP either  
496 on a request-by-request, challenge/response mechanism, or using a lower layer scheme as discussed in  
497 Section 20.

498 The complete set of SIP message details for this registration example is in Section 24.2.

499 Additional operations in SIP, such as querying for the capabilities of a SIP server or client using **OP-**  
500 **TIONS**, canceling a pending request using **CANCEL**, or supporting reliability of provisional responses  
501 using **PRACK** will be introduced in later sections.

## 5 Structure of the Protocol

The SIP protocol is structured as a layered protocol, which means that its behavior is described in terms of a set of fairly independent processing stages, with only a loose coupling between each stage. The structuring of the protocols into layers is for the purpose of presentation and conciseness; it allows the grouping of functions common across elements into a single place. It does not dictate an implementation in any way. When we say that an element “contains” a layer, that means it is compliant to the set of rules defined by that layer.

Not every element specified by the protocol contains every layer. Furthermore, the elements specified by SIP are logical elements, not physical ones. A physical realization can choose to act as different logical elements, perhaps even on a transaction by transaction basis.

The lowest layer of the SIP protocol is its syntax and encoding. Its encoding is specified using a BNF. The complete BNF is specified in Section 25. However, a basic overview of the structure of a SIP message can be found in Section 7. This section introduces enough of an understanding of the format of a SIP message to facilitate understanding the remainder of the protocol.

The next higher layer is the transport layer. This layer defines how a client takes a request, and physically sends it over the network, and how a response is sent by a server, and then received by a client. All SIP elements contain a transport layer. The transport layer is described in Section 19.

The next higher layer is the transaction layer. Transactions are a fundamental component of SIP. A transaction is a request, sent by a client transaction (using the transport layer), to a server transaction, along with all responses to that request sent from the server transaction back to the client. The transaction layer handles retransmissions, matching of responses to requests, and timeouts. Any task that a UAC wishes to accomplish takes place using a series of transactions. Discussion of transactions can be found in Section 17. User agents contain a transaction layer, as do stateful proxies. Stateless proxies do not contain a transaction layer.

The transaction layer has a client component (referred to as a client transaction), and a server component (referred to as a server transaction), each of which are represented by an FSM that is constructed to process a particular request. The layer on top of the transaction layer is called the transaction user (TU), of which there are several types. When a TU wishes to send a request, it creates a client transaction instance and passes it the request, along with the destination IP address, port, and transport to send the request to.

SIP provides the ability for a transaction to be canceled by the client which initiated it. When a client cancels a transaction, it requests that the server give up on further processing, revert to the state that existed before the transaction was initiated, and generate a specific error response to that transaction. This is done with a CANCEL request, which constitutes its own transaction, but references the transaction to be cancelled. Cancellation is described in Section 9.

There are several different types of transaction users. A UAC contains a UAC core, a UAS contains a UAS core, and a proxy contains a proxy core. The behavior of the UAC and UAS cores depend largely on the method. However, there are some common rules for all methods. These rules are captured in Section 8. The primarily deal with construction of a request, in the case of a UAC, and processing of that request, and generation of a response, in the case of a UAS.

UAC and UAS core behavior for the REGISTER method is described in Section 10. Registrations play an important role in SIP. In fact, a UAS that handles a REGISTER is given a special name - a registrar - and it is described in that section.

UAC and UAS core behavior for the OPTIONS method, used for determining the capabilities of a UAC, are described in Section 11.

546 Certain other requests are sent within a *dialog*. A dialog is a peer-to-peer SIP relationship between a  
547 two user agents that persists for some time. The dialog facilitates sequencing of messages between the user  
548 agents, and proper routing of requests between both them. One way to setup a dialog is with the INVITE  
549 method. When a UAC sends a request that is within the context of a dialog, it follows the common UAC  
550 rules as discussed in Section 8, but also the rules for mid-dialog requests. Section 12 discusses dialogs,  
551 and presents the procedures for their construction, and maintenance, in addition to construction of requests  
552 within a dialog.

553 The UAS core can generate provisional responses to requests, which are responses that provide ad-  
554 ditional information about the request processing, but do not indicate completion. Normally, provisional  
555 responses are not transmitted reliably. However, an optional mechanism exists for them to be transmitted  
556 reliably. This mechanism makes use of a method called PRACK, sent as a separate transaction within the  
557 dialog between the UAC and UAS, which is used to acknowledge a reliable provisional response.

558 The most important method in SIP is the INVITE method, which is used to establish a session between  
559 participants. A session is a collection of participants, and streams of media between them, for the purposes  
560 of communication. Section 13 discusses how sessions are initiated, resulting in one or more SIP dialogs.  
561 Section 14 discusses how characteristics of that session are modified, through the use of an INVITE request  
562 within a dialog. Finally, section 15 discusses how a session is terminated.

563 The procedures of Sections 8, 10, 11, 12, 13, 14, and 15 deal entirely with the UA core. Section 16  
564 discusses the proxy element, which facilitates routing of messages between user agents.

## 565 6 Definitions

566 This specification uses a number of terms to refer to the roles played by participants in SIP communications.  
567 The definitions of client, server and proxy are similar to those used by the Hypertext Transport Protocol  
568 (HTTP) (RFC 2616 [9]). The terms and generic syntax of URI and URL are defined in RFC 2396 [10]. The  
569 following terms have special significance for SIP.

570 **Back-to-Back user agent:** A back-to-back user agent (B2BUA) is a logical entity that receives a request  
571 and processes it as an user agent server (UAS). In order to determine how the request should be  
572 answered, it acts as an user agent client (UAC) and generates requests. Unlike a proxy server, it  
573 maintains dialog state and must participate in all requests sent on the dialogs it has established. Since  
574 it is a concatenation of a UAC and UAS, no explicit definitions are needed for its behavior.

575 **Call:** A call is an informal term that refers to a dialog between peers generally set up for the purposes of a  
576 multimedia conversation.

577 **Call leg:** Another name for a dialog.

578 **Call stateful:** A proxy is call stateful if it retains state for a dialog from the initiating INVITE to the termi-  
579 nating BYE request. A call stateful proxy is always stateful, but the converse is not true.

580 **Client:** A client is any network element that sends SIP requests and receives SIP responses. Clients may or  
581 may not interact directly with a human user. *User agent clients* and *proxies* are clients.

582 **Conference:** A multimedia session (see below) that contains multiple participants.

- 583 **Dialog:** A dialog is a peer-to-peer SIP relationship between a UAC and UAS that persists for some time.  
584 A dialog is established by SIP messages, such as a 2xx response to an INVITE request. A dialog is  
585 identified by a call identifier, local address, and remote address. A dialog was formerly known as a  
586 call leg in RFC 2543.
- 587 **Downstream:** A direction of message forwarding within a transaction that refers to the direction that re-  
588 quests flow from the user agent client to user agent server.
- 589 **Final response:** A response that terminates a SIP transaction, as opposed to a *provisional response* that  
590 does not. All 2xx, 3xx, 4xx, 5xx and 6xx responses are final.
- 591 **Informational Response:** Same as a provisional response.
- 592 **Initiator, calling party, caller:** The party initiating a session with an INVITE request. A caller retains this  
593 role from the time it sends the INVITE until the termination of any dialogs established by the INVITE.
- 594 **Invitation:** An INVITE request.
- 595 **Invitee, invited user, called party, callee:** The party that receives an INVITE request for the purposes of  
596 establishing a new session. A callee retains this role from the time it receives the INVITE until the  
597 termination of the dialog established by that INVITE.
- 598 **Location server:** See *location service*.
- 599 **Location service:** A location service is used by a SIP redirect or proxy server to obtain information about  
600 a callee's possible location(s). It is an abstract database, sometimes referred to as a location server.  
601 The contents of the database can be populated in many ways, including being written by registrars.
- 602 **Loop:** A request that arrives at a proxy, is forwarded, and later arrives back at the same proxy. When it  
603 arrives the second time, its Request-URI is identical to the first time, and other headers that affect  
604 proxy operation are unchanged, so that the proxy would make the same processing decision on the  
605 request it made the first time around. Looped requests are errors, and the procedures for detecting  
606 them and handling them are described by the protocol.
- 607 **Method:** The method is the primary function that a request is meant to invoke on a server. The method is  
608 carried in the request message itself. Example methods are INVITE and BYE.
- 609 **Outbound proxy:** A *proxy* that receives all requests from a client, even though it is not the server resolved  
610 by the Request-URI. The outbound proxy sends these requests, after any local processing, to the  
611 address indicated in the Request-URI, or to another outbound proxy.
- 612 **Parallel search:** In a parallel search, a proxy issues several requests to possible user locations upon receiv-  
613 ing an incoming request. Rather than issuing one request and then waiting for the final response before  
614 issuing the next request as in a *sequential search*, a parallel search issues requests without waiting for  
615 the result of previous requests.
- 616 **Provisional response:** A response used by the server to indicate progress, but that does not terminate a SIP  
617 transaction. 1xx responses are provisional, other responses are considered *final*. Normally, provisional  
618 responses are not sent reliably. A provisional response that is sent reliably is referred to as a *reliable*  
619 *provisional response*.



- 620 **Proxy, proxy server:** An intermediary entity that acts as both a server and a client for the purpose of making  
621 requests on behalf of other clients. A proxy server primarily plays the role of routing, which means  
622 its job is to ensure that a request is passed on to another entity that can further process the request.  
623 Proxies are also useful for enforcing policy and for firewall traversal. A proxy interprets, and, if  
624 necessary, rewrites parts of a request message before forwarding it.
- 625 **Registrar:** A registrar is a server that accepts REGISTER requests, and places the information it receives  
626 in those requests into the location service for the domain it handles.
- 627 **Regular Transaction:** A regular transaction is any transaction with a method other than INVITE, ACK, or  
628 CANCEL.
- 629 **Reliable Provisional Response:** A provisional response that is sent reliably from the UAS to UAC.
- 630 **Ringback:** Ringback is the signaling tone produced by the calling party's application indicating that a  
631 called party is being alerted (ringing).
- 632 **Server:** A server is a network element that receives requests in order to service them and sends back re-  
633 sponses to those requests. Examples of servers are proxies, user agent servers, redirect servers, and  
634 registrars.
- 635 **Sequential search:** In a sequential search, a proxy server attempts each contact address in sequence, pro-  
636 ceeding to the next one only after the previous has generated a non-2xx final response.
- 637 **Session:** From the SDP specification: "A multimedia session is a set of multimedia senders and receivers  
638 and the data streams flowing from senders to receivers. A multimedia conference is an example of a  
639 multimedia session." (RFC 2327 [6]) (A session as defined for SDP can comprise one or more RTP  
640 sessions.) As defined, a callee can be invited several times, by different calls, to the same session.  
641 If SDP is used, a session is defined by the concatenation of the *user name*, *session id*, *network type*,  
642 *address type*, and *address* elements in the origin field.
- 643 **(SIP) transaction:** A SIP transaction occurs between a client and a server and comprises all messages from  
644 the first request sent from the client to the server up to a final (non-1xx) response sent from the server  
645 to the client, and the ACK for the response in the case the response was a non-2xx. The ACK for a  
646 2xx response is a separate transaction.
- 647 **Spiral:** A spiral is a SIP request that is routed to a proxy, forwarded onwards, and arrives once again at that  
648 proxy, but this time, differs in a way that will result in a different processing decision than the original  
649 request. Typically, this means that the request's Request-URI differs from its previous arrival. A  
650 spiral is not an error condition, unlike a loop.
- 651 **Stateful proxy:** A logical entity that maintains the client and server transaction state machines defined by  
652 this specification during the processing of a request. Also known as a transaction stateful proxy. The  
653 behavior of a stateful proxy is further defined in Section 16. A stateful proxy is not the same as a call  
654 stateful proxy.
- 655 **Stateless proxy:** A logical entity that does not maintain the client or server transaction state machines  
656 defined in this specification when it processes requests. A stateless proxy forwards every request it  
657 receives downstream and every response it receives upstream.

658 **Transaction User (TU):** The layer of protocol processing that resides above the transaction layer. Trans-  
659 action users include the UAC core, UAS core, and proxy core.

660 **Upstream:** A direction of message forwarding within a transaction that refers to the direction that responses  
661 flow from the user agent server to user agent client.

662 **URL-encoded:** A character string encoded according to RFC 1738, Section 2.2 [11].

663 **User agent client (UAC):** A user agent client is a logical entity that creates a new request, and then uses  
664 the client transaction state machinery to send it. The role of UAC lasts only for the duration of that  
665 transaction. In other words, if a piece of software initiates a request, it acts as a UAC for the duration  
666 of that transaction. If it receives a request later on, it assumes the role of a user agent server for the  
667 processing of that transaction.

668 **UAC Core:** The set of processing functions required of a UAC that reside above the transaction and trans-  
669 port layers.

670 **User agent server (UAS):** A user agent server is a logical entity that generates a response to a SIP request.  
671 The response accepts, rejects or redirects the request. This role lasts only for the duration of that  
672 transaction. In other words, if a piece of software responds to a request, it acts as a UAS for the  
673 duration of that transaction. If it generates a request later on, it assumes the role of a user agent client  
674 for the processing of that transaction.

675 **UAS Core:** The set of processing functions required at a UAS that reside above the transaction and transport  
676 layers.

677 **User agent (UA):** A logical entity that can act as both a user agent client and user agent server for the  
678 duration of a dialog.

679 The role of UAC and UAS as well as proxy and redirect servers are defined on a transaction-by-  
680 transaction basis. For example, the user agent initiating a call acts as a UAC when sending the initial  
681 INVITE request and as a UAS when receiving a BYE request from the callee. Similarly, the same software  
682 can act as a proxy server for one request and as a redirect server for the next request.

683 Proxy, location, and registrar servers defined above are *logical* entities; implementations MAY combine  
684 them into a single application.

## 685 7 SIP Messages

686 SIP is a text-based protocol and uses the ISO 10646 character set in UTF-8 encoding (RFC 2279 [12]).

687 A SIP message is either a request from a client to a server, or a response from a server to a client.

688 Both Request (section 7.1) and Response (section 7.2) messages use the **generic-message** format  
689 of RFC 822 [13]. Both types of messages consist of a **start-line**, one or more header fields (also known as  
690 “headers”), an empty line indicating the end of the header fields, and an optional **message-body**.

```
691 generic-message = start-line  
                  *message-header  
                  CRLF  
                  [ message-body ]
```

692 The start-line, each message-header line, and the empty line **MUST** be terminated by a carriage-return  
693 line-feed sequence (CRLF). Note that the empty line **MUST** be present even if the message-body is not.

694 Except for the above difference in character sets, much of SIP's message and header field syntax is  
695 identical to HTTP/1.1. Rather than repeating the syntax and semantics here we use [HX.Y] to refer to  
696 Section X.Y of the current HTTP/1.1 specification (RFC 2616 [9]).

697 Note, however, that SIP is not an extension of HTTP.

## 698 7.1 Requests

699 SIP Requests are distinguished by having a Request-Line for a start-line. A Request-Line begins with  
700 a method token, followed by the Request-URI and the protocol version, and ending with CRLF. The ele-  
701 ments are separated by SP characters. No CR or LF are allowed except in the end-of-line CRLF sequence.  
702 No LWS is allowed in any of the elements.

703 Method Request-URI SIP-Version

### 704 • Method

705 This specification defines seven methods : REGISTER for registering contact information, INVITE,  
706 ACK, PRACK and CANCEL for setting up sessions, BYE for terminating sessions and OPTIONS  
707 for querying servers about their capabilities. SIP extensions may define additional methods.

### 708 • Request-URI

709 The Request-URI is a SIP URI as described in Section 21.1 or a general URI (RFC 2396 [10]). It  
710 indicates the user or service to which this request is being addressed. The Request-URI **MUST NOT**  
711 contain unescaped spaces or control characters and **MUST NOT** be enclosed in "<>".

712 SIP servers **MAY** support Request-URIs with schemes other than "sip", for example the "tel" URI  
713 scheme of RFC 2806 [14]. It **MAY** translate non-SIP URIs using any mechanism at its disposal,  
714 resulting in either a SIP URI or some other scheme.

### 715 • SIP Version

716 Both request and response messages include the version of SIP in use, and follow [H3.1] (with HTTP  
717 replaced by SIP, and HTTP/1.1 replaced by SIP/2.0) regarding version ordering, compliance require-  
718 ments, and upgrading of version numbers. To be compliant with this specification, applications send-  
719 ing SIP messages **MUST** include a SIP- Version of "SIP/2.0". The string is case-insensitive, but  
720 implementations **MUST** send upper-case.

721 Unlike HTTP/1.1, SIP treats the version number as a literal string. In practice, this should make no  
722 difference.

## 723 7.2 Responses

724 SIP Responses are distinguished by having a Status-Line for a start-line. A Status-Line, consists of the  
725 protocol version followed by a numeric Status-Code and its associated textual phrase, with each element  
726 separated by SP characters. No CR or LF is allowed except in the final CRLF sequence.

727 SIP-version Status-Code Reason-Phrase

728 The **Status-Code** is a 3-digit integer result code that indicates the outcome of an attempt to understand  
729 and satisfy a request. The **Reason-Phrase** is intended to give a short textual description of the **Status-**  
730 **Code**. The **Status-Code** is intended for use by automata, whereas the **Reason-Phrase** is intended for the  
731 human user. A client is not required to examine or display the **Reason-Phrase**.

732 The first digit of the **Status-Code** defines the class of response. The last two digits do not have any  
733 categorization role. For this reason, any response with a status code between 100 and 199 is referred to as  
734 a “1xx response”, any response with a status code between 200 and 299 as a “2xx response”, and so on.  
735 SIP/2.0 allows 6 values for the first digit:

736 **1xx**: Provisional – request received, continuing to process the request;

737 **2xx**: Success – the action was successfully received, understood, and accepted;

738 **3xx**: Redirection – further action needs to be taken in order to complete the request;

739 **4xx**: Client Error – the request contains bad syntax or cannot be fulfilled at this server;

740 **5xx**: Server Error – the server failed to fulfill an apparently valid request;

741 **6xx**: Global Failure – the request cannot be fulfilled at any server.

742 Full definitions of these classes and each registered code appear in Section 23.

## 743 7.3 Header Fields

744 SIP header fields are similar to HTTP header fields in both syntax and semantics. In particular, SIP header  
745 fields follow the [H4.2] definitions of syntax for message-header, the rules for extending header fields over  
746 multiple lines, the use of multiple message-header fields with the same field-name, and the rules regarding  
747 ordering of header fields.

### 748 7.3.1 Header Field Format

749 Header fields follow the same generic header format as that given in Section 3.1 of RFC 822 [13]. Each  
750 header field consists of a field name followed by a colon (":") and the field value.

751 field-name: field-value

752 Note that the formal grammar for a **message-header** specified in Section 25 allow for an arbitrary amount  
753 of whitespace on either side of the colon. No space before the colon and a single space (SP) between the  
754 colon and the field-value is preferred. That is,

755 Subject: lunch

756 Subject : lunch

757 Subject :lunch

758 Subject: lunch

759 are all valid, and equivalent, but the last is the preferred form.

760 Header fields can be extended over multiple lines by preceding each extra line with at least one SP or  
761 horizontal tab (HT). The line break and the whitespace at the beginning of the next line are treated as a  
762 single SP character. Thus the following are equivalent:

763 Subject: I know you're there, pick up the phone and talk to me!  
764 Subject: I know you're there,  
765       pick up the phone  
766       and talk to me!

767       The relative order of header fields with different field names is not significant. The relative order of those  
768 with the same field name is important. Multiple header fields with the same field-name may be present in a  
769 message if and only if the entire field-value for that header field is defined as a comma-separated list (i.e.,  
770 #(values)). It MUST be possible to combine the multiple header fields into one "field-name: field-value"  
771 pair, without changing the semantics of the message, by appending each subsequent field-value to the first,  
772 each separated by a comma.

773       Implementations MUST be able to process multiple header fields with the same name in any combination  
774 of the single-value-per-line or comma-separated value forms.

775       The following blocks of headers are valid and equivalent:

```
776 Route: <sip:alice@atlanta.com>
777 Subject: Lunch
778 Route: <sip:bob@biloxi.com>
779 Route: <sip:carol@chicago.com>
780
781 Route: <sip:alice@atlanta.com>, <sip:bob@biloxi.com>
782 Route: <sip:carol@chicago.com>
783 Subject: Lunch
784
785 Subject: Lunch
786 Route: <sip:alice@atlanta.com>, <sip:bob@biloxi.com>, <sip:carol@chicago.com>
```

787       Each of the following blocks is valid but not equivalent to the others:

```
788 Route: <sip:alice@atlanta.com>
789 Route: <sip:bob@biloxi.com>
790 Route: <sip:carol@chicago.com>
791
792 Route: <sip:bob@biloxi.com>
793 Route: <sip:alice@atlanta.com>
794 Route: <sip:carol@chicago.com>
795
796 Route: <sip:alice@atlanta.com>, <sip:carol@chicago.com>, <sip:bob@biloxi.com>
```

797       The format of a header field-value is defined per header-name. It will always be either an opaque  
798 sequence of TEXT-UTF8 octets, or a combination of whitespace, tokens, separators, and quoted strings.  
799 Many of them will adhere to the general form of a value followed by a semi-colon separated sequence of  
800 parameter-name, parameter-value pairs:

801               field-name: field-value \*(;parameter-name=parameter-value)

802 When comparing headers, field names are always case-insensitive. Unless otherwise stated in the def-  
803 inition of a particular header field, field values, parameter names, and parameter values (tokens in general)  
804 are case-insensitive. Unless specified otherwise, values expressed as quoted strings are case-sensitive.

805 The following are equivalent:

```
806 Contact: <sip:alice@atlanta.com>;expires=3600  
807 CONTACT: <sip:alice@atlanta.com>;EXPIRES=3600  
808  
809 Content-Disposition: session;handling=optional  
810 content-disposition: Session;HANDLING=OPTIONAL  
811
```

812 The following are not equivalent;

```
813 Warning: 370 devnull "Choose a bigger pipe"  
814 Warning: 370 devnull "CHOOSE A BIGGER PIPE"
```

### 815 7.3.2 Header Field Classification

816 Some header fields only make sense in requests or responses. These are called Request Header Fields and  
817 Response Header fields respectively. Those header fields that can appear in either a request or response are  
818 called General Header Fields. If a header appears in a message not matching its category (such as a request  
819 header in a response), it **MUST** be ignored. Section 22 defines the classification of each header.

### 820 7.3.3 Compact Form

821 SIP provides a mechanism to represent common header fields in an abbreviated form. This may be useful  
822 when messages would otherwise become too large to be carried on the transport available to it (exceeding  
823 the MTU when using UDP for example). These compact forms are defined in Section 22. A compact form  
824 **MAY** be substituted for the longer form of a header name at any time without changing the semantics of a  
825 the message. Multiple header fields in a message with the same header name **MAY** appear with an arbitrary  
826 mix of its long and short field name form. Implementations **MUST** accept both the long and short forms of  
827 each header name.

## 828 7.4 Bodies

829 Requests, including new requests defined in extensions to this specification, **MAY** contain message bodies  
830 unless otherwise noted.

831 For response messages, the request method and the response status code determine the type and inter-  
832 pretation of any message body. All responses **MAY** include a body.

### 833 7.4.1 Message Body Type

834 The Internet media type of the message body **MUST** be given by the **Content-Type** header field. If the body  
835 has undergone any encoding (such as compression) then this **MUST** be indicated by the **Content-Encoding**  
836 header field, otherwise **Content-Encoding** **MUST** be omitted. If applicable, the character set of the message  
837 body is indicated as part of the **Content-Type** header-field value.

838 The “multipart” MIME type defined in RFC 2046 [15] MAY be used within the body of the message.  
839 Implementations that send requests containing multipart message bodies MUST be able to send a session  
840 description as a non-multipart message body if the remote implementation requests this through an Accept  
841 header field.

#### 842 7.4.2 Message Body Length

843 The body length in bytes is provided by the Content-Length header field. Section 22.14 describes the  
844 necessary contents of this header in detail.

845 The “chunked” transfer encoding of HTTP/1.1 MUST NOT be used for SIP. (Note: The chunked encoding  
846 modifies the body of a message in order to transfer it as a series of chunks, each with its own size indicator.)

### 847 7.5 Framing SIP messages

848 Unlike HTTP, SIP MAY use UDP or other unreliable datagram protocols. Each such datagram carries one  
849 request or response. Datagrams, including all headers, SHOULD NOT be larger than the path maximum  
850 transmission unit (MTU) if the MTU is known, or 1500 bytes if the MTU is unknown. However, implemen-  
851 tations MUST be able to handle messages up to the maximum datagram packet size. For UDP, this size is  
852 65,535 bytes, including headers.

853 The MTU of 1500 bytes accommodates encapsulation within the “typical” ethernet MTU without IP fragmen-  
854 tation. Recent studies [16, p. 154] indicate that an MTU of 1500 bytes is a reasonable assumption. The next lower  
855 common MTU values are 1006 bytes for SLIP and 296 for low-delay PPP (RFC 1191 [17]). Thus, another reason-  
856 able value would be a message size of 950 bytes, to accommodate packet headers within the SLIP MTU without  
857 fragmentation.

858 In the interest of robustness, any leading empty line(s) MUST be ignored. In other words, if the Request  
859 or Response message begins with one or more CRLF, CR, or LFs, these characters MUST be ignored.

860 Likewise, Implementations processing SIP messages over stream oriented transports MUST ignore noise  
861 between messages.

## 862 8 General User Agent Behavior

863 A user agent represents an end system. It contains a User Agent Client (UAC), which generates requests,  
864 and a User Agent Server (UAS) which responds to them. A UAC is capable of generating a request based on  
865 some external stimulus (the user clicking a button, or a signal on a PSTN line), and processing a response.  
866 A UAS is capable of receiving a request, and generating response, based on user input, external stimulus,  
867 the result of a program execution, or some other mechanism.

868 When a UAC sends a request, it will pass through some number of proxy servers, which forward the  
869 request towards the UAS. When the UAS generates a response, the response is forwarded towards the UAC.

870 UAC and UAS procedures depend strongly on two factors. First, whether the request or response is  
871 inside or outside of a dialog, and second, based on the method of a request. Dialogs are discussed thoroughly  
872 in Section 12; they represent a peer-to-peer relationship between user agents, and are established by specific  
873 SIP methods, such as INVITE.

874 In this section, we discuss the method independent rules for UAC and UAS behavior when processing  
875 of requests that are outside of a dialog. This includes, of course, the requests which themselves establish a  
876 dialog.

## 877 8.1 UAC Behavior

### 878 8.1.1 Generating the Request

879 A valid SIP request formulated by a UAC **MUST** at a minimum contain the following headers: **To**, **From**,  
880 **CSeq**, **Call-ID**, and **Via**; all of these headers are mandatory in all SIP messages. These five headers are  
881 the fundamental building blocks of a SIP message, as they jointly provide for most of the critical message  
882 routing services including the addressing of messages, the routing of responses, ordering of messages, and  
883 the unique identification of transactions.

884 Examples of requests send outside of a dialog include an **INVITE** to establish a session (Section 13) and  
885 an **OPTIONS** to query for capabilities (Section 11).

886 **8.1.1.1 To** The **To** general-header field first and foremost specifies the desired “logical” recipient of the  
887 request, or the address of record of the user or resource that is the target of this request. This may or may  
888 not be the ultimate recipient of the request. The **To** header **MAY** contain a SIP URI, but it may also make  
889 use of other URI schemes (for example as the tel URL [14]) when appropriate. The **To** header field allows  
890 for a display name; this is meant to contain a descriptive version of the URI, and is intended to be displayed  
891 to a user interface.

892 A UAC may learn how to populate the **To** header field for a particular request in a number of ways.  
893 Usually the user will suggest the **To** header field through a human interface, perhaps inputting the URI  
894 manually or selecting it from some sort of address book.

895 A request outside of a dialog **MUST NOT** contain a tag; the tag in the **To** field of a request identifies the  
896 peer of the dialog. Since no dialog is established, no tag is present.

897 For further information on the **To** header see Section 22.39.

898 The following is an example of valid **To** header:

899 **To:** Carol <sip:carol@chicago.com>

900 **8.1.1.2 From** The **From** general-header field indicates the logical identity of the initiator of the request,  
901 possibly the user’s address of record. Like the **To** field, it contains a URI and optionally a display name.  
902 It is used by SIP elements to determine processing rules to apply to a request (for example, automatic call  
903 rejection). As such, it is very important that the URI not contain IP addresses or host names, since these are  
904 not logical names.

905 The **From** header field allows for a display name; this is meant to contain a descriptive version of the  
906 URI, and is intended to be displayed to a user interface. A UAC **SHOULD** use the display name “Anonymous”  
907 if the identity of the client is to remain hidden.

908 Usually the value that populates the **From** header field in requests generated by a particular user agent  
909 is pre-provisioned by the user or by the administrators of the user’s local domain. If a particular user agent  
910 is used by multiple users, it might have switchable profiles that include a URI corresponding to the identity  
911 of the profiled user. Recipients of requests can authenticate the originator of a request in order to ascertain  
912 that they are who their **From** header field claims they are (see Section 20.2 for more on authentication).

913 The **From** field **MUST** contain a new “tag” parameter, chosen by the UAC. See Section 21.3 for details  
914 on choosing a tag.

915 For further information on the **From** header see Section 22.20.

916 Examples:



917 From: "Bob" <sip:bob@biloxi.com> ;tag=a48s  
918 From: sip:+12125551212@server.phone2net.com;tag=887s  
919 From: Anonymous <sip:c8oqz84zk7z@privacy.org>;tag=hyh8

920 **8.1.1.3 Call-ID** The Call-ID general-header field acts as a unique identifier to group together series of  
921 messages. It is always the same for all requests and responses sent by either UA in a dialog. It is also the  
922 same in each registration from a UA within a single boot cycle.

923 In a new request created by a UAC outside of any dialog, unless overridden by method specific behavior,  
924 it MUST be selected by the UAC as a globally unique identifier over space and time; all SIP user agents  
925 must have a means to guarantee that the Call-ID headers they produce will not be inadvertently generated  
926 by any other user agent.

927 Use of cryptographically random identifiers [18] in the generation of Call-IDs is RECOMMENDED. Im-  
928 plementations MAY use the form "localid@host". Call-IDs are case-sensitive and are simply compared  
929 byte-by-byte.

930 Using cryptographically random identifiers provides some protection against session hijacking, and reduces the  
931 likelihood of unintentional Call-ID collisions.

932 No provisioning or human interface is required for the selection of the Call-ID header field value for a  
933 request.

934 For further information on the Call-ID header see Section 22.8.

935 Example:

936 Call-ID: f81d4fae-7dec-11d0-a765-00a0c91e6bf6@foo.bar.com

937 **8.1.1.4 CSeq** The Cseq header serves as a way to identify and order transactions. It consists of a  
938 sequence number and a method. The method MUST match that of the request. The sequence number value  
939 is arbitrary, but MUST be expressible as a 32-bit unsigned integer and MUST be less than 2\*\*31.

940 As long as it follows the above guidelines, a client may use any mechanism it would like to select CSeq  
941 header field values.

942 For further information on the CSeq header see Section 22.16.

943 Example:

944 CSeq: 4711 INVITE

945 **8.1.1.5 Via** The Via header is used to determine the transport to use for sending a request, and for  
946 identifying the IP address and port where the response is to be sent. Rules for setting and using the values  
947 in this header are described in Section 19.

948 For further information on the Via header see Section 22.42.

949 **8.1.1.6 Contact** The Contact header provides a SIP URI that can be used to contact that specific in-  
950 stance of the user agent for subsequent requests. The Contact header MUST be present in any request that  
951 can result in the establishment of a dialog. For the methods defined in this specification, that includes only  
952 the INVITE request. For these requests, the scope of the Contact is the dialog. That is, the Contact header  
953 refers to the URL that the UA would like to receive requests at, for requests that are part of that dialog only.  
954 Only a single URI MUST be present.

955 For further information on the Contact header, see Section 22.10.

956 **8.1.1.7 Request-URI** The initial Request-URI of the message SHOULD be set to the value of the URI  
957 in the To field. One notable exception is the REGISTER method; behavior for setting the Request-URI  
958 of register is given in Section 10. Another exception is the case of pre-existing Route headers; in that case,  
959 the procedures of Section 12.2.1.1 as they pertain to the Request- URI are followed, even though there is  
960 no dialog.

961 **8.1.1.8 Supported and Require** If the UAC supports extensions to SIP that can be applied by the  
962 server to the response, the UAC SHOULD include a Supported header in the request listing the option tags  
963 for those extensions. This includes support for reliability for provisional responses, which is an extension  
964 even though it is defined within this specification. The option tag for reliability of provisional responses is  
965 100rel.

966 The option-tags listed MUST only refer to extensions defined in standards track RFCs. This is to prevent  
967 servers from insisting that clients implement non-standard, vendor defined features in order to receive ser-  
968 vice. Extensions defined by experimental and informational RFCs are explicitly excluded from usage with  
969 the Supported header in a request, since they too are often used to document vendor defined extensions.

970 If the UAC wishes to insist that a UAS understand an extension that the UAC will apply to the request in  
971 order to process the request, it MUST insert a Require header into the request listing the option tag for that  
972 extension. If the UAC wishes to apply an extension to the request and insist that a proxy understand that  
973 extension, it MUST insert a Proxy-Require header into the request listing the option tag for that extension.

974 A Require header in a request with the option tag 100rel means that the UAC wishes for all provi-  
975 sional responses to this request to be transmitted reliably. This header MUST NOT be present in any requests  
976 excepting INVITE, although extensions to SIP may allow its usage with other request methods.

977 **8.1.1.9 Additional Message Components** After a new request has been created, the headers described  
978 above have been properly constructed, any additional optional headers are added, as are any headers specific  
979 to the method.

980 SIP requests MAY contain a MIME-encoded message-body. Regardless of the type of body that a request  
981 contains, certain headers must be formulated to characterize the contents of the body. For further information  
982 on these headers see Section 7.4.

## 983 **8.1.2 Sending the Request**

984 The destination for the request is then computed. This can be a preconfigured IP address, port and transport  
985 of an outbound proxy, or it can be determined through DNS procedures applied to the Request-URI. These  
986 procedures are described in [8], which yield an ordered set of address, port and transports to attempt.

987 The UAC SHOULD follow the procedures defined there for stateful elements, trying each address until a  
988 server is contacted. Each try constitutes a new transaction, and therefore a new client transaction MUST be  
989 constructed for each.

## 990 **8.1.3 Processing Responses**

991 Responses are first processed by the transport layer, and then passed up to the transaction layer. The trans-  
992 action layer performs its processing, and then passes it up to the TU. The majority of response processing  
993 in the TU is method specific. However, there are some general behaviors independent of the method.

994 **8.1.3.1 Unrecognized Responses** A UAC MUST treat any response they do not recognize as being  
995 equivalent to the x00 response code of that class, and MUST be able to process the x00 response code for  
996 all classes. For example, if a UAC receives an unrecognized response code of 431, it can safely assume that  
997 there was something wrong with its request and treat the response as if it had received a 400 (Bad Request)  
998 response code.

999 **8.1.3.2 Vias** If more than one Via header field is present in a response, the UAC SHOULD discard the  
1000 message.

1001         The presence of additional Via header fields that precede the originator of the request suggests that the message  
1002         was misrouted or possibly corrupted.

1003 **8.1.3.3 Processing Reliable 1xx Responses** A 1xx response that contains a Require header with the  
1004 option tag 100rel is a reliable provisional response. The UA core follows the procedures in Section 18.2  
1005 to process the response, which will result in the generation of a PRACK request to acknowledge the reliable  
1006 provisional response.

1007 **8.1.3.4 Processing 3xx responses** Upon receipt of a redirection response (e.g. a 3xx response status  
1008 code), clients SHOULD use the URI(s) in the Contact header field to formulate a new request.

1009         To do that, the client copies all but the “method-param” and “header” elements of the addr-spec part  
1010 of the Contact header field into the Request-URI of the request. It uses the “header” parameters to create  
1011 headers for the request, replacing any default headers normally used.

1012         In all other respects, requests sent upon receipt of a redirect response SHOULD re-use the headers and  
1013 bodies of the original request.

1014         The Contact values present in redirection responses SHOULD NOT be cached across calls, as they may  
1015 not represent the most desirable location for a particular destination address.

1016 **8.1.3.5 Processing 4xx responses** Certain 4xx response codes require specific UA processing, indepen-  
1017 dent of the method.

1018         If a 401 or 407 response is received, the UAC SHOULD follow the authorization procedures of Section  
1019 20.2.2 and Section 20.2.3 to retry the request with credentials.

1020         If a 413 response is received (Section 23.4.11), it means that the request contained a body that was  
1021 longer than the UAS was willing to accept. If possible, the UAC SHOULD retry the request, either omitting  
1022 the body or using one of a smaller length.

1023         If a 415 response is received (Section 23.4.13), it means the request contained media types not supported  
1024 by the UAS. The UAC SHOULD retry sending the request, this time only using content with types listed in  
1025 the Accept header in the response, with encodings listed in the Accept-Encoding header in the response,  
1026 and with languages listed in the Accept-Language in the response.

1027         If a 420 response is received (Section 23.4.14), it means the request contained a Require or Proxy-  
1028 Require header listing an option-tag for a feature not supported by a proxy or UAS. The UAC SHOULD  
1029 retry the request, this time omitting any extensions listed in the Unsupported header in the response.

1030         In all of the above cases, retrying the request is accomplished by creating a new request with the appro-  
1031 priate modifications. This new request SHOULD have the same value of the Call-ID, To, and From of the  
1032 previous request, but the CSeq should contain a new sequence number that is one higher than the previous.

1033         With other 4xx responses, a retry may or may not be possible depending on the method and the use case.

## 1034 8.2 UAS Behavior

1035 When a request outside of a dialog is processed by a UAS, there are a set of processing rules which are  
1036 followed, independent of the method. Section 12 gives guidance on how a UAS can tell whether a request  
1037 is inside or outside of a dialog.

### 1038 8.2.1 Authentication/Authorization

1039 A UAS MAY authenticate the originator of a request, and this process may require the server to issue a  
1040 challenge for credentials. The required behavior is independent of the method of the request, and is detailed  
1041 in Section 20.2.

### 1042 8.2.2 Method Inspection

1043 Once a request is authenticated (or no authentication was desired), the UAS MUST inspect the method of the  
1044 request. If the UAS does not support the method of a request it MUST generate a 405 (Method Not Allowed)  
1045 response. Procedures for generation of responses are described in Section 8.2.7. The UAS MUST also add  
1046 an Allow header to the 405 response. The Allow header field MUST list the set of methods supported by the  
1047 UAS generating the message.

1048 The Allow header is presented in Section 22.5.

1049 If the method is one supported by the server, processing continues.

### 1050 8.2.3 Header Inspection

1051 If a UAS does not understand a header field in a request (i.e. the header is not defined in this specification  
1052 or in any supported extension), the server MUST ignore that header and continue processing the message. A  
1053 UAS SHOULD ignore any malformed headers which are not necessary for processing requests.

1054 **8.2.3.1 To and Request-URI** The To header field identifies the original recipient of the request desig-  
1055 nated by the user identified in the From field. The original recipient may or may not be the UAS processing  
1056 the request, due to call forwarding or other proxy operations. A UAS MAY apply any policy it wishes in  
1057 determination of whether to accept requests when the To field is not the identity of the UAS. However, it is  
1058 RECOMMENDED that a UAS accept requests even if they do not recognize the URI scheme (e.g., a tel:  
1059 URI) in the To header, or if the To header does not address a known or current user of this UAS. If, on the  
1060 other hand, the UAS decides to reject the request, it SHOULD generate a response with a 403 status code and  
1061 send it to the server transaction for transmission.

1062 However, the Request-URI identifies the UAS that is to process the request. If the Request-URI does  
1063 not identify an address that the UAS is willing to accept requests for, it SHOULD reject the request with  
1064 a 404 (Not Found) response. If the Request-URI does not provide sufficient information for the UAS to  
1065 determine whether it is willing to process the request, it SHOULD return a 485 (Ambiguous) response. This  
1066 response SHOULD contain a Contact header field containing URIs of new addresses to be tried. Typically,  
1067 a UA which uses the REGISTER method to bind its address of record to a specific contact address, will see  
1068 requests whose Request-URI equals those contact addresses.

1069 **8.2.3.2 Require** Assuming the UAS decides that it is the proper element to process the request, it ex-  
1070 amines the Require header field, if present.

1071 The **Require** general-header field is used by UAC to tell UAS about SIP extensions that the UAC expects  
1072 the UAS to support in order to properly process the request. If a UAS does not understand an option listed  
1073 in a **Require** header field, it **MUST** respond by generating a response with status code 420 (Bad Extension).  
1074 The UAS **MUST** add a **Unsupported**, and list in it those options it does not understand amongst those in  
1075 the **Require** header of the request. Upon receipt of the 420 the client **SHOULD** retry the request, this time  
1076 without using those extensions listed in the **Unsupported** header in the response.

1077 Example:

```
1078 UACC->UAS:   INVITE sip:watson@bell-telephone.com SIP/2.0
1079              Require: com.example.billing
1080              Payment: sheep_skins, conch_shells
1081
1082 UASS->UAC:   SIP/2.0 420 Bad Extension
1083              Unsupported: com.example.billing
```

1084 This is to make sure that the client-server interaction will proceed without delay when all options are understood  
1085 by both sides, and only slow down if options are not understood (as in the example above). For a well-matched  
1086 client-server pair, the interaction proceeds quickly, saving a round-trip often required by negotiation mechanisms.  
1087 In addition, it also removes ambiguity when the client requires features that the server does not understand. Some  
1088 features, such as call handling fields, are only of interest to end systems.

#### 1089 8.2.4 Content Processing

1090 Assuming the UAS understands any extensions required by the client, the UAS examines the body of the  
1091 message, and the headers that describe it. If there are any bodies whose type (indicated by the **Content-**  
1092 **Type**), language (indicated by the **Content-Language**) or encoding (indicated by the **Content-Encoding**)  
1093 are not understood, and that body part is not optional (as indicated by the **Content-Disposition**) header, the  
1094 UAS **MUST** reject the request with a 415 (Unsupported Media Type) response. The response **MUST** contain  
1095 a **Accept** header listing the types of all bodies it understands, in the event the request contained bodies of  
1096 types not supported by the UAS. If the request contained content encodings not understood by the UAS,  
1097 the response **MUST** contain an **Accept-Encoding** header listing the encodings understood by the UAS. If  
1098 the request contained content with languages not understood by the UAS, the response **MUST** contain an  
1099 **Accept-Language** header indicating the languages understood by the UAS.

1100 Beyond these checks, body handling is method and type specific.

1101 For further information on the processing of Content-specific headers see Section 7.4.

#### 1102 8.2.5 Applying Extensions

1103 A UAS that wishes to apply some extension when generating the response **MUST** only do so if support for  
1104 that extension is indicated in the **Supported** header in the request. If the desired extension is not supported,  
1105 the server **SHOULD** rely only on baseline SIP and any other extensions supported by the client. To ensure  
1106 that the **SHOULD** can be fulfilled, any specification of a new extension **MUST** include discussion of how  
1107 to gracefully return to baseline SIP when the extension is not present. In rare circumstances, where the  
1108 server cannot process the request without the extension, the server **MAY** send a 421 (Extension Required)  
1109 response. This response indicates that the proper response cannot be generated without support of a specific  
1110 extension. The needed extension(s) **MUST** be included in a **Require** header in the response. This behavior  
1111 is **NOT RECOMMENDED**, as it will generally break interoperability.

1112 Any extensions applied to a non-421 response MUST be listed in a **Require** header included in the  
1113 response. Of course, the server MUST NOT apply extensions not listed in the **Supported** header in the  
1114 request. As a result of this, the **Require** header in a response will only ever contain option tags defined in  
1115 standards track RFCs.

### 1116 **8.2.6 Processing the Request**

1117 Assuming all of the checks in the previous subsections are passed, the UAS processing becomes method  
1118 specific. Section 10 deals with the **REGISTER** request, section 11 deals with the **OPTIONS** request,  
1119 section 13 deals with the **INVITE** request, and section 15 deals with the **BYE** request.

### 1120 **8.2.7 Generating the Response**

1121 When a UAS wishes to construct a response to a request, it follows these procedures. Additional procedures  
1122 may be needed depending on the status code of the response and the circumstances of its construction. These  
1123 additional procedures are documented elsewhere.

1124 The **From** field of the response MUST equal the **From** field of the request. The **Call-ID** field of the  
1125 response MUST equal the **Call-ID** field of the request. The **Cseq** field of the response MUST equal the **Cseq**  
1126 field of the request. The **Via** headers in the response MUST equal the **Via** headers in the request, and MUST  
1127 maintain the same ordering.

1128 If a request contained a **To** tag in the request, the **To** field in the response MUST equal that of the request.  
1129 However, if the **To** field in the request did not contain a tag, the URI in the **To** field in the response MUST  
1130 equal the URI in the **To** field in the request. Additionally, the UAS MUST add a tag to the **To** field in the  
1131 response. This serves to identify the UAS that is responding, possibly resulting in a component of a dialog  
1132 ID. The same tag MUST be used for all responses to that request, both provisional and final. Procedures for  
1133 generation of tags are defined in Section 21.3.

### 1134 **8.2.8 Stateless UAS Behavior**

1135 A stateless UAS is a UAS that doesn't maintain transaction state. It replies to requests normally, but discards  
1136 any state that would ordinarily be retained by a UAS after a response has been sent. If a stateless UAS  
1137 receives a retransmission of a request, it regenerates the response and resends it, just as if it were the replying  
1138 to the first instance of the request. Stateless UASs do not use a transaction layer; they receive requests  
1139 directly from the transport layer and send responses directly to the transport layer.

1140 The stateless UAS role is needed primarily to handle unauthenticated requests for which a challenge  
1141 response is issued. If unauthenticated requests were handled statefully, then malicious floods of unauthenti-  
1142 cated requests could create massive amounts of transaction state that might slow or complete halt call pro-  
1143 cessing in a UAS, effectively creating a denial of service condition; for more information see Section 20.4.

1144 The most important behaviors of a stateless UAS are the following:

- 1145 • A stateless UAS MUST NOT send provisional (1xx) responses.
- 1146 • A stateless UAS MUST NOT retransmit responses.
- 1147 • A stateless UAS MUST ignore **ACK** requests.
- 1148 • A stateless UAS MUST ignore **CANCEL** requests.

- 1149 • To header tags **MUST** be generated for responses in a stateless manner - in a manner that will generate  
1150 the same tag for the same request consistently. For information on tag construction see Section 21.3.

1151 In all other respects, a stateless UAS behaves in the same manner as a stateful UAS. A UAS can operate  
1152 in either a stateful or stateless mode for each new request.

### 1153 8.3 Redirect Servers

1154 In some architectures it may be desirable to reduce the processing load on proxy servers that are responsible  
1155 for routing requests by relying on redirection. Redirection allows servers to push routing information for a  
1156 request back in a response to the client, thereby taking themselves out of the loop of further messaging for  
1157 this transaction while still aiding in locating the target of the request. When the originator of the request  
1158 receives the redirection it will send a new request based on the routing information it has received. By  
1159 propagating routing information from the core of the network to its edges, redirection allows for considerable  
1160 network scalability.

1161 A redirect server is logically constituted of a server transaction layer and a transaction user that has  
1162 access to a location service of some kind (see Section 10 for more on registrars and location services). This  
1163 location service is effectively a database containing mappings between a single URI and a set of one or more  
1164 alternative locations at which the target of that URI can be found.

1165 A redirect server does not issue any SIP requests of its own. After receiving a request other than **CAN-**  
1166 **CEL**, the server gathers the list of alternative locations from the location service and either returns a final  
1167 response of class 3xx or it refuses the request. For well-formed **CANCEL** requests, it **SHOULD** return a  
1168 2xx response. This response ends the SIP transaction. The redirect server maintains transaction state for an  
1169 entire SIP transaction. It is the responsibility of clients to detect forwarding loops between redirect servers.

1170 When a redirect server returns a 3xx response to a request, it populates the list of (one or more) alterna-  
1171 tive locations into **Contact** headers. An “**expires**” parameter to the **Contact** header may also be supplied  
1172 to indicate the lifetime of the **Contact** data.

1173 The **Contact** header field contains URIs giving the new locations or user names to try, or may simply  
1174 specify additional transport parameters. A 301 or 302 response may also give the same location and user-  
1175 name that was targeted by the initial request but specify additional transport parameters such as a different  
1176 server or multicast address to try, or a change of SIP transport from UDP to TCP or vice versa.

1177 Note that the **Contact** header field **MAY** also refer to a different entity than the one originally called. For  
1178 example, a SIP call connected to GSTN gateway may need to deliver a special informational announcement  
1179 such as “The number you have dialed has been changed.”

1180 A **Contact** response header field can contain any suitable URI indicating where the called party can be  
1181 reached, not limited to SIP URIs. For example, it could contain URL’s for phones, fax, or irc (if they were  
1182 defined) or a mailto: (RFC 2368, [19]) URL.

1183 The “**expires**” parameter of the **Contact** header field indicates how long the URI is valid. The parameter  
1184 is either a number indicating seconds or a quoted string containing a **SIP-date**. If this parameter is not  
1185 provided, the value of the **Expires** header field determines how long the URI is valid. Implementations  
1186 **MAY** treat values larger than 2\*\*32-1 (4294967295 seconds or 136 years) as equivalent to 2\*\*32-1.

