

SIP: Session Initiation Protocol

Status of this Memo

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

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

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

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

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

Copyright Notice

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

Abstract

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

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

Contents

1	Introduction	8
2	Overview of SIP Functionality	8
3	Terminology	9
4	Overview of Operation	9
5	Structure of the Protocol	15
6	Definitions	16
7	SIP Messages	20
7.1	Requests	21
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	8	General User Agent Behavior	25
43	8.1	UAC Behavior	26
44	8.1.1	Generating the Request	26
45	8.1.2	Sending the Request	29
46	8.1.3	Loose Routing Policies	30
47	8.1.4	Processing Responses	31
48	8.2	UAS Behavior	33
49	8.2.1	Method Inspection	33
50	8.2.2	Header Inspection	33
51	8.2.3	Content Processing	35
52	8.2.4	Applying Extensions	35
53	8.2.5	Processing the Request	35
54	8.2.6	Generating the Response	35
55	8.2.7	Stateless UAS Behavior	36
56	8.3	Redirect Servers	37
57	9	Canceling a Request	38
58	9.1	Client Behavior	38
59	9.2	Server Behavior	39
60	10	Registrations	39
61	10.1	Overview	39
62	10.2	Constructing the REGISTER Request	40
63	10.2.1	Adding Bindings	41
64	10.2.2	Removing Bindings	42
65	10.2.3	Fetching Bindings	43
66	10.2.4	Refreshing Bindings	43
67	10.2.5	Setting the Internal Clock	43
68	10.2.6	Discovering a Registrar	43
69	10.2.7	Transmitting a Request	43
70	10.2.8	Error Responses	44
71	10.3	Processing REGISTER Requests	44
72	11	Querying for Capabilities	46
73	11.1	Construction of OPTIONS Request	46
74	11.2	Processing of OPTIONS Request	47

75	12 Dialogs	48
76	12.1 Creation of a Dialog	48
77	12.1.1 UAS behavior	48
78	12.1.2 UAC behavior	49
79	12.2 Requests within a Dialog	50
80	12.2.1 UAC Behavior	50
81	12.2.2 UAS behavior	51
82	12.3 Termination of a Dialog	53
83	13 Initiating a Session	53
84	13.1 Overview	53
85	13.2 Caller Processing	53
86	13.2.1 Creating the Initial INVITE	53
87	13.2.2 Processing INVITE Responses	54
88	13.3 Callee Processing	56
89	13.3.1 Processing of the INVITE	56
90	14 Modifying an Existing Session	58
91	14.1 UAC Behavior	58
92	14.2 UAS Behavior	59
93	15 Terminating a Session	60
94	15.1 Terminating a Dialog with a BYE Request	60
95	15.1.1 UAC Behavior	60
96	15.1.2 UAS Behavior	61
97	16 Proxy Behavior	61
98	16.1 Overview	61
99	16.2 Stateful Proxy	62
100	16.3 Request Validation	62
101	16.4 Making a Routing Decision	65
102	16.5 Request Processing	67
103	16.6 Response Processing	71
104	16.7 Processing Timer C	76
105	16.8 Handling Transport Errors	76
106	16.9 CANCEL Processing	76
107	16.10 Stateless Proxy	77
108	16.11 Record-Route Example	78
109	17 Transactions	79
110	17.1 Client Transaction	80
111	17.1.1 INVITE Client Transaction	81
112	17.1.2 non-INVITE Client Transaction	84
113	17.1.3 Matching Responses to Client Transactions	86
114	17.1.4 Handling Transport Errors	87
115	17.2 Server Transaction	87

116	17.2.1 INVITE Server Transaction	87
117	17.2.2 non-INVITE Server Transaction	89
118	17.2.3 Matching Requests to Server Transactions	89
119	17.2.4 Handling Transport Errors	91
120	17.3 RTT Estimation	92
121	18 Reliability of Provisional Responses	92
122	18.1 UAS Behavior	93
123	18.2 UAC Behavior	94
124	19 Transport	95
125	19.1 Clients	95
126	19.1.1 Sending Requests	95
127	19.1.2 Receiving Responses	96
128	19.2 Servers	97
129	19.2.1 Receiving Requests	97
130	19.2.2 Sending Responses	97
131	19.3 Framing	98
132	19.4 Error Handling	98
133	20 Usage of HTTP Authentication	98
134	20.1 Framework	99
135	20.2 User-to-User Authentication	100
136	20.3 Proxy to User Authentication	101
137	20.4 The Digest Authentication Scheme	102
138	21 S/MIME	104
139	21.1 S/MIME Certificates	104
140	21.2 S/MIME Key Exchange	104
141	21.3 Securing MIME bodies	106
142	21.4 Tunneling SIP in MIME	106
143	21.4.1 Tunneling Integrity and Authentication	106
144	21.4.2 Tunneling Encryption	108
145	22 Security Considerations	109
146	22.1 Threat Models	109
147	22.1.1 Registration Hijacking	110
148	22.1.2 Impersonating a Server	110
149	22.1.3 Tampering with Message Bodies	110
150	22.1.4 Tearing Down Sessions	111
151	22.1.5 Denial of Service and Amplification	111
152	22.2 Security Mechanisms	112
153	22.2.1 Transport and Network Layer Security	112
154	22.2.2 HTTP Authentication	113
155	22.2.3 S/MIME	113
156	22.3 Implementing Security Mechanisms	114

157	22.3.1 Requirements for Implementers of SIP	114
158	22.3.2 Security Solutions	114
159	22.4 Limitations	117
160	22.4.1 HTTP Digest	117
161	22.4.2 S/MIME	118
162	22.4.3 TLS	119
163	22.5 Privacy	119
164	23 Common Message Components	119
165	23.1 SIP Uniform Resource Indicators	119
166	23.1.1 SIP URI Components	120
167	23.1.2 Character Escaping Requirements	122
168	23.1.3 Example SIP URIs	123
169	23.1.4 SIP URI Comparison	123
170	23.1.5 Forming Requests from a SIP URI	124
171	23.1.6 Relating SIP URIs and tel URLs	125
172	23.2 Option Tags	126
173	23.3 Tags	127
174	24 Header Fields	127
175	24.1 Accept	130
176	24.2 Accept-Encoding	130
177	24.3 Accept-Language	131
178	24.4 Alert-Info	131
179	24.5 Allow	131
180	24.6 Authentication-Info	132
181	24.7 Authorization	132
182	24.8 Call-ID	132
183	24.9 Call-Info	132
184	24.10 Contact	133
185	24.11 Content-Disposition	133
186	24.12 Content-Encoding	134
187	24.13 Content-Language	134
188	24.14 Content-Length	134
189	24.15 Content-Type	135
190	24.16 CSeq	135
191	24.17 Date	135
192	24.18 Error-Info	135
193	24.19 Expires	136
194	24.20 From	136
195	24.21 In-Reply-To	137
196	24.22 Max-Forwards	137
197	24.23 Min-Expires	137
198	24.24 MIME-Version	137
199	24.25 Organization	137

200	24.26 Priority	138
201	24.27 Proxy-Authenticate	138
202	24.28 Proxy-Authorization	138
203	24.29 Proxy-Require	139
204	24.30 RACK	139
205	24.31 Record-Route	139
206	24.32 Reply-To	139
207	24.33 Require	140
208	24.34 Retry-After	140
209	24.35 Route	140
210	24.36 RSeq	141
211	24.37 Server	141
212	24.38 Subject	141
213	24.39 Supported	141
214	24.40 Timestamp	141
215	24.41 To	142
216	24.42 Unsupported	142
217	24.43 User-Agent	142
218	24.44 Via	142
219	24.45 Warning	143
220	24.46 WWW-Authenticate	144
221	25 Response Codes	144
222	25.1 Provisional 1xx	145
223	25.1.1 100 Trying	145
224	25.1.2 180 Ringing	145
225	25.1.3 181 Call Is Being Forwarded	145
226	25.1.4 182 Queued	145
227	25.1.5 183 Session Progress	145
228	25.2 Successful 2xx	145
229	25.2.1 200 OK	145
230	25.3 Redirection 3xx	146
231	25.3.1 300 Multiple Choices	146
232	25.3.2 301 Moved Permanently	146
233	25.3.3 302 Moved Temporarily	146
234	25.3.4 305 Use Proxy	146
235	25.3.5 380 Alternative Service	147
236	25.4 Request Failure 4xx	147
237	25.4.1 400 Bad Request	147
238	25.4.2 401 Unauthorized	147
239	25.4.3 402 Payment Required	147
240	25.4.4 403 Forbidden	147
241	25.4.5 404 Not Found	147
242	25.4.6 405 Method Not Allowed	147
243	25.4.7 406 Not Acceptable	147

244	25.4.8	407 Proxy Authentication Required	148
245	25.4.9	408 Request Timeout	148
246	25.4.10	410 Gone	148
247	25.4.11	413 Request Entity Too Large	148
248	25.4.12	414 Request-URI Too Long	148
249	25.4.13	415 Unsupported Media Type	148
250	25.4.14	416 Unsupported URI Scheme	148
251	25.4.15	420 Bad Extension	148
252	25.4.16	421 Extension Required	149
253	25.4.17	423 Registration Too Brief	149
254	25.4.18	480 Temporarily Unavailable	149
255	25.4.19	481 Call/Transaction Does Not Exist	149
256	25.4.20	482 Loop Detected	149
257	25.4.21	483 Too Many Hops	149
258	25.4.22	484 Address Incomplete	149
259	25.4.23	485 Ambiguous	150
260	25.4.24	486 Busy Here	150
261	25.4.25	487 Request Terminated	150
262	25.4.26	488 Not Acceptable Here	150
263	25.4.27	491 Request Pending	150
264	25.4.28	493 Undecipherable	150
265	25.5	Server Failure 5xx	151
266	25.5.1	500 Server Internal Error	151
267	25.5.2	501 Not Implemented	151
268	25.5.3	502 Bad Gateway	151
269	25.5.4	503 Service Unavailable	151
270	25.5.5	504 Server Time-out	151
271	25.5.6	505 Version Not Supported	151
272	25.5.7	513 Message Too Large	152
273	25.6	Global Failures 6xx	152
274	25.6.1	600 Busy Everywhere	152
275	25.6.2	603 Decline	152
276	25.6.3	604 Does Not Exist Anywhere	152
277	25.6.4	606 Not Acceptable	152
278	26	Examples	153
279	26.1	Registration	153
280	26.2	Session Setup	154
281	27	Augmented BNF for the SIP Protocol	159
282	27.1	Basic Rules	160
283	28	IANA Considerations	174
284	28.1	Option Tags	174
285	28.1.1	Registration of 100rel	175

286	28.2 Warn-Codes	175
287	28.3 Header Field Names	175
288	28.4 Method and Response Codes	176
289	29 Changes Made in Version 00	176
290	30 Changes Made in Version 01	181
291	31 Changes Made in Version 02	181
292	32 Changes Made in Version 03	183
293	33 Changes Made in Version 04	185
294	34 Changes Made in Version 05	186
295	35 Changes Made in Version 06	189
296	36 Acknowledgments	198
297	37 Authors' Addresses	198

298 **1 Introduction**

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

311 **2 Overview of SIP Functionality**

312 The Session Initiation Protocol (SIP) is an application-layer control protocol that can establish, modify, and
 313 terminate multimedia sessions (conferences) such as Internet telephony calls. SIP can also invite participants
 314 to already existing sessions, such as multicast conferences. Media can be added to (and removed from)
 315 an existing session. SIP transparently supports name mapping and redirection services, which supports
 316 *personal mobility* [1, p. 44] - users can maintain a single externally visible identifier (SIP URI) regardless
 317 of their network location.

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

- 319 **User location:** determination of the end system to be used for communication;
- 320 **User availability:** determination of the willingness of the called party to engage in communications;
- 321 **User capabilities:** determination of the media and media parameters to be used;
- 322 **Session setup:** “ringing”, establishment of session parameters at both called and calling party;
- 323 **Session management:** including transfer and termination of sessions, modifying session parameters, and
324 invoking services.

325 SIP is not a vertically integrated communications system. SIP is rather a component that can be used with
326 other IETF protocols to build a complete multimedia architecture. Typically, these architectures will include
327 protocols such as the real-time transport protocol (RTP) (RFC 1889 [2]) for transporting real-time data and
328 providing QoS feedback, the real-time streaming protocol (RTSP) (RFC 2326 [3]) for controlling delivery of
329 streaming media, the Media Gateway Control Protocol (MEGACO) (RFC 3015 [4]) for controlling gateways
330 to the Public Switched Telephone Network (PSTN), and the session description protocol (SDP) (RFC 2327
331 [5]) for describing multimedia sessions. Therefore, SIP should be used in conjunction with other protocols
332 in order to provide complete services to the users. However, the basic functionality and operation of SIP
333 does not depend on any of these protocols.

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

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

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

347 SIP works with both IPv4 and IPv6.

348 **3 Terminology**

349 In this document, the key words “MUST”, “MUST NOT”, “REQUIRED”, “SHALL”, “SHALL NOT”, “SHOULD”,
350 “SHOULD NOT”, “RECOMMENDED”, “MAY”, and “OPTIONAL” are to be interpreted as described in RFC
351 2119 [6] and indicate requirement levels for compliant SIP implementations.

352 **4 Overview of Operation**

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

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

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

364 Alice “calls” Bob using his SIP identity, a type of Uniform Resource Identifier (URI) called a SIP URI
365 and defined in Section 23.1. It has a similar form to an email address, typically containing a username and
366 a host name. In this case, it is sip:bob@biloxi.com, where biloxi.com is the domain of Bob’s SIP service
367 provider (which can be an enterprise, retail provider, etc). Alice also has a SIP URI of sip:alice@atlanta.com.
368 Alice might have typed in Bob’s URI or perhaps clicked on a hyperlink or an entry in an address book.

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

```
377 INVITE sip:bob@biloxi.com SIP/2.0  
378 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bK776asdhd  
379 To: Bob <sip:bob@biloxi.com>  
380 From: Alice <sip:alice@atlanta.com>;tag=1928301774  
381 Call-ID: a84b4c76e66710  
382 CSeq: 314159 INVITE  
383 Contact: <sip:alice@pc33.atlanta.com>  
384 Content-Type: application/sdp  
385 Content-Length: 142  
386  
387 (Alice’s SDP not shown)
```

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

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

393 To contains a display name (Bob) and a SIP URI (sip:bob@biloxi.com) towards which the request was
394 originally directed. Display names are described in RFC 2822 [7].

395 From also contains a display name (Alice) and a SIP URI (sip:alice@atlanta.com) that indicate the
396 originator of the request. This header field also has a tag parameter containing a pseudorandom string
397 (1928301774) that was added to the URI by the softphone. It is used for identification purposes.




Figure 1: SIP session setup example with SIP trapezoid

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

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

403 **Contact** contains a SIP URI that represents a direct route to reach or contact Alice, usually composed
404 of a username at an FQDN. While a FQDN is preferred, many end systems do not have registered domain
405 names, so IP addresses are permitted. While the **Via** header field tells other elements where to send the
406 response, the **Contact** header field tells other elements where to send future requests for this dialog.

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

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

409 The complete set of SIP header fields is defined in Section 24.

410 The details of the session, type of media, codec, sampling rate, etc. are not described using SIP. Rather,
411 the body of a SIP message contains a description of the session, encoded in some other protocol format.
412 One such format is Session Description Protocol (SDP) [5]. This SDP message (not shown in the example)
413 is carried by the SIP message in a way that is analogous to a document attachment being carried by an email
414 message, or a web page being carried in an HTTP message.

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

419 The atlanta.com SIP server is a type of SIP server known as a proxy server. A proxy server receives
420 SIP requests and forwards them on behalf of the requestor. In this example, the proxy server receives the
421 **INVITE** request and sends a 100 (Trying) response back to Alice's softphone. The 100 (Trying) response
422 indicates that the **INVITE** has been received and that the proxy is working on her behalf to route the **INVITE**
423 to the destination. Responses in SIP use a three-digit code followed by a descriptive phrase. This response
424 contains the same **To**, **From**, **Call-ID**, and **CSeq** as the **INVITE**, which allows Alice's softphone to correlate
425 this response to the sent **INVITE**. The atlanta.com proxy server locates the proxy server at biloxi.com,
426 possibly by performing a particular type of DNS (Domain Name Service) lookup to find the SIP server
427 that serves the biloxi.com domain. This is described in [8]. As a result, it obtains the IP address of the
428 biloxi.com proxy server and forwards, or proxies, the **INVITE** request there. Before forwarding the request,
429 the atlanta.com proxy server adds an additional **Via** header field that contains its own IP address (the **INVITE**
430 already contains Alice's IP address in the first **Via**). The biloxi.com proxy server receives the **INVITE** and
431 responds with a 100 (Trying) response back to the Atlanta.com proxy server to indicate that it has received
432 the **INVITE** and is processing the request. The proxy server consults a database, generically called a location
433 service, that contains the current IP address of Bob. (We shall see in the next section how this database can
434 be populated.) The biloxi.com proxy server adds another **Via** header with its own IP address to the **INVITE**
435 and proxies it to Bob's SIP phone.

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

443 responses to the INVITE.

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

446 In this example, Bob decides to answer the call. When he picks up the handset, his SIP phone sends a
447 200 (OK) response to indicate that the call has been answered. The 200 (OK) contains a message body with
448 the SDP media description of the type of session that Bob is willing to establish with Alice. As a result, there
449 is a two-phase exchange of SDP messages; Alice sent one to Bob, and Bob sent one back to Alice. This
450 two-phase exchange provides basic negotiation capabilities and is based on a simple offer/answer model of
451 SDP exchange. If Bob did not wish to answer the call or was busy on another call, an error response would
452 have been sent instead of the 200 (OK), which would have resulted in no media session being established.
453 The complete list of SIP response codes is in Section 25. The 200 (OK) (message F9 in Figure 1) might
454 look like this as Bob sends it out:

```
455 SIP/2.0 200 OK
456 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bK776asdhds
457 Via: SIP/2.0/UDP bigbox3.site3.atlanta.com;branch=z9hG4bK77ef4c2312983.1
458 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
459 To: Bob <sip:bob@biloxi.com>;tag=a6c85cf
460 From: Alice <sip:alice@atlanta.com>;tag=1928301774
461 Call-ID: a84b4c76e66710
462 CSeq: 314159 INVITE
463 Contact: <sip:bob@192.0.2.8>
464 Content-Type: application/sdp
465 Content-Length: 131
466
467 (Bob's SDP not shown)
```

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

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

481 In this case, the 200 (OK) is routed back through the two proxies and is received by Alice's softphone
482 which then stops the ringback tone and indicates that the call has been answered. Finally, an acknowledge-
483 ment message, ACK, is sent by Alice to Bob to confirm the reception of the final response (200 (OK)). In this
484 example, the ACK is sent directly from Alice to Bob, bypassing the two proxies. This is because, through
485 the INVITE/200 (OK) exchange, the two SIP user agents have learned each other's IP address through the

486 **Contact** header fields, which was not known when the initial **INVITE** was sent. The lookups performed by
487 the two proxies are no longer needed, so they drop out of the call flow. This completes the **INVITE/200/ACK**
488 three-way handshake used to establish SIP sessions and is the end of the transaction. Full details on session
489 setup are in Section 13.

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

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

501 At the end of the call, Bob disconnects (hangs up) first, and generates a **BYE** message. This **BYE** is
502 routed directly to Alice's softphone, again bypassing the proxies. Alice confirms receipt of the **BYE** with a
503 200 (OK) response, which terminates the session and the **BYE** transaction. No **ACK** is sent - an **ACK** is only
504 sent in response to a response to an **INVITE** request. The reasons for this special handling for **INVITE** will
505 be discussed later, but relate to the reliability mechanisms in SIP, the length of time it can take for a ringing
506 phone to be answered, and forking. For this reason, request handling in SIP is often classified as either
507 **INVITE** or non-**INVITE**, referring to all other methods besides **INVITE**. Full details on session termination
508 are in Section 15.

509 Full details of all the messages shown in the example of Figure 1 are shown in Section 26.2.

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

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

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

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

535 Finally, it is important to note that in SIP, registration is used for routing incoming SIP requests and
536 has no role in authorizing outgoing requests. Authorization and authentication are handled in SIP either
537 on a request-by-request, challenge/response mechanism, or using a lower layer scheme as discussed in
538 Section 22.

539 The complete set of SIP message details for this registration example is in Section 26.1.

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

543 5 Structure of the Protocol

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

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

552 The lowest layer of SIP is its syntax and encoding. Its encoding is specified using a BNF. The complete
553 BNF is specified in Section 27. However, a basic overview of the structure of a SIP message can be found
554 in Section 7. This section provides enough understanding of the format of a SIP message to facilitate
555 understanding the remainder of the protocol.

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

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

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

571 A TU which creates a client transaction can also cancel it. When a client cancels a transaction, it requests
572 that the server stop further processing, revert to the state that existed before the transaction was initiated,
573 and generate a specific error response to that transaction. This is done with a CANCEL request, which

574 constitutes its own transaction, but references the transaction to be cancelled. Cancellation is described in
575 Section 9.

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

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

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

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

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

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

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

606 6 Definitions

607 This specification uses a number of terms to refer to the roles played by participants in SIP communications.
608 The terms and generic syntax of URI and URL are defined in RFC 2396 [9]. The following terms have
609 special significance for SIP.

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

- 615 **Call:** A call is an informal term that refers to a dialog between peers generally set up for the purposes of a
616 multimedia conversation.
- 617 **Call leg:** Another name for a dialog.
- 618 **Call stateful:** A proxy is call stateful if it retains state for a dialog from the initiating INVITE to the termi-
619 nating BYE request. A call stateful proxy is always stateful, but the converse is not true.
- 620 **Client:** A client is any network element that sends SIP requests and receives SIP responses. Clients may or
621 may not interact directly with a human user. *User agent clients* and *proxies* are clients.
- 622 **Conference:** A multimedia session (see below) that contains multiple participants.
- 623 **Dialog:** A dialog is a peer-to-peer SIP relationship between a UAC and UAS that persists for some time.
624 A dialog is established by SIP messages, such as a 2xx response to an INVITE request. A dialog is
625 identified by a call identifier, local address, and remote address. A dialog was formerly known as a
626 call leg in RFC 2543.
- 627 **Downstream:** A direction of message forwarding within a transaction that refers to the direction that re-
628 quests flow from the user agent client to user agent server.
- 629 **Final response:** A response that terminates a SIP transaction, as opposed to a *provisional response* that
630 does not. All 2xx, 3xx, 4xx, 5xx and 6xx responses are final.
- 631 **Header:** A header is a component of a sip message that conveys information about the message. It is
632 structured as a header name, followed by a colon, followed by its value.
- 633 **Home Domain:** The domain providing service to a SIP user. Typically, this is the domain present in the
634 URI in the address-of-record of a registration.
- 635 **Informational Response:** Same as a provisional response.
- 636 **Initiator, calling party, caller:** The party initiating a session (and dialog) with an INVITE request. A caller
637 retains this role from the time it sends the initial INVITE which established a dialog, until the termi-
638 nation of that dialog.
- 639 **Invitation:** An INVITE request.
- 640 **Invitee, invited user, called party, callee:** The party that receives an INVITE request for the purposes of
641 establishing a new session. A callee retains this role from the time it receives the INVITE until the
642 termination of the dialog established by that INVITE.
- 643 **Location service:** A location service is used by a SIP redirect or proxy server to obtain information about
644 a callee's possible location(s). It contains a list of bindings of address-of-record keys to zero or more
645 contact addresses. The bindings can be created and removed in many ways; this specification defines
646 a REGISTER method that updates the bindings.
- 647 **Loop:** A request that arrives at a proxy, is forwarded, and later arrives back at the same proxy. When it
648 arrives the second time, its Request-URI is identical to the first time, and other headers that affect
649 proxy operation are unchanged, so that the proxy would make the same processing decision on the
650 request it made the first time around. Looped requests are errors, and the procedures for detecting
651 them and handling them are described by the protocol.

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

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

656 **Outbound proxy:** A *proxy* that receives all requests from a client, even though it is not the server resolved
657 by the **Request-URI**. The outbound proxy sends these requests, after any local processing, to the
658 address indicated in the **Request-URI**, or to another outbound proxy. Typically, a UA is manually
659 configured with its outbound proxy, or can learn it through auto-configuration protocols.

660 **Parallel search:** In a parallel search, a proxy issues several requests to possible user locations upon receiv-
661 ing an incoming request. Rather than issuing one request and then waiting for the final response before
662 issuing the next request as in a *sequential search*, a parallel search issues requests without waiting for
663 the result of previous requests.

664 **Provisional response:** A response used by the server to indicate progress, but that does not terminate a SIP
665 transaction. 1xx responses are provisional, other responses are considered *final*. Normally, provisional
666 responses are not sent reliably. A provisional response that is sent reliably is referred to as a *reliable*
667 *provisional response*.

668 **Proxy, proxy server:** An intermediary entity that acts as both a server and a client for the purpose of making
669 requests on behalf of other clients. A proxy server primarily plays the role of routing, which means
670 its job is to ensure that a request is passed on to another entity “closer” to the targeted user. Proxies
671 are also useful for enforcing policy (for example, making sure a user is allowed to make a call). A
672 proxy interprets, and, if necessary, rewrites specific parts of a request message before forwarding it.

673 **Recursion:** A client recurses on a 3xx response when it generates a new request to the URIs in the **Contact**
674 headers in the response.

675 **Redirect Server:** A redirect server is a server that generates 3xx responses to requests it receives, directing
676 the client to contact an alternate URI.

677 **Registrar:** A registrar is a server that accepts REGISTER requests, and places the information it receives
678 in those requests into the location service for the domain it handles.

679 **Regular Transaction:** A regular transaction is any transaction with a method other than INVITE, ACK, or
680 CANCEL.

681 **Reliable Provisional Response:** A provisional response that is sent reliably from the UAS to UAC.

682 **Request:** A SIP message sent from a client to a server, for the purpose of invoking a particular operation.

683 **Response:** A SIP message sent from a server to a client, for indicating the status of a request sent from the
684 client to the server.

685 **Ringback:** Ringback is the signaling tone produced by the calling party’s application indicating that a
686 called party is being alerted (ringing).

687 **Route Refresh Request:** A route refresh request sent within a dialog is defined as a request that can modify
688 the *route set* of the dialog.

689 **Server:** A server is a network element that receives requests in order to service them and sends back re-
690 sponses to those requests. Examples of servers are proxies, user agent servers, redirect servers, and
691 registrars.

692 **Sequential search:** In a sequential search, a proxy server attempts each contact address in sequence, pro-
693 ceeding to the next one only after the previous has generated a non-2xx final response.

694 **Session:** From the SDP specification: “A multimedia session is a set of multimedia senders and receivers
695 and the data streams flowing from senders to receivers. A multimedia conference is an example of a
696 multimedia session.” (RFC 2327 [5]) (A session as defined for SDP can comprise one or more RTP
697 sessions.) As defined, a callee can be invited several times, by different calls, to the same session.
698 If SDP is used, a session is defined by the concatenation of the *user name*, *session id*, *network type*,
699 *address type*, and *address* elements in the origin field.

700 **(SIP) transaction:** A SIP transaction occurs between a client and a server and comprises all messages from
701 the first request sent from the client to the server up to a final (non-1xx) response sent from the server
702 to the client, and the ACK for the response in the case the response was a non-2xx. The ACK for a
703 2xx response is a separate transaction.

704 **Spiral:** A spiral is a SIP request that is routed to a proxy, forwarded onwards, and arrives once again at that
705 proxy, but this time, differs in a way that will result in a different processing decision than the original
706 request. Typically, this means that the request’s Request-URI differs from its previous arrival. A
707 spiral is not an error condition, unlike a loop. A typical cause for this is call forwarding. A user calls
708 joe@example.com. The example.com proxy forwards it to Joe’s PC, which in turn, forwards it to
709 bob@example.com. This request is proxied back to the example.com proxy. However, this is not a
710 loop. Since the request is targeted at a different user, it is considered a spiral, and is a valid condition.
711

712 **Stateful proxy:** A logical entity that maintains the client and server transaction state machines defined by
713 this specification during the processing of a request. Also known as a transaction stateful proxy. The
714 behavior of a stateful proxy is further defined in Section 16. A stateful proxy is not the same as a call
715 stateful proxy.

716 **Stateless proxy:** A logical entity that does not maintain the client or server transaction state machines
717 defined in this specification when it processes requests. A stateless proxy forwards every request it
718 receives downstream and every response it receives upstream.

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

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

723 **URL-encoded:** A character string encoded according to RFC 1738, Section 2.2 [10].

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

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

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

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

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

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

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

746 7 SIP Messages

747 SIP is a text-based protocol and uses the ISO 10646 character set in UTF-8 encoding (RFC 2279 [11]).

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

749 Both **Request** (section 7.1) and **Response** (section 7.2) messages use the basic format of RFC 2822
750 [7], even though the syntax differs in character set and syntax specifics. (SIP allows header fields that would
751 not be valid RFC 2822 header fields, for example.) Both types of messages consist of a **start-line**, one or
752 more header fields (also known as “headers”), an empty line indicating the end of the header fields, and an
753 optional message-body.

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

754
755 The start-line, each message-header line, and the empty line **MUST** be terminated by a carriage-return
756 line-feed sequence (CRLF). Note that the empty line **MUST** be present even if the message-body is not.

757 Except for the above difference in character sets, much of SIP’s message and header field syntax is
758 identical to HTTP/1.1. Rather than repeating the syntax and semantics here, we use [HX.Y] to refer to
759 Section X.Y of the current HTTP/1.1 specification (RFC 2616 [12]).

760 However, SIP is not an extension of HTTP.

761 7.1 Requests

762 SIP requests are distinguished by having a Request-Line for a start-line. A Request-Line contains a
763 method name, a Request-URI, and the protocol version separated by a single space (SP) character. The
764 Request-Line ends with CRLF. No CR or LF are allowed except in the end-of-line CRLF sequence. No
765 LWS is allowed in any of the elements.

766 Method Request-URI SIP-Version

767 **Method:** This specification defines seven methods: REGISTER for registering contact information, IN-
768 VITE, ACK, PRACK and CANCEL for setting up sessions, BYE for terminating sessions and OP-
769 TIONS for querying servers about their capabilities. SIP extensions, documented in standards track
770 RFCs, may define additional methods.

771 **Request-URI:** The Request-URI is a SIP URI as described in Section 23.1 or a general URI (RFC 2396 [9]).
772 It indicates the user or service to which this request is being addressed. The Request-URI MUST NOT
773 contain unescaped spaces or control characters and MUST NOT be enclosed in "<>".

774 SIP elements MAY support Request-URIs with schemes other than "sip", for example the "tel" URI
775 scheme of RFC 2806 [13]. SIP elements MAY translate non-SIP URIs using any mechanism at their
776 disposal, resulting in either a SIP URI or some other scheme.

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

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

784 7.2 Responses

785 SIP responses are distinguished from requests by having a Status-Line as their start-line. A Status-Line
786 consists of the protocol version followed by a numeric Status-Code and its associated textual phrase, with
787 each element separated by a single SP character. No CR or LF is allowed except in the final CRLF
788 sequence.

789 SIP-version Status-Code Reason-Phrase

790 The Status-Code is a 3-digit integer result code that indicates the outcome of an attempt to understand
791 and satisfy a request. The Reason-Phrase is intended to give a short textual description of the Status-
792 Code. The Status-Code is intended for use by automata, whereas the Reason-Phrase is intended for
793 the human user. A client is not required to examine or display the Reason-Phrase. While this specifica-
794 tion suggests specific wording for the reason phrase, implementations MAY choose other text, e.g., in the
795 language indicated in the Accept-Language header field of the request.

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

800 **1xx:** Provisional – request received, continuing to process the request;

801 **2xx:** Success – the action was successfully received, understood, and accepted;

802 **3xx:** Redirection – further action needs to be taken in order to complete the request;

803 **4xx:** Client Error – the request contains bad syntax or cannot be fulfilled at this server;

804 **5xx:** Server Error – the server failed to fulfill an apparently valid request;

805 **6xx:** Global Failure – the request cannot be fulfilled at any server.

806 Section 25 defines these classes and describes the individual codes.

807 7.3 Header Fields

808 SIP header fields are similar to HTTP header fields in both syntax and semantics. In particular, SIP header
809 fields follow the [H4.2] definitions of syntax for **message-header**, the rules for extending header fields
810 over multiple lines, the use of multiple message-header fields with the same field-name, and the rules re-
811 garding ordering of header fields.

812 7.3.1 Header Field Format

813 Header fields follow the same generic header format as that given in Section 2.2 of RFC 2822 [7]. Each
814 header field consists of a field name followed by a colon (“:”) and the field value.

815 `field-name: field-value`

816 The formal grammar for a **message-header** specified in Section 27 allows for an arbitrary amount of
817 whitespace on either side of the colon; however, implementations should avoid spaces between the field
818 name and the colon and use a single space (SP) between the colon and the **field-value**. Thus,

```
819 Subject:          lunch
820 Subject      :    lunch
821 Subject      :lunch
822 Subject: lunch
```

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

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

827 Subject: I know you're there, pick up the phone and talk to me!
828 Subject: I know you're there,
829 pick up the phone
830 and talk to me!

831 The relative order of header fields with different field names is not significant. However, it is RECOM-
832 MENDED that headers which are needed for proxy processing (Via, Route, Record-Route, Proxy-Require,
833 Max-Forwards, and Proxy-Authorization, for example) appear towards the top of the message, to facil-
834 itate rapid parsing. The relative order of header fields with the same field name is important. Multiple
835 header fields with the same field-name MAY be present in a message if and only if the entire field-value for
836 that header field is defined as a comma-separated list (that is, #(values)). It MUST be possible to combine
837 the multiple header fields into one "field-name: field-value" pair, without changing the semantics of the
838 message, by appending each subsequent field-value to the first, each separated by a comma.

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

841 The following groups of header fields are valid and equivalent:

842 Route: <sip:alice@atlanta.com>
843 Subject: Lunch
844 Route: <sip:bob@biloxi.com>
845 Route: <sip:carol@chicago.com>
846
847 Route: <sip:alice@atlanta.com>, <sip:bob@biloxi.com>
848 Route: <sip:carol@chicago.com>
849 Subject: Lunch
850
851 Subject: Lunch
852 Route: <sip:alice@atlanta.com>, <sip:bob@biloxi.com>, <sip:carol@chicago.com>

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

854 Route: <sip:alice@atlanta.com>
855 Route: <sip:bob@biloxi.com>
856 Route: <sip:carol@chicago.com>
857
858 Route: <sip:bob@biloxi.com>
859 Route: <sip:alice@atlanta.com>
860 Route: <sip:carol@chicago.com>
861
862 Route: <sip:alice@atlanta.com>, <sip:carol@chicago.com>, <sip:bob@biloxi.com>

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

867 field-name: field-value *(;parameter-name=parameter-value)

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

870 All new header fields MUST follow this generic format unless they have been inherited from other RFC
871 2822-like specifications.

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

876 For example,

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

878 is equivalent to

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

880 and

881 Content-Disposition: session;handling=optional

882 is equivalent to

883 content-disposition: Session;HANDLING=OPTIONAL

884 The following two header fields are not equivalent:

885 Warning: 370 devnull "Choose a bigger pipe"

886 Warning: 370 devnull "CHOOSE A BIGGER PIPE"

887 7.3.2 Header Field Classification

888 Some header fields only make sense in requests or responses. These are called request header fields and
889 response header fields, respectively. If a header appears in a message not matching its category (such as a
890 request header field in a response), it MUST be ignored. Section 24 defines the classification of each header
891 field.

892 7.3.3 Compact Form

893 SIP provides a mechanism to represent common header fields in an abbreviated form. This may be useful
894 when messages would otherwise become too large to be carried on the transport available to it (exceeding
895 the maximum transmission unit (MTU) when using UDP, for example). These compact forms are defined
896 in Section 24. A compact form MAY be substituted for the longer form of a header name at any time without
897 changing the semantics of the message. The same type of header field MAY appear in both long and short
898 forms within the same message. Implementations MUST accept both the long and short forms of each header
899 name.

900 7.4 Bodies

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

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

905 7.4.1 Message Body Type

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

910 The “multipart” MIME type defined in RFC 2046 [14] MAY be used within the body of the message.
911 Implementations that send requests containing multipart message bodies MUST send a session description
912 as a non-multipart message body if the remote implementation requests this through an Accept header field
913 that does not contain multipart.

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

915 7.4.2 Message Body Length

916 The body length in bytes is provided by the Content-Length header field. Section 24.14 describes the
917 necessary contents of this header in detail.

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

920 7.5 Framing SIP messages

921 Unlike HTTP, SIP implementations can use UDP or other unreliable datagram protocols. Each such data-
922 gram carries one request or response. See Section 19 on constraints on usage of unreliable transports.

923 Likewise, implementations processing SIP messages over stream-oriented transports MUST ignore any
924 CRLF appearing before the start-line [H4.1]

925 8 General User Agent Behavior

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

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

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

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

940 Security procedures for requests and responses outside of a dialog are described in Section 22. Specif-
941 ically, mechanisms exist for the UAS and UAC to mutually authenticate. A limited set of privacy features
942 are also supported through encryption of bodies using S/MIME.

943 8.1 UAC Behavior

944 This section covers UAC behavior outside of a dialog.

945 8.1.1 Generating the Request

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

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

954 **8.1.1.1 Request-URI** The initial Request-URI of the message SHOULD be set to the value of the URI
955 in the To field. One notable exception is the REGISTER method; behavior for setting the Request-URI of
956 register is given in Section 10.

957 Another exception is the case of pre-existing Route headers; in that case, the procedures of Sec-
958 tion 12.2.1.1 as they pertain to the Request-URI are followed, even though there is no dialog. Pre-existing
959 Route headers are an ordered set of URIs that identify a chain of servers to which outgoing requests from a
960 UAC will be sent. Commonly, they are configured on the user agent by a user or service provider manually,
961 or through some non-SIP mechanism. They are most often used to identify a local outbound proxy server
962 through which a UAC will send all requests, which in turn allows service providers to maintain a common
963 point of policy enforcement for requests.

964 **8.1.1.2 To** The To general-header field first and foremost specifies the desired “logical” recipient of the
965 request, or the address of record of the user or resource that is the target of this request. This may or may
966 not be the ultimate recipient of the request. The To header MAY contain a SIP URI, but it may also make use
967 of other URI schemes (the tel URL [13], for example) when appropriate. All SIP implementations MUST
968 support the SIP URI. The To header field allows for a display name.

969 A UAC may learn how to populate the To header field for a particular request in a number of ways.
970 Usually the user will suggest the To header field through a human interface, perhaps inputting the URI
971 manually or selecting it from some sort of address book. Frequently, the user will not enter a complete URI,
972 but rather, a string of digits or letters (i.e., “bob”). It is at the discretion of the UA to choose how to interpret
973 this input. Using it to form the user part of a SIP URL implies that the UA wishes the name to be resolved in
974 the domain the right hand side (RHS) of the at-sign in the SIP URI (i.e., sip:bob@example.com). The RHS
975 will frequently be the home domain of the user, which allows for the home domain to process the outgoing

976 request. This is useful for features like “speed dial” which require interpretation of the user part in the home
977 domain. The tel URL is used when the UA does not wish to specify the domain that should interpret the
978 user input. Rather, each domain that the request passes through would be given that opportunity. As an
979 example, a user in an airport might log in, and send requests through an outbound proxy in the airport. If
980 they enter “411” (this is the phone number for local directory assistance in the United States), that needs to
981 be interpreted and processed by the outbound proxy in the airport, not the user’s home domain. In this case,
982 tel:411 would be the right choice.

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

985 For further information on the **To** header see Section 24.41.

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

987 To: Carol <sip:carol@chicago.com>

988 **8.1.1.3 From** The **From** general-header field indicates the logical identity of the initiator of the request,
989 possibly the user’s address of record. Like the **To** field, it contains a URI and optionally a display name.
990 It is used by SIP elements to determine processing rules to apply to a request (for example, automatic call
991 rejection). As such, it is very important that the **From** URI not contain IP addresses or the FQDN of the host
992 the UA is running on, since these are not logical names.

993 The **From** header field allows for a display name. A UAC **SHOULD** use the display name “Anony-
994 mous”, along with a syntactically correct, but otherwise meaningless URI (like sip:988776a@ahhs.aa), if
995 the identity of the client is to remain hidden.

996 Usually the value that populates the **From** header field in requests generated by a particular user agent
997 is pre-provisioned by the user or by the administrators of the user’s local domain. If a particular user agent
998 is used by multiple users, it might have switchable profiles that include a URI corresponding to the identity
999 of the profiled user. Recipients of requests can authenticate the originator of a request in order to ascertain
1000 that they are who their **From** header field claims they are (see Section 20 for more on authentication).

1001 The **From** field **MUST** contain a new “tag” parameter, chosen by the UAC. See Section 23.3 for details
1002 on choosing a tag.

1003 For further information on the **From** header see Section 24.20.

1004 Examples:

1005 From: "Bob" <sip:bob@biloxi.com> ;tag=a48s

1006 From: sip:+12125551212@server.phone2net.com;tag=887s

1007 From: Anonymous <sip:c8oqz84zk7z@privacy.org>;tag=hyh8

1008 **8.1.1.4 Call-ID** The **Call-ID** general-header field acts as a unique identifier to group together a series of
1009 messages. It **MUST** be the same for all requests and responses sent by either UA in a dialog. It **SHOULD** be
1010 the same in each registration from a UA.

1011 In a new request created by a UAC outside of any dialog, the **Call-ID** header **MUST** be selected by the
1012 UAC as a globally unique identifier over space and time unless overridden by method specific behavior.
1013 All SIP user agents must have a means to guarantee that the **Call-ID** headers they produce will not be
1014 inadvertently generated by any other user agent. Note that when requests are retried after certain failure
1015 responses that solicit an amendment to a request (for example, a challenge for authentication), these retried
1016 requests are not considered new requests, and therefore do not need new **Call-ID** headers; see Section 8.1.4.6.

1017

1018 Use of cryptographically random identifiers [15] in the generation of Call-IDs is RECOMMENDED. Im-
1019 plementations MAY use the form "localid@host". Call-IDs are case-sensitive and are simply compared
1020 byte-by-byte.

1021

Using cryptographically random identifiers provides some protection against session hijacking and reduces the
1022 likelihood of unintentional Call-ID collisions.

1023

No provisioning or human interface is required for the selection of the Call-ID header field value for a
1024 request.

1025

For further information on the Call-ID header see Section 24.8.

1026

Example:

1027

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

1028

8.1.1.5 CSeq The Cseq header serves as a way to identify and order transactions. It consists of a
1029 sequence number and a method. The method MUST match that of the request. For requests outside of a
1030 dialog, the sequence number value is arbitrary, but MUST be expressible as a 32-bit unsigned integer and
1031 MUST be less than 2**31. As long as it follows the above guidelines, a client may use any mechanism it
1032 would like to select CSeq header field values.

1033

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

1034

Example:

1035

```
CSeq: 4711 INVITE
```

1036

8.1.1.6 Max-Forwards The Max-Forwards header serves to limit the number of hops a request can
1037 transit on the way to its destination. It consists of an integer that is decremented by one at each hop.
1038 If the Max-Forwards value reaches 0 before the request reaches its destination, it will be rejected with a
1039 483 Too Many Hops error response.

1040

A UAC MUST insert a Max-Forwards header field into each request it originates with a value of 70.

1041

8.1.1.7 Via The Via header is used to indicate the transport used for the transaction, and to identify the
1042 location where the response is to be sent.

1043

When the UAC creates a request, it MUST insert a Via into that request. The protocol and version in
1044 the header MUST be SIP and 2.0, respectively. The Via header it inserts MUST contain a branch parameter.
1045 This parameter is used to uniquely identify the transaction created by that request. This parameter is used
1046 by both the client, and the server.

1047

The branch parameter value MUST be unique across time for all requests sent by the UA. The exception
1048 to this rule is CANCEL. As discussed below, a CANCEL request will have the same value of the branch
1049 parameter as the request it cancels.

1050

1051

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

1052

1053

1054

The branch ID inserted by an element compliant with this specification MUST always begin with the
characters "z9hG4bK". These 7 characters are used as a magic cookie (7 is deemed sufficient to ensure that
an older RFC 2543 implementation would not pick such a value), so that servers receiving the request can

1055 determine that the branch ID was constructed in the fashion described by this specification (i.e., globally
1056 unique). Beyond this requirement, the precise format of the branch token is implementation-defined.

1057 The Via header maddr, ttl, and sent-by components will be set when the request is processed by the
1058 transport layer (Section 19).

1059 Via processing for proxies is described in Sections 3 and sec:proxy-response-processing-via.

1060 **8.1.1.8 Contact** The Contact header provides a SIP URI that can be used to contact that specific in-
1061 stance of the user agent for subsequent requests. The Contact header MUST be present in any request that
1062 can result in the establishment of a dialog. For the methods defined in this specification, that includes only
1063 the INVITE request. For these requests, the scope of the Contact is the dialog. That is, the Contact header
1064 refers to the URI at which the UA would like to receive requests, for requests that are part of that dialog
1065 only. Only a single URI MUST be present.

1066 For further information on the Contact header, see Section 24.10.

1067 **8.1.1.9 Supported and Require** If the UAC supports extensions to SIP that can be applied by the
1068 server to the response, the UAC SHOULD include a Supported header in the request listing the option tags
1069 (Section 23.2) for those extensions. This includes support for reliability for provisional responses, which is
1070 an extension even though it is defined within this specification. The option tag for reliability of provisional
1071 responses is 100rel.

1072 The option-tags listed MUST only refer to extensions defined in standards-track RFCs. This is to prevent
1073 servers from insisting that clients implement non-standard, vendor-defined features in order to receive ser-
1074 vice. Extensions defined by experimental and informational RFCs are explicitly excluded from usage with
1075 the Supported header in a request, since they too are often used to document vendor-defined extensions.

1076 If the UAC wishes to insist that a UAS understand an extension that the UAC will apply to the request
1077 in order to process the request, it MUST insert a Require header into the request listing the option tag for
1078 that extension. If the UAC wishes to apply an extension to the request and insist that any proxies that are
1079 traversed understand that extension, it MUST insert a Proxy-Require header into the request listing the
1080 option tag for that extension.

1081 As with the Supported header, the option-tags in the Require header MUST only refer to extensions
1082 defined in standards-track RFCs.

1083 A Require header in a request with the option tag 100rel means that the UAC wishes for all provi-
1084 sional responses to this request to be transmitted reliably. This header MUST NOT be present in any requests
1085 excepting INVITE, although extensions to SIP may allow its usage with other request methods.

1086 **8.1.1.10 Additional Message Components** After a new request has been created, and the headers de-
1087 scribed above have been properly constructed, any additional optional headers are added, as are any headers
1088 specific to the method.

1089 SIP requests MAY contain a MIME-encoded message-body. Regardless of the type of body that a request
1090 contains, certain headers must be formulated to characterize the contents of the body. For further information
1091 on these headers see Sections 24.14, 24.15 and 24.12.

1092 **8.1.2 Sending the Request**

1093 The destination for the request is then computed. A loose-routing element MAY use local policy to determine
1094 the IP address, port, and transport used to reach the destination. One example of such a policy is an element

1095 configured to send requests to a default outbound proxy. Section 8.1.3 discusses restrictions on loose-
1096 routing policies. For other elements, the destination can be determined by applying the DNS procedures
1097 described in [8] to the Request-URI. These procedures yield an ordered set of address, port, and transports
1098 to attempt. The UAC SHOULD follow the procedures defined there for stateful elements, trying each address
1099 until a server is contacted. Each try constitutes a new transaction, and therefore each carries a different Via
1100 header with a new branch parameter. Furthermore, the transport value in the Via header is set to whatever
1101 transport was determined for the target server.

1102 **8.1.3 Loose Routing Policies**

1103 An element MAY apply a local loose-routing policy when preparing and sending a request. This policy MAY
1104 affect the Request-URI and Route header field values in the request as well as where the request is sent,
1105 and what transport mechanism is used to send it.

1106 Elements SHOULD use the strict-routing policy of removing the topmost value from a route set, placing
1107 it in the Request-URI and sending the request to the location indicated by that URI.

1108 This is the behavior of elements implementing earlier strict versions of Route/Record-Route.

1109 Where appropriate, elements MAY deviate from the strict-routing policy as long as the following restric-
1110 tions are met:

1111 **8.1.3.1 Modifying the Route header field** A loose-routing element MAY remove the topmost Route
1112 header field value. It MUST remove the topmost Route header field value if that value indicates a resource
1113 this element is responsible for. The element MUST NOT modify or remove any subsequent Route header
1114 field values. The element MAY place additional Route header field values into the Route header field before
1115 any existing values (effectively pushing values onto the top of the Route set).

1116 A loose-routing element may chose to not remove the first Route header field value. For example, elements
1117 configured to use default outbound proxies in lieu of using the DNS resolution procedures will leave the topmost
1118 Route header field value in the message.

1119 When the topmost Route header field value indicates a resource this element is responsible for, the message
1120 has reached the element indicated by the route, and that value must be removed from the Route header field. This
1121 assures that Route header field values are consumed when the destination they indicate has been reached.

1122 **8.1.3.2 Modifying the Request-URI** If the Request-URI identifies a resource for which this element
1123 is responsible, the loose-route policy SHOULD include modifying the Request-URI before sending the
1124 request.

1125 This restriction ensures that a Request-URI is modified once the resource it indicates has been reached.

1126 **8.1.3.3 Destination Choice** A loose-routing policy MUST direct the request to or the resource indicated
1127 in the first Route header field value, or to a proxy it trusts to ensure this property.

1128 This restriction ensures the resource indicated by the topmost Route header field value is actually visited.

1129 **8.1.3.4 Loop Avoidance** The Request-URI of a request emitted by a loose-routing element MUST differ
1130 from the URI in the first Route header field value.

1131 This restriction is necessary to avoid triggering false loop detections in older systems. The following
1132 algorithm can be applied to ensure sufficient difference in otherwise matching Request-URIs and first
1133 Route header field values.

1134 For each of these items, D is the address of the next hop (which may or may not be equivalent to A).
1135 If the topmost element in the received Route header field is `∫sip:a@A∫`, the outgoing request will contain

```
1136     METHOD sip:a@A;maddr=D  
1137     Route: <∫sip:a@A>
```

1138 If the topmost element in the received Route header field is `∫sip:a@A;maddr=D∫`, the outgoing request
1139 will contain

```
1140     METHOD sip:a@A  
1141     Route: <∫sip:a@A;maddr=D>
```

1142 If the topmost element in the received Route header field is `∫sip:a@A;maddr=B∫` and $D \neq B$, the outgoing
1143 request will contain

```
1144     METHOD sip:a@A;maddr=D  
1145     Route: <∫sip:a@A;maddr=B>
```

1146 8.1.4 Processing Responses

1147 Responses are first processed by the transport layer and then passed up to the transaction layer. The trans-
1148 action layer performs its processing and then passes it up to the TU. The majority of response processing in
1149 the TU is method specific. However, there are some general behaviors independent of the method.

1150 **8.1.4.1 Transaction Layer Errors** In some cases, the response returned by the transaction layer will
1151 not be a SIP message, but rather a transaction layer event. The only event that the TU will encounter is the
1152 timeout event. When the timeout event is received from the transaction layer, it **MUST** be treated as if a 408
1153 (Request Timeout) status code has been received.

1154 **8.1.4.2 Unrecognized Responses** A UAC **MUST** treat any response it does not recognize as being equiv-
1155 alent to the x00 response code of that class, and **MUST** be able to process the x00 response code for all
1156 classes. For example, if a UAC receives an unrecognized response code of 431, it can safely assume that
1157 there was something wrong with its request and treat the response as if it had received a 400 (Bad Request)
1158 response code.

1159 **8.1.4.3 Vias** If more than one *Via* header field is present in a response, the UAC **SHOULD** discard the
1160 message.

1161 The presence of additional *Via* header fields that precede the originator of the request suggests that the message
1162 was misrouted or possibly corrupted.

1163 **8.1.4.4 Processing Reliable 1xx Responses** A 1xx response that contains a **Require** header with the
1164 option tag `100rel` is a reliable provisional response. The UA core follows the procedures in Section 18.2
1165 to process the response, which will result in the generation of a **PRACK** request to acknowledge the reliable
1166 provisional response.

1167 **8.1.4.5 Processing 3xx responses** Upon receipt of a redirection response (for example, a 3xx response
1168 status code), clients SHOULD use the URI(s) in the Contact header field to formulate one or more new
1169 requests based on the redirected request.

1170 If more than one URI is present in Contact header fields within the 3xx response, the UA MUST deter-
1171 mine an order in which these contact addresses should be processed. UAs MUST consult the “q” parameter
1172 value of the Contact header fields (see Section 22.10) if available. Contact addresses MUST be ordered from
1173 highest qvalue to lowest. If no qvalue is present, a contact address is considered to have a qvalue of 1.0.
1174 Note that two or more contact addresses might have an equal qvalue - these URIs are eligible to be tried in
1175 parallel.

1176 Once an ordered list has been established, UACs MUST try to contact each URI in the ordered list in turn
1177 until a server responds. If there are contact addresses with an equal qvalue, the UAC MAY decide randomly
1178 on an order in which to process these addresses, or it MAY attempt to process contact addresses of equal
1179 qvalue in parallel.

1180 Note that for example, the UAC may effectively divide the ordered list into groups, processing the groups
1181 serially and processing the destinations in each group in parallel.

1182 If contacting an address in the list results in a failure, as defined in the next paragraph, the element moves
1183 to the next address in the list, until the list is exhausted. If the list is exhausted, then the request has failed.

1184 Failures SHOULD be detected through failure response codes (codes greater than 399) or network time-
1185 outs. Client transaction will report any transport layer failures to the transaction user.

1186 When a failure for a particular contact address is received, the client SHOULD try the next contact
1187 address. This will involve creating a new client transaction to deliver a new request.

1188 In order to create a request based on a contact address in a 3xx response, a UAC MUST copy the entire
1189 URI from the Contact header into the Request-URI, except for the “method-param” and “header” URI
1190 parameters (see Section 23.1.1 for a definition of these parameters). It uses the “header” parameters to
1191 create headers for the new request, overwriting headers associated with the redirected request in accordance
1192 with the guidelines in Section 23.1.5.

1193 Note that in some instances, headers that have been communicated in the contact address may instead
1194 append to existing request headers in the original redirected request. As a general rule, if the header can
1195 accept a comma-separated list of values, then the new header value MAY be appended to any existing values
1196 in the original redirected request. If the header does not accept multiple values, the value in the original
1197 redirected request MAY be overwritten by the header value communicated in the contact address.

1198 For example, if a contact address is returned with the following value:

```
1199 sip:user@host?Subject=foo&Call-Info=<http://www.foo.com>
```

1200 Then any Subject header in the original redirected request is overwritten, but the HTTP URL is merely
1201 appended to any existing Call-Info header field values.

1202 It is RECOMMENDED that the UAC reuse the same To, From, and Call-ID used in the original redirected
1203 request, but the UAC MAY also choose to update for example the Call-ID header field value for new requests.

1204 Finally, once the new request has been constructed, it is sent using a new client transaction, and therefore
1205 MUST have a new branch ID in the top Via field as discussed in Section 8.1.1.7.

1206 In all other respects, requests sent upon receipt of a redirect response SHOULD re-use the headers and
1207 bodies of the original request.

1208 In some instances, Contact header values may be cached at UAC temporarily or permanently depending
1209 on the status code received and the presence of an expiration interval; see Sections 25.3.2 and 25.3.3.

1210 **8.1.4.6 Processing 4xx responses** Certain 4xx response codes require specific UA processing, indepen-
1211 dent of the method.

1212 If a 401 (Unauthorized) or 407 (Proxy Authentication Required) response is received, the UAC SHOULD
1213 follow the authorization procedures of Section 20.2 and Section 20.3 to retry the request with credentials.

1214 If a 413 (Request Entity Too Large) response is received (Section 25.4.11), the request contained a body
1215 that was longer than the UAS was willing to accept. If possible, the UAC SHOULD retry the request, either
1216 omitting the body or using one of a smaller length.

1217 If a 415 (Unsupported Media Type) response is received (Section 25.4.13), the request contained media
1218 types not supported by the UAS. The UAC SHOULD retry sending the request, this time only using content
1219 with types listed in the `Accept` header in the response, with encodings listed in the `Accept-Encoding` header
1220 in the response, and with languages listed in the `Accept-Language` in the response.

1221 If a 416 (Unsupported URI Scheme) response is received (Section 25.4.14, the `Request-URI` used a
1222 URI scheme not supported by the server. The client SHOULD retry the request, this time, using a SIP URI.

1223 If a 420 (Bad Extension) response is received (Section 25.4.15), the request contained a `Require` or
1224 `Proxy-Require` header listing an option-tag for a feature not supported by a proxy or UAS. The UAC
1225 SHOULD retry the request, this time omitting any extensions listed in the `Unsupported` header in the re-
1226 sponse.

1227 In all of the above cases, the request is retried by creating a new request with the appropriate modifica-
1228 tions. This new request SHOULD have the same value of the `Call-ID`, `To`, and `From` of the previous request,
1229 but the `CSeq` should contain a new sequence number that is one higher than the previous.

1230 With other 4xx responses, including those yet to be defined, a retry may or may not be possible depend-
1231 ing on the method and the use case.

1232 **8.2 UAS Behavior**

1233 When a request outside of a dialog is processed by a UAS, there is a set of processing rules which are
1234 followed, independent of the method. Section 12 gives guidance on how a UAS can tell whether a request
1235 is inside or outside of a dialog.

1236 Note that request processing is atomic. If a request is accepted, all state changes associated with it MUST
1237 be performed. If it is rejected, all state changes MUST NOT be performed.

1238 **8.2.1 Method Inspection**

1239 Once a request is authenticated (or no authentication was desired), the UAS MUST inspect the method of the
1240 request. If the UAS does not support the method of a request it MUST generate a 405 (Method Not Allowed)
1241 response. Procedures for generation of responses are described in Section 8.2.6. The UAS MUST also add
1242 an `Allow` header to the 405 (Method Not Allowed) response. The `Allow` header field MUST list the set of
1243 methods supported by the UAS generating the message.

1244 The `Allow` header field is presented in Section 24.5.

1245 If the method is one supported by the server, processing continues.

1246 **8.2.2 Header Inspection**

1247 If a UAS does not understand a header field in a request (that is, the header is not defined in this specification
1248 or in any supported extension), the server MUST ignore that header and continue processing the message. A
1249 UAS SHOULD ignore any malformed headers that are not necessary for processing requests.

1250 **8.2.2.1 To and Request-URI** The **To** header field identifies the original recipient of the request desig-
1251 nated by the user identified in the **From** field. The original recipient may or may not be the UAS processing
1252 the request, due to call forwarding or other proxy operations. A UAS MAY apply any policy it wishes in
1253 determination of whether to accept requests when the **To** field is not the identity of the UAS. However, it is
1254 RECOMMENDED that a UAS accept requests even if they do not recognize the URI scheme (for example,
1255 a `tel:` URI) in the **To** header, or if the **To** header field does not address a known or current user of this
1256 UAS. If, on the other hand, the UAS decides to reject the request, it SHOULD generate a response with a 403
1257 (Forbidden) status code and pass it to the server transaction layer for transmission.

1258 However, the **Request-URI** identifies the UAS that is to process the request. If the **Request-URI** uses
1259 a scheme not supported by the UAS, it SHOULD reject the request with a 416 (Unsupported URI Scheme)
1260 response. If the **Request-URI** does not identify an address that the UAS is willing to accept requests for,
1261 it SHOULD reject the request with a 404 (Not Found) response. Typically, a UA that uses the **REGISTER**
1262 method to bind its address of record to a specific contact address will see requests whose **Request-URI**
1263 equals those contact addresses. Other potential sources of received **Request-URIs** include the **Contact**
1264 headers of requests and responses sent by the UA that establish or refresh dialogs.

1265 **8.2.2.2 Merged Requests** If the request has no tag in the **To**, the TU checks ongoing transactions. If the
1266 **To**, **From**, **Call-ID**, **CSeq** exactly match (including tags) those of any request received previously, but the
1267 **branch-ID** in the topmost **Via** is different from those received previously, the TU SHOULD generate a 482
1268 (Loop Detected) response and pass it to the server transaction.

