Internet Engineering Task Force
INTERNET-DRAFT
draft-draft-ietf-mmusic-sip-04.ps

MMUSIC WG
Handley/Schulzrinne/Schooler
ISI/Columbia U./Caltech
November 11, 1997
Expires: April 1, 1998

# SIP: Session Initiation Protocol

## Status of this Memo



**Abstract**

Many styles of multimedia conferencing are likely to co-exist on the Internet, and many of them share the need to invite users to participate. The Session Initiation Protocol (SIP) is a simple protocol designed to enable the invitation of users to participate in such multimedia sessions. It is not tied to any specific conference control scheme. In particular, it aims to enable user mobility by relaying and redirecting invitations to a user's current location.

This document is a product of the Multi-party Multimedia Session Control (MMUSIC) working group of the Internet Engineering Task Force. Comments are solicited and should be addressed to the working group's mailing list at confctrl@isi.edu and/or the authors.

## Contents

**1 Introduction** 7

1.1 Overview of SIP Functionality . . . 8

1.2 Terminology . . . 9

1.3 Definitions . . . 9

1.4 Summary of SIP Operation . . . . . 11
1.4.1 SIP Addressing . . . . . 11
1.4.2 Locating a SIP Server . . . . . 12
1.4.3 SIP Transaction . . . . . 13
1.4.4 SIP Invitation . . . . . 13
1.4.5 Locating a User . . . . . 15
1.4.6 Changing an Existing Session . . . . . 16
1.4.7 Registration Services . . . . . 16
1.5 Protocol Properties . . . . . 17
1.5.1 Minimal State . . . . . 17
1.5.2 Transport-Protocol Neutral . . . . . 17
1.5.3 Text-Based . . . . . 17

**2 SIP Uniform Resource Locators 17**

**3 SIP Message Overview 20**

**4 Request 22**

4.1 Request-Line . . . . . 22
4.2 Methods . . . . . 22
4.2.1 INVITE . . . . . 22
4.2.2 ACK . . . . . 23
4.2.3 OPTIONS . . . . . 23
4.2.4 BYE . . . . . 23
4.2.5 REGISTER . . . . . 24
4.2.6 UNREGISTER . . . . . 24
4.3 Request-URI . . . . . 24

4.3.1 SIP Version . . . . . 24
4.4 Option Tags . . . . . 25
4.4.1 Registering New Option Tags with IANA . . . . . 25

**5 Response 25**

5.1 Status-Line . . . . . 26
5.1.1 Status Codes and Reason Phrases . . . . . 26

**6 Header Field Definitions 27**

6.1 General Header Fields . . . . . 29
6.2 Entity Header Fields . . . . . 29
6.3 Request Header Fields . . . . . 29
6.4 Response Header Fields . . . . . 31
6.5 Header Field Format . . . . . 31
6.6 Accept . . . . . 32
6.7 Accept-Language . . . . . 32
6.8 Allow . . . . . 32
6.9 Also . . . . . 32
6.10 Authorization . . . . . 32
6.11 Call-Disposition . . . . . 33
6.12 Call-ID . . . . . 33
6.13 Content-Length . . . . . 34
6.14 Content-Type . . . . . 34
6.15 Date . . . . . 34
6.16 Expires . . . . . 35
6.17 From . . . . . 35

6.18 Location . . . . . . . . . . 36
6.19 Organization . . . . . . . . . . 38
6.20 Priority . . . . . . . . . . 38
6.21 Proxy-Authenticate . . . . . . . . . . 38
6.22 Proxy-Authorization . . . . . . . . . . 38
6.23 Public . . . . . . . . . . 38
6.24 Require . . . . . . . . . . 38
6.25 Retry-After . . . . . . . . . . 39
6.26 CSeq . . . . . . . . . . 40
6.27 Server . . . . . . . . . . 40
6.28 Subject . . . . . . . . . . 40
6.29 Unsupported . . . . . . . . . . 40
6.30 Timestamp . . . . . . . . . . 41
6.31 To . . . . . . . . . . 41
6.32 User-Agent . . . . . . . . . . 41
6.33 Via . . . . . . . . . . 41
6.34 Warning . . . . . . . . . . 42
6.35 WWW-Authenticate . . . . . . . . . . 43

**7 Status Code Definitions 43**

7.1 Informational 1xx . . . . . . . . . . 43
7.1.1 100 Trying . . . . . . . . . . 44
7.1.2 180 Ringing . . . . . . . . . . 44
7.1.3 181 Queued . . . . . . . . . . 44
7.2 Successful 2xx . . . . . . . . . . 44
7.2.1 200 OK . . . . . . . . . . 44

7.3 Redirection 3xx . . . . . . . . . . 44

7.3.1 300 Multiple Choices . . . . . . . . . . 45

7.3.2 301 Moved Permanently . . . . . . . . . . 45

7.3.3 302 Moved Temporarily . . . . . . . . . . 45

7.3.4 380 Alternative Service . . . . . . . . . . 45

7.4 Request Failure 4xx . . . . . . . . . . 45

7.4.1 400 Bad Request . . . . . . . . . . 46

7.4.2 401 Unauthorized . . . . . . . . . . 46

7.4.3 402 Payment Required . . . . . . . . . . 46

7.4.4 403 Forbidden . . . . . . . . . . 46

7.4.5 404 Not Found . . . . . . . . . . 46

7.4.6 405 Method Not Allowed . . . . . . . . . . 46

7.4.7 407 Proxy Authentication Required . . . . . . . . . . 46

7.4.8 408 Request Timeout . . . . . . . . . . 47

7.4.9 420 Bad Extension . . . . . . . . . . 47

7.4.10 480 Temporarily Unavailable . . . . . . . . . . 47

7.4.11 481 Invalid Call-ID . . . . . . . . . . 47

7.4.12 482 Loop Detected . . . . . . . . . . 47

7.5 Server Failure 5xx . . . . . . . . . . 47

7.5.1 500 Server Internal Error . . . . . . . . . . 47

7.5.2 501 Not implemented . . . . . . . . . . 48

7.5.3 502 Bad Gateway . . . . . . . . . . 48

7.5.4 503 Service Unavailable . . . . . . . . . . 48

7.5.5 504 Gateway Timeout . . . . . . . . . . 48

7.6 Global Failures 6xx . . . . . . . . . . 48

7.6.1 600 Busy . . . 48

7.6.2 603 Decline . . . 49

7.6.3 604 Does not exist anywhere . . . 49

7.6.4 606 Not Acceptable . . . 49

**8 SIP Message Body** **49**

8.1 Body Inclusion . . . 50

8.2 Message Body Length . . . 50

**9 Examples** **50**

9.1 Invitation to Multimedia Conference . . . 50

9.1.1 Request . . . 50

9.1.2 Reply . . . 51

9.2 Two-party Call . . . 52

9.3 Aborting a Call . . . 53

9.3.1 Redirects . . . 53

9.3.2 Alternative Services . . . 54

9.3.3 Negotiation . . . 55

9.4 OPTIONS Request . . . 56

**10 Compact Form** **56**

**11 SIP Transport** **57**

11.1 Achieving Reliability For UDP Transport . . . 58

11.1.1 General Operation . . . 58

11.1.2 INVITE . . . 58

11.2 Connection Management for TCP . . . 59

**12 Behavior of SIP Servers** **59**

12.1 Redirect Server . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 62

12.2 User Agent Server . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 62

12.3 Proxies Issuing Single Unicast Requests . . . . . . . . . . . . . . . . . . . . . . . . . . . . 62

12.4 Proxy Server Issuing Several Requests . . . . . . . . . . . . . . . . . . . . . . . . . . . . 62

**13 Third-Party Call Initiation** **65**

**14 ISDN and Intelligent Network Services** **66**

**15 Security Considerations** **69**

15.1 Confidentiality . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 69

15.2 Integrity . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 69

15.3 Access Control . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 69

15.4 Privacy . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 70

**A Minimal Implementation** **70**

A.1 Client . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 70

A.2 Server . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 70

**B Summary of Augmented BNF** **71**

**C Open Issues** **72**

**D Acknowledgments** **73**

**E Authors' Addresses** **73**

# 1 Introduction

## 1.1 Overview of SIP Functionality

The Session Initiation Protocol (SIP) is an application-layer protocol that can establish and control multimedia sessions or calls. These multimedia sessions include multimedia conferences, distance learning, Internet telephony and similar applications. SIP can invite a person to both unicast and multicast sessions; the initiator does not necessarily have to be a member of the session it is inviting to. Media and participants can be added to an existing session. SIP can be used to "call" both persons and "robots", for example, to invite a media storage device to record an ongoing conference or to invite a video-on-demand server to play a video into a conference. (SIP does not directly control these services, however; see RTSP [1].)

SIP can be used to initiate sessions as well as invite members to sessions that have been advertised and established by other means. (Sessions may be advertised using multicast protocols such as SAP [2], electronic mail, news groups, web pages or directories (LDAP), among others.)

SIP transparently supports name mapping and redirection services, allowing the implementation of ISDN and Intelligent Network telephony subscriber services. Section 14 discusses these services in detail.

SIP supports *personal mobility*. In the parlance of telecommunications intelligent network services, this is defined as: "Personal mobility is the ability of end users to originate and receive calls and access subscribed telecommunication services on any terminal in any location, and the ability of the network to identify end users as they move. Personal mobility is based on the use of a unique personal identity (i.e., 'personal number')." [3, p. 44]. Personal mobility complements terminal mobility, i.e., the ability to maintain communications when moving a single end system from one network to another.