1187 Redirect servers **MUST** ignore features that are not understood (including unrecognized headers, **Re-**  
1188 **quired** extensions, or even method names) and proceed with the redirection of the session in question. If  
1189 a particular extension requires that intermediate devices support it, the extension **MUST** be tagged in the  
1190 **Proxy-Require** field as well (see Section 22.28).

## 9 Canceling a Request

The previous section has discussed general UA behavior for generating requests, and processing responses, for requests of all methods. In this section, we discuss a general purpose method, called CANCEL.

The CANCEL request, as the name implies, is used to cancel a previous request sent by a client. Specifically, it asks the user agent server to cease processing the request, and generate an error response to that request. CANCEL has no effect on a request that has already been responded to. Because of this, it is most useful to CANCEL requests which can take a long time to respond to. For this reason, CANCEL is most useful for INVITE requests, which can take a long time to generate a response. In that usage, a UAS that receives a CANCEL request for an INVITE, but has not yet sent a response, would “stop ringing”, and then respond to the INVITE with a specific error response (a 487).

Cancel requests can be constructed and sent by any type of client, including both proxies and user agent servers. Section 15 discusses under what conditions a UAC would CANCEL an INVITE request, and Section 16 discusses proxy usage of INVITE.

Because a stateful proxy can generate its own CANCEL, a stateful proxy also responds to a CANCEL, rather than simply forwarding a response it would receive from a downstream element. For that reason, CANCEL is referred to as a “hop-by-hop” request, since it is responded to at each stateful proxy hop.

### 9.1 Client Behavior

A CANCEL request SHOULD NOT be sent to cancel a request other than INVITE.

Since requests other than INVITE are responded immediately, sending a CANCEL for a non-INVITE request would always create a race condition.

The following procedures are used to construct a CANCEL request. The Request-URI, Call-ID, To, the numeric part of CSeq and From header fields in the CANCEL request MUST be identical to those in the request being cancelled, including tags. A CANCEL constructed by a client MUST have only a single Via header, whose value matches the top Via in the request being cancelled. Using the same values for these headers allows the CANCEL to be matched with the request it cancels (Section 9.2 indicates how such matching occurs). However, the method part of the CSeq header MUST have a value of CANCEL. This allows it to be identified and processed as a transaction in its own right (See Section 17). If the request being cancelled contained Route header fields the CANCEL request MUST include these Route header fields.

This is needed so that stateless proxies are able to route CANCEL requests properly.

Once the CANCEL is constructed, the client SHOULD check whether any response (provisional or final) has been received for the request being cancelled (herein referred to as the “original request”). The CANCEL request MUST NOT be sent if no provisional response has been received, rather, the client MUST wait for the arrival of a provisional response before sending the request. If the original request has generated a final response, the CANCEL SHOULD NOT be sent, as it is an effective no-op, since CANCEL has no effect on requests which have already generated a final response. When the client decides to send the CANCEL, it creates a client transaction for the CANCEL, and passes it the CANCEL request along with the destination address, port and transport. The destination address, port, and transport for the CANCEL MUST be identical to those used to send the original request.

If it was allowed to send the CANCEL before receiving a response for the previous request the server could receive the CANCEL before the original request.

Note that both the transaction corresponding to the original request and the CANCEL transaction will complete independently. However, a UAC canceling a request cannot rely on receiving a 487 (Request Terminated) response for the original request, as an RFC 2543-compliant UAS will not generate such a



1234 response. If there is no final response for the original request in 64\*T1 seconds for an INVITE transaction,  
1235 and T3 seconds for a non-INVITE transaction, the client SHOULD then consider the original transaction  
1236 cancelled and SHOULD destroy the client transaction handling the original request.

## 1237 9.2 Server Behavior

1238 The CANCEL method requests that the TU at the server side cancel a pending request with the same Call-  
1239 ID, To, From, top Via header and Request-URI and CSeq (sequence number only) header field values.

1240 The processing of a CANCEL request at a server depends on the type of server. A stateless proxy will  
1241 forward it, a stateful proxy might respond to it and generate some CANCEL requests of its own, and a UAS  
1242 will respond to it. See Section 16.8 for proxy treatment of CANCEL.

1243 When a UAS receives a CANCEL, it looks for any server transactions which were created by requests  
1244 with the same To, From, Call-ID, Cseq numeric value, Request-URI and top Via header. If no matching  
1245 transactions are found, the CANCEL SHOULD be responded to with a 481 (Call Leg/Transaction Does Not  
1246 Exist). If the transaction for the original request still exists, the behavior of the UAS on receiving a CANCEL  
1247 request depends on whether it has already sent a final response for original request. If it has, the CANCEL  
1248 request has no effect on the processing of the original request, no effect on any session state, and no effect  
1249 on the responses generated for the original request. If the UAS has not issued a final response for the original  
1250 request, its behavior depends on the method of the original request. If the original request was an INVITE,  
1251 the UAS SHOULD immediately respond to the INVITE with a 487 (Request Terminated). The behavior upon  
1252 reception of a CANCEL request for any other method defined in this spec is effectively no-op. Extensions  
1253 to this spec that define new methods MUST define the behavior of a UAS upon reception of a CANCEL for  
1254 those methods.

1255 Regardless of the method of the original request, the CANCEL request itself is answered with a 200  
1256 (OK) response in either case. Once the response is constructed it is passed to the server transaction for the  
1257 CANCEL request.

## 1258 10 Registrations

### 1259 10.1 Overview of Usage

1260 SIP is a protocol that offers a discovery capability. For one user to initiate a session with another, SIP must  
1261 discover the current host(s) that the called user is reachable at. This discovery process is accomplished  
1262 by SIP proxy servers, which are responsible for receiving a request, determining where to send it based  
1263 on knowledge of the location of the user, and then sending it there. To do this, proxies consult an abstract  
1264 service known as a *location service*, which provides address bindings for a particular domain. These address  
1265 bindings map an incoming SIP URL, sip:bob@Biloxi.com, for example, to one or more SIP URLs  
1266 which are somehow “closer” to the desired user, sip:bob@engineering.Biloxi.com, for example.  
1267 Ultimately, a proxy will consult a location service which maps a received URL to the current host(s) that a  
1268 user is logged in to.

1269 There are many ways by which the contents of the location service can be established. One way is  
1270 administratively. In the above example, Bob is known to be a member of the engineering department through  
1271 access to a corporate database. SIP provides a mechanism, however, for a user agent to explicitly create a  
1272 binding in the location service of a proxy. This mechanism is known as registration.

1273 The process of registration entails sending a REGISTER message to a special type of UAS known as a

1274 registrar. The registrar acts as a front end to the location service for a domain, reading and writing mappings  
 1275 based on the contents of the REGISTER messages. This location service will then be consulted by a proxy  
 1276 server that is responsible for routing requests for that domain.

1277 SIP does not mandate a particular mechanism for implementing the location service. The only require-  
 1278 ment is that a registrar for some domain MUST be capable of reading and writing data to the location service,  
 1279 and a proxy for that domain MUST be capable of reading that same data. A registrar MAY be co-located with  
 1280 a particular SIP proxy server for the same domain, allowing usage of an in memory database for the location  
 1281 service. Usage of a shared database is another implementation choice. The choice depends entirely on the  
 1282 architectural requirements (redundancy, scalability, etc) of a particular deployment.

1283 Registration creates bindings in a location service for a particular domain that associate an “address of  
 1284 record” URI with one or more “contact addresses”. This means that when a proxy for that domain receives a  
 1285 request whose request URI matches the address of record, the proxy will forward the request to the contact  
 1286 addresses registered to that address of record. Generally, it only makes sense to register an address of record  
 1287 at a location service for a domain when requests for that address of record would be routed to that domain.  
 1288 In most cases, this means that the domain of the registration will need to match the domain in the URI of  
 1289 the address of record.

1290 The most important usage of the registration mechanism is to inform a proxy of the mapping between  
 1291 the address of record and the current host on which the UA resides. However, the registration process is a  
 1292 general mechanism for establishing bindings, and can be used for other purposes (for example, to set up call  
 1293 forwarding).

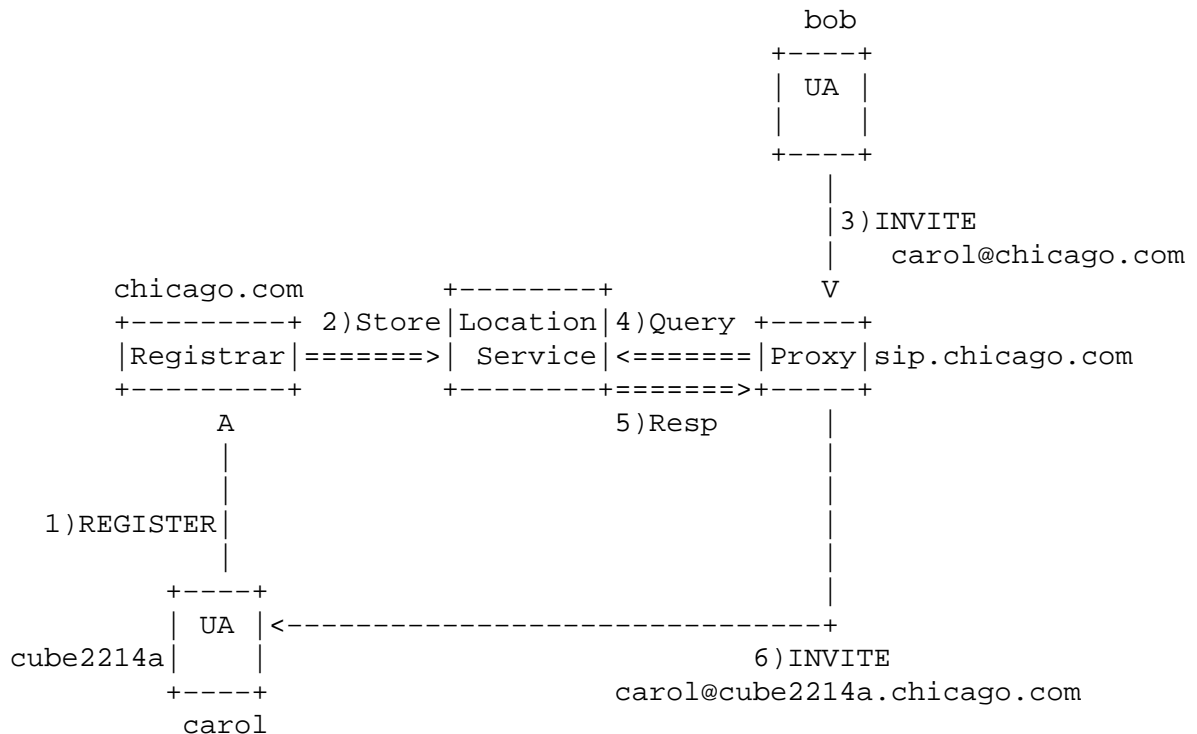


Figure 2: REGISTER example

## 1294 10.2 Construction of the REGISTER request

1295 Several operations can be performed with a REGISTER method with respect to a registrar. One of these is  
1296 the basic registration operation that is described above, which provides a new binding between an address  
1297 of record and one or more contact addresses. Registration on behalf of a particular address of record may be  
1298 performed by a third party if they are authorized to do so. A client may also remove previous bindings, or  
1299 query to determine which bindings are currently in place for an address of record.

1300 Aside from the exceptions noted in this and the following sections, the construction of the REGISTER  
1301 method, and behavior of clients sending a REGISTER is identical to the general UAC behavior described in  
1302 Section 8.1 and Section 17.1. Regardless of the operation that is performed by a REGISTER, the following  
1303 header fields MUST be formulated as follows:

1304 **Request-URI:** The Request-URI names the domain of the location service that the registration is meant  
1305 for (e.g. "chicago.com"). The user name MUST be empty.

1306 **To:** The To header field contains the address of record whose registration is to be created or modified.  
1307 Note that the initial To header field and the Request-URI field SHOULD therefore be different in a  
1308 REGISTER message.

1309 **From:** The From header field contains the address of record of the person responsible for the registration,  
1310 which MAY be identical to the value of the To header field. For third-party registrations the From  
1311 header field and To header field are different.

1312 **Call-ID:** All registrations from a user agent client SHOULD use the same Call-ID header value, at least  
1313 within the same reboot cycle.

1314 If different Call-IDs were used for overlapping REGISTER messages coming from the same client, the  
1315 registrar might have trouble determining their ordering.

1316 **Contact:** REGISTER requests MAY contain one or more Contact header fields. Contact addresses are  
1317 presented in the Contact header fields of REGISTER requests.

1318 Note that user agents MUST NOT send a new registration (containing new Contact header fields, as  
1319 opposed to a retransmission) until they have received a response from the registrar for the previous one.

1320 The following optional Contact header parameters also contain behavior specific to the registration  
1321 process.

1322 **action:** The "action" parameter has been deprecated. UACs SHOULD NOT use the "action" parameter.

1323 **expires:** The "expires" parameter indicates how long the UAC would like the binding to be valid. The  
1324 parameter is either a number indicating seconds or a quoted string containing a SIP-date. If this  
1325 parameter is not provided, the value of the Expires header field determines how long the binding is  
1326 valid. Implementations MAY treat values larger than 2\*\*32-1 (4294967295 seconds or 136 years) as  
1327 equivalent to 2\*\*32-1.

### 1328 10.2.1 Adding Bindings with REGISTER

1329 For a simple registration, a REGISTER request sent to a registrar includes contact addresses to which  
1330 requests should be forward for the originating user's address of record. The address of record itself (i.e.

1331 'sip:carol@chicago.com') MUST populate the To header of the REGISTER. The Contact header fields of  
1332 the request typically contain SIP URIs that identify particular SIP endpoints (i.e. 'sip:carol@cube2214a.chicago.com'),  
1333 but they MAY use any URI scheme; this way a SIP UA can choose to register telephone numbers (with the  
1334 tel URL, [14]) or email addresses (with a mailto URL, [19]) as Contacts for an address of record.

1335 For example, if Carol, whose address of record is 'sip:carol@chicago.com', needed to register, she would  
1336 typically want to register with the registrar associated with the location service of chicago.com. This location  
1337 service would then be accessed by a proxy server that receives requests targeting users in the chicago.com  
1338 domain, and hence new requests for Carol's address of record will be routed to her SIP endpoint.

1339 Once a client has established bindings at a registrar, it MAY send subsequent registrations containing  
1340 new bindings or modifications to pre-existing bindings as necessary. The 2xx response to the REGISTER  
1341 message will contain (in Contact header fields) a complete list of bindings that have been registered for this  
1342 address of record at this registrar.

1343 **10.2.1.1 Setting the Expiration Interval of Contact Addresses** When a client sends a REGISTER  
1344 request, it MAY suggest an expiration interval that indicates how long the client would like the registration  
1345 to be valid (although as is detailed in Section 10.3, the registrar has the ultimate say).

1346 There are two ways in which a client can suggest an expiration interval for a binding: through an Expires  
1347 header, or an "expires" Contact header parameter. The latter allows expiration intervals to be suggested  
1348 on a per-binding basis when more than one binding is given in a single REGISTER, whereas the former  
1349 suggests an expiration interval for all Contact header fields that do not contain the "expires" parameter.

1350 If neither mechanism for expressing a suggested expiration time is present in a REGISTER, a default  
1351 suggestion of one hour is assumed.

1352 **10.2.1.2 Setting Preference among Contact Addresses** If more than one Contact is sent in a REGIS-  
1353 TER, then the registering UA intends to associate all of the URIs given in these Contact headers with the  
1354 address of record present in the To field. This list can be prioritized with the "q" mechanism.

1355 **q:** The "q" parameter indicates a relative preference for the particular Contact header field compared to  
1356 other bindings present in this REGISTER message or existing within the location service of the  
1357 registrar. For an example of how a proxy server uses "q" values, see Section 16.5.

## 1358 **10.2.2 Removing Bindings with REGISTER**

1359 Registrations are removed from the registrar through an expiration process; registrations are soft state and  
1360 need to be refreshed periodically. A client may attempt to influence the expiration intervals selected by the  
1361 registrar as described in Section 10.2.1.

1362 A registering user agent requests the immediate removal of a binding by specifying an expiration in-  
1363 terval of "0" for that contact address in a REGISTER. It is RECOMMENDED that user agents support this  
1364 mechanism so that bindings can be removed (for whatever reason) before their expiration interval has passed.

1365 The REGISTER-specific Contact header field value of "\*" applies to all registrations, but it MUST only  
1366 be used when the Expires header is present with a value of "0".

1367 Use of the "\*" Contact header field value allows a registering user agent to remove all of its bindings expediently.

### 1368 **10.2.3 Fetching Bindings with REGISTER**

1369 If no **Contact** headers are present in a **REGISTER**, then the UA is not in fact registering any new bindings,  
1370 and the list of bindings is therefore left unchanged. As noted above, in a successful response to this **REG-**  
1371 **ISTER** message, the complete list of existing bindings is returned, and thus a **REGISTER** without **Contact**  
1372 headers serves as a fetch operation.

### 1373 **10.2.4 Refreshing Registrations**

1374 When a 2xx response has been received by the client for a **REGISTER** request, the client **MUST** determine  
1375 when each of the bindings enumerated in the response needs to be refreshed. This may include bindings that  
1376 were registered in previous **REGISTER** transactions.

1377 Since the list of bindings returned in the response to a **REGISTER** may contain bindings that were not  
1378 included in this **REGISTER** transaction, the client must correlate **Contact** header fields in the response  
1379 with the **Contact** header fields it sent in the request in order to establish proper expiration timers. This  
1380 correlation should be performed in accordance with the URI comparison rules given in Section 21.1.4.

1381 The registering UA **MUST** re-register each contact address at least as often as the mandated expiration  
1382 interval. A **REGISTER** that refreshes a binding **SHOULD** have the same **Call-ID** as the request which  
1383 created the binding. The **CSeq** header **SHOULD** have a numeric sequence number that is one higher than  
1384 the value sent in the last request with the same **Call-ID**.

1385 Note that a UA **MUST** update its expiration timers for refreshing each binding every time it receives  
1386 a response to a registration request.

1387 Registration refreshes **SHOULD** be sent to the same address as the original registration, unless redirected.

### 1388 **10.2.5 Discovering a Registrar**

1389 Depending on the policy of their administrative domain, SIP UAs can be configured with the address of a  
1390 local registrar. Some UAs may be equipped with protocol tools (outside the scope of SIP) that allow them  
1391 to discover their local registrar dynamically.

1392 Note that as an alternate means of discovering a registrar if no local registrar is configured in the user  
1393 agent, clients **MAY** register via multicast. Multicast registrations are addressed to the well-known "all SIP  
1394 servers" multicast address "sip.mcast.net" (224.0.1.75). This request **MUST** be scoped to ensure it is not  
1395 forwarded beyond the boundaries of the administrative system. This **MAY** be done with either TTL or  
1396 administrative scopes (see [20]), depending on what is implemented in the network. SIP user agents **MAY**  
1397 listen to that address and use it to become aware of the location of other local users (see [21]); however, they  
1398 do not respond to the request.

1399 Multicast registration may be inappropriate in some environments, for example, if multiple businesses share the  
1400 same local area network.

1401 If a SIP UA knows of an appropriate registrar it **SHOULD** attempt to register with this server periodically  
1402 - management of registration intervals is detailed below.

## 1403 **10.3 Processing of REGISTER at the Registrar**

1404 A registrar is a UAS that responds to a **REGISTER** request, and stores the information gathered from that  
1405 request in a location service that is in turn accessible to proxy servers within its administrative domain. A  
1406 registrar handles requests as a UAS (in conformity with Section 8.2 and Section 17.2) but it accepts only the

1407 REGISTER method and generates only the responses detailed in this section. Note that the REGISTER  
1408 method also does not support the Record-Route or Route header, and that proxy servers MUST NOT add  
1409 Record-Route headers to REGISTER requests.

1410 A registrar must know (through provisioning or some other mechanism) the set of administrative do-  
1411 main(s) for which its associated location service(s) are responsible. REGISTER requests MUST be pro-  
1412 cessed by a registrar in the order that they are received.

1413 Upon the arrival of a REGISTER message, the registrar MUST inspect the Request-URI to determine  
1414 whether it has access to a location service responsible for the domain to which this request is addressed.  
1415 If this message is for some other administrative domain, then if the registrar can act as a proxy server, it  
1416 SHOULD forward the request to the addressed domain (following the general behavior for proxying messages  
1417 described in Section 16).

1418 When a registrar receives a REGISTER message, it is RECOMMENDED that the registrar authenticate  
1419 the user agent client. Mechanisms for the authentication of SIP user agents are described in Section 20.2;  
1420 registration behavior in no way overrides the generic authentication framework for SIP. If no authentication  
1421 mechanism is available, the registrar MAY take the From address as the asserted identity of the originator of  
1422 the request.

1423 Once the identity of the registering user has been ascertained, it is RECOMMENDED that the registrar  
1424 determine if the authenticated user agent is authorized to request and/or modify registrations for this address  
1425 of record. For example, a registrar might consult a authorization database (directly or through an appropriate  
1426 protocol) that maps credentials or other tokens of identity resulting from authentication to one or more  
1427 addresses of record for which this identity is responsible.

1428 Note that in architectures that support third-party registration, one entity may be responsible for updating the  
1429 registrations associated with multiple addresses of record.

1430 When the registrar has determined that the client is permitted to make the request, the registrar MUST  
1431 extract the address of record from the To header field of the REGISTER. Note that the registrar MUST  
1432 extract the entire To header field URI in order to use it as an index in the location service.

1433 Next, the registrar MUST query its location service (the repository of previously registered bindings)  
1434 for the set of bindings associated with this address of record. If the address of record is not valid for this  
1435 administrative domain (for example, because the username is not assigned), then the registration attempt  
1436 fails (see below). A full URI comparison (as described in Section 21.1.4) MUST be performed to determine  
1437 whether a given binding matches this address of record.

1438 The registrar now MUST extract all the Contact header fields from the REGISTER message (note that  
1439 there may be no Contact header field).

1440 Each contact address in a REGISTER MUST now be compared to all existing registrations at this loca-  
1441 tion service according to the rules in Section 21.1.4. Note that URIs other than SIP URIs in contact addresses  
1442 MUST be compared according to the standard URI equivalency rules for the URI schema in question.

1443 If a match is found among pre-existing registrations, the registrar MUST copy all parameters associated  
1444 with the current Contact header field from the REGISTER message into the pre-existing binding in its  
1445 location service (overwriting with changed values any existing parameters as necessary, with the exception  
1446 of “expires”). Expiration intervals for this contact address MUST also be reset, based on any suggested  
1447 expiration in the REGISTER (remember that this can be “0”).

1448 If no match is found among the set of pre-existing registrations, the registrar MUST create a new binding  
1449 in its location service between the address of record and the current Contact header field. All Contact  
1450 header field parameters are copied verbatim into this new binding (again with the exception of “expires”).  
1451 An expiration interval MUST be selected by the registrar, taking into account any suggested expiration for

1452 this contact address in the REGISTER.

1453           Allowing the registrar to set the registration interval protects it against excessively frequent registration refreshes  
1454 while limiting the state that it needs to maintain and decreasing the likelihood of registrations going stale.

1455       The expiration interval mandated by the registrar may be either longer or shorter than the interval sug-  
1456 gested by the sender of the REGISTER, though the registrar SHOULD abide by the registering client's  
1457 suggestion.

1458           A server MAY decide to lengthen the expiration interval if the refresh rate of a particular client exceeds a thresh-  
1459 old, for example.

1460       After the expiration interval selected by the registrar for a binding has passed, if the binding has not been  
1461 refreshed (increasing the expiration interval), the registrar SHOULD silently discard the binding.

1462       Once all bindings in the location service have been updated to reflect any changes present to contact  
1463 addresses in the REGISTER message, the registrar MUST remove any bindings that expire immediately.

1464           The REGISTER might have set the expiration interval for some bindings to "0" to remove them before their  
1465 expiration interval passes.

1466       Finally, the registrar must generate a response. If the address of record given in the To header field of  
1467 the REGISTER method is valid for its administrative domain, then a 200 response MUST be sent, which  
1468 MUST contain a complete list (within Contact header fields) of the currently valid bindings in the location  
1469 service associated with the address of record contained in the To field of the REGISTER request. This list  
1470 MAY be empty (in which case the 200 would not contain any Contact headers).

1471       In a successful response to a REGISTER, wherein the bindings for this address of record are enumerated  
1472 as described above, the registrar MUST supply an expiration interval for each contact address in either an  
1473 "expires" parameter of a Contact header or an Expires header. This interval specifies the expiration interval  
1474 that has been mandated by the registrar (taking into account the registering UA's suggestion).

1475       If the registration failed because the address of record contained in the To field of the REGISTER is not  
1476 valid for this domain, then a 404 MUST be sent.

## 1477 11 Querying for Capabilities

1478       The SIP method OPTIONS allows a UA to query another UA or a proxy server as to its capabilities. This  
1479 allows a client to discover information about the methods, content types, extensions, codecs etc. supported  
1480 without actually "ringing" the other party. For example, before a client inserts a Require header field into  
1481 an INVITE listing an option that it is not certain the destination UAS supports, the client can query the  
1482 destination UAS with an OPTIONS to see if this option is returned in a Supported header field.

1483       The target of the OPTIONS request is identified by the Request-URI, which could identify another  
1484 User Agent or a SIP Server. If the OPTIONS is addressed to a proxy server, the Request-URI is set  
1485 without a user part, similar to the way a Request-URI is set for a REGISTER request. Alternatively, a  
1486 server receiving an OPTIONS request with a Max-Forwards header value of 0 MAY respond to the request  
1487 regardless of the Request-URI.

1488           This behavior is common with HTTP/1.1. This behavior can be used as a "traceroute" functionality to check the  
1489 capabilities of individual hop servers by sending a series of OPTIONS requests with incremented Max-Forwards  
1490 values.

1491       An OPTIONS request sent as part of an established dialog does not have any impact on the dialog.

## 1492 11.1 Construction of OPTIONS Request

1493 An OPTIONS request is constructed using the standard rules for a SIP request as discussed Section 8.1.1.

1494 A Contact header field MAY be present in an OPTIONS.

1495 An Accept header field SHOULD be included to indicate the type of message body the UAC wishes to  
1496 receive in the response.

1497 Example OPTIONS request:

```
1498 OPTIONS sip:carol@chicago.com SIP/2.0
1499 Via: SIP/2.0/UDP 10.1.1.1:5060;branch=23411513a6
1500 Via: SIP/2.0/UDP 10.1.3.3:5060
1501 To: <sip:carol@chicago.com>
1502 From: Alice <sip:alice@atlanta.com>;tag=1928301774
1503 Call-ID: a84b4c76e66710@10.1.3.3
1504 CSeq: 63104 OPTIONS
1505 Contact: <sip:alice@10.1.3.3>
1506 Accept: application/sdp
1507 Content-Length: 0
```

## 1508 11.2 Processing of OPTIONS Request

1509 The response to an OPTIONS is constructed using the standard rules for a SIP response as discussed in  
1510 Section 8.2.7. The response code chosen is the same that would have been chosen had the request been an  
1511 INVITE. That is, a 200 (OK) would be returned if the UAS is ready to accept a call, a 486 (Busy Here)  
1512 would be returned if the UAS is busy, etc. This allows an OPTIONS request to be used to determine the  
1513 basic state of a UAS, which can be an indication of whether the UAC will accept an INVITE request.

1514 Note that this use of OPTIONS has limitations due the differences in proxy handling of OPTIONS and  
1515 INVITE requests. While a forked INVITE can result in multiple 200 OK responses being returned, a forked  
1516 OPTIONS will only result in a single 200 OK response, since it is treated by proxies using the non-INVITE  
1517 handling. See Section 13.2.1 for the normative details.

1518 If the response to an OPTIONS is generated by a proxy server, the proxy returns a 200 (OK) listing the  
1519 capabilities of the server. The response does not contain a message body.

1520 Allow, Accept, Accept-Encoding, Accept-Language, and Supported header fields SHOULD be  
1521 present in a 200 OK response to an OPTIONS request.

1522 A Contact header field MAY be present in a 200 OK response.

1523 A Warning header field MAY be present.

1524 A message body MAY be sent, the type of which is determined by the Accept header in the OPTIONS  
1525 request.

1526 Example OPTIONS response generated by a UAS (corresponding to the request in Section 11.1):

```
1527 SIP/2.0 200 OK
1528 Via: SIP/2.0/UDP 10.1.1.1:5060;branch=23411513a6
1529 Via: SIP/2.0/UDP 10.1.3.3:5060
1530 To: <sip:carol@chicago.com>;tag=93810874
```



```
1531 From: Alice <sip:alice@atlanta.com>;tag=1928301774
1532 Call-ID: a84b4c76e66710@10.1.3.3
1533 CSeq: 63104 OPTIONS
1534 Contact: <sip:carol@10.3.6.6>
1535 Allow: INVITE, ACK, CANCEL, OPTIONS, BYE
1536 Accept: application/sdp
1537 Accept-Encoding: gzip
1538 Accept-Language: en
1539 Supported: foo
1540 Content-Type: application/sdp
1541 Content-Length: 274
1542
1543 v=0
1544 o=carol 28908764872 28908764872 IN IP4 10.3.6.6
1545 s=-
1546 t=0 0
1547 c=IN IP4 10.3.6.6
1548 m=audio 0 RTP/AVP 0 1 3 99
1549 a=rtpmap:0 PCMU/8000
1550 a=rtpmap:1 1016/8000
1551 a=rtpmap:3 GSM/8000
1552 a=rtpmap:99 SX7300/8000
1553 m=video 0 RTP/AVP 31 34
1554 a=rtpmap:31 H261/90000
1555 a=rtpmap:34 H263/90000
```

## 1556 12 Dialogs

1557 A key concept for a user agent is that of a dialog. A dialog represents a peer- to-peer SIP relationship between  
1558 a two user agents that persists for some time. The dialog facilitates sequencing of messages between the  
1559 user agents, and proper routing of requests between both them. The dialog represents a context in which to  
1560 interpret SIP messages. The previous section discussed method independent UA processing for requests and  
1561 responses outside of a dialog. This section discusses how those requests and responses are used to construct  
1562 a dialog, and then how subsequent requests and responses are sent within a dialog.

1563 A dialog is identified at each UA with a dialog ID, which consists of a Call-ID value, a local URI and  
1564 local tag (together called the local address), and a remote URI and remote tag (together called the remote  
1565 address). The dialog ID at each UA involved in the dialog is not the same. Specifically, the local URI and  
1566 local tag at one UA are identical to the remote URI and remote tag at the peer UA. The tags are opaque  
1567 tokens that facilitate the generation of unique dialog IDs.

1568 A dialog ID is also associated with all responses, and with any request that contains a tag in the To field.  
1569 The rules for computing the dialog ID of a message depend on whether the entity is a UAC or UAS. For a  
1570 UAC, the Call-ID value of the dialog ID is set to the Call-ID of the message, the remote address is set to the  
1571 To field of the message, and the local address is set to the From field of the message (these rules apply to

1572 both requests and responses). As one would expect, for a UAS, the Call-ID value of the dialog ID is set to  
1573 the Call-ID of the message, the remote address is set to the From field of the message, and the local address  
1574 is set to the To field of the message.

1575 A dialog contains certain pieces of state needed for further message transmissions within the dialog.  
1576 This state consists of the Call-ID, a local sequence number (used to order requests from the UA to its peer),  
1577 a remote sequence number (used to order requests from its peer to the UA), and a route set, which is an  
1578 ordered list of URIs. The route set is the set of servers that need to be traversed to send a request to the peer.  
1579 A dialog can also be in the “early” state, which occurs when it is created with a provisional response, and  
1580 then transition to the “confirmed” state when the final response comes.

## 1581 12.1 Creation of a Dialog

1582 Dialogs are created through the generation of non-failure responses to requests with specific methods.  
1583 Within this specification, only 2xx and 1xx responses with a To tag to INVITE establish a dialog. A di-  
1584 alog established by a non-final response to a request is in the “early” state and it is called an early dialog.  
1585 Extensions MAY define other means for creating dialogs. Section 13 gives more details that are specific to  
1586 the INVITE method. Here, we describe the process for creation of dialog state that is not dependent on the  
1587 method.

1588 A dialog is identified by a dialog ID. A dialog ID consists of three components, namely a call identifier  
1589 component, a local address component and a remote address component. UAs MUST assign values to these  
1590 components as described below.

### 1591 12.1.1 UAS behavior

1592 When a UAS responds to a request with a response that establishes a dialog (such as a 2xx to INVITE), the  
1593 UAS MUST copy all Record-Route headers from the request into the response (including the URIs, URI  
1594 parameters, and any Record-Route header parameters, whether they are known or unknown to the UAS)  
1595 and MUST maintain the order of those headers. The UAS MUST add a Contact header field to the response.  
1596 The Contact header field contains an address where the UAS would like to be contacted for subsequent  
1597 requests in the dialog (which includes the ACK for a 2xx response in the case of an INVITE). Generally, the  
1598 host portion of this URI is the IP address of the host, or its FQDN. The URI provided in the Contact header  
1599 MUST be a SIP URI.

1600 The UAS then constructs the state of the dialog. This state MUST be maintained for the duration of the  
1601 dialog. First, the route set MUST be computed by following these steps:

- 1602 1. The list of URIs in the Record-Route headers in the request, if present, are taken, including any URI  
1603 parameters.
- 1604 2. The URI in the Contact header from the request if present, is taken, including any URI parameters.  
1605 The URI is appended to the bottom of the list of URIs from the previous step.

1606 Contact was not mandatory in RFC2543. Thus, if the UAS is talking to an older UAC, the UAC might not  
1607 have inserted the Contact header.

- 1608 3. The resulting list of URIs is called the *route set*.

1609 These rules clearly imply that a UA MUST be able to parse and process Record-Route header fields. This is a  
1610 change from RFC2543, where all record-route and route processing was optional for user agents.

1611 It is possible for the *route set* to be empty. This will occur if neither **Record-Route** headers nor a  
1612 **Contact** header were present in the request. The UAS **MUST** also remember whether the bottom-most entry  
1613 in the *route set* was constructed from a **Contact** header or not. This is effectively a boolean value, which we  
1614 refer to as **CONTACT\_SET**. This is needed in order for the UA to determine whether the bottom most value  
1615 can be updated from subsequent requests; if it was constructed from a **Contact**, it can be updated.

1616 The remote sequence number **MUST** be set to the value of the sequence number in the **Cseq** header of  
1617 the request. The local sequence number **MUST** be empty. The call identifier component of the dialog ID  
1618 **MUST** be set to the value of the **Call-ID** in the request. The local address component of the dialog ID **MUST**  
1619 be set to the **To** field in the response to the request (which therefore includes the tag), and the remote address  
1620 component of the dialog ID **MUST** be set to the **From** field in the request. A UAS **MUST** be prepared to  
1621 receive a request without a tag in the **From** field, in which case the tag is considered to effectively have a  
1622 value of null.

1623 This is to maintain backwards compatibility with RFC2543, which did not mandate **From** tags.

### 1624 **12.1.2 UAC behavior**

1625 When a UAC receives a response that establishes a dialog, it constructs the state of the dialog. This state  
1626 **MUST** be maintained for the duration of the dialog. First, the route set **MUST** be computed by following  
1627 these steps:

- 1628 1. The list of URIs present in the **Record-Route** headers in the response are taken, if present, including  
1629 all URI parameters, and their order is reversed.
- 1630 2. The URI in the **Contact** header from the response, if present, is taken, including all URI parameters,  
1631 and appended to the end of the list from the previous step.
- 1632 3. The list of URIs resulting from the above two operations is referred to as the *route set*.

1633 It is possible for the *route set* to be empty. This will occur if neither **Record-Route** headers nor a  
1634 **Contact** header were present in the response. The UAC **MUST** also remember whether the bottom-most  
1635 entry in the *route set* was constructed from a **Contact** header or not. This is effectively a boolean value,  
1636 which we refer to as **CONTACT\_SET**. This is needed in order for the UA to determine whether the bottom  
1637 most value can be updated from subsequent requests; if it was constructed from a **Contact**, it can be updated.

1638 The local sequence number **MUST** be set to the value of the sequence number in the  
1639 **Cseq** header of the request. The remote sequence number **MUST** be empty (it is established when the UA  
1640 sends a request within the dialog). The call identifier component of the dialog ID **MUST** be set to the value  
1641 of the **Call-ID** in the request. The local address component of the dialog ID **MUST** be set to the **From**  
1642 field in the request, and the remote address component of the dialog ID **MUST** be set to the **To** field of the  
1643 response. A UAC **MUST** be prepared to receive a response without a tag in the **To** field, in which case the  
1644 tag is considered to effectively have a value of null.

1645 This is to maintain backwards compatibility with RFC2543, which did not mandate **To** tags.

## 1646 **12.2 Requests within a Dialog**

1647 Once a dialog has been established between two UAs either of them **MAY** initiate new transactions as needed  
1648 within the dialog. However, a dialog imposes some restrictions on the use of simultaneous transactions.

1649 A TU **MUST NOT** initiate a new regular transaction within a dialog while a regular transaction is in  
1650 progress (in either direction) within that dialog.

1651 OPEN ISSUE #113: Should we relax the constraint on non-overlapping regular transactions?

1652 A route refresh request sent within a dialog is defined as a request that can modify the *route set* of  
1653 the dialog. For dialogs that have been established with an INVITE, the only route refresh request defined  
1654 is re-INVITE (see Section 14). Other extensions may define different route refresh requests for dialogs  
1655 established in other ways.

1656 Note that an ACK is *NOT* a route refresh request.

## 1657 12.2.1 UAC Behavior

1658 **12.2.1.1 Generating the Request** A request within a dialog is constructed by using many of the com-  
1659 ponents of the state stored as part of the dialog.

1660 The **To** header field of the request **MUST** be set to the remote address, and the **From** header field **MUST**  
1661 be set to the local address (both including tags, assuming the tags are not null).

1662 The **Call-ID** of the request **MUST** be set to the **Call-ID** of the dialog. Requests within a dialog **MUST**  
1663 contain strictly monotonically increasing and contiguous **CSeq** sequence numbers (increasing-by-one) in  
1664 each direction. Therefore, if the local sequence number is not empty, the value of the local sequence number  
1665 **MUST** be incremented by one, and this value **MUST** be placed into the **Cseq** header. If the local sequence  
1666 number is empty, an initial value **MUST** be chosen using the guidelines of Section 8.1.1.4. The method field  
1667 in the **Cseq** header **MUST** match the method of the request.

1668 With a length of 32 bits, a client could generate, within a single call, one request a second for about 136 years  
1669 before needing to wrap around. The initial value of the sequence number is chosen so that subsequent requests within  
1670 the same call will not wrap around. A non-zero initial value allows clients to use a time-based initial sequence  
1671 number. A client could, for example, choose the 31 most significant bits of a 32-bit second clock as an initial  
1672 sequence number.

1673 The **Request-URI** of requests is determined according to the following rules:

1674 The UAC takes the list of URI in the *route set*. The top URI **MUST** be inserted into the request URI of  
1675 the request, including all URI parameters. Any URI parameters not allowed in the request URI **MUST** then  
1676 be stripped. Each of the remaining URIs (if any) from the *route set*, including all URI parameters, **MUST** be  
1677 placed into a **Route** header field into the request, in order.

1678 A TU **SHOULD** follow the rules just mentioned to build the **Request-URI** of the request, regardless of  
1679 whether the UA uses an outbound proxy server or not. However, in some instances, a UA may not be willing  
1680 or capable of sending the request to the top element in the *route set*. One example is a UA that is not capable  
1681 of DNS, and therefore may not be able to follow those procedures. In these cases, the UA **MAY** send the  
1682 request to a local outbound server. In this case, it **MUST NOT** remove the top **Route** header.

1683 In dialogs created by an INVITE, if the UA is the caller, it sets the **Request-URI** to the same value it used for  
1684 the initial request, and sends it to its local outbound server.

1685 Bug#161: Which **Request-URI** does the callee use?

1686 A UAC **SHOULD** include a **Contact** header in any route refresh requests within a dialog, and unless  
1687 there is a need to change it, the URI **SHOULD** be the same as used in previous requests within the dialog. As  
1688 discussed in Section 12.2.2, a **Contact** header in a route refresh request updates the route set. This allows a  
1689 UA to provide a new contact address, should its address change during the duration of the dialog.

1690 However, requests that are not route refresh requests do not affect the *route set* for the dialog.

1691 Once the request has been constructed, the address of the server is computed and the request is sent,  
1692 using the same procedures for requests outside of a dialog (Section 8.1.1).

1693 **12.2.1.2 Processing the Responses** The UAC will receive responses to the request from the transaction  
1694 layer.

1695 The behavior of a UAC that receives a 3xx response for a request sent within a dialog is the same as if  
1696 the request would have been sent outside a dialog. This behavior is described in Section 13.2.2.

1697 Note however that when the UAC tries alternative locations it still uses the *route set* for the dialog to build the  
1698 Route header of the request.

1699 If a UAC has a *route set* for a dialog, and receives a 2xx response to a route refresh it sent, the **Contact**  
1700 header field of the response is examined. If not present, the *route set* remains unchanged. If the response had  
1701 a **Contact** header field, and the boolean variable CONTACT\_SET is false, the URI in the **Contact** header  
1702 field in the response is added to the bottom of the *route set*, and CONTACT\_SET is set to true. If the route  
1703 refresh request response had a **Contact** header field, and CONTACT\_SET is true, the URI in the **Contact**  
1704 header field of the response to the route refresh request replaces the bottom value in the *route set*. If a route  
1705 refresh request is responded with a non-2xx final response the *route set* remains unchanged as if no route  
1706 refresh request had been issued.

1707 If the response for the a request within a dialog is a 481 (Call/Transaction Does Not Exist) or a 408  
1708 (Request Timeout) the UAC SHOULD terminate the dialog. A UAC SHOULD also terminate a dialog if no  
1709 response at all is received for the request (the client transaction would inform the TU about the timeout.)

1710 For INVITE initiated dialogs terminating the dialog consists of sending a BYE.

## 1711 12.2.2 UAS behavior

1712 The UAS will receive the request from the transaction layer. If the request has a tag in the **To** header field,  
1713 the UAS core computes the dialog identifier corresponding to the request and compares it with existing  
1714 dialogs. If there is a match, this is a mid-dialog request. In that case, the same processing rules for requests  
1715 outside of a dialog, discussed in Section 8.2, are applied by the UAS once the request is received from the  
1716 transaction layer.

1717 If the request has a tag in the **To** header field but the dialog identifier does not match any of the existing  
1718 dialogs, the UAS may have crashed and restarted, or may have received a request for a different (possibly  
1719 failed) UAS. The UAS MAY either accept or reject the request. Accepting the request provides robustness, so  
1720 that dialogs can persist even through crashes. UAs wishing to support this capability must choose monoton-  
1721 ically increasing CSeq sequence numbers even across reboots. This is because subsequent requests from  
1722 the crashed-and-rebooted UA towards the other UA need to have a CSeq sequence number higher than  
1723 previous requests in that direction.

1724 Note also that the crashed-and-rebooted UA will have lost any **Route** headers which would need to be  
1725 inserted into a subsequent request. Therefore, it is possible that the requests may not be properly forwarded  
1726 by proxies.

1727 RTP media agents allowing restarts need to be robust by accepting out-of-range timestamps and sequence num-  
1728 bers.

1729 If the UAS wishes to reject the request, because it does not wish to recreate the dialog, it MUST respond  
1730 to the request with a 481 (Call/Transaction Does Not exist) status code and pass that to the server transaction.  
1731

1732 Requests that do not change in any way the state of a dialog may be received within a dialog (e.g., an  
1733 OPTIONS request). They are processed as if they had been received outside the dialog.

1734 Requests within a dialog MAY contain **Record-Route** and **Contact** header fields. However, requests  
1735 that are not route refresh requests do not update the *route set* for the dialog. This specification only defines  
1736 one route refresh request: re-INVITE (see Section 14).

1737 Special rules apply when updated **Record-Route** or **Contact** header fields are received inside a route  
1738 refresh request. If a UAS has a *route set* for a dialog, and receives a route refresh for that dialog containing

1739 Record-Route header fields, it MUST copy those header fields into any 2xx response to that request. If the  
1740 boolean variable CONTACT\_SET is true, the Contact header field in the request (if present) replaces the  
1741 last entry in the *route set*. If the boolean variable CONTACT\_SET is false, the UAS MUST add the URI in the  
1742 Contact header field in the route refresh request to the bottom of the *route set*, and then set CONTACT\_SET  
1743 to true. If the request did not contain a Contact header field, the route-set at the UAS remains unchanged.

1744           Route refresh requests only update the Contact of the *route set* and not the elements formed from Record-  
1745           Route. Updating the latter would introduce severe backwards compatibility problems with RFC 2543 compliant  
1746           systems.

1747           If the remote sequence number is empty, it MUST be set to the value of the sequence number in the Cseq  
1748 header in the request. If the remote sequence number was not empty, but the sequence number of the request  
1749 is lower than the remote sequence number, the request is out of order and MUST be rejected with a 500  
1750 response. If the remote sequence number was not empty, and the sequence number of the request is greater  
1751 than the remote sequence number, the request is in order. It is possible for the CSeq header to be higher  
1752 than the remote sequence number by more than one. This is not an error condition, and a UAS SHOULD be  
1753 prepared to receive and process requests with CSeq values more than one higher than the previous received  
1754 request. The UAS MUST then set the remote sequence number to the value of the sequence number in the  
1755 Cseq header in the request.

### 1756 12.3 Termination of a Dialog

1757 Dialogs can end in several different ways, depending on the method. When a dialog is established with  
1758 INVITE, it is terminated with a BYE. No other means to terminate a dialog are described in this specification,  
1759 but extensions can define other ways.

## 1760 13 Initiating a Session

### 1761 13.1 Overview

1762 When a user agent client desires to initiate a session (for example, audio, video, or a game), it formulates  
1763 an INVITE request. The INVITE request asks a server to establish a session. This request is forwarded by  
1764 proxies, eventually arriving at one or more UAS which can potentially accept the invitation. These UAS's  
1765 will frequently need to query the user about whether to accept the invitation. After some time, those UAS can  
1766 accept the invitation (meaning the session is to be established) by sending a 2xx response. If the invitation  
1767 is not accepted, a 3xx,4xx,5xx or 6xx response is sent, depending on the reason for the rejection. Before  
1768 sending a final response, the UAS can also send a provisional response (1xx), either reliably or unreliably,  
1769 to advise the UAC of progress in contacting the called user.

1770           After possibly receiving one or more provisional responses, the UA will get one or more 2xx responses or  
1771 one non-2xx final response. Because of the protracted amount of time it can take to receive final responses  
1772 to INVITE, the reliability mechanisms for INVITE transactions differ from those of other requests (like  
1773 OPTIONS). Once it receives a final response, the UAC needs send an ACK for every final response it  
1774 receives. The procedure for sending this ACK depends on the type of response. For final responses between  
1775 300 and 699, the ACK processing is done in the transaction layer, and follows one set of rules (See Section  
1776 17). For 2xx responses, the ACK is generated by the UAC core.

1777           A 2xx response to an INVITE establishes a session, and it also creates a dialog between the UA that  
1778 issued the INVITE and the UA that generated the 2xx response. Therefore, when multiple 2xx responses are

1779 received from different remote UAs (because the INVITE forked), each 2xx establishes a different dialog.  
1780 All these dialogs are part of the same call.

1781 This section provides details on the establishment of a session using INVITE.

## 1782 13.2 Caller Processing

### 1783 13.2.1 Creating the Initial INVITE

1784 Since the initial INVITE represents a request outside of a dialog, its construction follows the procedures of  
1785 Section 8.1.1. Additional processing is required for the specific case of INVITE.

1786 An Allow header field (Section 22.5) SHOULD be present in the INVITE. It indicates what methods can  
1787 be invoked within a dialog, on the UA sending the INVITE, for the duration of the dialog. For example, a  
1788 UA capable of receiving INFO requests within a dialog [22] SHOULD include an Allow header listing the  
1789 INFO method.

1790 A Supported header field (Section 22.37) SHOULD be present in the INVITE. It enumerates all the  
1791 extensions understood by the UAC.

1792 An Accept (Section 22.1) header field MAY be present in the INVITE. It indicates which content-types  
1793 are acceptable to the UA, in both the response received by it, and in any subsequent requests sent to it within  
1794 dialogs established by the INVITE. The Accept header is especially useful for indicating support of various  
1795 session description formats.

1796 The UA MAY add an Expires header field (Section 22.19) to limit the validity of the invitation. If the  
1797 time indicated in the Expires header field is reached and no final answer for the INVITE has been received  
1798 the UAC core SHOULD generate a CANCEL request for the original INVITE.

1799 A UAC MAY also find useful to add, among others, Subject (Section 22.36), Organization (Section  
1800 22.24) and User-Agent (Section 22.41) header fields. They all contain information related to the INVITE.

1801 The UAC MAY choose to add a message body to the INVITE. Section 8.1.1.9 deals with how to construct  
1802 the header fields- Content-Type among others- needed to describe the message body.