1269 The same request has arrived at the UAS more than once, following different paths, most likely due to forking.
1270 The UAS processes the first such request received and responds with a 482 (Loop Detected) to the rest of them.

1271 **8.2.2.3 Require** Assuming the UAS decides that it is the proper element to process the request, it ex-
1272 amines the **Require** header field, if present.

1273 The **Require** general-header field is used by a UAC to tell a UAS about SIP extensions that the UAC
1274 expects the UAS to support in order to process the request properly. Its format is described in Section 24.33.
1275 If a UAS does not understand an option-tag listed in a **Require** header field, it MUST respond by generating a
1276 response with status code 420 (Bad Extension). The UAS MUST add an **Unsupported** header field, and list
1277 in it those options it does not understand amongst those in the **Require** header of the request. Upon receipt
1278 of the 420 (Bad Extension) the client SHOULD retry the request, this time without using those extensions
1279 listed in the **Unsupported** header field in the response.

1280 Note that **Require** and **Proxy-Require** MUST NOT be used in a SIP **CANCEL** request, or in an **ACK**
1281 request sent for a non-2xx response. These headers should be ignored if they are present in these requests.

1282 An **ACK** request for a 2xx response MUST contain only those **Require** and **Proxy-Require** values that
1283 were present in the initial request.

1284 Example:

```
1285 UAC->UAS:  INVITE sip:watson@bell-telephone.com SIP/2.0  
1286             Require: 100rel  
1287  
1288  
1289 UAS->UAC:  SIP/2.0 420 Bad Extension  
1290             Unsupported: 100rel
```

1291 This is to make sure that the client-server interaction will proceed without delay when all options are understood

1292 by both sides, and only slow down if options are not understood (as in the example above). For a well-matched
1293 client-server pair, the interaction proceeds quickly, saving a round-trip often required by negotiation mechanisms.
1294 In addition, it also removes ambiguity when the client requires features that the server does not understand. Some
1295 features, such as call handling fields, are only of interest to end systems.

1296 **8.2.3 Content Processing**

1297 Assuming the UAS understands any extensions required by the client, the UAS examines the body of the
1298 message, and the headers that describe it. If there are any bodies whose type (indicated by the **Content-**
1299 **Type**), language (indicated by the **Content-Language**) or encoding (indicated by the **Content-Encoding**)
1300 are not understood, and that body part is not optional (as indicated by the **Content-Disposition** header), the
1301 UAS **MUST** reject the request with a 415 (Unsupported Media Type) response. The response **MUST** contain
1302 an **Accept** header listing the types of all bodies it understands, in the event the request contained bodies
1303 of types not supported by the UAS. If the request contained content encodings not understood by the UAS,
1304 the response **MUST** contain an **Accept-Encoding** header listing the encodings understood by the UAS. If
1305 the request contained content with languages not understood by the UAS, the response **MUST** contain an
1306 **Accept-Language** header indicating the languages understood by the UAS.

1307 Beyond these checks, body handling depends on the method and type.

1308 For further information on the processing of Content-specific headers see Section 7.4 as well as Sec-
1309 tion 24.11 through 24.15.

1310 **8.2.4 Applying Extensions**

1311 A UAS that wishes to apply some extension when generating the response **MUST** only do so if support for
1312 that extension is indicated in the **Supported** header in the request. If the desired extension is not supported,
1313 the server **SHOULD** rely only on baseline SIP and any other extensions supported by the client. To ensure
1314 that the **SHOULD** can be fulfilled, any specification of a new extension **MUST** include discussion of how
1315 to return gracefully to baseline SIP when the extension is not present. In rare circumstances, where the
1316 server cannot process the request without the extension, the server **MAY** send a 421 (Extension Required)
1317 response. This response indicates that the proper response cannot be generated without support of a specific
1318 extension. The needed extension(s) **MUST** be included in a **Require** header in the response. This behavior
1319 is **NOT RECOMMENDED**, as it will generally break interoperability.

1320 Any extensions applied to a non-421 response **MUST** be listed in a **Require** header included in the
1321 response. Of course, the server **MUST NOT** apply extensions not listed in the **Supported** header in the
1322 request. As a result of this, the **Require** header in a response will only ever contain option tags defined in
1323 standards-track RFCs.

1324 **8.2.5 Processing the Request**

1325 Assuming all of the checks in the previous subsections are passed, the UAS processing becomes method-
1326 specific. Section 10 covers the **REGISTER** request, section 11 covers the **OPTIONS** request, section 13
1327 covers the **INVITE** request, and section 15 covers the **BYE** request.

1328 **8.2.6 Generating the Response**

1329 When a UAS wishes to construct a response to a request, it follows these procedures. Additional procedures
1330 may be needed depending on the status code of the response and the circumstances of its construction. These

1331 additional procedures are documented elsewhere.

1332 **8.2.6.1 Sending a Provisional Response** One largely non-method-specific guideline for the generation
1333 of responses is that UASs SHOULD NOT issue a provisional response for a non-INVITE request. Rather,
1334 UASs SHOULD generate a final response to a non-INVITE request as soon as possible.

1335 When a 100 (Trying) response is generated, any Timestamp header present in the request MUST be
1336 copied into this 100 (Trying) response.

1337 **8.2.6.2 Headers and Tags** The From field of the response MUST equal the From field of the request.
1338 The Call-ID field of the response MUST equal the Call-ID field of the request. The Cseq field of the response
1339 MUST equal the Cseq field of the request. The Via headers in the response MUST equal the Via headers in
1340 the request and MUST maintain the same ordering.

1341 If a request contained a To tag in the request, the To field in the response MUST equal that of the request.
1342 However, if the To field in the request did not contain a tag, the URI in the To field in the response MUST
1343 equal the URI in the To field in the request; additionally, the UAS MUST add a tag to the To field in the
1344 response (with the exception of the 100 (Trying) response, in which a tag MAY be present). This serves to
1345 identify the UAS that is responding, possibly resulting in a component of a dialog ID. The same tag MUST
1346 be used for all responses to that request, both final and provisional (again excepting the 100 (Trying)).
1347 Procedures for generation of tags are defined in Section 23.3.

1348 **8.2.7 Stateless UAS Behavior**

1349 A stateless UAS is a UAS that does not maintain transaction state. It replies to requests normally, but
1350 discards any state that would ordinarily be retained by a UAS after a response has been sent. If a stateless
1351 UAS receives a retransmission of a request, it regenerates the response and resends it, just as if it were the
1352 replying to the first instance of the request. Stateless UASs do not use a transaction layer; they receive
1353 requests directly from the transport layer and send responses directly to the transport layer.

1354 The stateless UAS role is needed primarily to handle unauthenticated requests for which a challenge
1355 response is issued. If unauthenticated requests were handled statefully, then malicious floods of unauthenti-
1356 cated requests could create massive amounts of transaction state that might slow or complete halt call pro-
1357 cessing in a UAS, effectively creating a denial of service condition; for more information see Section 22.1.5.

1358 The most important behaviors of a stateless UAS are the following:

- 1359 ● A stateless UAS MUST NOT send provisional (1xx) responses.
- 1360 ● A stateless UAS MUST NOT retransmit responses.
- 1361 ● A stateless UAS MUST ignore ACK requests.
- 1362 ● A stateless UAS MUST ignore CANCEL requests.
- 1363 ● To header tags MUST be generated for responses in a stateless manner - in a manner that will generate
1364 the same tag for the same request consistently. For information on tag construction see Section 23.3.

1365 In all other respects, a stateless UAS behaves in the same manner as a stateful UAS. A UAS can operate
1366 in either a stateful or stateless mode for each new request.

1367 8.3 Redirect Servers

1368 In some architectures it may be desirable to reduce the processing load on proxy servers that are responsible
1369 for routing requests, and improve signaling path robustness, by relying on redirection. Redirection allows
1370 servers to push routing information for a request back in a response to the client, thereby taking themselves
1371 out of the loop of further messaging for this transaction while still aiding in locating the target of the request.
1372 When the originator of the request receives the redirection, it will send a new request based on the URI it has
1373 received. By propagating URIs from the core of the network to its edges, redirection allows for considerable
1374 network scalability.

1375 A redirect server is logically constituted of a server transaction layer and a transaction user that has
1376 access to a location service of some kind (see Section 10 for more on registrars and location services). This
1377 location service is effectively a database containing mappings between a single URI and a set of one or more
1378 alternative locations at which the target of that URI can be found.

1379 A redirect server does not issue any SIP requests of its own. After receiving a request other than CAN-
1380 CEL, the server gathers the list of alternative locations from the location service and either returns a final
1381 response of class 3xx or it refuses the request. For well-formed CANCEL requests, it SHOULD return a
1382 2xx response. This response ends the SIP transaction. The redirect server maintains transaction state for an
1383 entire SIP transaction. It is the responsibility of clients to detect forwarding loops between redirect servers.

1384 When a redirect server returns a 3xx response to a request, it populates the list of (one or more) alterna-
1385 tive locations into Contact headers. An "expires" parameter to the Contact header may also be supplied
1386 to indicate the lifetime of the Contact data.

1387 The Contact header field contains URIs giving the new locations or user names to try, or may simply
1388 specify additional transport parameters. A 301 (Moved Permanently) or 302 (Moved Temporarily) response
1389 may also give the same location and username that was targeted by the initial request but specify additional
1390 transport parameters such as a different server or multicast address to try, or a change of SIP transport from
1391 UDP to TCP or vice versa.

1392 However, redirect servers MUST NOT redirect a request to a URI equal to the one in the Request-URI;
1393 instead, provided that the URI does not point to itself, the redirect server SHOULD proxy the request to the
1394 destination URI.

1395 If a client is using an outbound proxy, and that proxy actually redirects requests, a potential arises for infinite
1396 redirection loops.

1397 Note that the Contact header field MAY also refer to a different entity than the one originally called. For
1398 example, a SIP call connected to GSTN gateway may need to deliver a special informational announcement
1399 such as "The number you have dialed has been changed."

1400 A Contact response header field can contain any suitable URI indicating where the called party can be
1401 reached, not limited to SIP URIs. For example, it could contain URIs for phones, fax, or irc (if they were
1402 defined) or a mailto: (RFC 2368, [16]) URL.

1403 The "expires" parameter of the Contact header field indicates how long the URI is valid. The value of
1404 the parameter is a number indicating seconds. If this parameter is not provided, the value of the Expires
1405 header field determines how long the URI is valid. Implementations MAY treat values larger than 2**32-
1406 1 (4294967295 seconds or 136 years) as equivalent to 2**32-1. Malformed values should be treated as
1407 equivalent to 3600.

1408 Redirect servers MUST ignore features that are not understood (including unrecognized headers, Re-
1409 quired extensions, or even method names) and proceed with the redirection of the session in question. If
1410 a particular extension requires that intermediate devices support it, the extension MUST be tagged in the
1411 Proxy-Require field as well (see Section 24.29).

1412 9 Canceling a Request

1413 The previous section has discussed general UA behavior for generating requests, and processing responses,
1414 for requests of all methods. In this section, we discuss a general purpose method, called CANCEL.

1415 The CANCEL request, as the name implies, is used to cancel a previous request sent by a client. Specif-
1416 ically, it asks the UAS to cease processing the request and to generate an error response to that request.
1417 CANCEL has no effect on a request to which a UAS has already responded. Because of this, it is most
1418 useful to CANCEL requests to which can take a long time to respond. For this reason, CANCEL is most
1419 useful for INVITE requests, which can take a long time to generate a response. In that usage, a UAS that
1420 receives a CANCEL request for an INVITE, but has not yet sent a response, would “stop ringing”, and then
1421 respond to the INVITE with a specific error response (a 487).

1422 CANCEL requests can be constructed and sent by any type of client, including both proxies and user
1423 agent clients. Section 15 discusses under what conditions a UAC would CANCEL an INVITE request, and
1424 Section 16.9 discusses proxy usage of CANCEL.

1425 Because a stateful proxy can generate its own CANCEL, a stateful proxy also responds to a CANCEL,
1426 rather than simply forwarding a response it would receive from a downstream element. For that reason,
1427 CANCEL is referred to as a “hop-by-hop” request, since it is responded to at each stateful proxy hop.

1428 9.1 Client Behavior

1429 A CANCEL request SHOULD NOT be sent to cancel a request other than INVITE.

1430 Since requests other than INVITE are responded to immediately, sending a CANCEL for a non-INVITE request
1431 would always create a race condition.

1432 The following procedures are used to construct a CANCEL request. The Request-URI, Call-ID, To,
1433 the numeric part of CSeq and From header fields in the CANCEL request MUST be identical to those in
1434 the request being cancelled, including tags. A CANCEL constructed by a client MUST have only a single
1435 Via header, whose value matches the top Via in the request being cancelled. Using the same values for
1436 these headers allows the CANCEL to be matched with the request it cancels (Section 9.2 indicates how such
1437 matching occurs). However, the method part of the CSeq header MUST have a value of CANCEL. This
1438 allows it to be identified and processed as a transaction in its own right (See Section 17). If the request being
1439 cancelled contains Route header fields, the CANCEL request MUST include these Route header fields.

1440 This is needed so that stateless proxies are able to route CANCEL requests properly.

1441 The CANCEL request MUST NOT contain any Require or Proxy-Require header fields.

1442 Once the CANCEL is constructed, the client SHOULD check whether any response (provisional or final)
1443 has been received for the request being cancelled (herein referred to as the “original request”). The CANCEL
1444 request MUST NOT be sent if no provisional response has been received, rather, the client MUST wait for the
1445 arrival of a provisional response before sending the request. If the original request has generated a final
1446 response, the CANCEL SHOULD NOT be sent, as it is an effective no-op, since CANCEL has no effect
1447 on requests that have already generated a final response. When the client decides to send the CANCEL, it
1448 creates a client transaction for the CANCEL and passes it the CANCEL request along with the destination
1449 address, port, and transport. The destination address, port, and transport for the CANCEL MUST be identical
1450 to those used to send the original request.

1451 If it was allowed to send the CANCEL before receiving a response for the previous request, the server could
1452 receive the CANCEL before the original request.

1453 Note that both the transaction corresponding to the original request and the CANCEL transaction will
1454 complete independently. However, a UAC canceling a request cannot rely on receiving a 487 (Request

1455 Terminated) response for the original request, as an RFC 2543-compliant UAS will not generate such a
1456 response. If there is no final response for the original request in $64 * T1$ seconds ($T1$ is defined in Section
1457 17.1.1.1), the client SHOULD then consider the original transaction cancelled and SHOULD destroy the client
1458 transaction handling the original request.

1459 9.2 Server Behavior

1460 The CANCEL method requests that the TU at the server side cancel a pending transaction. The transaction
1461 to be canceled is determined by taking the CANCEL request, and then assuming that the request method
1462 were anything but CANCEL, apply the transaction matching procedures of Section 17.2.3. The matching
1463 transaction is the one to be canceled.

1464 The processing of a CANCEL request at a server depends on the type of server. A stateless proxy will
1465 forward it, a stateful proxy might respond to it and generate some CANCEL requests of its own, and a UAS
1466 will respond to it. See Section 16.9 for proxy treatment of CANCEL.

1467 A UAS first processes the CANCEL request according to the general UAS processing described in
1468 Section 8.2. However, since CANCEL requests are hop-by-hop and cannot be resubmitted, they cannot be
1469 challenged by the server in order to get proper credentials in an Authorization header field. Note also that
1470 CANCEL requests do not contain Require header fields.

1471 If the CANCEL did not find a matching transaction according to the procedure above, the CANCEL
1472 SHOULD be responded to with a 481 (Call Leg/Transaction Does Not Exist). If the transaction for the
1473 original request still exists, the behavior of the UAS on receiving a CANCEL request depends on whether it
1474 has already sent a final response for the original request. If it has, the CANCEL request has no effect on the
1475 processing of the original request, no effect on any session state, and no effect on the responses generated
1476 for the original request. If the UAS has not issued a final response for the original request, its behavior
1477 depends on the method of the original request. If the original request was an INVITE, the UAS SHOULD
1478 immediately respond to the INVITE with a 487 (Request Terminated). The behavior upon reception of a
1479 CANCEL request for any other method defined in this specification is effectively no-op. Extensions to this
1480 specification that define new methods MUST define the behavior of a UAS upon reception of a CANCEL for
1481 those methods.

1482 Regardless of the method of the original request, as long as the CANCEL matched an existing trans-
1483 action, the CANCEL request itself is answered with a 200 (OK) response. This response is constructed
1484 following the procedures described in Section 8.2.6 noting that the To tag of the response to the CANCEL
1485 and the To tag in the response to the original request SHOULD be the same. The response to CANCEL is
1486 passed to the server transaction for transmission.

1487 10 Registrations

1488 10.1 Overview

1489 SIP offers a discovery capability. If a user wants to initiate a session with another user, SIP must discover the
1490 current host(s) that the destination user is reachable at. This discovery process is accomplished by SIP proxy
1491 servers, which are responsible for receiving a request, determining where to send it based on knowledge of
1492 the location of the user, and then sending it there. To do this, proxies consult an abstract service known as a
1493 *location service*, which provides address bindings for a particular domain. These address bindings map an
1494 incoming SIP URI, `sip:bob@Biloxi.com`, for example, to one or more SIP URIs which are somehow

1495 “closer” to the desired user, sip:bob@engineering.Biloxi.com, for example. Ultimately, a proxy
1496 will consult a location service which maps a received URI to the current host(s) that a user is logged in to.

1497 Registration creates bindings in a location service for a particular domain that associate an address-of-
1498 record URI with one or more contact addresses. This means that when a proxy for that domain receives a
1499 request whose request URI matches the address-of-record, the proxy will forward the request to the contact
1500 addresses registered to that address-of-record. Generally, it only makes sense to register an address-of-
1501 record at a location service for a domain when requests for that address-of-record would be routed to that
1502 domain. In most cases, this means that the domain of the registration will need to match the domain in the
1503 URI of the address-of-record.

1504 There are many ways by which the contents of the location service can be established. One way is
1505 administratively. In the above example, Bob is known to be a member of the engineering department through
1506 access to a corporate database. SIP provides a mechanism, however, for a user agent to explicitly create a
1507 binding. This mechanism is known as registration.

1508 Registration entails sending a REGISTER request to a special type of UAS known as a registrar. The
1509 registrar acts as a front end to the location service for a domain, reading and writing mappings based on the
1510 contents of the REGISTER requests. This location service will then be consulted by a proxy server that is
1511 responsible for routing requests for that domain.

1512 SIP does not mandate a particular mechanism for implementing the location service. The only require-
1513 ment is that a registrar for some domain MUST be able to read and write data to the location service, and
1514 a proxy for that domain MUST be capable of reading that same data. A registrar MAY be co-located with a
1515 particular SIP proxy server for the same domain.

1516 10.2 Constructing the REGISTER Request

1517 REGISTER requests add, remove and query bindings. A REGISTER request may add a new binding
1518 between an address-of-record and one or more contact addresses. Registration on behalf of a particular
1519 address-of-record may be performed by a suitably authorized third party. A client may also remove previous
1520 bindings, or query to determine which bindings are currently in place for an address-of-record.

1521 Except as noted, the construction of the REGISTER request and the behavior of clients sending a
1522 REGISTER request is identical to the general UAC behavior described in Section 8.1 and Section 17.1.
1523 The following header fields MUST be included:

1524 **Request-URI:** The Request-URI names the domain of the location service that the registration is meant
1525 for (e.g., “sip:chicago.com”). The “userinfo” and “@” components of the SIP URI MUST NOT be
1526 present.

1527 **To:** The To header field contains the address of record whose registration is to be created, queried or mod-
1528 ified. The To header field and the Request-URI field typically differ, as the former contains a user
1529 name. This address-of-record MUST be a SIP URI.

1530 **From:** The From header field contains the address-of-record of the person responsible for the registration.
1531 The value is the same as the To header field unless the request is a third-party registration.

1532 **Call-ID:** All registrations from a user agent client SHOULD use the same Call-ID header value for registra-
1533 tions sent to a particular registrar.

1534 If the same client were to use different Call-ID values, a registrar could not detect whether a delayed
1535 REGISTER request might have arrived out of order.

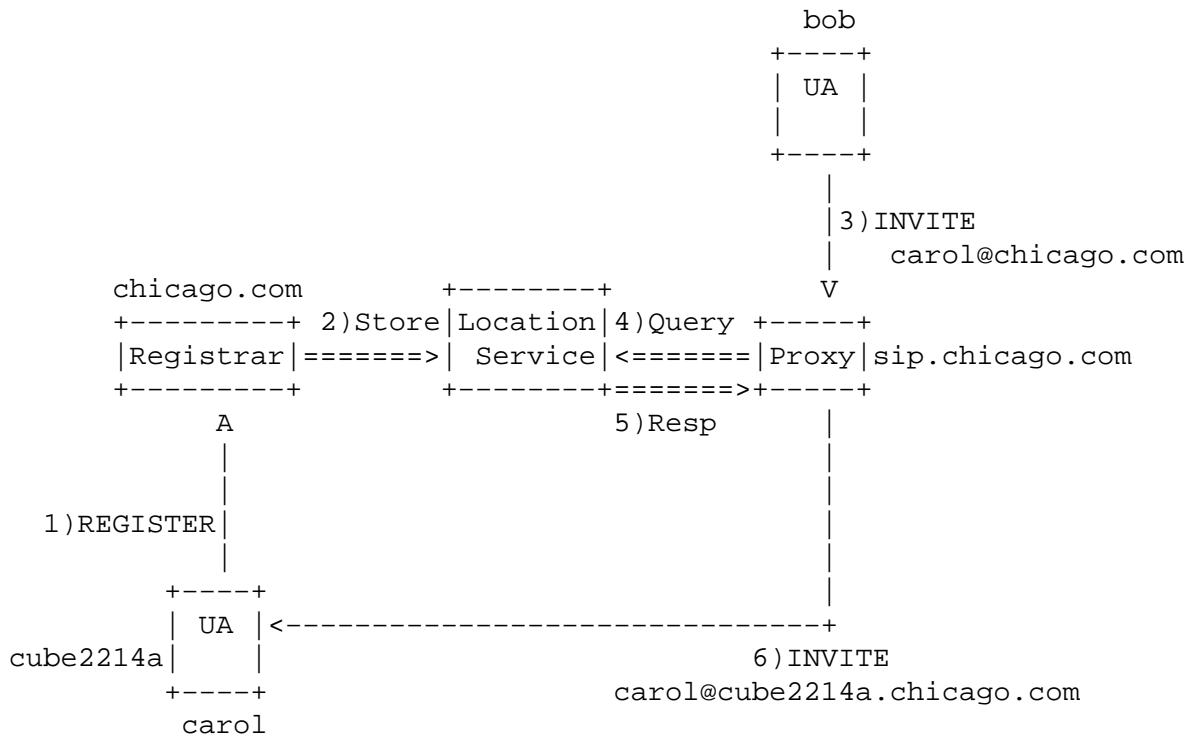


Figure 2: REGISTER example

1536 **CSeq:** The CSeq value guarantees proper ordering of REGISTER requests. A UA MUST increment the
 1537 CSeq value by one for each REGISTER request with the same Call-ID.

1538 **Contact:** REGISTER requests contain zero or more Contact header fields, containing address bindings.

1539 User agents MUST NOT send a new registration (i.e., containing new Contact header fields, as opposed
 1540 to a retransmission) until they have received a final response from the registrar for the previous one or the
 1541 previous REGISTER request has timed out.

1542 The following Contact header parameters have a special meaning in REGISTER requests:

1543 **action:** The “action” parameter from RFC 2543 has been deprecated. UACs SHOULD NOT use the
 1544 “action” parameter.

1545 **expires:** The “expires” parameter indicates how long the UA would like the binding to be valid. The value
 1546 is a number indicating seconds. If this parameter is not provided, the value of the Expires header field
 1547 is used instead. Implementations MAY treat values larger than $2^{32}-1$ (4294967295 seconds or 136
 1548 years) as equivalent to $2^{32}-1$. Malformed values should be treated as equivalent to 3600.

1549 10.2.1 Adding Bindings

1550 The REGISTER request sent to a registrar includes contact addresses to which SIP requests for the address-
 1551 of-record should be forwarded. The address-of-record is included in the To header field of the REGISTER

1552 request.

1553 The **Contact** header fields of the request typically contain SIP URIs that identify particular SIP end-
1554 points (for example, "sip:carol@cube2214a.chicago.com"), but they MAY use any URI scheme. A SIP UA
1555 can choose to register telephone numbers (with the tel URL, [13]) or email addresses (with a mailto URL,
1556 [16]) as **Contacts** for an address-of-record.

1557 For example, Carol, with address-of-record "sip:carol@chicago.com", would register with the SIP reg-
1558 istrar of the domain chicago.com. Her registrations would then be used by a proxy server in the chicago.com
1559 domain to route requests for Carol's address-of-record to her SIP endpoint.

1560 Once a client has established bindings at a registrar, it MAY send subsequent registrations containing
1561 new bindings or modifications to existing bindings as necessary. The 2xx response to the **REGISTER**
1562 request will contain, in **Contact** header fields, a complete list of bindings that have been registered for this
1563 address-of-record at this registrar.

1564 Registrations do not need to update all bindings. Typically, a UA only updates its own SIP URI as well
1565 as any non-SIP URIs.

1566 **10.2.1.1 Setting the Expiration Interval of Contact Addresses** When a client sends a **REGISTER**
1567 request, it MAY suggest an expiration interval that indicates how long the client would like the registration
1568 to be valid. (As described in Section 10.3, the registrar selects the actual time interval based on its local
1569 policy.)

1570 There are two ways in which a client can suggest an expiration interval for a binding: through an
1571 **Expires** header field, or an "expires" **Contact** header parameter. The latter allows expiration intervals to
1572 be suggested on a per-binding basis when more than one binding is given in a single **REGISTER** request,
1573 whereas the former suggests an expiration interval for all **Contact** header fields that do not contain the
1574 "expires" parameter.

1575 If neither mechanism for expressing a suggested expiration time is present in a **REGISTER**, a default
1576 suggestion of one hour is assumed.

1577 **10.2.1.2 Preferences among Contact Addresses** If more than one **Contact** is sent in a **REGISTER**
1578 request, the registering UA intends to associate all of the URIs given in these **Contact** headers with the
1579 address-of-record present in the **To** field. This list can be prioritized with the "q" parameter in the **Contact**
1580 header fields. The "q" parameter indicates a relative preference for the particular **Contact** header field
1581 compared to other bindings present in this **REGISTER** message or existing within the location service of
1582 the registrar. Section 16.5 describes how a proxy server uses this preference indication.

1583 **10.2.2 Removing Bindings**

1584 Registrations are soft state and expire unless refreshed, but can also be explicitly removed. A client can
1585 attempt to influence the expiration interval selected by the registrar as described in Section 10.2.1. A user
1586 agent requests the immediate removal of a binding by specifying an expiration interval of "0" for that
1587 contact address in a **REGISTER** request. User agents SHOULD support this mechanism so that bindings
1588 can be removed before their expiration interval has passed.

1589 The **REGISTER**-specific **Contact** header field value of "*" applies to all registrations, but it MUST only
1590 be used when the **Expires** header field is present with a value of "0".

1591 Use of the "*" **Contact** header field value allows a registering user agent to remove all of its bindings without
1592 knowing their precise values.

1593 If no **Contact** header fields are present in a **REGISTER** request, the list of bindings is left unchanged.

1594 **10.2.3 Fetching Bindings**

1595 A success response to any **REGISTER** request contains the complete list of existing bindings, regardless of
1596 whether the request contained a **Contact** header field or not.

1597 **10.2.4 Refreshing Bindings**

1598 Each UA is responsible to refresh the bindings that it has previously established. A UA **SHOULD NOT** refresh
1599 bindings set up by other UAs.

1600 The 200 (OK) response from the registrar contains a list of **Contact** fields enumerating all current
1601 bindings. The UA compares each contact address to see if it created the contact address, using comparison
1602 rules in Section 23.1.4. If so, it updates the expiration time interval according to the **expires** parameter or,
1603 if absent, the **Expires** field value. The UA then issues a **REGISTER** request for each of its bindings before
1604 the expiration interval has elapsed. It **MAY** combine several updates into one **REGISTER** request.

1605 A UA **SHOULD** use the same **Call-ID** for all registrations during a single boot cycle. Registration re-
1606 freshes **SHOULD** be sent to the same network address as the original registration, unless redirected.

1607 **10.2.5 Setting the Internal Clock**

1608 If the response for **REGISTER** request contains a **Date** header, the client **MAY** use this header field to learn
1609 the current time in order to set any internal clocks.

1610 **10.2.6 Discovering a Registrar**

1611 UAs can use three ways to determine the address to send registrations to: by configuration, using the address-
1612 of-record and multicast. A UA can be configured, in ways beyond the scope of this specification, with
1613 a registrar address. If there is no configured registrar address, the UA **SHOULD** use the host part of the
1614 address-of-record as the **Request-URI** and address the request there, using the normal SIP server location
1615 mechanisms [8]. For example, the UA for the user "sip:carol@chicago.com" addresses the **REGISTER**
1616 request to "chicago.com".

1617 Finally, a UA can be configured to use multicast. Multicast registrations are addressed to the well-known
1618 "all SIP servers" multicast address "sip.mcast.net" (224.0.1.75 for IPv4). No well-known IPv6 multicast
1619 address has been allocated; such an allocation will be documented separately when needed. This request
1620 **MUST** be scoped to ensure it is not forwarded beyond the boundaries of the administrative system. This
1621 **MAY** be done with either **TTL** or administrative scopes (see [17]), depending on what is implemented in the
1622 network. SIP user agents **MAY** listen to that address and use it to become aware of the location of other local
1623 users (see [18]); however, they do not respond to the request.

1624 Multicast registration may be inappropriate in some environments, for example, if multiple businesses share the
1625 same local area network.

1626 **10.2.7 Transmitting a Request**

1627 Once the **REGISTER** method has been constructed, and the destination of the message identified, UACs
1628 should follow the procedures described in Section 8.1.2 to hand off the **REGISTER** to the transaction layer.

1629 If the transaction layer returns a timeout error because the REGISTER yielded no response, the UAC
1630 SHOULD wait some reasonable time interval before re-attempting a registration to the same registrar; no
1631 specific interval is mandated.

1632 10.2.8 Error Responses

1633 If a UA receives a 423 (Registration Too Brief) response, it MAY retry the registration after making the
1634 expiration interval of all contact addresses in the REGISTER request equal to or greater than the expiration
1635 interval within the Min-Expires header of the 423 (Registration Too Brief) response.

1636 10.3 Processing REGISTER Requests

1637 A registrar is a UAS that responds to REGISTER requests and maintains a list of bindings that are accessible
1638 to proxy servers within its administrative domain. A registrar handles requests according to Section 8.2 and
1639 Section 17.2, but it accepts only REGISTER requests. A registrar does not generate 6xx responses. If a
1640 registrar listens at a multicast interface, it MAY redirect multicast REGISTER requests to its own unicast
1641 interface with a 302 (Moved Temporarily) response.

1642 A REGISTER request MUST NOT contain Record-Route or Route header fields; registrars MUST
1643 ignore them if they appear.

1644 A registrar must know (e.g., through configuration) the set of domain(s) for which it maintains bindings.
1645 REGISTER requests MUST be processed by a registrar in the order that they are received. REGISTER
1646 requests MUST also be processed atomically, meaning that REGISTER requests are either processed com-
1647 pletely or not at all. Each REGISTER message must be processed independently of any other registration
1648 or binding changes.

1649 When receiving a REGISTER request, a registrar follows these steps:

- 1650 1. The registrar inspects the Request-URI to determine whether it has access to bindings for the domain
1651 identified in the Request-URI. If not and if the server also acts as a proxy server, the server SHOULD
1652 forward the request to the addressed domain, following the general behavior for proxying messages
1653 described in Section 16.
- 1654 2. To guarantee that the registrar supports any necessary extensions, the registrar processes Require
1655 header fields as described for UASs in Section 8.2.2.
- 1656 3. A registrar SHOULD authenticate the UAC. Mechanisms for the authentication of SIP user agents are
1657 described in Section 20; registration behavior in no way overrides the generic authentication frame-
1658 work for SIP. If no authentication mechanism is available, the registrar MAY take the From address as
1659 the asserted identity of the originator of the request.
- 1660 4. The registrar SHOULD determine if the authenticated user is authorized to modify registrations for
1661 this address-of-record. For example, a registrar might consult a authorization database that maps user
1662 names to a list of addresses-of-record for which this identity is authorized to modify bindings. If not,
1663 the registrar returns 403 (Forbidden) and skips the remaining steps.

1664 In architectures that support third-party registration, one entity may be responsible for updating the regis-
1665 trations associated with multiple addresses-of-record.

1666 5. The registrar extracts the address-of-record from the **To** header field of request. If the address-of-
1667 record is not valid for the domain in the **Request-URI**, the registrar sends a 404 (Not Found) response
1668 and skips the remaining steps. The URI **MUST** then converted to a canonical form. To do that, all URI
1669 parameters are removed (including the user param), and any escaped characters are converted to their
1670 unescaped form. The result serves as an index into the list of bindings.

1671 6. The registrar checks whether the request contains any **Contact** header fields. If not, it skips to the last
1672 step.

1673 Next, the registrar checks if there is one **Contact** field that contains the special value “*” and a
1674 **Expires** field. If the request has additional **Contact** fields or an expiration time other than zero, the
1675 request is invalid and the server returns 400 (Invalid Request) and skips the remaining steps. If not, the
1676 registrar checks whether the **Call-ID** agrees with the value stored for each binding. If not, it removes
1677 the binding. If it does agree, it only removes the binding if the **CSeq** in the request is higher than the
1678 value stored for that binding and leaves the binding as is otherwise. It then skips to the last step.

1679 7. The registrar now processes each contact address in the **Contact** header field in turn. For each address,
1680 it determines the expiration interval as follows:

- 1681 • If the field value has an “expires” parameter, that value is used.
- 1682 • If there is no such parameter, but the request has an **Expires** header field, that value is used.
- 1683 • If there is neither, a locally-configured default value is used.

1684 The registrar **MAY** shorten the expiration interval. If and only if the expiration interval is greater than
1685 zero **AND** smaller than one hour **AND** less than a registrar-configured minimum, the registrar **MAY**
1686 reject the registration with a response of 423 (Registration Too Brief). This response **MUST** contain a
1687 **Min-Expires** header field that states the minimum expiration interval the registrar is willing to honor.
1688 It then skips the remaining steps.

1689 Allowing the registrar to set the registration interval protects it against excessively frequent registration
1690 refreshes while limiting the state that it needs to maintain and decreasing the likelihood of registrations going
1691 stale. The expiration interval of a registration is frequently used in the creation of services. An example is a
1692 follow-me service, where the user may only be available at a terminal for a brief period. Therefore, registrars
1693 should accept brief registrations; a request should only be rejected if the interval is so short that the refreshes
1694 would degrade registrar performance.

1695 For each address, it then searches the list of current bindings using the URI comparison rules. If
1696 the binding does not exist, it is tentatively added. If the binding does exist, the registrar checks the
1697 **Call-ID** value. If the existing binding has the same **Call-ID** value differs from the request, the binding
1698 is removed if the expiration time is zero and updated otherwise. If they are the same, the registrar
1699 compares the **CSeq** value. If the value is higher than that of the existing binding, it updates or
1700 removes the binding as above. If not, the update is aborted and the request fails.

1701 This algorithm ensures that out-of-order requests from the same UA are ignored.

1702 Each binding record records the **Call-ID** and **CSeq** values from the request.

1703 The binding updates are committed (i.e., made visible to the proxy) if and only if all binding updates
1704 and additions succeed. If any one of them fails, the request fails with 500 (Server Error) response and
1705 all tentative binding updates are removed.

- 1706 8. The registrar returns a 200 (OK) response. The response MUST contain Contact header fields enu-
1707 merating all current bindings. Each Contact value MUST feature an "expires" parameter indicating
1708 its expiration interval chosen by the registrar. The response SHOULD include a Date header field.

1709 11 Querying for Capabilities

1710 The SIP method OPTIONS allows a UA to query another UA or a proxy server as to its capabilities. This
1711 allows a client to discover information about the methods, content types, extensions, codecs etc. supported
1712 without actually "ringing" the other party. For example, before a client inserts a Require header field into
1713 an INVITE listing an option that it is not certain the destination UAS supports, the client can query the
1714 destination UAS with an OPTIONS to see if this option is returned in a Supported header field.

1715 The target of the OPTIONS request is identified by the Request-URI, which could identify another
1716 User Agent or a SIP Server. If the OPTIONS is addressed to a proxy server, the Request-URI is set
1717 without a user part, similar to the way a Request-URI is set for a REGISTER request. Alternatively, a
1718 server receiving an OPTIONS request with a Max-Forwards header value of 0 MAY respond to the request
1719 regardless of the Request-URI.

1720 This behavior is common with HTTP/1.1. This behavior can be used as a "traceroute" functionality to check the
1721 capabilities of individual hop servers by sending a series of OPTIONS requests with incremented Max-Forwards
1722 values.

1723 As is the case for general UA behavior, the transaction layer can return a timeout error if the OPTIONS
1724 yields no response. This may indicate that the target is unreachable and hence unavailable.

1725 An OPTIONS request MAY be sent as part of an established dialog to query the peer on capabilities that
1726 may be utilized later in the dialog.

1727 11.1 Construction of OPTIONS Request

1728 An OPTIONS request is constructed using the standard rules for a SIP request as discussed Section 8.1.1.

1729 A Contact header field MAY be present in an OPTIONS.

1730 An Accept header field SHOULD be included to indicate the type of message body the UAC wishes to
1731 receive in the response. Typically, this is set to a format that is used to describe the media capabilities of a
1732 UA, such as SDP (application/sdp).

1733 The response to an OPTIONS request is assumed to be scoped to the Request-URI in the original
1734 request. However, only when an OPTIONS is sent as part of an established dialog is it guaranteed that
1735 future requests will be received by the server which generated the OPTIONS response.

1736 Example OPTIONS request:

```
1737 OPTIONS sip:carol@chicago.com SIP/2.0
1738 Via: SIP/2.0/UDP 192.0.2.4;branch=z9hG4bKhjhs8ass877
1739 To: <sip:carol@chicago.com>
1740 From: Alice <sip:alice@atlanta.com>;tag=1928301774
1741 Call-ID: a84b4c76e66710
1742 CSeq: 63104 OPTIONS
1743 Contact: <sip:alice@192.0.2.4>
1744 Accept: application/sdp
1745 Content-Length: 0
```

1746 11.2 Processing of OPTIONS Request

1747 The response to an OPTIONS is constructed using the standard rules for a SIP response as discussed in
1748 Section 8.2.6. The response code chosen is the same that would have been chosen had the request been an
1749 INVITE. That is, a 200 (OK) would be returned if the UAS is ready to accept a call, a 486 (Busy Here)
1750 would be returned if the UAS is busy, etc. This allows an OPTIONS request to be used to determine the
1751 basic state of a UAS, which can be an indication of whether the UAC will accept an INVITE request.

1752 An OPTIONS request received within a dialog generates a 200 (OK) response which is identical to
1753 one constructed outside a dialog and does not have any impact on the dialog. This use of OPTIONS
1754 has limitations due the differences in proxy handling of OPTIONS and INVITE requests. While a forked
1755 INVITE can result in multiple 200 (OK) responses being returned, a forked OPTIONS will only result in a
1756 single 200 (OK) response, since it is treated by proxies using the non-INVITE handling. See Section 13.2.1
1757 for the normative details.

1758 If the response to an OPTIONS is generated by a proxy server, the proxy returns a 200 (OK) listing the
1759 capabilities of the server. The response does not contain a message body.

1760 Allow, Accept, Accept-Encoding, Accept-Language, and Supported header fields SHOULD be
1761 present in a 200 (OK) response to an OPTIONS request. If the response is generated by a proxy, the
1762 Allow header field SHOULD be omitted as it is ambiguous since a proxy is method agnostic.

1763 Contact header fields MAY be present in a 200 (OK) response and have the same semantics as in a
1764 redirect. That is, they may list a set of alternative names and methods of reaching the user.

1765 A Warning header field MAY be present.

1766 A message body MAY be sent, the type of which is determined by the Accept header in the OPTIONS
1767 request (application/sdp if the Accept header was not present). If the types include one that can describe
1768 media capabilities, the UA SHOULD include a body in the response for that purpose. Details on construction
1769 of such a body in the case of application/sdp are described in [19].

1770 Example OPTIONS response generated by a UAS (corresponding to the request in Section 11.1):

```
1771 SIP/2.0 200 OK
1772 Via: SIP/2.0/UDP 192.0.2.4;branch=z9hG4bKhjhs8ass877
1773 To: <sip:carol@chicago.com>;tag=93810874
1774 From: Alice <sip:alice@atlanta.com>;tag=1928301774
1775 Call-ID: a84b4c76e66710@100.1.3.3
1776 CSeq: 63104 OPTIONS
1777 Contact: <sip:carol@chicago.com>
1778 Contact: <mailto:carol@chicago.com>
1779 Allow: INVITE, ACK, CANCEL, OPTIONS, BYE
1780 Accept: application/sdp
1781 Accept-Encoding: gzip
1782 Accept-Language: en
1783 Supported: foo
1784 Content-Type: application/sdp
1785 Content-Length: 274
1786
1787 (SDP not shown)
```

1788 12 Dialogs

1789 A key concept for a user agent is that of a dialog. A dialog represents a peer-to-peer SIP relationship between
1790 a two user agents that persists for some time. The dialog facilitates sequencing of messages between the
1791 user agents and proper routing of requests between both of them. The dialog represents a context in which to
1792 interpret SIP messages. Section 8 discussed method- independent UA processing for requests and responses
1793 outside of a dialog. This section discusses how those requests and responses are used to construct a dialog,
1794 and then how subsequent requests and responses are sent within a dialog.

1795 A dialog is identified at each UA with a dialog ID, which consists of a Call-ID value, a local URI and
1796 local tag (together called the local address), and a remote URI and remote tag (together called the remote
1797 address). The dialog ID at each UA involved in the dialog is not the same. Specifically, the local URI and
1798 local tag at one UA are identical to the remote URI and remote tag at the peer UA. The tags are opaque
1799 tokens that facilitate the generation of unique dialog IDs.

1800 A dialog ID is also associated with all responses and with any request that contains a tag in the To field.
1801 The rules for computing the dialog ID of a message depend on whether the entity is a UAC or UAS. For a
1802 UAC, the Call-ID value of the dialog ID is set to the Call-ID of the message, the remote address is set to the
1803 To field of the message, and the local address is set to the From field of the message (these rules apply to
1804 both requests and responses). As one would expect, for a UAS, the Call-ID value of the dialog ID is set to
1805 the Call-ID of the message, the remote address is set to the From field of the message, and the local address
1806 is set to the To field of the message.

1807 A dialog contains certain pieces of state needed for further message transmissions within the dialog.
1808 This state consists of the dialog ID, a local sequence number (used to order requests from the UA to its
1809 peer), a remote sequence number (used to order requests from its peer to the UA), and a route set, which is
1810 an ordered list of URIs. The *route set* is the set of servers that need to be traversed to send a request to the
1811 peer. A dialog can also be in the “early” state, which occurs when it is created with a provisional response,
1812 and then transition to the “confirmed” state when the final response comes.

1813 12.1 Creation of a Dialog

1814 Dialogs are created through the generation of non-failure responses to requests with specific methods.
1815 Within this specification, only 2xx and 101-199 responses with a To tag to INVITE establish a dialog.
1816 A dialog established by a non-final response to a request is in the “early” state and it is called an early dia-
1817 log. Extensions MAY define other means for creating dialogs. Section 13 gives more details that are specific
1818 to the INVITE method. Here, we describe the process for creation of dialog state that is not dependent on
1819 the method.

1820 A dialog is identified by a dialog ID. A dialog ID consists of three components, namely a call identifier
1821 component, a local address component and a remote address component. UAs MUST assign values to these
1822 components as described below.

1823 12.1.1 UAS behavior

1824 When a UAS responds to a request with a response that establishes a dialog (such as a 2xx to INVITE), the
1825 UAS MUST copy all Record-Route headers from the request into the response (including the URIs, URI
1826 parameters, and any Record-Route header parameters, whether they are known or unknown to the UAS)
1827 and MUST maintain the order of those headers. The UAS MUST add a Contact header field to the response.
1828 The Contact header field contains an address where the UAS would like to be contacted for subsequent

1829 requests in the dialog (which includes the ACK for a 2xx response in the case of an INVITE). Generally,
1830 the host portion of this URI is the IP address or FQDN of the host. The URI provided in the Contact header
1831 field MUST be a SIP URI and have global scope (i.e., the same SIP URI can be used outside this dialog to
1832 contact the UAS). The same way, the scope of the SIP URI in the Contact header field of the INVITE is not
1833 limited to this dialog either. It can therefore be used to contact the UAC even outside this dialog.

1834 The UAS then constructs the state of the dialog. This state MUST be maintained for the duration of the
1835 dialog. First, the *route set* MUST be computed by following these steps:

- 1836 1. The list of URIs in the Record-Route headers in the request, if present, are taken, including any URI
1837 parameters.
- 1838 2. The URI in the Contact header from the request if present, is taken, including any URI parameters.
1839 The URI is appended to the bottom of the list of URIs from the previous step.

1840 Contact was not mandatory in RFC 2543. Thus, if the UAS is communicating with an older UAC, the
1841 UAC might not have inserted the Contact header field.

- 1842 3. The resulting list of URIs is called the *route set*.

1843 These rules clearly imply that a UA MUST be able to parse and process Record-Route header fields. This is a
1844 change from RFC 2543, where all record-route and route processing was optional for user agents.

1845 It is possible for the *route set* to be empty. This will occur if neither Record-Route headers nor a
1846 Contact header were present in the request. The UAS MUST also remember whether the bottom-most entry
1847 in the *route set* was constructed from a Contact header. This is effectively a boolean value, which we refer
1848 to as CONTACT_SET. From this value the UA can determine whether the bottom-most value can be updated
1849 from subsequent requests; if it was constructed from a Contact, it can be updated.

1850 The remote sequence number MUST be set to the value of the sequence number in the Cseq header of
1851 the request. The local sequence number MUST be empty. The call identifier component of the dialog ID
1852 MUST be set to the value of the Call-ID in the request. The local address component of the dialog ID MUST
1853 be set to the To field in the response to the request (which therefore includes the tag), and the remote address
1854 component of the dialog ID MUST be set to the From field in the request. A UAS MUST be prepared to
1855 receive a request without a tag in the From field, in which case the tag is considered to have a value of null.

1856 This is to maintain backwards compatibility with RFC 2543, which did not mandate From tags.

1857 12.1.2 UAC behavior

1858 When a UAC receives a response that establishes a dialog, it constructs the state of the dialog. This state
1859 MUST be maintained for the duration of the dialog. First, the *route set* MUST be computed by following
1860 these steps:

- 1861 1. The list of URIs present in the Record-Route headers in the response are taken, if present, including
1862 all URI parameters, and their order is reversed.
- 1863 2. The URI in the Contact header from the response, if present, is taken, including all URI parameters,
1864 and appended to the end of the list from the previous step.
- 1865 3. The list of URIs resulting from the above two operations is referred to as the *route set*.

1866 It is possible for the *route set* to be empty. This will occur if neither **Record-Route** headers nor a
1867 **Contact** header were present in the response. The UAC **MUST** also remember whether the bottom-most
1868 entry in the *route set* was constructed from a **Contact** header. This is effectively a boolean value, which we
1869 refer to as **CONTACT_SET**. From this value the UA can determine whether the bottom-most value can be
1870 updated from subsequent requests; if it was constructed from a **Contact**, it can be updated.

1871 The local sequence number **MUST** be set to the value of the sequence number in the **Cseq** header of the
1872 request. The remote sequence number **MUST** be empty (it is established when the UA sends a request within
1873 the dialog). The call identifier component of the dialog ID **MUST** be set to the value of the **Call-ID** in the
1874 request. The local address component of the dialog ID **MUST** be set to the **From** field in the request, and
1875 the remote address component of the dialog ID **MUST** be set to the **To** field of the response. A UAC **MUST**
1876 be prepared to receive a response without a tag in the **To** field, in which case the tag is considered to have a
1877 value of null.

1878 This is to maintain backwards compatibility with RFC 2543, which did not mandate **To** tags.

1879 **12.2 Requests within a Dialog**

1880 Once a dialog has been established between two UAs, either of them **MAY** initiate new transactions as needed
1881 within the dialog. However, a dialog imposes some restrictions on the use of simultaneous transactions.

1882 A TU **MUST NOT** initiate a new regular transaction within a dialog while a regular transaction is in
1883 progress (in either direction) within that dialog. If there is a non-**INVITE** client or server transaction in
1884 progress the TU **MUST** wait until this transaction enters the completed or the terminated state to initiate the
1885 new transaction.

1886 OPEN ISSUE #113: Should we relax the constraint on non-overlapping regular transactions?

1887 A route refresh request sent within a dialog is defined as a request that can modify the *route set* of
1888 the dialog. For dialogs that have been established with an **INVITE**, the only route refresh request defined
1889 is re-**INVITE** (see Section 14). Other extensions may define different route refresh requests for dialogs
1890 established in other ways.

1891 Note that an **ACK** is *NOT* a route refresh request.

1892 **12.2.1 UAC Behavior**

1893 **12.2.1.1 Generating the Request** A request within a dialog is constructed by using many of the com-
1894 ponents of the state stored as part of the dialog.

1895 The **To** header field of the request **MUST** be set to the remote address, and the **From** header field **MUST**
1896 be set to the local address (both including tags, assuming the tags are not null).

1897 The **Call-ID** of the request **MUST** be set to the **Call-ID** of the dialog. Requests within a dialog **MUST**
1898 contain strictly monotonically increasing and contiguous **CSeq** sequence numbers (increasing-by-one) in
1899 each direction. Therefore, if the local sequence number is not empty, the value of the local sequence number
1900 **MUST** be incremented by one, and this value **MUST** be placed into the **Cseq** header. If the local sequence
1901 number is empty, an initial value **MUST** be chosen using the guidelines of Section 8.1.1.5. The method field
1902 in the **Cseq** header **MUST** match the method of the request.

1903 With a length of 32 bits, a client could generate, within a single call, one request a second for about 136 years
1904 before needing to wrap around. The initial value of the sequence number is chosen so that subsequent requests within
1905 the same call will not wrap around. A non-zero initial value allows clients to use a time-based initial sequence
1906 number. A client could, for example, choose the 31 most significant bits of a 32-bit second clock as an initial
1907 sequence number.

1908 The Request-URI of requests is determined according to the following rules:

1909 The UAC takes the list of URI in the *route set*. The top URI MUST be inserted into the Request-URI
1910 of the request, including all URI parameters. Any URI parameters not allowed in the Request-URI MUST
1911 then be stripped. Each of the remaining URIs (if any) from the *route set*, including all URI parameters,
1912 MUST be placed into a Route header field into the request, in order.

1913 A TU SHOULD follow the rules just mentioned to build the Request-URI of the request, regardless of
1914 whether the UA uses an outbound proxy server or not. However, in some instances, a UA may not be willing
1915 or capable of sending the request to the top element in the *route set*. One example is a UA that is not capable
1916 of DNS, and therefore may not be able to follow those procedures. to use a loose-routing policy to send
1917 the request to its outbound proxy server (see section 8.1.3). This policy MUST include placing the topmost
1918 element in the *route set* as the first value in the message's Route header field as well as in the Request-
1919 URI. The loop-detection avoidance algorithm described in section 8.1.3 SHOULD be applied to the message
1920 before sending.

1921 A UAC SHOULD include a Contact header in any route refresh requests within a dialog, and unless
1922 there is a need to change it, the URI SHOULD be the same as used in previous requests within the dialog. As
1923 discussed in Section 12.2.2, a Contact header in a route refresh request updates the *route set*. This allows a
1924 UA to provide a new contact address, should its address change during the duration of the dialog.

1925 However, requests that are not route refresh requests do not affect the *route set* for the dialog.

1926 Once the request has been constructed, the address of the server is computed and the request is sent,
1927 using the same procedures for requests outside of a dialog (Section 8.1.1).

1928 **12.2.1.2 Processing the Responses** The UAC will receive responses to the request from the transaction
1929 layer. If the client transaction returns a timeout this is treated as a 408 (Request Timeout) response.

1930 The behavior of a UAC that receives a 3xx response for a request sent within a dialog is the same as if
1931 the request had been sent outside a dialog. This behavior is described in Section 13.2.2.

1932 Note, however, that when the UAC tries alternative locations, it still uses the *route set* for the dialog to build the
1933 Route header of the request.

1934 If a UAC has a *route set* for a dialog and receives a 2xx response to a route refresh it sent, the Contact
1935 header field of the response is examined. If not present, the *route set* remains unchanged. If the response had
1936 a Contact header field, and the boolean variable CONTACT_SET is false, the URI in the Contact header
1937 field in the response is added to the bottom of the *route set*, and CONTACT_SET is set to true. If the route
1938 refresh request response had a Contact header field, and CONTACT_SET is true, the URI in the Contact
1939 header field of the response to the route refresh request replaces the bottom value in the *route set*. If a route
1940 refresh request is responded with a non-2xx final response the *route set* remains unchanged as if no route
1941 refresh request had been issued.

1942 If the response for the a request within a dialog is a 481 (Call/Transaction Does Not Exist) or a 408
1943 (Request Timeout), the UAC SHOULD terminate the dialog. A UAC SHOULD also terminate a dialog if no
1944 response at all is received for the request (the client transaction would inform the TU about the timeout.)

1945 For INVITE initiated dialogs, terminating the dialog consists of sending a BYE.

1946 12.2.2 UAS behavior

1947 Requests sent within a dialog, as any other requests, are atomic. If a particular request is accepted by the
1948 UAS, *all* the state changes associated with it are performed. If the request is rejected, *none* of the state
1949 changes is performed.

1950 Note that some requests such as INVITEs affect several pieces of state.

1951 The UAS will receive the request from the transaction layer. If the request has a tag in the **To** header
1952 field, the UAS core computes the dialog identifier corresponding to the request and compares it with existing
1953 dialogs. If there is a match, this is a mid-dialog request. In that case, the UAS applies the same processing
1954 rules for requests outside of a dialog, discussed in Section 8.2.

1955 If the request has a tag in the **To** header field, but the dialog identifier does not match any existing di-
1956 alogs, the UAS may have crashed and restarted, or it may have received a request for a different (possibly
1957 failed) UAS (the UASs can construct the **To** tags so that a UAS can identify that the tag was for a UAS
1958 for which it is providing recovery). Another possibility is that the incoming request has been simply mis-
1959 souted. Based on the **To** tag, the UAS MAY either accept or reject the request. Accepting the request for
1960 acceptable **To** tags provides robustness, so that dialogs can persist even through crashes. UAs wishing to
1961 support this capability must take into consideration some issues such as choosing monotonically increasing
1962 **CSeq** sequence numbers even across reboots, reconstructing the *route set*, and accepting out-of-range RTP
1963 timestamps and sequence numbers.

1964 If the UAS wishes to reject the request, because it does not wish to recreate the dialog, it MUST respond
1965 to the request with a 481 (Call/Transaction Does Not Exist) status code and pass that to the server transaction.

1966
1967 Requests that do not change in any way the state of a dialog may be received within a dialog (for
1968 example, an **OPTIONS** request). They are processed as if they had been received outside the dialog.

1969 Requests within a dialog MAY contain **Record-Route** and **Contact** header fields. However, requests
1970 that are not route refresh requests do not update the *route set* for the dialog. This specification only defines
1971 one route refresh request: re-**INVITE** (see Section 14).

1972 Special rules apply when updated **Record-Route** or **Contact** header fields are received inside a route
1973 refresh request. If a UAS has a *route set* for a dialog and receives a route refresh for that dialog containing
1974 **Record-Route** header fields, it MUST copy those header fields into any 2xx response to that request. If the
1975 boolean variable **CONTACT_SET** is true, the **Contact** header field in the request (if present) replaces the
1976 last entry in the *route set*. If the boolean variable **CONTACT_SET** is false, the UAS MUST add the URI in the
1977 **Contact** header field in the route refresh request to the bottom of the *route set*, and then set **CONTACT_SET**
1978 to true. If the request did not contain a **Contact** header field, the route-set at the UAS remains unchanged.

1979 Route refresh requests only update the **Contact** of the *route set* and not the elements formed from **Record-**
1980 **Route**. Updating the latter would introduce severe backwards compatibility problems with RFC 2543-compliant
1981 systems.

1982 If the remote sequence number is empty, it MUST be set to the value of the sequence number in the
1983 **CSeq** header in the request. If the remote sequence number was not empty, but the sequence number of the
1984 request is lower than the remote sequence number, the request is out of order and MUST be rejected with
1985 a 500 (Server Internal Error) response. If the remote sequence number was not empty, and the sequence
1986 number of the request is greater than the remote sequence number, the request is in order. It is possible
1987 for the **CSeq** header to be higher than the remote sequence number by more than one. This is not an error
1988 condition, and a UAS SHOULD be prepared to receive and process requests with **CSeq** values more than
1989 one higher than the previous received request. The UAS MUST then set the remote sequence number to the
1990 value of the sequence number in the **CSeq** header in the request.

1991 If a proxy challenges a request generated by the UAC, the UAC has to resubmit the request with credentials. The
1992 resubmitted request will have a new **Cseq** number. The UAS will never see the first request, and thus, it will notice
1993 a gap in the **Cseq** number space. Such a gap does not represent any error condition.

1994 **12.3 Termination of a Dialog**

1995 Dialogs can end in several different ways, depending on the method. When a dialog is established with
1996 INVITE, it is terminated with a BYE. No other means to terminate a dialog are described in this specification,
1997 but extensions can define other ways.

1998 **13 Initiating a Session**

1999 **13.1 Overview**

2000 When a user agent client desires to initiate a session (for example, audio, video, or a game), it formulates
2001 an INVITE request. The INVITE request asks a server to establish a session. This request is forwarded by
2002 proxies, eventually arriving at one or more UAS that can potentially accept the invitation. These UASs will
2003 frequently need to query the user about whether to accept the invitation. After some time, those UAS can
2004 accept the invitation (meaning the session is to be established) by sending a 2xx response. If the invitation
2005 is not accepted, a 3xx, 4xx, 5xx or 6xx response is sent, depending on the reason for the rejection. Before
2006 sending a final response, the UAS can also send a provisional response (1xx), either reliably or unreliably,
2007 to advise the UAC of progress in contacting the called user.

2008 After possibly receiving one or more provisional responses, the UA will get one or more 2xx responses or
2009 one non-2xx final response. Because of the protracted amount of time it can take to receive final responses
2010 to INVITE, the reliability mechanisms for INVITE transactions differ from those of other requests (like
2011 OPTIONS). Once it receives a final response, the UAC needs to send an ACK for every final response it
2012 receives. The procedure for sending this ACK depends on the type of response. For final responses between
2013 300 and 699, the ACK processing is done in the transaction layer and follows one set of rules (See Section
2014 17). For 2xx responses, the ACK is generated by the UAC core.

2015 A 2xx response to an INVITE establishes a session, and it also creates a dialog between the UA that
2016 issued the INVITE and the UA that generated the 2xx response. Therefore, when multiple 2xx responses are
2017 received from different remote UAs (because the INVITE forked), each 2xx establishes a different dialog.
2018 All these dialogs are part of the same call.

2019 This section provides details on the establishment of a session using INVITE.

2020 **13.2 Caller Processing**

2021 **13.2.1 Creating the Initial INVITE**

2022 Since the initial INVITE represents a request outside of a dialog, its construction follows the procedures of
2023 Section 8.1.1. Additional processing is required for the specific case of INVITE.

2024 An Allow header field (Section 24.5) SHOULD be present in the INVITE. It indicates what methods can
2025 be invoked within a dialog, on the UA sending the INVITE, for the duration of the dialog. For example, a
2026 UA capable of receiving INFO requests within a dialog [20] SHOULD include an Allow header listing the
2027 INFO method.

2028 A Supported header field (Section 24.39) SHOULD be present in the INVITE. It enumerates all the
2029 extensions understood by the UAC.