SIP supports some or all of five facets of establishing and terminating multimedia communications:

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

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

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

**Call setup:** "ringing", establishment of call parameters at both called and calling party;

**Call handling:** including transfer and termination of calls.

SIP may also be used in conjunction with other call setup and signaling protocols. In that mode, an end system uses SIP protocol exchanges to determine the appropriate end system address and protocol from a given address that is protocol-independent. For example, SIP may be used to determine that the party may be reached via H.323, obtain the H.245 gateway and user address and then use H.225.0 to establish the call [4]. In another example, it may be used to determine that the callee is reachable via the public switched telephone network (PSTN) and indicate the phone number to be called, possibly suggesting an Internet-to-PSTN gateway to be used.

SIP can also initiate multi-party calls using a multipoint control unit (MCU) or fully-meshed interconnection instead of multicast.

Internet telephony gateways that connect PSTN parties may also use SIP to set up calls between them.

SIP does not offer conference control services such as floor control or voting and does not prescribe how a conference is to be managed, but SIP can be used to introduce conference control protocols.

SIP does not allocate multicast addresses, leaving this functionality to protocols such as SAP [2].

SIP can invite users to sessions with and without resource reservation. SIP does not reserve resources, but may convey to the invited system the information necessary to do this. Quality-of-service guarantees, if required, may depend on knowing the full membership of the session; this information may or may not be known to the agent performing session invitation.

SIP is designed as part of the overall IETF multimedia data and control architecture [5] currently incorporating protocols such as RSVP [6] for reserving network resources, the real-time transport protocol (RTP) [7] for transporting real-time data and providing QOS feedback, the real-time streaming protocol (RTSP) [8] for controlling delivery of streaming media, the session announcement protocol (SAP) [2] for advertising multimedia sessions via multicast and the session description protocol (SDP) [9] for describing multimedia sessions, but the functionality and operation of SIP does not depend on any of these protocols.

## 1.2 Terminology

In this document, the key words “MUST”, “MUST NOT”, “REQUIRED”, “SHALL”, “SHALL NOT”, “SHOULD”, “SHOULD NOT”, “RECOMMENDED”, “MAY”, and “OPTIONAL” are to be interpreted as described in RFC 2119 [10] and indicate requirement levels for compliant SIP implementations.

## 1.3 Definitions

This specification uses a number of terms to refer to the roles played by participants in SIP communications. The definitions of client, server and proxy are similar to those used by the Hypertext Transport Protocol (HTTP) [11]. The following terms have special significance for SIP.

**Call:** A call consists of a single invitation attempt from a single user. A SIP call is identified by a globally unique call-id (Section 6.12. Thus, if a user is, for example, invited to the same multicast session by several people, each of these invitations will be a unique call. A point-to-point Internet telephony conversation maps into a single SIP call. In a MCU-based conference, each participant uses a separate call to invite himself to the MCU.

**Client:** An application program that establishes connections for the purpose of sending requests. Clients may or may not interact directly with a human user.

**Final response:** A response that terminates a SIP transaction, as opposed to a *provisional response*. All 2xx, 3xx, 4xx, and 5xx responses are final.

**Initiator, calling party:** The party initiating a conference invitation. Note that the calling party does not have to be the same as the one creating a conference.

**Invitation:** A request sent to a user (or service) requesting participation in a session. A successful SIP invitation consists of two transactions: an INVITE request followed by a ACK request.

**Invitee, invited user, called party:** The person or service that the calling party is trying to invite to a conference.

**Location server:** See *location service*.

**Location service:** A service used by a SIP redirect or proxy server to obtain information about a callee's possible location(s). Location services are offered by location servers. Location servers may be co-located with a SIP server, but the manner in which a SIP server requests location services is beyond the scope of the document.

**Provisional response:** A response used by the server to indicate progress, but that does not terminate a SIP transaction. All 1xx and 6xx responses are provisional. Other responses are considered final.

**Proxy, proxy server:** An intermediary program that acts as both a server and a client for the purpose of making requests on behalf of other clients. Requests are serviced internally or by passing them on, possibly after translation, to other servers. A proxy must interpret, and, if necessary, rewrite a request message before forwarding it.

**Redirect server:** A server that accepts a SIP request, maps the address into zero or more new addresses and returns these addresses to the client. Unlike a proxy server, it does not initiate its own SIP request. Unlike a user agent server, it does not accept calls.

**Server:** An application program that accepts requests in order to service requests and sends back responses to those requests. Servers are either proxy, redirect or user agent servers. An application program may act as both server and client.

**Session:** "A multimedia session is a set of multimedia senders and receivers and the data streams flowing from senders to receivers. A multimedia conference is an example of a multimedia session." [9] (Note: a session as defined here may comprise one or more RTP sessions.) Since the word session is used differently by protocols relevant to SIP, this document avoids the term altogether.

**(SIP) transaction:** A SIP transaction occurs between a client and a server and comprises all messages from the first request sent from the client to the server up to a final (non-1xx) response sent from the server to the client. A transaction is for a single call (identified by a Call-ID, Section 6.12). There can only be one pending transaction between a server and client for each call id.

**User agent server, called user agent:** The server application that contacts the user when a SIP request is received and that returns a reply on behalf of the user. The reply may accept, reject or redirect the call. (Note: in SIP, user agents can be both clients and servers.)

An application program may be capable of acting both as a client and a server. For example, a typical multimedia conference control application would act as a client to initiate calls or to invite others to conferences and as a user agent server to accept invitations. The properties of the different SIP server types are summarized in Table 1.

| property | redirect server | proxy server | user agent server |
|---|---|---|---|
| also acts as client | no | yes | no |
| return 1xx status | yes | yes | yes |
| return 2xx status | no | yes | yes |
| return 3xx status | yes | yes | yes |
| return 4xx status | yes | yes | yes |
| return 5xx status | yes | yes | yes |
| return 6xx status | no | yes | yes |
| insert Via header | no | yes | no |
| accept ACK | no | yes | yes |

Table 1: Properties of the different SIP server types

## 1.4 Summary of SIP Operation

This section explains the basic protocol functionality and operation. Callers and callees are identified by SIP addresses, described in Section 1.4.1. When making a SIP call, a caller first locates the appropriate server (Section 1.4.2) and then sends a SIP request (Section 1.4.3). The most common SIP operation is the invitation (Section 1.4.4). Instead of directly reaching the intended callee, a SIP request may be redirected or trigger a chain of new SIP requests by proxies (Section 1.4.5). Users can register with SIP servers (Section 4.2.5).

### 1.4.1 SIP Addressing

SIP addresses contain a *user* and *host* part. The *user* part is an operating-system user name. The *host* part is either a domain name having a DNS A (address) record, or a numeric network address. Examples include:

```
mjh@metro.isi.edu
hgs@erlang.cs.columbia.edu
root@[193.175.132.42]
root@193.175.132.42
```

A user's address can be obtained out-of-band, can be learned via existing media agents, can be included in some mailers' message headers, or can be recorded during previous invitation interactions.

SIP addresses may contain a moniker (such as a civil name) or user name and domain name that may not map into a single host. [1].

SIP addresses may use any unambiguous user name, including aliases, identifying the called party as the *user* part of the address. They may use a domain name having an MX [12], SRV [13] or A [14] record for

[1] We avoid the term *location-independent*, since the address may indeed refer to a specific location, e.g., a company department.

the *host* part. These addresses may have different degrees of location- and provider-independence and are often chosen to be mnemonic. In many cases, the SIP address can be the same as a user's electronic mail address, but this is not required. SIP can thus leverage off the domain name system (DNS) to provide a first-stage location mechanisms. Examples of SIP names include

```
M.Handley@cs.ucl.ac.uk
H.G.Schulzrinne@ieee.org
info@ietf.org
```

An address can designate an individual (possibly located at one of several end systems), the first available person from a group of individuals or a whole group. The form of the address, e.g., `sales@example.com`, is not sufficient, in general, to determine the intent of the caller.

If a user or service chooses to be reachable at an address that is guessable from the person's name and organizational affiliation, the traditional method of ensuring privacy by having an unlisted "phone" number is compromised. However, unlike traditional telephony, SIP offers authentication and access control mechanisms and can avail itself of lower-layer security mechanisms, so that client software can reject unauthorized or undesired call attempts.

When used within SIP, SIP addresses are written as SIP URLs (Section sec:url), e.g., `sip://info@ietf.org` as SIP requests and responses may also contain non-SIP addresses, e.g., telephone numbers.

### 1.4.2 Locating a SIP Server

A SIP client MUST follow the following steps to resolve the *host* part of a callee address. If a client only supports TCP or UDP, but not both, the respective address type is omitted. If the SIP address contains a port number, that number is to be used, otherwise, the the default port number. The default port number for UDP and TCP is the same.

1. If the SIP address is a numeric IP address, contact a SIP server at that address.
2. If the SIP address does not contain a port number and if there is a SRV DNS resource record [13] of type sip.udp, contact the listed SIP servers in the order of the preference values contained in those resource records, using UDP as a transport protocol at the port number listed in the DNS resource record. [TBD: What if the SIP URL contains a port number?]
3. If the SIP address does not contain a port number and if there is a SRV DNS resource record [13] of type sip.tcp, contact the listed SIP servers in the order of the preference value contained in those resource records, using TCP as a transport protocol at the port number listed in the DNS resource record.
4. If there is a DNS MX record [12], contact the hosts listed in their order of preference at the default port number (TBD). For each host listed, first try to contact the SIP server using UDP, then TCP.
5. Finally, check if there is a DNS CNAME or A record for the given *host* and try to contact a SIP server at the one or more addresses listed, again trying first UDP, then TCP.

6. If all of the above methods fail, the caller MAY contact an SMTP server at the user's *host* and use the SMTP EXPN command to obtain an alternate address and repeat the steps above. As a last resort, a client MAY choose to deliver the session description to the callee using electronic mail.

If a server is found using one of the methods below, the other methods are not tried. A client SHOULD rely on ICMP "Port Unreachable" messages rather than time-outs to determine that a server is not reachable at a particular address.

A client MAY cache the result of the reachability steps for a particular address and retry that host address for the next call. If the client does not find a SIP server at the cached address, it MUST start the search at the beginning of the sequence.