1803 There are special rules for message bodies that contain a session description - their corresponding  
1804 Content-Disposition is "session". SIP uses an offer/answer model where one UA sends a session de-  
1805 scription, called the offer, which contains a proposed description of the session. The offer indicates the  
1806 desired communications means (audio, video, games), parameters of those means (such as codec types) and  
1807 addresses for receiving media from the answerer. The other UA responds with another session description,  
1808 called the answer, which indicates which communications means are accepted, the parameters which apply  
1809 to those means, and addresses for receiving media from the offerer. The offer/answer model can be mapped  
1810 into the INVITE transaction in two ways. The first, which is the most intuitive, is that the INVITE contains  
1811 the offer, the 2xx response contains the answer, and no session description is provided in the ACK. In this  
1812 model, the UAC is the offerer, and the UAS is the answerer. A second model is that the INVITE contains no  
1813 session description, the 2xx response contains the offer, and the ACK contains the answer. In this model, the  
1814 UAS is the offerer, and the UAC is the answerer. The second model is useful for gateways from H.323v1  
1815 to SIP, where the H.323 media characteristics are not known until the call is established. This is also useful  
1816 for sessions that use third-party call control. As a result of these models, if the INVITE contains a session  
1817 description, the ACK MUST NOT contain one. Conversely, if the caller chooses to omit the session descrip-  
1818 tion in the INVITE, the ACK MUST contain one (if a 2xx response is received). 2xx responses to an INVITE  
1819 MUST always contain a session description. All user agents that support INVITE MUST support both models.

1820 The Session Description Protocol (SDP) [6] MUST be supported by all user agents as a means to describe  
1821 sessions, and its usage for construction offers and answers MUST follow the procedures defined in [23].

1822 Note that the restrictions of the offer-answer model (session description only in the INVITE *OR* in  
1823 the ACK, but not in both) just described only apply to bodies whose Content-Disposition header field  
1824 is "session". Therefore, it is possible that both the INVITE and the ACK contain a body message (e.g.,  
1825 the INVITE carries a photo (Content-Disposition: render) and the ACK a session description (Content-  
1826 Disposition: session) ).

1827 If the Content-Disposition header field is missing, bodies of Content-Type application/sdp imply the  
1828 disposition "session", while other content types imply "render".

1829 Once the INVITE has been created, the UAC follows the procedures defined for sending requests outside  
1830 of a dialog (Section 8). This results in the construction of a client transaction that will ultimately send the  
1831 request and deliver responses to the UAC.

1832 If a UA *A* sends an INVITE request to *B* and receives an INVITE request from *B* before it has received  
1833 the response to its request from *B*, *A* MAY return a 500 (Internal Server Error), which SHOULD include a  
1834 Retry-After header field specifying when the request should be resubmitted.

### 1835 13.2.2 Processing INVITE Responses

1836 Once the INVITE has been passed to the INVITE client transaction, the UAC waits for responses for the IN-  
1837 VITE. Responses are matched to their corresponding INVITE because they have the same Call-ID, the same  
1838 From header field, the same To header field, excluding the tag, and the same CSeq. Rules for comparisons  
1839 of these headers are described in Section 22.

1840 **13.2.2.1 1xx responses** Zero, one or multiple provisional responses may arrive before one or more  
1841 final responses are received. Provisional responses for an INVITE request can create "early dialogs". If a  
1842 provisional response has a tag in the To field, and if the dialog ID of the response does not match an existing  
1843 dialog, one is constructed using the procedures defined in Section 12.1.2.

1844 The early dialog will only be needed if the UAC needs to send a request to its peer within the dialog be-  
1845 fore the initial INVITE transaction completes. Header fields present in a provisional response are applicable  
1846 as long as the dialog is in the early state (e.g., an Allow header field in a provisional response contains the  
1847 methods that can be used in the dialog while this is in the early state).

1848 **13.2.2.2 3xx responses** A 3xx response may contain a Contact header field providing new addresses  
1849 where the callee might be reachable. Depending on the status code of the 3xx response (see Section 23.3)  
1850 the UAC MAY choose to try those new addresses.

1851 **13.2.2.3 4xx, 5xx and 6xx responses** A single non-2xx final response may be received for the IN-  
1852 VITE. 4xx, 5xx and 6xx responses may contain a Contact header field indicating the location where addi-  
1853 tional information about the error can be found.

1854 All early dialogs are considered terminated upon reception of the non-2xx final response.

1855 After having received the non-2xx final response the UAC core considers the INVITE transaction com-  
1856 pleted. The INVITE client transaction handles generation of ACKs for the response (see Section 17).

1857 **13.2.2.4 2xx responses** Multiple 2xx responses may arrive at the UAC for a single INVITE request  
1858 due to a forking proxy. Each response is distinguished by the tag parameter in the To header field, and each  
1859 represents a distinct dialog, with a distinct dialog identifier.



1860 If the dialog identifier in the 2xx response matches the dialog identifier of an existing dialog, the dialog  
1861 MUST be transitioned to the “confirmed” state, and the route set for the dialog MUST be recomputed based  
1862 on the 2xx response using the procedures of Section 12.1.2. Otherwise, a new dialog in the “confirmed”  
1863 state is constructed in the same fashion.

1864 The route set only is recomputed for backwards compatibility. RFC 2543 did not mandate mirroring of Record-  
1865 Route headers in a 1xx, only 2xx. However, we cannot update the entire state of the dialog, since mid-dialog  
1866 requests may have been sent within the early call leg, modifying the sequence numbers, for example.

1867 The UAC core MUST generate an ACK request for each 2xx received from the transaction layer. The  
1868 header fields of the ACK are constructed in the same way as for any request sent within a dialog (see Section  
1869 12) with the exception of the CSeq. The sequence number of the CSeq header field MUST be the same as  
1870 the INVITE being acknowledged, but the CSeq method MUST be ACK. If the INVITE did not contain an  
1871 offer, the 2xx will contain one, and therefore the ACK MUST carry an answer in its body.

1872 Once the ACK has been constructed, the procedures of [8] are used to determine the destination address,  
1873 port and transport. However, the request is passed to the transport layer directly for transmission, rather than  
1874 a client transaction. This is because the UAC core handles retransmissions of the ACK, not the transaction  
1875 layer. The ACK MUST be passed to the client transport every time a retransmission of the 2xx final response  
1876 that triggered the ACK arrives.

1877 The UAC core considers the INVITE transaction completed 64\*T1 seconds after the reception of the  
1878 first 2xx response. At this point all the early dialogs that have not transitioned to established dialogs are  
1879 terminated. Once the INVITE transaction is considered completed by the UAC core, no more new 2xx  
1880 responses are expected to arrive.

1881 If, after acknowledging any 2xx response to an INVITE, the caller does not want to continue with that  
1882 dialog, then the caller MUST terminate the dialog by sending a BYE request as described in Section 15.

### 1883 13.3 Callee Processing

#### 1884 13.3.1 Processing of the INVITE

1885 The UAS core will receive INVITE requests from the transaction layer. It first performs the request process-  
1886 ing procedures of Section 8.2, which are applied for both requests inside and outside of a dialog.

1887 Assuming these processing states complete without generating a response, the UAS core performs the  
1888 additional processing steps:

- 1889 1. If the request is an INVITE that contains an Expires header field the UAS core inspects this header  
1890 field. If the INVITE has already expired a 487 response SHOULD be generated. In any case, if the  
1891 INVITE expires before the UAS has generated a final response a 487 response SHOULD be generated.
- 1892 2. If the request has no tag in the To the UAS core checks ongoing transactions. If the To, From, Call-ID,  
1893 CSeq exactly match (including tags) those of any request received previously, but the branch-ID in  
1894 the topmost Via is different from those received previously, the UAS core SHOULD generate a 482  
1895 (Loop detected) response and pass it to the server transaction.

1896 The same request that was generated by the UAC has arrived to the UAS more than once following different  
1897 paths. The UAS processes the request that was received first and responds with 482 (Loop detected) to the rest  
1898 of them.

1899 If no match is found, the request does not belong to any existing dialog. If the request is an INVITE  
1900 the UAS core follows the procedures described in this section.

- 1901 3. If the request is a mid-dialog request, the method-independent processing described in Section 12.2.2  
1902 is first applied. It might also modify the session; Section 14 provides details.
- 1903 4. If the request has a tag in the **To** header field but the dialog identifier does not match any of the existing  
1904 dialogs, the UAS may have crashed and restarted, or may have received a request for a different  
1905 (possibly failed) UAS. Section 12.2.2 provides guidelines to achieve a robust behaviour under such a  
1906 situation.

1907 Processing from here forward assumes that the **INVITE** is outside of a dialog, and is thus for the purposes  
1908 of establishing a new session.

1909 The **INVITE** may contain a session description, in which case the UAS is being presented with an offer  
1910 for that session. It is possible that the user is already a participant in that session, even though the **INVITE**  
1911 is outside of a dialog. This can happen when a user is invited to the same multicast conference by multiple  
1912 other participants. If desired, the UAS **MAY** use identifiers within the session description to detect this  
1913 duplication. For example, **SDP** contains a session id and version number in the origin (**o**) field. If the user  
1914 is already a member of the session and the session parameters contained in the session description have not  
1915 changed, the UAS **MAY** silently accept the **INVITE** (i.e., send a 2xx response without prompting the user).

1916 The **INVITE** may not contain a session description at all, in which case the UAS is being asked to  
1917 participate in a session, but the UAC has asked that the UAS provide the offer of the session.

1918 The callee can indicate progress, accept, redirect, or reject the invitation. In all of these cases, it formu-  
1919 lates a response using the procedures described in Section 8.2.7.

1920 **13.3.1.1 Progress** The UAS may not be able to answer the invitation immediately, and might choose  
1921 to indicate some kind of progress to the caller (for example, an indication that a phone is ringing). This is  
1922 accomplished with a provisional response between 101 and 199. These provisional responses establish early  
1923 dialogs and therefore follow the procedures of Section 12.1.1 in addition to those of Section 8.2.7. A UAS  
1924 **MAY** send as many provisional responses as it likes. Each of these **MUST** indicate the same dialog ID. **SIP**,  
1925 however, does not guarantee that these provisional responses are reliably delivered to the UAC.

1926 **13.3.1.2 The INVITE is redirected** If the UAS decides to redirect the call, a 3xx response is sent. A  
1927 300 (Multiple Choices), 301 (Moved Permanently) or 302 (Moved Temporarily) response **SHOULD** contain  
1928 a **Contact** header field containing URIs of new addresses to be tried. The response is passed to the **INVITE**  
1929 server transaction, which will deal with its retransmissions.

1930 **13.3.1.3 The INVITE is rejected** A common scenario occurs when the callee is currently not willing  
1931 or able to take additional calls at this end system. A 486 (Busy Here) **SHOULD** be returned in such scenario.  
1932 If the UAS knows that no other end system will be able to accept this call a 600 (Busy Everywhere) response  
1933 **SHOULD** be sent instead. However, it is unlikely that a UAS will be able to know this in general, and thus  
1934 this response will not usually be used. The response is passed to the **INVITE** server transaction, which will  
1935 deal with its retransmissions.

1936 **13.3.1.4 The INVITE is accepted** The UAS core generates a 2xx response. This response establishes  
1937 a dialog, and therefore follows the procedures of Section 12.1.1 in addition to those of Section 8.2.7.

1938 A 2xx response to an INVITE SHOULD contain the Allow header field and the Supported header field,  
1939 and MAY contain the Accept header field. Including these header fields allows the UAC to determine the  
1940 features and extensions supported by the UAS for the duration of the call, without probing.

1941 If the INVITE request contained an offer, the 2xx MUST contain an answer. If the INVITE did not contain  
1942 an offer, the 2xx MUST contain an offer.

1943 Once the response has been constructed it is passed to the INVITE server transaction. Note, however, that  
1944 the INVITE server transaction will be destroyed as soon as it receives this final response. Therefore, it is  
1945 necessary to pass periodically the response to the transport until the ACK arrives. The 2xx response is passed  
1946 to the transport with an interval that starts at T1 seconds and doubles for each retransmission until it reaches  
1947 T2 seconds (T1 and T2 are defined in Section 17). Response retransmissions cease when an ACK request is  
1948 received with the same dialog ID as the response. This is independent of whatever transport protocols are  
1949 used to send the response.

1950 Since 2xx is retransmitted end-to-end, there may be hops between UAS and UAC which are UDP. To ensure  
1951 reliable delivery across these hops, the response is retransmitted periodically even if the transport at the UAS is  
1952 reliable.

1953 If the server retransmits the 2xx response for  $64 * T1$  seconds without receiving an ACK, it considers the  
1954 dialog completed, the session terminated, and therefore it SHOULD send a BYE.

## 1955 14 Modifying an Existing Session

1956 A successful INVITE request (see Section 13) establishes both a dialog between two user agents and a  
1957 session (using the offer/answer model). Section 12 explains how to modify an existing dialog using a route  
1958 refresh request (e.g., changing the *route set* of the dialog). This section describes how to modify the actual  
1959 session. This modification can involve changing addresses or ports, adding a media stream, deleting a media  
1960 stream, and so on. This is accomplished by sending a new INVITE request within the same dialog that  
1961 established the session. An INVITE request sent within an existing dialog is known as a re-INVITE.

1962 Note that a single re-INVITE can modify at the same time the dialog and the parameters of the session.

1963 Either the caller or callee can modify an existing session.

### 1964 14.1 UAC Behavior

1965 The same offer-answer model that applies to session descriptions in INVITEs (Section 13.2.1) applies to  
1966 re-INVITEs. As a result, a UAC that wants to add a media stream, for example, will create a new offer that  
1967 contains this media stream, and send that in an INVITE request to its peer. It is important to note that the  
1968 full description of the session, not just the change, is sent. This maintains the idempotency of SIP, supports  
1969 stateless session processing in various elements, and supports failover and recovery capabilities. Of course,  
1970 a UAC MAY send a re-INVITE with no session description, in which case the response to the re-INVITE will  
1971 contain the offer.

1972 If the session description format has the capability for version numbers, the offerer SHOULD indicate  
1973 that the version of the session description has changed.

1974 The To, From, Call-ID, CSeq, and Request-URI of a re-INVITE are set following the same rules as  
1975 for regular requests within an existing dialog, described in Section 12.

1976 Note that, as opposed to initial INVITEs (see Section 13), re-INVITEs contain tags in the To header  
1977 field and are sent using the *route set* for the dialog. Therefore, a single final (2xx or non-2xx) response is  
1978 received for re-INVITEs.

1979 Note that a UAC MUST NOT initiate a new INVITE transaction within a dialog while another transaction  
1980 (INVITE or non-INVITE) is in progress. However, a UA MAY initiate a regular transaction within an early  
1981 dialog - while an INVITE transaction is in progress.

1982 If a re-INVITE is responded with a non-2xx final response the session parameters MUST remain un-  
1983 changed, as if no re-INVITE had been issued. Note that, as stated in Section 12.2.1.2, if the non-2xx final  
1984 response is a 481 (Call/Transaction Does Not Exist) or a 408 (Request Timeout) or no response at all is  
1985 received for the re-INVITE the UAC will terminate the dialog.

1986 The rules for transmitting a re-INVITE and for generating an ACK for a 2xx response to re-INVITE are  
1987 the same as for an INVITE (Section 13.2.1).

## 1988 14.2 UAS Behavior

1989 Section 13.3.1 describes the steps to follow in order to distinguish incoming re-INVITEs from incoming  
1990 initial INVITEs. This Section describes the procedures to follow upon reception of a re-INVITE for an  
1991 existing dialog.

1992 A UAS that receives a second INVITE before it sent the final response to a first INVITE with a lower  
1993 CSeq sequence number on the same dialog MUST return a 500 response to the second INVITE and MUST  
1994 include a Retry-After header field with a randomly chosen value of between 0 and 10 seconds. Similarly,  
1995 a UAS the receives an INVITE on a dialog while an INVITE it had sent on that dialog is in progress MUST  
1996 return a 500 response to the received INVITE and MUST include a Retry-After header field with a randomly  
1997 chosen value of between 0 and 10 seconds.

1998 If a user agent receives a re-INVITE for an existing dialog it MUST check any version identifiers in the  
1999 session description or, if there are no version identifiers, the content of the session description to see if it has  
2000 changed. If the session description has changed, the user agent server MUST adjust the session parameters  
2001 accordingly, possibly after asking the user for confirmation.

2002 Versioning of the session description can be used to accommodate the capabilities of new arrivals to a conference,  
2003 add or delete media or change from a unicast to a multicast conference.

2004 If a UAS generates a 2xx response and never receives an ACK, it SHOULD generate a BYE to terminate  
2005 the dialog.

2006 A UAS providing an offer in a 2xx (because the INVITE did not contain an offer) MUST offer the same  
2007 session description as last provided to the peer, with the exception of being able to change the IP address/port  
2008 if so desired.

2009 Under error conditions (e.g., the UAS has crashed and restarted) the session description in the 2xx response for  
2010 an empty re-INVITE may be different than the one in use at that moment. If the new session description is not  
2011 acceptable for the UAC it SHOULD then send a BYE (after ACKing the 2xx response).

## 2012 15 Terminating a Session

2013 This section describes the procedures to be followed in order to terminate a SIP dialog. For two-party  
2014 sessions that are otherwise unbound in time the termination of the dialog implies the termination of the  
2015 session. Other types of sessions such as multicast sessions are not terminated when a participant terminates  
2016 the SIP dialog that he used to join the session. However, the SIP dialog SHOULD be terminated even  
2017 though its termination does not imply the termination of the session. A UA joining a multicast session MAY  
2018 terminate the SIP dialog immediately after the INVITE transaction used to join the session has completed.

2019 Either the caller or callee may terminate a dialog for any reason. A caller terminates a dialog either with  
2020 BYE or CANCEL depending on the state of the dialog. A callee uses BYE to terminate a confirmed dialog.

2021 Note that if the callee wants to terminate an early dialog it just returns a non-2xx final response for the INVITE.

2022 Sections 13 and 12 document some cases where dialog termination is normative behavior. As a general  
2023 rule, if a UA decides that the dialog is to be terminated, it MUST follow the procedures here to initiate  
2024 signaling action to convey that.

2025 When a UAC sends an INVITE request to create a session, if a 1xx response with a tag in the To field  
2026 is received, an early dialog is created. When a 2xx response is received, the dialog becomes confirmed.  
2027 For either state of the dialog, if the UAC desires to terminate the session, the UAC SHOULD follow the  
2028 procedures described in Section 15.1.1 to terminate the session. If the callee for a new session wishes to  
2029 terminate the dialog, it uses the procedures of Section 15.1.1, but MUST NOT do so until it has receive an  
2030 ACK or until the server transaction times out.

2031 This does not mean a user can't hang up right away; it just means that the software in their phone needs to  
2032 maintain state for a short while in order to properly clean up.

2033 OPEN ISSUE #202: Is this the right solution.

2034 If the UAC desires to end the session before any type of dialog has been created, it SHOULD send a  
2035 CANCEL for the INVITE request that requested establishment of the session that is to be terminated. The  
2036 UAC constructs and sends the CANCEL following the procedures described in Section 9. This CANCEL  
2037 will normally result in a 487 response to be returned to the INVITE, indicating successful cancellation.  
2038 However, it is possible that the CANCEL and a 2xx response to the INVITE "pass on the wire". In this case,  
2039 the UAC will receive a 2xx to the INVITE. It SHOULD then terminate the call by following the procedures  
2040 described in Section 15.1.1.

## 2041 15.1 Terminating a Dialog with a BYE

### 2042 15.1.1 UAC Behavior

2043 A user agent client uses BYE request, sent within a dialog, to indicate to the server that it wishes to terminate  
2044 the session. This will also terminate the dialog. A BYE request MAY be issued by either caller or callee. A  
2045 BYE request SHOULD NOT be sent before the creation of a dialog (either early or confirmed). In that case  
2046 the UAC SHOULD follow the procedures described in Section 9 instead.

2047 Proxies ensure that a CANCEL request is routed in the same way as the INVITE was. However, a proxy  
2048 performing load balancing may route a BYE without a Route header field in a different way than the INVITE, since  
2049 both requests have different CSeq sequence numbers.

2050 The To, From, Call-ID, CSeq, and Request-URI of a BYE are set following the same rules as for  
2051 regular requests sent within a dialog, described in Section 12.

2052 Once the BYE is constructed, it creates a new non-INVITE client transaction, and passes it the BYE  
2053 request. The user agent SHOULD stop sending media as soon as the BYE request is passed to the client  
2054 transaction.

### 2055 15.1.2 UAS Behavior

2056 A UAS core receiving a BYE request checks to see if it matches an existing dialog. If the BYE does  
2057 not match an existing dialog, the UAS core SHOULD generate a 481 response and pass that to the server  
2058 transaction.

2059 A UAS core receiving a BYE request for an existing dialog MUST follow the procedures of Section  
2060 12.2.2 to process the request. Once done, the UAS MUST cease transmitting media streams for the session

2061 being terminated. The UAS core MUST generate a 2xx response to the BYE, and MUST pass that to the  
2062 server transaction for transmission.

2063 The UAS MUST still respond to any pending requests received for that dialog, (which can only be an  
2064 INVITE). It is RECOMMENDED that a 487 (Request Terminated) response is generated to those pending  
2065 requests.

## 2066 **16 Proxy Behavior**

### 2067 **16.1 Overview**

2068 SIP proxies are elements that route SIP requests to user agent servers and SIP responses to user agent clients.  
2069 A request may traverse several proxies on its way to a UAS. Each will make routing decisions, modifying  
2070 the request before forwarding it to the next element. Responses will route through the same set of proxies  
2071 traversed by the request in the reverse order.

2072 It is important to note that being a proxy is a logical role for a SIP element. When a request arrives, an  
2073 element that can play the role of a proxy must first decide if it needs to respond to the request on its own.  
2074 For instance, the request could be malformed or the element may need credentials from the client before  
2075 acting as a proxy. The element MAY respond with any appropriate error code. When responding directly to  
2076 a request, the element is playing the role of a UAS and MUST behave as described in Section 8.2.

2077 A proxy can operate in either a stateful or stateless mode for each new request.

2078 When stateless, a proxy acts as a simple forwarding element. It forwards each request downstream to  
2079 a single element determined by making a routing decision based on the request. It simply forwards every  
2080 response it receives upstream. A stateless proxy discards information about a message once it has been  
2081 forwarded.

2082 On the other hand, a stateful proxy remembers information (specifically, transaction state) about each  
2083 incoming request and any requests it sends as a result of processing the incoming request. It uses this  
2084 information to affect the processing of future messages associated with that request. A stateful proxy MAY  
2085 chose to “fork” a request, routing it to multiple destinations. Any request that is forwarded to more than  
2086 one location MUST be handled statefully. Any request processed using TCP (or any other mechanism that is  
2087 inherently stateful), MUST be handled statefully.

2088 Much of the processing involved when acting statelessly or statefully for a request is identical. The next  
2089 several subsections are written from the point of view of a stateful proxy. The last section calls out those  
2090 places where a stateless proxy behaves differently.

### 2091 **16.2 Stateful Proxy**

2092 When stateful, a proxy is purely a SIP transaction processing engine. Its behavior is modeled here in terms  
2093 of the Server and Client Transactions defined in Section 17. A stateful proxy has a server transaction  
2094 associated with one or more client transactions by a higher layer proxy processing component (see figure 3),  
2095 known as a proxy core. An incoming request is processed by a server transaction. Requests from the server  
2096 transaction are passed to a proxy core. The proxy core determines where to route the request, choosing  
2097 one or more next-hop locations. An outgoing request for each next-hop location is processed by its own  
2098 associated client transaction. The proxy core collects the responses from the client transactions and uses  
2099 them to send responses to the server transaction.

2100 A stateful proxy creates a new server transaction for each new request received. Any retransmissions of

2101 the request will then be handled by that server transaction per Section 17.

2102 Note that this is a model of proxy behavior, not of software. An implementation is free to take any  
2103 approach that replicates the external behavior this model defines.




Figure 3: Stateful Proxy Model

2104 For all new requests, including any with unknown methods, an element intending to proxy the request  
2105 MUST:

- 2106 1. Validate the request (Section 16.3)
- 2107 2. Make a routing decision (Section 16.4)
- 2108 3. Forward the request to each chosen destination (Section 16.5)
- 2109 4. Process all responses (Section 16.6)

### 2110 **16.3 Request Validation**

2111 Before an element can proxy a request, it MUST verify the message's validity. A valid message must pass  
2112 the following checks:

- 2113 1. Reasonable Syntax

2114 2. Max-Forwards

2115 3. Loop Detection

2116 4. Proxy-Require

2117 5. Proxy-Authorization

2118 If any of these checks fail, the element **MUST** behave as a user agent server (see Section 8.2) and respond  
2119 with an error code.

2120 1. Reasonable Syntax check

2121 The request **MUST** be well-formed enough to be handled with a server transaction. Any components  
2122 involved in the remainder of these Request Validation steps or the Request Processing section **MUST**  
2123 be well-formed. Any other components, well-formed or not, **SHOULD** be ignored. For instance, an  
2124 element **SHOULD NOT** reject a request because of a malformed **Date** header field.

2125 This protocol is designed to be extended. Future extensions may define new methods and header fields  
2126 at any time. An element **MUST NOT** refuse to proxy a request because it contains a method or header  
2127 field it does not know about.

2128 2. Max-Forwards check

2129 The **Max-Forwards** header (Section 22.22) is used to limit the number of elements a SIP request can  
2130 traverse.

2131 If the request does not contain a **Max-Forwards** header field, this check is passed.

2132 If the request contains a **Max-Forwards** header field with a field value greater than zero, the check is  
2133 passed.

2134 If the request contains a **Max-Forwards** header field with a field value of zero (0), the element **MUST**  
2135 **NOT** forward the request. If the request was for **OPTIONS**, the element **MAY** act as the final recipient  
2136 and respond per Section 11. Otherwise, the element **MUST** return a 483 (Too many hops) response.

2137 3. Loop Detection check

2138 An element **MUST** check for forwarding loops before forwarding a request. If the request contains a  
2139 **Via** header field value with A sent-by value that equals a value placed into previous requests by the  
2140 proxy, the request has been forwarded by this element before. The request has either looped or is  
2141 legitimately spiraling through the element. To determine if the request has looped, the element **MUST**  
2142 perform the **branch** parameter calculation described in Section 3 on this message and compare it to  
2143 the parameter received in that **Via** field value. If the parameters match, the request has looped. If  
2144 they differ, the request is spiraling, and processing continues. If a loop is detected, the element **MUST**  
2145 return a 482 (Loop Detected) response.

2146 An element **MUST NOT** forward a request to a multicast group which already appears in any of the  
2147 **Via** headers.

2148 4. Proxy-Require check

2149 Future extensions to this protocol may introduce features that require special handling by proxies.  
2150 Endpoints will include a **Proxy-Require** header in requests that use these features, telling the proxy  
2151 it should not process the request unless the feature is understood.



2152 If the request contains a Proxy-Require header (Section 22.28) with one or more option-tags this  
2153 element does not understand, the element MUST return a 420 (Bad Extension) response. The response  
2154 MUST include an Unsupported (Section 22.40) header field listing those option-tags the element did  
2155 not understand.

#### 2156 5. Proxy-Authorization check

2157 If an element requires credentials before forwarding a request, the request MUST be inspected as  
2158 described in Section 20.2.3. That section also defines what the element must do if the inspection fails.

### 2159 16.4 Making a Routing Decision

2160 At this point, the proxy must decide where to forward the request. This can be modeled as computing a set  
2161 of destinations for the request. This set will either be predetermined by the contents of the request or will  
2162 be obtained from an abstract location service. Each destination is represented as a URI and an optional IP  
2163 address, port and transport. This combination is referred to as a "next-hop location".

2164 First, the proxy core checks the received request for Route headers. If any Route header fields are  
2165 present in the request, the element MUST use the URI (including all of its parameters) from the topmost  
2166 Route header field as only next hop URI in the destination set, with no IP address, port and transport set for  
2167 that next hop. The destination set is complete, containing **only** this URI, and the proxy MUST proceed to the  
2168 Request Processing of Section 16.5.

2169 The Route mechanism is used to control the path a request takes through SIP elements, much like strict  
2170 IP source routing. The UAC will insert Route header fields (see Section 12), usually based on information  
2171 provided by proxies through Record-Route header fields (see Section 6).

2172 Assuming there were no Route headers in the received request, the proxy checks the Request-URI of  
2173 the received request. If it has a maddr parameter, and that parameter does not indicate an interface the  
2174 proxy is listening on, the Request-URI MUST be placed into the destination set as the only next hop URI,  
2175 with no IP address, port and transport set for that next hop, and the proxy MUST proceed to Section 16.5.  
2176 If the maddr parameter was present, but did indicate an interface the proxy is listening on, the proxy MUST  
2177 strip the maddr and continue processing as if no maddr were present.

2178 OPEN ISSUE #213: Do we strip just the maddr, or the port and transport as well?

2179 OPEN ISSUE #218: Are we really sure this ordering of precedence of Route, maddr, and domain is correct??  
2180 It is not yet clear. This needs resolution asap finally, since it affects things like loose source routing, outbound proxy  
2181 processing at a UA, and so on.

2182 If the domain of the Request-URI indicates a domain this element is not responsible for, it SHOULD set  
2183 the next hop URI to the Request-URI, and leave the IP address, port and transport of the next hop empty.  
2184 That next hops MUST be placed into the destination set as the only next hop, and the element MUST proceed  
2185 to the task of Request Processing (Section 16.5).

2186 There are many circumstances in which a proxy might receive a request for a domain it is not responsible for.  
2187 A firewall proxy handling outgoing calls (the way HTTP proxies handle outgoing requests) is an example of where  
2188 this is likely to occur.

2189 If the destination set for the request has not been predetermined as described above, this implies that the  
2190 element is responsible for the domain in the Request-URI, and the element MAY use whatever mechanism  
2191 it desires to determine where to send the request. Any of these mechanisms can be modeled as accessing  
2192 an abstract Location Service. This may consist of obtaining information from a location service created

2193 by a SIP Registrar, reading a database, consulting a presence server, utilizing other protocols, or simply  
2194 performing an algorithmic substitution on the Request-URI. The output of these mechanisms is used to  
2195 construct the destination set.

2196 Any information in or about the request or the current environment of the element MAY be used in the  
2197 construction of the destination set. For instance, different sets may be constructed depending contents or  
2198 presence of header fields and bodies, the time of day of the request's arrival, the interface on which the  
2199 request arrived, failure of previous requests, or even the element's current level of utilization.

2200 As potential destinations are located through these services, their next hops are added to the destination  
2201 set. Next-hop locations may only be placed in the destination set once. If a next-hop location is already  
2202 present in the set (based on the definition of equality for the URI type and equality of the optional parame-  
2203 ters), it MUST NOT be added again.

2204 A proxy MAY continue to add destinations to the set after beginning Request Processing. It MAY use any  
2205 information obtained during that processing to determine new locations. For instance, a proxy may choose  
2206 to incorporate contacts obtained in a redirect response (3xx class) into the destination set. If a proxy uses a  
2207 dynamic source of information while building the destination set (for instance, if it consults a SIP Registrar),  
2208 it SHOULD monitor that source for the duration of processing the request. New locations SHOULD be added  
2209 to the destination set as they become available. As above, any given URI MUST NOT be added to the set  
2210 more than once.

2211           Allowing a URI to be added to the set only once reduces unnecessary network traffic, and in the case of incor-  
2212           porating contacts from redirect requests prevents infinite recursion.

2213 An example trivial location service is achieved by configuring an element with a default outbound des-  
2214 tination. All requests are forwarded to this location. The Request-URI of the request is placed in the  
2215 destination set with the optional next-hop IP address, port and transport parameters set to the default out-  
2216 bound destination. The destination set is complete, containing **only** this URI, and the element proceeds to  
2217 the task of Request Processing.

2218 If the Request-URI indicates a resource at this proxy that does not exist, the proxy MUST return a 404  
2219 (Not Found) response.

2220 If the destination set remains empty after applying all of the above, the proxy MUST return an error  
2221 response, which SHOULD be the 480 (Temporarily Unavailable) response.

## 2222 **16.5 Request Processing**

2223 As soon as the destination set is non-empty, a proxy MAY begin forwarding the request. A stateful proxy  
2224 MAY process the set in any order. It MAY process multiple destinations serially, allowing each client transac-  
2225 tion to complete before starting the next. It MAY start client transactions with every destination in parallel. It  
2226 also MAY arbitrarily divide the set into groups, processing the groups serially and processing the destinations  
2227 in each group in parallel.

2228 A common ordering mechanism is to use the qvalue parameter of destinations obtained from Contact  
2229 header fields (see Section 22.10). Destinations are processed from highest qvalue to lowest. Destinations  
2230 with equal qvalues may be processed in parallel.

2231 A stateful proxy must have a mechanism to maintain the destination set as responses are received and  
2232 associate the responses to each forwarded request with the original request. For the purposes of this model,  
2233 this mechanism is a "response context" created by the proxy layer before forwarding the first request.

2234 For each destination, the proxy forwards the request following these steps:

- 2235 1. Make a copy of the received request
- 2236 2. Update the Request-URI
- 2237 3. Add a Via header field value
- 2238 4. Update the Max-Forwards field if present
- 2239 5. Update the Route header field if present
- 2240 6. Optionally add a Record-route header field value
- 2241 7. Optionally add additional headers
- 2242 8. send the new request

2243 Each of these steps is detailed below:

#### 2244 1. Copy request

2245 The proxy starts with a copy of the received request. The copy **MUST** initially contain all of the header  
2246 fields from the received request. Only those fields detailed in the processing described below may be  
2247 removed. The copy **SHOULD** maintain the ordering of the header fields as in the received request. The  
2248 proxy **MUST NOT** reorder field values with a common field name (See Section 7.3.1).

2249 An actual implementation need not perform a copy; the primary requirement is that the processing of each  
2250 next hop begin with the same request.

#### 2251 2. Request-URI

2252 The **Request-URI** in the copy's start line **MUST** be replaced with the URI for this destination. If the  
2253 URI contains any parameters not allowed in a Request-URI, they **MUST** be removed.

2254 This is the essence of a proxy's role. This is the mechanism through which a proxy routes a request  
2255 toward its destination.

#### 2256 3. Via

2257 The proxy **MUST** insert a **Via** header field into the copy before the existing **Via** header fields. The **Via**  
2258 header maddr, ttl, and sent-by components will be set when the request is processed by the transport  
2259 layer (Section 19). The **Via** headers ensure that responses will follow the same set of elements that  
2260 the request traversed.

2261 The proxy **MUST** include a "branch" parameter (Section 22.42) in the **Via** header. When the path of  
2262 a request through one or more forking proxies is graphed, the result is a tree. The branch parameter  
2263 identifies the "branch" each request was forwarded on. The branch parameter value **MUST** be unique  
2264 for each client transaction to which the request is forwarded. The precise format of the branch. token  
2265 is implementation-defined. In order to be able to both detect loops and associate responses with the  
2266 corresponding request, the parameter **SHOULD** consist of two parts separable by the implementation.  
2267 The first part is used to detect loops and distinguish loops from spirals. The second is used to match  
2268 responses to requests.

2269 Loop detection is performed by verifying that those fields having an impact on the routing decision  
2270 have not changed. The value placed in the this part of the **branch** parameter SHOULD reflect all of  
2271 those fields (which include any **Proxy-Require** and **Proxy-Authorization** headers). This is to ensure  
2272 that if the request is routed back to the proxy, and one of those fields changes, it is treated as a spiral  
2273 and not a loop (Section 3). A common way to create this value is to compute a cryptographic hash  
2274 of the **To**, **From**, **Call-ID** header fields, the **Request-URI** of the request received (before translation)  
2275 and the sequence number from the **CSeq** header field, in addition to any **Proxy-Require** and **Proxy-**  
2276 **Authorization** fields that may be present. The algorithm used to compute the hash is implementation-  
2277 dependent, but MD5 [24], expressed in hexadecimal, is a reasonable choice. (Note that base64 is not  
2278 permissible for a token.)

2279 In order to correctly match responses to requests (Section 17.1.3), the value SHOULD also contain a  
2280 part that is a globally unique function of of the branch on which this request will be forwarded. One  
2281 example is a hash of a sequence number, local IP address and **request-URI** of the request.

2282 For example: 7a83e5750418bce23d5106b4c06cc632.1

2283 The “branch” parameter MUST depend on all information used for routing decisions, including the incom-  
2284 ing **request-URI** and any header values affecting the routing choices. This is necessary to distinguish looped  
2285 requests from requests whose routing parameters have changed before returning to this server.

2286 Note that the request method MUST NOT be included in the calculation of the **branch** parameter.  
2287 In particular, **CANCEL** and **ACK** requests MUST have the same **branch** value as the corresponding  
2288 request they cancel or acknowledge. The **branch** parameter is used in correlating those requests at  
2289 server handling them (see Section 17.2.3 and 9.2).

#### 2290 4. Max-Forwards

2291 If the copy contains a **Max-Forwards** header field, the proxy must decrement its value by one (1).

#### 2292 5. Route

2293 If the copy contains a **Route** header field, the proxy must remove the first (topmost) value. Note that  
2294 this value was placed in the destination set and then into the **Request-URI** of this copy in previous  
2295 steps.

#### 2296 6. Record-Route

2297 If this proxy wishes to request to remain on the path of future requests in this dialog, it MUST insert a  
2298 **Record-Route** header value (Section refsec:record-route) into the copy before any existing **Record-**  
2299 **Route** header values. See Section 12 for details on whether this request will be honored. Each proxy  
2300 in the path of a request makes this request independently - the presence of a **Record-Route** header  
2301 does not obligate this proxy to add a value.

2302 If the request is honored, the information the proxy places in the **Record-Route** header value will be  
2303 used at the endpoints to construct **Route** headers. As shown in the processing steps above, **Route**  
2304 headers determine forwarding destinations much like strict IP source routing.

2305 The URI placed in the **Record-Route** header value MUST be a SIP URI. This URI MAY be different  
2306 for each destination the request is forwarded to. The URI SHOULD NOT contain the transport param-  
2307 eter unless the proxy has knowledge (such as in a private network) that the next downstream element  
2308 that will be in the path of subsequent requests supports that transport.

2309           The URI this proxy provides will be used by some other element to make a routing decision. This proxy, in  
2310           general, has no way to know what the capabilities of that element are, so it must restrict itself to the mandatory  
2311           elements of a SIP implementation: SIP URIs and UDP transports.

2312           The URI placed in the **Record-Route** header value **MUST** resolve to this element when the server  
2313           location procedures of [8] are applied to it. This ensures subsequent requests are routed back to this  
2314           element.

2315           The URI placed in the **Record-Route** header value **SHOULD** be such that if a subsequent request is  
2316           received with this URI in the **Request-URI**, the proxy's normal request processing will cause it to be  
2317           forwarded to one of the previous elements, including the originating client, traversed by the original  
2318           request. This improves robustness, ensuring that the **Request-URI** contains enough information to  
2319           forward subsequent requests to a reasonable destination even in the absence of **Route** headers.

2320           The URI placed in the **Record-Route** header value **MUST** vary with the **Request-URI** in the received  
2321           request. A request may legitimately pass through this proxy more than once on the way to its final  
2322           destination (this is called a spiraling request). The **Request-URI** will be different each time the  
2323           request passes through. If this proxy places the same URI in the **Record-Route** header field each time,  
2324           subsequent requests will be rejected as looped requests. It is insufficient to simply copy the **Request-**  
2325           **URI** from each request into the **Record-Route** header. Some modification, such as adding an **maddr**  
2326           parameter, is necessary.

2327           URIs satisfying the above paragraphs can be constructed in many ways. One way is to use a URI that  
2328           is nearly the same as the **Contact** header in the initial request (if present, else the **From** field), but with  
2329           the **maddr** and port set to resolve to the proxy, and with a transaction identifier added to the user part of  
2330           the request-URI (in order to meet the requirement that the URI in the **Record-Route** be different for  
2331           each distinct **Request-URI**). A call stateful proxy could use a URI of the form sip:proxy.example.com  
2332           and use information from the stored call state to meet the requirements.

2333           The proxy **MAY** include **Record-Route** header parameters in the value it provides. These will be  
2334           returned in some responses to the request (200 responses to **INVITE** for example) and may be useful  
2335           for pushing state into the message.

2336           The **Record-Route** process is designed to work for any SIP request that initiates a dialog. The only  
2337           such request in this specification is **INVITE**. Extensions to the protocol **MAY** define others, and the  
2338           mechanisms described here will apply. The request that initiates a dialog and all refreshes (re-**INVITE**  
2339           for example) **MUST** have **Record-Route** header values added to them if the proxy wishes to remain  
2340           in the request path. This means a proxy will often need to record-route requests that contain **Route**  
2341           headers. Section 12 describes how this will affect a dialog.

2342           Including **Record-Route** even when **Route** headers already exist in a request improves robustness in the  
2343           presence of a preloaded **Route** header field and recovery from endpoint failure.

2344           A proxy **MAY** insert a **Record-Route** header into any request. If a proxy needs to be in the path of  
2345           any type of dialog (such as one straddling a firewall), it **SHOULD** add a **Record-Route** header value  
2346           to every request with a method it does not understand since that method may have dialog semantics.

2347           Generally, the choice about whether to record-route or not is a tradeoff of features vs. performance.  
2348           Faster request processing and higher scalability is achieved when proxies do not record route. How-  
2349           ever, provision of certain services may require a proxy to observe all messages in a dialog. It is

2350 RECOMMENDED that proxies do not automatically record route. They should do so only if speci-  
2351 cally required.

#### 2352 7. Adding Additional Headers

2353 The proxy MAY add any other appropriate headers to the copy at this point.

#### 2354 8. Forward Request

2355 A stateful proxy creates a new client transaction for this request as described in Section 17.1. If  
2356 the next-hop location used in building this request contains the optional addressing parameters, the  
2357 transaction is instructed to send the request based on those parameters. Otherwise, the proxy uses  
2358 the procedures of Section [8] to compute an ordered set of addresses from the Request-URI, and  
2359 as described there, attempts to contact the first one by instructing the client transaction to send the  
2360 request there. If this fails, the stateful proxy continues down the list. Each attempt is a new client  
2361 transaction, and therefore represents a new branch, so that the processing described above for each  
2362 branch would need to be repeated. This results in a requirement to use a different branch ID parameter  
2363 for each attempt.

### 2364 16.6 Response Processing

2365 When a response is received by an element, it first tries to locate a client transaction (Section 17.1.3) match-  
2366 ing the response. If none is found, the element MUST process the response (even if it is an informational  
2367 response) as a stateless proxy (described below). If a match is found, the response is handed to the client  
2368 transaction.

2369 Forwarding responses for which a client transaction (or more generally any knowledge of having sent an asso-  
2370 ciated request) is not found improves robustness. In particular, it ensures that “late” 2xx class responses to INVITE  
2371 requests are forwarded properly.

2372 As client transactions pass responses to the proxy layer, the following processing MUST take place:

- 2373 1. Find the appropriate response context
- 2374 2. Remove the topmost Via
- 2375 3. Add the response to the response context
- 2376 4. Check to see if this response should be forwarded

2377 The following processing MUST be performed on each response that is forwarded. Note that more than  
2378 one response to each request will likely be forwarded - each provisional and one final at the least.

- 2379 1. Aggregate authorization header fields if necessary
- 2380 2. Forward the response
- 2381 3. Generate any necessary CANCEL requests

2382 If no final response has been forwarded after every client transaction associated with the response context  
2383 has been terminated, the proxy must choose and forward the “best” response from those it has seen so far.

2384 Each of the above steps are detailed below:

## 2385 1. Find Context

2386 The proxy locates the “response context” it created before forwarding the original request using the  
2387 key described in Section 16.5. The remaining processing steps take place in this context.

## 2388 2. Via

2389 The proxy removes the topmost Via field value from the response. The address in this value necessar-  
2390 ily matches the proxy since the response matched a client transaction above. The branch parameter  
2391 from this value can be used to determine which branch the response corresponds to.

2392 If no Via field values remain in the response, the response was meant for this element and MUST  
2393 NOT be forwarded. The remainder of the processing described in this section is not performed on this  
2394 message. This will happen, for instance, when the element generates CANCEL requests as described  
2395 in Section sec:proxy-response-processing-cancel.

## 2396 3. Add response to context

2397 Final responses received are stored in the response context until a final response is generated on  
2398 the server transaction associated with this context. The response may a candidate for the best final  
2399 response to be returned on that server transaction. Information from this response may be needed in  
2400 forming the best response even if this response is not chosen.

2401 If the proxy chooses to recurse on a 3xx class response, it MUST NOT add the response to the response  
2402 context

## 2403 4. Check response for forwarding

2404 Until a final response has been sent on the server transaction, the following responses MUST be for-  
2405 forwarded immediately:

- 2406 • Any provisional response other than 100 Trying
- 2407 • Any 2xx response

2408 If a 6xx response is received, it is not immediately forwarded, but the stateful proxy SHOULD cancel  
2409 all pending transactions as described in Section 9.

2410 This is a change from RFC2543, which mandated that the 6xx be forwarded immediately. The problem  
2411 with this is that it is possible for a 2xx to arrive on another branch, in which case the proxy would have to  
2412 forward that in the case of an INVITE transaction. The result is that the UAC could receive a 6xx followed by  
2413 a 2xx, which should never be allowed to happen. So, instead, upon receiving a 6xx, a proxy will CANCEL,  
2414 which will generally result in 487s to all outstanding client transactions, and then at that point the 6xx is  
2415 forwarded upstream.

2416 After a final response has been sent on the server transaction, the following responses MUST be for-  
2417 forwarded immediately:

- 2418 • Any 2xx class response to an INVITE request

2419 A stateful proxy MUST NOT immediately forward any other responses. In particular, a stateful proxy  
2420 MUST NOT forward any 100 Trying response. Those responses that are candidates for forwarding later  
2421 as the “best” response have been gathered as described in step “Add Response to Context”.

2422 Any response chosen for immediate forwarding MUST be processed as described in steps “Aggregate  
2423 authorization headers” through “Record-Route”.

## 2424 5. Choosing the best response

2425 A stateful proxy MUST send a final response to a response context's server transaction if no final  
2426 responses have been immediately forwarded by the above rules and all client transactions in this  
2427 response context have been terminated.

2428 The stateful proxy MUST choose the "best" final response among those received and stored in the  
2429 response context.

2430 If there are no final responses in the context, the proxy MUST send a 408 (Request Timeout) response  
2431 to the server transaction.

2432 Otherwise, the proxy MUST forward one of the responses from the lowest response class stored in the  
2433 response context. The proxy MAY select any response within that lowest class. The proxy SHOULD  
2434 give preference to responses that provide information affecting resubmission of this request, such as  
2435 401, 407, 415, 420, and 484.

2436 A proxy which receives a 503 response SHOULD NOT forward it upstream unless it can determine that  
2437 any subsequent requests it might proxy will also generate a 503. In other words, forwarding a 503  
2438 means that the proxy knows it cannot service any requests, not just the one for the Request-URI in  
2439 the request which generated the 503.

2440 The forwarded response MUST be processed as described in steps "Aggregate authorization headers"  
2441 through "Record-Route".

2442 For example, if a proxy forwarded a request to 4 locations, and received 503, 407, 501, and 404  
2443 responses, it may choose to forward the 407 response.

2444 1xx and 2xx class responses may be involved in the establishment dialogs. When a request does not  
2445 contain a To tag, the To tag in the response is used by the UAC to distinguish multiple responses to  
2446 a dialog creating request. A proxy MUST NOT insert a tag into the To header of a 1xx or 2xx class  
2447 response if the request did not contain one. A proxy MUST NOT modify the tag in the To header of a  
2448 1xx or 2xx class response.

2449 3-6xx class responses are delivered hop-hop. When issuing a 3-6xx class response, the element is  
2450 effectively acting as a UAS, issuing its own response, usually based on the responses received from  
2451 downstream elements. An element SHOULD preserve the To tag when simply forwarding a 3-6xx  
2452 class response to a request that did not contain a To tag.

2453 A proxy MUST NOT modify the To tag in any forwarded response to a request that contains a To tag.

2454 While it makes no difference to the upstream elements if the proxy replaced the To tag in a forwarded  
2455 3-6xx class response, preserving the original tag may assist with debugging.

2456 When the proxy is aggregating information from several responses, choosing a To tag from among them  
2457 is arbitrary, and generating a new To tag may make debugging easier. This happens, for instance, when  
2458 combining 401 and 407 challenges, or combining Contact values from unencrypted and unauthenticated 3xx  
2459 class responses.

## 2460 6. Aggregate authorization headers

2461 If the selected response is a 401 or 407, the proxy MUST collect any WWW-Authenticate and Proxy-  
2462 Authenticate header fields from all other 401 and 407 responses received so far in this response  
2463 context and add them to this response before forwarding.



2464 This is necessary because any or all of the destinations the request was forwarded to may have re-  
2465 requested credentials. The client must receive all of those challenges and supply credentials for each of  
2466 them when it retries the request. Motivation for this behavior is provided in Section 20.

## 2467 7. Record-Route

2468 If the selected response contains a **Record-Route** header field value originally provided by this proxy,  
2469 the proxy *MAY* chose to rewrite the value before forwarding the response. This allows the proxy to  
2470 provide different URIs for itself to the next upstream and downstream elements. A proxy may choose  
2471 to use this mechanism for any reason. For instance, it is useful for multi-homed hosts.

2472 The new URI provided by the proxy *MUST* satisfy the same constraints on URIs placed in **Record-**  
2473 **Route** header fields in requests (see Section 6) with the following modifications:

2474 The URI *SHOULD NOT* contain the transport parameter unless the proxy has knowledge that the next  
2475 upstream (as opposed to downstream) element that will be in the path of subsequent requests supports  
2476 that transport.