2030 An Accept (Section 24.1) header field MAY be present in the INVITE. It indicates which content-types
2031 are acceptable to the UA, in both the response received by it, and in any subsequent requests sent to it within
2032 dialogs established by the INVITE. The Accept header is especially useful for indicating support of various

2033 session description formats.

2034 The UA MAY add an Expires header field (Section 24.19) to limit the validity of the invitation. If the
2035 time indicated in the Expires header field is reached and no final answer for the INVITE has been received
2036 the UAC core SHOULD generate a CANCEL request for the original INVITE.

2037 A UAC MAY also find useful to add, among others, Subject (Section 24.38), Organization (Section
2038 24.25) and User-Agent (Section 24.43) header fields. They all contain information related to the INVITE.

2039 The UAC MAY choose to add a message body to the INVITE. Section 8.1.1.10 deals with how to con-
2040 struct the header fields – Content-Type among others – needed to describe the message body.

2041 There are special rules for message bodies that contain a session description - their corresponding
2042 Content-Disposition is “session”. SIP uses an offer/answer model where one UA sends a session de-
2043 scription, called the offer, which contains a proposed description of the session. The offer indicates the
2044 desired communications means (audio, video, games), parameters of those means (such as codec types) and
2045 addresses for receiving media from the answerer. The other UA responds with another session description,
2046 called the answer, which indicates which communications means are accepted, the parameters which apply
2047 to those means, and addresses for receiving media from the offerer. The offer/answer model can be mapped
2048 into the INVITE transaction in two ways. The first, which is the most intuitive, is that the INVITE contains
2049 the offer, the 2xx response contains the answer, and no session description is provided in the ACK. In this
2050 model, the UAC is the offerer, and the UAS is the answerer. A second model is that the INVITE contains no
2051 session description, the 2xx response contains the offer, and the ACK contains the answer. In this model, the
2052 UAS is the offerer, and the UAC is the answerer. The second model is useful for gateways from H.323v1
2053 to SIP, where the H.323 media characteristics are not known until the call is established. This is also useful
2054 for sessions that use third-party call control. As a result of these models, if the INVITE contains a session
2055 description, the ACK MUST NOT contain one. Conversely, if the caller chooses to omit the session descrip-
2056 tion in the INVITE, the ACK MUST contain one (if a 2xx response is received). 2xx responses to an INVITE
2057 MUST always contain a session description. All user agents that support INVITE MUST support both models.

2058 The Session Description Protocol (SDP) [5] MUST be supported by all user agents as a means to describe
2059 sessions, and its usage for construction offers and answers MUST follow the procedures defined in [19].

2060 The restrictions of the offer-answer model (session description only in the INVITE OR in the ACK,
2061 but not in both) just described only apply to bodies whose Content-Disposition header field is “session”.
2062 Therefore, it is possible that both the INVITE and the ACK contain a body message (e.g., the INVITE carries
2063 a photo (Content-Disposition: render) and the ACK a session description (Content-Disposition: session)
2064).

2065 If the Content-Disposition header field is missing, bodies of Content-Type application/sdp imply the
2066 disposition “session”, while other content types imply “render”.

2067 Once the INVITE has been created, the UAC follows the procedures defined for sending requests outside
2068 of a dialog (Section 8). This results in the construction of a client transaction that will ultimately send the
2069 request and deliver responses to the UAC.

2070 13.2.2 Processing INVITE Responses

2071 Once the INVITE has been passed to the INVITE client transaction, the UAC waits for responses for the IN-
2072 VITE. Responses are matched to their corresponding INVITE because they have the same Call-ID, the same
2073 From header field, the same To header field, excluding the tag, and the same CSeq. Rules for comparisons
2074 of these headers are described in Section 24. If the INVITE client transaction returns a timeout rather than a
2075 response the TU acts as if a 408 (Request Timeout) response had been received.

2076 **13.2.2.1 1xx responses** Zero, one or multiple provisional responses may arrive before one or more
2077 final responses are received. Provisional responses for an INVITE request can create “early dialogs”. If a
2078 provisional response has a tag in the To field, and if the dialog ID of the response does not match an existing
2079 dialog, one is constructed using the procedures defined in Section 12.1.2.

2080 The early dialog will only be needed if the UAC needs to send a request to its peer within the dialog be-
2081 fore the initial INVITE transaction completes. Header fields present in a provisional response are applicable
2082 as long as the dialog is in the early state (e.g., an Allow header field in a provisional response contains the
2083 methods that can be used in the dialog while this is in the early state).

2084 **13.2.2.2 3xx responses** A 3xx response may contain a Contact header field providing new addresses
2085 where the callee might be reachable. Depending on the status code of the 3xx response (see Section 25.3)
2086 the UAC MAY choose to try those new addresses.

2087 **13.2.2.3 4xx, 5xx and 6xx responses** A single non-2xx final response may be received for the IN-
2088 VITE. 4xx, 5xx and 6xx responses may contain a Contact header field indicating the location where addi-
2089 tional information about the error can be found.

2090 All early dialogs are considered terminated upon reception of the non-2xx final response.

2091 After having received the non-2xx final response the UAC core considers the INVITE transaction com-
2092 pleted. The INVITE client transaction handles generation of ACKs for the response (see Section 17).

2093 **13.2.2.4 2xx responses** Multiple 2xx responses may arrive at the UAC for a single INVITE request
2094 due to a forking proxy. Each response is distinguished by the tag parameter in the To header field, and each
2095 represents a distinct dialog, with a distinct dialog identifier.

2096 If the dialog identifier in the 2xx response matches the dialog identifier of an existing dialog, the dialog
2097 MUST be transitioned to the “confirmed” state, and the route set for the dialog MUST be recomputed based
2098 on the 2xx response using the procedures of Section 12.1.2. Otherwise, a new dialog in the “confirmed”
2099 state is constructed in the same fashion.

2100 The route set only is recomputed for backwards compatibility. RFC 2543 did not mandate mirroring of Record-
2101 Route headers in a 1xx, only 2xx. However, we cannot update the entire state of the dialog, since mid-dialog
2102 requests may have been sent within the early call leg, modifying the sequence numbers, for example.

2103 The UAC core MUST generate an ACK request for each 2xx received from the transaction layer. The
2104 header fields of the ACK are constructed in the same way as for any request sent within a dialog (see Section
2105 12) with the exception of the CSeq and the header fields related to authentication. The sequence number
2106 of the CSeq header field MUST be the same as the INVITE being acknowledged, but the CSeq method
2107 MUST be ACK. The ACK MUST contain the same credentials as the INVITE. If the INVITE did not contain
2108 an offer, the 2xx will contain one, and therefore the ACK MUST carry an answer in its body. If the offer in
2109 the 2xx response is not acceptable the UAC core MUST generate a valid answer in the ACK and then send a
2110 BYE immediately.

2111 Once the ACK has been constructed, the procedures of [8] are used to determine the destination address,
2112 port and transport. However, the request is passed to the transport layer directly for transmission, rather than
2113 a client transaction. This is because the UAC core handles retransmissions of the ACK, not the transaction
2114 layer. The ACK MUST be passed to the client transport every time a retransmission of the 2xx final response
2115 that triggered the ACK arrives.

2116 The UAC core considers the INVITE transaction completed 64*T1 seconds after the reception of the
2117 first 2xx response. At this point all the early dialogs that have not transitioned to established dialogs are
2118 terminated. Once the INVITE transaction is considered completed by the UAC core, no more new 2xx
2119 responses are expected to arrive.

2120 If, after acknowledging any 2xx response to an INVITE, the caller does not want to continue with that
2121 dialog, then the caller MUST terminate the dialog by sending a BYE request as described in Section 15.

2122 **13.3 Callee Processing**

2123 **13.3.1 Processing of the INVITE**

2124 The UAS core will receive INVITE requests from the transaction layer. It first performs the request process-
2125 ing procedures of Section 8.2, which are applied for both requests inside and outside of a dialog.

2126 Assuming these processing states complete without generating a response, the UAS core performs the
2127 additional processing steps:

- 2128 1. If the request is an INVITE that contains an Expires header field the UAS core inspects this header
2129 field. If the INVITE has already expired a 487 (Request Terminated) response SHOULD be generated.
2130 In any case, if the INVITE expires before the UAS has generated a final response a 487 (Request
2131 Terminated) response SHOULD be generated.
- 2132 2. If the request is a mid-dialog request, the method-independent processing described in Section 12.2.2
2133 is first applied. It might also modify the session; Section 14 provides details.
- 2134 3. If the request has a tag in the To header field but the dialog identifier does not match any of the existing
2135 dialogs, the UAS may have crashed and restarted, or may have received a request for a different
2136 (possibly failed) UAS. Section 12.2.2 provides guidelines to achieve a robust behaviour under such a
2137 situation.

2138 Processing from here forward assumes that the INVITE is outside of a dialog, and is thus for the purposes
2139 of establishing a new session.

2140 The INVITE may contain a session description, in which case the UAS is being presented with an offer
2141 for that session. It is possible that the user is already a participant in that session, even though the INVITE
2142 is outside of a dialog. This can happen when a user is invited to the same multicast conference by multiple
2143 other participants. If desired, the UAS MAY use identifiers within the session description to detect this
2144 duplication. For example, SDP contains a session id and version number in the origin (o) field. If the user
2145 is already a member of the session, and the session parameters contained in the session description have
2146 not changed, the UAS MAY silently accept the INVITE (that is, send a 2xx response without prompting the
2147 user).

2148 The INVITE may not contain a session description at all, in which case the UAS is being asked to
2149 participate in a session, but the UAC has asked that the UAS provide the offer of the session.

2150 The callee can indicate progress, accept, redirect, or reject the invitation. In all of these cases, it formu-
2151 lates a response using the procedures described in Section 8.2.6.

2152 **13.3.1.1 Progress** The UAS may not be able to answer the invitation immediately, and might choose
2153 to indicate some kind of progress to the caller (for example, an indication that a phone is ringing). This
2154 is accomplished with a provisional response between 101 and 199. These provisional responses establish

2155 early dialogs and therefore follow the procedures of Section 12.1.1 in addition to those of Section 8.2.6. A
2156 UAS MAY send as many provisional responses as it likes. Each of these MUST indicate the same dialog ID.
2157 However, these will not be delivered reliably unless reliable provisional responses are used.

2158 If the UAS will require an extended period of time to answer the INVITE, it will need to ask for an
2159 "extension" in order to prevent proxies from cancelling the transaction. A proxy has the option of canceling
2160 a transaction when there is a gap of 3 minutes between messages in a transaction. To prevent cancellation,
2161 the UAS MUST send a non-100 provisional response at least that often. This response SHOULD be sent
2162 reliably, if supported by the UAC. If not, the UAS SHOULD send provisional responses every minute, to
2163 handle the possibility of lost provisional responses.

2164 An INVITE transaction can go on for extended durations when the user is placed on hold, or when interworking
2165 with PSTN systems which allow communications to take place without answering the call. The latter is common in
2166 Interactive Voice Response (IVR) systems.

2167 **13.3.1.2 The INVITE is redirected** If the UAS decides to redirect the call, a 3xx response is sent. A
2168 300 (Multiple Choices), 301 (Moved Permanently) or 302 (Moved Temporarily) response SHOULD contain
2169 a Contact header field containing URIs of new addresses to be tried. The response is passed to the INVITE
2170 server transaction, which will deal with its retransmissions.

2171 **13.3.1.3 The INVITE is rejected** A common scenario occurs when the callee is currently not willing
2172 or able to take additional calls at this end system. A 486 (Busy Here) SHOULD be returned in such scenario.
2173 If the UAS knows that no other end system will be able to accept this call a 600 (Busy Everywhere) response
2174 SHOULD be sent instead. However, it is unlikely that a UAS will be able to know this in general, and thus
2175 this response will not usually be used. The response is passed to the INVITE server transaction, which will
2176 deal with its retransmissions.

2177 A UAS rejecting an offer contained in an INVITE SHOULD return a 488 (Not Acceptable Here) response.
2178 Such a response SHOULD include a Warning header field explaining why the offer was rejected.

2179 **13.3.1.4 The INVITE is accepted** The UAS core generates a 2xx response. This response establishes
2180 a dialog, and therefore follows the procedures of Section 12.1.1 in addition to those of Section 8.2.6.

2181 A 2xx response to an INVITE SHOULD contain the Allow header field and the Supported header field,
2182 and MAY contain the Accept header field. Including these header fields allows the UAC to determine the
2183 features and extensions supported by the UAS for the duration of the call, without probing.

2184 If the INVITE request contained an offer, the 2xx MUST contain an answer. If the INVITE did not contain
2185 an offer, the 2xx MUST contain an offer.

2186 Once the response has been constructed it is passed to the INVITE server transaction. Note, however, that
2187 the INVITE server transaction will be destroyed as soon as it receives this final response. Therefore, it is
2188 necessary to pass periodically the response to the transport until the ACK arrives. The 2xx response is passed
2189 to the transport with an interval that starts at T1 seconds and doubles for each retransmission until it reaches
2190 T2 seconds (T1 and T2 are defined in Section 17). Response retransmissions cease when an ACK request is
2191 received with the same dialog ID as the response. This is independent of whatever transport protocols are
2192 used to send the response.

2193 Since 2xx is retransmitted end-to-end, there may be hops between UAS and UAC which are UDP. To ensure
2194 reliable delivery across these hops, the response is retransmitted periodically even if the transport at the UAS is
2195 reliable.

2196 If the server retransmits the 2xx response for 64*T1 seconds without receiving an ACK, it considers the
2197 dialog completed, the session terminated, and therefore it SHOULD send a BYE.

2198 14 Modifying an Existing Session

2199 A successful INVITE request (see Section 13) establishes both a dialog between two user agents and a
2200 session (using the offer/answer model). Section 12 explains how to modify an existing dialog using a route
2201 refresh request (e.g., changing the *route set* of the dialog). This section describes how to modify the actual
2202 session. This modification can involve changing addresses or ports, adding a media stream, deleting a media
2203 stream, and so on. This is accomplished by sending a new INVITE request within the same dialog that
2204 established the session. An INVITE request sent within an existing dialog is known as a re-INVITE.

2205 Note that a single re-INVITE can modify at the same time the dialog and the parameters of the session.

2206 Either the caller or callee can modify an existing session.

2207 The behaviour of a UA on detection of media failure is a matter of local policy. However, automated
2208 generation of re-INVITE or BYE is NOT RECOMMENDED to avoid flooding the network with traffic when
2209 there is congestion. In any case, if these messages are sent automatically, they SHOULD be sent after some
2210 randomized interval.

2211 Note that the paragraph above refers to automatically generated BYEs and re-INVITEs. If the user hangs up
2212 upon media failure the UA would send a BYE request as usual.

2213 14.1 UAC Behavior

2214 The same offer-answer model that applies to session descriptions in INVITEs (Section 13.2.1) applies to
2215 re-INVITEs. As a result, a UAC that wants to add a media stream, for example, will create a new offer that
2216 contains this media stream, and send that in an INVITE request to its peer. It is important to note that the
2217 full description of the session, not just the change, is sent. This maintains the idempotency of SIP, supports
2218 stateless session processing in various elements, and supports failover and recovery capabilities. Of course,
2219 a UAC MAY send a re-INVITE with no session description, in which case the response to the re-INVITE will
2220 contain the offer.

2221 If the session description format has the capability for version numbers, the offerer SHOULD indicate
2222 that the version of the session description has changed.

2223 The To, From, Call-ID, CSeq, and Request-URI of a re-INVITE are set following the same rules as
2224 for regular requests within an existing dialog, described in Section 12.

2225 A UAC MAY choose not to add Alert-Info header fields or bodies with Content-Disposition "alert" to
2226 re-INVITEs because UASs do not typically alert the user upon reception of a re-INVITE.

2227 Note that, as opposed to initial INVITEs (see Section 13), re-INVITEs contain tags in the To header
2228 field and are sent using the *route set* for the dialog. Therefore, a single final (2xx or non-2xx) response is
2229 received for re-INVITEs.

2230 Note that a UAC MUST NOT initiate a new INVITE transaction within a dialog while another transaction
2231 (INVITE or non-INVITE) is in progress in either direction.

- 2232 1. If there is an ongoing INVITE client transaction the TU MUST wait until the transaction reaches the
2233 *completed* or *terminated* state before initiating the new INVITE.
- 2234 2. If there is an ongoing INVITE server transaction the TU MUST wait until the transaction reaches the
2235 *confirmed* or *terminated* state before initiating the new INVITE.

- 2236 3. If there is an ongoing non-INVITE client or server transaction the TU MUST wait until the transaction
2237 reaches the *completed* or *terminated* state before initiating the new INVITE.

2238 However, a UA MAY initiate a regular transaction while an INVITE transaction is in progress.

2239 If a re-INVITE is responded with a non-2xx final response the session parameters MUST remain un-
2240 changed, as if no re-INVITE had been issued. Note that, as stated in Section 12.2.1.2, if the non-2xx final
2241 response is a 481 (Call/Transaction Does Not Exist) or a 408 (Request Timeout) or no response at all is
2242 received for the re-INVITE (a timeout is returned by the INVITE client transaction) the UAC will terminate
2243 the dialog.

2244 The rules for transmitting a re-INVITE and for generating an ACK for a 2xx response to re-INVITE are
2245 the same as for an INVITE (Section 13.2.1).

2246 14.2 UAS Behavior

2247 Section 13.3.1 describes the steps to follow in order to distinguish incoming re-INVITEs from incoming
2248 initial INVITEs. This Section describes the procedures to follow upon reception of a re-INVITE for an
2249 existing dialog.

2250 A UAS that receives a second INVITE before it sent the final response to a first INVITE with a lower
2251 CSeq sequence number on the same dialog MUST return a 500 (Server Internal Error) response to the second
2252 INVITE and MUST include a **Retry-After** header field with a randomly chosen value of between 0 and 10
2253 seconds.

2254 A UAS that receives an INVITE on a dialog while an INVITE it had sent on that dialog is in progress
2255 MUST return a 491 (Request Pending) response to the received INVITE and MUST include a **Retry-After**
2256 header field with a value chosen as follows:

- 2257 1. If the UAS is the owner of the **Call-ID** of the dialog ID the **Retry-After** header field has a randomly
2258 chosen value of between 2.1 and 4 seconds in units of 10 ms.
- 2259 2. If the UAS is *not* the owner of the **Call-ID** of the dialog ID the **Retry-After** header field has a randomly
2260 chosen value of between 0 and 2 seconds in units of 10 ms.

2261 If a user agent receives a re-INVITE for an existing dialog it MUST check any version identifiers in the
2262 session description or, if there are no version identifiers, the content of the session description to see if it has
2263 changed. If the session description has changed, the user agent server MUST adjust the session parameters
2264 accordingly, possibly after asking the user for confirmation.

2265 Versioning of the session description can be used to accommodate the capabilities of new arrivals to a conference,
2266 add or delete media or change from a unicast to a multicast conference.

2267 If the new session description is not acceptable the UAS can reject it by returning a 488 (Not Acceptable
2268 Here) response for the re-INVITE. This response SHOULD include a **Warning** header field.

2269 If a UAS generates a 2xx response and never receives an ACK, it SHOULD generate a **BYE** to terminate
2270 the dialog.

2271 A UAS MAY choose not to generate 180 (Ringing) responses for a re-INVITE because UACs do not
2272 typically render this information to the user. For the same reason UASs MAY choose not to use **Alert-Info**
2273 header fields or bodies with **Content-Disposition** "alert" in responses to a re-INVITE either.

2274 A UAS providing an offer in a 2xx (because the INVITE did not contain an offer) MUST offer the same
2275 session description as last provided to the peer, with the exception of being able to change the IP address/port
2276 if so desired.

2277 Under error conditions (e.g., the UAS has crashed and restarted) the session description in the 2xx response for
2278 an empty re-INVITE may be different than the one in use at that moment. If the new session description is not
2279 acceptable for the UAC it SHOULD then send a BYE (after ACKing the 2xx response).

2280 15 Terminating a Session

2281 This section describes the procedures to be followed in order to terminate a SIP dialog. For two-party
2282 sessions that are otherwise unbound in time the termination of the dialog implies the termination of the
2283 session. Other types of sessions such as multicast sessions are not terminated when a participant terminates
2284 the SIP dialog that he used to join the session. However, the SIP dialog SHOULD be terminated even
2285 though its termination does not imply the termination of the session. A UA joining a multicast session MAY
2286 terminate the SIP dialog immediately after the INVITE transaction used to join the session has completed.

2287 Either the caller or callee may terminate a dialog for any reason. A caller terminates a dialog either with
2288 BYE or CANCEL depending on the state of the dialog. A callee uses BYE to terminate a confirmed dialog.

2289 If the callee wants to terminate an early dialog it just returns a non-2xx final response for the INVITE.

2290 Sections 13 and 12 document some cases where dialog termination is normative behavior. As a general
2291 rule, if a UA decides that the dialog is to be terminated, it MUST follow the procedures here to initiate
2292 signaling action to convey that.

2293 When a UAC sends an INVITE request to create a session, if a 1xx response with a tag in the To field
2294 is received, an early dialog is created. When a 2xx response is received, the dialog becomes confirmed. For
2295 a confirmed dialog, if the UAC desires to terminate the session, the UAC SHOULD follow the procedures
2296 described in Section 15.1.1 to terminate the session. If the callee for a new session wishes to terminate the
2297 dialog, it uses the procedures of Section 15.1.1, but MUST NOT do so until it has received an ACK or until
2298 the server transaction times out.

2299 This does not mean a user can't hang up right away; it just means that the software in their phone needs to
2300 maintain state for a short while in order to properly clean up.

2301 If the UAC desires to end the session before a confirmed dialog has been created, it SHOULD send a
2302 CANCEL for the INVITE request that requested establishment of the session that is to be terminated. The
2303 UAC constructs and sends the CANCEL following the procedures described in Section 9. This CANCEL
2304 will normally result in a 487 (Request Terminated) response to be returned to the INVITE, indicating suc-
2305 cessful cancellation. However, it is possible that the CANCEL and a 2xx response to the INVITE "pass on
2306 the wire". In this case, the UAC will receive a 2xx to the INVITE. It SHOULD then terminate the call by
2307 following the procedures described in Section 15.1.1.

2308 A UAC can terminate a specific early dialog by following the procedures described in Section 15.1.1.
2309 This would only terminate one particular early dialog.

2310 15.1 Terminating a Dialog with a BYE Request

2311 15.1.1 UAC Behavior

2312 A user agent client uses BYE request, sent within a dialog, to indicate to the server that it wishes to terminate
2313 the session. This will also terminate the dialog. A BYE request MAY be issued by either caller or callee. A
2314 BYE request SHOULD NOT be sent before the creation of a dialog (either early or confirmed). In that case
2315 the UAC SHOULD follow the procedures described in Section 9 instead.

2316 Proxies ensure that a CANCEL request is routed in the same way as the INVITE was. However, a proxy
2317 performing load balancing may route a BYE without a Route header field in a different way than the INVITE, since
2318 both requests have different CSeq sequence numbers.

2319 The To, From, Call-ID, CSeq, and Request-URI of a BYE are set following the same rules as for
2320 regular requests sent within a dialog, described in Section 12.

2321 Once the BYE is constructed, it creates a new non-INVITE client transaction, and passes it the BYE
2322 request. The user agent SHOULD stop sending media as soon as the BYE request is passed to the client
2323 transaction. If the response for the BYE is a 481 (Call/Transaction Does Not Exist) or a 408 (Request
2324 Timeout) or no response at all is received for the BYE (a timeout is returned by the client transaction) the
2325 UAC considers the dialog down anyway.

2326 15.1.2 UAS Behavior

2327 A UAS first processes the BYE request according to the general UAS processing described in Section 8.2.

2328 A UAS core receiving a BYE request checks to see if it matches an existing dialog. If the BYE does not
2329 match an existing dialog, the UAS core SHOULD generate a 481 (Call/Transaction Does Not Exist) response
2330 and pass that to the server transaction.

2331 This rule means that a BYE sent without tags by a UAC will be rejected. This is a change from RFC 2543, which
2332 allowed BYE without tags.

2333 A UAS core receiving a BYE request for an existing dialog MUST follow the procedures of Section
2334 12.2.2 to process the request. Once done, the UAS MUST cease transmitting media streams for the session
2335 being terminated. The UAS core MUST generate a 2xx response to the BYE, and MUST pass that to the
2336 server transaction for transmission.

2337 The UAS MUST still respond to any pending requests received for that dialog, (which can only be an
2338 INVITE). It is RECOMMENDED that a 487 (Request Terminated) response is generated to those pending
2339 requests.

2340 16 Proxy Behavior

2341 16.1 Overview

2342 SIP proxies are elements that route SIP requests to user agent servers and SIP responses to user agent clients.
2343 A request may traverse several proxies on its way to a UAS. Each will make routing decisions, modifying
2344 the request before forwarding it to the next element. Responses will route through the same set of proxies
2345 traversed by the request in the reverse order.

2346 Being a proxy is a logical role for a SIP element. When a request arrives, an element that can play the
2347 role of a proxy must first decide if it needs to respond to the request on its own. For instance, the request
2348 could be malformed or the element may need credentials from the client before acting as a proxy. The
2349 element MAY respond with any appropriate error code. When responding directly to a request, the element
2350 is playing the role of a UAS and MUST behave as described in Section 8.2.

2351 A proxy can operate in either a stateful or stateless mode for each new request. When stateless, a proxy
2352 acts as a simple forwarding element. It forwards each request downstream to a single element determined
2353 by making a routing decision based on the request. It simply forwards every response it receives upstream.
2354 A stateless proxy discards information about a message once it has been forwarded.

2355 On the other hand, a stateful proxy remembers information (specifically, transaction state) about each
2356 incoming request and any requests it sends as a result of processing the incoming request. It uses this
2357 information to affect the processing of future messages associated with that request. A stateful proxy MAY
2358 chose to "fork" a request, routing it to multiple destinations. Any request that is forwarded to more than

2359 one location **MUST** be handled statefully. Any request processed using TCP (or any other mechanism that is
2360 inherently stateful), **MUST** be handled statefully.

2361 A stateful proxy **MAY** transition to stateless operation at any time during the processing of a request,
2362 so long as it did not do anything that would otherwise prevent it from being stateless initially (forking, for
2363 example, or generation of a 100 response). When performing such a transition, all state is simply discarded.
2364 The proxy **SHOULD NOT** send a **CANCEL**.

2365 Much of the processing involved when acting statelessly or statefully for a request is identical. The next
2366 several subsections are written from the point of view of a stateful proxy. The last section calls out those
2367 places where a stateless proxy behaves differently.

2368 **16.2 Stateful Proxy**

2369 When stateful, a proxy is purely a SIP transaction processing engine. Its behavior is modeled here in terms
2370 of the Server and Client Transactions defined in Section 17. A stateful proxy has a server transaction
2371 associated with one or more client transactions by a higher layer proxy processing component (see figure 3),
2372 known as a proxy core. An incoming request is processed by a server transaction. Requests from the server
2373 transaction are passed to a proxy core. The proxy core determines where to route the request, choosing
2374 one or more next-hop locations. An outgoing request for each next-hop location is processed by its own
2375 associated client transaction. The proxy core collects the responses from the client transactions and uses
2376 them to send responses to the server transaction.

2377 A stateful proxy creates a new server transaction for each new request received. Any retransmissions of
2378 the request will then be handled by that server transaction per Section 17.

2379 This is a model of proxy behavior, not of software. An implementation is free to take any approach that
2380 replicates the external behavior this model defines.


2381 For all new requests, including any with unknown methods, an element intending to proxy the request
2382 **MUST**:

- 2383 1. Validate the request (Section 16.3)
- 2384 2. Make a routing decision (Section 16.4)
- 2385 3. Forward the request to each chosen destination (Section 16.5)
- 2386 4. Process all responses (Section 16.6)

2387 **16.3 Request Validation**

2388 Before an element can proxy a request, it **MUST** verify the message's validity. A valid message must pass
2389 the following checks:

- 2390 1. Reasonable Syntax
- 2391 2. Max-Forwards
- 2392 3. (Optional) Loop Detection
- 2393 4. Proxy-Require



2394 5. Proxy-Authorization

2395 If any of these checks fail, the element **MUST** behave as a user agent server (see Section 8.2) and respond
2396 with an error code.

2397 Notice that a proxy is not required to detect merged requests and **MUST NOT** treat merged requests as an
2398 error condition. The endpoints receiving the requests will resolve the merge as described in Section 8.2.2.2.
2399

2400 1. Reasonable Syntax check

2401 The request **MUST** be well-formed enough to be handled with a server transaction. Any components
2402 involved in the remainder of these Request Validation steps or the Request Processing section **MUST** be
2403 well-formed. Any other components, well-formed or not, **SHOULD** be ignored and remain unchanged
2404 when the message is forwarded. For instance, an element **SHOULD NOT** reject a request because of
2405 a malformed **Date** header field. Likewise, a proxy **SHOULD NOT** remove a malformed **Date** header
2406 before forwarding a request.

2407 This protocol is designed to be extended. Future extensions may define new methods and header fields
2408 at any time. An element **MUST NOT** refuse to proxy a request because it contains a method or header
2409 field it does not know about.

2410 2. Max-Forwards check

2411 The **Max-Forwards** header (Section 24.22) is used to limit the number of elements a SIP request can
2412 traverse.

2413 If the request does not contain a **Max-Forwards** header field, this check is passed.

2414 If the request contains a **Max-Forwards** header field with a field value greater than zero, the check is
2415 passed.

2416 If the request contains a **Max-Forwards** header field with a field value of zero (0), the element **MUST**
2417 **NOT** forward the request. If the request was for **OPTIONS**, the element **MAY** act as the final recipient
2418 and respond per Section 11. Otherwise, the element **MUST** return a 483 (Too many hops) response.

2419 3. Optional Loop Detection check

2420 An element **MAY** check for forwarding loops before forwarding a request. If the request contains a
2421 **Via** header field value with A sent-by value that equals a value placed into previous requests by the
2422 proxy, the request has been forwarded by this element before. The request has either looped or is
2423 legitimately spiraling through the element. To determine if the request has looped, the element **MAY**
2424 perform the **branch** parameter calculation described in Step 3 of Section 16.5 on this message and
2425 compare it to the parameter received in that **Via** field value. If the parameters match, the request
2426 has looped. If they differ, the request is spiraling, and processing continues. If a loop is detected, the
2427 element **MAY** return a 482 (Loop Detected) response.

2428 In earlier versions of this memo, loop detection was **REQUIRED**. This requirement has been relaxed in
2429 favor of the **Max-Forwards** mechanism.

2430 4. Proxy-Require check

2431 Future extensions to this protocol may introduce features that require special handling by proxies.
2432 Endpoints will include a **Proxy-Require** header in requests that use these features, telling the proxy
2433 it should not process the request unless the feature is understood.

2434 If the request contains a **Proxy-Require** header (Section 24.29) with one or more option-tags this
2435 element does not understand, the element **MUST** return a 420 (Bad Extension) response. The response
2436 **MUST** include an **Unsupported** (Section 24.42) header field listing those option-tags the element did
2437 not understand.

2438 5. Proxy-Authorization check

2439 If an element requires credentials before forwarding a request, the request **MUST** be inspected as
2440 described in Section 20.3. That section also defines what the element must do if the inspection fails.

2441 16.4 Making a Routing Decision

2442 At this point, the proxy must decide where to forward the request. This can be modeled as computing a set
2443 of destinations for the request. This set will either be predetermined by the contents of the request or will
2444 be obtained from an abstract location service. Each destination is represented as a URI and an optional IP
2445 address, port and transport. This combination is referred to as a “next-hop location”.

2446 First, the proxy core checks the received request for **Route** headers. If any **Route** header fields are
2447 present in the request, the proxy **MUST** choose a single next-hop location to place in the destination set. The
2448 proxy **SHOULD** choose to use a strict-routing policy, placing the URI (including all of its parameters) from
2449 the topmost **Route** header field as the only next hop URI in the destination set, with no IP address, port
2450 and transport set for that next hop. The proxy **MAY** choose to use a loose-routing policy, selecting a URI,
2451 address, port and transport based on that policy. A loose-routing policy **MAY** use any information in or about
2452 the request in determining where to route it. Restrictions on the a loose-routing proxy’s policy are discussed
2453 in Section 8.1.3.

2454 Once the single next-hop location is placed into the destination set, the set is complete, and the proxy
2455 **MUST** proceed to the Request Processing of Section 16.5.

2456 The **Route** mechanism is used to affect the path a request takes through SIP elements. A strict-routing
2457 policy results in behaviour much like strict IP source routing. Loose-routing policies will result in the
2458 specified URIs being reached, possibly visiting additional elements in the process. A UAC will insert
2459 **Route** header fields (see Section 12), based on information provided by proxies through **Record-Route**
2460 header fields or by policy obtained through configuration. (see Step 6 of Section 16.5).

2461 Assuming there were no **Route** headers in the received request, the proxy checks the **Request-URI** of
2462 the received request. If the **Request-URI** has a URI whose scheme is not understood by the proxy, the
2463 proxy **SHOULD** reject the request with a 416 (Unsupported URI Scheme) response. If the **Request-URI**
2464 contains an **maddr** parameter, the proxy **MUST** check to see if its value is in the set of addresses or domains
2465 the proxy is configured to be responsible for. If the **Request-URI** has an **maddr** parameter with a value
2466 the proxy is responsible for, and the request was received using the port and transport indicated (explicitly
2467 or by default) in the **Request-URI**, the proxy **MUST** strip the **maddr** and any non-default port or transport
2468 parameter and continue processing as if those values had not been present in the request. Otherwise, if the
2469 **Request-URI** contains an **maddr** parameter, the **Request-URI** **MUST** be placed into the destination set as
2470 the only next hop URI, with no IP address, port and transport set for that next hop, and the proxy **MUST**
2471 proceed to Section 16.5.

2472 A request may arrive with an `maddr` matching the proxy, but on a port or transport different from that indicated
2473 in the URI. Such a request needs to be forwarded to the proxy using the indicated port and transport.

2474 If the domain of the `Request-URI` indicates a domain this element is not responsible for, it SHOULD set
2475 the next hop URI to the `Request-URI`, and leave the IP address, port and transport of the next hop empty.
2476 That next hop MUST be placed into the destination set as the only next hop, and the element MUST proceed
2477 to the task of Request Processing (Section 16.5).

2478 There are many circumstances in which a proxy might receive a request for a domain it is not responsible for.
2479 A firewall proxy handling outgoing calls (the way HTTP proxies handle outgoing requests) is an example of where
2480 this is likely to occur.

2481 If the destination set for the request has not been predetermined as described above, this implies that the
2482 element is responsible for the domain in the `Request-URI`, and the element MAY use whatever mechanism
2483 it desires to determine where to send the request. Any of these mechanisms can be modeled as accessing an
2484 abstract Location Service. This may consist of obtaining information from a location service created by a SIP
2485 Registrar, reading a database, consulting a presence server, utilizing other protocols, or simply performing
2486 an algorithmic substitution on the `Request-URI`. When accessing the location service constructed by the
2487 registrar, the `Request-URI` MUST first be canonicalized as described in Section 10.3 before being used as
2488 an index. The output of these mechanisms is used to construct the destination set.

2489 If the `Request-URI` does not provide sufficient information for the proxy to determine the destination
2490 set, it SHOULD return a 485 (Ambiguous) response. This response SHOULD contain a `Contact` header field
2491 containing URIs of new addresses to be tried. For example, an `INVITE` to `sip:John.Smith@company.com`
2492 may be ambiguous at a proxy whose location service has multiple John Smiths listed. See Section 25.4.23
2493 for details.

2494 Any information in or about the request or the current environment of the element MAY be used in the
2495 construction of the destination set. For instance, different sets may be constructed depending on contents or
2496 the presence of header fields and bodies, the time of day of the request's arrival, the interface on which the
2497 request arrived, failure of previous requests, or even the element's current level of utilization.

2498 As potential destinations are located through these services, their next hops are added to the destination
2499 set. Next-hop locations may only be placed in the destination set once. If a next-hop location is already
2500 present in the set (based on the definition of equality for the URI type and equality of the optional parame-
2501 ters), it MUST NOT be added again.

2502 If the recieved request contained no `Route` headers, a proxy MAY continue to add destinations to the
2503 set after beginning Request Processing. It MAY use any information obtained during that processing to
2504 determine new locations. For instance, a proxy may choose to incorporate contacts obtained in a redirect
2505 response (3xx class) into the destination set. If a proxy uses a dynamic source of information while building
2506 the destination set (for instance, if it consults a SIP Registrar), it SHOULD monitor that source for the duration
2507 of processing the request. New locations SHOULD be added to the destination set as they become available.
2508 As above, any given URI MUST NOT be added to the set more than once.

2509 Allowing a URI to be added to the set only once reduces unnecessary network traffic, and in the case of incor-
2510 porating contacts from redirect requests prevents infinite recursion.

2511 An example trivial location service is achieved by configuring an element with a default outbound des-
2512 tination. All requests are forwarded to this location. The `Request-URI` of the request is placed in the
2513 destination set with the optional next-hop IP address, port and transport parameters set to the default out-
2514 bound destination. The destination set is complete, containing **only** this URI, and the element proceeds to
2515 the task of Request Processing.

2516 If the Request-URI indicates a resource at this proxy that does not exist, the proxy MUST return a 404
2517 (Not Found) response.

2518 If the destination set remains empty after applying all of the above, the proxy MUST return an error
2519 response, which SHOULD be the 480 (Temporarily Unavailable) response.

2520 **16.5 Request Processing**

2521 As soon as the destination set is non-empty, a proxy MAY begin forwarding the request. A stateful proxy
2522 MAY process the set in any order. It MAY process multiple destinations serially, allowing each client transac-
2523 tion to complete before starting the next. It MAY start client transactions with every destination in parallel. It
2524 also MAY arbitrarily divide the set into groups, processing the groups serially and processing the destinations
2525 in each group in parallel.

2526 A common ordering mechanism is to use the qvalue parameter of destinations obtained from Contact
2527 header fields (see Section 24.10). Destinations are processed from highest qvalue to lowest. Destinations
2528 with equal qvalues may be processed in parallel.

2529 A stateful proxy must have a mechanism to maintain the destination set as responses are received and
2530 associate the responses to each forwarded request with the original request. For the purposes of this model,
2531 this mechanism is a “response context” created by the proxy layer before forwarding the first request.

2532 For each destination, the proxy forwards the request following these steps:

- 2533 1. Make a copy of the received request
- 2534 2. Update the Request-URI
- 2535 3. Add a Via header field value
- 2536 4. Update the Max-Forwards field
- 2537 5. Update the Route header field if present
- 2538 6. Optionally add a Record-route header field value
- 2539 7. Optionally add additional headers
- 2540 8. send the new request
- 2541 9. Set timer C

2542 Each of these steps is detailed below:

2543 1. Copy request

2544 The proxy starts with a copy of the received request. The copy MUST initially contain all of the header
2545 fields from the received request. Only those fields detailed in the processing described below may be
2546 removed. The copy SHOULD maintain the ordering of the header fields as in the received request. The
2547 proxy MUST NOT reorder field values with a common field name (See Section 7.3.1).

2548 An actual implementation need not perform a copy; the primary requirement is that the processing of each
2549 next hop begin with the same request.

2550 2. Request-URI

2551 The Request-URI in the copy's start line MUST be replaced with the URI for this destination. If the
2552 URI contains any parameters not allowed in a Request-URI, they MUST be removed.

2553 This is the essence of a proxy's role. This is the mechanism through which a proxy routes a request
2554 toward its destination.

2555 3. Via

2556 The proxy MUST insert a Via header field into the copy before the existing Via header fields. The
2557 construction of this header follows the same guidelines of Section 8.1.1.7. This implies that the proxy
2558 will compute its own branch parameter, which will be globally unique for that branch, and contain the
2559 requisite magic cookie.

2560 Proxies choosing to detect loops have an additional constraint in the value they use for construction of
2561 the branch parameter. A proxy choosing to detect loops SHOULD create a branch parameter separable
2562 into two parts by the implementation. The first part MUST satisfy the constraints of Section 8.1.1.7 as
2563 described above. The second is used to perform loop detection and distinguish loops from spirals.

2564 Loop detection is performed by verifying that, when a request returns to a proxy, those fields having an
2565 impact on the processing of the request have not changed. The value placed in this part of the branch
2566 parameter SHOULD reflect all of those fields (including any Proxy-Require and Proxy-Authorization
2567 headers). This is to ensure that if the request is routed back to the proxy and one of those fields
2568 changes, it is treated as a spiral and not a loop (Section 16.3 item 2) A common way to create this
2569 value is to compute a cryptographic hash of the To, From, Call-ID header fields, the Request-URI
2570 of the request received (before translation) and the sequence number from the CSeq header field, in
2571 addition to any Proxy-Require and Proxy-Authorization fields that may be present. The algorithm
2572 used to compute the hash is implementation-dependent, but MD5 [21], expressed in hexadecimal, is
2573 a reasonable choice. (Base64 is not permissible for a token.)

2574 If a proxy wishes to detect loops, the "branch" parameter it supplies MUST depend on all information
2575 affecting processing of a request, including the incoming request-URI and any header values affecting the
2576 request's admission or routing. This is necessary to distinguish looped requests from requests whose routing
2577 parameters have changed before returning to this server.

2578 The request method MUST NOT be included in the calculation of the branch parameter. In particular,
2579 CANCEL and ACK requests (for non-2xx responses) MUST have the same branch value as the cor-
2580 responding request they cancel or acknowledge. The branch parameter is used in correlating those
2581 requests at the server handling them (see Section 17.2.3 and 9.2).

2582 4. Max-Forwards

2583 If the copy does not contain a Max-Forwards header field, the proxy must add one with a field value
2584 of 70.

2585 Some existing UAs will not provide a Max-Forwards header field in a request.

2586 If the copy contains a Max-Forwards header field, the proxy must decrement its value by one (1).

2587 5. Route

2588 If the copy contains a **Route** header field, the proxy's routing policy will determine whether that field
2589 should be modified. A proxy with a strict-routing policy **MUST** remove the first (topmost) **Route**
2590 header field value. (The strict-routing policy would have already placed that value into the Request-
2591 URI of this copy.) A proxy with a loose-routing policy **MAY** remove the topmost value. Restrictions on
2592 a loose-routing proxy's policy with respect to the topmost **Route** header are described in Section 8.1.3.
2593

2594 6. Record-Route

2595 If this proxy wishes remain on the path of future requests in a dialog created by this request, it **MUST**
2596 insert a **Record-Route** header value into the copy before any existing **Record-Route** header values,
2597 even if a **Route** field is already present.

2598 Requests establishing a dialog may contain preloaded **Route** header fields.

2599 If this request is already part of a dialog, the proxy **SHOULD** insert a **Record-Route** header field value
2600 if it wishes to remain on the path of future requests in the dialog. In normal endpoint operation as
2601 described in Section 12 these **Record-Route** header field values will not have any effect on the route
2602 sets used by the endpoints.

2603 The proxy will remain on the path if it chooses to not insert a **Record-Route** header field value into requests
2604 that are already part of a dialog. However, it would be removed from the path when an endpoint that has failed
2605 reconstitutes the dialog.

2606 A proxy **MAY** insert a **Record-Route** header value into any request. If the request does not initiate
2607 a dialog, the endpoints will ignore the value. See Section 12 for details on how endpoints use the
2608 **Record-Route** header field values to construct **Route** header fields.

2609 Each proxy in the path of a request chooses whether to add a **Record-Route** header field value
2610 independently - the presence of a **Record-Route** header field in a request does not obligate this proxy
2611 to add a value.

2612 The URI placed in the **Record-Route** header value **MUST** be a SIP URI. This URI **MAY** be different
2613 for each destination the request is forwarded to. The URI **SHOULD NOT** contain the transport param-
2614 eter unless the proxy has knowledge (such as in a private network) that the next downstream element
2615 that will be in the path of subsequent requests supports that transport.

2616 The URI this proxy provides will be used by some other element to make a routing decision. This proxy, in
2617 general, has no way to know what the capabilities of that element are, so it must restrict itself to the mandatory
2618 elements of a SIP implementation: SIP URIs and UDP transports.

2619 The URI placed in the **Record-Route** header value **MUST** resolve to this element when the server
2620 location procedures of [8] are applied to it. This ensures subsequent requests are routed back to this
2621 element.

2622 The URI placed in the **Record-Route** header value **SHOULD** be such that if a subsequent request is
2623 received with this URI in the **Request-URI**, the proxy's normal request processing will cause it to be
2624 forwarded to one of the previous elements, including the originating client, traversed by the original

2625 request. This improves robustness, ensuring that the **Request-URI** contains enough information to
2626 forward subsequent requests to a reasonable destination even in the absence of **Route** headers.

2627 The URI placed in the **Record-Route** header value *MUST* vary with the **Request-URI** in the received
2628 request. A request may legitimately pass through this proxy more than once on the way to its final
2629 destination (this is called a spiraling request). The **Request-URI** will be different each time the
2630 request passes through. If this proxy places the same URI in the **Record-Route** header field each time,
2631 subsequent requests will be rejected as looped requests. It is insufficient to simply copy the **Request-**
2632 **URI** from each request into the **Record-Route** header. Some modification, such as adding an **maddr**
2633 parameter, is necessary.

2634 URIs satisfying the above paragraphs can be constructed in many ways. One way is to use a URI that
2635 is nearly the same as the **Contact** header in the initial request (if present, else the **From** field), but with
2636 the **maddr** and port set to resolve to the proxy, and with a transaction identifier added to the user part of
2637 the request-URI (in order to meet the requirement that the URI in the **Record-Route** be different for
2638 each distinct **Request-URI**). A call stateful proxy could use a URI of the form sip:proxy.example.com
2639 and use information from the stored call state to meet the requirements.

2640 The proxy *MAY* include **Record-Route** header parameters in the value it provides. These will be
2641 returned in some responses to the request (200 (OK) responses to **INVITE** for example) and may be
2642 useful for pushing state into the message.

2643 The **Record-Route** process is designed to work for any SIP request that initiates a dialog. The only
2644 such request in this specification is **INVITE**. Extensions to the protocol *MAY* define others, and the
2645 mechanisms described here will apply.

2646 If a proxy needs to be in the path of any type of dialog (such as one straddling a firewall), it *SHOULD*
2647 add a **Record-Route** header value to every request with a method it does not understand since that
2648 method may have dialog semantics.

2649 The URI a proxy places into a **Record-Route** value is only valid for the lifetime of any dialog created
2650 by transaction in which it occurs. A dialog-stateful proxy, for example, *MAY* refuse to accept future
2651 requests with that value in the **Request-URI** after the dialog has terminated. Non-dialog-stateful
2652 proxies, of course, have no concept of when the dialog has terminated, but they *MAY* encode enough
2653 information in the value to compare it against the dialog identifier of future requests and *MAY* reject
2654 requests not matching that information. Endpoints *MUST NOT* use a URI obtained from a **Record-**
2655 **Route** header value outside the dialog in which it was provided. See Section 12 for more information
2656 on an endpoint's use of **Record-Route** header values.

2657 Generally, the choice about whether to record-route or not is a tradeoff of features vs. performance.
2658 Faster request processing and higher scalability is achieved when proxies do not record route. How-
2659 ever, provision of certain services may require a proxy to observe all messages in a dialog. It is
2660 **RECOMMENDED** that proxies do not automatically record route. They should do so only if specifi-
2661 cally required.

2662 7. Adding Additional Headers

2663 The proxy *MAY* add any other appropriate headers to the copy at this point.

2664 8. Forward Request

2665 A stateful proxy creates a new client transaction for this request as described in Section 17.1. If
2666 the next-hop location used in building this request contains the optional addressing parameters, the
2667 transaction is instructed to send the request based on those parameters. Otherwise, the proxy uses
2668 the procedures of Section [8] to compute an ordered set of addresses from the Request-URI, and
2669 as described there, attempts to contact the first one by instructing the client transaction to send the
2670 request there. If the client transaction reports failure to send the request or a timeout from its state
2671 machine, the stateful proxy continues to the next address that ordered set. Each attempt is a new client
2672 transaction, and therefore represents a new branch, so that the processing described above for each
2673 branch would need to be repeated. This results in a requirement to use a different branch ID parameter
2674 for each attempt. If the ordered set is exhausted, the request cannot be forwarded to this element in
2675 the destination set. The proxy does not need to place anything in the response context, but otherwise
2676 acts as if this element of the destination set returned a 408 (Request Timeout) final response.

2677 9. Set timer C

2678 In order to handle the case where an INVITE request never generates a final response, a transaction
2679 timeout value is used. This is accomplished through a timer, called timer C, which MUST set for each
2680 client transaction when an INVITE request is proxied. The timer MUST be larger than 3 minutes.
2681 Section 16.6 bullet 2 discusses how this timer is updated with provisional responses, and Section 16.7
2682 discusses processing when it fires.

2683 16.6 Response Processing

2684 When a response is received by an element, it first tries to locate a client transaction (Section 17.1.3) match-
2685 ing the response. If none is found, the element MUST process the response (even if it is an informational
2686 response) as a stateless proxy (described below). If a match is found, the response is handed to the client
2687 transaction.

2688 Forwarding responses for which a client transaction (or more generally any knowledge of having sent an asso-
2689 ciated request) is not found improves robustness. In particular, it ensures that "late" 2xx class responses to INVITE
2690 requests are forwarded properly.

2691 As client transactions pass responses to the proxy layer, the following processing MUST take place:

- 2692 1. Find the appropriate response context
- 2693 2. Update timer C for provisional responses
- 2694 3. Remove the topmost Via
- 2695 4. Add the response to the response context
- 2696 5. Check to see if this response should be forwarded

2697 The following processing MUST be performed on each response that is forwarded. It is likely that more
2698 than one response to each request will be forwarded: at least each provisional and one final response.

- 2699 1. Aggregate authorization header fields if necessary;
- 2700 2. forward the response;

2701 3. generate any necessary CANCEL requests.

2702 If no final response has been forwarded after every client transaction associated with the response context
2703 has been terminated, the proxy must choose and forward the “best” response from those it has seen so far.

2704 Each of the above steps are detailed below:

2705 1. Find Context

2706 The proxy locates the “response context” it created before forwarding the original request using the
2707 key described in Section 16.5. The remaining processing steps take place in this context.

2708 2. Update timer C for provisional responses

2709 For an INVITE transaction, if the response is a provisional response with status codes 101 to 199
2710 inclusive (i.e., anything but 100), the proxy MUST reset timer C for that client transaction. The timer
2711 MAY be reset to a different value, but this value MUST be greater than 3 minutes.

2712 3. Via

2713 The proxy removes the topmost Via field value from the response.

2714 If no Via field values remain in the response, the response was meant for this element and MUST
2715 NOT be forwarded. The remainder of the processing described in this section is not performed on this
2716 message, the UAC processing rules described in Section 8.1.4 are followed instead (transport layer
2717 processing has already occurred).

2718 This will happen, for instance, when the element generates CANCEL requests as described in Sec-
2719 tion 10.

2720 4. Add response to context ;

2721 Final responses received are stored in the response context until a final response is generated on the
2722 server transaction associated with this context. The response may be a candidate for the best final
2723 response to be returned on that server transaction. Information from this response may be needed in
2724 forming the best response even if this response is not chosen.

2725 If the proxy chooses to recurse on any contacts in a 3xx class response by adding them to the destina-
2726 tion set, it MUST remove them from the response before adding the response to the response context.
2727 If the proxy recurses on all of the contacts in a 3xx class response, the proxy SHOULD NOT add the
2728 resulting contactless response to the response context.

2729 Removing the contact before adding the response to the response contact prevents the next element up-
2730 stream from retrying a location this proxy has already attempted.

2731 3xx class responses may contain a mixture of SIP and non-SIP URIs. A proxy may choose to recurse on
2732 the SIP URIs and place the remainder into the response context to be returned potentially in the final response.

2733 If a proxy receives a 416 (Unsupported URI Scheme) response to a request whose Request-URI
2734 scheme was not SIP, but the scheme in the original received request was SIP (that is, the proxy changed
2735 the scheme from SIP to something else when it proxied a request), the proxy SHOULD add a new URI
2736 to the destination set. This URI SHOULD be a SIP URI version of the non-SIP URI that was just tried.
2737 In the case of the tel URL, this is accomplished by placing the telephone-subscriber part of the tel
2738 URL into the user part of the SIP URI, and setting the hostpart to the domain where the prior request
2739 was sent.

2740 As with a 3xx response, if a proxy “recurses” on the 416 by trying a SIP URI instead, the 416 response
2741 SHOULD NOT be added to the response context.

2742 5. Check response for forwarding

2743 Until a final response has been sent on the server transaction, the following responses MUST be for-
2744 warded immediately:

- 2745 • Any provisional response other than 100 (Trying)
- 2746 • Any 2xx response

2747 If a 6xx response is received, it is not immediately forwarded, but the stateful proxy SHOULD cancel
2748 all pending transactions as described in Section 10.

2749 This is a change from RFC 2543, which mandated that the proxy was to forward the 6xx response imme-
2750 diately. For an INVITE transaction, this approach had the problem that a 2xx response could arrive on another
2751 branch, in which case the proxy would have to forward the 2xx. The result was that the UAC could receive
2752 a 6xx response followed by a 2xx response, which should never be allowed to happen. Under the new rules,
2753 upon receiving a 6xx, a proxy will issue a CANCEL request, which will generally result in 487 responses from
2754 all outstanding client transactions, and then at that point the 6xx is forwarded upstream.

2755 After a final response has been sent on the server transaction, the following responses MUST be for-
2756 warded immediately:

- 2757 • Any 2xx class response to an INVITE request

2758 A stateful proxy MUST NOT immediately forward any other responses. In particular, a stateful proxy
2759 MUST NOT forward any 100 (Trying) response. Those responses that are candidates for forwarding
2760 later as the “best” response have been gathered as described in step “Add Response to Context”.

2761 Any response chosen for immediate forwarding MUST be processed as described in steps “Aggregate
2762 authorization headers” through “Record-Route”.

2763 This step, combined with the next, ensures that a stateful proxy will forward exactly one final response
2764 to a non-INVITE request, and either exactly one non-2xx class response or one or more 2xx-class
2765 responses to an INVITE request.

2766 6. Choosing the best response

2767 A stateful proxy MUST send a final response to a response context’s server transaction if no final
2768 responses have been immediately forwarded by the above rules and all client transactions in this
2769 response context have been terminated.

2770 The stateful proxy MUST choose the “best” final response among those received and stored in the
2771 response context.

2772 If there are no final responses in the context, the proxy MUST send a 408 (Request Timeout) response
2773 to the server transaction.

2774 Otherwise, the proxy MUST forward one of the responses from the lowest response class stored in the
2775 response context. The proxy MAY select any response within that lowest class. The proxy SHOULD
2776 give preference to responses that provide information affecting resubmission of this request, such as
2777 401, 407, 415, 420, and 484.

2778 A proxy which receives a 503 (Service Unavailable) response SHOULD NOT forward it upstream
2779 unless it can determine that any subsequent requests it might proxy will also generate a 503. In other
2780 words, forwarding a 503 means that the proxy knows it cannot service any requests, not just the one
2781 for the Request-URI in the request which generated the 503.

2782 The forwarded response MUST be processed as described in steps "Aggregate authorization headers"
2783 through "Record-Route".

2784 For example, if a proxy forwarded a request to 4 locations, and received 503, 407, 501, and 404
2785 responses, it may choose to forward the 407 (Proxy Authentication Required) response.

2786 1xx and 2xx class responses may be involved in the establishment dialogs. When a request does not
2787 contain a To tag, the To tag in the response is used by the UAC to distinguish multiple responses to
2788 a dialog creating request. A proxy MUST NOT insert a tag into the To header of a 1xx or 2xx class
2789 response if the request did not contain one. A proxy MUST NOT modify the tag in the To header of a
2790 1xx or 2xx class response.

2791 Since a proxy may not insert a tag into the To header of a 1xx class response to a request that did
2792 not contain one, it cannot issue non-100 provisional responses on its own. However, it can branch the
2793 request to a UAS sharing the same element as the proxy. This UAS can return its own provisional
2794 responses, entering into an early dialog with the initiator of the request. The UAS does not have to be
2795 a discreet process from the proxy. It could be a virtual UAS implemented in the same code space as
2796 the proxy.

2797 3-6xx class responses are delivered hop-hop. When issuing a 3-6xx class response, the element is
2798 effectively acting as a UAS, issuing its own response, usually based on the responses received from
2799 downstream elements. An element SHOULD preserve the To tag when simply forwarding a 3-6xx
2800 class response to a request that did not contain a To tag.

2801 A proxy MUST NOT modify the To tag in any forwarded response to a request that contains a To tag.

2802 While it makes no difference to the upstream elements if the proxy replaced the To tag in a forwarded
2803 3-6xx class response, preserving the original tag may assist with debugging.

2804 When the proxy is aggregating information from several responses, choosing a To tag from among them
2805 is arbitrary, and generating a new To tag may make debugging easier. This happens, for instance, when
2806 combining 401 (Unauthorized) and 407 (Proxy Authentication Required) challenges, or combining Contact
2807 values from unencrypted and unauthenticated 3xx class responses.

2808 7. Aggregate authorization headers

2809 If the selected response is a 401 (Unauthorized) or 407 (Proxy Authentication Required), the proxy
2810 MUST collect any WWW-Authenticate and Proxy-Authenticate header fields from all other 401
2811 (Unauthorized) and 407 (Proxy Authentication Required) responses received so far in this response
2812 context and add them to this response before forwarding. Each WWW-Authenticate and Proxy-
2813 Authenticate header field added to the response MUST preserve that header field value. The result-
2814 ing 401 (Unauthorized) or 407 (Proxy Authentication Required) response may have several WWW-
2815 Authenticate AND Proxy-Authenticate headers.

2816 This is necessary because any or all of the destinations the request was forwarded to may have re-
2817 quested credentials. The client must receive all of those challenges and supply credentials for each of
2818 them when it retries the request. Motivation for this behavior is provided in Section 22.

2819 8. Record-Route

2820 If the selected response contains a **Record-Route** header field value originally provided by this proxy,
2821 the proxy MAY chose to rewrite the value before forwarding the response. This allows the proxy to
2822 provide different URIs for itself to the next upstream and downstream elements. A proxy may choose
2823 to use this mechanism for any reason. For instance, it is useful for multi-homed hosts.

2824 The new URI provided by the proxy MUST satisfy the same constraints on URIs placed in **Record-Route**
2825 header fields in requests (see Step 6 of Section 16.5) with the following modifications:

2826 The URI SHOULD NOT contain the transport parameter unless the proxy has knowledge that the next
2827 upstream (as opposed to downstream) element that will be in the path of subsequent requests supports
2828 that transport.

2829 The URI placed in the **Record-Route** header value SHOULD be such that if a subsequent request is
2830 received with this URI in the **Request-URI**, the proxy's normal request processing will cause it to
2831 be forwarded to the same next-hop element (as opposed to some previous element) as the originally
2832 forwarded request.

2833 When a proxy does decide to modify the **Record-Route** header in the response, one of the operations
2834 it must perform is to locate the **Record-Route** that it had inserted. If the request spiraled, and the
2835 proxy inserted a **Record-Route** in each iteration of the spiral, locating the correct header in the
2836 response (which must be the proper iteration in the reverse direction) is tricky. The rules above dictate
2837 that a proxy insert a different URI into the **Record-Route** for each distinct **Request-URI** received.
2838 The two issues can be solved jointly. A RECOMMENDED mechanism is for the proxy to append a
2839 piece of data to the user portion of the URI. This piece of data is a hash of the transaction key (those
2840 peices of data used to match a request against existing transactions as discussed in section 17.2.3)
2841 for the incoming request, concatenated with a unique identifier for the proxy instance. Since the
2842 transaction key either contains **Request-URI** or depends on it (when the key is encoded in the branch
2843 parameter of the topmost **Via** header), this key will be unique for each distinct **Request-URI**. When
2844 the response arrives, the proxy modifies the first **Record-Route** whose identifier matches the proxy
2845 instance. The modification results in a URI without this piece of data appended to the user portion of
2846 the URI. Upon the next iteration, the same algorithm (find the topmost **Record-Route** header with
2847 the parameter) will correctly extract the next **Record-Route** header inserted by that proxy.