Implementation note for socket-based programs: For TCP, `connect()` returns `ECONNREFUSED` if there is no server at the designated address; for UDP, the socket should be bound to the destination address using `connect()` rather than `sendto()` or similar.

This sequence is modeled after that described for SMTP, where MX records are to be checked before A records [15].

### 1.4.3 SIP Transaction

Once the *host* part has been resolved to a SIP server, the client sends one or more SIP requests to that server and receives one or more responses from the server. A request (and its retransmissions) together with the responses triggered by that request make up a SIP transaction.

If TCP is used, request and responses within a single SIP transaction are carried over the same TCP connection. Thus, the client SHOULD maintain the connection until a final response has been received. Several SIP requests from the same client to the same server may use the same TCP connection or may open a new connection for each request. If the client sent the request sends via unicast UDP, the response is sent to the source address of the UDP request. (Implementation note: use `recvfrom()` to obtain the source address and port of the request.) If the request is sent via multicast UDP, the response is directed to the same multicast address and destination port. For UDP, reliability is achieved using retransmission (Section 11).

Need motivation why we ALWAYS want to have a multicast return.

The SIP message format and operation is independent of the transport protocol.

### 1.4.4 SIP Invitation

A successful SIP invitation consists of two requests, INVITE followed by ACK. The INVITE (Section 4.2.1) request asks the callee to join a particular conference or establish a two-party conversation. After the callee has agreed to participate in the call, the caller confirms that it has received that response by sending an ACK (Section 4.2.2) request. If the call is rejected or otherwise unsuccessful, the caller sends a BYE request instead of an ACK.

The INVITE request typically contains a session description, for example written in SDP format, that provides the called party with enough information to join the session. For multicast sessions, the session description enumerates the media types and formats that may be distributed to that session. For unicast session, the session description enumerates the media types and formats that the caller is willing to receive and where it wishes the media data to be sent. In either case, if the callee wishes to accept the call, it responds to the invitation by returning a similar description listing the media it wishes to receive. For a multicast session, the callee should only return a session description if it is unable to receive the media indicated in the caller's description. The caller may ignore the session description returned or use it to change the global session description.

The session description may refer to a session start time in the future. Actual transmission of data SHOULD not start until the time indicated in the session description.

The protocol exchanges for the INVITE method are shown in Fig. 1 for a proxy server and in Fig. 2 for a redirect server. The proxy server accepts the INVITE request (step 1), contacts the location service with all or parts of the address (step 2) and obtains a more precise location (step 3). The proxy server then issues a SIP INVITE request to the address(es) returned by the location service (step 4). The user agent server alerts the user (step 5) and returns a success indication to the proxy server (step 6). The proxy server then returns the success result to the original caller (step 7). The receipt of this message is confirmed by the caller using an ACK message, which is forwarded to the callee (steps 8 and 9), with a response returned (steps 10 and 11). All requests have the same Call-ID.

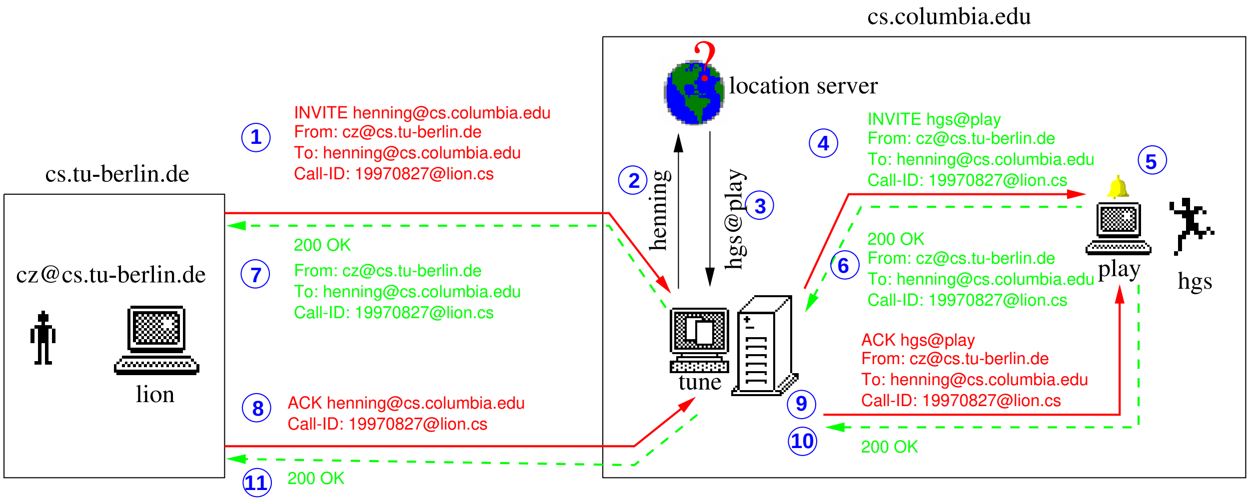


Figure 1: Example of SIP proxy server

The redirect server accepts the INVITE request (step 1), contacts the location service as before (steps 2 and 3) and, instead of contacting the newly found address itself, returns the address to the caller (step 4). The caller issues a new request, with a new call-ID, to the address returned by the first server (step 6). In the example, the call succeeds (step 7). The caller and callee complete the handshanke with an ACK (steps 8 and 9).

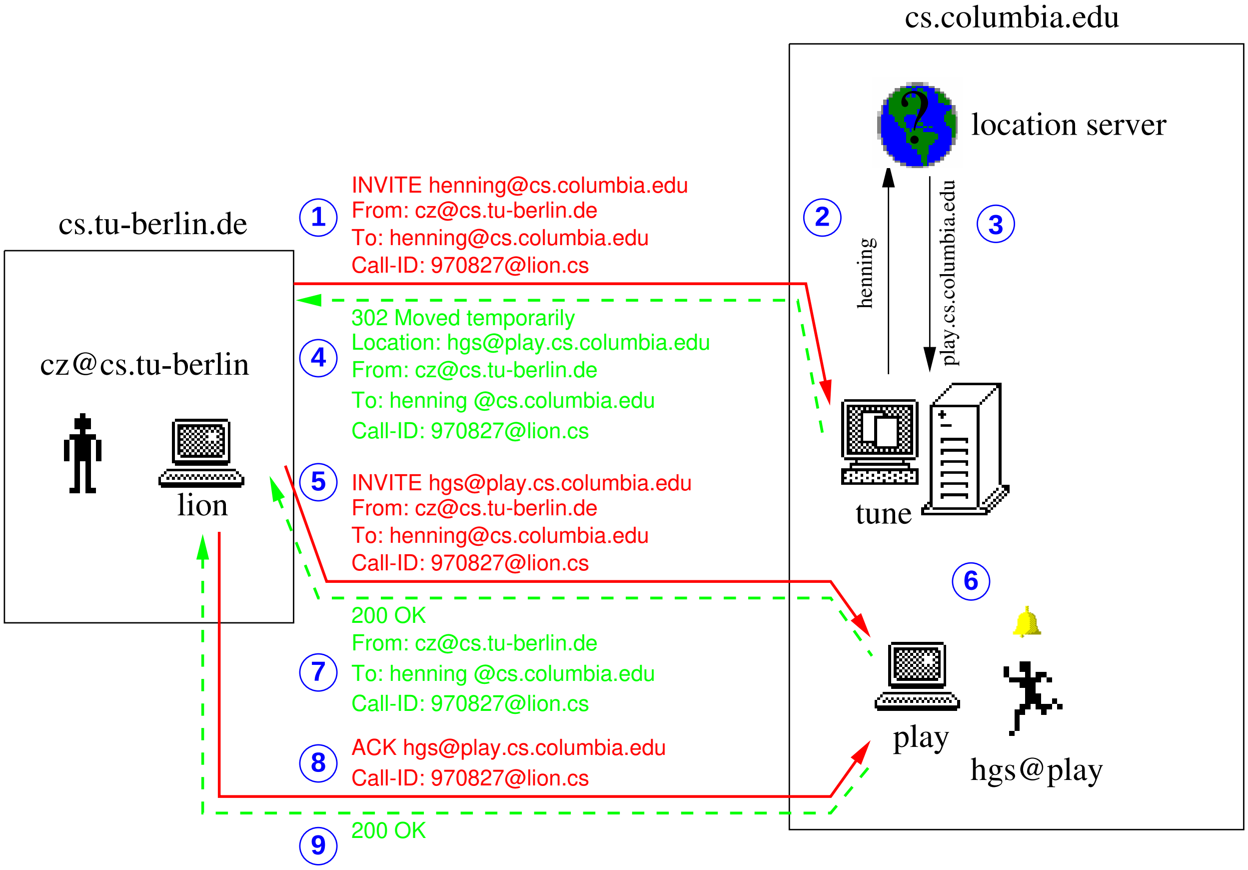


Figure 2: Example of SIP redirect server

The next section discusses what happens if the location service returns more than one possible alternative.

### 1.4.5 Locating a User

A callee may move between a number of different end systems over time. These locations can be dynamically registered with the SIP server (Section 4.2.5) or a location server, typically for a single administrative domain, or a location server may use other protocols, such as finger [16], rwho, multicast-based protocols or operating-system dependent mechanism to actively determine the end system where a user might be reachable. The location services yield a list of a zero or more possible locations, possibly even sorted in order of likelihood of success.

The location server can be part of the SIP server or the SIP server may use a different protocol (e.g., finger [16] or LDAP [17]) to map addresses. A single user may be registered at different locations, either because she is logged in at several hosts simultaneously or because the location server has (temporarily) inaccurate information.

The action taken on receiving a list of locations varies with the type of SIP server. A SIP redirect server simply returns the list to the client sending the request as Location headers (Section 6.18). A SIP proxy server can sequentially or in parallel try the addresses until the call is successful (2xx response) or the callee has declined the call (60x response). With sequential attempts, a proxy server can implement an "anycast" service.