2477 The URI placed in the **Record-Route** header value *SHOULD* be such that if a subsequent request is  
2478 received with this URI in the **Request-URI**, the proxy's normal request processing will cause it to  
2479 be forwarded to the same next-hop element (as opposed to some previous element) as the originally  
2480 forwarded request.

2481 When a proxy does decide to modify the **Record-Route** header in the response, one of the operations  
2482 it must perform is to locate the **Record-Route** that it had inserted. If the request spiraled, and the  
2483 proxy inserted a **Record-Route** in each iteration of the spiral, locating the correct header in the  
2484 response (which must be the proper iteration in the reverse direction) is tricky. Note that the rules  
2485 above dictate that a proxy insert a different URI into the **Record-Route** for each distinct **Request-**  
2486 **URI** received. The two issues can be solved jointly. A *RECOMMENDED* mechanism is for the proxy  
2487 to append a piece of data to the user portion of the URI. This piece of data is a hash of the transaction  
2488 key for the incoming request, concatenated with a unique identifier for the proxy instance. Since the  
2489 transaction key includes the **Request-URI**, this key will be unique for each distinct **Request-URI**.  
2490 When the response arrives, the proxy modifies the first **Record-Route** whose identifier matches the  
2491 proxy instance. The modification results in a URI without this piece of data appended to the user  
2492 portion of the URI. Upon the next iteration, the same algorithm (find the topmost **Record-Route**  
2493 header with the parameter) will correctly extract the next **Record-Route** header inserted by that  
2494 proxy.

## 2495 8. Forward response

2496 After performing the processing described in steps "Aggregate authorization headers" through "Record-  
2497 Route", the proxy may perform any feature specific manipulations on the selected response. Unless  
2498 otherwise specified, the proxy *MUST NOT* remove the message body or any header values other than  
2499 the **Via** header value discussed in Section refsec:proxy-response-processing-via. The proxy *MUST*  
2500 pass the response to the server transaction associated with the response context. This will result in  
2501 the response being sent to the location now indicated in the topmost **Via** field value. If the server  
2502 transaction is no longer available to handle the transmission, the element *MUST* forward the response  
2503 statelessly by sending it to the server transport.

2504 Even after forwarding a final response, the proxy *MUST* maintain the response context until all of its  
2505 associated transactions have been terminated.

## 2506 9. Generate CANCELs

2507 OPEN ISSUE #7: If CANCEL is restricted to INVITE only, this behavior must restrict itself to  
2508 INVITE requests.

2509 If the forwarded response was a final response, the proxy MUST generate a CANCEL request for all  
2510 pending client transactions associated with this response context. A proxy SHOULD also generate a  
2511 CANCEL request for all pending client transactions associated with this response context when it  
2512 receives a 6xx response. A pending client transaction is one that has received a provisional response,  
2513 but no final response and has not had an associated CANCEL generated for it. Generating CANCEL  
2514 requests is described in Section 9.1.

## 2515 16.7 Handling transport errors

2516 If the transport layer notifies a proxy of an error when it tries to forward a request (see Section 19.4), the  
2517 proxy MUST behave as if the forwarded request received a 400 response.

2518 If the proxy is notified of an error when forwarding a response, it drops the response. The proxy SHOULD  
2519 NOT cancel any outstanding client transactions associated with this response context due to this notification.

2520 If a proxy cancels its outstanding client transactions, a single malicious or misbehaving client can cause all  
2521 transactions to fail through its Via header field.

## 2522 16.8 CANCEL Processing

2523 A stateful proxy may generate a CANCEL to any other request it has generated at any time. For instance,  
2524 it may choose to generate CANCELs based on having a transaction exceed the time specified in the Ex-  
2525 pired header of certain requests, or as a result of any logic it applies while forwarding requests. A proxy  
2526 MUST cancel any pending client transactions associated with a response context when it receives a matching  
2527 CANCEL request.

2528 OPEN ISSUE #185: Should generating CANCEL at a proxy based on Expires in INVITE be deprecated?

2529 While a CANCEL request is handled in a stateful proxy by its own server transaction, a new response  
2530 context is not created for it. Instead, the proxy layer searches its existing response contexts for the server  
2531 transaction handling the request associated with this CANCEL. If a matching response context is found, the  
2532 element MUST immediately return a 200 OK response to the CANCEL request. In this case, the element is  
2533 acting as a user agent server as defined in Section 8.2. Furthermore, the element MUST generate CANCEL  
2534 requests for all pending client transactions in the context as described in Section 9.

2535 If a response context is not found, the element does not have any knowledge of the request to apply  
2536 the CANCEL to. It MUST forward the CANCEL request (it may have statelessly forwarded the associated  
2537 request previously).

## 2538 16.9 Stateless proxy

2539 When acting statelessly, a proxy is a simple message forwarder. Much of the processing performed when  
2540 acting statelessly is the same as when behaving statefully. The differences are detailed here.

2541 A stateless proxy does not have any notion of a transaction, or of the response context used to describe  
2542 stateful proxy behavior. Instead, the stateless proxy takes messages, both requests and responses, directly  
2543 from the transport layer (See section 19). As a result, stateless proxies do not retransmit messages on their

2544 own. They do, however, forward all retransmission they receive (they do not have the ability to distinguish  
2545 a retransmission from the original message). Furthermore, when handling a request statelessly, an element  
2546 MUST NOT generate its own 100 Trying (or any other provisional) response.

2547 A stateless proxy must validate a request as described in Section 16.3

2548 A stateless proxy must make a routing decision as described in Section 16.4 with the following excep-  
2549 tion:

- 2550 • A stateless proxy MUST choose one and only one destination from the destination set. This choice  
2551 MUST only rely on fields in the message and time-invariant properties of the server. In particular, a  
2552 retransmitted request MUST be forwarded to the same destination each time it is processed. Further-  
2553 more, CANCEL and non-Routed ACK requests MUST generate the same choice as their associated  
2554 INVITE.

2555 A stateless proxy must process the request before forwarding as described in Section 16.5 with the  
2556 following exceptions:

- 2557 • The **branch** parameter on the inserted **Via** header field MUST be the same each time a retransmitted  
2558 request is forwarded. Thus for a stateless proxy, the **branch** parameter calculation MUST **only** depend  
2559 on message parameters affecting the routing of the request which are invariant on retransmission.
- 2560 • The **branch** parameter MUST vary with the value of the **branch** parameter of the topmost **Via** field  
2561 value in the original request. If two requests arrive with different topmost **Via** field values, the top-  
2562 most **Via** field values in the resulting forwarded requests MUST be different. This is necessary to avoid  
2563 merging requests as they traverse the proxy. One way to ensure this when forwarding requests state-  
2564 lessly is to include the original request's topmost **branch** in the hash calculation forming the second  
2565 part (used to match requests and responses) of the **branch** parameter discussed in Section 16.5 step 3.
- 2566 • The request is sent directly to the transport layer instead of through a client transaction. If the next-  
2567 hop destination parameters don't provide an explicit destination, the element applies the procedures  
2568 of [8] to the **Request-URI** to determine where to send the request.

2569 Since a stateless proxy must forward retransmitted requests to the same destination and add identical **branch**  
2570 parameters to each of them, it can only use information from the message itself and time-invariant configuration  
2571 data for those calculations. If the configuration state is not time-invariant (for example, if a routing table is updated)  
2572 any requests that could be affected by the change may not be forwarded statelessly during an interval equal to the  
2573 transaction timeout window before or after the change. The method of processing the affected requests in that  
2574 interval is an implementation decision. A common solution is to forward them transaction statefully.

2575 Stateless proxies MUST NOT perform special processing for **CANCEL** requests. They are processed by  
2576 the above rules as any other requests. In particular, a stateless proxy applies normal **Route** header processing  
2577 to **CANCEL** requests.

2578 Response processing as described in Section 16.6 does not apply to a proxy behaving statelessly. When  
2579 a response arrives at a stateless proxy, the proxy inspects the address in the first (topmost) **Via** header value.  
2580 If that address matches the proxy, the proxy MUST remove that value from the response and forward the  
2581 result to the location indicated in the next **Via** header value. Unless specified otherwise, the proxy MUST  
2582 NOT remove any other header values or the message body. If the address does not match the proxy, the  
2583 message MUST be silently discarded.

2584 **17 Transactions**

2585 SIP is fundamentally a transactional protocol. This means that interactions between components take place  
 2586 in a series of independent message exchanges. Specifically, a SIP transaction consists of a single request,  
 2587 and any responses to that request (which include zero or more provisional responses and one or more final  
 2588 responses). In the case of a transaction where the request was an INVITE (known as an INVITE transaction),  
 2589 the transaction also includes the ACK only if the final response was not a 2xx response. If the response was  
 2590 a 2xx, the ACK is not considered part of the transaction.

2591 The reason for this separation is rooted in the importance of delivering all 200 OK responses to an INVITE to  
 2592 the UAC. To deliver them all to the UAC, the UAS alone takes responsibility for retransmitting them, and the UAC  
 2593 alone takes responsibility for acknowledging them with ACK. Since this ACK is retransmitted only by the UAC, it  
 2594 is effectively considered its own transaction.

2595 Transactions have a client side and a server side. The client side is known as a client transaction, and the  
 2596 server side, as a server transaction. The client transaction sends the request, and the server transaction sends  
 2597 the response. The client and server transactions are logical functions that are embedded in any number of  
 2598 elements. Specifically, they exist within user agents and stateful proxy servers. Consider the example of  
 2599 Section 4. In this example, the UAC executes the client transaction, and its outbound proxy executes the  
 2600 server transaction. The outbound proxy also executes a client transaction, which sends the request to a  
 2601 server transaction in the inbound proxy. That proxy also executes a client transaction, which in turn, sends  
 2602 the request to a server transaction in the UAS. This is shown pictorially in Figure 4.

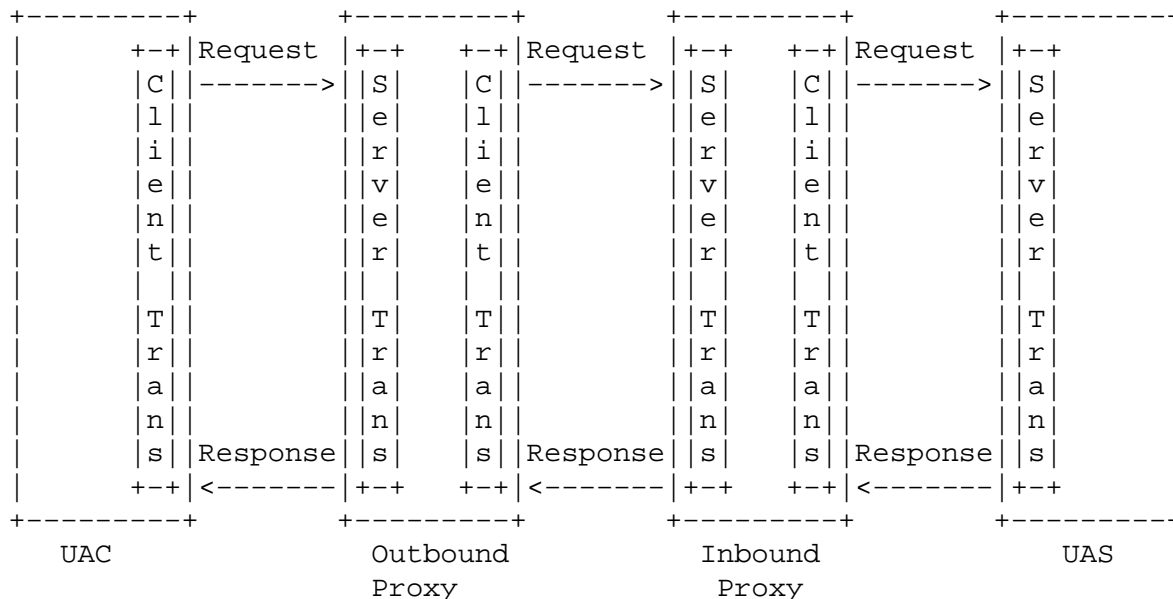


Figure 4: Transaction relationships

2603 A stateless proxy does not contain a client or server transaction. The transaction exists between the  
 2604 UA or stateful proxy on one side of the stateless proxy, and the UA or stateful proxy on the other side.  
 2605 As far as SIP transactions are concerned, stateless proxies are effectively transparent. The purpose of the  
 2606 client transaction is to receive a request from the element the client is embedded in (call this element the  
 2607 "Transaction User" or TU; it can be a UA or a stateful proxy), and reliably deliver the request to that server

2608 transaction. The client transaction is also responsible for receiving responses, and delivering them to the  
2609 TU, filtering out any retransmissions or disallowed responses (such as a response to ACK). In the case of  
2610 an INVITE transaction, that includes generation of the ACK request for any final response excepting a 2xx  
2611 response.

2612 Similarly, the purpose of the server transaction is to receive requests from the transport layer, and deliver  
2613 them to the TU. The server transaction filters any request retransmissions from the network. The server  
2614 transaction accepts responses from the TU, and delivers them to the transport layer for transmission over the  
2615 network. In the case of an INVITE transaction, it absorbs the ACK request for any final response excepting  
2616 a 2xx response.

2617 The 2xx response, and the ACK for it, have special treatment. This response is retransmitted only by a  
2618 UAS, and its ACK generated only by the UAC. This end-to-end treatment is needed so that a caller knows  
2619 the entire set of users that have accepted the call. Because of this special handling, retransmissions of the  
2620 2xx response are handled by the UA core, not the transaction layer. Similarly, generation of the ACK for the  
2621 2xx is handled by the UA core. Each proxy along the path merely forwards each 2xx response to INVITE,  
2622 and its corresponding ACK.

2623 A reliable provisional response, and the PRACK for it, also have special treatment. Reliable provisional  
2624 responses are also only retransmitted by the UAS core, and the PRACK generated by the UAC core. Unlike  
2625 ACK, however, PRACK is a normal non-INVITE transaction, which means that it will generate its own final  
2626 response. The reason for this seemingly inexplicable difference between PRACK and ACK is that reliability  
2627 of provisional responses was added on later as an extra feature, and therefore needed to be done within the  
2628 confines of SIP extensibility. SIP extensibility only allowed the additions of new methods which behaved  
2629 like any other non-INVITE method.

## 2630 17.1 Client transaction

2631 The client transaction provides its functionality through the maintenance of a state machine.

2632 The TU communicates with the client transaction through a simple interface. When the TU wishes to  
2633 initiate a new transaction, it creates a client transaction, and passes it the SIP request to send, a value for  
2634 timer C (described below), and an IP address, port, and transport to send it to. The client transaction begins  
2635 execution of its state machine. Valid responses are passed up to the TU from the client transaction.

2636 There are two types of client transaction state machines, depending on the method the request passed  
2637 by the TU. One handles client transactions for INVITE request. This type of machine is referred to as an  
2638 INVITE client transaction. Another type handles client transactions for all requests except INVITE and  
2639 ACK. This is referred to as a non-INVITE client transaction. There is no client transaction for ACK. If the  
2640 TU wishes to send an ACK, it passes one directly to the transport layer for transmission.

2641 The INVITE transaction is different from those of other methods because of its extended duration. Nor-  
2642 mally, human input is required in order to respond to an INVITE. The long delays expected for sending a  
2643 response argue for a three way handshake. Requests of other methods, on the other hand, are expected to  
2644 completely rapidly. In fact, because of its reliance on just a two way handshake, TUs SHOULD respond  
2645 immediately to non-INVITE requests. Protocol extensions which require longer durations for generation of  
2646 a response (such as a new method that does require human interaction) SHOULD instead use two transactions  
2647 - one to send the request, and another in the reverse direction to convey the result of the request.

### 2648 17.1.1 INVITE Client Transaction

2649 **17.1.1.1 Overview of INVITE Transaction** The INVITE transaction consists of a three-way handshake.  
2650 The client transaction sends an INVITE, the server transaction sends responses, and the client transaction  
2651 sends an ACK. For unreliable transports (such as UDP), the client transaction will retransmit requests at an  
2652 interval that starts at T1 seconds and doubles after every retransmission. The request is not retransmitted over  
2653 reliable transports. After receiving a 1xx response, any retransmissions cease altogether, and the client waits  
2654 for further responses. The server transaction can send additional 1xx responses, which are not transmitted  
2655 reliably by the server transaction. If the provisional response needs to be sent reliably, this is handled by  
2656 the TU. Eventually, the server transaction decides to send a final response. For unreliable transports, that  
2657 response is retransmitted periodically, and for reliable transports, its sent once. For each final response that  
2658 is received at the client transaction, the client transaction sends an ACK, the purpose of which is to quench  
2659 retransmissions of the response.

2660 **17.1.1.2 Formal Description** The state machine for the INVITE client transaction is shown in Figure 5.  
2661 The initial state, "calling", MUST be entered when the TU initiates a new client transaction with an INVITE  
2662 request. The client transaction MUST pass the request to the transport layer for transmission (see Section  
2663 19). If an unreliable transport is being used, the client transaction SHOULD start timer A with a value  
2664 of T1, and SHOULD NOT start timer A when a reliable transport is being used (Timer A controls request  
2665 retransmissions). For any transport, the client transaction MUST start timer B with a value of 64\*T1 seconds  
2666 (Timer B controls transaction timeouts).

2667 When timer A fires, the client transaction SHOULD retransmit the request by passing it to the transport  
2668 layer, and SHOULD reset the timer with a value of 2\*T1. When the timer fires 2\*T1 seconds later, the  
2669 request SHOULD be retransmitted again (assuming the client transaction is still in this state). This process  
2670 SHOULD continue, so that the request is retransmitted with intervals that double after each transmission.  
2671 These retransmissions SHOULD only be done while the client transaction is in the "calling" state.

2672 The default value for T1 is 500ms. T1 is an estimate of the RTT between the client and server transac-  
2673 tions. The optional RTT estimation procedure of Section 17.3 MAY be followed, in which case the resulting  
2674 estimate MAY be used instead of 500ms. If no RTT estimation is used, other values MAY be used in private  
2675 networks where it is known that RTT has a different value. On the public Internet, T1 MAY be chosen larger,  
2676 but SHOULD NOT be smaller.

2677 If the client transaction is still in the "calling" when timer B fires, the client transaction SHOULD inform  
2678 the TU that a timeout has occurred. The client transaction MUST NOT generate an ACK. The value of 64\*T1  
2679 is equal to the amount of time required to send seven requests in the case of an unreliable transport.

2680 If the client transaction receives a provisional response while in the "calling" state, it transitions to  
2681 the "proceeding" state. Upon entering this state, the client transaction MUST start timer C with the value  
2682 provided by the TU when the client transaction was created. This timeout dictates how long the client  
2683 transaction waits for a final response before giving up (i.e., roughly how long does it "let the phone ring"). In  
2684 the "proceeding" state, the client transaction SHOULD NOT retransmit the request any longer. Furthermore,  
2685 the provisional response MUST be passed to the TU. Any further provisional responses MUST be passed up  
2686 to the TU while in the "proceeding" state. Passing of all provisional responses is necessary since the TU  
2687 will handle reliability of these messages, and therefore even retransmissions of a provisional response must  
2688 be passed upwards. When timer C fires, the client transaction MUST transition to the terminated state, and  
2689 it MUST inform the TU of the timeout.

2690 When in either the "calling" or "proceeding" states, reception of a response with status code from 300-  
2691 699 MUST cause the client transaction to transition to "completed". The client transaction MUST pass the  
2692 received response up to the TU, and it MUST generate an ACK request, even if the transport is reliable

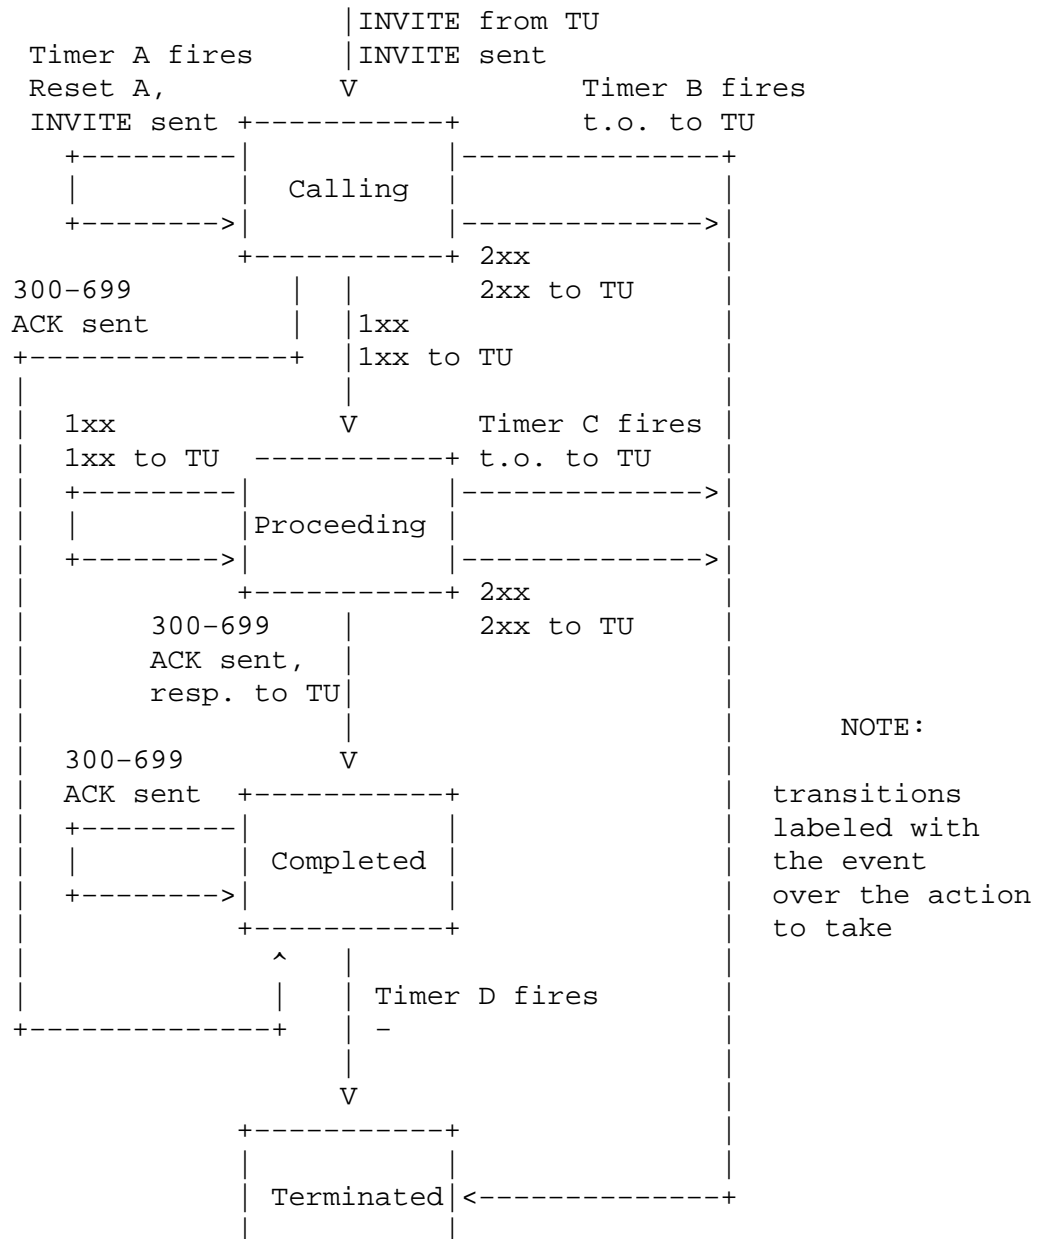


Figure 5: INVITE client transaction

2693 (guidelines for constructing the ACK from the response are given in Section 17.1.1.3) and then pass the ACK  
 2694 to the transport layer for transmission. The ACK MUST be sent to the same address, port and transport that  
 2695 the original request was sent to. The client transaction SHOULD start timer D when it enters the “completed”  
 2696 state, with a value of T3 seconds for unreliable transports, and zero seconds for reliable transports. T3 is  
 2697 the total amount of time that the server transaction can remain in the “completed” state when unreliable  
 2698 transports are used. For the default values of the timers below, this is 16 seconds.

2699 OPEN ISSUE #210: Timer D should be based on the values of the timers selected at the server, but these values

2700 aren't known by the client. We could alternatively specify an absolute minimum.

2701 Any retransmissions of the final response that are received while in the "completed" state SHOULD cause  
2702 the ACK to be re-passed to the transport layer for retransmission, but the newly received response MUST  
2703 NOT be passed up to the TU. A retransmission of the response is defined as any response which would match  
2704 the same client transaction, based on the rules of Section 17.1.3.

2705 If timer D fires while the client transaction is in the "completed" state, the client transaction MUST move  
2706 to the terminated state, and it MUST inform the TU of the timeout.

2707 When in either the "calling" or "proceeding" states, reception of a 2xx response MUST cause the client  
2708 transaction to enter the terminated state, and the response MUST be passed up to the TU. The handling of  
2709 this response depends on whether the TU is a proxy core or a UAC core. A UAC core will handle generation  
2710 of the ACK for this response, while a proxy core will always forward the 200 OK upstream. The differing  
2711 treatment of 200 OK between proxy and UAC is the reason that handling of it does not take place in the  
2712 transaction layer.

2713 The client transaction MUST be destroyed the instant it enters the terminated state. This is actually nec-  
2714 essary to guarantee correct operation. The reason is that 2xx responses to an INVITE are treated differently;  
2715 each one is forwarded by proxies, and the ACK handling in a UAC is different. Thus, each 2xx needs to be  
2716 passed to a proxy core (so that it can be forwarded) and to a UAC core (so it can be acknowledged). No  
2717 transaction layer processing takes place. Whenever a response is received by the transport, if the transport  
2718 layer finds no matching client transaction (using the rules of Section 17.1.3, the response is passed directly  
2719 to the core. Since the matching client transaction is destroyed by the first 2xx, subsequent 2xx will find no  
2720 match and therefore be passed to the core.

2721 **17.1.1.3 Construction of the ACK Request** The ACK request constructed by the client transaction  
2722 MUST contain values for the Call-ID, From, and Request-URI which are equal to the values of those  
2723 headers in the request that created the client transaction (call this the "original request"). The To field in the  
2724 ACK MUST equal the To field in the response being acknowledged, and will therefore usually differ from  
2725 the To field in the original request by the addition of the tag parameter. The ACK MUST contain a single  
2726 Via header, and this MUST be equal to the top Via header of the original request. The ACK request MUST  
2727 contain the same Route headers as the request whose response it is acknowledging. The CSeq header in  
2728 the ACK MUST contain the same value for the sequence number as was present in the original request, but  
2729 the method parameter MUST be equal to "ACK".

2730 These rules for construction of ACK only apply to the client transaction. A UAC core which generates  
2731 an ACK for 2xx MUST instead follow the rules described in Section 13.

2732 For example, consider the following request:

```
2733 INVITE sip:bob@biloxi.com SIP/2.0
2734 Via: SIP/2.0/UDP 10.1.3.3
2735 To: Bob <sip:bob@biloxi.com>
2736 From: Alice <sip:alice@atlanta.com>;tag=88sja8x
2737 Call-ID: 987asjd97y7atg@10.1.3.3
2738 CSeq: 986759 INVITE
```

2739 The ACK request for a non-2xx final response to this request would look like:



2740 ACK sip:bob@biloxi.com SIP/2.0  
2741 Via: SIP/2.0/UDP 10.1.3.3  
2742 To: Bob <sip:bob@biloxi.com>;tag=99sa0xk  
2743 From: Alice <sip:alice@atlanta.com>;tag=88sja8x  
2744 Call-ID: 987asjd97y7atg@10.1.3.3  
2745 CSeq: 986759 ACK

## 2746 17.1.2 non-INVITE Client Transaction

2747 **17.1.2.1 Overview of the non-INVITE Transaction** non-INVITE transactions do not make use of ACK.  
2748 They are a simple request-response interaction. For unreliable transports, requests are retransmitted at an  
2749 interval which starts at T1, and doubles until it hits T2. If a provisional response is received, retransmis-  
2750 sions continue for unreliable transports, but at an interval of T2. The server transaction retransmits the last  
2751 response it sent (which can be a provisional or final response) only when a retransmission of the request is  
2752 received. This is why request retransmissions need to continue even after a provisional response, they are  
2753 what ensure reliable delivery of the final response.

2754 Unlike an INVITE transaction, a non-INVITE transaction has no special handling for the 2xx response.  
2755 The result is that only a single 2xx response to a non-INVITE is ever delivered to a UAC.

2756 **17.1.2.2 Formal Description** The state machine for the non-INVITE client transaction is shown in Fig-  
2757 ure 6. It is very similar to the state machine for INVITE.

2758 The “Trying” state is entered when the TU initiates a new client transaction with a request. When  
2759 entering this state, the client transaction SHOULD set Timer F to fire in T3 seconds. The request MUST be  
2760 passed to the transport layer for transmission. If an unreliable transport is in use, the client transaction MUST  
2761 set timer E to fire in T1 seconds. If timer E fires while still in this state, the timer is reset, but this time with a  
2762 value of  $\text{MIN}(2 * T1, T2)$ . When the timer fires again, it is reset to a  $\text{MIN}(4 * T1, T2)$ . This process continues,  
2763 so that retransmissions occur with an exponentially increasing interval that caps at T2. The default value  
2764 of T2 is 4s, and it represents the amount of time a non-INVITE server transaction will take to respond to a  
2765 request, if it does not respond immediately. For the default values of T1 and T2, this results in intervals of  
2766 500 ms, 1 s, 2 s, 4 s, 4 s, 4s, etc.

2767 If Timer F fires while the client transaction is still in the “Trying” state, the client transaction SHOULD  
2768 inform the TU about the timeout, and then it SHOULD enter the “Terminated” state. If a provisional response  
2769 is received while in the “Trying” state, the response MUST be passed to the TU, and then the client transaction  
2770 SHOULD move to the “Proceeding” state. If a final response (status codes 200-699) is received while in the  
2771 “Trying” state, the response MUST be passed to the TU, and the client transaction MUST transition to the  
2772 “Completed” state.

2773 If Timer E fires while in the “Proceeding” state, the request MUST be passed to the transport layer  
2774 for retransmission, and Timer E MUST be reset with a value of T2 seconds. If timer F fires while in the  
2775 “Proceeding” state, the TU MUST be informed of a timeout, and the client transaction MUST transition to the  
2776 terminated state. If a final response (status codes 200-699) is received while in the “Proceeding” state, the  
2777 response MUST be passed to the TU, and the client transaction MUST transition to the “Completed” state.

2778 Once the client transaction enters the “Completed” state, it MUST set Timer K to fire in T4 seconds for  
2779 unreliable transports, and zero seconds for reliable transports. The “Completed” state exists to buffer any  
2780 additional response retransmissions that may be received (which is why the client transaction remains there  
2781 only for unreliable transports). T4 represents the amount of time the network will take to clear messages



2786 to terminated, is new.

2787 Once the transaction is in the terminated state, it **MUST** be destroyed. As with client transactions, this is  
2788 needed to ensure reliability of the 2xx responses to INVITE.

### 2789 **17.1.3 Matching Responses to Client Transactions**

2790 When the transport layer in the client receives a response, it has to figure out which client transaction will  
2791 handle the response, so that the processing of Sections 17.1.1 and 17.1.2 can take place.

2792 A response matches a client transaction through a comparison process with fields in the request that  
2793 created the transaction. Specifically, the From, Call-ID, CSeq, and the topmost Via header **MUST** match  
2794 the same fields in the request, using the matching operations for those headers defined in Section 22. If  
2795 the To field in the request had a tag, the To field in the response **MUST** match the To field in the request,  
2796 as described in Section 22.39. However, if the To field in the request did not contain a tag, the To field in  
2797 the response **MUST** match that in the request, except that the tag **MUST NOT** be considered as part of the  
2798 matching process. This is needed since a UAS will add a tag to the To field of the response.

2799 A response which matches a transaction match by a previous response is considered a retransmission of  
2800 that response.

### 2801 **17.1.4 Handling Transport Errors**

2802 When the client transaction sends a request to the transport layer to be sent, the following procedures are  
2803 followed if the transport layer indicates a failure.

2804 The client transaction **SHOULD** inform the TU that a transport failure has occurred, and the client trans-  
2805 action **SHOULD** transition directly to the terminated state.

## 2806 **17.2 Server Transaction**

2807 The server transaction is responsible for the delivery of requests to the TU, and the reliable transmission of  
2808 responses. It accomplishes this through a state machine. Server transactions are created by the core when a  
2809 request is received, and transaction handling is desired for that request (this won't always be the case).

2810 As with the client transactions, the state machine depends on whether the received request is an INVITE  
2811 request or not.

### 2812 **17.2.1 INVITE Server Transaction**

2813 The state diagram for the INVITE server transaction is shown in Figure 7.

2814 When a server transaction is constructed with a request, it enters the "Proceeding" state. The server  
2815 transaction **MUST** generate a 100 response (not any status code - the specific value of 100) unless it knows  
2816 that the TU will generate a provisional or final response within 200 ms, in which case it **MAY** generate a 100  
2817 response. This provisional response is needed to rapidly quench request retransmissions in order to avoid  
2818 network congestion. The request **MUST** be passed to the TU.

2819 The TU passes any number of provisional responses to the server transaction. So long as the server  
2820 transaction is in the "Proceeding" state, each of these **MUST** be passed to the transport layer for transmission.  
2821 They are not sent reliably by the transaction layer (they are not retransmitted by it), and do not cause a  
2822 change in the state of the server transaction. When provisional responses need to be delivered reliably,  
2823 it is handled by the TU, which will retransmit the provisional responses itself, and pass downwards each



2830 server transaction; retransmissions of 2xx responses are handled by the TU. The server transaction MUST  
2831 then transition to the “terminated” state.

2832 While in the “Proceeding” state, if the TU passes a response with status code from 300 to 699 to the  
2833 server transaction, the response MUST be passed to the transport layer for transmission, and the state machine  
2834 MUST enter the “Completed” state. For unreliable transports, timer G is set to fire in T1 seconds, and is not  
2835 set to fire for reliable transports.

2836 This is a change from RFC2543, where responses were always retransmitted, even over reliable transports.

2837 When the “Completed” state is entered, timer H MUST be set to fire in  $64 * T1$  seconds, for all transports.  
2838 Timer H determines when the server transaction gives up retransmitting the response. Its value is chosen to  
2839 equal Timer B, the amount of time a client transaction will continue to retry sending a request. If timer G  
2840 fires, the response is passed to the transport layer once more for retransmission, and timer G is set to fire in  
2841  $\text{MIN}(2 * T1, T2)$  seconds. From then on, when timer G fires, the response is passed to the transport again for  
2842 transmission, and timer G is reset with a value that doubles, unless that value exceeds T2, in which case it  
2843 is reset with the value of T2. This is identical to the retransmit behavior for requests in the “Trying” state of  
2844 the non- INVITE client transaction. Furthermore, while in the “completed” state, if a request retransmission  
2845 is received, the server SHOULD pass the response to the transport for retransmission.

2846 If an ACK is received while the server transaction is in the “Completed” state, the server transaction  
2847 MUST transition to the “confirmed” state. As Timer G is ignored in this state, any retransmissions of the  
2848 response will cease.

2849 If timer H fires while in the “Completed” state, it implies that the ACK was never received. In this case,  
2850 the server transaction MUST transition to the terminated state, and MUST indicate to the TU that a transaction  
2851 failure has occurred.

2852 The purpose of the “confirmed” state is to absorb any additional ACK messages that arrive, triggered  
2853 from retransmissions of the final response. When this state is entered, timer I is set to fire in T4 seconds for  
2854 unreliable transports, and zero seconds for reliable transports. Once timer I fires, the server MUST transition  
2855 to the “Terminated” state.

2856 Once the transaction is in the terminated state, it MUST be destroyed. As with client transactions, this is  
2857 needed to ensure reliability of the 2xx responses to INVITE.

### 2858 17.2.2 non-INVITE Server Transaction

2859 The state machine for the non-INVITE server transaction is shown in Figure 8.

2860 The state machine is initialized in the “Trying” state, and is passed a request other than INVITE or  
2861 ACK when initialized. This request is passed up to the TU. Once in the “Trying” state, any further request  
2862 retransmissions are discarded. A request is a retransmission if it matches the same server transaction, using  
2863 the rules specified in Section 17.2.3.

2864 While in the “Trying” state, if the TU passes a provisional response to the server transaction, the server  
2865 transaction MUST enter the “Proceeding” state. The response MUST be passed to the transport layer for  
2866 transmission. Any further provisional responses that are received from the TU while in the “Proceeding”  
2867 state MUST be passed to the transport layer for transmission. If a retransmission of the request is received  
2868 while in the “Proceeding” state, the most recently sent provisional response MUST be passed to the transport  
2869 layer for retransmission. If the TU passes a final response (status codes 200-699) to the server while in the  
2870 “Proceeding” state, the transaction MUST enter the “Completed” state, and the response MUST be passed to  
2871 the transport layer for transmission.

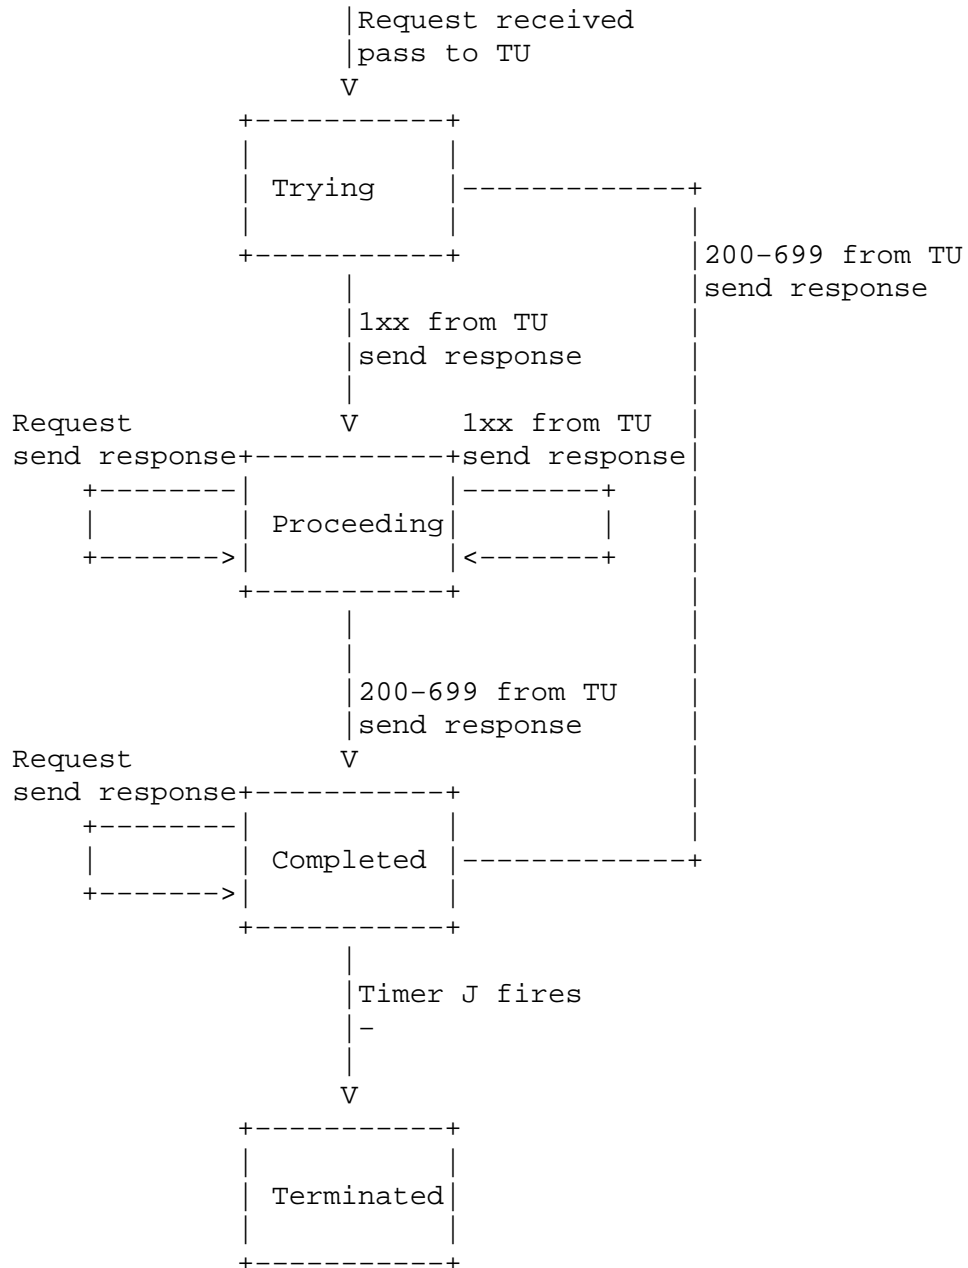


Figure 8: non-INVITE server transaction

2872 When the server transaction enters the “Completed” state, it MUST set Timer J to fire in T3 seconds for  
 2873 unreliable transports, and zero seconds for reliable transports. While in the “Completed” state, the server  
 2874 transaction MUST pass the final response to the transport layer for retransmission whenever a retransmission  
 2875 of the request is received. Any other final responses passed by the TU to the server transaction MUST be  
 2876 discarded while in the “Completed” state. The server transaction remains in this state until Timer J fires, at  
 2877 which point it MUST transition to the “Terminated” state.

2878 The server transaction MUST be destroyed the instant it enters the “Terminated” state.

### 2879 **17.2.3 Matching Requests to Server Transactions**

2880 When an INVITE or ACK request is received from the network by the server, it has to be matched to an  
2881 existing INVITE transaction. The INVITE request matches a transaction if the Request-URI, To, From,  
2882 Call-ID, CSeq, and top Via header match those of the INVITE request which created the transaction. In  
2883 this case, the INVITE is a retransmission of the original one that created the transaction. The ACK request  
2884 matches a transaction if the Request-URI, From, Call-ID, CSeq number (not the method), and top Via  
2885 header match those of the INVITE request which created the transaction, and the To field of the ACK  
2886 matches the To field of the response sent by the server transaction (which then includes the tag). Matching  
2887 is done based on the matching rules defined for each of those headers. The usage of the tag in the To field  
2888 helps disambiguate ACK for 2xx from ACK for other responses at a proxy which may have forwarded both  
2889 responses (which can occur in unusual conditions). An ACK request that matches an INVITE transaction  
2890 matched by a previous ACK is considered a retransmission of that previous ACK.

2891 For all other request methods, a request is matched to a transaction if the Request-URI, To, From,  
2892 Call-ID and Cseq (including the method) and top Via header match those of the request which created  
2893 the transaction. Matching is done based on the matching rules defined for each of those headers. When a  
2894 non-INVITE request matches an existing transaction, it is a retransmission of the request which created that  
2895 transaction.

2896 Because the matching rules include the Request-URI, the server cannot match a response to a transac-  
2897 tion. When the TU passes a response to the server, it must inform the TU which transaction the response is  
2898 for.

## 2899 **17.3 RTT Estimation**

2900 Most of the timeouts used in the transaction state machines derive from T1, which is an estimate of the RTT  
2901 between the client and server transactions. This subsection defines optional procedures that a client can use  
2902 to build up estimates of the RTT to a particular IP address. To perform this procedure, the client MUST  
2903 maintain a table of variables for each destination IP address to which an RTT estimate is being made.

2904 OPEN ISSUE #212: Is destination IP address the right index for an RTT estimate? How about Request-URI?

2905 If a client wishes to measure RTT for a particular IP address, it MUST include a Timestamp header into  
2906 a request containing the time when the request is initially created and passed to a new client transaction,  
2907 which transmits the request. If a 100 response (not any 1xx, only the 100 response) is received before the  
2908 client transaction generates a retransmission, an RTT estimate is made. This is consistent with the RFC  
2909 2988 requirements on TCP for using Karn’s algorithm in RTT estimation.

2910 The estimate, called R, is made by computing the difference between the current time and the value of  
2911 Timestamp header in the 100 response. The value of R is applied to the estimation of RTO as described  
2912 in Section 2 of RFC 2988 [25], with the following differences. First, the initial value of RTO is 500 ms for  
2913 SIP, not 3 s as is used for TCP. Second, there is no minimum value for the RTO, as there is for TCP, if SIP  
2914 is being run on a private network. When run on the public Internet, the minimum is 500 ms, as opposed to  
2915 1 s for TCP. This difference is because of the expected usage of SIP in private networks where rapid call  
2916 setup times are service critical. Once RTO is computed, the timer T1 is set to the value of RTO, and all other  
2917 timers scale proportionally as described above.

## 2918 18 Reliability of Provisional Responses

2919 Normally, provisional responses are not transmitted reliably. The TU generates a single provisional re-  
2920 sponse, and passes it to the server transaction, which sends it once. RFC 2543 provided no means for  
2921 reliable transmission of these messages.

2922 It was later observed that reliability was important in several cases, including interoperability scenarios  
2923 with the PSTN. Therefore, an optional capability was added in this specification to support reliable trans-  
2924 mission of provisional responses.

2925 The reliability mechanism works by mirroring the current reliability mechanisms for 2xx final responses  
2926 to INVITE. Those requests are transmitted periodically by the TU, until a separate transaction, ACK, is  
2927 received, that indicates reception of the 2xx by the UAC. The reliability for the 2xx to INVITE and ACK  
2928 messages are end-to-end. In order to achieve reliability for provisional responses, we do nearly the same  
2929 thing. Reliable provisional responses are retransmitted by the TU with an exponential backoff. Those  
2930 retransmissions cease when a PRACK message is received. The PRACK request plays the same role as  
2931 ACK, but for provisional responses. There is an important difference, however. PRACK is a normal SIP  
2932 message, like BYE. As such, its own reliability is ensured hop-by-hop through each stateful proxy. Similarly,  
2933 PRACK has its own response. If this were not the case, the PRACK message could not traverse existing  
2934 proxy servers.

2935 Each provisional response is given a sequence number, carried in the RSeq header in the response.  
2936 The PRACK messages contain an RACK header, which indicates the sequence number of the provisional  
2937 response which is being acknowledged. The acknowledgements are not cumulative, and the specifications  
2938 recommend a single outstanding provisional response at a time, for purposes of congestion control.

### 2939 18.1 UAS Behavior

2940 A UAS MAY send any non-100 provisional response to INVITE reliably, so long as the initial INVITE request  
2941 (the request whose provisional response is being sent reliably) contained a Supported header with the op-  
2942 tion tag 100rel. While this specification does not allow reliable provisional responses for any method but  
2943 INVITE, extensions that define new methods which can establish dialogs may make use of the mechanism.

2944 The UAS MUST send any non-100 provisional response reliably if the initial request contained a Require  
2945 header with the option tag 100rel. If the UAS is unwilling to do so, it MUST reject the initial request with  
2946 a 420 (Bad Extension) and include a Unsupported header containing the option tag 100rel.

2947 A UAS MUST NOT attempt to send a 100 response reliably. Only provisional responses numbered 101  
2948 to 199 may be sent reliably. If the request did not include either a Supported or Require header indicating  
2949 this feature, the UAS MUST NOT send the provisional response reliably.

2950 100 responses are hop-by-hop only. For this reason, the reliability mechanisms described here, which are end-  
2951 to-end, cannot be used.

2952 An element which can act as a proxy can also send reliable provisional responses; in that case, it acts as  
2953 a UAS for purposes of that transaction. However, it MUST NOT attempt to do so for any request that contains  
2954 a tag in the To field. That is, a proxy cannot generate reliable provisional responses to requests sent within  
2955 the context of a dialog. Of course, unlike a UAS, when the proxy element receives a PRACK that does not  
2956 match any outstanding reliable provisional response, the PRACK MUST be proxied.

2957 The rest of this discussion assumes that the initial request contained a Supported or Require header  
2958 listing 100rel, and that there is a provisional response to be sent reliably.



2959 The provisional response to be sent reliably is constructed by the UAS core according to the procedures  
2960 of Section 8.2.7 and Section 12. Specifically, the provisional response MUST establish a dialog if one  
2961 is not yet created. In addition, it MUST contain **Require** header containing the option tag `100rel`, and  
2962 MUST include an **RSeq** header. The value of the header for the first reliable provisional response in a  
2963 transaction MUST be between 1 and  $2^{31} - 1$ . It is RECOMMENDED that it be chosen uniformly in this  
2964 range. The **RSeq** numbering space is within a single transaction. This means that provisional responses for  
2965 different requests MAY use the same values for the **RSeq** number.

2966 The reliable provisional response is passed to the transaction layer periodically with an interval that  
2967 starts at  $T_1$  seconds and doubles for each retransmission ( $T_1$  is defined in Section 17). Once passed to the  
2968 server transaction, it is added to an internal list of unacknowledged reliable provisional responses.

2969 This differs from retransmissions of `2xx` responses, which cap at  $T_2$  seconds. This is because retransmissions of  
2970 **ACK** are triggered on receipt of a `2xx`, but retransmissions of **PRACK** take place independently of reception of `1xx`.

2971 Retransmissions cease when a matching **PRACK** is received. **PRACK** is like any other request within a  
2972 dialog, and the UAS core processes it according to the procedures of Sections 8.2 and 12.2.2. A matching  
2973 **PRACK** is defined as one within the same dialog as the response, and whose method, **CSeq**-num, and  
2974 response-num in the **RAck** header match, respectively, the method and sequence number from the **CSeq**  
2975 and sequence number from the **RSeq** of the reliable provisional response.