2848 9. Forward response

2849 After performing the processing described in steps "Aggregate authorization headers" through "Record-Route",
2850 the proxy may perform any feature specific manipulations on the selected response. Unless
2851 otherwise specified, the proxy MUST NOT remove the message body or any header values other than
2852 the **Via** header value discussed in Section 3. In particular, the proxy MUST NOT remove any "received"
2853 parameter it may have added to the next **Via** header value while processing the request associated with
2854 this response. The proxy MUST pass the response to the server transaction associated with the re-
2855 sponse context. This will result in the response being sent to the location now indicated in the topmost
2856 **Via** field value. If the server transaction is no longer available to handle the transmission, the element
2857 MUST forward the response statelessly by sending it to the server transport. The server transaction
2858 may indicate failure to send the response or signal a timeout in its state machine. These errors should
2859 be logged for diagnostic purposes as appropriate, but the protocol requires no remedial action from
2860 the proxy.

2861 The proxy MUST maintain the response context until all of its associated transactions have been ter-
2862 minated, even after forwarding a final response.

2863 10. Generate CANCELs

2864 OPEN ISSUE #7: If CANCEL is restricted to INVITE only, this behavior must restrict itself to
2865 INVITE requests.

2866 If the forwarded response was a final response, the proxy MUST generate a CANCEL request for all
2867 pending client transactions associated with this response context. A proxy SHOULD also generate a
2868 CANCEL request for all pending client transactions associated with this response context when it
2869 receives a 6xx response. A pending client transaction is one that has received a provisional response,
2870 but no final response and has not had an associated CANCEL generated for it. Generating CANCEL
2871 requests is described in Section 9.1.

2872 The requirement to CANCEL pending client transactions upon forwarding a final response does not
2873 guarantee that an endpoint will not receive multiple 200 (OK) responses to an INVITE. 200 (OK)
2874 responses on more than one branch may be generated before the CANCEL requests can be sent and
2875 processed. Further, it is reasonable to expect that a future extension may override this requirement to
2876 issue CANCEL requests.

2877 16.7 Processing Timer C

2878 If timer C should fire, the proxy MUST either reset the timer with any value it chooses, or generate a CAN-
2879 CEL for that particular request.

2880 16.8 Handling Transport Errors

2881 If the transport layer notifies a proxy of an error when it tries to forward a request (see Section 19.4), the
2882 proxy MUST behave as if the forwarded request received a 400 (Bad Request) response.

2883 If the proxy is notified of an error when forwarding a response, it drops the response. The proxy SHOULD
2884 NOT cancel any outstanding client transactions associated with this response context due to this notification.

2885 If a proxy cancels its outstanding client transactions, a single malicious or misbehaving client can cause all
2886 transactions to fail through its Via header field.

2887 16.9 CANCEL Processing

2888 A stateful proxy may generate a CANCEL to any other request it has generated at any time (subject to receiv-
2889 ing a provisional response to that request as described in section 9.1). A proxy MUST cancel any pending
2890 client transactions associated with a response context when it receives a matching CANCEL request.

2891 A stateful proxy MAY generate CANCEL requests for pending INVITE client transactions based on the
2892 period specified in the INVITEs Expires header field elapsing. However, this is generally unnecessary since
2893 the endpoints involved will take care of signaling the end of the transaction.

2894 While a CANCEL request is handled in a stateful proxy by its own server transaction, a new response
2895 context is not created for it. Instead, the proxy layer searches its existing response contexts for the server
2896 transaction handling the request associated with this CANCEL. If a matching response context is found, the
2897 element MUST immediately return a 200 (OK) response to the CANCEL request. In this case, the element is

2898 acting as a user agent server as defined in Section 8.2. Furthermore, the element MUST generate CANCEL
2899 requests for all pending client transactions in the context as described in Section 10.

2900 If a response context is not found, the element does not have any knowledge of the request to apply
2901 the CANCEL to. It MUST forward the CANCEL request (it may have statelessly forwarded the associated
2902 request previously).

2903 16.10 Stateless Proxy

2904 When acting statelessly, a proxy is a simple message forwarder. Much of the processing performed when
2905 acting statelessly is the same as when behaving statefully. The differences are detailed here.

2906 A stateless proxy does not have any notion of a transaction, or of the response context used to describe
2907 stateful proxy behavior. Instead, the stateless proxy takes messages, both requests and responses, directly
2908 from the transport layer (See section 19). As a result, stateless proxies do not retransmit messages on their
2909 own. They do, however, forward all retransmission they receive (they do not have the ability to distinguish
2910 a retransmission from the original message). Furthermore, when handling a request statelessly, an element
2911 MUST NOT generate its own 100 (Trying) or any other provisional response.

2912 A stateless proxy must validate a request as described in Section 16.3

2913 A stateless proxy must make a routing decision as described in Section 16.4 with the following excep-
2914 tion:

- 2915 • A stateless proxy MUST choose one and only one destination from the destination set. This choice
2916 MUST only rely on fields in the message and time-invariant properties of the server. In particular, a
2917 retransmitted request MUST be forwarded to the same destination each time it is processed. Further-
2918 more, CANCEL and non-Routed ACK requests MUST generate the same choice as their associated
2919 INVITE.

2920 A stateless proxy must process the request before forwarding as described in Section 16.5 with the
2921 following exceptions:

- 2922 • The requirement for unique branch IDs across time applies to stateless proxies as well. However, a
2923 stateless proxy cannot simply use a random number generator to compute the first component of the
2924 branch ID, as described in Section 16.5 bullet 3. This is because retransmissions of a request need
2925 to have the same value, and a stateless proxy cannot tell a retransmission from the original request.
2926 Therefore, the component of the branch parameter that makes it unique MUST be the same each time
2927 a retransmitted request is forwarded. Thus for a stateless proxy, the branch parameter MUST be
2928 computed as a combinatoric function of message parameters which are invariant on retransmission.
- 2929 • The stateless proxy MAY use any technique it likes to guarantee uniqueness of its branch IDs across
2930 transactions. However, the following procedure is RECOMMENDED. The proxy examines the branch
2931 ID of the received request. If it begins with the magic cookie, the first component of the branch ID of
2932 the outgoing request is computed as a hash of the received branch ID. Otherwise, the first component
2933 of the branch ID is computed as a hash of the topmost Via, the To header, the From header, the
2934 Call-ID header, the CSeq number (but not method), and the Request-URI from the received request.
2935 One of these fields will always vary across two different transactions.
- 2936 • The request is sent directly to the transport layer instead of through a client transaction. If the next-
2937 hop destination parameters don't provide an explicit destination, the element applies the procedures
2938 of [8] to the Request-URI to determine where to send the request.

2939 Since a stateless proxy must forward retransmitted requests to the same destination and add identical branch
2940 parameters to each of them, it can only use information from the message itself and time-invariant configuration
2941 data for those calculations. If the configuration state is not time-invariant (for example, if a routing table is updated)
2942 any requests that could be affected by the change may not be forwarded statelessly during an interval equal to the
2943 transaction timeout window before or after the change. The method of processing the affected requests in that
2944 interval is an implementation decision. A common solution is to forward them transaction statefully.

2945 Stateless proxies **MUST NOT** perform special processing for **CANCEL** requests. They are processed
2946 by the above rules as any other requests. In particular, a stateless proxy applies the same Route header
2947 processing to **CANCEL** requests that it applies to any other request.

2948 Response processing as described in Section 16.6 does not apply to a proxy behaving statelessly. When
2949 a response arrives at a stateless proxy, the proxy inspects the sent-by value in the first (topmost) Via header
2950 value. If that address matches the proxy (it equals a value this proxy has inserted into previous requests)
2951 the proxy **MUST** remove that value from the response and forward the result to the location indicated in the
2952 next Via header value. Unless specified otherwise, the proxy **MUST NOT** remove any other header values or
2953 the message body. If the address does not match the proxy, the message **MUST** be silently discarded.

2954 **16.11 Record-Route Example**

2955 This example demonstrates one way Record-Route header values can be constructed to satisfy the require-
2956 ments described in section 16.5 item 6 and section 16.6 item 8.

2957 Consider a proxy at server12.atlanta.com listening on port 5061 which receives the following request
2958 (many headers are omitted for brevity):

```
2959 INVITE sip:user@example.com SIP/2.0  
2960 Via: SIP/2.0/UDP callerspc.univ.edu  
2961 Contact: sip:caller@callerspc.univ.edu
```

2962 The proxy forwards this request to a UAS at sip:j_user@div11.example.com, and record-
2963 routes:

```
2964 INVITE sip:j_user@div11.example.com SIP/2.0  
2965 Via: SIP/2.0/UDP server12.atlanta.com:5061  
2966 Via: SIP/2.0/UDP callerspc.univ.edu  
2967 Record-Route: <sip:caller.8jjs@callerspc.univ.edu:5061;  
2968           maddr=server12.atlanta.com>  
2969 Contact: sip:caller@callerspc.univ.edu
```

2970 The 200 (OK) response received by the proxy will look like, in part:

```
2971 SIP/2.0 200 OK  
2972 Via: SIP/2.0/UDP server12.atlanta.com:5061  
2973 Via: SIP/2.0/UDP callerspc.univ.edu  
2974 Record-Route: <sip:caller.8jjs@callerspc.univ.edu:5061;  
2975           maddr=server12.atlanta.com>  
2976 Contact: sip:j_user@host32.div11.example.com
```

2977 The proxy modifies its Record-Route header in the response, resulting in the new response forwarded
2978 upstream:

```
2979 SIP/2.0 200 OK
2980 Via: SIP/2.0/UDP callerspc.univ.edu
2981 Record-Route: <sip:user@example.com:5061;maddr=server12.atlanta.com>
2982 Contact: sip:j_user@host32.div11.example.com
```

2983 The route set computed by the UAS is:

```
2984 sip:caller.8jjs@callerspc.univ.edu:5061;maddr=server12.atlanta.com
2985 sip:caller@callerspc.univ.edu
```

2986 and the route set computed by the UAC is:

```
2987 sip:j_user@example.com:5061;maddr=server12.atlanta.com
2988 sip:j_user@host32.div11.example.com
```

2989 17 Transactions

2990 SIP is a transactional protocol: interactions between components take place in a series of independent
2991 message exchanges. Specifically, a SIP transaction consists of a single request, and any responses to that
2992 request (which include zero or more provisional responses and one or more final responses). In the case
2993 of a transaction where the request was an INVITE (known as an INVITE transaction), the transaction also
2994 includes the ACK only if the final response was not a 2xx response. If the response was a 2xx, the ACK is
2995 not considered part of the transaction.

2996 The reason for this separation is rooted in the importance of delivering all 200 (OK) responses to an INVITE to
2997 the UAC. To deliver them all to the UAC, the UAS alone takes responsibility for retransmitting them, and the UAC
2998 alone takes responsibility for acknowledging them with ACK. Since this ACK is retransmitted only by the UAC, it
2999 is effectively considered its own transaction.

3000 Transactions have a client side and a server side. The client side is known as a client transaction, and the
3001 server side, as a server transaction. The client transaction sends the request, and the server transaction sends
3002 the response. The client and server transactions are logical functions that are embedded in any number of
3003 elements. Specifically, they exist within user agents and stateful proxy servers. Consider the example of
3004 Section 4. In this example, the UAC executes the client transaction, and its outbound proxy executes the
3005 server transaction. The outbound proxy also executes a client transaction, which sends the request to a
3006 server transaction in the inbound proxy. That proxy also executes a client transaction, which in turn, sends
3007 the request to a server transaction in the UAS. This is shown pictorially in Figure 4.

3008 A stateless proxy does not contain a client or server transaction. The transaction exists between the
3009 UA or stateful proxy on one side of the stateless proxy, and the UA or stateful proxy on the other side.
3010 As far as SIP transactions are concerned, stateless proxies are effectively transparent. The purpose of the
3011 client transaction is to receive a request from the element the client is embedded in (call this element the
3012 "Transaction User" or TU; it can be a UA or a stateful proxy), and reliably deliver the request to that server
3013 transaction. The client transaction is also responsible for receiving responses, and delivering them to the
3014 TU, filtering out any retransmissions or disallowed responses (such as a response to ACK). In the case of

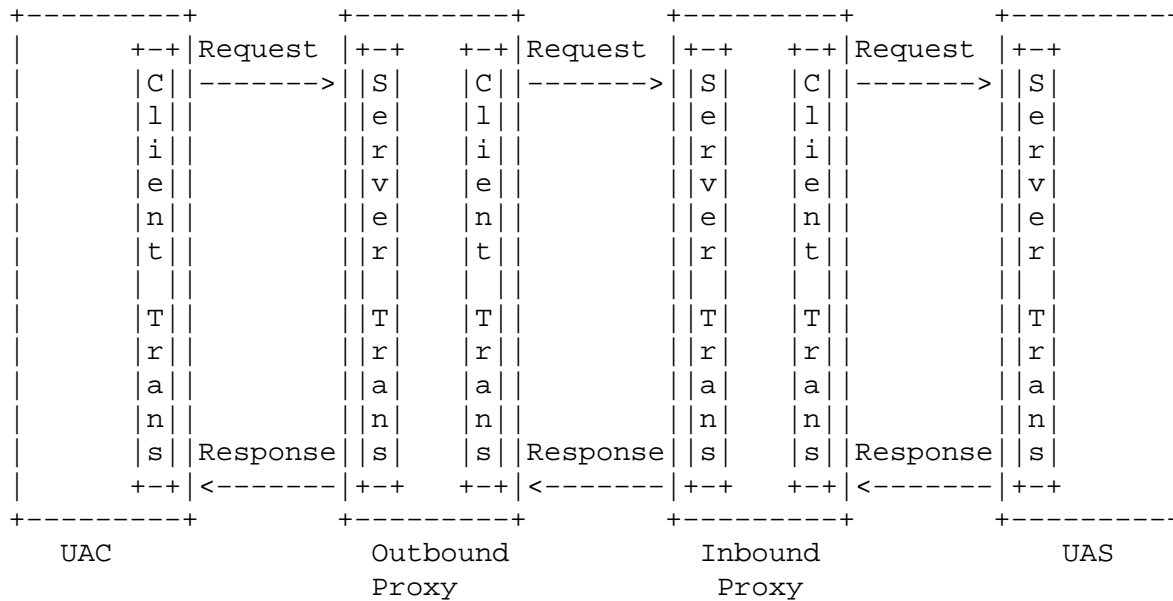


Figure 4: Transaction relationships

3015 an INVITE transaction, that includes generation of the ACK request for any final response excepting a 2xx
 3016 response.

3017 Similarly, the purpose of the server transaction is to receive requests from the transport layer, and deliver
 3018 them to the TU. The server transaction filters any request retransmissions from the network. The server
 3019 transaction accepts responses from the TU, and delivers them to the transport layer for transmission over the
 3020 network. In the case of an INVITE transaction, it absorbs the ACK request for any final response excepting
 3021 a 2xx response.

3022 The 2xx response, and the ACK for it, have special treatment. This response is retransmitted only by a
 3023 UAS, and its ACK generated only by the UAC. This end-to-end treatment is needed so that a caller knows
 3024 the entire set of users that have accepted the call. Because of this special handling, retransmissions of the
 3025 2xx response are handled by the UA core, not the transaction layer. Similarly, generation of the ACK for the
 3026 2xx is handled by the UA core. Each proxy along the path merely forwards each 2xx response to INVITE,
 3027 and its corresponding ACK.

3028 A reliable provisional response, and the PRACK for it, also have special treatment. Reliable provisional
 3029 responses are also only retransmitted by the UAS core, and the PRACK generated by the UAC core. Unlike
 3030 ACK, however, PRACK is a normal non-INVITE transaction, which means that it will generate its own final
 3031 response. The reason for this seemingly inexplicable difference between PRACK and ACK is that reliability
 3032 of provisional responses was added on later as an extra feature, and therefore needed to be done within the
 3033 confines of SIP extensibility. SIP extensibility only allowed the additions of new methods which behaved
 3034 like any other non-INVITE method.

3035 17.1 Client Transaction

3036 The client transaction provides its functionality through the maintenance of a state machine.

3037 The TU communicates with the client transaction through a simple interface. When the TU wishes to
3038 initiate a new transaction, it creates a client transaction, and passes it the SIP request to send, and an IP
3039 address, port, and transport to send it to. The client transaction begins execution of its state machine. Valid
3040 responses are passed up to the TU from the client transaction.

3041 There are two types of client transaction state machines, depending on the method of the request passed
3042 by the TU. One handles client transactions for INVITE request. This type of machine is referred to as an
3043 INVITE client transaction. Another type handles client transactions for all requests except INVITE and
3044 ACK. This is referred to as a non-INVITE client transaction. There is no client transaction for ACK. If the
3045 TU wishes to send an ACK, it passes one directly to the transport layer for transmission.

3046 The INVITE transaction is different from those of other methods because of its extended duration. Nor-
3047 mally, human input is required in order to respond to an INVITE. The long delays expected for sending a
3048 response argue for a three way handshake. Requests of other methods, on the other hand, are expected to
3049 completely rapidly. In fact, because of its reliance on just a two way handshake, TUs SHOULD respond
3050 immediately to non-INVITE requests. Protocol extensions which require longer durations for generation of
3051 a response (such as a new method that does require human interaction) SHOULD instead use two transactions
3052 - one to send the request, and another in the reverse direction to convey the result of the request.

3053 17.1.1 INVITE Client Transaction

3054 **17.1.1.1 Overview of INVITE Transaction** The INVITE transaction consists of a three-way handshake.
3055 The client transaction sends an INVITE, the server transaction sends responses, and the client transaction
3056 sends an ACK. For unreliable transports (such as UDP), the client transaction will retransmit requests at an
3057 interval that starts at T1 seconds and doubles after every retransmission. T1 is an estimate of the RTT, and
3058 it defaults to 500 ms. Nearly all of the transaction timers described here scale with T1, and changing T1 is
3059 how their values are adjusted. The request is not retransmitted over reliable transports. After receiving a
3060 1xx response, any retransmissions cease altogether, and the client waits for further responses. The server
3061 transaction can send additional 1xx responses, which are not transmitted reliably by the server transaction.
3062 If the provisional response needs to be sent reliably, this is handled by the TU. Eventually, the server trans-
3063 action decides to send a final response. For unreliable transports, that response is retransmitted periodically,
3064 and for reliable transports, its sent once. For each final response that is received at the client transaction, the
3065 client transaction sends an ACK, the purpose of which is to quench retransmissions of the response.

3066 **17.1.1.2 Formal Description** The state machine for the INVITE client transaction is shown in Figure 5.
3067 The initial state, "calling", MUST be entered when the TU initiates a new client transaction with an INVITE
3068 request. The client transaction MUST pass the request to the transport layer for transmission (see Section
3069 19). If an unreliable transport is being used, the client transaction SHOULD start timer A with a value
3070 of T1, and SHOULD NOT start timer A when a reliable transport is being used (Timer A controls request
3071 retransmissions). For any transport, the client transaction MUST start timer B with a value of 64*T1 seconds
3072 (Timer B controls transaction timeouts).

3073 When timer A fires, the client transaction SHOULD retransmit the request by passing it to the transport
3074 layer, and SHOULD reset the timer with a value of 2*T1. The formal definition of *retransmit* within the
3075 context of the transaction layer, is to take the message previously sent to the transport layer, and pass it to
3076 the transport layer once more.

3077 When timer A fires 2*T1 seconds later, the request SHOULD be retransmitted again (assuming the client
3078 transaction is still in this state). This process SHOULD continue, so that the request is retransmitted with




Figure 5: INVITE client transaction

3079 intervals that double after each transmission. These retransmissions SHOULD only be done while the client
 3080 transaction is in the “calling” state.

3081 The default value for T1 is 500 ms. T1 is an estimate of the RTT between the client and server transac-
 3082 tions. The optional RTT estimation procedure of Section 17.3 MAY be followed, in which case the resulting
 3083 estimate MAY be used instead of 500 ms. If no RTT estimation is used, other values MAY be used in private
 3084 networks where it is known that RTT has a different value. On the public Internet, T1 MAY be chosen larger,

3085 but SHOULD NOT be smaller.

3086 If the client transaction is still in the “calling” state when timer B fires, the client transaction SHOULD
3087 inform the TU that a timeout has occurred. The client transaction MUST NOT generate an ACK. The value of
3088 $64 * T1$ is equal to the amount of time required to send seven requests in the case of an unreliable transport.

3089 If the client transaction receives a provisional response while in the “calling” state, it transitions to the
3090 “proceeding” state. In the “proceeding” state, the client transaction SHOULD NOT retransmit the request any
3091 longer. Furthermore, the provisional response MUST be passed to the TU. Any further provisional responses
3092 MUST be passed up to the TU while in the “proceeding” state. Passing of all provisional responses is
3093 necessary since the TU will handle reliability of these messages, and therefore even retransmissions of a
3094 provisional response must be passed upwards.

3095 When in either the “calling” or “proceeding” states, reception of a response with status code from 300-
3096 699 MUST cause the client transaction to transition to “completed”. The client transaction MUST pass the
3097 received response up to the TU, and the client transaction MUST generate an ACK request, even if the
3098 transport is reliable (guidelines for constructing the ACK from the response are given in Section 17.1.1.3)
3099 and then pass the ACK to the transport layer for transmission. The ACK MUST be sent to the same address,
3100 port and transport that the original request was sent to. The client transaction SHOULD start timer D when it
3101 enters the “completed” state, with a value of at least 32 seconds for unreliable transports, and a value of zero
3102 seconds for reliable transports. Timer D is a reflection of the amount of time that the server transaction can
3103 remain in the “completed” state when unreliable transports are used. This is equal to Timer H in the INVITE
3104 server transaction, whose default is $64 * T1$. However, the client transaction does not know the value of $T1$
3105 in use by the server transaction, so an absolute minimum of 32s is used instead of basing Timer D on $T1$.

3106 Any retransmissions of the final response that are received while in the “completed” state SHOULD cause
3107 the ACK to be re-passed to the transport layer for retransmission, but the newly received response MUST
3108 NOT be passed up to the TU. A retransmission of the response is defined as any response which would match
3109 the same client transaction, based on the rules of Section 17.1.3.

3110 If timer D fires while the client transaction is in the “completed” state, the client transaction MUST move
3111 to the terminated state, and it MUST inform the TU of the timeout.

3112 When in either the “calling” or “proceeding” states, reception of a 2xx response MUST cause the client
3113 transaction to enter the terminated state, and the response MUST be passed up to the TU. The handling of
3114 this response depends on whether the TU is a proxy core or a UAC core. A UAC core will handle generation
3115 of the ACK for this response, while a proxy core will always forward the 200 (OK) upstream. The differing
3116 treatment of 200 (OK) between proxy and UAC is the reason that handling of it does not take place in the
3117 transaction layer.

3118 The client transaction MUST be destroyed the instant it enters the terminated state. This is actually nec-
3119 essary to guarantee correct operation. The reason is that 2xx responses to an INVITE are treated differently;
3120 each one is forwarded by proxies, and the ACK handling in a UAC is different. Thus, each 2xx needs to be
3121 passed to a proxy core (so that it can be forwarded) and to a UAC core (so it can be acknowledged). No
3122 transaction layer processing takes place. Whenever a response is received by the transport, if the transport
3123 layer finds no matching client transaction (using the rules of Section 17.1.3), the response is passed directly
3124 to the core. Since the matching client transaction is destroyed by the first 2xx, subsequent 2xx will find no
3125 match and therefore be passed to the core.

3126 **17.1.1.3 Construction of the ACK Request** The ACK request constructed by the client transaction
3127 MUST contain values for the Call-ID, From, and Request-URI which are equal to the values of those
3128 headers in the request passed to the transport by the client transaction (call this the “original request”). The

3129 To field in the ACK MUST equal the To field in the response being acknowledged, and will therefore usually
3130 differ from the To field in the original request by the addition of the tag parameter. The ACK MUST contain
3131 a single Via header, and this MUST be equal to the top Via header of the original request. The ACK request
3132 MUST contain the same Route headers as the request whose response it is acknowledging . The CSeq
3133 header in the ACK MUST contain the same value for the sequence number as was present in the original
3134 request, but the method parameter MUST be equal to "ACK".

3135 If the INVITE request whose response is being acknowledged had Route headers, those headers MUST
3136 appear in the ACK. This is to ensure that the ACK can be routed properly through any downstream stateless
3137 proxies.

3138 Although any request MAY contain a body, a body in an ACK is special since the request cannot be
3139 rejected if the body is not understood. Therefore, placement of bodies in ACK for non-2xx is NOT REC-
3140 OMMENDED, but if done, the body types are restricted to any that appeared in the INVITE, assuming that
3141 that the response to the INVITE was not 415. If it was, the body in the ACK MAY be any type listed in the
3142 Accept header in the 415.

3143 These rules for construction of ACK only apply to the client transaction. A UAC core which generates
3144 an ACK for 2xx MUST instead follow the rules described in Section 13.

3145 For example, consider the following request:

```
3146 INVITE sip:bob@biloxi.com SIP/2.0
3147 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKkjshdyff
3148 To: Bob <sip:bob@biloxi.com>
3149 From: Alice <sip:alice@atlanta.com>;tag=88sja8x
3150 Call-ID: 987asjd97y7atg
3151 CSeq: 986759 INVITE
```

3152 The ACK request for a non-2xx final response to this request would look like this:

```
3153 ACK sip:bob@biloxi.com SIP/2.0
3154 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKkjshdyff
3155 To: Bob <sip:bob@biloxi.com>;tag=99sa0xk
3156 From: Alice <sip:alice@atlanta.com>;tag=88sja8x
3157 Call-ID: 987asjd97y7atg
3158 CSeq: 986759 ACK
```

3159 17.1.2 non-INVITE Client Transaction

3160 **17.1.2.1 Overview of the non-INVITE Transaction** Non-INVITE transactions do not make use of ACK.
3161 They are a simple request-response interaction. For unreliable transports, requests are retransmitted at an
3162 interval which starts at T1, and doubles until it hits T2. If a provisional response is received, retransmis-
3163 sions continue for unreliable transports, but at an interval of T2. The server transaction retransmits the last
3164 response it sent (which can be a provisional or final response) only when a retransmission of the request is
3165 received. This is why request retransmissions need to continue even after a provisional response, they are
3166 what ensure reliable delivery of the final response.

3167 Unlike an INVITE transaction, a non-INVITE transaction has no special handling for the 2xx response.

3171 The "Trying" state is entered when the TU initiates a new client transaction with a request. When
3172 entering this state, the client transaction SHOULD set timer F to fire in $64 * T1$ seconds. The request MUST
3173 be passed to the transport layer for transmission. If an unreliable transport is in use, the client transaction
3174 MUST set timer E to fire in $T1$ seconds. If timer E fires while still in this state, the timer is reset, but this
3175 time with a value of $\text{MIN}(2 * T1, T2)$. When the timer fires again, it is reset to a $\text{MIN}(4 * T1, T2)$. This
3176 process continues, so that retransmissions occur with an exponentially increasing interval that caps at $T2$.
3177 The default value of $T2$ is 4s, and it represents the amount of time a non-INVITE server transaction will take
3178 to respond to a request, if it does not respond immediately. For the default values of $T1$ and $T2$, this results
3179 in intervals of 500 ms, 1 s, 2 s, 4 s, 4 s, 4s, etc.

3180 If Timer F fires while the client transaction is still in the "Trying" state, the client transaction SHOULD
3181 inform the TU about the timeout, and then it SHOULD enter the "Terminated" state. If a provisional response
3182 is received while in the "Trying" state, the response MUST be passed to the TU, and then the client transaction
3183 SHOULD move to the "Proceeding" state. If a final response (status codes 200-699) is received while in the
3184 "Trying" state, the response MUST be passed to the TU, and the client transaction MUST transition to the
3185 "Completed" state.

3186 If Timer E fires while in the "Proceeding" state, the request MUST be passed to the transport layer
3187 for retransmission, and Timer E MUST be reset with a value of $T2$ seconds. If timer F fires while in the
3188 "Proceeding" state, the TU MUST be informed of a timeout, and the client transaction MUST transition to the
3189 terminated state. If a final response (status codes 200-699) is received while in the "Proceeding" state, the
3190 response MUST be passed to the TU, and the client transaction MUST transition to the "Completed" state.

3191 Once the client transaction enters the "Completed" state, it MUST set Timer K to fire in $T4$ seconds for
3192 unreliable transports, and zero seconds for reliable transports. The "Completed" state exists to buffer any
3193 additional response retransmissions that may be received (which is why the client transaction remains there
3194 only for unreliable transports). $T4$ represents the amount of time the network will take to clear messages
3195 between client and server transactions. The default value of $T4$ is 5s. A response is a retransmission when it
3196 matches the same transaction, using the rules specified in Section 17.1.3. If Timer K fires while in this state,
3197 the client transaction MUST transition to the "Terminated" state.

3198 Once the transaction is in the terminated state, it MUST be destroyed. As with client transactions, this is
3199 needed to ensure reliability of the 2xx responses to INVITE.

3200 17.1.3 Matching Responses to Client Transactions

3201 When the transport layer in the client receives a response, it has to figure out which client transaction will
3202 handle the response, so that the processing of Sections 17.1.1 and 17.1.2 can take place.

3203 The branch parameter in the top *Via* header is used for this purpose. A response matches a client
3204 transaction under two conditions. First, if the response has the same value of the branch parameter in the top
3205 *Via* header as the branch parameter in the top *Via* header of the request that created the transaction. Second,
3206 if the method parameter in the *CSeq* header matches the method of the request that created the transaction.
3207 The method is needed since a CANCEL request constitutes a different transaction, but shares the same value
3208 of the branch parameter.

3209 A response which matches a transaction matched by a previous response is considered a retransmission
3210 of that response.

3211 **17.1.4 Handling Transport Errors**

3212 When the client transaction sends a request to the transport layer to be sent, the following procedures are
3213 followed if the transport layer indicates a failure.

3214 The client transaction SHOULD inform the TU that a transport failure has occurred, and the client trans-
3215 action SHOULD transition directly to the terminated state.

3216 **17.2 Server Transaction**

3217 The server transaction is responsible for the delivery of requests to the TU, and the reliable transmission of
3218 responses. It accomplishes this through a state machine. Server transactions are created by the core when a
3219 request is received, and transaction handling is desired for that request (this won't always be the case).

3220 As with the client transactions, the state machine depends on whether the received request is an INVITE
3221 request or not.

3222 **17.2.1 INVITE Server Transaction**

3223 The state diagram for the INVITE server transaction is shown in Figure 7.

3224 When a server transaction is constructed with a request, it enters the "Proceeding" state. The server
3225 transaction MUST generate a 100 response (not any status code – the specific value of 100) unless it knows
3226 that the TU will generate a provisional or final response within 200 ms, in which case it MAY generate a
3227 100 (Trying) response. This provisional response is needed to rapidly quench request retransmissions in
3228 order to avoid network congestion. The 100 response is constructed according to the procedures in Section
3229 8.2.6, except that insertion of tags in the To field of the response (when none was present in the request), is
3230 downgraded from MAY to SHOULD NOT. The request MUST be passed to the TU.

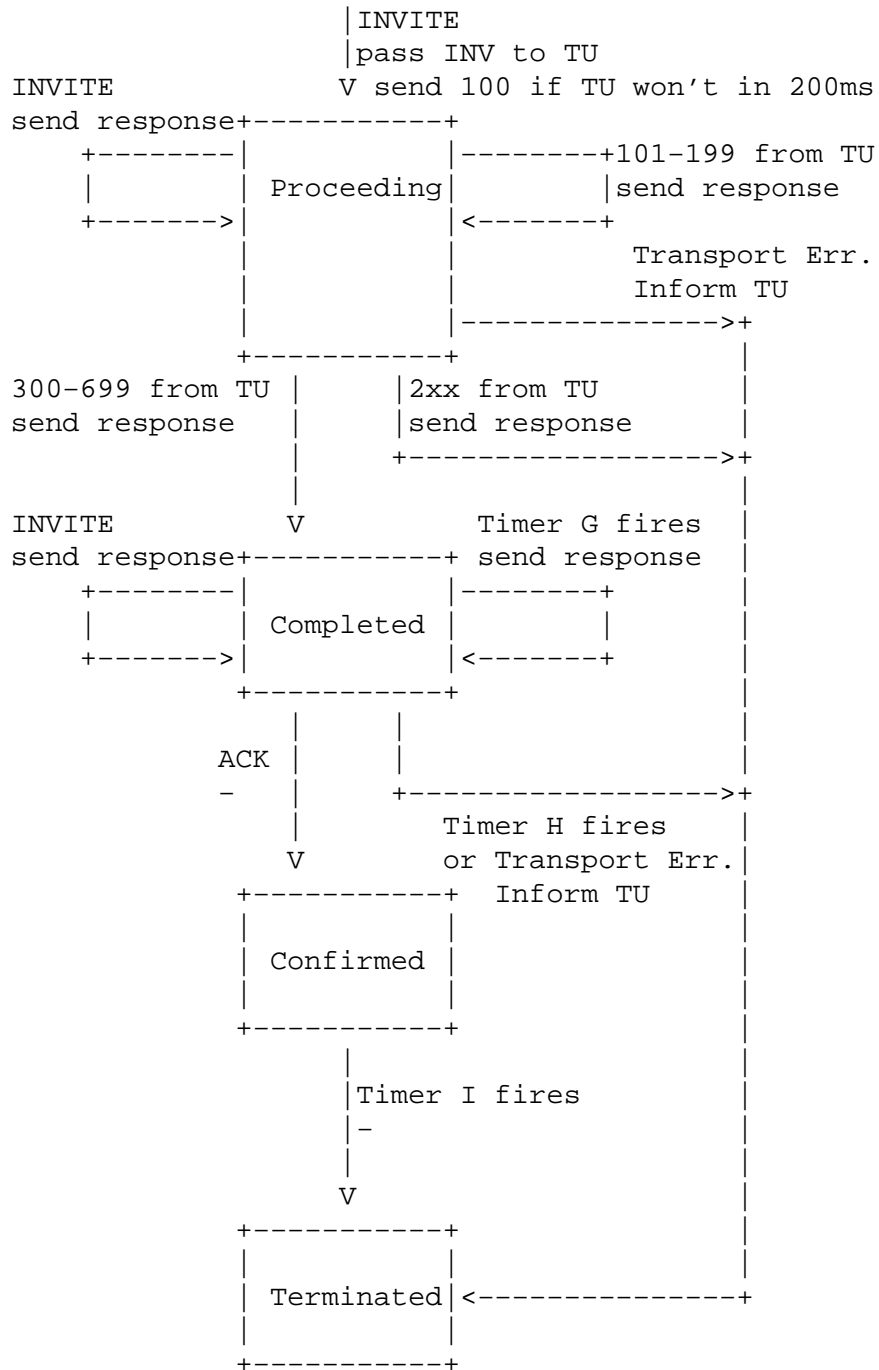
3231 The TU passes any number of provisional responses to the server transaction. So long as the server
3232 transaction is in the "Proceeding" state, each of these MUST be passed to the transport layer for transmission.
3233 They are not sent reliably by the transaction layer (they are not retransmitted by it), and do not cause a
3234 change in the state of the server transaction. When provisional responses need to be delivered reliably,
3235 it is handled by the TU, which will retransmit the provisional responses itself, and pass downwards each
3236 retransmission to the server transaction. If a request retransmission is received while in the "Proceeding"
3237 state, the most recent provisional response that was received from the TU MUST be passed to the transport
3238 layer for retransmission. A request is a retransmission if it matches the same server transaction based on the
3239 rules of Section 17.2.3.

3240 If, while in the "proceeding" state, the TU passes a 2xx Response to the server transaction, the server
3241 transaction MUST pass this response to the transport layer for transmission. It is not retransmitted by the
3242 server transaction; retransmissions of 2xx responses are handled by the TU. The server transaction MUST
3243 then transition to the "terminated" state.

3244 While in the "Proceeding" state, if the TU passes a response with status code from 300 to 699 to the
3245 server transaction, the response MUST be passed to the transport layer for transmission, and the state machine
3246 MUST enter the "Completed" state. For unreliable transports, timer G is set to fire in T1 seconds, and is not
3247 set to fire for reliable transports.

3248 This is a change from RFC 2543, where responses were always retransmitted, even over reliable transports.

3249 When the "Completed" state is entered, timer H MUST be set to fire in 64*T1 seconds, for all transports.
3250 Timer H determines when the server transaction gives up retransmitting the response. Its value is chosen to



3251 equal Timer B, the amount of time a client transaction will continue to retry sending a request. If timer G
3252 fires, the response is passed to the transport layer once more for retransmission, and timer G is set to fire in
3253 $\text{MIN}(2*T1, T2)$ seconds. From then on, when timer G fires, the response is passed to the transport again for
3254 transmission, and timer G is reset with a value that doubles, unless that value exceeds T2, in which case it
3255 is reset with the value of T2. This is identical to the retransmit behavior for requests in the “Trying” state of
3256 the non- INVITE client transaction. Furthermore, while in the “completed” state, if a request retransmission
3257 is received, the server SHOULD pass the response to the transport for retransmission.

3258 If an ACK is received while the server transaction is in the “Completed” state, the server transaction
3259 MUST transition to the “confirmed” state. As Timer G is ignored in this state, any retransmissions of the
3260 response will cease.

3261 If timer H fires while in the “Completed” state, it implies that the ACK was never received. In this case,
3262 the server transaction MUST transition to the terminated state, and MUST indicate to the TU that a transaction
3263 failure has occurred.

3264 The purpose of the “confirmed” state is to absorb any additional ACK messages that arrive, triggered
3265 from retransmissions of the final response. When this state is entered, timer I is set to fire in T4 seconds for
3266 unreliable transports, and zero seconds for reliable transports. Once timer I fires, the server MUST transition
3267 to the “Terminated” state.

3268 Once the transaction is in the terminated state, it MUST be destroyed. As with client transactions, this is
3269 needed to ensure reliability of the 2xx responses to INVITE.

3270 17.2.2 non-INVITE Server Transaction

3271 The state machine for the non-INVITE server transaction is shown in Figure 8.


3272 The state machine is initialized in the “Trying” state, and is passed a request other than INVITE or
3273 ACK when initialized. This request is passed up to the TU. Once in the “Trying” state, any further request
3274 retransmissions are discarded. A request is a retransmission if it matches the same server transaction, using
3275 the rules specified in Section 17.2.3.

3276 While in the “Trying” state, if the TU passes a provisional response to the server transaction, the server
3277 transaction MUST enter the “Proceeding” state. The response MUST be passed to the transport layer for
3278 transmission. Any further provisional responses that are received from the TU while in the “Proceeding”
3279 state MUST be passed to the transport layer for transmission. If a retransmission of the request is received
3280 while in the “Proceeding” state, the most recently sent provisional response MUST be passed to the transport
3281 layer for retransmission. If the TU passes a final response (status codes 200-699) to the server while in the
3282 “Proceeding” state, the transaction MUST enter the “Completed” state, and the response MUST be passed to
3283 the transport layer for transmission.

3284 When the server transaction enters the “Completed” state, it MUST set Timer J to fire in $64*T1$ seconds
3285 for unreliable transports, and zero seconds for reliable transports. While in the “Completed” state, the server
3286 transaction MUST pass the final response to the transport layer for retransmission whenever a retransmission
3287 of the request is received. Any other final responses passed by the TU to the server transaction MUST be
3288 discarded while in the “Completed” state. The server transaction remains in this state until Timer J fires, at
3289 which point it MUST transition to the “Terminated” state.

3290 The server transaction MUST be destroyed the instant it enters the “Terminated” state.

3291 17.2.3 Matching Requests to Server Transactions



3292 When a request is received from the network by the server, it has to be matched to an existing transaction.
3293 This is accomplished in the following manner.

3294 The branch parameter in the topmost *Via* header the request is examined. If it is present, and begins
3295 with the magic cookie "z9hG4bK", the request was generated by a client transaction compliant to this
3296 specification. Therefore, the branch parameter will be unique across all transactions sent by that client. The
3297 request matches a transaction if the branch parameter in the request is equal to the one in the top *Via* header
3298 of the request that created the transaction, the source address and port of the request are the same as the
3299 source address and port of the the request that created the transaction, and in the case of a CANCEL request,
3300 the method of the request that created the transaction was also CANCEL. This matching rule applies to both
3301 INVITE and non-INVITE transactions alike.

3302 Source address and port are used as part of the matching process because there could be duplication of branch pa-
3303 rameters from different clients; uniqueness in time is mandated for construction of the parameter, but not uniqueness
3304 in space.

3305 If the branch parameter in the top *Via* header is not present, or does not contain the magic cookie, the
3306 following procedures are used. These exist to handle backwards compatibility with RFC 2543 compliant
3307 implementations.

3308 The INVITE request matches a transaction if the Request-URI, To, From, Call-ID, CSeq, and top
3309 *Via* header match those of the INVITE request which created the transaction. In this case, the INVITE is
3310 a retransmission of the original one that created the transaction. The ACK request matches a transaction
3311 if the Request-URI, From, Call-ID, CSeq number (not the method), and top *Via* header match those of
3312 the INVITE request which created the transaction, and the To field of the ACK matches the To field of
3313 the response sent by the server transaction (which then includes the tag). Matching is done based on the
3314 matching rules defined for each of those headers. The usage of the tag in the To field helps disambiguate
3315 ACK for 2xx from ACK for other responses at a proxy which may have forwarded both responses (which can
3316 occur in unusual conditions). An ACK request that matches an INVITE transaction matched by a previous
3317 ACK is considered a retransmission of that previous ACK.

3318 For all other request methods, a request is matched to a transaction if the Request-URI, To, From,
3319 Call-ID and Cseq (including the method) and top *Via* header match those of the request which created
3320 the transaction. Matching is done based on the matching rules defined for each of those headers. When a
3321 non-INVITE request matches an existing transaction, it is a retransmission of the request which created that
3322 transaction.

3323 Because the matching rules include the Request-URI, the server cannot match a response to a transac-
3324 tion. When the TU passes a response to the server transaction, it must pass it to the specific server transaction
3325 for which the response is targeted.

3326 17.2.4 Handling Transport Errors

3327 When the server transaction sends a response to the transport layer to be sent, the following procedures are
3328 followed if the transport layer indicates a failure.

3329 First, the procedures in [8] are followed, which attempt to deliver the response to a backup. If those
3330 should all fail, such that all elements generate ICMP errors, or no SRV records are present, the server
3331 transaction SHOULD inform the TU that a failure has occurred, and SHOULD transition to the terminated
3332 state.

3333 17.3 RTT Estimation

3334 Most of the timeouts used in the transaction state machines derive from T1, which is an estimate of the RTT
3335 between the client and server transactions. This subsection defines optional procedures that a client can use
3336 to build up estimates of the RTT to a particular IP address. To perform this procedure, the client MUST
3337 maintain a table of variables for each destination IP address to which an RTT estimate is being made.

3338 If a client wishes to measure RTT for a particular IP address, it MUST include a `Timestamp` header into a
3339 request containing the time when the request is initially created and passed to a new client transaction, which
3340 transmits the request. If a 100 (Trying) response (not any 1xx, only the 100 (Trying) response) is received
3341 before the client transaction generates a retransmission, an RTT estimate is made. This is consistent with
3342 the RFC 2988 requirements on TCP for using Karn's algorithm in RTT estimation.

3343 The estimate, called R, is made by computing the difference between the current time and the value of
3344 `Timestamp` header in the 100 response. The value of R is applied to the estimation of RTO as described
3345 in Section 2 of RFC 2988 [22], with the following differences. First, the initial value of RTO is 500 ms for
3346 SIP, not 3 s as is used for TCP. Second, there is no minimum value for the RTO, as there is for TCP, if SIP
3347 is being run on a private network. When run on the public Internet, the minimum is 500 ms, as opposed to
3348 1 s for TCP. This difference is because of the expected usage of SIP in private networks where rapid call
3349 setup times are service critical. Once RTO is computed, the timer T1 is set to the value of RTO, and all other
3350 timers scale proportionally as described above.

3351 This value of T1 would be used for scaling all of the client and server transaction timers described above,
3352 when a request or response, respectively, is sent to that IP address.

3353 If the IP address is that of a stateless proxy, the actual round trip time that is measured will be the average
3354 to all transaction stateful proxies or UAs that are reached through the stateless proxy. This estimate may
3355 therefore be too low or too high for a specific transactional element being communicated with through the
3356 stateless proxy.

3357 18 Reliability of Provisional Responses

3358 Normally, provisional responses are not transmitted reliably. The TU generates a single provisional re-
3359 sponse, and passes it to the server transaction, which sends it once. RFC 2543 provided no means for
3360 reliable transmission of these messages.

3361 It was later observed that reliability was important in several cases, including interoperability scenarios
3362 with the PSTN. Therefore, an optional capability was added in this specification to support reliable trans-
3363 mission of provisional responses.

3364 The reliability mechanism works by mirroring the current reliability mechanisms for 2xx final responses
3365 to INVITE. Those requests are transmitted periodically by the TU, until a separate transaction, ACK, is
3366 received, that indicates reception of the 2xx by the UAC. The reliability for the 2xx to INVITE and ACK
3367 messages are end-to-end. In order to achieve reliability for provisional responses, we do nearly the same
3368 thing. Reliable provisional responses are retransmitted by the TU with an exponential backoff. Those
3369 retransmissions cease when a PRACK message is received. The PRACK request plays the same role as
3370 ACK, but for provisional responses. There is an important difference, however. PRACK is a normal SIP
3371 message, like BYE. As such, its own reliability is ensured hop-by-hop through each stateful proxy. Similarly,
3372 PRACK has its own response. If this were not the case, the PRACK message could not traverse existing
3373 proxy servers.

3374 Each provisional response is given a sequence number, carried in the `RSeq` header in the response.

3375 The PRACK messages contain an RACK header, which indicates the sequence number of the provisional
3376 response which is being acknowledged. The acknowledgements are not cumulative, and the specifications
3377 recommend a single outstanding provisional response at a time, for purposes of congestion control.

3378 18.1 UAS Behavior

3379 A UAS MAY send any non-100 provisional response to INVITE reliably, so long as the initial INVITE request
3380 (the request whose provisional response is being sent reliably) contained a Supported header with the op-
3381 tion tag 100rel. While this specification does not allow reliable provisional responses for any method but
3382 INVITE, extensions that define new methods which can establish dialogs may make use of the mechanism.

3383 The UAS MUST send any non-100 provisional response reliably if the initial request contained a Require
3384 header with the option tag 100rel. If the UAS is unwilling to do so, it MUST reject the initial request with
3385 a 420 (Bad Extension) and include a Unsupported header containing the option tag 100rel.

3386 A UAS MUST NOT attempt to send a 100 (Trying) response reliably. Only provisional responses num-
3387 bered 101 to 199 may be sent reliably. If the request did not include either a Supported or Require header
3388 indicating this feature, the UAS MUST NOT send the provisional response reliably.

3389 100 (Trying) responses are hop-by-hop only. For this reason, the reliability mechanisms described here, which
3390 are end-to-end, cannot be used.

3391 An element which can act as a proxy can also send reliable provisional responses; in that case, it acts as
3392 a UAS for purposes of that transaction. However, it MUST NOT attempt to do so for any request that contains
3393 a tag in the To field. That is, a proxy cannot generate reliable provisional responses to requests sent within
3394 the context of a dialog. Of course, unlike a UAS, when the proxy element receives a PRACK that does not
3395 match any outstanding reliable provisional response, the PRACK MUST be proxied.

3396 The rest of this discussion assumes that the initial request contained a Supported or Require header
3397 listing 100rel, and that there is a provisional response to be sent reliably.

3398 The provisional response to be sent reliably is constructed by the UAS core according to the procedures
3399 of Section 8.2.6 and Section 12. Specifically, the provisional response MUST establish a dialog if one
3400 is not yet created. In addition, it MUST contain Require header containing the option tag 100rel, and
3401 MUST include an RSeq header. The value of the header for the first reliable provisional response in a
3402 transaction MUST be between 1 and $2^{31} - 1$. It is RECOMMENDED that it be chosen uniformly in this
3403 range. The RSeq numbering space is within a single transaction. This means that provisional responses for
3404 different requests MAY use the same values for the RSeq number.

3405 The reliable provisional response is passed to the transaction layer periodically with an interval that
3406 starts at T1 seconds and doubles for each retransmission (T1 is defined in Section 17). Once passed to the
3407 server transaction, it is added to an internal list of unacknowledged reliable provisional responses.

3408 This differs from retransmissions of 2xx responses, which cap at T2 seconds. This is because retransmissions of
3409 ACK are triggered on receipt of a 2xx, but retransmissions of PRACK take place independently of reception of 1xx.

3410 Retransmissions cease when a matching PRACK is received. PRACK is like any other request within a
3411 dialog, and the UAS core processes it according to the procedures of Sections 8.2 and 12.2.2. A matching
3412 PRACK is defined as one within the same dialog as the response, and whose method, CSeq-num, and
3413 response-num in the RACK header match, respectively, the method and sequence number from the CSeq
3414 and sequence number from the RSeq of the reliable provisional response.

3415 If a PRACK request is received that does not match any unacknowledged reliable provisional response,
3416 the UAS MUST respond to the PRACK with a 481 response. If the PRACK does match an unacknowledged

3417 reliable provisional response, it MUST be responded to with a 2xx response. The UAS can be certain at
3418 this point that the provisional response has been received in order. It SHOULD cease retransmissions of the
3419 reliable provisional response, and MUST remove it from the list of unacknowledged provisional responses.

3420 If a reliable provisional response is retransmitted for 64*T1 seconds without reception of a correspond-
3421 ing PRACK, the UAS SHOULD reject the original request with a 5xx response.

3422 If the PRACK contained a body, the body is treated in the same way a body in an ACK is treated.

3423 After the first reliable provisional response for a request has been acknowledged, the UAS MAY send
3424 additional reliable provisional responses. The UAS MUST NOT send a second reliable provisional response
3425 until the first is acknowledged. After the first, it is RECOMMENDED that the UAS not send an additional
3426 reliable provisional response until the previous is acknowledged. The first reliable provisional response
3427 receives special treatment because it conveys the initial sequence number. If additional reliable provisional
3428 responses were sent before the first was acknowledged, the UAS could not be certain these were received in
3429 order.

3430 The value of the RSeq in each subsequent reliable provisional response for the same request MUST be
3431 greater by exactly one. RSeq numbers MUST NOT wrap around. Because the initial one is chosen to be less
3432 than $2^{31} - 1$, but the maximum is $2^{32} - 1$, there can be up to 2^{31} reliable provisional responses per
3433 request, which is more than sufficient.

3434 Note that the UAS MAY send a final response to the initial request before having received PRACKs for
3435 all unacknowledged reliable provisional responses. In that case, it SHOULD NOT continue to retransmit the
3436 unacknowledged reliable provisional responses, but it MUST be prepared to process PRACK requests for
3437 those outstanding responses. A UAS MUST NOT send new reliable provisional responses (as opposed to
3438 retransmissions of unacknowledged ones) after sending a final response to a request.

3439 18.2 UAC Behavior

3440 If a provisional response is received for the initial request, and that response contains a Require header
3441 containing the option tag 100rel, the response is to be sent reliably. If the response is a 100 (Trying) (as
3442 opposed to 101 to 199), this option tag MUST be ignored, and the procedures below MUST NOT be used.

3443 Assuming the response is to be transmitted reliably, the UAC MUST create a new request with method
3444 PRACK. This request is sent within the dialog associated with the provisional response (indeed, the provi-
3445 sional response may have created the dialog). PRACK requests MAY contain bodies, which are interpreted
3446 according to their type and disposition.

3447 Note that the PRACK is like any other non-INVITE request within a dialog. In particular, a UAC
3448 SHOULD NOT retransmit the PRACK request when it receives a retransmission of the provisional response
3449 being acknowledged, although doing so does not create a protocol error.

3450 Once a reliable provisional response is received, retransmissions of that response MUST be discarded.
3451 A response is a retransmission when its dialog ID, CSeq and RSeq match the original response. The UAC
3452 MUST maintain a sequence number which indicates the most recently received in-order reliable provisional
3453 response for the initial request. This sequence number MUST be maintained until a final response is received
3454 for the initial request. Its value MUST be initialized to the RSeq header in the first reliable provisional
3455 response received for the initial request.

3456 Handling of subsequent reliable provisional responses for the same initial request follows the same rules
3457 as above, with the following difference. Reliable provisional responses are guaranteed to be in order. As
3458 a result, if the UAC receives another reliable provisional response to the same request, and its RSeq value
3459 isn't one higher than the value of the sequence number, that response MUST NOT be acknowledged with a

3460 PRACK, and MUST NOT be processed further by the TU. An implementation MAY discard the response, or
3461 MAY cache the response in the hopes of receiving the missing responses.

3462 The UAC MAY acknowledge reliable provisional responses received after the final response, or MAY
3463 discard them.

3464 19 Transport

3465 The transport layer is responsible for the actual transmission of requests and responses over network trans-
3466 ports. This includes determination of the connection to use for a request or response, in the case of connec-
3467 tion oriented transports.

3468 The transport layer is responsible for managing any persistent connections (for transports like TCP, TLS
3469 and SCTP) including ones it opened, as well as ones opened to it. This includes connections opened by
3470 the client or server transports, so that connections are shared between client and server transport functions.
3471 These connections are indexed by the [address, port, transport] at the far end of the connection. When a
3472 connection is opened by the transport layer, this index is set to the destination IP, port and transport. When
3473 the connection is accepted by the transport layer, this index is set to the source IP, port and transport. Note
3474 that, because the source port is often ephemeral, connections accepted by the transport layer will frequently
3475 not be reused. The result is that two proxies in a “peering” relationship using a connection oriented transport
3476 will frequently have two connections in use, one for transactions initiated in each direction.

3477 It is RECOMMENDED that connections be kept open for some implementation defined duration after the
3478 last message was sent or received over that connection. This duration SHOULD at least equal the longest
3479 amount of time the element would need in order to bring a transaction from instantiation to the terminated
3480 state. This is to insure that transactions complete over the same connection they are initiated on (i.e., re-
3481 quest, response, and in the case of INVITE, ACK for non-2xx responses)). This usually means at least the
3482 maximum of T3 and 64*T1. However, it could be larger in an element that has a TU that is using a large
3483 value for timer C, for example.

3484 All SIP elements MUST implement UDP and TCP. Other transports MAY be implemented by any entity.

3485 Making TCP mandatory for UA is a substantial change from RFC 2543. It has arisen out of the need to handle
3486 larger messages, which MUST use TCP, as discussed below. Thus, even if an element never sends large messages, it
3487 may receive one, and needs to be able to do that.

3488 19.1 Clients

3489 19.1.1 Sending Requests

3490 The client side of the transport layer is responsible for sending the request and receiving responses. The
3491 user of the transport layer passes the client transport the request, an IP address, port, transport, and possibly
3492 TTL for multicast destinations.

3493 If a request is within 500 bytes of the path MTU, or if it is larger than 1000 bytes when the path MTU is
3494 unknown, it MUST be sent using TCP. This is to prevent fragmentation of messages over UDP, and to provide
3495 congestion control for larger messages. However, implementations MUST be able to handle messages up to
3496 the maximum datagram packet size. For UDP, this size is 65,535 bytes, including headers.

3497 The 500 byte “buffer” between the message size and the MTU accomodates the fact that the response in SIP can
3498 be larger than the request. This happens due to the addition of Record-Route headers to the responses to INVITE,
3499 for example. With the extra buffer, the response can be 500 bytes larger than the request, and still not be fragmented.
3500 1000 is chosen when path MTU is not known, based on the assumption of a 1500 byte ethernet MTU

3501 A client that sends a request to a multicast address MUST add the “maddr” parameter to its Via header
3502 field, and SHOULD add the “ttl” parameter. (In that case, the maddr parameter SHOULD contain the des-
3503 tination multicast address, although under exceptional circumstances it MAY contain a unicast address.)
3504 Requests sent to multicast groups SHOULD be scoped to ensure that they are not forwarded beyond the
3505 administrative domain to which they were targeted. This scoping MAY be done with either TTL or adminis-
3506 trative scopes [17], depending on what is implemented in the network.

3507 It is important to note that the layers above the transport layer do not operate differently for multicast
3508 as opposed to unicast requests. This means that SIP treats multicast more like anycast, assuming that there
3509 is a single recipient generating responses to requests. If this is not the case, the first response will end
3510 up “winning”, based on the client transaction rules. Any other responses from different UA will appear
3511 as retransmissions and be discarded. This limits the utility of multicast to cases where an anycast type of
3512 function is desired, such as registrations.

3513 Before a request is sent, the client transport MUST insert a value of the sent-by field into the Via header.
3514 This field contains an IP address or host name, and port. The usage of an FQDN is RECOMMENDED. This
3515 field is used for sending responses under certain conditions.

3516 For reliable transports, the response is normally sent on the connection the request was received on.
3517 Therefore, the client transport MUST be prepared to receive the response on the same connection used to
3518 send the request. Under error conditions, the server may attempt to open a new connection to send the
3519 response. To handle this case, the transport layer MUST also be prepared to receive an incoming connection
3520 on the source IP address that the request was sent from, and port number in the sent-by field. It also MUST
3521 be prepared to receiving incoming connections on any address and port which would be selected by a server
3522 based on the procedures described in Section 5 of [8].

3523 For unreliable unicast transports, the client transport MUST be prepared to receive responses on the
3524 source IP address that the request is sent from (as responses are sent back to the source address), but the port
3525 number in the sent-by field. Furthermore, as with reliable transports, in certain cases the response will be
3526 sent elsewhere. The client MUST be prepared to receive responses on any address and port which would be
3527 selected by a server based on the procedures described in Section 5 of [8].

3528 For multicast, the client transport MUST be prepared to receive responses on the same multicast group
3529 and port that the request is sent to (e.g., it needs to be a member of the multicast group it sent the request
3530 to.)

3531 If a request is destined to an IP address, port, and transport to which an existing connection is open, it
3532 is RECOMMENDED that this connection be used to send the request, but another connection MAY be opened
3533 and used.

3534 If a request is sent using multicast, it is sent to the group address, port, and TTL provided by the transport
3535 user. If a request is sent using unicast unreliable transports, it is sent to the IP address and port provided by
3536 the transport user.

3537 19.1.2 Receiving Responses

3538 When a response is received, the client transport examines the top Via header. If the value of the sent-by
3539 parameter in that header does not correspond to a value that the client transport is configured to insert into
3540 requests, the response MUST be rejected.

3541 If there are any client transactions in existence, the client transport uses the matching procedures of Sec-
3542 tion 17.1.3 to attempt to match the response to an existing transaction. If there is a match, the response MUST
3543 be passed to that transaction. Otherwise, the response MUST be passed to the core (whether it be stateless

3544 proxy, stateful proxy, or UA) for further processing. Handling of these “stray” responses is dependent on
3545 the core (a stateless proxy will forward all responses, for example).

3546 19.2 Servers

3547 19.2.1 Receiving Requests

3548 When the server transport receives a request over any transport, it MUST examine the value of the sent-by
3549 parameter in the top Via header field. If the host portion of the sent-by parameter contains a domain name,
3550 or if it contains an IP address that differs from the packet source address, the server MUST add a “received”
3551 attribute to that Via header field. This attribute MUST contain the source address that the packet was received
3552 from. This is to assist the server transport layer in sending the response, since it must be sent to the source
3553 IP address that the request came from.

3554 Consider a request received by the server transport which looks like, in part:

```
3555 INVITE sip:bob@Biloxi.com SIP/2.0  
3556 Via: SIP/2.0/UDP bobspc.biloxi.com:5060
```

3557 The request is received with a source IP address of 1.2.3.4. Before passing the request up, the transport
3558 would add a received parameter, so that the request would look like, in part:

```
3559 INVITE sip:bob@Biloxi.com SIP/2.0  
3560 Via: SIP/2.0/UDP bobspc.biloxi.com:5060;received=1.2.3.4
```

3561 Next, the server transport attempts to match the request to the server transaction. It does so using
3562 the matching rules described in Section 17.2.3. If a matching server transaction is found, the request is
3563 passed to that transaction for processing. If no match is found, the request is passed to the core, which
3564 may decide to construct a new server transaction for that request. Note that when a UAS core sends a 2xx
3565 response to INVITE, the server transaction is destroyed. This means that when the ACK arrives, there will
3566 be no matching server transaction, and based on this rule, the ACK is passed to the UAS core, where it is
3567 processed.

