

# 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 that 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 user’s current location, authenticate and authorize users for services, implement provider call-routing policies, and provide features to users. SIP also provides a registration function that allows users to upload their current locations for use by proxy servers. SIP runs on top of several different transport protocols.

## Contents

<b>1</b>	<b>Introduction</b>	<b>8</b>
<b>2</b>	<b>Overview of SIP Functionality</b>	<b>8</b>
<b>3</b>	<b>Terminology</b>	<b>9</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>16</b>
<b>7</b>	<b>SIP Messages</b>	<b>20</b>
7.1	Requests . . . . .	20
7.2	Responses . . . . .	21

34	7.3	Header Fields . . . . .	22
35	7.3.1	Header Field Format . . . . .	22
36	7.3.2	Header Field Classification . . . . .	24
37	7.3.3	Compact Form . . . . .	24
38	7.4	Bodies . . . . .	25
39	7.4.1	Message Body Type . . . . .	25
40	7.4.2	Message Body Length . . . . .	25
41	7.5	Framing SIP messages . . . . .	25
42	<b>8</b>	<b>General User Agent Behavior</b>	<b>25</b>
43	8.1	UAC Behavior . . . . .	26
44	8.1.1	Generating the Request . . . . .	26
45	8.1.2	Sending the Request . . . . .	30
46	8.1.3	Processing Responses . . . . .	30
47	8.2	UAS Behavior . . . . .	33
48	8.2.1	Method Inspection . . . . .	33
49	8.2.2	Header Inspection . . . . .	33
50	8.2.3	Content Processing . . . . .	34
51	8.2.4	Applying Extensions . . . . .	35
52	8.2.5	Processing the Request . . . . .	35
53	8.2.6	Generating the Response . . . . .	35
54	8.2.7	Stateless UAS Behavior . . . . .	36
55	8.3	Redirect Servers . . . . .	36
56	<b>9</b>	<b>Canceling a Request</b>	<b>37</b>
57	9.1	Client Behavior . . . . .	38
58	9.2	Server Behavior . . . . .	39
59	<b>10</b>	<b>Registrations</b>	<b>39</b>
60	10.1	Overview . . . . .	39
61	10.2	Constructing the REGISTER Request . . . . .	41
62	10.2.1	Adding Bindings . . . . .	42
63	10.2.2	Removing Bindings . . . . .	43
64	10.2.3	Fetching Bindings . . . . .	43
65	10.2.4	Refreshing Bindings . . . . .	43
66	10.2.5	Setting the Internal Clock . . . . .	43
67	10.2.6	Discovering a Registrar . . . . .	43
68	10.2.7	Transmitting a Request . . . . .	44
69	10.2.8	Error Responses . . . . .	44
70	10.3	Processing REGISTER Requests . . . . .	44
71	<b>11</b>	<b>Querying for Capabilities</b>	<b>46</b>
72	11.1	Construction of OPTIONS Request . . . . .	47
73	11.2	Processing of OPTIONS Request . . . . .	47
74	<b>12</b>	<b>Dialogs</b>	<b>48</b>

75	12.1	Creation of a Dialog . . . . .	49
76	12.1.1	UAS behavior . . . . .	49
77	12.1.2	UAC Behavior . . . . .	50
78	12.2	Requests within a Dialog . . . . .	50
79	12.2.1	UAC Behavior . . . . .	51
80	12.2.2	UAS Behavior . . . . .	52
81	12.3	Termination of a Dialog . . . . .	53
82	<b>13</b>	<b>Initiating a Session</b>	<b>53</b>
83	13.1	Overview . . . . .	53
84	13.2	UAC Processing . . . . .	54
85	13.2.1	Creating the Initial INVITE . . . . .	54
86	13.2.2	Processing INVITE Responses . . . . .	55
87	13.3	UAS Processing . . . . .	57
88	13.3.1	Processing of the INVITE . . . . .	57
89	<b>14</b>	<b>Modifying an Existing Session</b>	<b>59</b>
90	14.1	UAC Behavior . . . . .	59
91	14.2	UAS Behavior . . . . .	60
92	<b>15</b>	<b>Terminating a Session</b>	<b>61</b>
93	15.1	Terminating a Session with a BYE Request . . . . .	61
94	15.1.1	UAC Behavior . . . . .	61
95	15.1.2	UAS Behavior . . . . .	62
96	<b>16</b>	<b>Proxy Behavior</b>	<b>62</b>
97	16.1	Overview . . . . .	62
98	16.2	Stateful Proxy . . . . .	63
99	16.3	Request Validation . . . . .	63
100	16.4	Route Information Preprocessing . . . . .	66
101	16.5	Determining request targets . . . . .	66
102	16.6	Request Forwarding . . . . .	68
103	16.7	Response Processing . . . . .	73
104	16.8	Processing Timer C . . . . .	78
105	16.9	Handling Transport Errors . . . . .	78
106	16.10	CANCEL Processing . . . . .	78
107	16.11	Stateless Proxy . . . . .	78
108	16.12	Summary of Proxy Route Processing . . . . .	80
109	16.12.1	Examples . . . . .	80
110	<b>17</b>	<b>Transactions</b>	<b>84</b>
111	17.1	Client Transaction . . . . .	85
112	17.1.1	INVITE Client Transaction . . . . .	85
113	17.1.2	Non-INVITE Client Transaction . . . . .	89
114	17.1.3	Matching Responses to Client Transactions . . . . .	91
115	17.1.4	Handling Transport Errors . . . . .	91

116	17.2 Server Transaction . . . . .	91
117	17.2.1 INVITE Server Transaction . . . . .	91
118	17.2.2 Non-INVITE Server Transaction . . . . .	95
119	17.2.3 Matching Requests to Server Transactions . . . . .	95
120	17.2.4 Handling Transport Errors . . . . .	96
121	<b>18 Transport</b>	<b>96</b>
122	18.1 Clients . . . . .	97
123	18.1.1 Sending Requests . . . . .	97
124	18.1.2 Receiving Responses . . . . .	98
125	18.2 Servers . . . . .	98
126	18.2.1 Receiving Requests . . . . .	98
127	18.2.2 Sending Responses . . . . .	99
128	18.3 Framing . . . . .	100
129	18.4 Error Handling . . . . .	100
130	<b>19 Common Message Components</b>	<b>100</b>
131	19.1 SIP and SIPS Uniform Resource Indicators . . . . .	100
132	19.1.1 SIP and SIPS URI Components . . . . .	101
133	19.1.2 Character Escaping Requirements . . . . .	103
134	19.1.3 Example SIP and SIPS URIs . . . . .	104
135	19.1.4 URI Comparison . . . . .	104
136	19.1.5 Forming Requests from a URI . . . . .	106
137	19.1.6 Relating SIP URIs and tel URLs . . . . .	107
138	19.2 Option Tags . . . . .	108
139	19.3 Tags . . . . .	109
140	<b>20 Header Fields</b>	<b>109</b>
141	20.1 Accept . . . . .	112
142	20.2 Accept-Encoding . . . . .	112
143	20.3 Accept-Language . . . . .	113
144	20.4 Alert-Info . . . . .	113
145	20.5 Allow . . . . .	113
146	20.6 Authentication-Info . . . . .	114
147	20.7 Authorization . . . . .	114
148	20.8 Call-ID . . . . .	114
149	20.9 Call-Info . . . . .	114
150	20.10 Contact . . . . .	115
151	20.11 Content-Disposition . . . . .	115
152	20.12 Content-Encoding . . . . .	116
153	20.13 Content-Language . . . . .	116
154	20.14 Content-Length . . . . .	116
155	20.15 Content-Type . . . . .	117
156	20.16 CSeq . . . . .	117
157	20.17 Date . . . . .	117

158	20.18 Error-Info . . . . .	118
159	20.19 Expires . . . . .	118
160	20.20 From . . . . .	118
161	20.21 In-Reply-To . . . . .	119
162	20.22 Max-Forwards . . . . .	119
163	20.23 Min-Expires . . . . .	119
164	20.24 MIME-Version . . . . .	119
165	20.25 Organization . . . . .	119
166	20.26 Priority . . . . .	120
167	20.27 Proxy-Authenticate . . . . .	120
168	20.28 Proxy-Authorization . . . . .	120
169	20.29 Proxy-Require . . . . .	121
170	20.30 Record-Route . . . . .	121
171	20.31 Reply-To . . . . .	121
172	20.32 Require . . . . .	121
173	20.33 Retry-After . . . . .	122
174	20.34 Route . . . . .	122
175	20.35 Server . . . . .	122
176	20.36 Subject . . . . .	122
177	20.37 Supported . . . . .	123
178	20.38 Timestamp . . . . .	123
179	20.39 To . . . . .	123
180	20.40 Unsupported . . . . .	123
181	20.41 User-Agent . . . . .	124
182	20.42 Via . . . . .	124
183	20.43 Warning . . . . .	124
184	20.44 WWW-Authenticate . . . . .	126
185	<b>21 Response Codes</b> . . . . .	<b>126</b>
186	21.1 Provisional 1xx . . . . .	126
187	21.1.1 100 Trying . . . . .	126
188	21.1.2 180 Ringing . . . . .	126
189	21.1.3 181 Call Is Being Forwarded . . . . .	126
190	21.1.4 182 Queued . . . . .	127
191	21.1.5 183 Session Progress . . . . .	127
192	21.2 Successful 2xx . . . . .	127
193	21.2.1 200 OK . . . . .	127
194	21.3 Redirection 3xx . . . . .	127
195	21.3.1 300 Multiple Choices . . . . .	127
196	21.3.2 301 Moved Permanently . . . . .	127
197	21.3.3 302 Moved Temporarily . . . . .	128
198	21.3.4 305 Use Proxy . . . . .	128
199	21.3.5 380 Alternative Service . . . . .	128
200	21.4 Request Failure 4xx . . . . .	128
201	21.4.1 400 Bad Request . . . . .	128

202	21.4.2	401 Unauthorized . . . . .	128
203	21.4.3	402 Payment Required . . . . .	129
204	21.4.4	403 Forbidden . . . . .	129
205	21.4.5	404 Not Found . . . . .	129
206	21.4.6	405 Method Not Allowed . . . . .	129
207	21.4.7	406 Not Acceptable . . . . .	129
208	21.4.8	407 Proxy Authentication Required . . . . .	129
209	21.4.9	408 Request Timeout . . . . .	129
210	21.4.10	410 Gone . . . . .	129
211	21.4.11	413 Request Entity Too Large . . . . .	130
212	21.4.12	414 Request-URI Too Long . . . . .	130
213	21.4.13	415 Unsupported Media Type . . . . .	130
214	21.4.14	416 Unsupported URI Scheme . . . . .	130
215	21.4.15	420 Bad Extension . . . . .	130
216	21.4.16	421 Extension Required . . . . .	130
217	21.4.17	423 Interval Too Brief . . . . .	130
218	21.4.18	480 Temporarily Unavailable . . . . .	131
219	21.4.19	481 Call/Transaction Does Not Exist . . . . .	131
220	21.4.20	482 Loop Detected . . . . .	131
221	21.4.21	483 Too Many Hops . . . . .	131
222	21.4.22	484 Address Incomplete . . . . .	131
223	21.4.23	485 Ambiguous . . . . .	131
224	21.4.24	486 Busy Here . . . . .	132
225	21.4.25	487 Request Terminated . . . . .	132
226	21.4.26	488 Not Acceptable Here . . . . .	132
227	21.4.27	491 Request Pending . . . . .	132
228	21.4.28	493 Undecipherable . . . . .	132
229	21.5	Server Failure 5xx . . . . .	132
230	21.5.1	500 Server Internal Error . . . . .	132
231	21.5.2	501 Not Implemented . . . . .	133
232	21.5.3	502 Bad Gateway . . . . .	133
233	21.5.4	503 Service Unavailable . . . . .	133
234	21.5.5	504 Server Time-out . . . . .	133
235	21.5.6	505 Version Not Supported . . . . .	133
236	21.5.7	513 Message Too Large . . . . .	133
237	21.6	Global Failures 6xx . . . . .	133
238	21.6.1	600 Busy Everywhere . . . . .	134
239	21.6.2	603 Decline . . . . .	134
240	21.6.3	604 Does Not Exist Anywhere . . . . .	134
241	21.6.4	606 Not Acceptable . . . . .	134
242	<b>22</b>	<b>Usage of HTTP Authentication</b>	<b>134</b>
243	22.1	Framework . . . . .	135
244	22.2	User-to-User Authentication . . . . .	136
245	22.3	Proxy-to-User Authentication . . . . .	137

246	22.4 The Digest Authentication Scheme . . . . .	139
247	<b>23 S/MIME</b>	<b>140</b>
248	23.1 S/MIME Certificates . . . . .	140
249	23.2 S/MIME Key Exchange . . . . .	141
250	23.3 Securing MIME bodies . . . . .	142
251	23.4 SIP Header Privacy and Integrity using S/MIME: Tunneling SIP . . . . .	143
252	23.4.1 Integrity and Confidentiality Properties of SIP Headers . . . . .	144
253	23.4.2 Tunneling Integrity and Authentication . . . . .	145
254	23.4.3 Tunneling Encryption . . . . .	146
255	<b>24 Examples</b>	<b>148</b>
256	24.1 Registration . . . . .	148
257	24.2 Session Setup . . . . .	149
258	<b>25 Augmented BNF for the SIP Protocol</b>	<b>154</b>
259	25.1 Basic Rules . . . . .	155
260	<b>26 Security Considerations: Threat Model and Security Usage Recommendations</b>	<b>169</b>
261	26.1 Attacks and Threat Models . . . . .	169
262	26.1.1 Registration Hijacking . . . . .	170
263	26.1.2 Impersonating a Server . . . . .	170
264	26.1.3 Tampering with Message Bodies . . . . .	171
265	26.1.4 Tearing Down Sessions . . . . .	171
266	26.1.5 Denial of Service and Amplification . . . . .	172
267	26.2 Security Mechanisms . . . . .	172
268	26.2.1 Transport and Network Layer Security . . . . .	173
269	26.2.2 SIPS URI scheme . . . . .	173
270	26.2.3 HTTP Authentication . . . . .	174
271	26.2.4 S/MIME . . . . .	174
272	26.3 Implementing Security Mechanisms . . . . .	175
273	26.3.1 Requirements for Implementers of SIP . . . . .	175
274	26.3.2 Security Solutions . . . . .	175
275	26.4 Limitations . . . . .	179
276	26.4.1 HTTP Digest . . . . .	179
277	26.4.2 S/MIME . . . . .	179
278	26.4.3 TLS . . . . .	180
279	26.4.4 SIPS URIs . . . . .	180
280	26.5 Privacy . . . . .	181
281	<b>27 IANA Considerations</b>	<b>181</b>
282	27.1 Option Tags . . . . .	181
283	27.2 Warn-Codes . . . . .	182
284	27.3 Header Field Names . . . . .	182
285	27.4 Method and Response Codes . . . . .	183
286	27.5 The “application/sip” MIME type. . . . .	183

287	<b>28 Changes From RFC 2543</b>	<b>183</b>
288	28.1 Major Functional Changes . . . . .	184
289	28.2 Minor Functional Changes . . . . .	186
290	<b>29 Acknowledgments</b>	<b>187</b>
291	<b>30 Authors' Addresses</b>	<b>187</b>
292	<b>A Table of Timer Values</b>	<b>191</b>

## 293 1 Introduction

294 There are many applications of the Internet that require the creation and management of a session, where  
 295 a session is considered an exchange of data between an association of participants. The implementation of  
 296 these applications is complicated by the practices of participants: users may move between endpoints, they  
 297 may be addressable by multiple names, and they may communicate in several different media - sometimes  
 298 simultaneously. Numerous protocols have been authored that carry various forms of real-time multimedia  
 299 session data such as voice, video, or text messages. SIP works in concert with these protocols by enabling  
 300 Internet endpoints (called *user agents*) to discover one another and to agree on a characterization of a ses-  
 301 sion they would like to share. For locating prospective session participants, and for other functions, SIP  
 302 enables creation of an infrastructure of network hosts (called *proxy servers*) to which user agents can send  
 303 registrations, invitations to sessions, and other requests. SIP is an agile, general-purpose tool for creating,  
 304 modifying, and terminating sessions that works independently of underlying transport protocols and without  
 305 dependency on the type of session that is being established.

## 306 2 Overview of SIP Functionality

307 SIP is an application-layer control protocol that can establish, modify, and terminate multimedia sessions  
 308 (conferences) such as Internet telephony calls. SIP can also invite participants to already existing sessions,  
 309 such as multicast conferences. Media can be added to (and removed from) an existing session. SIP trans-  
 310 parently supports name mapping and redirection services, which supports *personal mobility* [26] - users can  
 311 maintain a single externally visible identifier regardless of their network location.

312 SIP supports five facets of establishing and terminating multimedia communications:

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

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

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

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

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

319 SIP is not a vertically integrated communications system. SIP is rather a component that can be used with  
 320 other IETF protocols to build a complete multimedia architecture. Typically, these architectures will include

321 protocols such as the real-time transport protocol (RTP) (RFC 1889 [27]) for transporting real-time data and  
322 providing QoS feedback, the real-time streaming protocol (RTSP) (RFC 2326 [28]) for controlling delivery  
323 of streaming media, the Media Gateway Control Protocol (MEGACO) (RFC 3015 [29]) for controlling  
324 gateways to the Public Switched Telephone Network (PSTN), and the session description protocol (SDP)  
325 (RFC 2327 [1]) for describing multimedia sessions. Therefore, SIP should be used in conjunction with other  
326 protocols in order to provide complete services to the users. However, the basic functionality and operation  
327 of SIP does not depend on any of these protocols.

328 SIP does not provide services. SIP rather provides primitives that can be used to implement different  
329 services. For example, SIP can locate a user and deliver an opaque object to his current location. If this  
330 primitive is used to deliver a session description written in SDP, for instance, the endpoints can agree on the  
331 parameters of a session. If the same primitive is used to deliver a photo of the caller as well as the session  
332 description, a "caller ID" service can be easily implemented. As this example shows, a single primitive is  
333 typically used to provide several different services.

334 SIP does not offer conference control services such as floor control or voting and does not prescribe how  
335 a conference is to be managed. SIP can be used to initiate a session that uses some other conference control  
336 protocol. Since SIP messages and the sessions they establish can pass through entirely different networks,  
337 SIP cannot, and does not, provide any kind of network resource reservation capabilities.

338 The nature of the services provided make security particularly important. To that end, SIP provides a  
339 suite of security services, which include denial-of-service prevention, authentication (both user to user and  
340 proxy to user), integrity protection, and encryption and privacy services.

341 SIP works with both IPv4 and IPv6.

### 342 **3 Terminology**

343 In this document, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD",  
344 "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" are to be inter-  
345 preted as described in RFC 2119 [2] and indicate requirement levels for compliant SIP implementations.

### 346 **4 Overview of Operation**

347 This section introduces the basic operations of SIP using simple examples. This section is tutorial in nature  
348 and does not contain any normative statements.

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

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

358 Alice "calls" Bob using his SIP identity, a type of Uniform Resource Identifier (URI) called a *SIP*  
359 *URI* and which is defined in Section 19.1. It has a similar form to an email address, typically containing  
360 a username and a host name. In this case, it is sip:bob@biloxi.com, where biloxi.com is the domain of

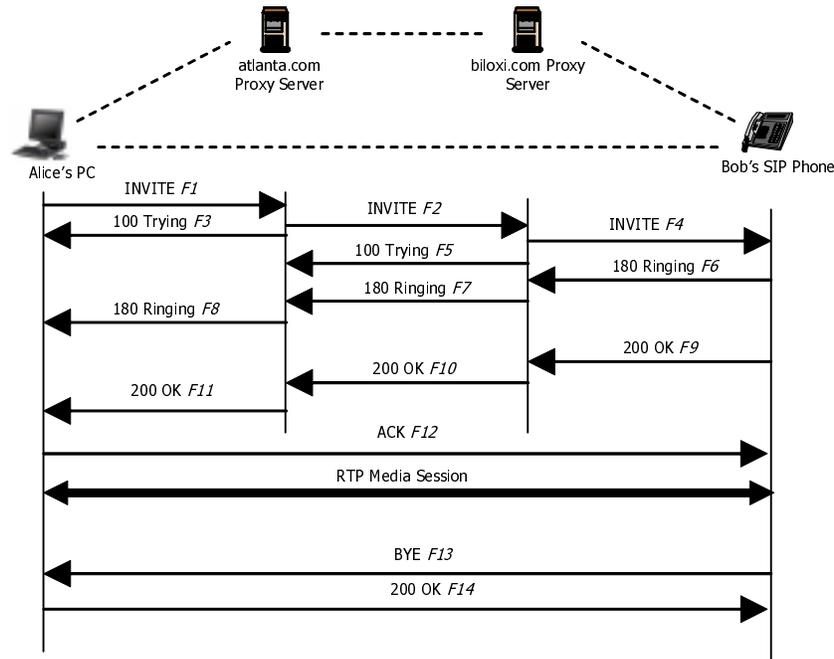


Figure 1: SIP session setup example with SIP trapezoid

361 Bob's SIP service provider (which can be an enterprise, retail provider, etc). Alice also has a SIP URI  
 362 of sip:alice@atlanta.com. Alice might have typed in Bob's URI or perhaps clicked on a hyperlink or  
 363 an entry in an address book. SIP also provides a secure URI, called a SIPS URI. An example would be  
 364 sips:bob@biloxi.com. A call made to a SIPS URI guarantees that secure, encrypted transport (namely TLS)  
 365 is used to carry all SIP messages at every hop between the caller and callee.

366 SIP is based on an HTTP-like request/response transaction model. Each transaction consists of a request  
 367 that invokes a particular *method*, or function, on the server and at least one response. In this example, the  
 368 transaction begins with Alice's softphone sending an INVITE request addressed to Bob's SIP URI. INVITE  
 369 is an example of a SIP method that specifies the action that the requestor (Alice) wants the server (Bob)  
 370 to take. The INVITE request contains a number of header fields. Header fields are named attributes that  
 371 provide additional information about a message. The ones present in an INVITE include a unique identifier  
 372 for the call, the destination address, Alice's address, and information about the type of session that Alice  
 373 wishes to establish with Bob. The INVITE (message F1 in Figure 1) might look like this:

```

374 INVITE sip:bob@biloxi.com SIP/2.0
375 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bK776asdhds
376 Max-Forwards: 70
377 To: Bob <sip:bob@biloxi.com>
378 From: Alice <sip:alice@atlanta.com>;tag=1928301774
379 Call-ID: a84b4c76e66710@pc33.atlanta.com
380 CSeq: 314159 INVITE
381 Contact: <sip:alice@pc33.atlanta.com>
382 Content-Type: application/sdp

```

383 Content-Length: 142  
384  
385 (Alice's SDP not shown)

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

389 Via contains the address (pc33.atlanta.com) at which Alice is expecting to receive responses to this  
390 request. It also contains a branch parameter that contains an identifier for this transaction.

391 To contains a display name (Bob) and a SIP or SIPS URI (sip:bob@biloxi.com) towards which the  
392 request was originally directed. Display names are described in RFC 2822 [3].

393 From also contains a display name (Alice) and a SIP or SIPS URI (sip:alice@atlanta.com) that indicate  
394 the originator of the request. This header field also has a tag parameter containing a pseudorandom string  
395 (1928301774) that was added to the URI by the softphone. It is used for identification purposes.

396 Call-ID contains a globally unique identifier for this call, generated by the combination of a pseudoran-  
397 dom string and the softphone's IP address. The combination of the To tag, From tag, and Call-ID completely  
398 define a peer-to-peer SIP relationship between Alice and Bob and is referred to as a *dialog*.

399 CSeq or Command Sequence contains an integer and a method name. The CSeq number is incremented  
400 for each new request within a dialog and is a traditional sequence number.

401 Contact contains a SIP or SIPS URI that represents a direct route to contact Alice, usually composed  
402 of a username at a fully qualified domain name (FQDN). While an FQDN is preferred, many end systems  
403 do not have registered domain names, so IP addresses are permitted. While the Via header field tells other  
404 elements where to send the response, the Contact header field tells other elements where to send future  
405 requests.

406 Max-Forwards serves to limit the number of hops a request can make on the way to its destination. It  
407 consists of an integer that is decremented by one at each hop.

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

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

410 The complete set of SIP header fields is defined in Section 20.

411 The details of the session, type of media, codec, sampling rate, etc. are not described using SIP. Rather,  
412 the body of a SIP message contains a description of the session, encoded in some other protocol format.  
413 One such format is Session Description Protocol (SDP) [1]. This SDP message (not shown in the example)  
414 is carried by the SIP message in a way that is analogous to a document attachment being carried by an email  
415 message, or a web page being carried in an HTTP message.

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

420 The atlanta.com SIP server is a type of SIP server known as a proxy server. A proxy server receives  
421 SIP requests and forwards them on behalf of the requestor. In this example, the proxy server receives the  
422 INVITE request and sends a 100 (Trying) response back to Alice's softphone. The 100 (Trying) response  
423 indicates that the INVITE has been received and that the proxy is working on her behalf to route the INVITE  
424 to the destination. Responses in SIP use a three-digit code followed by a descriptive phrase. This response  
425 contains the same To, From, Call-ID, CSeq and branch parameter in the Via as the INVITE, which allows  
426 Alice's softphone to correlate this response to the sent INVITE. The atlanta.com proxy server locates the

427 proxy server at biloxi.com, possibly by performing a particular type of DNS (Domain Name Service) lookup  
428 to find the SIP server that serves the biloxi.com domain. This is described in [4]. As a result, it obtains  
429 the IP address of the biloxi.com proxy server and forwards, or proxies, the INVITE request there. Before  
430 forwarding the request, the atlanta.com proxy server adds an additional Via header field that contains its own  
431 address (the INVITE already contains Alice's address in the first Via). The biloxi.com proxy server receives  
432 the INVITE and responds with a 100 (Trying) response back to the atlanta.com proxy server to indicate that  
433 it has received the INVITE and is processing the request. The proxy server consults a database, generically  
434 called a location service, that contains the current IP address of Bob. (We shall see in the next section how  
435 this database can be populated.) The biloxi.com proxy server adds another Via header field value with its  
436 own address to the INVITE and proxies it to Bob's SIP phone.

437 Bob's SIP phone receives the INVITE and alerts Bob to the incoming call from Alice so that Bob can  
438 decide whether to answer the call, that is, Bob's phone rings. Bob's SIP phone indicates this in a 180  
439 (Ringing) response, which is routed back through the two proxies in the reverse direction. Each proxy uses  
440 the Via header field to determine where to send the response and removes its own address from the top.  
441 As a result, although DNS and location service lookups were required to route the initial INVITE, the 180  
442 (Ringing) response can be returned to the caller without lookups or without state being maintained in the  
443 proxies. This also has the desirable property that each proxy that sees the INVITE will also see all responses  
444 to the INVITE.

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

447 In this example, Bob decides to answer the call. When he picks up the handset, his SIP phone sends a  
448 200 (OK) response to indicate that the call has been answered. The 200 (OK) contains a message body with  
449 the SDP media description of the type of session that Bob is willing to establish with Alice. As a result, there  
450 is a two-phase exchange of SDP messages: Alice sent one to Bob, and Bob sent one back to Alice. This  
451 two-phase exchange provides basic negotiation capabilities and is based on a simple offer/answer model of  
452 SDP exchange. If Bob did not wish to answer the call or was busy on another call, an error response would  
453 have been sent instead of the 200 (OK), which would have resulted in no media session being established.  
454 The complete list of SIP response codes is in Section 21. The 200 (OK) (message F9 in Figure 1) might  
455 look like this as Bob sends it out:

```
456 SIP/2.0 200 OK
457 Via: SIP/2.0/UDP server10.biloxi.com;branch=z9hG4bKnashds8
458     ;received=10.2.1.1
459 Via: SIP/2.0/UDP bigbox3.site3.atlanta.com;branch=z9hG4bK77ef4c2312983.1
460     ;received=10.1.1.1
461 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bK776asdhds
462     ;received=10.1.3.3
463 To: Bob <sip:bob@biloxi.com>;tag=a6c85cf
464 From: Alice <sip:alice@atlanta.com>;tag=1928301774
465 Call-ID: a84b4c76e66710
466 CSeq: 314159 INVITE
467 Contact: <sip:bob@192.0.2.8>
468 Content-Type: application/sdp
469 Content-Length: 131
470
```

471 (Bob's SDP not shown)

472 The first line of the response contains the response code (200) and the reason phrase (OK). The remain-  
473 ing lines contain header fields. The Via, To, From, Call-ID, and CSeq header fields are copied from the  
474 INVITE request. (There are three Via header field values - one added by Alice's SIP phone, one added by  
475 the atlanta.com proxy, and one added by the biloxi.com proxy.) Bob's SIP phone has added a tag parameter  
476 to the To header field. This tag will be incorporated by both endpoints into the dialog and will be included  
477 in all future requests and responses in this call. The Contact header field contains a URI at which Bob can  
478 be directly reached at his SIP phone. The Content-Type and Content-Length refer to the message body  
479 (not shown) that contains Bob's SDP media information.

480 In addition to DNS and location service lookups shown in this example, proxy servers can make flexible  
481 "routing decisions" to decide where to send a request. For example, if Bob's SIP phone returned a 486 (Busy  
482 Here) response, the biloxi.com proxy server could proxy the INVITE to Bob's voicemail server. A proxy  
483 server can also send an INVITE to a number of locations at the same time. This type of parallel search is  
484 known as *forking*.

485 In this case, the 200 (OK) is routed back through the two proxies and is received by Alice's softphone,  
486 which then stops the ringback tone and indicates that the call has been answered. Finally, Alice's softphone  
487 sends an acknowledgement message, ACK to Bob's SIP phone to confirm the reception of the final response  
488 (200 (OK)). In this example, the ACK is sent directly from Alice's softphone to Bob's SIP phone, bypassing  
489 the two proxies. This occurs because the endpoints have learned each other's address from the Contact  
490 header fields through the INVITE/200 (OK) exchange, which was not known when the initial INVITE was  
491 sent. The lookups performed by the two proxies are no longer needed, so the proxies drop out of the call  
492 flow. This completes the INVITE/200/ACK three-way handshake used to establish SIP sessions. Full details  
493 on session setup are in Section 13.

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

497 During the session, either Alice or Bob may decide to change the characteristics of the media session.  
498 This is accomplished by sending a re-INVITE containing a new media description. This re-INVITE refer-  
499 ences the existing dialog so that the other party knows that it is to modify an existing session instead of  
500 establishing a new session. The other party sends a 200 (OK) to accept the change. The requestor responds  
501 to the 200 (OK) with an ACK. If the other party does not accept the change, he sends an error response such  
502 as 406 (Not Acceptable), which also receives an ACK. However, the failure of the re-INVITE does not cause  
503 the existing call to fail - the session continues using the previously negotiated characteristics. Full details on  
504 session modification are in Section 14.

505 At the end of the call, Bob disconnects (hangs up) first and generates a BYE message. This BYE is  
506 routed directly to Alice's softphone, again bypassing the proxies. Alice confirms receipt of the BYE with a  
507 200 (OK) response, which terminates the session and the BYE transaction. No ACK is sent - an ACK is only  
508 sent in response to a response to an INVITE request. The reasons for this special handling for INVITE will  
509 be discussed later, but relate to the reliability mechanisms in SIP, the length of time it can take for a ringing  
510 phone to be answered, and forking. For this reason, request handling in SIP is often classified as either  
511 INVITE or non-INVITE, referring to all other methods besides INVITE. Full details on session termination  
512 are in Section 15.

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

514 In some cases, it may be useful for proxies in the SIP signaling path to see all the messaging between the

515 endpoints for the duration of the session. For example, if the biloxi.com proxy server wished to remain in the  
516 SIP messaging path beyond the initial INVITE, it would add to the INVITE a required routing header field  
517 known as **Record-Route** that contained a URI resolving to the hostname or IP address of the proxy. This  
518 information would be received by both Bob's SIP phone and (due to the **Record-Route** header field being  
519 passed back in the 200 (OK)) Alice's softphone and stored for the duration of the dialog. The biloxi.com  
520 proxy server would then receive and proxy the ACK, BYE, and 200 (OK) to the BYE. Each proxy can  
521 independently decide to receive subsequent messaging, and that messaging will go through all proxies that  
522 elect to receive it. This capability is frequently used for proxies that are providing mid-call features.

523 Registration is another common operation in SIP. Registration is one way that the biloxi.com server  
524 can learn the current location of Bob. Upon initialization, and at periodic intervals, Bob's SIP phone sends  
525 REGISTER messages to a server in the biloxi.com domain known as a SIP registrar. The REGISTER mes-  
526 sages associate Bob's SIP or SIPS URI (sip:bob@biloxi.com) with the machine into which he is currently  
527 logged (conveyed as a SIP or SIPS URI in the **Contact** header field). The registrar writes this association,  
528 also called a binding, to a database, called the *location service*, where it can be used by the proxy in the  
529 biloxi.com domain. Often, a registrar server for a domain is co-located with the proxy for that domain. It is  
530 an important concept that the distinction between types of SIP servers is logical, not physical.

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

535 The location service is just an abstract concept. It generally contains information that allows a proxy to  
536 input a URI and receive a set of zero or more URIs that tell the proxy where to send the request. Registrations  
537 are one way to create this information, but not the only way. Arbitrary mapping functions can be configured  
538 at the discretion of the administrator.

539 Finally, it is important to note that in SIP, registration is used for routing incoming SIP requests and  
540 has no role in authorizing outgoing requests. Authorization and authentication are handled in SIP either  
541 on a request-by-request basis with a challenge/response mechanism, or by using a lower layer scheme as  
542 discussed in Section 26.

543 The complete set of SIP message details for this registration example is in Section 24.1.

544 Additional operations in SIP, such as querying for the capabilities of a SIP server or client using **OP-**  
545 **TIONS**, or canceling a pending request using **CANCEL**, will be introduced in later sections.

## 546 **5 Structure of the Protocol**

547 SIP is structured as a layered protocol, which means that its behavior is described in terms of a set of fairly  
548 independent processing stages with only a loose coupling between each stage. The protocol behavior is  
549 described as layers for the purpose of presentation, allowing the description of functions common across  
550 elements in a single section. It does not dictate an implementation in any way. When we say that an element  
551 "contains" a layer, we mean it is compliant to the set of rules defined by that layer.

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

555 The lowest layer of SIP is its syntax and encoding. Its encoding is specified using an augmented Backus-  
556 Naur Form grammar (BNF). The complete BNF is specified in Section 25; an overview of a SIP message's  
557 structure can be found in Section 7.

558 The second layer is the transport layer. It defines how a client sends requests and receives responses and  
559 how a server receives requests and sends responses over the network. All SIP elements contain a transport  
560 layer. The transport layer is described in Section 18.

561 The third layer is the transaction layer. Transactions are a fundamental component of SIP. A transaction  
562 is a request sent by a client transaction (using the transport layer) to a server transaction, along with all  
563 responses to that request sent from the server transaction back to the client. The transaction layer handles  
564 application-layer retransmissions, matching of responses to requests, and application-layer timeouts. Any  
565 task that a user agent client (UAC) accomplishes takes place using a series of transactions. Discussion of  
566 transactions can be found in Section 17. User agents contain a transaction layer, as do stateful proxies.  
567 Stateless proxies do not contain a transaction layer. The transaction layer has a client component (referred  
568 to as a client transaction) and a server component (referred to as a server transaction), each of which are  
569 represented by a finite state machine that is constructed to process a particular request.

570 The layer above the transaction layer is called the transaction user (TU). Each of the SIP entities, except  
571 the stateless proxy, is a transaction user. When a TU wishes to send a request, it creates a client transaction  
572 instance and passes it the request along with the destination IP address, port, and transport to which to send  
573 the request. A TU that creates a client transaction can also cancel it. When a client cancels a transaction,  
574 it requests that the server stop further processing, revert to the state that existed before the transaction was  
575 initiated, and generate a specific error response to that transaction. This is done with a CANCEL request,  
576 which constitutes its own transaction, but references the transaction to be cancelled (Section 9).

577 The SIP elements, that is, user agent clients and servers, stateless and stateful proxies and registrars,  
578 contain a *core* that distinguishes them from each other. Cores, except for the stateless proxy, are transaction  
579 users. While the behavior of the UAC and UAS cores depends on the method, there are some common rules  
580 for all methods (Section 8). For a UAC, these rules govern the construction of a request; for a UAS, they  
581 govern the processing of a request and generating a response. Since registrations play an important role in  
582 SIP, a UAS that handles a REGISTER is given the special name registrar. Section 10 describes UAC and  
583 UAS core behavior for the REGISTER method. Section 11 describes UAC and UAS core behavior for the  
584 OPTIONS method, used for determining the capabilities of a UA.

585 Certain other requests are sent within a dialog. A dialog is a peer-to-peer SIP relationship between two  
586 user agents that persists for some time. The dialog facilitates sequencing of messages and proper routing  
587 of requests between the user agents. The INVITE method is the only way defined in this specification to  
588 establish a dialog. When a UAC sends a request that is within the context of a dialog, it follows the common  
589 UAC rules as discussed in Section 8 but also the rules for mid-dialog requests. Section 12 discusses dialogs  
590 and presents the procedures for their construction and maintenance, in addition to construction of requests  
591 within a dialog.

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

597 The procedures of Sections 8, 10, 11, 12, 13, 14, and 15 deal entirely with the UA core (Section 9  
598 describes cancellation, which applies to both UA core and proxy core). Section 16 discusses the proxy  
599 element, which facilitates routing of messages between user agents.

## 6 Definitions

This specification uses a number of terms to refer to the roles played by participants in SIP communications. The terms and generic syntax of URI and URL are defined in RFC 2396 [5]. The following terms have special significance for SIP.

**Address-of-Record:** An address-of-record (AOR) is a SIP or SIPS URI that points to a domain with a location service that can map the URI to another URI where the user might be available. Typically, the location service is populated through registrations. An AOR is frequently thought of as the “public address” of the user.

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

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

**Call Leg:** Another name for a dialog [30]; no longer used in this specification.

**Call Stateful:** A proxy is call stateful if it retains state for a dialog from the initiating INVITE to the terminating BYE request. A call stateful proxy is always transaction stateful, but the converse is not necessarily true.

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

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

**Core:** Core designates the functions specific to a particular type of SIP entity, i.e., specific to either a stateful or stateless proxy, a user agent or registrar. All cores except those for the stateless proxy are transaction users.

**Dialog:** A dialog is a peer-to-peer SIP relationship between two UAs that persists for some time. A dialog is established by SIP messages, such as a 2xx response to an INVITE request. A dialog is identified by a call identifier, local address, and remote address. A dialog was formerly known as a call leg in RFC 2543.

**Downstream:** A direction of message forwarding within a transaction that refers to the direction that requests flow from the user agent client to user agent server.

**Final Response:** A response that terminates a SIP transaction, as opposed to a *provisional response* that does not. All 2xx, 3xx, 4xx, 5xx and 6xx responses are final.

**Header:** A header is a component of a SIP message that conveys information about the message. It is structured as a sequence of header fields.

635 **Header field:** A header field is a component of the SIP message header. It consists of one or more header  
636 field values separated by comma or having the same header field name.

637 **Header field value:** A header field value consists of a field name and a field value, separated by a colon.

638 **Home Domain:** The domain providing service to a SIP user. Typically, this is the domain present in the  
639 URI in the address-of-record of a registration.

640 **Informational Response:** Same as a provisional response.

641 **Initiator, Calling Party, Caller:** The party initiating a session (and dialog) with an INVITE request. A  
642 caller retains this role from the time it sends the initial INVITE that established a dialog until the  
643 termination of that dialog.

644 **Invitation:** An INVITE request.

645 **Invitee, Invited User, Called Party, Callee:** The party that receives an INVITE request for the purposes of  
646 establishing a new session. A callee retains this role from the time it receives the INVITE until the  
647 termination of the dialog established by that INVITE.

648 **Location Service:** A location service is used by a SIP redirect or proxy server to obtain information about  
649 a callee's possible location(s). It contains a list of bindings of address-of-record keys to zero or more  
650 contact addresses. The bindings can be created and removed in many ways; this specification defines  
651 a REGISTER method that updates the bindings.

652 **Loop:** A request that arrives at a proxy, is forwarded, and later arrives back at the same proxy. When it  
653 arrives the second time, its Request-URI is identical to the first time, and other header fields that  
654 affect proxy operation are unchanged, so that the proxy would make the same processing decision on  
655 the request it made the first time. Looped requests are errors, and the procedures for detecting them  
656 and handling them are described by the protocol.

657 **Loose Routing:** A proxy is said to be loose routing if it follows the procedures defined in this specification  
658 for processing of the Route header field. These procedures separate the destination of the request  
659 (present in the Request-URI) from the set of proxies that need to be visited along the way (present  
660 in the Route header field). A proxy compliant to these mechanisms is also known as a loose router.

661 **Message:** Data sent between SIP elements as part of the protocol. SIP messages are either requests or  
662 responses.

663 **Method:** The method is the primary function that a request is meant to invoke on a server. The method is  
664 carried in the request message itself. Example methods are INVITE and BYE.

665 **Outbound Proxy:** A *proxy* that receives requests from a client, even though it may not be the server re-  
666 solved by the Request-URI. Typically, a UA is manually configured with an outbound proxy, or can  
667 learn about one through auto-configuration protocols.

668 **Parallel Search:** In a parallel search, a proxy issues several requests to possible user locations upon re-  
669 ceiving an incoming request. Rather than issuing one request and then waiting for the final response  
670 before issuing the next request as in a *sequential search*, a parallel search issues requests without  
671 waiting for the result of previous requests.

- 672 **Provisional Response:** A response used by the server to indicate progress, but that does not terminate a SIP  
673 transaction. 1xx responses are provisional, other responses are considered *final*. Provisional responses  
674 are not sent reliably.
- 675 **Proxy, Proxy Server:** An intermediary entity that acts as both a server and a client for the purpose of  
676 making requests on behalf of other clients. A proxy server primarily plays the role of routing, which  
677 means its job is to ensure that a request is sent to another entity “closer” to the targeted user. Proxies  
678 are also useful for enforcing policy (for example, making sure a user is allowed to make a call). A  
679 proxy interprets, and, if necessary, rewrites specific parts of a request message before forwarding it.
- 680 **Recursion:** A client recurses on a 3xx response when it generates a new request to one or more of the URIs  
681 in the Contact header field in the response.
- 682 **Redirect Server:** A redirect server is a user agent server that generates 3xx responses to requests it receives,  
683 directing the client to contact an alternate set of URIs.
- 684 **Registrar:** A registrar is a server that accepts REGISTER requests and places the information it receives  
685 in those requests into the location service for the domain it handles.
- 686 **Regular Transaction:** A regular transaction is any transaction with a method other than INVITE, ACK, or  
687 CANCEL.
- 688 **Request:** A SIP message sent from a client to a server, for the purpose of invoking a particular operation.
- 689 **Response:** A SIP message sent from a server to a client, for indicating the status of a request sent from the  
690 client to the server.
- 691 **Ringback:** Ringback is the signaling tone produced by the calling party’s application indicating that a  
692 called party is being alerted (ringing).
- 693 **Route Set:** A route set is a collection of ordered SIP or SIPS URI which represent a list of proxies that  
694 must be traversed when sending a particular request. A route set can be learned, through headers like  
695 Record-Route, or it can be configured.
- 696 **Server:** A server is a network element that receives requests in order to service them and sends back re-  
697 sponses to those requests. Examples of servers are proxies, user agent servers, redirect servers, and  
698 registrars.
- 699 **Sequential Search:** In a sequential search, a proxy server attempts each contact address in sequence, pro-  
700 ceeding to the next one only after the previous has generated a final response. A 2xx or 6xx class final  
701 response always terminates a sequential search.
- 702 **Session:** From the SDP specification: “A multimedia session is a set of multimedia senders and receivers  
703 and the data streams flowing from senders to receivers. A multimedia conference is an example of a  
704 multimedia session.” (RFC 2327 [1]) (A session as defined for SDP can comprise one or more RTP  
705 sessions.) As defined, a callee can be invited several times, by different calls, to the same session. If  
706 SDP is used, a session is defined by the concatenation of the *SDP user name*, *session id*, *network type*,  
707 *address type*, and *address* elements in the origin field.

- 708 **SIP Transaction:** A SIP transaction occurs between a client and a server and comprises all messages from  
709 the first request sent from the client to the server up to a final (non-1xx) response sent from the server  
710 to the client. If the request is INVITE and the final response is a non-2xx, the transaction also includes  
711 an ACK to the response. The ACK for a 2xx response to an INVITE request is a separate transaction.
- 712 **Spiral:** A spiral is a SIP request that is routed to a proxy, forwarded onwards, and arrives once again at that  
713 proxy, but this time differs in a way that will result in a different processing decision than the original  
714 request. Typically, this means that the request's Request-URI differs from its previous arrival. A  
715 spiral is not an error condition, unlike a loop. A typical cause for this is call forwarding. A user calls  
716 joe@example.com. The example.com proxy forwards it to Joe's PC, which in turn, forwards it to  
717 bob@example.com. This request is proxied back to the example.com proxy. However, this is not a  
718 loop. Since the request is targeted at a different user, it is considered a spiral, and is a valid condition.
- 719 **Stateful Proxy:** A logical entity that maintains the client and server transaction state machines defined by  
720 this specification during the processing of a request. Also known as a transaction stateful proxy. The  
721 behavior of a stateful proxy is further defined in Section 16. A (transaction) stateful proxy is not the  
722 same as a call stateful proxy.
- 723 **Stateless Proxy:** A logical entity that does not maintain the client or server transaction state machines  
724 defined in this specification when it processes requests. A stateless proxy forwards every request it  
725 receives downstream and every response it receives upstream.
- 726 **Strict Routing:** A proxy is said to be strict routing if it follows the Route processing rules of RFC 2543  
727 and many prior Internet Draft versions of this RFC. That rule caused proxies to destroy the contents of  
728 the Request-URI when a Route header field was present. Strict routing behavior is not used in this  
729 specification, in favor of a loose routing behavior. Proxies that perform strict routing are also known  
730 as strict routers.
- 731 **Target Refresh Request:** A target refresh request sent within a dialog is defined as a request that can  
732 modify the remote target of the dialog.
- 733 **Transaction User (TU):** The layer of protocol processing that resides above the transaction layer. Trans-  
734 action users include the UAC core, UAS core, and proxy core.
- 735 **Upstream:** A direction of message forwarding within a transaction that refers to the direction that responses  
736 flow from the user agent server back to the user agent client.
- 737 **URL-encoded:** A character string encoded according to RFC 1738, Section 2.2 [6].
- 738 **User Agent Client (UAC):** A user agent client is a logical entity that creates a new request, and then uses  
739 the client transaction state machinery to send it. The role of UAC lasts only for the duration of that  
740 transaction. In other words, if a piece of software initiates a request, it acts as a UAC for the duration  
741 of that transaction. If it receives a request later, it assumes the role of a user agent server for the  
742 processing of that transaction.
- 743 **UAC Core:** The set of processing functions required of a UAC that reside above the transaction and trans-  
744 port layers.

745 **User Agent Server (UAS):** A user agent server is a logical entity that generates a response to a SIP request.  
 746 The response accepts, rejects, or redirects the request. This role lasts only for the duration of that  
 747 transaction. In other words, if a piece of software responds to a request, it acts as a UAS for the  
 748 duration of that transaction. If it generates a request later, it assumes the role of a user agent client for  
 749 the processing of that transaction.

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

752 **User Agent (UA):** A logical entity that can act as both a user agent client and user agent server.

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

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

## 759 7 SIP Messages

760 SIP is a text-based protocol and uses the ISO 10646 character set in UTF-8 encoding (RFC 2279 [7]).

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

762 Both **Request** (section 7.1) and **Response** (section 7.2) messages use the basic format of RFC 2822  
 763 [3], even though the syntax differs in character set and syntax specifics. (SIP allows header fields that would  
 764 not be valid RFC 2822 header fields, for example.) Both types of messages consist of a **start-line**, one or  
 765 more header fields, an empty line indicating the end of the header fields, and an optional **message-body**.

```

766     generic-message = start-line
                        *message-header
                        CRLF
                        [ message-body ]
     start-line      = Request-Line / Status-Line
```

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

769 Except for the above difference in character sets, much of SIP's message and header field syntax is  
 770 identical to HTTP/1.1. Rather than repeating the syntax and semantics here, we use [HX.Y] to refer to  
 771 Section X.Y of the current HTTP/1.1 specification (RFC 2616 [8]).

772 However, SIP is not an extension of HTTP.

### 773 7.1 Requests

774 SIP requests are distinguished by having a **Request-Line** for a **start-line**. A **Request-Line** contains a  
 775 method name, a **Request-URI**, and the protocol version separated by a single space (SP) character.

776 The **Request-Line** ends with CRLF. No CR or LF are allowed except in the end-of-line CRLF se-  
 777 quence. No linear whitespace (LWS) is allowed in any of the elements.

778 Request-Line = Method SP Request-URI SP SIP-Version CRLF

779 **Method:** This specification defines six methods: REGISTER for registering contact information, INVITE,  
780 ACK, and CANCEL for setting up sessions, BYE for terminating sessions, and OPTIONS for query-  
781 ing servers about their capabilities. SIP extensions, documented in standards track RFCs, may define  
782 additional methods.

783 **Request-URI:** The Request-URI is a SIP or SIPS URI as described in Section 19.1 or a general URI  
784 (RFC 2396 [5]). It indicates the user or service to which this request is being addressed. The Request-  
785 URI MUST NOT contain unescaped spaces or control characters and MUST NOT be enclosed in "<>".  
786 SIP elements MAY support Request-URIs with schemes other than "sip" and "sips", for example the  
787 "tel" URI scheme of RFC 2806 [9]. SIP elements MAY translate non-SIP URIs using any mechanism  
788 at their disposal, resulting in either SIP URI, SIPS URI, or some other scheme.

789 **SIP-Version:** Both request and response messages include the version of SIP in use, and follow [H3.1] (with  
790 HTTP replaced by SIP, and HTTP/1.1 replaced by SIP/2.0) regarding version ordering, compliance  
791 requirements, and upgrading of version numbers. To be compliant with this specification, applications  
792 sending SIP messages MUST include a SIP-Version of "SIP/2.0". The SIP-Version string is case-  
793 insensitive, but implementations MUST send upper-case.

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

## 796 7.2 Responses

797 SIP responses are distinguished from requests by having a Status-Line as their start-line. A Status-Line  
798 consists of the protocol version followed by a numeric Status-Code and its associated textual phrase, with  
799 each element separated by a single SP character.

800 No CR or LF is allowed except in the final CRLF sequence.

801 Status-Line = SIP-Version SP Status-Code SP Reason-Phrase CRLF

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

806 While this specification suggests specific wording for the reason phrase, implementations MAY choose  
807 other text, for example, in the language indicated in the Accept-Language header field of the request.

808 The first digit of the Status-Code defines the class of response. The last two digits do not have any  
809 categorization role. For this reason, any response with a status code between 100 and 199 is referred to as  
810 a "1xx response", any response with a status code between 200 and 299 as a "2xx response", and so on.  
811 SIP/2.0 allows six values for the first digit:

812 **1xx:** Provisional – request received, continuing to process the request;

813 **2xx:** Success – the action was successfully received, understood, and accepted;

814 **3xx:** Redirection – further action needs to be taken in order to complete the request;

815 **4xx:** Client Error – the request contains bad syntax or cannot be fulfilled at this server;

816 **5xx:** Server Error – the server failed to fulfill an apparently valid request;

817 **6xx:** Global Failure – the request cannot be fulfilled at any server.

818 Section 21 defines these classes and describes the individual codes.

### 819 7.3 Header Fields

820 SIP header fields are similar to HTTP header fields in both syntax and semantics. In particular, SIP header  
821 fields follow the [H4.2] definitions of syntax for **message-header** and the rules for extending header fields  
822 over multiple lines. However, the latter is specified in HTTP with implicit whitespace and folding. This  
823 specification conforms with RFC 2234 [10] and uses only explicit whitespace and folding as an integral part  
824 of the grammar.

825 [H4.2] also specifies that multiple header fields of the same field name whose value is a comma-separated  
826 list can be combined into one header field. That applies to SIP as well, but the specific rule is different  
827 because of the different grammars. Specifically, any SIP header whose grammar is of the form:

828 `header = "header-name" HCOLON header-value *(COMMA header-value)`

829 allows for combining header fields of the same name into a comma-separated list. This is also true for  
830 the **Contact** header, as long as none of the header field values are “\*”.

#### 831 7.3.1 Header Field Format

832 Header fields follow the same generic header format as that given in Section 2.2 of RFC 2822 [3]. Each  
833 header field consists of a field name followed by a colon (":") and the field value.

834 `field-name: field-value`

835 The formal grammar for a **message-header** specified in Section 25 allows for an arbitrary amount of  
836 whitespace on either side of the colon; however, implementations should avoid spaces between the field  
837 name and the colon and use a single space (SP) between the colon and the **field-value**. Thus,

838 `Subject: lunch`

839 `Subject : lunch`

840 `Subject :lunch`

841 `Subject: lunch`

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

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

846 `Subject: I know you're there, pick up the phone and talk to me!`

847 `Subject: I know you're there,`

848 `pick up the phone`

849 `and talk to me!`

850 The relative order of header fields with different field names is not significant. However, it is RECOM-  
851 MENDED that header fields which are needed for proxy processing (Via, Route, Record-Route, Proxy-  
852 Require, Max-Forwards, and Proxy-Authorization, for example) appear towards the top of the message  
853 to facilitate rapid parsing. The relative order of header field rows with the same field name is important.  
854 Multiple header field rows with the same field-name MAY be present in a message if and only if the entire  
855 field-value for that header field is defined as a comma-separated list (that is, if follows the grammar defined  
856 in Section 7.3). It MUST be possible to combine the multiple header field rows into one "field-name: field-  
857 value" pair, without changing the semantics of the message, by appending each subsequent field-value to  
858 the first, each separated by a comma. The exceptions to this rule are the WWW-Authenticate, Authoriza-  
859 tion, Proxy-Authenticate, and Proxy-Authorization header fields. Multiple header field rows with these  
860 names MAY be present in a message, but since their grammar does not follow the general form listed in  
861 Section 7.3, they MUST NOT be combined into a single header field row.

862 Implementations MUST be able to process multiple header field rows with the same name in any combi-  
863 nation of the single-value-per-line or comma-separated value forms.

864 The following groups of header field rows are valid and equivalent:

```
865 Route: <sip:alice@atlanta.com>
866 Subject: Lunch
867 Route: <sip:bob@biloxi.com>
868 Route: <sip:carol@chicago.com>
869
870 Route: <sip:alice@atlanta.com>, <sip:bob@biloxi.com>
871 Route: <sip:carol@chicago.com>
872 Subject: Lunch
873
874 Subject: Lunch
875 Route: <sip:alice@atlanta.com>, <sip:bob@biloxi.com>, <sip:carol@chicago.com>
```

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

```
877 Route: <sip:alice@atlanta.com>
878 Route: <sip:bob@biloxi.com>
879 Route: <sip:carol@chicago.com>
880
881 Route: <sip:bob@biloxi.com>
882 Route: <sip:alice@atlanta.com>
883 Route: <sip:carol@chicago.com>
884
885 Route: <sip:alice@atlanta.com>, <sip:carol@chicago.com>, <sip:bob@biloxi.com>
```

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

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

891 Even though an arbitrary number of parameter pairs may be attached to a header field value, any given  
892 **parameter-name** MUST NOT appear more than once.

893 When comparing header fields, field names are always case-insensitive. Unless otherwise stated in  
894 the definition of a particular header field, field values, parameter names, and parameter values are case-  
895 insensitive. Tokens are always case-insensitive. Unless specified otherwise, values expressed as quoted  
896 strings are case-sensitive.

897 For example,

898 Contact: <sip:alice@atlanta.com>;expires=3600

899 is equivalent to

900 CONTACT: <sip:alice@atlanta.com>;EXPIRES=3600

901 and

902 Content-Disposition: session;handling=optional

903 is equivalent to

904 content-disposition: Session;HANDLING=OPTIONAL

905 The following two header fields are not equivalent:

906 Warning: 370 devnull "Choose a bigger pipe"

907 Warning: 370 devnull "CHOOSE A BIGGER PIPE"

### 908 **7.3.2 Header Field Classification**

909 Some header fields only make sense in requests or responses. These are called request header fields and  
910 response header fields, respectively. If a header field appears in a message not matching its category (such  
911 as a request header field in a response), it MUST be ignored. Section 20 defines the classification of each  
912 header field.

### 913 **7.3.3 Compact Form**

914 SIP provides a mechanism to represent common header field names in an abbreviated form. This may  
915 be useful when messages would otherwise become too large to be carried on the transport available to it  
916 (exceeding the maximum transmission unit (MTU) when using UDP, for example). These compact forms  
917 are defined in Section 20. A compact form MAY be substituted for the longer form of a header field name at  
918 any time without changing the semantics of the message. A header field name MAY appear in both long and  
919 short forms within the same message. Implementations MUST accept both the long and short forms of each  
920 header name.

## 921 7.4 Bodies

922 Requests, including new requests defined in extensions to this specification, MAY contain message bodies  
923 unless otherwise noted. The interpretation of the body depends on the request method.

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

### 926 7.4.1 Message Body Type

927 The Internet media type of the message body MUST be given by the Content-Type header field. If the body  
928 has undergone any encoding such as compression, then this MUST be indicated by the Content-Encoding  
929 header field; otherwise, Content-Encoding MUST be omitted. If applicable, the character set of the message  
930 body is indicated as part of the Content-Type header-field value.

931 The “multipart” MIME type defined in RFC 2046 [11] MAY be used within the body of the message.  
932 Implementations that send requests containing multipart message bodies MUST send a session description  
933 as a non-multipart message body if the remote implementation requests this through an Accept header field  
934 that does not contain multipart.

935 Note that SIP messages MAY contain binary bodies or body parts.

### 936 7.4.2 Message Body Length

937 The body length in bytes is provided by the Content-Length header field. Section 20.14 describes the  
938 necessary contents of this header field in detail.

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

## 941 7.5 Framing SIP messages

942 Unlike HTTP, SIP implementations can use UDP or other unreliable datagram protocols. Each such data-  
943 gram carries one request or response. See Section 18 on constraints on usage of unreliable transports.

944 Implementations processing SIP messages over stream-oriented transports MUST ignore any CRLF ap-  
945 pearing before the start-line [H4.1].

946 The Content-Length header field value is used to locate the end of each SIP message in a stream. It will always  
947 be present when SIP messages are sent over stream-oriented transports.

## 948 8 General User Agent Behavior

949 A user agent represents an end system. It contains a user agent client (UAC), which generates requests, and  
950 a user agent server (UAS), which responds to them. A UAC is capable of generating a request based on  
951 some external stimulus (the user clicking a button, or a signal on a PSTN line) and processing a response. A  
952 UAS is capable of receiving a request and generating a response based on user input, external stimulus, the  
953 result of a program execution, or some other mechanism.

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

956 UAC and UAS procedures depend strongly on two factors. First, based on whether the request or  
957 response is inside or outside of a dialog, and second, based on the method of a request. Dialogs are discussed

958 thoroughly in Section 12; they represent a peer-to-peer relationship between user agents and are established  
959 by specific SIP methods, such as INVITE.

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

963 Security procedures for requests and responses outside of a dialog are described in Section 26. Specif-  
964 ically, mechanisms exist for the UAS and UAC to mutually authenticate. A limited set of privacy features  
965 are also supported through encryption of bodies using S/MIME.

## 966 8.1 UAC Behavior

967 This section covers UAC behavior outside of a dialog.

### 968 8.1.1 Generating the Request

969 A valid SIP request formulated by a UAC MUST at a minimum contain the following header fields: To, From,  
970 CSeq, Call-ID, Max-Forwards, and Via; all of these header fields are mandatory in all SIP messages.  
971 These six header fields are the fundamental building blocks of a SIP message, as they jointly provide for  
972 most of the critical message routing services including the addressing of messages, the routing of responses,  
973 limiting message propagation, ordering of messages, and the unique identification of transactions. These  
974 header fields are in addition to the mandatory request line, which contains the method, Request-URI, and  
975 SIP version.

976 Examples of requests sent outside of a dialog include an INVITE to establish a session (Section 13) and  
977 an OPTIONS to query for capabilities (Section 11).

978 **8.1.1.1 Request-URI** The initial Request-URI of the message SHOULD be set to the value of the URI  
979 in the To field. One notable exception is the REGISTER method; behavior for setting the Request-URI  
980 of REGISTER is given in Section 10. It may also be undesirable for privacy reasons or convenience to  
981 set these fields to the same value (especially if the originating UA expects that the Request-URI will be  
982 changed during transit).

983 In some special circumstances, the presence of a pre-existing route set can affect the Request-URI of  
984 the message. A pre-existing route set is an ordered set of URIs that identify a chain of servers, to which a  
985 UAC will send outgoing requests that are outside of a dialog. Commonly, they are configured on the UA by  
986 a user or service provider manually, or through some other non-SIP mechanism. When a provider wishes  
987 to configure a UA with an outbound proxy, it is RECOMMENDED that this be done by providing it with a  
988 pre-existing route set with a single URI, that of the outbound proxy.

989 When a pre-existing route set is present, the procedures for populating the Request-URI and Route  
990 header field detailed in Section 12.2.1.1 MUST be followed, even though there is no dialog. If the Request-  
991 URI specifies a SIPS URI, all the SIP URI in the route set MUST be converted to SIPS URI (by changing  
992 the scheme to SIPS) before performing the processing of Section 12.2.1.1.

993 **8.1.1.2 To** The To header field first and foremost specifies the desired “logical” recipient of the request,  
994 or the address-of-record of the user or resource that is the target of this request. This may or may not be  
995 the ultimate recipient of the request. The To header field MAY contain a SIP or SIPS URI, but it may also  
996 make use of other URI schemes (the tel URL [9], for example) when appropriate. All SIP implementations

997 MUST support the SIP and URI scheme. Any implementation that supports TLS MUST support the SIPS  
998 URI scheme. The To header field allows for a display name.

999 A UAC may learn how to populate the To header field for a particular request in a number of ways.  
1000 Usually the user will suggest the To header field through a human interface, perhaps inputting the URI  
1001 manually or selecting it from some sort of address book. Frequently, the user will not enter a complete  
1002 URI, but rather, a string of digits or letters (for example, "bob"). It is at the discretion of the UA to choose  
1003 how to interpret this input. Using it to form the user part of a SIP URI implies that the UA wishes the  
1004 name to be resolved in the domain to the right-hand side (RHS) of the at-sign in the SIP URI (for instance,  
1005 sip:bob@example.com). Using it to form the user part of a SIPS URI implies that the UA wishes to securely  
1006 communicate, and that the name is to be resolved in the domain to the RHS of the at-sign. The RHS  
1007 will frequently be the home domain of the user, which allows for the home domain to process the outgoing  
1008 request. This is useful for features like "speed dial" that require interpretation of the user part in the home  
1009 domain. The tel URL may be used when the UA does not wish to specify the domain that should interpret a  
1010 telephone number that has been inputted by the user. Rather, each domain through which the request passes  
1011 would be given that opportunity. As an example, a user in an airport might log in and send requests through  
1012 an outbound proxy in the airport. If they enter "411" (this is the phone number for local directory assistance  
1013 in the United States), that needs to be interpreted and processed by the outbound proxy in the airport, not  
1014 the user's home domain. In this case, tel:411 would be the right choice.