2976 If a **PRACK** request is received that does not match any unacknowledged reliable provisional response,  
2977 the UAS MUST respond to the **PRACK** with a `481` response. If the **PRACK** does match an unacknowledged  
2978 reliable provisional response, it MUST be responded to with a `2xx` response. The UAS can be certain at  
2979 this point that the provisional response has been received in order. It SHOULD cease retransmissions of the  
2980 reliable provisional response, and MUST remove it from the list of unacknowledged provisional responses.

2981 If a reliable provisional response is retransmitted for  $64 * T_1$  seconds without reception of a correspond-  
2982 ing **PRACK**, the UAS SHOULD reject the original request with a `500` class response.

2983 If the **PRACK** contained a body, the body is treated in the same way a body in an **ACK** is treated.

2984 After the first reliable provisional response for a request has been acknowledged, the UAS MAY send  
2985 additional reliable provisional responses. The UAS MUST NOT send a second reliable provisional response  
2986 until the first is acknowledged. After the first, it is RECOMMENDED that the UAS not send an additional  
2987 reliable provisional response until the previous is acknowledged. The first reliable provisional response  
2988 receives special treatment because it conveys the initial sequence number. If additional reliable provisional  
2989 responses were sent before the first was acknowledged, the UAS could not be certain these were received in  
2990 order.

2991 The value of the **RSeq** in each subsequent reliable provisional response for the same request MUST be  
2992 greater by exactly one. **RSeq** numbers MUST NOT wrap around. Because the initial one is chosen to be less  
2993 than  $2^{31} - 1$ , but the maximum is  $2^{32} - 1$ , there can be up to  $2^{31}$  reliable provisional responses per  
2994 request, which is more than sufficient.

2995 Note that the UAS MAY send a final response to the initial request before having received **PRACK**s for  
2996 all unacknowledged reliable provisional responses. In that case, it SHOULD NOT continue to retransmit the  
2997 unacknowledged reliable provisional responses, but it MUST be prepared to process **PRACK** requests for  
2998 those outstanding responses. A UAS MUST NOT send new reliable provisional responses (as opposed to  
2999 retransmissions of unacknowledged ones) after sending a final response to a request.

## 3000 18.2 UAC Behavior

3001 If a provisional response is received for the initial request, and that response contains a **Require** header  
3002 containing the option tag `100rel`, the response is to be sent reliably. If the response is a 100 (as opposed  
3003 to 101 to 199), this option tag **MUST** be ignored, and the procedures below **MUST NOT** be used.

3004 Assuming the response is to be transmitted reliably, the UAC **MUST** create a new request with method  
3005 **PRACK**. This request is sent within the dialog associated with the provisional response (indeed, the provi-  
3006 sional response may have created the dialog). **PRACK** requests **MAY** contain bodies, which are interpreted  
3007 according to their type and disposition.

3008 Note that the **PRACK** is like any other non-**INVITE** request within a dialog. In particular, a UAC  
3009 **SHOULD NOT** retransmit the **PRACK** request when it receives a retransmission of the provisional response  
3010 being acknowledged, although doing so does not create a protocol error.

3011 Once a reliable provisional response is received, retransmissions of that response **MUST** be discarded.  
3012 A response is a retransmission when its dialog ID, **CSeq** and **RSeq** match the original response. The UAC  
3013 **MUST** maintain a sequence number which indicates the most recently received in-order reliable provisional  
3014 response for the initial request. This sequence number **MUST** be maintained until a final response is received  
3015 for the initial request. Its value **MUST** be initialized to the **RSeq** header in the first reliable provisional  
3016 response received for the initial request.

3017 Handling of subsequent reliable provisional responses for the same initial request follows the same rules  
3018 as above, with the following difference. Reliable provisional responses are guaranteed to be in order. As  
3019 a result, if the UAC receives another reliable provisional response to the same request, and its **RSeq** value  
3020 isn't one higher than the value of the sequence number, that response **MUST NOT** be acknowledged with a  
3021 **PRACK**, and **MUST NOT** be processed further by the TU. An implementation **MAY** discard the response, or  
3022 **MAY** cache the response in the hopes of receiving the missing responses.

3023 The UAC **MAY** acknowledge reliable provisional responses received after the final response, or **MAY**  
3024 discard them.

## 3025 19 Transport

3026 The transport layer is responsible for the actual transmission of requests and responses over network trans-  
3027 ports. This includes determination of the connection to use for a request or response, in the case of connec-  
3028 tion oriented transports.

3029 The transport layer is responsible for managing any persistent connections (for transports like TCP, TLS  
3030 and SCTP) including ones it opened, as well as ones opened to it. This includes connections opened by the  
3031 client or server transports, so that connections are shared between client and server transport functions. It is  
3032 **RECOMMENDED** that connections be kept open for some implementation defined time after the last message  
3033 was sent or received over that connection. This time **SHOULD** be at least 16 seconds in order to ensure with  
3034 high probability that responses can be sent over the same connection a request was sent.

3035 All SIP elements **MUST** support UDP at a minimum.

### 3036 19.1 Clients

#### 3037 19.1.1 Sending Requests

3038 The client side of the transport layer is responsible for sending the request and receiving responses. The  
3039 user of the transport layer passes the client transport the request, an IP address, port, transport, and possibly

3040 TTL for multicast destinations.

3041 A client that sends a request to a multicast address MUST add the “maddr” parameter to its Via header  
3042 field, and SHOULD add the “ttl” parameter. (In that case, the maddr parameter SHOULD contain the des-  
3043 tination multicast address, although under exceptional circumstances it MAY contain a unicast address.)  
3044 Requests sent to multicast groups SHOULD be scoped to ensure that they are not forwarded beyond the  
3045 administrative domain to which they were targeted. This scoping MAY be done with either TTL or admin-  
3046 istrative scopes [20], depending on what is implemented in the network.

3047 It is important to note that the layers above the transport layer do not operate differently for multicast  
3048 as opposed to unicast requests. This means that SIP treats multicast more like anycast, assuming that there  
3049 is a single recipient generating responses to requests. If this is not the case, the first response will end  
3050 up “winning”, based on the client transaction rules. Any other responses from different UA will appear  
3051 as retransmissions and be discarded. This limits the utility of multicast to cases where an anycast type of  
3052 function is desired, such as registrations.

3053 OPEN ISSUE #7: This is a proposed resolution to whether or not multicast should be removed entirely.

3054 Before a request is sent, the client transport MUST insert a value of the sent-by field into the Via header.  
3055 This field contains an IP address or host name, and port. In certain cases discussed in Section 19.2.2, this  
3056 IP address and port are used to construct a SIP URI for sending the response. The transport layer MUST  
3057 be prepared to receive incoming connections (and receive responses sent over such connections) on any IP  
3058 addresses and ports that this SIP URI might resolve to using the procedures defined in [8]. The transport  
3059 layer MUST also be prepared to receive an incoming connection on the source IP address that the request  
3060 was sent from, and port number in the sent-by field. The client transport MUST also be prepared to receive  
3061 the response on the same connection used to send the request.

3062 For unreliable unicast transports, the client transport MUST be prepared to receive responses on the  
3063 source IP address that the request is sent from (as responses are sent back to the source address), but the  
3064 port number in the sent-by field. Furthermore, as with reliable transports, in certain cases the IP address and  
3065 port are used to construct a URI for sending the response. The client transport MUST be prepared to receive  
3066 responses on any IP address/port combinations that this SIP URI might resolve to using the procedures of  
3067 [8].

3068 For multicast, the client transport MUST be prepared to receive responses on the same multicast group  
3069 and port that the request is sent to (e.g., it needs to be a member of the multicast group it sent the request  
3070 to.)

3071 If a request is destined to an IP address, port, and transport to which an existing connection is open, it  
3072 is RECOMMENDED that this connection be used to send the request, but another connection MAY be opened  
3073 and used.

3074 If a request is sent using multicast, it is sent to the group address, port, and TTL provided by the transport  
3075 user. If a request is sent using unicast unreliable transports, it is sent to the IP address and port provided by  
3076 the transport user.

### 3077 19.1.2 Receiving Responses

3078 When a response is received, the client transport examines the top Via header. If the value of the sent-by  
3079 parameter in that header does not correspond to a value that the client transport is configured to insert into  
3080 requests, the response MUST be rejected.

3081 If there are any client transactions in existence, the client transport uses the matching procedures of Sec-  
3082 tion 17.1.3 to attempt to match the response to an existing transaction. If there is a match, the response MUST

3083 be passed to that transaction. Otherwise, the response **MUST** be passed to the core (whether it be stateless  
3084 proxy, stateful proxy, or UA) for further processing. Handling of these “stray” responses is dependent on  
3085 the core (a stateless proxy will forward all responses, for example).

## 3086 19.2 Servers

### 3087 19.2.1 Receiving Requests

3088 When the server transport receives a request over any transport, it **MUST** examine the value of the sent-by  
3089 parameter in the top *Via* header field. If the host portion of the sent-by parameter contains a domain name,  
3090 or if it contains an IP address that differs from the packet source address, the server **MUST** add a “received”  
3091 attribute to that *Via* header field. This attribute **MUST** contain the source address that the packet was received  
3092 from. This is to assist the server transport layer in sending the response, since it must be sent to the source  
3093 IP address that the request came from.

3094 Consider a request received by the server transport which looks like, in part:

```
3095 INVITE sip:bob@biloxi.com SIP/2.0  
3096 Via: SIP/2.0/UDP bobspc.biloxi.com:5060
```

3097 The request is received with a source IP address of 1.2.3.4. Before passing the request up, the transport  
3098 would add a received parameter, so that the request would look like, in part:

```
3099 INVITE sip:bob@biloxi.com SIP/2.0  
3100 Via: SIP/2.0/UDP bobspc.biloxi.com:5060;received=1.2.3.4
```

3101 Next, the client transport attempts to match the request to the client transaction. It does so using the  
3102 matching rules described in Section 17.2.3. If a matching server transaction is found, the request is passed  
3103 to that transaction for processing. If no match is found, the request is passed to the core, which may  
3104 decide to construct a new server transaction for that request. Note that when a UAS core sends a 2xx  
3105 response to *INVITE*, the server transaction is destroyed. This means that when the *ACK* arrives, there will  
3106 be no matching server transaction, and based on this rule, the *ACK* is passed to the UAS core, where it is  
3107 processed.

### 3108 19.2.2 Sending Responses

3109 The server transport uses the value of the top *Via* header in order to determine where to send a response. It  
3110 **MUST** follow the following process:

- 3111 • If the “sent-protocol” is a reliable transport protocol such as TCP, TLS or SCTP, the response **MUST**  
3112 be sent using the existing connection to the source of the original request that created the transaction, if  
3113 that connection is still open. This does require the server transport to maintain an association between  
3114 server transactions and transport connections. If that connection is no longer open, the server **MAY**  
3115 open a connection to the IP address in the *received* parameter, if present, using the port in the *sent-*  
3116 *by* value, or the default port for that transport, if no port is specified (5060 for UDP and TCP, 5061  
3117 for TLS and SSL). If that connection attempt fails, the server **SHOULD** use the procedures in [8] for

3118 servers in order to determine the IP address and port to open the connection and send the response to.

- 3119
- 3120 • Otherwise, if the *Via* header field contains a “*maddr*” parameter, forward the response to the address  
3121 listed there, using the port indicated in “*sent-by*”, or port 5060 if none is present. If the address is  
3122 a multicast address, the response SHOULD be sent using the TTL indicated in the “*ttl*” parameter, or  
3123 with a TTL of 1 if that parameter is not present.
  - 3124 • Otherwise (for unreliable unicast transports), if the top *Via* has a *received* parameter, send the re-  
3125 sponse to the address in the “*received*” parameter, using the port indicated in the “*sent-by*” value, or  
3126 using port 5060 if none is specified explicitly. If this fails, e.g., elicits an ICMP “port unreachable”  
3127 response, send the response to the address in the “*sent-by*” parameter. The address to send to is  
3128 determined by following the procedures defined in [8] for servers.
  - 3129 • Otherwise, if it is not receiver-tagged, send the response to the address indicated by the “*sent-by*”  
3130 value.

### 3131 19.3 Framing

3132 In the case of message oriented transports (such as UDP), if the message has a *Content-Length* header, the  
3133 message body is assumed to contain that many bytes. If there are additional bytes in the transport packet  
3134 below the end of the body, they MUST be discarded. If the transport packet ends before the end of the  
3135 message body, this is considered an error. If the message is a response, it MUST be discarded. If its a  
3136 request, the element SHOULD generate a 400 class response. If the message has no *Content-Length* header,  
3137 the message body is assumed to end at the end of the transport packet.

3138 In the case of stream oriented transports (such as TCP), the *Content-Length* header indicates the size  
3139 of the body. The *Content-Length* header MUST be used with stream oriented transports.

### 3140 19.4 Error Handling

3141 Error handling is independent of whether the message was a request or response.

3142 If the transport user asks for a message to be sent over an unreliable transport, and the result is an ICMP  
3143 error, the behavior depends on the type of ICMP error. A host, network, port or protocol unreachable errors,  
3144 or parameter problem errors SHOULD cause the transport layer to inform the transport user of a failure in  
3145 sending. Source quench and TTL exceeded ICMP errors SHOULD be ignored.

3146 If the transport user asks for a request to be sent over a reliable transport, and the result is a connection  
3147 failure, the transport layer SHOULD inform the transport user of a failure in sending.

## 3148 20 Security Considerations

3149 The fundamental security issues confronting SIP are: preserving the confidentiality and integrity of mes-  
3150 saging, preventing replay attacks or message spoofing, providing for the authentication and privacy of the  
3151 participants in a session, and preventing denial of service attacks.

3152 SIP messages frequently contain sensitive information about their senders - not just what they have to  
3153 say, but with whom they communicate, when they communicate and for how long, and from where they

3154 participate in sessions. Many applications and their users require that this sort of private information be  
3155 hidden from any parties that do not need to know it.

3156 Encryption provides the best means to preserve the confidentiality of signaling - it can also guarantee  
3157 that messages are not modified by any malicious intermediaries. However, SIP requests and responses  
3158 cannot be encrypted end-to-end (that is, between a pair of distinct user agents who share encryption keys)  
3159 in their entirety because message fields such as the Request-URI, Route and Via need, in most network  
3160 architectures, to be visible to proxies so that SIP requests are routed correctly. Note that proxy servers need  
3161 to modify signaling as well (adding Via headers) in order for SIP to function. Proxy servers must therefore  
3162 be a part of trust relationships in SIP networks.

3163 Note that there are also less direct ways in which private information can be divulged. If a user or service  
3164 chooses to be reachable at an address that is guessable from the person's name and organizational affiliation  
3165 (which describes most addresses of record), the traditional method of ensuring privacy by having an unlisted  
3166 "phone number" is compromised. A user location service can infringe on the privacy of the recipient of a  
3167 session invitation by divulging their specific whereabouts to the caller; an implementation consequently  
3168 SHOULD be able to restrict, on a per-user basis, what kind of location and availability information is given  
3169 out to certain classes of callers.

3170 SIP entities also have a need to identify one another in a secure fashion. Ordinarily a SIP UA asserts  
3171 an identity for the initiator of a request in the From header field, but in many systems this information  
3172 is controlled directly by the end user, and thus spoofing the contents of the From is trivial. When a SIP  
3173 endpoint asserts the identity of its user to a peer user agent or to a proxy server, that identity should in some  
3174 way be verifiable. A cryptographic authentication mechanism is provided in SIP to address this requirement.

3175 The most comprehensive mechanisms for securing SIP messages (providing confidentiality and integrity  
3176 guarantees for signaling as well as authentication) make use of transport or network layer encryption. en-  
3177 cryptation encrypts the entire SIP request or response on the wire so that packet sniffers or other eavesdroppers  
3178 cannot see who is calling whom.

3179 Note that the security of SIP signaling itself has no bearing on the security of protocols used in concert  
3180 with SIP such as RTP, or with any MIME types carried as SIP bodies, such as SDP. Any media associated  
3181 with a session can be encrypted end-to-end without any of the problems associated with encrypting SIP  
3182 signaling. Media encryption is outside the scope of this document.

## 3183 **20.1 Transport and Network Layer Security**

3184 SIP requests and responses MAY be protected by security mechanisms at the transport or network layer.  
3185 No particular mechanism is mandated by this document, but two popular alternatives are briefly examined:  
3186 protection at the transport layer can be afforded by TLS [26], and network layer security is provided by  
3187 IPsec [27].

3188 Transport or network layer security encrypts signaling traffic, guaranteeing message confidentiality and  
3189 integrity (note however that the originator and recipient of a session may be deducible by observers per-  
3190 forming a network traffic analysis). The keys used to establish encrypt traffic can also be used to verify an  
3191 asserted identity in many architectures, and therefore provide a means of authentication.

3192 IPsec is a network layer protocol - essentially, a secure replacement for traditional IP (Internet Protocol).  
3193 IPsec is most suited to VPN (virtual private network) architectures in which a set of SIP hosts (mingled user  
3194 agents and proxy servers) or bridged administrative domains have a trust relationship with one another.

3195 TLS is a transport protocol and hence, like TCP and UDP, TLS can be specified as the desired transport  
3196 protocol within a Via header field or a SIP-URI. TLS is most suited to architectures in which a chain of trust

3197 joins together a set of hosts (e.g. Alice trusts her local proxy server, which in turn trust Bob's local proxy  
3198 server, which Bob trusts, hence Bob and Alice can communicate securely).

3199 TLS must be tightly coupled with a SIP application. Note that transport mechanisms are specified on  
3200 a hop-by-hop basis in SIP, and that in some networks TLS might be used for only certain portions of the  
3201 signaling path.

3202 It is RECOMMENDED that SIP endpoints support TLS as a secure transport for SIP.

## 3203 20.2 SIP Authentication

3204 SIP provides a stateless challenged-based mechanism for authentication. Any time that a proxy server or  
3205 user agent receives a request, they MAY challenge the initiator of the request to provide assurance of their  
3206 identity. Once the originator has been identified, the recipient of the request SHOULD ascertain whether or  
3207 not this user is authorized to make the request in question. No authorization systems are recommended or  
3208 discussed in this document.

3209 The "Digest" authentication mechanism described in this section provide message authentication only,  
3210 without message integrity or confidentiality. Protective measures above and beyond authentication need to  
3211 be taken to prevent active attackers from modifying and/or replaying SIP requests and responses.

3212 Note that due to its weak security, the usage of "basic" authentication has been deprecated, and that  
3213 servers MUSTNOT accept credentials using the "basic" authorization scheme, and servers also MUSTNOT  
3214 challenge with "basic". This is a change from RFC 2543.

### 3215 20.2.1 Framework

3216 The framework for SIP authentication closely parallels that of HTTP (RFC 2617 [28]). In particular, the  
3217 BNF for auth- scheme, auth-param, challenge, realm, realm-value, and credentials is identical. The  
3218 401 response is used by user agent servers in SIP to challenge the identity of a user agent client. Additionally,  
3219 registrars and redirect servers MAY make use of 401 (Unauthorized) responses for authentication, but proxies  
3220 MUST NOT, and instead MAY use the 407 (Proxy Authentication Required) response. The requirements for  
3221 inclusion of the Proxy-Authenticate, Proxy- Authorization, WWW-Authenticate, and Authorization in  
3222 the various messages are identical to those described in RFC 2617 [28].

3223 Since SIP does not have the concept of a canonical root URL, the notion of protection spaces is in-  
3224 terpreted differently in SIP. The realm string alone defines the protection domain. This is a change from  
3225 RFC2543, in which the Request-URI and the realm together defined the protection domain; this definition  
3226 gave rise to some amount of confusion since the Request-URI sent by the UAC and the Request-URI re-  
3227 ceived by the server issuing a challenge might be different, and indeed the final form of the Request-URI  
3228 might not be known to the UAC. Also, the previous definition depended on the presence of a SIP URI in the  
3229 Request-URI, and seemed to rule out alternative URI schemes (like the tel URL).

3230 Operators of user agents or proxy servers that will authenticate received requests MUST adhere to the  
3231 following guidelines for creation of a realm string representing their server:

- 3232 • Realm strings MUST be globally unique. It is RECOMMENDED that a realm string contain a hostname  
3233 or domain name, following the recommendation in Section 3.2.1 of RFC2617 [[28]].
- 3234 • Realm strings SHOULD present a human-readable identifier that can be rendered to a user.

3235 For example:

3236        INVITE sip:bob@biloxi.com SIP/2.0  
3237        WWW-Authenticate: Digest realm="biloxi.com"

3238        Generally, SIP authentication is for a specific realm, a protection domain. Thus, for Digest authentica-  
3239        tion, each such protection domain has its own set of user names and secrets. If a server does not care about  
3240        authenticating individual users, it may make sense to establish a "global" user name and secret for its realm  
3241        as a default challenge if a particular Request-URI does not have its own realm or set of user names (e.g.  
3242        an INVITE to 'sip:10.3.6.6'). Similarly, UACs representing many users, such as PSTN gateways, MAY have  
3243        their own device-specific credentials for particular realms.

3244        While a server can legitimately challenge most SIP requests, there are two requests defined by the SIP  
3245        standard today that require special handling for authentication: ACK and CANCEL.

3246        Complications with the ACK method arise because it requires no response. Under an authentication  
3247        scheme that uses responses to carry nonces (such as Digest), some problems come up for any requests that  
3248        take no response (including ACK). For this reason any credentials in the INVITE that were accepted by  
3249        a server MUST be accepted by that server for the ACK. UACs creating an ACK message should duplicate  
3250        all of the Authorization and Proxy-Authorization headers that appeared in the INVITE to which the ACK  
3251        corresponds.

3252        Although the CANCEL method does take a response (a 2xx), servers MUSTNOT attempt to challenge  
3253        CANCEL requests since these requests cannot be resubmitted. Generally, a CANCEL request SHOULD be  
3254        accepted by a server if it comes from the same host that sent the request being cancelled (provided that some  
3255        sort of transport or network layer security association, as described above, is in place).

3256        When a challenge is received by a UAC, it SHOULD render to the user the contents of the "realm"  
3257        parameter in the challenge (which appears in either a WWW-Authenticate header or Proxy-Authenticate  
3258        header) if the UAC device does not already know of a credential for the realm in question. A service  
3259        provider that pre-configures UAs with credentials for its realm should be aware that users will not have the  
3260        opportunity to present distinct credentials for this realm when challenged at a pre-configured device.

3261        Finally, note that even if a UAC can locate credentials that are associated with the proper realm, there is  
3262        always a potential that these credentials may no longer be valid, or that for whatever reason the challenging  
3263        server will not accept these credentials. In this instance a server will commonly repeat its challenge. A  
3264        UAC MUSTNOT reattempt requests with the credentials that have just been rejected (unless the request was  
3265        rejected because of a stale nonce).

## 3266        20.2.2 User to User Authentication

3267        When a UAS receives a request from a UAC, the UAS MAY authenticate the originator before the request  
3268        is processed. If no credentials (in the Authorization header field are provided in the request, the UAS can  
3269        challenge the originator to provide credentials by rejecting the request with a 401 (Unauthorized) status  
3270        code.

3271        The WWW-Authenticate response-header field MUST be included in 401 (Unauthorized) response mes-  
3272        sages. The field value consists of at least one challenge that indicates the authentication scheme(s) and  
3273        parameters applicable to the Request-URI. See [H14.47] for a definition of the syntax.

3274        An example of the WWW-Authenticate in a 401 challenge is:

3275        WWW-Authenticate: Digest realm="biloxi.com"



3276 When the originating UAC receives the 401 it SHOULD, if it is able, re-originate the request with the  
3277 proper credentials. The UAC may require input from the originating user before proceeding. Once authenti-  
3278 cation credentials have been supplied (either directly by the user, or discovered in an internal keyring), user  
3279 agents SHOULD cache the credentials for a given value of the To header and "realm" and attempt to re-use  
3280 these values on the next request for that destination.

3281 UAs MAY cache credentials in any way they would like. The following rule MAY be followed when  
3282 caching user credentials:

- 3283 • If a UA receives a WWW-Authenticate in a 401/407 to a request with a particular To header URI, it  
3284 MAY reuse that credential in any subsequent request to the same To header URI.

3285 Any user agent that wishes to authenticate itself with a UAS or registrar – usually, but not necessarily,  
3286 after receiving a 401 response – MAY do so by including an Authorization header field with the request.  
3287 The Authorization field value consists of credentials containing the authentication information of the user  
3288 agent for the realm of the resource being requested.

3289 An example of the Authorization header is:

```
3290 Authorization: Digest username="bob",  
3291                 realm="biloxi.com",  
3292                 nonce="dcd98b7102dd2f0e8b11d0f600bfb0c093",  
3293                 uri=sip:alice@atlanta.com,  
3294                 qop=auth,  
3295                 nc=00000001,  
3296                 cnonce="0a4f113b",  
3297                 response="6629fae49393a05397450978507c4ef1",  
3298                 opaque="5ccc069c403ebaf9f0171e9517f40e41"  
3299
```

3300 When a UAC resubmits a request with its credentials after receiving a 401 (or 407) response, it MUST  
3301 increment the CSeq header field as it would normally do when sending an updated request.

### 3302 20.2.3 Proxy to User Authentication

3303 Similarly, when a UAC sends a request to a proxy server, the proxy server MAY authenticate the originator  
3304 before the request is processed. If no credentials (in the Proxy-Authorization header field) are provided  
3305 in the request, the UAS can challenge the originator to provide credentials by rejecting the request with a  
3306 407 (Proxy Authentication Required) status code. The proxy MUST populate the 407 (Proxy Authentication  
3307 Required) message with a Proxy-Authenticate header applicable to the proxy for the requested resource.

3308 The use of the Proxy-Authentication and Proxy-Authorization parallel that described in [28, Sec-  
3309 tion 3.6], with one difference. Proxies MUST NOT add the Proxy-Authorization header. 407 (Proxy Au-  
3310 thentication Required) responses MUST be forwarded upstream towards the UAC following the procedures  
3311 for any other response. It is the client's responsibility to add the Proxy-Authorization header containing  
3312 credentials for the realm of the proxy which has asked for authentication.

3313 If a proxy were to resubmit a request with a Proxy-Authorization header field, it would need to increment the  
3314 CSeq in the new request. However, this would mean that the UAC which submitted the original request would  
3315 discard a response from the UAS, as the CSeq value would be different.

3316 When the originating UAC receives the 407 it SHOULD, if it is able, re-originate the request with the  
3317 proper credentials. It should follow the same procedures for the display of the “realm” parameter that are  
3318 given above for responding to 401. The UAC SHOULD also cache the credentials used in the re-originated  
3319 request.

3320 The following rule is RECOMMENDED for proxy credential caching:

- 3321 • If a UA receives a Proxy-Authenticate in a 401/407 to a request with a particular Call-ID, it includes  
3322 credentials for that realm in all subsequent requests that contain the same Call-ID. In accordance with  
3323 the definition of a Call-ID, these credentials MUST NOT be cached across dialogs. However, this  
3324 does mean a request could contain credentials that are not needed at any proxy along the path.

3325 Additionally, if a UA is configured with the realm of its local outbound proxy, when one exists, then the  
3326 UA MAY cache credentials for that realm across dialogs.

3327 Any user agent that wishes to authenticate itself to a proxy server – usually, but not necessarily, after  
3328 receiving a 407 response – MAY do so by including an Proxy-Authorization header field with the request.  
3329 The Proxy-Authorization request-header field allows the client to identify itself (or its user) to a proxy  
3330 which requires authentication. The Proxy-Authorization field value consists of credentials containing the  
3331 authentication information of the user agent for the proxy and/or realm of the resource being requested.

3332 A Proxy-Authorization header field applies only to the proxy whose realm is identifier in the “realm”  
3333 parameter (this proxy may previously have demanded authentication using the Proxy-Authenticate field).  
3334 When multiple proxies are used in a chain, the Proxy-Authorization header field MUST NOT be consumed  
3335 by any proxy whose realm does not match the “realm” parameter specified in the Proxy-Authorization  
3336 header.

3337 Note that if an authentication scheme is used in the Proxy- Authorization that does not support realms,  
3338 a proxy server MUST attempt to parse all Proxy-Authorization headers to determine whether or not one  
3339 of them has what it considers to be valid credentials. Because this is potentially very time consuming in  
3340 large networks, proxy servers SHOULD use an authentication scheme that supports realms in the Proxy-  
3341 Authorization header.

3342 If a request is forked (as described in Section 16.6, various proxy servers and/or user agents may wish  
3343 to challenge the UAC. In this case the forking proxy server is responsible for aggregating these challenges  
3344 into a single response. Each WWW-Authenticate and Proxy-Authenticate received in responses to the  
3345 forked request MUST be placed into the single response that is sent by the forking proxy to the user agent;  
3346 the ordering of these headers is not significant.

3347 When a proxy issues a challenge in response to a request, it will not forward the request until the UAC has  
3348 provided valid credentials. A forking proxy may forward a request simultaneously to multiple proxy servers that  
3349 require authentication, each of which in turn will not forward the request until the originating UAC has authenticated  
3350 itself in their respective realm. If the UAC does not provide credentials for each of these challenges, then the proxy  
3351 servers that issued the challenges will not forward requests to user agents where the destination user might be  
3352 located, and therefore, the virtues of forking are largely lost.

3353 If at least one UAS responds to a forked request with a challenge, than a 401 MUST be sent as the  
3354 aggregated response by the forking proxy to the UAC; otherwise, if only proxy servers respond, a 407 MUST  
3355 be used.

3356 When resubmitting its request in response to a 401 or 407 that contains multiple challenges, a UAC MAY  
3357 include an Authorization for each WWW-Authenticate and Proxy-Authorization for each Proxy-Authenticate  
3358 for which the UAC wishes to supply a credential. As noted above, multiple credentials in a request SHOULD  
3359 be differentiated by the “realm” parameter.

3360 Note that it is possible for multiple challenges associated with the same realm to appear in the same 401  
3361 or 407 (for example, when multiple proxies within the same administrative domain, which use a common  
3362 realm, are reached by a forking request).

3363 See [H14.34] for a definition of the syntax of Proxy- Authentication and Proxy-Authorization.

#### 3364 20.2.4 Authentication Schemes

3365 SIP implementations MAY use HTTP's Digest authentication scheme ([28]) to provide a rudimentary form of  
3366 security. This section overviews usage of these mechanisms in SIP. The scheme usage is almost completely  
3367 identical to that for HTTP [28]. This section outlines this operation, pointing to RFC 2617 ([28]) for details  
3368 and noting the differences that arise when using SIP. Since RFC 2543 is based on HTTP Digest as defined in  
3369 RFC 2069 [29], SIP servers supporting RFC 2617 MUST ensure they are backwards compatible with RFC  
3370 2069. Procedures for this backwards compatibility are specified in RFC 2617. Note however that servers  
3371 MUSTNOT accept or request Basic authentication.

3372 **20.2.4.1 HTTP Digest** The rules for Digest authentication follow those defined in [28, Section 3], with  
3373 "HTTP 1.1" replaced by "SIP/2.0" in addition to the following differences:

3374 1. The URI included in the challenge has the following BNF:

3375 URI = SIP-URL

3376 2. The example in Section 3.5 of RFC 2617 has an error in that the 'uri' parameter of the Authorization  
3377 header for Digest authentication is enclosed in quotation marks. Usage in SIP follows the BNF in  
3378 RFC 2617 for Authorization (and by extension Proxy-Authorization) in that the value of the URI  
3379 MUSTNOT be enclosed in quotation marks.

3380 3. The BNF for digest-uri-value is:

3381 digest-uri-value = Request-URI ; as defined in Section 25

3382 4. The example procedure for choosing a nonce based on Etag does not work for SIP.

3383 5. The text in RFC 2617 [28] regarding cache operation does not apply to SIP.

3384 6. RFC 2617 [28] requires that a server check that the URI in the request line, and the URI included in  
3385 the Authorization header, point to the same resource. In a SIP context, these two URI's may actually  
3386 refer to different users, due to forwarding at some proxy. Therefore, in SIP, a server MAY check  
3387 that the Request-URI in the Authorization header corresponds to a user for whom that the server is  
3388 willing to accept forwarded or direct calls.

3389 7. As a clarification to the calculation of the A2 value for message integrity assurance in the Digest  
3390 authentication scheme, implementers should assume, when the entity-body is empty (i.e. when SIP  
3391 messages have no body) that the hash of the entity-body resolves to the MD5 hash of an empty string,  
3392 or:

3393 8. RFC 2617 notes that a cnonce value MUSTNOT be sent in an Authorization (and by extension Proxy-  
3394 Authorization) header if no qop directive as been sent. Therefore, any algorithms that have a depen-  
3395 dency on the cnonce (including "MD5-Sess") require that the qop directive be sent. Use of the "qop"

3396 parameter is optional in RFC 2617 for the purposes of backwards compatibility with RFC 2069; since  
3397 RFC 2543 was based on RFC 2069, the “qop” parameter must unfortunately remain optional for  
3398 clients and servers to receive. However, servers MUST always send a “qop” parameter in WWW-  
3399 Authenticate and Proxy-Authenticate headers. If a client receives a “qop” parameter in a challenge  
3400 header, it MUST send the “qop” parameter in any resulting authorization header.

3401  $H(\text{entity-body}) = \text{MD5}(\text{""}) = \text{"d41d8cd98f00b204e9800998ecf8427e"}$

3402 RFC2543 did not allow usage of the Authentication-Info header (it effectively used RFC 2069). How-  
3403 ever, we now allow usage of this header, since it provides integrity checks over the bodies and provides  
3404 mutual authentication. RFC2617 [28] defines mechanisms for backwards compatibility using the qop at-  
3405 tribute in the request. These mechanisms MUST be used by a server to determine if the client supports the  
3406 new mechanisms in RFC 2617 that were not specified in RFC 2069.

### 3407 20.3 SIP Encryption

3408 No mechanism is currently specified for encrypting entire SIP messages end-to-end for the purpose of con-  
3409 fidentiality. This is a hard problem because network intermediaries (like proxy servers) need to view certain  
3410 headers in order to route messages correctly, and if these intermediaries are excluded from security associa-  
3411 tions then SIP messages will essentially be unroutable.

3412 That much said, SIP messages carry MIME bodies and the MIME standard includes mechanisms for  
3413 securing MIME contents to ensure both integrity and confidentiality (including the 'multipart/encrypted'  
3414 MIME type, see [30]), but detailed description of the use of secure MIME types are outside the scope of this  
3415 document. Implementors should note, however, that there may be rare network intermediaries (not typical  
3416 proxy servers) that rely on viewing or modifying the bodies of SIP messages (especially SDP), and that  
3417 secure MIME may prevent these sorts of intermediaries from functioning.

3418 This applies particularly to certain types of firewalls.

3419 End-to-end encryption relies on keys shared by the two user agents involved in the request. Typically,  
3420 the message is sent encrypted with the public key of the recipient, so that only that recipient can read the  
3421 message. SIP does not define any mechanism for end-to-end key exchange.

3422 Note that the PGP mechanism for encrypting the headers and bodies of SIP messages described in RFC2543 has  
3423 been deprecated.

### 3424 20.4 Denial of Service

3425 Denial of service attacks focus on rendering a particular network element unavailable, usually by directing  
3426 an excessive amount of network traffic at its interfaces. A distributed denial of service attack allows one  
3427 network user to cause multiple network hosts to flood a target host with a large amount of network traffic.

3428 In many architectures SIP proxy servers face the public Internet in order to accept requests from world-  
3429 wide IP endpoints. When the host on which a SIP proxy server is operating is routable from the public  
3430 Internet, it should be deployed in an administrative domain with secure routing policies (blocking source-  
3431 routed traffic, preferably filtering ping traffic).

3432 SIP creates a number of potential opportunities for distributed denial of service attacks that must be  
3433 recognized and addressed by the implementers and operators of SIP systems.

3434 Floods of messages directed at proxy servers can lock up proxy server resources and prevent desirable  
3435 traffic from reaching its destination. There is a computational expense associated with processing a SIP

3436 transaction at a proxy server, and that expense is greater for stateful proxy servers than it is for stateless  
3437 proxy servers. Therefore stateful proxies are more susceptible to flooding than stateless proxy servers.

3438 Attackers can create bogus requests that contain a falsified source IP address and a corresponding *Via*  
3439 header field which identify a targeted host as the originator of the request and then send this request to a large  
3440 number of SIP network elements, thereby using hapless SIP UAs or proxies to generate denial of service  
3441 traffic aimed at the target.

3442 Similarly, attackers might use falsified *Route* headers in a request that identify the target host and then  
3443 send such messages to forking proxies that will amplify messaging sent to the target. *Record-Route* could  
3444 be used to similar effect when the attacker is certain that the SIP dialog initiated by the request will result in  
3445 numerous transactions originating in the backwards direction.

3446 One could prevent one's host from being commandeered for such an attack by disallowing requests that  
3447 do not make use of a persistent security association established through a transport or network layer security  
3448 instrument such as TLS or IPsec. This could be an appropriate security solution for two proxy servers that  
3449 trust one another and exchange significant amounts of signaling traffic with one another, or between a user  
3450 agent and its outbound proxy.

3451 Both TLS and IPsec can also make use of bastion hosts at the edges of administrative domains that  
3452 participate in the security associations to aggregate secure tunnels and sockets. These bastion hosts can also  
3453 take the brunt of denial of service attacks, ensuring that SIP hosts within the administrative domain are not  
3454 encumbered with superfluous messaging.

3455 If such a persistent security association is not feasible, user agents and proxy servers SHOULD chal-  
3456 lenge questionable requests with only a *single* 401 (Unauthorized) or 407 (Proxy Authentication Required)  
3457 - forgoing the normal response retransmission algorithm, and behaving statelessly towards unauthenticated  
3458 requests.

3459 Retransmitting the 401 or 407 status response amplifies the problem of an attacker using a falsified header (such  
3460 as *Via*) to direct traffic to a third party.

3461 A number of denial of service attacks open up if *REGISTER* requests are not properly authenticated  
3462 and authorized by registrars. Attackers could de-register some or all users in an administrative domain,  
3463 thereby preventing these users from being invited to new sessions. An attacker could also register a large  
3464 number of contacts designating the same host for a given address of record in order to use the registrar and  
3465 any associated proxy servers as amplifiers in a denial of service attack. Attackers might also attempt to  
3466 deplete available memory and disk resources of a registrar by registering huge numbers of bindings.

3467 With either TCP or UDP, a denial of service attack exists by a rogue proxy sending 6xx responses.  
3468 Although a client SHOULD choose to ignore such responses if it requested authentication, a proxy cannot do  
3469 so. It is obliged to forward the 6xx response back to the client. The client can then ignore the response, but  
3470 if it repeats the request it will probably reach the same rogue proxy again, and the process will repeat.

3471 The use of multicast to transmit SIP requests can greatly increase the potential for denial of service  
3472 attacks.

## 3473 **21 Common Message Components**

3474 There are certain components of SIP messages that appear in various places within SIP messages (and  
3475 sometimes, outside of them), which merit separate discussion.

## 3476 21.1 SIP Uniform Resource Indicators

3477 A SIP URI identifies a communications resource. Like all URIs, SIP URIs may be placed in web pages,  
3478 email messages or printed literature. They contain sufficient information to initiate and maintain a commu-  
3479 nication session with the resource.

3480 Examples of communications resources include

- 3481 ● a user of an online service
- 3482 ● an appearance on a multiline phone
- 3483 ● a mailbox on a messaging system
- 3484 ● a PSTN phone number at a gateway service
- 3485 ● a group (such as “sales” or “helpdesk”) in an organization

### 3486 21.1.1 SIP URI components

3487 The “sip:” scheme follows the guidelines in RFC 2396 [10]. It uses a form similar to the `mailto` URL, al-  
3488 lowing the specification of SIP request-header fields and the SIP message- body. This makes it possible  
3489 to specify the subject, media type, or urgency of sessions initiated by using a URI on a web page or in an  
3490 email message. The formal syntax for a SIP URI is presented in Section 25. Its general form is

3491 `sip:user:password@host:port:url-parameters?headers`

3492 These tokens, and some of the tokens in their expansion, have the following meanings.

3493 **user:** The identifier of a particular resource at the host being addressed. Note that “host” as used here may,  
3494 and frequently does, refer to a domain.

3495 The “userpart” of a URI consists of this user field, the password field and the @ sign following them.  
3496 The userpart of a URI is optional and MAY be absent when the destination host does not have a notion  
3497 of users or when the host itself is the resource being identified. If the @ sign is present in a SIP URI,  
3498 the user field MUST NOT be empty.

3499 If the host being addressed is capable of processing telephone numbers, an Internet telephony gateway  
3500 for instance, a `telephone- subscriber` field defined in RFC 2806 [14] MAY be used to populate the  
3501 user field. There are special escaping rules for encoding `telephone-subscriber` fields in SIP URIs  
3502 described in Section 21.1.2.

3503 **password:** A password associated with the user

3504 While the SIP URI syntax allows this field to be present, its use is NOT RECOMMENDED, because  
3505 the passing of authentication information in clear text (such as URIs) has proven to be a security risk  
3506 in almost every case where it has been used. For instance, transporting a PIN number in this field  
3507 exposes the PIN.

3508 **host:** The entity hosting the SIP resource

3509 The `host` part contains either a fully-qualified domain name or numeric IPv4 or IPv6 address. Using  
3510 the fully-qualified domain name form is RECOMMENDED whenever possible.

3511 **port:** The port number where the request is to be sent.

3512 **URI parameters:** Parameters affecting a request constructed from the URI.

3513 URI parameters are added after the `hostport` component and are separated by semi-colons. This  
3514 extensible mechanism includes the `transport`, `maddr`, `ttl`, `user`, and `method` parameters.

3515 The `transport` parameter determines the transport mechanism to be used for sending SIP messages.  
3516 SIP can use any network transport protocol. Parameter names are defined for UDP [31], TCP [32],  
3517 TLS [26], and SCTP [33].

3518 The `maddr` parameter indicates the server address to be contacted for this user, overriding any address  
3519 derived from the `host` field. [8] describes the proper interpretation of the `transport`, `maddr` and  
3520 `hostport` in order to obtain the destination address, port and transport for sending a request.

3521 The `maddr` field can be used as a simple form of loose source routing. It allows a URI to specify a specific  
3522 proxy that must be traversed en-route to the destination. This capability is useful for a roaming user that is  
3523 forced to use an outbound proxy, but wishes to force requests through their home proxy.

3524 The `ttl` parameter determines the time-to-live value of the UDP multicast packet and **MUST** only  
3525 be used if `maddr` is a multicast address and the transport protocol is UDP. The `user` parameter  
3526 was described above. For example, to specify to call `alice@atlanta.com` using multicast to  
3527 `239.255.255.1` with a `ttl` of 15, the following URI would be used:

```
3528 sip:alice@atlanta.com;maddr=239.255.255.1;ttl=15
```

3529 The set of valid `telephone-subscriber` strings is a subset of valid `user` strings. The `user` URI pa-  
3530 rameter exists to distinguish telephone numbers from user names that happen to look like telephone  
3531 numbers. If the user string contains a telephone number formatted as a `telephone-subscriber`, the  
3532 `user` parameter value “`phone`” **SHOULD** be present. Even without this parameter, recipients of SIP  
3533 URIs **MAY** interpret the pre-@ part as a telephone number if local restrictions on the name space for  
3534 user name allow it.

3535 The method of the SIP request constructed from the URI can be specified with the `method` parameter.  
3536 Since the url-parameter mechanism is extensible, SIP elements **MUST** silently ignore any url-parameters  
3537 that they do not understand.

3538 **Headers:** Headers to be included in a request constructed from the URI.

3539 Headers fields in the SIP request can be specified with the “?” mechanism within a SIP URI. The  
3540 header names and values are encoded in ampersand separated `hname = hvalue` pairs. The special  
3541 `hname` “`body`” indicates that the associated `hvalue` is the `message-body` of the SIP request.

3542 Table 1 summarizes the use of SIP URI components based on the context in which the URI appears. The  
3543 external column describes URIs appearing anywhere outside of a SIP message, for instance on a web page  
3544 or business card. Entries marked “`m`” are mandatory, those marked “`o`” are optional, and those marked “`-`”  
3545 are not allowed. Elements processing URIs **SHOULD** ignore any disallowed components if they are present.  
3546 The second column indicates the default value of an optional element if it is not present. “`-`” indicates that  
3547 the element is either not optional, or has no default value.

3548 SIP URIs in Contact header fields have different restrictions depending on the context in which the  
 3549 header field appears. One set applies to messages that establish and maintain dialogs (INVITE and its 200  
 3550 OK response). The other applies to registration and redirection messages (REGISTER, its 200 OK response,  
 3551 and 3xx class responses to any method).

3552 OPEN ISSUE #203: maddr is disallowed in To/From, but not port. Should port be disallowed?

3553 OPEN ISSUE #204: Password is disallowed in From, but not To. Why?

3554 OPEN ISSUE #205: Should we allow method and header URI components in registration/redirect Con-  
 3555 tacts. What do they mean?

	default	Req.-URI	To	From	reg./redir. Contact	dialog Contact/ R-R/Route	external
user	–	o	o	o	o	o	o
password	–	o	o	-	o	o	o
host	–	m	m	m	m	m	m
port	5060	o	o	o	o	o	o
user-param	ip	o	o	o	o	o	o
method	INVITE	-	-	-	o	-	o
maddr-param	–	o	-	-	o	o	o
ttl-param	1	o	-	-	o	-	o
transp.-param	udp	o	-	-	o	o	o
other-param	–	o	o	o	o	o	o
headers	–	-	-	-	o	-	o

Table 1: Use and default values of URI components for SIP headers, Request-URI and references

### 3556 21.1.2 Character escaping requirements

3557 SIP follows the requirements and guidelines of RFC 2396 when defining the set of characters that must be  
 3558 escaped in a SIP URI, and uses its “”%” HEX HEX” mechanism for escaping. From RFC 2396:

3559 The set of characters actually reserved within any given URI component is defined by that com-  
 3560 ponent. In general, a character is reserved if the semantics of the URI changes if the character  
 3561 is replaced with its escaped US-ASCII encoding. [10].

3562 Excluded US-ASCII characters [10, Sec. 2.4.3], such as space and control characters and characters used as  
 3563 URI delimiters, also MUST be escaped. URIs MUST NOT contain unescaped space and control characters.

3564 For each component, the set of valid BNF expansions defines exactly which characters may appear  
 3565 unescaped. All other characters MUST be escaped.

3566 For example, “@” is not in the set of characters in the user component, so the user “j@s0n” must have  
 3567 at least the @ sign encoded, as in “j%40s0n”.

3568 Expanding the hname and hvalue tokens in Section 25 show that all URI reserved characters in header  
 3569 names and values MUST be escaped.

3570 The telephone-subscriber subset of the user component has special escaping considerations. The set  
 3571 of characters not reserved in the RFC 2806 [14] description of telephone-subscriber contains a number



3572 of characters in various syntax elements that need to be escaped when used in SIP URIs. Any characters  
3573 occurring in a `telephone-subscriber` that do not appear in an expansion of the BNF for the `user` rule MUST  
3574 be escaped.

3575 Note that character escaping is not allowed in the host component of a SIP URI (the % character is not  
3576 valid in its expansion). This is likely to change in the future as requirements for Internationalized Domain  
3577 Names are finalized. Current implementations MUST NOT attempt to improve robustness by treating received  
3578 escaped characters in the host component as literally equivalent to their unescaped counterpart. The behavior  
3579 required to meet the requirements of IDN may be significantly different.

### 3580 21.1.3 Example SIP URIs

```
3581 sip:alice@atlanta.com
3582 sip:alice:secretword@atlanta.com;transport=tcp
3583 sip:alice@atlanta.com?subject=project%20x&priority=urgent
3584 sip:+1-212-555-1212:1234@gateway.com;user=phone
3585 sip:1212@gateway.com
3586 sip:alice@10.1.1.1
3587 sip:atlanta.com;method=REGISTER?to=alice%40atlanta.com
3588 sip:alice;day=tuesday@atlanta.com
```

3589 The last example URI above has a `user` field value of “alice;day=tuesday”. The escaping rules defined  
3590 above allow a semicolon to appear unescaped in this field. Note, however, that for the purposes of this  
3591 protocol, the field is opaque. The apparent structure in that value is only useful to the entity responsible for  
3592 the resource.

### 3593 21.1.4 SIP URI Comparison

3594 SIP URIs are compared for equality according to the following rules:

- 3595 • Comparisons of scheme name (“sip”), domain names, parameter names and header names are case-  
3596 insensitive, all other comparisons are case-sensitive. (OPEN ISSUE #100 : There is a proposal to  
3597 make only quoted string comparisons case-sensitive.)
- 3598 • The ordering of parameters and headers is not significant in comparing SIP URIs.
- 3599 • Characters other than those in the “reserved” and “unsafe” sets (see RFC 2396 [10]) are equivalent to  
3600 their “”%” HEX HEX” encoding.
- 3601 • An IP address that is the result of a DNS lookup of a host name does **not** match that host name.
- 3602 • For two URIs to be equal, the `user`, `password`, `host`, and `port` components must match. A URI  
3603 omitting the optional port component will match a URI explicitly declaring port 5060. A URI omitting  
3604 the user component will **not** match a URI that includes one. A URI omitting the password component  
3605 will **not** match a URI that includes one.
- 3606 • URI uri-parameter components are compared as follows  
3607 – Any uri-parameter appearing in both URIs must match.



### 3638 **21.1.5 Forming Requests from a SIP URI**

3639 An implementation must take care when forming requests directly from a URI. URIs from business cards,  
3640 web pages, and even from sources inside the protocol such as registered contacts may contain inappropriate  
3641 header fields or body parts.