3568 19.2.2 Sending Responses

3569 The server transport uses the value of the top Via header in order to determine where to send a response. It
3570 MUST follow the following process:

- 3571 • If the “sent-protocol” is a reliable transport protocol such as TCP, TLS or SCTP, the response MUST
3572 be sent using the existing connection to the source of the original request that created the transaction, if
3573 that connection is still open. This does require the server transport to maintain an association between
3574 server transactions and transport connections. If that connection is no longer open, the server MAY
3575 open a connection to the IP address in the received parameter, if present, using the port in the sent-
3576 by value, or the default port for that transport, if no port is specified (5060 for UDP and TCP, 5061
3577 for TLS and SSL). If that connection attempt fails, the server SHOULD use the procedures in [8] for
3578 servers in order to determine the IP address and port to open the connection and send the response to.

3579

- 3580 • Otherwise, if the *Via* header field contains a “*maddr*” parameter, forward the response to the address
3581 listed there, using the port indicated in “*sent-by*”, or port 5060 if none is present. If the address is
3582 a multicast address, the response SHOULD be sent using the TTL indicated in the “*ttl*” parameter, or
3583 with a TTL of 1 if that parameter is not present.

- 3584 • Otherwise (for unreliable unicast transports), if the top *Via* has a *received* parameter, send the re-
3585 sponse to the address in the “*received*” parameter, using the port indicated in the “*sent-by*” value, or
3586 using port 5060 if none is specified explicitly. If this fails, e.g., elicits an ICMP “port unreachable”
3587 response, send the response to the address in the “*sent-by*” parameter. The address to send to is
3588 determined by following the procedures defined in Section 5 of [8].

- 3589 • Otherwise, if it is not receiver-tagged, send the response to the address indicated by the “*sent-by*”
3590 value, using the procedures in Section 5 of [8].

3591 **19.3 Framing**

3592 In the case of message oriented transports (such as UDP), if the message has a *Content-Length* header, the
3593 message body is assumed to contain that many bytes. If there are additional bytes in the transport packet
3594 below the end of the body, they MUST be discarded. If the transport packet ends before the end of the
3595 message body, this is considered an error. If the message is a response, it MUST be discarded. If its a
3596 request, the element SHOULD generate a 400 class response. If the message has no *Content-Length* header,
3597 the message body is assumed to end at the end of the transport packet.

3598 In the case of stream oriented transports (such as TCP), the *Content-Length* header indicates the size
3599 of the body. The *Content-Length* header MUST be used with stream oriented transports.

3600 **19.4 Error Handling**

3601 Error handling is independent of whether the message was a request or response.

3602 If the transport user asks for a message to be sent over an unreliable transport, and the result is an ICMP
3603 error, the behavior depends on the type of ICMP error. A host, network, port or protocol unreachable errors,
3604 or parameter problem errors SHOULD cause the transport layer to inform the transport user of a failure in
3605 sending. Source quench and TTL exceeded ICMP errors SHOULD be ignored.

3606 If the transport user asks for a request to be sent over a reliable transport, and the result is a connection
3607 failure, the transport layer SHOULD inform the transport user of a failure in sending.

3608 **20 Usage of HTTP Authentication**

3609 SIP provides a stateless challenged-based mechanism for authentication that is based on authentication in
3610 HTTP. Any time that a proxy server or user agent receives a request (with the exceptions given in Sec-
3611 tion 20.1), it MAY challenge the initiator of the request to provide assurance of its identity. Once the
3612 originator has been identified, the recipient of the request SHOULD ascertain whether or not this user is au-
3613 thorized to make the request in question. No authorization systems are recommended or discussed in this
3614 document.

3615 The “Digest” authentication mechanism described in this section provides message authentication and
3616 replay protection only, without message integrity or confidentiality. Protective measures above and beyond

3617 those provided by Digest need to be taken to prevent active attackers from modifying SIP requests and
3618 responses.

3619 Note that due to its weak security, the usage of "Basic" authentication has been deprecated. Servers
3620 MUST NOT accept credentials using the "Basic" authorization scheme, and servers also MUST NOT challenge
3621 with "Basic". This is a change from RFC 2543.

3622 20.1 Framework

3623 The framework for SIP authentication closely parallels that of HTTP (RFC 2617 [23]). In particular,
3624 the BNF for auth- scheme, auth-param, challenge, realm, realm-value, and credentials is identical
3625 (although the usage of "Basic" as a scheme is not permitted). The 401 (Unauthorized) response is used
3626 by user agent servers in SIP to challenge the identity of a user agent client. Additionally, registrars and
3627 redirect servers MAY make use of 401 (Unauthorized) responses for authentication, but proxies MUST NOT,
3628 and instead MAY use the 407 (Proxy Authentication Required) response. The requirements for inclusion of
3629 the Proxy-Authenticate, Proxy- Authorization, WWW-Authenticate, and Authorization in the various
3630 messages are identical to those described in RFC 2617 [23].

3631 Since SIP does not have the concept of a canonical root URL, the notion of protection spaces is in-
3632 terpreted differently in SIP. The realm string alone defines the protection domain. This is a change from
3633 RFC 2543, in which the Request-URI and the realm together defined the protection domain; this definition
3634 gave rise to some amount of confusion since the Request-URI sent by the UAC and the Request-URI
3635 received by the server issuing a challenge might be different, and indeed the final form of the Request-URI
3636 might not be known to the UAC. Also, the previous definition depended on the presence of a SIP URI in the
3637 Request-URI, and seemed to rule out alternative URI schemes (like for example the tel URL).

3638 Operators of user agents or proxy servers that will authenticate received requests MUST adhere to the
3639 following guidelines for creation of a realm string for their server:

- 3640 • Realm strings MUST be globally unique. It is RECOMMENDED that a realm string contain a hostname
3641 or domain name, following the recommendation in Section 3.2.1 of RFC 2617 [[23]].
- 3642 • Realm strings SHOULD present a human-readable identifier that can be rendered to a user.

3643 For example:

```
3644 INVITE sip:bob@biloxi.com SIP/2.0  
3645 WWW-Authenticate: Digest realm="biloxi.com", <...>
```

3646 Generally, SIP authentication is meaningful for a specific realm, a protection domain. Thus, for Digest
3647 authentication, each such protection domain has its own set of user names and secrets. If a server does not
3648 care about authenticating individual users, it may make sense to establish a "global" user name and secret
3649 for its realm as a default challenge if a particular Request-URI does not have its own realm or set of user
3650 names, For example, an INVITE to 'sip:100.3.6.6'. Similarly, UACs representing many users, such as PSTN
3651 gateways, MAY have their own device-specific credentials for particular realms.

3652 While a server can legitimately challenge most SIP requests, there are two requests defined by the SIP
3653 standard today that require special handling for authentication: ACK and CANCEL.

3654 Complications of the ACK method arise because it requires no response. Under an authentication
3655 scheme that uses responses to carry values used to compute nonces (such as Digest), some problems come

3656 up for any requests that take no response (including ACK). For this reason any credentials in the INVITE that
3657 were accepted by a server MUST be accepted by that server for the ACK. UACs creating an ACK message
3658 should duplicate all of the Authorization and Proxy-Authorization headers that appeared in the INVITE to
3659 which the ACK corresponds. Servers MUST NOT attempt to challenge an ACK.

3660 Although the CANCEL method does take a response (a 2xx), servers MUST NOT attempt to challenge
3661 CANCEL requests since these requests cannot be resubmitted. Generally, a CANCEL request SHOULD be
3662 accepted by a server if it comes from the same host that sent the request being cancelled (provided that some
3663 sort of transport or network layer security association, as described in Section 22.2.1, is in place).

3664 When a challenge is received by a UAC, it SHOULD render to the user the contents of the "realm"
3665 parameter in the challenge (which appears in either a WWW-Authenticate header or Proxy-Authenticate
3666 header) if the UAC device does not already know of a credential for the realm in question. A service
3667 provider that pre-configures UAs with credentials for its realm should be aware that users will not have the
3668 opportunity to present their own credentials for this realm when challenged at a pre-configured device.

3669 Finally, note that even if a UAC can locate credentials that are associated with the proper realm, there is
3670 always a potential that these credentials may no longer be valid, or that for whatever reason the challenging
3671 server will not accept these credentials. In this instance a server will commonly repeat its challenge. A
3672 UAC MUST NOT reattempt requests with the credentials that have just been rejected (unless the request was
3673 rejected because of a stale nonce).

3674 20.2 User-to-User Authentication

3675 When a UAS receives a request from a UAC, the UAS MAY authenticate the originator before the request
3676 is processed. If no credentials (in the Authorization header field) are provided in the request, the UAS
3677 can challenge the originator to provide credentials by rejecting the request with a 401 (Unauthorized) status
3678 code.

3679 The WWW-Authenticate response-header field MUST be included in 401 (Unauthorized) response mes-
3680 sages. The field value consists of at least one challenge that indicates the authentication scheme(s) and
3681 parameters applicable to the Request-URI. See [H14.47] for a definition of the syntax.

3682 An example of the WWW-Authenticate header field in a 401 challenge is:

```
3683 WWW-Authenticate: Digest
3684     realm="biloxi.com",
3685     qop="auth,auth-int",
3686     nonce="dcd98b7102dd2f0e8b11d0f600bfb0c093",
3687     opaque="5ccc069c403ebaf9f0171e9517f40e41"
```

3688 When the originating UAC receives the 401 (Unauthorized), it SHOULD, if it is able, re-originate the
3689 request with the proper credentials. The UAC may require input from the originating user before proceeding.
3690 Once authentication credentials have been supplied (either directly by the user, or discovered in an internal
3691 keyring), user agents SHOULD cache the credentials for a given value of the To header and "realm" and
3692 attempt to re-use these values on the next request for that destination. UAs MAY cache credentials in any
3693 way they would like.

3694 Once credentials have been located, any user agent that wishes to authenticate itself with a UAS or reg-
3695 istrar – usually, but not necessarily, after receiving a 401 (Unauthorized) response – MAY do so by including
3696 an Authorization header field with the request. The Authorization field value consists of credentials con-

3697 taining the authentication information of the user agent for the realm of the resource being requested as well
3698 as parameters required in support of authentication and replay protection.

3699 An example of the Authorization header is:

```
3700 Authorization: Digest username="bob",  
3701     realm="biloxi.com",  
3702     nonce="dcd98b7102dd2f0e8b11d0f600bfb0c093",  
3703     uri=sip:alice@atlanta.com,  
3704     qop=auth,  
3705     nc=00000001,  
3706     cnonce="0a4f113b",  
3707     response="6629fae49393a05397450978507c4ef1",  
3708     opaque="5ccc069c403ebaf9f0171e9517f40e41"  
3709
```

3710 When a UAC resubmits a request with its credentials after receiving a 401 (Unauthorized) or 407 (Proxy
3711 Authentication Required) response, it MUST increment the CSeq header field as it would normally when
3712 sending an updated request.

3713 20.3 Proxy to User Authentication

3714 Similarly, when a UAC sends a request to a proxy server, the proxy server MAY authenticate the originator
3715 before the request is processed. If no credentials (in the Proxy-Authorization header field) are provided
3716 in the request, the UAS can challenge the originator to provide credentials by rejecting the request with a
3717 407 (Proxy Authentication Required) status code. The proxy MUST populate the 407 (Proxy Authentication
3718 Required) message with a Proxy-Authenticate header applicable to the proxy for the requested resource.

3719 The use of the Proxy-Authentication and Proxy-Authorization parallel that described in [23, Sec-
3720 tion 3.6], with one difference. Proxies MUST NOT add the Proxy-Authorization header. 407 (Proxy Au-
3721 thentication Required) responses MUST be forwarded upstream towards the UAC following the procedures
3722 for any other response. It is the client's responsibility to add the Proxy-Authorization header containing
3723 credentials for the realm of the proxy which has asked for authentication.

3724 If a proxy were to resubmit a request with a Proxy-Authorization header field, it would need to increment the
3725 CSeq in the new request. However, this would mean that the UAC which submitted the original request would
3726 discard a response from the UAS, as the CSeq value would be different.

3727 When the originating UAC receives the 407 (Proxy Authentication Required) it SHOULD, if it is able,
3728 re-originate the request with the proper credentials. It should follow the same procedures for the display
3729 of the "realm" parameter that are given above for responding to 401. The UAC SHOULD also cache the
3730 credentials used in the re-originated request.

3731 The following rule is RECOMMENDED for proxy credential caching:

3732 If a UA receives a Proxy-Authenticate header in a 401/407 response to a request with a particular Call-
3733 ID, it should incorporate credentials for that realm in all subsequent requests that contain the same Call-ID.
3734 These credentials MUST NOT be cached across dialogs; however, if a UA is configured with the realm of its
3735 local outbound proxy, when one exists, then the UA MAY cache credentials for that realm across dialogs.
3736 Note that this does mean a future requests in a dialog could contain credentials that are not needed by any
3737 proxy along the Route header path.

3738 Any user agent that wishes to authenticate itself to a proxy server – usually, but not necessarily, after
3739 receiving a 407 (Proxy Authentication Required) response – MAY do so by including a Proxy-Authorization
3740 header field with the request. The Proxy-Authorization request-header field allows the client to identify
3741 itself (or its user) to a proxy which requires authentication. The Proxy-Authorization header field value
3742 consists of credentials containing the authentication information of the user agent for the proxy and/or realm
3743 of the resource being requested.

3744 A Proxy-Authorization header field applies only to the proxy whose realm is identifier in the “realm”
3745 parameter (this proxy may previously have demanded authentication using the Proxy-Authenticate field).
3746 When multiple proxies are used in a chain, the Proxy-Authorization header field MUST NOT be consumed
3747 by any proxy whose realm does not match the “realm” parameter specified in the Proxy-Authorization
3748 header.

3749 Note that if an authentication scheme is used in the Proxy- Authorization that does not support realms,
3750 a proxy server MUST attempt to parse all Proxy-Authorization headers to determine whether or not one
3751 of them has what it considers to be valid credentials. Because this is potentially very time consuming in
3752 large networks, proxy servers SHOULD use an authentication scheme that supports realms in the Proxy-
3753 Authorization header.

3754 If a request is forked (as described in Section 16.6, various proxy servers and/or user agents may wish
3755 to challenge the UAC. In this case the forking proxy server is responsible for aggregating these challenges
3756 into a single response. Each WWW-Authenticate and Proxy-Authenticate received in responses to the
3757 forked request MUST be placed into the single response that is sent by the forking proxy to the user agent;
3758 the ordering of these headers is not significant.

3759 When a proxy server issues a challenge in response to a request, it will not proxy the request until the UAC has
3760 provided valid credentials. A forking proxy may forward a request simultaneously to multiple proxy servers that
3761 require authentication, each of which in turn will not forward the request until the originating UAC has authenticated
3762 itself in their respective realm. If the UAC does not provide credentials for each of these challenges, then the proxy
3763 servers that issued the challenges will not forward requests to user agents where the destination user might be
3764 located, and therefore, the virtues of forking are largely lost.

3765 If at least one UAS responds to a forked request with a challenge, than a 401 (Unauthorized) MUST be
3766 sent as the aggregated response by the forking proxy to the UAC; otherwise, if only proxy servers respond,
3767 a 407 MUST be used.

3768 When resubmitting its request in response to a 401 (Unauthorized) or 407 (Proxy Authentication Re-
3769 quired) that contains multiple challenges, a UAC MAY include an Authorization for each WWW-Authenticate
3770 and Proxy-Authorization for each Proxy-Authenticate for which the UAC wishes to supply a credential.
3771 As noted above, multiple credentials in a request SHOULD be differentiated by the “realm” parameter.

3772 It is possible for multiple challenges associated with the same realm to appear in the same 401 (Unautho-
3773 rized) or 407 (Proxy Authentication Required). This can occur, for example, when multiple proxies within
3774 the same administrative domain, which use a common realm, are reached by a forking request.

3775 See [H14.34] for a definition of the syntax of Proxy- Authentication and Proxy-Authorization.

3776 20.4 The Digest Authentication Scheme

3777 This section describes the modifications and clarifications required to apply the HTTP Digest authentication
3778 scheme to SIP. The SIP scheme usage is almost completely identical to that for HTTP [23]. Since RFC
3779 2543 is based on HTTP Digest as defined in RFC 2069 [24], SIP servers supporting RFC 2617 MUST ensure
3780 they are backwards compatible with RFC 2069. Procedures for this backwards compatibility are specified
3781 in RFC 2617. Note however that servers MUST NOT accept or request Basic authentication.

3782 **20.4.0.1 HTTP Digest** The rules for Digest authentication follow those defined in [23, Section 3], with
3783 “HTTP 1.1” replaced by “SIP/2.0” in addition to the following differences:

3784 1. The URI included in the challenge has the following BNF:

3785 URI = SIP-URI

3786 2. The BNF in RFC 2617 has an error in that the 'uri' parameter of the Authorization header for HTTP
3787 Digest authentication is not enclosed in quotation marks. (The example in Section 3.5 of RFC 2617
3788 is correct.) For SIP, the 'uri' MUST be enclosed in quotation marks.

3789 3. The BNF for digest-uri-value is:

3790 digest-uri-value = Request-URI ; as defined in Section 27

3791 4. The example procedure for choosing a nonce based on Etag does not work for SIP.

3792 5. The text in RFC 2617 [23] regarding cache operation does not apply to SIP.

3793 6. RFC 2617 [23] requires that a server check that the URI in the request line, and the URI included in
3794 the Authorization header, point to the same resource. In a SIP context, these two URI's may actually
3795 refer to different users, due to forwarding at some proxy. Therefore, in SIP, a server MAY check
3796 that the Request-URI in the Authorization header corresponds to a user for whom that the server is
3797 willing to accept forwarded or direct requests.

3798 7. As a clarification to the calculation of the A2 value for message integrity assurance in the Digest
3799 authentication scheme, implementers should assume, when the entity-body is empty (i.e. when SIP
3800 messages have no body) that the hash of the entity-body resolves to the MD5 hash of an empty string,
3801 or:

3802 H(entity-body) = MD5("") = "d41d8cd98f00b204e9800998ecf8427e"

3803 8. RFC 2617 notes that a nonce value MUST NOT be sent in an Authorization (and by extension Proxy-
3804 Authorization) header if no qop directive as been sent. Therefore, any algorithms that have a de-
3805 dependency on the nonce (including “MD5-Sess”) require that the qop directive be sent. Use of the
3806 “qop” parameter is optional in RFC 2617 for the purposes of backwards compatibility with RFC 2069;
3807 since RFC 2543 was based on RFC 2069, the “qop” parameter must unfortunately remain optional
3808 for clients and servers to receive. However, servers MUST always send a “qop” parameter in WWW-
3809 Authenticate and Proxy-Authenticate headers. If a client receives a “qop” parameter in a challenge
3810 header, it MUST send the “qop” parameter in any resulting authorization header.

3811 RFC 2543 did not allow usage of the Authentication-Info header (it effectively used RFC 2069). How-
3812 ever, we now allow usage of this header, since it provides integrity checks over the bodies and provides
3813 mutual authentication. RFC 2617 [23] defines mechanisms for backwards compatibility using the qop at-
3814 tribute in the request. These mechanisms MUST be used by a server to determine if the client supports the
3815 new mechanisms in RFC 2617 that were not specified in RFC 2069.

3816 21 S/MIME

3817 SIP messages carry MIME bodies and the MIME standard includes mechanisms for securing MIME con-
3818 tents to ensure both integrity and confidentiality (including the 'multipart/signed/' and 'application/pkcs7-
3819 mime' MIME types, see RFC 1847 [25], RFC 2630 [26] and RFC 2633 [27]). Implementers should note,
3820 however, that there may be rare network intermediaries (not typical proxy servers) that rely on viewing or
3821 modifying the bodies of SIP messages (especially SDP), and that secure MIME may prevent these sorts of
3822 intermediaries from functioning.

3823 This applies particularly to certain types of firewalls.

3824 Note that the PGP mechanism for encrypting the headers and bodies of SIP messages described in RFC 2543
3825 has been deprecated.

3826 21.1 S/MIME Certificates

3827 The certificates that are used to identify an end-user for the purposes of S/MIME differ from those used
3828 by servers in one important respect - rather than asserting that the identity of the holder corresponds to
3829 a particular hostname, these certificates assert that the holder is identified by an end-user address - this
3830 address is composed of the concatenation of the "userinfo" "@" and "domainname" portions of a SIP
3831 URI (in other words, an email address of the form "bob@biloxi.com"), most commonly corresponding to a
3832 user's address of record.

3833 These certificates are used to sign or encrypt bodies of SIP messages. Bodies are signed with the pri-
3834 vate key of the sender (who may include their public key with the message as appropriate), but bodies are
3835 encrypted with the public key of the intended recipient. Obviously, senders must have foreknowledge of the
3836 public key of recipients in order to encrypt message bodies. Public keys can be stored within a user agent
3837 on a virtual keyring.

3838 Each user agent that supports S/MIME MUST contain a keyring specifically for end-users certificates.
3839 This keyring should map between addresses of record and corresponding certificates, including any associ-
3840 ated with the owner or operator of the user agent, when appropriate. Over time, users SHOULD use the same
3841 certificate when they populate the originating URI of signaling (the From header) with the same address of
3842 record.

3843 Any mechanisms that depend on the existence of end-user certificates, however, have a serious limitation
3844 in that there is virtually no consolidated authority today that provides certificates for end-user applications.
3845 But if at all possible, users SHOULD acquire certificates from known public certificate authorities. As an al-
3846 ternative, users MAY create self-signed certificates. The implications of self-signed certificates are explored
3847 further in Section 22.4.2.

3848 Above and beyond the problem of acquiring an end-user certificate, there are few well-known central-
3849 ized directories that distribute end-user certificates. However, the holder of a certificate SHOULD publish
3850 their certificate in any public directories as appropriate. Similarly, UACs SHOULD support a mechanism
3851 for importing (manually or automatically) certificates discovered in public directories corresponding to the
3852 target URIs of SIP requests.

3853 21.2 S/MIME Key Exchange

3854 SIP itself can also be used as a means to distribute public keys in the following manner.

3855 Whenever the CMS SignedData message is used in S/MIME for SIP, it MUST contain the certificate
3856 bearing the public key necessary to verify the signature.

3857 When a UAC sends a request containing an S/MIME body that initiates a dialog, or sends a non-
3858 INVITE request outside the context of a dialog, the UAC SHOULD structure the body as an S/MIME 'mul-
3859 ticast/signed' CMS SignedData body; if the desired CMS service is EnvelopedData, the UAC should send
3860 the EnvelopedData message encapsulated within a SignedData message.

3861 When a UAS receives a request containing an S/MIME CMS body which includes a certificate, the UAS
3862 SHOULD first verify the certificate, if possible, with any available certificate authority. The UAS SHOULD
3863 also determine the subject of the certificate and compare this value to the From field of the request. If the
3864 certificate cannot be verified, because it is self-signed, or signed by no known authority, the UAS SHOULD
3865 notify the user of the status of the certificate (including the subject of the certificate, its signator, and any key
3866 fingerprint information) and request explicit permission before proceeding. If the certificate was successfully
3867 verified and the subject of the certificate corresponds to the From header field of the SIP request, or if the
3868 user (after notification) explicitly authorizes the use of the certificate, the UAS SHOULD add this certificate
3869 to a local keyring, indexed by the address of record of the holder of the certificate.

3870 When a UAS sends a response containing an S/MIME body that answers the first request in a dialog, or
3871 a response to a non-INVITE request outside the context of a dialog, the UAS SHOULD structure the body
3872 as a S/MIME 'multipart/signed' CMS SignedData body; if the desired CMS service is EnvelopedData, the
3873 UAS SHOULD send the EnvelopedData message encapsulated within a SignedData message. If the S/MIME
3874 body received by the UAS was encrypted with a public key recognized by the UAS, it MAY opt not to sign
3875 its response when appropriate.

3876 When a UAC receives a response containing an S/MIME CMS body which includes a certificate, the
3877 UAC SHOULD first verify the certificate, if possible, with any available certificate authority. The UAC
3878 SHOULD also determine the subject of the certificate and compare this value to the To field of the response;
3879 although the two may very well be different, and this is not necessarily indicative of a security breach.
3880 If the certificate cannot be verified, because it is self-signed, or signed by no known authority, the UAC
3881 SHOULD notify the user of the status of the certificate (including the subject of the certificate, its signator,
3882 and any key fingerprint information) and request explicit permission before proceeding. If the certificate was
3883 successfully verified and the subject of the certificate corresponds to the To header in the response, or if the
3884 user (after notification) explicitly authorizes the use of the certificate, the UAC SHOULD add this certificate
3885 to a local keyring, indexed by the address of record of the holder of the certificate. If the UAC had not
3886 transmitted its own certificate to the UAS in any previous transaction, it SHOULD use a CMS SignedData
3887 body for its next request or response.

3888 On future occasions, when the UA receives requests or responses that contain a From header field
3889 corresponding to a value in its keyring, the UA SHOULD compare the certificate offered in these messages
3890 with the existing certificate in its keyring. If there is a discrepancy, the UA SHOULD notify the user of a
3891 change of the certificate (preferably in terms that indicate that this is a potential security breach) and acquire
3892 the user's permission before continuing to process the signaling. If the user authorizes this certificate, it
3893 MUST be added to the keyring alongside any previous value(s) for this address of record.

3894 Note well however, that this key exchange mechanism does not guarantee the secure exchange of keys
3895 when self-signed certificates, or certificates signed by an obscure authority, are used - it is vulnerable to
3896 well-known attacks. In the opinion of the authors, however, the security it provides is proverbially better
3897 than nothing; it is in fact comparable to the widely used SSH application. These limitations are explored in
3898 greater detail in Section 22.4.2.

3899 If a user agent receives an S/MIME body that has been encrypted with a public key unknown to the
3900 recipient, it MUST reject the request with a 493 (Undecipherable) response. This response SHOULD contain
3901 a valid certificate for the respondent (corresponding, if possible, to any address of record given in the To

3902 header of the rejected request) within a MIME body. A 493 (Undecipherable) sent without any certificate
3903 indicates that the respondent cannot or will not utilize S/MIME.

3904 Finally, if during the course of a dialog a user agent receives a certificate in a CMS SignedData message
3905 that does not correspond with the certificates previously exchanged during a dialog, the user agent MUST
3906 notify its user of the change, preferably in terms that indicate that this is a potential security breach.

3907 **21.3 Securing MIME bodies**

3908 There are two types of secure MIME bodies that are of interest to SIP: 'multipart/signed' and 'application/pkcs7-
3909 mime'. The procedures for the use of these bodies should follow the S/MIME specification ([27]) with a
3910 few variations.

- 3911 • 'multipart/signed' MUST be used only with CMS detached signatures.

3912 This allows backwards compatibility with non-S/MIME-compliant recipients.

- 3913 • If a UAC has no certificate on its keyring associated with the address of record to which it wants to
3914 send a request, it cannot send an encrypted 'application/pkcs7-mime' MIME message. UACs MAY
3915 send an initial request such as an OPTIONS message with a CMS detached signature in order to
3916 solicit the certificate of the remote side (the signature SHOULD be over a 'message/sip' body of the
3917 type described in Section 21.4).

- 3918 • Senders of S/MIME bodies SHOULD use the 'SMIMECapabilities' (see Section 2.5.2 of [27]) attribute
3919 to express their capabilities and preferences for further communications. Note especially that senders
3920 MAY use the 'preferSignedData' capability to encourage receivers to respond with CMS SignedData
3921 messages (for example, when sending an OPTIONS request as described above).

- 3922 • S/MIME implementations MUST at a minimum support SHA1 as a digital signature algorithm, and
3923 3DES as an encryption algorithm; all other signature and encryption algorithms MAY be supported.
3924 Implementations can negotiate support for these algorithms with the 'SMIMECapabilities' attribute.

3925 **21.4 Tunneling SIP in MIME**

3926 As a means of providing some degree of end-to-end authentication, integrity or confidentiality for SIP head-
3927 ers, S/MIME can encapsulate entire SIP messages within MIME bodies of type "message/sip" and then
3928 apply MIME security to these bodies in the same manner employed for typical SIP bodies.

3929 Note that these "message/sip" bodies can be sent as a part of a MIME "multipart/mixed" body if another
3930 MIME types (such as SDP) should also be used in the request.

3931 **21.4.1 Tunneling Integrity and Authentication**

3932 Tunneling SIP messages within S/MIME bodies can provide integrity for SIP headers if the headers which
3933 the sender wishes to secure are replicated in a "message/sip" MIME body signed with a CMS detached
3934 signature.

3935 Provided that the "message/sip" body contains at least the fundamental dialog identifiers (To, From,
3936 Call-ID, CSeq), then a signed MIME body can provide limited authentication. At the very least, if the
3937 certificate used to sign the body is unknown to the recipient and cannot be verified, the signature can be used

3938 to ascertain that a later request in a dialog was transmitted by the same certificate-holder that initiated the
3939 dialog. If the recipient of the signed MIME body has some stronger incentive to trust the certificate (they
3940 were able to verify it, acquire it from a trusted repository, or they've used it frequently) then the signature
3941 can be taken as a stronger assertion of the identity of the subject of the certificate.

3942 In order to eliminate possible confusions about the addition or subtraction of entire headers, senders
3943 SHOULD replicate all headers from the request within the signed body. Any message bodies that require
3944 integrity protection SHOULD be attached to the "inner" message.

3945 Upon receipt of a SIP message with a signed "message/sip" body, recipients may compare headers in
3946 the "outer" message with headers in the "inner" message. At the discretion of the recipient, if significant
3947 discrepancies between the two exist, the message MAY be rejected with a 403 (Forbidden) response if it
3948 is a request, or any existing dialog MAY be terminated if a security violation has occurred. User agents
3949 SHOULD notify users of this circumstance and request explicit guidance on how to proceed. Provided that
3950 the signature is valid for the "inner" message, headers in the inner message SHOULD be preferred to headers
3951 in the "outer" message.

3952 Many SIP headers are altered of necessity as messages are routed through proxy servers. These include,
3953 but are not necessarily limited to, the Request-URI, Via headers, Record-Route and Route headers, the
3954 Max-Forwards header, and the Proxy-Authorization header; note that extensions to SIP, or nonstandard
3955 (X-) headers, may also result in headers that are added or subtracted from messages as they traverse the
3956 network. A variation in these headers SHOULD NOT be interpreted as a breach of integrity by the recipient
3957 of a signed message.

3958 The following is an example of the use of a tunneled "message/sip" body:

```
3959     INVITE sip:bob@biloxi.com SIP/2.0
3960     Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
3961     To: Bob <bob@biloxi.com>
3962     From: Alice <alice@atlanta.com>;tag=1928301774
3963     Call-ID: a84b4c76e66710
3964     CSeq: 314159 INVITE
3965     Contact: <sip:alice@pc33.atlanta.com>
3966     Content-Type: multipart/signed;
3967         protocol="application/pkcs7-signature";
3968         micalg=sha1; boundary=boundary42
3969
3970     --boundary42
3971     Content-Type: message/sip
3972
3973     INVITE sip:bob@biloxi.com SIP/2.0
3974     Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
3975     To: Bob <bob@biloxi.com>
3976     From: Alice <alice@atlanta.com>;tag=1928301774
3977     Call-ID: a84b4c76e66710
3978     CSeq: 314159 INVITE
3979     Contact: <sip:alice@pc33.atlanta.com>
3980     Content-Type: application/sdp
3981     Content-Length: 147
```

```
3982
3983     v=0
3984     o=UserA 2890844526 2890844526 IN IP4 here.com
3985     s=Session SDP
3986     c=IN IP4 pc33.atlanta.com
3987     t=0 0
3988     m=audio 49172 RTP/AVP 0
3989     a=rtpmap:0 PCMU/8000
3990
3991     --boundary42
3992     Content-Type: application/pkcs7-signature; name=smime.p7s
3993     Content-Transfer-Encoding: base64
3994     Content-Disposition: attachment; filename=smime.p7s
3995
3996     ghyHhHUujhJhjH77n8HHGTrfvbnj756tbB9HG4VQpfyF467GhIGfHfYT6
3997     4VQpfyF467GhIGfHfYT6jH77n8HHGghyHhHUujhJh756tbB9HGTrfvbnj
3998     n8HHGTrfvhJhjH776tbB9HG4VQbnj7567GhIGfHfYT6ghyHhHUujpfyF4
3999     7GhIGfHfYT64VQbnj756
4000
4001     --boundary42-
```

4002 21.4.2 Tunneling Encryption

4003 It may also be desirable to use this mechanism to encrypt a “message/sip” MIME body within a CMS
4004 EnvelopedData message S/MIME body, but in practice, most headers are of at least some use to the network;
4005 the general use of encryption with S/MIME is to secure message bodies like SDP rather than message
4006 headers. Some informational headers, such as the **Subject** or **Organization** could perhaps warrant end-to-
4007 end security. Headers defined by future SIP applications might also require obfuscation.

4008 Another possible application of encrypting headers is selective anonymity. A request could be con-
4009 structed with a **From** header field that contains no personal information (e.g., sip:anonymous@anonymizer.com).
4010 However, a second **From** header field containing the genuine address of record of the originator could be
4011 encrypted within a “message/sip” MIME body where it will only be visible to the endpoints of a dialog.

4012 In the following example, the text boxed in asterisks (“*”) is encrypted:

```
4013     INVITE sip:bob@biloxi.com SIP/2.0
4014     Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
4015     To: Bob <bob@biloxi.com>
4016     From: Alice <alice@atlanta.com>;tag=1928301774
4017     Call-ID: a84b4c76e66710
4018     CSeq: 314159 INVITE
4019     Contact: <sip:alice@pc33.atlanta.com>
4020     Content-Type: application/pkcs7-mime; smime-type=enveloped-data;
4021     name=smime.p7m
4022     Content-Transfer-Encoding: base64
4023     Content-Disposition: attachment; filename=smime.p7m
```

```

4024
4025 *****
4026 * Content-Type: application/sdp *
4027 * *
4028 * v=0 *
4029 * o=alice 53655765 2353687637 IN IP4 pc33.atlanta.com *
4030 * s=- *
4031 * t=0 0 *
4032 * c=IN IP4 pc33.atlanta.com *
4033 * m=audio 3456 RTP/AVP 0 1 3 99 *
4034 * a=rtpmap:0 PCMU/8000 *
4035 *****
4036

```

4037 22 Security Considerations

4038 SIP is not an easy protocol to secure. Its use of intermediaries, its multi-faceted trust relationships, its
4039 expected usage between elements with no trust at all, and its user-to-user operation make security far from
4040 trivial. Security solutions are needed that are deployable today, without extensive coordination, in a wide
4041 variety of environments and usages. In order to meet these diverse needs, several distinct mechanisms
4042 applicable to different aspects and usages of SIP will be required.

4043 Note that the security of SIP signaling itself has no bearing on the security of protocols used in concert
4044 with SIP such as RTP, or with the security implications of any specific bodies SIP might carry (although
4045 MIME security plays a substantial role in securing SIP). Any media associated with a session can be en-
4046 crypted end-to-end independently of any associated SIP signaling. Media encryption is outside the scope of
4047 this document.

4048 The considerations that follow first examine a set of classic threat models which broadly identify the
4049 security needs of the SIP protocol. The set of security services required to address these threats is then
4050 detailed, followed by an explanation of several security mechanisms that can be used to provide these ser-
4051 vices. Next, the requirements for implementers of SIP are enumerated, along with exemplary deployments
4052 in which these security mechanisms could be used to improve the security of SIP. Some notes on privacy
4053 conclude this section.

4054 22.1 Threat Models

4055 This section details some threats that should be common to most deployments of SIP. These threats have
4056 been chosen specifically to illustrate each of the security services that SIP requires.

4057 The following examples by no means provide an exhaustive list of the threats against the SIP proto-
4058 col; rather, these are "classic" threats that demonstrate the need for particular security services which can
4059 potentially prevent whole categories of threats.

4060 **22.1.1 Registration Hijacking**

4061 The SIP registration mechanism allows a user agent to identify itself to a registrar as a device at which a user
4062 (designated by an address of record) is located. A registrar assesses the identity asserted in the From header
4063 field of a REGISTER message to determine whether or not this request can modify the contact addresses
4064 associated with the address of record in the To header field; while these two fields are frequently the same,
4065 there are many valid deployments in which a third-party may register contacts on a user's behalf.

4066 The From header of a SIP request, however, can essentially be modified arbitrarily by the owner of a
4067 user agent, and this opens the door to malicious registrations. An attacker that successfully impersonates a
4068 party authorized to change contacts associated with an address of record could, for example, de-register all
4069 existing contacts for a URI and then register their own device as the appropriate contact address, thereby
4070 directing all requests for the affected user to the attacker's device.

4071 This threat belongs to a family of threats that rely on the absence of cryptographic assurance of a re-
4072 quest's originator. Any SIP UAS that represents a valuable service (a gateway that interworks SIP requests
4073 with traditional telephone calls, for example) might want to control access to its resources by authenticating
4074 requests that it receives. Even end-user UAs, for example SIP phones, have an interest in ascertaining the
4075 identities of originators of requests.

4076 This threat demonstrates the need for security services that enable SIP entities to authenticate the origi-
4077 nators of requests.

4078 **22.1.2 Impersonating a Server**

4079 The domain to which a request is destined is generally specified in the Request-URI; user agents commonly
4080 contact a server in this domain directly in order to deliver a request. However, there is always a possibility
4081 that an attacker could impersonate the remote server, and that the user agent's request could be intercepted
4082 by some other party.

4083 For example, consider a case in which a redirect server at one domain, chicago.com, impersonates a
4084 redirect server at another domain, biloxi.com. A user agent sends a request to biloxi.com, but the redirect
4085 server at chicago.com answers with a forged response that has appropriate SIP headers for a response from
4086 biloxi.com. The forged contact addresses in the redirection response could direct the originating user agent
4087 to inappropriate or insecure resources, or simply prevent requests for biloxi.com from succeeding.

4088 This family of threats has a vast membership, many of which are critical. As a converse to the registration
4089 hijacking threat, consider the case in which a registration sent to biloxi.com is intercepted by chicago.com,
4090 which replies to the intercepted registration with a forged 301 (Moved Permanently) response. This response
4091 might seem to come from biloxi.com yet designate chicago.com as the appropriate registrar. All future
4092 REGISTER requests from the originating user agent would then go to chicago.com.

4093 Prevention of this threat requires a means by which user agents can authenticate the servers to whom
4094 they send requests.

4095 **22.1.3 Tampering with Message Bodies**

4096 As a matter of course, SIP user agents route requests through trusted proxy servers. Regardless of how that
4097 trust is established (authentication of proxies is discussed elsewhere in this section), a user agent may trust
4098 a proxy server to route a request, but not to inspect or possibly modify the bodies contained in that request.

4099 Consider a UA that is using SIP message bodies to communicate session encryption keys for a media
4100 session. Although it trusts the proxy server of the domain it is contacting to deliver signaling properly, it

4101 may not be desirable for the administrators of that domain to be capable of decrypting any subsequent media
4102 session. Worse yet, if the proxy server were actively malicious, it could modify the session key, either acting
4103 as a man-in-the-middle, or perhaps changing the security characteristics requested by the originating user
4104 agent.

4105 This family of threats applies not only to session keys, but to most conceivable forms of content car-
4106 ried end-to-end in SIP. These might include MIME bodies that should be rendered to the user, SDP, or
4107 encapsulated telephony signals among others.

4108 Also note that some headers in SIP are meaningful end-to-end, for example, the **Subject**. User agents
4109 might be protective of these headers as well as bodies (a malicious intermediary changing the **Subject**
4110 header might make an important request appear to be spam, for example). However, since many headers are
4111 legitimately inspected or altered by proxy servers as a request is routed, not all headers should be secured
4112 end-to-end.

4113 For these reasons, the UA might want to secure SIP message bodies, and in some limited cases headers,
4114 end-to-end. The security services required for bodies include confidentiality, integrity, and authentication.
4115 These end-to-end services should be independent of the means used to secure interactions with intermedi-
4116 aries such as proxy servers.

4117 **22.1.4 Tearing Down Sessions**

4118 Once a dialog has been established by initial messaging, subsequent requests can be sent that modify the
4119 state of the dialog and/or session. It is critical that principals in a session can be certain that such requests
4120 are not forged by attackers.

4121 Consider a case in which a third-party attacker captures some initial messages in a dialog shared by
4122 two parties in order to learn the parameters of the session (**To**, **From**, and so forth) and then inserts a **BYE**
4123 request into the session. The attacker could opt to forge the request such that it seemed to come from either
4124 participant. Once the **BYE** is received by its target, the session will be torn down prematurely.

4125 Similar mid-session threats include the transmission of forged re-**INVITE**s that alter the session (possibly
4126 to reduce session security or redirect media streams as part of a wiretapping attack).

4127 The most effective countermeasure to this threat is the authentication of the sender of the **BYE** - in
4128 this instance, the recipient needs only know that the **BYE** came from the same party with whom the corre-
4129 sponding dialog was established (as opposed to ascertaining the absolute identity of the sender). Also, if the
4130 attacker is unable to learn the parameters of the session due to confidentiality, it would not be possible to
4131 forge the **BYE**; however, some intermediaries (like proxy servers) will need to inspect those parameters as
4132 the session is established.

4133 **22.1.5 Denial of Service and Amplification**

4134 Denial of service attacks focus on rendering a particular network element unavailable, usually by directing
4135 an excessive amount of network traffic at its interfaces. A distributed denial of service attack allows one
4136 network user to cause multiple network hosts to flood a target host with a large amount of network traffic.

4137 In many architectures SIP proxy servers face the public Internet in order to accept requests from world-
4138 wide IP endpoints. SIP creates a number of potential opportunities for distributed denial of service attacks
4139 that must be recognized and addressed by the implementers and operators of SIP systems.

4140 Attackers can create bogus requests that contain a falsified source IP address and a corresponding **Via**
4141 header field which identify a targeted host as the originator of the request and then send this request to a large

4142 number of SIP network elements, thereby using hapless SIP UAs or proxies to generate denial of service
4143 traffic aimed at the target.

4144 Similarly, attackers might use falsified **Route** headers in a request that identify the target host and then
4145 send such messages to forking proxies that will amplify messaging sent to the target. **Record-Route** could
4146 be used to similar effect when the attacker is certain that the SIP dialog initiated by the request will result in
4147 numerous transactions originating in the backwards direction.

4148 A number of denial of service attacks open up if **REGISTER** requests are not properly authenticated
4149 and authorized by registrars. Attackers could de-register some or all users in an administrative domain,
4150 thereby preventing these users from being invited to new sessions. An attacker could also register a large
4151 number of contacts designating the same host for a given address of record in order to use the registrar and
4152 any associated proxy servers as amplifiers in a denial of service attack. Attackers might also attempt to
4153 deplete available memory and disk resources of a registrar by registering huge numbers of bindings.

4154 The use of multicast to transmit SIP requests can greatly increase the potential for denial of service
4155 attacks.

4156 These problems demonstrate a general need to define architectures that minimize the risks of denial of
4157 service, and the need to be mindful in recommendations for security mechanisms of this class of attacks.

4158 **22.2 Security Mechanisms**

4159 From the threats described above, we gather that the fundamental security services required for the SIP
4160 protocol are: preserving the confidentiality and integrity of messaging, preventing replay attacks or message
4161 spoofing, providing for the authentication and privacy of the participants in a session, and preventing denial
4162 of service attacks. Bodies within SIP messages separately require the security services of: confidentiality,
4163 integrity, and authentication.

4164 Rather than defining new security mechanisms specific to SIP, SIP reuses wherever possible existing
4165 security models derived from the HTTP and SMTP space.

4166 Full encryption of messages provides the best means to preserve the confidentiality of signaling - it can
4167 also guarantee that messages are not modified by any malicious intermediaries. However, SIP requests and
4168 responses cannot be naively encrypted end-to-end in their entirety because, in most network architectures,
4169 message fields such as the **Request-URI**, **Route** and **Via** need to be visible to proxies so that SIP requests
4170 are routed correctly. Note that proxy servers need to modify some features of messages as well (such as
4171 adding **Via** headers) in order for SIP to function. Proxy servers must therefore be trusted, to some degree,
4172 by SIP user agents. To this purpose, low layer security mechanisms for SIP are recommended, which encrypt
4173 the entire SIP requests or responses on the wire on a hop-by-hop basis, and which allow endpoints to verify
4174 the identity of proxy servers to whom they send requests.

4175 SIP entities also have a need to identify one another in a secure fashion. When a SIP endpoint asserts the
4176 identity of its user to a peer user agent or to a proxy server, that identity should in some way be verifiable.
4177 A cryptographic authentication mechanism is provided in SIP to address this requirement.

4178 An independent security mechanism for SIP message bodies supplies an alternative means of end-to-end
4179 mutual authentication, as well as providing a limit on the degree to which user agents must trust intermedi-
4180 aries.

4181 **22.2.1 Transport and Network Layer Security**

4182 Transport or network layer security encrypts signaling traffic, guaranteeing message confidentiality and in-
4183 tegrity (note however that the originator and recipient of a session may be deducible by observers performing

4184 a network traffic analysis). The certificates used to encrypt traffic can also be used to provide a means of
4185 authentication in many architectures.

4186 Two popular alternatives for providing security at the transport and network layer are, respectively, TLS
4187 [28] and IPSec [29].

4188 IPSec is a set of network-layer protocol tools that can collectively be used as a secure replacement for
4189 traditional IP (Internet Protocol). IPSec is most suited to architectures in which a set of SIP hosts (mingled
4190 user agents and proxy servers) or bridged administrative domains (possibly using security gateways) have
4191 an existing trust relationship with one another, although IPSec can also be used on a per-hop basis.

4192 IPSec is generally implemented at an operating-system level within a host, and in many architectures
4193 it does not require integration with SIP applications. Any deployment of IPSec for SIP would require an
4194 IPSec profile describing the protocols tools that would be required to secure SIP and the modes in which
4195 they would operate. No such profile is given in this document.

4196 TLS provides transport-layer security over connection-oriented protocols (for the purposes of this doc-
4197 ument, TCP); "tls" (signifying TLS over TCP) can be specified as the desired transport protocol within a
4198 Via header field or a SIP-URI. TLS is most suited to architectures in which a chain of trust joins together a
4199 set of hosts. For example, Alice trusts her local proxy server, which in turn trust Bob's local proxy server,
4200 which Bob trusts, hence Bob and Alice can communicate securely.

4201 TLS must be tightly coupled with a SIP application. Note that transport mechanisms are specified on a
4202 hop-by-hop basis in SIP, and that in some deployments TLS might be used for only certain portions of the
4203 signaling path.

4204 When TLS is used in a SIP application, implementers MUST minimally support the TLS_RSA_WITH_AES_128_CBC_SHA
4205 ciphersuite. For purposes of backwards compatibility, proxy servers, redirect servers and registrars SHOULD
4206 support TLS_RSA_WITH_3DES_EDE_CBC_SHA. Implementers MAY also support any other ciphersuite.

4207 22.2.2 HTTP Authentication

4208 SIP provides a challenge capability, based on HTTP authentication, that relies on the 401 and 407 response
4209 codes as well as headers for carrying challenges and credentials. Without significant modification, the reuse
4210 of the HTTP Digest authentication scheme in SIP allows for replay protection and one-way authentication.

4211 The usage of Digest authentication in SIP is detailed in Section 20.

4212 22.2.3 S/MIME

4213 As is discussed above, encrypting entire SIP messages end-to-end for the purpose of confidentiality is not ap-
4214 propriate because network intermediaries (like proxy servers) need to view certain headers in order to route
4215 messages correctly, and if these intermediaries are excluded from security associations then SIP messages
4216 will essentially be unroutable.

4217 However, S/MIME allows SIP user agents to encrypt MIME bodies within SIP, securing these bodies
4218 end-to-end without affecting message headers. S/MIME can provide end-to-end confidentiality and integrity
4219 for message bodies, as well as mutual authentication. It is also possible to use S/MIME to provide a form of
4220 integrity and confidentiality for SIP headers through SIP message tunneling.

4221 The usage of S/MIME in SIP is detailed in Section 21.

4222 **22.3 Implementing Security Mechanisms**

4223 **22.3.1 Requirements for Implementers of SIP**

4224 Proxy servers, redirect servers, and registrars **MUST** implement TLS, and **MUST** support both mutual and
4225 one-way authentication. It is strongly **RECOMMENDED** that user agents be capable initiating TLS; user
4226 agents **MAY** also be capable of acting as a TLS server. Proxy servers, redirect servers, and registrars **SHOULD**
4227 possess a site certificate whose subject corresponds to their hostname. User agents **MAY** have certificates of
4228 their own for mutual authentication with TLS, but no provisions are set forth in this document for their use.
4229 User agents **MUST** support a mechanism for verifying certificates they receive during TLS negotiation.

4230 Proxy servers, redirect servers, registrars and user agents **MAY** also implement IPSec, or other lower-
4231 layer security protocols.

4232 When a user agent attempts to contact a proxy server, redirect server or registrar, the UAC **SHOULD**
4233 initiate a TLS connection over which it will send SIP messages. In some architectures UACs **MAY** receive
4234 requests over such TLS connections as well.

4235 Proxy servers, redirect servers, registrars and user agents **MUST** implement Digest Authorization. Proxy
4236 servers, redirect servers and registrars **SHOULD** be configured with at least one Digest realm, and at least one
4237 "realm" string supported by a given server **SHOULD** correspond to the server's hostname or domainname.

4238 Proxy servers, redirect servers, registrars and user agents **MAY** also implement enhancements to Digest
4239 or alternate header-level security mechanisms.

4240 User agents **SHOULD** support S/MIME encryption and signing of SIP message MIME bodies.

4241 **22.3.2 Security Solutions**

4242 The operation of these security mechanisms in concert can follow, to some degree, the existing web and
4243 email security models. At a high level, user agents authenticate themselves to servers (proxy servers, redirect
4244 servers and registrars) with a Digest username and password; servers authenticate themselves to user agents,
4245 and to one another, with a site certificate delivered by TLS.

4246 On a peer-to-peer level, user agents ordinarily transitively trust the network to authenticate one another;
4247 however, S/MIME can also be used to provide direct authentication when the network does not or if the
4248 network itself is not trusted.

4249 The following is an illustrative example in which these security mechanisms are used by various user
4250 agents and servers to prevent the sorts of threats described in Section 22.1. While implementers and network
4251 administrators **MAY** follow the normative guidelines given in the remainder of this section, these are provided
4252 only as example implementations.

4253 **22.3.2.1 Registration** When a user agent comes on line and registers with its local administrative do-
4254 main, it **SHOULD** establish a TLS connection with its registrar (the means by which the user agent determines
4255 how to reach its registrar are described in Section 10). The registrar **SHOULD** offer a certificate to the user
4256 agent, and the site identified by the certificate **MUST** correspond with the domain in which the user agent in-
4257 tends to register; for example, if the user agent intends to register the address of record 'alice@atlanta.com',
4258 the site certificate must identify a host within the atlanta.com domain (such as 'sip.atlanta.com'). When
4259 it receives the TLS Certificate message, the user agent **SHOULD** verify the certificate and inspect the site
4260 identified by the certificate. If the certificate is invalid, revoked, or if it does not identify the appropriate
4261 party, the user agent **MUSTNOT** send the REGISTER message and otherwise proceed with the registration.

4262 When a valid certificate has been provided by the registrar, the user agent knows that the registrar is not an
4263 attacker who might redirect the user agent, steal passwords, or attempt any similar attacks.

4264 The user agent then creates a REGISTER request which SHOULD be addressed to a Request-URI
4265 corresponding to the site certificate received from the registrar. When the REGISTER request is sent by the
4266 user agent over the existing TLS connection, the registrar SHOULD challenge the request with a 407 (Proxy
4267 Authentication Required) response; the “realm” parameter within the Proxy-Authenticate header of the
4268 response SHOULD correspond to the domain previously given by the site certificate. When the UAC receives
4269 the challenge, it SHOULD either prompt the user for credentials or take an appropriate credential from a
4270 keyring corresponding to the “realm” parameter in the challenge. The username of this credential SHOULD
4271 correspond with the “userinfo” portion of the URI in the To header of the REGISTER request. Once the
4272 Digest credentials have been inserted into an appropriate Proxy-Authorization header, the REGISTER
4273 should be resubmitted to the registrar.

4274 Since the registrar requires the user agent to authenticate itself, it would be difficult for an attacker to forge
4275 REGISTER requests for the user’s address of record. Also note that since the REGISTER is sent over a confidential
4276 TLS connection, attackers will not be able to intercept the REGISTER to record credentials for any possible replay
4277 attack.

4278 Once the registration has been accepted by the registrar, the user agent SHOULD leave this TLS connec-
4279 tion open provided that the registrar also acts as the proxy server to which requests are sent for users in this
4280 administrative domain. The existing TLS connection will be reused to deliver incoming requests to the user
4281 agent that has just completed registration.

4282 Because the user agent has already authenticated the server on the other side of the TLS connection, all requests
4283 that come over this connection are known to have passed through the proxy server - attackers cannot create spoofed
4284 requests that appear to have been sent through that proxy server.

4285 **22.3.2.2 Requests and Transitive Trust** Now let’s say that Alice’s user agent would like to initiate a
4286 session with a user in a remote administrative domain, namely ‘bob@biloxi.com’. We’ll also say that the
4287 local administrative domain (‘atlanta.com’) has a local outbound proxy.

4288 The proxy server that handles inbound requests for an administrative domain MAY also act as a local
4289 outbound proxy; for simplicity’s sake we’ll assume this to be the case for ‘atlanta.com’ (otherwise the user
4290 agent would initiate a new TLS connection to a separate server at this point). Assuming that the client has
4291 completed the registration process described in the preceding section, it SHOULD reuse the TLS connection
4292 to the local proxy server when it wishes to send an INVITE request to another user. The user agent SHOULD
4293 reuse cached credentials in the INVITE to avoid prompting the user unnecessarily.

4294 When the local outbound proxy server has validated the credentials presented by the user agent in the
4295 INVITE, it SHOULD inspect the Request-URI to determine how the message should be routed (see [8]).
4296 If the “domainname” portion of the Request-URI had corresponded to the local domain (‘atlanta.com’),
4297 rather the “biloxi.com”, then the proxy server would have consulted its location service to determine how
4298 best to reach the requested user.

4299 Had ‘alice@atlanta.com’ been attempting to contact, say, ‘alex@atlanta.com’, the local proxy would have prox-
4300 ied to the request to the TLS connection Alex had established with the register when he registered. Since Alex would
4301 receive this request over his authenticated channel, he would be assured that Alice’s request had been authorized by
4302 the proxy server of the local administrative domain.

4303 However, in this instance the Request-URI designates a remote domain. The local outbound proxy
4304 server at ‘atlanta.com’ SHOULD therefore establish a TLS connection with the remote proxy server at

4305 'biloxi.com'. Since both of the participants in this TLS connection are servers that possess site certifi-
4306 cates, mutual TLS authentication SHOULD occur. Each side of the connection SHOULD verify and inspect
4307 the certificate of the other, noting the domain name that appears in the certificate for comparison with the
4308 headers of SIP messages. The 'atlanta.com' proxy server, for example, SHOULD verify at this stage that the
4309 certificate received from the remote side corresponds with the 'biloxi.com' domain. Once it has done so,
4310 and TLS negotiation has completed, resulting in a secure channel between the two proxies, the 'atlanta.com'
4311 proxy can forward the INVITE request to 'biloxi.com'.

4312 The proxy server at 'biloxi.com' SHOULD in turn inspect the certificate of the proxy server at 'at-
4313 lanta.com' and compare the domain asserted by the certificate with the "domainname" portion of the From
4314 header in the INVITE request. The biloxi proxy can thereby ascertain whether or not it should consider Alice
4315 to be transitively authenticated. The biloxi proxy MAY have a strict security policy that requires it to reject
4316 requests that do not match the administrative domain from which they have been proxied, or perhaps even
4317 more strictly, requests that originate from administrative domains that do not have some policy agreement
4318 with biloxi.

4319 Such security policies could be instituted to prevent the SIP equivalent of SMTP 'open relays' which are fre-
4320 quently exploited to generate spam.

4321 Once the INVITE has been approved by the biloxi proxy, the proxy server SHOULD identify the existing
4322 TLS channel, if any, associated with the user targeted by this request (in this case 'bob@biloxi.com'). The
4323 INVITE should be proxied through this channel to Bob; since the request is received over a TLS connection
4324 which had previously been authenticated as the biloxi proxy, Bob transitively trusts the identity asserted in
4325 the From header.

4326 Before they forward the request, both proxy servers SHOULD add Record-Route headers to the request
4327 so that all future requests in this dialog will pass through the proxy servers. The proxy servers can thereby
4328 continue to provide transitive authentication, confidentiality, replay protection, and so forth for lifetime of
4329 this dialog. If the proxy servers do not add themselves to the Record-Route, future messages will pass
4330 directly end-to-end between Alice and Bob without any security services (unless the two parties agree on
4331 some independent end-to-end security).

4332 An attacker preying on this architecture would, for example, be unable to forge a BYE request and insert it into
4333 the signaling stream between Bob and Alice because the attacker has no way of ascertaining the parameters of the
4334 session because of the use of confidentiality, and moreover because the integrity mechanism transitively protects the
4335 traffic all the way from Alice to Bob.

4336 **22.3.2.3 Peer to Peer Requests** Alternatively, consider a user agent asserting the identity 'carol@chicago.com'
4337 that has no local outbound proxy. When Carol wishes to send an INVITE to 'bob@biloxi.com', her user
4338 agent SHOULD initiate a TLS connection with the biloxi proxy directly (using the mechanism described in
4339 [8] to determine how to best to reach the given Request-URI). When her user agent receives a certificate
4340 from the biloxi proxy, it SHOULD be verified normally before she passes her INVITE across the TLS con-
4341 nection. However, 'carol@chicago.com' has no means of proving her identity to the biloxi proxy; but she
4342 does have a CMS detached signature over a "message/sip" body in the INVITE. It is unlikely in this instance
4343 that Carol would have any credentials in the 'biloxi.com' realm, since she has no formal association with
4344 biloxi.com. The biloxi proxy MAY also have a strict policy that precludes it from even bothering to challenge
4345 requests that do not have 'biloxi.com' in the "domainname" portion of the From header - it treats these users
4346 as unauthenticated.

4347 The biloxi proxy has a policy for Bob that all non-authenticated requests should be redirected to the
4348 appropriate contact address registered against 'bob@biloxi.com', namely ;sip:bob@192.0.2.4;. Carol re-

4349 ceives the redirection response over the TLS connection she established with the biloxi proxy, so she trusts
4350 the veracity of the contact address.

4351 Carol SHOULD then established a TCP connection with the designated address and send a new INVITE
4352 with a Request-URI containing the received contact address (recomputing the signature in the body as
4353 the request is readied). Bob receives this INVITE on an insecure interface, but his user agent inspects
4354 and in this instance recognizes the From header of the request and subsequently matches a locally cached
4355 certificate with the one presented in the signature of the body of the INVITE. He replies in similar fashion,
4356 authenticating himself to Carol, and a secure dialog begins.

4357 Sometimes firewalls or NATs in an administrative domain could preclude the establishment of a direct TCP
4358 connection to a user agent. In these cases, proxy servers could also potentially relay requests to user agents in a way
4359 that has no trust implications (for example, forgoing an existing TLS connection and forwarding the request over
4360 cleartext TCP) as local policy dictates.

4361 **22.3.2.4 DoS Protection** In order to minimize the risk of a denial of service attack against architectures
4362 using these security solutions, implementers should take note of the following guidelines.

4363 When the host on which a SIP proxy server is operating is routable from the public Internet, it SHOULD
4364 be deployed in an administrative domain with secure routing policies (blocking source-routed traffic, prefer-
4365 ably filtering ping traffic). Both TLS and IPsec can also make use of bastion hosts at the edges of ad-
4366 ministrative domains that participate in the security associations to aggregate secure tunnels and sockets.
4367 These bastion hosts can also take the brunt of denial of service attacks, ensuring that SIP hosts within the
4368 administrative domain are not encumbered with superfluous messaging.

4369 No matter what security solutions are deployed, floods of messages directed at proxy servers can lock up
4370 proxy server resources and prevent desirable traffic from reaching its destination. There is a computational
4371 expense associated with processing a SIP transaction at a proxy server, and that expense is greater for
4372 stateful proxy servers than it is for stateless proxy servers. Therefore stateful proxies are more susceptible
4373 to flooding than stateless proxy servers.