1015 A request outside of a dialog MUST NOT contain a tag; the tag in the To field of a request identifies the  
1016 peer of the dialog. Since no dialog is established, no tag is present.

1017 For further information on the To header field, see Section 20.39. The following is an example of valid  
1018 To header field:

```
1019 To: Carol <sip:carol@chicago.com>
```

1020 **8.1.1.3 From** The From header field indicates the logical identity of the initiator of the request, possibly  
1021 the user's address-of-record. Like the To header field, it contains a URI and optionally a display name. It  
1022 is used by SIP elements to determine which processing rules to apply to a request (for example, automatic  
1023 call rejection). As such, it is very important that the From URI not contain IP addresses or the FQDN of the  
1024 host on which the UA is running, since these are not logical names.

1025 The From header field allows for a display name. A UAC SHOULD use the display name "Anonymous",  
1026 along with a syntactically correct, but otherwise meaningless URI (like sip:thisis@anonymous.invalid), if  
1027 the identity of the client is to remain hidden.

1028 Usually the value that populates the From header field in requests generated by a particular UA is pre-  
1029 provisioned by the user or by the administrators of the user's local domain. If a particular UA is used by  
1030 multiple users, it might have switchable profiles that include a URI corresponding to the identity of the  
1031 profiled user. Recipients of requests can authenticate the originator of a request in order to ascertain that  
1032 they are who their From header field claims they are (see Section 22 for more on authentication).

1033 The From field MUST contain a new "tag" parameter, chosen by the UAC. See Section 19.3 for details  
1034 on choosing a tag.

1035 For further information on the From header field, see Section 20.20. Examples:

```
1036 From: "Bob" <sips:bob@biloxi.com> ;tag=a48s  
1037 From: sip:+12125551212@phone2net.com;tag=887s  
1038 From: Anonymous <sip:c8oqz84zk7z@privacy.org>;tag=hyh8
```

1039 **8.1.1.4 Call-ID** The Call-ID header field acts as a unique identifier to group together a series of mes-  
1040 sages. It MUST be the same for all requests and responses sent by either UA in a dialog. It SHOULD be the  
1041 same in each registration from a UA.

1042 In a new request created by a UAC outside of any dialog, the Call-ID header field MUST be selected by  
1043 the UAC as a globally unique identifier over space and time unless overridden by method-specific behavior.  
1044 All SIP UAs must have a means to guarantee that the Call-ID header fields they produce will not be inad-  
1045 vertently generated by any other UA. Note that when requests are retried after certain failure responses that  
1046 solicit an amendment to a request (for example, a challenge for authentication), these retried requests are  
1047 not considered new requests, and therefore do not need new Call-ID header fields; see Section 8.1.3.5.

1048 Use of cryptographically random identifiers [12] in the generation of Call-IDs is RECOMMENDED. Im-  
1049 plementations MAY use the form "localid@host". Call-IDs are case-sensitive and are simply compared  
1050 byte-by-byte.

1051 Using cryptographically random identifiers provides some protection against session hijacking and reduces the  
1052 likelihood of unintentional Call-ID collisions.

1053 No provisioning or human interface is required for the selection of the Call-ID header field value for a  
1054 request.

1055 For further information on the Call-ID header field, see Section 20.8.

1056 Example:

1057 Call-ID: f81d4fae-7dec-11d0-a765-00a0c91e6bf6@foo.bar.com

1058 **8.1.1.5 CSeq** The CSeq header field serves as a way to identify and order transactions. It consists  
1059 of a sequence number and a method. The method MUST match that of the request. For non-REGISTER  
1060 requests outside of a dialog, the sequence number value is arbitrary. The sequence number value MUST  
1061 be expressible as a 32-bit unsigned integer and MUST be less than 2\*\*31. As long as it follows the above  
1062 guidelines, a client may use any mechanism it would like to select CSeq header field values.

1063 Section 12.2.1.1 discusses construction of the CSeq for requests within a dialog.

1064 Example:

1065 CSeq: 4711 INVITE

1066 **8.1.1.6 Max-Forwards** The Max-Forwards header field serves to limit the number of hops a request  
1067 can transit on the way to its destination. It consists of an integer that is decremented by one at each hop.  
1068 If the Max-Forwards value reaches 0 before the request reaches its destination, it will be rejected with a  
1069 483(Too Many Hops) error response.

1070 A UAC MUST insert a Max-Forwards header field into each request it originates with a value which  
1071 SHOULD be 70. This number was chosen to be sufficiently large to guarantee that a request would not be  
1072 dropped in any SIP network when there were no loops, but not so large as to consume proxy resources when  
1073 a loop does occur. Lower values should be used with caution and only in networks where topologies are  
1074 known by the UA.

1075 **8.1.1.7 Via** The Via header field is used to indicate the transport used for the transaction and to identify  
1076 the location where the response is to be sent. A Via header field value is added only after the transport that  
1077 will be used to reach the next hop has been selected (which may involve the usage of the procedures in [4]).

1078

1079 When the UAC creates a request, it MUST insert a **Via** into that request. The protocol name and protocol  
1080 version in the header field MUST be SIP and 2.0, respectively. The **Via** header field value MUST contain a  
1081 branch parameter. This parameter is used to identify the transaction created by that request. This parameter  
1082 is used by both the client and the server.

1083 The branch parameter value MUST be unique across space and time for all requests sent by the UA.  
1084 The exceptions to this rule are **CANCEL** and **ACK** for non-2xx responses. As discussed below, a **CANCEL**  
1085 request will have the same value of the branch parameter as the request it cancels. As discussed in Section  
1086 17.1.1.3, an **ACK** for a non-2xx response will also have the same branch ID as the **INVITE** whose response  
1087 it acknowledges.

1088 The uniqueness property of the branch ID parameter, to facilitate its use as a transaction ID, was not part of RFC  
1089 2543

1090 The branch ID inserted by an element compliant with this specification MUST always begin with the  
1091 characters "z9hG4bK". These 7 characters are used as a magic cookie (7 is deemed sufficient to ensure that  
1092 an older RFC 2543 implementation would not pick such a value), so that servers receiving the request can  
1093 determine that the branch ID was constructed in the fashion described by this specification (that is, globally  
1094 unique). Beyond this requirement, the precise format of the branch token is implementation-defined.

1095 The **Via** header **maddr**, **ttl**, and **sent-by** components will be set when the request is processed by the  
1096 transport layer (Section 18).

1097 **Via** processing for proxies is described in Section 16.6 Item 8 and Section 16.7 Item 3.

1098 **8.1.1.8 Contact** The **Contact** header field provides a SIP URI that can be used to contact that specific  
1099 instance of the UA for subsequent requests. The **Contact** header field MUST be present and contain exactly  
1100 one SIP URI in any request that can result in the establishment of a dialog. For the methods defined in this  
1101 specification, that includes only the **INVITE** request. For these requests, the scope of the **Contact** is global.  
1102 That is, the **Contact** header field value contains the URI at which the UA would like to receive requests,  
1103 and this URI MUST be valid even if used in subsequent requests outside of any dialogs.

1104 If the **Request-URI** contains a SIPS URI, the **Contact** header field MUST contain a SIPS URI as well.

1105 For further information on the **Contact** header field, see Section 20.10.

1106 **8.1.1.9 Supported and Require** If the UAC supports extensions to SIP that can be applied by the  
1107 server to the response, the UAC SHOULD include a **Supported** header field in the request listing the option  
1108 tags (Section 19.2) for those extensions.

1109 The option tags listed MUST only refer to extensions defined in standards-track RFCs. This is to pre-  
1110 vent servers from insisting that clients implement non-standard, vendor-defined features in order to receive  
1111 service. Extensions defined by experimental and informational RFCs are explicitly excluded from usage  
1112 with the **Supported** header field in a request, since they too are often used to document vendor-defined  
1113 extensions.

1114 If the UAC wishes to insist that a UAS understand an extension that the UAC will apply to the request  
1115 in order to process the request, it MUST insert a **Require** header field into the request listing the option tag  
1116 for that extension. If the UAC wishes to apply an extension to the request and insist that any proxies that are  
1117 traversed understand that extension, it MUST insert a **Proxy-Require** header field into the request listing the  
1118 option tag for that extension.

1119 As with the **Supported** header field, the option tags in the **Require** and **Proxy-Require** header fields  
1120 MUST only refer to extensions defined in standards-track RFCs.

1121 **8.1.1.10 Additional Message Components** After a new request has been created, and the header fields  
1122 described above have been properly constructed, any additional optional header fields are added, as are any  
1123 header fields specific to the method.

1124 SIP requests MAY contain a MIME-encoded message-body. Regardless of the type of body that a request  
1125 contains, certain header fields must be formulated to characterize the contents of the body. For further  
1126 information on these header fields, see Sections 20.11 through 20.15.

## 1127 **8.1.2 Sending the Request**

1128 The destination for the request is then computed. Unless there is local policy specifying otherwise, then  
1129 the destination MUST be determined by applying the DNS procedures described in [4] as follows. If the  
1130 first element in the route set indicated a strict router (resulting in forming the request as described in Sec-  
1131 tion 12.2.1.1), the procedures MUST be applied to the Request-URI of the request. Otherwise, the pro-  
1132 cedures are applied to the first Route header field value in the request (if one exists), or to the request's  
1133 Request-URI if there is no Route header field present. These procedures yield an ordered set of address,  
1134 port, and transports to attempt.

1135 Local policy MAY specify an alternate set of destinations to attempt. If the Request-URI contains a  
1136 SIPS URI, any alternate destinations MUST be contacted with TLS. Beyond that, there are no restrictions on  
1137 the alternate destinations if the request contains no Route header field. This provides a simple alternative  
1138 to a pre-existing route set as a way to specify an outbound proxy. However, that approach for configuring  
1139 an outbound proxy is NOT RECOMMENDED; a pre-existing route set with a single URI SHOULD be used  
1140 instead. If the request contains a Route header field, the request SHOULD be sent to the locations derived  
1141 from its topmost value, but MAY be sent to any server that the UA is certain will honor the Route and  
1142 Request-URI policies specified in this document (as opposed to those in RFC 2543). In particular, a UAC  
1143 configured with an outbound proxy SHOULD attempt to send the request to the location indicated in the first  
1144 Route header field value instead of adopting the policy of sending all messages to the outbound proxy.

1145 This ensures that outbound proxies choosing not to add Record-Route header field values will drop out of the  
1146 path of subsequent requests. It allows endpoints that cannot resolve the first Route URI to delegate that task to an  
1147 outbound proxy.

1148 The UAC SHOULD follow the procedures defined in [4] for stateful elements, trying each address until  
1149 a server is contacted. Each try constitutes a new transaction, and therefore each carries a different topmost  
1150 Via header field value with a new branch parameter. Furthermore, the transport value in the Via header field  
1151 is set to whatever transport was determined for the target server.

## 1152 **8.1.3 Processing Responses**

1153 Responses are first processed by the transport layer and then passed up to the transaction layer. The trans-  
1154 action layer performs its processing and then passes the response up to the TU. The majority of response  
1155 processing in the TU is method specific. However, there are some general behaviors independent of the  
1156 method.

1157 **8.1.3.1 Transaction Layer Errors** In some cases, the response returned by the transaction layer will not  
1158 be a SIP message, but rather a transaction layer error. When a timeout error is received from the transaction  
1159 layer, it MUST be treated as if a 408 (Request Timeout) status code has been received. If a fatal transport

1160 error is reported by the transport layer (generally, due to fatal ICMP errors in UDP or connection failures in  
1161 TCP), the condition **MUST** be treated as a 503 (Service Unavailable) status code.

1162 **8.1.3.2 Unrecognized Responses** A UAC **MUST** treat any final response it does not recognize as being  
1163 equivalent to the x00 response code of that class, and **MUST** be able to process the x00 response code for  
1164 all classes. For example, if a UAC receives an unrecognized response code of 431, it can safely assume that  
1165 there was something wrong with its request and treat the response as if it had received a 400 (Bad Request)  
1166 response code. A UAC **MUST** treat any provisional response different than 100 that it does not recognize as  
1167 183 (Session Progress). A UAC **MUST** be able to process 100 and 183 responses.

1168 **8.1.3.3 Vias** If more than one *Via* header field value is present in a response, the UAC **SHOULD** discard  
1169 the message.

1170           The presence of additional *Via* header field values that precede the originator of the request suggests that the  
1171 message was misrouted or possibly corrupted.

1172 **8.1.3.4 Processing 3xx Responses** Upon receipt of a redirection response (for example, a 301 response  
1173 status code), clients **SHOULD** use the URI(s) in the *Contact* header field to formulate one or more new  
1174 requests based on the redirected request. If the original request had a SIPS URI in the *Request-URI*, the  
1175 client **MUST** discard any *Contact* header fields which do not contain SIPS URIs.

1176           If more than one URI is present in *Contact* header field within the 3xx response, the UA **MUST** determine  
1177 an order in which these contact addresses should be processed. UAs **MUST** consult the “q” parameter value  
1178 of the *Contact* header field value (see Section 20.10) if available. Contact addresses **MUST** be ordered from  
1179 highest qvalue to lowest. If no qvalue is present, a contact address is considered to have a qvalue of 1.0.  
1180 Note that two or more contact addresses might have an equal qvalue - these URIs are eligible to be tried in  
1181 parallel.

1182           Once an ordered list has been established, UACs **MAY** remove from the list any entry that they do not  
1183 want to try. After this, UACs **MUST** try to contact each URI in the ordered list in turn by sending a request  
1184 for a single contact address at a time, continuing down the ordered list only when a final response to the  
1185 current request has been received. If there are contact addresses with an equal qvalue, the UAC **MAY** decide  
1186 randomly on an order in which to process these addresses, or it **MAY** attempt to process contact addresses of  
1187 equal qvalue in parallel.

1188           Note that, for example, the UAC may effectively divide the ordered list into groups, processing the  
1189 groups serially and processing the destinations in each group in parallel.

1190           If contacting an address in the list results in a failure, as defined in the next paragraph, the element moves  
1191 to the next address in the list, until the list is exhausted. If the list is exhausted, then the request has failed.

1192           Failures **SHOULD** be detected through failure response codes (codes greater than 399); for network errors  
1193 the client transaction will report any transport layer failures to the transaction user. Note that some response  
1194 codes (detailed in 8.1.3.5) indicate that the request can be retried; requests that are reattempted should not  
1195 be considered failures.

1196           When a failure for a particular contact address is received, the client **SHOULD** try the next contact  
1197 address. This will involve creating a new client transaction to deliver a new request.

1198           In order to create a request based on a contact address in a 3xx response, a UAC **MUST** copy the entire  
1199 URI from the *Contact* header field value into the *Request-URI*, except for the “method-param” and  
1200 “header” URI parameters (see Section 19.1.1 for a definition of these parameters). It uses the “header”

1201 parameters to create header field values for the new request, overwriting header field values associated with  
1202 the redirected request in accordance with the guidelines in Section 19.1.5.

1203 Note that in some instances, header fields that have been communicated in the contact address may  
1204 instead append to existing request header fields in the original redirected request. As a general rule, if the  
1205 header field can accept a comma-separated list of values, then the new header field value MAY be appended  
1206 to any existing values in the original redirected request. If the header field does not accept multiple values,  
1207 the value in the original redirected request MAY be overwritten by the header field value communicated in  
1208 the contact address. For example, if a contact address is returned with the following value:

```
1209 sip:user@host?Subject=foo&Call-Info=<http://www.foo.com>
```

1210 Then any **Subject** header field in the original redirected request is overwritten, but the HTTP URL is  
1211 merely appended to any existing **Call-Info** header field values.

1212 It is RECOMMENDED that the UAC reuse the same **To**, **From**, and **Call-ID** used in the original redirected  
1213 request, but the UAC MAY also choose to update the **Call-ID** header field value for new requests, for example.

1214 Finally, once the new request has been constructed, it is sent using a new client transaction, and therefore  
1215 MUST have a new branch ID in the top **Via** field as discussed in Section 8.1.1.7.

1216 In all other respects, requests sent upon receipt of a redirect response SHOULD re-use the header fields  
1217 and bodies of the original request.

1218 Redirections can result in requests that are in turn redirected. For example, if an initial 3xx response  
1219 contains multiple contacts, and the retry of the request to the first of these contacts is in turn redirected,  
1220 UACs must reconcile the two resulting sets of URIs. UAs MUST combine the two sets of contact addresses  
1221 and recompute the ordering of the elements following the steps described above. However, if any two URIs  
1222 in the set are equivalent, the less preferred URI, meaning the URI with the numerically highest “q” value,  
1223 MUST be discarded.

1224 In some instances, **Contact** header field values may be cached at UAC temporarily or permanently  
1225 depending on the status code received and the presence of an expiration interval; see Sections 21.3.2 and  
1226 21.3.3.

1227 **8.1.3.5 Processing 4xx Responses** Certain 4xx response codes require specific UA processing, independ-  
1228 ent of the method.

1229 If a 401 (Unauthorized) or 407 (Proxy Authentication Required) response is received, the UAC SHOULD  
1230 follow the authorization procedures of Section 22.2 and Section 22.3 to retry the request with credentials.

1231 If a 413 (Request Entity Too Large) response is received (Section 21.4.11), the request contained a body  
1232 that was longer than the UAS was willing to accept. If possible, the UAC SHOULD retry the request, either  
1233 omitting the body or using one of a smaller length.

1234 If a 415 (Unsupported Media Type) response is received (Section 21.4.13), the request contained media  
1235 types not supported by the UAS. The UAC SHOULD retry sending the request, this time only using content  
1236 with types listed in the **Accept** header field in the response, with encodings listed in the **Accept-Encoding**  
1237 header field in the response, and with languages listed in the **Accept-Language** in the response.

1238 If a 416 (Unsupported URI Scheme) response is received (Section 21.4.14), the **Request-URI** used a  
1239 URI scheme not supported by the server. The client SHOULD retry the request, this time, using a SIP URI.

1240 If a 420 (Bad Extension) response is received (Section 21.4.15), the request contained a **Require** or  
1241 **Proxy-Require** header field listing an option-tag for a feature not supported by a proxy or UAS. The UAC  
1242 SHOULD retry the request, this time omitting any extensions listed in the **Unsupported** header field in the

1243 response.

1244 In all of the above cases, the request is retried by creating a new request with the appropriate modifica-  
1245 tions. This new request SHOULD have the same value of the Call-ID, To, and From of the previous request,  
1246 but the CSeq should contain a new sequence number that is one higher than the previous.

1247 With other 4xx responses, including those yet to be defined, a retry may or may not be possible depend-  
1248 ing on the method and the use case.

## 1249 8.2 UAS Behavior

1250 When a request outside of a dialog is processed by a UAS, there is a set of processing rules that are followed,  
1251 independent of the method. Section 12 gives guidance on how a UAS can tell whether a request is inside or  
1252 outside of a dialog.

1253 Note that request processing is atomic. If a request is accepted, all state changes associated with it MUST  
1254 be performed. If it is rejected, all state changes MUST NOT be performed.

1255 UASs SHOULD process the requests in the order of the steps that follow in this section (that is, starting  
1256 with authentication, then inspecting the method, the header fields, and so on throughout the remainder of  
1257 this section).

### 1258 8.2.1 Method Inspection

1259 Once a request is authenticated (or authentication is skipped), the UAS MUST inspect the method of the  
1260 request. If the UAS recognizes but does not support the method of a request, it MUST generate a 405  
1261 (Method Not Allowed) response. Procedures for generating responses are described in Section 8.2.6. The  
1262 UAS MUST also add an Allow header field to the 405 (Method Not Allowed) response. The Allow header  
1263 field MUST list the set of methods supported by the UAS generating the message. The Allow header field is  
1264 presented in Section 20.5.

1265 If the method is one supported by the server, processing continues.

### 1266 8.2.2 Header Inspection

1267 If a UAS does not understand a header field in a request (that is, the header field is not defined in this spec-  
1268 ification or in any supported extension), the server MUST ignore that header field and continue processing  
1269 the message. A UAS SHOULD ignore any malformed header fields that are not necessary for processing  
1270 requests.

1271 **8.2.2.1 To and Request-URI** The To header field identifies the original recipient of the request desig-  
1272 nated by the user identified in the From field. The original recipient may or may not be the UAS processing  
1273 the request, due to call forwarding or other proxy operations. A UAS MAY apply any policy it wishes to  
1274 determine whether to accept requests when the To header field is not the identity of the UAS. However, it is  
1275 RECOMMENDED that a UAS accept requests even if they do not recognize the URI scheme (for example, a  
1276 tel: URI) in the To header field, or if the To header field does not address a known or current user of this  
1277 UAS. If, on the other hand, the UAS decides to reject the request, it SHOULD generate a response with a 403  
1278 (Forbidden) status code and pass it to the server transaction for transmission.

1279 However, the Request-URI identifies the UAS that is to process the request. If the Request-URI uses  
1280 a scheme not supported by the UAS, it SHOULD reject the request with a 416 (Unsupported URI Scheme)  
1281 response. If the Request-URI does not identify an address that the UAS is willing to accept requests for,

1282 it SHOULD reject the request with a 404 (Not Found) response. Typically, a UA that uses the REGISTER  
1283 method to bind its address-of-record to a specific contact address will see requests whose Request-URI  
1284 equals that contact address. Other potential sources of received Request-URIs include the Contact header  
1285 fields of requests and responses sent by the UA that establish or refresh dialogs.

1286 **8.2.2.2 Merged Requests** If the request has no tag in the To header field, the UAS core MUST check  
1287 the request against ongoing transactions. If the To tag, From tag, Call-ID, CSeq exactly match (including  
1288 tags) those associated with an ongoing transaction, but the branch-ID in the topmost Via does not match ,  
1289 the UAS core SHOULD generate a 482 (Loop Detected) response and pass it to the server transaction.

1290 The same request has arrived at the UAS more than once, following different paths, most likely due to forking.  
1291 The UAS processes the first such request received and responds with a 482 (Loop Detected) to the rest of them.

1292 **8.2.2.3 Require** Assuming the UAS decides that it is the proper element to process the request, it ex-  
1293 amines the Require header field, if present.

1294 The Require header field is used by a UAC to tell a UAS about SIP extensions that the UAC expects  
1295 the UAS to support in order to process the request properly. Its format is described in Section 20.32. If a  
1296 UAS does not understand an option-tag listed in a Require header field, it MUST respond by generating a  
1297 response with status code 420 (Bad Extension). The UAS MUST add an Unsupported header field, and list  
1298 in it those options it does not understand amongst those in the Require header field of the request.

1299 Note that Require and Proxy-Require MUST NOT be used in a SIP CANCEL request, or in an ACK  
1300 request sent for a non-2xx response. These header fields MUST be ignored if they are present in these  
1301 requests.

1302 An ACK request for a 2xx response MUST contain only those Require and Proxy-Require values that  
1303 were present in the initial request.

1304 Example:

```
1305 UAC->UAS:  INVITE sip:watson@bell-telephone.com SIP/2.0  
1306             Require: 100rel
```

1307

1308

```
1309 UAS->UAC:  SIP/2.0 420 Bad Extension  
1310             Unsupported: 100rel
```

1311 This behavior ensures that the client-server interaction will proceed without delay when all options are under-  
1312 stood by both sides, and only slow down if options are not understood (as in the example above). For a well-matched  
1313 client-server pair, the interaction proceeds quickly, saving a round-trip often required by negotiation mechanisms.  
1314 In addition, it also removes ambiguity when the client requires features that the server does not understand. Some  
1315 features, such as call handling fields, are only of interest to end systems.

## 1316 8.2.3 Content Processing

1317 Assuming the UAS understands any extensions required by the client, the UAS examines the body of the  
1318 message, and the header fields that describe it. If there are any bodies whose type (indicated by the Content-  
1319 Type), language (indicated by the Content-Language) or encoding (indicated by the Content-Encoding)  
1320 are not understood, and that body part is not optional (as indicated by the Content-Disposition header  
1321 field), the UAS MUST reject the request with a 415 (Unsupported Media Type) response. The response MUST  
1322 contain an Accept header field listing the types of all bodies it understands, in the event the request contained

1323 bodies of types not supported by the UAS. If the request contained content encodings not understood by the  
1324 UAS, the response MUST contain an **Accept-Encoding** header field listing the encodings understood by  
1325 the UAS. If the request contained content with languages not understood by the UAS, the response MUST  
1326 contain an **Accept-Language** header field indicating the languages understood by the UAS. Beyond these  
1327 checks, body handling depends on the method and type. For further information on the processing of  
1328 content-specific header fields, see Section 7.4 as well as Section 20.11 through 20.15.

#### 1329 **8.2.4 Applying Extensions**

1330 A UAS that wishes to apply some extension when generating the response MUST NOT do so unless support  
1331 for that extension is indicated in the **Supported** header field in the request. If the desired extension is not  
1332 supported, the server SHOULD rely only on baseline SIP and any other extensions supported by the client. In  
1333 rare circumstances, where the server cannot process the request without the extension, the server MAY send  
1334 a 421 (Extension Required) response. This response indicates that the proper response cannot be generated  
1335 without support of a specific extension. The needed extension(s) MUST be included in a **Require** header  
1336 field in the response. This behavior is NOT RECOMMENDED, as it will generally break interoperability.

1337 Any extensions applied to a non-421 response MUST be listed in a **Require** header field included in the  
1338 response. Of course, the server MUST NOT apply extensions not listed in the **Supported** header field in the  
1339 request. As a result of this, the **Require** header field in a response will only ever contain option tags defined  
1340 in standards-track RFCs.

#### 1341 **8.2.5 Processing the Request**

1342 Assuming all of the checks in the previous subsections are passed, the UAS processing becomes method-  
1343 specific. Section 10 covers the **REGISTER** request, section 11 covers the **OPTIONS** request, section 13  
1344 covers the **INVITE** request, and section 15 covers the **BYE** request.

#### 1345 **8.2.6 Generating the Response**

1346 When a UAS wishes to construct a response to a request, it follows the general procedures detailed in the  
1347 following subsections. Additional behaviors specific to the response code in question, which are not detailed  
1348 in this section, may also be required.

1349 Once all procedures associated with the creation of a response have been completed, the UAS hands the  
1350 response back to the server transaction from which it received the request.

1351 **8.2.6.1 Sending a Provisional Response** One largely non-method-specific guideline for the generation  
1352 of responses is that UASs SHOULD NOT issue a provisional response for a non-**INVITE** request. Rather,  
1353 UASs SHOULD generate a final response to a non-**INVITE** request as soon as possible.

1354 When a 100 (Trying) response is generated, any **Timestamp** header field present in the request MUST be  
1355 copied into this 100 (Trying) response. If there is a delay in generating the response, the UAS SHOULD add  
1356 a delay value into the **Timestamp** value in the response. This value MUST contain the difference between  
1357 time of sending of the response and receipt of the request, measured in seconds.

1358 **8.2.6.2 Headers and Tags** The **From** field of the response MUST equal the **From** header field of the  
1359 request. The **Call-ID** header field of the response MUST equal the **Call-ID** header field of the request. The

1360 CSeq header field of the response MUST equal the CSeq field of the request. The Via header field values in  
1361 the response MUST equal the Via header field values in the request and MUST maintain the same ordering.

1362 If a request contained a To tag in the request, the To header field in the response MUST equal that of  
1363 the request. However, if the To header field in the request did not contain a tag, the URI in the To header  
1364 field in the response MUST equal the URI in the To header field; additionally, the UAS MUST add a tag to  
1365 the To header field in the response (with the exception of the 100 (Trying) response, in which a tag MAY be  
1366 present). This serves to identify the UAS that is responding, possibly resulting in a component of a dialog  
1367 ID. The same tag MUST be used for all responses to that request, both final and provisional (again excepting  
1368 the 100 (Trying)). Procedures for generation of tags are defined in Section 19.3.

### 1369 8.2.7 Stateless UAS Behavior

1370 A stateless UAS is a UAS that does not maintain transaction state. It replies to requests normally, but  
1371 discards any state that would ordinarily be retained by a UAS after a response has been sent. If a stateless  
1372 UAS receives a retransmission of a request, it regenerates the response and resends it, just as if it were  
1373 replying to the first instance of the request. Stateless UASs do not use a transaction layer; they receive  
1374 requests directly from the transport layer and send responses directly to the transport layer.

1375 The stateless UAS role is needed primarily to handle unauthenticated requests for which a challenge  
1376 response is issued. If unauthenticated requests were handled statefully, then malicious floods of unau-  
1377 thenticated requests could create massive amounts of transaction state that might slow or completely halt  
1378 call processing in a UAS, effectively creating a denial of service condition; for more information see Sec-  
1379 tion 26.1.5.

1380 The most important behaviors of a stateless UAS are the following:

- 1381 • A stateless UAS MUST NOT send provisional (1xx) responses.
- 1382 • A stateless UAS MUST NOT retransmit responses.
- 1383 • A stateless UAS MUST ignore ACK requests.
- 1384 • A stateless UAS MUST ignore CANCEL requests.
- 1385 • To header tags MUST be generated for responses in a stateless manner - in a manner that will generate  
1386 the same tag for the same request consistently. For information on tag construction see Section 19.3.

1387 In all other respects, a stateless UAS behaves in the same manner as a stateful UAS. A UAS can operate  
1388 in either a stateful or stateless mode for each new request.

## 1389 8.3 Redirect Servers

1390 In some architectures it may be desirable to reduce the processing load on proxy servers that are responsible  
1391 for routing requests, and improve signaling path robustness, by relying on redirection. Redirection allows  
1392 servers to push routing information for a request back in a response to the client, thereby taking themselves  
1393 out of the loop of further messaging for this transaction while still aiding in locating the target of the request.  
1394 When the originator of the request receives the redirection, it will send a new request based on the URI(s)  
1395 it has received. By propagating URIs from the core of the network to its edges, redirection allows for  
1396 considerable network scalability.

1397 A redirect server is logically constituted of a server transaction layer and a transaction user that has  
1398 access to a location service of some kind (see Section 10 for more on registrars and location services). This  
1399 location service is effectively a database containing mappings between a single URI and a set of one or more  
1400 alternative locations at which the target of that URI can be found.

1401 A redirect server does not issue any SIP requests of its own. After receiving a request other than **CAN-**  
1402 **CEL**, the server either refuses the request or gathers the list of alternative locations from the location service  
1403 and returns a final response of class 3xx. For well-formed **CANCEL** requests, it **SHOULD** return a 2xx re-  
1404 sponse. This response ends the SIP transaction. The redirect server maintains transaction state for an entire  
1405 SIP transaction. It is the responsibility of clients to detect forwarding loops between redirect servers.

1406 When a redirect server returns a 3xx response to a request, it populates the list of (one or more) alter-  
1407 native locations into the **Contact** header field. An “**expires**” parameter to the **Contact** header field values  
1408 may also be supplied to indicate the lifetime of the **Contact** data.

1409 The **Contact** header field contains URIs giving the new locations or user names to try, or may simply  
1410 specify additional transport parameters. A 301 (Moved Permanently) or 302 (Moved Temporarily) response  
1411 may also give the same location and username that was targeted by the initial request but specify additional  
1412 transport parameters such as a different server or multicast address to try, or a change of SIP transport from  
1413 UDP to TCP or vice versa.

1414 However, redirect servers **MUST NOT** redirect a request to a URI equal to the one in the **Request-URI**;  
1415 instead, provided that the URI does not point to itself, the redirect server **SHOULD** proxy the request to the  
1416 destination URI.

1417 If a client is using an outbound proxy, and that proxy actually redirects requests, a potential arises for infinite  
1418 redirection loops.

1419 Note that a **Contact** header field value **MAY** also refer to a different resource than the one originally  
1420 called. For example, a SIP call connected to PSTN gateway may need to deliver a special informational  
1421 announcement such as “The number you have dialed has been changed.”

1422 A **Contact** response header field can contain any suitable URI indicating where the called party can be  
1423 reached, not limited to SIP URIs. For example, it could contain URIs for phones, fax, or irc (if they were  
1424 defined) or a **mailto:** (RFC 2368, [31]) URL. However, if the **Request-URI** of the request contained a SIPS  
1425 URI, the **Contact** header fields in the 3xx response **MUST** all be SIPS URIs.

1426 The “**expires**” parameter of a **Contact** header field value indicates how long the URI is valid. The value  
1427 of the parameter is a number indicating seconds. If this parameter is not provided, the value of the **Expires**  
1428 header field determines how long the URI is valid. Malformed values **SHOULD** be treated as equivalent to  
1429 3600.

1430 This provides a modest level of backwards compatibility with RFC 2543, which allowed absolute times in this  
1431 header field. If an absolute time is received, it will be treated as malformed, and then default to 3600.

1432 Redirect servers **MUST** ignore features that are not understood (including unrecognized header fields, any  
1433 unknown option tags in **Require**, or even method names) and proceed with the redirection of the request in  
1434 question.

## 1435 9 Canceling a Request

1436 The previous section has discussed general UA behavior for generating requests and processing responses  
1437 for requests of all methods. In this section, we discuss a general purpose method, called **CANCEL**.

1438 The **CANCEL** request, as the name implies, is used to cancel a previous request sent by a client. Specif-  
1439 ically, it asks the UAS to cease processing the request and to generate an error response to that request.

1440 CANCEL has no effect on a request to which a UAS has already given a final response. Because of this,  
1441 it is most useful to CANCEL requests to which it can take a server long time to respond. For this reason,  
1442 CANCEL is best for INVITE requests, which can take a long time to generate a response. In that usage,  
1443 a UAS that receives a CANCEL request for an INVITE, but has not yet sent a final response, would “stop  
1444 ringing”, and then respond to the INVITE with a specific error response (a 487).

1445 CANCEL requests can be constructed and sent by both proxies and user agent clients. Section 15  
1446 discusses under what conditions a UAC would CANCEL an INVITE request, and Section 16.10 discusses  
1447 proxy usage of CANCEL.

1448 A stateful proxy responds to a CANCEL, rather than simply forwarding a response it would receive  
1449 from a downstream element. For that reason, CANCEL is referred to as a “hop-by-hop” request, since it is  
1450 responded to at each stateful proxy hop.

## 1451 9.1 Client Behavior

1452 A CANCEL request SHOULD NOT be sent to cancel a request other than INVITE.

1453 Since requests other than INVITE are responded to immediately, sending a CANCEL for a non-INVITE request  
1454 would always create a race condition.

1455 The following procedures are used to construct a CANCEL request. The Request-URI, Call-ID, To,  
1456 the numeric part of CSeq, and From header fields in the CANCEL request MUST be identical to those in  
1457 the request being cancelled, including tags. A CANCEL constructed by a client MUST have only a single  
1458 Via header field value matching the top Via value in the request being cancelled. Using the same values  
1459 for these header fields allows the CANCEL to be matched with the request it cancels (Section 9.2 indicates  
1460 how such matching occurs). However, the method part of the CSeq header field MUST have a value of  
1461 CANCEL. This allows it to be identified and processed as a transaction in its own right (See Section 17).

1462 If the request being cancelled contains a Route header field, the CANCEL request MUST include that  
1463 Route header field's values.

1464 This is needed so that stateless proxies are able to route CANCEL requests properly.

1465 The CANCEL request MUST NOT contain any Require or Proxy-Require header fields.

1466 Once the CANCEL is constructed, the client SHOULD check whether it has received any response (pro-  
1467 visional or final) for the request being cancelled (herein referred to as the “original request”).

1468 If no provisional response has been received, the CANCEL request MUST NOT be sent; rather, the client  
1469 MUST wait for the arrival of a provisional response before sending the request. If the original request has  
1470 generated a final response, the CANCEL SHOULD NOT be sent, as it is an effective no-op, since CANCEL  
1471 has no effect on requests that have already generated a final response. When the client decides to send the  
1472 CANCEL, it creates a client transaction for the CANCEL and passes it the CANCEL request along with  
1473 the destination address, port, and transport. The destination address, port, and transport for the CANCEL  
1474 MUST be identical to those used to send the original request.

1475 If it was allowed to send the CANCEL before receiving a response for the previous request, the server could  
1476 receive the CANCEL before the original request.

1477 Note that both the transaction corresponding to the original request and the CANCEL transaction will  
1478 complete independently. However, a UAC canceling a request cannot rely on receiving a 487 (Request  
1479 Terminated) response for the original request, as an RFC 2543-compliant UAS will not generate such a  
1480 response. If there is no final response for the original request in 64\*T1 seconds (T1 is defined in Section  
1481 17.1.1.1), the client SHOULD then consider the original transaction cancelled and SHOULD destroy the client  
1482 transaction handling the original request.

## 1483 9.2 Server Behavior

1484 The CANCEL method requests that the TU at the server side cancel a pending transaction. The TU deter-  
1485 mines the transaction to be cancelled by taking the CANCEL request, and then assuming that the request  
1486 method is anything but CANCEL and applying the transaction matching procedures of Section 17.2.3. The  
1487 matching transaction is the one to be cancelled.

1488 The processing of a CANCEL request at a server depends on the type of server. A stateless proxy will  
1489 forward it, a stateful proxy might respond to it and generate some CANCEL requests of its own, and a UAS  
1490 will respond to it. See Section 16.10 for proxy treatment of CANCEL.

1491 A UAS first processes the CANCEL request according to the general UAS processing described in  
1492 Section 8.2. However, since CANCEL requests are hop-by-hop and cannot be resubmitted, they cannot be  
1493 challenged by the server in order to get proper credentials in an Authorization header field. Note also that  
1494 CANCEL requests do not contain a Require header field.

1495 If the UAS did not find a matching transaction for the CANCEL according to the procedure above, it  
1496 SHOULD respond to the CANCEL with a 481 (Call Leg/Transaction Does Not Exist). If the transaction  
1497 for the original request still exists, the behavior of the UAS on receiving a CANCEL request depends on  
1498 whether it has already sent a final response for the original request. If it has, the CANCEL request has no  
1499 effect on the processing of the original request, no effect on any session state, and no effect on the responses  
1500 generated for the original request. If the UAS has not issued a final response for the original request, its  
1501 behavior depends on the method of the original request. If the original request was an INVITE, the UAS  
1502 SHOULD immediately respond to the INVITE with a 487 (Request Terminated). The behavior upon reception  
1503 of a CANCEL request for any other method defined in this specification is effectively no-op.

1504 Regardless of the method of the original request, as long as the CANCEL matched an existing transac-  
1505 tion, the UAS answers the CANCEL request itself with a 200 (OK) response. This response is constructed  
1506 following the procedures described in Section 8.2.6 noting that the To tag of the response to the CANCEL  
1507 and the To tag in the response to the original request SHOULD be the same. The response to CANCEL is  
1508 passed to the server transaction for transmission.

## 1509 10 Registrations

### 1510 10.1 Overview

1511 SIP offers a discovery capability. If a user wants to initiate a session with another user, SIP must discover the  
1512 current host(s) at which the destination user is reachable. This discovery process is frequently accomplished  
1513 by SIP network elements such as proxy servers and redirect servers which are responsible for receiving a  
1514 request, determining where to send it based on knowledge of the location of the user, and then sending it  
1515 there. To do this, SIP network elements consult an abstract service known as a *location service*, which  
1516 provides address bindings for a particular domain. These address bindings map an incoming SIP or SIPS  
1517 URI, sip:bob@biloxi.com, for example, to one or more URIs that are somehow “closer” to the desired  
1518 user, sip:bob@engineering.biloxi.com, for example. Ultimately, a proxy will consult a location  
1519 service that maps a received URI to the user agent(s) at which the desired recipient is currently residing.

1520 Registration creates bindings in a location service for a particular domain that associate an address-of-  
1521 record URI with one or more contact addresses. Thus, when a proxy for that domain receives a request whose  
1522 Request-URI matches the address-of-record, the proxy will forward the request to the contact addresses  
1523 registered to that address-of-record. Generally, it only makes sense to register an address-of-record at a

1524 domain's location service when requests for that address-of-record would be routed to that domain. In  
 1525 most cases, this means that the domain of the registration will need to match the domain in the URI of the  
 1526 address-of-record.

1527 There are many ways by which the contents of the location service can be established. One way is  
 1528 administratively. In the above example, Bob is known to be a member of the engineering department through  
 1529 access to a corporate database. However, SIP provides a mechanism for a UA to create a binding explicitly.  
 1530 This mechanism is known as registration.

1531 Registration entails sending a REGISTER request to a special type of UAS known as a registrar. A  
 1532 registrar acts as the front end to the location service for a domain, reading and writing mappings based on  
 1533 the contents of REGISTER requests. This location service is then typically consulted by a proxy server that  
 1534 is responsible for routing requests for that domain.

1535 An illustration of the overall registration process is given in 2. Note that the registrar and proxy server  
 1536 are logical roles that can be played by a single device in a network; for purposes of clarity the two are  
 1537 separated in this illustration. Also note that UAs may send requests through a proxy server in order to reach  
 1538 a registrar if the two are separate elements.

1539 SIP does not mandate a particular mechanism for implementing the location service. The only require-  
 1540 ment is that a registrar for some domain MUST be able to read and write data to the location service, and  
 1541 a proxy or redirect server for that domain MUST be capable of reading that same data. A registrar MAY be  
 1542 co-located with a particular SIP proxy server for the same domain.

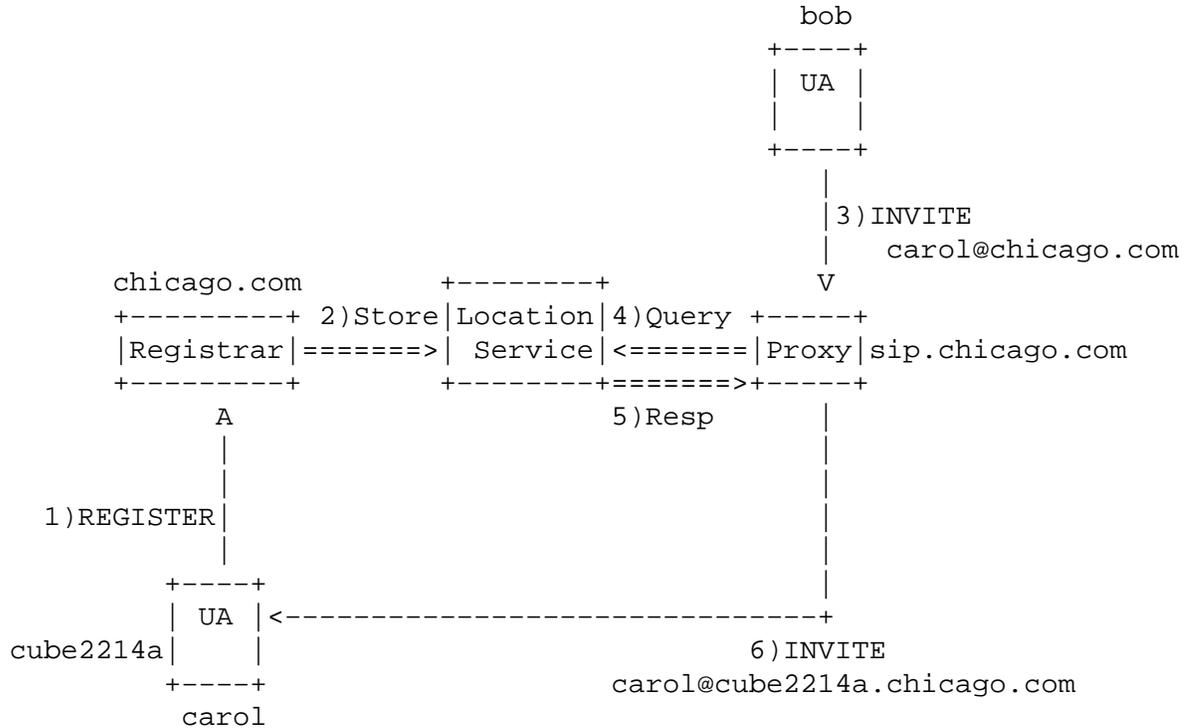


Figure 2: REGISTER example

## 1543 10.2 Constructing the REGISTER Request

1544 REGISTER requests add, remove, and query bindings. A REGISTER request can add a new binding  
1545 between an address-of-record and one or more contact addresses. Registration on behalf of a particular  
1546 address-of-record can be performed by a suitably authorized third party. A client can also remove previous  
1547 bindings or query to determine which bindings are currently in place for an address-of-record.

1548 Except as noted, the construction of the REGISTER request and the behavior of clients sending a  
1549 REGISTER request is identical to the general UAC behavior described in Section 8.1 and Section 17.1.

1550 A REGISTER request does *not* establish a dialog. A UAC MAY include a Route header field in a  
1551 REGISTER request based on a pre-existing route set as described in Section 8.1. The Record-Route  
1552 header field has no meaning in REGISTER requests or responses, and MUST be ignored if present. In  
1553 particular, the UAC MUST NOT create a new route set based on the presence or absence of a Record-Route  
1554 header field in any response to a REGISTER request.

1555 The following header fields, except Contact, MUST be included in a REGISTER request. A Contact  
1556 header field MAY be included:

1557 **Request-URI:** The Request-URI names the domain of the location service for which the registration is  
1558 meant (for example, "sip:chicago.com"). The "userinfo" and "@" components of the SIP URI MUST  
1559 NOT be present.

1560 **To:** The To header field contains the address of record whose registration is to be created, queried, or  
1561 modified. The To header field and the Request-URI field typically differ, as the former contains a  
1562 user name. This address-of-record MUST be a SIP URI or SIPS URI.

1563 **From:** The From header field contains the address-of-record of the person responsible for the registration.  
1564 The value is the same as the To header field unless the request is a third-party registration.

1565 **Call-ID:** All registrations from a UAC SHOULD use the same Call-ID header field value for registrations  
1566 sent to a particular registrar.

1567 If the same client were to use different Call-ID values, a registrar could not detect whether a delayed  
1568 REGISTER request might have arrived out of order.

1569 **CSeq:** The CSeq value guarantees proper ordering of REGISTER requests. A UA MUST increment the  
1570 CSeq value by one for each REGISTER request with the same Call-ID.

1571 **Contact:** REGISTER requests MAY contain a Contact header field with zero or more values containing  
1572 address bindings.

1573 UAs MUST NOT send a new registration (that is, containing new Contact header field values, as opposed  
1574 to a retransmission) until they have received a final response from the registrar for the previous one or the  
1575 previous REGISTER request has timed out.

1576 The following Contact header parameters have a special meaning in REGISTER requests:

1577 **action:** The "action" parameter from RFC 2543 has been deprecated. UACs SHOULD NOT use the  
1578 "action" parameter.

1579 **expires:** The “expires” parameter indicates how long the UA would like the binding to be valid. The value  
1580 is a number indicating seconds. If this parameter is not provided, the value of the Expires header field  
1581 is used instead. Implementations MAY treat values larger than 2\*\*32-1 (4294967295 seconds or 136  
1582 years) as equivalent to 2\*\*32-1. Malformed values SHOULD be treated as equivalent to 3600.

### 1583 10.2.1 Adding Bindings

1584 The REGISTER request sent to a registrar includes the contact address(es) to which SIP requests for the  
1585 address-of-record should be forwarded. The address-of-record is included in the To header field of the  
1586 REGISTER request.

1587 The Contact header field values of the request typically consist of SIP or SIPS URIs that identify  
1588 particular SIP endpoints (for example, “sip:carol@cube2214a.chicago.com”), but they MAY use any URI  
1589 scheme. A SIP UA can choose to register telephone numbers (with the tel URL, [9]) or email addresses  
1590 (with a mailto URL, [31]) as Contacts for an address-of-record, for example.

1591 For example, Carol, with address-of-record “sip:carol@chicago.com”, would register with the SIP reg-  
1592 istrar of the domain chicago.com. Her registrations would then be used by a proxy server in the chicago.com  
1593 domain to route requests for Carol’s address-of-record to her SIP endpoint.

1594 Once a client has established bindings at a registrar, it MAY send subsequent registrations containing  
1595 new bindings or modifications to existing bindings as necessary. The 2xx response to the REGISTER  
1596 request will contain, in a Contact header field, a complete list of bindings that have been registered for this  
1597 address-of-record at this registrar.

1598 If the address-of-record in the To header field of a REGISTER request is a SIPS URI, then any Contact  
1599 header field values in the request MUST also be a SIPS URIs.

1600 Registrations do not need to update all bindings. Typically, a UA only updates its own contact addresses.

1601 **10.2.1.1 Setting the Expiration Interval of Contact Addresses** When a client sends a REGISTER  
1602 request, it MAY suggest an expiration interval that indicates how long the client would like the registration  
1603 to be valid. (As described in Section 10.3, the registrar selects the actual time interval based on its local  
1604 policy.)

1605 There are two ways in which a client can suggest an expiration interval for a binding: through an  
1606 Expires header field or an “expires” Contact header parameter. The latter allows expiration intervals to  
1607 be suggested on a per-binding basis when more than one binding is given in a single REGISTER request,  
1608 whereas the former suggests an expiration interval for all Contact header field values that do not contain  
1609 the “expires” parameter.

1610 If neither mechanism for expressing a suggested expiration time is present in a REGISTER, a default  
1611 suggestion of one hour SHOULD be assumed.

1612 **10.2.1.2 Preferences among Contact Addresses** If more than one Contact is sent in a REGISTER  
1613 request, the registering UA intends to associate all of the URIs in these Contact header field values with the  
1614 address-of-record present in the To field. This list can be prioritized with the “q” parameter in the Contact  
1615 header field. The “q” parameter indicates a relative preference for the particular Contact header field value  
1616 compared to other bindings present in this REGISTER message or existing within the location service of  
1617 the registrar. Section 16.6 describes how a proxy server uses this preference indication.

### 1618 **10.2.2 Removing Bindings**

1619 Registrations are soft state and expire unless refreshed, but can also be explicitly removed. A client can  
1620 attempt to influence the expiration interval selected by the registrar as described in Section 10.2.1. A UA  
1621 requests the immediate removal of a binding by specifying an expiration interval of "0" for that contact  
1622 address in a REGISTER request. UAs SHOULD support this mechanism so that bindings can be removed  
1623 before their expiration interval has passed.

1624 The REGISTER-specific Contact header field value of "\*" applies to all registrations, but it MUST NOT  
1625 be used unless the Expires header field is present with a value of "0".

1626 Use of the "\*" Contact header field value allows a registering UA to remove all of its bindings without knowing  
1627 their precise values.

### 1628 **10.2.3 Fetching Bindings**

1629 A success response to any REGISTER request contains the complete list of existing bindings, regardless of  
1630 whether the request contained a Contact header field. If no Contact header field is present in a REGISTER  
1631 request, the list of bindings is left unchanged.

### 1632 **10.2.4 Refreshing Bindings**

1633 Each UA is responsible for refreshing the bindings that it has previously established. A UA SHOULD NOT  
1634 refresh bindings set up by other UAs.

1635 The 200 (OK) response from the registrar contains a list of Contact fields enumerating all current  
1636 bindings. The UA compares each contact address to see if it created the contact address, using comparison  
1637 rules in Section 19.1.4. If so, it updates the expiration time interval according to the expires parameter or,  
1638 if absent, the Expires field value. The UA then issues a REGISTER request for each of its bindings before  
1639 the expiration interval has elapsed. It MAY combine several updates into one REGISTER request.

1640 A UA SHOULD use the same Call-ID for all registrations during a single boot cycle. Registration re-  
1641 freshes SHOULD be sent to the same network address as the original registration, unless redirected.

### 1642 **10.2.5 Setting the Internal Clock**

1643 If the response for a REGISTER request contains a Date header field, the client MAY use this header field  
1644 to learn the current time in order to set any internal clocks.

### 1645 **10.2.6 Discovering a Registrar**

1646 UAs can use three ways to determine the address to which to send registrations: by configuration, using the  
1647 address-of-record, and multicast. A UA can be configured, in ways beyond the scope of this specification,  
1648 with a registrar address. If there is no configured registrar address, the UA SHOULD use the host part of the  
1649 address-of-record as the Request-URI and address the request there, using the normal SIP server location  
1650 mechanisms [4]. For example, the UA for the user "sip:carol@chicago.com" addresses the REGISTER  
1651 request to "sip:chicago.com".

1652 Finally, a UA can be configured to use multicast. Multicast registrations are addressed to the well-known  
1653 "all SIP servers" multicast address "sip.mcast.net" (224.0.1.75 for IPv4). No well-known IPv6 multicast  
1654 address has been allocated; such an allocation will be documented separately when needed. SIP UAs MAY

1655 listen to that address and use it to become aware of the location of other local users (see [32]); however, they  
1656 do not respond to the request.

1657           Multicast registration may be inappropriate in some environments, for example, if multiple businesses share the  
1658 same local area network.

### 1659 **10.2.7 Transmitting a Request**

1660 Once the REGISTER method has been constructed, and the destination of the message identified, UACs  
1661 follow the procedures described in Section 8.1.2 to hand off the REGISTER to the transaction layer.

1662           If the transaction layer returns a timeout error because the REGISTER yielded no response, the UAC  
1663 SHOULD NOT immediately re-attempt a registration to the same registrar.

1664           An immediate re-attempt is likely to also timeout. Waiting some reasonable time interval for the conditions  
1665 causing the timeout to be corrected reduces unnecessary load on the network. No specific interval is mandated.

### 1666 **10.2.8 Error Responses**

1667 If a UA receives a 423 (Interval Too Brief) response, it MAY retry the registration after making the expiration  
1668 interval of all contact addresses in the REGISTER request equal to or greater than the expiration interval  
1669 within the Min-Expires header field of the 423 (Interval Too Brief) response.

## 1670 **10.3 Processing REGISTER Requests**

1671 A registrar is a UAS that responds to REGISTER requests and maintains a list of bindings that are accessible  
1672 to proxy servers and redirect servers within its administrative domain. A registrar handles requests according  
1673 to Section 8.2 and Section 17.2, but it accepts only REGISTER requests. A registrar MUST not generate  
1674 6xx responses.

1675           A registrar MAY redirect REGISTER requests as appropriate. One common usage would be for a  
1676 registrar listening on a multicast interface to redirect multicast REGISTER requests to its own unicast  
1677 interface with a 302 (Moved Temporarily) response.

1678           Registrars MUST ignore the Record-Route header field if it is included in a REGISTER request. Reg-  
1679 istrars MUST NOT include a Record-Route header field in any response to a REGISTER request.

1680           A registrar might receive a request that traversed a proxy which treats REGISTER as an unknown request and  
1681 which added a Record-Route header field value.

1682           A registrar has to know (for example, through configuration) the set of domain(s) for which it maintains  
1683 bindings. REGISTER requests MUST be processed by a registrar in the order that they are received. REG-  
1684 ISTER requests MUST also be processed atomically, meaning that a particular REGISTER request is either  
1685 processed completely or not at all. Each REGISTER message MUST be processed independently of any  
1686 other registration or binding changes.

1687           When receiving a REGISTER request, a registrar follows these steps:

- 1688 1. The registrar inspects the Request-URI to determine whether it has access to bindings for the domain  
1689 identified in the Request-URI. If not, and if the server also acts as a proxy server, the server SHOULD  
1690 forward the request to the addressed domain, following the general behavior for proxying messages  
1691 described in Section 16.

- 1692 2. To guarantee that the registrar supports any necessary extensions, the registrar MUST process the  
1693 **Require** header field values as described for UASs in Section 8.2.2.
- 1694 3. A registrar SHOULD authenticate the UAC. Mechanisms for the authentication of SIP user agents  
1695 are described in Section 22. Registration behavior in no way overrides the generic authentication  
1696 framework for SIP. If no authentication mechanism is available, the registrar MAY take the **From**  
1697 address as the asserted identity of the originator of the request.
- 1698 4. The registrar SHOULD determine if the authenticated user is authorized to modify registrations for  
1699 this address-of-record. For example, a registrar might consult a authorization database that maps user  
1700 names to a list of addresses-of-record for which that user has authorization to modify bindings. If the  
1701 authenticated user is not authorized to modify bindings, the registrar MUST return a 403 (Forbidden)  
1702 and skip the remaining steps.
- 1703 In architectures that support third-party registration, one entity may be responsible for updating the regis-  
1704 trations associated with multiple addresses-of-record.
- 1705 5. The registrar extracts the address-of-record from the **To** header field of the request. If the address-of-  
1706 record is not valid for the domain in the **Request-URI**, the registrar MUST send a 404 (Not Found)  
1707 response and skip the remaining steps. The URI MUST then be converted to a canonical form. To do  
1708 that, all URI parameters MUST be removed (including the **user-param**), and any escaped characters  
1709 MUST be converted to their unescaped form. The result serves as an index into the list of bindings.
- 1710 6. The registrar checks whether the request contains the **Contact** header field. If not, it skips to the last  
1711 step. If the **Contact** header field is present, the registrar checks if there is one **Contact** field value  
1712 that contains the special value "\*" and an **Expires** field. If the request has additional **Contact** fields  
1713 or an expiration time other than zero, the request is invalid, and the server MUST return a 400 Invalid  
1714 Request and skip the remaining steps. If not, the registrar checks whether the **Call-ID** agrees with the  
1715 value stored for each binding. If not, it MUST remove the binding. If it does agree, it MUST remove  
1716 the binding only if the **CSeq** in the request is higher than the value stored for that binding. Otherwise  
1717 the registrar MUST leave the binding as is. It then skips to the last step.
- 1718 7. If the address-of-record in the **To** header field of the request represents a SIPS URI, then the registrar  
1719 MUST discard any **Contact** header field values that do not use the SIPS URI scheme before performing  
1720 any further processing.
- 1721 8. The registrar now processes each contact address in the **Contact** header field in turn. For each address,  
1722 it determines the expiration interval as follows:
- 1723 • If the field value has an "expires" parameter, that value MUST be used.
  - 1724 • If there is no such parameter, but the request has an **Expires** header field, that value MUST be  
1725 used.
  - 1726 • If there is neither, a locally-configured default value MUST be used.

1727 The registrar MAY shorten the expiration interval. If and only if the expiration interval is greater than  
1728 zero AND smaller than one hour AND less than a registrar-configured minimum, the registrar MAY  
1729 reject the registration with a response of 423 (Registration Too Brief). This response MUST contain a  
1730 **Min-Expires** header field that states the minimum expiration interval the registrar is willing to honor.  
1731 It then skips the remaining steps.

1732           Allowing the registrar to set the registration interval protects it against excessively frequent registration  
1733 refreshes while limiting the state that it needs to maintain and decreasing the likelihood of registrations going  
1734 stale. The expiration interval of a registration is frequently used in the creation of services. An example is a  
1735 follow-me service, where the user may only be available at a terminal for a brief period. Therefore, registrars  
1736 should accept brief registrations; a request should only be rejected if the interval is so short that the refreshes  
1737 would degrade registrar performance.

1738           For each address, the registrar then searches the list of current bindings using the URI comparison  
1739 rules. If the binding does not exist, it is tentatively added. If the binding does exist, the registrar  
1740 checks the Call-ID value. If the Call-ID value in the existing binding differs from the Call-ID value in  
1741 the request, the binding **MUST** be removed if the expiration time is zero and updated otherwise. If they  
1742 are the same, the registrar compares the CSeq value. If the value is higher than that of the existing  
1743 binding, it **MUST** update or remove the binding as above. If not, the update **MUST** be aborted and the  
1744 request fails.

1745           This algorithm ensures that out-of-order requests from the same UA are ignored.

1746           Each binding record records the Call-ID and CSeq values from the request.

1747           The binding updates **MUST** be committed (that is, made visible to the proxy or redirect server) if and  
1748 only if all binding updates and additions succeed. If any one of them fails (for example, because the  
1749 back-end database commit failed), the request **MUST** fail with a 500 (Server Error) response and all  
1750 tentative binding updates **MUST** be removed.

1751           9. The registrar returns a 200 (OK) response. The response **MUST** contain **Contact** header field values  
1752 enumerating all current bindings. Each **Contact** value **MUST** feature an “expires” parameter indi-  
1753 cating its expiration interval chosen by the registrar. The response **SHOULD** include a **Date** header  
1754 field.

## 1755 **11 Querying for Capabilities**

1756           The SIP method **OPTIONS** allows a UA to query another UA or a proxy server as to its capabilities. This  
1757 allows a client to discover information about the supported methods, content types, extensions, codecs, etc.  
1758 without “ringing” the other party. For example, before a client inserts a **Require** header field into an **INVITE**  
1759 listing an option that it is not certain the destination UAS supports, the client can query the destination UAS  
1760 with an **OPTIONS** to see if this option is returned in a **Supported** header field.

1761           The target of the **OPTIONS** request is identified by the **Request-URI**, which could identify another  
1762 UA or a SIP server. If the **OPTIONS** is addressed to a proxy server, the **Request-URI** is set without a user  
1763 part, similar to the way a **Request-URI** is set for a **REGISTER** request.

1764           Alternatively, a server receiving an **OPTIONS** request with a **Max-Forwards** header field value of 0  
1765 **MAY** respond to the request regardless of the **Request-URI**.

1766           This behavior is common with HTTP/1.1. This behavior can be used as a “traceroute” functionality to check the  
1767 capabilities of individual hop servers by sending a series of **OPTIONS** requests with incremented **Max-Forwards**  
1768 values.

1769           As is the case for general UA behavior, the transaction layer can return a timeout error if the **OPTIONS**  
1770 yields no response. This may indicate that the target is unreachable and hence unavailable.

1771           An **OPTIONS** request **MAY** be sent as part of an established dialog to query the peer on capabilities that  
1772 may be utilized later in the dialog.

## 1773 11.1 Construction of OPTIONS Request

1774 An OPTIONS request is constructed using the standard rules for a SIP request as discussed Section 8.1.1.

1775 A Contact header field MAY be present in an OPTIONS.

1776 An Accept header field SHOULD be included to indicate the type of message body the UAC wishes to  
1777 receive in the response. Typically, this is set to a format that is used to describe the media capabilities of a  
1778 UA, such as SDP (application/sdp).

1779 The response to an OPTIONS request is assumed to be scoped to the Request-URI in the original  
1780 request. However, only when an OPTIONS is sent as part of an established dialog is it guaranteed that  
1781 future requests will be received by the server that generated the OPTIONS response.

1782 Example OPTIONS request:

```
1783 OPTIONS sip:carol@chicago.com SIP/2.0
1784 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKhjhs8ass877
1785 Max-Forwards: 70
1786 To: <sip:carol@chicago.com>
1787 From: Alice <sip:alice@atlanta.com>;tag=1928301774
1788 Call-ID: a84b4c76e66710
1789 CSeq: 63104 OPTIONS
1790 Contact: <sip:alice@pc33.atlanta.com>
1791 Accept: application/sdp
1792 Content-Length: 0
```

## 1793 11.2 Processing of OPTIONS Request

1794 The response to an OPTIONS is constructed using the standard rules for a SIP response as discussed in  
1795 Section 8.2.6. The response code chosen MUST be the same that would have been chosen had the request  
1796 been an INVITE. That is, a 200 (OK) would be returned if the UAS is ready to accept a call, a 486 (Busy  
1797 Here) would be returned if the UAS is busy, etc. This allows an OPTIONS request to be used to determine  
1798 the basic state of a UAS, which can be an indication of whether the UAC will accept an INVITE request.