3642 The policies to apply during message formation are an implementation decision. An implementation  
3643 SHOULD treat the presence of any headers or body parts in the URI as a request to include them in the  
3644 message, and choose to honor the request on an per-component basis.

3645 An implementation SHOULD NOT honor these obviously dangerous header fields: From, Call-ID, CSeq,  
3646 Via, and Record-Route.

3647 An implementation SHOULD take special care in honoring any requested Route header field values in  
3648 order to not be used as an unwitting agent in malicious attacks.

3649 An implementation SHOULD NOT honor requests to include headers that may cause it to falsely advertise  
3650 its location or capabilities. These include: Accept, Accept-Encoding, Accept-Language, Allow, Contact  
3651 (in its dialog usage), Organization, Supported, and User-Agent.

3652 An implementation SHOULD verify the accuracy of any requested descriptive headers, including: Content-  
3653 Disposition, Content-Encoding, Content-Language, Content-Length, Content-Type, Date, Mime-  
3654 Version, and Timestamp.

### 3655 **21.2 Option Tags**

3656 Option tags are unique identifiers used to designate new options (extensions) in SIP. These tags are used in  
3657 Require (Section 22.31), Proxy-Require (Section 22.28, Supported (Section 22.37) and Unsupported  
3658 (Section 22.40) header fields. Note that these options appear as parameters in those headers in an option-tag  
3659 = token form (see Section 25 for the definition of token).

3660 The creator of a new SIP option MUST either prefix the option with their reverse domain name or register  
3661 the new option with the Internet Assigned Numbers Authority (IANA) (See Section 26).

3662 An example of a reverse-domain-name option is "com.foo.mynewfeature", whose inventor can be reached  
3663 at "foo.com". For these features, individual organizations are responsible for ensuring that option names do  
3664 not collide within the same domain. The host name part of the option MUST use lower-case; the option name  
3665 is case-sensitive.

3666 Options registered with IANA do not contain periods and are globally unique. IANA option tags are  
3667 case-sensitive.

### 3668 **21.3 Tags**

3669 The "tag" parameter is used in the To and From fields of SIP messages. It serves as a general mechanism  
3670 to identify a particular instance of a user agent for a particular SIP URI.

3671 As proxies can fork requests, the same request can reach multiple instances of a user (mobile and home  
3672 phones, for example). Since each can respond, there needs to be a means for the originator of a session to  
3673 distinguish the responses. Tag fields in the To and From disambiguate these multiple instances of the same  
3674 user.

3675 This situation also arises with multicast requests.

3676 When a tag is generated by a UA for insertion into a request or response, it MUST be globally unique and  
3677 cryptographically random with at least 32 bits of randomness. A property of this selection requirement is  
3678 that a UA will place a different tag into the From header of an INVITE as it would place into the To header

3679 of the response to the same INVITE. This is needed in order for a UA to invite itself to a session, a common  
3680 case for “hairpinning” of calls in PSTN gateways.

3681 Besides the requirement for global uniqueness, the algorithm for generating a tag is implementation  
3682 specific. Tags are helpful in fault tolerant systems, where a dialog is to be recovered on an alternate server  
3683 after a failure. A UAS can select the tag in such a way that a backup can recognize a request as part of a  
3684 dialog on the failed server, and therefore determine that it should attempt to recover the dialog and any other  
3685 state associated with it.

## 3686 22 Header Fields

3687 The general syntax for header fields is covered in Section 7.3. This section lists the full set of header fields  
3688 along with notes on syntax, meaning, and usage. Throughout this section, we use [HX.Y] to refer to Section  
3689 X.Y of the current HTTP/1.1 specification RFC 2616 [9]. Examples of each header field are given.

3690 Information about header fields in relation to methods and proxy processing is summarized in Ta-  
3691 bles 2 and 3.

3692 The “where” column describes the request and response types in which the header field can be used.  
3693 Values in this column are:

3694 **R:** refers to header fields that can be used in requests.

3695 **r:** designates a header field as applicable to all responses, while a list of numeric values indicates the status  
3696 codes with which the header field can be used.

3697 **c:** indicates a header field is copied from the request to the response.

3698 The “proxy” column describes the operations a proxy may perform on a header.

3699 **c:** indicates that a proxy can add (concatenate) comma-separated elements to the header

3700 **m:** indicates that a proxy can modify the header

3701 **a:** indicates that a proxy can add the header if not present

3702 **r:** indicates that a proxy must be able to read the header. Headers that need to be read cannot be en-  
3703 crypted.

3704 The next six columns relate to the presence of a header field in a method, with the contents indicating:

3705 **o:** for optional

3706 **m:** for mandatory

3707 **m\*:** indicates a header that SHOULD be sent, but servers need to be prepared to receive messages without  
3708 that header field.

3709 **\*:** indicates that the header fields are required if the message body is not empty. See sections 22.14, 22.15  
3710 and 7.4 for details.

3711 **-:** for not applicable.

3712 **c:** for conditional. The header field is either mandatory or optional, depending on the presence of a route  
3713 set or the response code.

3714 “Optional” means that a UA MAY include the header field in a request or response, and a UA MAY ignore  
3715 the header field if present in the request or response (The exception to this rule is the **Require** header field  
3716 discussed in 22.31). A “mandatory” header field MUST be present in a request, and MUST be understood  
3717 by the UAS receiving the request. A mandatory response header field MUST be present in the response,  
3718 and the header field MUST be understood by the UAC processing the response. “Not applicable” means that  
3719 the header field MUST NOT be present in a request. If one is placed in a request by mistake, it MUST be  
3720 ignored by the UAS receiving the request. Similarly, a header field labeled “not applicable” for a response  
3721 means that the UAS MUST NOT place the header in the response, and the UAC MUST ignore the header in  
3722 the response.

3723 A compact form of some common header fields is also defined for use when overall message size is an  
3724 issue.

3725 The **Contact**, **From**, and **To** header fields contain a URI. If the URI contains a comma, question mark  
3726 or semicolon, the URI MUST be enclosed in angle brackets (< and >). Any URI parameters are contained  
3727 within these brackets. If the URI is not enclosed in angle brackets, any semicolon-delimited parameters are  
3728 header-parameters, not URI parameters.

## 3729 **22.1 Accept**

3730 The **Accept** header follows the syntax defined in [H14.1]. The semantics are also identical, with the excep-  
3731 tion that if no **Accept** header is present, the server SHOULD assume a default value of `application/sdp`.

3732 Example:

3733 `Accept: application/sdp;level=1, application/x-private, text/html`

## 3734 **22.2 Accept-Encoding**

3735 The **Accept-Encoding** header field is similar to **Accept**, but restricts the content-codings [H3.5] that are  
3736 acceptable in the response. See [H14.3]. The syntax of this header is defined in [H14.3]. The semantics in  
3737 SIP are identical to those defined in [H14.3].

3738 An empty **Accept-Encoding** header field is permissible, even though the syntax in [H14.3] does not  
3739 provide for it. It is equivalent to **Accept-Encoding: identity**, that is, only the identity encoding, meaning  
3740 no encoding, is permissible. If this header is not present, the default value is **identity**. This differs slightly  
3741 from the HTTP definition, which indicates that when not present, any encoding can be used, but the identity  
3742 encoding is preferred.

3743 Example:

3744 `Accept-Encoding: gzip`

## 3745 **22.3 Accept-Language**

3746 The **Accept-Language** header follows the syntax defined in [H14.4]. The rules for ordering the languages  
3747 based on the “q” parameter apply to SIP as well.

Header field	where	proxy	ACK	BYE	CAN	INV	OPT	REG	PRA
Accept	R		-	o	-	m*	o	o	o
Accept	2xx		-	-	-	m*	o	o	-
Accept	415		-	o	-	o	o	o	o
Accept-Encoding	R		-	o	-	m*	o	o	o
Accept-Encoding	2xx		-	-	-	m*	o	o	-
Accept-Encoding	415		-	o	-	o	o	o	o
Accept-Language	R		-	o	-	m*	o	o	o
Accept-Language	2xx		-	-	-	m*	o	o	-
Accept-Language	415		-	o	-	o	o	o	o
Alert-Info	R	am	-	-	-	o	-	-	-
Alert-Info	180	am	-	-	-	o	-	-	-
Allow	R		o	o	o	o	o	o	o
Allow	2xx		-	o	o	m*	m*	o	o
Allow	r		-	o	o	o	o	o	o
Allow	405		-	m	m	m	m	m	m
Authentication-Info	2xx		-	o	-	o	o	o	o
Authorization	R		o	o	o	o	o	o	o
Call-ID	c	r	m	m	m	m	m	m	m
Call-Info		am	-	-	-	o	o	o	-
Contact	R		o	-	-	m	o	o	-
Contact	1xx		-	-	-	o	o	-	-
Contact	2xx		-	-	-	m	o	o	-
Contact	3xx		-	o	-	o	o	o	o
Contact	485		-	o	-	o	o	o	o
Content-Disposition			o	o	-	o	o	o	o
Content-Encoding			o	o	-	o	o	o	o
Content-Language			o	o	-	o	o	o	o
Content-Length		r	m*	m*	m*	m*	m*	m*	m*
Content-Type			*	*	-	*	*	*	*
CSeq	c	r	m	m	m	m	m	m	m
Date		a	o	o	o	o	o	o	o
Error-Info	300-699		-	o	o	o	o	o	o
Expires			-	-	-	o	-	o	-
From	c	r	m	m	m	m	m	m	m
In-Reply-To	R		-	-	-	o	-	-	-
Max-Forwards	R	rm	o	o	o	o	o	o	o
MIME-Version			o	o	o	o	o	o	o
Organization		am	-	-	-	o	o	o	-

Table 2: Summary of header fields, A–O

Header field	where	proxy	ACK	BYE	CAN	INV	OPT	REG	PRA
Priority	R	a	-	-	-	o	-	-	-
Proxy-Authenticate	407		-	m	m	m	m	m	m
Proxy-Authorization	R	r	o	o	o	o	o	o	o
Proxy-Require	R	r	o	o	o	o	o	o	o
RAck	R		-	-	-	-	-	-	m
Record-Route	R	amr	o	o	o	o	o	-	o
Record-Route	2xx,401,484		-	o	o	o	o	-	o
Require	g	acr	o	o	o	o	o	o	o
Retry-After	404,413,480,486		-	o	o	o	o	o	o
	500,503		-	o	o	o	o	o	o
	600,603		-	o	o	o	o	o	o
Route	R	r	c	c	c	c	c	-	c
RSeq	1xx		-	o	-	o	o	o	-
Server	r		-	o	o	o	o	o	o
Subject	R		-	-	-	o	-	-	-
Supported			-	o	o	o	o	o	o
Timestamp			o	o	o	o	o	o	o
To	gc(1)	r	m	m	m	m	m	m	m
Unsupported	420		-	o	o	o	o	o	o
User-Agent			o	o	o	o	o	o	o
Via	c	acmr	m	m	m	m	m	m	m
Warning	r		o	o	o	o	o	o	o
WWW-Authenticate	401		-	m	m	m	m	m	m

Table 3: Summary of header fields, P-Z; (1): copied with possible addition of tag

3748 The **Accept-Language** header is used in requests to indicate the preferred languages for reason phrases,  
 3749 session descriptions, or status responses carried as message bodies in the response. If no **Accept-Language**  
 3750 header field is present in a request, the server assumes all languages are acceptable to the client.

3751 Example:

3752 `Accept-Language: da, en-gb;q=0.8, en;q=0.7`

## 3753 22.4 Alert-Info

3754 When present in an **INVITE** request, the **Alert-Info** header field specifies an alternative ring tone to the UAS.  
 3755 When present in a 180 (Ringing) response, the **Alert-Info** header field specifies an alternative ringback tone  
 3756 to the UAC. A typical usage is for a proxy to insert this header to provide a distinctive ring feature.

3757 The **Alert-Info** header can introduce security risks. These risks and the ways to handle them are dis-  
 3758 cussed in Section 22.9, which discusses the **Call-Info** header since the risks are identical.

3759 In addition, a user **SHOULD** be able to disable this feature selectively.

3760 This helps prevent disruptions that could result from the use of this header by untrusted elements.

3761 Example:

3762 Alert-Info: <http://www.example.com/sounds/moo.wav>

## 3763 22.5 Allow

3764 The Allow header field lists the set of methods supported by the UA generating the message.

3765 All methods, including ACK and CANCEL, understood by the UA MUST be included in the list of  
3766 methods in the Allow header, when present. The absence of an Allow header MUST NOT be interpreted to  
3767 mean that the UA sending the message supports no methods. Rather, it implies that the UA is not providing  
3768 any information on what methods it supports.

3769 Supplying an Allow header in responses to methods other than OPTIONS reduces the number of mes-  
3770 sages needed.

3771 Example:

3772 Allow: INVITE, ACK, OPTIONS, CANCEL, BYE

## 3773 22.6 Authentication-Info

3774 The Authentication-Info header provides for mutual authentication with HTTP Digest. A UAS MAY include  
3775 this header in a 2xx response to a request that was successfully authenticated using digest based on the  
3776 Authorization header.

3777 Syntax and semantics follow those specified in RFC2617 [28].

3778 Example:

3779 Authentication-Info: nextnonce="47364c23432d2e131a5fb210812c"

## 3780 22.7 Authorization

3781 The Authorization header field contains authentication credentials of a UA. Section 20.2.2 overviews the  
3782 use of the Authorization header field, and Section 20.2.4 describes the syntax and semantics when used  
3783 with HTTP Basic and Digest authentication.

3784 Note that this header field, along with Proxy-Authorization, breaks the general rules about multiple  
3785 header fields. Although not a comma-separated list, this header field may be present multiple times, and  
3786 MUST NOT be combined into a single header using the usual rules described in Section 7.3.

3787 Example:

3788 Authorization: Digest username="Alice", realm="Bob's Friends",  
3789 nonce="84a4cc6f3082121f32b42a2187831a9e",  
3790 response="7587245234b3434cc3412213e5f113a5432"

## 3791 22.8 Call-ID

3792 The Call-ID header field uniquely identifies a particular invitation or all registrations of a particular client.  
3793 Note that a single multimedia conference can give rise to several calls with different Call-IDs, for example,  
3794 if a user invites a single individual several times to the same (long-running) conference. Call-IDs are case-  
3795 sensitive and are simply compared byte-by-byte.

3796 The compact form of the Call-ID header field is i.



3797 Examples:

3798 Call-ID: f81d4fae-7dec-11d0-a765-00a0c91e6bf6@biloxi.com  
3799 i:f81d4fae-7dec-11d0-a765-00a0c91e6bf6@10.4.1.4

## 3800 22.9 Call-Info

3801 The **Call-Info** header field provides additional information about the caller or callee, depending on whether  
3802 it is found in a request or response. The purpose of the URI is described by the “purpose” parameter.  
3803 The “icon” parameter designates an image suitable as an iconic representation of the caller or callee. The  
3804 “info” parameter describes the caller or callee in general, for example, through a web page. The “card”  
3805 parameter provides a business card, for example, in vCard [34] or LDIF [35] formats. Additional tokens can  
3806 be registered using IANA and the procedures in Section 26.

3807 Use of the **Call-Info** header field can pose a security risk. If a callee fetches the URIs provided by a  
3808 malicious caller, the callee may be at risk for displaying inappropriate or offensive content, dangerous or  
3809 illegal content, and so on. Therefore, it is RECOMMENDED that a UA only render the information in the  
3810 **Call-Info** header if it can verify the authenticity of the element that originated the header and trusts that  
3811 element. This need not be the peer UA; a proxy can insert this header into requests.

3812 The use of this header is important in converged applications.

3813 Example:

3814 Call-Info: <http://www.example.com/alice/photo.jpg> ;purpose=icon,  
3815 <http://www.example.com/alice/> ;purpose=info

## 3816 22.10 Contact

3817 The **Contact** header field provides a URI whose meaning depends on the the type of request or response it  
3818 is in.

3819 Parameters defined for **Contact** include “q” and “expires”. Additional parameters may be defined in  
3820 other specifications. Even if the “display-name” is empty, the “name-addr” form MUST be used if the  
3821 “addr-spec” contains a comma, semicolon, or question mark. Note that there may or may not be LWS  
3822 between the display-name and the “<”.

3823 The **Contact** header field fulfills functionality similar to the **Location** header field in HTTP. However, the  
3824 HTTP header field only allows one address, unquoted. Since URIs can contain commas and semicolons as reserved  
3825 characters, they can be mistaken for header or parameter delimiters, respectively. The current syntax corresponds to  
3826 that for the **To** and **From** header fields, which also allow the use of display names.

3827 The compact form of the **Contact** header field is **m** (for “moved”).

3828 Examples:

3829 Contact: "Mr. Watson" <sip:watson@worchester.bell-telephone.com>  
3830 ;q=0.7; expires=3600,  
3831 "Mr. Watson" <mailto:watson@bell-telephone.com> ;q=0.1  
3832 m: <sip:bob@10.5.1.5>

## 3833 22.11 Content-Disposition

3834 The Content-Disposition header field describes how the message body or, for multipart messages, a mes-  
3835 sage body part is to be interpreted by the UAC or UAS. This SIP header field extends the MIME Content-  
3836 Type (RFC 1806 [36]).

3837 The value “**session**” indicates that the body part describes a session, for either calls or early (pre-call)  
3838 media. The value “**render**” indicates that the body part should be displayed or otherwise rendered to the  
3839 user. For backward-compatibility, if the Content-Disposition header is missing, bodies of Content-Type  
3840 `application/sdp` imply the disposition “**session**”, while other content types imply “**render**”.

3841 The disposition type “**icon**” indicates that the body part contains an image suitable as an iconic repre-  
3842 sentation of the caller or callee. The value “**alert**” indicates that the body part contains information, such as  
3843 an audio clip, that should be rendered instead of ring tone.

3844 The handling parameter, `handling-param`, describes how the UAS should react if it receives a message  
3845 body whose content type or disposition type it does not understand. The parameter has defined values of  
3846 “**optional**” and “**required**”. If the handling parameter is missing, the value “**required**” is to be assumed.  
3847 If this header field is missing, the MIME type determines the default content disposition. If there is none,  
3848 “**render**” is assumed.

3849 Example:

```
3850 Content-Disposition: session
```

## 3851 22.12 Content-Encoding

3852 The Content-Encoding header field is used as a modifier to the “**media-type**”. When present, its value  
3853 indicates what additional content codings have been applied to the entity-body, and thus what decoding  
3854 mechanisms **MUST** be applied in order to obtain the media-type referenced by the Content-Type header  
3855 field. Content-Encoding is primarily used to allow a body to be compressed without losing the identity of  
3856 its underlying media type.

3857 If multiple encodings have been applied to an entity, the content codings **MUST** be listed in the order in  
3858 which they were applied.

3859 All content-coding values are case-insensitive. IANA acts as a registry for content-coding value tokens.  
3860 See [H3.5] for a definition of the syntax for content-coding.

3861 Clients **MAY** apply content encodings to the body in requests. A server **MAY** apply content encodings to  
3862 the bodies in responses. The server **MUST** only use encodings listed in the Accept-Encoding header in the  
3863 request.

3864 The compact form of the Content-Encoding header field is **e**.

3865 Examples:

```
3866 Content-Encoding: gzip  
3867 e: tar
```

## 3868 22.13 Content-Language

3869 See [H14.12].

3870 Example:

3871 Content-Language: fr

## 3872 22.14 Content-Length

3873 The Content-Length header field indicates the size of the message-body, in decimal number of octets, sent  
3874 to the recipient.

3875 Applications SHOULD use this field to indicate the size of the message-body to be transferred, regardless  
3876 of the media type of the entity. The size of the message-body does *not* include the CRLF separating headers  
3877 and body. Any Content-Length greater than or equal to zero is a valid value. If no body is present in a  
3878 message, then the Content-Length header field MUST be set to zero.

3879 The ability to omit Content-Length simplifies the creation of cgi-like scripts that dynamically generate re-  
3880 sponses.

3881 The compact form of the header is l.

3882 Examples:

3883 Content-Length: 349

3884 l: 173

## 3885 22.15 Content-Type

3886 The Content-Type header field indicates the media type of the message-body sent to the recipient. The  
3887 “media-type” element is defined in [H3.7]. The Content-Type header MUST be present if the body is not  
3888 empty. If the body is empty, and a Content-Type header is present, it indicates that the body of the specific  
3889 type has zero length (for example, an empty audio file).

3890 The compact form of the header is c.

3891 Examples:

3892 Content-Type: application/sdp

3893 c: text/html; charset=ISO-8859-4

## 3894 22.16 CSeq

3895 A CSeq header field in a request contains a single decimal sequence number and the request method. The  
3896 sequence number MUST be expressible as a 32-bit unsigned integer. The CSeq header serves to order trans-  
3897 actions within a dialog, to provide a means to uniquely identify transactions, and to differentiate between  
3898 new requests and request retransmissions.

3899 Example:

3900 CSeq: 4711 INVITE

## 3901 22.17 Date

3902 The Date header field contains an RFC 1123 date (see [H14.18]). Note that unlike HTTP/1.1, SIP only  
3903 supports the most recent RFC 1123 [37] formatting for dates. As in [H3.3], SIP restricts the timezone in  
3904 SIP-date to “GMT”, while RFC 1123 allows any timezone.

3905 The consistent use of GMT between Date, Expires and Retry-After headers allows implementation of simple  
3906 clients that do not have a notion of absolute time.

3907 Note that `rfc1123-date` is case-sensitive.

3908 The `Date` header field reflects the time when the request or response is first sent.

3909 The `Date` header field can be used by simple end systems without a battery-backed clock to acquire a notion of  
3910 current time. However, in its GMT form, it requires clients to know their offset from GMT.

3911 Example:

3912 `Date: Sat, 13 Nov 2010 23:29:00 GMT`

## 3913 22.18 Error-Info

3914 The `Error-Info` header field provides a pointer to additional information about the error status response.

3915 SIP UACs have user interface capabilities ranging from pop-up windows and audio on PC softclients to audio-  
3916 only on "black" phones or endpoints connected via gateways. Rather than forcing a server generating an error to  
3917 choose between sending an error status code with a detailed reason phrase and playing an audio recording, the  
3918 `Error-Info` header field allows both to be sent. The UAC then has the choice of which error indicator to render to the  
3919 caller.

3920 A UAC MAY treat a SIP URI in an `Error-Info` header field as if it were a `Contact` in a redirect and  
3921 generate a new `INVITE`, resulting in a recorded announcement session being established. A non-SIP URI  
3922 MAY be rendered to the user.

3923 Examples:

3924 `SIP/2.0 404 The number you have dialed is not in service`  
3925 `Error-Info: <sip:not-in-service-recording@atlanta.com>`

## 3926 22.19 Expires

3927 The `Expires` header field gives the date and time after which the message (or content) expires. The precise  
3928 meaning of this is method dependent.

3929 Note that the expiration time in an `INVITE` does *not* affect the duration of the actual session that may  
3930 result from the invitation. Session description protocols may offer the ability to express time limits on the  
3931 session duration, however.

3932 The value of this field can be either a date (see the `Date` header field) or an integer number of seconds  
3933 (in decimal), measured from the receipt of the request. The latter approach is preferable for short durations,  
3934 as it does not depend on clients and servers sharing a synchronized clock.

3935 Examples:

3936 `Expires: Thu, 01 Dec 1994 16:00:00 GMT`  
3937 `Expires: 5`

## 3938 22.20 From

3939 The `From` header field indicates the initiator of the request. Note that this may be different from the initiator  
3940 of the dialog. Requests sent by the callee to the caller use the callee's address in the `From` header field.

3941 The optional "`display-name`" is meant to be rendered by a human user interface. A system SHOULD  
3942 use the display name "Anonymous" if the identity of the client is to remain hidden.

3943 Even if the “display-name” is empty, the “name-addr” form MUST be used if the “addr-spec” con-  
3944 tains a comma, question mark, or semicolon. Syntax issues are discussed in Section 7.3.1.

3945 The compact form of the header is f.

3946 Examples:

3947 From: "A. G. Bell" <sip:agb@bell-telephone.com> ;tag=a48s

3948 From: sip:+12125551212@server.phone2net.com;tag=887s

3949 f: Anonymous <sip:c8oqz84zk7z@privacy.org>;tag=hyh8

## 3950 22.21 In-Reply-To

3951 The In-Reply-To header field enumerates the Call-IDs that this call references or returns. These Call-IDs  
3952 may have been cached by the client then included in this header in a return call.

3953 This allows automatic call distribution systems to route return calls to the originator of the first call. This also  
3954 allows callees to filter calls, so that only return calls for calls they originated will be accepted. This field is not a  
3955 substitute for request authentication.

3956 Example:

3957 In-Reply-To: 70710@saturn.bell-tel.com, 17320@saturn.bell-tel.com

## 3958 22.22 Max-Forwards

3959 The Max-Forwards header field may be used with any SIP method to limit the number of proxies or gate-  
3960 ways that can forward the request to the next downstream server. This can also be useful when the client is  
3961 attempting to trace a request chain that appears to be failing or looping in mid-chain.

3962 The Max-Forwards value is a decimal integer indicating the remaining number of times this request  
3963 message is allowed to be forwarded. This count is decremented by each server that forwards the request.

3964 Example:

3965 Max-Forwards: 6

## 3966 22.23 MIME-Version

3967 See [H19.4.1].

3968 Example:

3969 MIME-Version: 1.0

## 3970 22.24 Organization

3971 The Organization header field conveys the name of the organization to which the entity issuing the request  
3972 or response belongs.

3973 The field MAY be used by client software to filter calls.

3974 Example:

3975 Organization: Boxes by Bob

### 3976 **22.25 Priority**

3977 The **Priority** header field indicates the urgency of the request as perceived by the client. Defined values  
3978 include "non-urgent", "normal", "urgent", and "emergency".

3979 It is RECOMMENDED that the value of "emergency" only be used when life, limb, or property are in  
3980 imminent danger. Otherwise, there are no semantics defined for this header field.

3981 These are the values of RFC 2076 [38], with the addition of "emergency".

3982 Examples:

3983 Subject: A tornado is heading our way!

3984 Priority: emergency

3985 or

3986 Subject: Weekend plans

3987 Priority: non-urgent

### 3988 **22.26 Proxy-Authenticate**

3989 The **Proxy-Authenticate** header field consists of a challenge that indicates the authentication scheme and  
3990 parameters applicable to the proxy for this **Request-URI**.

3991 The syntax for this header and its use is defined in [H14.33]. See 20.2.3 for further details on its usage.

3992 Example:

3993 Proxy-Authenticate: Digest realm="Carrier SIP",

3994 domain="sip:ssl.carrier.com",

3995 nonce="f84f1cec41e6cbe5aea9c8e88d359",

3996 opaque="", stale=FALSE, algorithm=MD5

### 3997 **22.27 Proxy-Authorization**

3998 The **Proxy-Authorization** header field allows the client to identify itself (or its user) to a proxy that requires  
3999 authentication. The **Proxy-Authorization** field value consists of credentials containing the authentication  
4000 information of the user agent for the proxy and/or realm of the resource being requested.

4001 See [H14.34] for a definition of the syntax, and section 20.2.3 for a discussion of its usage.

4002 Note that this header field, along with **Authorization**, breaks the general rules about multiple header  
4003 fields. Although not a comma-separated list, this header field may be present multiple times, and **MUST NOT**  
4004 be combined into a single header using the usual rules described in Section 7.3.1.

4005 Example:

4006 Proxy-Authorization: Digest username="Alice", realm="Atlanta ISP",

4007 nonce="c60f3082ee1212b402a21831ae",

4008 response="245f23415f11432b3434341c022"

## 4009 **22.28 Proxy-Require**

4010 The **Proxy-Require** header field is used to indicate proxy-sensitive features that must be supported by the  
4011 proxy. See Section 22.31 for more details on the mechanics of this message and a usage example.

4012 Example:

4013 `Proxy-Require: foo`

## 4014 **22.29 RACK**

4015 The **RACK** header is sent in a **PRACK** request to support reliability of provisional responses. It contains two  
4016 numbers and a method tag. The first number is the value from the **RSeq** header in the provisional response  
4017 that is being acknowledged. The next number, and the method, are copied from the **CSeq** in the response  
4018 that is being acknowledged. The method name in the **RACK** header is case sensitive.

4019 Example:

4020 `RAck: 776656 1 INVITE`

## 4021 **22.30 Record-Route**

4022 The **Record-Route** is inserted by proxies in a request to force future requests in the session to route through  
4023 the proxy.

4024 Details of its use with the **Route** header field are described in Section 16.4.

4025 Example:

4026 `Record-Route: <sip:bob@biloxi.com;maddr=10.1.1.1> ,`  
4027 `<sip:bob@biloxi.com;maddr=10.2.1.1>`

## 4028 **22.31 Require**

4029 The **Require** header field is used by UACs to tell UASs about options that the UAC expects the UAS to  
4030 support in order to process the request. Although an optional header, the **Require** **MUST NOT** be ignored if  
4031 it is present.

4032 This is to ensure that the client-server interaction will proceed without delay when all options are understood  
4033 by both sides, and only slow down if options are not understood (as in the example above). For a well-matched  
4034 client-server pair, the interaction proceeds quickly, saving a round-trip often required by negotiation mechanisms.  
4035 In addition, it also removes ambiguity when the client requires features that the server does not understand. Some  
4036 features, such as call handling fields, are only of interest to end systems.

4037 Example:

4038 `Require: com.example.billing`

## 4039 **22.32 Retry-After**

4040 The **Retry-After** header field can be used with a 503 (Service Unavailable) response to indicate how long  
4041 the service is expected to be unavailable to the requesting client and with a 404 (Not Found), 600 (Busy), or

4042 603 (Decline) response to indicate when the called party anticipates being available again. The value of this  
4043 field can be either a **SIP-date** or an integer number of seconds (in decimal) after the time of the response.

4044 An optional comment can be used to indicate additional information about the time of callback. An  
4045 optional “duration” parameter indicates how long the called party will be reachable starting at the initial  
4046 time of availability. If no duration parameter is given, the service is assumed to be available indefinitely.

4047 Examples:

4048 `Retry-After: Mon, 21 Jul 1997 18:48:34 GMT (I'm in a meeting)`

4049 `Retry-After: Mon, 01 Jan 9999 00:00:00 GMT`

4050 `(Dear John: Don't call me back, ever)`

4051 `Retry-After: Fri, 26 Sep 1997 21:00:00 GMT;duration=3600`

4052 `Retry-After: 120`

4053 In the third example, the callee is reachable for one hour starting at 21:00 GMT. In the last example, the  
4054 delay is 2 minutes.

### 4055 **22.33 Route**

4056 The **Route** is used to force routing for a request through the listed set of proxies. Details of its use with the  
4057 **Record-Route** header field are described in Section 13.

4058 Example:

4059 `Route: <sip:bob@biloxi.com;maddr=10.1.1.1>, <sip:bob@10.4.1.4>`

### 4060 **22.34 RSeq**

4061 The **RSeq** header is used in provisional responses in order to transmit them reliably. It contains a single  
4062 numeric value from 1 to 2\*\*32 - 1. For details on its usage, see Section 18.1.

4063 Example:

4064 `RSeq: 988789`

### 4065 **22.35 Server**

4066 The **Server** header field contains information about the software used by the UAS to handle the request.  
4067 The syntax for this field is defined in [H14.38].

4068 Example:

4069 `Server: HomeProxy v2`

### 4070 **22.36 Subject**

4071 This header field provides a summary or indicates the nature of the call, allowing call filtering without  
4072 having to parse the session description. Note that the session description does not have to use the same  
4073 subject indication as the invitation.



4074 The compact form of the header is **s**.

4075 Example:

4076 Subject: Need more boxes

4077 s: Tech Support

### 4078 **22.37 Supported**

4079 The **Supported** header field enumerates all the extensions supported by the UAC or UAS. If empty, it means  
4080 that no extensions are supported.

4081 Example:

4082 Supported: foo, bar

### 4083 **22.38 Timestamp**

4084 The **Timestamp** header field describes when the UAC sent the request to the UAS. The use of the **Times-**  
4085 **tamp** is covered in Section 13.

4086 Example:

4087 Timestamp: 54

### 4088 **22.39 To**

4089 The **To** header field specifies the logical recipient of the request.

4090 The optional “display-name” is meant to be rendered by a human-user interface. The “tag” parameter  
4091 serves as a general mechanism to distinguish multiple instances of a user identified by a single SIP URI.

4092 See Section 13 for details of the “tag” parameter.

4093 Section 22.20 describes how **To** and **From** header fields are compared for the purpose of matching  
4094 requests to dialogs. Even if the “display-name” is empty, the “name-addr” form **MUST** be used if the  
4095 “addr-spec” contains a comma, question mark, or semicolon. Note that LWS is common, but **not** manda-  
4096 tory between the **display-name** and the “<”.

4097 The compact form of the header is **t**.

4098 The following are examples of valid **To** headers:

4099 To: The Operator <sip:operator@cs.columbia.edu>;tag=287447

4100 t: sip:+12125551212@server.phone2net.com

### 4101 **22.40 Unsupported**

4102 The **Unsupported** header field lists the features not supported by the UAS. See Section 22.31 for motivation.

4103 Example:

4104 Unsupported: foo

## 4105 22.41 User-Agent

4106 The **User-Agent** header field contains information about the UAC originating the request. The syntax and  
4107 semantics are defined in [H14.43].

4108 Example:

```
4109 User-Agent: Softphone Beta1.5
```

## 4110 22.42 Via

4111 The **Via** field indicates the path taken by the request so far and indicates the path that should be followed in  
4112 routing responses.

4113 The **Via** header field contains the transport protocol used to send the message, the client's host name or  
4114 network address and, if not the default port number, the port number at which it wishes to receive responses.  
4115 The **Via** header field can also contain parameters such as "maddr", "ttl", "received", and "branch", whose  
4116 meaning and use are described in other sections.

4117 The compact form of the header is **v**.

4118 Example:

```
4119 Via: SIP/2.0/UDP erlang.bell-telephone.com:5060  
4120 Via: SIP/2.0/UDP 128.59.16.1:5060 ;received=128.59.19.3
```

4121 In this example, the message originated from a multi-homed host with two addresses, 128.59.16.1  
4122 and 128.59.19.3. The sender guessed wrong as to which network interface would be used. Erlang.bell-  
4123 telephone.com noticed the mismatch and added a parameter to the previous hop's **Via** header field, contain-  
4124 ing the address that the packet actually came from.

4125 Another example:

```
4126 Via: SIP/2.0/UDP first.example.com:4000;ttl=16  
4127 ;maddr=224.2.0.1 ;branch=a7c6a8dlze.1
```

## 4128 22.43 Warning

4129 The **Warning** header field is used to carry additional information about the status of a response. **Warning**  
4130 headers are sent with responses and contain a three-digit warning code, host name, and warning text.

4131 The "warn-text" should be in a natural language that is most likely to be intelligible to the human user  
4132 receiving the response. This decision can be based on any available knowledge, such as the location of the  
4133 cache or user, the **Accept-Language** field in a request, or the **Content-Language** field in a response. The  
4134 default language is **i-default** [39].

4135 The first digit of warning codes beginning with "3" indicates warnings specific to SIP.

4136 This is a list of the currently-defined "warn-code"s, each with a recommended warn-text in English, and  
4137 a description of its meaning. Note that these warnings describe failures induced by the session description.

4138 Warnings 300 through 329 are reserved for indicating problems with keywords in the session description,  
4139 330 through 339 are warnings related to basic network services requested in the session description, 370  
4140 through 379 are warnings related to quantitative QoS parameters requested in the session description, and  
4141 390 through 399 are miscellaneous warnings that do not fall into one of the above categories.

4142 **300 Incompatible network protocol:** One or more network protocols contained in the session description  
4143 are not available.

4144 **301 Incompatible network address formats:** One or more network address formats contained in the ses-  
4145 sion description are not available.

4146 **302 Incompatible transport protocol:** One or more transport protocols described in the session descrip-  
4147 tion are not available.

4148 **303 Incompatible bandwidth units:** One or more bandwidth measurement units contained in the session  
4149 description were not understood.

4150 **304 Media type not available:** One or more media types contained in the session description are not avail-  
4151 able.

4152 **305 Incompatible media format:** One or more media formats contained in the session description are not  
4153 available.

4154 **306 Attribute not understood:** One or more of the media attributes in the session description are not sup-  
4155 ported.

4156 **307 Session description parameter not understood:** A parameter other than those listed above was not  
4157 understood.

4158 **330 Multicast not available:** The site where the user is located does not support multicast.

4159 **331 Unicast not available:** The site where the user is located does not support unicast communication (usu-  
4160 ally due to the presence of a firewall).

4161 **370 Insufficient bandwidth:** The bandwidth specified in the session description or defined by the media  
4162 exceeds that known to be available.

4163 **399 Miscellaneous warning:** The warning text can include arbitrary information to be presented to a hu-  
4164 man user or logged. A system receiving this warning MUST NOT take any automated action.

4165 1xx and 2xx have been taken by HTTP/1.1.

4166 If the warning is caused by the session description, the status response SHOULD include a session de-  
4167 scription similar to that included in OPTIONS responses indicating the capabilities of the UAS. Additional  
4168 "warn-code"s, as in the example below, can be defined through IANA.

4169 Examples:

4170 Warning: 307 isi.edu "Session parameter 'foo' not understood"  
4171 Warning: 301 isi.edu "Incompatible network address type 'E.164'"

## 4172 22.44 WWW-Authenticate

4173 The WWW-Authenticate header field consists of a challenge that indicates the authentication scheme and  
4174 parameters applicable for this Request-URI.

4175 The syntax for this header and use is defined in [H14.47]. See 20.2.2 for further details on its usage.

4176 Example:

4177 WWW-Authenticate: Digest realm="Bob's Friends",  
4178 domain="sip:boxesbybob.com",  
4179 nonce="f84f1cec41e6cbe5aea9c8e88d359",  
4180 opaque="", stale=FALSE, algorithm=MD5

## 4181 **23 Response Codes**

4182 The response codes are consistent with, and extend, HTTP/1.1 response codes. Not all HTTP/1.1 response  
4183 codes are appropriate, and only those that are appropriate are given here. Other HTTP/1.1 response codes  
4184 SHOULD NOT be used. Response codes not defined by HTTP/1.1 have codes x80 upwards to avoid clashes  
4185 with future HTTP response codes. Also, SIP defines a new class, 6xx. The default behavior for unknown  
4186 response codes is given for each category of codes.

### 4187 **23.1 Provisional 1xx**

4188 Provisional responses indicate that the server or proxy contacted is performing some further action and does  
4189 not yet have a definitive response. A server typically sends a 1xx response if it expects to take more than  
4190 200 ms to obtain a final response. Note that 1xx responses are not transmitted reliably, that is, they do not  
4191 cause the client to send an ACK.

4192 Provisional (1xx) responses MAY contain message bodies, including session descriptions.

4193 Provisional responses are also known as informational responses.

#### 4194 **23.1.1 100 Trying**

4195 This response indicates that the request has been received by the next hop server and that some unspeci-  
4196 fied action is being taken on behalf of this call (e.g., a database is being consulted). This response stops  
4197 retransmissions of an INVITE by a UAC.

#### 4198 **23.1.2 180 Ringing**

4199 The user agent receiving the INVITE is trying to alert the user. This response MAY be used to initiate local  
4200 ringback.

#### 4201 **23.1.3 181 Call Is Being Forwarded**

4202 A proxy server MAY use this status code to indicate that the call is being forwarded to a different set of  
4203 destinations.

#### 4204 **23.1.4 182 Queued**

4205 The called party is temporarily unavailable, but the callee has decided to queue the call rather than reject it.  
4206 When the callee becomes available, it will return the appropriate final status response. The reason phrase  
4207 MAY give further details about the status of the call, e.g., "5 calls queued; expected waiting time is 15  
4208 minutes". The server MAY issue several 182 responses to update the caller about the status of the queued  
4209 call.

### 4210 **23.1.5 183 Session Progress**

4211 The 183 (Session Progress) response is used to convey information about the progress of the call which is  
4212 not otherwise classified. The Reason-Phrase, header fields, or message body MAY be used to convey more  
4213 details about the call progress.

## 4214 **23.2 Successful 2xx**

4215 The request was successful.

### 4216 **23.2.1 200 OK**

4217 The request has succeeded. The information returned with the response depends on the method used in the  
4218 request.

## 4219 **23.3 Redirection 3xx**

4220 3xx responses give information about the user's new location, or about alternative services that might be  
4221 able to satisfy the call.

### 4222 **23.3.1 300 Multiple Choices**

4223 The address in the request resolved to several choices, each with its own specific location, and the user (or  
4224 user agent) can select a preferred communication end point and redirect its request to that location.

4225 The response MAY include a message body containing a list of resource characteristics and location(s)  
4226 from which the user or user agent can choose the one most appropriate, if allowed by the Accept request  
4227 header.

4228 The choices SHOULD also be listed as Contact fields (Section 22.10). Unlike HTTP, the SIP response  
4229 MAY contain several Contact fields or a list of addresses in a Contact field. User agents MAY use the  
4230 Contact header field value for automatic redirection or MAY ask the user to confirm a choice. However, this  
4231 specification does not define any standard for such automatic selection.

4232 This status response is appropriate if the callee can be reached at several different locations and the server cannot  
4233 or prefers not to proxy the request.

### 4234 **23.3.2 301 Moved Permanently**

4235 The user can no longer be found at the address in the Request-URI and the requesting client SHOULD retry  
4236 at the new address given by the Contact header field (Section 22.10). The caller SHOULD update any local  
4237 directories, address books and user location caches with this new value and redirect future requests to the  
4238 address(es) listed.

### 4239 **23.3.3 302 Moved Temporarily**

4240 The requesting client SHOULD retry the request at the new address(es) given by the Contact header field  
4241 (Section 22.10). The Request-URI of the new request uses the value of the Contact header in the response.  
4242 The new request can take two different forms. In the first approach, the To, From, Call-ID, and CSeq  
4243 header fields in the new request are the same as in the original request, with a new branch identifier in the

4244 Via header field. Proxies MUST follow this behavior and UACs MAY. In the second approach, UAs MAY  
4245 also use the Contact information for the To header field, as well as a new Call-ID value.

4246 The duration of the redirection can be indicated through an Expires (Section 22.19) header. If there is  
4247 no explicit expiration time, the address is only valid for this call and MUST NOT be cached for future calls.

#### 4248 **23.3.4 305 Use Proxy**

4249 The requested resource MUST be accessed through the proxy given by the Contact field. The Contact  
4250 field gives the URI of the proxy. The recipient is expected to repeat this single request via the proxy. 305  
4251 responses MUST only be generated by user agent servers.

#### 4252 **23.3.5 380 Alternative Service**

4253 The call was not successful, but alternative services are possible. The alternative services are described in  
4254 the message body of the response. Formats for such bodies are not defined here, and may be the subject of  
4255 future standardization.

### 4256 **23.4 Request Failure 4xx**

4257 4xx responses are definite failure responses from a particular server. The client SHOULD NOT retry the  
4258 same request without modification (e.g., adding appropriate authorization). However, the same request to a  
4259 different server might be successful.

#### 4260 **23.4.1 400 Bad Request**

4261 The request could not be understood due to malformed syntax. The Reason-Phrase SHOULD identify the  
4262 syntax problem in more detail, e.g., "Missing Call-ID header".

#### 4263 **23.4.2 401 Unauthorized**

4264 The request requires user authentication. This response is issued by user agent servers and registrars, while  
4265 407 (Proxy Authentication Required) is used by proxy servers.

#### 4266 **23.4.3 402 Payment Required**

4267 Reserved for future use.

#### 4268 **23.4.4 403 Forbidden**

4269 The server understood the request, but is refusing to fulfill it. Authorization will not help, and the request  
4270 SHOULD NOT be repeated.

#### 4271 **23.4.5 404 Not Found**

4272 The server has definitive information that the user does not exist at the domain specified in the Request-  
4273 URI. This status is also returned if the domain in the Request-URI does not match any of the domains  
4274 handled by the recipient of the request.

4275 **23.4.6 405 Method Not Allowed**

4276 The method specified in the Request-Line is not allowed for the address identified by the Request-URI.  
4277 The response MUST include an Allow header field containing a list of valid methods for the indicated address.

4278 **23.4.7 406 Not Acceptable**

4279 The resource identified by the request is only capable of generating response entities which have content  
4280 characteristics not acceptable according to the accept headers sent in the request.

4281 **23.4.8 407 Proxy Authentication Required**

4282 This code is similar to 401 (Unauthorized), but indicates that the client MUST first authenticate itself with  
4283 the proxy. SIP access authentication is explained in section 20 and 20.2.3.

4284 This status code can be used for applications where access to the communication channel (e.g., a tele-  
4285 phony gateway) rather than the callee requires authentication.

4286 **23.4.9 408 Request Timeout**

4287 The server could not produce a response within a suitable amount of time, for example, if it could not  
4288 determine the location of the user in time. The client MAY repeat the request without modifications at any  
4289 later time.

4290 **23.4.10 410 Gone**

4291 The requested resource is no longer available at the server and no forwarding address is known. This  
4292 condition is expected to be considered permanent. If the server does not know, or has no facility to determine,  
4293 whether or not the condition is permanent, the status code 404 (Not Found) SHOULD be used instead.

4294 **23.4.11 413 Request Entity Too Large**

4295 The server is refusing to process a request because the request entity is larger than the server is willing or  
4296 able to process. The server MAY close the connection to prevent the client from continuing the request.

4297 If the condition is temporary, the server SHOULD include a Retry-After header field to indicate that it is  
4298 temporary and after what time the client MAY try again.

4299 **23.4.12 414 Request-URI Too Long**

4300 The server is refusing to service the request because the Request-URI is longer than the server is willing to  
4301 interpret.

4302 **23.4.13 415 Unsupported Media Type**

4303 The server is refusing to service the request because the message body of the request is in a format not sup-  
4304 ported by the server for the requested method. The server SHOULD return a list of acceptable formats using  
4305 the Accept, Accept-Encoding and Accept-Language header fields. UAC processing of this response is  
4306 described in Section 8.1.3.5.

4307 **23.4.14 420 Bad Extension**

4308 The server did not understand the protocol extension specified in a **Proxy-Require** (Section 22.28) or **Re-**  
4309 **quire** (Section 22.31) header field. The server **SHOULD** include a list of the unsupported extensions in an  
4310 **Unsupported** header in the response. UAC processing of this response is described in Section 8.1.3.5.

4311 **23.4.15 421 Extension Required**

4312 The UAS needs a particular extension to process the request, but this extension is not listed in a **Supported**  
4313 header in the request. Responses with this status code **MUST** contain a **Require** header listing the required  
4314 extensions.

4315 In general, a UAS **SHOULD NOT** use this response when it wishes to apply an extension to a request. The  
4316 end result will often be no service at all, and a break in interoperability. Rather, servers **SHOULD** process the  
4317 request using baseline SIP capabilities and any extensions supported by the client.

4318 **23.4.16 480 Temporarily Unavailable**

4319 The callee's end system was contacted successfully but the callee is currently unavailable (e.g., not logged  
4320 in, logged in in such a manner as to preclude communication with the callee or activated the "do not disturb"  
4321 feature). The response **MAY** indicate a better time to call in the **Retry-After** header. The user could also be  
4322 available elsewhere (unbeknownst to this host). The reason phrase **SHOULD** indicate a more precise cause  
4323 as to why the callee is unavailable. This value **SHOULD** be setable by the user agent. Status 486 (Busy Here)  
4324 **MAY** be used to more precisely indicate a particular reason for the call failure.

4325 This status is also returned by a redirect server that recognizes the user identified by the **Request-URI**,  
4326 but does not currently have a valid forwarding location for that user.

4327 **23.4.17 481 Call/Transaction Does Not Exist**

4328 This status indicates that the UAS received a request that does not match any existing dialog or transaction.

4329 **23.4.18 482 Loop Detected**

4330 The server has detected a loop (Section 3).

4331 **23.4.19 483 Too Many Hops**

4332 The server received a request that contains a **Max-Forwards** (Section 22.22) header with the value zero.

4333 **23.4.20 484 Address Incomplete**

4334 The server received a request with a **Request-URI** that was incomplete. Additional information **SHOULD**  
4335 be provided.

4336 This status code allows overlapped dialing. With overlapped dialing, the client does not know the length of the  
4337 dialing string. It sends strings of increasing lengths, prompting the user for more input, until it no longer receives a  
4338 484 status response.



**4339 23.4.21 485 Ambiguous**

4340 The callee address provided in the request was ambiguous. The response MAY contain a listing of possible  
4341 unambiguous addresses in **Contact** headers.

4342 Revealing alternatives can infringe on privacy concerns of the user or the organization. It MUST be  
4343 possible to configure a server to respond with status 404 (Not Found) or to suppress the listing of possible  
4344 choices if the request address was ambiguous.

4345 Example response to a request with the URL `lee@example.com`:

4346 485 Ambiguous SIP/2.0

4347 Contact: Carol Lee <sip:carol.lee@example.com>

4348 Contact: Ping Lee <sip:p.lee@example.com>

4349 Contact: Lee M. Foote <sip:lee.foote@example.com>

4350 Some email and voice mail systems provide this functionality. A status code separate from 3xx is used since  
4351 the semantics are different: for 300, it is assumed that the same person or service will be reached by the choices  
4352 provided. While an automated choice or sequential search makes sense for a 3xx response, user intervention is  
4353 required for a 485 response.