4374 User agents and proxy servers SHOULD challenge questionable requests with only a *single* 401 (Unau-
4375 thorized) or 407 (Proxy Authentication Required), forgoing the normal response retransmission algorithm,
4376 and behaving statelessly towards unauthenticated requests.

4377 Retransmitting the 401 (Unauthorized) or 407 (Proxy Authentication Required) status response amplifies the
4378 problem of an attacker using a falsified header (such as Via) to direct traffic to a third party.

4379 With either TCP or UDP, a denial of service attack exists by a rogue proxy sending 6xx responses.
4380 Although a client SHOULD choose to ignore such responses if it requested authentication, a proxy cannot do
4381 so. It is obliged to forward the 6xx response back to the client. The client can then ignore the response, but
4382 if it repeats the request it will probably reach the same rogue proxy again, and the process will repeat.

4383 22.4 Limitations

4384 Although these security mechanisms, when applied in a judicious manner, can thwart many threats, there are
4385 limitations in the scope of the mechanisms that must be understood by implementers and network operators.

4386 22.4.1 HTTP Digest

4387 One of the primary limitations of using HTTP Digest in SIP is that the integrity mechanisms in Digest do
4388 not work very well for SIP. Specifically, they offer protection of the Request-URI and the method of a
4389 message, but not for any of the headers that user agents would most likely wish to secure.

4390 The existing replay protection mechanisms described in RFC 2617 also have some limitations for SIP.
4391 The next-nonce mechanism, for example, does not support pipelined requests. The nonce-count mechanism
4392 should be used for replay protection.

4393 Another limitation of HTTP Digest is the scope of realms. Digest is valuable when a user wants to
4394 authenticate themselves to a resource with which they have a pre-existing association, like a service provider
4395 of which the user is a customer. Consider that by contrast, the scope of TLS is global, since certificates are
4396 globally verifiable regardless of any pre-existing association between the user agent and the server.

4397 Future enhancements to HTTP Digest could conceivably resolve some or all of these limitations.

4398 **22.4.2 S/MIME**

4399 The largest outstanding defect with the S/MIME mechanism is the lack of prevalent public key infrastructure
4400 for end users. If self-signed certificates (or certificates that cannot be verified by one of the participants in
4401 a dialog) are used, the SIP-based key exchange mechanism described in Section 21.2 is susceptible to a
4402 man-in-the-middle attack with which an attacker can potentially inspect and modify S/MIME bodies. The
4403 attacker needs to intercept the first exchange of keys between the two parties in a dialog, remove the existing
4404 CMS detached signatures from the request and response, and insert a different CMS detached signature
4405 containing a certificate supplied by the attacker (but which seems to be a certificate for the proper address
4406 of record). Each party will think they have exchanged keys with the other, when in fact each has the public
4407 key of the attacker.

4408 It is important to note that the attacker can only leverage this vulnerability on the first exchange of keys
4409 between two parties - on subsequent occasions, the alteration of the key would be noticeable to user agents.
4410 It would also be difficult for the attacker to remain in the path of all future dialogs between the two parties
4411 over time (as potentially days, weeks, or years pass).

4412 SSH is susceptible to the same man-in-the-middle attack on the first exchange of keys; however, it is
4413 widely acknowledged that while SSH is not perfect, it does improve the security of connections. The use of
4414 key fingerprints could provide some assistance to SIP, just as it does for SSH. For example, if two parties use
4415 SIP to establish a voice communications session, each could read off the fingerprint of the key they received
4416 from the other, which could be compared against the original; it would certainly be more difficult for the
4417 man-in-the-middle to emulate the voices of the participants than their signaling.

4418 The S/MIME mechanism allows user agents to send encrypted requests without preamble if they possess
4419 a certificate for the destination address of record on their keyring. However, it is also possible that a device
4420 which does not hold certificates, or at least not that particular certificate, will be currently registered as
4421 the sole contact address for that address of record, and it will therefore be unable to properly process the
4422 encrypted request, which could lead to some avoidable error signaling. This is especially likely when an
4423 encrypted request is forked.

4424 The keys associated with S/MIME are most useful when associated with a particular user (an address
4425 of record) rather than a device (a user agent). When users move between devices, it may be difficult to
4426 transport private keys securely between user agents; how such keys might be acquired by a device is outside
4427 the scope of this document.

4428 Another, more prosaic difficulty with the S/MIME mechanism is that it can result in very large messages,
4429 especially when the SIP tunneling mechanism described in Section 21.4 is used. For that reason, it is
4430 RECOMMENDED that TCP should be used as a transport protocol when S/MIME tunneling is employed.

4431 **22.4.3 TLS**

4432 The most commonly voiced concern about TLS is that it cannot run over UDP; TLS requires a connection-
4433 oriented underlying transport protocol, which for the purposes of this document means TCP. Even running
4434 TCP, regardless of any additional overhead incurred by TLS, is argued to be too intensive for some embedded
4435 devices.

4436 It may also be arduous for a local outbound proxy server and/or registrar to maintain many simultane-
4437 ous long-lived TLS connections with numerous user agents might. This introduces some valid scalability
4438 concerns, especially for intensive ciphersuites. Maintaining redundancy of long-lived TLS connections,
4439 especially when a user agent is solely responsible for their establishment, could also be cumbersome.

4440 TLS only allows SIP entities to authenticate servers to which they are adjacent; TLS offers strictly
4441 hop-by-hop security. Neither TLS, nor any other mechanism specified in this document, allows clients to
4442 authenticate proxy servers to whom they cannot form a direct TCP connection.

4443 **22.5 Privacy**

4444 SIP messages frequently contain sensitive information about their senders - not just what they have to
4445 say, but with whom they communicate, when they communicate and for how long, and from where they
4446 participate in sessions. Many applications and their users require that this sort of private information be
4447 hidden from any parties that do not need to know it.

4448 Note that there are also less direct ways in which private information can be divulged. If a user or service
4449 chooses to be reachable at an address that is guessable from the person's name and organizational affiliation
4450 (which describes most addresses of record), the traditional method of ensuring privacy by having an unlisted
4451 "phone number" is compromised. A user location service can infringe on the privacy of the recipient of a
4452 session invitation by divulging their specific whereabouts to the caller; an implementation consequently
4453 SHOULD be able to restrict, on a per-user basis, what kind of location and availability information is given
4454 out to certain classes of callers.

4455 **23 Common Message Components**

4456 There are certain components of SIP messages that appear in various places within SIP messages (and
4457 sometimes, outside of them), which merit separate discussion.

4458 **23.1 SIP Uniform Resource Indicators**

4459 A SIP URI identifies a communications resource. Like all URIs, SIP URIs may be placed in web pages,
4460 email messages or printed literature. They contain sufficient information to initiate and maintain a commu-
4461 nication session with the resource.

4462 Examples of communications resources include

- 4463 ● a user of an online service;
- 4464 ● an appearance on a multiline phone;
- 4465 ● a mailbox on a messaging system
- 4466 ● a PSTN number at a gateway service;
- 4467 ● a group (such as "sales" or "helpdesk") in an organization.

4468 **23.1.1 SIP URI Components**

4469 The “sip:” scheme follows the guidelines in RFC 2396 [9]. It uses a form similar to the `mailto` URL,
4470 allowing the specification of SIP `request-header` fields and the `SIP message-body`. This makes it possible
4471 to specify the subject, media type, or urgency of sessions initiated by using a URI on a web page or in an
4472 email message. The formal syntax for a SIP URI is presented in Section 27. Its general form is

4473 `sip:user:password@host:port;url-parameters?headers`

4474 These tokens, and some of the tokens in their expansion, have the following meaning.

4475 **user:** The identifier of a particular resource at the host being addressed. Note that “host” as used here
4476 may, and frequently does, refer to a domain. The “userpart” of a URI consists of this user field, the
4477 password field and the @ sign following them. The userpart of a URI is optional and MAY be absent
4478 when the destination host does not have a notion of users or when the host itself is the resource being
4479 identified. If the @ sign is present in a SIP URI, the user field MUST NOT be empty.

4480 If the host being addressed is capable of processing telephone numbers, an Internet telephony gateway
4481 for instance, a `telephone-subscriber` field defined in RFC 2806 [13] MAY be used to populate the
4482 `user` field. There are special escaping rules for encoding `telephone-subscriber` fields in SIP URIs
4483 described in Section 23.1.2.

4484 **password:** A password associated with the user. While the SIP URI syntax allows this field to be present,
4485 its use is NOT RECOMMENDED, because the passing of authentication information in clear text (such
4486 as URIs) has proven to be a security risk in almost every case where it has been used. For instance,
4487 transporting a PIN number in this field exposes the PIN. Note that the password field is just an exten-
4488 sion of user portion. Implementations not wishing to give special significance to the password portion
4489 of the field MAY simply treat “user:password” as a single string.

4490 **host:** The entity hosting the SIP resource. The `host` part contains either a fully-qualified domain name
4491 or numeric IPv4 or IPv6 address. Using the fully-qualified domain name form is RECOMMENDED
4492 whenever possible.

4493 **port:** The port number where the request is to be sent.

4494 **URI parameters:** Parameters affecting a request constructed from the URI.

4495 URI parameters are added after the `hostport` component and are separated by semi-colons. URI
4496 parameters take the form:

4497 `parameter-name “=” parameter-value`

4498 Even though an arbitrary number of URI parameters may be included in a URI, any given parameter-
4499 name MUST NOT appear more than once.

4500 This extensible mechanism includes the `transport`, `maddr`, `ttl`, `user`, and `method` parameters.

4501 The `transport` parameter determines the transport mechanism to be used for sending SIP messages,
4502 as specified in [8]. SIP can use any network transport protocol. Parameter names are defined for
4503 UDP [30], TCP [31], TLS [28] (note that this is specifically TLS over TCP), and SCTP [32].

4504 The `maddr` parameter indicates the server address to be contacted for this user, overriding any address
4505 derived from the `host` field. When an `maddr` parameter is present, the `port` and `transport` components
4506 of the URI apply to the address indicated in the `maddr` parameter value. [8] describes the proper
4507 interpretation of the `transport`, `maddr` and `hostport` in order to obtain the destination address, port
4508 and `transport` for sending a request.

4509 The `maddr` field can be used as a simple form of loose source routing. It allows a URI to specify a specific
4510 proxy that must be traversed en-route to the destination. This capability is useful for a roaming user that
4511 is forced to use an outbound proxy, but wishes to force requests through their home proxy. Alternatively,
4512 preloaded Route values can be used to provide this capability (see item 8.1.1.1 in section 8.1.1).

4513 The `ttl` parameter determines the time-to-live value of the UDP multicast packet and **MUST** only be
4514 used if `maddr` is a multicast address and the transport protocol is UDP. For example, to specify to call
4515 `alice@atlanta.com` using multicast to `239.255.255.1` with a `ttl` of 15, the following URI would
4516 be used:

```
4517 sip:alice@atlanta.com;maddr=239.255.255.1;ttl=15
```

4518 The set of valid `telephone-subscriber` strings is a subset of valid `user` strings. The `user` URI pa-
4519 rameter exists to distinguish telephone numbers from user names that happen to look like telephone
4520 numbers. If the user string contains a telephone number formatted as a `telephone-subscriber`, the
4521 `user` parameter value “`phone`” **SHOULD** be present. Even without this parameter, recipients of SIP
4522 URIs **MAY** interpret the `pre-@` part as a telephone number if local restrictions on the name space for
4523 user name allow it.

4524 The method of the SIP request constructed from the URI can be specified with the `method` parameter.
4525 Since the `url-parameter` mechanism is extensible, SIP elements **MUST** silently ignore any `url-parameters`
4526 that they do not understand.

4527 **Headers:** Headers to be included in a request constructed from the URI. Headers fields in the SIP request
4528 can be specified with the “?” mechanism within a SIP URI. The header names and values are en-
4529 coded in ampersand separated `hname = hvalue` pairs. The special `hname` “`body`” indicates that the
4530 associated `hvalue` is the `message-body` of the SIP request.

4531 Table 1 summarizes the use of SIP URI components based on the context in which the URI appears. The
4532 external column describes URIs appearing anywhere outside of a SIP message, for instance on a web page
4533 or business card. Entries marked “`m`” are mandatory, those marked “`o`” are optional, and those marked “`-`”
4534 are not allowed. Elements processing URIs **SHOULD** ignore any disallowed components if they are present.
4535 The second column indicates the default value of an optional element if it is not present. “`-`” indicates that
4536 the element is either not optional, or has no default value.

4537 SIP URIs in `Contact` header fields have different restrictions depending on the context in which the
4538 header field appears. One set applies to messages that establish and maintain dialogs (`INVITE` and its 200
4539 (OK) response). The other applies to registration and redirection messages (`REGISTER`, its 200 (OK)
4540 response, 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	5060	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	udp	o	-	-	o	o	o
other-param	–	o	o	o	o	o	o
headers	–	-	-	-	o	-	o

Table 1: Use and default values of URI components for SIP headers, Request-URI and references

4541 23.1.2 Character Escaping Requirements

4542 SIP follows the requirements and guidelines of RFC 2396 [9] when defining the set of characters that must
4543 be escaped in a SIP URI, and uses its “”%” HEX HEX” mechanism for escaping. From RFC 2396:

4544 The set of characters actually reserved within any given URI component is defined by that com-
4545 ponent. In general, a character is reserved if the semantics of the URI changes if the character
4546 is replaced with its escaped US-ASCII encoding. [9].

4547 Excluded US-ASCII characters [9, Sec. 2.4.3], such as space and control characters and characters used as
4548 URI delimiters, also MUST be escaped. URIs MUST NOT contain unescaped space and control characters.

4549 For each component, the set of valid BNF expansions defines exactly which characters may appear
4550 unescaped. All other characters MUST be escaped.

4551 For example, “@” is not in the set of characters in the user component, so the user “j@s0n” must have
4552 at least the @ sign encoded, as in “j%40s0n”.

4553 Expanding the hname and hvalue tokens in Section 27 show that all URI reserved characters in header
4554 names and values MUST be escaped.

4555 The telephone-subscriber subset of the user component has special escaping considerations. The set
4556 of characters not reserved in the RFC 2806 [13] description of telephone-subscriber contains a number
4557 of characters in various syntax elements that need to be escaped when used in SIP URIs. Any characters
4558 occurring in a telephone-subscriber that do not appear in an expansion of the BNF for the user rule MUST
4559 be escaped.

4560 Note that character escaping is not allowed in the host component of a SIP URI (the % character is not
4561 valid in its expansion). This is likely to change in the future as requirements for Internationalized Domain
4562 Names are finalized. Current implementations MUST NOT attempt to improve robustness by treating received
4563 escaped characters in the host component as literally equivalent to their unescaped counterpart. The behavior
4564 required to meet the requirements of IDN may be significantly different.

4565 23.1.3 Example SIP URIs

4566 sip:alice@atlanta.com
4567 sip:alice:secretword@atlanta.com;transport=tcp
4568 sip:alice@atlanta.com?subject=project%20x&priority=urgent
4569 sip:+1-212-555-1212:1234@gateway.com;user=phone
4570 sip:1212@gateway.com
4571 sip:alice@192.0.2.4
4572 sip:atlanta.com;method=REGISTER?to=alice%40atlanta.com
4573 sip:alice;day=tuesday@atlanta.com

4574 The last example URI above has a **user** field value of “alice;day=tuesday”. The escaping rules defined
4575 above allow a semicolon to appear unescaped in this field. Note, however, that for the purposes of this
4576 protocol, the field is opaque. The apparent structure in that value is only useful to the entity responsible for
4577 the resource.

4578 23.1.4 SIP URI Comparison

4579 SIP URIs are compared for equality according to the following rules:

- 4580 • Comparison of the userpart of sip URIs is case-sensitive. This includes userparts containing pass-
4581 words or formatted as telephone-subscribers. Comparison of all other components of the URI is
4582 case-insensitive unless explicitly defined otherwise.
- 4583 • The ordering of parameters and headers is not significant in comparing SIP URIs.
- 4584 • Characters other than those in the “reserved” and “unsafe” sets (see RFC 2396 [9]) are equivalent to
4585 their “”%” HEX HEX” encoding.
- 4586 • An IP address that is the result of a DNS lookup of a host name does **not** match that host name.
- 4587 • For two URIs to be equal, the **user**, **password**, **host**, and **port** components must match. A URI
4588 omitting the optional port component will match a URI explicitly declaring port 5060. A URI omitting
4589 the user component will **not** match a URI that includes one. A URI omitting the password component
4590 will **not** match a URI that includes one.
- 4591 • URI **uri-parameter** components are compared as follows
 - 4592 – Any **uri-parameter** appearing in both URIs must match.
 - 4593 – A **user**, **transport**, **ttl**, or **method url-parameter** appearing in only one URI must contain its
4594 default value or the URIs do not match.
4595 A URI that includes an **maddr** parameter will **not** match a URI that contains no **maddr** parameter.
 - 4596
 - 4597 – All other **url-parameters** appearing in only one URI are ignored when comparing the URIs.
- 4598 • URI **header** components are never ignored. Any present **header** component **MUST** be present in
4599 both URIs and match for the URIs to match. The matching rules are defined for each header in
4600 Section sec:header-fields.

4601 The URIs within each of the following sets are equivalent:

4602 sip:%61lice@atlanta.com:5060

4603 sip:alice@AtLanTa.CoM;Transport=udp

4604 sip:carol@chicago.com

4605 sip:carol@chicago.com;newparam=5

4606 sip:carol@chicago.com;security=on

4607 sip:biloxi.com;transport=tcp;method=REGISTER?to=sip:bob%40biloxi.com

4608 sip:biloxi.com;method=REGISTER;transport=tcp?to=sip:bob%40biloxi.com

4609 sip:alice@atlanta.com?subject=project%20x&priority=urgent

4610 sip:alice@atlanta.com?priority=urgent&subject=project%20x

4611 The URIs within each of the following sets are **not** equivalent:

4612 SIP:ALICE@AtLanTa.CoM;Transport=udp (different usernames)

4613 sip:alice@AtLanTa.CoM;Transport=UDP

4614 sip:bob@biloxi.com (different port and transport)

4615 sip:bob@biloxi.com:6000;transport=tcp

4616 sip:carol@chicago.com (different header component)

4617 sip:carol@chicago.com?Subject=next%20meeting

4618 sip:bob@phone21.bboxesbybob.com (even though that's what

4619 sip:bob@192.0.2.4 phone21.bboxesbybob.com resolves to)

4620 Note that equality is not transitive:

4621 sip:carol@chicago.com and sip:carol@chicago.com;security=on are equivalent

4622 and sip:carol@chicago.com and sip:carol@chicago.com;security=off are equivalent

4623 But sip:carol@chicago.com;security=on and sip:carol@chicago.com;security=off are **not** equivalent

4624 Comparing URIs is a major part of comparing several SIP headers (see Section 24).

4625 23.1.5 Forming Requests from a SIP URI

4626 An implementation must take care when forming requests directly from a URI. URIs from business cards,
4627 web pages, and even from sources inside the protocol such as registered contacts may contain inappropriate
4628 header fields or body parts.

4629 An implementation **MUST** include any provided transport, maddr, ttl, or user parameter in the Request-
4630 URI of the formed request. If the URI contains a method parameter, its value **MUST** be used as the method

4631 of the request. The `method` parameter MUST NOT be placed in the Request-URI. Unknown URI parameters
4632 MUST be placed in the message's Request-URI.

4633 An implementation SHOULD treat the presence of any headers or body parts in the URI as a request to
4634 include them in the message, and choose to honor the request on an per-component basis.

4635 An implementation SHOULD NOT honor these obviously dangerous header fields: `From`, `Call-ID`, `CSeq`,
4636 `Via`, and `Record-Route`.

4637 An implementation SHOULD take special care in honoring any requested `Route` header field values in
4638 order to not be used as an unwitting agent in malicious attacks.

4639 An implementation SHOULD NOT honor requests to include headers that may cause it to falsely advertise
4640 its location or capabilities. These include: `Accept`, `Accept-Encoding`, `Accept-Language`, `Allow`, `Contact`
4641 (in its dialog usage), `Organization`, `Supported`, and `User-Agent`.

4642 An implementation SHOULD verify the accuracy of any requested descriptive headers, including: `Content-`
4643 `Disposition`, `Content-Encoding`, `Content-Language`, `Content-Length`, `Content-Type`, `Date`, `Mime-`
4644 `Version`, and `Timestamp`.

4645 If the request formed from constructing a message from a given URI is not a valid SIP request, the URI
4646 is invalid. An implementation MUST NOT proceed with transmitting the request. It should instead pursue
4647 the course of action due an invalid URI in the context it occurs.

4648 The constructed request can be invalid in many ways. These include, but are not limited to, syntax error in
4649 header fields, invalid combinations of URI parameters, or an incorrect description of the message body.

4650 Sending a request formed from a given URI may require capabilities unavailable to the implementation.
4651 The URI might indicate use of an unimplemented transport or extension for example. An implementation
4652 SHOULD refuse to send these requests rather than modifying them to match their capabilities. An imple-
4653 mentation MUST NOT send a request requiring an extension that it does not support.

4654 For example, such a request can be formed through the presence of a `headerRequire` header parameter or a
4655 `method` URI parameter with an unknown or explicitly unsupported value.

4656 **23.1.6 Relating SIP URIs and tel URLs**

4657 When a tel URL [13] is converted to a SIP URI, the entire telephone-subscriber portion of the tel URL,
4658 including any paramters,is placed into the userpart of the SIP URI.

4659 Thus, `tel:+358-555-1234567;postd=pp22` becomes

4660 `sip:+358-555-1234567;postd=pp22@foo.com`

4661 not

4662 `sip:+358-555-1234567@foo.com;postd=pp22`

4663 In general, equivalent "tel" URLs converted to SIP URIs in this fashion may not produce equivalent SIP
4664 URIs. The userpart of SIP URIs is compared as a case-sensitive string. Variance in case-insensitive portions
4665 of tel URLs and reordering of tel URL parameters does not affect tel URL equivalence, but does affect the
4666 equivalence of SIP URIs formed from them.

4667 For example,

4668 `tel:+358-555-1234567;postd=pp22`

4669 tel:+358-555-1234567;POSTD=PP22

4670 are equivalent, while

4671 sip:+358-555-1234567;postd=pp22@foo.com

4672 sip:+358-555-1234567;POSTD=PP22@foo.com

4673 are not.

4674 Likewise,

4675 tel:+358-555-1234567;postd=pp22;isub=1411

4676 tel:+358-555-1234567;isub=1411;postd=pp22

4677 are equivalent, while

4678 sip:+358-555-1234567;postd=pp22;isub=1411@foo.com

4679 sip:+358-555-1234567;isub=1411;postd=pp22@foo.com

4680 are not.

4681 To mitigate this problem, elements constructing telephone-subscriber fields to place in the userpart of
4682 a SIP URI SHOULD fold any case-insensitive portion of telephone-subscriber to lower case, and order the
4683 telephone-subscriber parameters lexically by parameter name. (All components of a tel URL except for
4684 future-extension parameters are defined to be compared case-insensitive.)

4685 Following this suggestion, both

4686 tel:+358-555-1234567;postd=pp22

4687 tel:+358-555-1234567;POSTD=PP22

4688 become

4689 sip:+358-555-1234567;postd=pp22@foo.com

4690 and both

4691 tel:+358-555-1234567;postd=pp22;isub=1411

4692 tel:+358-555-1234567;isub=1411;postd=pp22

4693 become

4694 sip:+358-555-1234567;isub=1411;postd=pp22

4695 23.2 Option Tags

4696 Option tags are unique identifiers used to designate new options (extensions) in SIP. These tags are used in
4697 Require (Section 24.33), Proxy-Require (Section 24.29), Supported (Section 24.39) and Unsupported
4698 (Section 24.42) header fields. Note that these options appear as parameters in those headers in an option-tag

4699 = token form (see Section 27 for the definition of token).

4700 The creator of a new SIP option **MUST** either prefix the option with their reverse domain name or register
4701 the new option with the Internet Assigned Numbers Authority (IANA) (See Section 28).

4702 An example of a reverse-domain-name option is “com.foo.mynewfeature”, whose inventor can be reached
4703 at “foo.com”. For these features, individual organizations are responsible for ensuring that option names do
4704 not collide within the same domain. The host name part of the option **MUST** use lower-case; the option name
4705 is case-insensitive.

4706 Options registered with IANA do not contain periods and are globally unique. IANA option tags are
4707 case-insensitive.

4708 **23.3 Tags**

4709 The “tag” parameter is used in the **To** and **From** fields of SIP messages. It serves as a general mechanism
4710 to identify a particular instance of a user agent for a particular SIP URI.

4711 As proxies can fork requests, the same request can reach multiple instances of a user (mobile and home
4712 phones, for example). Since each can respond, there needs to be a means for the originator of a session to
4713 distinguish the responses. Tag fields in the **To** and **From** disambiguate these multiple instances of the same
4714 user.

4715 This situation also arises with multicast requests.

4716 When a tag is generated by a UA for insertion into a request or response, it **MUST** be globally unique
4717 and cryptographically random with at least 32 bits of randomness. A property of this selection requirement
4718 is that a UA will place a different tag into the **From** header of an **INVITE** as it would place into the **To**
4719 header of the response to the same **INVITE**. This is needed in order for a UA to invite itself to a session, a
4720 common case for “hairpinning” of calls in PSTN gateways. Similarly, two **INVITE**s for different calls will
4721 have different **From** tags.

4722 Besides the requirement for global uniqueness, the algorithm for generating a tag is implementation
4723 specific. Tags are helpful in fault tolerant systems, where a dialog is to be recovered on an alternate server
4724 after a failure. A UAS can select the tag in such a way that a backup can recognize a request as part of a
4725 dialog on the failed server, and therefore determine that it should attempt to recover the dialog and any other
4726 state associated with it.

4727 **24 Header Fields**

4728 The general syntax for header fields is covered in Section 7.3. This section lists the full set of header fields
4729 along with notes on syntax, meaning, and usage. Throughout this section, we use [HX.Y] to refer to Section
4730 X.Y of the current HTTP/1.1 specification RFC 2616 [12]. Examples of each header field are given.

4731 Information about header fields in relation to methods and proxy processing is summarized in Tables 2
4732 and 3.

4733 The “where” column describes the request and response types in which the header field can be used.
4734 Values in this column are:

4735 **R:** header fields may only appear in requests;

4736 **r:** header field may only appear in responses;

4737 **2xx, 4xx, etc.:** A numerical value or range indicates response codes with which the header field can be
4738 used;

4739 **c:** header field is copied from the request to the response.

4740 An empty entry in the “where” column indicates that the header may be present in all requests and re-
4741 sponses.

4742 The “proxy” column describes the operations a proxy may perform on a header:

4743 **c:** A proxy can add (concatenate) comma-separated elements to the header.

4744 **m:** A proxy can modify the header.

4745 **a:** A proxy can add the header if not present.

4746 **r:** A proxy must be able to read the header and thus this header cannot be encrypted.

4747 The next six columns relate to the presence of a header field in a method:

4748 **o:** The header field is optional.

4749 **m:** The header field is mandatory.

4750 **m*:** The header field SHOULD be sent, but servers need to be prepared to receive messages without that
4751 header field.

4752 **t:** The header field SHOULD be sent, but servers need to be prepared to receive messages without that header
4753 field. If TCP is used as transport, then the header field MUST be sent.

4754 ***:** The header field is required if the message body is not empty. See sections 24.14, 24.15 and 7.4 for
4755 details.

4756 **-:** The header field is ignored.

4757 **c:** Conditional; the header field is either mandatory or optional, depending on the presence of a route set or
4758 the response code.

4759 “Optional” means that a UA MAY include the header field in a request or response, and a UA MAY ignore
4760 the header field if present in the request or response (The exception to this rule is the **Require** header field
4761 discussed in 24.33). A “mandatory” header field MUST be present in a request, and MUST be understood
4762 by the UAS receiving the request. A mandatory response header field MUST be present in the response,
4763 and the header field MUST be understood by the UAC processing the response. “Not applicable” means that
4764 the header field MUST NOT be present in a request. If one is placed in a request by mistake, it MUST be
4765 ignored by the UAS receiving the request. Similarly, a header field labeled “not applicable” for a response
4766 means that the UAS MUST NOT place the header in the response, and the UAC MUST ignore the header in
4767 the response. A UA SHOULD ignore extension header parameters that are not understood.

4768 A compact form of some common header fields is also defined for use when overall message size is an
4769 issue.

4770 The **Contact**, **From**, and **To** header fields contain a URI. If the URI contains a comma, question mark
4771 or semicolon, the URI MUST be enclosed in angle brackets (< and >). Any URI parameters are contained
4772 within these brackets. If the URI is not enclosed in angle brackets, any semicolon-delimited parameters are
4773 header-parameters, not URI parameters.

Header field	where	proxy	ACK	BYE	CAN	INV	OPT	REG	PRA
Accept	R		-	o	-	m*	m*	o	o
Accept	2xx		-	-	-	m*	m*	o	-
Accept	415		-	o	-	o	o	o	o
Accept-Encoding	R		-	o	-	m*	o	o	o
Accept-Encoding	2xx		-	-	-	m*	m*	o	-
Accept-Encoding	415		-	o	-	o	o	o	o
Accept-Language	R		-	o	-	m*	o	o	o
Accept-Language	2xx		-	-	-	m*	m*	o	-
Accept-Language	415		-	o	-	o	o	o	o
Alert-Info	R	am	-	-	-	o	-	-	-
Alert-Info	180	am	-	-	-	o	-	-	-
Allow	R		o	o	o	o	o	o	o
Allow	2xx		-	o	o	m*	m*	o	o
Allow	r		-	o	o	o	o	o	o
Allow	405		-	m	m	m	m	m	m
Authentication-Info	2xx		-	o	-	o	o	o	o
Authorization	R		o	o	o	o	o	o	o
Call-ID	c	r	m	m	m	m	m	m	m
Call-Info		am	-	-	-	o	o	o	-
Contact	R		o	-	-	m	o	o	-
Contact	1xx		-	-	-	o	o	-	-
Contact	2xx		-	-	-	m	o	o	-
Contact	3xx		-	o	-	o	o	o	o
Contact	485		-	o	-	o	o	o	o
Content-Disposition			o	o	-	o	o	o	o
Content-Encoding			o	o	-	o	o	o	o
Content-Language			o	o	-	o	o	o	o
Content-Length		r	t	t	t	t	t	t	t
Content-Type			*	*	-	*	*	*	*
CSeq	c	r	m	m	m	m	m	m	m
Date		a	o	o	o	o	o	o	o
Error-Info	300-699		-	o	o	o	o	o	o
Expires			-	-	-	o	-	o	-
From	c	r	m	m	m	m	m	m	m
In-Reply-To	R		-	-	-	o	-	-	-
Max-Forwards	R	amr	m	m	m	m	m	m	m
Min-Expires	423		-	-	-	-	-	m	-
MIME-Version			o	o	o	o	o	o	o
Organization		am	-	-	-	o	o	o	-

Table 2: Summary of header fields, A–O

Header field	where	proxy	ACK	BYE	CAN	INV	OPT	REG	PRA
Priority	R	a	-	-	-	o	-	-	-
Proxy-Authenticate	407		-	m	m	m	m	m	m
Proxy-Authorization	R	r	o	o	o	o	o	o	o
Proxy-Require	R	r	-	o	-	o	o	o	o
RAck	R		-	-	-	-	-	-	m
Record-Route	R	amr	o	o	o	o	o	-	o
Record-Route	2xx,401,484		-	o	o	o	o	-	o
Reply-To			-	-	-	o	-	-	-
Require		acr	-	o	-	o	o	o	o
Retry-After	404,413,480,486		-	o	o	o	o	o	o
	500,503		-	o	o	o	o	o	o
	600,603		-	o	o	o	o	o	o
Route	R	r	c	c	c	c	c	-	c
RSeq	1xx		-	o	-	o	o	o	-
Server	r		-	o	o	o	o	o	o
Subject	R		-	-	-	o	-	-	-
Supported	R		-	o	o	o	o	o	o
Supported	2xx		-	o	o	o	m*	o	o
Timestamp			o	o	o	o	o	o	o
To	c(1)	r	m	m	m	m	m	m	m
Unsupported	420		-	o	o	o	o	o	o
User-Agent			o	o	o	o	o	o	o
Via	c	acmr	m	m	m	m	m	m	m
Warning	r		-	o	o	o	o	o	o
WWW-Authenticate	401		-	m	m	m	m	m	m

Table 3: Summary of header fields, P-Z; (1): copied with possible addition of tag

4774 24.1 Accept

4775 The Accept header follows the syntax defined in [H14.1]. The semantics are also identical, with the excep-
 4776 tion that if no Accept header is present, the server SHOULD assume a default value of application/sdp.
 4777 An empty Accept header means that no formats are acceptable.

4778 Example:

4779 Accept: application/sdp;level=1, application/x-private, text/html

4780 24.2 Accept-Encoding

4781 The Accept-Encoding header field is similar to Accept, but restricts the content-codings [H3.5] that are
 4782 acceptable in the response. See [H14.3]. The syntax of this header is defined in [H14.3]. The semantics in
 4783 SIP are identical to those defined in [H14.3].

4784 An empty Accept-Encoding header field is permissible, even though the syntax in [H14.3] does not
 4785 provide for it. It is equivalent to Accept-Encoding: identity, that is, only the identity encoding, meaning

4786 no encoding, is permissible. If no **Accept-Encoding** header is present, the server SHOULD assume a default
4787 value of **identity**. This differs slightly from the HTTP definition, which indicates that when not present,
4788 any encoding can be used, but the **identity** encoding is preferred.

4789 Example:

4790 `Accept-Encoding: gzip`

4791 **24.3 Accept-Language**

4792 The **Accept-Language** header is used in requests to indicate the preferred languages for reason phrases,
4793 session descriptions, or status responses carried as message bodies in the response. If no **Accept-Language**
4794 header is present, the server SHOULD assume all languages are acceptable to the client. The **Accept-**
4795 **Language** header follows the syntax defined in [H14.4]. The rules for ordering the languages based on the
4796 “q” parameter apply to SIP as well.

4797 Example:

4798 `Accept-Language: da, en-gb;q=0.8, en;q=0.7`

4799 **24.4 Alert-Info**

4800 When present in an **INVITE** request, the **Alert-Info** header field specifies an alternative ring tone to the UAS.
4801 When present in a 180 (Ringing) response, the **Alert-Info** header field specifies an alternative ringback tone
4802 to the UAC. A typical usage is for a proxy to insert this header to provide a distinctive ring feature.

4803 The **Alert-Info** header can introduce security risks. These risks and the ways to handle them are dis-
4804 cussed in Section 24.9, which discusses the **Call-Info** header since the risks are identical.

4805 In addition, a user SHOULD be able to disable this feature selectively.

4806 This helps prevent disruptions that could result from the use of this header by untrusted elements.

4807 Example:

4808 `Alert-Info: <http://www.example.com/sounds/moo.wav>`

4809 **24.5 Allow**

4810 The **Allow** header field lists the set of methods supported by the UA generating the message.

4811 All methods, including **ACK** and **CANCEL**, understood by the UA MUST be included in the list of
4812 methods in the **Allow** header, when present. The absence of an **Allow** header MUST NOT be interpreted to
4813 mean that the UA sending the message supports no methods. Rather, it implies that the UA is not providing
4814 any information on what methods it supports.

4815 Supplying an **Allow** header in responses to methods other than **OPTIONS** reduces the number of mes-
4816 sages needed.

4817 Example:

4818 `Allow: INVITE, ACK, OPTIONS, CANCEL, BYE`

4819 24.6 Authentication-Info

4820 The Authentication-Info header provides for mutual authentication with HTTP Digest. A UAS MAY include
4821 this header in a 2xx response to a request that was successfully authenticated using digest based on the
4822 Authorization header.

4823 Syntax and semantics follow those specified in RFC 2617 [23].

4824 Example:

```
4825 Authentication-Info: nextnonce="47364c23432d2e131a5fb210812c"
```

4826 24.7 Authorization

4827 The Authorization header field contains authentication credentials of a UA. Section 20.2 overviews the use
4828 of the Authorization header field, and Section 20.4 describes the syntax and semantics when used with
4829 HTTP authentication. This header field, along with Proxy-Authorization, breaks the general rules about
4830 multiple header fields. Although not a comma-separated list, this header field may be present multiple times,
4831 and MUST NOT be combined into a single header using the usual rules described in Section 7.3.

4832 In the example below, there are no quotes around the Digest parameter:

```
4833 Authorization: Digest username="Alice", realm="Bob's Friends",  
4834 nonce="84a4cc6f3082121f32b42a2187831a9e",  
4835 response="7587245234b3434cc3412213e5f113a5432"
```

4836 24.8 Call-ID

4837 The Call-ID header field uniquely identifies a particular invitation or all registrations of a particular client.
4838 A single multimedia conference can give rise to several calls with different Call-IDs, for example, if a user
4839 invites a single individual several times to the same (long-running) conference. Call-IDs are case-sensitive
4840 and are simply compared byte-by-byte.

4841 The compact form of the Call-ID header field is i.

4842 Examples:

```
4843 Call-ID: f81d4fae-7dec-11d0-a765-00a0c91e6bf6@biloxi.com  
4844 i:f81d4fae-7dec-11d0-a765-00a0c91e6bf6@192.0.2.4
```

4845 24.9 Call-Info

4846 The Call-Info header field provides additional information about the caller or callee, depending on whether
4847 it is found in a request or response. The purpose of the URI is described by the "purpose" parameter.
4848 The "icon" parameter designates an image suitable as an iconic representation of the caller or callee. The
4849 "info" parameter describes the caller or callee in general, for example, through a web page. The "card"
4850 parameter provides a business card, for example, in vCard [33] or LDIF [34] formats. Additional tokens can
4851 be registered using IANA and the procedures in Section 28.

4852 Use of the Call-Info header field can pose a security risk. If a callee fetches the URIs provided by a
4853 malicious caller, the callee may be at risk for displaying inappropriate or offensive content, dangerous or
4854 illegal content, and so on. Therefore, it is RECOMMENDED that a UA only render the information in the

4855 Call-Info header if it can verify the authenticity of the element that originated the header and trusts that
4856 element. This need not be the peer UA; a proxy can insert this header into requests.

4857 Example:

```
4858 Call-Info: <http://www.example.com/alice/photo.jpg> ;purpose=icon,  
4859 <http://www.example.com/alice/> ;purpose=info
```

4860 24.10 Contact

4861 The Contact header field provides a URI whose meaning depends on the the type of request or response it
4862 is in.

4863 A Contact header field can contain a display name, a URI with URI parameters, and header parameters.

4864 This document defines the Contact parameters “q” and “expires”. These parameters are only used
4865 when the Contact is present in a REGISTER request or response, or in a 3xx response. Additional param-
4866 eters may be defined in other specifications.

4867 When the header field contains a display name, the URI including all URI parameters is enclosed in
4868 “<” and “>”. If no “<” and “>” are present, all parameters after the URI are header parameters, not URI
4869 parameters. The display name can be tokens, or a quoted string, if a larger character set is desired. Even if
4870 the “display-name” is empty, the “name-addr” form MUST be used if the “addr-spec” contains a comma,
4871 semicolon, or question mark. There may or may not be LWS between the display-name and the “<”. These
4872 rules for parsing a display name, URI and URI parameters, and header parameters also apply for the header
4873 fields To and From.

4874 The Contact header has a role similar to the Location header field in HTTP. However, the HTTP header field
4875 only allows one address, unquoted. Since URIs can contain commas and semicolons as reserved characters, they
4876 can be mistaken for header or parameter delimiters, respectively.

4877 The compact form of the Contact header field is m (for “moved”).

4878 The second example below shows a Contact header field containing both a URI parameter (transport)
4879 and a header parameter (expires).

```
4880 Contact: "Mr. Watson" <sip:watson@worchester.bell-telephone.com>  
4881 ;q=0.7; expires=3600,  
4882 "Mr. Watson" <mailto:watson@bell-telephone.com> ;q=0.1  
4883 m: <sip:bob@192.0.2.4;transport=tcp>;expires=60
```

4884 24.11 Content-Disposition

4885 The Content-Disposition header field describes how the message body or, for multipart messages, a mes-
4886 sage body part is to be interpreted by the UAC or UAS. This SIP header field extends the MIME Content-
4887 Type (RFC 1806 [35]).

4888 The value “session” indicates that the body part describes a session, for either calls or early (pre-call)
4889 media. The value “render” indicates that the body part should be displayed or otherwise rendered to the
4890 user. For backward-compatibility, if the Content-Disposition header is missing, the server SHOULD assume
4891 bodies of Content-Type application/sdp are the disposition “session”, while other content types are
4892 “render”.

4893 The disposition type “icon” indicates that the body part contains an image suitable as an iconic repre-
4894 sentation of the caller or callee. The value “alert” indicates that the body part contains information, such as
4895 an audio clip, that should be rendered instead of ring tone.

4896 The handling parameter, `handling-param`, describes how the UAS should react if it receives a message
4897 body whose content type or disposition type it does not understand. The parameter has defined values
4898 of “optional” and “required”. If the handling parameter is missing, the value “required” SHOULD be
4899 assumed. If this header field is missing, the MIME type determines the default content disposition. If there
4900 is none, “render” is assumed.

4901 Example:

```
4902 Content-Disposition: session
```

4903 24.12 Content-Encoding

4904 The Content-Encoding header field is used as a modifier to the “media-type”. When present, its value
4905 indicates what additional content codings have been applied to the entity-body, and thus what decoding
4906 mechanisms MUST be applied in order to obtain the media-type referenced by the Content-Type header
4907 field. Content-Encoding is primarily used to allow a body to be compressed without losing the identity of
4908 its underlying media type.

4909 If multiple encodings have been applied to an entity, the content codings MUST be listed in the order in
4910 which they were applied.

4911 All content-coding values are case-insensitive. IANA acts as a registry for content-coding value tokens.
4912 See [H3.5] for a definition of the syntax for content-coding.

4913 Clients MAY apply content encodings to the body in requests. A server MAY apply content encodings to
4914 the bodies in responses. The server MUST only use encodings listed in the Accept-Encoding header in the
4915 request.

4916 The compact form of the Content-Encoding header field is `e`. Examples:

```
4917 Content-Encoding: gzip  
4918 e: tar
```

4919 24.13 Content-Language

4920 See [H14.12]. Example:

```
4921 Content-Language: fr
```

4922 24.14 Content-Length

4923 The Content-Length header field indicates the size of the message-body, in decimal number of octets, sent
4924 to the recipient.

4925 Applications SHOULD use this field to indicate the size of the message-body to be transferred, regardless
4926 of the media type of the entity. If TCP is used as transport, the header field MUST be used. The size of the
4927 message-body does *not* include the CRLF separating headers and body. Any Content-Length greater than
4928 or equal to zero is a valid value. If no body is present in a message, then the Content-Length header field
4929 MUST be set to zero.

4930 The ability to omit `Content-Length` simplifies the creation of cgi-like scripts that dynamically generate re-
4931 sponses.

4932 The compact form of the header is l.

4933 Examples:

4934 Content-Length: 349

4935 l: 173

4936 **24.15 Content-Type**

4937 The `Content-Type` header field indicates the media type of the message-body sent to the recipient. The
4938 “media-type” element is defined in [H3.7]. The `Content-Type` header **MUST** be present if the body is not
4939 empty. If the body is empty, and a `Content-Type` header is present, it indicates that the body of the specific
4940 type has zero length (for example, an empty audio file).

4941 The compact form of the header is c.

4942 Examples:

4943 Content-Type: application/sdp

4944 c: text/html; charset=ISO-8859-4

4945 **24.16 CSeq**

4946 A `CSeq` header field in a request contains a single decimal sequence number and the request method. The
4947 sequence number **MUST** be expressible as a 32-bit unsigned integer. The `CSeq` header serves to order trans-
4948 actions within a dialog, to provide a means to uniquely identify transactions, and to differentiate between
4949 new requests and request retransmissions.

4950 Example:

4951 CSeq: 4711 INVITE

4952 **24.17 Date**

4953 The `Date` header field contains an RFC 1123 date (see [H14.18]). Unlike HTTP/1.1, SIP only supports
4954 the most recent RFC 1123 [36] format for dates. As in [H3.3], SIP restricts the timezone in `SIP-date` to
4955 “GMT”, while RFC 1123 allows any timezone. `rfc1123-date` is case-sensitive.

4956 The `Date` header field reflects the time when the request or response is first sent.

4957 The `Date` header field can be used by simple end systems without a battery-backed clock to acquire a notion of
4958 current time. However, in its GMT form, it requires clients to know their offset from GMT.

4959 Example:

4960 Date: Sat, 13 Nov 2010 23:29:00 GMT

4961 **24.18 Error-Info**

4962 The `Error-Info` header field provides a pointer to additional information about the error status response.

4963 SIP UACs have user interface capabilities ranging from pop-up windows and audio on PC softclients to audio-
4964 only on "black" phones or endpoints connected via gateways. Rather than forcing a server generating an error to
4965 choose between sending an error status code with a detailed reason phrase and playing an audio recording, the
4966 Error-Info header field allows both to be sent. The UAC then has the choice of which error indicator to render to the
4967 caller.

4968 A UAC MAY treat a SIP URI in an Error-Info header field as if it were a Contact in a redirect and
4969 generate a new INVITE, resulting in a recorded announcement session being established. A non-SIP URI
4970 MAY be rendered to the user.

4971 Examples:

```
4972 SIP/2.0 404 The number you have dialed is not in service  
4973 Error-Info: <sip:not-in-service-recording@atlanta.com>
```

4974 24.19 Expires

4975 The Expires header field gives the relative time after which the message (or content) expires. The precise
4976 meaning of this is method dependent.

4977 The expiration time in an INVITE does *not* affect the duration of the actual session that may result
4978 from the invitation. Session description protocols may offer the ability to express time limits on the session
4979 duration, however.

4980 The value of this field is an integer number of seconds (in decimal), measured from the receipt of the
4981 request.

4982 Examples:

```
4983 Expires: 5
```

4984 24.20 From

4985 The From header field indicates the initiator of the request. This may be different from the initiator of the
4986 dialog. Requests sent by the callee to the caller use the callee's address in the From header field.

4987 The optional "display-name" is meant to be rendered by a human user interface. A system SHOULD use
4988 the display name "Anonymous" if the identity of the client is to remain hidden. Even if the "display-name"
4989 is empty, the "name-addr" form MUST be used if the "addr-spec" contains a comma, question mark, or
4990 semicolon. Syntax issues are discussed in Section 7.3.1.

4991 Section 12 describes how From header fields are compared for the purpose of matching requests to
4992 dialogs. See Section 24.10 for the rules for parsing a display name, URI and URI parameters, and header
4993 parameters.

4994 The compact form of the header is f.

4995 Examples:

```
4996 From: "A. G. Bell" <sip:agb@bell-telephone.com> ;tag=a48s  
4997 From: sip:+12125551212@server.phone2net.com;tag=887s  
4998 f: Anonymous <sip:c8oqz84zk7z@privacy.org>;tag=hyh8
```


4999 **24.21 In-Reply-To**

5000 The **In-Reply-To** header field enumerates the **Call-IDs** that this call references or returns. These **Call-IDs**
5001 may have been cached by the client then included in this header in a return call.

5002 This allows automatic call distribution systems to route return calls to the originator of the first call. This also
5003 allows callees to filter calls, so that only return calls for calls they originated will be accepted. This field is not a
5004 substitute for request authentication.

5005 Example:

5006 `In-Reply-To: 70710@saturn.bell-tel.com, 17320@saturn.bell-tel.com`

5007 **24.22 Max-Forwards**

5008 The **Max-Forwards** header field must be used with any SIP method to limit the number of proxies or
5009 gateways that can forward the request to the next downstream server. This can also be useful when the client
5010 is attempting to trace a request chain that appears to be failing or looping in mid-chain.

5011 The **Max-Forwards** value is a decimal integer indicating the remaining number of times this request
5012 message is allowed to be forwarded. This count is decremented by each server that forwards the request.

5013 This header field should be inserted by elements that can not otherwise guarantee loop detection. For
5014 example, a B2BUA should insert a **Max-Forwards** header field.

5015 Example:

5016 `Max-Forwards: 6`

5017 **24.23 Min-Expires**

5018 The **Min-Expires** header field conveys the minimum registration expiration interval to a registrar. The
5019 header field contains a decimal integer number of seconds. The use of the header field in a 423 (Registration
5020 Too Brief) response is described in Sections 10.2.8, 10.3, and 25.4.17.

5021 Example:

5022 `Min-Expires: 60`

5023 **24.24 MIME-Version**

5024 See [H19.4.1].

5025 Example:

5026 `MIME-Version: 1.0`

5027 **24.25 Organization**

5028 The **Organization** header field conveys the name of the organization to which the entity issuing the request
5029 or response belongs.

5030 The field MAY be used by client software to filter calls.

5031 Example:

5032 Organization: Boxes by Bob

5033 **24.26 Priority**

5034 The Priority header field indicates the urgency of the request as perceived by the client. The Priority header
5035 field describes the priority that the SIP request should have to the receiving human or its agent. For example,
5036 it may be factored into decisions about call routing and acceptance. It does not influence the use of com-
5037 munications resources such as packet forwarding priority in routers or access to circuits in PSTN gateways.
5038 The header field can have the values "non-urgent", "normal", "urgent", and "emergency", but additional
5039 values can be defined elsewhere. It is RECOMMENDED that the value of "emergency" only be used when
5040 life, limb, or property are in imminent danger. Otherwise, there are no semantics defined for this header
5041 field.

5042 These are the values of RFC 2076 [37], with the addition of "emergency".

5043 Examples:

5044 Subject: A tornado is heading our way!

5045 Priority: emergency

5046 or

5047 Subject: Weekend plans

5048 Priority: non-urgent

5049 **24.27 Proxy-Authenticate**

5050 The Proxy-Authenticate header field contains an authentication challenge. The syntax for this header and
5051 its use is defined in [H14.33]. See 20.3 for further details on its usage.

5052 Example:

5053 Proxy-Authenticate: Digest realm="Carrier SIP",
5054 domain="sip:ss1.carrier.com",
5055 nonce="f84f1cec41e6cbe5aea9c8e88d359",
5056 opaque="", stale=FALSE, algorithm=MD5

5057 **24.28 Proxy-Authorization**

5058 The Proxy-Authorization header field allows the client to identify itself (or its user) to a proxy that requires
5059 authentication. The Proxy-Authorization field value consists of credentials containing the authentication
5060 information of the user agent for the proxy and/or realm of the resource being requested.

5061 See [H14.34] for a definition of the syntax, and section 20.3 for a discussion of its usage.

5062 This header field, along with **Authorization**, breaks the general rules about multiple header fields. Al-
5063 though not a comma-separated list, this header field may be present multiple times, and **MUST NOT** be
5064 combined into a single header using the usual rules described in Section 7.3.1.

5065 Example:

```
5066 Proxy-Authorization: Digest username="Alice", realm="Atlanta ISP",  
5067 nonce="c60f3082ee1212b402a21831ae",  
5068 response="245f23415f11432b3434341c022"
```

5069 24.29 Proxy-Require

5070 The **Proxy-Require** header field is used to indicate proxy-sensitive features that must be supported by the
5071 proxy. See Section 24.33 for more details on the mechanics of this message and a usage example.

5072 Example:

```
5073 Proxy-Require: foo
```

5074 24.30 RACK

5075 The **RACK** header is sent in a **PRACK** request to support reliability of provisional responses. It contains two
5076 numbers and a method tag. The first number is the value from the **RSeq** header in the provisional response
5077 that is being acknowledged. The next number, and the method, are copied from the **CSeq** in the response
5078 that is being acknowledged. The method name in the **RACK** header is case sensitive.

5079 Example:

```
5080 RACK: 776656 1 INVITE
```

5081 24.31 Record-Route

5082 The **Record-Route** is inserted by proxies in a request to force future requests in the session to be routed
5083 through the proxy.

5084 Details of its use with the **Route** header field are described in Section 16.4.

5085 Example:

```
5086 Record-Route: <sip:bob@biloxi.com;maddr=192.0.2.4>,  
5087 <sip:bob@biloxi.com;maddr=192.0.6.1>
```

5088 24.32 Reply-To

5089 The **Reply-To** header field contains a logical return URI which may be different from the **From** header field.
5090 For example, the URI **MAY** be used to return missed calls or unestablished sessions.

5091 If the user wished to remain anonymous, the header field **SHOULD** either be omitted from the request or
5092 populated in such a way that does not reveal any private information.

5093 Even if the “display-name” is empty, the “name-addr” form MUST be used if the “addr-spec” con-
5094 tains a comma, question mark, or semicolon. Syntax issues are discussed in Section 7.3.1.

5095 Example:

5096 Reply-To: Bob <sip:bob@biloxi.com>

5097 **24.33 Require**

5098 The **Require** header field is used by UACs to tell UASs about options that the UAC expects the UAS to
5099 support in order to process the request. Although an optional header, the **Require** MUST NOT be ignored if
5100 it is present.

5101 The **Require** header contains a list of option tags, described in Section 23.2. Each option tag defines
5102 a SIP extension that MUST be understood to process the request. Frequently, this is used to indicate that a
5103 specific set of extension headers need to be understood. A UAC compliant to this specification MUST only
5104 include option tags corresponding to standards-track RFCs.

5105 Example:

5106 Require: 100rel

5107 **24.34 Retry-After**

5108 The **Retry-After** header field can be used with a 503 (Service Unavailable) response to indicate how long
5109 the service is expected to be unavailable to the requesting client and with a 404 (Not Found), 600 (Busy), or
5110 603 (Decline) response to indicate when the called party anticipates being available again. The value of this
5111 field is a positive integer number of seconds (in decimal) after the time of the response.

5112 An optional comment can be used to indicate additional information about the time of callback. An
5113 optional “duration” parameter indicates how long the called party will be reachable starting at the initial
5114 time of availability. If no duration parameter is given, the service is assumed to be available indefinitely.

5115 Examples:

5116 Retry-After: 18000;duration=3600

5117 Retry-After: 120 (I'm in a meeting)

5118 **24.35 Route**

5119 The **Route** is used to force routing for a request through the listed set of proxies. Details of its use with the
5120 **Record-Route** header field are described in Section 13.

5121 Example:

5122 Route: <sip:bob@biloxi.com;maddr=192.0.2.4>, <sip:bob@pc33.atlanta.com>

5123 **24.36 RSeq**

5124 The RSeq header is used in provisional responses in order to transmit them reliably. It contains a single
5125 numeric value from 1 to $2^{32} - 1$. For details on its usage, see Section 18.1.

5126 Example:

5127 RSeq: 988789

5128 **24.37 Server**

5129 The Server header field contains information about the software used by the UAS to handle the request.
5130 The syntax for this field is defined in [H14.38].

5131 Revealing the specific software version of the server might allow the server to become more vulnerable
5132 to attacks against software that is known to contain security holes. Implementors SHOULD make the Server
5133 header field a configurable option.

5134 Example:

5135 Server: HomeProxy v2

5136 **24.38 Subject**

5137 The Subject header field provides a summary or indicates the nature of the call, allowing call filtering
5138 without having to parse the session description. The session description does not have to use the same
5139 subject indication as the invitation.

5140 The compact form of the header is s.

5141 Example:

5142 Subject: Need more boxes

5143 s: Tech Support

5144 **24.39 Supported**

5145 The Supported header field enumerates all the extensions supported by the UAC or UAS.

5146 The Supported header contains a list of option tags, described in Section 23.2, that are understood by
5147 the UAC or UAS. A UA compliant to this specification MUST only include option tags corresponding to
5148 standards-track RFCs. If empty, it means that no extensions are supported.

5149 Example:

5150 Supported: 100rel

5151 **24.40 Timestamp**

5152 The Timestamp header field describes when the UAC sent the request to the UAS. See Section 8.2.6 for
5153 details on how to generate a response to a request that contains the header field, and Section 17.3 for usage
5154 in RTT estimation.

5155 Example:

5156 Timestamp: 54

5157 **24.41 To**

5158 The **To** header field specifies the logical recipient of the request.

5159 The optional “display-name” is meant to be rendered by a human-user interface. The “tag” parameter
5160 serves as a general mechanism to distinguish multiple instances of a user identified by a single SIP URI.

5161 See Section 13 for details of the “tag” parameter.

5162 Section 12 describes how **To** and **From** header fields are compared for the purpose of matching requests
5163 to dialogs. See Section 24.10 for the rules for parsing a display name, URI and URI parameters, and header
5164 parameters.

5165 The compact form of the header is t.

5166 The following are examples of valid **To** headers:

```
5167     To: The Operator <sip:operator@cs.columbia.edu>;tag=287447  
5168     t: sip:+12125551212@server.phone2net.com
```

5169 **24.42 Unsupported**

5170 The **Unsupported** header field lists the features not supported by the UAS. See Section 24.33 for motivation.

5171 Example:

```
5172     Unsupported: foo
```

5173 **24.43 User-Agent**

5174 The **User-Agent** header field contains information about the UAC originating the request. The syntax and
5175 semantics are defined in [H14.43].

5176 Revealing the specific software version of the user agent might allow the user agent to become more
5177 vulnerable to attacks against software that is known to contain security holes. Implementors SHOULD make
5178 the **User-Agent** header field a configurable option.

5179 Example:

```
5180     User-Agent: Softphone Beta1.5
```

5181 **24.44 Via**

5182 The **Via** field indicates the path taken by the request so far and indicates the path that should be followed in
5183 routing responses. The branch ID parameter in the **Via** header serves as a transaction identifier, and is used
5184 by proxies to detect loops.

5185 The **Via** header field contains the transport protocol used to send the message, the client’s host name or
5186 network address and, if not the default port number, the port number at which it wishes to receive responses.
5187 The **Via** header field can also contain parameters such as “maddr”, “ttl”, “received”, and “branch”, whose
5188 meaning and use are described in other sections.

5189 Transport protocols defined here are “UDP”, “TCP”, “TLS”, and “SCTP”. “TLS” means TLS over
5190 TCP.

5191 The host or network address and port number are not required to follow the SIP URI syntax. Specifically,
5192 LWS on either side of the “:” or “/” is allowed, as shown in the second example below.

```
5193 Via: SIP/2.0/UDP erlang.bell-telephone.com:5060;branch=z9hG4bK87asdks7  
5194 Via: SIP/2.0/UDP 128.59.16.1:5060 ;received=128.59.19.3;branch=z9hG4bK77asjd
```

5195 The compact form of the header is v.

5196 In this example, the message originated from a multi-homed host with two addresses, 128.59.16.1
5197 and 128.59.19.3. The sender guessed wrong as to which network interface would be used. Erlang.bell-
5198 telephone.com noticed the mismatch and added a parameter to the previous hop’s Via header field, contain-
5199 ing the address that the packet actually came from.

5200 Another example:

```
5201 Via: SIP / 2.0 / UDP first.example.com: 4000;ttl=16  
5202 ;maddr=224.2.0.1 ;branch=z9hG4bKa7c6a8dlze.1
```

5203 Even though this specification mandates that the branch parameter be present in all requests, the BNF
5204 for the header indicates that it is optional. This allows interoperation with RFC 2543 elements, which did
5205 not have to insert the branch parameter.

5206 24.45 Warning

5207 The Warning header field is used to carry additional information about the status of a response. Warning
5208 headers are sent with responses and contain a three-digit warning code, host name, and warning text.

5209 The “warn-text” should be in a natural language that is most likely to be intelligible to the human user
5210 receiving the response. This decision can be based on any available knowledge, such as the location of the
5211 user, the Accept-Language field in a request, or the Content-Language field in a response. The default
5212 language is i-default [38].

5213 The currently-defined “warn-code”s are listed below, with a recommended warn-text in English and a
5214 description of their meaning. These warnings describe failures induced by the session description. The first
5215 digit of warning codes beginning with “3” indicates warnings specific to SIP. Warnings 300 through 329 are
5216 reserved for indicating problems with keywords in the session description, 330 through 339 are warnings
5217 related to basic network services requested in the session description, 370 through 379 are warnings related
5218 to quantitative QoS parameters requested in the session description, and 390 through 399 are miscellaneous
5219 warnings that do not fall into one of the above categories.