1799 An OPTIONS request received within a dialog generates a 200 (OK) response that is identical to one  
1800 constructed outside a dialog and does not have any impact on the dialog.

1801 This use of OPTIONS has limitations due the differences in proxy handling of OPTIONS and INVITE  
1802 requests. While a forked INVITE can result in multiple 200 (OK) responses being returned, a forked OP-  
1803 TIONS will only result in a single 200 (OK) response, since it is treated by proxies using the non-INVITE  
1804 handling. See Section 16.7 for the normative details.

1805 If the response to an OPTIONS is generated by a proxy server, the proxy returns a 200 (OK) listing the  
1806 capabilities of the server. The response does not contain a message body.

1807 Allow, Accept, Accept-Encoding, Accept-Language, and Supported header fields SHOULD be  
1808 present in a 200 (OK) response to an OPTIONS request. If the response is generated by a proxy, the  
1809 Allow header field SHOULD be omitted as it is ambiguous since a proxy is method agnostic. Contact header  
1810 fields MAY be present in a 200 (OK) response and have the same semantics as in a 3xx response. That is,  
1811 they may list a set of alternative names and methods of reaching the user. A Warning header field MAY be  
1812 present.

1813 A message body MAY be sent, the type of which is determined by the **Accept** header field in the **OP-**  
1814 **TIONS** request (application/sdp is the default if the **Accept** header field is not present). If the types include  
1815 one that can describe media capabilities, the UAS SHOULD include a body in the response for that purpose.  
1816 Details on construction of such a body in the case of application/sdp are described in [13].

1817 Example **OPTIONS** response generated by a UAS (corresponding to the request in Section 11.1):

```
1818 SIP/2.0 200 OK
1819 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKhjhs8ass877
1820 ;received=192.0.2.4
1821 To: <sip:carol@chicago.com>;tag=93810874
1822 From: Alice <sip:alice@atlanta.com>;tag=1928301774
1823 Call-ID: a84b4c76e66710
1824 CSeq: 63104 OPTIONS
1825 Contact: <sip:carol@chicago.com>
1826 Contact: <mailto:carol@chicago.com>
1827 Allow: INVITE, ACK, CANCEL, OPTIONS, BYE
1828 Accept: application/sdp
1829 Accept-Encoding: gzip
1830 Accept-Language: en
1831 Supported: foo
1832 Content-Type: application/sdp
1833 Content-Length: 274
1834
1835 (SDP not shown)
```

## 1836 12 Dialogs

1837 A key concept for a user agent is that of a dialog. A dialog represents a peer-to-peer SIP relationship between  
1838 two user agents that persists for some time. The dialog facilitates sequencing of messages between the user  
1839 agents and proper routing of requests between both of them. The dialog represents a context in which to  
1840 interpret SIP messages. Section 8 discussed method independent UA processing for requests and responses  
1841 outside of a dialog. This section discusses how those requests and responses are used to construct a dialog,  
1842 and then how subsequent requests and responses are sent within a dialog.

1843 A dialog is identified at each UA with a dialog ID, which consists of a **Call-ID** value, a local tag and a  
1844 remote tag. The dialog ID at each UA involved in the dialog is not the same. Specifically, the local tag at one  
1845 UA is identical to the remote tag at the peer UA. The tags are opaque tokens that facilitate the generation of  
1846 unique dialog IDs.

1847 A dialog ID is also associated with all responses and with any request that contains a tag in the **To** field.  
1848 The rules for computing the dialog ID of a message depend on whether the SIP element is a UAC or UAS.  
1849 For a UAC, the **Call-ID** value of the dialog ID is set to the **Call-ID** of the message, the remote tag is set to  
1850 the tag in the **To** field of the message, and the local tag is set to the tag in the **From** field of the message  
1851 (these rules apply to both requests and responses). As one would expect, for a UAS, the **Call-ID** value of the  
1852 dialog ID is set to the **Call-ID** of the message, the remote tag is set to the tag in the **From** field of the

1853 message, and the local tag is set to the tag in the **To** field of the message.

1854 A dialog contains certain pieces of state needed for further message transmissions within the dialog.  
1855 This state consists of the dialog ID, a local sequence number (used to order requests from the UA to its  
1856 peer), a remote sequence number (used to order requests from its peer to the UA), a local URI, a remote  
1857 URI, the Contact URI of the peer, a boolean flag called “secure”, and a route set, which is an ordered list of  
1858 URIs. The route set is the list of servers that need to be traversed to send a request to the peer. A dialog can  
1859 also be in the “early” state, which occurs when it is created with a provisional response, and then transition  
1860 to the “confirmed” state when a 2xx final response arrives. For other responses, or if no response arrives at  
1861 all on that dialog, the early dialog terminates.

## 1862 **12.1 Creation of a Dialog**

1863 Dialogs are created through the generation of non-failure responses to requests with specific methods.  
1864 Within this specification, only 2xx and 101-199 responses with a **To tag** to INVITE establish a dialog.  
1865 A dialog established by a non-final response to a request is in the “early” state and it is called an early dia-  
1866 log. Extensions MAY define other means for creating dialogs. Section 13 gives more details that are specific  
1867 to the INVITE method. Here, we describe the process for creation of dialog state that is not dependent on  
1868 the method.

1869 UAs MUST assign values to the dialog ID components as described below.

### 1870 **12.1.1 UAS behavior**

1871 When a UAS responds to a request with a response that establishes a dialog (such as a 2xx to INVITE),  
1872 the UAS MUST copy all Record-Route header field values from the request into the response (including  
1873 the URIs, URI parameters, and any Record-Route header field parameters, whether they are known or  
1874 unknown to the UAS) and MUST maintain the order of those values. The UAS MUST add a Contact header  
1875 field to the response. The Contact header field contains an address where the UAS would like to be con-  
1876 tacted for subsequent requests in the dialog (which includes the ACK for a 2xx response in the case of an  
1877 INVITE). Generally, the host portion of this URI is the IP address or FQDN of the host. The URI provided  
1878 in the Contact header field MUST be a SIP or SIPS URI. If the request which initiated the dialog contained  
1879 a SIPS URI in the Request-URI, the Contact header field MUST be a SIPS URI. In either case, the URI  
1880 SHOULD have global scope (that is, the same URI can be used in messages outside this dialog). The same  
1881 way, the scope of the URI in the Contact header field of the INVITE is not limited to this dialog either. It  
1882 can therefore be used in messages to the UAC even outside this dialog.

1883 The UAS then constructs the state of the dialog. This state MUST be maintained for the duration of the  
1884 dialog.

1885 If the request arrived over TLS, and the Request-URI contained a SIPS URI, the “secure” flag is set to  
1886 TRUE.

1887 The route set MUST be set to the list of URIs in the Record-Route header field from the request, taken  
1888 in order and preserving all URI parameters. If no Record-Route header field is present in the request, the  
1889 route set MUST be set to the empty set. This route set, even if empty, overrides any pre-existing route set for  
1890 future requests in this dialog. The remote target MUST be set to the URI from the Contact header field of  
1891 the request. If the “secure” flag is true, the UA MUST convert any SIP URI in the route set and remote target  
1892 to SIPS URI (this is done by just changing the scheme).

1893 The remote sequence number MUST be set to the value of the sequence number in the CSeq header field  
1894 of the request. The local sequence number MUST be empty. The call identifier component of the dialog ID

1895 MUST be set to the value of the **Call-ID** in the request. The local tag component of the dialog ID MUST be  
1896 set to the tag in the **To** field in the response to the request (which always includes a tag), and the remote tag  
1897 component of the dialog ID MUST be set to the tag from the **From** field in the request. A UAS MUST be  
1898 prepared to receive a request without a tag in the **From** field, in which case the tag is considered to have a  
1899 value of null.

1900 This is to maintain backwards compatibility with RFC 2543, which did not mandate **From** tags.

1901 The remote URI MUST be set to the URI in the **From** field, and the local URI MUST be set to the URI in  
1902 the **To** field.

### 1903 12.1.2 UAC Behavior

1904 When a UAC sends a request that can establish a dialog (such as an **INVITE**) it MUST provide a SIP or SIPS  
1905 URI with global scope (i.e., the same SIP URI can be used in messages outside this dialog) in the **Contact**  
1906 header field of the request. If the request is sent to a **Request-URI** with a SIPS URI, the **Contact** header  
1907 MUST be a SIPS URI.

1908 When a UAC receives a response that establishes a dialog, it constructs the state of the dialog. This state  
1909 MUST be maintained for the duration of the dialog.

1910 If the request was sent over TLS, and the **Request-URI** contained a SIPS URI, the “secure” flag is set  
1911 to **TRUE**.

1912 The route set MUST be set to the list of URIs in the **Record-Route** header field from the response,  
1913 taken in reverse order and preserving all URI parameters. If no **Record-Route** header field is present in  
1914 the response, the route set MUST be set to the empty set. This route set, even if empty, overrides any pre-  
1915 existing route set for future requests in this dialog. The remote target MUST be set to the URI from the  
1916 **Contact** header field of the response. If the “secure” flag is true, the UA MUST convert any SIP URI in the  
1917 route set and remote target to SIPS URI (this is done by just changing the scheme).

1918 The local sequence number MUST be set to the value of the sequence number in the **CSeq** header field  
1919 of the request. The remote sequence number MUST be empty (it is established when the remote UA sends  
1920 a request within the dialog). The call identifier component of the dialog ID MUST be set to the value of the  
1921 **Call-ID** in the request. The local tag component of the dialog ID MUST be set to the tag in the **From** field  
1922 in the request, and the remote tag component of the dialog ID MUST be set to the tag in the **To** field of the  
1923 response. A UAC MUST be prepared to receive a response without a tag in the **To** field, in which case the  
1924 tag is considered to have a value of null.

1925 This is to maintain backwards compatibility with RFC 2543, which did not mandate **To** tags.

1926 The remote URI MUST be set to the URI in the **To** field, and the local URI MUST be set to the URI in  
1927 the **From** field.

## 1928 12.2 Requests within a Dialog

1929 Once a dialog has been established between two UAs, either of them MAY initiate new transactions as needed  
1930 within the dialog. The UA sending the request will take the UAC role for the transaction. The UA receiving  
1931 the request will take the UAS role. Note that these may be different roles than the UAs held during the  
1932 transaction that established the dialog.

1933 Requests within a dialog MAY contain **Record-Route** and **Contact** header fields. However, these re-  
1934 quests do not cause the dialog’s route set to be modified, although they may modify the remote target URI.  
1935 Specifically, requests that are not target refresh requests do not modify the dialog’s remote target URI, and  
1936 requests that are target refresh requests do. For dialogs that have been established with an **INVITE**, the only  
1937 target refresh request defined is re-**INVITE** (see Section 14). Other extensions may define different target  
1938 refresh requests for dialogs established in other ways.

1939 Note that an ACK is *NOT* a target refresh request.  
1940 Target refresh requests only update the dialog's remote target URI, and not the route set formed from **Record-**  
1941 **Route**. Updating the latter would introduce severe backwards compatibility problems with RFC 2543-compliant  
1942 systems.

## 1943 12.2.1 UAC Behavior

1944 **12.2.1.1 Generating the Request** A request within a dialog is constructed by using many of the com-  
1945 ponents of the state stored as part of the dialog.

1946 The URI in the **To** field of the request **MUST** be set to the remote URI from the dialog state. The tag  
1947 in the **To** header field of the request **MUST** be set to the remote tag of the dialog ID. The **From** URI of the  
1948 request **MUST** be set to the local URI from the dialog state. The tag in the **From** header field of the request  
1949 **MUST** be set to the local tag of the dialog ID. If the value of the remote or local tags is null, the tag parameter  
1950 **MUST** be omitted from the **To** or **From** header fields, respectively.

1951 Usage of the URI from the **To** and **From** fields in the original request within subsequent requests is done for  
1952 backwards compatibility with RFC 2543, which used the URI for dialog identification. In this specification, only  
1953 the tags are used for dialog identification. It is expected that mandatory reflection of the original **To** and **From** URI  
1954 in mid-dialog requests will be deprecated in a subsequent revision of this specification.

1955 The **Call-ID** of the request **MUST** be set to the **Call-ID** of the dialog. Requests within a dialog **MUST**  
1956 contain strictly monotonically increasing and contiguous **CSeq** sequence numbers (increasing-by-one) in  
1957 each direction (excepting **ACK** and **CANCEL** of course, whose numbers equal the requests being acknowl-  
1958 edged or cancelled). Therefore, if the local sequence number is not empty, the value of the local sequence  
1959 number **MUST** be incremented by one, and this value **MUST** be placed into the **CSeq** header field. If the  
1960 local sequence number is empty, an initial value **MUST** be chosen using the guidelines of Section 8.1.1.5.  
1961 The method field in the **CSeq** header field value **MUST** match the method of the request.

1962 With a length of 32 bits, a client could generate, within a single call, one request a second for about 136 years  
1963 before needing to wrap around. The initial value of the sequence number is chosen so that subsequent requests within  
1964 the same call will not wrap around. A non-zero initial value allows clients to use a time-based initial sequence  
1965 number. A client could, for example, choose the 31 most significant bits of a 32-bit second clock as an initial  
1966 sequence number.

1967 The UAC uses the remote target and route set to build the **Request-URI** and **Route** header field of the  
1968 request.

1969 If the route set is empty, the UAC **MUST** place the remote target URI into the **Request-URI**. The UAC  
1970 **MUST NOT** add a **Route** header field to the request.

1971 If the route set is not empty, and the first URI in the route set contains the **lr** parameter (see Sec-  
1972 tion 19.1.1), the UAC **MUST** place the remote target URI into the **Request-URI** and **MUST** include a **Route**  
1973 header field containing the route set values in order, including all parameters.

1974 If the route set is not empty, and its first URI does not contain the **lr** parameter, the UAC **MUST** place  
1975 the first URI from the route set into the **Request-URI**, stripping any parameters that are not allowed in a  
1976 **Request-URI**. The UAC **MUST** add a **Route** header field containing the remainder of the route set values  
1977 in order, including all parameters. The UAC **MUST** then place the remote target URI into the **Route** header  
1978 field as the last value.

1979 For example, if the remote target is sip:user@remoteua and the route set contains

1980 <sip:proxy1> , <sip:proxy2> , <sip:proxy3;lr> , <sip:proxy4>

1981 The request will be formed with the following Request-URI and Route header field:

1982 METHOD sip:proxyl

1983 Route: < sip:proxy2> , < sip:proxy3;lr> , < sip:proxy4> , < sip:user@remoteua>

1984 If the first URI of the route set does not contain the lr parameter, the proxy indicated does not understand the  
1985 routing mechanisms described in this document and will act as specified in RFC 2543, replacing the Request-URI  
1986 with the first Route header field value it receives while forwarding the message. Placing the Request-URI at the  
1987 end of the Route header field preserves the information in that Request-URI across the strict router (it will be  
1988 returned to the Request-URI when the request reaches a loose-router).

1989 A UAC SHOULD include a Contact header field in any target refresh requests within a dialog, and unless  
1990 there is a need to change it, the URI SHOULD be the same as used in previous requests within the dialog. If  
1991 the "secure" flag is true, that URI MUST be a SIPS URI. As discussed in Section 12.2.2, a Contact header  
1992 field in a target refresh request updates the remote target URI. This allows a UA to provide a new contact  
1993 address, should its address change during the duration of the dialog.

1994 However, requests that are not target refresh requests do not affect the remote target URI for the dialog.  
1995 The rest of the request is formed as described in Section 8.1.1.

1996 Once the request has been constructed, the address of the server is computed and the request is sent,  
1997 using the same procedures for requests outside of a dialog (Section 8.1.2).

1998 The procedures in Section 8.1.2 will normally result in the request being sent to the address indicated by the  
1999 topmost Route header field value or the Request-URI if no Route header field is present. Subject to certain  
2000 restrictions, they allow the request to be sent to an alternate address (such as a default outbound proxy not represented  
2001 in the route set).

2002 **12.2.1.2 Processing the Responses** The UAC will receive responses to the request from the transaction  
2003 layer. If the client transaction returns a timeout this is treated as a 408 (Request Timeout) response.

2004 The behavior of a UAC that receives a 3xx response for a request sent within a dialog is the same as if  
2005 the request had been sent outside a dialog. This behavior is described in Section 8.1.3.4.

2006 Note, however, that when the UAC tries alternative locations, it still uses the route set for the dialog to build the  
2007 Route header of the request.

2008 When a UAC receives a 2xx response to a target refresh request, it MUST replace the dialog's remote  
2009 target URI with the URI from the Contact header field in that response, if present. If the "secure" flag is  
2010 true, the UAC MUST convert the URI to a SIPS URI if it is not one already.

2011 If the response for a request within a dialog is a 481 (Call/Transaction Does Not Exist) or a 408 (Request  
2012 Timeout), the UAC SHOULD terminate the dialog. A UAC SHOULD also terminate a dialog if no response  
2013 at all is received for the request (the client transaction would inform the TU about the timeout.)

2014 For INVITE initiated dialogs, terminating the dialog consists of sending a BYE.

## 2015 12.2.2 UAS Behavior

2016 Requests sent within a dialog, as any other requests, are atomic. If a particular request is accepted by the  
2017 UAS, *all* the state changes associated with it are performed. If the request is rejected, *none* of the state  
2018 changes is performed.

2019 Note that some requests such as INVITEs affect several pieces of state.

2020 The UAS will receive the request from the transaction layer. If the request has a tag in the To header  
2021 field, the UAS core computes the dialog identifier corresponding to the request and compares it with existing  
2022 dialogs. If there is a match, this is a mid-dialog request. In that case, the UAS first applies the same  
2023 processing rules for requests outside of a dialog, discussed in Section 8.2.

2024 If the request has a tag in the To header field, but the dialog identifier does not match any existing di-  
2025 alogs, the UAS may have crashed and restarted, or it may have received a request for a different (possibly  
2026 failed) UAS (the UASs can construct the To tags so that a UAS can identify that the tag was for a UAS  
2027 for which it is providing recovery). Another possibility is that the incoming request has been simply mis-  
2028 routed. Based on the To tag, the UAS MAY either accept or reject the request. Accepting the request for  
2029 acceptable To tags provides robustness, so that dialogs can persist even through crashes. UAs wishing to  
2030 support this capability must take into consideration some issues such as choosing monotonically increasing  
2031 CSeq sequence numbers even across reboots, reconstructing the route set, and accepting out-of-range RTP  
2032 timestamps and sequence numbers.

2033 If the UAS wishes to reject the request, because it does not wish to recreate the dialog, it MUST respond  
2034 to the request with a 481 (Call/Transaction Does Not Exist) status code and pass that to the server transaction.

2035 Requests that do not change in any way the state of a dialog may be received within a dialog (for  
2036 example, an OPTIONS request). They are processed as if they had been received outside the dialog.

2037 If the remote sequence number is empty, it MUST be set to the value of the sequence number in the CSeq  
2038 header field value in the request. If the remote sequence number was not empty, but the sequence number of  
2039 the request is lower than the remote sequence number, the request is out of order and MUST be rejected with  
2040 a 500 (Server Internal Error) response. If the remote sequence number was not empty, and the sequence  
2041 number of the request is greater than the remote sequence number, the request is in order. It is possible for  
2042 the CSeq sequence number to be higher than the remote sequence number by more than one. This is not  
2043 an error condition, and a UAS SHOULD be prepared to receive and process requests with CSeq values more  
2044 than one higher than the previous received request. The UAS MUST then set the remote sequence number to  
2045 the value of the sequence number in the CSeq header field value in the request.

2046 If a proxy challenges a request generated by the UAC, the UAC has to resubmit the request with credentials. The  
2047 resubmitted request will have a new CSeq number. The UAS will never see the first request, and thus, it will notice  
2048 a gap in the CSeq number space. Such a gap does not represent any error condition.

2049 When a UAS receives a target refresh request, it MUST replace the dialog's remote target URI with the  
2050 URI from the Contact header field in that request, if present. If the "secure" flag is true, the UAC MUST  
2051 convert the URI to a SIPS URI if it is not one already.

## 2052 12.3 Termination of a Dialog

2053 Independent of the method, if a request outside of a dialog generates a non-2xx final response, any early  
2054 dialogs created through provisional responses to that request are terminated. The mechanism for terminating  
2055 confirmed dialogs is method specific. In this specification, the BYE method terminates a session and the  
2056 dialog associated with it. See Section 15 for details.

## 2057 13 Initiating a Session

### 2058 13.1 Overview

2059 When a user agent client desires to initiate a session (for example, audio, video, or a game), it formulates an  
2060 INVITE request. The INVITE request asks a server to establish a session. This request may be forwarded by  
2061 proxies, eventually arriving at one or more UAS that can potentially accept the invitation. These UASs will

2062 frequently need to query the user about whether to accept the invitation. After some time, those UAS can  
2063 accept the invitation (meaning the session is to be established) by sending a 2xx response. If the invitation  
2064 is not accepted, a 3xx, 4xx, 5xx or 6xx response is sent, depending on the reason for the rejection. Before  
2065 sending a final response, the UAS can also send provisional responses (1xx) to advise the UAC of progress  
2066 in contacting the called user.

2067 After possibly receiving one or more provisional responses, the UAC will get one or more 2xx responses  
2068 or one non-2xx final response. Because of the protracted amount of time it can take to receive final responses  
2069 to INVITE, the reliability mechanisms for INVITE transactions differ from those of other requests (like  
2070 OPTIONS). Once it receives a final response, the UAC needs to send an ACK for every final response  
2071 it receives. The procedure for sending this ACK depends on the type of response. For final responses  
2072 between 300 and 699, the ACK processing is done in the transaction layer and follows one set of rules (See  
2073 Section 17). For 2xx responses, the ACK is generated by the UAC core.

2074 A 2xx response to an INVITE establishes a session, and it also creates a dialog between the UA that  
2075 issued the INVITE and the UA that generated the 2xx response. Therefore, when multiple 2xx responses are  
2076 received from different remote UAs (because the INVITE forked), each 2xx establishes a different dialog.  
2077 All these dialogs are part of the same call.

2078 This section provides details on the establishment of a session using INVITE. A UA that supports IN-  
2079 VITE MUST also support ACK, CANCEL and BYE.

## 2080 13.2 UAC Processing

### 2081 13.2.1 Creating the Initial INVITE

2082 Since the initial INVITE represents a request outside of a dialog, its construction follows the procedures of  
2083 Section 8.1.1. Additional processing is required for the specific case of INVITE.

2084 An Allow header field (Section 20.5) SHOULD be present in the INVITE. It indicates what methods can  
2085 be invoked within a dialog, on the UA sending the INVITE, for the duration of the dialog. For example, a  
2086 UA capable of receiving INFO requests within a dialog [33] SHOULD include an Allow header field listing  
2087 the INFO method.

2088 A Supported header field (Section 20.37) SHOULD be present in the INVITE. It enumerates all the  
2089 extensions understood by the UAC.

2090 An Accept (Section 20.1) header field MAY be present in the INVITE. It indicates which Content-Types  
2091 are acceptable to the UA, in both the response received by it, and in any subsequent requests sent to it within  
2092 dialogs established by the INVITE. The Accept header field is especially useful for indicating support of  
2093 various session description formats.

2094 The UAC MAY add an Expires header field (Section 20.19) to limit the validity of the invitation. If the  
2095 time indicated in the Expires header field is reached and no final answer for the INVITE has been received  
2096 the UAC core SHOULD generate a CANCEL request for the INVITE, as per Section 9.

2097 A UAC MAY also find it useful to add, among others, Subject (Section 20.36), Organization (Sec-  
2098 tion 20.25) and User-Agent (Section 20.41) header fields. They all contain information related to the  
2099 INVITE.

2100 The UAC MAY choose to add a message body to the INVITE. Section 8.1.1.10 deals with how to con-  
2101 struct the header fields – Content-Type among others – needed to describe the message body.

2102 There are special rules for message bodies that contain a session description - their corresponding  
2103 Content-Disposition is “session”. SIP uses an offer/answer model where one UA sends a session de-  
2104 scription, called the offer, which contains a proposed description of the session. The offer indicates the

2105 desired communications means (audio, video, games), parameters of those means (such as codec types) and  
2106 addresses for receiving media from the answerer. The other UA responds with another session description,  
2107 called the answer, which indicates which communications means are accepted, the parameters that apply to  
2108 those means, and addresses for receiving media from the offerer. The offer/answer model defines restric-  
2109 tions on when offers and answers can be made. This results in restrictions on where the offers and answers  
2110 can appear in SIP messages. In this specification, offers and answers can only appear in INVITE requests  
2111 and responses, and ACK. The usage of offers and answers is further restricted. For the initial INVITE  
2112 transaction, the rules are:

- 2113 • The initial offer MUST be in either an INVITE or, if not there, in the first reliable non-failure message  
2114 from the UAS back to the UAC. In this specification, that is the final 2xx response.
- 2115 • If the initial offer is in an INVITE, the answer MUST be in a reliable non-failure message from UAS  
2116 back to UAC which is correlated to that INVITE. For this specification, that is only the final 2xx  
2117 response to that INVITE.
- 2118 • If the initial offer is in the first reliable non-failure message from the UAS back to UAC, the answer  
2119 MUST be in the acknowledgement for that message (in this specification, ACK for a 2xx response).
- 2120 • After having sent or received an answer to the first offer, the UAC MAY generate subsequent offers  
2121 in requests, but only if it has received answers to any previous offers, and has not sent any offers to  
2122 which it hasn't gotten an answer.
- 2123 • Once the UAS has sent or received an answer to the initial offer, it MUST NOT generate subsequent  
2124 offers in any responses to the initial INVITE. This means that a UAS based on this specification alone  
2125 can never generate subsequent offers until completion of the initial transaction.

2126 Concretely, the above rules specify two exchanges - the offer is in the INVITE, and the answer in the  
2127 2xx, or the offer is in the 2xx, and the answer is in the ACK. All user agents that support INVITE MUST  
2128 support these two exchanges.

2129 The Session Description Protocol (SDP) [1] MUST be supported by all user agents as a means to describe  
2130 sessions, and its usage for constructing offers and answers MUST follow the procedures defined in [13].

2131 The restrictions of the offer-answer model just described only apply to bodies whose Content-Disposition  
2132 header field value is "session". Therefore, it is possible that both the INVITE and the ACK contain a body  
2133 message (for example, the INVITE carries a photo (Content-Disposition: render) and the ACK a session  
2134 description (Content-Disposition: session)).

2135 If the Content-Disposition header field is missing, bodies of Content-Type application/sdp imply the  
2136 disposition "session", while other content types imply "render".

2137 Once the INVITE has been created, the UAC follows the procedures defined for sending requests outside  
2138 of a dialog (Section 8). This results in the construction of a client transaction that will ultimately send the  
2139 request and deliver responses to the UAC.

### 2140 13.2.2 Processing INVITE Responses

2141 Once the INVITE has been passed to the INVITE client transaction, the UAC waits for responses for the  
2142 INVITE. If the INVITE client transaction returns a timeout rather than a response the TU acts as if a 408  
2143 (Request Timeout) response had been received, as described in Section 8.1.3.

2144 **13.2.2.1 1xx responses** Zero, one or multiple provisional responses may arrive before one or more  
2145 final responses are received. Provisional responses for an INVITE request can create “early dialogs”. If a  
2146 provisional response has a tag in the To field, and if the dialog ID of the response does not match an existing  
2147 dialog, one is constructed using the procedures defined in Section 12.1.2.

2148 The early dialog will only be needed if the UAC needs to send a request to its peer within the dialog  
2149 before the initial INVITE transaction completes. Header fields present in a provisional response are appli-  
2150 cable as long as the dialog is in the early state (for example, an Allow header field in a provisional response  
2151 contains the methods that can be used in the dialog while this is in the early state).

2152 **13.2.2.2 3xx responses** A 3xx response may contain one or more Contact header field values provid-  
2153 ing new addresses where the callee might be reachable. Depending on the status code of the 3xx response  
2154 (see Section 21.3) the UAC MAY choose to try those new addresses.

2155 **13.2.2.3 4xx, 5xx and 6xx responses** A single non-2xx final response may be received for the IN-  
2156 VITE. 4xx, 5xx and 6xx responses may contain a Contact header field value indicating the location where  
2157 additional information about the error can be found.

2158 All early dialogs are considered terminated upon reception of the non-2xx final response.

2159 After having received the non-2xx final response the UAC core considers the INVITE transaction com-  
2160 pleted. The INVITE client transaction handles generation of ACKs for the response (see Section 17).

2161 **13.2.2.4 2xx responses** Multiple 2xx responses may arrive at the UAC for a single INVITE request  
2162 due to a forking proxy. Each response is distinguished by the tag parameter in the To header field, and each  
2163 represents a distinct dialog, with a distinct dialog identifier.

2164 If the dialog identifier in the 2xx response matches the dialog identifier of an existing dialog, the dialog  
2165 MUST be transitioned to the “confirmed” state, and the route set for the dialog MUST be recomputed based  
2166 on the 2xx response using the procedures of Section 12.2.1.2. Otherwise, a new dialog in the “confirmed”  
2167 state MUST be constructed using the procedures of Section 12.1.2.

2168 Note that the only piece of state that is recomputed is the route set. Other pieces of state such as the highest  
2169 sequence numbers (remote and local) sent within the dialog are not recomputed. The route set only is recomputed  
2170 for backwards compatibility. RFC 2543 did not mandate mirroring of the Record-Route header field in a 1xx, only  
2171 2xx. However, we cannot update the entire state of the dialog, since mid-dialog requests may have been sent within  
2172 the early dialog, modifying the sequence numbers, for example.

2173 The UAC core MUST generate an ACK request for each 2xx received from the transaction layer. The  
2174 header fields of the ACK are constructed in the same way as for any request sent within a dialog (see  
2175 Section 12) with the exception of the CSeq and the header fields related to authentication. The sequence  
2176 number of the CSeq header field MUST be the same as the INVITE being acknowledged, but the CSeq  
2177 method MUST be ACK. The ACK MUST contain the same credentials as the INVITE. If the 2xx contains  
2178 an offer (based on the rules above), the ACK MUST carry an answer in its body. If the offer in the 2xx  
2179 response is not acceptable, the UAC core MUST generate a valid answer in the ACK and then send a BYE  
2180 immediately.

2181 Once the ACK has been constructed, the procedures of [4] are used to determine the destination address,  
2182 port and transport. However, the request is passed to the transport layer directly for transmission, rather than  
2183 a client transaction. This is because the UAC core handles retransmissions of the ACK, not the transaction  
2184 layer. The ACK MUST be passed to the client transport every time a retransmission of the 2xx final response  
2185 that triggered the ACK arrives.

2186 The UAC core considers the INVITE transaction completed 64\*T1 seconds after the reception of the  
2187 first 2xx response. At this point all the early dialogs that have not transitioned to established dialogs are  
2188 terminated. Once the INVITE transaction is considered completed by the UAC core, no more new 2xx  
2189 responses are expected to arrive.

2190 If, after acknowledging any 2xx response to an INVITE, the UAC does not want to continue with that  
2191 dialog, then the UAC MUST terminate the dialog by sending a BYE request as described in Section 15.

## 2192 13.3 UAS Processing

### 2193 13.3.1 Processing of the INVITE

2194 The UAS core will receive INVITE requests from the transaction layer. It first performs the request process-  
2195 ing procedures of Section 8.2, which are applied for both requests inside and outside of a dialog.

2196 Assuming these processing states complete without generating a response, the UAS core performs the  
2197 additional processing steps:

- 2198 1. If the request is an INVITE that contains an Expires header field the UAS core sets a timer for  
2199 the number of seconds indicated in the header field value. When the timer fires, the invitation is  
2200 considered to be expired. If the invitation expires before the UAS has generated a final response, a  
2201 487 (Request Terminated) response SHOULD be generated.
- 2202 2. If the request is a mid-dialog request, the method-independent processing described in Section 12.2.2  
2203 is first applied. It might also modify the session; Section 14 provides details.
- 2204 3. If the request has a tag in the To header field but the dialog identifier does not match any of the  
2205 existing dialogs, the UAS may have crashed and restarted, or may have received a request for a  
2206 different (possibly failed) UAS. Section 12.2.2 provides guidelines to achieve a robust behavior under  
2207 such a situation.

2208 Processing from here forward assumes that the INVITE is outside of a dialog, and is thus for the purposes  
2209 of establishing a new session.

2210 The INVITE may contain a session description, in which case the UAS is being presented with an offer  
2211 for that session. It is possible that the user is already a participant in that session, even though the INVITE  
2212 is outside of a dialog. This can happen when a user is invited to the same multicast conference by multiple  
2213 other participants. If desired, the UAS MAY use identifiers within the session description to detect this  
2214 duplication. For example, SDP contains a session id and version number in the origin (o) field. If the user  
2215 is already a member of the session, and the session parameters contained in the session description have  
2216 not changed, the UAS MAY silently accept the INVITE (that is, send a 2xx response without prompting the  
2217 user).

2218 If the INVITE does not contain a session description, the UAS is being asked to participate in a session,  
2219 and the UAC has asked that the UAS provide the offer of the session. It MUST provide the offer in its first  
2220 non-failure reliable message back to the UAC. In this specification, that is a 2xx response to the INVITE.

2221 The UAS can indicate progress, accept, redirect, or reject the invitation. In all of these cases, it formu-  
2222 lates a response using the procedures described in Section 8.2.6.

2223 **13.3.1.1 Progress** If the UAS is not able to answer the invitation immediately, it can choose to indicate  
2224 some kind of progress to the UAC (for example, an indication that a phone is ringing). This is accomplished

2225 with a provisional response between 101 and 199. These provisional responses establish early dialogs and  
2226 therefore follow the procedures of Section 12.1.1 in addition to those of Section 8.2.6. A UAS MAY send  
2227 as many provisional responses as it likes. Each of these MUST indicate the same dialog ID. However, these  
2228 will not be delivered reliably.

2229 If the UAS desires an extended period of time to answer the INVITE, it will need to ask for an “ex-  
2230 tension” in order to prevent proxies from canceling the transaction. A proxy has the option of canceling a  
2231 transaction when there is a gap of 3 minutes between messages in a transaction. To prevent cancellation, the  
2232 UAS MUST send a non-100 provisional response at every minute, to handle the possibility of lost provisional  
2233 responses.

2234 An INVITE transaction can go on for extended durations when the user is placed on hold, or when interworking  
2235 with PSTN systems which allow communications to take place without answering the call. The latter is common in  
2236 Interactive Voice Response (IVR) systems.

2237 **13.3.1.2 The INVITE is redirected** If the UAS decides to redirect the call, a 3xx response is sent. A  
2238 300 (Multiple Choices), 301 (Moved Permanently) or 302 (Moved Temporarily) response SHOULD contain  
2239 a **Contact** header field containing one or more URIs of new addresses to be tried. The response is passed to  
2240 the INVITE server transaction, which will deal with its retransmissions.

2241 **13.3.1.3 The INVITE is rejected** A common scenario occurs when the callee is currently not willing  
2242 or able to take additional calls at this end system. A 486 (Busy Here) SHOULD be returned in such scenario.  
2243 If the UAS knows that no other end system will be able to accept this call a 600 (Busy Everywhere) response  
2244 SHOULD be sent instead. However, it is unlikely that a UAS will be able to know this in general, and thus  
2245 this response will not usually be used. The response is passed to the INVITE server transaction, which will  
2246 deal with its retransmissions.

2247 A UAS rejecting an offer contained in an INVITE SHOULD return a 488 (Not Acceptable Here) response.  
2248 Such a response SHOULD include a **Warning** header field value explaining why the offer was rejected.

2249 **13.3.1.4 The INVITE is accepted** The UAS core generates a 2xx response. This response establishes  
2250 a dialog, and therefore follows the procedures of Section 12.1.1 in addition to those of Section 8.2.6.

2251 A 2xx response to an INVITE SHOULD contain the **Allow** header field and the **Supported** header field,  
2252 and MAY contain the **Accept** header field. Including these header fields allows the UAC to determine the  
2253 features and extensions supported by the UAS for the duration of the call, without probing.

2254 If the INVITE request contained an offer, and the UAS had not yet sent an answer, the 2xx MUST contain  
2255 an answer. If the INVITE did not contain an offer, the 2xx MUST contain an offer if the UAS had not yet  
2256 sent an offer.

2257 Once the response has been constructed it is passed to the INVITE server transaction. Note, however,  
2258 that the INVITE server transaction will be destroyed as soon as it receives this final response and passes it  
2259 to the transport. Therefore, it is necessary to pass periodically the response directly to the transport until  
2260 the **ACK** arrives. The 2xx response is passed to the transport with an interval that starts at T1 seconds and  
2261 doubles for each retransmission until it reaches T2 seconds (T1 and T2 are defined in Section 17). Response  
2262 retransmissions cease when an **ACK** request for the response is received. This is independent of whatever  
2263 transport protocols are used to send the response.

2264 Since 2xx is retransmitted end-to-end, there may be hops between UAS and UAC that are UDP. To ensure reliable  
2265 delivery across these hops, the response is retransmitted periodically even if the transport at the UAS is reliable.

2266 If the server retransmits the 2xx response for 64\*T1 seconds without receiving an ACK, the dialog is  
2267 confirmed, but the session SHOULD be terminated. This is accomplished with a BYE as described in Section  
2268 15.

## 2269 14 Modifying an Existing Session

2270 A successful INVITE request (see Section 13) establishes both a dialog between two user agents and a  
2271 session using the offer-answer model. Section 12 explains how to modify an existing dialog using a target  
2272 refresh request (for example, changing the remote target URI of the dialog). This section describes how  
2273 to modify the actual session. This modification can involve changing addresses or ports, adding a media  
2274 stream, deleting a media stream, and so on. This is accomplished by sending a new INVITE request within  
2275 the same dialog that established the session. An INVITE request sent within an existing dialog is known as  
2276 a re-INVITE.

2277 Note that a single re-INVITE can modify the dialog and the parameters of the session at the same time.

2278 Either the caller or callee can modify an existing session.

2279 The behavior of a UA on detection of media failure is a matter of local policy. However, automated  
2280 generation of re-INVITE or BYE is NOT RECOMMENDED to avoid flooding the network with traffic when  
2281 there is congestion. In any case, if these messages are sent automatically, they SHOULD be sent after some  
2282 randomized interval.

2283 Note that the paragraph above refers to automatically generated BYEs and re-INVITEs. If the user hangs up  
2284 upon media failure the UA would send a BYE request as usual.

### 2285 14.1 UAC Behavior

2286 The same offer-answer model that applies to session descriptions in INVITEs (Section 13.2.1) applies to  
2287 re-INVITEs. As a result, a UAC that wants to add a media stream, for example, will create a new offer that  
2288 contains this media stream, and send that in an INVITE request to its peer. It is important to note that the full  
2289 description of the session, not just the change, is sent. This supports stateless session processing in various  
2290 elements, and supports failover and recovery capabilities. Of course, a UAC MAY send a re-INVITE with no  
2291 session description, in which case the first reliable non-failure response to the re-INVITE will contain the  
2292 offer (in this specification, that is a 2xx response).

2293 If the session description format has the capability for version numbers, the offerer SHOULD indicate  
2294 that the version of the session description has changed.

2295 The To, From, Call-ID, CSeq, and Request-URI of a re-INVITE are set following the same rules as  
2296 for regular requests within an existing dialog, described in Section 12.

2297 A UAC MAY choose not to add an Alert-Info header field or a body with Content-Disposition "alert"  
2298 to re-INVITEs because UASs do not typically alert the user upon reception of a re-INVITE.

2299 Unlike an INVITE, which can fork, a re-INVITE will never fork, and therefore, only ever generate a  
2300 single final response. The reason a re-INVITE will never fork is that the Request-URI identifies the target  
2301 as the UA instance it established the dialog with, rather than identifying an address-of-record for the user.

2302 Note that a UAC MUST NOT initiate a new INVITE transaction within a dialog while another INVITE  
2303 transaction is in progress in either direction.

2304 1. If there is an ongoing INVITE client transaction, the TU MUST wait until the transaction reaches the  
2305 *completed* or *terminated* state before initiating the new INVITE.

- 2306 2. If there is an ongoing INVITE server transaction, the TU MUST wait until the transaction reaches the  
2307 *confirmed* or *terminated* state before initiating the new INVITE.

2308 However, a UA MAY initiate a regular transaction while an INVITE transaction is in progress. A UA  
2309 MAY also initiate an INVITE transaction while a regular transaction is in progress.

2310 If a UA receives a non-2xx final response to a re-INVITE, the session parameters MUST remain un-  
2311 changed, as if no re-INVITE had been issued. Note that, as stated in Section 12.2.1.2, if the non-2xx final  
2312 response is a 481 (Call/Transaction Does Not Exist), or a 408 (Request Timeout), or no response at all is  
2313 received for the re-INVITE (that is, a timeout is returned by the INVITE client transaction), the UAC will  
2314 terminate the dialog.

2315 The rules for transmitting a re-INVITE and for generating an ACK for a 2xx response to re-INVITE are  
2316 the same as for the initial INVITE (Section 13.2.1).

## 2317 14.2 UAS Behavior

2318 Section 13.3.1 describes the procedure for distinguishing incoming re-INVITEs from incoming initial IN-  
2319 VITEs and handling a re-INVITE for an existing dialog.

2320 A UAS that receives a second INVITE before it sends the final response to a first INVITE with a lower  
2321 CSeq sequence number on the same dialog MUST return a 500 (Server Internal Error) response to the second  
2322 INVITE and MUST include a Retry-After header field with a randomly chosen value of between 0 and 10  
2323 seconds.

2324 A UAS that receives an INVITE on a dialog while an INVITE it had sent on that dialog is in progress  
2325 MUST return a 491 (Request Pending) response to the received INVITE and MUST include a Retry-After  
2326 header field with a value chosen as follows:

- 2327 1. If the UAS is the owner of the Call-ID of the dialog ID (meaning it generated the value), the Retry-  
2328 After header field has a randomly chosen value of between 2.1 and 4 seconds in units of 10 ms.
- 2329 2. If the UAS is *not* the owner of the Call-ID of the dialog ID, the Retry-After header field has a ran-  
2330 domly chosen value of between 0 and 2 seconds in units of 10 ms.

2331 If a UA receives a re-INVITE for an existing dialog, it MUST check any version identifiers in the session  
2332 description or, if there are no version identifiers, the content of the session description to see if it has changed.  
2333 If the session description has changed, the UAS MUST adjust the session parameters accordingly, possibly  
2334 after asking the user for confirmation.

2335 Versioning of the session description can be used to accommodate the capabilities of new arrivals to a conference,  
2336 add or delete media, or change from a unicast to a multicast conference.

2337 If the new session description is not acceptable, the UAS can reject it by returning a 488 (Not Acceptable  
2338 Here) response for the re-INVITE. This response SHOULD include a Warning header field.

2339 If a UAS generates a 2xx response and never receives an ACK, it SHOULD generate a BYE to terminate  
2340 the dialog.

2341 A UAS MAY choose not to generate 180 (Ringing) responses for a re-INVITE because UACs do not  
2342 typically render this information to the user. For the same reason, UASs MAY choose not to use an Alert-  
2343 Info header field or a body with Content-Disposition "alert" in responses to a re-INVITE.

2344 A UAS providing an offer in a 2xx (because the INVITE did not contain an offer) SHOULD construct  
2345 the offer as if the UAS were making a brand new call, subject to the constraints of sending an offer that  
2346 updates an existing session, as described in [13] in the case of SDP. Specifically, this means that it SHOULD

2347 include as many media formats and media types that the UA is willing to support. The UAS MUST ensure  
2348 that the session description overlaps with its previous session description in media formats, transports, or  
2349 other parameters that require support from the peer. This is to avoid the need for the peer to reject the session  
2350 description. If, however, it is unacceptable to the UAC, the UAC SHOULD generate an answer with a valid  
2351 session description, and then send a BYE to terminate the session.

## 2352 15 Terminating a Session

2353 This section describes the procedures for terminating a session established by SIP. The state of the session  
2354 and the state of the dialog are very closely related. When a session is initiated with an INVITE, each 1xx or  
2355 2xx response from a distinct UAS creates a dialog, and if that response completes the offer/answer exchange,  
2356 it also creates a session. As a result, each session is "associated" with a single dialog - the one which resulted  
2357 in its creation. If an initial INVITE generates a non-2xx final response, that terminates all sessions (if any)  
2358 and all dialogs (if any) that were created through responses to the request. By virtue of completing the  
2359 transaction, a non-2xx final response also prevents further sessions from being created as a result of the  
2360 INVITE. The BYE request is used to terminate a specific session or attempted session. In this case, the  
2361 specific session is the one with the peer UA on the other side of the dialog. When a BYE is received on a  
2362 dialog, any session associated with that dialog SHOULD terminate. A UA MUST NOT send a BYE outside of  
2363 a dialog. The caller's UA MAY send a BYE for either confirmed or early dialogs, and the callee's UA MAY  
2364 send a BYE on confirmed dialogs, but MUST NOT send a BYE on early dialogs. However, the callee's UA  
2365 MUST NOT send a BYE on a confirmed dialog until it has received an ACK for its 2xx response or until the  
2366 server transaction times out. If no SIP extensions have defined other application layer state associated with  
2367 the dialog, the BYE also terminates the dialog.

2368 The impact of a non-2xx final response to INVITE on dialogs and sessions makes the use of CANCEL  
2369 attractive. The CANCEL attempts to force a non-2xx response to the INVITE (in particular, a 487). There-  
2370 fore, if a UAC wishes to give up on its call attempt entirely, it can send a CANCEL. If the INVITE results in  
2371 2xx final response(s) to the INVITE, this means that a UAS accepted the invitation while the CANCEL was  
2372 in progress. The UAC MAY continue with the sessions established by any 2xx responses, or MAY terminate  
2373 them with BYE.

2374 The notion of "hanging up" is not well defined within SIP. It is specific to a particular, albeit common, user  
2375 interface. Typically, when the user hangs up, it indicates a desire to terminate the attempt to establish a session, and  
2376 to terminate any sessions already created. For the caller's UA, this would imply a CANCEL request if the initial  
2377 INVITE has not generated a final response, and a BYE to all confirmed dialogs after a final response. For the callee's  
2378 UA, it would typically imply a BYE; presumably, when the user picked up the phone, a 2xx was generated, and so  
2379 hanging up would result in a BYE after the ACK is received. This does not mean a user cannot hang up before  
2380 receipt of the ACK, it just means that the software in his phone needs to maintain state for a short while in order to  
2381 clean up properly. If the particular UI allows for the user to reject a call before its answered, a 403 (Forbidden) is a  
2382 good way to express that. As per the rules above, a BYE can't be sent.

### 2383 15.1 Terminating a Session with a BYE Request

#### 2384 15.1.1 UAC Behavior

2385 A BYE request is constructed as would any other request within a dialog, as described in Section 12.

2386 Once the BYE is constructed, the UAC core creates a new non-INVITE client transaction, and passes it  
2387 the BYE request. The UAC MUST consider the session terminated (and therefore stop sending or listening  
2388 for media) as soon as the BYE request is passed to the client transaction. If the response for the BYE is a

2389 481 (Call/Transaction Does Not Exist) or a 408 (Request Timeout) or no response at all is received for the  
2390 BYE (that is, a timeout is returned by the client transaction), the UAC MUST consider the session and the  
2391 dialog terminated.

### 2392 15.1.2 UAS Behavior

2393 A UAS first processes the BYE request according to the general UAS processing described in Section 8.2.  
2394 A UAS core receiving a BYE request checks if it matches an existing dialog. If the BYE does not match an  
2395 existing dialog, the UAS core SHOULD generate a 481 (Call/Transaction Does Not Exist) response and pass  
2396 that to the server transaction.

2397 This rule means that a BYE sent without tags by a UAC will be rejected. This is a change from RFC 2543, which  
2398 allowed BYE without tags.

2399 A UAS core receiving a BYE request for an existing dialog MUST follow the procedures of Sec-  
2400 tion 12.2.2 to process the request. Once done, the UAS SHOULD terminate the session (and therefore stop  
2401 sending and listening for media). The only case where it can elect not to are multicast sessions, where par-  
2402 ticipation is possible even if the other participant in the dialog has terminated its involvement in the session.  
2403 Whether or not it ends its participation on the session, the UAS core MUST generate a 2xx response to the  
2404 BYE, and MUST pass that to the server transaction for transmission.

2405 The UAS MUST still respond to any pending requests received for that dialog. It is RECOMMENDED that  
2406 a 487 (Request Terminated) response is generated to those pending requests.

## 2407 16 Proxy Behavior

### 2408 16.1 Overview

2409 SIP proxies are elements that route SIP requests to user agent servers and SIP responses to user agent clients.  
2410 A request may traverse several proxies on its way to a UAS. Each will make routing decisions, modifying  
2411 the request before forwarding it to the next element. Responses will route through the same set of proxies  
2412 traversed by the request in the reverse order.

2413 Being a proxy is a logical role for a SIP element. When a request arrives, an element that can play the  
2414 role of a proxy first decides if it needs to respond to the request on its own. For instance, the request may be  
2415 malformed or the element may need credentials from the client before acting as a proxy. The element MAY  
2416 respond with any appropriate error code. When responding directly to a request, the element is playing the  
2417 role of a UAS and MUST behave as described in Section 8.2.

2418 A proxy can operate in either a stateful or stateless mode for each new request. When stateless, a proxy  
2419 acts as a simple forwarding element. It forwards each request downstream to a single element determined by  
2420 making a targeting and routing decision based on the request. It simply forwards every response it receives  
2421 upstream. A stateless proxy discards information about a message once the message has been forwarded.  
2422 A stateful proxy remembers information (specifically, transaction state) about each incoming request and  
2423 any requests it sends as a result of processing the incoming request. It uses this information to affect the  
2424 processing of future messages associated with that request. A stateful proxy MAY choose to “fork” a request,  
2425 routing it to multiple destinations. Any request that is forwarded to more than one location MUST be handled  
2426 statefully.

2427 In some circumstances, a proxy MAY forward requests using stateful transports (such as TCP) without  
2428 being transaction-stateful. For instance, a proxy MAY forward a request from one TCP connection to another

2429 transaction statelessly as long as it places enough information in the message to be able to forward the  
2430 response down the same connection the request arrived on. Requests forwarded between different types of  
2431 transports where the proxy's TU must take an active role in ensuring reliable delivery on one of the transports  
2432 MUST be forwarded transaction statefully.

2433 A stateful proxy MAY transition to stateless operation at any time during the processing of a request,  
2434 so long as it did not do anything that would otherwise prevent it from being stateless initially (forking, for  
2435 example, or generation of a 100 response). When performing such a transition, all state is simply discarded.  
2436 The proxy SHOULD NOT initiate a CANCEL request.

2437 Much of the processing involved when acting statelessly or statefully for a request is identical. The next  
2438 several subsections are written from the point of view of a stateful proxy. The last section calls out those  
2439 places where a stateless proxy behaves differently.

## 2440 16.2 Stateful Proxy

2441 When stateful, a proxy is purely a SIP transaction processing engine. Its behavior is modeled here in terms of  
2442 the server and client transactions defined in Section 17. A stateful proxy has a server transaction associated  
2443 with one or more client transactions by a higher layer proxy processing component (see figure 3), known as  
2444 a proxy core. An incoming request is processed by a server transaction. Requests from the server transaction  
2445 are passed to a proxy core. The proxy core determines where to route the request, choosing one or more  
2446 next-hop locations. An outgoing request for each next-hop location is processed by its own associated  
2447 client transaction. The proxy core collects the responses from the client transactions and uses them to send  
2448 responses to the server transaction.

2449 A stateful proxy creates a new server transaction for each new request received. Any retransmissions  
2450 of the request will then be handled by that server transaction per Section 17. The proxy core MUST behave  
2451 as a UAS with respect to sending an immediate provisional on that server transaction (such as 100 Trying)  
2452 as described in Section 8.2.6. Thus, a stateful proxy SHOULD NOT generate 100 Trying responses to non-  
2453 INVITE requests.

2454 This is a model of proxy behavior, not of software. An implementation is free to take any approach that  
2455 replicates the external behavior this model defines.

2456 For all new requests, including any with unknown methods, an element intending to proxy the request  
2457 MUST:

- 2458 1. Validate the request (Section 16.3)
- 2459 2. Preprocess routing information (Section 16.4)
- 2460 3. Determine target(s) for the request (Section 16.5)
- 2461 4. Forward the request to each target (Section 16.6)
- 2462 5. Process all responses (Section 16.7)

## 2463 16.3 Request Validation

2464 Before an element can proxy a request, it MUST verify the message's validity. A valid message must pass  
2465 the following checks:

- 2466 1. Reasonable Syntax

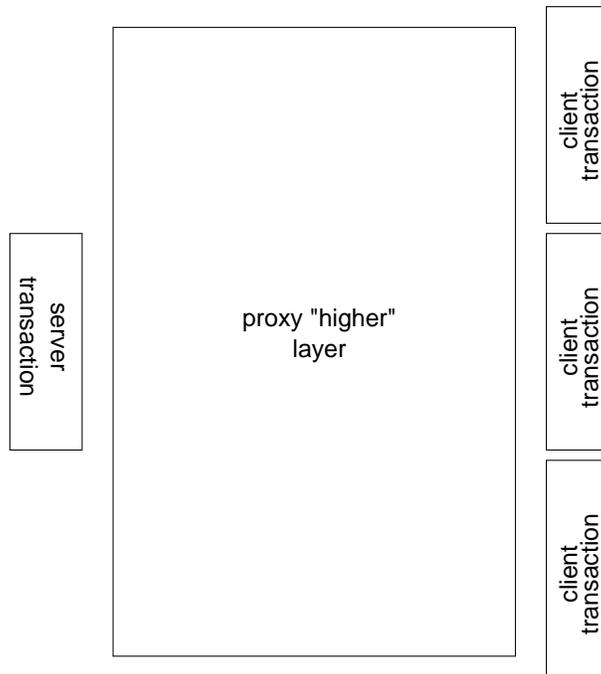


Figure 3: Stateful Proxy Model

- 2467 2. URI scheme
- 2468 3. Max-Forwards
- 2469 4. (Optional) Loop Detection
- 2470 5. Proxy-Require
- 2471 6. Proxy-Authorization

2472 If any of these checks fail, the element MUST behave as a user agent server (see Section 8.2) and respond  
2473 with an error code.

2474 Notice that a proxy is not required to detect merged requests and MUST NOT treat merged requests as an  
2475 error condition. The endpoints receiving the requests will resolve the merge as described in Section 8.2.2.2.

2476 1. Reasonable syntax check

2477 The request MUST be well-formed enough to be handled with a server transaction. Any components  
2478 involved in the remainder of these Request Validation steps or the Request Forwarding section MUST  
2479 be well-formed. Any other components, well-formed or not, SHOULD be ignored and remain un-  
2480 changed when the message is forwarded. For instance, an element would not reject a request because  
2481 of a malformed Date header field. Likewise, a proxy would not remove a malformed Date header  
2482 field before forwarding a request.

2483 This protocol is designed to be extended. Future extensions may define new methods and header fields  
2484 at any time. An element MUST NOT refuse to proxy a request because it contains a method or header  
2485 field it does not know about.

2486 2. URI scheme check

2487 If the Request-URI has a URI whose scheme is not understood by the proxy, the proxy SHOULD  
2488 reject the request with a 416 (Unsupported URI Scheme) response.

2489 3. Max-Forwards check

2490 The Max-Forwards header field (Section 20.22) is used to limit the number of elements a SIP request  
2491 can traverse.

2492 If the request does not contain a Max-Forwards header field, this check is passed.

2493 If the request contains a Max-Forwards header field with a field value greater than zero, the check is  
2494 passed.

2495 If the request contains a Max-Forwards header field with a field value of zero (0), the element MUST  
2496 NOT forward the request. If the request was for OPTIONS, the element MAY act as the final recipient  
2497 and respond per Section 11. Otherwise, the element MUST return a 483 (Too many hops) response.

2498 4. Optional Loop Detection check

2499 An element MAY check for forwarding loops before forwarding a request. If the request contains a  
2500 Via header field with a sent-by value that equals a value placed into previous requests by the proxy,  
2501 the request has been forwarded by this element before. The request has either looped or is legitimately  
2502 spiraling through the element. To determine if the request has looped, the element MAY perform the  
2503 branch parameter calculation described in Step 8 of Section 16.6 on this message and compare it to  
2504 the parameter received in that Via header field. If the parameters match, the request has looped. If  
2505 they differ, the request is spiraling, and processing continues. If a loop is detected, the element MAY  
2506 return a 482 (Loop Detected) response.

2507 5. Proxy-Require check

2508 Future extensions to this protocol may introduce features that require special handling by proxies.  
2509 Endpoints will include a Proxy-Require header field in requests that use these features, telling the  
2510 proxy not to process the request unless the feature is understood.

2511 If the request contains a Proxy-Require header field (Section 20.29) with one or more option-tags this  
2512 element does not understand, the element MUST return a 420 (Bad Extension) response. The response

2513 MUST include an **Unsupported** (Section 20.40) header field listing those option-tags the element did  
2514 not understand.

#### 2515 6. Proxy-Authorization check

2516 If an element requires credentials before forwarding a request, the request **MUST** be inspected as  
2517 described in Section 22.3. That section also defines what the element must do if the inspection fails.

### 2518 16.4 Route Information Preprocessing

2519 The proxy **MUST** inspect the **Request-URI** of the request. If the **Request-URI** of the request contains a  
2520 value this proxy previously placed into a **Record-Route** header field (see Section 16.6 item 4), the proxy  
2521 **MUST** replace the **Request-URI** in the request with the last value from the **Route** header field, and remove  
2522 that value from the **Route** header field. The proxy **MUST** then proceed as if it received this modified request.

2523 This will only happen when the element sending the request to the proxy (which may have been an endpoint)  
2524 is a strict router. This rewrite on receive is necessary to enable backwards compatibility with those elements. It  
2525 also allows elements following this specification to preserve the **Request-URI** through strict-routing proxies (see  
2526 Section 12.2.1.1).

2527 This requirement does not obligate a proxy to keep state in order to detect URIs it previously placed in **Record-  
2528 Route** header fields. Instead, a proxy need only place enough information in those URIs to recognize them as values  
2529 it provided when they later appear.

2530 If the **Request-URI** contains an **maddr** parameter, the proxy **MUST** check to see if its value is in the set  
2531 of addresses or domains the proxy is configured to be responsible for. If the **Request-URI** has an **maddr**  
2532 parameter with a value the proxy is responsible for, and the request was received using the port and transport  
2533 indicated (explicitly or by default) in the **Request-URI**, the proxy **MUST** strip the **maddr** and any non-default  
2534 port or transport parameter and continue processing as if those values had not been present in the request.

2535 A request may arrive with an **maddr** matching the proxy, but on a port or transport different from that indicated  
2536 in the URI. Such a request needs to be forwarded to the proxy using the indicated port and transport.

2537 If the first value in the **Route** header field indicates this proxy, the proxy **MUST** remove that value from  
2538 the request.

### 2539 16.5 Determining request targets

2540 Next, the proxy calculates the target(s) of the request. The set of targets will either be predetermined  
2541 by the contents of the request or will be obtained from an abstract location service. Each target in the set is  
2542 represented as a URI.

2543 If the **Request-URI** of the request contains an **maddr** parameter, the **Request-URI** **MUST** be placed  
2544 into the target set as the only target URI, and the proxy **MUST** proceed to Section 16.6.

2545 If the domain of the **Request-URI** indicates a domain this element is not responsible for, the **Request-  
2546 URI** **MUST** be placed into the target set as the only target, and the element **MUST** proceed to the task of  
2547 Request Forwarding (Section 16.6).

2548 There are many circumstances in which a proxy might receive a request for a domain it is not responsible for.  
2549 A firewall proxy handling outgoing calls (the way HTTP proxies handle outgoing requests) is an example of where  
2550 this is likely to occur.

2551 If the target set for the request has not been predetermined as described above, this implies that the  
2552 element is responsible for the domain in the Request-URI, and the element MAY use whatever mechanism  
2553 it desires to determine where to send the request. Any of these mechanisms can be modeled as accessing an  
2554 abstract Location Service. This may consist of obtaining information from a location service created by a SIP  
2555 Registrar, reading a database, consulting a presence server, utilizing other protocols, or simply performing  
2556 an algorithmic substitution on the Request-URI. When accessing the location service constructed by a  
2557 registrar, the Request-URI MUST first be canonicalized as described in Section 10.3 before being used as  
2558 an index. The output of these mechanisms is used to construct the target set. If the Request-URI contains  
2559 a SIPS URI, all elements in the target set MUST be SIPS URIs.

2560 If the Request-URI does not provide sufficient information for the proxy to determine the target set,  
2561 it SHOULD return a 485 (Ambiguous) response. This response SHOULD contain a Contact header field  
2562 containing URIs of new addresses to be tried. For example, an INVITE to sip:John.Smith@company.com  
2563 may be ambiguous at a proxy whose location service has multiple John Smiths listed. See Section 21.4.23  
2564 for details.

2565 Any information in or about the request or the current environment of the element MAY be used in the  
2566 construction of the target set. For instance, different sets may be constructed depending on contents or the  
2567 presence of header fields and bodies, the time of day of the request's arrival, the interface on which the  
2568 request arrived, failure of previous requests, or even the element's current level of utilization.

2569 As potential targets are located through these services, their URIs are added to the target set. Targets can  
2570 only be placed in the target set once. If a target URI is already present in the set (based on the definition of  
2571 equality for the URI type), it MUST NOT be added again.

2572 A proxy MUST NOT add additional targets to the target set if the Request-URI of the original request  
2573 does not indicate a resource this proxy is responsible for.

2574 A proxy can only change the Request-URI of a request during forwarding if it is responsible for that URI. If  
2575 the proxy is not responsible for that URI, it will not recurse on 3xx or 416 responses as described below.

2576 If the Request-URI of the original request indicates a resource this proxy is responsible for, the proxy  
2577 MAY continue to add targets to the set after beginning Request Forwarding. It MAY use any information  
2578 obtained during that processing to determine new targets. For instance, a proxy may choose to incorporate  
2579 contacts obtained in a redirect response (3xx) into the target set. If a proxy uses a dynamic source of  
2580 information while building the target set (for instance, if it consults a SIP Registrar), it SHOULD monitor  
2581 that source for the duration of processing the request. New locations SHOULD be added to the target set as  
2582 they become available. As above, any given URI MUST NOT be added to the set more than once.

2583 Allowing a URI to be added to the set only once reduces unnecessary network traffic, and in the case of incor-  
2584 porating contacts from redirect requests prevents infinite recursion.

2585 For example, a trivial location service is a "no-op", where the target URI is equal to the incoming request  
2586 URI. The request is sent to a specific next hop proxy for further processing. During request forwarding of  
2587 Section 16.6, Item 6, the identity of that next hop, expressed as a SIP or SIPS URI, is inserted as the top-most  
2588 Route header field value into the request.

2589 If the Request-URI indicates a resource at this proxy that does not exist, the proxy MUST return a 404  
2590 (Not Found) response.

2591 If the target set remains empty after applying all of the above, the proxy MUST return an error response,  
2592 which SHOULD be the 480 (Temporarily Unavailable) response.

## 2593 16.6 Request Forwarding

2594 As soon as the target set is non-empty, a proxy MAY begin forwarding the request. A stateful proxy MAY  
2595 process the set in any order. It MAY process multiple targets serially, allowing each client transaction to  
2596 complete before starting the next. It MAY start client transactions with every target in parallel. It also MAY  
2597 arbitrarily divide the set into groups, processing the groups serially and processing the targets in each group  
2598 in parallel.