**4354 23.4.22 486 Busy Here**

4355 The callee's end system was contacted successfully but the callee is currently not willing or able to take  
4356 additional calls at this end system. The response MAY indicate a better time to call in the **Retry-After**  
4357 header. The user could also be available elsewhere, such as through a voice mail service. Status 600 (Busy  
4358 Everywhere) SHOULD be used if the client knows that no other end system will be able to accept this call.

**4359 23.4.23 487 Request Terminated**

4360 The request was terminated by a BYE or CANCEL request. This response is never returned for a CANCEL  
4361 request itself.

**4362 23.4.24 488 Not Acceptable Here**

4363 The response has the same meaning as 606 (Not Acceptable), but only applies to the specific entity addressed  
4364 by the Request-URI and the request may succeed elsewhere.

**4365 23.5 Server Failure 5xx**

4366 5xx responses are failure responses given when a server itself has erred.

**4367 23.5.1 500 Server Internal Error**

4368 The server encountered an unexpected condition that prevented it from fulfilling the request. The client MAY  
4369 display the specific error condition, and MAY retry the request after several seconds.

4370 If the condition is temporary, the server MAY indicate when the client may retry the request using the  
4371 **Retry-After** header.

### 4372 **23.5.2 501 Not Implemented**

4373 The server does not support the functionality required to fulfill the request. This is the appropriate response  
4374 when a UAS does not recognize the request method and is not capable of supporting it for any user. (Proxies  
4375 forward all requests regardless of method.)

### 4376 **23.5.3 502 Bad Gateway**

4377 The server, while acting as a gateway or proxy, received an invalid response from the downstream server it  
4378 accessed in attempting to fulfill the request.

### 4379 **23.5.4 503 Service Unavailable**

4380 The server is currently unable to handle the request due to a temporary overloading (i.e., congestion) or  
4381 maintenance of the server. The implication is that this is a temporary condition which will be alleviated  
4382 after some delay. If known, the length of the delay MAY be indicated in a **Retry-After** header. If no **Retry-**  
4383 **After** is given, the client MUST handle the response as it would for a 500 response.

4384 A client (proxy or UAC) receiving a 503 SHOULD attempt to forward the request to an alternate server. It  
4385 SHOULD NOT forward any other requests to that server for the duration specified in the **Retry-After** header,  
4386 if present.

4387 Note: The existence of the 503 status code does not imply that a server has to use it when becoming  
4388 overloaded. Some servers MAY wish to simply refuse the connection.

### 4389 **23.5.5 504 Server Time-out**

4390 The server did not receive a timely response from the server (e.g., a location server) it accessed in attempting  
4391 to process the request. Note that 408 (Request Timeout) should be used if there was no response within the  
4392 period specified in the **Expires** header field from the upstream server.

### 4393 **23.5.6 505 Version Not Supported**

4394 The server does not support, or refuses to support, the SIP protocol version that was used in the request  
4395 message. The server is indicating that it is unable or unwilling to complete the request using the same major  
4396 version as the client, other than with this error message. The response MAY contain an entity describing why  
4397 that version is not supported and what other protocols are supported by that server. The format for such an  
4398 entity is not defined here and may be the subject of future standardization.

### 4399 **23.5.7 513 Message Too Large**

4400 The server was unable to process the request since the message length exceeded its capabilities.

## 4401 **23.6 Global Failures 6xx**

4402 6xx responses indicate that a server has definitive information about a particular user, not just the particular  
4403 instance indicated in the **Request-URI**.

### 4404 **23.6.1 600 Busy Everywhere**

4405 The callee's end system was contacted successfully but the callee is busy and does not wish to take the call  
4406 at this time. The response MAY indicate a better time to call in the **Retry-After** header. If the callee does  
4407 not wish to reveal the reason for declining the call, the callee uses status code 603 (Decline) instead. This  
4408 status response is returned only if the client knows that no other end point (such as a voice mail system) will  
4409 answer the request. Otherwise, 486 (Busy Here) should be returned.

### 4410 **23.6.2 603 Decline**

4411 The callee's machine was successfully contacted but the user explicitly does not wish to or cannot partici-  
4412 pate. The response MAY indicate a better time to call in the **Retry-After** header.

### 4413 **23.6.3 604 Does Not Exist Anywhere**

4414 The server has authoritative information that the user indicated in the **Request-URI** does not exist anywhere.

### 4415 **23.6.4 606 Not Acceptable**

4416 The user's agent was contacted successfully but some aspects of the session description such as the requested  
4417 media, bandwidth, or addressing style were not acceptable.

4418 A 606 (Not Acceptable) response means that the user wishes to communicate, but cannot adequately sup-  
4419 port the session described. The 606 (Not Acceptable) response MAY contain a list of reasons in a **Warning**  
4420 header field describing why the session described cannot be supported. Reasons are listed in Section 22.43.  
4421 It is hoped that negotiation will not frequently be needed, and when a new user is being invited to join an  
4422 already existing conference, negotiation may not be possible. It is up to the invitation initiator to decide  
4423 whether or not to act on a 606 (Not Acceptable) response.

## 4424 **24 Examples**

4425 In the following examples, we often omit the message body and the corresponding **Content-Length** and  
4426 **Content-Type** headers for brevity.

### 4427 **24.1 Registration**

4428 Bob registers on start-up. The message flow is shown in Figure 9.

```
4429  
4430 Fl REGISTER Bob -> Registrar  
4431  
4432 REGISTER sip:registrar.biloxi.com SIP/2.0  
4433 Via: SIP/2.0/UDP 10.4.1.4:5060  
4434 To: Bob <sip:bob@biloxi.com>  
4435 From: Bob <sip:bob@biloxi.com>;tag=456248  
4436 Call-ID: 843817637684230@phone21.bboxesbybob.com  
4437 CSeq: 1826 REGISTER
```

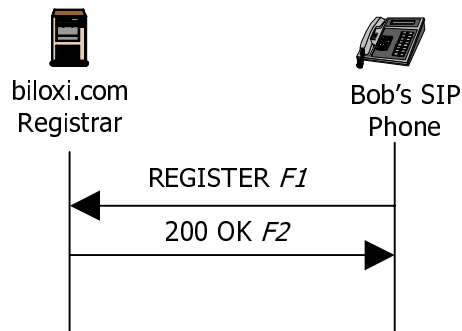


Figure 9: SIP Registration Example

```

4438 Contact: <sip:bob@10.4.1.4>
4439 Expires: 7200
4440 Content-Length: 0

```

4441 The registration expires after two hours. The registrar responds with a 200 OK:

```

4442
4443 F2 200 OK Registrar -> Bob
4444
4445 SIP/2.0 200 OK
4446 Via: SIP/2.0/UDP 10.4.1.4:5060
4447 To: Bob <sip:bob@biloxi.com>
4448 From: Bob <sip:bob@biloxi.com>;tag=456248
4449 Call-ID: 843817637684230@phone21.bboxesbybob.com
4450 CSeq: 1826 REGISTER
4451 Contact: <sip:bob@10.4.1.4>
4452 Expires: 7200
4453 Content-Length: 0
4454

```

## 4455 24.2 Session Setup

4456 This example contains the full details of the example session setup in Section 4. The message flow is shown  
4457 in Figure 1.

```

4458
4459 F1 INVITE Alice -> atlanta.com proxy
4460
4461 INVITE sip:bob@biloxi.com SIP/2.0
4462 Via: SIP/2.0/UDP 10.1.3.3:5060
4463 To: Bob <sip:bob@biloxi.com>

```

4464 From: Alice <sip:alice@atlanta.com>;tag=1928301774  
4465 Call-ID: a84b4c76e66710@10.1.3.3  
4466 CSeq: 314159 INVITE  
4467 Contact: <sip:alice@10.1.3.3>  
4468 Content-Type: application/sdp  
4469 Content-Length: 142  
4470  
4471 (Alice's SDP not shown)

4472  
4473 F2 100 Trying atlanta.com proxy -> Alice  
4474  
4475 SIP/2.0 100 Trying  
4476 Via: SIP/2.0/UDP 10.1.3.3:5060  
4477 To: Bob <sip:bob@biloxi.com>  
4478 From: Alice <sip:alice@atlanta.com>;tag=1928301774  
4479 Call-ID: a84b4c76e66710@10.1.3.3  
4480 CSeq: 314159 INVITE  
4481 Content-Length: 0

4482  
4483 F3 INVITE atlanta.com proxy -> biloxi.com proxy  
4484  
4485 INVITE sip:bob@biloxi.com SIP/2.0  
4486 Via: SIP/2.0/UDP 10.1.1.1:5060;branch=77ef4c2312983.1  
4487 Via: SIP/2.0/UDP 10.1.3.3:5060  
4488 To: Bob <sip:bob@biloxi.com>  
4489 From: Alice <sip:alice@atlanta.com>;tag=1928301774  
4490 Call-ID: a84b4c76e66710@10.1.3.3  
4491 CSeq: 314159 INVITE  
4492 Contact: <sip:alice@10.1.3.3>  
4493 Content-Type: application/sdp  
4494 Content-Length: 142  
4495  
4496 (Alice's SDP not shown)

4497  
4498 F4 100 Trying biloxi.com proxy -> atlanta.com proxy  
4499  
4500 SIP/2.0 100 Trying  
4501 Via: SIP/2.0/UDP 10.1.1.1:5060;branch=77ef4c2312983.1  
4502 Via: SIP/2.0/UDP 10.1.3.3:5060  
4503 To: Bob <sip:bob@biloxi.com>  
4504 From: Alice <sip:alice@atlanta.com>;tag=1928301774  
4505 Call-ID: a84b4c76e66710@10.1.3.3

4506 CSeq: 314159 INVITE  
4507 Content-Length: 0

4508  
4509 F5 INVITE biloxi.com proxy -> Bob  
4510  
4511 INVITE sip:bob@10.4.1.4 SIP/2.0  
4512 Via: SIP/2.0/UDP 10.2.1.1:5060;branch=4b43c2ff8.1  
4513 Via: SIP/2.0/UDP 10.1.1.1:5060;branch=77ef4c2312983.1  
4514 Via: SIP/2.0/UDP 10.1.3.3:5060  
4515 To: Bob <sip:bob@biloxi.com>  
4516 From: Alice <sip:alice@atlanta.com>;tag=1928301774  
4517 Call-ID: a84b4c76e66710@10.1.3.3  
4518 CSeq: 314159 INVITE  
4519 Contact: <sip:alice@10.1.3.3>  
4520 Content-Type: application/sdp  
4521 Content-Length: 142  
4522  
4523 (Alice's SDP not shown)

4524  
4525 F6 180 Ringing Bob -> biloxi.com proxy  
4526  
4527 SIP/2.0 180 Ringing  
4528 Via: SIP/2.0/UDP 10.2.1.1:5060;branch=4b43c2ff8.1  
4529 Via: SIP/2.0/UDP 10.1.1.1:5060;branch=77ef4c2312983.1  
4530 Via: SIP/2.0/UDP 10.1.3.3:5060  
4531 To: Bob <sip:bob@biloxi.com>;tag=a6c85cf  
4532 From: Alice <sip:alice@atlanta.com>;tag=1928301774  
4533 Call-ID: a84b4c76e66710@10.1.3.3  
4534 CSeq: 314159 INVITE  
4535 Content-Length: 0

4536  
4537 F7 180 Ringing biloxi.com proxy -> atlanta.com proxy  
4538  
4539 SIP/2.0 180 Ringing  
4540 Via: SIP/2.0/UDP 10.1.1.1:5060;branch=77ef4c2312983.1  
4541 Via: SIP/2.0/UDP 10.1.3.3:5060  
4542 To: Bob <sip:bob@biloxi.com>;tag=a6c85cf  
4543 From: Alice <sip:alice@atlanta.com>;tag=1928301774  
4544 Call-ID: a84b4c76e66710@10.1.3.3  
4545 CSeq: 314159 INVITE  
4546 Content-Length: 0

4547

4548 F8 180 Ringing atlanta.com proxy -&gt; Alice

4549

4550 SIP/2.0 180 Ringing

4551 Via: SIP/2.0/UDP 10.1.3.3:5060

4552 To: Bob &lt;sip:bob@biloxi.com&gt;;tag=a6c85cf

4553 From: Alice &lt;sip:alice@atlanta.com&gt;;tag=1928301774

4554 Call-ID: a84b4c76e66710@10.1.3.3

4555 CSeq: 314159 INVITE

4556 Content-Length: 0

4557

4558 F9 200 OK Bob -&gt; biloxi.com proxy

4559

4560 SIP/2.0 200 OK

4561 Via: SIP/2.0/UDP 10.2.1.1:5060;branch=4b43c2ff8.1

4562 Via: SIP/2.0/UDP 10.1.1.1:5060;branch=77ef4c2312983.1

4563 Via: SIP/2.0/UDP 10.1.3.3:5060

4564 To: Bob &lt;sip:bob@biloxi.com&gt;;tag=a6c85cf

4565 From: Alice &lt;sip:alice@atlanta.com&gt;;tag=1928301774

4566 Call-ID: a84b4c76e66710@10.1.3.3

4567 CSeq: 314159 INVITE

4568 Contact: &lt;sip:bob@10.4.1.4&gt;

4569 Content-Type: application/sdp

4570 Content-Length: 131

4571

4572 (Bob's SDP not shown)

4573

4574 F10 200 OK biloxi.com proxy -&gt; atlanta.com proxy

4575

4576 SIP/2.0 200 OK

4577 Via: SIP/2.0/UDP 10.1.1.1:5060;branch=77ef4c2312983.1

4578 Via: SIP/2.0/UDP 10.1.3.3:5060

4579 To: Bob &lt;sip:bob@biloxi.com&gt;;tag=a6c85cf

4580 From: Alice &lt;sip:alice@atlanta.com&gt;;tag=1928301774

4581 Call-ID: a84b4c76e66710@10.1.3.3

4582 CSeq: 314159 INVITE

4583 Contact: &lt;sip:bob@10.4.1.4&gt;

4584 Content-Type: application/sdp

4585 Content-Length: 131

4586

4587 (Bob's SDP not shown)

4588

4589 F11 200 OK atlanta.com proxy -> Alice

4590

4591 SIP/2.0 200 OK

4592 Via: SIP/2.0/UDP 10.1.3.3:5060

4593 To: Bob <sip:bob@biloxi.com>;tag=a6c85cf

4594 From: Alice <sip:alice@atlanta.com>;tag=1928301774

4595 Call-ID: a84b4c76e66710@10.1.3.3

4596 CSeq: 314159 INVITE

4597 Contact: <sip:bob@10.4.1.4>

4598 Content-Type: application/sdp

4599 Content-Length: 131

4600

4601 (Bob's SDP not shown)

4602

4603 F12 ACK Alice -> Bob

4604

4605 ACK sip:bob@10.4.1.4 SIP/2.0

4606 Via: SIP/2.0/UDP 10.1.3.3:5060

4607 To: Bob <sip:bob@biloxi.com>;tag=a6c85cf

4608 From: Alice <sip:alice@atlanta.com>;tag=1928301774

4609 Call-ID: a84b4c76e66710@10.1.3.3

4610 CSeq: 314159 ACK

4611 Content-Length: 0

4612 The media session between Alice and Bob is now established.

4613 Bob hangs up first. Note that Bob's SIP phone maintains its own CSeq numbering space, which, in this  
4614 example, begins with 231. Also not that since Bob is making the request, the To and From URLs and tags  
4615 have been swapped.

4616

4617 F13 BYE Bob -> Alice

4618

4619 BYE sip:alice@10.1.3.3 SIP/2.0

4620 Via: SIP/2.0/UDP 10.4.1.4:5060

4621 From: Bob <sip:bob@biloxi.com>;tag=a6c85cf

4622 To: Alice <sip:alice@atlanta.com>;tag=1928301774

4623 Call-ID: a84b4c76e66710@10.1.3.3

4624 CSeq: 231 BYE

4625 Content-Length: 0

4626

4627 F14 200 OK Alice -> Bob

4628

4629 SIP/2.0 200 OK



4630 Via: SIP/2.0/UDP 10.4.1.4:5060  
4631 From: Bob <sip:bob@biloxi.com>;tag=a6c85cf  
4632 To: Alice <sip:alice@atlanta.com>;tag=1928301774  
4633 Call-ID: a84b4c76e66710@10.1.3.3  
4634 CSeq: 231 BYE  
4635 Content-Length: 0

4636 The SIP Call Flows document [40] contains further examples of SIP messages.  
4637 ;; This buffer is for notes you don't want to save, and for Lisp evaluation. ;; If you want to create a file,  
4638 first visit that file with C-x C-f, ;; then enter the text in that file's own buffer.

## 4639 **25 Augmented BNF for the SIP Protocol**

4640 All of the mechanisms specified in this document are described in both prose and an augmented Backus-  
4641 Naur Form (BNF) similar to that used by RFC 2234 [41]. Implementors need to be familiar with the notation  
4642 in order to understand this specification. The augmented BNF includes the following constructs:

4643 name = definition

4644 The name of a rule is simply the name itself (without any enclosing "<" and ">") and is separated from  
4645 its definition by the equal "=" character. White space is only significant in that the indentation of continua-  
4646 tion lines indicates a rule definition that spans more than one line. Certain basic rules are in uppercase, such  
4647 as SP, LWS, HT, CRLF, DIGIT, ALPHA, etc. Angle brackets are used within definitions to clarify the use  
4648 of rule names.

4649 "literal"

4650 Quotation marks surround literal text. Unless stated otherwise, the text is case-insensitive.

4651 rule1 | rule2

4652 Elements separated by a bar ("|") are alternatives, that is, "yes | no" will accept yes or no.

4653 (rule1 rule2)

4654 Elements enclosed in parentheses are treated as a single element. Thus, "(elem (foo | bar) elem)" allows the  
4655 token sequences "elem foo elem" and "elem bar elem".

4656 \*rule

4657 The character "\*" preceding an element indicates repetition. The full form is "< n >\* < m >element"  
4658 indicating at least < n > and at most < m > occurrences of element. Default values are 0 and infinity so  
4659 that "(element)" allows any number, including zero; "1\*element" requires at least one; and "1\*2element"  
4660 allows one or two.

4661 [rule]

4662 Square brackets enclose optional elements; "[foo bar]" is equivalent to "\*1(foo bar)".

4663 N rule

4664 Specific repetition: "<n>(element)" is equivalent to "<n>\*<n>(element)"; that is, exactly <n> occur-  
 4665 rences of (element). Thus 2DIGIT is a 2-digit number, and 3ALPHA is a string of three alphabetic charac-  
 4666 ters.

4667 ; comment

4668 A semi-colon, set off some distance to the right of rule text, starts a comment that continues to the end of  
 4669 line. This is a simple way of including useful notes in parallel with the specifications.

## 4670 25.1 Basic Rules

4671 The following rules are used throughout this specification to describe basic parsing constructs. The US-  
 4672 ASCII coded character set is defined by ANSI X3.4-1986.

```

OCTET    = %x00-ff ; any 8-bit sequence of data
CHAR      = %x00-7f ; any US-ASCII character (octets 0 - 127)
upalpha   = "A" | "B" | "C" | "D" | "E" | "F" | "G" | "H" | "I" |
           "J" | "K" | "L" | "M" | "N" | "O" | "P" | "Q" | "R" |
           "S" | "T" | "U" | "V" | "W" | "X" | "Y" | "Z"
lowalpha  = "a" | "b" | "c" | "d" | "e" | "f" | "g" | "h" | "i" |
           "j" | "k" | "l" | "m" | "n" | "o" | "p" | "q" | "r" |
           "s" | "t" | "u" | "v" | "w" | "x" | "y" | "z"
alpha     = lowalpha | upalpha
DIGIT     = "0" | "1" | "2" | "3" | "4" | "5" | "6" | "7" |
           "8" | "9"
alphanum  = alpha | DIGIT
CTL       = %x00-1f | %x7f ; (octets 0 - 31) and DEL (127)
CR        = %d13 ; US-ASCII CR, carriage return character
LF        = %d10 ; US-ASCII LF, line feed character
SP        = %d32 ; US-ASCII SP, space character
HT        = %d09 ; US-ASCII HT, horizontal tab character
CRLF     = CR LF ; typically the end of a line

```

4673

4674 The following are defined in RFC 2396 [10] for the SIP URI:

```

reserved  = ", " | "/" | "?" | "." | "@" | " " | "|" | "+"
           | "$" | ";"
unreserved = alphanum | mark
mark      = "-" | "_" | "." | "!" | "~" | "*" | "'"
           | "(" | ")"
escaped   = "%" hex hex

```

4675

4676 SIP header field values can be folded onto multiple lines if the continuation line begins with a space or  
 4677 horizontal tab. All linear white space, including folding, has the same semantics as SP. A recipient MAY  
 4678 replace any linear white space with a single SP before interpreting the field value or forwarding the message  
 4679 downstream. This is intended to behave exactly as HTTP 1.1 as described in RFC2615 [9]. The SWS  
 4680 construct is similar to LWS but allows zero instances of space or tab

4681 LWS = \*( SP | HT ) [CRLF] 1\*( SP | HT ) ; linear whitespace  
 SWS = \*( SP | HT ) [CRLF] \*( SP | HT ) ; sep whitespace

4682 To separate the header name from the rest of value, a colon is used, which, by the above rule, allows  
 4683 whitespace before, but no line break, and whitespace after, including a linebreak. The HCOLON defines  
 4684 this construct.

4685 HCOLON = \*( SP | HT ) ":" SWS

4686 The TEXT-UTF8 rule is only used for descriptive field contents and values that are not intended to be  
 4687 interpreted by the message parser. Words of \*TEXT-UTF8 contain characters from the UTF-8 character  
 4688 set (RFC 2279 [12]). The TEXT-UTF8-TRIM rule is used for descriptive field contents that are *not* quoted  
 4689 strings, where leading and trailing LWS is not meaningful. In this regard, SIP differs from HTTP, which  
 4690 uses the ISO 8859-1 character set.

TEXT-UTF8 = \*(TEXT-UTF8char | LWS)  
 TEXT-UTF8-TRIM = \*TEXT-UTF8char \*(LWS TEXT-UTF8char)  
 TEXT-UTF8char = %x21-7e | UTF8-NONASCII  
 UTF8-NONASCII = %xc0-df 1UTF8-CONT  
 | %xe0-ef 2UTF8-CONT  
 | %xf0-f7 3UTF8-CONT  
 | %xf8-fb 4UTF8-CONT  
 | %xfc-fd 5UTF8-CONT  
 4691 UTF8-CONT = %x80-bf

4692 A CRLF is allowed in the definition of TEXT-UTF8 only as part of a header field continuation. It is  
 4693 expected that the folding LWS will be replaced with a single SP before interpretation of the TEXT-UTF8  
 4694 value.

4695 Hexadecimal numeric characters are used in several protocol elements. Some elements (authentication)  
 4696 force hex alphas to be lower case.

4697 LHEX = digit | "a" | "b" | "c" | "d" | "e" | "f"

4698 Others allow mixed upper and lower case

4699 hex = LHEX | "A" | "B" | "C" | "D" | "E" | "F"

4700 Many SIP header field values consist of words separated by LWS or special characters. Unless otherwise  
 4701 stated, tokens are case-insensitive. These special characters MUST be in a quoted string to be used within a  
 4702 parameter value.

token = 1\*(alphanum | "-" | "." | "!" | "%" | "\*" | "\_" | "+" | "'" | "" | """)  
 separators = "(" | ")" | "<" | ">" | "@" |  
 ";" | "," | ":" | "\" | "<>" |  
 "/" | "[" | "]" | "?" | "=" |  
 "{" | "}" | SP | HT

4703

4704 When tokens are used or separators are used between elements, whitespace is often allowed before or  
 4705 after these characters:

MINUS = SWS "-" SWS ; minus  
 DOT = SWS "." SWS ; period  
 PERCENT = SWS "%" SWS ; percent  
 BANG = SWS "!" SWS ; exclamation  
 PLUS = SWS "+" SWS ; plus  
 STAR = SWS "\*" SWS ; asterisk  
 TILDE = SWS "~" SWS ; tilde  
 EQUAL = SWS "=" SWS ; equal  
 LPAREN = SWS "(" SWS ; left parenthesis  
 RPAREN = SWS ")" SWS ; right parenthesis  
 LANGLE = SWS "<" SWS ; left angle bracket  
 RAQUOT = SWS ">" SWS ; right angle quote  
 LAQUOT = SWS "<"; left angle quote  
 RANGLE = SWS ">" SWS ; right angle bracket  
 BAR = SWS "|" SWS ; vertical bar  
 ATSIGN = SWS "@" SWS ; atsign  
 COMMA = SWS "," SWS ; comma  
 SEMI = SWS ";" SWS ; semicolon  
 COLON = SWS ":" SWS ; colon  
 DQUOT = SWS "<>" SWS ; double quotation mark  
 LDQUOT = SWS "<"; open double quotation mark  
 RDQUOT = SWS ">" SWS ; close double quotation mark  
 LBRACK = SWS "{" SWS ; left square bracket  
 RBRACK = SWS "}" SWS ; right square bracket

4706

4707 Comments can be included in some SIP header fields by surrounding the comment text with parentheses.  
 4708 Comments are only allowed in fields containing "comment" as part of their field value definition. In all other  
 4709 fields, parentheses are considered part of the field value.

comment = LPAREN \*(ctext | quoted-pair | comment) RPAREN  
 ; ctext includes all chars except left and right parens  
 4710 ctext = %x21-27 | %x2a-7e | UTF8-NONASCII | LW

4710

4711 A string of text is parsed as a single word if it is quoted using double-quote marks. In quoted strings,  
 4712 quotation marks (") and backslashes (\) need to be escaped.

quoted-string = ( SWS "<>" \*(qdtex | quoted-pair) "<>" )  
 qdtex = LWS | %x21 | %x23-5b | %x5d-7e  
 4713 | UTF8-NONASCII

4713

4714 The backslash character ("") MAY be used as a single-character quoting mechanism only within quoted-  
 4715 string and comment constructs. Unlike HTTP/1.1, the characters CR and LF cannot be escaped by this  
 4716 mechanism to avoid conflict with line folding and header separation.

```

4717 quoted-pair    = "\" (%x00 - %x09 | %x0b | %x0c
                | %x0e - %x7f)

SIP-URL         = "sip:" [ userinfo "@" ] hostport
                url-parameters [ headers ]
userinfo        = [ user | telephone-subscriber [ ":" password ] ]
user            = *( unreserved | escaped | user-unreserved )
user-unreserved = "&" | "=" | "+" | "$" | "," | ";" | "?" | "/"
password       = *( unreserved | escaped |
                "&" | "=" | "+" | "$" | "," )

hostport        = host [ ":" port ]
host            = hostname | IPv4address | IPv6reference
hostname       = *( domainlabel "." ) toplabel [ "." ]
domainlabel    = alphanum
                | alphanum *( alphanum | "-" ) alphanum
4718 toplabel     = alpha | alpha *( alphanum | "-" ) alphanum

IPv4address     = 1*3DIGIT "." 1*3DIGIT "." 1*3DIGIT "." 1*3DIGIT
IPv6reference   = "[" IPv6address "]"
IPv6address     = hexpart [ ":" IPv4address ]
hexpart        = hexseq | hexseq ":" [ hexseq ] | ":" [ hexseq ]
hexseq         = hex4 *( ":" hex4 )
hex4           = 1*4HEX
4719 port        = 1*DIGIT

url-parameters = *( "," url-parameter )
url-parameter   = transport-param | user-param | method-param
                | ttl-param | maddr-param | other-param
transport-param = "transport="
                ( "udp" | "tcp" | "sctp" | "tls"
                | other-transport )
other-transport = token
user-param     = "user=" ( "phone" | "ip" | other-user )
other-user     = token
method-param   = "method=" Method
ttl-param      = "ttl=" ttl
maddr-param    = "maddr=" host
other-param    = pname [ "=" pvalue ]
pname         = 1*paramchar
pvalue        = 1*paramchar
paramchar     = param-unreserved | unreserved | escaped
4720 param-unreserved = "[" | "]" | "/" | "." | "&" | "+" | "$"

```

headers = "?" header \*( "&" header )  
 header = hname "=" hvalue  
 hname = 1\*( hnv-unreserved | unreserved | escaped )  
 hvalue = \*( hnv-unreserved | unreserved | escaped )  
 hnv-unreserved = "[ | "]" | "/" | "?" | ":" | "+" | "\$"

SIP-message = Request | Response  
 Request = Request-Line  
           \*( message-header )  
           CRLF  
           [ message-body ]  
 Request-Line = Method SP Request-URI SP SIP-Version CRLF  
 Request-URI = SIP-URL | absoluteURI  
 absoluteURI = scheme COLON ( hier-part | opaque-part )  
 hier-part = ( net-path | abs-path ) [ "?" query ]  
 net-path = "/" authority [ abs-path ]  
 abs-path = "/" path-segments  
 opaque-part = uric-no-slash \*uric  
 uric = reserved — unreserved — escaped  
 uric-no-slash = unreserved | escaped | ";" | "?" | ":" | "@"  
               | "&" | "=" | "+" | "\$" | ","  
 scheme = alpha \*( alpha | digit | "+" | "-" | "." )  
 authority = server | reg-name  
 server = [ [ userinfo "@" ] hostport ]  
 reg-name = 1\*( unreserved | escaped | "\$" | "  
               | ";" | "." | "@" | "&" | "=" | "+" )  
 query = \*uric  
 SIP-Version = "SIP/2.0"

## message-header

```
= Accept
| Accept-Encoding
| Accept-Language
| Alert-Info
| Allow
| Authentication-Info
| Authorization
| Call-ID
| Call-Info
| Contact
| Content-Disposition
| Content-Encoding
| Content-Language
| Content-Length
| Content-Type
| CSeq
| Date
| Error-Info
| Expires
| From
| In-Reply-To
| Max-Forwards
| MIME-Version
| Organization
| Priority
| Proxy-Authenticate
| Proxy-Authorization
| Proxy-Require
| RAck
| Record-Route
| Require
| Retry-After
| Route
| RSeq
| Server
| Subject
| Supported
| Timestamp
| To
| Unsupported
| User-Agent
| Via
| Warning
| WWW-Authenticate
```

4723

Method = "INVITE" | "ACK" | "OPTIONS" | "BYE"  
 | "CANCEL" | "REGISTER" | "PRACK"  
 | extension-method  
 extension-method = token  
 option-tag = token  
 Response = Status-Line  
 \*( message-header )  
 CRLF  
 [ message-body ]

4724

Status-Line = SIP-version SP Status-Code SP Reason-Phrase CRLF  
 Status-Code = Informational  
 | Redirection  
 | Success  
 | Client-Error  
 | Server-Error  
 | Global-Failure  
 | extension-code  
 extension-code = 3DIGIT  
 Reason-Phrase = \*(reserved — unreserved — escaped — SP — HT)

4725

Informational = "100" ; Trying  
 | "180" ; Ringing  
 | "181" ; Call Is Being Forwarded  
 | "182" ; Queued  
 | "183" ; Session Progress

4726

Success = "200" ; OK

4727

Redirection = "300" ; Multiple Choices  
 | "301" ; Moved Permanently  
 | "302" ; Moved Temporarily  
 | "305" ; Use Proxy  
 | "380" ; Alternative Service

4728



Client-Error = "400" ; Bad Request  
| "401" ; Unauthorized  
| "402" ; Payment Required  
| "403" ; Forbidden  
| "404" ; Not Found  
| "405" ; Method Not Allowed  
| "406" ; Not Acceptable  
| "407" ; Proxy Authentication Required  
| "408" ; Request Timeout  
| "409" ; Conflict  
| "410" ; Gone  
| "413" ; Request Entity Too Large  
| "414" ; Request-URI Too Large  
| "415" ; Unsupported Media Type  
| "420" ; Bad Extension  
| "480" ; Temporarily not available  
| "481" ; Call Leg/Transaction Does Not Exist  
| "482" ; Loop Detected  
| "483" ; Too Many Hops  
| "484" ; Address Incomplete  
| "485" ; Ambiguous  
| "486" ; Busy Here  
| "487" ; Request Terminated  
4729 | "488" ; Not Acceptable Here

Server-Error = "500" ; Internal Server Error  
| "501" ; Not Implemented  
| "502" ; Bad Gateway  
| "503" ; Service Unavailable  
| "504" ; Server Time-out  
4730 | "505" ; SIP Version not supported

Global-Failure = "600" ; Busy Everywhere  
| "603" ; Decline  
| "604" ; Does not exist anywhere  
4731 | "606" ; Not Acceptable

Accept = "Accept" HCOLON ( accept-range \*(COMMA accept-range) )  
 accept-range = media-range [ accept-params ]  
 media-range = ( "\*"/\*"  
               | ( m-type SWS "/" "\*" SWS )  
               | ( m-type SLASH m-subtype )  
               ) \*( SEMI parameter )  
 accept-params = SEMI "q" EQUAL qvalue \*( accept-extension )  
 accept-extension = SEMI ae-name [ EQUAL ae-value ]  
 ae-name = token  
 4732 ae-value = token | quoted-string

Accept-Encoding = "Accept-Encoding" HCOLON ( encoding \*(COMMA encoding) )  
 encoding = codings [ SEMI "q" EQUAL qvalue ]  
 codings = content-coding | "\*"

4733 content-coding = token  
 qvalue = ( "0" [ "." 0\*3DIGIT ] )  
           | ( "1" [ "." 0\*3("0") ] )

Accept-Language = "Accept-Language" HCOLON ( language \*(COMMA language) )  
 language = language-range [ SEMI "q" EQUAL qvalue ]  
 4734 language-range = ( ( 1\*8ALPHA \*( MINUS 1\*8ALPHA ) ) | "\*" )

Alert-Info = "Alert-Info" HCOLON alert-param \*(COMMA alert-param)  
 alert-param = LAQUOT URI RAQUOT \*( COLON generic-param )  
 generic-param = token [ EQUAL gen-value ]  
 4735 gen-value = token | host | quoted-string

4736 Allow = "Allow" HCOLON Method \*(COMMA Method)

Authorization = "Authorization" HCOLON credentials  
 credentials = "Digest" digest-response  
 digest-response = dig-resp \*(COMMA dig-resp)  
 dig-resp = username | realm | nonce | digest-uri  
           | dresponse | [ algorithm ] | [ cnonce ]  
           | [ opaque ] | [ message-qop ]  
           | [ nonce-count ] | [ auth-param ]  
 username = "username" EQUAL username-value  
 username-value = quoted-string  
 digest-uri = "uri" EQUAL digest-uri-value  
 digest-uri-value = request-uri ; As specified by HTTP/1.1  
 message-qop = "qop" EQUAL qop-value  
 cnonce = "cnonce" EQUAL cnonce-value  
 cnonce-value = nonce-value  
 nonce-count = "nc" EQUAL nc-value  
 nc-value = 8LHEX  
 dresponse = "response" EQUAL request-digest  
 request-digest = LDQUOT 32LHEX RDQUOT  
 auth-param = auth-param-name EQUAL ( token | quoted-string )  
 auth-param-name = token

4737

AuthenticationInfo = "Authentication-Info" COLON auth-info  
 auth-info = auth-inf \*(COMMA auth-inf)  
 auth-inf = nextnonce | [ message-qop ]  
           | [ response-auth ] | [ cnonce ]  
           | [ nonce-count ]  
 nextnonce = "nextnonce" EQUAL nonce-value  
 response-auth = "rspauth" EQUAL response-digest  
 response-digest = LDQUOT \*LHEX RDQUOT

4738

Call-ID = ( "Call-ID" | "" ) HCOLON callid  
 callid = token [ ATSIGN token ]

4739

Call-Info = "Call-Info" HCOLON info \*(COMMA info)  
 info = LAQUOT URI RAQUOT \*( SEMI info-param)  
 info-param = "purpose" EQUAL ( "icon" | "info"  
           | "card" | token ) | generic-param

4740

Contact = ("Contact" | "m" ) HCOLON  
           (STAR | contact-param \*(COMMA contact-param))  
 contact-param = name-addr | addr-spec \*(SEMI contact-params)  
 name-addr = [ display-name ] LAQUOT addr-spec RAQUOT  
 addr-spec = SIP-URL | URI  
 display-name = \*(token LWS) | quoted-string

4741

contact-params = c-p-q | c-p-action | c-p-expires  
                   | contact-extension  
 c-p-q          = "q" EQUAL qvalue  
 c-p-action     = "action" EQUAL ("proxy" | "redirect")  
 c-p-expires    = "expires" EQUAL ( delta-seconds  
                   | LDQUOT SIP-date RDQUOT)  
 contact-extension = generic-param  
 qvalue          = ( "0" [ "." 0\*3DIGIT ] )  
 4742            | ( "1" [ "." 0\*3("0") ] )  
 4743  
 delta-seconds = 1\*DIGIT  
  
 Content-Disposition = "Content-Disposition" HCOLON  
                       disposition-type \*( SEMI disposition-param )  
 disposition-type = "render" | "session" | "icon" | "alert"  
                   | disp-extension-token  
 disposition-param = "handling" EQUAL  
                   ( "optional" | "required"  
                   | other-handling ) | generic-param  
 other-handling = token  
 4744 disp-extension-token = token  
  
 Content-Encoding = ( "Content-Encoding" | "e" ) HCOLON  
 4745                  content-coding \*(COMMA content-coding)  
  
 Content-Language = "Content-Language" HCOLON  
                   language-tag \*(COMMA language-tag)  
 language-tag     = primary-tag \*( MINUS subtag )  
 primary-tag      = 1\*8ALPHA  
 4746 subtag          = 1\*8ALPHA  
  
 4747 Content-Length = ( "Content-Length" | "l" ) HCOLON 1\*DIGIT  
  
 Content-Type     = ( "Content-Type" | "c" ) HCOLON media-type  
 media-type      = m-type SLASH m-subtype \*(SEMI m-parameter)  
 m-type          = discrete-type | composite-type  
 discrete-type   = "text" | "image" | "audio" | "video"  
                   | "application" | extension-token  
 composite-type  = "message" | "multipart" | extension-token  
 extension-token = ietf-token | x-token  
 ietf-token      = token  
 x-token         = ("X" | "x") "-" token  
 m-subtype       = extension-token | iana-token  
 iana-token      = token  
 parameter       = m-attribute EQUAL m-value  
 m-attribute     = token  
 4748 m-value         = token | quoted-string

4749 CSeq = "CSeq" HCOLON 1\*DIGIT LWS Method

Date = "Date" HCOLON SIP-date  
SIP-date = rfc1123-date  
rfc1123-date = wkday COMMA date1 SP time SP "GMT"  
date1 = 2DIGIT SP month SP 4DIGIT  
; day month year (e.g., 02 Jun 1982)  
time = 2DIGIT ":" 2DIGIT ":" 2DIGIT  
; 00:00:00 - 23:59:59  
wkday = "Mon" | "Tue" | "Wed"  
| "Thu" | "Fri" | "Sat" | "Sun"  
month = "Jan" | "Feb" | "Mar" | "Apr"  
| "May" | "Jun" | "Jul" | "Aug"  
4750 | "Sep" | "Oct" | "Nov" | "Dec"

Error-Info = "Error-Info" HCOLON error-uri \*(COMMA error-uri)  
4751 error-uri = LAQUOT URI RAQUOT \*( SEMI generic-param )

Expires = "Expires" HCOLON ( SIP-date | delta-seconds )  
From = ( "From" | "f" ) HCOLON from-spec  
from-spec = ( name-addr | addr-spec )  
\*( SEMI from-param )  
from-param = tag-param | generic-param  
4752 tag-param = "tag" EQUAL token

4753 In-Reply-To = "In-Reply-To" HCOLON called \*(COMMA called)

4754 Max-Forwards = "Max-Forwards" HCOLON 1\*DIGIT

4755 MIME-Version = "MIME-Version" HCOLON 1\*DIGIT ":" 1\*DIGIT

4756 Organization = "Organization" HCOLON TEXT-UTF8-TRIM

Priority = "Priority" HCOLON priority-value  
priority-value = "emergency" | "urgent" | "normal"  
| "non-urgent" | other-priority  
4757 other-priority = token

Proxy-Authenticate = "Proxy-Authenticate" HCOLON  
 challenge \*(COMMA challenge)  
 challenge = "Digest" digest-challenge  
 digest-challenge = digest-chlng \*(COMMA digest-chlng)  
 digest-chlng = realm | [ domain ] | nonce  
 | [ opaque ] | [ stale ] | [ algorithm ]  
 | [ qop-options ] | [auth-param]  
 realm = "realm" EQUALS realm-value  
 realm-value = quoted-string  
 domain = "domain" EQUAL LDQUOT URI  
 ( 1\*SP URI ) RDQUOT  
 URI = absoluteURI | abs\_path  
 nonce = "nonce" EQUAL nonce-value  
 nonce-value = quoted-string  
 opaque = "opaque" EQUAL quoted-string  
 stale = "stale" EQUAL ( "true" | "false" )  
 algorithm = "algorithm" EQUAL ( "MD5" | "MD5-sess"  
 | token )  
 qop-options = "qop" EQUAL LDQUOT qop-value \*(COMMA qop-value) RDQUOT  
 qop-value = "auth" | "auth-int" | token

4758

4759 Proxy-Authorization = "Proxy-Authorization" HCOLON credentials

4760 Proxy-Require = "Proxy-Require" HCOLON option-tag \*(COMMA option-tag)

RAck = "RAck" HCOLON response-num LWS CSeq-num LWS Method  
 response-num = 1\*DIGIT  
 CSeq-num = 1\*DIGIT  
 response-num = 1\*DIGIT

4761

Record-Route = "Record-Route" HCOLON rec-route \*(COMMA rec-route)  
 rec-route = name-addr \*( SEMI rr-param )  
 rr-param = generic-param  
 Require = "Require" HCOLON option-tag \*(COMMA option-tag)

4762

Retry-After = "Retry-After" HCOLON  
 ( SIP-date | delta-seconds )  
 [ comment ] \*( SEMI retry-param )  
 retry-param = "duration" EQUAL delta-seconds  
 | generic-param

4763

Route = "Route" HCOLON route=param \*(COMMA route-param)  
 route-param = name-addr \*( SEMI rr-param )

4764

4765 RSeq = "RSeq" HCOLON response-num

Server = "Server" HCOLON 1\*( product | comment )  
 product = token [SLASH product-version]  
 4766 product-version = token  
 4767 Subject = ( "Subject" | "s" ) HCOLON TEXT-UTF8-TRIM  
 Supported = ( "Supported" | "k" ) HCOLON  
 4768 (option-tag \*(COMMA option-tag)  
 Timestamp = "Timestamp" HCOLON \*(DIGIT  
 [ "." \*(DIGIT) ] [ delay ]  
 4769 delay = \*(DIGIT) [ "." \*(DIGIT) ]  
 To = ( "To" | "t" ) HCOLON ( name-addr  
 | addr-spec ) \*( SEMI to-param )  
 4770 to-param = tag-param | generic-param  
 4771 Unsupported = "Unsupported" HCOLON option-tag \*(COMMA option-tag)  
 4772 User-Agent = "User-Agent" HCOLON 1\*( product | comment )  
 Via = ( "Via" | "v" ) HCOLON via-parm \*(COMMA via-parm)  
 via-parm sent-protocol sent-by \*( SEMI via-params ) [ comment ] )  
 via-params = via-hidden | via-ttl | via-maddr  
 | via-received | via-branch  
 | via-extension  
 via-hidden = "hidden"  
 via-ttl = "ttl" EQUAL ttl  
 via-maddr = "maddr" EQUAL host  
 via-received = "received" EQUAL host  
 via-branch = "branch" EQUAL token  
 via-extension = generic-param  
 sent-protocol = protocol-name SLASH protocol-version  
 SLASH transport  
 protocol-name = "SIP" | token  
 protocol-version = token  
 transport = "UDP" | "TCP" | "TLS" | "SCTP"  
 | other-transport  
 sent-by = host [ COLON port ]  
 4773 ttl = 1\*3DIGIT ; 0 to 255  
 Warning = "Warning" HCOLON warning-value \*(COMMA warning-value)  
 warning-value = warn-code SP warn-agent SP warn-text  
 warn-code = 3DIGIT  
 warn-agent = ( host [ COLON port ] ) | pseudonym  
 ; the name or pseudonym of the server adding  
 ; the Warning header, for use in debugging  
 warn-text = quoted-string  
 4774 pseudonym = token

4775 WWW-Authenticate = "WWW-Authenticate" HCOLON challenge  
4776 message-body = \*OCTET

## 4777 26 IANA Considerations

4778 All new or experimental method names, header field names, and status codes used in SIP applications  
4779 SHOULD be registered with IANA in order to prevent potential naming conflicts. It is RECOMMENDED that  
4780 new "option-tag"s and "warn-code"s also be registered. Before IANA registration, new protocol elements  
4781 SHOULD be characterized in an Internet-Draft or, preferably, an RFC.

4782 For Internet-Drafts, IANA is requested to make the draft available as part of the registration database.

4783 By the time an RFC is published, colliding names may have already been implemented.

4784 When a registration for either a new header field, new method, or new status code is created based on  
4785 an Internet-Draft, and that Internet-Draft becomes an RFC, the person that performed the registration MUST  
4786 notify IANA to change the registration to point to the RFC instead of the Internet-Draft.

4787 Registrations should be sent to `iana@iana.org`.

### 4788 26.1 Option Tags

4789 Option tags are used in header fields such as Require, Supported, Proxy-Require, and Unsupported in  
4790 support of SIP compatibility mechanisms for extensions. For more on the use of option tags in these header  
4791 fields, see Section 21.2. The option tag itself is a string that is associated with a particular SIP option (that  
4792 is, an extension) that identifies the option in signaling between SIP endpoints.

4793 When registering a new SIP option with IANA, the following information MUST be provided:

- 4794 • Name and description of option. The name MAY be of any length, but SHOULD be no more than  
4795 twenty characters long. The name MUST consist of alphanum (See Section 25) characters only.
- 4796 • A listing of any new SIP header fields, header parameter fields, or parameter values defined by this  
4797 option. A SIP option MUST NOT redefine header fields or parameters defined in either RFC 2543, any  
4798 standards-track extensions to RFC 2543, or other extensions registered through IANA.
- 4799 • Indication of who has change control over the option (for example, IETF, ISO, ITU-T, other interna-  
4800 tional standardization bodies, a consortium, or a particular company or group of companies).
- 4801 • A reference to a further description if available, for example (in order of preference) an RFC, a pub-  
4802 lished paper, a patent filing, a technical report, documented source code, or a computer manual.
- 4803 • Contact information (postal and email address).

4804 This procedure has been borrowed from RTSP [4] and the RTP AVP [42].

#### 4805 26.1.1 Registration of 100rel

4806 This specification registers a single option tag, "100rel". The required information is:

4807 **Name:** "100rel"



4808 **Description:** This option tag is for reliability of provisional responses. When present in a **Supported**  
4809 header, it indicates that the UA can send or receive reliable provisional responses. When present in a  
4810 **Require** header in a request, it indicates that the UAS **MUST** send all provisional responses reliably.  
4811 When present in a **Require** header in a reliable provisional response, it indicates that the response is  
4812 to be sent reliably.

4813 **New Headers:** The **RSeq** and **RAck** header fields are defined by this option.

4814 **Change Control:** IETF.

4815 **Reference:** RFCXXXX [Note to IANA: Fill in with the RFC number of this specification.

4816 **Contact Information:** Jonathan Rosenberg, jdrosen@jdrosen.net. 72 Eagle Rock Avenue, First Floor, East  
4817 Hanover, NJ, 07936.

## 4818 26.2 Warn-Codes

4819 Warning codes provide information supplemental to the status code in SIP response messages when the  
4820 failure of the transaction results from a Session Description Protocol (SDP, [6]). New “warn-code” values  
4821 can be registered with IANA as they arise.

4822 The “warn-code” consists of three digits. A first digit of “3” indicates warnings specific to SIP.

4823 Warnings 300 through 329 are reserved for indicating problems with keywords in the session description,  
4824 330 through 339 are warnings related to basic network services requested in the session description, 370  
4825 through 379 are warnings related to quantitative QoS parameters requested in the session description, and  
4826 390 through 399 are miscellaneous warnings that do not fall into one of the above categories.

4827 1xx and 2xx have been taken by HTTP/1.1.

## 4828 26.3 Header Field Names

4829 Header field names do not require working group or working group chair review prior to IANA registration,  
4830 but **SHOULD** be documented in an RFC or Internet-Draft before IANA is consulted.

4831 The following information needs to be provided to IANA in order to register a new header field name:

- 4832 ● The name and email address of the individual performing the registration.
- 4833 ● The name of the header field being registered.
- 4834 ● A compact form version for that header field, if one is defined.
- 4835 ● The name of the draft or RFC where the header field is defined.
- 4836 ● A copy of the draft or RFC where the header field is defined.

4837 Header fields **SHOULD NOT** use the X- prefix notation and **MUST NOT** duplicate the names of header  
4838 fields used by SMTP or HTTP unless the syntax is a compatible superset and the semantics are similar.  
4839 Some common and widely used header fields **MAY** be assigned one-letter compact forms (Section 7.3.3).  
4840 Compact forms can only be assigned after SIP working group review. In the absence of this working group,  
4841 a designated expert reviews the request.

## 4842 **26.4 Method and Response Codes**

4843 Because the status code space is limited, they do require working group or working group chair review, and  
4844 MUST be documented in an RFC or Internet draft. The same procedures apply to new method names.

4845 The following information needs to be provided to IANA in order to register a new response code or  
4846 method:

- 4847 • The name and email address of the individual performing the registration.
- 4848 • The number of the response code or name of the method being registered.
- 4849 • The default reason phrase for that status code, if applicable.
- 4850 • The name of the draft or RFC where the method or status code is defined.
- 4851 • A copy of the draft or RFC where the method or status code is defined.