5220 **300 Incompatible network protocol:** One or more network protocols contained in the session description
5221 are not available.

5222 **301 Incompatible network address formats:** One or more network address formats contained in the ses-
5223 sion description are not available.

5224 **302 Incompatible transport protocol:** One or more transport protocols described in the session descrip-
5225 tion are not available.

5226 **303 Incompatible bandwidth units:** One or more bandwidth measurement units contained in the session
5227 description were not understood.

5228 **304 Media type not available:** One or more media types contained in the session description are not avail-
5229 able.

5230 **305 Incompatible media format:** One or more media formats contained in the session description are not
5231 available.

5232 **306 Attribute not understood:** One or more of the media attributes in the session description are not sup-
5233 ported.

5234 **307 Session description parameter not understood:** A parameter other than those listed above was not
5235 understood.

5236 **330 Multicast not available:** The site where the user is located does not support multicast.

5237 **331 Unicast not available:** The site where the user is located does not support unicast communication (usu-
5238 ally due to the presence of a firewall).

5239 **370 Insufficient bandwidth:** The bandwidth specified in the session description or defined by the media
5240 exceeds that known to be available.

5241 **399 Miscellaneous warning:** The warning text can include arbitrary information to be presented to a hu-
5242 man user or logged. A system receiving this warning MUST NOT take any automated action.

5243 1xx and 2xx have been taken by HTTP/1.1.

5244 Additional "warn-code"s, as in the example below, can be defined through IANA.

5245 Examples:

5246 Warning: 307 isi.edu "Session parameter 'foo' not understood"

5247 Warning: 301 isi.edu "Incompatible network address type 'E.164'"

5248 24.46 WWW-Authenticate

5249 The WWW-Authenticate header field contains an authentication challenge. The syntax for this header
5250 field and use is defined in [H14.47]. See 20.2 for further details on its usage.

5251 Example:

```
5252 WWW-Authenticate: Digest realm="Bob's Friends",  
5253 domain="sip:boxesbybob.com",  
5254 nonce="f84f1cec41e6cbe5aea9c8e88d359",  
5255 opaque="", stale=FALSE, algorithm=MD5
```

5256 25 Response Codes

5257 The response codes are consistent with, and extend, HTTP/1.1 response codes. Not all HTTP/1.1 response
5258 codes are appropriate, and only those that are appropriate are given here. Other HTTP/1.1 response codes
5259 SHOULD NOT be used. Response codes not defined by HTTP/1.1 have codes x80 upwards to avoid clashes
5260 with future HTTP response codes. Also, SIP defines a new class, 6xx.

5261 **25.1 Provisional 1xx**

5262 Provisional responses, also known as informational responses, indicate that the server or proxy contacted is
5263 performing some further action and does not yet have a definitive response. A server typically sends a 1xx
5264 response if it expects to take more than 200 ms to obtain a final response. Note that 1xx responses are not
5265 transmitted reliably, that is, they do not cause the client to send an ACK. Provisional (1xx) responses MAY
5266 contain message bodies, including session descriptions.

5267 **25.1.1 100 Trying**

5268 This response indicates that the request has been received by the next hop server and that some unspecified
5269 action is being taken on behalf of this call (e.g., a database is being consulted). This response, like all other
5270 provisional responses, stops retransmissions of an INVITE by a UAC. The 100 (Trying) response is different
5271 from other provisional responses, in that it is never forwarded upstream by a stateful proxy.

5272 **25.1.2 180 Ringing**

5273 The user agent receiving the INVITE is trying to alert the user. This response MAY be used to initiate local
5274 ringback.

5275 **25.1.3 181 Call Is Being Forwarded**

5276 A proxy server MAY use this status code to indicate that the call is being forwarded to a different set of
5277 destinations.

5278 **25.1.4 182 Queued**

5279 The called party is temporarily unavailable, but the callee has decided to queue the call rather than reject it.
5280 When the callee becomes available, it will return the appropriate final status response. The reason phrase
5281 MAY give further details about the status of the call, e.g., "5 calls queued; expected waiting time is 15
5282 minutes". The server MAY issue several 182 (Queued) responses to update the caller about the status of the
5283 queued call.

5284 **25.1.5 183 Session Progress**

5285 The 183 (Session Progress) response is used to convey information about the progress of the call which is
5286 not otherwise classified. The Reason-Phrase, header fields, or message body MAY be used to convey more
5287 details about the call progress.

5288 **25.2 Successful 2xx**

5289 The request was successful.

5290 **25.2.1 200 OK**

5291 The request has succeeded. The information returned with the response depends on the method used in the
5292 request.

5293 **25.3 Redirection 3xx**

5294 3xx responses give information about the user's new location, or about alternative services that might be
5295 able to satisfy the call.

5296 **25.3.1 300 Multiple Choices**

5297 The address in the request resolved to several choices, each with its own specific location, and the user (or
5298 user agent) can select a preferred communication end point and redirect its request to that location.

5299 The response *MAY* include a message body containing a list of resource characteristics and location(s)
5300 from which the user or user agent can choose the one most appropriate, if allowed by the *Accept* request
5301 header. However, no MIME types have been defined for this message body.

5302 The choices *SHOULD* also be listed as *Contact* fields (Section 24.10). Unlike HTTP, the SIP response
5303 *MAY* contain several *Contact* fields or a list of addresses in a *Contact* field. User agents *MAY* use the
5304 *Contact* header field value for automatic redirection or *MAY* ask the user to confirm a choice. However, this
5305 specification does not define any standard for such automatic selection.

5306 This status response is appropriate if the callee can be reached at several different locations and the server cannot
5307 or prefers not to proxy the request.

5308 **25.3.2 301 Moved Permanently**

5309 The user can no longer be found at the address in the *Request-URI* and the requesting client *SHOULD* retry
5310 at the new address given by the *Contact* header field (Section 24.10). The requestor *SHOULD* update any
5311 local directories, address books and user location caches with this new value and redirect future requests to
5312 the address(es) listed.

5313 **25.3.3 302 Moved Temporarily**

5314 The requesting client *SHOULD* retry the request at the new address(es) given by the *Contact* header field
5315 (Section 24.10). The *Request-URI* of the new request uses the value of the *Contact* header in the response.

5316 The duration of the validity of the *Contact* URI can be indicated through an *Expires* (Section 24.19)
5317 header field or an *expires* parameter in the *Contact* header field. Both proxies and UAs *MAY* cache this
5318 URI for the duration of the expiration time. If there is no explicit expiration time, the address is only valid
5319 once for recursing, and *MUST NOT* be cached for future transactions.

5320 If the URI cached from the *Contact* header field fails, the *Request-URI* from the redirected request
5321 *MAY* be tried again a single time.

5322 The temporary URI may have become out of date sooner than the expiration time, and a new temporary URI
5323 may be available.

5324 **25.3.4 305 Use Proxy**

5325 The requested resource *MUST* be accessed through the proxy given by the *Contact* field. The *Contact* field
5326 gives the URI of the proxy. The recipient is expected to repeat this single request via the proxy. 305 (Use
5327 Proxy) responses *MUST* only be generated by user agent servers.

5328 **25.3.5 380 Alternative Service**

5329 The call was not successful, but alternative services are possible. The alternative services are described in
5330 the message body of the response. Formats for such bodies are not defined here, and may be the subject of
5331 future standardization.

5332 **25.4 Request Failure 4xx**

5333 4xx responses are definite failure responses from a particular server. The client SHOULD NOT retry the
5334 same request without modification (e.g., adding appropriate authorization). However, the same request to a
5335 different server might be successful.

5336 **25.4.1 400 Bad Request**

5337 The request could not be understood due to malformed syntax. The Reason-Phrase SHOULD identify the
5338 syntax problem in more detail, e.g., "Missing Call-ID header".

5339 **25.4.2 401 Unauthorized**

5340 The request requires user authentication. This response is issued by user agent servers and registrars, while
5341 407 (Proxy Authentication Required) is used by proxy servers.

5342 **25.4.3 402 Payment Required**

5343 Reserved for future use.

5344 **25.4.4 403 Forbidden**

5345 The server understood the request, but is refusing to fulfill it. Authorization will not help, and the request
5346 SHOULD NOT be repeated.

5347 **25.4.5 404 Not Found**

5348 The server has definitive information that the user does not exist at the domain specified in the Request-
5349 URI. This status is also returned if the domain in the Request-URI does not match any of the domains
5350 handled by the recipient of the request.

5351 **25.4.6 405 Method Not Allowed**

5352 The method specified in the Request-Line is understood, but not allowed for the address identified by the
5353 Request-URI. The response MUST include an Allow header field containing a list of valid methods for the
5354 indicated address.

5355 **25.4.7 406 Not Acceptable**

5356 The resource identified by the request is only capable of generating response entities which have content
5357 characteristics not acceptable according to the accept headers sent in the request.

5358 **25.4.8 407 Proxy Authentication Required**

5359 This code is similar to 401 (Unauthorized), but indicates that the client **MUST** first authenticate itself with
5360 the proxy. SIP access authentication is explained in section 22 and 20.3.

5361 This status code can be used for applications where access to the communication channel (e.g., a tele-
5362 phony gateway) rather than the callee requires authentication.

5363 **25.4.9 408 Request Timeout**

5364 The server could not produce a response within a suitable amount of time, for example, if it could not
5365 determine the location of the user in time. The client **MAY** repeat the request without modifications at any
5366 later time.

5367 **25.4.10 410 Gone**

5368 The requested resource is no longer available at the server and no forwarding address is known. This
5369 condition is expected to be considered permanent. If the server does not know, or has no facility to determine,
5370 whether or not the condition is permanent, the status code 404 (Not Found) **SHOULD** be used instead.

5371 **25.4.11 413 Request Entity Too Large**

5372 The server is refusing to process a request because the request entity is larger than the server is willing or
5373 able to process. The server **MAY** close the connection to prevent the client from continuing the request.

5374 If the condition is temporary, the server **SHOULD** include a **Retry-After** header field to indicate that it is
5375 temporary and after what time the client **MAY** try again.

5376 **25.4.12 414 Request-URI Too Long**

5377 The server is refusing to service the request because the **Request-URI** is longer than the server is willing to
5378 interpret.

5379 **25.4.13 415 Unsupported Media Type**

5380 The server is refusing to service the request because the message body of the request is in a format not sup-
5381 ported by the server for the requested method. The server **SHOULD** return a list of acceptable formats using
5382 the **Accept**, **Accept-Encoding** and **Accept-Language** header fields. UAC processing of this response is
5383 described in Section 8.1.4.6.

5384 **25.4.14 416 Unsupported URI Scheme**

5385 The server cannot process the request because the scheme of the URI in the **Request-URI** is unknown to
5386 the server. Client processing of this response is described in Section 8.1.4.6.

5387 **25.4.15 420 Bad Extension**

5388 The server did not understand the protocol extension specified in a **Proxy-Require** (Section 24.29) or **Re-**
5389 **quire** (Section 24.33) header field. The server **SHOULD** include a list of the unsupported extensions in an
5390 **Unsupported** header in the response. UAC processing of this response is described in Section 8.1.4.6.

5391 **25.4.16 421 Extension Required**

5392 The UAS needs a particular extension to process the request, but this extension is not listed in a **Supported**
5393 header in the request. Responses with this status code **MUST** contain a **Require** header field listing the
5394 required extensions.

5395 A UAS **SHOULD NOT** use this response unless it truly cannot provide any useful service to the client.
5396 Instead, if a desirable extension is not listed in the **Supported** header field, servers **SHOULD** process the
5397 request using baseline SIP capabilities and any extensions supported by the client.

5398 **25.4.17 423 Registration Too Brief**

5399 The registrar is rejecting a registration request because a **Contact** header field expiration time was too small.
5400 The use of this response and the related **Min-Expires** header field are described in Sections 10.2.8, 10.3,
5401 and 24.23.

5402 **25.4.18 480 Temporarily Unavailable**

5403 The callee's end system was contacted successfully but the callee is currently unavailable (e.g., is not logged
5404 in, logged in in such a manner as to preclude communication with the callee or has activated the "do not
5405 disturb" feature). The response **MAY** indicate a better time to call in the **Retry-After** header. The user could
5406 also be available elsewhere (unbeknownst to this host). The reason phrase **SHOULD** indicate a more precise
5407 cause as to why the callee is unavailable. This value **SHOULD** be setable by the user agent. Status 486 (Busy
5408 Here) **MAY** be used to more precisely indicate a particular reason for the call failure.

5409 This status is also returned by a redirect or proxy server that recognizes the user identified by the
5410 **Request-URI**, but does not currently have a valid forwarding location for that user.

5411 **25.4.19 481 Call/Transaction Does Not Exist**

5412 This status indicates that the UAS received a request that does not match any existing dialog or transaction.

5413 **25.4.20 482 Loop Detected**

5414 The server has detected a loop (Section 2).

5415 **25.4.21 483 Too Many Hops**

5416 The server received a request that contains a **Max-Forwards** (Section 24.22) header with the value zero.

5417 **25.4.22 484 Address Incomplete**

5418 The server received a request with a **Request-URI** that was incomplete. Additional information **SHOULD**
5419 be provided in the reason phrase.

5420 This status code allows overlapped dialing. With overlapped dialing, the client does not know the length of the
5421 dialing string. It sends strings of increasing lengths, prompting the user for more input, until it no longer receives a
5422 484 (Address Incomplete) status response.

5423 25.4.23 485 Ambiguous

5424 The Request-URI was ambiguous. The response MAY contain a listing of possible unambiguous addresses
5425 in Contact header fields. Revealing alternatives can infringe on privacy of the user or the organization. It
5426 MUST be possible to configure a server to respond with status 404 (Not Found) or to suppress the listing of
5427 possible choices for ambiguous Request-URIs.

5428 Example response to a request with the Request-URI `sip:lee@example.com`:

```
5429 485 Ambiguous SIP/2.0
5430 Contact: Carol Lee <sip:carol.lee@example.com>
5431 Contact: Ping Lee <sip:p.lee@example.com>
5432 Contact: Lee M. Foote <sip:lee.foote@example.com>
```

5433 Some email and voice mail systems provide this functionality. A status code separate from 3xx is used since
5434 the semantics are different: for 300, it is assumed that the same person or service will be reached by the choices
5435 provided. While an automated choice or sequential search makes sense for a 3xx response, user intervention is
5436 required for a 485 (Ambiguous) response.

5437 25.4.24 486 Busy Here

5438 The callee's end system was contacted successfully but the callee is currently not willing or able to take
5439 additional calls at this end system. The response MAY indicate a better time to call in the **Retry-After**
5440 header. The user could also be available elsewhere, such as through a voice mail service. Status 600 (Busy
5441 Everywhere) SHOULD be used if the client knows that no other end system will be able to accept this call.

5442 25.4.25 487 Request Terminated

5443 The request was terminated by a BYE or CANCEL request. This response is never returned for a CANCEL
5444 request itself.

5445 25.4.26 488 Not Acceptable Here

5446 The response has the same meaning as 606 (Not Acceptable), but only applies to the specific entity addressed
5447 by the Request-URI and the request may succeed elsewhere. A message body containing a description of
5448 media capabilities MAY be present in the response, which is formatted according to the **Accept** header field
5449 in the INVITE (or application/sdp if not present), the same as a message body in a 200 (OK) response to
5450 an OPTIONS request.

5451 25.4.27 491 Request Pending

5452 The request was received by a UAS which had a pending request within the same dialog. Section 14.2
5453 describes how such "glare" situations are resolved.

5454 25.4.28 493 Undecipherable

5455 The request was received by a UAS which contained an encrypted MIME body for which the recipient
5456 does not possess or will not provide an appropriate decryption key. This response MAY have a single body

5457 containing an appropriate public key that should be used to encrypt MIME bodies sent to this user agent.
5458 Details of the usage of this response code can be found in Section 21.2.

5459 **25.5 Server Failure 5xx**

5460 5xx responses are failure responses given when a server itself has erred.

5461 **25.5.1 500 Server Internal Error**

5462 The server encountered an unexpected condition that prevented it from fulfilling the request. The client MAY
5463 display the specific error condition, and MAY retry the request after several seconds.

5464 If the condition is temporary, the server MAY indicate when the client may retry the request using the
5465 **Retry-After** header.

5466 **25.5.2 501 Not Implemented**

5467 The server does not support the functionality required to fulfill the request. This is the appropriate response
5468 when a UAS does not recognize the request method and is not capable of supporting it for any user. (Proxies
5469 forward all requests regardless of method.) Note that a 405 (Method Not Allowed) is sent when the server
5470 recognizes the request method, but that method is not allowed or supported.

5471 **25.5.3 502 Bad Gateway**

5472 The server, while acting as a gateway or proxy, received an invalid response from the downstream server it
5473 accessed in attempting to fulfill the request.

5474 **25.5.4 503 Service Unavailable**

5475 The server is temporarily unable to process the request due to a temporary overloading or maintenance of
5476 the server. The server MAY indicate when the client should retry the request in a **Retry-After** header. If no
5477 **Retry-After** is given, the client MUST act as if it had received a 500 (Server Internal Error) response.

5478 A client (proxy or UAC) receiving a 503 (Service Unavailable) SHOULD attempt to forward the request
5479 to an alternate server. It SHOULD NOT forward any other requests to that server for the duration specified in
5480 the **Retry-After** header field, if present.

5481 Servers MAY refuse the connection or drop the request instead of responding with 503 (Service Unavail-
5482 able).

5483 **25.5.5 504 Server Time-out**

5484 The server did not receive a timely response from an external server it accessed in attempting to process the
5485 request. 408 (Request Timeout) should be used instead if there was no response within the period specified
5486 in the **Expires** header field from the upstream server.

5487 **25.5.6 505 Version Not Supported**

5488 The server does not support, or refuses to support, the SIP protocol version that was used in the request. The
5489 server is indicating that it is unable or unwilling to complete the request using the same major version as the

5490 client, other than with this error message.

5491 **25.5.7 513 Message Too Large**

5492 The server was unable to process the request since the message length exceeded its capabilities.

5493 **25.6 Global Failures 6xx**

5494 6xx responses indicate that a server has definitive information about a particular user, not just the particular
5495 instance indicated in the Request-URI.

5496 **25.6.1 600 Busy Everywhere**

5497 The callee's end system was contacted successfully but the callee is busy and does not wish to take the call
5498 at this time. The response MAY indicate a better time to call in the **Retry-After** header. If the callee does
5499 not wish to reveal the reason for declining the call, the callee uses status code 603 (Decline) instead. This
5500 status response is returned only if the client knows that no other end point (such as a voice mail system) will
5501 answer the request. Otherwise, 486 (Busy Here) should be returned.

5502 **25.6.2 603 Decline**

5503 The callee's machine was successfully contacted but the user explicitly does not wish to or cannot partic-
5504 ipate. The response MAY indicate a better time to call in the **Retry-After** header. This status response is
5505 returned only if the client knows that no other end point will answer the request.

5506 **25.6.3 604 Does Not Exist Anywhere**

5507 The server has authoritative information that the user indicated in the Request-URI does not exist anywhere.

5508 **25.6.4 606 Not Acceptable**

5509 The user's agent was contacted successfully but some aspects of the session description such as the requested
5510 media, bandwidth, or addressing style were not acceptable.

5511 A 606 (Not Acceptable) response means that the user wishes to communicate, but cannot adequately
5512 support the session described. The 606 (Not Acceptable) response MAY contain a list of reasons in a **Warn-**
5513 **ing** header field describing why the session described cannot be supported. A message body containing a
5514 description of media capabilities MAY be present in the response, which is formatted according to the **Ac-**
5515 **cept** header field in the INVITE (or **application/sdp** if not present), the same as a message body in a 200
5516 (OK) response to an OPTIONS request. Reasons are listed in Section 24.45. It is hoped that negotiation
5517 will not frequently be needed, and when a new user is being invited to join an already existing conference,
5518 negotiation may not be possible. It is up to the invitation initiator to decide whether or not to act on a 606
5519 (Not Acceptable) response. This status response is returned only if the client knows that no other end point
5520 will answer the request.

5521 26 Examples

5522 In the following examples, we often omit the message body and the corresponding Content-Length and
5523 Content-Type headers for brevity.

5524 26.1 Registration

5525 Bob registers on start-up. The message flow is shown in Figure 9.

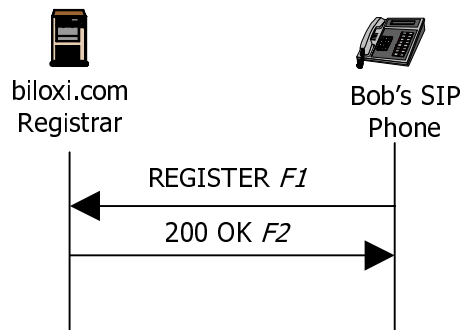


Figure 9: SIP Registration Example

5526

5527 F1 REGISTER Bob -> Registrar

5528

```

5529 REGISTER sip:registrar.biloxi.com SIP/2.0
5530 Via: SIP/2.0/UDP 192.0.2.4:5060;branch=z9hG4bKnashds7
5531 To: Bob <sip:bob@biloxi.com>
5532 From: Bob <sip:bob@biloxi.com>;tag=456248
5533 Call-ID: 843817637684230@998sdasdh09
5534 CSeq: 1826 REGISTER
5535 Contact: <sip:bob@192.0.2.4>
5536 Expires: 7200
5537 Content-Length: 0
  
```

5538 The registration expires after two hours. The registrar responds with a 200 OK:

5539

5540 F2 200 OK Registrar -> Bob

5541

```

5542 SIP/2.0 200 OK
5543 Via: SIP/2.0/UDP 192.0.2.4:5060;branch=z9hG4bKnashds7
5544 To: Bob <sip:bob@biloxi.com>
5545 From: Bob <sip:bob@biloxi.com>;tag=456248
  
```

5546 Call-ID: 843817637684230@998sdasdh09
5547 CSeq: 1826 REGISTER
5548 Contact: <sip:bob@192.0.2.4>
5549 Expires: 7200
5550 Content-Length: 0
5551

5552 26.2 Session Setup

5553 This example contains the full details of the example session setup in Section 4. The message flow is shown
5554 in Figure 1.

5555
5556 F1 INVITE Alice -> atlanta.com proxy
5557
5558 INVITE sip:bob@biloxi.com SIP/2.0
5559 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
5560 To: Bob <sip:bob@biloxi.com>
5561 From: Alice <sip:alice@atlanta.com>;tag=1928301774
5562 Call-ID: a84b4c76e66710
5563 CSeq: 314159 INVITE
5564 Contact: <sip:alice@pc33.atlanta.com>
5565 Content-Type: application/sdp
5566 Content-Length: 142
5567
5568 (Alice's SDP not shown)

5569
5570 F2 100 Trying atlanta.com proxy -> Alice
5571
5572 SIP/2.0 100 Trying
5573 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
5574 To: Bob <sip:bob@biloxi.com>
5575 From: Alice <sip:alice@atlanta.com>;tag=1928301774
5576 Call-ID: a84b4c76e66710
5577 CSeq: 314159 INVITE
5578 Content-Length: 0

5579
5580 F3 INVITE atlanta.com proxy -> biloxi.com proxy
5581
5582 INVITE sip:bob@biloxi.com SIP/2.0

5583 Via: SIP/2.0/UDP bigbox3.site3.atlanta.com;branch=z9hG4bK77ef4c2312983.1
5584 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
5585 To: Bob <sip:bob@biloxi.com>
5586 From: Alice <sip:alice@atlanta.com>;tag=1928301774
5587 Call-ID: a84b4c76e66710
5588 CSeq: 314159 INVITE
5589 Contact: <sip:alice@pc33.atlanta.com>
5590 Content-Type: application/sdp
5591 Content-Length: 142
5592
5593 (Alice's SDP not shown)

5594
5595 F4 100 Trying biloxi.com proxy -> atlanta.com proxy
5596
5597 SIP/2.0 100 Trying
5598 Via: SIP/2.0/UDP bigbox3.site3.atlanta.com;branch=z9hG4bK77ef4c2312983.1
5599 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
5600 To: Bob <sip:bob@biloxi.com>
5601 From: Alice <sip:alice@atlanta.com>;tag=1928301774
5602 Call-ID: a84b4c76e66710
5603 CSeq: 314159 INVITE
5604 Content-Length: 0

5605
5606 F5 INVITE biloxi.com proxy -> Bob
5607
5608 INVITE sip:bob@192.0.2.4 SIP/2.0
5609 Via: SIP/2.0/UDP server10.biloxi.com;branch=z9hG4bK4b43c2ff8.1
5610 Via: SIP/2.0/UDP bigbox3.site3.atlanta.com;branch=z9hG4bK77ef4c2312983.1
5611 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
5612 To: Bob <sip:bob@biloxi.com>
5613 From: Alice <sip:alice@atlanta.com>;tag=1928301774
5614 Call-ID: a84b4c76e66710
5615 CSeq: 314159 INVITE
5616 Contact: <sip:alice@pc33.atlanta.com>
5617 Content-Type: application/sdp
5618 Content-Length: 142
5619
5620 (Alice's SDP not shown)

5621

5622 F6 180 Ringing Bob -> biloxi.com proxy
5623
5624 SIP/2.0 180 Ringing
5625 Via: SIP/2.0/UDP server10.biloxi.com;branch=z9hG4bK4b43c2ff8.1
5626 Via: SIP/2.0/UDP bigbox3.site3.atlanta.com;branch=z9hG4bK77ef4c2312983.1
5627 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
5628 To: Bob <sip:bob@biloxi.com>;tag=a6c85cf
5629 From: Alice <sip:alice@atlanta.com>;tag=1928301774
5630 Call-ID: a84b4c76e66710
5631 CSeq: 314159 INVITE
5632 Content-Length: 0

5633
5634 F7 180 Ringing biloxi.com proxy -> atlanta.com proxy
5635
5636 SIP/2.0 180 Ringing
5637 Via: SIP/2.0/UDP bigbox3.site3.atlanta.com;branch=z9hG4bK77ef4c2312983.1
5638 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
5639 To: Bob <sip:bob@biloxi.com>;tag=a6c85cf
5640 From: Alice <sip:alice@atlanta.com>;tag=1928301774
5641 Call-ID: a84b4c76e66710
5642 CSeq: 314159 INVITE
5643 Content-Length: 0

5644
5645 F8 180 Ringing atlanta.com proxy -> Alice
5646
5647 SIP/2.0 180 Ringing
5648 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
5649 To: Bob <sip:bob@biloxi.com>;tag=a6c85cf
5650 From: Alice <sip:alice@atlanta.com>;tag=1928301774
5651 Call-ID: a84b4c76e66710
5652 CSeq: 314159 INVITE
5653 Content-Length: 0

5654
5655 F9 200 OK Bob -> biloxi.com proxy
5656
5657 SIP/2.0 200 OK
5658 Via: SIP/2.0/UDP server10.biloxi.com;branch=z9hG4bK4b43c2ff8.1
5659 Via: SIP/2.0/UDP bigbox3.site3.atlanta.com;branch=z9hG4bK77ef4c2312983.1
5660 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8

5661 To: Bob <sip:bob@biloxi.com>;tag=a6c85cf
5662 From: Alice <sip:alice@atlanta.com>;tag=1928301774
5663 Call-ID: a84b4c76e66710
5664 CSeq: 314159 INVITE
5665 Contact: <sip:bob@192.0.2.4>
5666 Content-Type: application/sdp
5667 Content-Length: 131
5668
5669 (Bob's SDP not shown)

5670
5671 F10 200 OK biloxi.com proxy -> atlanta.com proxy
5672
5673 SIP/2.0 200 OK
5674 Via: SIP/2.0/UDP bigbox3.site3.atlanta.com;branch=z9hG4bK77ef4c2312983.1
5675 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
5676 To: Bob <sip:bob@biloxi.com>;tag=a6c85cf
5677 From: Alice <sip:alice@atlanta.com>;tag=1928301774
5678 Call-ID: a84b4c76e66710
5679 CSeq: 314159 INVITE
5680 Contact: <sip:bob@192.0.2.4>
5681 Content-Type: application/sdp
5682 Content-Length: 131
5683
5684 (Bob's SDP not shown)

5685
5686 F11 200 OK atlanta.com proxy -> Alice
5687
5688 SIP/2.0 200 OK
5689 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds8
5690 To: Bob <sip:bob@biloxi.com>;tag=a6c85cf
5691 From: Alice <sip:alice@atlanta.com>;tag=1928301774
5692 Call-ID: a84b4c76e66710
5693 CSeq: 314159 INVITE
5694 Contact: <sip:bob@192.0.2.4>
5695 Content-Type: application/sdp
5696 Content-Length: 131
5697
5698 (Bob's SDP not shown)

5699

5700 F12 ACK Alice -> Bob

5701

5702 ACK sip:bob@192.0.2.4 SIP/2.0

5703 Via: SIP/2.0/UDP pc33.atlanta.com;branch=z9hG4bKnashds9

5704 To: Bob <sip:bob@biloxi.com>;tag=a6c85cf

5705 From: Alice <sip:alice@atlanta.com>;tag=1928301774

5706 Call-ID: a84b4c76e66710

5707 CSeq: 314159 ACK

5708 Content-Length: 0

5709 The media session between Alice and Bob is now established.

5710 Bob hangs up first. Note that Bob's SIP phone maintains its own CSeq numbering space, which, in
5711 this example, begins with 231. Since Bob is making the request, the To and From URIs and tags have been
5712 swapped.

5713

5714 F13 BYE Bob -> Alice

5715

5716 BYE sip:alice@pc33.atlanta.com SIP/2.0

5717 Via: SIP/2.0/UDP 192.0.2.4;branch=z9hG4bKnashds10

5718 From: Bob <sip:bob@biloxi.com>;tag=a6c85cf

5719 To: Alice <sip:alice@atlanta.com>;tag=1928301774

5720 Call-ID: a84b4c76e66710

5721 CSeq: 231 BYE

5722 Content-Length: 0

5723

5724 F14 200 OK Alice -> Bob

5725

5726 SIP/2.0 200 OK

5727 Via: SIP/2.0/UDP 192.0.2.4;branch=z9hG4bKnashds10

5728 From: Bob <sip:bob@biloxi.com>;tag=a6c85cf

5729 To: Alice <sip:alice@atlanta.com>;tag=1928301774

5730 Call-ID: a84b4c76e66710

5731 CSeq: 231 BYE

5732 Content-Length: 0

5733 The SIP Call Flows document [39] contains further examples of SIP messages.

5734 ;; This buffer is for notes you don't want to save, and for Lisp evaluation. ;; If you want to create a file,
5735 first visit that file with C-x C-f, ;; then enter the text in that file's own buffer.

5736 **27 Augmented BNF for the SIP Protocol**

5737 All of the mechanisms specified in this document are described in both prose and an augmented Backus-
5738 Naur Form (BNF) similar to that used by RFC 2234 [40]. Implementors need to be familiar with the notation
5739 in order to understand this specification. The augmented BNF includes the following constructs:

5740 name = definition

5741 The name of a rule is simply the name itself (without any enclosing “<” and “>”) and is separated from
5742 its definition by the equal “=” character. White space is only significant in that the indentation of continua-
5743 tion lines indicates a rule definition that spans more than one line. Certain basic rules are in uppercase, such
5744 as SP, LWS, HT, CRLF, DIGIT, ALPHA, etc. Angle brackets are used within definitions to clarify the use
5745 of rule names.

5746 "literal"

5747 Quotation marks surround literal text. Unless stated otherwise, the text is case-insensitive.

5748 rule1 | rule2

5749 Elements separated by a bar (“|”) are alternatives, that is, “yes | no” will accept yes or no.

5750 (rule1 rule2)

5751 Elements enclosed in parentheses are treated as a single element. Thus, “(elem (foo | bar) elem)” allows the
5752 token sequences “elem foo elem” and “elem bar elem”.

5753 *rule

5754 The character “*” preceding an element indicates repetition. The full form is “< n >* < m >element”
5755 indicating at least < n > and at most < m > occurrences of element. Default values are 0 and infinity so
5756 that “*(element)” allows any number, including zero; “1*element” requires at least one; and “1*2element”
5757 allows one or two.

5758 [rule]

5759 Square brackets enclose optional elements; “[foo bar]” is equivalent to “1*(foo bar)”.

5760 N rule

5761 Specific repetition: “<n>(element)” is equivalent to “<n>* <n>(element)”; that is, exactly <n> occur-
5762 rences of (element). Thus 2DIGIT is a 2-digit number, and 3ALPHA is a string of three alphabetic charac-
5763 ters.

5764 ; comment

5765 A semi-colon, set off some distance to the right of rule text, starts a comment that continues to the end of
5766 line. This is a simple way of including useful notes in parallel with the specifications.

5767 27.1 Basic Rules

5768 The following rules are used throughout this specification to describe basic parsing constructs. The US-
5769 ASCII coded character set is defined by ANSI X3.4-1986.

```
OCTET    = %x00-ff ; any 8-bit sequence of data
CHAR     = %x00-7f ; any US-ASCII character (octets 0 - 127)
upalpha  = "A" | "B" | "C" | "D" | "E" | "F" | "G" | "H" | "I" |
          "J" | "K" | "L" | "M" | "N" | "O" | "P" | "Q" | "R" |
          "S" | "T" | "U" | "V" | "W" | "X" | "Y" | "Z"
lowalpha = "a" | "b" | "c" | "d" | "e" | "f" | "g" | "h" | "i" |
          "j" | "k" | "l" | "m" | "n" | "o" | "p" | "q" | "r" |
          "s" | "t" | "u" | "v" | "w" | "x" | "y" | "z"
alpha    = lowalpha | upalpha
DIGIT    = "0" | "1" | "2" | "3" | "4" | "5" | "6" | "7" |
          "8" | "9"
alphanum = alpha | DIGIT
CTL      = %x00-1f | %x7f ; (octets 0 – 31) and DEL (127)
CR       = %d13 ; US-ASCII CR, carriage return character
LF       = %d10 ; US-ASCII LF, line feed character
SP       = %d32 ; US-ASCII SP, space character
HT       = %d09 ; US-ASCII HT, horizontal tab character
5770 CRLF   = CR LF ; typically the end of a line
```

5771 The following are defined in RFC 2396 [9] for the SIP URI:

```
reserved = "," | "/" | "?" | "." | "@" | " " | ";" | "+"
          | "$" | ","
unreserved = alphanum | mark
mark       = "-" | "_" | "!" | "~" | "*" | "'"
          | "(" | ")"
5772 escaped = "%" hex hex
```

5773 SIP header field values can be folded onto multiple lines if the continuation line begins with a space or
5774 horizontal tab. All linear white space, including folding, has the same semantics as SP. A recipient MAY
5775 replace any linear white space with a single SP before interpreting the field value or forwarding the message
5776 downstream. This is intended to behave exactly as HTTP 1.1 as described in RFC2615 [12]. The SWS
5777 construct is similar to LWS but allows zero instances of space or tab

```
LWS = *( SP | HT ) [CRLF] 1*( SP | HT ) ; linear whitespace
5778 SWS = *( SP | HT ) [CRLF] *( SP | HT ) ; sep whitespace
```

5779 To separate the header name from the rest of value, a colon is used, which, by the above rule, allows
5780 whitespace before, but no line break, and whitespace after, including a linebreak. The HCOLON defines
5781 this construct.

5782 HCOLON = *(SP | HT) ":" SWS

5783 The TEXT-UTF8 rule is only used for descriptive field contents and values that are not intended to be
 5784 interpreted by the message parser. Words of *TEXT-UTF8 contain characters from the UTF-8 character
 5785 set (RFC 2279 [11]). The TEXT-UTF8-TRIM rule is used for descriptive field contents that are *not* quoted
 5786 strings, where leading and trailing LWS is not meaningful. In this regard, SIP differs from HTTP, which
 5787 uses the ISO 8859-1 character set.

TEXT-UTF8 = *(TEXT-UTF8char | LWS)
 TEXT-UTF8-TRIM = *TEXT-UTF8char>(*LWS TEXT-UTF8char)
 TEXT-UTF8char = %x21-7e | UTF8-NONASCII
 UTF8-NONASCII = %xc0-df 1UTF8-CONT
 | %xe0-ef 2UTF8-CONT
 | %xf0-f7 3UTF8-CONT
 | %xf8-fb 4UTF8-CONT
 | %xfc-fd 5UTF8-CONT
 5788 UTF8-CONT = %x80-bf

5789 A CRLF is allowed in the definition of TEXT-UTF8 only as part of a header field continuation. It is
 5790 expected that the folding LWS will be replaced with a single SP before interpretation of the TEXT-UTF8
 5791 value.

5792 Hexadecimal numeric characters are used in several protocol elements. Some elements (authentication)
 5793 force hex alphas to be lower case.

5794 LHEX = digit | "a" | "b" | "c" | "d" | "e" | "f"

5795 Others allow mixed upper and lower case

5796 hex = LHEX | "A" | "B" | "C" | "D" | "E" | "F"

5797 Many SIP header field values consist of words separated by LWS or special characters. Unless otherwise
 5798 stated, tokens are case-insensitive. These special characters *MUST* be in a quoted string to be used within a
 5799 parameter value. The word construct is used in Call-ID to allow most separators to be used.

token = 1*(alphanumeric | "-" | "." | "!" | "%" | "*" | "_" | "+" | "'" | "''" | "``")
 separators = "(" | ")" | "<" | ">" | "@" |
 ";" | "," | "." | "\" | "<"> |
 "/" | "[" | "]" | "?" | "=" |
 "{" | "}" | SP | HT
 word = 1*(alphanumeric | "-" | "." | "!" | "%" | "*" |
 | "_" | "+" | "'" | "''" | "``" |
 "(" | ")" | "<" | ">" |
 "." | "\" | "<"> |
 "/" | "[" | "]" | "?" |
 5800 "{" | "}" | SP | HT)

5801 When tokens are used or separators are used between elements, whitespace is often allowed before or
 5802 after these characters:

MINUS = SWS "-" SWS ; minus
 DOT = SWS "." SWS ; period
 PERCENT = SWS "%" SWS ; percent
 BANG = SWS "!" SWS ; exclamation
 PLUS = SWS "+" SWS ; plus
 STAR = SWS "*" SWS ; asterisk
 SLASH = SWS "/" SWS ; slash
 TILDE = SWS "~" SWS ; tilde
 EQUAL = SWS "=" SWS ; equal
 LPAREN = SWS "(" SWS ; left parenthesis
 RPAREN = SWS ")" SWS ; right parenthesis
 LANGLE = SWS "<" SWS ; left angle bracket
 RAQUOT = SWS ">" SWS ; right angle quote
 LAQUOT = SWS "<" SWS ; left angle quote
 RANGLE = SWS ">" SWS ; right angle bracket
 BAR = SWS "|" SWS ; vertical bar
 ATSIGN = SWS "@" SWS ; atsign
 COMMA = SWS "," SWS ; comma
 SEMI = SWS ";" SWS ; semicolon
 COLON = SWS ":" SWS ; colon
 DQUOT = SWS "<>" SWS ; double quotation mark
 LDQUOT = SWS "<" SWS ; open double quotation mark
 RDQUOT = SWS ">" SWS ; close double quotation mark
 LBRECK = SWS "{" SWS ; left square bracket
 RBRECK = SWS "}" SWS ; right square bracket

5803

5804 Comments can be included in some SIP header fields by surrounding the comment text with parentheses.
 5805 Comments are only allowed in fields containing "comment" as part of their field value definition. In all other
 5806 fields, parentheses are considered part of the field value.

comment = LPAREN *(ctext | quoted-pair)
 ; ctext includes all chars except left and right parens and backslash
 ctext = %x21-27 | %x2a-5b | %x5d-7e
 | LWS

5807

5808 A string of text is parsed as a single word if it is quoted using double-quote marks. In quoted strings,
 5809 quotation marks (") and backslashes (\) need to be escaped.

quoted-string = (SWS "<" *(qdtex | quoted-pair) ">")
 qdtex = LWS | %x21 | %x23-5b | %x5d-7e
 | UTF8-NONASCII

5810

5811 The backslash character ("") MAY be used as a single-character quoting mechanism only within quoted-
 5812 string and comment constructs. Unlike HTTP/1.1, the characters CR and LF cannot be escaped by this
 5813 mechanism to avoid conflict with line folding and header separation.

quoted-pair = "\" (%x00 - %x09 | %x0b | %x0c
 | %x0e - %x7f)

5814


```

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 | absoluteURI
absoluteURI = scheme COLON ( hier-part | opaque-part )
hier-part   = ( net-path | abs-path ) [ "?" query ]
net-path    = "//" authority [ abs-path ]
abs-path    = "/" path-segments
opaque-part = uric-no-slash *uric
uric        = reserved — unreserved — escaped
uric-no-slash = unreserved | escaped | "," | "?" | ":" | "@"
              | "&" | "=" | "+" | "$" | ";"
path-segments = segment *( "/" segment )
segment       = *pchar *( SEMI param )
param         = *pchar
pchar        = unreserved | escaped |
              "." | "@" | " " | "=" | "+" | "$" | ";"
scheme       = alpha *( alpha | digit | "+" | "-" | "." )
authority    = server | reg-name
server       = [ [ userinfo "@" ] hostport ]
reg-name     = 1*( unreserved | escaped | "$" | ";"
              | "," | ":" | "@" | "&" | "=" | "+" )
query        = *uric
SIP-Version  = "SIP/2.0"

```

message-header

= Accept
| Accept-Encoding
| Accept-Language
| Alert-Info
| Allow
| Authentication-Info
| Authorization
| Call-ID
| Call-Info
| Contact
| Content-Disposition
| Content-Encoding
| Content-Language
| Content-Length
| Content-Type
| CSeq
| Date
| Error-Info
| Expires
| From
| In-Reply-To
| Max-Forwards
| MIME-Version
| Min-Expires
| Organization
| Priority
| Proxy-Authenticate
| Proxy-Authorization
| Proxy-Require
| RAck
| Record-Route
| Reply-To
| Require
| Retry-After
| Route
| RSeq
| Server
| Subject
| Supported
| Timestamp
| To
| Unsupported
| User-Agent
| Via
| Warning
| WWW-Authenticate

Method = "INVITE" | "ACK" | "OPTIONS" | "BYE"
 | "CANCEL" | "REGISTER" | "PRACK"
 | extension-method

extension-method = token

option-tag = token

Response

= Status-Line
 *(message-header)
 CRLF
 [message-body]

5821

Status-Line = SIP-version SP Status-Code SP Reason-Phrase CRLF
 Status-Code

= Informational
 | Redirection
 | Success
 | Client-Error
 | Server-Error
 | Global-Failure
 | extension-code

extension-code = 3DIGIT

Reason-Phrase = *(reserved — unreserved — escaped — SP — HT)

5822

Informational

= "100" ; Trying
 | "180" ; Ringing
 | "181" ; Call Is Being Forwarded
 | "182" ; Queued
 | "183" ; Session Progress

5823

Success = "200" ; OK

5824

Redirection = "300" ; Multiple Choices
 | "301" ; Moved Permanently
 | "302" ; Moved Temporarily
 | "305" ; Use Proxy
 | "380" ; Alternative Service

5825

Client-Error = "400" ; Bad Request
| "401" ; Unauthorized
| "402" ; Payment Required
| "403" ; Forbidden
| "404" ; Not Found
| "405" ; Method Not Allowed
| "406" ; Not Acceptable
| "407" ; Proxy Authentication Required
| "408" ; Request Timeout
| "409" ; Conflict
| "410" ; Gone
| "413" ; Request Entity Too Large
| "414" ; Request-URI Too Large
| "415" ; Unsupported Media Type
| "416" ; Unsupported URI Scheme
| "420" ; Bad Extension
| "423" ; Registration 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

5826

Server-Error = "500" ; Internal Server Error
| "501" ; Not Implemented
| "502" ; Bad Gateway
| "503" ; Service Unavailable
| "504" ; Server Time-out
| "505" ; SIP Version not supported

5827

Global-Failure = "600" ; Busy Everywhere
| "603" ; Decline
| "604" ; Does not exist anywhere
| "606" ; Not Acceptable

5828

Accept = "Accept" HCOLON (accept-range *(COMMA accept-range))
 accept-range = media-range [accept-params]
 media-range = ("*"/*"
 | (m-type SWS "/" "*" SWS)
 | (m-type SLASH m-subtype)
) *(SEMI parameter)
 accept-params = SEMI "q" EQUAL qvalue *(accept-extension)
 accept-extension = SEMI ae-name [EQUAL ae-value]
 ae-name = token
 5829 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
 5830 qvalue = ("0" ["." 0*3DIGIT])
 | ("1" ["." 0*3("0")])

Accept-Language = "Accept-Language" HCOLON (language *(COMMA language))
 language = language-range [SEMI "q" EQUAL qvalue]
 5831 language-range = ((1*8ALPHA *(MINUS 1*8ALPHA)) | "*")

Alert-Info = "Alert-Info" HCOLON alert-param *(COMMA alert-param)
 alert-param = LAQUOT URI RAQUOT *(SEMI generic-param)
 generic-param = token [EQUAL gen-value]
 5832 gen-value = token | host | quoted-string

5833 Allow = "Allow" HCOLON Method *(COMMA Method)

Authorization = "Authorization" HCOLON credentials
 credentials = ("Digest" digest-response) | (token gen-resp)
 digest-response = dig-resp *(COMMA dig-resp)
 dig-resp = username | realm | nonce | digest-uri
 | dresponse | [algorithm] | [cnonce]
 | [opaque] | [message-qop]
 | [nonce-count] | [auth-param]
 username = "username" EQUAL username-value
 username-value = quoted-string
 digest-uri = "uri" EQUAL digest-uri-value
 digest-uri-value = request-uri ; As specified by HTTP/1.1
 message-qop = "qop" EQUAL qop-value
 cnonce = "cnonce" EQUAL cnonce-value
 cnonce-value = nonce-value
 nonce-count = "nc" EQUAL nc-value
 nc-value = 8LHEX
 dresponse = "response" EQUAL request-digest
 request-digest = LDQUOT 32LHEX RDQUOT
 auth-param = auth-param-name EQUAL (token | quoted-string)
 auth-param-name = token
 gen-resp = *token *((COMMA *token) | (EQUAL
 5834 (*token | quoted-string))

AuthenticationInfo = "Authentication-Info" COLON ainfo *(COMMA ainfo)
 ainfo = nextnonce | [message-qop]
 | [response-auth] | [cnonce]
 | [nonce-count]
 nextnonce = "nextnonce" EQUAL nonce-value
 response-auth = "rspauth" EQUAL response-digest
 5835 response-digest = LDQUOT *LHEX RDQUOT

Call-ID = ("Call-ID" | "i") HCOLON callid
 5836 callid = word ["@" word]

Call-Info = "Call-Info" HCOLON info *(COMMA info)
 info = LAQUOT URI RAQUOT *(SEMI info-param)
 info-param = "purpose" EQUAL ("icon" | "info"
 5837 | "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 | URI
 5838 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
 qvalue = ("0" ["." 0*3DIGIT])
 5839 | ("1" ["." 0*3("0")])

5840 delta-seconds = 1*DIGIT

Content-Disposition = "Content-Disposition" HCOLON
 disposition-type *(SEMI disposition-param)
 disposition-type = "render" | "session" | "icon" | "alert"
 | disp-extension-token
 disposition-param = "handling" EQUAL
 ("optional" | "required"
 | other-handling) | generic-param
 other-handling = token
 5841 disp-extension-token = token

5842 Content-Encoding = ("Content-Encoding" | "e") HCOLON
 content-coding *(COMMA content-coding)

Content-Language = "Content-Language" HCOLON
 language-tag *(COMMA language-tag)
 language-tag = primary-tag *(MINUS subtag)
 primary-tag = 1*8ALPHA
 5843 subtag = 1*8ALPHA

5844 Content-Length = ("Content-Length" | "l") HCOLON 1*DIGIT

Content-Type = ("Content-Type" | "c") HCOLON media-type
 media-type = m-type SLASH m-subtype *(SEMI m-parameter)
 m-type = discrete-type | composite-type
 discrete-type = "text" | "image" | "audio" | "video"
 | "application" | extension-token
 composite-type = "message" | "multipart" | extension-token
 extension-token = ietf-token | x-token
 ietf-token = token
 x-token = ("X" | "x") "-" token
 m-subtype = extension-token | iana-token
 iana-token = token
 m-parameter = m-attribute EQUAL m-value
 m-attribute = token
 5845 m-value = token | quoted-string

5846 CSeq = "CSeq" HCOLON 1*DIGIT LWS Method

Date = "Date" HCOLON SIP-date
SIP-date = rfc1123-date
rfc1123-date = wkday COMMA date1 SP time SP "GMT"
date1 = 2DIGIT SP month SP 4DIGIT
; day month year (e.g., 02 Jun 1982)
time = 2DIGIT ":" 2DIGIT ":" 2DIGIT
; 00:00:00 - 23:59:59
wkday = "Mon" | "Tue" | "Wed"
| "Thu" | "Fri" | "Sat" | "Sun"
month = "Jan" | "Feb" | "Mar" | "Apr"
| "May" | "Jun" | "Jul" | "Aug"
5847 | "Sep" | "Oct" | "Nov" | "Dec"

Error-Info = "Error-Info" HCOLON error-uri *(COMMA error-uri)
5848 error-uri = LAQUOT URI 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
5849 tag-param = "tag" EQUAL token

5850 In-Reply-To = "In-Reply-To" HCOLON called *(COMMA called)

5851 Max-Forwards = "Max-Forwards" HCOLON 1*DIGIT

5852 MIME-Version = "MIME-Version" HCOLON 1*DIGIT ":" 1*DIGIT

5853 Min-Expires = "Min-Expires" HCOLON delta-seconds

5854 Organization = "Organization" HCOLON TEXT-UTF8-TRIM

Priority = "Priority" HCOLON priority-value
priority-value = "emergency" | "urgent" | "normal"
| "non-urgent" | other-priority
5855 other-priority = token

Proxy-Authenticate = "Proxy-Authenticate" HCOLON
 challenge *(COMMA challenge)
 challenge = "Digest" digest-challenge
 digest-challenge = digest-chlng *(COMMA digest-chlng)
 digest-chlng = realm | [domain] | nonce
 | [opaque] | [stale] | [algorithm]
 | [qop-options] | [auth-param]
 realm = "realm" EQUALS realm-value
 realm-value = quoted-string
 domain = "domain" EQUAL LDQUOT URI
 (1*SP URI) RDQUOT
 URI = absoluteURI | abs_path
 nonce = "nonce" EQUAL nonce-value
 nonce-value = quoted-string
 opaque = "opaque" EQUAL quoted-string
 stale = "stale" EQUAL ("true" | "false")
 algorithm = "algorithm" EQUAL ("MD5" | "MD5-sess"
 | token)
 qop-options = "qop" EQUAL LDQUOT qop-value *(COMMA qop-value) RDQUOT
 qop-value = "auth" | "auth-int" | token
 5856
 5857 Proxy-Authorization = "Proxy-Authorization" HCOLON credentials
 5858 Proxy-Require = "Proxy-Require" HCOLON option-tag *(COMMA option-tag)
 RAck = "RAck" HCOLON response-num LWS CSeq-num LWS Method
 response-num = 1*DIGIT
 CSeq-num = 1*DIGIT
 5859 response-num = 1*DIGIT
 Record-Route = "Record-Route" HCOLON rec-route *(COMMA rec-route)
 rec-route = name-addr *(SEMI rr-param)
 5860 rr-param = generic-param
 Reply-To = ("Reply-To" | "f") HCOLON rplyto-spec
 rplyto-spec = (name-addr | addr-spec)
 *(SEMI rplyto-param)
 rplyto-param = generic-param
 5861 Require = "Require" HCOLON option-tag *(COMMA option-tag)
 Retry-After = "Retry-After" HCOLON delta-seconds
 [comment] *(SEMI retry-param)
 5862 retry-param = "duration" EQUAL delta-seconds
 | generic-param
 Route = "Route" HCOLON route=param *(COMMA route-param)
 5863 route-param = name-addr *(SEMI rr-param)

5864 RSeq = "RSeq" HCOLON response-num

Server = "Server" HCOLON 1*(product | comment)
product = token [SLASH product-version]
5865 product-version = token

5866 Subject = ("Subject" | "s") HCOLON TEXT-UTF8-TRIM

Supported = ("Supported" | "k") HCOLON
5867 (option-tag *(COMMA option-tag)

Timestamp = "Timestamp" HCOLON 1*(DIGIT
["." *(DIGIT)] [delay]
5868 delay = *(DIGIT) ["." *(DIGIT)]

To = ("To" | "t") HCOLON (name-addr
| addr-spec) *(SEMI to-param)
5869 to-param = tag-param | generic-param

5870 Unsupported = "Unsupported" HCOLON option-tag *(COMMA option-tag)

5871 User-Agent = "User-Agent" HCOLON 1*(product | comment)

Via = ("Via" | "v") HCOLON via-param *(COMMA via-param)
via-param = sent-protocol 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]
5872 ttl = 1*3DIGIT ; 0 to 255

Warning = "Warning" HCOLON warning-value *(COMMA warning-value)
warning-value = warn-code SP warn-agent SP warn-text
warn-code = 3DIGIT
warn-agent = (host [COLON port]) | pseudonym
; the name or pseudonym of the server adding
; the Warning header, for use in debugging
warn-text = quoted-string
pseudonym = token

WWW-Authenticate = "WWW-Authenticate" HCOLON challenge

message-body = *OCTET

28 IANA Considerations

All new or experimental method names, header field names, and status codes used in SIP applications SHOULD be registered with IANA in order to prevent potential naming conflicts. It is RECOMMENDED that new "option-tag"s and "warn-code"s also be registered. Before IANA registration, new protocol elements SHOULD be described in an Internet-Draft or, preferably, an RFC.

For Internet-Drafts, IANA is requested to make the draft available as part of the registration database.

By the time an RFC is published, colliding names may have already been implemented.

When a registration for either a new header field, new method, or new status code is created based on an Internet-Draft, and that Internet-Draft becomes an RFC, the person that performed the registration MUST notify IANA to change the registration to point to the RFC instead of the Internet-Draft.

Registrations should be sent to iana@iana.org.

28.1 Option Tags

Option tags are used in header fields such as Require, Supported, Proxy-Require, and Unsupported in support of SIP compatibility mechanisms for extensions (Section 23.2). The option tag itself is a string that is associated with a particular SIP option (that is, an extension). It identifies the option to SIP endpoints.

When registering a new SIP option with IANA, the following information MUST be provided:

- Name and description of option. The name MAY be of any length, but SHOULD be no more than twenty characters long. The name MUST consist of alphanum (Section 27) characters only.
- A listing of any new SIP header fields, header parameter fields, or parameter values defined by this option. A SIP option MUST NOT redefine header fields or parameters defined in either RFC 2543, any standards-track extensions to RFC 2543, or other extensions registered through IANA.
- Indication of who has change control over the option (for example, IETF, ISO, ITU-T, other international standardization bodies, a consortium, or a particular company or group of companies).
- A reference to a further description if available, for example (in order of preference) an RFC, a published paper, a patent filing, a technical report, documented source code, or a computer manual.
- Contact information (postal and email address).

This procedure has been borrowed from RTSP [3] and the RTP AVP [41].

5903 **28.1.1 Registration of 100rel**

5904 This specification registers a single option tag, "100rel". The required information is:

5905 **Name:** "100rel"

5906 **Description:** This option tag is for reliability of provisional responses. When present in a **Supported**
5907 header, it indicates that the UA can send or receive reliable provisional responses. When present in a
5908 **Require** header in a request, it indicates that the UAS **MUST** send all provisional responses reliably.
5909 When present in a **Require** header in a reliable provisional response, it indicates that the response is
5910 to be sent reliably.

5911 **New Headers:** The RSeq and RACK header fields are defined by this option.

5912 **Change Control:** IETF.

5913 **Reference:** RFCXXXX [Note to IANA: Fill in with the RFC number of this specification.]

5914 **Contact Information:** Jonathan Rosenberg, jdrosen@jdrosen.net. 72 Eagle Rock Avenue, First Floor, East
5915 Hanover, NJ, 07936, USA.

5916 **28.2 Warn-Codes**

5917 Warning codes provide information supplemental to the status code in SIP response messages when the
5918 failure of the transaction results from a Session Description Protocol (SDP, [5]). New "warn-code" values
5919 can be registered with IANA as they arise.

5920 The "warn-code" consists of three digits. A first digit of "3" indicates warnings specific to SIP.

5921 Warnings 300 through 329 are reserved for indicating problems with keywords in the session description,
5922 330 through 339 are warnings related to basic network services requested in the session description, 370
5923 through 379 are warnings related to quantitative QoS parameters requested in the session description, and
5924 390 through 399 are miscellaneous warnings that do not fall into one of the above categories.

5925 1xx and 2xx have been taken by HTTP/1.1.

5926 **28.3 Header Field Names**

5927 Header field names do not require working group or working group chair review prior to IANA registration,
5928 but **SHOULD** be documented in an RFC or Internet-Draft before IANA is consulted.

5929 The following information needs to be provided to IANA in order to register a new header field name:

- 5930 ● The name and email address of the individual performing the registration;
- 5931 ● the name of the header field being registered;
- 5932 ● a compact form version for that header field, if one is defined;
- 5933 ● the name of the draft or RFC where the header field is defined;
- 5934 ● a copy of the draft or RFC where the header field is defined.

5935 Header fields SHOULD NOT use the X- prefix notation and MUST NOT duplicate the names of header
5936 fields used by SMTP or HTTP unless the syntax is a compatible superset and the semantics are similar.
5937 Some common and widely used header fields MAY be assigned one-letter compact forms (Section 7.3.3).
5938 Compact forms can only be assigned after SIP working group review. In the absence of this working group,
5939 a designated expert reviews the request.

5940 **28.4 Method and Response Codes**

5941 Because the status code space is limited, they do require working group or working group chair review, and
5942 MUST be documented in an RFC or Internet draft. The same procedures apply to new method names.

5943 The following information needs to be provided to IANA in order to register a new response code or
5944 method:

- 5945 • The name and email address of the individual performing the registration;
- 5946 • the number of the response code or name of the method being registered;
- 5947 • the default reason phrase for that status code, if applicable;
- 5948 • the name of the draft or RFC where the method or status code is defined;
- 5949 • a copy of the draft or RFC where the method or status code is defined.

5950 **29 Changes Made in Version 00**

- 5951 • Indicated that UAC should send both CANCEL and BYE after a retransmission fails.
- 5952 • Added semicolon and question mark to the list of unreserved characters for the user part of SIP URLs
5953 to handle tel: URLs properly.
- 5954 • Uniform handling of if hop count Max-Forwards: return 483. Note that this differs from HTTP/1.1
5955 behavior, where only OPTIONS and TRACE allow this header, but respond as the final recipient
5956 when the value reaches zero.
- 5957 • Clarified that a forking proxy sends ACKs only for INVITE requests.
- 5958 • Clarified wording of DNS caching. Added paragraph on “negative caching”, i.e., what to do if one
5959 of the hosts failed. It is probably not a good idea to simply drop this host from the list if the DNS ttl
5960 value is more than a few minutes, since that would mean that load balancing may not work for quite a
5961 while after a server is brought back on line. This will be true in particular if a server group receives a
5962 large number of requests from a small number of upstream servers, as is likely to be the case for calls
5963 between major consumer ISPs. However, without getting into arbitrary and complicated retry rules, it
5964 seems hard to specify any general algorithm. Might it be worthwhile to simply limit the “black list”
5965 interval to a few minutes?
- 5966 • Added optional Call-Info and Alert-Info header fields that describe the caller and information to be
5967 used in alerting. (Currently, avoided use of “purpose” qualification since it is not yet clear whether
5968 rendering content without understanding its meaning is always appropriate. For example, if a UAS

- 5969 does not understand that this header is to replace ringing, it would mix both local ring tone and the
5970 indicated sound URL.) TBD!
- 5971 ● SDP “s=” lines can’t be empty, unfortunately.
 - 5972 ● Noted that **maddr** could also contain a unicast address, but **SHOULD** contain the multicast address if
5973 the request is sent via multicast (Section 24.44.
 - 5974 ● Clarified that responses are sent to port in **Via sent-by** value.
 - 5975 ● Added “other-*” to the **user** URL parameter and the **Hide** and **Content-Disposition** headers.
 - 5976 ● Clarified generation of timeout (408) responses in forking proxies and mention the **Expires** header.
 - 5977 ● Clarified that **CANCEL** and **INVITE** are separate transactions (Fig. 7). Thus, the **INVITE** request
5978 generates a 487 (Request Terminated) if a **CANCEL** or **BYE** arrives.
 - 5979 ● Clarified that **Record-Route** **SHOULD** be inserted in every request, but that the route, once estab-
5980 lished, persists. This provides robustness if the called UAS crashes.
 - 5981 ● Emphasized that proxy, redirect, registrar and location servers are logical, not physical entities and
5982 that UAC and UAS roles are defined on a request-by-request basis. (Section 6)
 - 5983 ● In Section 24.44, noted that the **maddr** and **received** parameters also need to be encrypted when
5984 doing **Via** hiding.
 - 5985 ● Simplified Fig. 7 to only show **INVITE** transaction.
 - 5986 ● Added definition of the use of **Contact** (Section 24.10) for **OPTIONS**.
 - 5987 ● Added HTTP/RFC 822 headers **Content-Language** and **MIME-Version**.
 - 5988 ● Added note in minimal section indicating that UAs need to support UDP.
 - 5989 ● Added explanation explaining what a UA should do when receiving an initial **INVITE** with a tag.
 - 5990 ● Clarified UA and proxy behavior for 302 responses.
 - 5991 ● Added details on what a UAS should do when receiving a tagged **INVITE** request for an unknown call
5992 leg. This could occur if the UAS had crashed and the UAC sends a re-**INVITE** or if the **BYE** got lost
5993 and the UAC still believes to be in the call.
 - 5994 ● Added definition of **Contact** in 4xx, 5xx and 6xx to “redirect” to more error details.
 - 5995 ● Added note to forking proxy description to gather ***-Authenticate** from responses. This allows several
5996 branches to be authenticated simultaneously.
 - 5997 ● Changed URI syntax to use URL escaping instead of quotation marks.
 - 5998 ● Changed SIP URL definition to reference RFC 2806 for telephone-subscriber part.