2599 A common ordering mechanism is to use the qvalue parameter of targets obtained from Contact header  
2600 fields (see Section 20.10). Targets are processed from highest qvalue to lowest. Targets with equal qvalues  
2601 may be processed in parallel.

2602 A stateful proxy must have a mechanism to maintain the target set as responses are received and associate  
2603 the responses to each forwarded request with the original request. For the purposes of this model, this  
2604 mechanism is a "response context" created by the proxy layer before forwarding the first request.

2605 For each target, the proxy forwards the request following these steps:

- 2606 1. Make a copy of the received request
- 2607 2. Update the Request-URI
- 2608 3. Update the Max-Forwards header field
- 2609 4. Optionally add a Record-route header field value
- 2610 5. Optionally add additional header fields
- 2611 6. Postprocess routing information
- 2612 7. Determine the next-hop address, port, and transport
- 2613 8. Add a Via header field value
- 2614 9. Add a Content-Length header field if necessary
- 2615 10. Forward the new request
- 2616 11. Set timer C

2617 Each of these steps is detailed below:

### 2618 1. Copy request

2619 The proxy starts with a copy of the received request. The copy MUST initially contain all of the header  
2620 fields from the received request. Fields not detailed in the processing described below MUST NOT be  
2621 removed. The copy SHOULD maintain the ordering of the header fields as in the received request.  
2622 The proxy MUST NOT reorder field values with a common field name (See Section 7.3.1). The proxy  
2623 MUST NOT add to, modify, or remove the message body.

2624 An actual implementation need not perform a copy; the primary requirement is that the processing for each  
2625 next hop begin with the same request.

## 2626 2. Request-URI

2627 The **Request-URI** in the copy's start line **MUST** be replaced with the URI for this target. If the URI  
2628 contains any parameters not allowed in a Request-URI, they **MUST** be removed.

2629 This is the essence of a proxy's role. This is the mechanism through which a proxy routes a request  
2630 toward its destination.

2631 In some circumstances, the received **Request-URI** is placed into the target set without being modified.  
2632 For that target, the replacement above is effectively a no-op.

## 2633 3. Max-Forwards

2634 If the copy contains a **Max-Forwards** header field, the proxy **MUST** decrement its value by one (1).

2635 If the copy does not contain a **Max-Forwards** header field, the proxy **MUST** add one with a field value  
2636 which **SHOULD** be 70.

2637 Some existing UAs will not provide a **Max-Forwards** header field in a request.

## 2638 4. Record-Route

2639 If this proxy wishes to remain on the path of future requests in a dialog created by this request (as-  
2640 suming the request creates a dialog), it **MUST** insert a **Record-Route** header field value into the copy  
2641 before any existing **Record-Route** header field values, even if a **Route** header field is already present.

2642 Requests establishing a dialog may contain a preloaded **Route** header field.

2643 If this request is already part of a dialog, the proxy **SHOULD** insert a **Record-Route** header field value  
2644 if it wishes to remain on the path of future requests in the dialog. In normal endpoint operation as  
2645 described in Section 12 these **Record-Route** header field values will not have any effect on the route  
2646 sets used by the endpoints.

2647 The proxy will remain on the path if it chooses to not insert a **Record-Route** header field value into  
2648 requests that are already part of a dialog. However, it would be removed from the path when an endpoint that  
2649 has failed reconstitutes the dialog.

2650 A proxy **MAY** insert a **Record-Route** header field value into any request. If the request does not  
2651 initiate a dialog, the endpoints will ignore the value. See Section 12 for details on how endpoints use  
2652 the **Record-Route** header field values to construct **Route** header fields.

2653 Each proxy in the path of a request chooses whether to add a **Record-Route** header field value  
2654 independently - the presence of a **Record-Route** header field in a request does not obligate this proxy  
2655 to add a value.

2656 The URI placed in the **Record-Route** header field value **MUST** be a SIP URI. This URI **MUST** contain  
2657 an **lr** parameter (see Section 19.1.1). This URI **MAY** be different for each destination the request is  
2658 forwarded to. The URI **SHOULD NOT** contain the transport parameter unless the proxy has knowledge  
2659 (such as in a private network) that the next downstream element that will be in the path of subsequent  
2660 requests supports that transport.

2661 The URI this proxy provides will be used by some other element to make a routing decision. This proxy, in  
2662 general, has no way to know what the capabilities of that element are, so it must restrict itself to the mandatory  
2663 elements of a SIP implementation: SIP URIs and either the TCP or UDP transports.

2664 The URI placed in the **Record-Route** header field **MUST** resolve to the element inserting it (or a  
2665 suitable stand-in) when the server location procedures of [4] are applied to it, so that subsequent  
2666 requests reach the same SIP element. If the **Request-URI** contains a SIPS URI, the URI placed into  
2667 the **Record-Route** header field **MUST** be a SIPS URI.

2668 If the URI placed in the **Record-Route** header field needs to be rewritten when it passes back through  
2669 in a response, the URI **MUST** be distinct enough to locate at that time. (The request may spiral through  
2670 this proxy, resulting in more than one **Record-Route** header field value being added). Item 8 of  
2671 Section 16.7 recommends a mechanism to make the URI sufficiently distinct.

2672 The proxy **MAY** include parameters in the **Record-Route** header field value. These will be echoed in  
2673 some responses to the request such as the 200 (OK) responses to **INVITE**. Such parameters may be  
2674 useful for keeping state in the message rather than the proxy.

2675 If a proxy needs to be in the path of any type of dialog (such as one straddling a firewall), it **SHOULD**  
2676 add a **Record-Route** header field value to every request with a method it does not understand since  
2677 that method may have dialog semantics.

2678 The URI a proxy places into a **Record-Route** header field is only valid for the lifetime of any dialog  
2679 created by the transaction in which it occurs. A dialog-stateful proxy, for example, **MAY** refuse to  
2680 accept future requests with that value in the **Request-URI** after the dialog has terminated. Non-  
2681 dialog-stateful proxies, of course, have no concept of when the dialog has terminated, but they **MAY**  
2682 encode enough information in the value to compare it against the dialog identifier of future requests  
2683 and **MAY** reject requests not matching that information. Endpoints **MUST NOT** use a URI obtained  
2684 from a **Record-Route** header field outside the dialog in which it was provided. See Section 12 for  
2685 more information on an endpoint's use of **Record-Route** header fields.

2686 **Record-routing** may be required by certain services where the proxy needs to observe all messages  
2687 in a dialog. However, it slows down processing and impairs scalability and thus proxies should only  
2688 record-route if required for a particular service.

2689 The **Record-Route** process is designed to work for any SIP request that initiates a dialog. **INVITE** is  
2690 the only such request in this specification, but extensions to the protocol **MAY** define others.

## 2691 5. Add Additional Header Fields

2692 The proxy **MAY** add any other appropriate header fields to the copy at this point.

## 2693 6. Postprocess routing information

2694 A proxy **MAY** have a local policy that mandates that a request visit a specific set of proxies before being  
2695 delivered to the destination. A proxy **MUST** ensure that all such proxies are loose routers. Generally,  
2696 this can only be known with certainty if the proxies are within the same administrative domain. This  
2697 set of proxies is represented by a set of URIs (each of which contains the **lr** parameter). This set **MUST**  
2698 be pushed into the **Route** header field of the copy ahead of any existing values, if present. If the  
2699 **Route** header field is absent, it **MUST** be added, containing that list of URIs. If the **Request-URI**  
2700 specifies a SIPS URI, the set of URIs **MUST** all be converted to SIPS URI, if they were not already  
2701 SIPS URI.

2702 If the proxy has a local policy that mandates that the request visit one specific proxy, an alternative to  
2703 pushing a **Route** value into the **Route** header field is to bypass the forwarding logic of item 10 below,  
2704 and instead just send the request to the address, port, and transport for that specific proxy. If the

2705 request has a **Route** header field, this alternative **MUST NOT** be used unless it is known that next hop  
2706 proxy is a loose router. Otherwise, this approach **MAY** be used, but the **Route** insertion mechanism  
2707 above is preferred for its robustness, flexibility, generality and consistency of operation. Furthermore,  
2708 if the **Request-URI** contains a SIPS URI, TLS **MUST** be used to communicate with that proxy.

2709 If the copy contains a **Route** header field, the proxy **MUST** inspect the URI in its first value. If that  
2710 URI does not contain a **lr** parameter, the proxy **MUST** modify the copy as follows:

- 2711 • The proxy **MUST** place the **Request-URI** into the **Route** header field as the last value.
- 2712 • The proxy **MUST** then place the first **Route** header field value into the **Request-URI** and remove  
2713 that value from the **Route** header field.

2714 Appending the **Request-URI** to the **Route** header field is part of a mechanism used to pass the information  
2715 in that **Request-URI** through strict-routing elements. "Popping" the first **Route** header field value into the  
2716 **Request-URI** formats the message the way a strict-routing element expects to receive it (with its own URI in  
2717 the **Request-URI** and the next location to visit in the first **Route** header field value).

## 2718 7. Determine Next-Hop Address, Port, and Transport

2719 The proxy **MAY** have a local policy to send the request to a specific IP address, port, and transport,  
2720 independent of the values of the **Route** and **Request-URI**. Such a policy **MUST NOT** be used if the  
2721 proxy is not certain that the IP address, port, and transport correspond to a server that is a loose router.  
2722 However, this mechanism for sending the request through a specific next hop is **NOT RECOMMENDED**;  
2723 instead a **Route** header field should be used for that purpose as described above.

2724 In the absence of such an overriding mechanism, the proxy applies the procedures listed in [4] as  
2725 follows to determine where to send the request. If the proxy has reformatted the request to send to  
2726 a strict-routing element as described in step 6 above, the proxy **MUST** apply those procedures to the  
2727 **Request-URI** of the request. Otherwise, the proxy **MUST** apply the procedures to the first value in  
2728 the **Route** header field, if present, else the **Request-URI**. The procedures will produce an ordered set  
2729 of (address, port, transport) tuples.

2730 As described in [4], the proxy **MUST** attempt to deliver the message to the first tuple in that set, and  
2731 proceed through the set in order until the delivery attempt succeeds.

2732 For each tuple attempted, the proxy **MUST** format the message as appropriate for the tuple and send  
2733 the request using a new client transaction as detailed in steps 8 through 10. Since each attempt uses a  
2734 new client transaction, it represents a new branch. Thus, the branch parameter provided with the **Via**  
2735 header field inserted in step 8 **MUST** be different for each attempt.

2736 If the client transaction reports failure to send the request or a timeout from its state machine, the  
2737 proxy continues to the next address in that ordered set. If the ordered set is exhausted, the request  
2738 cannot be forwarded to this element in the target set. The proxy does not need to place anything in  
2739 the response context, but otherwise acts as if this element of the target set returned a 408 (Request  
2740 Timeout) final response.

## 2741 8. Add a **Via** header field value

2742 The proxy **MUST** insert a **Via** header field value into the copy before the existing **Via** header field  
2743 values. The construction of this value follows the same guidelines of Section 8.1.1.7. This implies

2744 that the proxy will compute its own branch parameter, which will be globally unique for that branch,  
2745 and contain the requisite magic cookie.

2746 Proxies choosing to detect loops have an additional constraint in the value they use for construction of  
2747 the branch parameter. A proxy choosing to detect loops SHOULD create a branch parameter separable  
2748 into two parts by the implementation. The first part MUST satisfy the constraints of Section 8.1.1.7 as  
2749 described above. The second is used to perform loop detection and distinguish loops from spirals.

2750 Loop detection is performed by verifying that, when a request returns to a proxy, those fields hav-  
2751 ing an impact on the processing of the request have not changed. The value placed in this part of  
2752 the branch parameter SHOULD reflect all of those fields (including any Route, Proxy-Require and  
2753 Proxy-Authorization header fields). This is to ensure that if the request is routed back to the proxy  
2754 and one of those fields changes, it is treated as a spiral and not a loop (Section 16.3 A common  
2755 way to create this value is to compute a cryptographic hash of the To tag, From tag, Call-ID header  
2756 field, the Request-URI of the request received (before translation) and the sequence number from  
2757 the CSeq header field, in addition to any Proxy-Require and Proxy-Authorization header fields that  
2758 may be present. The algorithm used to compute the hash is implementation-dependent, but MD5  
2759 [34], expressed in hexadecimal, is a reasonable choice. (Base64 is not permissible for a token.)

2760         If a proxy wishes to detect loops, the “branch” parameter it supplies MUST depend on all information  
2761 affecting processing of a request, including the incoming Request-URI and any header fields affecting the  
2762 request’s admission or routing. This is necessary to distinguish looped requests from requests whose routing  
2763 parameters have changed before returning to this server.

2764 The request method MUST NOT be included in the calculation of the branch parameter. In particular,  
2765 CANCEL and ACK requests (for non-2xx responses) MUST have the same branch value as the cor-  
2766 responding request they cancel or acknowledge. The branch parameter is used in correlating those  
2767 requests at the server handling them (see Sections 17.2.3 and 9.2).

#### 2768 9. Add a Content-Length header field if necessary

2769 If the request will be sent to the next hop using a stream-based transport and the copy contains no  
2770 Content-Length header field, the proxy MUST insert one with the correct value for the body of the  
2771 request (see Section 20.14).

#### 2772 10. Forward Request

2773 A stateful proxy MUST create a new client transaction for this request as described in Section 17.1 and  
2774 instructs the transaction to send the request using the address, port and transport determined in step 7.  
2775

#### 2776 11. Set timer C

2777 In order to handle the case where an INVITE request never generates a final response, the TU uses  
2778 a timer which is called timer C. Timer C MUST be set for each client transaction when an INVITE  
2779 request is proxied. The timer MUST be larger than 3 minutes. Section 16.7 bullet 2 discusses how this  
2780 timer is updated with provisional responses, and Section 16.8 discusses processing when it fires.

2781 **16.7 Response Processing**

2782 When a response is received by an element, it first tries to locate a client transaction (Section 17.1.3) match-  
2783 ing the response. If none is found, the element **MUST** process the response (even if it is an informational  
2784 response) as a stateless proxy (described below). If a match is found, the response is handed to the client  
2785 transaction.

2786 Forwarding responses for which a client transaction (or more generally any knowledge of having sent an associ-  
2787 ated request) is not found improves robustness. In particular, it ensures that “late” 2xx responses to INVITE requests  
2788 are forwarded properly.

2789 As client transactions pass responses to the proxy layer, the following processing **MUST** take place:

- 2790 1. Find the appropriate response context
- 2791 2. Update timer C for provisional responses
- 2792 3. Remove the topmost Via
- 2793 4. Add the response to the response context
- 2794 5. Check to see if this response should be forwarded immediately
- 2795 6. When necessary, choose the best final response from the response context

2796 If no final response has been forwarded after every client transaction associated with the response  
2797 context has been terminated, the proxy must choose and forward the “best” response from those it has  
2798 seen so far.

2799 The following processing **MUST** be performed on each response that is forwarded. It is likely that  
2800 more than one response to each request will be forwarded: at least each provisional and one final  
2801 response.

- 2802 7. Aggregate authorization header field values if necessary
- 2803 8. Optionally rewrite Record-Route header field values
- 2804 9. Forward the response
- 2805 10. Generate any necessary CANCEL requests

2806 Each of the above steps are detailed below:

- 2807 1. Find Context

2808 The proxy locates the “response context” it created before forwarding the original request using the  
2809 key described in Section 16.6. The remaining processing steps take place in this context.

- 2810 2. Update timer C for provisional responses

2811 For an INVITE transaction, if the response is a provisional response with status codes 101 to 199  
2812 inclusive (i.e., anything but 100), the proxy **MUST** reset timer C for that client transaction. The timer  
2813 **MAY** be reset to a different value, but this value **MUST** be greater than 3 minutes.

## 2814 3. Via

2815 The proxy removes the topmost *Via* header field value from the response.

2816 If no *Via* header field values remain in the response, the response was meant for this element and  
2817 MUST NOT be forwarded. The remainder of the processing described in this section is not performed  
2818 on this message, the UAC processing rules described in Section 8.1.3 are followed instead (transport  
2819 layer processing has already occurred).

2820 This will happen, for instance, when the element generates CANCEL requests as described in Sec-  
2821 tion 10.

## 2822 4. Add response to context

2823 Final responses received are stored in the response context until a final response is generated on the  
2824 server transaction associated with this context. The response may be a candidate for the best final  
2825 response to be returned on that server transaction. Information from this response may be needed in  
2826 forming the best response even if this response is not chosen.

2827 If the proxy chooses to recurse on any contacts in a 3xx response by adding them to the target set, it  
2828 MUST remove them from the response before adding the response to the response context. However,  
2829 a proxy MUST NOT recurse to a non-SIPS URI if the Request-URI of the original request was a SIPS  
2830 URI. If the proxy recurses on all of the contacts in a 3xx response, the proxy SHOULD NOT add the  
2831 resulting contactless response to the response context.

2832 Removing the contact before adding the response to the response context prevents the next element up-  
2833 stream from retrying a location this proxy has already attempted.

2834 3xx responses may contain a mixture of SIP, SIPS, and non-SIP URIs. A proxy may choose to recurse on  
2835 the SIP and SIPS URIs and place the remainder into the response context to be returned potentially in the final  
2836 response.

2837 If a proxy receives a 416 (Unsupported URI Scheme) response to a request whose Request-URI  
2838 scheme was not SIP, but the scheme in the original request was SIP or SIPS (that is, the  
2839 proxy changed the scheme from SIP or SIPS to something else when it proxied a request), the proxy  
2840 SHOULD add a new URI to the target set. This URI SHOULD be a SIP URI version of the non-SIP URI  
2841 that was just tried. In the case of the tel URL, this is accomplished by placing the telephone-subscriber  
2842 part of the tel URL into the user part of the SIP URI, and setting the hostpart to the domain where the  
2843 prior request was sent. See Section 19.1.6 for more detail on forming SIP URIs from tel URLs.

2844 As with a 3xx response, if a proxy “recurses” on the 416 by trying a SIP or SIPS URI instead, the 416  
2845 response SHOULD NOT be added to the response context.

## 2846 5. Check response for forwarding

2847 Until a final response has been sent on the server transaction, the following responses MUST be for-  
2848 forwarded immediately:

- 2849 ● Any provisional response other than 100 (Trying)
- 2850 ● Any 2xx response

2851 If a 6xx response is received, it is not immediately forwarded, but the stateful proxy SHOULD cancel  
2852 all client pending transactions as described in Section 10, and it MUST NOT create any new branches  
2853 in this context.

2854           This is a change from RFC 2543, which mandated that the proxy was to forward the 6xx response imme-  
2855           diately. For an INVITE transaction, this approach had the problem that a 2xx response could arrive on another  
2856           branch, in which case the proxy would have to forward the 2xx. The result was that the UAC could receive  
2857           a 6xx response followed by a 2xx response, which should never be allowed to happen. Under the new rules,  
2858           upon receiving a 6xx, a proxy will issue a CANCEL request, which will generally result in 487 responses from  
2859           all outstanding client transactions, and then at that point the 6xx is forwarded upstream.

2860           After a final response has been sent on the server transaction, the following responses MUST be for-  
2861           warded immediately:

- 2862           • Any 2xx response to an INVITE request

2863           A stateful proxy MUST NOT immediately forward any other responses. In particular, a stateful proxy  
2864           MUST NOT forward any 100 (Trying) response. Those responses that are candidates for forwarding  
2865           later as the “best” response have been gathered as described in step “Add Response to Context”.

2866           Any response chosen for immediate forwarding MUST be processed as described in steps “Aggregate  
2867           Authorization Header Field Values” through “Record-Route”.

2868           This step, combined with the next, ensures that a stateful proxy will forward exactly one final response  
2869           to a non-INVITE request, and either exactly one non-2xx response or one or more 2xx responses to  
2870           an INVITE request.

## 2871 6. Choosing the best response

2872           A stateful proxy MUST send a final response to a response context’s server transaction if no final  
2873           responses have been immediately forwarded by the above rules and all client transactions in this  
2874           response context have been terminated.

2875           The stateful proxy MUST choose the “best” final response among those received and stored in the  
2876           response context.

2877           If there are no final responses in the context, the proxy MUST send a 408 (Request Timeout) response  
2878           to the server transaction.

2879           Otherwise, the proxy MUST forward a response from the responses stored in the response context.  
2880           It MUST choose from the 6xx class responses if any exist in the context. If no 6xx class responses  
2881           are present, the proxy SHOULD choose from the lowest response class stored in the response context.  
2882           The proxy MAY select any response within that chosen class. The proxy SHOULD give preference to  
2883           responses that provide information affecting resubmission of this request, such as 401, 407, 415, 420,  
2884           and 484 if the 4xx class is chosen.

2885           A proxy which receives a 503 (Service Unavailable) response SHOULD NOT forward it upstream  
2886           unless it can determine that any subsequent requests it might proxy will also generate a 503. In other  
2887           words, forwarding a 503 means that the proxy knows it cannot service any requests, not just the one  
2888           for the Request-URI in the request which generated the 503.

2889           The forwarded response MUST be processed as described in steps “Aggregate Authorization Header  
2890           Field Values” through “Record-Route”.

2891           For example, if a proxy forwarded a request to 4 locations, and received 503, 407, 501, and 404  
2892           responses, it may choose to forward the 407 (Proxy Authentication Required) response.

2893           1xx and 2xx responses may be involved in the establishment of dialogs. When a request does not  
2894           contain a To tag, the To tag in the response is used by the UAC to distinguish multiple responses to

2895 a dialog creating request. A proxy **MUST NOT** insert a tag into the **To** header field of a 1xx or 2xx  
2896 response if the request did not contain one. A proxy **MUST NOT** modify the tag in the **To** header field  
2897 of a 1xx or 2xx response.

2898 Since a proxy may not insert a tag into the **To** header field of a 1xx response to a request that did not  
2899 contain one, it cannot issue non-100 provisional responses on its own. However, it can branch the  
2900 request to a UAS sharing the same element as the proxy. This UAS can return its own provisional  
2901 responses, entering into an early dialog with the initiator of the request. The UAS does not have to be  
2902 a discreet process from the proxy. It could be a virtual UAS implemented in the same code space as  
2903 the proxy.

2904 3-6xx responses are delivered hop-hop. When issuing a 3-6xx response, the element is effectively  
2905 acting as a UAS, issuing its own response, usually based on the responses received from downstream  
2906 elements. An element **SHOULD** preserve the **To** tag when simply forwarding a 3-6xx response to a  
2907 request that did not contain a **To** tag.

2908 A proxy **MUST NOT** modify the **To** tag in any forwarded response to a request that contains a **To** tag.

2909 While it makes no difference to the upstream elements if the proxy replaced the **To** tag in a forwarded  
2910 3-6xx response, preserving the original tag may assist with debugging.

2911 When the proxy is aggregating information from several responses, choosing a **To** tag from among them  
2912 is arbitrary, and generating a new **To** tag may make debugging easier. This happens, for instance, when  
2913 combining 401 (Unauthorized) and 407 (Proxy Authentication Required) challenges, or combining Contact  
2914 values from unencrypted and unauthenticated 3xx responses.

## 2915 7. Aggregate Authorization Header Field Values

2916 If the selected response is a 401 (Unauthorized) or 407 (Proxy Authentication Required), the proxy  
2917 **MUST** collect any **WWW-Authenticate** and **Proxy-Authenticate** header field values from all other  
2918 401 (Unauthorized) and 407 (Proxy Authentication Required) responses received so far in this re-  
2919 sponse context and add them to this response without modification before forwarding. The resulting  
2920 401 (Unauthorized) or 407 (Proxy Authentication Required) response could have several **WWW-**  
2921 **Authenticate** AND **Proxy-Authenticate** header field values.

2922 This is necessary because any or all of the destinations the request was forwarded to may have re-  
2923 quested credentials. The client needs to receive all of those challenges and supply credentials for each  
2924 of them when it retries the request. Motivation for this behavior is provided in Section 26.

## 2925 8. Record-Route

2926 If the selected response contains a **Record-Route** header field value originally provided by this proxy,  
2927 the proxy **MAY** choose to rewrite the value before forwarding the response. This allows the proxy to  
2928 provide different URIs for itself to the next upstream and downstream elements. A proxy may choose  
2929 to use this mechanism for any reason. For instance, it is useful for multi-homed hosts.

2930 The new URI provided by the proxy **MUST** satisfy the same constraints on URIs placed in **Record-**  
2931 **Route** header fields in requests (see Step 4 of Section 16.6) with the following modifications:

2932 The URI **SHOULD NOT** contain the transport parameter unless the proxy has knowledge that the next  
2933 upstream (as opposed to downstream) element that will be in the path of subsequent requests supports  
2934 that transport.

2935 When a proxy does decide to modify the **Record-Route** header field in the response, one of the  
2936 operations it performs is locating the **Record-Route** value that it had inserted. If the request spiraled,  
2937 and the proxy inserted a **Record-Route** value in each iteration of the spiral, locating the correct value  
2938 in the response (which must be the proper iteration in the reverse direction) is tricky. The rules above  
2939 recommend that a proxy wishing to rewrite **Record-Route** header field values insert sufficiently  
2940 distinct URIs into the **Record-Route** header field so that the right one may be selected for rewriting.  
2941 A RECOMMENDED mechanism to achieve this is for the proxy to append a unique identifier for the  
2942 proxy instance to the user portion of the URI.

2943 When the response arrives, the proxy modifies the first **Record-Route** whose identifier matches the  
2944 proxy instance. The modification results in a URI without this piece of data appended to the user  
2945 portion of the URI. Upon the next iteration, the same algorithm (find the topmost **Record-Route**  
2946 header field value with the parameter) will correctly extract the next **Record-Route** header field  
2947 value inserted by that proxy.

2948 Not every response to a request to which a proxy adds a **Record-Route** header field value will contain  
2949 a **Record-Route** header field. If the response does contain a **Record-Route** header field, it will contain the  
2950 value the proxy added.

## 2951 9. Forward response

2952 After performing the processing described in steps “Aggregate Authorization Header Field Values”  
2953 through “Record-Route”, the proxy MAY perform any feature specific manipulations on the selected  
2954 response. The proxy MUST NOT add to, modify, or remove the message body. Unless otherwise  
2955 specified, the proxy MUST NOT remove any header field values other than the **Via** header field value  
2956 discussed in Section 16.7 Item 3. In particular, the proxy MUST NOT remove any “received” pa-  
2957 rameter it may have added to the next **Via** header field value while processing the request associated  
2958 with this response. The proxy MUST pass the response to the server transaction associated with the  
2959 response context. This will result in the response being sent to the location now indicated in the top-  
2960 most **Via** header field value. If the server transaction is no longer available to handle the transmission,  
2961 the element MUST forward the response statelessly by sending it to the server transport. The server  
2962 transaction might indicate failure to send the response or signal a timeout in its state machine. These  
2963 errors would be logged for diagnostic purposes as appropriate, but the protocol requires no remedial  
2964 action from the proxy.

2965 The proxy MUST maintain the response context until all of its associated transactions have been ter-  
2966 minated, even after forwarding a final response.

## 2967 10. Generate CANCELs

2968 If the forwarded response was a final response, the proxy MUST generate a **CANCEL** request for all  
2969 pending client transactions associated with this response context. A proxy SHOULD also generate a  
2970 **CANCEL** request for all pending client transactions associated with this response context when it  
2971 receives a 6xx response. A pending client transaction is one that has received a provisional response,  
2972 but no final response (it is in the proceeding state) and has not had an associated **CANCEL** generated  
2973 for it. Generating **CANCEL** requests is described in Section 9.1.

2974 The requirement to **CANCEL** pending client transactions upon forwarding a final response does not  
2975 guarantee that an endpoint will not receive multiple 200 (OK) responses to an **INVITE**. 200 (OK)

2976 responses on more than one branch may be generated before the CANCEL requests can be sent and  
2977 processed. Further, it is reasonable to expect that a future extension may override this requirement to  
2978 issue CANCEL requests.

## 2979 **16.8 Processing Timer C**

2980 If timer C should fire, the proxy MUST either reset the timer with any value it chooses, or terminate the  
2981 client transaction. If the client transaction has received a provisional response, the proxy MUST generate a  
2982 CANCEL request matching that transaction. If the client transaction has not received a provisional response,  
2983 the proxy MUST behave as if the transaction received a 408 (Request Timeout) response.

2984 Allowing the proxy to reset the timer allows the proxy to dynamically extend the transaction's lifetime  
2985 based on current conditions (such as utilization) when the timer fires.

## 2986 **16.9 Handling Transport Errors**

2987 If the transport layer notifies a proxy of an error when it tries to forward a request (see Section 18.4), the  
2988 proxy MUST behave as if the forwarded request received a 400 (Bad Request) response.

2989 If the proxy is notified of an error when forwarding a response, it drops the response. The proxy SHOULD  
2990 NOT cancel any outstanding client transactions associated with this response context due to this notification.

2991 If a proxy cancels its outstanding client transactions, a single malicious or misbehaving client can cause all  
2992 transactions to fail through its Via header field.

## 2993 **16.10 CANCEL Processing**

2994 A stateful proxy MAY generate a CANCEL to any other request it has generated at any time (subject to re-  
2995 ceiving a provisional response to that request as described in section 9.1). A proxy MUST cancel any pending  
2996 client transactions associated with a response context when it receives a matching CANCEL request.

2997 A stateful proxy MAY generate CANCEL requests for pending INVITE client transactions based on the  
2998 period specified in the INVITE's Expires header field elapsing. However, this is generally unnecessary  
2999 since the endpoints involved will take care of signaling the end of the transaction.

3000 While a CANCEL request is handled in a stateful proxy by its own server transaction, a new response  
3001 context is not created for it. Instead, the proxy layer searches its existing response contexts for the server  
3002 transaction handling the request associated with this CANCEL. If a matching response context is found, the  
3003 element MUST immediately return a 200 (OK) response to the CANCEL request. In this case, the element is  
3004 acting as a user agent server as defined in Section 8.2. Furthermore, the element MUST generate CANCEL  
3005 requests for all pending client transactions in the context as described in Section 16.7 step 10.

3006 If a response context is not found, the element does not have any knowledge of the request to apply  
3007 the CANCEL to. It MUST statelessly forward the CANCEL request (it may have statelessly forwarded the  
3008 associated request previously).

## 3009 **16.11 Stateless Proxy**

3010 When acting statelessly, a proxy is a simple message forwarder. Much of the processing performed when  
3011 acting statelessly is the same as when behaving statefully. The differences are detailed here.

3012 A stateless proxy does not have any notion of a transaction, or of the response context used to describe  
3013 stateful proxy behavior. Instead, the stateless proxy takes messages, both requests and responses, directly

3014 from the transport layer (See section 18). As a result, stateless proxies do not retransmit messages on their  
3015 own. They do, however, forward all retransmission they receive (they do not have the ability to distinguish  
3016 a retransmission from the original message). Furthermore, when handling a request statelessly, an element  
3017 MUST NOT generate its own 100 (Trying) or any other provisional response.

3018 A stateless proxy MUST validate a request as described in Section 16.3

3019 A stateless proxy MUST follow the request processing steps described in Sections 16.4 through 16.5 with  
3020 the following exception:

- 3021 • A stateless proxy MUST choose one and only one target from the target set. This choice MUST only  
3022 rely on fields in the message and time-invariant properties of the server. In particular, a retransmitted  
3023 request MUST be forwarded to the same destination each time it is processed. Furthermore, CANCEL  
3024 and non-Routed ACK requests MUST generate the same choice as their associated INVITE.

3025 A stateless proxy MUST follow the request processing steps described in Section 16.6 with the following  
3026 exceptions:

- 3027 • The requirement for unique branch IDs across space and time applies to stateless proxies as well.  
3028 However, a stateless proxy cannot simply use a random number generator to compute the first com-  
3029 ponent of the branch ID, as described in Section 16.6 bullet 8. This is because retransmissions of  
3030 a request need to have the same value, and a stateless proxy cannot tell a retransmission from the  
3031 original request. Therefore, the component of the branch parameter that makes it unique MUST be  
3032 the same each time a retransmitted request is forwarded. Thus for a stateless proxy, the **branch** pa-  
3033 rameter MUST be computed as a combinatoric function of message parameters which are invariant on  
3034 retransmission.

3035 The stateless proxy MAY use any technique it likes to guarantee uniqueness of its branch IDs across  
3036 transactions. However, the following procedure is RECOMMENDED. The proxy examines the branch  
3037 ID in the topmost **Via** header field of the received request. If it begins with the magic cookie, the first  
3038 component of the branch ID of the outgoing request is computed as a hash of the received branch ID.  
3039 Otherwise, the first component of the branch ID is computed as a hash of the topmost **Via**, the tag in  
3040 the **To** header field, the tag in the **From** header field, the **Call-ID** header field, the **CSeq** number (but  
3041 not method), and the **Request-URI** from the received request. One of these fields will always vary  
3042 across two different transactions.

- 3043 • All other message transformations specified in Section 16.6 MUST result in the same transformation  
3044 of a retransmitted request. In particular, if the proxy inserts a **Record-Route** value or pushes URIs  
3045 into the **Route** header field, it MUST place the same values in retransmissions of the request. As  
3046 for the **Via** branch parameter, this implies that the transformations MUST be based on time-invariant  
3047 configuration or retransmission-invariant properties of the request.
- 3048 • A stateless proxy determines where to forward the request as described for stateful proxies in Sec-  
3049 tion 16.6 Item 10. The request is sent directly to the transport layer instead of through a client trans-  
3050 action.

3051 Since a stateless proxy must forward retransmitted requests to the same destination and add identical branch  
3052 parameters to each of them, it can only use information from the message itself and time-invariant configuration  
3053 data for those calculations. If the configuration state is not time-invariant (for example, if a routing table is updated)  
3054 any requests that could be affected by the change may not be forwarded statelessly during an interval equal to the  
3055 transaction timeout window before or after the change. The method of processing the affected requests in that  
3056 interval is an implementation decision. A common solution is to forward them transaction statefully.

3057 Stateless proxies MUST NOT perform special processing for CANCEL requests. They are processed by  
3058 the above rules as any other requests. In particular, a stateless proxy applies the same Route header field  
3059 processing to CANCEL requests that it applies to any other request.

3060 Response processing as described in Section 16.7 does not apply to a proxy behaving statelessly. When  
3061 a response arrives at a stateless proxy, the proxy MUST inspect the sent-by value in the first (topmost) Via  
3062 header field value. If that address matches the proxy (it equals a value this proxy has inserted into previous  
3063 requests) the proxy MUST remove that header field value from the response and forward the result to the  
3064 location indicated in the next Via header field value. The proxy MUST NOT add to, modify, or remove the  
3065 message body. Unless specified otherwise, the proxy MUST NOT remove any other header field values. If  
3066 the address does not match the proxy, the message MUST be silently discarded.

## 3067 16.12 Summary of Proxy Route Processing

3068 In the absence of local policy to the contrary, the processing a proxy performs on a request containing a  
3069 Route header field can be summarized in the following steps.

- 3070 1. The proxy will inspect the Request-URI. If it indicates a resource owned by this proxy, the proxy  
3071 will replace it with the results of running a location service. Otherwise, the proxy will not change the  
3072 Request-URI.
- 3073 2. The proxy will inspect the URI in the topmost Route header field value. If it indicates this proxy, the  
3074 proxy removes it from the Route header field (this route node has been reached).
- 3075 3. The proxy will forward the request to the resource indicated by the URI in the topmost Route header  
3076 field value or in the Request-URI if no Route header field is present. The proxy determines the  
3077 address, port and transport to use when forwarding the request by applying the procedures in [4] to  
3078 that URI.

3079 If no strict-routing elements are encountered on the path of the request, the Request-URI will always  
3080 indicate the target of the request.

### 3081 16.12.1 Examples

3082 16.12.1.1 Basic SIP Trapezoid This scenario is the basic SIP trapezoid, U1 -> P1 -> P2 -> U2, with  
3083 both proxies record-routing. Here is the flow.

3084 U1 sends:

```
3085 INVITE sip:callee@domain.com SIP/2.0  
3086 Contact: sip:caller@u1.example.com
```

3087 to P1. P1 is an outbound proxy. P1 is not responsible for domain.com, so it looks it up in DNS and  
3088 sends it there. It also adds a Record-Route header field value:

```
3089 INVITE sip:callee@domain.com SIP/2.0  
3090 Contact: sip:caller@u1.example.com  
3091 Record-Route: <sip:p1.example.com;lr>
```

3092 P2 gets this. It is responsible for domain.com so it runs a location service and rewrites the Request-  
3093 URI. It also adds a Record-Route header field value. There is no Route header field, so it resolves the new  
3094 Request-URI to determine where to send the request:

```
3095 INVITE sip:callee@u2.domain.com SIP/2.0
3096 Contact: sip:caller@u1.example.com
3097 Record-Route: <sip:p2.domain.com;lr>
3098 Record-Route: <sip:p1.example.com;lr>
```

3099 The callee at u2.domain.com gets this and responds with a 200 OK:

```
3100 SIP/2.0 200 OK
3101 Contact: sip:callee@u2.domain.com
3102 Record-Route: <sip:p2.domain.com;lr>
3103 Record-Route: <sip:p1.example.com;lr>
```

3104 The callee at u2 also sets its dialog state's remote target URI to sip:caller@u1.example.com and its route  
3105 set to

```
3106 (<sip:p2.domain.com;lr>, <sip:p1.example.com;lr>)
```

3107 This is forwarded by P2 to P1 to U1 as normal. Now, U1 sets its dialog state's remote target URI to  
3108 sip:callee@u2.domain.com and its route set to

```
3109 (<sip:p1.example.com;lr>, <sip:p2.domain.com;lr>)
```

3110 Since all the route set elements contain the lr parameter, U1 constructs the following BYE request:

```
3111 BYE sip:callee@u2.domain.com SIP/2.0
3112 Route: <sip:p1.example.com;lr>, <sip:p2.domain.com;lr>
```

3113 As any other element (including proxies) would do, it resolves the URI in the topmost Route header  
3114 field value using DNS to determine where to send the request. This goes to P1. P1 notices that it is not  
3115 responsible for the resource indicated in the Request-URI so it doesn't change it. It does see that it is the  
3116 first value in the Route header field, so it removes that value, and forwards the request to P2:

```
3117 BYE sip:callee@u2.domain.com SIP/2.0
3118 Route: <sip:p2.domain.com;lr>
```

3119 P2 also notices it is not responsible for the resource indicated by the Request-URI (it is responsible for  
3120 domain.com, not u2.domain.com), so it doesn't change it. It does see itself in the first Route header field  
3121 value, so it removes it and forwards the following to u2.domain.com based on a DNS lookup against the  
3122 Request-URI:

```
3123 BYE sip:callee@u2.domain.com SIP/2.0
```

3124 **16.12.1.2 Traversing a strict-routing proxy** In this scenario, a dialog is established across four prox-  
3125 ies, each of which adds **Record-Route** header field values. The third proxy implements the strict-routing  
3126 procedures specified in RFC 2543 and the bis drafts up to bis-05.

3127 U1->P1->P2->P3->P4->U2

3128 The INVITE arriving at U2 contains

```
3129 INVITE sip:callee@u2.domain.com SIP/2.0
3130 Contact: sip:caller@u1.example.com
3131 Record-Route: <sip:p4.domain.com;lr>
3132 Record-Route: <sip:p3.middle.com>
3133 Record-Route: <sip:p2.example.com;lr>
3134 Record-Route: <sip:p1.example.com;lr>
```

3135 Which U2 responds to with a 200 OK. Later, U2 sends the following **BYE** request to P4 based on the  
3136 first **Route** header field value.

```
3137 BYE sip:caller@u1.example.com SIP/2.0
3138 Route: <sip:p4.domain.com;lr>
3139 Route: <sip:p3.middle.com>
3140 Route: <sip:p2.example.com;lr>
3141 Route: <sip:p1.example.com;lr>
```

3142 P4 is not responsible for the resource indicated in the **Request-URI** so it will leave it alone. It notices  
3143 that it is the element in the first **Route** header field value so it removes it. It then prepares to send the request  
3144 based on the now first **Route** header field value of sip:p3.middle.com, but it notices that this URI does not  
3145 contain the lr parameter, so before sending, it reformats the request to be:

```
3146 BYE sip:p3.middle.com SIP/2.0
3147 Route: <sip:p2.example.com;lr>
3148 Route: <sip:p1.example.com;lr>
3149 Route: <sip:caller@u1.example.com>
```

3150 P3 is a strict router, so it forwards the following to P2:

```
3151 BYE sip:p2.example.com;lr SIP/2.0
3152 Route: <sip:p1.example.com;lr>
3153 Route: <sip:caller@u1.example.com>
```

3154 P2 sees the request-URI is a value it placed into a **Record-Route** header field, so before further pro-  
3155 cessing, it rewrites the request to be

```
3156 BYE sip:caller@u1.example.com SIP/2.0
3157 Route: <sip:p1.example.com;lr>
```

3158 P2 is not responsible for u1.example.com so it sends the request to P1 based on the resolution of the  
3159 Route header field value.

3160 P1 notices itself in the topmost Route header field value, so it removes it, resulting in:

3161 BYE sip:caller@u1.example.com SIP/2.0

3162 Since P1 is not responsible for u1.example.com and there is no Route header field, P1 will forward the  
3163 request to u1.example.com based on the Request-URI.

3164 **16.12.1.3 Rewriting Record-Route header field values** In this scenario, U1 and U2 are in different  
3165 private namespaces and they enter a dialog through a proxy P1, which acts as a gateway between the names-  
3166 paces.

3167 U1->P1->U2

3168 U1 sends:

3169 INVITE sip:callee@gateway.leftprivatespace.com SIP/2.0  
3170 Contact: <sip:caller@u1.leftprivatespace.com>

3171 P1 uses its location service and sends the following to U2:

3172 INVITE sip:callee@rightprivatespace.com SIP/2.0  
3173 Contact: <sip:caller@u1.leftprivatespace.com>  
3174 Record-Route: <sip:gateway.rightprivatespace.com;lr>

3175 U2 sends this 200 (OK) back to P1:

3176 SIP/2.0 200 OK  
3177 Contact: <sip:callee@u2.rightprivatespace.com>  
3178 Record-Route: <sip:gateway.rightprivatespace.com;lr>

3179 P1 rewrites its Record-Route header parameter to provide a value that U1 will find useful, and sends  
3180 the following to U1:

3181 SIP/2.0 200 OK  
3182 Contact: <sip:callee@u2.rightprivatespace.com>  
3183 Record-Route: <sip:gateway.leftprivatespace.com;lr>

3184 Later, U1 sends the following BYE request to P1:

3185 BYE sip:callee@u2.rightprivatespace.com SIP/2.0  
3186 Route: <sip:gateway.leftprivatespace.com;lr>

3187 which P1 forwards to U2 as

3188 BYE sip:callee@u2.rightprivatespace.com SIP/2.0

3189 **17 Transactions**

3190 SIP is a transactional protocol: interactions between components take place in a series of independent  
 3191 message exchanges. Specifically, a SIP transaction consists of a single request and any responses to that  
 3192 request, which include zero or more provisional responses and one or more final responses. In the case  
 3193 of a transaction where the request was an INVITE (known as an INVITE transaction), the transaction also  
 3194 includes the ACK only if the final response was not a 2xx response. If the response was a 2xx, the ACK is  
 3195 not considered part of the transaction.

3196 The reason for this separation is rooted in the importance of delivering all 200 (OK) responses to an INVITE  
 3197 to the UAC. To deliver them all to the UAC, the UAS alone takes responsibility for retransmitting them (see Sec-  
 3198 tion 13.3.1.4), and the UAC alone takes responsibility for acknowledging them with ACK (see Section 13.2.2.4).  
 3199 Since this ACK is retransmitted only by the UAC, it is effectively considered its own transaction.

3200 Transactions have a client side and a server side. The client side is known as a client transaction and the  
 3201 server side as a server transaction. The client transaction sends the request, and the server transaction sends  
 3202 the response. The client and server transactions are logical functions that are embedded in any number of  
 3203 elements. Specifically, they exist within user agents and stateful proxy servers. Consider the example in  
 3204 Section 4. In this example, the UAC executes the client transaction, and its outbound proxy executes the  
 3205 server transaction. The outbound proxy also executes a client transaction, which sends the request to a  
 3206 server transaction in the inbound proxy. That proxy also executes a client transaction, which in turn sends  
 3207 the request to a server transaction in the UAS. This is shown in Figure 4.

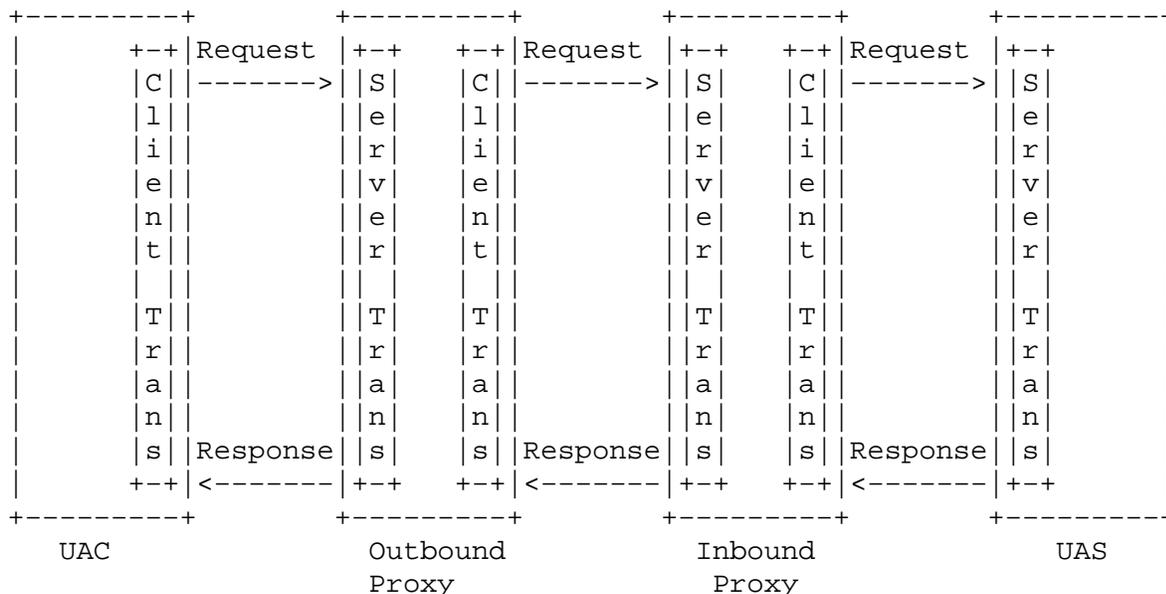


Figure 4: Transaction relationships

3208 A stateless proxy does not contain a client or server transaction. The transaction exists between the UA  
3209 or stateful proxy on one side, and the UA or stateful proxy on the other side. As far as SIP transactions are  
3210 concerned, stateless proxies are effectively transparent. The purpose of the client transaction is to receive  
3211 a request from the element in which the client is embedded (call this element the "Transaction User" or  
3212 TU; it can be a UA or a stateful proxy), and reliably deliver the request to a server transaction. The client  
3213 transaction is also responsible for receiving responses and delivering them to the TU, filtering out any re-  
3214 sponse retransmissions or disallowed responses (such as a response to ACK). Additionally, in the case of an  
3215 INVITE request, the client transaction is responsible for generating the ACK request for any final response  
3216 excepting a 2xx response.

3217 Similarly, the purpose of the server transaction is to receive requests from the transport layer and deliver  
3218 them to the TU. The server transaction filters any request retransmissions from the network. The server  
3219 transaction accepts responses from the TU and delivers them to the transport layer for transmission over the  
3220 network. In the case of an INVITE transaction, it absorbs the ACK request for any final response excepting  
3221 a 2xx response.

3222 The 2xx response and its ACK receive special treatment. This response is retransmitted only by a UAS,  
3223 and its ACK generated only by the UAC. This end-to-end treatment is needed so that a caller knows the  
3224 entire set of users that have accepted the call. Because of this special handling, retransmissions of the 2xx  
3225 response are handled by the UA core, not the transaction layer. Similarly, generation of the ACK for the 2xx  
3226 is handled by the UA core. Each proxy along the path merely forwards each 2xx response to INVITE and  
3227 its corresponding ACK.

## 3228 17.1 Client Transaction

3229 The client transaction provides its functionality through the maintenance of a state machine.

3230 The TU communicates with the client transaction through a simple interface. When the TU wishes to  
3231 initiate a new transaction, it creates a client transaction and passes it the SIP request to send and an IP  
3232 address, port, and transport to which to send it. The client transaction begins execution of its state machine.  
3233 Valid responses are passed up to the TU from the client transaction.

3234 There are two types of client transaction state machines, depending on the method of the request passed  
3235 by the TU. One handles client transactions for INVITE requests. This type of machine is referred to as  
3236 an INVITE client transaction. Another type handles client transactions for all requests except INVITE and  
3237 ACK. This is referred to as a non-INVITE client transaction. There is no client transaction for ACK. If the  
3238 TU wishes to send an ACK, it passes one directly to the transport layer for transmission.

3239 The INVITE transaction is different from those of other methods because of its extended duration. Nor-  
3240 mally, human input is required in order to respond to an INVITE. The long delays expected for sending a  
3241 response argue for a three-way handshake. On the other hand, requests of other methods are expected to  
3242 complete rapidly. Because of the non-INVITE transaction's reliance on a two-way handshake, TUs SHOULD  
3243 respond immediately to non-INVITE requests.

### 3244 17.1.1 INVITE Client Transaction

3245 **17.1.1.1 Overview of INVITE Transaction** The INVITE transaction consists of a three-way handshake.  
3246 The client transaction sends an INVITE, the server transaction sends responses, and the client transaction  
3247 sends an ACK. For unreliable transports (such as UDP), the client transaction retransmits requests at an  
3248 interval that starts at T1 seconds and doubles after every retransmission. T1 is an estimate of the round-  
3249 trip time (RTT), and it defaults to 500 ms. Nearly all of the transaction timers described here scale with

3250 T1, and changing T1 adjusts their values. The request is not retransmitted over reliable transports. After  
3251 receiving a 1xx response, any retransmissions cease altogether, and the client waits for further responses.  
3252 The server transaction can send additional 1xx responses, which are not transmitted reliably by the server  
3253 transaction. Eventually, the server transaction decides to send a final response. For unreliable transports,  
3254 that response is retransmitted periodically, and for reliable transports, it is sent once. For each final response  
3255 that is received at the client transaction, the client transaction sends an ACK, the purpose of which is to  
3256 quench retransmissions of the response.

3257 **17.1.1.2 Formal Description** The state machine for the INVITE client transaction is shown in Figure 5.  
3258 The initial state, "calling", MUST be entered when the TU initiates a new client transaction with an INVITE  
3259 request. The client transaction MUST pass the request to the transport layer for transmission (see Section 18).  
3260 If an unreliable transport is being used, the client transaction MUST start timer A with a value of T1. If a  
3261 reliable transport is being used, the client transaction SHOULD NOT start timer A (Timer A controls request  
3262 retransmissions). For any transport, the client transaction MUST start timer B with a value of 64\*T1 seconds  
3263 (Timer B controls transaction timeouts).

3264 When timer A fires, the client transaction MUST retransmit the request by passing it to the transport  
3265 layer, and MUST reset the timer with a value of 2\*T1. The formal definition of *retransmit* within the context  
3266 of the transaction layer is to take the message previously sent to the transport layer and pass it to the transport  
3267 layer once more.

3268 When timer A fires 2\*T1 seconds later, the request MUST be retransmitted again (assuming the client  
3269 transaction is still in this state). This process MUST continue so that the request is retransmitted with intervals  
3270 that double after each transmission. These retransmissions SHOULD only be done while the client transaction  
3271 is in the "calling" state.

3272 The default value for T1 is 500 ms. T1 is an estimate of the RTT between the client and server trans-  
3273 actions. Elements MAY (though it is NOT RECOMMENDED) use smaller values of T1 within closed, private  
3274 networks that do not permit general Internet connection. T1 MAY be chosen larger, and this is RECOM-  
3275 MENDED if it is known in advance (such as on high latency access links) that the RTT is larger. Whatever  
3276 the value of T1, the exponential backoffs on retransmissions described in this section MUST be used.

3277 If the client transaction is still in the "calling" state when timer B fires, the client transaction SHOULD  
3278 inform the TU that a timeout has occurred. The client transaction MUST NOT generate an ACK. The value of  
3279 64\*T1 is equal to the amount of time required to send seven requests in the case of an unreliable transport.

3280 If the client transaction receives a provisional response while in the "Calling" state, it transitions to the  
3281 "proceeding" state. In the "proceeding" state, the client transaction SHOULD NOT retransmit the request any  
3282 longer. Furthermore, the provisional response MUST be passed to the TU. Any further provisional responses  
3283 MUST be passed up to the TU while in the "proceeding" state.

3284 When in either the "Calling" or "Proceeding" states, reception of a response with status code from  
3285 300-699 MUST cause the client transaction to transition to "Completed". The client transaction MUST pass  
3286 the received response up to the TU, and the client transaction MUST generate an ACK request, even if the  
3287 transport is reliable (guidelines for constructing the ACK from the response are given in Section 17.1.1.3)  
3288 and then pass the ACK to the transport layer for transmission. The ACK MUST be sent to the same address,  
3289 port, and transport to which the original request was sent. The client transaction SHOULD start timer D  
3290 when it enters the "Completed" state, with a value of at least 32 seconds for unreliable transports, and a  
3291 value of zero seconds for reliable transports. Timer D reflects the amount of time that the server transaction  
3292 can remain in the "Completed" state when unreliable transports are used. This is equal to Timer H in the  
3293 INVITE server transaction, whose default is 64\*T1. However, the client transaction does not know the value

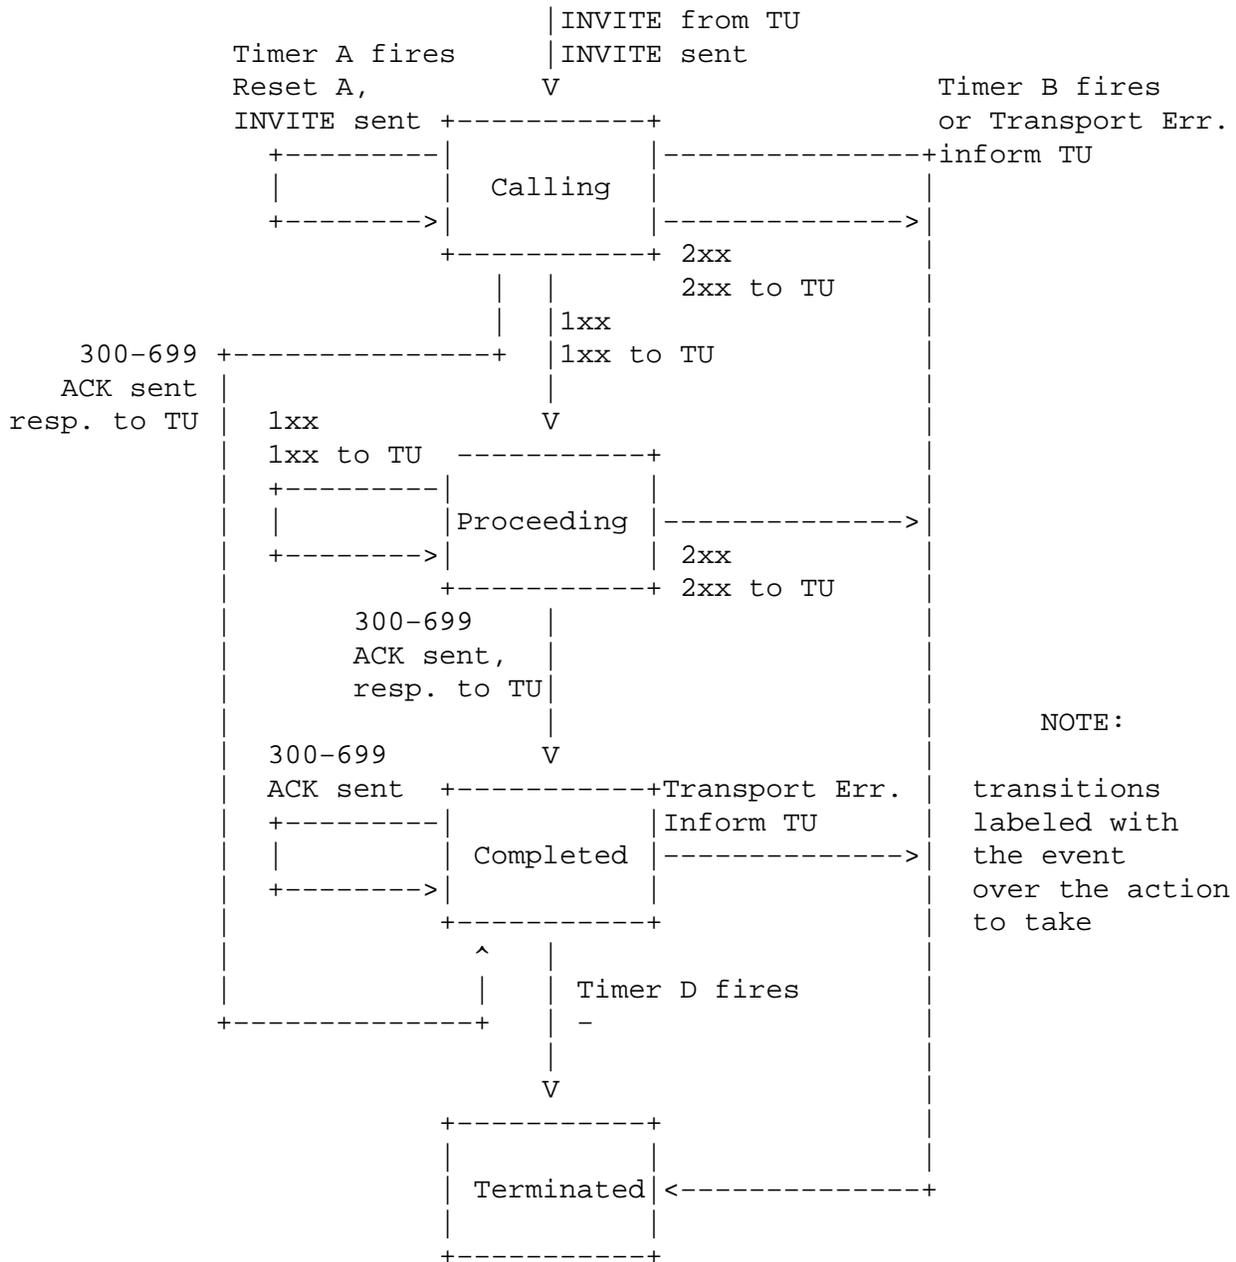


Figure 5: INVITE client transaction

3294 of T1 in use by the server transaction, so an absolute minimum of 32s is used instead of basing Timer D on  
 3295 T1.

3296 Any retransmissions of the final response that are received while in the “Completed” state MUST cause  
 3297 the ACK to be re-passed to the transport layer for retransmission, but the newly received response MUST  
 3298 NOT be passed up to the TU. A retransmission of the response is defined as any response which would match  
 3299 the same client transaction based on the rules of Section 17.1.3.

3300 If timer D fires while the client transaction is in the “Completed” state, the client transaction MUST move  
3301 to the terminated state, and it MUST inform the TU of the timeout.

3302 When in either the “Calling” or “Proceeding” states, reception of a 2xx response MUST cause the client  
3303 transaction to enter the “Terminated” state, and the response MUST be passed up to the TU. The handling of  
3304 this response depends on whether the TU is a proxy core or a UAC core. A UAC core will handle generation  
3305 of the ACK for this response, while a proxy core will always forward the 200 (OK) upstream. The differing  
3306 treatment of 200 (OK) between proxy and UAC is the reason that handling of it does not take place in the  
3307 transaction layer.

3308 The client transaction MUST be destroyed the instant it enters the “Terminated” state. This is actually  
3309 necessary to guarantee correct operation. The reason is that 2xx responses to an INVITE are treated differ-  
3310 ently; each one is forwarded by proxies, and the ACK handling in a UAC is different. Thus, each 2xx needs  
3311 to be passed to a proxy core (so that it can be forwarded) and to a UAC core (so it can be acknowledged). No  
3312 transaction layer processing takes place. Whenever a response is received by the transport, if the transport  
3313 layer finds no matching client transaction (using the rules of Section 17.1.3), the response is passed directly  
3314 to the core. Since the matching client transaction is destroyed by the first 2xx, subsequent 2xx will find no  
3315 match and therefore be passed to the core.

3316 **17.1.1.3 Construction of the ACK Request** This section specifies the construction of ACK requests  
3317 sent within the client transaction. A UAC core that generates an ACK for 2xx MUST instead follow the rules  
3318 described in Section 13.

3319 The ACK request constructed by the client transaction MUST contain values for the Call-ID, From, and  
3320 Request-URI that are equal to the values of those header fields in the request passed to the transport by  
3321 the client transaction (call this the “original request”). The To header field in the ACK MUST equal the To  
3322 header field in the response being acknowledged, and therefore will usually differ from the To header field  
3323 in the original request by the addition of the tag parameter. The ACK MUST contain a single Via header  
3324 field, and this MUST be equal to the top Via header field of the original request. The CSeq header field in  
3325 the ACK MUST contain the same value for the sequence number as was present in the original request, but  
3326 the method parameter MUST be equal to “ACK”.

3327 If the INVITE request whose response is being acknowledged had Route header fields, those header  
3328 fields MUST appear in the ACK. This is to ensure that the ACK can be routed properly through any down-  
3329 stream stateless proxies.

3330 Although any request MAY contain a body, a body in an ACK is special since the request cannot be  
3331 rejected if the body is not understood. Therefore, placement of bodies in ACK for non-2xx is NOT RECOM-  
3332 MENDED, but if done, the body types are restricted to any that appeared in the INVITE, assuming that the  
3333 response to the INVITE was not 415. If it was, the body in the ACK MAY be any type listed in the Accept  
3334 header field in the 415.

3335 For example, consider the following request:

```
3336 INVITE sip:bob@biloxi.com SIP/2.0
3337 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKkjshdyff
3338 To: Bob <sip:bob@biloxi.com>
3339 From: Alice <sip:alice@atlanta.com>;tag=88sja8x
3340 Max-Forwards: 70
3341 Call-ID: 987asjd97y7atg
3342 CSeq: 986759 INVITE
```

3343 The ACK request for a non-2xx final response to this request would look like this:

```
3344 ACK sip:bob@biloxi.com SIP/2.0
3345 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKkjshdyff
3346 To: Bob <sip:bob@biloxi.com>;tag=99sa0xk
3347 From: Alice <sip:alice@atlanta.com>;tag=88sja8x
3348 Max-Forwards: 70
3349 Call-ID: 987asjd97y7atg
3350 CSeq: 986759 ACK
```

## 3351 17.1.2 Non-INVITE Client Transaction

3352 **17.1.2.1 Overview of the non-INVITE Transaction** Non-INVITE transactions do not make use of ACK.  
3353 They are simple request-response interactions. For unreliable transports, requests are retransmitted at an interval which starts at T1 and doubles until it hits T2. If a provisional response is received, retransmissions  
3354 continue for unreliable transports, but at an interval of T2. The server transaction retransmits the last response it sent, which can be a provisional or final response, only when a retransmission of the request is  
3355 received. This is why request retransmissions need to continue even after a provisional response, they are to  
3356 ensure reliable delivery of the final response.