## 4852 **27 Changes Made in Version 00**

- 4853 • Indicated that UAC should send both CANCEL and BYE after a retransmission fails.
- 4854 • Added semicolon and question mark to the list of unreserved characters for the user part of SIP URLs  
4855 to handle tel: URLs properly.
- 4856 • Uniform handling of if hop count Max-Forwards: return 483. Note that this differs from HTTP/1.1  
4857 behavior, where only OPTIONS and TRACE allow this header, but respond as the final recipient when  
4858 the value reaches zero.
- 4859 • Clarified that a forking proxy sends ACKs only for INVITE requests.
- 4860 • Clarified wording of DNS caching. Added paragraph on “negative caching”, i.e., what to do if one  
4861 of the hosts failed. It is probably not a good idea to simply drop this host from the list if the DNS ttl  
4862 value is more than a few minutes, since that would mean that load balancing may not work for quite a  
4863 while after a server is brought back on line. This will be true in particular if a server group receives a  
4864 large number of requests from a small number of upstream servers, as is likely to be the case for calls  
4865 between major consumer ISPs. However, without getting into arbitrary and complicated retry rules, it  
4866 seems hard to specify any general algorithm. Might it be worthwhile to simply limit the “black list”  
4867 interval to a few minutes?
- 4868 • Added optional Call-Info and Alert-Info header fields that describe the caller and information to be  
4869 used in alerting. (Currently, avoided use of “purpose” qualification since it is not yet clear whether  
4870 rendering content without understanding its meaning is always appropriate. For example, if a UAS  
4871 does not understand that this header is to replace ringing, it would mix both local ring tone and the  
4872 indicated sound URL.) TBD!
- 4873 • SDP “s=” lines can’t be empty, unfortunately.
- 4874 • Noted that maddr could also contain a unicast address, but SHOULD contain the multicast address if  
4875 the request is sent via multicast (Section 22.42).

- 4876 • Clarified that responses are sent to port in *Via sent-by* value.
- 4877 • Added “other-\*” to the *user* URL parameter and the *Hide* and *Content-Disposition* headers.
- 4878 • Clarified generation of timeout (408) responses in forking proxies and mention the *Expires* header.
- 4879 • Clarified that *CANCEL* and *INVITE* are separate transactions (Fig. 7). Thus, the *INVITE* request  
4880 generates a 487 (Request Terminated) if a *CANCEL* or *BYE* arrives.
- 4881 • Clarified that *Record-Route* SHOULD be inserted in every request, but that the route, once estab-  
4882 lished, persists. This provides robustness if the called UAS crashes.
- 4883 • Emphasized that proxy, redirect, registrar and location servers are logical, not physical entities and  
4884 that UAC and UAS roles are defined on a request-by-request basis. (Section 6)
- 4885 • In Section 22.42, noted that the *maddr* and *received* parameters also need to be encrypted when  
4886 doing *Via* hiding.
- 4887 • Simplified Fig. 7 to only show *INVITE* transaction.
- 4888 • Added definition of the use of *Contact* (Section 22.10) for *OPTIONS*.
- 4889 • Added HTTP/RFC822 headers *Content-Language* and *MIME-Version*.
- 4890 • Added note in minimal section indicating that UAs need to support UDP.
- 4891 • Added explanation explaining what a UA should do when receiving an initial *INVITE* with a tag.
- 4892 • Clarified UA and proxy behavior for 302 responses.
- 4893 • Added details on what a UAS should do when receiving a tagged *INVITE* request for an unknown call  
4894 leg. This could occur if the UAS had crashed and the UAC sends a re-*INVITE* or if the *BYE* got lost  
4895 and the UAC still believes to be in the call.
- 4896 • Added definition of *Contact* in 4xx, 5xx and 6xx to “redirect” to more error details.
- 4897 • Added note to forking proxy description to gather \*-*Authenticate* from responses. This allows several  
4898 branches to be authenticated simultaneously.
- 4899 • Changed URI syntax to use URL escaping instead of quotation marks.
- 4900 • Changed SIP URL definition to reference RFC 2806 for *telephone-subscriber* part.
- 4901 • Clarified that the *To* URI should basically be ignored by the receiving UAS except for matching  
4902 requests to call legs. In particular, *To* headers with a scheme or name unknown to the callee should  
4903 be accepted.
- 4904 • Clarified that *maddr* is to be added by any client, either proxy or UAC.
- 4905 • Added response code 488 to indicate that there was no common media at the particular destination.  
4906 (606 indicates such failure globally.)

- 4907     • In Section 22.19, noted that registration updates can shorten the validity period.
- 4908     • Added note to enclose the URI for digest in quotation marks. The BNF in RFC 2617 is in error.
- 4909     • Clarified that registrars use Authorization and WWW-Authenticate, not proxy authentication.
- 4910     • Added note in Section 22.10 that “headers” are copied from Contact into the new request.
- 4911     • Changed URL syntax so that port specifications have to have at least one digit, in line with other URL  
4912       formats such as “http”. Previously, an empty port number was permissible.
- 4913     • In SDP section, added a section on how to add and delete streams in re-INVITEs.
- 4914     • IETF-blessed extensions now have short names, without org.ietf. prefix.
- 4915     • Cseq is unique within a call leg, not just within a call (Section 22.16).
- 4916     • Added IPv6 literal addresses to the SIP URL definition, according to RFC 2732 [43]. Modified the  
4917       IPv4 address to limit segments to at most three digits.
- 4918     • modify registration procedure so that it explicitly references the URL comparison. Updates with  
4919       shorter expiration time are now allowed.
- 4920     • For send-only media, SDP still must indicate the address and port, since these are needed as destina-  
4921       tions for RTCP messages.
- 4922     • Changed references regarding DNS SRV records from RFC 2052 to RFC 2782, which is now a Pro-  
4923       posed Standard. Integrated SRV into the search procedure and removed the SRV appendix. The only  
4924       visible change is that protocol and service names are now prefixed by an underscore. Added wording  
4925       that incorporates the precedence of maddr.
- 4926     • Allow parameters in Record-Route and Route headers.
- 4927     • In Table 1, list udp as the default value for the transport parameter in SIP URI.
- 4928     • Removed sentence that From can be encrypted. It cannot, since the header is needed for call-leg  
4929       identification.
- 4930     • Added note that a UAC only copies a To tag into subsequent transactions if it arrives in a 200 OK to  
4931       an INVITE. This avoids the problem that occurs when requests get resubmitted after receiving, say,  
4932       a 407 (or possibly 500, 503, 504, 305, 400, 411, 413, maybe even 408). Under the old rules, these  
4933       requests would have a tag, which would force the called UAS to reject the request, since it doesn't  
4934       have an entry for this tag.
- 4935     • Loop detection has been modified to take the request-URI into account. This allows the same request  
4936       to visit the server twice, but with different request URIs (“spiral”).
- 4937     • Elaborated on URL comparison and comparison of From/To fields.
- 4938     • Added np-queried user parameter.
- 4939     • Changed tag syntax from UUID to token, since there's no reason to restrict it to hex.

- 4940 ● Added Content-Disposition header based on earlier discussions about labeling what to do with a  
4941 message body (part).
- 4942 ● Clarification: proxies must insert To tags for locally generated responses.
- 4943 ● Clarification: multicast may be used for subsequent registrations.
- 4944 ● Feature: Added Supported header. Needed if client wants to indicate things the server can usefully  
4945 return in the response.
- 4946 ● Bug: The From, To, and Via headers were missing extension parameters. The Encryption and  
4947 Response-Key header fields now “officially” allow parameters consisting only of a token, rather  
4948 than just “token = value”.
- 4949 ● Bug: Allow was listed as optional in 405 responses in Table 2. It is mandatory.
- 4950 ● Added: “A BYE request from either called or calling party terminates any pending INVITE, but the  
4951 INVITE request transaction MUST be completed with a final response.”
- 4952 ● Clarified: “If an INVITE request for an existing session fails, the session description agreed upon in  
4953 the last successful INVITE transaction remains in force.”
- 4954 ● Clarified what happens if two INVITE requests meet each other on the wire, either traveling the same  
4955 or in opposite directions:
  - 4956 A UAC MUST NOT issue another INVITE request for the same call leg before the pre-  
4957 vious transaction has completed. A UAS that receives an INVITE before it sent the final  
4958 response to an INVITE with a lower CSeq number MUST return a 400 (Bad Request)  
4959 response and MUST include a Retry-After header field with a randomly chosen value of  
4960 between 0 and 10 seconds. A UA that receives an INVITE while it has an INVITE transac-  
4961 tion pending, returns a 500 (Internal Server Error) and also includes a Retry-After header  
4962 field.
- 4963 ● Expires header clarified: limits only duration of INVITE transaction, not the actual session. SDP  
4964 does the latter.
- 4965 ● The In-Reply-To header was added.
- 4966 ● There were two incompatible BNFs for WWW-Authenticate. One defined for PGP, and the other  
4967 borrowed from HTTP. For basic or digest:

4968 WWW-Authenticate: basic realm="Wallyworld"

4969 and for pgp:

4970 WWW-Authenticate: pgp; realm="Wallyworld"

4971 The latter is incorrect and the semicolon has been removed.

- 4972 • Added rules for **Route** construction from called to calling UA.
- 4973 • We now allow **Accept** and **Accept-Encoding** in **BYE** and **CANCEL** requests. There is no particular  
4974 reason not to allow them, as both requests could theoretically return responses, particularly when  
4975 interworking with other signaling systems.
- 4976 • PGP “pgp-pubalgorithm” allows server to request the desired public-key algorithm.
- 4977 • ABNF rules now describe tokens explicitly rather than by subtraction; explicit character enumeration  
4978 for **CTL**, etc.
- 4979 • Registrars should be careful to check the **Date** header as the expiration time may well be in the past,  
4980 as seen by the client.
- 4981 • **Content-Length** is mandatory; Table 2 erroneously marked it as optional.
- 4982 • **User-Agent** was classified in a syntax definition as a request header rather than a general header.
- 4983 • Clarified ordering of items to be signed and include realm in list.
- 4984 • Allow **Record-Route** in 401 and 484 responses.
- 4985 • Hop-by-hop headers need to precede end-to-end headers only if authentication is used.
- 4986 • 1xx message bodies MAY now contain session descriptions.
- 4987 • Changed references to HTTP/1.1 and authentication to point to the latest RFCs.
- 4988 • Added 487 (Request terminated) status response. It is issued if the original request was terminated  
4989 via **CANCEL** or **BYE**.
- 4990 • The spec was not clear on the identification of a call leg. Section 1.3 says it’s the combination of **To**,  
4991 **From**, and **Call-ID**. However, requests from the callee to the caller have the **To** and **From** reversed, so  
4992 this definition is not quite accurate. Additionally, the “tag” field should be included in the definition  
4993 of call leg. The spec now says that a call leg is defined as the combination of local-address, remote-  
4994 address, and call-id, where these addresses include tags.  
Text was added to Section 6.21 to emphasize that the **From** and **To** headers designate the originator  
4995 of the request, not that of the call leg.  
4996
- 4997 • All URI parameters, except **method**, are allowed in a **Request-URI**. Consequently, also updated the  
4998 description of which parameters are copied from 3xx responses in Sec. 22.10.
- 4999 • The use of CRLF, CR, or LF to terminate lines was confusing. Basically, each header line can be  
5000 terminated by a CR, LF, or CRLF. Furthermore, the end of the headers is signified by a “double  
5001 return”. Simplified to require sending of CRLF, but require senders to receive CR and LF as well and  
5002 only allow CR CR, LF LF in addition to double CRLF as a header-body separator.
- 5003 • Round brackets in **Contact** header were part of the HTTP legacy, and very hard to implement. They  
5004 are also not that useful and were removed.

- 5005 • The spec said that a proxy is a back-to-back UAS/UAC. This is almost, but not quite, true. For  
5006 example, a UAS should insert a tag into a provisional response, but a proxy should not. This was  
5007 clarified.
- 5008 • Section 6.13 in the RFC begins mid-paragraph after the BNF. The following text was misplaced in the  
5009 conversion to ASCII:  
5010       Even if the “display-name” is empty, the “name-addr” form MUST be used if the “addr-  
5011 spec” contains a comma, semicolon or question mark.

## 5012 **28 Changes Made in Version 01**

- 5013 • Uniform syntax specification for semicolon parameters:  
5014       Foo        = "Foo" ":" something \*( ";" foo-param )  
5014       foo-param = "bar" "=" token  
5014                | generic-param
- 5015 • Removed np-queried user parameter since this is now part of a tel URL extension parameter.
- 5016 • In SDP section, noted that if the capabilities intersection is empty, a dummy format list still has to be  
5017 returned due to SDP syntax constraints. Previously, the text had required that no formats be listed.  
5018 (Brian Rosen)
- 5019 • Reorganized tables 2 and 3 to show proxy interaction with headers rather than “end-to-end” or “hop-  
5020 by-hop”.

## 5021 **29 Changes Made in Version 02**

- 5022 • Added “or UAS” in description of received headers in Section 22.42. This makes the response  
5023 algorithm work even if the last IP address in the Via is incorrect.
- 5024 • Tentatively removed restriction that CANCEL requests cannot have Route headers. (Billy Biggs)
- 5025 • Tentatively added Also header for BYE requests, as it is widely implemented and a simple means to  
5026 implement unsupervised call transfer. Subject to removal if there is protest. (Billy Biggs)
- 5027 • If a proxy sends a request by UDP (TCP), the spec did not disallow placing TCP (UDP) in the transport  
5028 parameter of the Via field, which it should. Added a note that the transport protocol actually used is  
5029 included.
- 5030 • No default value for the q parameter in Contact is defined. This is not strictly needed, but is useful for  
5031 consistent behaviors at recursive proxies and at UAC’s. Now 0.5.
- 5032 • Clarified that To and From tag values should be different to simplify request matching when calling  
5033 oneself.
- 5034 • Removed ability to carry multiple requests in a single UDP packet (Section 22.14).

5035 ● Added note that Allow MAY be included in requests, to indicate requestor capabilities for the same  
5036 call ID.

5037 ● Added note to Section 22.17 indicating that registrars MUST include the Date header to accomodate  
5038 UAs that do not have a notion of absolute time.

5039 ● Added note emphasizing that non-SIP URIs are permissible in REGISTER.

5040 ● Rewrote the server lookup section to be more precise and more like pseudo-code, with nesting instead  
5041 of “gotos”.

5042 ● Removed note

5043 Note that the two URLs example.com and example.com:5060, while considered equal,  
5044 may not lead to the same server, as the former causes a DNS SRV lookup, while the latter  
5045 only uses the A record.

5046 since that is no longer the case.

5047 ● Emphasized that proxies have to forward requests with unknown methods.

5048 ● Aligned definition of call leg with URI comparison rules.

5049 ● Required that second branch parameter be globally unique, so that a proxy can distinguish different  
5050 branches in spiral scenarios similar to the following, with record-routing in place:

```
5051           B  ---> P1  -----> P2  -----> P1  -----> A
5052 BYE B    B/1      P1/2 , B/1    P2/3 , P1/2 , B/1    P1/4 , P2/3 , P1/2 , B/1
```

5053 Here, A/1 denotes the Via entry with host A and branch parameter 1. Also, this requires updating the  
5054 definition of isomorphic requests, since the Request-URI is the same for all BYE that are record-  
5055 routed.

5056 ● Removed Via hiding from spec, for the following reasons:

- 5057 – complexity, particularly hidden “gotchas” that surface at various points (as in this instance);
- 5058 – interference with loop detection and debugging;
- 5059 – Unlike HTTP, where via-hiding makes sense since all data is contained in the request or re-  
5060 sponse, Via-hiding in SIP by itself does nothing to hide the caller or callee, as address informa-  
5061 tion is revealed in a number of places:
  - 5062 \* Contact;
  - 5063 \* Route/Record-Route;
  - 5064 \* SDP, including the o= and c= lines;
  - 5065 \* possibly accidental leakage in User-Agent header and Call-ID headers.
- 5066 – Unless this is implemented everywhere, the feature is not likely to be very useful, without the  
5067 sender having any recourse such as “don’t route this request unless you can hide”. It appears  
5068 that almost all existing proxies simply ignore the Hide header.

5069 ● Added Error-Info header field.



### 30 Changes Made in Version 03

5070

- 5071 • Description of **Route** and **Record-Route** moved to separate section, which is new. All UAs must  
5072 now support this mechanism.
- 5073 • Removed status code 411, since it cannot occur (Jonathan Rosenberg, James Jack).
- 5074 • Rewrote **Record-Route** section to reflect new mechanism. In particular, requests from callee to caller  
5075 now use the same path as in the opposite direction, without substituting the **From** header field values.  
5076 The **maddr** parameter is now optional.
- 5077 • Disallowed SIP URLs that only have a password, without a user name. The prototype from RFC 1738  
5078 also doesn't allow this.
- 5079 • Allow registrar to set the expiration time.
- 5080 • **CSeq** (Section 22.16) is counted within a call leg, not a call.
- 5081 • Removed wording that connection closing is equivalent to **CANCEL** or 500. This does not work for  
5082 connections that are used for multiple transactions and has other problems.
- 5083 • Cleaned up **CSeq** section. Removed text about inserting **CSeq** method when it is absent. Clarified  
5084 that **CSeq** increments for all requests, not just **invite**. Clarified that all out of order requests, not  
5085 just out of order **INVITE**, are rejected with a 400 class response. Clarified the meaning of "initial"  
5086 sequence number. Clarified that after a request forks, each 200 OK is a separate call leg, and thus,  
5087 separate **CSeq** space. Clarified that **CSeq** numbers are independent for each direction of a call leg.
- 5088 • Massive reorganization and cleanup of the **SDP** section. Introduced the concept of the offer-answer  
5089 model. Clarified that set of codecs in **m** line are usable all at the same time. Inserted size restriction  
5090 on representation of values in **o** line. Explicitly describe forked media. New media lines for adding  
5091 streams appear at the bottom of the **SDP** (used to say **append**).
- 5092 • Removed **Also**.
- 5093 • Added text to **Require** and **Proxy-Require** sections, making it a **SHOULD** to retry the request without  
5094 the unsupported extension.
- 5095 • Added text to section on 415, saying that UAC **SHOULD** retry the request without the unsupported  
5096 body.
- 5097 • Added text to section on **CANCEL** and **ACK**, clarifying much of the behavior.
- 5098 • Modified **Content-Type** to indicate that it can be present even if the body is empty.
- 5099 • **From** tags mandatory
- 5100 • Old text said that if you hang up before sending an **ACK**, you need not send the **ACK**. That is wrong.  
5101 Text fixed so that an **ACK** is always sent.
- 5102 • Old text said that if you never got a response to an **INVITE**, the UAC should send both an **INVITE** and  
5103 **CANCEL**. This doesn't make sense. Rahter, it should do nothing and consider the call terminated.

- 5104      • Added text that says pending requests are responded to with a 487 if a BYE is received.
- 5105      • Updated section 2.2, so that its clear that **Contact** is not used with BYE.
- 5106      • Clarified Via processing rules. Added text on handling loops when proxies route on headers besides  
5107          the request URI. Added text on handling case when sent-by contains a domain name. Added text to  
5108          6.47 on opening TCP connections to send responses upstream.
- 5109      • Clarified that a 1xx with an unknown xx is not the same as the 100 response.
- 5110      • Removed usage of **Retry-After** in REGISTER.
- 5111      • Clarified usage of persistent connections.
- 5112      • Clarified that servers supporting HTTP basic or digest in rfc2617 **MUST** be backwards compatible  
5113          with RFC 2069.
- 5114      • Clarified that **ACK** contains the same branch ID as the request its acknowledging.
- 5115      • Added definitions for spiral, B2BUA.
- 5116      • Rephrased definitions for UAC, UAS, Call, call-leg, caller, callee, making them more concrete.
- 5117      • URL comparison ignores parameters not present in both URLs only for unknown parameters.
- 5118      • Clarified that \* in **Contact** is used only in REGISTER with Expires header zero. Mentioned \* case  
5119          in section on **Contact** syntax.
- 5120      • Removed text that says a UA can insert a **Contact** in 2xx that indicates the address of a proxy. Not  
5121          likely to work in general.
- 5122      • Removed SDP text about aligning media streams within a media type to handle certain crash and  
5123          restart cases.
- 5124      • Receiving a 481 to a mid-call request terminates that call leg. Agreed upon at IETF 49.
- 5125      • Introduced definition of regular transaction - non-INVITE excepting **ACK** and **CANCEL**.
- 5126      • Clarified rules for overlapping transactions.
- 5127      • Forking proxies **MUST** be stateful (used to say SHOULD). Proxies that send requests on multicast  
5128          **MUST** be stateful (used to say nothing)
- 5129      • Text added recommending that registrars authorize that entity in **From** field can register address-of-  
5130          record in the **To** field.
- 5131      • Forwarding of non-100 provisionals upstream in a proxy changed from SHOULD to MUST.
- 5132      • Removed PGP.

## 5133 31 Changes Made in Version 04

- 5134 • Removed Unsupported as a request header from Table 3.
- 5135 • Clarified SDP procedures for changing IP address and port. Specifically, spelled out the duration for  
5136 which a UA needs to received media on the old port and address.
- 5137 • Added text in the SDP session which recommends that the answerer use the same ordering of codecs  
5138 as used on the offer, in order to help ensure symmetric codec operation under normal conditions.
- 5139 • Fixed bug in the example in the SDP section, where the new media line was listed at the top. Should  
5140 have been the bottom.
- 5141 • Authorization credentials are cached based on the URL of the To header, not the entire To header as  
5142 10.48 implied.
- 5143 • Section 10.31, on Proxy-Authenticate, indicated that a server responds with a 401 if the client  
5144 guessed wrong. This is incorrect. It should be 407.
- 5145 • Section 10.14, removed motivational text about Contact allowing an INVITE to be routed directly  
5146 between end systems, since its confusing. Some have interpreted to mean that Record-Route is  
5147 ignored when Contact is present.
- 5148 • Added reference to SCTP RFC.
- 5149 • Updated 2.2 to allow non-SIP URLs in OPTIONS and 2xx to OPTIONS.
- 5150 • Fixed example in 20.5. Added ACK for 487, and added To tag to 487 response.
- 5151 • Clarified further URL comparisons. Its only URL parameters without defaults that are ignored if not  
5152 present in both URLs.
- 5153 • Section 1.5.2, UDP mandatory for all. TCP is a SHOULD for UA, MUST for proxy, registrar, redirect  
5154 servers.
- 5155 • Brought syntax for Contact, Via, and the SIP URL into alignment between the text and postscript  
5156 versions.
- 5157 • Updated the text in section 6 which said that the ordering of header fields follows HTTP, with the  
5158 exception of Via, where order matters. However, the HTTP spec says that order matters, so this  
5159 sentence is redundant and confusing. The sentence was removed.
- 5160 • Added e lines to SDP examples in the Examples section.
- 5161 • Rewrote Allow discussion, more formally defining its semantics and usage cases.
- 5162 • Updated text on 604 status, to indicate that its based on the Request-URI, not the To.
- 5163 • Added response registrations to IANA considerations. Provided more details on registration process.
- 5164 • Clarified that only a UAS rejects a request because the To tag doesn't match a local value.

- 5165 • Clarified that stateless proxies need to route based on static criteria only.
- 5166 • Proxy and UAC CANCEL generation upon 2xx, 6xx if it forked is now a SHOULD; used to be a MAY.
- 5167 • Added text saying that a UAS SHOULD send a BYE if it never gets an ACK for a 2xx establishing a  
5168 call leg.
- 5169 • Added text saying that a UAS SHOULD send a re-INVITE if it never gets an ACK for a 2xx to a  
5170 re-INVITE.
- 5171 • Added text on 503 processing, indicating that a client should try a different server when receiving a  
5172 503, and that a proxy shouldn't forward a 503 upstream unless it can't service any other requests.
- 5173 • Removed motivational text in Section 10.43 on Via headers since its not consistent with the text before  
5174 it.
- 5175 • Changed IPSec reference to RFC2401, from RFC1825.
- 5176 • Updated retransmission definition in 17.3.4 to be consistent with the rest of the spec.
- 5177 • Softened the language for insertion of the transport param in the record-route. Specifically, it can be  
5178 inserted in private networks where it is known apriori that the specific transport is supported.
- 5179 • Updated definition of B2BUA.
- 5180 • Added text to section on 420 processing, which mandates that the client retry the request without  
5181 extensions listed in the Unsupported header in the response.
- 5182 • Allow Authentication-Info header to be used for HTTP digest.

## 5183 **32 Changes Made in Version 05**

- 5184 • Updated Table 2 to reflect that Error-Info is a response header in 3xx-6xx responses (it was previously  
5185 listed as a request header).
- 5186 • Removed WWW-Authenticate as a request header from Table 3. Authentication of responses is now  
5187 done according to RFC2617.
- 5188 • Updated the Accept, Accept-Encoding and Accept-Language sections. More details on precise  
5189 semantics for the various requests and responses is now provided. Presence of these headers is now  
5190 a SHOULD for INVITE and 2xx to INVITE when a non-default value is present. Extra emphasis is  
5191 placed on including the Accept-Language in INVITE and 2xx in order to support internationaliza-  
5192 tion. Usage of these three headers in CANCEL has been removed since it makes no sense.
- 5193 • Generalized local outbound processing rules in Section 16.4.1 to cover the case where the UAS is  
5194 using a local outbound proxy which was not in the initial call setup path.
- 5195 • Updated record-routing section, so that a proxy can insert a transport param if it knows that the proxy  
5196 on one side supports the specific transport (the previous text required the proxy to know whether the  
5197 proxies on both sides supported the specific transport).

- 5198      • Added Authentication-Info to Section 10.
- 5199      • Clarified the meaning of Table 2 for responses.
- 5200      • Updated Table 1 to reflect that maddr is no longer mandatory in Record-Route.
- 5201      • Updated Table 3 so that header fields in responses to ACK are never listed as optional, mandatory, etc.  
5202        - only not applicable. This is because responses to ACK are not allowed. Also improved wording in  
5203        Section 5.1.1 to clarify that there MUST NOT be responses to ACK.
- 5204      • Updated SRV procedures. Old text said to treat a failure to contact a server as a 4xx, which would  
5205        stop the SRV processing. But, this is not so. Sentence was stricken.
- 5206      • Updated 12.1 to clarify that 2xx INVITE responses MUST contain session descriptions.
- 5207      • Changed User-Agent to a request header in Table 3.
- 5208      • Updated SDP section, so that a UA cannot change the SDP when it gets a re-INVITE with no SDP.
- 5209      • Clarified Appendix B that a unicast offer MUST have a unicast response.
- 5210      • Clarified that any request can be record-routed, but it may not be used by the UA, depending on the  
5211        method.
- 5212      • non-2xx responses to INVITE no longer retransmitted over TCP.
- 5213      • Removed lower bound on T1 and T2 in private networks, which can use lower values. Furthermore,  
5214        T1 can be smaller on the public Internet if proper RTT estimation is used.
- 5215      • UAS Cannot send a BYE for a call leg until it receives ACK, in order to eliminate a race condition  
5216        between BYE and 200 OK.
- 5217      • Support of CR or LF alone as line terminators, as opposed to CRLF, is no longer required.
- 5218      • Client behavior on receipt of a 3xx to re-INVITE is now specified, and it is no longer forbidden to  
5219        generate a 3xx. This is needed to maintain the idempotency of INVITE, as a proxy might redirect  
5220        without knowing its a 3xx.
- 5221      • CANCEL cannot be sent before a 1xx is received, in order to eliminate race condition between request  
5222        and CANCEL.
- 5223      • Termination of the client and server transactions is now based entirely on timeouts, rather than re-  
5224        transmission counters, in order to unify TCP and UDP behavior. Timeout values scale as a function  
5225        of the RTT estimate, defined as T1. For reliable transports, many of these timers are now set to zero.  
5226        Many timeouts differ than in bis-04.
- 5227      • Added a working RTT estimation algorithm using the Timestamp header, and specified it to be  
5228        compliant to RFC 2988.
- 5229      • UAS accepting requests with unknown schemes in the URI in the To field is now a RECOMMENDED  
5230        instead of SHOULD. This reflects the fact that processing a request when the To field doesn't match is  
5231        a matter of policy.

- 5232 ● Bodies are now allowed in any request and response, including CANCEL, although there may not be  
5233 any semantics associated with that.
- 5234 ● Supporting of INVITE without SDP is now a MUST (no strength was previously specified).
- 5235 ● Registration procedures for visiting, which had a few sentences in bis-04, have been removed. Roam-  
5236 ing is a complex issue, and should be treated elsewhere.
- 5237 ● Bis-04 mandated that a 2xx response to REGISTER contain expires Contact parameters indicating  
5238 the expiration time of a contact. This behavior has now been made consistent with requests, so that  
5239 the expiration time of a contact is the same in either case: the expires param is used first if present,  
5240 then the Expires header if present, else one hour for SIP URLs.
- 5241 ● Action parameter in contact registrations is deprecated.
- 5242 ● 2xx to REGISTER MUST contain current contacts. This was just a SHOULD in bis-04.
- 5243 ● Multicast operation radically changed. Now, the treatment is no different than unicast. That is, only  
5244 the first non-1xx response to a multicast request will be used. This is a natural consequence of the  
5245 layering now applied to the protocol. This still enables anycast types of functions, mirroring the real  
5246 usage of registrar discovery.
- 5247 ● To completely separate transport rules from transaction rules, the rule in bis-04 that said a UAC  
5248 SHOULD keep a connection opened until a response is received, has been turned into a timer recom-  
5249 mendation. Specifically, the spec now says that it is RECOMMENDED that connections be kept opened  
5250 for a minimum interval of sufficient duration to guarantee, with high probability, that responses are  
5251 sent over the same connections as a request.
- 5252 ● Re-use of existing connections for new requests to the same address and port is now RECOMMENDED,  
5253 it was only a MAY in bis-04.
- 5254 ● Modification of headers below the Authorization header by proxies is no longer disallowed, since the  
5255 only mechanism that used Authorization in that way, PGP, has been deprecated previously.
- 5256 ● Authentication of registrations now RECOMMENDED; no strength was defined previously.
- 5257 ● Registering of new headers with IANA is now SHOULD; no strength was defined previously.
- 5258 ● Proxy aggregation of challenges now a SHOULD; no strength was defined previously.
- 5259 ● Server support of basic authentication downgraded from SHOULD to MAY.
- 5260 ● UAC resubmitting requests with credentials after a challenge upgraded from MAY to SHOULD.
- 5261 ● TLS is now RECOMMENDED as the transport layer security for SIP signaling.
- 5262 ● UA recursion on a redirect is now SHOULD; no strength was assigned previously.
- 5263 ● UA reuse of headers in a recursed request is now SHOULD; no strength was assigned previously.
- 5264 ● Security considerations added for Call-Info and Alert-Info.

- 5265 • Proxies no longer forward a 6xx immediately on receiving it. Instead, they CANCEL pending  
5266 branches immediately. This avoids a potential race condition that would result in a UAC getting a  
5267 6xx followed by a 2xx. In all cases except this race condition, the result will be the same - the 6xx is  
5268 forwarded upstream.
- 5269 • The term call-leg has been eliminated from the spec; a more generic term, dialog, is used in its place.
- 5270 • For SRV processing, subsequent requests with the same Call-ID (as opposed to the same transaction  
5271 in bis-04) are sent to the same server.
- 5272 • SRV processing generalized to deal with the fact that the default port is transport dependent.
- 5273 • Per IESG request, draft-ietf-sip-serverfeatures has been integrated into bis.
- 5274 • Per IESG request, draft-ietf-sip-100rel will be integrated into bis. This is marked with a placeholder  
5275 in this draft.
- 5276 • The BNF has been converted from implicit LWS to explicit LWS.
- 5277 • Caching of responses in a proxy to avoid redoing location server lookups used to be a SHOULD.  
5278 Caching behavior for responses is now fully encapsulated in the transaction processing.
- 5279 • Proxy usage of SRV in processing Route headers upgraded from SHOULD to MUST.

### 5280 **33 Changes Made in Version 06**

- 5281 • The two states of a dialog are now called early and confirmed.
- 5282 • CANCEL requests now carry Route header fields.
- 5283 • Changes section in -05 forgot to mention the removal of the Encryption and Response-Key headers.  
5284 These were removed since the only mechanism that used them, PGP, had already been deprecated. As  
5285 such, they were effectively “garbage collected”.
- 5286 • Updated error in transaction definition. ACK-2xx is a separate transaction, ACK for non-2xx is part  
5287 of the same transaction.
- 5288 • Changed “Contact-Length” typo to “Content-Length” in various sections, including throughout the  
5289 Examples section.
- 5290 • Changed Table 3 entry for Record-Route and Route for REGISTER from “o” for optional to “-” for  
5291 Not Allowed.
- 5292 • Changed Table 3 entry for Route for ACK, BYE, CANCEL, INVITE, and OPTIONS from “o” for  
5293 optional to “c” for conditional, depending on whether a route set has been defined for the dialog or  
5294 the response code.
- 5295 • Updated Figure 5 - adding missing label on “calling” to “completed” transition.
- 5296 • Fixed errored transport example from Section 19.2.1.

- 5297 • Clarified that 17.2.3 and 17.1.3 are rules that define retransmissions.
- 5298 • fixed reported bugs in bnf (missing productions, bad tex markup), etc. Added new SWS production
- 5299 to have an LWS which allows zero spaces, and used that With any separators. Removed the # rule.
- 5300 • ACK for non-2xx has to have the same Route as the request its acknowledging. The text formerly
- 5301 said that the ACK MUST NOT contain Route, this has now radically changed to MUST have Route if
- 5302 the request its cancelling had one.
- 5303 • Clarified that stateless proxies apply Route processing logic to CANCEL requests.
- 5304 • Emphasized that escaping in the hostname portion of SIP URIs is not currently allowed.
- 5305 • Added discussion on when configuration changes affect the ability of a proxy to forward requests
- 5306 stateful or statelessly.
- 5307 • Explicitly stated that a proxy may add a Record-Route header field value to any request
- 5308 • Added discussion on the use of To tags in hop-hop responses at a proxy
- 5309 • Relaxed text concerning proxies forwarding CANCELs when a matching response context can't be
- 5310 found to allow the CANCEL to be processed statefully.
- 5311 • Changed references to "short" form of SIP headers to "compact" form.
- 5312 • Changed Date example to a valid date.
- 5313 • Clarified how ACK gets from transport to UAS core.
- 5314 • Adding missing "SIP/2.0" to first REGISTER in the examples section.
- 5315 • Fixed bug in 17.2.3 which said that an ACK matched a server transaction if the CSeq method (not
- 5316 number) matched that of the INVITE. It should be the reverse - number, not method.
- 5317 • Fixed bug in 22.15 where it said Content-Length instead of Content-Type.
- 5318 • Incorporated draft-ietf-sip-100rel-04 into bis.
- 5319 • Reliability of provisional responses now only defined for provisional responses to INVITE, although
- 5320 extension methods can allow its usage. This is because PRACK needs to be sent within the context
- 5321 of a dialog, and only responses to INVITE establish dialogs.
- 5322 • Can no longer send a reliable provisional response after a final response; its not compatible with the
- 5323 transaction machines, which generally assume no provisionals after a final.
- 5324 • Proxy behavior for reliable provisional responses no longer defined separately; the spec states that it
- 5325 simply acts as a uas.
- 5326 • Scope of record-route headers for a reliable provisional response is now the dialog rather than the
- 5327 particular request.
- 5328 • Example PRACK flows were lost when incorporating into bis.



- 5329 • Formal IANA registration of “100rel” option tag.
- 5330 • If reliable provisional response gets no PRACK after 32\*T1, UAS sends 5xx to original request.
- 5331 • Recommended UA behavior for caching credentials.
- 5332 • Included guidelines for devices presenting pre-configured credentials vs. prompting end users to  
5333 provide credentials for a specific realm.
- 5334 • Added section on Stateless UAS Behavior, clarifying secure handling of unauthenticated requests to  
5335 prevent potential DoS threat.
- 5336 • Provided motivation for aggregation of challenges in the Security Considerations, and made the be-  
5337 havioral language there more specific.
- 5338 • Provided guidelines for the construction of realm strings for authentication.
- 5339 • Changed concept of protection domain for SIP so that it is no longer defined by both a Request-URI  
5340 and a realm - it is now only defined by a realm.
- 5341 • Reversed opinion on whether the URI parameter in a Digest Authorization header should be quoted  
5342 or not - we now assert that it should NOT be quoted.
- 5343 • Put in some text encouraging UACs not to resubmit rejected credentials when re-challenged.
- 5344 • Added falsification of source IP address to the Via denial of service attack case.
- 5345 • Provided canonical MD5 hash for an empty message body to be used in Digest integrity calculation.
- 5346 • Added security considerations for the CANCEL and ACK methods.
- 5347 • Deprecated and removed Basic auth scheme. Proxies MUST NOT accept or request Basic.
- 5348 • Strengthened language regarding the sending of the “qop” parameter - receipt of cnonce is based on  
5349 “qop”.

## 5350 **34 Acknowledgments**

5351 We wish to thank the members of the IETF MMUSIC and SIP WGs for their comments and suggestions.  
5352 Detailed comments were provided by Brian Bidulock, Jim Buller, Neil Deason, Dave Devanathan, Cdric  
5353 Fluckiger, Yaron Goland, Bernie Hneisen, Phil Hoffer, Christian Huitema, Jean Jervis, Gadi Karmi, Peter  
5354 Kjellerstedt, Anders Kristensen, Jonathan Lennox, Gethin Liddell, Keith Moore, Vern Paxson, Moshe J.  
5355 Sambol, Chip Sharp, Igor Slepchin, Eric Tremblay., and Rick Workman.

5356 Brian Rosen provided the compiled BNF.

5357 This work is based, inter alia, on [44, 45].

**5358 35 Authors' Addresses**

5359 Authors addresses are listed alphabetically for the editors, the writers, and then the original authors of RFC  
5360 2543.

5361 Jonathan Rosenberg  
5362 dynamicsoft  
5363 72 Eagle Rock Ave  
5364 East Hanover, NJ 07936  
5365 USA  
5366 electronic mail: jdrosen@dynamicsoft.com

5367 Henning Schulzrinne  
5368 Dept. of Computer Science  
5369 Columbia University  
5370 1214 Amsterdam Avenue  
5371 New York, NY 10027  
5372 USA  
5373 electronic mail: schulzrinne@cs.columbia.edu

5374 Gonzalo Camarillo  
5375 Ericsson  
5376 Advanced Signalling Research Lab.  
5377 FIN-02420 Jorvas  
5378 Finland  
5379 electronic mail: Gonzalo.Camarillo@ericsson.com

5380 Alan Johnston  
5381 WorldCom  
5382 100 South 4th Street  
5383 St. Louis, MO 63102  
5384 USA  
5385 electronic mail: alan.johnston@wcom.com

5386 Jon Peterson  
5387 NeuStar, Inc  
5388 1800 Sutter Street, Suite 570  
5389 Concord, CA 94520  
5390 USA  
5391 electronic mail: jon.peterson@neustar.com

5392 Robert Sparks  
5393 dynamicsoft, Inc.  
5394 5100 Tennyson Parkway  
5395 Suite 1200  
5396 Plano, Texas 75024  
5397 USA  
5398 electronic mail: rsparks@dynamicsoft.com

5399 Mark Handley  
5400 ACIRI  
5401 electronic mail: mjh@aciri.org

5402 Eve Schooler  
5403 Computer Science Department 256-80  
5404 California Institute of Technology  
5405 Pasadena, CA 91125  
5406 USA  
5407 electronic mail: schooler@cs.caltech.edu

## 5408 **References**

- 5409 [1] R. Pandya, "Emerging mobile and personal communication systems," *IEEE Communications Maga-*  
5410 *zine*, Vol. 33, pp. 44–52, June 1995.
- 5411 [2] R. Braden, Ed., L. Zhang, S. Berson, S. Herzog, and S. Jamin, "Resource ReSerVation protocol  
5412 (RSVP) – version 1 functional specification," Request for Comments 2205, Internet Engineering Task  
5413 Force, Sept. 1997.
- 5414 [3] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, "RTP: a transport protocol for real-time  
5415 applications," Request for Comments 1889, Internet Engineering Task Force, Jan. 1996.
- 5416 [4] H. Schulzrinne, A. Rao, and R. Lanphier, "Real time streaming protocol (RTSP)," Request for Com-  
5417 ments 2326, Internet Engineering Task Force, Apr. 1998.
- 5418 [5] M. Handley, C. Perkins, and E. Whelan, "Session announcement protocol," Request for Comments  
5419 2974, Internet Engineering Task Force, Oct. 2000.
- 5420 [6] M. Handley and V. Jacobson, "SDP: session description protocol," Request for Comments 2327, Inter-  
5421 net Engineering Task Force, Apr. 1998.
- 5422 [7] S. Bradner, "Key words for use in RFCs to indicate requirement levels," Request for Comments 2119,  
5423 Internet Engineering Task Force, Mar. 1997.
- 5424 [8] H. Schulzrinne and J. Rosenberg, "SIP: Session initiation protocol – locating SIP servers," Internet  
5425 Draft, Internet Engineering Task Force, Mar. 2001. Work in progress.
- 5426 [9] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, and T. Berners-Lee, "Hypertext  
5427 transfer protocol – HTTP/1.1," Request for Comments 2616, Internet Engineering Task Force, June  
5428 1999.
- 5429 [10] T. Berners-Lee, R. Fielding, and L. Masinter, "Uniform resource identifiers (URI): generic syntax,"  
5430 Request for Comments 2396, Internet Engineering Task Force, Aug. 1998.
- 5431 [11] T. Berners-Lee, L. Masinter, and M. McCahill, "Uniform resource locators (URL)," Request for Com-  
5432 ments 1738, Internet Engineering Task Force, Dec. 1994.

- 5433 [12] F. Yergeau, "UTF-8, a transformation format of ISO 10646," Request for Comments 2279, Internet  
5434 Engineering Task Force, Jan. 1998.
- 5435 [13] D. Crocker, "Standard for the format of ARPA internet text messages," Request for Comments 822,  
5436 Internet Engineering Task Force, Aug. 1982.
- 5437 [14] A. Vaha-Sipila, "URLs for telephone calls," Request for Comments 2806, Internet Engineering Task  
5438 Force, Apr. 2000.
- 5439 [15] N. Freed and N. Borenstein, "Multipurpose internet mail extensions (MIME) part two: Media types,"  
5440 Request for Comments 2046, Internet Engineering Task Force, Nov. 1996.
- 5441 [16] W. R. Stevens, *TCP/IP illustrated: the protocols*, Vol. 1. Reading, Massachusetts: Addison-Wesley,  
5442 1994.
- 5443 [17] J. C. Mogul and S. E. Deering, "Path MTU discovery," Request for Comments 1191, Internet Engi-  
5444 neering Task Force, Nov. 1990.
- 5445 [18] D. Eastlake, S. Crocker, and J. Schiller, "Randomness recommendations for security," Request for  
5446 Comments 1750, Internet Engineering Task Force, Dec. 1994.
- 5447 [19] P. Hoffman, L. Masinter, and J. Zawinski, "The mailto URL scheme," Request for Comments 2368,  
5448 Internet Engineering Task Force, July 1998.
- 5449 [20] D. Meyer, "Administratively scoped IP multicast," Request for Comments 2365, Internet Engineering  
5450 Task Force, July 1998.
- 5451 [21] E. M. Schooler, "A multicast user directory service for synchronous rendezvous," Master's Thesis CS-  
5452 TR-96-18, Department of Computer Science, California Institute of Technology, Pasadena, California,  
5453 Aug. 1996.
- 5454 [22] S. Donovan, "The SIP INFO method," Request for Comments 2976, Internet Engineering Task Force,  
5455 Oct. 2000.
- 5456 [23] J. Rosenberg and H. Schulzrinne, "An offer/answer model with SDP," Internet Draft, Internet Engi-  
5457 neering Task Force, Oct. 2001. Work in progress.
- 5458 [24] R. Rivest, "The MD5 message-digest algorithm," Request for Comments 1321, Internet Engineering  
5459 Task Force, Apr. 1992.
- 5460 [25] V. Paxson and M. Allman, "Computing TCP's retransmission timer," Request for Comments 2988,  
5461 Internet Engineering Task Force, Nov. 2000.
- 5462 [26] T. Dierks and C. Allen, "The TLS protocol version 1.0," Request for Comments 2246, Internet Engi-  
5463 neering Task Force, Jan. 1999.
- 5464 [27] S. Kent and R. Atkinson, "Security architecture for the internet protocol," Request for Comments 2401,  
5465 Internet Engineering Task Force, Nov. 1998.

- 5466 [28] J. Franks, P. Hallam-Baker, J. Hostetler, S. Lawrence, P. Leach, A. Luotonen, and L. Stewart, "HTTP  
5467 authentication: Basic and digest access authentication," Request for Comments 2617, Internet Engi-  
5468 neering Task Force, June 1999.
- 5469 [29] J. Franks, P. Hallam-Baker, J. Hostetler, P. Leach, A. Luotonen, E. Sink, and L. Stewart, "An exten-  
5470 sion to HTTP : Digest access authentication," Request for Comments 2069, Internet Engineering Task  
5471 Force, Jan. 1997.
- 5472 [30] J. Galvin, S. Murphy, S. Crocker, and N. Freed, "Security multipart for MIME: multipart/signed and  
5473 multipart/encrypted," Request for Comments 1847, Internet Engineering Task Force, Oct. 1995.
- 5474 [31] J. Postel, "User datagram protocol," Request for Comments 768, Internet Engineering Task Force,  
5475 Aug. 1980.
- 5476 [32] J. Postel, "DoD standard transmission control protocol," Request for Comments 761, Internet Engi-  
5477 neering Task Force, Jan. 1980.
- 5478 [33] R. Stewart, Q. Xie, K. Morneault, C. Sharp, H. Schwarzbauer, T. Taylor, I. Rytina, M. Kalla, L. Zhang,  
5479 and V. Paxson, "Stream control transmission protocol," Request for Comments 2960, Internet Engi-  
5480 neering Task Force, Oct. 2000.
- 5481 [34] F. Dawson and T. Howes, "vcard MIME directory profile," Request for Comments 2426, Internet  
5482 Engineering Task Force, Sept. 1998.
- 5483 [35] G. Good, "The LDAP data interchange format (LDIF) - technical specification," Request for Com-  
5484 ments 2849, Internet Engineering Task Force, June 2000.
- 5485 [36] R. Troost and S. Dorner, "Communicating presentation information in internet messages: The content-  
5486 disposition header," Request for Comments 1806, Internet Engineering Task Force, June 1995.
- 5487 [37] R. Braden and Ed, "Requirements for internet hosts - application and support," Request for Comments  
5488 1123, Internet Engineering Task Force, Oct. 1989.
- 5489 [38] J. Palme, "Common internet message headers," Request for Comments 2076, Internet Engineering  
5490 Task Force, Feb. 1997.
- 5491 [39] H. Alvestrand, "IETF policy on character sets and languages," Request for Comments 2277, Internet  
5492 Engineering Task Force, Jan. 1998.
- 5493 [40] A. Johnston, S. Donovan, R. Sparks, C. Cunningham, D. Willis, J. Rosenberg, K. Summers, and  
5494 H. Schulzrinne, "SIP telephony call flow examples," Internet Draft, Internet Engineering Task Force,  
5495 Apr. 2001. Work in progress.
- 5496 [41] D. Crocker, Ed., and P. Overell, "Augmented BNF for syntax specifications: ABNF," Request for  
5497 Comments 2234, Internet Engineering Task Force, Nov. 1997.
- 5498 [42] H. Schulzrinne, "RTP profile for audio and video conferences with minimal control," Request for  
5499 Comments 1890, Internet Engineering Task Force, Jan. 1996.
- 5500 [43] R. Hinden, B. Carpenter, and L. Masinter, "Format for literal IPv6 addresses in URL's," Request for  
5501 Comments 2732, Internet Engineering Task Force, Dec. 1999.

5502 [44] E. M. Schooler, "Case study: multimedia conference control in a packet-switched teleconferencing  
5503 system," *Journal of Internetworking: Research and Experience*, Vol. 4, pp. 99–120, June 1993. ISI  
5504 reprint series ISI/RS-93-359.

5505 [45] H. Schulzrinne, "Personal mobility for multimedia services in the Internet," in *European Workshop on*  
5506 *Interactive Distributed Multimedia Systems and Services (IDMS)*, (Berlin, Germany), Mar. 1996.

## 5507 **Full Copyright Statement**

5508 Copyright (c) The Internet Society (2002). All Rights Reserved.

5509 This document and translations of it may be copied and furnished to others, and derivative works that  
5510 comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and  
5511 distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and  
5512 this paragraph are included on all such copies and derivative works. However, this document itself may not  
5513 be modified in any way, such as by removing the copyright notice or references to the Internet Society or  
5514 other Internet organizations, except as needed for the purpose of developing Internet standards in which case  
5515 the procedures for copyrights defined in the Internet Standards process must be followed, or as required to  
5516 translate it into languages other than English.

5517 The limited permissions granted above are perpetual and will not be revoked by the Internet Society or  
5518 its successors or assigns.

5519 This document and the information contained herein is provided on an "AS IS" basis and THE IN-  
5520 TERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WAR-  
5521 RANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT  
5522 THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED  
5523 WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.