- 5999 • Clarified that the **To** URI should basically be ignored by the receiving UAS except for matching
6000 requests to call legs. In particular, **To** headers with a scheme or name unknown to the callee should
6001 be accepted.
- 6002 • Clarified that **maddr** is to be added by any client, either proxy or UAC.
- 6003 • Added response code 488 to indicate that there was no common media at the particular destination.
6004 (606 indicates such failure globally.)
- 6005 • In Section 24.19, noted that registration updates can shorten the validity period.
- 6006 • Added note to enclose the URI for digest in quotation marks. The BNF in RFC 2617 is in error.
- 6007 • Clarified that registrars use **Authorization** and **WWW-Authenticate**, not proxy authentication.
- 6008 • Added note in Section 24.10 that “headers” are copied from **Contact** into the new request.
- 6009 • Changed URL syntax so that port specifications have to have at least one digit, in line with other URL
6010 formats such as “http”. Previously, an empty port number was permissible.
- 6011 • In SDP section, added a section on how to add and delete streams in re-INVITEs.
- 6012 • IETF-blessed extensions now have short names, without **org.ietf.** prefix.
- 6013 • **Cseq** is unique within a call leg, not just within a call (Section 24.16).
- 6014 • Added IPv6 literal addresses to the SIP URL definition, according to RFC 2732 [42]. Modified the
6015 IPv4 address to limit segments to at most three digits.
- 6016 • Modified registration procedure so that it explicitly references the URL comparison. Updates with
6017 shorter expiration time are now allowed.
- 6018 • For send-only media, SDP still must indicate the address and port, since these are needed as destina-
6019 tions for RTCP messages.
- 6020 • Changed references regarding DNS SRV records from RFC 2052 to RFC 2782, which is now a Pro-
6021 posed Standard. Integrated SRV into the search procedure and removed the SRV appendix. The only
6022 visible change is that protocol and service names are now prefixed by an underscore. Added wording
6023 that incorporates the precedence of **maddr**.
- 6024 • Allow parameters in **Record-Route** and **Route** headers.
- 6025 • In Table 1, list **udp** as the default value for the transport parameter in SIP URI.
- 6026 • Removed sentence that **From** can be encrypted. It cannot, since the header is needed for call-leg
6027 identification.
- 6028 • Added note that a UAC only copies a **To** tag into subsequent transactions if it arrives in a 200 OK to
6029 an INVITE. This avoids the problem that occurs when requests get resubmitted after receiving, say,
6030 a 407 (or possibly 500, 503, 504, 305, 400, 411, 413, maybe even 408). Under the old rules, these
6031 requests would have a tag, which would force the called UAS to reject the request, since it doesn't
6032 have an entry for this tag.

- 6033 • Loop detection has been modified to take the request-URI into account. This allows the same request
6034 to visit the server twice, but with different request URIs (“spiral”).
- 6035 • Elaborated on URL comparison and comparison of From/To fields.
- 6036 • Added np-queried user parameter.
- 6037 • Changed tag syntax from UUID to token, since there’s no reason to restrict it to hex.
- 6038 • Added Content-Disposition header based on earlier discussions about labeling what to do with a
6039 message body (part).
- 6040 • Clarification: proxies must insert To tags for locally generated responses.
- 6041 • Clarification: multicast may be used for subsequent registrations.
- 6042 • Feature: Added Supported header. Needed if client wants to indicate things the server can usefully
6043 return in the response.
- 6044 • Bug: The From, To, and Via headers were missing extension parameters. The Encryption and
6045 Response-Key header fields now “officially” allow parameters consisting only of a token, rather
6046 than just “token = value”.
- 6047 • Bug: Allow was listed as optional in 405 responses in Table 2. It is mandatory.
- 6048 • Added: “A BYE request from either called or calling party terminates any pending INVITE, but the
6049 INVITE request transaction MUST be completed with a final response.”
- 6050 • Clarified: “If an INVITE request for an existing session fails, the session description agreed upon in
6051 the last successful INVITE transaction remains in force.”
- 6052 • Clarified what happens if two INVITE requests meet each other on the wire, either traveling the same
6053 or in opposite directions:

6054 A UAC MUST NOT issue another INVITE request for the same call leg before the pre-
6055 vious transaction has completed. A UAS that receives an INVITE before it sent the final
6056 response to an INVITE with a lower CSeq number MUST return a 400 (Bad Request)
6057 response and MUST include a Retry-After header field with a randomly chosen value of
6058 between 0 and 10 seconds. A UA that receives an INVITE while it has an INVITE transac-
6059 tion pending, returns a 500 (Internal Server Error) and also includes a Retry-After header
6060 field.
- 6061 • Expires header clarified: limits only duration of INVITE transaction, not the actual session. SDP
6062 does the latter.
- 6063 • The In-Reply-To header was added.
- 6064 • There were two incompatible BNFs for WWW-Authenticate. One defined for PGP, and the other
6065 borrowed from HTTP. For basic or digest:

6066 WWW-Authenticate: basic realm="Wallyworld"

6067 and for pgp:

6068 WWW-Authenticate: pgp; realm="Wallyworld"

6069 The latter is incorrect and the semicolon has been removed.

- 6070 • Added rules for **Route** construction from called to calling UA.
- 6071 • We now allow **Accept** and **Accept-Encoding** in **BYE** and **CANCEL** requests. There is no particular
6072 reason not to allow them, as both requests could theoretically return responses, particularly when
6073 interworking with other signaling systems.
- 6074 • PGP “pgp-pubalgorithm” allows server to request the desired public-key algorithm.
- 6075 • ABNF rules now describe tokens explicitly rather than by subtraction; explicit character enumeration
6076 for CTL, etc.
- 6077 • Registrars should be careful to check the **Date** header as the expiration time may well be in the past,
6078 as seen by the client.
- 6079 • **Content-Length** is mandatory; Table 2 erroneously marked it as optional.
- 6080 • **User-Agent** was classified in a syntax definition as a request header rather than a general header.
- 6081 • Clarified ordering of items to be signed and include realm in list.
- 6082 • Allow **Record-Route** in 401 and 484 responses.
- 6083 • Hop-by-hop headers need to precede end-to-end headers only if authentication is used.
- 6084 • 1xx message bodies MAY now contain session descriptions.
- 6085 • Changed references to HTTP/1.1 and authentication to point to the latest RFCs.
- 6086 • Added 487 (Request terminated) status response. It is issued if the original request was terminated
6087 via **CANCEL** or **BYE**.
- 6088 • The spec was not clear on the identification of a call leg. Section 1.3 says it's the combination of **To**,
6089 **From**, and **Call-ID**. However, requests from the callee to the caller have the **To** and **From** reversed, so
6090 this definition is not quite accurate. Additionally, the “tag” field should be included in the definition
6091 of call leg. The spec now says that a call leg is defined as the combination of local-address, remote-
6092 address, and call-id, where these addresses include tags.
- 6093 Text was added to Section 6.21 to emphasize that the **From** and **To** headers designate the originator
6094 of the request, not that of the call leg.
- 6095 • All URI parameters, except **method**, are allowed in a **Request-URI**. Consequently, also updated the
6096 description of which parameters are copied from 3xx responses in Sec. 24.10.

- 6097 • The use of CRLF, CR, or LF to terminate lines was confusing. Basically, each header line can be
6098 terminated by a CR, LF, or CRLF. Furthermore, the end of the headers is signified by a “double
6099 return”. Simplified to require sending of CRLF, but require senders to receive CR and LF as well and
6100 only allow CR CR, LF LF in addition to double CRLF as a header-body separator.
- 6101 • Round brackets in **Contact** header were part of the HTTP legacy, and very hard to implement. They
6102 are also not that useful and were removed.
- 6103 • The spec said that a proxy is a back-to-back UAS/UAC. This is almost, but not quite, true. For
6104 example, a UAS should insert a tag into a provisional response, but a proxy should not. This was
6105 clarified.
- 6106 • Section 6.13 in the RFC begins mid-paragraph after the BNF. The following text was misplaced in the
6107 conversion to ASCII:
6108 Even if the “display-name” is empty, the “name-addr” form **MUST** be used if the “addr-
6109 spec” contains a comma, semicolon or question mark.

6110 **30 Changes Made in Version 01**

- 6111 • Uniform syntax specification for semicolon parameters:

Foo	=	"Foo" ":" something *(";" foo-param)
foo-param	=	"bar" "=" token
		generic-param
- 6112
- 6113 • Removed np-queried user parameter since this is now part of a tel URL extension parameter.
- 6114 • In SDP section, noted that if the capabilities intersection is empty, a dummy format list still has to be
6115 returned due to SDP syntax constraints. Previously, the text had required that no formats be listed.
6116 (Brian Rosen)
- 6117 • Reorganized tables 2 and 3 to show proxy interaction with headers rather than “end-to-end” or “hop-
6118 by-hop”.

6119 **31 Changes Made in Version 02**

- 6120 • Added “or UAS” in description of received headers in Section 24.44. This makes the response
6121 algorithm work even if the last IP address in the **Via** is incorrect.
- 6122 • Tentatively removed restriction that **CANCEL** requests cannot have **Route** headers. (Billy Biggs)
- 6123 • Tentatively added **Also** header for **BYE** requests, as it is widely implemented and a simple means to
6124 implement unsupervised call transfer. Subject to removal if there is protest. (Billy Biggs)
- 6125 • If a proxy sends a request by UDP (TCP), the spec did not disallow placing TCP (UDP) in the transport
6126 parameter of the **Via** field, which it should. Added a note that the transport protocol actually used is
6127 included.

- 6128 • No default value for the **q** parameter in Contact is defined. This is not strictly needed, but is useful for
6129 consistent behaviors at recursive proxies and at UAC's. Now 0.5.
- 6130 • Clarified that **To** and **From** tag values should be different to simplify request matching when calling
6131 oneself.
- 6132 • Removed ability to carry multiple requests in a single UDP packet (Section 24.14).
- 6133 • Added note that **Allow** **MAY** be included in requests, to indicate requestor capabilities for the same
6134 call ID.
- 6135 • Added note to Section 24.17 indicating that registrars **MUST** include the **Date** header to accomodate
6136 UAs that do not have a notion of absolute time.
- 6137 • Added note emphasizing that non-SIP URIs are permissible in REGISTER.
- 6138 • Rewrote the server lookup section to be more precise and more like pseudo-code, with nesting instead
6139 of "gotos".
- 6140 • Removed note

6141 Note that the two URLs example.com and example.com:5060, while considered equal,
6142 may not lead to the same server, as the former causes a DNS SRV lookup, while the latter
6143 only uses the A record.

6144 since that is no longer the case.

- 6145 • Emphasized that proxies have to forward requests with unknown methods.
- 6146 • Aligned definition of call leg with URI comparison rules.
- 6147 • Required that second branch parameter be globally unique, so that a proxy can distinguish different
6148 branches in spiral scenarios similar to the following, with record-routing in place:

```
6149           B  ---> P1  -----> P2  -----> P1  -----> A
6150 BYE B    B/1      P1/2,B/1    P2/3,P1/2,B/1    P1/4,P2/3,P1/2,B/1
```

6151 Here, A/1 denotes the **Via** entry with host A and branch parameter 1. Also, this requires updating the
6152 definition of isomorphic requests, since the **Request-URI** is the same for all **BYE** that are record-
6153 routed.

- 6154 • Removed **Via** hiding from spec, for the following reasons:
 - 6155 – complexity, particularly hidden "gotchas" that surface at various points (as in this instance);
 - 6156 – interference with loop detection and debugging;
 - 6157 – Unlike HTTP, where via-hiding makes sense since all data is contained in the request or re-
6158 sponse, **Via**-hiding in SIP by itself does nothing to hide the caller or callee, as address informa-
6159 tion is revealed in a number of places:
 - 6160 * **Contact**;

- 6161 * Route/Record-Route;
6162 * SDP, including the o= and c= lines;
6163 * possibly accidental leakage in User-Agent header and Call-ID headers.
- 6164 – Unless this is implemented everywhere, the feature is not likely to be very useful, without the
6165 sender having any recourse such as “don’t route this request unless you can hide”. It appears
6166 that almost all existing proxies simply ignore the Hide header.
- 6167 • Added Error-Info header field.

6168 **32 Changes Made in Version 03**

- 6169 • Description of Route and Record-Route moved to separate section, which is new. All UAs must
6170 now support this mechanism.
- 6171 • Removed status code 411, since it cannot occur (Jonathan Rosenberg, James Jack).
- 6172 • Rewrote Record-Route section to reflect new mechanism. In particular, requests from callee to caller
6173 now use the same path as in the opposite direction, without substituting the From header field values.
6174 The maddr parameter is now optional.
- 6175 • Disallowed SIP URLs that only have a password, without a user name. The prototype from RFC 1738
6176 also doesn’t allow this.
- 6177 • Allow registrar to set the expiration time.
- 6178 • CSeq (Section 24.16) is counted within a call leg, not a call.
- 6179 • Removed wording that connection closing is equivalent to CANCEL or 500. This does not work for
6180 connections that are used for multiple transactions and has other problems.
- 6181 • Cleaned up CSeq section. Removed text about inserting CSeq method when it is absent. Clarified
6182 that CSeq increments for all requests, not just INVITE. Clarified that all out of order requests, not
6183 just out of order INVITE, are rejected with a 400 class response. Clarified the meaning of “initial”
6184 sequence number. Clarified that after a request forks, each 200 OK is a separate call leg, and thus,
6185 separate CSeq space. Clarified that CSeq numbers are independent for each direction of a call leg.
- 6186 • Massive reorganization and cleanup of the SDP section. Introduced the concept of the offer-answer
6187 model. Clarified that set of codecs in m line are usable all at the same time. Inserted size restriction
6188 on representation of values in o line. Explicitly describe forked media. New media lines for adding
6189 streams appear at the bottom of the SDP (used to say append).
- 6190 • Removed Also.
- 6191 • Added text to Require and Proxy-Require sections, making it a SHOULD to retry the request without
6192 the unsupported extension.
- 6193 • Added text to section on 415, saying that UAC SHOULD retry the request without the unsupported
6194 body.

- 6195 ● Added text to section on CANCEL and ACK, clarifying much of the behavior.
- 6196 ● Modified Content-Type to indicate that it can be present even if the body is empty.
- 6197 ● From tags mandatory
- 6198 ● Old text said that if you hang up before sending an ACK, you need not send the ACK. That is wrong.
6199 Text fixed so that an ACK is always sent.
- 6200 ● Old text said that if you never got a response to an INVITE, the UAC should send both an INVITE and
6201 CANCEL. This doesn't make sense. Rahter, it should do nothing and consider the call terminated.
- 6202 ● Added text that says pending requests are responded to with a 487 if a BYE is received.
- 6203 ● Updated section 2.2, so that its clear that Contact is not used with BYE.
- 6204 ● Clarified Via processing rules. Added text on handling loops when proxies route on headers besides
6205 the request URI. Added text on handling case when sent-by contains a domain name. Added text to
6206 6.47 on opening TCP connections to send responses upstream.
- 6207 ● Clarified that a 1xx with an unknown xx is not the same as the 100 response.
- 6208 ● Removed usage of Retry-After in REGISTER.
- 6209 ● Clarified usage of persistent connections.
- 6210 ● Clarified that servers supporting HTTP basic or digest in rfc2617 MUST be backwards compatible
6211 with RFC 2069.
- 6212 ● Clarified that ACK contains the same branch ID as the request its acknowledging.
- 6213 ● Added definitions for spiral, B2BUA.
- 6214 ● Rephrased definitions for UAC, UAS, Call, call-leg, caller, callee, making them more concrete.
- 6215 ● URL comparison ignores parameters not present in both URLs only for unknown parameters.
- 6216 ● Clarified that * in Contact is used only in REGISTER with Expires header zero. Mentioned * case
6217 in section on Contact syntax.
- 6218 ● Removed text that says a UA can insert a Contact in 2xx that indicates the address of a proxy. Not
6219 likely to work in general.
- 6220 ● Removed SDP text about aligning media streams within a media type to handle certain crash and
6221 restart cases.
- 6222 ● Receiving a 481 to a mid-call request terminates that call leg. Agreed upon at IETF 49.
- 6223 ● Introduced definition of regular transaction - non-INVITE excepting ACK and CANCEL.
- 6224 ● Clarified rules for overlapping transactions.

- 6225 • Forking proxies **MUST** be stateful (used to say **SHOULD**). Proxies that send requests on multicast
6226 **MUST** be stateful (used to say nothing)
- 6227 • Text added recommending that registrars authorize that entity in **From** field can register address-of-
6228 record in the **To** field.
- 6229 • Forwarding of non-100 provisionals upstream in a proxy changed from **SHOULD** to **MUST**.
- 6230 • Removed PGP.

6231 **33 Changes Made in Version 04**

- 6232 • Removed **Unsupported** as a request header from Table 3.
- 6233 • Clarified SDP procedures for changing IP address and port. Specifically, spelled out the duration for
6234 which a UA needs to received media on the old port and address.
- 6235 • Added text in the SDP session which recommends that the answerer use the same ordering of codecs
6236 as used on the offer, in order to help ensure symmetric codec operation under normal conditions.
- 6237 • Fixed bug in the example in the SDP section, where the new media line was listed at the top. Should
6238 have been the bottom.
- 6239 • **Authorization** credentials are cached based on the URL of the **To** header, not the entire **To** header as
6240 10.48 implied.
- 6241 • Section 10.31, on **Proxy-Authenticate**, indicated that a server responds with a 401 if the client
6242 guessed wrong. This is incorrect. It should be 407.
- 6243 • Section 10.14, removed motivational text about **Contact** allowing an **INVITE** to be routed directly
6244 between end systems, since its confusing. Some have interpreted to mean that **Record-Route** is
6245 ignored when **Contact** is present.
- 6246 • Added reference to SCTP RFC.
- 6247 • Updated 2.2 to allow non-SIP URLs in **OPTIONS** and 2xx to **OPTIONS**.
- 6248 • Fixed example in 20.5. Added **ACK** for 487, and added **To** tag to 487 response.
- 6249 • Clarified further URL comparisons. Its only URL parameters without defaults that are ignored if not
6250 present in both URLs.
- 6251 • Section 1.5.2, UDP mandatory for all. TCP is a **SHOULD** for UA, **MUST** for proxy, registrar, redirect
6252 servers.
- 6253 • Brought syntax for **Contact**, **Via**, and the SIP URL into alignment between the text and postscript
6254 versions.
- 6255 • Updated the text in section 6 which said that the ordering of header fields follows HTTP, with the
6256 exception of **Via**, where order matters. However, the HTTP spec says that order matters, so this
6257 sentence is redundant and confusing. The sentence was removed.

- 6258 • Added e lines to SDP examples in the Examples section.
- 6259 • Rewrote Allow discussion, more formally defining its semantics and usage cases.
- 6260 • Updated text on 604 status, to indicate that its based on the Request-URI, not the To.
- 6261 • Added response registrations to IANA considerations. Provided more details on registration process.
- 6262 • Clarified that only a UAS rejects a request because the To tag doesn't match a local value.
- 6263 • Clarified that stateless proxies need to route based on static criteria only.
- 6264 • Proxy and UAC CANCEL generation upon 2xx, 6xx if it forked is now a SHOULD; used to be a MAY.
- 6265 • Added text saying that a UAS SHOULD send a BYE if it never gets an ACK for a 2xx establishing a
6266 call leg.
- 6267 • Added text saying that a UAS SHOULD send a re-INVITE if it never gets an ACK for a 2xx to a
6268 re-INVITE.
- 6269 • Added text on 503 processing, indicating that a client should try a different server when receiving a
6270 503, and that a proxy shouldn't forward a 503 upstream unless it can't service any other requests.
- 6271 • Removed motivational text in Section 10.43 on Via headers since its not consistent with the text before
6272 it.
- 6273 • Changed IPSec reference to RFC 2401, from RFC 1825.
- 6274 • Updated retransmission definition in 17.3.4 to be consistent with the rest of the spec.
- 6275 • Softened the language for insertion of the transport param in the record-route. Specifically, it can be
6276 inserted in private networks where it is known apriori that the specific transport is supported.
- 6277 • Updated definition of B2BUA.
- 6278 • Added text to section on 420 processing, which mandates that the client retry the request without
6279 extensions listed in the Unsupported header in the response.
- 6280 • Allow Authentication-Info header to be used for HTTP digest.

6281 **34 Changes Made in Version 05**

- 6282 • Updated Table 2 to reflect that Error-Info is a response header in 3xx-6xx responses (it was previously
6283 listed as a request header).
- 6284 • Removed WWW-Authenticate as a request header from Table 3. Authentication of responses is now
6285 done according to RFC 2617.

- 6286 • Updated the **Accept**, **Accept-Encoding** and **Accept-Language** sections. More details on precise
6287 semantics for the various requests and responses is now provided. Presence of these headers is now
6288 a **SHOULD** for **INVITE** and **2xx** to **INVITE** when a non-default value is present. Extra emphasis is
6289 placed on including the **Accept-Language** in **INVITE** and **2xx** in order to support internationaliza-
6290 tion. Usage of these three headers in **CANCEL** has been removed since it makes no sense.
- 6291 • Generalized local outbound processing rules in Section 16.4.1 to cover the case where the UAS is
6292 using a local outbound proxy which was not in the initial call setup path.
- 6293 • Updated record-routing section, so that a proxy can insert a transport param if it knows that the proxy
6294 on one side supports the specific transport (the previous text required the proxy to know whether the
6295 proxies on both sides supported the specific transport).
- 6296 • Added **Authentication-Info** to Section 10.
- 6297 • Clarified the meaning of Table 2 for responses.
- 6298 • Updated Table 1 to reflect that **maddr** is no longer mandatory in **Record-Route**.
- 6299 • Updated Table 3 so that header fields in responses to **ACK** are never listed as optional, mandatory, etc.
6300 - only not applicable. This is because responses to **ACK** are not allowed. Also improved wording in
6301 Section 5.1.1 to clarify that there **MUST NOT** be responses to **ACK**.
- 6302 • Updated **SRV** procedures. Old text said to treat a failure to contact a server as a **4xx**, which would
6303 stop the **SRV** processing. But, this is not so. Sentence was stricken.
- 6304 • Updated 12.1 to clarify that **2xx INVITE** responses **MUST** contain session descriptions.
- 6305 • Changed **User-Agent** to a request header in Table 3.
- 6306 • Updated **SDP** section, so that a UA cannot change the **SDP** when it gets a re-**INVITE** with no **SDP**.
- 6307 • Clarified Appendix B that a unicast offer **MUST** have a unicast response.
- 6308 • Clarified that any request can be record-routed, but it may not be used by the UA, depending on the
6309 method.
- 6310 • non-**2xx** responses to **INVITE** no longer retransmitted over **TCP**.
- 6311 • Removed lower bound on **T1** and **T2** in private networks, which can use lower values. Furthermore,
6312 **T1** can be smaller on the public Internet if proper **RTT** estimation is used.
- 6313 • UAS Cannot send a **BYE** for a call leg until it receives **ACK**, in order to eliminate a race condition
6314 between **BYE** and **200 OK**.
- 6315 • Support of **CR** or **LF** alone as line terminators, as opposed to **CRLF**, is no longer required.
- 6316 • Client behavior on receipt of a **3xx** to re-**INVITE** is now specified, and it is no longer forbidden to
6317 generate a **3xx**. This is needed to maintain the idempotency of **INVITE**, as a proxy might redirect
6318 without knowing its a **3xx**.

- 6319 • CANCEL cannot be sent before a 1xx is received, in order to eliminate race condition between request
6320 and CANCEL.
- 6321 • Termination of the client and server transactions is now based entirely on timeouts, rather than re-
6322 transmission counters, in order to unify TCP and UDP behavior. Timeout values scale as a function
6323 of the RTT estimate, defined as T1. For reliable transports, many of these timers are now set to zero.
6324 Many timeouts differ than in bis-04.
- 6325 • Added a working RTT estimation algorithm using the Timestamp header, and specified it to be
6326 compliant to RFC 2988.
- 6327 • UAS accepting requests with unknown schemes in the URI in the To field is now a RECOMMENDED
6328 instead of SHOULD. This reflects the fact that processing a request when the To field doesn't match is
6329 a matter of policy.
- 6330 • Bodies are now allowed in any request and response, including CANCEL, although there may not be
6331 any semantics associated with that.
- 6332 • Supporting of INVITE without SDP is now a MUST (no strength was previously specified).
- 6333 • Registration procedures for visiting, which had a few sentences in bis-04, have been removed. Roam-
6334 ing is a complex issue, and should be treated elsewhere.
- 6335 • Bis-04 mandated that a 2xx response to REGISTER contain expires Contact parameters indicating
6336 the expiration time of a contact. This behavior has now been made consistent with requests, so that
6337 the expiration time of a contact is the same in either case: the expires param is used first if present,
6338 then the Expires header if present, else one hour for SIP URLs.
- 6339 • Action parameter in contact registrations is deprecated.
- 6340 • 2xx to REGISTER MUST contain current contacts. This was just a SHOULD in bis-04.
- 6341 • Multicast operation radically changed. Now, the treatment is no different than unicast. That is, only
6342 the first non-1xx response to a multicast request will be used. This is a natural consequence of the
6343 layering now applied to the protocol. This still enables anycast types of functions, mirroring the real
6344 usage of registrar discovery.
- 6345 • To completely separate transport rules from transaction rules, the rule in bis-04 that said a UAC
6346 SHOULD keep a connection opened until a response is received, has been turned into a timer recom-
6347 mendation. Specifically, the spec now says that it is RECOMMENDED that connections be kept opened
6348 for a minimum interval of sufficient duration to guarantee, with high probability, that responses are
6349 sent over the same connections as a request.
- 6350 • Re-use of existing connections for new requests to the same address and port is now RECOMMENDED,
6351 it was only a MAY in bis-04.
- 6352 • Modification of headers below the Authorization header by proxies is no longer disallowed, since the
6353 only mechanism that used Authorization in that way, PGP, has been deprecated previously.
- 6354 • Authentication of registrations now RECOMMENDED; no strength was defined previously.

- 6355 • Registering of new headers with IANA is now SHOULD; no strength was defined previously.
- 6356 • Proxy aggregation of challenges now a SHOULD; no strength was defined previously.
- 6357 • Server support of basic authentication downgraded from SHOULD to MAY.
- 6358 • UAC resubmitting requests with credentials after a challenge upgraded from MAY to SHOULD.
- 6359 • TLS is now RECOMMENDED as the transport layer security for SIP signaling.
- 6360 • UA recursion on a redirect is now SHOULD; no strength was assigned previously.
- 6361 • UA reuse of headers in a recursed request is now SHOULD; no strength was assigned previously.
- 6362 • Security considerations added for Call-Info and Alert-Info.
- 6363 • Proxies no longer forward a 6xx immediately on receiving it. Instead, they CANCEL pending
6364 branches immediately. This avoids a potential race condition that would result in a UAC getting a
6365 6xx followed by a 2xx. In all cases except this race condition, the result will be the same - the 6xx is
6366 forwarded upstream.
- 6367 • The term call-leg has been eliminated from the spec; a more generic term, dialog, is used in its place.
- 6368 • For SRV processing, subsequent requests with the same Call-ID (as opposed to the same transaction
6369 in bis-04) are sent to the same server.
- 6370 • SRV processing generalized to deal with the fact that the default port is transport dependent.
- 6371 • Per IESG request, draft-ietf-sip-serverfeatures has been integrated into bis.
- 6372 • Per IESG request, draft-ietf-sip-100rel will be integrated into bis. This is marked with a placeholder
6373 in this draft.
- 6374 • The BNF has been converted from implicit LWS to explicit LWS.
- 6375 • Caching of responses in a proxy to avoid redoing location server lookups used to be a SHOULD.
6376 Caching behavior for responses is now fully encapsulated in the transaction processing.
- 6377 • Proxy usage of SRV in processing Route headers upgraded from SHOULD to MUST.

6378 **35 Changes Made in Version 06**

- 6379 • Made TCP mandatory for user agents.
- 6380 • The two states of a dialog are now called early and confirmed.
- 6381 • CANCEL requests now carry Route header fields.
- 6382 • Changes section in -05 forgot to mention the removal of the Encryption and Response-Key headers.
6383 These were removed since the only mechanism that used them, PGP, had already been deprecated. As
6384 such, they were effectively “garbage collected”.

- 6385 • Updated error in transaction definition. ACK-2xx is a separate transaction, ACK for non-2xx is part
6386 of the same transaction.
- 6387 • Changed Contact-Length typo to Content-Length in various sections, including throughout the
6388 examples section.
- 6389 • Changed Table 3 entry for Record-Route and Route for REGISTER from "o" for optional to "-"
6390 for Not Allowed.
- 6391 • Changed Table 3 entry for Route for ACK, BYE, CANCEL, INVITE, and OPTIONS from "o" for
6392 optional to "c" for conditional, depending on whether a route set has been defined for the dialog or
6393 the response code.
- 6394 • Updated Figure 5; adding missing label on "calling" to "completed" transition.
- 6395 • Fixed errored transport example from Section 19.2.1.
- 6396 • Clarified that 17.2.3 and 17.1.3 are rules that define retransmissions.
- 6397 • fixed reported bugs in bnf (missing productions, bad tex markup), etc. Added new SWS production
6398 to have an LWS which allows zero spaces, and used that With any separators. Removed the # rule.
- 6399 • ACK for non-2xx has to have the same Route as the request its acknowledging. The text formerly
6400 said that the ACK MUST NOT contain Route, this has now radically changed to MUST have Route if
6401 the request its cancelling had one.
- 6402 • Clarified that stateless proxies apply Route processing logic to CANCEL requests.
- 6403 • Emphasized that escaping in the hostname portion of SIP URIs is not currently allowed.
- 6404 • Added discussion on when configuration changes affect the ability of a proxy to forward requests
6405 stateful or statelessly.
- 6406 • Explicitly stated that a proxy may add a Record-Route header field value to any request.
- 6407 • Added discussion on the use of To tags in hop-hop responses at a proxy.
- 6408 • Relaxed text concerning proxies forwarding CANCELs when a matching response context can't be
6409 found to allow the CANCEL to be processed statefully.
- 6410 • Changed references to "short" form of SIP headers to "compact" form.
- 6411 • Changed Date example to a valid date.
- 6412 • Clarified how ACK gets from transport to UAS core.
- 6413 • Adding missing "SIP/2.0" to first REGISTER in the examples section.
- 6414 • Fixed bug in 17.2.3 which said that an ACK matched a server transaction if the CSeq method (not
6415 number) matched that of the INVITE. It should be the reverse; number, not method.
- 6416 • Fixed bug in 22.15 where it said Content-Length instead of Content-Type.

- 6417 ● Incorporated draft-ietf-sip-100rel-04 into bis.
- 6418 ● Reliability of provisional responses now only defined for provisional responses to INVITE, although
6419 extension methods can allow its usage. This is because PRACK needs to be sent within the context
6420 of a dialog, and only responses to INVITE establish dialogs.
- 6421 ● Can no longer send a reliable provisional response after a final response; its not compatible with the
6422 transaction machines, which generally assume no provisionals after a final.
- 6423 ● Proxy behavior for reliable provisional responses no longer defined separately; the spec states that it
6424 simply acts as a uas.
- 6425 ● Scope of Record-Route header fields for a reliable provisional response is now the dialog rather than
6426 the particular request.
- 6427 ● Example PRACK flows were lost when incorporating into bis.
- 6428 ● Formal IANA registration of “100rel” option tag.
- 6429 ● If reliable provisional response gets no PRACK after 32*T1, UAS sends 5xx to original request.
- 6430 ● Recommended UA behavior for caching credentials.
- 6431 ● Included guidelines for devices presenting pre-configured credentials vs. prompting end users to
6432 provide credentials for a specific realm.
- 6433 ● Added section on stateless UAS Behavior, clarifying secure handling of unauthenticated requests to
6434 prevent potential DoS threat.
- 6435 ● Provided motivation for aggregation of challenges in the Security Considerations, and made the be-
6436 havioral language there more specific.
- 6437 ● Provided guidelines for the construction of realm strings for authentication.
- 6438 ● Changed concept of protection domain for SIP so that it is no longer defined by both a Request-URI
6439 and a realm; it is now only defined by a realm.
- 6440 ● Put in some text encouraging UACs not to resubmit rejected credentials when re-challenged.
- 6441 ● Added falsification of source IP address to the Via denial of service attack case.
- 6442 ● Provided canonical MD5 hash for an empty message body to be used in Digest integrity calculation.
- 6443 ● Added security considerations for the CANCEL and ACK methods.
- 6444 ● Deprecated and removed Basic auth scheme. Proxies MUST NOT accept or request Basic.
- 6445 ● Strengthened language regarding the sending of the “qop” parameter; receipt of cnonce is based on
6446 “qop”.
- 6447 ● Clarified the construction the URI in the Request-URI of REGISTER requests.

- 6448 ● Noted that registrars SHOULD provide **Date** headers in 200 (OK) responses to REGISTER, and that
6449 clients can use these Dates to set their internal clocks.
- 6450 ● Processing of REGISTERs at a registrar now must be with atomicity and isolation.
- 6451 ● Registrars now MUST process **Require** headers.
- 6452 ● Clarified **CSeq** increment over REGISTER messages for the same **Call-ID**, and necessity of tracking
6453 **Call-IDs** and **CSeqs** for contact addresses by a registrar
- 6454 ● Added registrar-side handling for
 - 6455 `Contact: *`
 - 6456 `Expires: 0`
- 6457 ● Added description generalizing processing of **OPTIONS** responses to include proxies as well as
6458 UAS. Included language describing use of **Max-Forwards** as a SIP capabilities traceroute. Described
6459 construction of a **Request-URI** for an **OPTIONS** sent to a proxy.
- 6460 ● Defined “Not Applicable” in Tables 2 and 3 to mean that the header field is undefined and should be
6461 ignored if present.
- 6462 ● Removed old references to general headers in Table 3.
- 6463 ● Allowed a proxy to insert a **Max-Forwards** header field in Table 2. Also added description of the use
6464 of the header by elements that can not otherwise guarantee loop detection.
- 6465 ● Fixed dialog matching reference in 22.37.
- 6466 ● Reinforced that all 6xx, including 603 and 606, are only sent if the UAS knows that no other endpoint
6467 will accept the call.
- 6468 ● Clarified that for 302 responses, the **Contact** is used just once to recurse a new transaction, unless an
6469 **Expires** header or **expires** parameter is present.
- 6470 ● Clarified that 405 is sent when the server knows the method, but the method is not allowed for the
6471 resource in the **Request-URI**. 501 is sent when the server has never heard of the method at all.
- 6472 ● Included note that no MIME types for message bodies of 3xx responses have been defined.
- 6473 ● Stated explicitly in Section 22.10 on **Contact** the rules for parsing display names, URI and URI
6474 parameters, and header parameters. Referenced this text in the sections on **To** and **From** header fields.
- 6475 ● Corrected references in **Timestamp** section.
- 6476 ● Noted in **Via** section that the host or network address and port part of the header does not follow
6477 the SIP URI syntax; spaces around **:** are permitted. Also noted that spaces are permitted around **/**.
6478 Modified an example to show this.
- 6479 ● Added text to describe the **Contact** header fields in a 2xx response to an **OPTIONS** as having redirect
6480 semantics. Modified example to show both a SIP and mailto **Contact** URI.

- 6481 ● Added text to describe the use of **OPTIONS** within a dialog to query a peer for capabilities, and noted
6482 that the request has no impact on the dialog.
- 6483 ● Added text to 302 (Moved Temporarily) section saying that if a cached **Contact URI** fails, the request
6484 may be retried with the original **Request-URI**. Removed recursion rules (moved to **UA** section) and
6485 “call” specific language. Specifically stated both proxies and **Uas** may cache **URI** for expiration
6486 interval.
- 6487 ● Added text to 488/606 section to allow **SDP** message bodies, formatted the same as **SDP** in 200 (OK)
6488 responses to **OPTIONS**. Removed text on **SDP** response message bodies from the **Warning** section.
- 6489 ● Outbound server is now called outbound proxy
- 6490 ● Clarified that a transaction in the completed state is not “in progress” when it comes to overlapping
6491 transactions.
- 6492 ● 488 response is used to reject an offer.
- 6493 ● Clarified how to reject an offer.
- 6494 ● Clarified that requests with **To** tag outside a dialog may have been simply missrouted.
- 6495 ● General **UAS** behaviour applies to **CANCEL** and **BYE**
- 6496 ● Clarified when to use **BYE** to terminate an early dialog.
- 6497 ● Explained when a **UAS** detects gaps in the **Cseq** space.
- 6498 ● Specified behavior for inclusion of bodies in **ACK** for non-2xx; **MUST** be same type as request, or one
6499 of the types in **Accept** if the response was 415.
- 6500 ● Updated the default value of timer **D** to be 32s, instead of **T3**.
- 6501 ● Clarified that **RTT** estimate of **T1** applies to all requests and responses sent to that **IP** address, and
6502 included a discussion of how this is not quite right when there are stateless proxies in the path.
- 6503 ● 180 (Ringing) responses for re-**INVITE**s are not typically useful.
- 6504 ● **ACK**s **MUST** contain the same credentials as the **INVITE**.
- 6505 ● **ACK** for non-2xx responses needs to contain the same **Route** headers as the request. Same reason
6506 **CANCEL** needs to.
- 6507 ● Increased minimum timer for holding persistent connections, and clarified the reasoning behind the
6508 timer.
- 6509 ● Clarified that persistent connections are indexed by address, port, transport, and that ephemeral source
6510 ports imply that peering relationships will ususally involve two connections.
- 6511 ● Timer **T3** no longer used; it was a dangling reference in bis-05.
- 6512 ● Clarified Figure 7 to indicate that 100 is only sent if **TU** won’t respond in 200ms.

- 6513 • Re-added text that said proxies **MUST** and UA **SHOULD** support TCP, which somehow got accidentally
6514 deleted from bis-05.
- 6515 • Clarified meaning of an empty **Accept** header field.
- 6516 • Added RFC 2616 security warning about **Server** header field to both **Server** and **User-Agent** header
6517 fields.
- 6518 • Added handling of transport failures to transaction state machines, and added a section for server
6519 transactions.
- 6520 • Disallowed port in **To/From** header URIs.
- 6521 • Allowed password in both **To** and **From** header URIs.
- 6522 • Disallowed the **method** URI parameter in **REGISTER** and **Redirect Contact** header URIs.
- 6523 • Absolved proxies from issuing **CANCEL**s based on the **Expires** header of an **INVITE**. Included text
6524 pointing out that they **MAY** do so, but it is unnecessary.
- 6525 • Clarified aggregating authentication challenges at a proxy.
- 6526 • Added notice that even though proxies are required to **CANCEL** outstanding client transactions upon
6527 forwarding a final response, an endpoint may still receive multiple 200 (OK) responses to an **INVITE**.
6528 Also noted that future extensions could override the requirement to **CANCEL**.
- 6529 • Reinforced that proxies must wait for provisional responses before generating **CANCEL** requests.
- 6530 • Request merging moved to general Ua behaviour section.
- 6531 • Request processing is atomic.
- 6532 • Clarified how to resolve glare conditions.
- 6533 • Added UAs should ignore unknown extension header parameters.
- 6534 • Clarified when quoted string vs. token can be used as a display name.
- 6535 • Explicitly stated that a header parameter name can appear at most once per header field value.
- 6536 • Noted that proxies no longer treat merged requests as an error.
- 6537 • Clarified that proxies can **Record-Route** header field values to requests already in dialogs to improve
6538 robustness, but that chosing not to do so will not normally cause them to be removed from the path.
- 6539 • Clarified that proxies do not remove any received parameters they may have added to **Via** header fields
6540 when forwarding responses.
- 6541 • Deprecated absolute time in **Expires** and **Retry-After**.
- 6542 • Added pointer to what to do with responses that were meant for a proxy
- 6543 • Summarized stateful proxy forwarding behavior with respect to what final responses get forwarded

- 6544 ● Clarified that elements on the start line of messages are separated by a single SP character
- 6545 ● Explicitly stated that a SIP URI parameter name to occur at most once in a URI.
- 6546 ● Changed Table 2 to show Accept, Accept-Encoding, Accept-Language, and Supported as for a
6547 2xx to an OPTIONS as m*
- 6548 ● Changed Table 2 to show Content-Length as “t”, which is defined to mean that it should be present,
6549 but must be present if TCP is used.
- 6550 ● Added the notion that registrars that accept registrations on a multicast interface might want to redirect
6551 registrations to a unicast interface.
- 6552 ● Request merging now a behavior of the UA, rather than the proxy server.
- 6553 ● Solidified the circumstances under which UAs should retry rejected requests with the same Call-ID
6554 but a different CSeq.
- 6555 ● Corrected erroneous statement that contact addresses were not cached across dialogs; now dependent
6556 on status code and expiration interval.
- 6557 ● Tags are a MUST for non-100 provisionals, a MAY for 100 (Trying).
- 6558 ● Discouraged generation of 1xx responses to non-INVITE requests.
- 6559 ● Fixed references to Content- handling headers in the UA section.
- 6560 ● Timestamp headers must be copied from requests into a 100 Trying for RTT calculation.
- 6561 ● Request processing is now said to be atomic.
- 6562 ● Potential infinite redirection loop problem fixed; redirect servers MUST NOT send a redirect to the
6563 same URI they received in the redirected request.
- 6564 ● Further specified which URIs servers can expect to see in Request-URIs of requests (relationship to
6565 contact headers).
- 6566 ● Defined pre-loaded route headers.
- 6567 ● Clarified normative language of Accept-Encoding, Accept-Language, and Content-Disposition
6568 in regard to no header being present.
- 6569 ● Noted that “transport=TLS” in a SIP URI refers to TLS over TCP.
- 6570 ● Refined discussion on forming requests based on a given SIP URI.
- 6571 ● Clarified “matching the topmost Via” for stateless proxies.
- 6572 ● Added discussion of how proxies respond to transaction failure and notification of state-machine
6573 timeouts.
- 6574 ● Corrected description of proxy behavior when recursing on 3xx contacts to account for contacts not
6575 recursed on (such as contacts containing non-SIP URIs).

- 6576 • Added Reply-To header field.
- 6577 • Clarified that responses to OPTIONS are scoped to the Request-URI of the request.
- 6578 • Added 491 (Request Pending) response code.
- 6579 • Proxies should not remove malformed headers that it doesn't care about when forwarding requests.
- 6580 • Noted that proxies can't generate their own 1xx provisional responses, but they can use a virtual
6581 collocated UAS to achieve the same effect.
- 6582 • Two SIP URIs which are identical with the exception of the presence of an maddr parameter in one,
6583 and no maddr parameter in the other are not equivalent.
- 6584 • Modified transaction, UA, and proxy sections so that branch ID is now a unique transaction identifier.
6585 Updated all example messages so that UAC insert branch ID, and magic cookie is present in all branch
6586 ID values.
- 6587 • CANCELs and ACKs MUST NOT contain Require or Proxy-Require headers.
- 6588 • A UA SHOULD NOT send re-INVITE or BYE upon media failure.
- 6589 • Only SIP URIs can be used as addresses of record in REGISTER requests.
- 6590 • Registrars MUST NOT increase the expiration interval of registrations. Intervals that are too short MAY
6591 be rejected with a 423 w/ Min-Expires.
- 6592 • Security Considerations substantially reorganized and expanded.
- 6593 • TLS support for proxy servers, registrars and redirect servers now a MUST.
- 6594 • Minimum ciphersuite for TLS now AES.
- 6595 • S/MIME now slightly more implementable. S/MIME support is now a SHOULD for UAs.
- 6596 • S/MIME now relies on RFC 2633 CMS messages.
- 6597 • Threat models against the SIP protocol are now provided.
- 6598 • Example architectures in which security mechanisms might be used are described.
- 6599 • Limitations of security mechanisms are described.
- 6600 • Added 493 (Undecipherable) response code.
- 6601 • Fixed ACK column in Table 3 entry for Warning.
- 6602 • Added text describing how to recurse on a 3xx as a UAC.
- 6603 • SIP URIs are compared case-sensitive across the userpart, case-insensitive everywhere else.
- 6604 • Proxies strip transport and port when stripping maddr.

- 6605 • Port and transport apply to maddr when maddr is present in a SIP URI.
- 6606 • Restored record-route example from bis-04.
- 6607 • Reinforced that SIP messages MAY contain binary bodies or body parts.
- 6608 • Added section discussing conversion of tel URLs to SIP URIs, focusing on issues with maintaining
6609 equivalence.
- 6610 • Clarified use of transaction key in building values to include in Record-Route values.
- 6611 • Clarified requirements on the inclusion of information in the loop-detection hash used in branch pa-
6612 rameters.
- 6613 • Noted in the proxy section that Record-Route values are only valid within the scope of the dialog in
6614 which they are provided.
- 6615 • Added definitions for redirect server, recursion, header, message, request, response, and route refresh
6616 request.
- 6617 • Placing headers needed by proxies (Via, Route, Record-Route, etc.) at the top of messages is now
6618 RECOMMENDED.
- 6619 • Reinforced that proxies processing messages do not fork, even by recursing on returned 3xx responses.
- 6620 • Removed restriction on proxies adding Record-Route to REGISTER requests. Added that registrars
6621 ignore Record-Route if it occurs.
- 6622 • Allowed for loose-route policies, capturing use of default outbound proxies as a loose route decision.
- 6623 • The scope of Contact header fields is not limited to the dialog.
- 6624 • Added text saying that when the caller wishes to be anonymous, the URI should be scrambled as well.
- 6625 • Moved 485 response generation from UAS to proxy.
- 6626 • Require MUST only reference standards track RFCs.
- 6627 • Removed requirement on proxies to not forward a request to a multicast group that had already been
6628 visited.
- 6629 • Deprecated loop-detection. Made Max-Forwards mandatory with an initial value of 70. Proxies
6630 insert a Max-Forwards of 70 if they find the header missing.
- 6631 • Placed HTTP Digest and S/MIME in sections independent of the security Considerations.
- 6632 • Added 416 (Unsupported URI Scheme) and discussion on its handling. Added guidance on how a
6633 UAC would select the URI in the To/Request-URI based on user input.
- 6634 • Noted that BYE without tags is now rejected, which is a backwards compatibility break with RFC
6635 2543.
- 6636 • Reference offer-answer for formatting of SDP in OPTIONS response, 488, 606.

- 6637 • Timer C now managed by the TU. Proxies have a minimum of 3 minutes, but it is extended through
6638 provisional responses.
- 6639 • Proxies can go stateless mid-transaction if they didn't do anything that would have otherwise pre-
6640 vented them from being stateless in the first place.

6641 **36 Acknowledgments**

6642 We wish to thank the members of the IETF MMUSIC and SIP WGs for their comments and suggestions.
6643 Detailed comments were provided by Brian Bidulock, Jim Buller, Neil Deason, Dave Devanathan, Cdric
6644 Fluckiger, Yaron Goland, Bernie Hneisen, Phil Hoffer, Christian Huitema, Jean Jervis, Gadi Karmi, Pe-
6645 ter Kjellerstedt, Anders Kristensen, Jonathan Lennox, Gethin Liddell, Alison Mankin, Keith Moore, Vern
6646 Paxson, Moshe J. Sambol, Chip Sharp, Igor Slepchin, Eric Tremblay., and Rick Workman.

6647 Brian Rosen provided the compiled BNF.

6648 This work is based, inter alia, on [43, 44].

6649 **37 Authors' Addresses**

6650 Authors addresses are listed alphabetically for the editors, the writers, and then the original authors of RFC
6651 2543. All listed authors actively contributed large amounts of text to this document.

6652 Jonathan Rosenberg
6653 dynamicsoft
6654 72 Eagle Rock Ave
6655 East Hanover, NJ 07936
6656 USA
6657 electronic mail: jdrosen@dynamicsoft.com

6658 Henning Schulzrinne
6659 Dept. of Computer Science
6660 Columbia University
6661 1214 Amsterdam Avenue
6662 New York, NY 10027
6663 USA
6664 electronic mail: schulzrinne@cs.columbia.edu

6665 Gonzalo Camarillo
6666 Ericsson
6667 Advanced Signalling Research Lab.
6668 FIN-02420 Jorvas
6669 Finland
6670 electronic mail: Gonzalo.Camarillo@ericsson.com

6671 Alan Johnston
6672 WorldCom

6673 100 South 4th Street
6674 St. Louis, MO 63102
6675 USA
6676 electronic mail: alan.johnston@wcom.com

6677 Jon Peterson
6678 NeuStar, Inc
6679 1800 Sutter Street, Suite 570
6680 Concord, CA 94520
6681 USA
6682 electronic mail: jon.peterson@neustar.com

6683 Robert Sparks
6684 dynamicsoft, Inc.
6685 5100 Tennyson Parkway
6686 Suite 1200
6687 Plano, Texas 75024
6688 USA
6689 electronic mail: rsparks@dynamicsoft.com

6690 Mark Handley
6691 ACIRI
6692 electronic mail: mjh@aciri.org

6693 Eve Schooler
6694 Computer Science Department 256-80
6695 California Institute of Technology
6696 Pasadena, CA 91125
6697 USA
6698 electronic mail: schooler@cs.caltech.edu

6699 **References**

- 6700 [1] R. Pandya, "Emerging mobile and personal communication systems," *IEEE Communications Maga-*
6701 *zine*, Vol. 33, pp. 44–52, June 1995.
- 6702 [2] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, "RTP: a transport protocol for real-time
6703 applications," Request for Comments 1889, Internet Engineering Task Force, Jan. 1996.
- 6704 [3] H. Schulzrinne, A. Rao, and R. Lanphier, "Real time streaming protocol (RTSP)," Request for Com-
6705 ments 2326, Internet Engineering Task Force, Apr. 1998.
- 6706 [4] F. Cuervo, N. Greene, A. Rayhan, C. Huitema, B. Rosen, and J. Segers, "Megaco protocol version
6707 1.0," Request for Comments 3015, Internet Engineering Task Force, Nov. 2000.
- 6708 [5] M. Handley and V. Jacobson, "SDP: session description protocol," Request for Comments 2327, Inter-
6709 net Engineering Task Force, Apr. 1998.

- 6710 [6] S. Bradner, "Key words for use in RFCs to indicate requirement levels," Request for Comments 2119,
6711 Internet Engineering Task Force, Mar. 1997.
- 6712 [7] P. Resnick and Editor, "Internet message format," Request for Comments 2822, Internet Engineering
6713 Task Force, Apr. 2001.
- 6714 [8] H. Schulzrinne and J. Rosenberg, "SIP: Session initiation protocol – locating SIP servers," Internet
6715 Draft, Internet Engineering Task Force, Mar. 2001. Work in progress.
- 6716 [9] T. Berners-Lee, R. Fielding, and L. Masinter, "Uniform resource identifiers (URI): generic syntax,"
6717 Request for Comments 2396, Internet Engineering Task Force, Aug. 1998.
- 6718 [10] T. Berners-Lee, L. Masinter, and M. McCahill, "Uniform resource locators (URL)," Request for Com-
6719 ments 1738, Internet Engineering Task Force, Dec. 1994.
- 6720 [11] F. Yergeau, "UTF-8, a transformation format of ISO 10646," Request for Comments 2279, Internet
6721 Engineering Task Force, Jan. 1998.
- 6722 [12] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, and T. Berners-Lee, "Hypertext
6723 transfer protocol – HTTP/1.1," Request for Comments 2616, Internet Engineering Task Force, June
6724 1999.
- 6725 [13] A. Vaha-Sipila, "URLs for telephone calls," Request for Comments 2806, Internet Engineering Task
6726 Force, Apr. 2000.
- 6727 [14] N. Freed and N. Borenstein, "Multipurpose internet mail extensions (MIME) part two: Media types,"
6728 Request for Comments 2046, Internet Engineering Task Force, Nov. 1996.
- 6729 [15] D. Eastlake, S. Crocker, and J. Schiller, "Randomness recommendations for security," Request for
6730 Comments 1750, Internet Engineering Task Force, Dec. 1994.
- 6731 [16] P. Hoffman, L. Masinter, and J. Zawinski, "The mailto URL scheme," Request for Comments 2368,
6732 Internet Engineering Task Force, July 1998.
- 6733 [17] D. Meyer, "Administratively scoped IP multicast," Request for Comments 2365, Internet Engineering
6734 Task Force, July 1998.
- 6735 [18] E. M. Schooler, "A multicast user directory service for synchronous rendezvous," Master's Thesis CS-
6736 TR-96-18, Department of Computer Science, California Institute of Technology, Pasadena, California,
6737 Aug. 1996.
- 6738 [19] J. Rosenberg and H. Schulzrinne, "An offer/answer model with SDP," Internet Draft, Internet Engi-
6739 neering Task Force, Oct. 2001. Work in progress.
- 6740 [20] S. Donovan, "The SIP INFO method," Request for Comments 2976, Internet Engineering Task Force,
6741 Oct. 2000.
- 6742 [21] R. Rivest, "The MD5 message-digest algorithm," Request for Comments 1321, Internet Engineering
6743 Task Force, Apr. 1992.

- 6744 [22] V. Paxson and M. Allman, "Computing TCP's retransmission timer," Request for Comments 2988,
6745 Internet Engineering Task Force, Nov. 2000.
- 6746 [23] J. Franks, P. Hallam-Baker, J. Hostetler, S. Lawrence, P. Leach, A. Luotonen, and L. Stewart, "HTTP
6747 authentication: Basic and digest access authentication," Request for Comments 2617, Internet Engi-
6748 neering Task Force, June 1999.
- 6749 [24] J. Franks, P. Hallam-Baker, J. Hostetler, P. Leach, A. Luotonen, E. Sink, and L. Stewart, "An exten-
6750 sion to HTTP : Digest access authentication," Request for Comments 2069, Internet Engineering Task
6751 Force, Jan. 1997.
- 6752 [25] J. Galvin, S. Murphy, S. Crocker, and N. Freed, "Security multipart for MIME: multipart/signed and
6753 multipart/encrypted," Request for Comments 1847, Internet Engineering Task Force, Oct. 1995.
- 6754 [26] R. Housley, "Cryptographic message syntax," Request for Comments 2630, Internet Engineering Task
6755 Force, June 1999.
- 6756 [27] B. Ramsdell and Ed, "S/MIME version 3 message specification," Request for Comments 2633, Internet
6757 Engineering Task Force, June 1999.
- 6758 [28] T. Dierks and C. Allen, "The TLS protocol version 1.0," Request for Comments 2246, Internet Engi-
6759 neering Task Force, Jan. 1999.
- 6760 [29] S. Kent and R. Atkinson, "Security architecture for the internet protocol," Request for Comments 2401,
6761 Internet Engineering Task Force, Nov. 1998.
- 6762 [30] J. Postel, "User datagram protocol," Request for Comments 768, Internet Engineering Task Force,
6763 Aug. 1980.
- 6764 [31] J. Postel, "DoD standard transmission control protocol," Request for Comments 761, Internet Engi-
6765 neering Task Force, Jan. 1980.
- 6766 [32] R. Stewart, Q. Xie, K. Morneault, C. Sharp, H. Schwarzbauer, T. Taylor, I. Rytina, M. Kalla, L. Zhang,
6767 and V. Paxson, "Stream control transmission protocol," Request for Comments 2960, Internet Engi-
6768 neering Task Force, Oct. 2000.
- 6769 [33] F. Dawson and T. Howes, "vcard MIME directory profile," Request for Comments 2426, Internet
6770 Engineering Task Force, Sept. 1998.
- 6771 [34] G. Good, "The LDAP data interchange format (LDIF) - technical specification," Request for Com-
6772 ments 2849, Internet Engineering Task Force, June 2000.
- 6773 [35] R. Troost and S. Dorner, "Communicating presentation information in internet messages: The content-
6774 disposition header," Request for Comments 1806, Internet Engineering Task Force, June 1995.
- 6775 [36] R. Braden and Ed, "Requirements for internet hosts - application and support," Request for Comments
6776 1123, Internet Engineering Task Force, Oct. 1989.
- 6777 [37] J. Palme, "Common internet message headers," Request for Comments 2076, Internet Engineering
6778 Task Force, Feb. 1997.

- 6779 [38] H. Alvestrand, "IETF policy on character sets and languages," Request for Comments 2277, Internet
6780 Engineering Task Force, Jan. 1998.
- 6781 [39] A. Johnston, S. Donovan, R. Sparks, C. Cunningham, D. Willis, J. Rosenberg, K. Summers, and
6782 H. Schulzrinne, "SIP telephony call flow examples," Internet Draft, Internet Engineering Task Force,
6783 Apr. 2001. Work in progress.
- 6784 [40] D. Crocker, Ed., and P. Overell, "Augmented BNF for syntax specifications: ABNF," Request for
6785 Comments 2234, Internet Engineering Task Force, Nov. 1997.
- 6786 [41] H. Schulzrinne, "RTP profile for audio and video conferences with minimal control," Request for
6787 Comments 1890, Internet Engineering Task Force, Jan. 1996.
- 6788 [42] R. Hinden, B. Carpenter, and L. Masinter, "Format for literal IPv6 addresses in URL's," Request for
6789 Comments 2732, Internet Engineering Task Force, Dec. 1999.
- 6790 [43] E. M. Schooler, "Case study: multimedia conference control in a packet-switched teleconferencing
6791 system," *Journal of Internetworking: Research and Experience*, Vol. 4, pp. 99–120, June 1993. ISI
6792 reprint series ISI/RS-93-359.
- 6793 [44] H. Schulzrinne, "Personal mobility for multimedia services in the Internet," in *European Workshop on
6794 Interactive Distributed Multimedia Systems and Services (IDMS)*, (Berlin, Germany), Mar. 1996.

6795 Full Copyright Statement

6796 Copyright (c) The Internet Society (2002). All Rights Reserved.

6797 This document and translations of it may be copied and furnished to others, and derivative works that
6798 comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and
6799 distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and
6800 this paragraph are included on all such copies and derivative works. However, this document itself may not
6801 be modified in any way, such as by removing the copyright notice or references to the Internet Society or
6802 other Internet organizations, except as needed for the purpose of developing Internet standards in which case
6803 the procedures for copyrights defined in the Internet Standards process must be followed, or as required to
6804 translate it into languages other than English.

6805 The limited permissions granted above are perpetual and will not be revoked by the Internet Society or
6806 its successors or assigns.

6807 This document and the information contained herein is provided on an "AS IS" basis and THE IN-
6808 TERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WAR-
6809 RANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT
6810 THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED
6811 WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.