3357 Unlike an INVITE transaction, a non-INVITE transaction has no special handling for the 2xx response.  
3358 The result is that only a single 2xx response to a non-INVITE is ever delivered to a UAC.  
3359

3360 **17.1.2.2 Formal Description** The state machine for the non-INVITE client transaction is shown in Figure 6. It is very similar to the state machine for INVITE.

3361 The “Trying” state is entered when the TU initiates a new client transaction with a request. When  
3362 entering this state, the client transaction SHOULD set timer F to fire in  $64 * T1$  seconds. The request MUST be  
3363 passed to the transport layer for transmission. If an unreliable transport is in use, the client transaction MUST  
3364 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  
3365 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  
3366 so that retransmissions occur with an exponentially increasing interval that caps at T2. The default value  
3367 of T2 is 4s, and it represents the amount of time a non-INVITE server transaction will take to respond to a  
3368 request, if it does not respond immediately. For the default values of T1 and T2, this results in intervals of  
3369 500 ms, 1 s, 2 s, 4 s, 4 s, etc.

3370 If Timer F fires while the client transaction is still in the “Trying” state, the client transaction SHOULD  
3371 inform the TU about the timeout, and then it SHOULD enter the “Terminated” state. If a provisional response  
3372 is received while in the “Trying” state, the response MUST be passed to the TU, and then the client transaction  
3373 SHOULD move to the “Proceeding” state. If a final response (status codes 200-699) is received while in the  
3374 “Trying” state, the response MUST be passed to the TU, and the client transaction MUST transition to the  
3375 “Completed” state.

3376 If Timer E fires while in the “Proceeding” state, the request MUST be passed to the transport layer  
3377 for retransmission, and Timer E MUST be reset with a value of T2 seconds. If timer F fires while in the  
3378 “Proceeding” state, the TU MUST be informed of a timeout, and the client transaction MUST transition to the  
3379 terminated state. If a final response (status codes 200-699) is received while in the “Proceeding” state, the  
3380 response MUST be passed to the TU, and the client transaction MUST transition to the “Completed” state.

3381 Once the client transaction enters the “Completed” state, it MUST set Timer K to fire in T4 seconds for  
3382

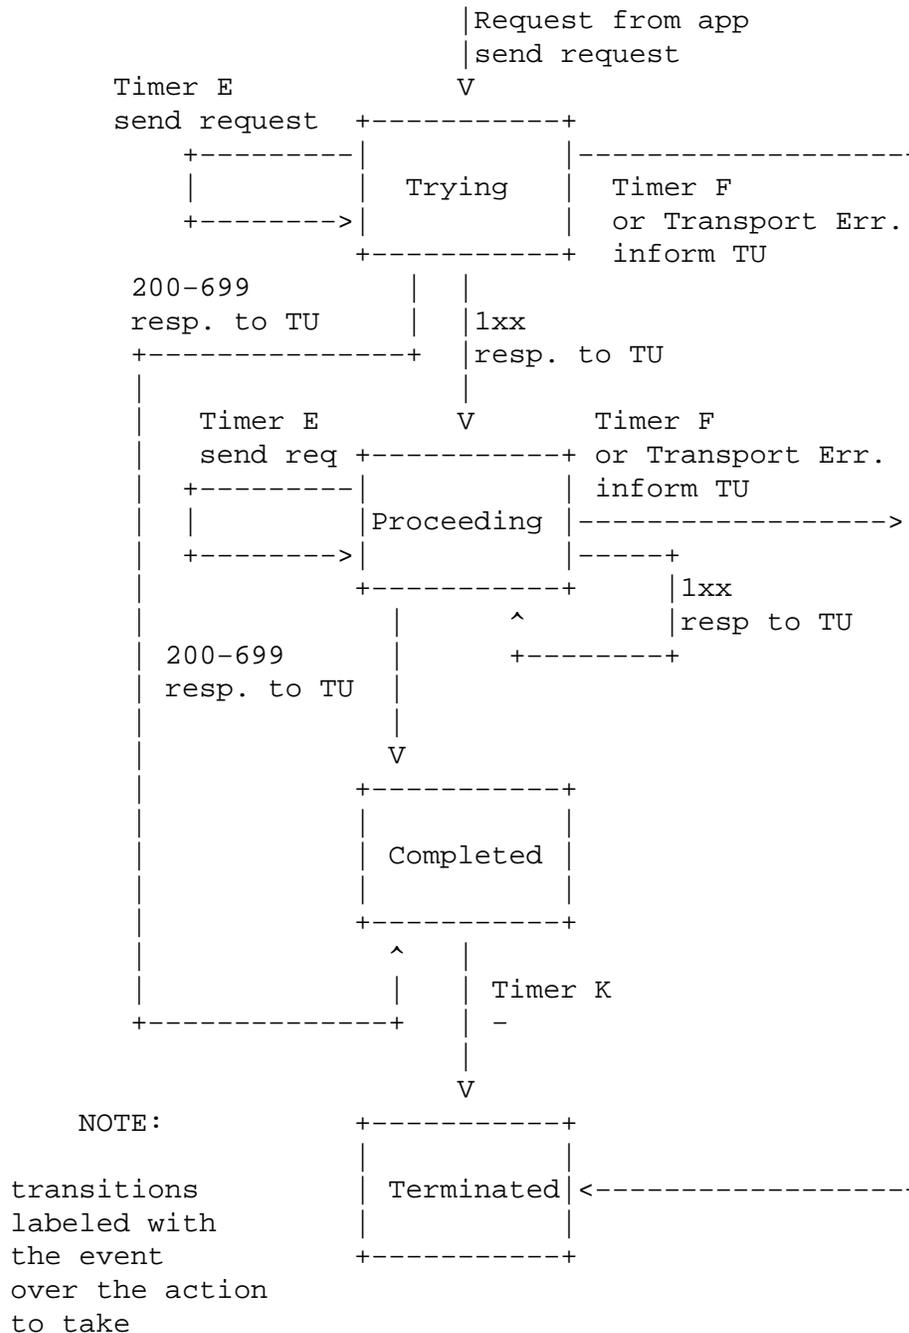


Figure 6: non-INVITE client transaction

3384 unreliable transports, and zero seconds for reliable transports. The “Completed” state exists to buffer any  
 3385 additional response retransmissions that may be received (which is why the client transaction remains there  
 3386 only for unreliable transports). T4 represents the amount of time the network will take to clear messages  
 3387 between client and server transactions. The default value of T4 is 5s. A response is a retransmission when it

3388 matches the same transaction, using the rules specified in Section 17.1.3. If Timer K fires while in this state,  
3389 the client transaction MUST transition to the "Terminated" state.

3390 Once the transaction is in the terminated state, it MUST be destroyed.

### 3391 **17.1.3 Matching Responses to Client Transactions**

3392 When the transport layer in the client receives a response, it has to determine which client transaction  
3393 will handle the response, so that the processing of Sections 17.1.1 and 17.1.2 can take place. The branch  
3394 parameter in the top Via header field is used for this purpose. A response matches a client transaction under  
3395 two conditions:

- 3396 1. If the response has the same value of the branch parameter in the top Via header field as the branch  
3397 parameter in the top Via header field of the request that created the transaction.
- 3398 2. If the method parameter in the CSeq header field matches the method of the request that created the  
3399 transaction. The method is needed since a CANCEL request constitutes a different transaction, but  
3400 shares the same value of the branch parameter.

3401 A response that matches a transaction matched by a previous response is considered a retransmission of  
3402 that response.

3403 If a request is sent via multicast, it is possible that it will generate multiple responses from different  
3404 servers. These responses will all have the same branch parameter in the topmost Via, but vary in the To  
3405 tag. The first response received, based on the rules above, will be used, and others will be viewed as  
3406 retransmissions. That is not an error; multicast SIP provides only a rudimentary "single-hop-discovery-  
3407 like" service that is limited to processing a single response. See Section 18.1.1 for details.

### 3408 **17.1.4 Handling Transport Errors**

3409 When the client transaction sends a request to the transport layer to be sent, the following procedures are  
3410 followed if the transport layer indicates a failure.

3411 The client transaction SHOULD inform the TU that a transport failure has occurred, and the client trans-  
3412 action SHOULD transition directly to the "Terminated" state. The TU will handle the failover mechanisms  
3413 described in [4].

## 3414 **17.2 Server Transaction**

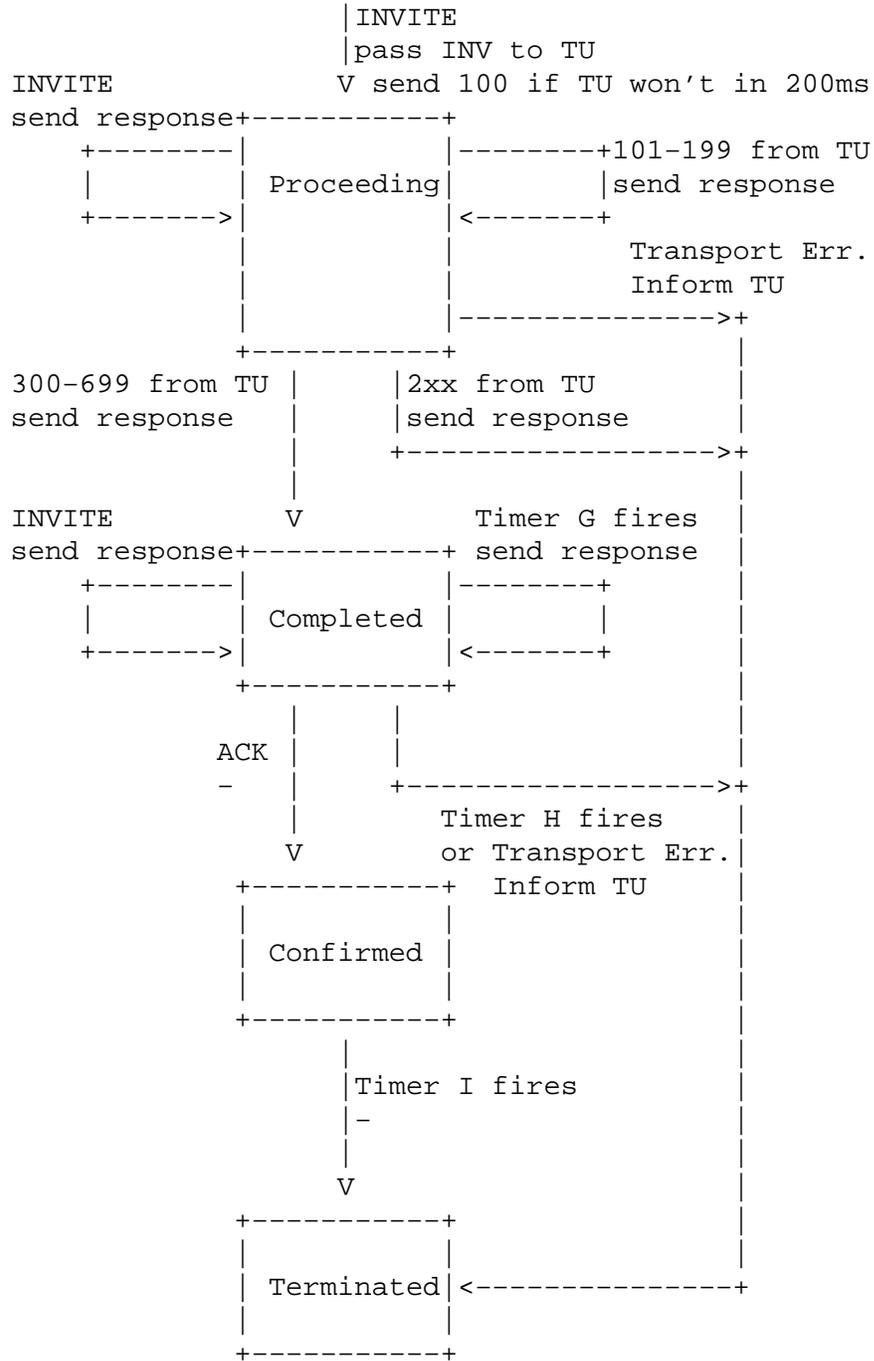
3415 The server transaction is responsible for the delivery of requests to the TU and the reliable transmission of  
3416 responses. It accomplishes this through a state machine. Server transactions are created by the core when a  
3417 request is received, and transaction handling is desired for that request (this is not always the case).

3418 As with the client transactions, the state machine depends on whether the received request is an INVITE  
3419 request.

### 3420 **17.2.1 INVITE Server Transaction**

3421 The state diagram for the INVITE server transaction is shown in Figure 7.

3422 When a server transaction is constructed with a request, it enters the "Proceeding" state. The server  
3423 transaction MUST generate a 100 (Trying) response unless it knows that the TU will generate a provisional



3424 or final response within 200 ms, in which case it MAY generate a 100 (Trying) response. This provisional  
3425 response is needed to quench request retransmissions rapidly in order to avoid network congestion. The 100  
3426 (Trying) response is constructed according to the procedures in Section 8.2.6, except that the insertion of  
3427 tags in the To header field of the response (when none was present in the request) is downgraded from MAY  
3428 to SHOULD NOT. The request MUST be passed to the TU.

3429 The TU passes any number of provisional responses to the server transaction. So long as the server  
3430 transaction is in the "Proceeding" state, each of these MUST be passed to the transport layer for transmission.  
3431 They are not sent reliably by the transaction layer (they are not retransmitted by it) and do not cause a change  
3432 in the state of the server transaction. If a request retransmission is received while in the "Proceeding" state,  
3433 the most recent provisional response that was received from the TU MUST be passed to the transport layer  
3434 for retransmission. A request is a retransmission if it matches the same server transaction based on the rules  
3435 of Section 17.2.3.

3436 If, while in the "Proceeding" state, the TU passes a 2xx response to the server transaction, the server  
3437 transaction MUST pass this response to the transport layer for transmission. It is not retransmitted by the  
3438 server transaction; retransmissions of 2xx responses are handled by the TU. The server transaction MUST  
3439 then transition to the "Terminated" state.

3440 While in the "Proceeding" state, if the TU passes a response with status code from 300 to 699 to the  
3441 server transaction, the response MUST be passed to the transport layer for transmission, and the state machine  
3442 MUST enter the "Completed" state. For unreliable transports, timer G is set to fire in T1 seconds, and is not  
3443 set to fire for reliable transports.

3444 This is a change from RFC 2543, where responses were always retransmitted, even over reliable transports.

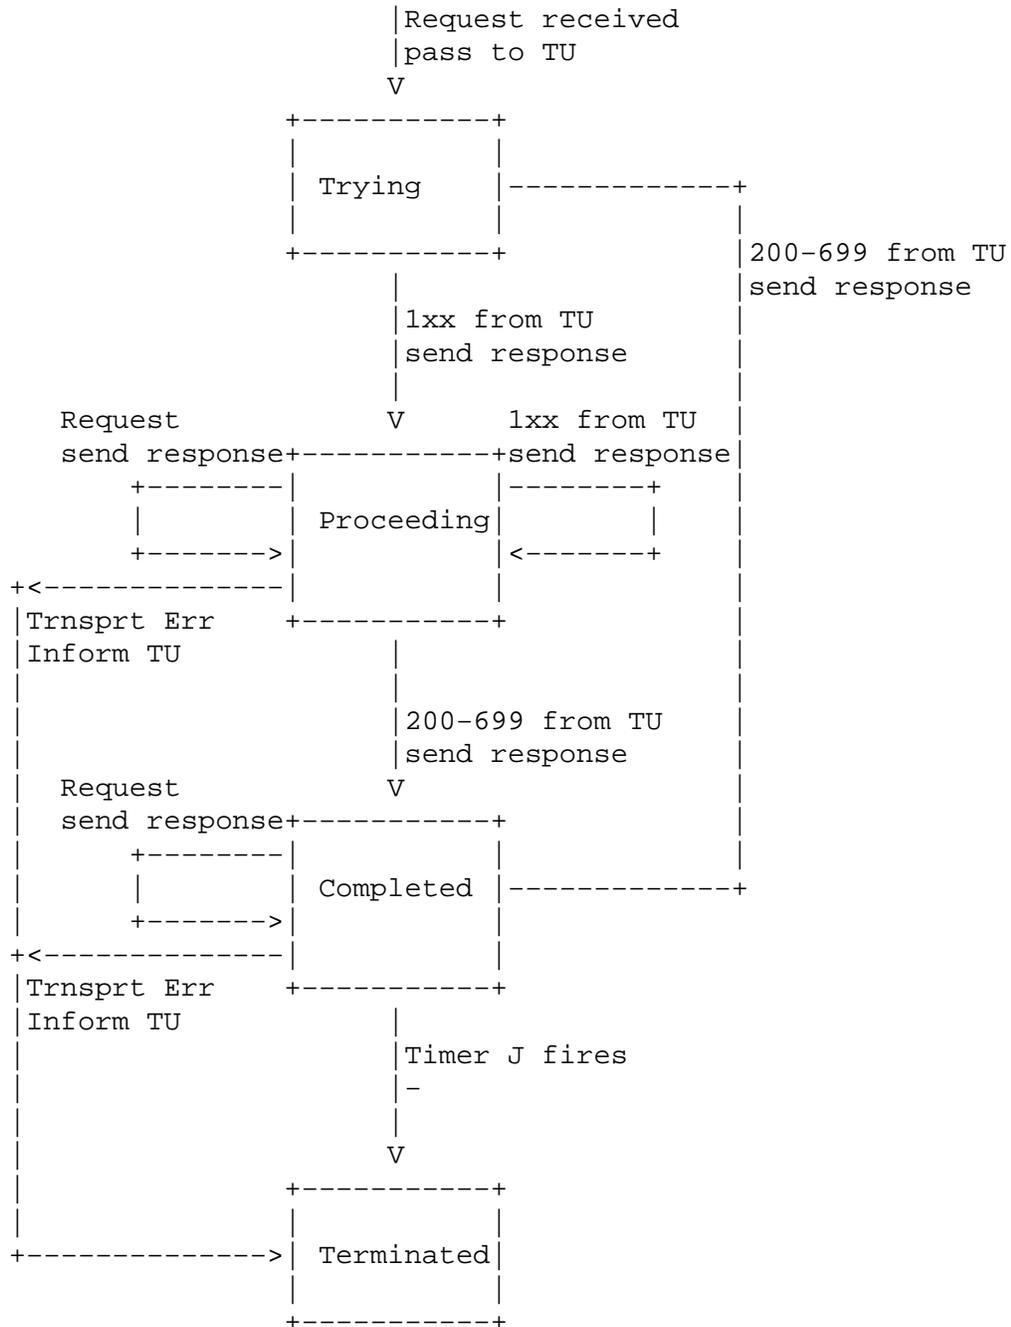
3445 When the "Completed" state is entered, timer H MUST be set to fire in  $64 * T1$  seconds for all transports.  
3446 Timer H determines when the server transaction abandons retransmitting the response. Its value is chosen  
3447 to equal Timer B, the amount of time a client transaction will continue to retry sending a request. If timer G  
3448 fires, the response is passed to the transport layer once more for retransmission, and timer G is set to fire in  
3449  $\text{MIN}(2 * T1, T2)$  seconds. From then on, when timer G fires, the response is passed to the transport again for  
3450 transmission, and timer G is reset with a value that doubles, unless that value exceeds T2, in which case it  
3451 is reset with the value of T2. This is identical to the retransmit behavior for requests in the "Trying" state of  
3452 the non-INVITE client transaction. Furthermore, while in the "Completed" state, if a request retransmission  
3453 is received, the server SHOULD pass the response to the transport for retransmission.

3454 If an ACK is received while the server transaction is in the "Completed" state, the server transaction  
3455 MUST transition to the "Confirmed" state. As Timer G is ignored in this state, any retransmissions of the  
3456 response will cease.

3457 If timer H fires while in the "Completed" state, it implies that the ACK was never received. In this  
3458 case, the server transaction MUST transition to the "Terminated" state, and MUST indicate to the TU that a  
3459 transaction failure has occurred.

3460 The purpose of the "Confirmed" state is to absorb any additional ACK messages that arrive, triggered  
3461 from retransmissions of the final response. When this state is entered, timer I is set to fire in T4 seconds for  
3462 unreliable transports, and zero seconds for reliable transports. Once timer I fires, the server MUST transition  
3463 to the "Terminated" state.

3464 Once the transaction is in the "Terminated" state, it MUST be destroyed. As with client transactions, this  
3465 is needed to ensure reliability of the 2xx responses to INVITE.



### 3466 17.2.2 Non-INVITE Server Transaction

3467 The state machine for the non-INVITE server transaction is shown in Figure 8.

3468 The state machine is initialized in the "Trying" state and is passed a request other than INVITE or  
3469 ACK when initialized. This request is passed up to the TU. Once in the "Trying" state, any further request  
3470 retransmissions are discarded. A request is a retransmission if it matches the same server transaction, using  
3471 the rules specified in Section 17.2.3.

3472 While in the "Trying" state, if the TU passes a provisional response to the server transaction, the server  
3473 transaction MUST enter the "Proceeding" state. The response MUST be passed to the transport layer for  
3474 transmission. Any further provisional responses that are received from the TU while in the "Proceeding"  
3475 state MUST be passed to the transport layer for transmission. If a retransmission of the request is received  
3476 while in the "Proceeding" state, the most recently sent provisional response MUST be passed to the transport  
3477 layer for retransmission. If the TU passes a final response (status codes 200-699) to the server while in the  
3478 "Proceeding" state, the transaction MUST enter the "Completed" state, and the response MUST be passed to  
3479 the transport layer for transmission.

3480 When the server transaction enters the "Completed" state, it MUST set Timer J to fire in  $64 * T1$  seconds  
3481 for unreliable transports, and zero seconds for reliable transports. While in the "Completed" state, the server  
3482 transaction MUST pass the final response to the transport layer for retransmission whenever a retransmission  
3483 of the request is received. Any other final responses passed by the TU to the server transaction MUST be  
3484 discarded while in the "Completed" state. The server transaction remains in this state until Timer J fires, at  
3485 which point it MUST transition to the "Terminated" state.

3486 The server transaction MUST be destroyed the instant it enters the "Terminated" state.

### 3487 17.2.3 Matching Requests to Server Transactions

3488 When a request is received from the network by the server, it has to be matched to an existing transaction.  
3489 This is accomplished in the following manner.

3490 The branch parameter in the topmost Via header field of the request is examined. If it is present and  
3491 begins with the magic cookie "z9hG4bK", the request was generated by a client transaction compliant to this  
3492 specification. Therefore, the branch parameter will be unique across all transactions sent by that client. The  
3493 request matches a transaction if the branch parameter in the request is equal to the one in the top Via header  
3494 field of the request that created the transaction, the sent-by value in the top Via of the request is equal to  
3495 the one in the request that created the transaction, and in the case of a CANCEL request, the method of  
3496 the request that created the transaction was also CANCEL. This matching rule applies to both INVITE and  
3497 non-INVITE transactions alike.

3498 The sent-by value is used as part of the matching process because there could be duplication of branch param-  
3499 eters from different clients; uniqueness in time is mandated for construction of the parameter, but not uniqueness in  
3500 space.

3501 If the branch parameter in the top Via header field is not present, or does not contain the magic cookie,  
3502 the following procedures are used. These exist to handle backwards compatibility with RFC 2543 compliant  
3503 implementations.

3504 The INVITE request matches a transaction if the Request-URI, To tag, From tag, Call-ID, CSeq, and  
3505 top Via header field match those of the INVITE request which created the transaction. In this case, the  
3506 INVITE is a retransmission of the original one that created the transaction. The ACK request matches a  
3507 transaction if the Request-URI, From tag, Call-ID, CSeq number (not the method), and top Via header

3508 field match those of the INVITE request which created the transaction, and the To tag of the ACK matches  
3509 the To tag of the response sent by the server transaction. Matching is done based on the matching rules  
3510 defined for each of those header fields. The usage of the tag in the To header field helps disambiguate ACK  
3511 for 2xx from ACK for other responses at a proxy, which may have forwarded both responses (which can  
3512 occur in unusual conditions). An ACK request that matches an INVITE transaction matched by a previous  
3513 ACK is considered a retransmission of that previous ACK.

3514 For all other request methods, a request is matched to a transaction if the Request-URI, To tag, From  
3515 tag, Call-ID Cseq (including the method), and top Via header field match those of the request that created  
3516 the transaction. Matching is done based on the matching rules defined for each of those header fields. When  
3517 a non-INVITE request matches an existing transaction, it is a retransmission of the request that created that  
3518 transaction.

3519 Because the matching rules include the Request-URI, the server cannot match a response to a transac-  
3520 tion. When the TU passes a response to the server transaction, it must pass it to the specific server transaction  
3521 for which the response is targeted.

#### 3522 17.2.4 Handling Transport Errors

3523 When the server transaction sends a response to the transport layer to be sent, the following procedures are  
3524 followed if the transport layer indicates a failure.

3525 First, the procedures in [4] are followed, which attempt to deliver the response to a backup. If those  
3526 should all fail, based on the definition of failure in [4], the server transaction SHOULD inform the TU that a  
3527 failure has occurred, and SHOULD transition to the terminated state.

## 3528 18 Transport

3529 The transport layer is responsible for the actual transmission of requests and responses over network trans-  
3530 ports. This includes determination of the connection to use for a request or response in the case of connection-  
3531 oriented transports.

3532 The transport layer is responsible for managing persistent connections for transport protocols like TCP  
3533 and SCTP, or TLS over those, including ones opened to the transport layer. This includes connections  
3534 opened by the client or server transports, so that connections are shared between client and server transport  
3535 functions. These connections are indexed by the tuple formed from the address, port, and transport protocol  
3536 at the far end of the connection. When a connection is opened by the transport layer, this index is set to the  
3537 destination IP, port and transport. When the connection is accepted by the transport layer, this index is set to  
3538 the source IP address, port number, and transport. Note that, because the source port is often ephemeral, but  
3539 it cannot be known whether it is ephemeral or selected through procedures in [4], connections accepted by  
3540 the transport layer will frequently not be reused. The result is that two proxies in a “peering” relationship  
3541 using a connection-oriented transport frequently will have two connections in use, one for transactions  
3542 initiated in each direction.

3543 It is RECOMMENDED that connections be kept open for some implementation-defined duration after the  
3544 last message was sent or received over that connection. This duration SHOULD at least equal the longest  
3545 amount of time the element would need in order to bring a transaction from instantiation to the terminated  
3546 state. This is to make it likely that transactions complete over the same connection on which they are  
3547 initiated (for example, request, response, and in the case of INVITE, ACK for non-2xx responses). This  
3548 usually means at least 64\*T1 (see Section 17.1.1.1 for a definition of T1). However, it could be larger in an

3549 element that has a TU using a large value for timer C (bullet 11 of Section 16.6), for example.

3550 All SIP elements **MUST** implement UDP and TCP. SIP elements **MAY** implement other protocols.

3551 Making TCP mandatory for the UA is a substantial change from RFC 2543. It has arisen out of the need to  
3552 handle larger messages, which **MUST** use TCP, as discussed below. Thus, even if an element never sends large  
3553 messages, it may receive one and needs to be able to handle them.

## 3554 18.1 Clients

### 3555 18.1.1 Sending Requests

3556 The client side of the transport layer is responsible for sending the request and receiving responses. The  
3557 user of the transport layer passes the client transport the request, an IP address, port, transport, and possibly  
3558 TTL for multicast destinations.

3559 If a request is within 200 bytes of the path MTU, or if it is larger than 1300 bytes and the path MTU  
3560 is unknown, the request **MUST** be sent using TCP. This prevents fragmentation of messages over UDP  
3561 and provides congestion control for larger messages. However, implementations **MUST** be able to handle  
3562 messages up to the maximum datagram packet size. For UDP, this size is 65,535 bytes, including IP and  
3563 UDP headers.

3564 The 200 byte "buffer" between the message size and the MTU accommodates the fact that the response in  
3565 SIP can be larger than the request. This happens due to the addition of **Record-Route** header field values to the  
3566 responses to **INVITE**, for example. With the extra buffer, the response can be about 170 bytes larger than the request,  
3567 and still not be fragmented on IPv4 (about 30 bytes is consumed by IP/UDP, assuming no IPsec). 1300 is chosen  
3568 when path MTU is not known, based on the assumption of a 1500 byte Ethernet MTU.

3569 If an element sends a request over TCP because of these message size constraints, and that request  
3570 would have otherwise been sent over UDP, if the attempt to establish the connection generates either an  
3571 ICMP Protocol Not Supported, or results in a TCP reset, the element **SHOULD** retry the request, using UDP.  
3572 This is only to provide backwards compatibility with RFC 2543 compliant implementations that do not  
3573 support UDP. It is anticipated that this behavior will be deprecated in a future revision of this specification.

3574 A client that sends a request to a multicast address **MUST** add the "maddr" parameter to its **Via** header  
3575 field value containing the destination multicast address, and for IPv4, **SHOULD** add the "ttl" parameter with  
3576 a value of 1. Usage of IPv6 multicast is not defined in this specification, and will be a subject of future  
3577 standardization when the need arises.

3578 These rules result in a purposeful limitation of multicast in SIP. Its primary function is to provide an  
3579 "single-hop-discovery-like" service, delivering a request to a group of homogeneous servers, where it is only  
3580 required to process the response from any one of them. This functionality is most useful for registrations.  
3581 In fact, based on the transaction processing rules in Section 17.1.3, the client transaction will accept the first  
3582 response, and view any others as retransmissions because they all contain the same **Via** branch identifier.

3583 Before a request is sent, the client transport **MUST** insert a value of the "sent-by" field into the **Via** header  
3584 field. This field contains an IP address or host name, and port. The usage of an FQDN is **RECOMMENDED**.  
3585 This field is used for sending responses under certain conditions, described below. If the port is absent, the  
3586 default value depends on the transport. It is 5060 for UDP, TCP and SCTP, 5061 for TLS.

3587 For reliable transports, the response is normally sent on the connection on which the request was re-  
3588 ceived. Therefore, the client transport **MUST** be prepared to receive the response on the same connection  
3589 used to send the request. Under error conditions, the server may attempt to open a new connection to send  
3590 the response. To handle this case, the transport layer **MUST** also be prepared to receive an incoming con-  
3591 nection on the source IP address from which the request was sent and port number in the "sent-by" field. It

3592 also MUST be prepared to receive incoming connections on any address and port that would be selected by  
3593 a server based on the procedures described in Section 5 of [4].

3594 For unreliable unicast transports, the client transport MUST be prepared to receive responses on the  
3595 source IP address from which the request is sent (as responses are sent back to the source address) and the  
3596 port number in the "sent-by" field. Furthermore, as with reliable transports, in certain cases the response  
3597 will be sent elsewhere. The client MUST be prepared to receive responses on any address and port that would  
3598 be selected by a server based on the procedures described in Section 5 of [4].

3599 For multicast, the client transport MUST be prepared to receive responses on the same multicast group  
3600 and port to which the request is sent (that is, it needs to be a member of the multicast group it sent the request  
3601 to.)

3602 If a request is destined to an IP address, port, and transport to which an existing connection is open, it  
3603 is RECOMMENDED that this connection be used to send the request, but another connection MAY be opened  
3604 and used.

3605 If a request is sent using multicast, it is sent to the group address, port, and TTL provided by the transport  
3606 user. If a request is sent using unicast unreliable transports, it is sent to the IP address and port provided by  
3607 the transport user.

## 3608 18.1.2 Receiving Responses

3609 When a response is received, the client transport examines the top Via header field value. If the value of  
3610 the "sent-by" parameter in that header field value does not correspond to a value that the client transport is  
3611 configured to insert into requests, the response MUST be silently discarded.

3612 If there are any client transactions in existence, the client transport uses the matching procedures of Sec-  
3613 tion 17.1.3 to attempt to match the response to an existing transaction. If there is a match, the response MUST  
3614 be passed to that transaction. Otherwise, the response MUST be passed to the core (whether it be stateless  
3615 proxy, stateful proxy, or UA) for further processing. Handling of these "stray" responses is dependent on  
3616 the core (a proxy will forward them, while a UA will discard, for example).

## 3617 18.2 Servers

### 3618 18.2.1 Receiving Requests

3619 A server SHOULD be prepared to received requests on any IP address, port and transport combination that can  
3620 be the result of a DNS lookup on a SIP or SIPS URI [4] that is handed out for the purposes of communicating  
3621 with that server. In this context, "handing out" includes placing a URI in a Contact header field in a  
3622 REGISTER request or a any redirect response, or in a Record-Route header field in a request or response.  
3623 A URI can also be "handed out" by placing it on a web page or business card. It is also RECOMMENDED that  
3624 a server listen for requests on the default SIP ports on all public interfaces. The typical exception would be  
3625 private networks, or when multiple server instances are running on the same host. For any port and interface  
3626 that a server listens on for UDP, it MUST listen on that same port and interface for TCP. This is because  
3627 a message may need to be sent using TCP, rather than UDP, if it is too large. As a result, the converse is  
3628 not true. A server need not, and indeed SHOULD NOT listen for UDP on a particular address and port just  
3629 because it is listening on that same address and port for UDP. There may, of course, be other reasons why a  
3630 server needs to listen for UDP on a particular address and port.

3631 When the server transport receives a request over any transport, it MUST examine the value of the "sent-  
3632 by" parameter in the top Via header field value. If the host portion of the "sent-by" parameter contains a

3633 domain name, or if it contains an IP address that differs from the packet source address, the server MUST  
3634 add a "received" parameter to that Via header field value. This parameter MUST contain the source address  
3635 from which the packet was received. This is to assist the server transport layer in sending the response, since  
3636 it must be sent to the source IP address from which the request came.

3637 Consider a request received by the server transport which looks like, in part:

```
3638 INVITE sip:bob@Biloxi.com SIP/2.0
3639 Via: SIP/2.0/UDP bobspc.biloxi.com:5060
```

3640 The request is received with a source IP address of 1.2.3.4. Before passing the request up, the transport  
3641 adds a "received" parameter, so that the request would look like, in part:

```
3642 INVITE sip:bob@Biloxi.com SIP/2.0
3643 Via: SIP/2.0/UDP bobspc.biloxi.com:5060;received=1.2.3.4
```

3644 Next, the server transport attempts to match the request to a server transaction. It does so using the  
3645 matching rules described in Section 17.2.3. If a matching server transaction is found, the request is passed  
3646 to that transaction for processing. If no match is found, the request is passed to the core, which may  
3647 decide to construct a new server transaction for that request. Note that when a UAS core sends a 2xx  
3648 response to INVITE, the server transaction is destroyed. This means that when the ACK arrives, there will  
3649 be no matching server transaction, and based on this rule, the ACK is passed to the UAS core, where it is  
3650 processed.

## 3651 18.2.2 Sending Responses

3652 The server transport uses the value of the top Via header field in order to determine where to send a response.  
3653 It MUST follow the following process:

- 3654 • If the "sent-protocol" is a reliable transport protocol such as TCP or SCTP, or TLS over those,  
3655 the response MUST be sent using the existing connection to the source of the original request that  
3656 created the transaction, if that connection is still open. This requires the server transport to maintain  
3657 an association between server transactions and transport connections. If that connection is no longer  
3658 open, the server SHOULD open a connection to the IP address in the "received" parameter, if present,  
3659 using the port in the "sent-by" value, or the default port for that transport, if no port is specified.  
3660 If that connection attempt fails, the server SHOULD use the procedures in [4] for servers in order to  
3661 determine the IP address and port to open the connection and send the response to.
- 3662 • Otherwise, if the Via header field value contains a "maddr" parameter, the response MUST be for-  
3663 forwarded to the address listed there, using the port indicated in "sent-by", or port 5060 if none is  
3664 present. If the address is a multicast address, the response SHOULD be sent using the TTL indicated  
3665 in the "ttl" parameter, or with a TTL of 1 if that parameter is not present.
- 3666 • Otherwise (for unreliable unicast transports), if the top Via has a "received" parameter, the response  
3667 MUST be sent to the address in the "received" parameter, using the port indicated in the "sent-by"  
3668 value, or using port 5060 if none is specified explicitly. If this fails, for example, elicits an ICMP

3669 “port unreachable” response, the procedures of Section 5 of [4] SHOULD be used to determine where  
3670 to send the response.

3671 • Otherwise, if it is not receiver-tagged, the response MUST be sent to the address indicated by the  
3672 “sent-by” value, using the procedures in Section 5 of [4].

### 3673 **18.3 Framing**

3674 In the case of message-oriented transports (such as UDP), if the message has a Content-Length header  
3675 field, the message body is assumed to contain that many bytes. If there are additional bytes in the transport  
3676 packet beyond the end of the body, they MUST be discarded. If the transport packet ends before the end  
3677 of the message body, this is considered an error. If the message is a response, it MUST be discarded. If its  
3678 a request, the element SHOULD generate a 400 (Bad Request) response. If the message has no Content-  
3679 Length header field, the message body is assumed to end at the end of the transport packet.

3680 In the case of stream-oriented transports such as TCP, the Content-Length header field indicates the  
3681 size of the body. The Content-Length header field MUST be used with stream oriented transports.

### 3682 **18.4 Error Handling**

3683 Error handling is independent of whether the message was a request or response.

3684 If the transport user asks for a message to be sent over an unreliable transport, and the result is an ICMP  
3685 error, the behavior depends on the type of ICMP error. Host, network, port or protocol unreachable errors,  
3686 or parameter problem errors SHOULD cause the transport layer to inform the transport user of a failure in  
3687 sending. Source quench and TTL exceeded ICMP errors SHOULD be ignored.

3688 If the transport user asks for a request to be sent over a reliable transport, and the result is a connection  
3689 failure, the transport layer SHOULD inform the transport user of a failure in sending.

## 3690 **19 Common Message Components**

3691 There are certain components of SIP messages that appear in various places within SIP messages (and  
3692 sometimes, outside of them) that merit separate discussion.

### 3693 **19.1 SIP and SIPS Uniform Resource Indicators**

3694 A SIP or SIPS URI identifies a communications resource. Like all URIs, SIP and SIPS URIs may be placed  
3695 in web pages, email messages, or printed literature. They contain sufficient information to initiate and  
3696 maintain a communication session with the resource.

3697 Examples of communications resources include the following:

- 3698 • a user of an online service
- 3699 • an appearance on a multi-line phone
- 3700 • a mailbox on a messaging system
- 3701 • a PSTN number at a gateway service
- 3702 • a group (such as “sales” or “helpdesk”) in an organization

3703 A SIPS URI specifies that the resource be contacted securely. This means, in particular, that TLS is to  
3704 be used between all elements, starting from the UAC, and ending at the UAS. Any resource described by a  
3705 SIP URI can be “upgraded” to a SIPS URI by just changing the scheme, if it is desired to communicate with  
3706 that resource securely.

### 3707 19.1.1 SIP and SIPS URI Components

3708 The “sip:” and “sips:” schemes follow the guidelines in RFC 2396 [5]. They use a form similar to the mailto  
3709 URL, allowing the specification of SIP request-header fields and the SIP message-body. This makes it  
3710 possible to specify the subject, media type, or urgency of sessions initiated by using a URI on a web page or  
3711 in an email message. The formal syntax for a SIP or SIPS URI is presented in Section 25. Its general form,  
3712 in the case of a SIP URI, is

3713 sip:user:password@host:port;uri-parameters?headers

3714 The format for a SIPS URI is the same, except that the scheme is “sips” instead of sip. These tokens, and  
3715 some of the tokens in their expansions, have the following meanings:

3716 **user:** The identifier of a particular resource at the host being addressed. The term “host” in this context  
3717 frequently refers to a domain. The “userinfo” of a URI consists of this user field, the password field,  
3718 and the @ sign following them. The userinfo part of a URI is optional and MAY be absent when the  
3719 destination host does not have a notion of users or when the host itself is the resource being identified.  
3720 If the @ sign is present in a SIP or SIPS URI, the user field MUST NOT be empty.

3721 If the host being addressed can process telephone numbers, for instance, an Internet telephony gate-  
3722 way, a telephone-subscriber field defined in RFC 2806 [9] MAY be used to populate the user field.  
3723 There are special escaping rules for encoding telephone-subscriber fields in SIP and SIPS URIs  
3724 described in Section 19.1.2.

3725 **password:** A password associated with the user. While the SIP and SIPS URI syntax allows this field to  
3726 be present, its use is NOT RECOMMENDED, because the passing of authentication information in clear  
3727 text (such as URIs) has proven to be a security risk in almost every case where it has been used. For  
3728 instance, transporting a PIN number in this field exposes the PIN.

3729 Note that the password field is just an extension of user portion. Implementations not wishing to give  
3730 special significance to the password portion of the field MAY simply treat “user:password” as a single  
3731 string.

3732 **host:** The host providing the SIP resource. The host part contains either a fully-qualified domain name  
3733 or numeric IPv4 or IPv6 address. Using the fully-qualified domain name form is RECOMMENDED  
3734 whenever possible.

3735 **port:** The port number where the request is to be sent.

3736 **URI parameters:** Parameters affecting a request constructed from the URI.

3737 URI parameters are added after the hostport component and are separated by semi-colons.

3738 URI parameters take the form:

3739 parameter-name “=” parameter-value

3740 Even though an arbitrary number of URI parameters may be included in a URI, any given parameter-  
3741 name MUST NOT appear more than once.

3742 This extensible mechanism includes the `transport`, `maddr`, `ttl`, `user`, `method` and `lr` parameters.

3743 The `transport` parameter determines the transport mechanism to be used for sending SIP messages,  
3744 as specified in [4]. SIP can use any network transport protocol. Parameter names are defined for  
3745 UDP [14], TCP [15], and SCTP [17]. For a SIPS URI, the `transport` parameter MUST indicate a  
3746 reliable transport.

3747 The `maddr` parameter indicates the server address to be contacted for this user, overriding any address  
3748 derived from the `host` field. When an `maddr` parameter is present, the `port` and `transport` components  
3749 of the URI apply to the address indicated in the `maddr` parameter value. [4] describes the proper  
3750 interpretation of the `transport`, `maddr`, and `hostport` in order to obtain the destination address, `port`,  
3751 and `transport` for sending a request.

3752 The `maddr` field has been used as a simple form of loose source routing. It allows a URI to specify a proxy  
3753 that must be traversed en-route to the destination. Continuing to use the `maddr` parameter this way is strongly  
3754 discouraged (the mechanisms that enable it are deprecated). Implementations should instead use the `Route`  
3755 mechanism described in this document, establishing a pre-existing route set if necessary (see Section 8.1.1.1).  
3756 This provides a full URI to describe the node to be traversed.

3757 The `ttl` parameter determines the time-to-live value of the UDP multicast packet and MUST only be  
3758 used if `maddr` is a multicast address and the transport protocol is UDP. For example, to specify to call  
3759 `alice@atlanta.com` using multicast to `239.255.255.1` with a `ttl` of 15, the following URI would  
3760 be used:

3761 `sip:alice@atlanta.com;maddr=239.255.255.1;ttl=15`

3762 The set of valid `telephone-subscriber` strings is a subset of valid `user` strings. The `user` URI pa-  
3763 rameter exists to distinguish telephone numbers from user names that happen to look like telephone  
3764 numbers. If the user string contains a telephone number formatted as a `telephone-subscriber`, the  
3765 `user` parameter value “`phone`” SHOULD be present. Even without this parameter, recipients of SIP  
3766 and SIPS URIs MAY interpret the pre-@ part as a telephone number if local restrictions on the name  
3767 space for user name allow it.

3768 The method of the SIP request constructed from the URI can be specified with the `method` parameter.  
3769 The `lr` parameter, when present, indicates that the element responsible for this resource implements  
3770 the routing mechanisms specified in this document. This parameter will be used in the URIs proxies  
3771 place into `Record-Route` header field values, and may appear in the URIs in a pre-existing route set.

3772 This parameter is used to achieve backwards compatibility with systems implementing the strict-routing  
3773 mechanisms of RFC 2543 and the `rfc2543bis` drafts up to `bis-05`. An element preparing to send a request  
3774 based on a URI not containing this parameter can assume the receiving element implements strict-routing and  
3775 reformat the message to preserve the information in the `Request-URI`.

3776 Since the `uri`-parameter mechanism is extensible, SIP elements MUST silently ignore any `uri`-parameters  
3777 that they do not understand.

3778 **Headers:** Header fields to be included in a request constructed from the URI.

3779 Headers fields in the SIP request can be specified with the “?” mechanism within a URI. The header  
 3780 names and values are encoded in ampersand separated hname = hvalue pairs. The special hname  
 3781 “body” indicates that the associated hvalue is the message-body of the SIP request.

3782 Table 1 summarizes the use of SIP and SIPS URI components based on the context in which the URI  
 3783 appears. The external column describes URIs appearing anywhere outside of a SIP message, for instance on  
 3784 a web page or business card. Entries marked “m” are mandatory, those marked “o” are optional, and those  
 3785 marked “-” are not allowed. Elements processing URIs SHOULD ignore any disallowed components if they  
 3786 are present. The second column indicates the default value of an optional element if it is not present. “-”  
 3787 indicates that the element is either not optional, or has no default value.

3788 URIs in Contact header fields have different restrictions depending on the context in which the header  
 3789 field appears. One set applies to messages that establish and maintain dialogs (INVITE and its 200 (OK)  
 3790 response). The other applies to registration and redirection messages (REGISTER, its 200 (OK) response,  
 3791 and 3xx class responses to any method).

	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	o
host	-	m	m	m	m	m	m
port	(1)	o	-	-	o	o	o
user-param	ip	o	o	o	o	o	o
method	INVITE	-	-	-	-	-	o
maddr-param	-	o	-	-	o	o	o
ttl-param	1	o	-	-	o	-	o
transp.-param	(2)	o	-	-	o	o	o
lr-param	-	o	-	-	-	o	o
other-param	-	o	o	o	o	o	o
headers	-	-	-	-	o	-	o

(1): The default port value is transport and scheme dependent. The default is 5060 for sip: using UDP, TCP, or SCTP. The default is 5061 for sip: using TLS over TCP and sips: over TCP.

(2): The default transport is scheme dependent. For sip:, it is UDP. For sips:, it is TCP.

Table 1: Use and default values of URI components for SIP header field values, Request-URI and references

### 3792 19.1.2 Character Escaping Requirements

3793 SIP follows the requirements and guidelines of RFC 2396 [5] when defining the set of characters that must  
 3794 be escaped in a SIP URI, and uses its “”%” HEX HEX” mechanism for escaping. From RFC 2396:

3795 The set of characters actually reserved within any given URI component is defined by that com-  
 3796 ponent. In general, a character is reserved if the semantics of the URI changes if the character  
 3797 is replaced with its escaped US-ASCII encoding. [5].

3798 Excluded US-ASCII characters [5], such as space and control characters and characters used as URI delimiters, also MUST be escaped. URIs MUST NOT contain unescaped space and control characters.

3799 For each component, the set of valid BNF expansions defines exactly which characters may appear unescaped. All other characters MUST be escaped.

3800 For example, “@” is not in the set of characters in the user component, so the user “j@s0n” must have at least the @ sign encoded, as in “j%40s0n”.

3801 Expanding the `hname` and `hvalue` tokens in Section 25 show that all URI reserved characters in header field names and values MUST be escaped.

3802 The `telephone-subscriber` subset of the `user` component has special escaping considerations. The set of characters not reserved in the RFC 2806 [9] description of `telephone-subscriber` contains a number of characters in various syntax elements that need to be escaped when used in SIP URIs. Any characters occurring in a `telephone-subscriber` that do not appear in an expansion of the BNF for the `user` rule MUST be escaped.

3803 Note that character escaping is not allowed in the host component of a SIP or SIPS URI (the % character is not valid in its expansion). This is likely to change in the future as requirements for Internationalized Domain Names are finalized. Current implementations MUST NOT attempt to improve robustness by treating received escaped characters in the host component as literally equivalent to their unescaped counterpart. The behavior required to meet the requirements of IDN may be significantly different.

### 3816 19.1.3 Example SIP and SIPS URIs

```
3817 sip:alice@atlanta.com
3818 sip:alice:secretword@atlanta.com;transport=tcp
3819 sips:alice@atlanta.com?subject=project%20x&priority=urgent
3820 sip:+1-212-555-1212:1234@gateway.com;user=phone
3821 sips:1212@gateway.com
3822 sip:alice@192.0.2.4
3823 sip:atlanta.com;method=REGISTER?to=alice%40atlanta.com
3824 sip:alice;day=tuesday@atlanta.com
```

3825 The last sample URI above has a `user` field value of “alice;day=tuesday”. The escaping rules defined above allow a semicolon to appear unescaped in this field. For the purposes of this protocol, the field is opaque. The structure of that value is only useful to the SIP element responsible for the resource.

### 3828 19.1.4 URI Comparison

3829 Some operations in this specification require determining whether two SIP or SIPS URIs are equivalent. In this specification, registrars need to compare bindings in `Contact` URIs in `REGISTER` requests (see Section 10.3.) SIP and SIPS URIs are compared for equality according to the following rules:

- 3832 ● A SIP and SIPS URI are not equivalent, even if the rest of the URIs are equivalent.
- 3833 ● Comparison of the `userinfo` of SIP and SIPS URIs is case-sensitive. This includes `userinfo` containing passwords or formatted as `telephone-subscribers`. Comparison of all other components of the URI is case-insensitive unless explicitly defined otherwise.
- 3834 ● The ordering of parameters and header fields is not significant in comparing SIP and SIPS URIs.

3837 • Characters other than those in the “reserved” and “unsafe” sets (see RFC 2396 [5]) are equivalent to  
3838 their “”%” HEX HEX” encoding.

3839 • An IP address that is the result of a DNS lookup of a host name does **not** match that host name.

3840 • For two URIs to be equal, the **user**, **password**, **host**, and **port** components must match.

3841 A URI omitting the user component will *not* match a URI that includes one. A URI omitting the  
3842 password component will **not** match a URI that includes one.

3843 A URI omitting any component with a default value will *not* match a URI explicitly containing that  
3844 component with its default value. For instance, a URI omitting the optional port component will  
3845 *not* match a URI explicitly declaring port 5060. The same is true for the **transport-parameter**, **ttl-**  
3846 **parameter**, **user-parameter**, and **method** components.

3847 Defining sip:user@host to *not* be equivalent to sip:user@host:5060 is a change from RFC 2543. When de-  
3848 riving addresses from URIs, equivalent addresses are expected from equivalent URIs. The URI sip:user@host:5060  
3849 will always resolve to port 5060. The URI sip:user@host may resolve to other ports through the DNS SRV  
3850 mechanisms detailed in [4].

3851 • URI **uri-parameter** components are compared as follows

3852 – Any **uri-parameter** appearing in both URIs must match.

3853 – A **user**, **ttl**, or **method uri-parameter** appearing in only one URI never matches, even if it  
3854 contains the default value.

3855 – A URI that includes an **maddr** parameter will *not* match a URI that contains no **maddr** param-  
3856 eter.

3857 – All other **uri-parameters** appearing in only one URI are ignored when comparing the URIs.

3858 • URI **header** components are never ignored. Any present **header** component **MUST** be present in  
3859 both URIs and match for the URIs to match. The matching rules are defined for each header field in  
3860 Section 20.

3861 The URIs within each of the following sets are equivalent:

3862 sip:%61lice@atlanta.com;transport=TCP

3863 sip:alice@AtLanTa.CoM;Transport=tcp

3864 sip:carol@chicago.com

3865 sip:carol@chicago.com;newparam=5

3866 sip:carol@chicago.com;security=on

3867 sip:biloxi.com;transport=tcp;method=REGISTER?to=sip:bob%40biloxi.com

3868 sip:biloxi.com;method=REGISTER;transport=tcp?to=sip:bob%40biloxi.com

3869 sip:alice@atlanta.com?subject=project%20x&priority=urgent

3870 sip:alice@atlanta.com?priority=urgent&subject=project%20x

3871 The URIs within each of the following sets are **not** equivalent:

3872 SIP:ALICE@AtLanTa.CoM;Transport=udp (different usernames)  
3873 sip:alice@AtLanTa.CoM;Transport=UDP

3874 sip:bob@biloxi.com (can resolve to different ports)  
3875 sip:bob@biloxi.com:5060

3876 sip:bob@biloxi.com (can resolve to different transports)  
3877 sip:bob@biloxi.com;transport=udp

3878 sip:bob@biloxi.com (can resolve to different port and transports)  
3879 sip:bob@biloxi.com:6000;transport=tcp

3880 sip:carol@chicago.com (different header component)  
3881 sip:carol@chicago.com?Subject=next%20meeting

3882 sip:bob@phone21.bboxesbybob.com (even though that's what  
3883 sip:bob@192.0.2.4 phone21.bboxesbybob.com resolves to)

3884 Note that equality is not transitive:

3885 sip:carol@chicago.com and sip:carol@chicago.com;security=on are equivalent

3886 and sip:carol@chicago.com and sip:carol@chicago.com;security=off are equivalent

3887 But sip:carol@chicago.com;security=on and sip:carol@chicago.com;security=off are **not** equivalent

### 3888 19.1.5 Forming Requests from a URI

3889 An implementation needs to take care when forming requests directly from a URI. URIs from business cards,  
3890 web pages, and even from sources inside the protocol such as registered contacts may contain inappropriate  
3891 header fields or body parts.

3892 An implementation **MUST** include any provided transport, maddr, ttl, or user parameter in the Request-  
3893 URI of the formed request. If the URI contains a method parameter, its value **MUST** be used as the method  
3894 of the request. The method parameter **MUST NOT** be placed in the Request-URI. Unknown URI parameters  
3895 **MUST** be placed in the message's Request-URI.

3896 An implementation **SHOULD** treat the presence of any headers or body parts in the URI as a desire to  
3897 include them in the message, and choose to honor the request on a per-component basis.

3898 An implementation **SHOULD NOT** honor these obviously dangerous header fields: From, Call-ID, CSeq,  
3899 Via, and Record-Route.

3900 An implementation **SHOULD NOT** honor any requested Route header field values in order to not be used  
3901 as an unwitting agent in malicious attacks.

3902 An implementation SHOULD NOT honor requests to include header fields that may cause it to falsely ad-  
3903 vertise its location or capabilities. These include: Accept, Accept-Encoding, Accept-Language, Allow,  
3904 Contact (in its dialog usage), Organization, Supported, and User-Agent.

3905 An implementation SHOULD verify the accuracy of any requested descriptive header fields, including:  
3906 Content-Disposition, Content-Encoding, Content-Language, Content-Length, Content-Type, Date,  
3907 Mime-Version, and Timestamp.

3908 If the request formed from constructing a message from a given URI is not a valid SIP request, the URI  
3909 is invalid. An implementation MUST NOT proceed with transmitting the request. It should instead pursue  
3910 the course of action due an invalid URI in the context it occurs.

3911 The constructed request can be invalid in many ways. These include, but are not limited to, syntax error in  
3912 header fields, invalid combinations of URI parameters, or an incorrect description of the message body.

3913 Sending a request formed from a given URI may require capabilities unavailable to the implementation.  
3914 The URI might indicate use of an unimplemented transport or extension, for example. An implementation  
3915 SHOULD refuse to send these requests rather than modifying them to match their capabilities. An imple-  
3916 mentation MUST NOT send a request requiring an extension that it does not support.

3917 For example, such a request can be formed through the presence of a Require header parameter or a method  
3918 URI parameter with an unknown or explicitly unsupported value.

#### 3919 **19.1.6 Relating SIP URIs and tel URLs**

3920 When a tel URL [9] is converted to a SIP or SIPS URI, the entire telephone-subscriber portion of the tel  
3921 URL, including any parameters, is placed into the userinfo part of the SIP or SIPS URI.

3922 Thus, tel:+358-555-1234567;postd=pp22 becomes

3923 sip:+358-555-1234567;postd=pp22@foo.com;user=phone

3924 or

3925 sips:+358-555-1234567;postd=pp22@foo.com;user=phone

3926 not

3927 sip:+358-555-1234567@foo.com;postd=pp22;user=phone

3928 or

3929 sips:+358-555-1234567@foo.com;postd=pp22;user=phone

3930 In general, equivalent "tel" URLs converted to SIP or SIPS URIs in this fashion may not produce equiv-  
3931 alent SIP or SIPS URIs. The userinfo of SIP and SIPS URIs are compared as a case-sensitive string.  
3932 Variance in case-insensitive portions of tel URLs and reordering of tel URL parameters does not affect tel  
3933 URL equivalence, but does affect the equivalence of SIP URIs formed from them.

3934 For example,

3935 tel:+358-555-1234567;postd=pp22

3936 tel:+358-555-1234567;POSTD=PP22

3937 are equivalent, while

3938 sip:+358-555-1234567;postd=pp22@foo.com;user=phone

3939 sip:+358-555-1234567;POSTD=PP22@foo.com;user=phone

3940 are not.

3941 Likewise,

3942 tel:+358-555-1234567;postd=pp22;isub=1411

3943 tel:+358-555-1234567;isub=1411;postd=pp22

3944 are equivalent, while

3945 sip:+358-555-1234567;postd=pp22;isub=1411@foo.com;user=phone

3946 sip:+358-555-1234567;isub=1411;postd=pp22@foo.com;user=phone

3947 are not.

3948 To mitigate this problem, elements constructing telephone-subscriber fields to place in the userinfo part  
3949 of a SIP or SIPS URI SHOULD fold any case-insensitive portion of telephone-subscriber to lower case,  
3950 and order the telephone-subscriber parameters lexically by parameter name. (All components of a tel URL  
3951 except for future-extension parameters are defined to be compared case-insensitive.)

3952 Following this suggestion, both

3953 tel:+358-555-1234567;postd=pp22

3954 tel:+358-555-1234567;POSTD=PP22

3955 become

3956 sip:+358-555-1234567;postd=pp22@foo.com;user=phone

3957 and both

3958 tel:+358-555-1234567;postd=pp22;isub=1411

3959 tel:+358-555-1234567;isub=1411;postd=pp22

3960 become

3961 sip:+358-555-1234567;isub=1411;postd=pp22;user=phone

## 3962 19.2 Option Tags

3963 Option tags are unique identifiers used to designate new options (extensions) in SIP. These tags are used in  
3964 Require (Section 20.32), Proxy-Require (Section 20.29), Supported (Section 20.37) and Unsupported  
3965 (Section 20.40) header fields. Note that these options appear as parameters in those header fields in an  
3966 option-tag = token form (see Section 25 for the definition of token).

3967 The creator of a new SIP option MUST either prefix the option with their reverse domain name or register  
3968 the new option with the Internet Assigned Numbers Authority (IANA) (See Section 27).

3969 An example of a reverse-domain-name option is "com.foo.mynewfeature", whose inventor can be reached  
3970 at "foo.com". For these features, individual organizations are responsible for ensuring that option names do  
3971 not collide within the same domain. The domain name part of the option MUST use lower-case; the option

3972 name is case-insensitive.

3973 Options registered with IANA do not contain periods and are globally unique. IANA option tags are  
3974 case-insensitive.

### 3975 19.3 Tags

3976 The “tag” parameter is used in the To and From header fields of SIP messages. It serves as a general  
3977 mechanism to identify a dialog, which is the combination of the Call-ID along with two tags, one from  
3978 each participant in the dialog. When a UA sends a request outside of a dialog, it contains a From tag only,  
3979 providing “half” of the dialog ID. The dialog is completed from the response(s), each of which contributes  
3980 the second half in the To header field. The forking of SIP requests means that multiple dialogs can be  
3981 established from a single request. This also explains the need for the two-sided dialog identifier; without a  
3982 contribution from the recipients, the originator could not disambiguate the multiple dialogs established from  
3983 a single request.

3984 When a tag is generated by a UA for insertion into a request or response, it MUST be globally unique  
3985 and cryptographically random with at least 32 bits of randomness. A property of this selection requirement  
3986 is that a UA will place a different tag into the From header of an INVITE as it would place into the To  
3987 header of the response to the same INVITE. This is needed in order for a UA to invite itself to a session, a  
3988 common case for “hairpinning” of calls in PSTN gateways. Similarly, two INVITEs for different calls will  
3989 have different From tags.

3990 Besides the requirement for global uniqueness, the algorithm for generating a tag is implementation-  
3991 specific. Tags are helpful in fault tolerant systems, where a dialog is to be recovered on an alternate server  
3992 after a failure. A UAS can select the tag in such a way that a backup can recognize a request as part of a  
3993 dialog on the failed server, and therefore determine that it should attempt to recover the dialog and any other  
3994 state associated with it.

## 3995 20 Header Fields

3996 The general syntax for header fields is covered in Section 7.3. This section lists the full set of header fields  
3997 along with notes on syntax, meaning, and usage. Throughout this section, we use [HX.Y] to refer to Section  
3998 X.Y of the current HTTP/1.1 specification RFC 2616 [8]. Examples of each header field are given.

3999 Information about header fields in relation to methods and proxy processing is summarized in Tables 2  
4000 and 3.

4001 The “where” column describes the request and response types in which the header field can be used.  
4002 Values in this column are:

4003 **R:** header field may only appear in requests;

4004 **r:** header field may only appear in responses;

4005 **2xx, 4xx, etc.:** A numerical value or range indicates response codes with which the header field can be  
4006 used;

4007 **c:** header field is copied from the request to the response.

4008 An empty entry in the “where” column indicates that the header field may be present in all requests and  
4009 responses.

4010 The “proxy” column describes the operations a proxy may perform on a header field:

4011 **a:** A proxy can add or concatenate the header field if not present.

4012 **m:** A proxy can modify an existing header field value.

4013 **d:** A proxy can delete a header field value.

4014 **r:** A proxy must be able to read the header field, and thus this header field cannot be encrypted.

4015 The next six columns relate to the presence of a header field in a method:

4016 **c:** Conditional; the header field is either mandatory or optional, depending on the presence of a route set or  
4017 the response code.

4018 **m:** The header field is mandatory.

4019 **m\*:** The header field SHOULD be sent, but clients/servers need to be prepared to receive messages without  
4020 that header field.

4021 **o:** The header field is optional.

4022 **t:** The header field SHOULD be sent, but clients/servers need to be prepared to receive messages without  
4023 that header field. If a stream-based protocol (such as TCP) is used as a transport, then the header field  
4024 MUST be sent.

4025 **\*:** The header field is required if the message body is not empty. See sections 20.14, 20.15 and 7.4 for  
4026 details.

4027 **-:** The header field is not applicable.

4028 “Optional” means that a UA MAY include the header field in a request or response, and a UA MAY ignore  
4029 the header field if present in the request or response (The exception to this rule is the **Require** header field  
4030 discussed in 20.32). A “mandatory” header field MUST be present in a request, and MUST be understood  
4031 by the UAS receiving the request. A mandatory response header field MUST be present in the response, and  
4032 the header field MUST be understood by the UAC processing the response. “Not applicable” means that the  
4033 header field MUST NOT be present in a request. If one is placed in a request by mistake, it MUST be ignored  
4034 by the UAS receiving the request. Similarly, a header field labeled “not applicable” for a response means  
4035 that the UAS MUST NOT place the header field in the response, and the UAC MUST ignore the header field  
4036 in the response.