If a proxy server forwards a SIP request, it MUST add itself to the end of the list of forwarders noted in the Via (Section 6.33) headers. The Via trace ensures that replies can take the same path back, thus ensuring correct operation through compliant firewalls and loop-free requests. On the reply path, each host most remove its Via, so that routing internal information is hidden from the callee and outside networks. When a multicast request is made, first the host making the request, then the multicast address itself are added to the path. A proxy server MUST check that it does not generate a request to a host listed in the Via list. (Note: If a host has several names or network addresses, this may not always work. Thus, each host also checks if it is part of the Via list.)

A SIP invitation may traverse more than one SIP proxy server. If one of these "forks" the request, i.e., issues more than one request in response to receiving the invitation request, it is possible that a client is reached, independently, by more than one copy of the invitation request. Each of these copies bears the same Call-ID. The user agent MUST return the appropriate status response, but SHOULD NOT alert the user.

As discussed in Section 1.4.1, a SIP address may designate a group rather than an individual. A client indicates using the Reach request header whether it wants to reach the first available individual or treat the address as a group, to be invited as a whole. The default is to attempt to reach the first available callee. If the address is designated as a group address, a proxy server MUST return the list of individuals instead of attempting to connect to these. (Otherwise, the proxy cannot report errors, redirections and call status individually. For example, some may be contacted successfully, while one of the group may be reachable under a different address.)

### 1.4.6 Changing an Existing Session

In some circumstances, it may be necessary to change the parameters of an existing session. For example, two parties may have been conversing and then want to add a third party, switching to multicast for efficiency. One of the participants invites the third party with the new multicast address and simultaneously sends an INVITE to the second party, with the new multicast session description, but the old call identifier.

### 1.4.7 Registration Services

The REGISTER and UNREGISTER requests allow a client to let a proxy or redirect server know which address it may be reached under. A client may also use it to install call handling features at the server.

## 1.5 Protocol Properties

### 1.5.1 Minimal State

A single conference session or call may involve one or more SIP request-response transactions. Proxy server do not have to keep state for a particular call, however, they maintain state for a single SIP transaction, as discussed in Section 12.

For efficiency, a server may cache the results of location service requests.

### 1.5.2 Transport-Protocol Neutral

SIP is able to utilize both UDP and TCP as transport protocols. UDP allows the application to more carefully control the timing of messages and their retransmission, to perform parallel searches without requiring TCP connection state for each outstanding request, and to use multicast. Routers can more readily snoop SIP UDP packets. TCP allows easier passage through existing firewalls, and given the similar protocol design, allows common servers for SIP, HTTP and the Real Time Streaming Protocol (RTSP) [1].

When TCP is used, SIP can use one or more connections to attempt to contact a user or to modify parameters of an existing conference. Different SIP requests for the same SIP call may use different TCP connections or a single persistent connection, as appropriate.

Clients SHOULD implement both UDP and TCP transport, servers MUST.

### 1.5.3 Text-Based

SIP is text based. This allows easy implementation in languages such as Tcl and Perl, allows easy debugging, and most importantly, makes SIP flexible and extensible. As SIP is used for initiating multimedia conferences rather than delivering media data, it is believed that the additional overhead of using a text-based protocol is not significant.

# 2 SIP Uniform Resource Locators

SIP URLs are used within SIP messages to indicate the originator and recipient of a SIP request, and to specify redirection addresses. A SIP URL may also be embedded in web pages or other hyperlinks to indicate that a user or service may be called.

Because interaction with some resources may require message headers or message bodies to be specified as well as the SIP address, the sip URL scheme is defined to allow setting SIP request-header fields and the SIP message-body. (This is similar to the mailto: URL.)

A SIP URL follows the guidelines of RFC 1630 [18, 19] and takes the following form:

```
SIP-URL          =  short-sip-url | full-sip-url
full-sip-url     =  "sip://" ( user | phone ) [ ":" password ]
                    "@" [ host | nhost ]
                    url-parameters [ headers ]
short-sip-url    =  ( user | phone) [ ":" password ]
                    "@" [ host | nhost ] : port
user             =  ; defined in RFC 1738 [20]
phone            =  "+" DIGIT *( DIGIT | "-" | "." )
host             =  ; defined in RFC 1738
nhost            =  "[" hostnumber "]" | hostnumber
hostnumber       =  digits "." digits "." digits "." digits
port             =  *digit
url-parameters   =  *( ";" url-parameter)
url-parameter    =  transport-param |
                    ttl-param | maddr-param
transport-param  =  "transport=" ( "udp" | "tcp" )
ttl-param        =  "ttl=" ttl
ttl              =  1*3DIGIT                     ; 0 to 255
maddr-param      =  "maddr=" maddr
maddr            =  ; dotted decimal multicast address
headers          =  "?" header *( "                " header )
header           =  hname "=" hvalue
hname            =  *urlc
hvalue           =  *urlc
urlc             =  ; defined in [19]
digits           =  1*digit
```

Thus, a SIP URL can take either a short form or a full form. The short form MAY only be used within SIP messages where the scheme (SIP) can be assumed. In all other cases, and when parameters are required to be specified, the full form MUST be used.

Note that all URL reserved characters must be encoded. The special hname "body" indicates that the associated hvalue is the message-body of the SIP INVITE request. Within sip URLs, the characters "?", "=", "&" are reserved.

The mailto: URL and RFC 822 email addresses require that numeric host addresses ("host numbers") are enclosed in square brackets (presumably, since host names might be numeric), while host numbers without brackets are used for all other URLs. The SIP URL allows both forms.

The password parameter can be used for a basic authentication mechanism that takes the place of an unlisted telephone number. Also, for Internet telephony gateways, it may serve as a PIN. Including just the password in the URL is more convenient than including a whole authentication header. This approach may be reasonably secure if the URL is part of a secure web page. Unless the SIP transaction is carried over a secure network connection, this carries the same security risks as all URL-based passwords and should only be used when security requirements are low. In almost all circumstances, use of the Authorization (Section 6.10) header is preferred.

The phone identifier is to be used when connecting to a telephony gateway. The phone number follows the rules for international numbers in ITU Recommendation E.123, with only numbers and hyphens allowed.

Examples of short and full-form SIP URLs are:

```
j.doe@big.com
sip://j.doe@big.com
sip://j.doe:secret@big.com;transport=tcp
sip://j.doe@big.com?subject=project%20x&priority=urgent
sip://+1-212-555-1212:1234@gateway.com
sip://alice@[10.1.2.3]
sip://alice@10.1.2.3
```

Within a SIP message, URLs are used to indicate the source and intended destination of a request, redirection addresses and the current destination of a request. Normally all these fields will contain SIP URLs. When additional parameters are not required, the short form SIP URL can be used unambiguously.

In some circumstances a non-SIP URL may be used in a SIP message. An example might be making a call from a telephone which is relayed by a gateway onto the internet as a SIP request. In such a case, the source of the call is really the telephone number of the caller, and so a SIP URL is inappropriate and a phone URL might be used instead. Thus where SIP specifies user addresses it allows these addresses to be URLs.

Clearly not all URLs are appropriate to be used in a SIP message as a user address. The correct behavior when an unknown scheme is encountered by a SIP server is defined in the context of each of the header fields that use a SIP URL.

SIP URLs can define specific parameters of the request, including the transport mechanism (UDP or TCP) and the use of multicast to make a request. These parameters are added after the host and are separated by semi-colons. For example, to specify to call `j.doe@big.com` using multicast to 239.255.255.1 with a ttl of 15, the following URL would be used:

```
sip://j.doe@big.com;maddr=239.255.255.1;ttl=15
```

The transport protocol UDP is to be assumed when a multicast address is given.

# 3 SIP Message Overview

Since much of the message syntax is identical to HTTP/1.1, rather than repeating it here we use [HX.Y] to refer to Section X.Y of the current HTTP/1.1 specification [11]. In addition, we describe SIP in both prose and an augmented Backus-Naur form (BNF) [H2.1] described in detail in [21].

All SIP messages are text-based and use HTTP/1.1 conventions [H4.1], except for the additional ability of SIP to use UDP. When sent over TCP or UDP, multiple SIP transactions can be carried in a single TCP connection or UDP datagram. UDP datagrams, including all headers, should not normally be larger than the path maximum transmission unit (MTU) if the MTU is known, or 1500 bytes if the MTU is unknown.

> The 1400 bytes accommodates lower-layer packet headers within the "typical" MTU of around 1500 bytes. There are few MTU values around 1 kB; the next value is 1006 bytes for SLIP and 296 for low-delay PPP [22]. Recent studies [23, p. 154] indicate that an MTU of 1500 bytes is a reasonable assumption. Thus, another reasonable value would be a message size of 950 bytes, to accommodate packet headers within the SLIP MTU without fragmentation.

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

```
SIP-message  =  Request | Response ; SIP messages
```

Both Request (section 4) and Response (section 5) messages use the generic message format of RFC 822 [24] for transferring entities (the payload of the message). Both types of message consist of a start-line, one or more header fields (also known as "headers"), an empty line (i.e., a line with nothing preceding the carriage-return line-feed (CRLF)) indicating the end of the header fields, and an optional message-body. To avoid confusion with similar-named headers in HTTP, we refer to the header describing the message body as entity headers. These components are described in detail in the upcoming sections.