4037 A UA SHOULD ignore extension header parameters that are not understood.

4038 A compact form of some common header field names is also defined for use when overall message size  
4039 is an issue.

4040 The **Contact**, **From**, and **To** header fields contain a URI. If the URI contains a comma, question mark  
4041 or semicolon, the URI MUST be enclosed in angle brackets (< and >). Any URI parameters are contained  
4042 within these brackets. If the URI is not enclosed in angle brackets, any semicolon-delimited parameters are  
4043 header-parameters, not URI parameters.

Header field	where	proxy	ACK	BYE	CAN	INV	OPT	REG
Accept	R		-	o	-	o	m*	o
Accept	2xx		-	-	-	o	m*	o
Accept	415		-	o	-	o	o	o
Accept-Encoding	R		-	o	-	o	o	o
Accept-Encoding	2xx		-	-	-	o	m*	o
Accept-Encoding	415		-	o	-	o	o	o
Accept-Language	R		-	o	-	o	o	o
Accept-Language	2xx		-	-	-	o	m*	o
Accept-Language	415		-	o	-	o	o	o
Alert-Info	R	ar	-	-	-	o	-	-
Alert-Info	180	ar	-	-	-	o	-	-
Allow	R		-	o	-	o	o	o
Allow	2xx		-	o	-	m*	m*	o
Allow	r		-	o	-	o	o	o
Allow	405		-	m	-	m	m	m
Authentication-Info	2xx		-	o	-	o	o	o
Authorization	R		o	o	o	o	o	o
Call-ID	c	r	m	m	m	m	m	m
Call-Info		ar	-	-	-	o	o	o
Contact	R		o	-	-	m	o	o
Contact	1xx		-	-	-	o	-	-
Contact	2xx		-	-	-	m	o	o
Contact	3xx	d	-	o	-	o	o	o
Contact	485		-	o	-	o	o	o
Content-Disposition			o	o	-	o	o	o
Content-Encoding			o	o	-	o	o	o
Content-Language			o	o	-	o	o	o
Content-Length		ar	t	t	t	t	t	t
Content-Type			*	*	-	*	*	*
CSeq	c	r	m	m	m	m	m	m
Date		a	o	o	o	o	o	o
Error-Info	300-699	a	-	o	o	o	o	o
Expires			-	-	-	o	-	o
From	c	r	m	m	m	m	m	m
In-Reply-To	R		-	-	-	o	-	-
Max-Forwards	R	amr	m	m	m	m	m	m
Min-Expires	423		-	-	-	-	-	m
MIME-Version			o	o	-	o	o	o
Organization		ar	-	-	-	o	o	o

Table 2: Summary of header fields, A–O

Header field	where	proxy	ACK	BYE	CAN	INV	OPT	REG
Priority	R	ar	-	-	-	o	-	-
Proxy-Authenticate	407		-	m	-	m	m	m
Proxy-Authenticate	401		-	o	o	o	o	o
Proxy-Authorization	R	dr	o	o	-	o	o	o
Proxy-Require	R	ar	-	o	-	o	o	o
Record-Route	R	ar	o	o	o	o	o	-
Record-Route	2xx,18x	mr	-	o	o	o	o	-
Reply-To			-	-	-	o	-	-
Require		ar	-	o	-	o	o	o
Retry-After	404,413,480,486		-	o	o	o	o	o
	500,503		-	o	o	o	o	o
	600,603		-	o	o	o	o	o
Route	R	adr	c	c	c	c	c	c
Server	r		-	o	o	o	o	o
Subject	R		-	-	-	o	-	-
Supported	R		-	o	o	m*	o	o
Supported	2xx		-	o	o	m*	m*	o
Timestamp			o	o	o	o	o	o
To	c(1)	r	m	m	m	m	m	m
Unsupported	420		-	o	o	o	o	o
User-Agent			o	o	o	o	o	o
Via	R	amr	m	m	m	m	m	m
Via	rc	dr	m	m	m	m	m	m
Warning	r		-	o	o	o	o	o
WWW-Authenticate	401		-	m	-	m	m	m
WWW-Authenticate	407		-	o	-	o	o	o

Table 3: Summary of header fields, P-Z; (1): copied with possible addition of tag

## 4044 20.1 Accept

4045 The Accept header field follows the syntax defined in [H14.1]. The semantics are also identical, with  
 4046 the exception that if no Accept header field is present, the server SHOULD assume a default value of  
 4047 application/sdp.

4048 An empty Accept header field means that no formats are acceptable.

4049 Example:

4050 Accept: application/sdp;level=1, application/x-private, text/html

## 4051 20.2 Accept-Encoding

4052 The Accept-Encoding header field is similar to Accept, but restricts the content-codings [H3.5] that are  
 4053 acceptable in the response. See [H14.3]. The syntax of this header field is defined in [H14.3]. The semantics  
 4054 in SIP are identical to those defined in [H14.3].

4055 An empty **Accept-Encoding** header field is permissible, even though the syntax in [H14.3] does not  
4056 provide for it. It is equivalent to **Accept-Encoding: identity**, that is, only the identity encoding, meaning  
4057 no encoding, is permissible.

4058 If no **Accept-Encoding** header field is present, the server **SHOULD** assume a default value of **identity**.

4059 This differs slightly from the HTTP definition, which indicates that when not present, any encoding can  
4060 be used, but the identity encoding is preferred.

4061 Example:

4062 `Accept-Encoding: gzip`

### 4063 **20.3 Accept-Language**

4064 The **Accept-Language** header field is used in requests to indicate the preferred languages for reason  
4065 phrases, session descriptions, or status responses carried as message bodies in the response. If no **Accept-**  
4066 **Language** header field is present, the server **SHOULD** assume all languages are acceptable to the client.

4067 The **Accept-Language** header field follows the syntax defined in [H14.4]. The rules for ordering the  
4068 languages based on the “q” parameter apply to SIP as well.

4069 Example:

4070 `Accept-Language: da, en-gb;q=0.8, en;q=0.7`

### 4071 **20.4 Alert-Info**

4072 When present in an **INVITE** request, the **Alert-Info** header field specifies an alternative ring tone to the UAS.  
4073 When present in a 180 (Ringing) response, the **Alert-Info** header field specifies an alternative ringback tone  
4074 to the UAC. A typical usage is for a proxy to insert this header field to provide a distinctive ring feature.

4075 The **Alert-Info** header field can introduce security risks. These risks and the ways to handle them are  
4076 discussed in Section 20.9, which discusses the **Call-Info** header field since the risks are identical.

4077 In addition, a user **SHOULD** be able to disable this feature selectively.

4078 This helps prevent disruptions that could result from the use of this header field by untrusted elements.

4079 Example:

4080 `Alert-Info: <http://www.example.com/sounds/moo.wav>`

### 4081 **20.5 Allow**

4082 The **Allow** header field lists the set of methods supported by the UA generating the message.

4083 All methods, including **ACK** and **CANCEL**, understood by the UA **MUST** be included in the list of  
4084 methods in the **Allow** header field, when present. The absence of an **Allow** header field **MUST NOT** be  
4085 interpreted to mean that the UA sending the message supports no methods. Rather, it implies that the UA is  
4086 not providing any information on what methods it supports.

4087 Supplying an **Allow** header field in responses to methods other than **OPTIONS** reduces the number of  
4088 messages needed.

4089 Example:

4090 `Allow: INVITE, ACK, OPTIONS, CANCEL, BYE`

## 4091 20.6 Authentication-Info

4092 The Authentication-Info header field provides for mutual authentication with HTTP Digest. A UAS MAY  
4093 include this header field in a 2xx response to a request that was successfully authenticated using digest based  
4094 on the Authorization header field.

4095 Syntax and semantics follow those specified in RFC 2617 [18].

4096 Example:

```
4097 Authentication-Info: nextnonce="47364c23432d2e131a5fb210812c"
```

## 4098 20.7 Authorization

4099 The Authorization header field contains authentication credentials of a UA. Section 22.2 overviews the use  
4100 of the Authorization header field, and Section 22.4 describes the syntax and semantics when used with  
4101 HTTP authentication.

4102 This header field, along with Proxy-Authorization, breaks the general rules about multiple header field  
4103 values. Although not a comma-separated list, this header field name may be present multiple times, and  
4104 MUST NOT be combined into a single header line using the usual rules described in Section 7.3.

4105 In the example below, there are no quotes around the Digest parameter:

```
4106 Authorization: Digest username="Alice", realm="atlanta.com",  
4107 nonce="84a4cc6f3082121f32b42a2187831a9e",  
4108 response="7587245234b3434cc3412213e5f113a5432"
```

## 4109 20.8 Call-ID

4110 The Call-ID header field uniquely identifies a particular invitation or all registrations of a particular client.  
4111 A single multimedia conference can give rise to several calls with different Call-IDs, for example, if a user  
4112 invites a single individual several times to the same (long-running) conference. Call-IDs are case-sensitive  
4113 and are simply compared byte-by-byte.

4114 The compact form of the Call-ID header field is i.

4115 Examples:

```
4116 Call-ID: f81d4fae-7dec-11d0-a765-00a0c91e6bf6@biloxi.com  
4117 i:f81d4fae-7dec-11d0-a765-00a0c91e6bf6@192.0.2.4
```

## 4118 20.9 Call-Info

4119 The Call-Info header field provides additional information about the caller or callee, depending on whether  
4120 it is found in a request or response. The purpose of the URI is described by the "purpose" parameter.  
4121 The "icon" parameter designates an image suitable as an iconic representation of the caller or callee. The  
4122 "info" parameter describes the caller or callee in general, for example, through a web page. The "card"  
4123 parameter provides a business card, for example, in vCard [35] or LDIF [36] formats. Additional tokens can  
4124 be registered using IANA and the procedures in Section 27.

4125 Use of the **Call-Info** header field can pose a security risk. If a callee fetches the URIs provided by a  
4126 malicious caller, the callee may be at risk for displaying inappropriate or offensive content, dangerous or  
4127 illegal content, and so on. Therefore, it is RECOMMENDED that a UA only render the information in the  
4128 **Call-Info** header field if it can verify the authenticity of the element that originated the header field and  
4129 trusts that element. This need not be the peer UA; a proxy can insert this header field into requests.

4130 Example:

```
4131 Call-Info: <http://www.example.com/alice/photo.jpg> ;purpose=icon,  
4132 <http://www.example.com/alice/> ;purpose=info
```

## 4133 20.10 Contact

4134 A **Contact** header field value provides a URI whose meaning depends on the type of request or response it  
4135 is in.

4136 A **Contact** header field value can contain a display name, a URI with URI parameters, and header  
4137 parameters.

4138 This document defines the **Contact** parameters “q” and “expires”. These parameters are only used  
4139 when the **Contact** is present in a REGISTER request or response, or in a 3xx response. Additional param-  
4140 eters may be defined in other specifications.

4141 When the header field value contains a display name, the URI including all URI parameters is enclosed  
4142 in “<” and “>”. If no “<” and “>” are present, all parameters after the URI are header parameters, not URI  
4143 parameters. The display name can be tokens, or a quoted string, if a larger character set is desired.

4144 Even if the “display-name” is empty, the “name-addr” form MUST be used if the “addr-spec” con-  
4145 tains a comma, semicolon, or question mark. There may or may not be LWS between the display-name  
4146 and the “<”.

4147 These rules for parsing a display name, URI and URI parameters, and header parameters also apply for  
4148 the header fields **To** and **From**.

4149 The **Contact** header field has a role similar to the **Location** header field in HTTP. However, the HTTP header  
4150 field only allows one address, unquoted. Since URIs can contain commas and semicolons as reserved characters,  
4151 they can be mistaken for header or parameter delimiters, respectively.

4152 The compact form of the **Contact** header field is **m** (for “moved”).

4153 The second example below shows a **Contact** header field value containing both a URI parameter  
4154 (**transport**) and a header parameter (**expires**).

```
4155 Contact: "Mr. Watson" <sip:watson@worchester.bell-telephone.com>  
4156 ;q=0.7; expires=3600,  
4157 "Mr. Watson" <mailto:watson@bell-telephone.com> ;q=0.1  
4158 m: <sips:bob@192.0.2.4>;expires=60
```

## 4159 20.11 Content-Disposition

4160 The **Content-Disposition** header field describes how the message body or, for multipart messages, a mes-  
4161 sage body part is to be interpreted by the UAC or UAS. This SIP header field extends the MIME **Content-**  
4162 **Type** (RFC 2183 [19]).

4163 The value “**session**” indicates that the body part describes a session, for either calls or early (pre-call)  
4164 media. The value “**render**” indicates that the body part should be displayed or otherwise rendered to the

4165 user. For backward-compatibility, if the **Content-Disposition** header field is missing, the server SHOULD  
4166 assume bodies of **Content-Type** `application/sdp` are the disposition “**session**”, while other content  
4167 types are “**render**”.

4168 The disposition type “**icon**” indicates that the body part contains an image suitable as an iconic repre-  
4169 sentation of the caller or callee. The value “**alert**” indicates that the body part contains information, such as  
4170 an audio clip, that should be rendered instead of ring tone.

4171 The handling parameter, **handling-param**, describes how the UAS should react if it receives a message  
4172 body whose content type or disposition type it does not understand. The parameter has defined values  
4173 of “**optional**” and “**required**”. If the handling parameter is missing, the value “**required**” SHOULD be  
4174 assumed.

4175 If this header field is missing, the MIME type determines the default content disposition. If there is  
4176 none, “**render**” is assumed.

4177 Example:

```
4178 Content-Disposition: session
```

## 4179 20.12 Content-Encoding

4180 The **Content-Encoding** header field is used as a modifier to the “**media-type**”. When present, its value  
4181 indicates what additional content codings have been applied to the entity-body, and thus what decoding  
4182 mechanisms MUST be applied in order to obtain the media-type referenced by the **Content-Type** header  
4183 field. **Content-Encoding** is primarily used to allow a body to be compressed without losing the identity of  
4184 its underlying media type.

4185 If multiple encodings have been applied to an entity-body, the content codings MUST be listed in the  
4186 order in which they were applied.

4187 All content-coding values are case-insensitive. IANA acts as a registry for content-coding value tokens.  
4188 See [H3.5] for a definition of the syntax for content-coding.

4189 Clients MAY apply content encodings to the body in requests. A server MAY apply content encodings to  
4190 the bodies in responses. The server MUST only use encodings listed in the **Accept-Encoding** header field  
4191 in the request.

4192 The compact form of the **Content-Encoding** header field is **e**. Examples:

```
4193 Content-Encoding: gzip  
4194 e: tar
```

## 4195 20.13 Content-Language

4196 See [H14.12]. Example:

```
4197 Content-Language: fr
```

## 4198 20.14 Content-Length

4199 The **Content-Length** header field indicates the size of the message-body, in decimal number of octets,  
4200 sent to the recipient. Applications SHOULD use this field to indicate the size of the message-body to be

4201 transferred, regardless of the media type of the entity. If a stream-based protocol (such as TCP) is used as  
4202 transport, the header field **MUST** be used.

4203 The size of the message-body does *not* include the CRLF separating headers and body. Any **Content-**  
4204 **Length** greater than or equal to zero is a valid value. If no body is present in a message, then the **Content-**  
4205 **Length** header field value **MUST** be set to zero.

4206 The ability to omit **Content-Length** simplifies the creation of cgi-like scripts that dynamically generate re-  
4207 sponses.

4208 The compact form of the header field is l.

4209 Examples:

4210 Content-Length: 349

4211 l: 173

## 4212 20.15 Content-Type

4213 The **Content-Type** header field indicates the media type of the message-body sent to the recipient. The  
4214 “media-type” element is defined in [H3.7]. The **Content-Type** header field **MUST** be present if the body is  
4215 not empty. If the body is empty, and a **Content-Type** header field is present, it indicates that the body of the  
4216 specific type has zero length (for example, an empty audio file).

4217 The compact form of the header field is c.

4218 Examples:

4219 Content-Type: application/sdp

4220 c: text/html; charset=ISO-8859-4

## 4221 20.16 CSeq

4222 A **CSeq** header field in a request contains a single decimal sequence number and the request method.  
4223 The sequence number **MUST** be expressible as a 32-bit unsigned integer. The method part of **CSeq** is  
4224 case-sensitive. The **CSeq** header field serves to order transactions within a dialog, to provide a means to  
4225 uniquely identify transactions, and to differentiate between new requests and request retransmissions. Two  
4226 **CSeq** header fields are considered equal if the sequence number and the request method are identical.

4227 Example:

4228 CSeq: 4711 INVITE

## 4229 20.17 Date

4230 The **Date** header field contains a the date and time. Unlike HTTP/1.1, SIP only supports the most recent  
4231 RFC 1123 [20] format for dates. As in [H3.3], SIP restricts the time zone in **SIP-date** to “GMT”, while  
4232 RFC 1123 allows any time zone. **rfc1123-date** is case-sensitive.

4233 The **Date** header field reflects the time when the request or response is first sent.

4234 The **Date** header field can be used by simple end systems without a battery-backed clock to acquire a notion of  
4235 current time. However, in its GMT form, it requires clients to know their offset from GMT.

4236 Example:

4237 Date: Sat, 13 Nov 2010 23:29:00 GMT

## 4238 20.18 Error-Info

4239 The Error-Info header field provides a pointer to additional information about the error status response.

4240 SIP UACs have user interface capabilities ranging from pop-up windows and audio on PC softclients to audio-  
4241 only on "black" phones or endpoints connected via gateways. Rather than forcing a server generating an error to  
4242 choose between sending an error status code with a detailed reason phrase and playing an audio recording, the  
4243 Error-Info header field allows both to be sent. The UAC then has the choice of which error indicator to render to the  
4244 caller.

4245 A UAC MAY treat a SIP or SIPS URI in an Error-Info header field as if it were a Contact in a redirect  
4246 and generate a new INVITE, resulting in a recorded announcement session being established. A non-SIP  
4247 URI MAY be rendered to the user.

4248 Examples:

4249 SIP/2.0 404 The number you have dialed is not in service  
4250 Error-Info: <sip:not-in-service-recording@atlanta.com>

## 4251 20.19 Expires

4252 The Expires header field gives the relative time after which the message (or content) expires.

4253 The precise meaning of this is method dependent.

4254 The expiration time in an INVITE does *not* affect the duration of the actual session that may result  
4255 from the invitation. Session description protocols may offer the ability to express time limits on the session  
4256 duration, however.

4257 The value of this field is an integral number of seconds (in decimal) between 0 and  $(2^{31})-1$ , measured  
4258 from the receipt of the request.

4259 Example:

4260 Expires: 5

## 4261 20.20 From

4262 The From header field indicates the initiator of the request. This may be different from the initiator of the  
4263 dialog. Requests sent by the callee to the caller use the callee's address in the From header field.

4264 The optional "display-name" is meant to be rendered by a human user interface. A system SHOULD use  
4265 the display name "Anonymous" if the identity of the client is to remain hidden. Even if the "display-name"  
4266 is empty, the "name-addr" form MUST be used if the "addr-spec" contains a comma, question mark, or  
4267 semicolon. Syntax issues are discussed in Section 7.3.1.

4268 Section 12 describes how From header fields are compared for the purpose of matching requests to  
4269 dialogs. See Section 20.10 for the rules for parsing a display name, URI and URI parameters, and header  
4270 parameters.

4271 The compact form of the From header field is f.

4272 Examples:

4273 From: "A. G. Bell" <sip:agb@bell-telephone.com> ;tag=a48s  
4274 From: sip:+12125551212@server.phone2net.com;tag=887s  
4275 f: Anonymous <sip:c8oqz84zk7z@privacy.org>;tag=hyh8

## 4276 **20.21 In-Reply-To**

4277 The **In-Reply-To** header field enumerates the **Call-IDs** that this call references or returns. These **Call-IDs**  
4278 may have been cached by the client then included in this header field in a return call.

4279 This allows automatic call distribution systems to route return calls to the originator of the first call. This also  
4280 allows callees to filter calls, so that only return calls for calls they originated will be accepted. This field is not a  
4281 substitute for request authentication.

4282 Example:

4283 `In-Reply-To: 70710@saturn.bell-tel.com, 17320@saturn.bell-tel.com`

## 4284 **20.22 Max-Forwards**

4285 The **Max-Forwards** header field must be used with any SIP method to limit the number of proxies or  
4286 gateways that can forward the request to the next downstream server. This can also be useful when the client  
4287 is attempting to trace a request chain that appears to be failing or looping in mid-chain.

4288 The **Max-Forwards** value is an integer in the range 0-255 indicating the remaining number of times this  
4289 request message is allowed to be forwarded. This count is decremented by each server that forwards the  
4290 request. The recommended value is 70.

4291 This header field should be inserted by elements that can not otherwise guarantee loop detection. For  
4292 example, a B2BUA should insert a **Max-Forwards** header field.

4293 Example:

4294 `Max-Forwards: 6`

## 4295 **20.23 Min-Expires**

4296 The **Min-Expires** header field conveys the minimum refresh interval supported for soft-state elements man-  
4297 aged by that server. This includes **Contact** header fields that are stored by a registrar. The header field  
4298 contains a decimal integer number of seconds from 0 to (2\*\*32)-1. The use of the header field in a 423  
4299 (Registration Too Brief) response is described in Sections 10.2.8, 10.3, and 21.4.17.

4300 Example:

4301 `Min-Expires: 60`

## 4302 **20.24 MIME-Version**

4303 See [H19.4.1].

4304 Example:

4305 `MIME-Version: 1.0`

## 4306 **20.25 Organization**

4307 The **Organization** header field conveys the name of the organization to which the SIP element issuing the  
4308 request or response belongs.

4309           The field MAY be used by client software to filter calls.

4310       Example:

4311       Organization: Boxes by Bob

## 4312 **20.26 Priority**

4313       The Priority header field indicates the urgency of the request as perceived by the client. The Priority header  
4314       field describes the priority that the SIP request should have to the receiving human or its agent. For example,  
4315       it may be factored into decisions about call routing and acceptance. For these decisions, a message contain-  
4316       ing no Priority header field SHOULD be treated as if it specified a Priority of "non-urgent". The Priority  
4317       header field does not influence the use of communications resources such as packet forwarding priority in  
4318       routers or access to circuits in PSTN gateways. The header field can have the values "non-urgent", "normal",  
4319       "urgent", and "emergency", but additional values can be defined elsewhere. It is RECOMMENDED that the  
4320       value of "emergency" only be used when life, limb, or property are in imminent danger. Otherwise, there  
4321       are no semantics defined for this header field.

4322           These are the values of RFC 2076 [37], with the addition of "emergency".

4323       Examples:

4324       Subject: A tornado is heading our way!  
4325       Priority: emergency

4326       or

4327       Subject: Weekend plans  
4328       Priority: non-urgent

## 4329 **20.27 Proxy-Authenticate**

4330       A Proxy-Authenticate header field value contains an authentication challenge.

4331       The syntax for this header field and its use is defined in [H14.33]. See 22.3 for further details on its  
4332       usage.

4333       Example:

4334       Proxy-Authenticate: Digest realm="atlanta.com",  
4335       domain="sip:ssl.carrier.com",  
4336       nonce="f84f1cec41e6cbe5aea9c8e88d359",  
4337       opaque="", stale=FALSE, algorithm=MD5

## 4338 **20.28 Proxy-Authorization**

4339       The Proxy-Authorization header field allows the client to identify itself (or its user) to a proxy that requires  
4340       authentication. A Proxy-Authorization field value consists of credentials containing the authentication  
4341       information of the user agent for the proxy and/or realm of the resource being requested.

4342 See [H14.34] for a definition of the syntax, and section 22.3 for a discussion of its usage.

4343 This header field, along with **Authorization**, breaks the general rules about multiple header field names.  
4344 Although not a comma-separated list, this header field name may be present multiple times, and **MUST NOT**  
4345 be combined into a single header line using the usual rules described in Section 7.3.1.

4346 Example:

```
4347 Proxy-Authorization: Digest username="Alice", realm="atlanta.com",  
4348     nonce="c60f3082ee1212b402a21831ae",  
4349     response="245f23415f11432b3434341c022"
```

## 4350 20.29 Proxy-Require

4351 The **Proxy-Require** header field is used to indicate proxy-sensitive features that must be supported by the  
4352 proxy. See Section 20.32 for more details on the mechanics of this message and a usage example.

4353 Example:

```
4354 Proxy-Require: foo
```

## 4355 20.30 Record-Route

4356 The **Record-Route** header field is inserted by proxies in a request to force future requests in the dialog to  
4357 be routed through the proxy.

4358 Examples of its use with the **Route** header field are described in Sections 16.12.1.

4359 Example:

```
4360 Record-Route: <sip:server10.biloxi.com;lr>, <sip:bigbox3.site3.atlanta.com;lr>
```

## 4361 20.31 Reply-To

4362 The **Reply-To** header field contains a logical return URI that may be different from the **From** header field.  
4363 For example, the URI **MAY** be used to return missed calls or unestablished sessions. If the user wished to  
4364 remain anonymous, the header field **SHOULD** either be omitted from the request or populated in such a way  
4365 that does not reveal any private information.

4366 Even if the “display-name” is empty, the “name-addr” form **MUST** be used if the “addr-spec” con-  
4367 tains a comma, question mark, or semicolon. Syntax issues are discussed in Section 7.3.1.

4368 Example:

```
4369 Reply-To: Bob <sip:bob@biloxi.com>
```

## 4370 20.32 Require

4371 The **Require** header field is used by UACs to tell UASs about options that the UAC expects the UAS to  
4372 support in order to process the request. Although an optional header field, the **Require** **MUST NOT** be  
4373 ignored if it is present.

4374 The **Require** header field contains a list of option tags, described in Section 19.2. Each option tag  
4375 defines a SIP extension that **MUST** be understood to process the request. Frequently, this is used to indicate

4376 that a specific set of extension header fields need to be understood. A UAC compliant to this specification  
4377 MUST only include option tags corresponding to standards-track RFCs.

4378 Example:

4379 Require: 100rel

### 4380 20.33 Retry-After

4381 The Retry-After header field can be used with a 503 (Service Unavailable) response to indicate how long  
4382 the service is expected to be unavailable to the requesting client and with a 404 (Not Found), 413 (Request  
4383 Entity Too Large), 480 (Temporarily Unavailable), 486 (Busy Here), 600 (Busy), or 603 (Decline) response  
4384 to indicate when the called party anticipates being available again. The value of this field is a positive integer  
4385 number of seconds (in decimal) after the time of the response.

4386 An optional comment can be used to indicate additional information about the time of callback. An  
4387 optional "duration" parameter indicates how long the called party will be reachable starting at the initial  
4388 time of availability. If no duration parameter is given, the service is assumed to be available indefinitely.

4389 Examples:

4390 Retry-After: 18000;duration=3600

4391 Retry-After: 120 (I'm in a meeting)

### 4392 20.34 Route

4393 The Route header field is used to force routing for a request through the listed set of proxies. Examples of  
4394 the use of the Record-Route header field are in Section 16.12.1.

4395 Example:

4396 Route: <sip:bigbox3.site3.atlanta.com;lr>, <sip:server10.biloxi.com;lr>

### 4397 20.35 Server

4398 The Server header field contains information about the software used by the UAS to handle the request.  
4399 The syntax for this field is defined in [H14.38].

4400 Revealing the specific software version of the server might allow the server to become more vulnerable  
4401 to attacks against software that is known to contain security holes. Implementers SHOULD make the Server  
4402 header field a configurable option.

4403 Example:

4404 Server: HomeProxy v2

### 4405 20.36 Subject

4406 The Subject header field provides a summary or indicates the nature of the call, allowing call filtering  
4407 without having to parse the session description. The session description does not have to use the same  
4408 subject indication as the invitation.

4409 The compact form of the Subject header field is s.

4410 Example:

4411 Subject: Need more boxes  
4412 s: Tech Support

### 4413 **20.37 Supported**

4414 The **Supported** header field enumerates all the extensions supported by the UAC or UAS.

4415 The **Supported** header field contains a list of option tags, described in Section 19.2, that are understood  
4416 by the UAC or UAS. A UA compliant to this specification **MUST** only include option tags corresponding to  
4417 standards-track RFCs. If empty, it means that no extensions are supported.

4418 Example:

4419 Supported: 100rel

### 4420 **20.38 Timestamp**

4421 The **Timestamp** header field describes when the UAC sent the request to the UAS.

4422 See Section 8.2.6 for details on how to generate a response to a request that contains the header field.  
4423 Although there is no normative behavior defined here that makes use of the header, it allows for extensions  
4424 or SIP applications to obtain RTT estimates.

4425 Example:

4426 Timestamp: 54

### 4427 **20.39 To**

4428 The **To** header field specifies the logical recipient of the request.

4429 The optional “display-name” is meant to be rendered by a human-user interface. The “tag” parameter  
4430 serves as a general mechanism for dialog identification.

4431 See Section 13 for details of the “tag” parameter.

4432 Section 12 describes how **To** and **From** header fields are compared for the purpose of matching requests  
4433 to dialogs. See Section 20.10 for the rules for parsing a display name, URI and URI parameters, and header  
4434 parameters.

4435 The compact form of the **To** header field is t.

4436 The following are examples of valid **To** header fields:

4437 To: The Operator <sip:operator@cs.columbia.edu>;tag=287447  
4438 t: sip:+12125551212@server.phone2net.com

### 4439 **20.40 Unsupported**

4440 The **Unsupported** header field lists the features not supported by the UAS. See Section 20.32 for motivation.

4441 Example:

4442 Unsupported: foo

## 4443 20.41 User-Agent

4444 The **User-Agent** header field contains information about the UAC originating the request. The syntax and  
4445 semantics are defined in [H14.43].

4446 Revealing the specific software version of the user agent might allow the user agent to become more  
4447 vulnerable to attacks against software that is known to contain security holes. Implementers SHOULD make  
4448 the **User-Agent** header field a configurable option.

4449 Example:

```
4450 User-Agent: Softphone Beta1.5
```

## 4451 20.42 Via

4452 The **Via** header field indicates the path taken by the request so far and indicates the path that should be  
4453 followed in routing responses. The branch ID parameter in the **Via** header field values serves as a transaction  
4454 identifier, and is used by proxies to detect loops.

4455 A **Via** header field value contains the transport protocol used to send the message, the client's host name  
4456 or network address, and possibly the port number at which it wishes to receive responses. A **Via** header field  
4457 value can also contain parameters such as "maddr", "ttl", "received", and "branch", whose meaning and  
4458 use are described in other sections.

4459 Transport protocols defined here are "UDP", "TCP", "TLS", and "SCTP". "TLS" means TLS over  
4460 TCP. When a request is sent to a SIPS URI, the protocol still indicates "SIP", and the transport protocol is  
4461 TLS.

```
4462 Via: SIP/2.0/UDP erlang.bell-telephone.com:5060;branch=z9hG4bK87asdk7  
4463 Via: SIP/2.0/UDP 128.59.16.1:5060 ;received=128.59.19.3;branch=z9hG4bK77asjd
```

4464 The compact form of the **Via** header field is **v**.

4465 In this example, the message originated from a multi-homed host with two addresses, 128.59.16.1  
4466 and 128.59.19.3. The sender guessed wrong as to which network interface would be used. Erlang.bell-  
4467 telephone.com noticed the mismatch and added a parameter to the previous hop's **Via** header field value,  
4468 containing the address that the packet actually came from.

4469 The host or network address and port number are not required to follow the SIP URI syntax. Specifically,  
4470 LWS on either side of the ":" or "/" is allowed, as shown here:

```
4471 Via: SIP / 2.0 / UDP first.example.com: 4000;ttl=16  
4472 ;maddr=224.2.0.1 ;branch=z9hG4bKa7c6a8dlze.1
```

4473 Even though this specification mandates that the branch parameter be present in all requests, the BNF  
4474 for the header field indicates that it is optional. This allows interoperability with RFC 2543 elements, which  
4475 did not have to insert the branch parameter.

## 4476 20.43 Warning

4477 The **Warning** header field is used to carry additional information about the status of a response. **Warning**  
4478 header field values are sent with responses and contain a three-digit warning code, host name, and warning

4479 text.

4480 The “warn-text” should be in a natural language that is most likely to be intelligible to the human user  
4481 receiving the response. This decision can be based on any available knowledge, such as the location of the  
4482 user, the Accept-Language field in a request, or the Content-Language field in a response. The default  
4483 language is i-default [21].

4484 The currently-defined “warn-code”s are listed below, with a recommended warn-text in English and a  
4485 description of their meaning. These warnings describe failures induced by the session description. The first  
4486 digit of warning codes beginning with “3” indicates warnings specific to SIP. Warnings 300 through 329 are  
4487 reserved for indicating problems with keywords in the session description, 330 through 339 are warnings  
4488 related to basic network services requested in the session description, 370 through 379 are warnings related  
4489 to quantitative QoS parameters requested in the session description, and 390 through 399 are miscellaneous  
4490 warnings that do not fall into one of the above categories.

4491 **300 Incompatible network protocol:** One or more network protocols contained in the session description  
4492 are not available.

4493 **301 Incompatible network address formats:** One or more network address formats contained in the ses-  
4494 sion description are not available.

4495 **302 Incompatible transport protocol:** One or more transport protocols described in the session descrip-  
4496 tion are not available.

4497 **303 Incompatible bandwidth units:** One or more bandwidth measurement units contained in the session  
4498 description were not understood.

4499 **304 Media type not available:** One or more media types contained in the session description are not avail-  
4500 able.

4501 **305 Incompatible media format:** One or more media formats contained in the session description are not  
4502 available.

4503 **306 Attribute not understood:** One or more of the media attributes in the session description are not sup-  
4504 ported.

4505 **307 Session description parameter not understood:** A parameter other than those listed above was not  
4506 understood.

4507 **330 Multicast not available:** The site where the user is located does not support multicast.

4508 **331 Unicast not available:** The site where the user is located does not support unicast communication (usu-  
4509 ally due to the presence of a firewall).

4510 **370 Insufficient bandwidth:** The bandwidth specified in the session description or defined by the media  
4511 exceeds that known to be available.

4512 **399 Miscellaneous warning:** The warning text can include arbitrary information to be presented to a hu-  
4513 man user or logged. A system receiving this warning MUST NOT take any automated action.

4514 1xx and 2xx have been taken by HTTP/1.1.

4515 Additional "warn-code"s can be defined through IANA, as defined in Section 27.2.

4516 Examples:

4517 Warning: 307 isi.edu "Session parameter 'foo' not understood"

4518 Warning: 301 isi.edu "Incompatible network address type 'E.164'"

## 4519 20.44 WWW-Authenticate

4520 A WWW-Authenticate header field value contains an authentication challenge. The syntax for this header  
4521 field and use is defined in [H14.47]. See 22.2 for further details on its usage.

4522 Example:

```
4523 WWW-Authenticate: Digest realm="atlanta.com",  
4524 domain="sip:boxesbybob.com",  
4525 nonce="f84f1cec41e6cbe5aea9c8e88d359",  
4526 opaque="", stale=FALSE, algorithm=MD5
```

## 4527 21 Response Codes

4528 The response codes are consistent with, and extend, HTTP/1.1 response codes. Not all HTTP/1.1 response  
4529 codes are appropriate, and only those that are appropriate are given here. Other HTTP/1.1 response codes  
4530 SHOULD NOT be used. Also, SIP defines a new class, 6xx.

### 4531 21.1 Provisional 1xx

4532 Provisional responses, also known as informational responses, indicate that the server contacted is perform-  
4533 ing some further action and does not yet have a definitive response. A server sends a 1xx response if it  
4534 expects to take more than 200 ms to obtain a final response. Note that 1xx responses are not transmitted  
4535 reliably. They never cause the client to send an ACK. Provisional (1xx) responses MAY contain message  
4536 bodies, including session descriptions.

#### 4537 21.1.1 100 Trying

4538 This response indicates that the request has been received by the next-hop server and that some unspecified  
4539 action is being taken on behalf of this call (for example, a database is being consulted). This response, like  
4540 all other provisional responses, stops retransmissions of an INVITE by a UAC. The 100 (Trying) response  
4541 is different from other provisional responses, in that it is never forwarded upstream by a stateful proxy.

#### 4542 21.1.2 180 Ringing

4543 The UA receiving the INVITE is trying to alert the user. This response MAY be used to initiate local ringback.

#### 4544 21.1.3 181 Call Is Being Forwarded

4545 A server MAY use this status code to indicate that the call is being forwarded to a different set of destinations.

#### 4546 **21.1.4 182 Queued**

4547 The called party is temporarily unavailable, but the server has decided to queue the call rather than reject it.  
4548 When the callee becomes available, it will return the appropriate final status response. The reason phrase  
4549 MAY give further details about the status of the call, for example, "5 calls queued; expected waiting time is  
4550 15 minutes". The server MAY issue several 182 (Queued) responses to update the caller about the status of  
4551 the queued call.

#### 4552 **21.1.5 183 Session Progress**

4553 The 183 (Session Progress) response is used to convey information about the progress of the call that is not  
4554 otherwise classified. The Reason-Phrase, header fields, or message body MAY be used to convey more  
4555 details about the call progress.

### 4556 **21.2 Successful 2xx**

4557 The request was successful.

#### 4558 **21.2.1 200 OK**

4559 The request has succeeded. The information returned with the response depends on the method used in the  
4560 request.

### 4561 **21.3 Redirection 3xx**

4562 3xx responses give information about the user's new location, or about alternative services that might be  
4563 able to satisfy the call.

#### 4564 **21.3.1 300 Multiple Choices**

4565 The address in the request resolved to several choices, each with its own specific location, and the user (or  
4566 UA) can select a preferred communication end point and redirect its request to that location.

4567 The response MAY include a message body containing a list of resource characteristics and location(s)  
4568 from which the user or UA can choose the one most appropriate, if allowed by the Accept request header  
4569 field. However, no MIME types have been defined for this message body.

4570 The choices SHOULD also be listed as Contact fields (Section 20.10). Unlike HTTP, the SIP response  
4571 MAY contain several Contact fields or a list of addresses in a Contact field. UAs MAY use the Contact  
4572 header field value for automatic redirection or MAY ask the user to confirm a choice. However, this specifi-  
4573 cation does not define any standard for such automatic selection.

4574 This status response is appropriate if the callee can be reached at several different locations and the server cannot  
4575 or prefers not to proxy the request.

#### 4576 **21.3.2 301 Moved Permanently**

4577 The user can no longer be found at the address in the Request-URI, and the requesting client SHOULD retry  
4578 at the new address given by the Contact header field (Section 20.10). The requestor SHOULD update any

4579 local directories, address books, and user location caches with this new value and redirect future requests to  
4580 the address(es) listed.

### 4581 **21.3.3 302 Moved Temporarily**

4582 The requesting client SHOULD retry the request at the new address(es) given by the **Contact** header field  
4583 (Section 20.10). The **Request-URI** of the new request uses the value of the **Contact** header field in the  
4584 response.

4585 The duration of the validity of the **Contact** URI can be indicated through an **Expires** (Section 20.19)  
4586 header field or an **expires** parameter in the **Contact** header field. Both proxies and UAs MAY cache this  
4587 URI for the duration of the expiration time. If there is no explicit expiration time, the address is only valid  
4588 once for recursing, and MUST NOT be cached for future transactions.

4589 If the URI cached from the **Contact** header field fails, the **Request-URI** from the redirected request  
4590 MAY be tried again a single time.

4591 The temporary URI may have become out-of-date sooner than the expiration time, and a new temporary URI  
4592 may be available.

### 4593 **21.3.4 305 Use Proxy**

4594 The requested resource MUST be accessed through the proxy given by the **Contact** field. The **Contact** field  
4595 gives the URI of the proxy. The recipient is expected to repeat this single request via the proxy. 305 (Use  
4596 Proxy) responses MUST only be generated by UASs.

### 4597 **21.3.5 380 Alternative Service**

4598 The call was not successful, but alternative services are possible. The alternative services are described in  
4599 the message body of the response. Formats for such bodies are not defined here, and may be the subject of  
4600 future standardization.

## 4601 **21.4 Request Failure 4xx**

4602 4xx responses are definite failure responses from a particular server. The client SHOULD NOT retry the same  
4603 request without modification (for example, adding appropriate authorization). However, the same request to  
4604 a different server might be successful.

### 4605 **21.4.1 400 Bad Request**

4606 The request could not be understood due to malformed syntax. The **Reason-Phrase** SHOULD identify the  
4607 syntax problem in more detail, for example, "Missing Call-ID header field".

### 4608 **21.4.2 401 Unauthorized**

4609 The request requires user authentication. This response is issued by UASs and registrars, while 407 (Proxy  
4610 Authentication Required) is used by proxy servers.

4611 **21.4.3 402 Payment Required**

4612 Reserved for future use.

4613 **21.4.4 403 Forbidden**

4614 The server understood the request, but is refusing to fulfill it. Authorization will not help, and the request  
4615 SHOULD NOT be repeated.

4616 **21.4.5 404 Not Found**

4617 The server has definitive information that the user does not exist at the domain specified in the Request-  
4618 URI. This status is also returned if the domain in the Request-URI does not match any of the domains  
4619 handled by the recipient of the request.

4620 **21.4.6 405 Method Not Allowed**

4621 The method specified in the Request-Line is understood, but not allowed for the address identified by the  
4622 Request-URI.

4623 The response MUST include an Allow header field containing a list of valid methods for the indicated  
4624 address.

4625 **21.4.7 406 Not Acceptable**

4626 The resource identified by the request is only capable of generating response entities that have content  
4627 characteristics not acceptable according to the Accept header fields sent in the request.

4628 **21.4.8 407 Proxy Authentication Required**

4629 This code is similar to 401 (Unauthorized), but indicates that the client MUST first authenticate itself with  
4630 the proxy. SIP access authentication is explained in Sections 26 and 22.3.

4631 This status code can be used for applications where access to the communication channel (for example,  
4632 a telephony gateway) rather than the callee requires authentication.

4633 **21.4.9 408 Request Timeout**

4634 The server could not produce a response within a suitable amount of time, for example, if it could not  
4635 determine the location of the user in time. The client MAY repeat the request without modifications at any  
4636 later time.

4637 **21.4.10 410 Gone**

4638 The requested resource is no longer available at the server and no forwarding address is known. This  
4639 condition is expected to be considered permanent. If the server does not know, or has no facility to determine,  
4640 whether or not the condition is permanent, the status code 404 (Not Found) SHOULD be used instead.

4641 **21.4.11 413 Request Entity Too Large**

4642 The server is refusing to process a request because the request entity-body is larger than the server is willing  
4643 or able to process. The server MAY close the connection to prevent the client from continuing the request.

4644 If the condition is temporary, the server SHOULD include a **Retry-After** header field to indicate that it is  
4645 temporary and after what time the client MAY try again.

4646 **21.4.12 414 Request-URI Too Long**

4647 The server is refusing to service the request because the **Request-URI** is longer than the server is willing to  
4648 interpret.

4649 **21.4.13 415 Unsupported Media Type**

4650 The server is refusing to service the request because the message body of the request is in a format not sup-  
4651 ported by the server for the requested method. The server SHOULD return a list of acceptable formats using  
4652 the **Accept**, **Accept-Encoding** and **Accept-Language** header fields. UAC processing of this response is  
4653 described in Section 8.1.3.5.

4654 **21.4.14 416 Unsupported URI Scheme**

4655 The server cannot process the request because the scheme of the URI in the **Request-URI** is unknown to  
4656 the server. Client processing of this response is described in Section 8.1.3.5.

4657 **21.4.15 420 Bad Extension**

4658 The server did not understand the protocol extension specified in a **Proxy-Require** (Section 20.29) or **Re-**  
4659 **quire** (Section 20.32) header field. The server SHOULD include a list of the unsupported extensions in an  
4660 **Unsupported** header field in the response. UAC processing of this response is described in Section 8.1.3.5.

4661 **21.4.16 421 Extension Required**

4662 The UAS needs a particular extension to process the request, but this extension is not listed in a **Supported**  
4663 header field in the request. Responses with this status code MUST contain a **Require** header field listing the  
4664 required extensions.

4665 A UAS SHOULD NOT use this response unless it truly cannot provide any useful service to the client.  
4666 Instead, if a desirable extension is not listed in the **Supported** header field, servers SHOULD process the  
4667 request using baseline SIP capabilities and any extensions supported by the client.

4668 **21.4.17 423 Interval Too Brief**

4669 The server is rejecting the request because the expiration time of the resource refreshed by the request is too  
4670 short. This response can be used by a registrar to reject a registration whose **Contact** header field expiration  
4671 time was too small. The use of this response and the related **Min-Expires** header field are described in  
4672 Sections 10.2.8, 10.3, and 20.23.

4673 **21.4.18 480 Temporarily Unavailable**

4674 The callee's end system was contacted successfully but the callee is currently unavailable (for example, is  
4675 not logged in, logged in but in a state that precludes communication with the callee, or has activated the "do  
4676 not disturb" feature). The response MAY indicate a better time to call in the **Retry-After** header field. The  
4677 user could also be available elsewhere (unbeknownst to this server). The reason phrase SHOULD indicate a  
4678 more precise cause as to why the callee is unavailable. This value SHOULD be settable by the UA. Status  
4679 486 (Busy Here) MAY be used to more precisely indicate a particular reason for the call failure.

4680 This status is also returned by a redirect or proxy server that recognizes the user identified by the  
4681 **Request-URI**, but does not currently have a valid forwarding location for that user.

4682 **21.4.19 481 Call/Transaction Does Not Exist**

4683 This status indicates that the UAS received a request that does not match any existing dialog or transaction.

4684 **21.4.20 482 Loop Detected**

4685 The server has detected a loop (Section 16.3 Item 4).

4686 **21.4.21 483 Too Many Hops**

4687 The server received a request that contains a **Max-Forwards** (Section 20.22) header field with the value  
4688 zero.

4689 **21.4.22 484 Address Incomplete**

4690 The server received a request with a **Request-URI** that was incomplete. Additional information SHOULD  
4691 be provided in the reason phrase.

4692 This status code allows overlapped dialing. With overlapped dialing, the client does not know the length of the  
4693 dialing string. It sends strings of increasing lengths, prompting the user for more input, until it no longer receives a  
4694 484 (Address Incomplete) status response.

4695 **21.4.23 485 Ambiguous**

4696 The **Request-URI** was ambiguous. The response MAY contain a listing of possible unambiguous addresses  
4697 in **Contact** header fields. Revealing alternatives can infringe on privacy of the user or the organization. It  
4698 MUST be possible to configure a server to respond with status 404 (Not Found) or to suppress the listing of  
4699 possible choices for ambiguous **Request-URIs**.

4700 Example response to a request with the **Request-URI** `sip:lee@example.com`:

```
4701 SIP/2.0 485 Ambiguous
4702 Contact: Carol Lee <sip:carol.lee@example.com>
4703 Contact: Ping Lee <sip:p.lee@example.com>
4704 Contact: Lee M. Foote <sips:lee.foote@example.com>
```

4705 Some email and voice mail systems provide this functionality. A status code separate from 3xx is used since  
4706 the semantics are different: for 300, it is assumed that the same person or service will be reached by the choices

4707 provided. While an automated choice or sequential search makes sense for a 3xx response, user intervention is  
4708 required for a 485 (Ambiguous) response.

#### 4709 **21.4.24 486 Busy Here**

4710 The callee's end system was contacted successfully, but the callee is currently not willing or able to take  
4711 additional calls at this end system. The response MAY indicate a better time to call in the **Retry-After** header  
4712 field. The user could also be available elsewhere, such as through a voice mail service. Status 600 (Busy  
4713 Everywhere) SHOULD be used if the client knows that no other end system will be able to accept this call.

#### 4714 **21.4.25 487 Request Terminated**

4715 The request was terminated by a **BYE** or **CANCEL** request. This response is never returned for a **CANCEL**  
4716 request itself.

#### 4717 **21.4.26 488 Not Acceptable Here**

4718 The response has the same meaning as 606 (Not Acceptable), but only applies to the specific resource  
4719 addressed by the **Request-URI** and the request may succeed elsewhere.

4720 A message body containing a description of media capabilities MAY be present in the response, which is  
4721 formatted according to the **Accept** header field in the **INVITE** (or **application/sdp** if not present), the same  
4722 as a message body in a 200 (OK) response to an **OPTIONS** request.

#### 4723 **21.4.27 491 Request Pending**

4724 The request was received by a UAS that had a pending request within the same dialog. Section 14.2 describes  
4725 how such "glare" situations are resolved.

#### 4726 **21.4.28 493 Undecipherable**

4727 The request was received by a UAS that contained an encrypted MIME body for which the recipient does not  
4728 possess or will not provide an appropriate decryption key. This response MAY have a single body containing  
4729 an appropriate public key that should be used to encrypt MIME bodies sent to this UA. Details of the usage  
4730 of this response code can be found in Section 23.2.

### 4731 **21.5 Server Failure 5xx**

4732 5xx responses are failure responses given when a server itself has erred.

#### 4733 **21.5.1 500 Server Internal Error**

4734 The server encountered an unexpected condition that prevented it from fulfilling the request. The client MAY  
4735 display the specific error condition and MAY retry the request after several seconds.

4736 If the condition is temporary, the server MAY indicate when the client may retry the request using the  
4737 **Retry-After** header field.

### 4738 **21.5.2 501 Not Implemented**

4739 The server does not support the functionality required to fulfill the request. This is the appropriate response  
4740 when a UAS does not recognize the request method and is not capable of supporting it for any user. (Proxies  
4741 forward all requests regardless of method.)

4742 Note that a 405 (Method Not Allowed) is sent when the server recognizes the request method, but that  
4743 method is not allowed or supported.

### 4744 **21.5.3 502 Bad Gateway**

4745 The server, while acting as a gateway or proxy, received an invalid response from the downstream server it  
4746 accessed in attempting to fulfill the request.

### 4747 **21.5.4 503 Service Unavailable**

4748 The server is temporarily unable to process the request due to a temporary overloading or maintenance of  
4749 the server. The server MAY indicate when the client should retry the request in a **Retry-After** header field.  
4750 If no **Retry-After** is given, the client MUST act as if it had received a 500 (Server Internal Error) response.

4751 A client (proxy or UAC) receiving a 503 (Service Unavailable) SHOULD attempt to forward the request  
4752 to an alternate server. It SHOULD NOT forward any other requests to that server for the duration specified in  
4753 the **Retry-After** header field, if present.

4754 Servers MAY refuse the connection or drop the request instead of responding with 503 (Service Unavail-  
4755 able).

### 4756 **21.5.5 504 Server Time-out**

4757 The server did not receive a timely response from an external server it accessed in attempting to process the  
4758 request. 408 (Request Timeout) should be used instead if there was no response within the period specified  
4759 in the **Expires** header field from the upstream server.

### 4760 **21.5.6 505 Version Not Supported**

4761 The server does not support, or refuses to support, the SIP protocol version that was used in the request. The  
4762 server is indicating that it is unable or unwilling to complete the request using the same major version as the  
4763 client, other than with this error message.

### 4764 **21.5.7 513 Message Too Large**

4765 The server was unable to process the request since the message length exceeded its capabilities.

## 4766 **21.6 Global Failures 6xx**

4767 6xx responses indicate that a server has definitive information about a particular user, not just the particular  
4768 instance indicated in the **Request-URI**.

### 4769 **21.6.1 600 Busy Everywhere**

4770 The callee's end system was contacted successfully but the callee is busy and does not wish to take the call  
4771 at this time. The response MAY indicate a better time to call in the **Retry-After** header field. If the callee  
4772 does not wish to reveal the reason for declining the call, the callee uses status code 603 (Decline) instead.  
4773 This status response is returned only if the client knows that no other end point (such as a voice mail system)  
4774 will answer the request. Otherwise, 486 (Busy Here) should be returned.

### 4775 **21.6.2 603 Decline**

4776 The callee's machine was successfully contacted but the user explicitly does not wish to or cannot partici-  
4777 pate. The response MAY indicate a better time to call in the **Retry-After** header field. This status response  
4778 is returned only if the client knows that no other end point will answer the request.

### 4779 **21.6.3 604 Does Not Exist Anywhere**

4780 The server has authoritative information that the user indicated in the **Request-URI** does not exist anywhere.

### 4781 **21.6.4 606 Not Acceptable**

4782 The user's agent was contacted successfully but some aspects of the session description such as the requested  
4783 media, bandwidth, or addressing style were not acceptable.

4784 A 606 (Not Acceptable) response means that the user wishes to communicate, but cannot adequately  
4785 support the session described. The 606 (Not Acceptable) response MAY contain a list of reasons in a **Warn-**  
4786 **ing** header field describing why the session described cannot be supported. Warning reason codes are listed  
4787 in Section 20.43.

4788 A message body containing a description of media capabilities MAY be present in the response, which is  
4789 formatted according to the **Accept** header field in the **INVITE** (or **application/sdp** if not present), the same  
4790 as a message body in a 200 (OK) response to an **OPTIONS** request.

4791 It is hoped that negotiation will not frequently be needed, and when a new user is being invited to join  
4792 an already existing conference, negotiation may not be possible. It is up to the invitation initiator to decide  
4793 whether or not to act on a 606 (Not Acceptable) response.

4794 This status response is returned only if the client knows that no other end point will answer the request.

## 4795 **22 Usage of HTTP Authentication**

4796 SIP provides a stateless, challenge-based mechanism for authentication that is based on authentication in  
4797 HTTP. Any time that a proxy server or UA receives a request (with the exceptions given in Section 22.1), it  
4798 MAY challenge the initiator of the request to provide assurance of its identity. Once the originator has been  
4799 identified, the recipient of the request SHOULD ascertain whether or not this user is authorized to make the  
4800 request in question. No authorization systems are recommended or discussed in this document.

4801 The "Digest" authentication mechanism described in this section provides message authentication and  
4802 replay protection only, without message integrity or confidentiality. Protective measures above and beyond  
4803 those provided by Digest need to be taken to prevent active attackers from modifying SIP requests and  
4804 responses.

4805 Note that due to its weak security, the usage of “Basic” authentication has been deprecated. Servers  
4806 MUST NOT accept credentials using the “Basic” authorization scheme, and servers also MUST NOT challenge  
4807 with “Basic”. This is a change from RFC 2543.

## 4808 22.1 Framework

4809 The framework for SIP authentication closely parallels that of HTTP (RFC 2617 [18]). In particular, the  
4810 BNF for auth-scheme, auth-param, challenge, realm, realm-value, and credentials is identical (al-  
4811 though the usage of “Basic” as a scheme is not permitted). In SIP, a UAS uses the 401 (Unauthorized)  
4812 response to challenge the identity of a UAC. Additionally, registrars and redirect servers MAY make use  
4813 of 401 (Unauthorized) responses for authentication, but proxies MUST NOT, and instead MAY use the 407  
4814 (Proxy Authentication Required) response. The requirements for inclusion of the Proxy-Authenticate,  
4815 Proxy-Authorization, WWW-Authenticate, and Authorization in the various messages are identical to  
4816 those described in RFC 2617 [18].

4817 Since SIP does not have the concept of a canonical root URL, the notion of protection spaces is in-  
4818 terpreted differently in SIP. The realm string alone defines the protection domain. This is a change from  
4819 RFC 2543, in which the Request-URI and the realm together defined the protection domain.

4820 This previous definition of protection domain caused some amount of confusion since the Request-URI sent by  
4821 the UAC and the Request-URI received by the challenging server might be different, and indeed the final form of  
4822 the Request-URI might not be known to the UAC. Also, the previous definition depended on the presence of a SIP  
4823 URI in the Request-URI and seemed to rule out alternative URI schemes (for example, the tel URL).

4824 Operators of user agents or proxy servers that will authenticate received requests MUST adhere to the  
4825 following guidelines for creation of a realm string for their server:

- 4826 • Realm strings MUST be globally unique. It is RECOMMENDED that a realm string contain a hostname  
4827 or domain name, following the recommendation in Section 3.2.1 of RFC 2617 [18].
- 4828 • Realm strings SHOULD present a human-readable identifier that can be rendered to a user.

4829 For example:

```
4830 INVITE sip:bob@biloxi.com SIP/2.0  
4831 Authorization: Digest realm="biloxi.com", <...>
```

4832 Generally, SIP authentication is meaningful for a specific realm, a protection domain. Thus, for Digest  
4833 authentication, each such protection domain has its own set of usernames and passwords. If a server does  
4834 not require authentication for a particular request, it MAY accept a default username, “anonymous”, which  
4835 has no password (password of “”). Similarly, UACs representing many users, such as PSTN gateways, MAY  
4836 have their own device-specific username and password, rather than accounts for particular users, for their  
4837 realm.

4838 While a server can legitimately challenge most SIP requests, there are two requests defined by this  
4839 document that require special handling for authentication: ACK and CANCEL.

4840 Under an authentication scheme that uses responses to carry values used to compute nonces (such as  
4841 Digest), some problems come up for any requests that take no response, including ACK. For this reason,  
4842 any credentials in the INVITE that were accepted by a server MUST be accepted by that server for the ACK.  
4843 UACs creating an ACK message will duplicate all of the Authorization and Proxy-Authorization header

4844 field values that appeared in the INVITE to which the ACK corresponds. Servers MUST NOT attempt to  
4845 challenge an ACK.

4846 Although the CANCEL method does take a response (a 2xx), servers MUST NOT attempt to challenge  
4847 CANCEL requests since these requests cannot be resubmitted. Generally, a CANCEL request SHOULD be  
4848 accepted by a server if it comes from the same hop that sent the request being canceled (provided that some  
4849 sort of transport or network layer security association, as described in Section 26.2.1, is in place).

4850 When a UAC receives a challenge, it SHOULD render to the user the contents of the "realm" param-  
4851 eter in the challenge (which appears in either a WWW-Authenticate header field or Proxy-Authenticate  
4852 header field) if the UAC device does not already know of a credential for the realm in question. A service  
4853 provider that pre-configures UAs with credentials for its realm should be aware that users will not have the  
4854 opportunity to present their own credentials for this realm when challenged at a pre-configured device.

4855 Finally, note that even if a UAC can locate credentials that are associated with the proper realm, the  
4856 potential exists that these credentials may no longer be valid or that the challenging server will not accept  
4857 these credentials for whatever reason (especially when "anonymous" with no password is submitted). In  
4858 this instance a server may repeat its challenge, or it may respond with a 403 Forbidden. A UAC MUST NOT  
4859 re-attempt requests with the credentials that have just been rejected (though the request may be are retried if  
4860 the nonce was stale).

## 4861 22.2 User-to-User Authentication

4862 When a UAS receives a request from a UAC, the UAS MAY authenticate the originator before the request  
4863 is processed. If no credentials (in the Authorization header field) are provided in the request, the UAS  
4864 can challenge the originator to provide credentials by rejecting the request with a 401 (Unauthorized) status  
4865 code.

4866 The WWW-Authenticate response-header field MUST be included in 401 (Unauthorized) response mes-  
4867 sages. The field value consists of at least one challenge that indicates the authentication scheme(s) and  
4868 parameters applicable to the realm. See [H14.47] for a definition of the syntax.

4869 An example of the WWW-Authenticate header field in a 401 challenge is:

```
4870 WWW-Authenticate: Digest  
4871 realm="biloxi.com",  
4872 qop="auth,auth-int",  
4873 nonce="dcd98b7102dd2f0e8b11d0f600bfb0c093",  
4874 opaque="5ccc069c403ebaf9f0171e9517f40e41"
```

4875 When the originating UAC receives the 401 (Unauthorized), it SHOULD, if it is able, re-originate the  
4876 request with the proper credentials. The UAC may require input from the originating user before proceeding.  
4877 Once authentication credentials have been supplied (either directly by the user, or discovered in an internal  
4878 keyring), UAs SHOULD cache the credentials for a given value of the To header field and "realm" and  
4879 attempt to re-use these values on the next request for that destination. UAs MAY cache credentials in any  
4880 way they would like.

4881 If no credentials for a realm can be located, UACs MAY attempt to retry the request with a username of  
4882 "anonymous" and no password (a password of "").

4883 Once credentials have been located, any UA that wishes to authenticate itself with a UAS or registrar  
4884 – usually, but not necessarily, after receiving a 401 (Unauthorized) response – MAY do so by including an

4885 Authorization header field with the request. The Authorization field value consists of credentials containing  
4886 the authentication information of the UA for the realm of the resource being requested as well as parameters  
4887 required in support of authentication and replay protection.

4888 An example of the Authorization header field is:

```
4889 Authorization: Digest username="bob",  
4890     realm="biloxi.com",  
4891     nonce="dcd98b7102dd2f0e8b11d0f600bfb0c093",  
4892     uri="sip:bob@biloxi.com",  
4893     qop=auth,  
4894     nc=00000001,  
4895     cnonce="0a4f113b",  
4896     response="6629fae49393a05397450978507c4ef1",  
4897     opaque="5ccc069c403ebaf9f0171e9517f40e41"  
4898
```

4899 When a UAC resubmits a request with its credentials after receiving a 401 (Unauthorized) or 407 (Proxy  
4900 Authentication Required) response, it MUST increment the CSeq header field value as it would normally  
4901 when sending an updated request.

### 4902 22.3 Proxy-to-User Authentication

4903 Similarly, when a UAC sends a request to a proxy server, the proxy server MAY authenticate the originator  
4904 before the request is processed. If no credentials (in the Proxy-Authorization header field) are provided  
4905 in the request, the proxy can challenge the originator to provide credentials by rejecting the request with a  
4906 407 (Proxy Authentication Required) status code. The proxy MUST populate the 407 (Proxy Authentication  
4907 Required) message with a Proxy-Authenticate header field value applicable to the proxy for the requested  
4908 resource.

4909 The use of Proxy-Authentication and Proxy-Authorization parallel that described in [18], with one  
4910 difference. Proxies MUST NOT add values to the Proxy-Authorization header field. All 407 (Proxy Au-  
4911 thentication Required) responses MUST be forwarded upstream toward the UAC following the procedures  
4912 for any other response. It is the UAC's responsibility to add the Proxy-Authorization header field value  
4913 containing credentials for the realm of the proxy that has asked for authentication.

4914 If a proxy were to resubmit a request adding a Proxy-Authorization header field value, it would need to  
4915 increment the CSeq in the new request. However, this would cause the UAC that submitted the original request to  
4916 discard a response from the UAS, as the CSeq value would be different.

4917 When the originating UAC receives the 407 (Proxy Authentication Required) it SHOULD, if it is able,  
4918 re-originate the request with the proper credentials. It should follow the same procedures for the display of  
4919 the "realm" parameter that are given above for responding to 401.

4920 If no credentials for a realm can be located, UACs MAY attempt to retry the request with a username of  
4921 "anonymous" and no password (a password of "").

4922 The UAC SHOULD also cache the credentials used in the re-originated request.

4923 The following rule is RECOMMENDED for proxy credential caching:

4924 If a UA receives a Proxy-Authenticate header field value in a 401/407 response to a request with a  
4925 particular Call-ID, it should incorporate credentials for that realm in all subsequent requests that contain the

4926 same Call-ID. These credentials MUST NOT be cached across dialogs; however, if a UA is configured with  
4927 the realm of its local outbound proxy, when one exists, then the UA MAY cache credentials for that realm  
4928 across dialogs. Note that this does mean a future request in a dialog could contain credentials that are not  
4929 needed by any proxy along the Route header path.