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

start-line       =  Request-Line |      Section 4.1
                    Status-Line         Section 5.1

message-header   =  *( general-header
                    | request-header
                    | entity-header )
```

In the interest of robustness, any leading empty line(s) MUST be ignored. In other words, if the Request or Response message begins with a CRLF, the CRLF should be ignored.

| | | | |
|---|---|---|---|
| general-header | = | Call-ID | ; Section 6.12 |
| | \| | CSeq | ; Section 6.26 |
| | \| | Date | ; Section 6.15 |
| | \| | Expires | ; Section 6.16 |
| | \| | From | ; Section 6.17 |
| | \| | Via | ; Section 6.33 |
| entity-header | = | Content-Length | ; Section 6.13 |
| | \| | Content-Type | ; Section 6.14 |
| request-header | = | Accept | ; Section 6.6 |
| | \| | Accept-Language | ; Section 6.7 |
| | \| | Authorization | ; Section 6.10 |
| | \| | Call-Disposition | ; Section 6.11 |
| | \| | Organization | ; Section 6.19 |
| | \| | Priority | ; Section 6.20 |
| | \| | Proxy-Authorization | ; Section 6.22 |
| | \| | Require | ; Section 6.24 |
| | \| | Subject | ; Section 6.28 |
| | \| | To | ; Section 6.31 |
| | \| | User-Agent | ; Section 6.32 |
| response-header | = | Location | ; Section 6.18 |
| | \| | Proxy-Authenticate | ; Section 6.21 |
| | \| | Public | ; Section 6.23 |
| | \| | Retry-After | ; Section 6.25 |
| | \| | Server | ; Section 6.27 |
| | \| | Unsupported | ; Section 6.29 |
| | \| | Warning | ; Section 6.34 |
| | \| | WWW-Authenticate | ; Section 6.35 |

Table 2: SIP headers

# 4 Request

The Request message format is shown below:

```
Request  =  Request-Line          ; Section 4.1
            *( general-header
            | request-header
            | entity-header )
            CRLF
            [ message-body ]      ; Section 8
```

## 4.1 Request-Line

The Request-Line begins with a method token, followed by the Request-URI and the protocol version, and ending with CRLF. The elements are separated by SP characters. No CR or LF are allowed except in the final CRLF sequence.

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

## 4.2 Methods

The methods are defined below. Methods that are not supported by a proxy or redirect server SHOULD be treated by that server as if they were an INVITE method and forwarded accordingly.

Methods that are not supported by a user agent server should cause a "501 Not Implemented" response to be returned (Section 7).

```
method  =  "INVITE" | "ACK" | "OPTIONS"
        |  "BYE" | "REGISTER" | "UNREGISTER"
```

### 4.2.1 INVITE

The INVITE method indicates that the user or service is being invited to participate in a session. The message body contains a description of the session the callee is being invited to. For two-party calls, the caller indicates the type of media it is able to receive as well as their parameters such as network destination. If the session description format allows this, it may also indicate "send-only" media. A success response indicates in its message body which media the callee wishes to receive.

A server MAY automatically respond to an invitation for a conference the user is already participating in, identified either by the SIP Call-ID or a globally unique identifier within the session description, with a "200 OK" response.

A user agent MUST check any version identifiers in the session description to see if it has changed. If the version number has changed, the user agent server MUST adjust the session parameters accordingly, possibly after asking the user for confirmation. (Versioning of the session description may be used to accomodate the capabilities of new arrivals to a conference or change from a unicast to a multicast conference.)

This method MUST be supported by a SIP server.

#### 4.2.2 ACK

ACK request confirms that the client has received a final response to an INVITE request. See Section 11 for details. This method MUST be supported by a SIP server and client.

#### 4.2.3 OPTIONS

The client is being queried as to its capabilities. A server that believes it can contact the user, such as a user agent where the user is logged in and has been recently active, MAY respond to this request with a capability set. Support of this method is OPTIONAL.

#### 4.2.4 BYE

The client indicates to the server that it wishes to abort the call attempt. The leaving party can use a Location header field to indicate that the recipient of request should contact the named address. This implements the "call transfer" telephony functionality. A client SHOULD also use this method to indicate to the callee that it wishes to abort an on-going call attempt.

> With UDP, the caller has no other way to signal her intent to drop the call attempt and the callee side will keep "ringing".

When using TCP, a client MAY also close the connection to abort a call attempt. Support of this method is OPTIONAL.

Support of this method is OPTIONAL.

#### 4.2.5 REGISTER

A client uses the REGISTER method to register the address listed in the request line to a SIP server. The host part of the request-URI SHOULD correspond to (one of the aliases of) name of the server or to the domain that it represents, if location-independent. After registration, the server MAY forward incoming SIP requests to the the network source address and port from the registration request. A server SHOULD silently drop the registration after one hour, unless refreshed by the client. A client may request and a server may indicate or lower or higher refresh interval and indicate the interval through the Expires header (Section 6.16). A single address (if host-independent) may be registered from several different clients.

If the request contains a Location header, requests for the request-URI will be directed to the address(es) given.

Support of this method is OPTIONAL.

> Beyond its use as a simple location service, this method is needed if there are several SIP servers on a single host, so that some cannot use the default port number. Each such server would register with a server for the administrative domain.

#### 4.2.6 UNREGISTER

A client cancels an existing registration established for the Request-URI with REGISTER with the UNREGISTER method. If it unregisters a Request-URI unknown to the servers, the server returns a 200 (OK) response. Support of this method is OPTIONAL.

### 4.3 Request-URI

The Request-URI field is a SIP URL as described in Section 2 or a general URI. It indicates the user or service that this request is being addressed to. Unlike the To field, the Request-URI field may be re-written by proxies. For example, a proxy may perform a lookup on the contents of the To field to resolve a username from a mail alias, and then use this username as part of the Request-URI field of requests it generates.

If a SIP server receives a request contain a URI indicating a scheme other than SIP which that server does not understand, the server MUST return a "400 Bad Request" response. It MUST do this even if the To field contains a scheme it does understand.

#### 4.3.1 SIP Version

Both request and response messages include the version of SIP in use, and basically follow [H3.1], with HTTP replaced by SIP. To be conditionally compliant with this specification, applications sending SIP messages MUST include a SIP-Version of "SIP/2.0".

### 4.4 Option Tags

Option tags are unique identifiers used to designate new options in SIP. These tags are used in Require (Section 6.24) and Unsupported (Section 6.29) fields.

Syntax:

```
option-tag  =  1*OCTET   ; LWS must be URL-escaped
```

The creator of a new SIP option should either prefix the option with a reverse domain name (e.g., "com.foo.mynewfeature" is an apt name for a feature whose inventor can be reached at "foo.com"), or register the new option with the Internet Assigned Numbers Authority (IANA).

#### 4.4.1 Registering New Option Tags with IANA

When registering a new SIP option, the following information should 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 should not contain any spaces, control characters or periods.
- 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;
- For proprietary options, contact information (postal and email address);

Borrowed from RTSP and the RTP AVP.

## 5 Response

After receiving and interpreting a request message, the recipient responds with a SIP response message. The response message format is shown below:

```
Response  =  Status-Line          ; Section 5.1
             *( general-header
             | response-header
             | entity-header )
             CRLF
             [ message-body ]     ; Section 8
```

[H6] applies except that HTTP-Version is replaced by SIP-Version. Also, SIP defines additional response codes and does not use some HTTP codes.

## 5.1 Status-Line

The first line of a Response message is the Status-Line, consisting of the protocol version ((Section 4.3.1) followed by a numeric Status-Code and its associated textual phrase, with each element separated by SP characters. No CR or LF is allowed except in the final CRLF sequence.

Status-Line = SIP-version SP Status-Code SP Reason-Phrase CRLF

### 5.1.1 Status Codes and Reason Phrases

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

We provide an overview of the Status-Code below, and provide full definitions in section 7. The first digit of the Status-Code defines the class of response. The last two digits do not have any categorization role. SIP/2.0 allows 6 values for the first digit:

**1xx:** Informational – request received, continuing process;

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

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

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

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

**6xx:** Global Failure - the request is invalid at any server.

Presented below are the individual values of the numeric response codes, and an example set of corresponding reason phrases for SIP/2.0. These reason phrases are only recommended; they may be replaced by local equivalents without affecting the protocol. Note that SIP adopts many HTTP/1.1 response codes. SIP/2.0 adds response codes in the range starting at x80 to avoid conflicts with newly defined HTTP response codes, and extends these response codes in the 6xx range.

```
Informational  =   "100"   ; Trying
               |   "180"   ; Ringing
               |   "181"   ; Queued

Success  =  "200"   ; OK
```

Figure 3: Informational and success status codes

```
Redirection  =   "300"   ; Multiple Choices
             |   "301"   ; Moved Permanently
             |   "302"   ; Moved Temporarily
             |   "303"   ; See Other
             |   "305"   ; Use Proxy
             |   "380"   ; Alternative Service
```

Figure 4: Redirection status codes

```
Status-Code      =   Informational          Fig. 3
                 |   Success                Fig. 3
                 |   Redirection            Fig. 4
                 |   Client-Error           Fig. 5
                 |   Server-Error           Fig. 6
                 |   Global-Failure         Fig. 7
                 |   extension-code
extension-code   =   3DIGIT

Reason-Phrase    =   *<TEXT, excluding CR, LF>
```

SIP response codes are extensible. SIP applications are not required to understand the meaning of all registered response codes, though such understanding is obviously desirable. However, applications MUST understand the class of any response code, as indicated by the first digit, and treat any unrecognized response as being equivalent to the x00 response code of that class, with the exception that an unrecognized response MUST NOT be cached. For example, if a client receives an unrecognized response code of 431, it can safely assume that there was something wrong with its request and treat the response as if it had received a 400 response code. In such cases, user agents SHOULD present to the user the message body returned with the response, since that message body is likely to include human-readable information which will explain the unusual status.

# 6 Header Field Definitions

SIP header fields are similar to HTTP header fields in both syntax and semantics [H4.2], [H14]. In general the ordering of the header fields is not of importance (with the exception of Via fields, see below), but

```
Client-Error  =  "400"  ; Bad Request
              |  "401"  ; Unauthorized
              |  "402"  ; Payment Required
              |  "403"  ; Forbidden
              |  "404"  ; Not Found
              |  "405"  ; Method Not Allowed
              |  "407"  ; Proxy Authentication Required
              |  "408"  ; Request Timeout
              |  "409"  ; Conflict
              |  "410"  ; Gone
              |  "411"  ; Length Required
              |  "412"  ; Precondition Failed
              |  "413"  ; Request Message Body Too Large
              |  "414"  ; Request-URI Too Large
              |  "415"  ; Unsupported Media Type
              |  "420"  ; Bad Extension
              |  "480"  ; Temporarily not available
              |  "481"  ; Invalid Call-ID
              |  "482"  ; Loop Detected
```

Figure 5: Client error status codes

```
Server-Error  =  "500"  ; Internal Server Error
              |  "501"  ; Not Implemented
              |  "502"  ; Bad Gateway
              |  "503"  ; Service Unavailable
              |  "504"  ; Gateway Timeout
              |  "505"  ; SIP Version not supported
```

Figure 6: Server error status codes

```
Global-Failure  |  "600"  ; Busy
                |  "603"  ; Decline
                |  "604"  ; Does not exist anywhere
                |  "606"  ; Not Acceptable
```

Figure 7: Global failure status Codes

proxies MUST NOT reorder or otherwise modify header fields other than by adding a new Via field. This allows an authentication field to be added after the Via fields that will not be invalidated by proxies.

The header fields required, optional and not applicable for each method are listed in Table 3. The Content-Type and Content-Length headers are required when there is a valid message body (of non-zero length) associated with the message (Section 8).

Other headers may be added as required; a server MAY ignore headers that it does not understand. A compact form of these header fields is also defined in Section 10 for use over UDP when the request has to fit into a single packet and size is an issue.

## 6.1 General Header Fields

There are a few header fields that have general applicability for both request and response messages. These header fields apply only to the message being transmitted.

General-header field names can be extended reliably only in combination with a change in the protocol version. However, new or experimental header fields may be given the semantics of general header fields if all parties in the communication recognize them to be general-header fields.

## 6.2 Entity Header Fields

Entity-header fields define meta-information about the message-body or, if no body is present, about the resource identified by the request. The term "entity header" is an HTTP 1.1 term where the reply body may contain a transformed version of the message body. The original message body is referred to as the "entity". We retain the same terminology for header fields but usually refer to the "message body" rather then the entity as the two are the same in SIP.

## 6.3 Request Header Fields

The request-header fields allow the client to pass additional information about the request, and about the client itself, to the server. These fields act as request modifiers, with semantics equivalent to the parameters on a programming language method invocation.

Request-header field names can be extended reliably only in combination with a change in the protocol version. However, new or experimental header fields MAY be given the semantics of request-header fields if all parties in the communication recognize them to be request-header fields. Unrecognized header fields are treated as entity-header fields.

| | type | ACK | BYE | INV | OPT | REG | UNR |
|---|---|---|---|---|---|---|---|
| Accept | R | o | - | o | o | o | o |
| Accept-Language | R | o | o | o | o | o | o |
| Allow | 405 | o | o | o | o | o | o |
| Also | R | - | - | o | - | - | - |
| Authorization | R | o | o | o | o | o | o |
| Call-Disposition | R | - | o | o | - | - | - |
| Call-ID | g | m | m | m | o | - | - |
| Content-Length | g | - | - | * | * | - | - |
| Content-Type | g | - | - | * | * | - | - |
| CSeq | g | o | o | o | o | o | o |
| Date | g | o | o | o | o | o | o |
| Expires | g | - | - | o | o | o | - |
| From | R | m | m | m | m | o | o |
| Location | R | - | o | - | - | o | - |
| Location | r | - | - | o | o | - | - |
| Organization | R | - | - | o | o | - | - |
| Proxy-Authenticate | R | o | o | o | o | o | o |
| Proxy-Authorization | R | o | o | o | o | o | o |
| Priority | R | - | - | o | - | - | - |
| Public | r | - | - | - | o | - | - |
| Require | R | o | o | o | o | o | o |
| Retry-After | 600,603 | - | - | o | - | - | - |
| Server | r | o | o | o | o | o | o |
| Subject | R | - | - | o | - | - | - |
| Timestamp | g | o | o | o | o | o | o |
| To | g | m | m | m | m | m | m |
| Unsupported | r | o | o | o | o | o | o |
| User-Agent | R | o | o | o | o | o | o |
| Via | g | m | m | m | m | m | m |
| Warning | r | o | o | o | o | o | o |
| WWW-Authenticate | 401 | o | o | o | o | o | o |

Table 3: Summary of header fields. “o”: optional, “m”: mandatory, “-”: not applicable, “R’: request header, “r”: response header, “g”: general header, “*”: needed if message body is not empty. A numeric value in the “type” column indicates the status code the header field is used with.

## 6.4 Response Header Fields

The response-header fields allow the server to pass additional information about the response which cannot be placed in the Status- Line. These header fields give information about the server and about further access to the resource identified by the Request-URI.

Response-header field names can be extended reliably only in combination with a change in the protocol version. However, new or experimental header fields MAY be given the semantics of response-header fields if all parties in the communication recognize them to be response-header fields. Unrecognized header fields are treated as entity-header fields.

## 6.5 Header Field Format

Header fields (general-header, request-header, response-header, and entity-header) follow the same generic header format as that given in Section 3.1 of RFC 822 [24].

Each header field consists of a name followed by a colon (":") and the field value. Field names are case-insensitive. The field value may be preceded by any amount of leading white space (LWS), though a single space (SP) is preferred. Header fields can be extended over multiple lines by preceding each extra line with at least one SP or horizontal tab (HT). Applications SHOULD follow HTTP "common form" when generating these constructs, since there might exist some implementations that fail to accept anything beyond the common forms.

```
message-header  =  field-name ":" [ field-value ] CRLF
field-name      =  token
field-value     =  *( field-content — LWS )
field-content   =  <the OCTETs making up the field-value
                   and consisting of either *TEXT or combinations
                   of token, tspecials, and quoted-string>
```

The order in which header fields are received is not significant if the header fields have different field names. Multiple header fields with the same field-name may be present in a message if and only if the entire field-value for that header field is defined as a comma-separated list (i.e., `#(values)`). It MUST be possible to combine the multiple header fields into one "field-name: field-value" pair, without changing the semantics of the message, by appending each subsequent field-value to the first, each separated by a comma. The order in which header fields with the same field-name are received is therefore significant to the interpretation of the combined field value, and thus a proxy MUST NOT change the order of these field values when a message is forwarded.

Field names are not case-sensitive, although their values may be.

## 6.6 Accept

See [H14.1] for syntax. This request header field is used only with the OPTIONS and INVITE request methods to indicate what description formats are acceptable in the response.

Example:

```
Accept: application/sdp;level=1, application/x-private
```

## 6.7 Accept-Language

See [H14.4] for syntax. The Accept-Language request header can be used to allow the client to indicate to the server in which language it would prefer to receive reason phrases. This may also be used as a hint by the proxy as to which destination to connect the call to (e.g., for selecting a human operator).

Example:

```
Accept-Language: da, en-gb;q=0.8, en;q=0.7
```

## 6.8 Allow

See [H14.7].

## 6.9 Also

The Also request header advises the callee to send invitations to the addresses listed. This supports third-party call initiation (Section 13).

```
Also = "Also" ":" 1#( SIP-URL ) [ comment ]
```

Example:

```
Also: sip://jones@foo.com, sip://mueller@bar.edu
```

## 6.10 Authorization

See [H14.8].

### 6.11 Call-Disposition

The Call-Disposition request header field allows the client to indicate how the server is to handle the call. The following options can be used singly or in combination:

**all:** If the *user* part of the SIP request address identifies a group rather than an individual, the "all" feature indicates that all members of the group should be alerted rather than the default of locating the first available individual from that group. Section 1.4.1 describes the behavior of proxy servers when resolving group aliases.

**do-not-forward:** The "do-not-forward" request prohibits proxies from forwarding the call to another individual (e.g., the call is personal or the caller does not want to be shunted to a secretary if the line is busy.)

**queue:** If the called party is temporarily unreachable, e.g., because it is in another call, the caller can indicate that it wants to have its call queued rather than rejected immediately. If the call is queued, the server returns "181 Queued" (see Section 7.1.3). A pending call be terminated by a BYE request (Section 4.2.4).

```
Call-Disposition  =  "Call-Disposition" ":" 1#( "all" | "do-not-forward"
                  |  "queue" )
```

Example:

```
Call-Disposition: all, do-not-forward, queue
```

HS: This header is experimental. The name is based on the SMTP Content-Disposition header.

### 6.12 Call-ID

The Call-ID general header uniquely identifies a particular invitation. Note that a single multimedia conference may give rise to several calls with different Call-IDs, e.g., if a user invites several different people. Since the Call-ID is unique for each caller, a user may invited to the same conference using several different Call-IDs. If desired, it must use identifiers within the session description to detect this duplication. Calls to different callee MUST always use different Call-IDs unless they are the result of a proxy server "forking" a single request.

The Call-ID may be any URL-encoded string that can be guaranteed to be globally unique for the duration of the request. Using the initiator's IP-address, process id, and instance (if more than one request is being made simultaneously) satisfies this requirement.

The form local-id@host is recommended, where host is either the fully qualified domain name or a globally routable IP address, and local-id depends on the application and operating system of the host, but is an ID that can be guaranteed to be unique during this session initiation request.

Call-ID = ( ”Call-ID” | ”i” ) ”:” atom ”@” host

Example:

```
Call-ID: 9707211351.AA08181@foo.bar.com
```

### 6.13 Content-Length

The Content-Length entity-header field indicates the size of the message-body, in decimal number of octets, sent to the recipient.

Content-Length = ”Content-Length” ”:” 1*DIGIT

An example is