4930 Any UA that wishes to authenticate itself to a proxy server – usually, but not necessarily, after receiving  
4931 a 407 (Proxy Authentication Required) response – MAY do so by including a Proxy-Authorization header  
4932 field value with the request. The Proxy-Authorization request-header field allows the client to identify itself  
4933 (or its user) to a proxy that requires authentication. The Proxy-Authorization header field value consists of  
4934 credentials containing the authentication information of the UA for the proxy and/or realm of the resource  
4935 being requested.

4936 A Proxy-Authorization header field value applies only to the proxy whose realm is identified in the  
4937 “realm” parameter (this proxy may previously have demanded authentication using the Proxy-Authenticate  
4938 field). When multiple proxies are used in a chain, a Proxy-Authorization header field value MUST NOT be  
4939 consumed by any proxy whose realm does not match the “realm” parameter specified in that value.

4940 Note that if an authentication scheme that does not support realms is used in the Proxy-Authorization  
4941 header field, a proxy server MUST attempt to parse all Proxy-Authorization header field values to determine  
4942 whether one of them has what the proxy server considers to be valid credentials. Because this is potentially  
4943 very time-consuming in large networks, proxy servers SHOULD use an authentication scheme that supports  
4944 realms in the Proxy-Authorization header field.

4945 If a request is forked (as described in Section 16.7), various proxy servers and/or UAs may wish to  
4946 challenge the UAC. In this case, the forking proxy server is responsible for aggregating these challenges  
4947 into a single response. Each WWW-Authenticate and Proxy-Authenticate value received in responses to  
4948 the forked request MUST be placed into the single response that is sent by the forking proxy to the UA; the  
4949 ordering of these header field values is not significant.

4950 When a proxy server issues a challenge in response to a request, it will not proxy the request until the UAC has  
4951 retried the request with valid credentials. A forking proxy may forward a request simultaneously to multiple proxy  
4952 servers that require authentication, each of which in turn will not forward the request until the originating UAC has  
4953 authenticated itself in their respective realm. If the UAC does not provide credentials for each challenge, then the  
4954 proxy servers that issued the challenges will not forward requests to the UA where the destination user might be  
4955 located, and therefore, the virtues of forking are largely lost.

4956 When resubmitting its request in response to a 401 (Unauthorized) or 407 (Proxy Authentication Re-  
4957 quired) that contains multiple challenges, a UAC MAY include an Authorization value for each WWW-  
4958 Authenticate value and a Proxy-Authorization value for each Proxy-Authenticate value for which the  
4959 UAC wishes to supply a credential. As noted above, multiple credentials in a request SHOULD be differen-  
4960 tiated by the “realm” parameter.

4961 It is possible for multiple challenges associated with the same realm to appear in the same 401 (Unautho-  
4962 rized) or 407 (Proxy Authentication Required). This can occur, for example, when multiple proxies within  
4963 the same administrative domain, which use a common realm, are reached by a forking request. When it re-  
4964 tries a request, a UAC MAY therefore supply multiple credentials in Authorization or Proxy-Authorization  
4965 header fields with the same “realm” parameter value. The same credentials SHOULD be used for the same  
4966 realm.

4967 See [H14.34] for a definition of the syntax of Proxy-Authentication and Proxy-Authorization.

## 4968 22.4 The Digest Authentication Scheme

4969 This section describes the modifications and clarifications required to apply the HTTP Digest authentication  
4970 scheme to SIP. The SIP scheme usage is almost completely identical to that for HTTP [18].

4971 Since RFC 2543 is based on HTTP Digest as defined in RFC 2069 [38], SIP servers supporting RFC  
4972 2617 MUST ensure they are backwards compatible with RFC 2069. Procedures for this backwards com-  
4973 patibility are specified in RFC 2617. Note, however, that SIP servers MUST NOT accept or request Basic  
4974 authentication.

4975 The rules for Digest authentication follow those defined in [18], with “HTTP/1.1” replaced by “SIP/2.0”  
4976 in addition to the following differences:

- 4977 1. The URI included in the challenge has the following BNF:

4978 
$$\text{URI} = \text{SIP-URI} / \text{SIPS-URI}$$

- 4979 2. The BNF in RFC 2617 has an error in that the 'uri' parameter of the Authorization header field for  
4980 HTTP Digest authentication is not enclosed in quotation marks. (The example in Section 3.5 of RFC  
4981 2617 is correct.) For SIP, the 'uri' MUST be enclosed in quotation marks.

- 4982 3. The BNF for digest-uri-value is:

4983 
$$\text{digest-uri-value} = \text{Request-URI} ; \text{ as defined in Section 25}$$

- 4984 4. The example procedure for choosing a nonce based on Etag does not work for SIP.

- 4985 5. The text in RFC 2617 [18] regarding cache operation does not apply to SIP.

- 4986 6. RFC 2617 [18] requires that a server check that the URI in the request line and the URI included in  
4987 the Authorization header field point to the same resource. In a SIP context, these two URIs may refer  
4988 to different users, due to forwarding at some proxy. Therefore, in SIP, a server MAY check that the  
4989 Request-URI in the Authorization header field value corresponds to a user for whom the server is  
4990 willing to accept forwarded or direct requests, but it is not necessarily a failure if the two fields are  
4991 not equivalent.

- 4992 7. As a clarification to the calculation of the A2 value for message integrity assurance in the Digest  
4993 authentication scheme, implementers should assume, when the entity-body is empty (that is, when  
4994 SIP messages have no body) that the hash of the entity-body resolves to the MD5 hash of an empty  
4995 string, or:

4996 
$$H(\text{entity-body}) = \text{MD5}("") = \text{"d41d8cd98f00b204e9800998ecf8427e"}$$

- 4997 8. RFC 2617 notes that a cnonce value MUST NOT be sent in an Authorization (and by extension Proxy-  
4998 Authorization) header field if no qop directive has been sent. Therefore, any algorithms that have a  
4999 dependency on the cnonce (including “MD5-Sess”) require that the qop directive be sent. Use of the  
5000 “qop” parameter is optional in RFC 2617 for the purposes of backwards compatibility with RFC 2069;  
5001 since RFC 2543 was based on RFC 2069, the “qop” parameter must unfortunately remain optional  
5002 for clients and servers to receive. However, servers MUST always send a “qop” parameter in WWW-  
5003 Authenticate and Proxy-Authenticate header field values. If a client receives a “qop” parameter in a  
5004 challenge header field, it MUST send the “qop” parameter in any resulting authorization header field.

5005 RFC 2543 did not allow usage of the Authentication-Info header field (it effectively used RFC 2069).  
5006 However, we now allow usage of this header field, since it provides integrity checks over the bodies and  
5007 provides mutual authentication. RFC 2617 [18] defines mechanisms for backwards compatibility using the  
5008 qop attribute in the request. These mechanisms MUST be used by a server to determine if the client supports  
5009 the new mechanisms in RFC 2617 that were not specified in RFC 2069.

## 5010 **23 S/MIME**

5011 SIP messages carry MIME bodies and the MIME standard includes mechanisms for securing MIME con-  
5012 tents to ensure both integrity and confidentiality (including the 'multipart/signed' and 'application/pkcs7-  
5013 mime' MIME types, see RFC 1847 [22], RFC 2630 [23] and RFC 2633 [24]). Implementers should note,  
5014 however, that there may be rare network intermediaries (not typical proxy servers) that rely on viewing or  
5015 modifying the bodies of SIP messages (especially SDP), and that secure MIME may prevent these sorts of  
5016 intermediaries from functioning.

5017 This applies particularly to certain types of firewalls.

5018 The PGP mechanism for encrypting the header fields and bodies of SIP messages described in RFC 2543 has  
5019 been deprecated.

### 5020 **23.1 S/MIME Certificates**

5021 The certificates that are used to identify an end-user for the purposes of S/MIME differ from those used  
5022 by servers in one important respect - rather than asserting that the identity of the holder corresponds to a  
5023 particular hostname, these certificates assert that the holder is identified by an end-user address. This address  
5024 is composed of the concatenation of the "userinfo" "@" and "domainname" portions of a SIP or SIPS URI  
5025 (in other words, an email address of the form "bob@biloxi.com"), most commonly corresponding to a user's  
5026 address-of-record.

5027 These certificates are also associated with keys that are used to sign or encrypt bodies of SIP messages.

5028 Bodies are signed with the private key of the sender (who may include their public key with the message  
5029 as appropriate), but bodies are encrypted with the public key of the intended recipient. Obviously, senders  
5030 must have foreknowledge of the public key of recipients in order to encrypt message bodies. Public keys  
5031 can be stored within a UA on a virtual keyring.

5032 Each user agent that supports S/MIME MUST contain a keyring specifically for end-users' certificates.  
5033 This keyring should map between addresses of record and corresponding certificates. Over time, users  
5034 SHOULD use the same certificate when they populate the originating URI of signaling (the From header  
5035 field) with the same address-of-record.

5036 Any mechanisms depending on the existence of end-user certificates are seriously limited in that there is  
5037 virtually no consolidated authority today that provides certificates for end-user applications. However, users  
5038 SHOULD acquire certificates from known public certificate authorities. As an alternative, users MAY create  
5039 self-signed certificates. The implications of self-signed certificates are explored further in Section 26.4.2.  
5040 Implementations may also use pre-configured certificates in deployments in which a previous trust relation-  
5041 ship exists between all SIP entities.

5042 Above and beyond the problem of acquiring an end-user certificate, there are few well-known central-  
5043 ized directories that distribute end-user certificates. However, the holder of a certificate SHOULD publish  
5044 their certificate in any public directories as appropriate. Similarly, UACs SHOULD support a mechanism  
5045 for importing (manually or automatically) certificates discovered in public directories corresponding to the  
5046 target URIs of SIP requests.

## 5047 23.2 S/MIME Key Exchange

5048 SIP itself can also be used as a means to distribute public keys in the following manner.

5049 Whenever the CMS SignedData message is used in S/MIME for SIP, it MUST contain the certificate  
5050 bearing the public key necessary to verify the signature.

5051 When a UAC sends a request containing an S/MIME body that initiates a dialog, or sends a non-  
5052 INVITE request outside the context of a dialog, the UAC SHOULD structure the body as an S/MIME 'multi-  
5053 part/signed' CMS SignedData body. If the desired CMS service is EnvelopedData (and the public key of the  
5054 target user is known), the UAC SHOULD send the EnvelopedData message encapsulated within a SignedData  
5055 message.

5056 When a UAS receives a request containing an S/MIME CMS body that includes a certificate, the UAS  
5057 SHOULD first verify the certificate, if possible, with any available certificate authority. The UAS SHOULD  
5058 also determine the subject of the certificate and compare this value to the FROM header field of the request.  
5059 If the certificate cannot be verified, because it is self-signed, or signed by no known authority, or if it is  
5060 verifiable but its subject does not correspond to the FROM header field of request, the UAS MUST notify its  
5061 user of the status of the certificate (including the subject of the certificate, its signer, and any key fingerprint  
5062 information) and request explicit permission before proceeding. If the certificate was successfully verified  
5063 and the subject of the certificate corresponds to the From header field of the SIP request, or if the user (after  
5064 notification) explicitly authorizes the use of the certificate, the UAS SHOULD add this certificate to a local  
5065 keyring, indexed by the address-of-record of the holder of the certificate.

5066 When a UAS sends a response containing an S/MIME body that answers the first request in a dialog, or  
5067 a response to a non-INVITE request outside the context of a dialog, the UAS SHOULD structure the body  
5068 as an S/MIME 'multipart/signed' CMS SignedData body. If the desired CMS service is EnvelopedData, the  
5069 UAS SHOULD send the EnvelopedData message encapsulated within a SignedData message.

5070 When a UAC receives a response containing an S/MIME CMS body that includes a certificate, the UAC  
5071 SHOULD first verify the certificate, if possible, with any available certificate authority. The UAC SHOULD  
5072 also determine the subject of the certificate and compare this value to the To field of the response; although  
5073 the two may very well be different, and this is not necessarily indicative of a security breach. If the certificate  
5074 cannot be verified because it is self-signed, or signed by no known authority, the UAC MUST notify its user  
5075 of the status of the certificate (including the subject of the certificate, its signator, and any key fingerprint  
5076 information) and request explicit permission before proceeding. If the certificate was successfully verified,  
5077 and the subject of the certificate corresponds to the To header field in the response, or if the user (after  
5078 notification) explicitly authorizes the use of the certificate, the UAC SHOULD add this certificate to a local  
5079 keyring, indexed by the address-of-record of the holder of the certificate. If the UAC had not transmitted its  
5080 own certificate to the UAS in any previous transaction, it SHOULD use a CMS SignedData body for its next  
5081 request or response.

5082 On future occasions, when the UA receives requests or responses that contain a From header field  
5083 corresponding to a value in its keyring, the UA SHOULD compare the certificate offered in these messages  
5084 with the existing certificate in its keyring. If there is a discrepancy, the UA MUST notify its user of a change  
5085 of the certificate (preferably in terms that indicate that this is a potential security breach) and acquire the  
5086 user's permission before continuing to process the signaling. If the user authorizes this certificate, it SHOULD  
5087 be added to the keyring alongside any previous value(s) for this address-of-record.

5088 Note well however, that this key exchange mechanism does not guarantee the secure exchange of keys  
5089 when self-signed certificates, or certificates signed by an obscure authority, are used - it is vulnerable to  
5090 well-known attacks. In the opinion of the authors, however, the security it provides is proverbially better

5091 than nothing; it is in fact comparable to the widely used SSH application. These limitations are explored in  
5092 greater detail in Section 26.4.2.

5093 If a UA receives an S/MIME body that has been encrypted with a public key unknown to the recipient,  
5094 it MUST reject the request with a 493 (Undecipherable) response. This response SHOULD contain a valid  
5095 certificate for the respondent (corresponding, if possible, to any address of record given in the To header  
5096 field of the rejected request) within a MIME body with a 'certs-only' "smime-type" parameter.

5097 A 493 (Undecipherable) sent without any certificate indicates that the respondent cannot or will not  
5098 utilize S/MIME encrypted messages, though they may still support S/MIME signatures.

5099 Note that a user agent that receives a request containing an S/MIME body that is not optional (with  
5100 a Content-Disposition header "handling" parameter of "required") MUST reject the request with a 415  
5101 Unsupported Media Type response if the MIME type is not understood. A user agent that receives such a  
5102 response when S/MIME is sent SHOULD notify its user that the remote device does not support S/MIME,  
5103 and it MAY subsequently resend the request without S/MIME, if appropriate; however, this 415 response  
5104 may constitute a downgrade attack.

5105 If a user agent sends an S/MIME body in a request, but receives a response that contains a MIME body  
5106 that is not secured, the UAC SHOULD notify its user that the session could not be secured. However, if a  
5107 user agent that supports S/MIME receives a request with an unsecured body, it SHOULD NOT respond with  
5108 a secured body, but if it expects S/MIME from the sender (for example, because the sender's From header  
5109 field value corresponds to an identity on its keychain), the UAS SHOULD notify its user that the session  
5110 could not be secured.

5111 Finally, if during the course of a dialog a UA receives a certificate in a CMS SignedData message that  
5112 does not correspond with the certificates previously exchanged during a dialog, the UA MUST notify its user  
5113 of the change, preferably in terms that indicate that this is a potential security breach.

### 5114 23.3 Securing MIME bodies

5115 There are two types of secure MIME bodies that are of interest to SIP: 'multipart/signed' and 'application/pkcs7-  
5116 mime'. The procedures for the use of these bodies should follow the S/MIME specification [24] with a few  
5117 variations.

- 5118 • "multipart/signed" MUST be used only with CMS detached signatures.

5119 This allows backwards compatibility with non-S/MIME-compliant recipients.

- 5120 • S/MIME bodies SHOULD have a Content-Disposition header field, and the value of the "handling"  
5121 parameter SHOULD be "required."
- 5122 • If a UAC has no certificate on its keyring associated with the address-of-record to which it wants to  
5123 send a request, it cannot send an encrypted "application/pkcs7-mime" MIME message. UACs MAY  
5124 send an initial request such as an OPTIONS message with a CMS detached signature in order to  
5125 solicit the certificate of the remote side (the signature SHOULD be over a "application/sip" body of the  
5126 type described in Section 23.4).

5127 Note that future standardization work on S/MIME may define non-certificate based keys.

- 5128 • Senders of S/MIME bodies SHOULD use the "SMIMECapabilities" (see Section 2.5.2 of [24]) at-  
5129 tribute to express their capabilities and preferences for further communications. Note especially that

5130 senders MAY use the “preferSignedData” capability to encourage receivers to respond with CMS  
5131 SignedData messages (for example, when sending an OPTIONS request as described above).

- 5132 ● S/MIME implementations MUST at a minimum support SHA1 as a digital signature algorithm, and  
5133 3DES as an encryption algorithm. All other signature and encryption algorithms MAY be supported.  
5134 Implementations can negotiate support for these algorithms with the “SMIMECapabilities” attribute.
- 5135 ● Each S/MIME body in a SIP message SHOULD be signed with only one certificate. If a UA receives  
5136 a message with multiple signatures, the outermost signature should be treated as the single certificate  
5137 for this body. Parallel signatures SHOULD NOT be used.

5138 The following is an example of an encrypted S/MIME SDP body within a SIP message:

```

5139     INVITE sip:bob@biloxi.com SIP/2.0
5140     Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
5141     To: Bob <sip:bob@biloxi.com>
5142     From: Alice <sip:alice@atlanta.com>;tag=1928301774
5143     Call-ID: a84b4c76e66710
5144     CSeq: 314159 INVITE
5145     Max-Forwards: 70
5146     Contact: <sip:alice@pc33.atlanta.com>
5147     Content-Type: application/pkcs7-mime; smime-type=enveloped-data;
5148                 name=smime.p7m
5149     Content-Transfer-Encoding: base64
5150     Content-Disposition: attachment; filename=smime.p7m
5151                 handling=required
5152
5153     *****
5154     * Content-Type: application/sdp *
5155     * *
5156     * v=0 *
5157     * o=alice 53655765 2353687637 IN IP4 pc33.atlanta.com *
5158     * s=- *
5159     * t=0 0 *
5160     * c=IN IP4 pc33.atlanta.com *
5161     * m=audio 3456 RTP/AVP 0 1 3 99 *
5162     * a=rtpmap:0 PCMU/8000 *
5163     *****

```

## 5164 23.4 SIP Header Privacy and Integrity using S/MIME: Tunneling SIP

5165 As a means of providing some degree of end-to-end authentication, integrity or confidentiality for SIP header  
5166 fields, S/MIME can encapsulate entire SIP messages within MIME bodies of type “application/sip” and  
5167 then apply MIME security to these bodies in the same manner as typical SIP bodies. These encapsulated  
5168 SIP requests and responses do not constitute a separate dialog or transaction, they are a copy of the “outer”  
5169 message that is used to verify integrity or to supply additional information.

5170 If a UAS receives a request that contains a tunneled "application/sip" S/MIME body, it SHOULD include  
5171 a tunneled "application/sip" body in the response with the same smime-type.

5172 Any traditional MIME bodies (such as SDP) SHOULD be attached to the "inner" message so that they  
5173 can also benefit from S/MIME security. Note that "application/sip" bodies can be sent as a part of a MIME  
5174 "multipart/mixed" body if any unsecured MIME types should also be transmitted in a request.

### 5175 **23.4.1 Integrity and Confidentiality Properties of SIP Headers**

5176 When the S/MIME integrity or confidentiality mechanisms are used, there may be discrepancies between the  
5177 values in the "inner" message and values in the "outer" message. The rules for handling any such differences  
5178 for all of the header fields described in this document are given in this section.

5179 **23.4.1.1 Integrity** Whenever integrity checks are performed, the integrity of a header field should be  
5180 determined by matching the value of the header field in the signed body with that in the "outer" messages  
5181 using the comparison rules of SIP as described in 20.

5182 Header fields that can be legitimately modified by proxy servers are: Request-URI, Via, Record-  
5183 Route, Route, Max-Forwards, and Proxy-Authorization. If these header fields are not intact end-to-end,  
5184 implementations SHOULD NOT consider this a breach of security. Changes to any other header fields defined  
5185 in this document constitute an integrity violation; users MUST be notified of a discrepancy.

5186 **23.4.1.2 Confidentiality** When messages are encrypted, header fields may be included in the encrypted  
5187 body that are not present in the "outer" message.

5188 Some header fields must always have a plaintext version because they are required header fields in  
5189 requests and responses - these include: To, From, Call-ID, CSeq, Contact. While it is probably not  
5190 useful to provide an encrypted alternative for the Call-ID, Cseq, or Contact, providing an alternative to the  
5191 information in the "outer" To or From is permitted. Note that the values in an encrypted body are not used  
5192 for the purposes of identifying transactions or dialogs - they are merely informational. If the From header  
5193 field in an encrypted body differs from the value in the "outer" message, the value within the encrypted  
5194 body SHOULD be displayed to the user, but MUST NOT be used in the "outer" header fields of any future  
5195 messages.

5196 Primarily, a user agent will want to encrypt header fields that have an end-to-end semantic, including:  
5197 Subject, Reply-To, Organization, Accept, Accept-Encoding, Accept-Language, Alert-Info, Error-  
5198 Info, Authentication-Info, Expires, In-Reply-To, Require, Supported, Unsupported, Retry-After, User-  
5199 Agent, Server, and Warning. If any of these header fields are present in an encrypted body, they should be  
5200 used instead of any "outer" header fields, whether this entails displaying the header field values to users or  
5201 setting internal states in the UA. They SHOULD NOT however be used in the "outer" headers of any future  
5202 messages.

5203 Since MIME bodies are attached to the "inner" message, implementations will usually encrypt MIME-  
5204 specific header fields, including: MIME-Version, Content-Type, Content-Length, Content-Language,  
5205 Content-Encoding and Content-Disposition. The "outer" message will have the proper MIME header  
5206 fields for S/MIME bodies. These header fields (and any MIME bodies they preface) should be treated as  
5207 normal MIME header fields and bodies received in a SIP message.

5208 It is not particularly useful to encrypt the following header fields: Date, Min-Expires, Timestamp,  
5209 Authorization, Priority, and WWW-Authenticate. This category also includes those header fields that can  
5210 be changed by proxy servers (described in the preceding section). UAs SHOULD never include these in an

5211 “inner” message if they are not included in the “outer” message. UAs that receive any of these header fields  
5212 in an encrypted body SHOULD ignore the encrypted values.

5213 Note that extensions to SIP may define additional header fields; the authors of these extensions should  
5214 describe the integrity and confidentiality properties of such header fields. If a SIP UA encounters an un-  
5215 known header field with an integrity violation, it MUST ignore the header field.

#### 5216 23.4.2 Tunneling Integrity and Authentication

5217 Tunneling SIP messages within S/MIME bodies can provide integrity for SIP header fields if the header  
5218 fields that the sender wishes to secure are replicated in a “application/sip” MIME body signed with a CMS  
5219 detached signature.

5220 Provided that the “application/sip” body contains at least the fundamental dialog identifiers (To, From,  
5221 Call-ID, CSeq), then a signed MIME body can provide limited authentication. At the very least, if the  
5222 certificate used to sign the body is unknown to the recipient and cannot be verified, the signature can be used  
5223 to ascertain that a later request in a dialog was transmitted by the same certificate-holder that initiated the  
5224 dialog. If the recipient of the signed MIME body has some stronger incentive to trust the certificate (they  
5225 were able to verify it, acquire it from a trusted repository, or they have used it frequently) then the signature  
5226 can be taken as a stronger assertion of the identity of the subject of the certificate.

5227 In order to eliminate possible confusions about the addition or subtraction of entire header fields, senders  
5228 SHOULD replicate all header fields from the request within the signed body. Any message bodies that require  
5229 integrity protection MUST be attached to the “inner” message.

5230 If an integrity violation in a message is detected by its recipient, the message MAY be rejected with a  
5231 403 (Forbidden) response if it is a request, or any existing dialog MAY be terminated. UAs SHOULD notify  
5232 users of this circumstance and request explicit guidance on how to proceed.

5233 The following is an example of the use of a tunneled “application/sip” body:

```
5234     INVITE sip:bob@biloxi.com SIP/2.0
5235     Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
5236     To: Bob <sip:bob@biloxi.com>
5237     From: Alice <sip:alice@atlanta.com>;tag=1928301774
5238     Call-ID: a84b4c76e66710
5239     CSeq: 314159 INVITE
5240     Max-Forwards: 70
5241     Contact: <sip:alice@pc33.atlanta.com>
5242     Content-Type: multipart/signed;
5243         protocol="application/pkcs7-signature";
5244         micalg=sha1; boundary=boundary42
5245     Content-Length: 568
5246
5247     --boundary42
5248     Content-Type: application/sip
5249
5250     INVITE sip:bob@biloxi.com SIP/2.0
5251     Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
5252     To: Bob <bob@biloxi.com>
```

5253 From: Alice <alice@atlanta.com>;tag=1928301774  
5254 Call-ID: a84b4c76e66710  
5255 CSeq: 314159 INVITE  
5256 Max-Forwards: 70  
5257 Contact: <sip:alice@pc33.atlanta.com>  
5258 Content-Type: application/sdp  
5259 Content-Length: 147  
5260  
5261 v=0  
5262 o=UserA 2890844526 2890844526 IN IP4 here.com  
5263 s=Session SDP  
5264 c=IN IP4 pc33.atlanta.com  
5265 t=0 0  
5266 m=audio 49172 RTP/AVP 0  
5267 a=rtpmap:0 PCMU/8000  
5268  
5269 --boundary42  
5270 Content-Type: application/pkcs7-signature; name=smime.p7s  
5271 Content-Transfer-Encoding: base64  
5272 Content-Disposition: attachment; filename=smime.p7s;  
5273 handling=required  
5274  
5275 ghyHhHUujhJhjH77n8HHGTrfvbnj756tbB9HG4VQpfyF467GhIGfHfYT6  
5276 4VQpfyF467GhIGfHfYT6jH77n8HHGghyHhHUujhJh756tbB9HGTrfvbnj  
5277 n8HHGTrfvhJhjH776tbB9HG4VQbnj7567GhIGfHfYT6ghyHhHUujpFyF4  
5278 7GhIGfHfYT64VQbnj756  
5279  
5280 --boundary42-

### 5281 23.4.3 Tunneling Encryption

5282 It may also be desirable to use this mechanism to encrypt a “application/sip” MIME body within a CMS  
5283 EnvelopedData message S/MIME body, but in practice, most header fields are of at least some use to the  
5284 network; the general use of encryption with S/MIME is to secure message bodies like SDP rather than  
5285 message headers. Some informational header fields, such as the **Subject** or **Organization** could perhaps  
5286 warrant end-to-end security. Headers defined by future SIP applications might also require obfuscation.

5287 Another possible application of encrypting header fields is selective anonymity. A request could be con-  
5288 structed with a **From** header field that contains no personal information (for example, sip:anonymous@anonymizer.invalid).  
5289 However, a second **From** header field containing the genuine address-of-record of the originator could be  
5290 encrypted within a “application/sip” MIME body where it will only be visible to the endpoints of a dialog.  
5291 motivationNote that if this mechanism is used for anonymity, the **From** header field will no longer  
5292 be usable by the recipient of a message as an index to their certificate keychain for retrieving the proper  
5293 S/MIME key to associated with the sender. The message must first be decrypted, and the “inner” **From**  
5294 header field **MUST** be used as an index.

5295 In order to provide end-to-end integrity, encrypted “application/sip” MIME bodies **SHOULD** be signed by

5296 the sender. This creates a "multipart/signed" MIME body that contains an encrypted body and a signature,  
5297 both of type "application/pkcs7-mime".

5298 In the following example, of an encrypted and signed message, the text boxed in asterisks ("\*") is  
5299 encrypted:

```

5300     INVITE sip:bob@biloxi.com SIP/2.0
5301     Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
5302     To: Bob <sip:bob@biloxi.com>
5303     From: Anonymous <sip:anonymous@atlanta.com>;tag=1928301774
5304     Call-ID: a84b4c76e66710
5305     CSeq: 314159 INVITE
5306     Max-Forwards: 70
5307     Contact: <sip:pc33.atlanta.com>
5308     Content-Type: multipart/signed;
5309         protocol="application/pkcs7-signature";
5310         micalg=sha1; boundary=boundary42
5311     Content-Length: 568
5312
5313     --boundary42
5314     Content-Type: application/pkcs7-mime; smime-type=enveloped-data;
5315         name=smime.p7m
5316     Content-Transfer-Encoding: base64
5317     Content-Disposition: attachment; filename=smime.p7m
5318         handling=required
5319     Content-Length: 231
5320
5321     *****
5322     * Content-Type: application/sip *
5323     * * *
5324     * INVITE sip:bob@biloxi.com SIP/2.0 *
5325     * Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8 *
5326     * To: Bob <bob@biloxi.com> *
5327     * From: Alice <alice@atlanta.com>;tag=1928301774 *
5328     * Call-ID: a84b4c76e66710 *
5329     * CSeq: 314159 INVITE *
5330     * Max-Forwards: 70 *
5331     * Contact: <sip:alice@pc33.atlanta.com> *
5332     * * *
5333     * Content-Type: application/sdp *
5334     * * *
5335     * v=0 *
5336     * o=alice 53655765 2353687637 IN IP4 pc33.atlanta.com *
5337     * s=Session SDP *
5338     * t=0 0 *
5339     * c=IN IP4 pc33.atlanta.com *

```

```

5340 * m=audio 3456 RTP/AVP 0 1 3 99 *
5341 * a=rtpmap:0 PCMU/8000 *
5342 *****
5343
5344 --boundary42
5345 Content-Type: application/pkcs7-signature; name=smime.p7s
5346 Content-Transfer-Encoding: base64
5347 Content-Disposition: attachment; filename=smime.p7s;
5348     handling=required
5349
5350 ghyHhHUujhJhjH77n8HHGTrfvbnj756tbB9HG4VQpfyF467GhIGfHfYT6
5351 4VQpfyF467GhIGfHfYT6jH77n8HHGghyHhHUujhJh756tbB9HGTrfvbnj
5352 n8HHGTrfvhJhjH776tbB9HG4VQbnj7567GhIGfHfYT6ghyHhHUujpfyF4
5353 7GhIGfHfYT64VQbnj756
5354
5355 --boundary42-

```

## 5356 24 Examples

5357 In the following examples, we often omit the message body and the corresponding Content-Length and  
5358 Content-Type header fields for brevity.

### 5359 24.1 Registration

5360 Bob registers on start-up. The message flow is shown in Figure 9. Note that the authentication usually  
5361 required for registration is not shown for simplicity.

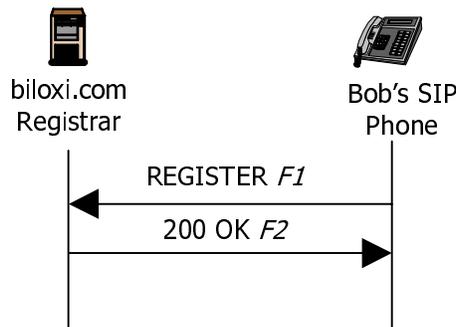


Figure 9: SIP Registration Example

```

5362
5363 F1 REGISTER Bob -> Registrar

```

5364  
5365 REGISTER sip:registrar.biloxi.com SIP/2.0  
5366 Via: SIP/2.0/UDP bobspc.biloxi.com:5060;branch=z9hG4bKnashds7  
5367 Max-Forwards: 70  
5368 To: Bob <sip:bob@biloxi.com>  
5369 From: Bob <sip:bob@biloxi.com>;tag=456248  
5370 Call-ID: 843817637684230@998sdasdh09  
5371 CSeq: 1826 REGISTER  
5372 Contact: <sip:bob@192.0.2.4>  
5373 Expires: 7200  
5374 Content-Length: 0

5375 The registration expires after two hours. The registrar responds with a 200 OK:

5376  
5377 F2 200 OK Registrar -> Bob  
5378  
5379 SIP/2.0 200 OK  
5380 Via: SIP/2.0/UDP bobspc.biloxi.com:5060;branch=z9hG4bKnashds7  
5381 ;received=192.0.2.4  
5382 To: Bob <sip:bob@biloxi.com>  
5383 From: Bob <sip:bob@biloxi.com>;tag=456248  
5384 Call-ID: 843817637684230@998sdasdh09  
5385 CSeq: 1826 REGISTER  
5386 Contact: <sip:bob@192.0.2.4>  
5387 Expires: 7200  
5388 Content-Length: 0  
5389

## 5390 24.2 Session Setup

5391 This example contains the full details of the example session setup in Section 4. The message flow is shown  
5392 in Figure 1. Note that these flows show the minimum required set of header fields - some other header fields  
5393 such as **Allow** and **Supported** would normally be present.

5394  
5395 F1 INVITE Alice -> atlanta.com proxy  
5396  
5397 INVITE sip:bob@biloxi.com SIP/2.0  
5398 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8  
5399 Max-Forwards: 70  
5400 To: Bob <sip:bob@biloxi.com>  
5401 From: Alice <sip:alice@atlanta.com>;tag=1928301774  
5402 Call-ID: a84b4c76e66710  
5403 CSeq: 314159 INVITE

5404 Contact: <sip:alice@pc33.atlanta.com>

5405 Content-Type: application/sdp

5406 Content-Length: 142

5407

5408 (Alice's SDP not shown)

5409

5410 F2 100 Trying atlanta.com proxy -> Alice

5411

5412 SIP/2.0 100 Trying

5413 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8

5414 ;received=10.1.3.3

5415 To: Bob <sip:bob@biloxi.com>

5416 From: Alice <sip:alice@atlanta.com>;tag=1928301774

5417 Call-ID: a84b4c76e66710

5418 CSeq: 314159 INVITE

5419 Content-Length: 0

5420

5421 F3 INVITE atlanta.com proxy -> biloxi.com proxy

5422

5423 INVITE sip:bob@biloxi.com SIP/2.0

5424 Via: SIP/2.0/UDP bigbox3.site3.atlanta.com;branch=z9hG4bK77ef4c2312983.1

5425 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8

5426 ;received=10.1.3.3

5427 Max-Forwards: 69

5428 To: Bob <sip:bob@biloxi.com>

5429 From: Alice <sip:alice@atlanta.com>;tag=1928301774

5430 Call-ID: a84b4c76e66710

5431 CSeq: 314159 INVITE

5432 Contact: <sip:alice@pc33.atlanta.com>

5433 Content-Type: application/sdp

5434 Content-Length: 142

5435

5436 (Alice's SDP not shown)

5437

5438 F4 100 Trying biloxi.com proxy -> atlanta.com proxy

5439

5440 SIP/2.0 100 Trying

5441 Via: SIP/2.0/UDP bigbox3.site3.atlanta.com;branch=z9hG4bK77ef4c2312983.1

5442 ;received=10.1.1.1

5443 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8

5444 ;received=10.1.3.3

5445 To: Bob <sip:bob@biloxi.com>  
5446 From: Alice <sip:alice@atlanta.com>;tag=1928301774  
5447 Call-ID: a84b4c76e66710  
5448 CSeq: 314159 INVITE  
5449 Content-Length: 0

5450  
5451 F5 INVITE biloxi.com proxy -> Bob  
5452  
5453 INVITE sip:bob@192.0.2.4 SIP/2.0  
5454 Via: SIP/2.0/UDP server10.biloxi.com;branch=z9hG4bK4b43c2ff8.1  
5455 Via: SIP/2.0/UDP bigbox3.site3.atlanta.com;branch=z9hG4bK77ef4c2312983.1  
5456 ;received=10.1.1.1  
5457 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8  
5458 ;received=10.1.3.3  
5459 Max-Forwards: 68  
5460 To: Bob <sip:bob@biloxi.com>  
5461 From: Alice <sip:alice@atlanta.com>;tag=1928301774  
5462 Call-ID: a84b4c76e66710  
5463 CSeq: 314159 INVITE  
5464 Contact: <sip:alice@pc33.atlanta.com>  
5465 Content-Type: application/sdp  
5466 Content-Length: 142  
5467  
5468 (Alice's SDP not shown)

5469  
5470 F6 180 Ringing Bob -> biloxi.com proxy  
5471  
5472 SIP/2.0 180 Ringing  
5473 Via: SIP/2.0/UDP server10.biloxi.com;branch=z9hG4bK4b43c2ff8.1  
5474 ;received=10.2.1.1  
5475 Via: SIP/2.0/UDP bigbox3.site3.atlanta.com;branch=z9hG4bK77ef4c2312983.1  
5476 ;received=10.1.1.1  
5477 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8  
5478 ;received=10.1.3.3  
5479 To: Bob <sip:bob@biloxi.com>;tag=a6c85cf  
5480 From: Alice <sip:alice@atlanta.com>;tag=1928301774  
5481 Call-ID: a84b4c76e66710  
5482 Contact: <sip:bob@192.0.2.4>  
5483 CSeq: 314159 INVITE  
5484 Content-Length: 0

5485  
5486 F7 180 Ringing biloxi.com proxy -> atlanta.com proxy

5487  
5488 SIP/2.0 180 Ringing  
5489 Via: SIP/2.0/UDP bigbox3.site3.atlanta.com;branch=z9hG4bK77ef4c2312983.1  
5490 ;received=10.1.1.1  
5491 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8  
5492 ;received=10.1.3.3  
5493 To: Bob <sip:bob@biloxi.com>;tag=a6c85cf  
5494 From: Alice <sip:alice@atlanta.com>;tag=1928301774  
5495 Call-ID: a84b4c76e66710  
5496 Contact: <sip:bob@192.0.2.4>  
5497 CSeq: 314159 INVITE  
5498 Content-Length: 0  
  
5499  
5500 F8 180 Ringing atlanta.com proxy -> Alice  
5501  
5502 SIP/2.0 180 Ringing  
5503 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8  
5504 ;received=10.1.3.3  
5505 To: Bob <sip:bob@biloxi.com>;tag=a6c85cf  
5506 From: Alice <sip:alice@atlanta.com>;tag=1928301774  
5507 Call-ID: a84b4c76e66710  
5508 Contact: <sip:bob@192.0.2.4>  
5509 CSeq: 314159 INVITE  
5510 Content-Length: 0  
  
5511  
5512 F9 200 OK Bob -> biloxi.com proxy  
5513  
5514 SIP/2.0 200 OK  
5515 Via: SIP/2.0/UDP server10.biloxi.com;branch=z9hG4bK4b43c2ff8.1  
5516 ;received=10.2.1.1  
5517 Via: SIP/2.0/UDP bigbox3.site3.atlanta.com;branch=z9hG4bK77ef4c2312983.1  
5518 ;received=10.1.1.1  
5519 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8  
5520 ;received=10.1.3.3  
5521 To: Bob <sip:bob@biloxi.com>;tag=a6c85cf  
5522 From: Alice <sip:alice@atlanta.com>;tag=1928301774  
5523 Call-ID: a84b4c76e66710  
5524 CSeq: 314159 INVITE  
5525 Contact: <sip:bob@192.0.2.4>  
5526 Content-Type: application/sdp  
5527 Content-Length: 131  
5528  
5529 (Bob's SDP not shown)

5530  
5531 F10 200 OK biloxi.com proxy -> atlanta.com proxy  
5532  
5533 SIP/2.0 200 OK  
5534 Via: SIP/2.0/UDP bigbox3.site3.atlanta.com;branch=z9hG4bK77ef4c2312983.1  
5535 ;received=10.1.1.1  
5536 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8  
5537 ;received=10.1.3.3  
5538 To: Bob <sip:bob@biloxi.com>;tag=a6c85cf  
5539 From: Alice <sip:alice@atlanta.com>;tag=1928301774  
5540 Call-ID: a84b4c76e66710  
5541 CSeq: 314159 INVITE  
5542 Contact: <sip:bob@192.0.2.4>  
5543 Content-Type: application/sdp  
5544 Content-Length: 131  
5545  
5546 (Bob's SDP not shown)  
5547  
5548 F11 200 OK atlanta.com proxy -> Alice  
5549  
5550 SIP/2.0 200 OK  
5551 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8  
5552 ;received=10.1.3.3  
5553 To: Bob <sip:bob@biloxi.com>;tag=a6c85cf  
5554 From: Alice <sip:alice@atlanta.com>;tag=1928301774  
5555 Call-ID: a84b4c76e66710  
5556 CSeq: 314159 INVITE  
5557 Contact: <sip:bob@192.0.2.4>  
5558 Content-Type: application/sdp  
5559 Content-Length: 131  
5560  
5561 (Bob's SDP not shown)  
5562  
5563 F12 ACK Alice -> Bob  
5564  
5565 ACK sip:bob@192.0.2.4 SIP/2.0  
5566 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds9  
5567 Max-Forwards: 70  
5568 To: Bob <sip:bob@biloxi.com>;tag=a6c85cf  
5569 From: Alice <sip:alice@atlanta.com>;tag=1928301774  
5570 Call-ID: a84b4c76e66710  
5571 CSeq: 314159 ACK  
5572 Content-Length: 0

5573 The media session between Alice and Bob is now established.

5574 Bob hangs up first. Note that Bob's SIP phone maintains its own CSeq numbering space, which, in  
5575 this example, begins with 231. Since Bob is making the request, the To and From URIs and tags have been  
5576 swapped.

5577

5578 F13 BYE Bob -> Alice

5579

5580 BYE sip:alice@pc33.atlanta.com SIP/2.0

5581 Via: SIP/2.0/UDP 192.0.2.4;branch=z9hG4bKnashds10

5582 Max-Forwards: 70

5583 From: Bob <sip:bob@biloxi.com>;tag=a6c85cf

5584 To: Alice <sip:alice@atlanta.com>;tag=1928301774

5585 Call-ID: a84b4c76e66710

5586 CSeq: 231 BYE

5587 Content-Length: 0

5588

5589 F14 200 OK Alice -> Bob

5590

5591 SIP/2.0 200 OK

5592 Via: SIP/2.0/UDP 192.0.2.4;branch=z9hG4bKnashds10

5593 ;received=10.1.3.3

5594 From: Bob <sip:bob@biloxi.com>;tag=a6c85cf

5595 To: Alice <sip:alice@atlanta.com>;tag=1928301774

5596 Call-ID: a84b4c76e66710

5597 CSeq: 231 BYE

5598 Content-Length: 0

5599 The SIP Call Flows document [39] contains further examples of SIP messages.

## 5600 **25 Augmented BNF for the SIP Protocol**

5601 All of the mechanisms specified in this document are described in both prose and an augmented Backus-  
5602 Naur Form (BNF) defined in RFC 2234 [10]. Section 6.1 of RFC 2234 defines a set of core rules that are  
5603 used by this specification, and not repeated here. Implementers need to be familiar with the notation and  
5604 content of RFC 2234 in order to understand this specification. Certain basic rules are in uppercase, such as  
5605 SP, LWS, HTAB, CRLF, DIGIT, ALPHA, etc. Angle brackets are used within definitions to clarify the use  
5606 of rule names.

5607 In some cases, the BNF for a choice will indicate that some elements are optional through angle brackets.  
5608 For example:

5609 `foo = bar / baz / [boo]`





5652 ctext includes all chars except left and right parens and backslash. A string of text is parsed as a single  
 5653 word if it is quoted using double-quote marks. In quoted strings, quotation marks (") and backslashes (\)  
 5654 need to be escaped.

```

quoted-string = SWS <"> *(qdtex / quoted-pair) <">
qdtex        = LWS / %x21 / %x23-5B / %x5D-7E
5655          / UTF8-NONASCII

```

5656 The backslash character ("") MAY be used as a single-character quoting mechanism only within quoted-  
 5657 string and comment constructs. Unlike HTTP/1.1, the characters CR and LF cannot be escaped by this  
 5658 mechanism to avoid conflict with line folding and header separation.

```

quoted-pair = "\" (%x00-09 / %x0B-0C
5659          / %x0E-7F)

```

```

SIP-URI      = "sip:" [ userinfo "@" ] hostport
              uri-parameters [ headers ]
SIPS-URI     = "sips:" [ userinfo "@" ] hostport
              uri-parameters [ headers ]
userinfo     = [ user / telephone-subscriber [ ":" password ] ]
user        = *( unreserved / escaped / user-unreserved )
user-unreserved = "&" / "=" / "+" / "$" / "," / ";" / "?" / "/"
password    = *( unreserved / escaped /
5660          "&" / "=" / "+" / "$" / "," )
hostport    = host [ ":" port ]
host        = hostname / IPv4address / IPv6reference
hostname    = *( domainlabel "." ) toplabel [ "." ]
domainlabel = alphanum
              / alphanum *( alphanum / "-" ) alphanum
5660 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*4HEXDIG
5661 port     = 1*DIGIT

```

5662 The BNF for telephone-subscriber can be found in RFC 2806 [9]. Note, however, that any characters  
 5663 allowed there that are not allowed in the user part of the SIP URI MUST be escaped.

uri-parameters = \*( ";" uri-parameter )  
 uri-parameter = transport-param / user-param / method-param  
                   / ttl-param / maddr-param / lr-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  
 lr-param = "lr"  
 other-param = pname [ "=" pvalue ]  
 pname = 1\*paramchar  
 pvalue = 1\*paramchar  
 paramchar = param-unreserved / unreserved / escaped  
 5664 param-unreserved = "[ / ]" / "/" / "?" / "." / "&" / "+" / "\$"  
  
 headers = "?" header \*( "&" header )  
 header = hname "=" hvalue  
 hname = 1\*( hnv-unreserved / unreserved / escaped )  
 hvalue = \*( hnv-unreserved / unreserved / escaped )  
 5665 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-URI / SIPS-URI / absoluteURI  
 absoluteURI = scheme ":" ( 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 / "," / "?" / ":" / "@"  
                 / "&" / "=" / "+" / "\$" / ";"  
 path-segments = segment \*( "/" segment )  
 segment = \*pchar \*( ";" param )  
 param = \*pchar  
 pchar = unreserved / escaped /  
           "." / "@" / "&" / "=" / "+" / "\$" / ";"  
 scheme = ALPHA \*( ALPHA / DIGIT / "+" / "-" / "." )  
 authority = srvr / reg-name  
 srvr = [ [ userinfo "@" ] hostport ]  
 reg-name = 1\*( unreserved / escaped / "\$" / "  
                 / ";" / ":" / "@" / "&" / "=" / "+" )  
 query = \*uric  
 SIP-Version = "SIP" "/" 1\*DIGIT "." 1\*DIGIT

5666

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  
/ Min-Expires  
/ Organization  
/ Priority  
/ Proxy-Authenticate  
/ Proxy-Authorization  
/ Proxy-Require  
/ Record-Route  
/ Reply-To  
/ Require  
/ Retry-After  
/ Route  
/ Server  
/ Subject  
/ Supported  
/ Timestamp  
/ To  
/ Unsupported  
/ User-Agent  
/ Via  
/ Warning  
/ WWW-Authenticate  
/ extension-header) CRLF

5667

INVITE<sub>m</sub> = %x49.4E.56.49.54.45 ; INVITE in caps  
 ACK<sub>m</sub> = %x41.43.4B ; ACK in caps  
 OPTIONSm = %x4F.50.54.49.4F.4E.53 ; OPTIONS in caps  
 BYE<sub>m</sub> = %x42.59.45 ; BYE in caps  
 CANCEL<sub>m</sub> = %x43.41.4E.43.45.4C ; CANCEL in caps  
 REGISTER<sub>m</sub> = %x52.45.47.49.53.54.45.52 ; REGISTER in caps  
 Method = INVITE<sub>m</sub> / ACK<sub>m</sub> / OPTIONSm / BYE<sub>m</sub>  
           / CANCEL<sub>m</sub> / REGISTER<sub>m</sub>  
           / extension-method  
 extension-method = token  
 Response = Status-Line  
           \*( message-header )  
           CRLF  
 5668           [ message-body ]

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  
 5669            / UTF8-NONASCII / UTF8-CONT / SP / HTAB)

Informational = "100" ; Trying  
               / "180" ; Ringing  
               / "181" ; Call Is Being Forwarded  
 5670           / "182" ; Queued  
               / "183" ; Session Progress

5671       Success = "200" ; OK

Redirection = "300" ; Multiple Choices  
               / "301" ; Moved Permanently  
               / "302" ; Moved Temporarily  
               / "305" ; Use Proxy  
 5672           / "380" ; Alternative Service

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  
/ "410" ; Gone  
/ "413" ; Request Entity Too Large  
/ "414" ; Request-URI Too Large  
/ "415" ; Unsupported Media Type  
/ "416" ; Unsupported URI Scheme  
/ "420" ; Bad Extension  
/ "421" ; Extension Required  
/ "423" ; Interval Too Brief  
/ "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  
/ "488" ; Not Acceptable Here  
/ "491" ; Request Pending  
/ "493" ; Undecipherable

5673

Server-Error = "500" ; Internal Server Error  
/ "501" ; Not Implemented  
/ "502" ; Bad Gateway  
/ "503" ; Service Unavailable  
/ "504" ; Server Time-out  
/ "505" ; SIP Version not supported  
/ "513" ; Message Too Large

5674

Global-Failure = "600" ; Busy Everywhere  
/ "603" ; Decline  
/ "604" ; Does not exist anywhere  
/ "606" ; Not Acceptable

5675

Accept = "Accept" HCOLON  
 ( accept-range \*(COMMA accept-range) )  
 accept-range = media-range [ accept-params ]  
 media-range = ( "\*"/\*"  
 / ( m-type SLASH "\*" )  
 / ( m-type SLASH m-subtype )  
 ) \*( SEMI m-parameter )  
 accept-params = SEMI "q" EQUAL qvalue \*( accept-extension )  
 accept-extension = SEMI ae-name [ EQUAL ae-value ]  
 ae-name = token  
 5676 ae-value = token / quoted-string

Accept-Encoding = "Accept-Encoding" HCOLON  
 ( encoding \*(COMMA encoding) )  
 encoding = codings [ SEMI "q" EQUAL qvalue ]  
 codings = content-coding / "\*" "  
 content-coding = token  
 5677 qvalue = ( "0" [ "." 0\*3DIGIT ] )  
 / ( "1" [ "." 0\*3("0") ] )

Accept-Language = "Accept-Language" HCOLON  
 ( language \*(COMMA language) )  
 language = language-range [ SEMI "q" EQUAL qvalue ]  
 5678 language-range = ( ( 1\*8ALPHA \*( "-" 1\*8ALPHA ) ) / "\*" )

Alert-Info = "Alert-Info" HCOLON alert-param \*(COMMA alert-param)  
 alert-param = LAQUOT absoluteURI RAQUOT \*( SEMI generic-param )  
 generic-param = token [ EQUAL gen-value ]  
 5679 gen-value = token / host / quoted-string

5680 Allow = "Allow" HCOLON Method \*(COMMA Method)

Authorization credentials = "Authorization" HCOLON credentials  
 = ("Digest" LWS digest-response)  
 / other-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 LDQUOT digest-uri-value RDQUOT  
 digest-uri-value = rquest-uri ; Equal to 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  
 other-response = auth-scheme LWS auth-param  
 \*(COMMA auth-param)  
 5681 auth-scheme = token  
  
 Authentication-Info = "Authentication-Info" HCOLON ainfo  
 \*(COMMA ainfo)  
 ainfo = [nextnonce] / [ message-qop ]  
 / [ response-auth ] / [ cnonce ]  
 / [nonce-count]  
 nextnonce = "nextnonce" EQUAL nonce-value  
 response-auth = "rspauth" EQUAL response-digest  
 5682 response-digest = LDQUOT \*LHEX RDQUOT  
  
 Call-ID = ( "Call-ID" / "i" ) HCOLON callid  
 5683 callid = word [ "@" word ]  
  
 Call-Info = "Call-Info" HCOLON info \*(COMMA info)  
 info = LAQUOT absoluteURI RAQUOT \*( SEMI info-param)  
 info-param = ( "purpose" EQUAL ( "icon" / "info"  
 5684 / "card" / token ) ) / generic-param

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-URI / SIPS-URI / absoluteURI

5685 display-name = \*(token LWS)/ quoted-string

contact-params = c-p-q / c-p-expires  
/ contact-extension

c-p-q = "q" EQUAL qvalue

c-p-expires = "expires" EQUAL delta-seconds

contact-extension = generic-param

5686 delta-seconds = 1\*DIGIT

Content-Disposition = "Content-Disposition" HCOLON  
disp-type \*( SEMI disp-param )

disp-type = "render" / "session" / "icon" / "alert"  
/ disp-extension-token

disp-param = handling-param / generic-param

handling-param = "handling" EQUAL  
( "optional" / "required"  
/ other-handling )

other-handling = token

5687 disp-extension-token = token

Content-Encoding = ( "Content-Encoding" / "e" ) HCOLON  
content-coding \*(COMMA content-coding)

5688

Content-Language = "Content-Language" HCOLON  
language-tag \*(COMMA language-tag)

language-tag = primary-tag \*( "-" subtag )

primary-tag = 1\*8ALPHA

5689 subtag = 1\*8ALPHA

5690 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-" token  
 m-subtype = extension-token / iana-token  
 iana-token = token  
 m-parameter = m-attribute EQUAL m-value  
 m-attribute = token  
 5691 m-value = token / quoted-string

5692 CSeq = "CSeq" HCOLON 1\*DIGIT LWS Method

Date = "Date" HCOLON SIP-date  
 SIP-date = rfc1123-date  
 rfc1123-date = wkday "," 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"  
 5693       / "Sep" / "Oct" / "Nov" / "Dec"

Error-Info = "Error-Info" HCOLON error-uri \*(COMMA error-uri)  
 5694 error-uri = LAQUOT absoluteURI RAQUOT \*( SEMI generic-param )

Expires = "Expires" HCOLON delta-seconds  
 From = ( "From" / "f" ) HCOLON from-spec  
 from-spec = ( name-addr / addr-spec )  
             \*( SEMI from-param )  
 from-param = tag-param / generic-param  
 5695 tag-param = "tag" EQUAL token

5696 In-Reply-To = "In-Reply-To" HCOLON callid \*(COMMA callid)

5697 Max-Forwards = "Max-Forwards" HCOLON 1\*DIGIT

5698 MIME-Version = "MIME-Version" HCOLON 1\*DIGIT "." 1\*DIGIT

5699           Min-Expires = "Min-Expires" HCOLON delta-seconds

5700           Organization = "Organization" HCOLON TEXT-UTF8-TRIM

              Priority = "Priority" HCOLON priority-value

              priority-value = "emergency" / "urgent" / "normal"

                              / "non-urgent" / other-priority

5701           other-priority = token

              Proxy-Authenticate = "Proxy-Authenticate" HCOLON challenge

              challenge = ("Digest" LWS digest-chn \*(COMMA digest-chn))

                              / other-challenge

              other-challenge = auth-scheme LWS auth-param

                              \*(COMMA auth-param)

              digest-chn = realm / [ domain ] / nonce

                              / [ opaque ] / [ stale ] / [ algorithm ]

                              / [ qop-options ] / [auth-param]

              realm = "realm" EQUAL 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

                              \*( " " qop-value ) RDQUOT

5702           qop-value = "auth" / "auth-int" / token

5703           Proxy-Authorization = "Proxy-Authorization" HCOLON credentials

              Proxy-Require = "Proxy-Require" HCOLON option-tag

                              \*(COMMA option-tag)

5704           option-tag = token

              Record-Route = "Record-Route" HCOLON rec-route \*(COMMA rec-route)

              rec-route = name-addr \*( SEMI rr-param )

5705           rr-param = generic-param

              Reply-To = "Reply-To" HCOLON rplyto-spec

              rplyto-spec = ( name-addr / addr-spec )

                              \*( SEMI rplyto-param )

              rplyto-param = generic-param

5706           Require = "Require" HCOLON option-tag \*(COMMA option-tag)

Retry-After = "Retry-After" HCOLON delta-seconds  
 [ comment ] \*( SEMI retry-param )  
 5707 retry-param = ("duration" EQUAL delta-seconds)  
 / generic-param

5708 Route = "Route" HCOLON route-param \*(COMMA route-param)  
 route-param = name-addr \*( SEMI rr-param )

5709 Server = "Server" HCOLON 1\*( product / comment )  
 product = token [SLASH product-version]  
 product-version = token

5710 Subject = ( "Subject" / "s" ) HCOLON TEXT-UTF8-TRIM

5711 Supported = ( "Supported" / "k" ) HCOLON  
 [option-tag \*(COMMA option-tag)]

5712 Timestamp = "Timestamp" HCOLON 1\*(DIGIT)  
 [ "." \*(DIGIT) ] [ delay ]  
 delay = \*(DIGIT) [ "." \*(DIGIT) ]

5713 To = ( "To" / "t" ) HCOLON ( name-addr  
 / addr-spec ) \*( SEMI to-param )  
 to-param = tag-param / generic-param

5714 Unsupported = "Unsupported" HCOLON option-tag \*(COMMA option-tag)

5715 User-Agent = "User-Agent" HCOLON 1\*( product / comment )

Via = ( "Via" / "v" ) HCOLON via-param \*(COMMA via-param)  
 via-param = sent-protocol LWS sent-by \*( SEMI via-params )  
 via-params = via-ttl / via-maddr  
 / via-received / via-branch  
 / via-extension  
 via-ttl = "ttl" EQUAL ttl  
 via-maddr = "maddr" EQUAL host  
 via-received = "received" EQUAL (IPv4address / IPv6address)  
 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 ]  
 5716 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 = hostport / pseudonym  
; the name or pseudonym of the server adding  
; the Warning header, for use in debugging  
warn-text = quoted-string  
pseudonym = token

WWW-Authenticate = "WWW-Authenticate" HCOLON challenge

extension-header = header-name HCOLON header-value  
header-name = token  
header-value = \*(TEXT-UTF8char / UTF8-CONT / LWS)

message-body = \*OCTET

## 26 Security Considerations: Threat Model and Security Usage Recommendations

SIP is not an easy protocol to secure. Its use of intermediaries, its multi-faceted trust relationships, its expected usage between elements with no trust at all, and its user-to-user operation make security far from trivial. Security solutions are needed that are deployable today, without extensive coordination, in a wide variety of environments and usages. In order to meet these diverse needs, several distinct mechanisms applicable to different aspects and usages of SIP will be required.

Note that the security of SIP signaling itself has no bearing on the security of protocols used in concert with SIP such as RTP, or with the security implications of any specific bodies SIP might carry (although MIME security plays a substantial role in securing SIP). Any media associated with a session can be encrypted end-to-end independently of any associated SIP signaling. Media encryption is outside the scope of this document.

The considerations that follow first examine a set of classic threat models that broadly identify the security needs of SIP. The set of security services required to address these threats is then detailed, followed by an explanation of several security mechanisms that can be used to provide these services. Next, the requirements for implementers of SIP are enumerated, along with exemplary deployments in which these security mechanisms could be used to improve the security of SIP. Some notes on privacy conclude this section.

### 26.1 Attacks and Threat Models

This section details some threats that should be common to most deployments of SIP. These threats have been chosen specifically to illustrate each of the security services that SIP requires.

The following examples by no means provide an exhaustive list of the threats against SIP; rather, these are "classic" threats that demonstrate the need for particular security services that can potentially prevent whole categories of threats.

5745       These attacks assume an environment in which attackers can potentially read any packet on the network  
5746 - it is anticipated that SIP will frequently be used on the public Internet. Attackers on the network may be  
5747 able to modify packets (perhaps at some compromised intermediary). Attackers may wish to steal services,  
5748 eavesdrop on communications, or disrupt sessions.

### 5749 **26.1.1 Registration Hijacking**

5750 The SIP registration mechanism allows a user agent to identify itself to a registrar as a device at which a  
5751 user (designated by an address of record) is located. A registrar assesses the identity asserted in the **From**  
5752 header field of a **REGISTER** message to determine whether this request can modify the contact addresses  
5753 associated with the address-of-record in the **To** header field. While these two fields are frequently the same,  
5754 there are many valid deployments in which a third-party may register contacts on a user's behalf.

5755       The **From** header field of a SIP request, however, can be modified arbitrarily by the owner of a UA, and  
5756 this opens the door to malicious registrations. An attacker that successfully impersonates a party authorized  
5757 to change contacts associated with an address-of-record could, for example, de-register all existing contacts  
5758 for a URI and then register their own device as the appropriate contact address, thereby directing all requests  
5759 for the affected user to the attacker's device.

5760       This threat belongs to a family of threats that rely on the absence of cryptographic assurance of a re-  
5761 quest's originator. Any SIP UAS that represents a valuable service (a gateway that interworks SIP requests  
5762 with traditional telephone calls, for example) might want to control access to its resources by authenticating  
5763 requests that it receives. Even end-user UAs, for example SIP phones, have an interest in ascertaining the  
5764 identities of originators of requests.

5765       This threat demonstrates the need for security services that enable SIP entities to authenticate the origi-  
5766 nators of requests.

### 5767 **26.1.2 Impersonating a Server**

5768 The domain to which a request is destined is generally specified in the **Request-URI**. UAs commonly  
5769 contact a server in this domain directly in order to deliver a request. However, there is always a possibility  
5770 that an attacker could impersonate the remote server, and that the UA's request could be intercepted by some  
5771 other party.

5772       For example, consider a case in which a redirect server at one domain, *chicago.com*, impersonates a  
5773 redirect server at another domain, *biloxi.com*. A user agent sends a request to *biloxi.com*, but the redirect  
5774 server at *chicago.com* answers with a forged response that has appropriate SIP header fields for a response  
5775 from *biloxi.com*. The forged contact addresses in the redirection response could direct the originating UA  
5776 to inappropriate or insecure resources, or simply prevent requests for *biloxi.com* from succeeding.

5777       This family of threats has a vast membership, many of which are critical. As a converse to the registration  
5778 hijacking threat, consider the case in which a registration sent to *biloxi.com* is intercepted by *chicago.com*,  
5779 which replies to the intercepted registration with a forged 301 (Moved Permanently) response. This response  
5780 might seem to come from *biloxi.com* yet designate *chicago.com* as the appropriate registrar. All future  
5781 **REGISTER** requests from the originating UA would then go to *chicago.com*.

5782       Prevention of this threat requires a means by which UAs can authenticate the servers to whom they send  
5783 requests.

### 5784 **26.1.3 Tampering with Message Bodies**

5785 As a matter of course, SIP UAs route requests through trusted proxy servers. Regardless of how that trust is  
5786 established (authentication of proxies is discussed elsewhere in this section), a UA may trust a proxy server  
5787 to route a request, but not to inspect or possibly modify the bodies contained in that request.

5788 Consider a UA that is using SIP message bodies to communicate session encryption keys for a media  
5789 session. Although it trusts the proxy server of the domain it is contacting to deliver signaling properly, it  
5790 may not want the administrators of that domain to be capable of decrypting any subsequent media session.  
5791 Worse yet, if the proxy server were actively malicious, it could modify the session key, either acting as a  
5792 man-in-the-middle, or perhaps changing the security characteristics requested by the originating UA.

5793 This family of threats applies not only to session keys, but to most conceivable forms of content car-  
5794 ried end-to-end in SIP. These might include MIME bodies that should be rendered to the user, SDP, or  
5795 encapsulated telephony signals, among others. Attackers might attempt to modify SDP bodies, for example,  
5796 in order to point RTP media streams to a wiretapping device in order to eavesdrop on subsequent voice  
5797 communications.

5798 Also note that some header fields in SIP are meaningful end-to-end, for example, **Subject**. UAs might  
5799 be protective of these header fields as well as bodies (a malicious intermediary changing the **Subject** header  
5800 field might make an important request appear to be spam, for example). However, since many header fields  
5801 are legitimately inspected or altered by proxy servers as a request is routed, not all header fields should be  
5802 secured end-to-end.