```
Content-Length: 3495
```

Applications SHOULD use this field to indicate the size of the message-body to be transferred, regardless of the media type of the entity. Any Content-Length greater than or equal to zero is a valid value. If no body is present in a message, then the Content-Length header MAY be omitted or set to zero. Section 8 describes how to determine the length of the message body.

### 6.14 Content-Type

The Content-Type entity-header field indicates the media type of the message-body sent to the recipient.

Content-Type = ”Content-Type” ”:” media-type

An example of the field is

```
Content-Type: application/sdp
```

### 6.15 Date

General header field. See [H14.19].

The Date header field is useful for simple devices without their own clock.

### 6.16 Expires

The Expires entity-header field gives the date and time after which the message content expires.

This header field is currently defined only for the REGISTER and INVITE methods. For REGISTER, it is a request and response-header field and allows the client to indicate how long the registration should be valid; the server uses it to indicate when the client has to re-register. The server's choice overrides that of the client. The server MAY choose a shorter time interval than that requested by the client, but SHOULD not choose a longer one.

For INVITE, it is a request and response-header field. In a request, the callee can limit the validity of an invitation. (For example, if a client wants to limit how long a search should take at most or when a conference being invited to is time-limited. A user interface may take this is as a hint to leave the invitation window on the screen even if the user is not currently at the workstation.) In a 302 response, a server can advise the client of the maximal duration of the redirection.

The value of this field can be either an HTTP-date or an integer number of seconds (in decimal), measured from the receipt of the request.

Expires = "Expires" ":" ( HTTP-date | delta-seconds )

Two example of its use are

```
Expires: Thu, 01 Dec 1994 16:00:00 GMT
Expires: 5
```

### 6.17 From

Requests MUST and responses SHOULD contain a From header field, indicating the invitation initiator. The field MUST be a SIP URL as defined in Section 2. Only a single initiator and a single invited user are allowed to be specified in a single SIP request. The sense of To and From header fields is maintained from request to response, i.e., if the From header is `sip://bob@example.edu` in the request then it is MUST also be `sip://bob@example.edu` in the response to that request.

The From field is a URL and not a simple SIP address (Section 1.4.1 address to allow a gateway to relay a call into a SIP request and still produce an appropriate From field.

From = ( "From" | "f" ) ":" *1( ( SIP-URL | URL ) [ comment ] )

Examples:

```
From: agb@bell-telephone.com (A. G. Bell)
From: +12125551212@server.phone2net.com
```

### 6.18 Location

The Location response header can be used with a 2xx or 3xx response codes to indicate a new location to try. It contains a URL giving the new location or username to try, or may simply specify additional transport parameters. A "301 Moved Permanently" or "302 Moved Temporarily" response SHOULD contain a Location field containing the URL giving a new address to try. A 301 or 302 response may also give the same location and username that was being tried but specify additional transport parameters such as a multicast address to try or a change of SIP transport from UDP to TCP or vice versa.

A user agent or redirect server sending a definitive, positive response (2xx), SHOULD insert a Location response header indicating the SIP address under which it is reachable most directly for future SIP requests. This may be the address of the server itself or that of a proxy (e.g., if the host is behind a firewall).

A Location response header may contain any suitable URL indicating where the called party may be reached, not limited to SIP URLs. For example, it may contain a phone or fax URL [25], a mailto: URL [26] or irc.

The following parameters are defined:

**q:** The qvalue indicates the relative preference among the locations given. qvalue values are decimal numbers from 0.0 to 1.0, with higher values indicating higher preference.

**class:** The class parameter whether this terminal is found in a residential or business setting. (A caller may defer a personal call if only a business line is available, for example.)

**description:** The description field further describes, as text, the terminal. It is expected that the user interface will render this text.

**duplex:** The duplex parameter lists whether the terminal can simultaneously send and receive ("full"), alternate between sending and receiving ("half"), can only receive ("receive-only") or only send ("send-only"). Typically, a caller will prefer a full-duplex terminal over a half-duplex terminal and these over receive-only or send-only terminals.

**features:** The feature list enumerates additional features of this terminal. Values for this field are for further study.

**language:** The language parameter lists, in order of preference, the languages spoken by the person answering. This feature may be used to have a caller automatically select the appropriate attendant or customer service representative, without having to declare its own language skills.

**media:** The media tag lists the media types supported by the terminal. Currently, the names defined in SDP may be used [9]: "audio", "video", "whiteboard", "text" and "data".

**mobility:** The mobility parameter indicates if the terminal is fixed or mobile. In some locales, this may affect voice quality or charges.

**priority:** The priority tag indicates the minimum priority level this terminal is to be used for. It can be used for automatically restricting the choice of terminals available to the user.

**service:** The service tag describes what service is being provided by the terminal.

```
Location            =  ( "Location" | "m" ) ( SIP-URL | URL )
                       *( ";" location-params )
extension-name      =  token
extension-value     =  *( token | quoted-string | LWS | extension-specials)
extension-specials  =  < any element of tspecials except <"> >
language-tag        =  < see [H3.10] >
priority-tag        =  "urgent" | "normal" | "non-urgent"
service-tag         =  "fax" | "IP" | "PSTN" | "ISDN" | "pager"
media-tag           =  < see SDP: "audio" | "video" | "email" ...
feature-list        =  "voice-mail" | "attendant"

location-params     =  "q"                  "="  qvalue
                    |  "class"              "="  ( "personal" | "business" )
                    |  "description"        "="  quoted-string
                    |  "duplex"             "="  ( "full" | "half" |
                                                 "receive-only" | "send-only" )
                    |  "features"           "="  1# feature-list
                    |  "language"           "="  1# language-tag
                    |  "media"              "="  1# media-tag
                    |  "mobility"           "="  ( "fixed" | "mobile" )
                    |  "priority"           "="  1# priority-tag
                    |  "service"            "="  1# service-tag
                    |  extension-attributes
extension-attribute =  extension-name       "="  extension-value
```

Examples:

```
Location: sip://watson@worcester.bell-telephone.com ;service=IP,voice-mail
          ;media=audio ;duplex=full ;q=0.7;
Location: phone://1-415-555-1212 ; service=ISDN;mobility=fixed;
          language=en,es,iw ;q=0.5
Location: phone://1-800-555-1212 ; service=pager;mobility=mobile;
          duplex=send-only;media=text; q=0.1; priority=urgent;
          description="For emergencies only"
Location: mailto:watson@bell-telephone.com
Location: http://www.bell-telephone.com/~watson
```

Attributes which are unknown should be omitted. New tags for class-tag and service-tag can be registered with IANA. The media tag uses Internet media types, e.g., audio, video, application/x-wb, etc. This is meant for indicating general communication capability, sufficient for the caller to choose an appropriate address.

### 6.19 Organization

The Organization request-header fields conveys the name of the organization to which the callee belongs. It may be inserted by proxies at the boundary of an organization and may be used by client software to filter calls.

### 6.20 Priority

The priority request header signals the urgency of the call to the callee.

```
Priority        =  "Priority" ":" priority-value
priority-value  =  "urgent" | "normal" | "non-urgent"
```

Example:

```
Subject: A tornado is heading our way!
Priority: urgent
```

### 6.21 Proxy-Authenticate

See [H14.33].

### 6.22 Proxy-Authorization

See [H14.34].

### 6.23 Public

See [H14.35].

### 6.24 Require

The Require header is used by clients to query the server about options that it may or may not support. The server MUST respond to this header by returning status code "420 Bad Extension" and list those options it does not understand in the Unsupported header.

```
Require  =  "Require" ":" 1#option-tag
```

Example:

```
C->S:   INVITE sip:watson@bell-telephone.com SIP/2.0
        Require: com.example.billing
        Payment: sheep_skins, conch_shells

S->C:   SIP/2.0 420 Bad Extension
        Unsupported: com.example.billing
```

This is to make sure that the client-server interaction will proceed optimally when all options are understood by both sides, and only slow down if options are not understood (as in the example above). For a well-matched client-server pair, the interaction proceeds quickly, saving a round-trip often required by negotiation mechanisms. In addition, it also removes ambiguity when the client requires features that the server does not understand.

> We explored using the W3C's PEP proposal [?] for this functionality. However, Require, Proxy-Require, and Unsupported allow the addition of extensions with far less complexity.
>
> This field roughly corresponds to the PEP field in the PEP draft.

### 6.25 Retry-After

The Retry-After response header field can be used with a "503 Service Unavailable" response to indicate how long the service is expected to be unavailable to the requesting client and with a "404 Not Found", "600 Busy", "603 Decline" response to indicate when the called party may be available again. The value of this field can be either an HTTP-date or an integer number of seconds (in decimal) after the time of the response. An optional comment can be used to indicate additional information about the time of callback. An optional duration parameter indicates how long the called party will be reachable starting at the initial time of availability.

```
Retry-After  =  "Retry-After" ":" ( HTTP-date | delta-seconds )
                [ comment ] [ ";duration" "=" delta-seconds
```

Examples of its use are

```
Retry-After: Mon, 21 Jul 1997 18:48:34 GMT (I'm in a meeting)
Retry-After: Mon,  1 Jan 9999 00:00:00 GMT
  (Dear John: Don't call me back, ever)
Retry-After: Fri, 26 Sep 1997 21:00:00 GMD;duration=3600
Retry-After: 120
```

In the third example, the callee is reachable for one hour starting at 21:00 GMT. In the last example, the delay is 2 minutes.

## 6.26 CSeq

The CSeq (command sequence) header field MAY be added by a SIP client making a request if it needs to distinguish responses to several consecutive requests sent with the same Call-ID. A CSeq field contains a single decimal sequence number chosen by the requesting client. Consecutive different requests made with the same Call-ID MUST contain strictly monotonically increasing sequence numbers; the sequence space MAY NOT be contiguous. Retransmissions of the same request carry the same sequence number. A server responding to a request containing a sequence number MUST echo the sequence number back in the response. The ACK request MUST contain the same CSeq value as the INVITE request that it refers to.

CSeq = "CSeq" ":" 1*DIGIT

CSeq header fields are NOT needed for SIP requests using the INVITE or OPTIONS methods but may be needed for future methods.

Example:

```
CSeq: 4711
```

## 6.27 Server

See [H14.39].

## 6.28 Subject

This is intended to provide a summary, or indicate the nature, of the call, allowing call filtering without having to parse the session description. (Also, the session description may not necessarily use the same subject indication as the invitation.)

Subject = ( "Subject" | "s" ) ":" *text

Example:

```
Subject: Tune in - they are talking about your work!
```

## 6.29 Unsupported

The Unsupported response header lists the features not supported by the server.

See Section 6.24 for a usage example and motivation.

### 6.30 Timestamp

The timestamp general header describes when the client sent the request to the server. The value of the timestamp is of significance only to the client and may use any timescale. The server MUST echo the exact same value and MAY, if it has accurate information about this, add a floating point number indicating the number of seconds that has elapsed since it has received the request. The timestamp is used by the client to compute the round-trip time to the server so that it can adjust the timeout value for retransmissions.