5803 For these reasons, the UA might want to secure SIP message bodies, and in some limited cases header  
5804 fields, end-to-end. The security services required for bodies include confidentiality, integrity, and authen-  
5805 tication. These end-to-end services should be independent of the means used to secure interactions with  
5806 intermediaries such as proxy servers.

### 5807 **26.1.4 Tearing Down Sessions**

5808 Once a dialog has been established by initial messaging, subsequent requests can be sent that modify the  
5809 state of the dialog and/or session. It is critical that principals in a session can be certain that such requests  
5810 are not forged by attackers.

5811 Consider a case in which a third-party attacker captures some initial messages in a dialog shared by two  
5812 parties in order to learn the parameters of the session (**To** tag, **From** tag, and so forth) and then inserts a  
5813 **BYE** request into the session. The attacker could opt to forge the request such that it seemed to come from  
5814 either participant. Once the **BYE** is received by its target, the session will be torn down prematurely.

5815 Similar mid-session threats include the transmission of forged re-**INVITE**s that alter the session (possibly  
5816 to reduce session security or redirect media streams as part of a wiretapping attack).

5817 The most effective countermeasure to this threat is the authentication of the sender of the **BYE**. In this  
5818 instance, the recipient needs only know that the **BYE** came from the same party with whom the correspond-  
5819 ing dialog was established (as opposed to ascertaining the absolute identity of the sender). Also, if the  
5820 attacker is unable to learn the parameters of the session due to confidentiality, it would not be possible to  
5821 forge the **BYE**. However, some intermediaries (like proxy servers) will need to inspect those parameters as  
5822 the session is established.

### 5823 **26.1.5 Denial of Service and Amplification**

5824 Denial-of-service attacks focus on rendering a particular network element unavailable, usually by directing  
5825 an excessive amount of network traffic at its interfaces. A distributed denial-of-service attack allows one  
5826 network user to cause multiple network hosts to flood a target host with a large amount of network traffic.

5827 In many architectures, SIP proxy servers face the public Internet in order to accept requests from world-  
5828 wide IP endpoints. SIP creates a number of potential opportunities for distributed denial-of-service attacks  
5829 that must be recognized and addressed by the implementers and operators of SIP systems.

5830 Attackers can create bogus requests that contain a falsified source IP address and a corresponding *Via*  
5831 header field that identify a targeted host as the originator of the request and then send this request to a large  
5832 number of SIP network elements, thereby using hapless SIP UAs or proxies to generate denial-of-service  
5833 traffic aimed at the target.

5834 Similarly, attackers might use falsified *Route* header field values in a request that identify the target  
5835 host and then send such messages to forking proxies that will amplify messaging sent to the target. *Record-  
5836 Route* could be used to similar effect when the attacker is certain that the SIP dialog initiated by the request  
5837 will result in numerous transactions originating in the backwards direction.

5838 A number of denial-of-service attacks open up if *REGISTER* requests are not properly authenticated  
5839 and authorized by registrars. Attackers could de-register some or all users in an administrative domain,  
5840 thereby preventing these users from being invited to new sessions. An attacker could also register a large  
5841 number of contacts designating the same host for a given address-of-record in order to use the registrar and  
5842 any associated proxy servers as amplifiers in a denial-of-service attack. Attackers might also attempt to  
5843 deplete available memory and disk resources of a registrar by registering huge numbers of bindings.

5844 The use of multicast to transmit SIP requests can greatly increase the potential for denial-of-service  
5845 attacks.

5846 These problems demonstrate a general need to define architectures that minimize the risks of denial-of-  
5847 service, and the need to be mindful in recommendations for security mechanisms of this class of attacks.

## 5848 **26.2 Security Mechanisms**

5849 From the threats described above, we gather that the fundamental security services required for the SIP  
5850 protocol are: preserving the confidentiality and integrity of messaging, preventing replay attacks or message  
5851 spoofing, providing for the authentication and privacy of the participants in a session, and preventing denial-  
5852 of-service attacks. Bodies within SIP messages separately require the security services of confidentiality,  
5853 integrity, and authentication.

5854 Rather than defining new security mechanisms specific to SIP, SIP reuses wherever possible existing  
5855 security models derived from the HTTP and SMTP space.

5856 Full encryption of messages provides the best means to preserve the confidentiality of signaling - it  
5857 can also guarantee that messages are not modified by any malicious intermediaries. However, SIP requests  
5858 and responses cannot be naively encrypted end-to-end in their entirety because message fields such as the  
5859 *Request-URI*, *Route*, and *Via* need to be visible to proxies in most network architectures so that SIP  
5860 requests are routed correctly. Note that proxy servers need to modify some features of messages as well (such  
5861 as adding *Via* header field values) in order for SIP to function. Proxy servers must therefore be trusted, to  
5862 some degree, by SIP UAs. To this purpose, low-layer security mechanisms for SIP are recommended, which  
5863 encrypt the entire SIP requests or responses on the wire on a hop-by-hop basis, and that allow endpoints to  
5864 verify the identity of proxy servers to whom they send requests.

5865 SIP entities also have a need to identify one another in a secure fashion. When a SIP endpoint asserts

5866 the identity of its user to a peer UA or to a proxy server, that identity should in some way be verifiable. A  
5867 cryptographic authentication mechanism is provided in SIP to address this requirement.

5868 An independent security mechanism for SIP message bodies supplies an alternative means of end-to-end  
5869 mutual authentication, as well as providing a limit on the degree to which user agents must trust intermedi-  
5870 aries.

### 5871 **26.2.1 Transport and Network Layer Security**

5872 Transport or network layer security encrypts signaling traffic, guaranteeing message confidentiality and  
5873 integrity.

5874 Oftentimes, certificates are used in the establishment of lower-layer security, and these certificates can  
5875 also be used to provide a means of authentication in many architectures.

5876 Two popular alternatives for providing security at the transport and network layer are, respectively, TLS  
5877 [16] and IPSec [25].

5878 IPSec is a set of network-layer protocol tools that collectively can be used as a secure replacement for  
5879 traditional IP (Internet Protocol). IPSec is most commonly used in architectures in which a set of hosts or  
5880 administrative domains have an existing trust relationship with one another. IPSec is usually implemented  
5881 at the operating system level in a host, or on a security gateway that provides confidentiality and integrity  
5882 for all traffic it receives from a particular interface (as in a VPN architecture). IPSec can also be used on a  
5883 hop-by-hop basis.

5884 In many architectures IPSec does not require integration with SIP applications; IPSec is perhaps best  
5885 suited to deployments in which adding security directly to SIP hosts would be arduous. UAs that have a  
5886 pre-shared keying relationship with their first-hop proxy server are also good candidates to use IPSec. Any  
5887 deployment of IPSec for SIP would require an IPSec profile describing the protocol tools that would be  
5888 required to secure SIP. No such profile is given in this document.

5889 TLS provides transport-layer security over connection-oriented protocols (for the purposes of this docu-  
5890 ment, TCP); "tls" (signifying TLS over TCP) can be specified as the desired transport protocol within a Via  
5891 header field value or a SIP-URI. TLS is most suited to architectures in which hop-by-hop security is required  
5892 between hosts with no pre-existing trust association. For example, Alice trusts her local proxy server, which  
5893 after a certificate exchange decides to trust Bob's local proxy server, which Bob trusts, hence Bob and Alice  
5894 can communicate securely.

5895 TLS must be tightly coupled with a SIP application. Note that transport mechanisms are specified on  
5896 a hop-by-hop basis in SIP, thus a UA that sends requests over TLS to a proxy server has no assurance that  
5897 TLS will be used end-to-end.

5898 The TLS\_RSA\_WITH\_AES\_128\_CBC\_SHA ciphersuite MUST be supported at a minimum by imple-  
5899 menters when TLS is used in a SIP application. For purposes of backwards compatibility, proxy servers,  
5900 redirect servers, and registrars SHOULD support TLS\_RSA\_WITH\_3DES\_EDE\_CBC\_SHA. Implementers  
5901 MAY also support any other ciphersuite.

### 5902 **26.2.2 SIPS URI scheme**

5903 The SIPS URI scheme adheres to the syntax of the SIP URI (described in 19), although the scheme string is  
5904 "sips" rather than "sip". The semantics of SIPS are very different from the SIP URI, however.

5905 A SIPS URI can be used as an address-of-record for a particular user - the URI by which the user is  
5906 canonically known (on their business cards, in the From header field of their requests, in the To header field  
5907 of REGISTER requests). When used as the Request-URI of a request, the SIPS scheme signifies that each

5908 hop over which the request is forwarded must be secured with TLS. When used by the originator of a request  
5909 (as would be the case if they encountered a SIPS URI as the address-of-record of the target), SIPS dictates  
5910 that the entire request path be so secured. No other mechanism in SIP allows the originator of a request to  
5911 specify security characteristics that are preferred for the entire request path.

5912 The SIPS scheme is also applicable to many of the other ways in which SIP URIs are used in SIP today,  
5913 including in the Request-URI, in addresses-of-record, contact addresses (populating Contact headers, in-  
5914 cluding those of REGISTER methods), and Route headers. The SIPS URI scheme allows these existing  
5915 fields to designate secure resources.

5916 In effect, using SIPS in the Request-URI ensures that TLS is used on every segment between the  
5917 originator of the request and the destination. This is a handy service, though one that is useful only in  
5918 architectures in which it is desirable to use TLS for every hop.

5919 The use of SIPS in particular entails that mutual TLS authentication SHOULD be employed, as SHOULD  
5920 the ciphersuite TLS\_RSA\_WITH\_AES\_128\_CBC\_SHA. Certificates received in the authentication process  
5921 SHOULD be verified against root certificates in the client; failure to verify a certificate SHOULD result in the  
5922 failure of the request.

5923 motivationNote that in the SIPS URI scheme, transport is independent of TLS, and thus “sips:alice@atlanta.com;transport=tl  
5924 and “sips:alice@atlanta.com;transport=sctp” are both valid (although note that UDP is not a valid transport  
5925 for SIPS). The use of “transport=tl” has consequently been deprecated, partly because it was specific to a  
5926 single hop of the request. This is a change since RFC 2543.

5927 Users that distribute a SIPS URI as an address-of-record may elect to operate devices that do not even  
5928 accept requests over insecure transports.

### 5929 **26.2.3 HTTP Authentication**

5930 SIP provides a challenge capability, based on HTTP authentication, that relies on the 401 and 407 response  
5931 codes as well as header fields for carrying challenges and credentials. Without significant modification, the  
5932 reuse of the HTTP Digest authentication scheme in SIP allows for replay protection and one-way authenti-  
5933 cation.

5934 The usage of Digest authentication in SIP is detailed in Section 22.

### 5935 **26.2.4 S/MIME**

5936 As is discussed above, encrypting entire SIP messages end-to-end for the purpose of confidentiality is not  
5937 appropriate because network intermediaries (like proxy servers) need to view certain header fields in order  
5938 to route messages correctly, and if these intermediaries are excluded from security associations, then SIP  
5939 messages will essentially be non-routable.

5940 However, S/MIME allows SIP UAs to encrypt MIME bodies within SIP, securing these bodies end-to-  
5941 end without affecting message headers. S/MIME can provide end-to-end confidentiality and integrity for  
5942 message bodies, as well as mutual authentication. It is also possible to use S/MIME to provide a form of  
5943 integrity and confidentiality for SIP header fields through SIP message tunneling.

5944 The usage of S/MIME in SIP is detailed in Section 23.

## 5945 **26.3 Implementing Security Mechanisms**

### 5946 **26.3.1 Requirements for Implementers of SIP**

5947 Proxy servers, redirect servers, and registrars **MUST** implement TLS, and **MUST** support both mutual and  
5948 one-way authentication. It is strongly **RECOMMENDED** that UAs be capable initiating TLS; UAs **MAY**  
5949 also be capable of acting as a TLS server. Proxy servers, redirect servers, and registrars **SHOULD** possess  
5950 a site certificate whose subject corresponds to their canonical hostname. UAs **MAY** have certificates of  
5951 their own for mutual authentication with TLS, but no provisions are set forth in this document for their  
5952 use. All SIP elements that support TLS **MUST** have a mechanism for verifying certificates received during  
5953 TLS negotiation; this entails possession of one or more root certificates issued by certificate authorities  
5954 (preferably well-known distributors of site certificates comparable to those issuing root certificates for web  
5955 browsers). All SIP elements that support TLS **MUST** also support the SIPS URI scheme.

5956 Proxy servers, redirect servers, registrars, and UAs **MAY** also implement IPsec or other lower-layer  
5957 security protocols.

5958 When a UA attempts to contact a proxy server, redirect server, or registrar, the UAC **SHOULD** initiate a  
5959 TLS connection over which it will send SIP messages. In some architectures, UASs **MAY** receive requests  
5960 over such TLS connections as well.

5961 Proxy servers, redirect servers, registrars, and UAs **MUST** implement Digest Authorization, encompassing  
5962 all of the aspects required in 22. Proxy servers, redirect servers, and registrars **SHOULD** be configured with  
5963 at least one Digest realm, and at least one “realm” string supported by a given server **SHOULD** correspond  
5964 to the server’s hostname or domainname.

5965 UAs **MAY** support the signing and encrypting of MIME bodies, and transference of credentials with  
5966 S/MIME as described in 23. If a UA holds one or more root certificates of certificate authorities in order to  
5967 verify certificates for TLS or IPsec, it **SHOULD** be capable of reusing these to verify S/MIME certificates,  
5968 as appropriate. A UA **MAY** hold root certificates specifically for verifying S/MIME certificates.

5969 Note that it is anticipated that future security extensions may upgrade the normative strength associated with  
5970 S/MIME as S/MIME implementations appear and the problem space becomes better understood.

### 5971 **26.3.2 Security Solutions**

5972 The operation of these security mechanisms in concert can follow the existing web and email security models  
5973 to some degree. At a high level, UAs authenticate themselves to servers (proxy servers, redirect servers, and  
5974 registrars) with a Digest username and password; servers authenticate themselves to UAs one hop away, or  
5975 to another server one hop away (and vice versa), with a site certificate delivered by TLS.

5976 On a peer-to-peer level, UAs trust the network to authenticate one another ordinarily; however, S/MIME  
5977 can also be used to provide direct authentication when the network does not, or if the network itself is not  
5978 trusted.

5979 The following is an illustrative example in which these security mechanisms are used by various UAs  
5980 and servers to prevent the sorts of threats described in Section 26.1. While implementers and network  
5981 administrators **MAY** follow the normative guidelines given in the remainder of this section, these are provided  
5982 only as example implementations.

5983 **26.3.2.1 Registration** When a UA comes online and registers with its local administrative domain, it  
5984 **SHOULD** establish a TLS connection with its registrar (Section 10 describes how the UA reaches its reg-  
5985 istrar). The registrar **SHOULD** offer a certificate to the UA, and the site identified by the certificate **MUST**

5986 correspond with the domain in which the UA intends to register; for example, if the UA intends to register  
5987 the address-of-record 'alice@atlanta.com', the site certificate must identify a host within the atlanta.com  
5988 domain (such as sip.atlanta.com). When it receives the TLS Certificate message, the UA SHOULD verify the  
5989 certificate and inspect the site identified by the certificate. If the certificate is invalid, revoked, or if it does  
5990 not identify the appropriate party, the UA MUST NOT send the REGISTER message and otherwise proceed  
5991 with the registration.

5992           When a valid certificate has been provided by the registrar, the UA knows that the registrar is not an attacker  
5993           who might redirect the UA, steal passwords, or attempt any similar attacks.

5994       The UA then creates a REGISTER request that SHOULD be addressed to a Request-URI correspond-  
5995       ing to the site certificate received from the registrar. When the UA sends the REGISTER request over  
5996       the existing TLS connection, the registrar SHOULD challenge the request with a 401 (Proxy Authentication  
5997       Required) response. The "realm" parameter within the Proxy-Authenticate header field of the response  
5998       SHOULD correspond to the domain previously given by the site certificate. When the UAC receives the  
5999       challenge, it SHOULD either prompt the user for credentials or take an appropriate credential from a keyring  
6000       corresponding to the "realm" parameter in the challenge. The username of this credential SHOULD corre-  
6001       spond with the "userinfo" portion of the URI in the To header field of the REGISTER request. Once the  
6002       Digest credentials have been inserted into an appropriate Proxy-Authorization header field, the REGIS-  
6003       TER should be resubmitted to the registrar.

6004           Since the registrar requires the user agent to authenticate itself, it would be difficult for an attacker to forge REG-  
6005       ISTER requests for the user's address-of-record. Also note that since the REGISTER is sent over a confidential  
6006       TLS connection, attackers will not be able to intercept the REGISTER to record credentials for any possible replay  
6007       attack.

6008       Once the registration has been accepted by the registrar, the UA SHOULD leave this TLS connection  
6009       open provided that the registrar also acts as the proxy server to which requests are sent for users in this  
6010       administrative domain. The existing TLS connection will be reused to deliver incoming requests to the UA  
6011       that has just completed registration.

6012           Because the UA has already authenticated the server on the other side of the TLS connection, all requests that  
6013       come over this connection are known to have passed through the proxy server - attackers cannot create spoofed  
6014       requests that appear to have been sent through that proxy server.

6015 **26.3.2.2 Interdomain Requests** Now let's say that Alice's UA would like to initiate a session with a user  
6016 in a remote administrative domain, namely "bob@biloxi.com". We will also say that the local administrative  
6017 domain (atlanta.com) has a local outbound proxy.

6018       The proxy server that handles inbound requests for an administrative domain MAY also act as a local  
6019       outbound proxy; for simplicity's sake we'll assume this to be the case for atlanta.com (otherwise the user  
6020       agent would initiate a new TLS connection to a separate server at this point). Assuming that the client has  
6021       completed the registration process described in the preceding section, it SHOULD reuse the TLS connection  
6022       to the local proxy server when it sends an INVITE request to another user. The UA SHOULD reuse cached  
6023       credentials in the INVITE to avoid prompting the user unnecessarily.

6024       When the local outbound proxy server has validated the credentials presented by the UA in the INVITE,  
6025       it SHOULD inspect the Request-URI to determine how the message should be routed (see [4]). If the  
6026       "domainname" portion of the Request-URI had corresponded to the local domain (atlanta.com) rather than  
6027       biloxi.com, then the proxy server would have consulted its location service to determine how best to reach  
6028       the requested user.

6029           Had "alice@atlanta.com" been attempting to contact, say, "alex@atlanta.com", the local proxy would have  
6030 proxied to the request to the TLS connection Alex had established with the registrar when he registered. Since  
6031 Alex would receive this request over his authenticated channel, he would be assured that Alice's request had been  
6032 authorized by the proxy server of the local administrative domain.

6033           However, in this instance the Request-URI designates a remote domain. The local outbound proxy  
6034 server at atlanta.com SHOULD therefore establish a TLS connection with the remote proxy server at biloxi.com.  
6035 Since both of the participants in this TLS connection are servers that possess site certificates, mutual TLS  
6036 authentication SHOULD occur. Each side of the connection SHOULD verify and inspect the certificate of  
6037 the other, noting the domain name that appears in the certificate for comparison with the header fields of  
6038 SIP messages. The atlanta.com proxy server, for example, SHOULD verify at this stage that the certificate  
6039 received from the remote side corresponds with the biloxi.com domain. Once it has done so, and TLS ne-  
6040 gotiation has completed, resulting in a secure channel between the two proxies, the atlanta.com proxy can  
6041 forward the INVITE request to biloxi.com.

6042           The proxy server at biloxi.com SHOULD inspect the certificate of the proxy server at atlanta.com in turn  
6043 and compare the domain asserted by the certificate with the "domainname" portion of the From header field  
6044 in the INVITE request. The biloxi proxy MAY have a strict security policy that requires it to reject requests  
6045 that do not match the administrative domain from which they have been proxied.

6046           Such security policies could be instituted to prevent the SIP equivalent of SMTP 'open relays' that are frequently  
6047 exploited to generate spam.

6048           This policy, however, only guarantees that the request came from the domain it ascribes to itself; it  
6049 does not allow biloxi.com to ascertain how atlanta.com authenticated Alice. Only if biloxi.com has some  
6050 other way of knowing atlanta.com's authentication policies could it possibly ascertain how Alice proved her  
6051 identity. biloxi.com might then institute an even stricter policy that forbids requests that come from domains  
6052 that are not known administratively to share a common authentication policy with biloxi.com.

6053           Once the INVITE has been approved by the biloxi proxy, the proxy server SHOULD identify the existing  
6054 TLS channel, if any, associated with the user targeted by this request (in this case "bob@biloxi.com"). The  
6055 INVITE should be proxied through this channel to Bob. Since the request is received over a TLS connection  
6056 that had previously been authenticated as the biloxi proxy, Bob knows that the From header field was not  
6057 tampered with and that atlanta.com has validated Alice, although not necessarily whether or not to trust  
6058 Alice's identity.

6059           Before they forward the request, both proxy servers SHOULD add a Record-Route header field to the  
6060 request so that all future requests in this dialog will pass through the proxy servers. The proxy servers can  
6061 thereby continue to provide security services for the lifetime of this dialog. If the proxy servers do not add  
6062 themselves to the Record-Route, future messages will pass directly end-to-end between Alice and Bob  
6063 without any security services (unless the two parties agree on some independent end-to-end security such  
6064 as S/MIME). In this respect the SIP trapezoid model can provide a nice structure where conventions of  
6065 agreement between the site proxies can provide a reasonably secure channel between Alice and Bob.

6066           An attacker preying on this architecture would, for example, be unable to forge a BYE request and insert it into  
6067 the signaling stream between Bob and Alice because the attacker has no way of ascertaining the parameters of the  
6068 session and also because the integrity mechanism transitively protects the traffic between Alice and Bob.

6069 **26.3.2.3 Peer to Peer Requests** Alternatively, consider a UA asserting the identity "carol@chicago.com"  
6070 that has no local outbound proxy. When Carol wishes to send an INVITE to "bob@biloxi.com", her UA  
6071 SHOULD initiate a TLS connection with the biloxi proxy directly (using the mechanism described in [4]  
6072 to determine how to best to reach the given Request-URI). When her UA receives a certificate from the

6073 biloxi proxy, it SHOULD be verified normally before she passes her INVITE across the TLS connection.  
6074 However, Carol has no means of proving her identity to the biloxi proxy, but she does have a CMS-detached  
6075 signature over a “message/sip” body in the INVITE. It is unlikely in this instance that Carol would have any  
6076 credentials in the biloxi.com realm, since she has no formal association with biloxi.com. The biloxi proxy  
6077 MAY also have a strict policy that precludes it from even bothering to challenge requests that do not have  
6078 biloxi.com in the “domainname” portion of the From header field - it treats these users as unauthenticated.

6079 The biloxi proxy has a policy for Bob that all non-authenticated requests should be redirected to the  
6080 appropriate contact address registered against 'bob@biloxi.com', namely <sip:bob@192.0.2.4>. Carol  
6081 receives the redirection response over the TLS connection she established with the biloxi proxy, so she  
6082 trusts the veracity of the contact address.

6083 Carol SHOULD then establish a TCP connection with the designated address and send a new INVITE  
6084 with a Request-URI containing the received contact address (recomputing the signature in the body as  
6085 the request is readied). Bob receives this INVITE on an insecure interface, but his UA inspects and, in  
6086 this instance, recognizes the From header field of the request and subsequently matches a locally cached  
6087 certificate with the one presented in the signature of the body of the INVITE. He replies in similar fashion,  
6088 authenticating himself to Carol, and a secure dialog begins.

6089 Sometimes firewalls or NATs in an administrative domain could preclude the establishment of a direct TCP  
6090 connection to a UA. In these cases, proxy servers could also potentially relay requests to UAs in a way that has no  
6091 trust implications (for example, forgoing an existing TLS connection and forwarding the request over cleartext TCP)  
6092 as local policy dictates.

6093 **26.3.2.4 DoS Protection** In order to minimize the risk of a denial-of-service attack against architectures  
6094 using these security solutions, implementers should take note of the following guidelines.

6095 When the host on which a SIP proxy server is operating is routable from the public Internet, it SHOULD  
6096 be deployed in an administrative domain with defensive operational policies (blocking source-routed traffic,  
6097 preferably filtering ping traffic). Both TLS and IPSec can also make use of bastion hosts at the edges of  
6098 administrative domains that participate in the security associations to aggregate secure tunnels and sockets.  
6099 These bastion hosts can also take the brunt of denial-of-service attacks, ensuring that SIP hosts within the  
6100 administrative domain are not encumbered with superfluous messaging.

6101 No matter what security solutions are deployed, floods of messages directed at proxy servers can lock up  
6102 proxy server resources and prevent desirable traffic from reaching its destination. There is a computational  
6103 expense associated with processing a SIP transaction at a proxy server, and that expense is greater for  
6104 stateful proxy servers than it is for stateless proxy servers. Therefore, stateful proxies are more susceptible  
6105 to flooding than stateless proxy servers.

6106 UAs and proxy servers SHOULD challenge questionable requests with only a *single* 401 (Unauthorized)  
6107 or 407 (Proxy Authentication Required), forgoing the normal response retransmission algorithm, and thus  
6108 behaving statelessly towards unauthenticated requests.

6109 Retransmitting the 401 (Unauthorized) or 407 (Proxy Authentication Required) status response amplifies the  
6110 problem of an attacker using a falsified header field value (such as Via) to direct traffic to a third party.

6111 In summary, the mutual authentication of proxy servers through mechanisms such as TLS significantly  
6112 reduces the potential for rogue intermediaries to introduce falsified requests or responses that can deny  
6113 service. This commensurately makes it harder for attackers to make innocent SIP nodes into agents of  
6114 amplification.

## 6115 **26.4 Limitations**

6116 Although these security mechanisms, when applied in a judicious manner, can thwart many threats, there are  
6117 limitations in the scope of the mechanisms that must be understood by implementers and network operators.

### 6118 **26.4.1 HTTP Digest**

6119 One of the primary limitations of using HTTP Digest in SIP is that the integrity mechanisms in Digest do  
6120 not work very well for SIP. Specifically, they offer protection of the Request-URI and the method of a  
6121 message, but not for any of the header fields that UAs would most likely wish to secure.

6122 The existing replay protection mechanisms described in RFC 2617 also have some limitations for SIP.  
6123 The next-nonce mechanism, for example, does not support pipelined requests. The nonce-count mechanism  
6124 should be used for replay protection.

6125 Another limitation of HTTP Digest is the scope of realms. Digest is valuable when a user wants to  
6126 authenticate themselves to a resource with which they have a pre-existing association, like a service provider  
6127 of which the user is a customer (which is quite a common scenario and thus Digest provides an extremely  
6128 useful function). By way of contrast, the scope of TLS is interdomain or multirealm, since certificates are  
6129 often globally verifiable, so that the UA can authenticate the server with no pre-existing association.

### 6130 **26.4.2 S/MIME**

6131 The largest outstanding defect with the S/MIME mechanism is the lack of a prevalent public key infrastruc-  
6132 ture for end users. If self-signed certificates (or certificates that cannot be verified by one of the participants  
6133 in a dialog) are used, the SIP-based key exchange mechanism described in Section 23.2 is susceptible to a  
6134 man-in-the-middle attack with which an attacker can potentially inspect and modify S/MIME bodies. The  
6135 attacker needs to intercept the first exchange of keys between the two parties in a dialog, remove the exist-  
6136 ing CMS-detached signatures from the request and response, and insert a different CMS-detached signature  
6137 containing a certificate supplied by the attacker (but which seems to be a certificate for the proper address-  
6138 of-record). Each party will think they have exchanged keys with the other, when in fact each has the public  
6139 key of the attacker.

6140 It is important to note that the attacker can only leverage this vulnerability on the first exchange of keys  
6141 between two parties - on subsequent occasions, the alteration of the key would be noticeable to the UAs. It  
6142 would also be difficult for the attacker to remain in the path of all future dialogs between the two parties  
6143 over time (as potentially days, weeks, or years pass).

6144 SSH is susceptible to the same man-in-the-middle attack on the first exchange of keys; however, it is  
6145 widely acknowledged that while SSH is not perfect, it does improve the security of connections. The use of  
6146 key fingerprints could provide some assistance to SIP, just as it does for SSH. For example, if two parties use  
6147 SIP to establish a voice communications session, each could read off the fingerprint of the key they received  
6148 from the other, which could be compared against the original. It would certainly be more difficult for the  
6149 man-in-the-middle to emulate the voices of the participants than their signaling (a practice that was used  
6150 with the Clipper chip-based secure telephone).

6151 The S/MIME mechanism allows UAs to send encrypted requests without preamble if they possess a  
6152 certificate for the destination address-of-record on their keyring. However, it is possible that any particular  
6153 device registered for an address-of-record will not hold the certificate that has been previously employed by  
6154 the device's current user, and that it will therefore be unable to process an encrypted request properly, which  
6155 could lead to some avoidable error signaling. This is especially likely when an encrypted request is forked.

6156 The keys associated with S/MIME are most useful when associated with a particular user (an address-  
6157 of-record) rather than a device (a UA). When users move between devices, it may be difficult to transport  
6158 private keys securely between UAs; how such keys might be acquired by a device is outside the scope of  
6159 this document.

6160 Another, more prosaic difficulty with the S/MIME mechanism is that it can result in very large messages,  
6161 especially when the SIP tunneling mechanism described in Section 23.4 is used. For that reason, it is  
6162 RECOMMENDED that TCP should be used as a transport protocol when S/MIME tunneling is employed.

### 6163 26.4.3 TLS

6164 The most commonly voiced concern about TLS is that it cannot run over UDP; TLS requires a connection-  
6165 oriented underlying transport protocol, which for the purposes of this document means TCP.

6166 It may also be arduous for a local outbound proxy server and/or registrar to maintain many simultaneous  
6167 long-lived TLS connections with numerous UAs. This introduces some valid scalability concerns, especially  
6168 for intensive ciphersuites. Maintaining redundancy of long-lived TLS connections, especially when a UA is  
6169 solely responsible for their establishment, could also be cumbersome.

6170 TLS only allows SIP entities to authenticate servers to which they are adjacent; TLS offers strictly  
6171 hop-by-hop security. Neither TLS, nor any other mechanism specified in this document, allows clients to  
6172 authenticate proxy servers to whom they cannot form a direct TCP connection.

### 6173 26.4.4 SIPS URIs

6174 Using TLS on every segment of a request path entails that the terminating UAS must be reachable over TLS.  
6175 This means that many hybrid architectures that use TLS to secure part of the request path, but rely on some  
6176 other mechanism for the final hop to a UAS, cannot make use of the SIPS AoR. Also, since many UAs will  
6177 not accept incoming TLS connections, even those UAs that do support TLS may be required to maintain  
6178 persistent TLS connections as described in the TLS limitations section above.

6179 It is very difficult to guarantee that TLS will be used end-to-end. It is possible that cryptographically  
6180 authenticated proxy servers that are non-compliant or compromised may choose to disregard the forwarding  
6181 rules associated with SIPS. These intermediaries may, for example, retarget a request from a SIPS URI to  
6182 a SIP URI. It is therefore recommended that recipients of a request to SIP URI inspect the To header field  
6183 value to see if it contains a SIPS URI. S/MIME may also be used to ensure that the original form of the To  
6184 header field is carried end-to-end. Entities that accept only SIPS request may also refuse connections on  
6185 insecure ports.

6186 End users will undoubtedly discern the difference between SIPS and SIP URIs, and they may manually  
6187 edit them in response to stimuli. This can either benefit or degrade security. For example, if an attacker  
6188 corrupts a DNS cache, inserting a fake record set that effectively removes all SIPS records for a proxy  
6189 server, then any SIPS requests that traverse this proxy server may fail. When a user, however, sees that  
6190 repeated calls to a SIPS AoR are failing, on some devices they could manually convert the scheme from  
6191 SIPS to SIP and retry. Of course, there are some safeguards against this (if the destination UA is truly  
6192 paranoid it could refuse all non-SIPS requests), but it is a limitation worth noting. On the bright side, users  
6193 might also divine that 'SIPS' would be valid even when they are presented only with a SIP URI.

## 6194 26.5 Privacy

6195 SIP messages frequently contain sensitive information about their senders - not just what they have to say, but  
6196 with whom they communicate, when they communicate and for how long, and from where they participate  
6197 in sessions. Many applications and their users require that this sort of private information be hidden from  
6198 any parties that do not need to know it.

6199 Note that there are also less direct ways in which private information can be divulged. If a user or service  
6200 chooses to be reachable at an address that is guessable from the person's name and organizational affiliation  
6201 (which describes most addresses-of-record), the traditional method of ensuring privacy by having an unlisted  
6202 "phone number" is compromised. A user location service can infringe on the privacy of the recipient of a  
6203 session invitation by divulging their specific whereabouts to the caller; an implementation consequently  
6204 SHOULD be able to restrict, on a per-user basis, what kind of location and availability information is given  
6205 out to certain classes of callers. This is a whole class of problem that is expected to be studied further in  
6206 ongoing SIP work.

6207 In some cases, users may want to conceal personal information in header fields that convey identity. This  
6208 can apply not only to the **From** and related headers representing the originator of the request, but also the  
6209 **To** - it may not be appropriate to convey to the final destination a speed-dialing nickname, or an unexpanded  
6210 identifier for a group of targets, either of which would be removed from the **Request-URI** as the request is  
6211 routed, but not changed in the **To** header field if the two were initially identical. Thus it MAY be desirable  
6212 for privacy reasons to create a **To** header field that differs from the **Request-URI**.

## 6213 27 IANA Considerations

6214 All new or experimental method names, header field names, and status codes used in SIP applications  
6215 SHOULD be registered with IANA in order to prevent potential naming conflicts. It is RECOMMENDED that  
6216 new "option-tag"s and "warn-code"s also be registered. Before IANA registration, new protocol elements  
6217 SHOULD be described in an Internet-Draft or, preferably, an RFC.

6218 For Internet-Drafts, IANA is requested to make the draft available as part of the registration database.

6219 By the time an RFC is published, colliding names may have already been implemented.

6220 When a registration for either a new header field, new method, or new status code is created based on  
6221 an Internet-Draft, and that Internet-Draft becomes an RFC, the person that performed the registration MUST  
6222 notify IANA to change the registration to point to the RFC instead of the Internet-Draft.

6223 Registrations should be sent to `iana@iana.org`.

### 6224 27.1 Option Tags

6225 Option tags are used in header fields such as **Require**, **Supported**, **Proxy-Require**, and **Unsupported** in  
6226 support of SIP compatibility mechanisms for extensions (Section 19.2). The option tag itself is a string that  
6227 is associated with a particular SIP option (that is, an extension). It identifies the option to SIP endpoints.

6228 When registering a new SIP option with IANA, the following information MUST be provided:

- 6229 ● Name and description of option. The name MAY be of any length, but SHOULD be no more than  
6230 twenty characters long. The name MUST consist of alphanum (Section 25) characters only.
- 6231 ● A listing of any new SIP header fields, header parameter fields, or parameter values defined by this  
6232 option. A SIP option MUST NOT redefine header fields or parameters defined in either RFC 2543, any

- 6233 standards-track extensions to RFC 2543, or other extensions registered through IANA.
- 6234 • Indication of who has change control over the option (for example, IETF, ISO, ITU-T, other interna-  
6235 tional standardization bodies, a consortium, or a particular company or group of companies).
  - 6236 • A reference to a further description if available, for example (in order of preference) an RFC, a pub-  
6237 lished paper, a patent filing, a technical report, documented source code, or a computer manual.
  - 6238 • Contact information (postal and email address).

6239 This procedure has been borrowed from RTSP [28] and the RTP AVP [40].

## 6240 27.2 Warn-Codes

6241 Warning codes provide information supplemental to the status code in SIP response messages when the  
6242 failure of the transaction results from a Session Description Protocol (SDP, [1]). New “warn-code” values  
6243 can be registered with IANA as they arise.

6244 The “warn-code” consists of three digits. A first digit of “3” indicates warnings specific to SIP.

6245 Warnings 300 through 329 are reserved for indicating problems with keywords in the session description,  
6246 330 through 339 are warnings related to basic network services requested in the session description, 370  
6247 through 379 are warnings related to quantitative QoS parameters requested in the session description, and  
6248 390 through 399 are miscellaneous warnings that do not fall into one of the above categories.

6249 1xx and 2xx have been taken by HTTP/1.1.

## 6250 27.3 Header Field Names

6251 Header field names do not require working group or working group chair review prior to IANA registration,  
6252 but SHOULD be documented in an RFC or Internet-Draft before IANA is consulted.

6253 The following information needs to be provided to IANA in order to register a new header field name:

- 6254 • The name and email address of the individual performing the registration;
- 6255 • the name of the header field being registered;
- 6256 • a compact form version for that header field, if one is defined;
- 6257 • the name of the draft or RFC where the header field is defined;
- 6258 • a copy of the draft or RFC where the header field is defined.

6259 Header fields SHOULD NOT use the X- prefix notation and MUST NOT duplicate the names of header  
6260 fields used by SMTP or HTTP unless the syntax is a compatible superset and the semantics are similar.  
6261 Some common and widely used header fields MAY be assigned one-letter compact forms (Section 7.3.3).  
6262 Compact forms can only be assigned after SIP working group review. In the absence of this working group,  
6263 a designated expert reviews the request.

## 6264 **27.4 Method and Response Codes**

6265 Because the status code space is limited, they do require working group or working group chair review, and  
6266 MUST be documented in an RFC or Internet draft. The same procedures apply to new method names.

6267 The following information needs to be provided to IANA in order to register a new response code or  
6268 method:

- 6269 • The name and email address of the individual performing the registration;
- 6270 • the number of the response code or name of the method being registered;
- 6271 • the default reason phrase for that status code, if applicable;
- 6272 • the name of the draft or RFC where the method or status code is defined;
- 6273 • a copy of the draft or RFC where the method or status code is defined.

## 6274 **27.5 The “application/sip” MIME type.**

6275 This document registers the “application/sip” MIME media type in order to allow SIP messages to be tun-  
6276 nelled as bodies within SIP, primarily for end-to-end security purposes. This media type is defined by the  
6277 following information:

6278 Media type name: application Media subtype name: sip Required parameters: none Optional parame-  
6279 ters: version

- 6280 • version: The SIP-Version number of the enclosed message (e.g., "2.0"). If not present, the version  
6281 can be determined from the first line of the body.

6282 Encoding scheme: see below Security considerations: see below

6283 SIP specifies UTF-8 encoding. While most header field names and data elements will lie in the 7-bit  
6284 ASCII compatible range, data elements and SIP bodies may contain 8-bit values. In order to preserve the  
6285 readability of SIP messages being carried as the body of other messages, “application/sip” bodies (including  
6286 any bodies they in turn contain) SHOULD be UTF-8 encoded. If transcoding a body to UTF-8 is not feasible,  
6287 the “application/sip” part MAY be binary encoded. If the transport is not 8-bit clean, encoding formats such  
6288 as base-64 can be used.

6289 Motivation and examples of this usage as a security mechanism in concert with S/MIME are given in  
6290 23.4.

## 6291 **28 Changes From RFC 2543**

6292 This RFC revises RFC 2543. It is mostly backwards compatible with RFC 2543. The changes described  
6293 here fix many errors discovered in RFC 2543 and provide information on scenarios not detailed in RFC  
6294 2543. The protocol has been presented in a more cleanly layered model here.

6295 We break the differences into functional behavior that is a substantial change from RFC 2543, which has  
6296 impact on interoperability or correct operation in some cases, and functional behavior that is different from  
6297 RFC 2543 but not a potential source of interoperability problems. There have been countless clarifications  
6298 as well, which are not documented here.

## 6299 28.1 Major Functional Changes

- 6300 ● When a UAC wishes to terminate a call before it has been answered, it sends **CANCEL**. If the original  
6301 **INVITE** still returns a 2xx, the UAC then sends **BYE**. **BYE** can only be sent on an existing call leg  
6302 (now called a dialog in this RFC), whereas it could be sent at any time in RFC 2543.
- 6303 ● The SIP BNF was converted to be RFC 2234 compliant.
- 6304 ● SIP URL BNF was made more general, allowing a greater set of characters in the user part. Fur-  
6305 thermore, comparison rules were simplified to be primarily case-insensitive, and detailed handling of  
6306 comparison in the presence of parameters was described. The most substantial change is that a URI  
6307 with a parameter with the default value does not match a URI without that parameter.
- 6308 ● Removed **Via** hiding. It had serious trust issues, since it relied on the next hop to perform the obfus-  
6309 cation process. Instead, **Via** hiding can be done as a local implementation choice in stateful proxies,  
6310 and thus is no longer documented.
- 6311 ● In RFC 2543, **CANCEL** and **INVITE** transactions were intermingled. They are separated now. When  
6312 a user sends an **INVITE** and then a **CANCEL**, the **INVITE** transaction still terminates normally. A  
6313 UAS needs to respond to the original **INVITE** request with a 487 response.
- 6314 ● Similarly, **CANCEL** and **BYE** transactions were intermingled; RFC 2543 allowed the UAS not to  
6315 send a response to **INVITE** when a **BYE** was received. That is disallowed here. The original **INVITE**  
6316 needs a response.
- 6317 ● In RFC 2543, UAs needed to support only UDP. In this RFC, UAs need to support both UDP and  
6318 TCP.
- 6319 ● In RFC 2543, a forking proxy only passed up one challenge from downstream elements in the event  
6320 of multiple challenges. In this RFC, proxies are supposed to collect all challenges and place them into  
6321 the forwarded response.
- 6322 ● In Digest credentials, the URI needs to be quoted; this is unclear from RFC 2617 and RFC 2069 which  
6323 are both inconsistent on it.
- 6324 ● SDP processing has been split off into a separate specification [13], and more fully specified as a  
6325 formal offer/answer exchange process that is effectively tunneled through SIP. SDP is allowed in  
6326 **INVITE/200** or **200/ACK** for baseline SIP implementations; RFC 2543 alluded to the ability to use it  
6327 in **INVITE**, **200**, and **ACK** in a single transaction, but this was not well specified. More complex SDP  
6328 usages are allowed in extensions.
- 6329 ● Added full support for IPv6 in URIs and in the **Via** header field. Support for IPv6 in **Via** has required  
6330 that its header field parameters allow the square bracket and colon characters. These characters were  
6331 previously not permitted. In theory, this could cause interop problems with older implementations.  
6332 However, we have observed that most implementations accept any non-control ASCII character in  
6333 these parameters.
- 6334 ● DNS SRV procedure is now documented in a separate specification [4]. This procedure uses both SRV  
6335 and NAPTR resource records and no longer combines data from across SRV records as described in  
6336 RFC 2543.

- 6337 ● Loop detection has been made optional, supplanted by a mandatory usage of **Max-Forwards**. The  
6338 loop detection procedure in RFC 2543 had a serious bug which would report “spirals” as an error  
6339 condition when it was not. The optional loop detection procedure is more fully and correctly specified  
6340 here.
- 6341 ● Usage of tags is now mandatory (they were optional in RFC 2543), as they are now the fundamental  
6342 building blocks of dialog identification.
- 6343 ● Added the **Supported** header field, allowing for clients to indicate what extensions are supported to  
6344 a server, which can apply those extensions to the response, and indicate their usage with a **Require** in  
6345 the response.
- 6346 ● Extension parameters were missing from the BNF for several header fields, and they have been added.
- 6347 ● Handling of **Route** and **Record-Route** construction was very underspecified in RFC 2543, and also  
6348 not the right approach. It has been substantially reworked in this specification (and made vastly  
6349 simpler), and this is arguably the largest change. Backwards compatibility is still provided for de-  
6350 ployments that do not use “pre-loaded routes”, where the initial request has a set of **Route** header  
6351 field values obtained in some way outside of **Record-Route**. In those situations, the new mechanism  
6352 is not interoperable.
- 6353 ● In RFC 2543, lines in a message could be terminated with CR, LF, or CRLF. This specification only  
6354 allows CRLF.
- 6355 ● Comments (expressed with rounded brackets) have been removed from the grammar of SIP.
- 6356 ● Usage of **Route** in **CANCEL** and **ACK** was not well defined in RFC 2543. It is now well specified; if  
6357 a request had a **Route** header field, its **CANCEL** or **ACK** for a non-2xx response to the request need  
6358 to carry the same **Route** header field values. **ACKs** for 2xx responses use the **Route** values learned  
6359 from the **Record-Route** of the 2xx responses.
- 6360 ● RFC 2543 allowed multiple requests in a single UDP packet. This usage has been removed.
- 6361 ● Usage of absolute time in the **Expires** header field and parameter has been removed. It caused inter-  
6362 operability problems in elements that were not time synchronized, a common occurrence. Relative  
6363 times are used instead.
- 6364 ● The branch parameter of the **Via** header field value is now mandatory for all elements to use. It now  
6365 plays the role of a unique transaction identifier. This avoids the complex and bug-laden transaction  
6366 identification rules from RFC 2543. A magic cookie is used in the parameter value to determine if  
6367 the previous hop has made the parameter globally unique, and comparison falls back to the old rules  
6368 when it is not present. Thus, interoperability is assured.
- 6369 ● In RFC 2543, closure of a TCP connection was made equivalent to a **CANCEL**. This was nearly  
6370 impossible to implement (and wrong) for TCP connections between proxies. This has been eliminated,  
6371 so that there is no coupling between TCP connection state and SIP processing.
- 6372 ● RFC 2543 was silent on whether a UA could initiate a new transaction to a peer while another was in  
6373 progress. That is now specified here. It is allowed for non-INVITE requests, disallowed for INVITE.

- 6374 ● PGP was removed. It was not sufficiently specified, and not compatible with the more complete PGP  
6375 MIME. It was replaced with S/MIME.
- 6376 ● Additional security features were added with TLS, and these are described in a much larger and  
6377 complete security considerations section.
- 6378 ● In RFC 2543, a proxy was not required to forward provisional responses from 101 to 199 upstream.  
6379 This was changed to MUST. This is important, since many subsequent features depend on delivery of  
6380 all provisional responses from 101 to 199.
- 6381 ● Little was said about the 503 response code in RFC 2543. It has since found substantial use in indicat-  
6382 ing failure or overload conditions in proxies. This requires somewhat special treatment. Specifically,  
6383 receipt of a 503 should trigger an attempt to contact the next element in the result of a DNS SRV  
6384 lookup. Also, 503 response is only forwarded upstream by a proxy under certain conditions.
- 6385 ● RFC 2543 defined, but did not sufficiently specify, a mechanism for UA authentication of a server.  
6386 That has been removed. Instead, the mutual authentication procedures of RFC 2617 are allowed.
- 6387 ● A UA cannot send a BYE for a call until it has received an ACK for the initial INVITE. This was  
6388 allowed in RFC 2543 but leads to a potential race condition.
- 6389 ● A UA or proxy cannot send CANCEL for a transaction until it gets a provisional response for the  
6390 request. This was allowed in RFC 2543 but leads to potential race conditions.
- 6391 ● The action parameter in registrations has been deprecated. It was insufficient for any useful services,  
6392 and caused conflicts when application processing was applied in proxies.
- 6393 ● RFC 2543 had a number of special cases for multicast. For example, certain responses were sup-  
6394 pressed, timers were adjusted, and so on. Multicast now plays a more limited role, and the protocol  
6395 operation is unaffected by usage of multicast as opposed to unicast. The limitations as a result of that  
6396 are documented.
- 6397 ● Basic authentication has been removed entirely and its usage forbidden.
- 6398 ● Proxies no longer forward a 6xx immediately on receiving it. Instead, they CANCEL pending  
6399 branches immediately. This avoids a potential race condition that would result in a UAC getting a  
6400 6xx followed by a 2xx. In all cases except this race condition, the result will be the same - the 6xx is  
6401 forwarded upstream.
- 6402 ● RFC 2543 did not address the problem of request merging. This occurs when a request forks at a  
6403 proxy and later rejoins at an element. Handling of merging is done only at a UA, and procedures are  
6404 defined for rejecting all but the first request.

## 6405 **28.2 Minor Functional Changes**

- 6406 ● Added the Alert-Info, Error-Info, and Call-Info header fields for optional content presentation to  
6407 users.
- 6408 ● Added the Content-Language, Content-Disposition and MIME-Version header fields.

- 6409     • Added a “glare handling” mechanism to deal with the case where both parties send each other a  
6410       re-INVITE simultaneously. It uses the new 491 (Request Pending) error code.
- 6411     • Added the In-Reply-To and Reply-To header fields for supporting the return of missed calls or mes-  
6412       sages at a later time.
- 6413     • Added TLS and SCTP as valid SIP transports.
- 6414     • There were a variety of mechanisms described for handling failures at any time during a call; those  
6415       are now generally unified. BYE is sent to terminate.
- 6416     • RFC 2543 mandated retransmission of INVITE responses over TCP, but noted it was really only  
6417       needed for 2xx. That was an artifact of insufficient protocol layering. With a more coherent transaction  
6418       layer defined here, that is no longer needed. Only 2xx responses to INVITEs are retransmitted over  
6419       TCP.
- 6420     • Client and server transaction machines are now driven based on timeouts rather than retransmit counts.  
6421       This allows the state machines to be properly specified for TCP and UDP.
- 6422     • The Date header field is used in REGISTER responses to provide a simple means for auto-configuration  
6423       of dates in user agents.
- 6424     • Allowed a registrar to reject registrations with expirations that are too short in duration. Defined the  
6425       423 response code and the Min-Expires for this purpose.
- 6426     • Added the “sips” URI scheme for end-to-end TLS. This scheme is not backwards compatible with  
6427       RFC 2543. Existing elements that receive a request with a SIPS URI scheme in the Request-URI  
6428       will likely reject the request. This is actually a feature; it ensures that a call to a SIPS URI is only  
6429       delivered if all path hops can be secured.

## 6430 **29 Acknowledgments**

6431 We wish to thank the members of the IETF MMUSIC and SIP WGs for their comments and suggestions.  
6432 Detailed comments were provided by Brian Bidulock, Jim Buller, Neil Deason, Dave Devanathan, Keith  
6433 Drage, Cédric Fluckiger, Yaron Goland, John Hearty, Bernie Höneisen, Jo Hornsby, Phil Hoffer, Christian  
6434 Huitema, Jean Jervis, Gadi Karimi, Peter Kjellerstedt, Anders Kristensen, Jonathan Lennox, Gethin Liddell,  
6435 Allison Mankin, William Marshall, Rohan Mahy, Keith Moore, Vern Paxson, Moshe J. Sambol, Chip Sharp,  
6436 Igor Slepchin, Eric Tremblay, and Rick Workman.

6437     Brian Rosen provided the compiled BNF.

6438     This work is based, inter alia, on [41, 42].

## 6439 **30 Authors’ Addresses**

6440 Authors addresses are listed alphabetically for the editors, the writers, and then the original authors of RFC  
6441 2543. All listed authors actively contributed large amounts of text to this document.

6442 Jonathan Rosenberg  
6443 dynamicsoft  
6444 72 Eagle Rock Ave  
6445 East Hanover, NJ 07936  
6446 USA  
6447 electronic mail: jdrosen@dynamicsoft.com

6448 Henning Schulzrinne  
6449 Dept. of Computer Science  
6450 Columbia University  
6451 1214 Amsterdam Avenue  
6452 New York, NY 10027  
6453 USA  
6454 electronic mail: schulzrinne@cs.columbia.edu

6455 Gonzalo Camarillo  
6456 Ericsson  
6457 Advanced Signalling Research Lab.  
6458 FIN-02420 Jorvas  
6459 Finland  
6460 electronic mail: Gonzalo.Camarillo@ericsson.com

6461 Alan Johnston  
6462 WorldCom  
6463 100 South 4th Street  
6464 St. Louis, MO 63102  
6465 USA  
6466 electronic mail: alan.johnston@wcom.com

6467 Jon Peterson  
6468 NeuStar, Inc  
6469 1800 Sutter Street, Suite 570  
6470 Concord, CA 94520  
6471 USA  
6472 electronic mail: jon.peterson@neustar.com

6473 Robert Sparks  
6474 dynamicsoft, Inc.  
6475 5100 Tennyson Parkway  
6476 Suite 1200  
6477 Plano, Texas 75024  
6478 USA  
6479 electronic mail: rsparks@dynamicsoft.com

6480 Mark Handley  
6481 ACIRI  
6482 electronic mail: mjh@aciri.org

6483 Eve Schooler  
6484 Computer Science Department 256-80  
6485 California Institute of Technology  
6486 Pasadena, CA 91125  
6487 USA  
6488 electronic mail: schooler@cs.caltech.edu

## 6489 Normative References

- 6490 [1] M. Handley and V. Jacobson, "SDP: session description protocol," Request for Comments 2327, Inter-  
6491 net Engineering Task Force, Apr. 1998.
- 6492 [2] S. Bradner, "Key words for use in RFCs to indicate requirement levels," Request for Comments 2119,  
6493 Internet Engineering Task Force, Mar. 1997.
- 6494 [3] P. Resnick and Editor, "Internet message format," Request for Comments 2822, Internet Engineering  
6495 Task Force, Apr. 2001.
- 6496 [4] H. Schulzrinne and J. Rosenberg, "SIP: Session initiation protocol – locating SIP servers," Internet  
6497 Draft, Internet Engineering Task Force, Mar. 2001. Work in progress.
- 6498 [5] T. Berners-Lee, R. Fielding, and L. Masinter, "Uniform resource identifiers (URI): generic syntax,"  
6499 Request for Comments 2396, Internet Engineering Task Force, Aug. 1998.
- 6500 [6] T. Berners-Lee, L. Masinter, and M. McCahill, "Uniform resource locators (URL)," Request for Com-  
6501 ments 1738, Internet Engineering Task Force, Dec. 1994.
- 6502 [7] F. Yergeau, "UTF-8, a transformation format of ISO 10646," Request for Comments 2279, Internet  
6503 Engineering Task Force, Jan. 1998.
- 6504 [8] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, and T. Berners-Lee, "Hypertext  
6505 transfer protocol – HTTP/1.1," Request for Comments 2616, Internet Engineering Task Force, June  
6506 1999.
- 6507 [9] A. Vaha-Sipila, "URLs for telephone calls," Request for Comments 2806, Internet Engineering Task  
6508 Force, Apr. 2000.
- 6509 [10] D. Crocker, Ed., and P. Overell, "Augmented BNF for syntax specifications: ABNF," Request for  
6510 Comments 2234, Internet Engineering Task Force, Nov. 1997.
- 6511 [11] N. Freed and N. Borenstein, "Multipurpose internet mail extensions (MIME) part two: Media types,"  
6512 Request for Comments 2046, Internet Engineering Task Force, Nov. 1996.
- 6513 [12] D. Eastlake, S. Crocker, and J. Schiller, "Randomness recommendations for security," Request for  
6514 Comments 1750, Internet Engineering Task Force, Dec. 1994.
- 6515 [13] J. Rosenberg and H. Schulzrinne, "An offer/answer model with SDP," Internet Draft, Internet Engi-  
6516 neering Task Force, Jan. 2002. Work in progress.

- 6517 [14] J. Postel, "User datagram protocol," Request for Comments 768, Internet Engineering Task Force,  
6518 Aug. 1980.
- 6519 [15] J. Postel, "DoD standard transmission control protocol," Request for Comments 761, Internet Engi-  
6520 neering Task Force, Jan. 1980.
- 6521 [16] R. Stewart, Q. Xie, K. Morneault, C. Sharp, H. Schwarzbauer, T. Taylor, I. Rytina, M. Kalla, L. Zhang,  
6522 and V. Paxson, "Stream control transmission protocol," Request for Comments 2960, Internet Engi-  
6523 neering Task Force, Oct. 2000.
- 6524 [17] J. Franks, P. Hallam-Baker, J. Hostetler, S. Lawrence, P. Leach, A. Luotonen, and L. Stewart, "HTTP  
6525 authentication: Basic and digest access authentication," Request for Comments 2617, Internet Engi-  
6526 neering Task Force, June 1999.
- 6527 [18] R. Troost, S. Dorner, and K. Moore, "Communicating presentation information in internet messages:  
6528 The content-disposition header field," Request for Comments 2183, Internet Engineering Task Force,  
6529 Aug. 1997.
- 6530 [19] R. Braden and Ed, "Requirements for internet hosts - application and support," Request for Comments  
6531 1123, Internet Engineering Task Force, Oct. 1989.
- 6532 [20] H. Alvestrand, "IETF policy on character sets and languages," Request for Comments 2277, Internet  
6533 Engineering Task Force, Jan. 1998.
- 6534 [21] J. Galvin, S. Murphy, S. Crocker, and N. Freed, "Security multipart for MIME: multipart/signed and  
6535 multipart/encrypted," Request for Comments 1847, Internet Engineering Task Force, Oct. 1995.
- 6536 [22] R. Housley, "Cryptographic message syntax," Request for Comments 2630, Internet Engineering Task  
6537 Force, June 1999.
- 6538 [23] B. Ramsdell and Ed, "S/MIME version 3 message specification," Request for Comments 2633, Internet  
6539 Engineering Task Force, June 1999.
- 6540 [24] T. Dierks and C. Allen, "The TLS protocol version 1.0," Request for Comments 2246, Internet Engi-  
6541 neering Task Force, Jan. 1999.
- 6542 [25] S. Kent and R. Atkinson, "Security architecture for the internet protocol," Request for Comments 2401,  
6543 Internet Engineering Task Force, Nov. 1998.

## 6544 **Non-Normative References**

- 6545 [26] R. Pandya, "Emerging mobile and personal communication systems," *IEEE Communications Maga-*  
6546 *zine*, Vol. 33, pp. 44–52, June 1995.
- 6547 [27] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, "RTP: a transport protocol for real-time  
6548 applications," Request for Comments 1889, Internet Engineering Task Force, Jan. 1996.
- 6549 [28] H. Schulzrinne, A. Rao, and R. Lanphier, "Real time streaming protocol (RTSP)," Request for Com-  
6550 ments 2326, Internet Engineering Task Force, Apr. 1998.

- 6551 [29] F. Cuervo, N. Greene, A. Rayhan, C. Huitema, B. Rosen, and J. Segers, "Megaco protocol version  
6552 1.0," Request for Comments 3015, Internet Engineering Task Force, Nov. 2000.
- 6553 [30] M. Handley, H. Schulzrinne, E. Schooler, and J. Rosenberg, "SIP: session initiation protocol," Request  
6554 for Comments 2543, Internet Engineering Task Force, Mar. 1999.
- 6555 [31] P. Hoffman, L. Masinter, and J. Zawinski, "The mailto URL scheme," Request for Comments 2368,  
6556 Internet Engineering Task Force, July 1998.
- 6557 [32] E. M. Schooler, "A multicast user directory service for synchronous rendezvous," Master's Thesis CS-  
6558 TR-96-18, Department of Computer Science, California Institute of Technology, Pasadena, California,  
6559 Aug. 1996.
- 6560 [33] S. Donovan, "The SIP INFO method," Request for Comments 2976, Internet Engineering Task Force,  
6561 Oct. 2000.
- 6562 [34] R. Rivest, "The MD5 message-digest algorithm," Request for Comments 1321, Internet Engineering  
6563 Task Force, Apr. 1992.
- 6564 [35] F. Dawson and T. Howes, "vcard MIME directory profile," Request for Comments 2426, Internet  
6565 Engineering Task Force, Sept. 1998.
- 6566 [36] G. Good, "The LDAP data interchange format (LDIF) - technical specification," Request for Com-  
6567 ments 2849, Internet Engineering Task Force, June 2000.
- 6568 [37] J. Palme, "Common internet message headers," Request for Comments 2076, Internet Engineering  
6569 Task Force, Feb. 1997.
- 6570 [38] J. Franks, P. Hallam-Baker, J. Hostetler, P. Leach, A. Luotonen, E. Sink, and L. Stewart, "An exten-  
6571 sion to HTTP : Digest access authentication," Request for Comments 2069, Internet Engineering Task  
6572 Force, Jan. 1997.
- 6573 [39] A. Johnston, S. Donovan, R. Sparks, C. Cunningham, D. Willis, J. Rosenberg, K. Summers, and  
6574 H. Schulzrinne, "SIP telephony call flow examples," Internet Draft, Internet Engineering Task Force,  
6575 Apr. 2001. Work in progress.
- 6576 [40] H. Schulzrinne, "RTP profile for audio and video conferences with minimal control," Request for  
6577 Comments 1890, Internet Engineering Task Force, Jan. 1996.
- 6578 [41] E. M. Schooler, "Case study: multimedia conference control in a packet-switched teleconferencing  
6579 system," *Journal of Internetworking: Research and Experience*, Vol. 4, pp. 99–120, June 1993. ISI  
6580 reprint series ISI/RS-93-359.
- 6581 [42] H. Schulzrinne, "Personal mobility for multimedia services in the Internet," in *European Workshop on*  
6582 *Interactive Distributed Multimedia Systems and Services (IDMS)*, (Berlin, Germany), Mar. 1996.

## 6583 **A Table of Timer Values**

6584 Table 4 summarizes the meaning and defaults of the various timers used by this specification.

Timer	Value	Section	Meaning
T1	500ms default	Section 17.1.1.1	RTT Estimate
T2	4s	Section 17.1.2.2	The maximum retransmit interval for non-INVITE requests and INVITE responses
T4	5s	Section 17.1.2.2	Maximum duration a message will remain in the network
Timer A	initially T1	Section 17.1.1.2	INVITE request retransmit interval, for UDP only
Timer B	64*T1	Section 17.1.1.2	INVITE transaction timeout timer
Timer C	> 3min	Section Section 16.6 bullet 11	proxy INVITE transaction timeout
Timer D	> 32s for UDP 0s for TCP/SCTP	Section 17.1.1.2	Wait time for response retransmits
Timer E	initially T1	Section 17.1.2.2	non-INVITE request retransmit interval, UDP only
Timer F	64*T1	Section 17.1.2.2	non-INVITE transaction timeout timer
Timer G	initially T1	Section 17.2.1	INVITE response retransmit interval
Timer H	64*T1	Section 17.2.1	Wait time for ACK receipt
Timer I	T4 for UDP 0s for TCP/SCTP	Section 17.2.1	Wait time for ACK retransmits
Timer J	64*T1 for UDP 0s for TCP/SCTP	Section 17.2.2	Wait time for non-INVITE request retransmits
Timer K	T4 for UDP 0s for TCP/SCTP	Section 17.1.2.2	Wait time for response retransmits

Table 4: Summary of timers

## 6585 Full Copyright Statement

6586 Copyright (c) The Internet Society (2002). All Rights Reserved.

6587 This document and translations of it may be copied and furnished to others, and derivative works that  
6588 comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and  
6589 distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and  
6590 this paragraph are included on all such copies and derivative works. However, this document itself may not  
6591 be modified in any way, such as by removing the copyright notice or references to the Internet Society or

6592 other Internet organizations, except as needed for the purpose of developing Internet standards in which case  
6593 the procedures for copyrights defined in the Internet Standards process must be followed, or as required to  
6594 translate it into languages other than English.

6595 The limited permissions granted above are perpetual and will not be revoked by the Internet Society or  
6596 its successors or assigns.

6597 This document and the information contained herein is provided on an "AS IS" basis and THE IN-  
6598 TERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WAR-  
6599 RANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT  
6600 THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED  
6601 WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.