```
Timestamp = "Timestamp" ":" *(DIGIT) [ "." *(DIGIT) ] [ delay ]
delay     = *(DIGIT) [ "." *(DIGIT) ]
```

### 6.31 To

The To request header field specifies the invited user, with the same SIP URL syntax as the From field.

```
To = ( "To" | "t" ) ":" ( SIP-URL | URL ) [ comment ]
```

If a SIP server receives a request destined To a URL indicating a scheme other than SIP and that is unknown to it, the server returns a “400 bad request” response.

Example:

```
To: sip://operator@cs.columbia.edu (The Operator)
```

### 6.32 User-Agent

See [H14.42].

### 6.33 Via

The Via field indicates the path taken by the request so far. This prevents request looping and ensures replies take the same path as the requests, which assists in firewall traversal and other unusual routing situations.

The client originating the request MUST insert a Via field containing its address to the request. Each subsequent proxy server that sends the request onwards MUST add its own additional Via field, which MUST be added before any existing Via fields. Additionally, if the message goes to a multicast address, an extra Via field is added before all the others giving the multicast address and TTL.

If a proxy server receives a request which contains its own address, it MUST respond with a “482 Loop Detected” status code. (This prevents a malfunctioning proxy server from causing loops. Also, it cannot be guaranteed that a proxy server can always detect that the address returned by a location service refers to a host listed in the Via list, as a single host may have aliases or several network interfaces.)

In the return path, Via fields are processed by a proxy or client according to the following rules:

- If the first Via field in the reply **received** is the client's or server's local address, remove the Via field and process the reply.
- If the first Via field in a reply is a multicast address, remove that Via field **before** sending to the multicast address.

These rules ensure that a proxy server only has to check the first Via field in a reply to see if it needs processing.

The format for a Via header is:

```
Via               =  ( "Via" | "v") ":" 1#( sent-protocol sent-by
                     *( ";" via-params ) [ comment ] )
via-params        =  "ttl" "=" ttl
sent-protocol     =  [ protocol-name "/" ] protocol-version
 [ "/" transport ]
protocol-name     =  "SIP" | token
protocol-version  =  token
transport         =  "UDP" | "TCP"
sent-by           =  host [ ":" port ]
ttl               =  1*3DIGIT                            ; 0 to 255
```

The "ttl" parameter is included only if the address is a multicast address.

Example:

```
Via: SIP/2.0/UDP first.example.com:4000
```

## 6.34 Warning

The Warning response-header field is used to carry additional information about the status of a response. Warning headers are sent with responses and have the following format:

```
Warning        =  "Warning" ":" 1#warning-value
warning-value  =  warn-code SP warn-agent SP warn-text
warn-code      =  2DIGIT
warn-agent     =  ( host [ ":" port ] ) | pseudonym
                  ; the name or pseudonym of the server adding
                  ; the Warning header, for use in debugging
warn-text      =  quoted-string
```

A response may carry more than one Warning header.

The warn-text should be in a natural language and character set that is most likely to be intelligible to the human user receiving the response. This decision may be based on any available knowledge, such as the location of the cache or user, the Accept-Language field in a request, the Content-Language field in a response, etc. The default language is English.

Any server may add Warning headers to a response. New Warning headers should be added after any existing Warning headers. A proxy server MUST NOT delete any Warning header that it received with a response.

When multiple Warning headers are attached to a response, the user agent SHOULD display as many of them as possible, in the order that they appear in the response. If it is not possible to display all of the warnings, the user agent should follow these heuristics:

- Warnings that appear early in the response take priority over those appearing later in the response.
- Warnings in the user's preferred character set take priority over warnings in other character sets but with identical warn-codes and warn-agents.

Systems that generate multiple Warning headers should order them with this user agent behavior in mind.

Example:

```
Warning: 606.4 isi.edu Multicast not available
Warning: 606.2 isi.edu Incompatible protocol (RTP/XXP)
```

### 6.35 WWW-Authenticate

See [H14.46].

## 7 Status Code Definitions

The response codes are consistent with, and extend, HTTP/1.1 response codes. Not all HTTP/1.1 response codes are appropriate, and only those that are appropriate are given here. Response codes not defined by HTTP/1.1 have codes x80 upwards to avoid clashes with future HTTP response codes. Also, SIP defines a new class, 6xx. The default behavior for unknown response codes is given for each category of codes.

### 7.1 Informational 1xx

Informational responses indicate that the server or proxy contacted is performing some further action and does not yet have a definitive response. The client SHOULD wait for a further response from the server, and the server SHOULD send such a response without further prompting. If UDP transport is being used, the client SHOULD periodically re-send the request in case the final response is lost. Typically a server should send a “1xx” response if it expects to take more than one second to obtain a final reply.

#### 7.1.1 100 Trying

Some further action is being taken (e.g., the request is being forwarded) but the user has not yet been located.

#### 7.1.2 180 Ringing

The user agent or conference server has located a possible location where the user has been recently and is trying to alert them.

#### 7.1.3 181 Queued

The called party was temporarily unavailable, but the caller indicated via a “Call-Disposition: Queue” directive (Section 6.11) to queue the call rather than reject it. When the callee becomes available, it will return the appropriate final status response. The reason phrase MAY give further details about the status of the call, e.g., “5 calls queued; expected waiting time is 15 minutes”. The server may issue several 181 responses to update the caller about the status of the queued call.

### 7.2 Successful 2xx

The request was successful and MUST terminate a search.

#### 7.2.1 200 OK

The request was successful in contacting the user, and the user has agreed to participate.

### 7.3 Redirection 3xx

3xx responses give information about the user’s new location, or about alternative services that may be able to satisfy the call. They SHOULD terminate an existing search, and MAY cause the initiator to begin a new search if appropriate.

#### 7.3.1 300 Multiple Choices

The requested resource corresponds to any one of a set of representations, each with its own specific location, and agent-driven negotiation (i.e., controlled by the SIP client) is being provided so that the user (or user agent) can select a preferred communication end point and redirect its request to that location.

The response SHOULD include an entity containing a list of resource characteristics and location(s) from which the user or user agent can choose the one most appropriate. The entity format is specified by the media type given in the Content-Type header field. Depending upon the format and the capabilities of the user agent, selection of the most appropriate choice may be performed automatically. However, this specification does not define any standard for such automatic selection.

The choices SHOULD also be listed as Location fields (Section 6.18). Unlike HTTP, the SIP response may contain several Location fields. User agents MAY use the Location field value for automatic redirection or MAY ask the user to confirm a choice.

#### 7.3.2 301 Moved Permanently

The requesting client should retry on the new address given by the Location field because the user has permanently moved and the address this response is in reply to is no longer a current address for the user. A 301 response MUST NOT suggest any of the hosts in the Via (Section 6.33) path of the request as the user's new location.

#### 7.3.3 302 Moved Temporarily

The requesting client should retry on the new address(es) given by the Location header. A 302 response MUST NOT suggest any of the hosts in the Via (Section 6.33) path of the request as the user's new location. The duration of the redirection can be indicated through an Expires (Section 6.16) header.

#### 7.3.4 380 Alternative Service

The call was not successful, but alternative services are possible. The alternative services are described in the message body of the response.

### 7.4 Request Failure 4xx

4xx responses are definite failure responses from a particular server. The client SHOULD NOT retry the same request without modification (e.g., adding appropriate authorization). However, the same request to a different server may be successful.

### 7.4.1 400 Bad Request

The request could not be understood due to malformed syntax.

### 7.4.2 401 Unauthorized

The request requires user authentication.

### 7.4.3 402 Payment Required

Reserved for future use.

### 7.4.4 403 Forbidden

The server understood the request, but is refusing to fulfill it. Authorization will not help, and the request should not be repeated.

### 7.4.5 404 Not Found

The server has definitive information that the user does not exist at the domain specified in the Request-URI.

### 7.4.6 405 Method Not Allowed

The method specified in the Request-Line is not allowed for the address identified by the Request-URI. The response MUST include an Allow header containing a list of valid methods for the indicated address.

### 7.4.7 407 Proxy Authentication Required

This code is similar to 401 (Unauthorized), but indicates that the client MUST first authenticate itself with the proxy. The proxy MUST return a Proxy-Authenticate header field (section 6.21) containing a challenge applicable to the proxy for the requested resource. The client MAY repeat the request with a suitable Proxy-Authorization header field (section 6.22). SIP access authentication is explained in section [H11].

This status code should be used for applications where access to the communication channel (e.g., a telephony gateway) rather than the callee herself requires authentication.

#### 7.4.8 408 Request Timeout

The client did not produce a request within the time that the server was prepared to wait. The client MAY repeat the request without modifications at any later time.

#### 7.4.9 420 Bad Extension

The server did not understand the protocol extension specified with strength “must”.

#### 7.4.10 480 Temporarily Unavailable

The callee’s end system was contacted successfully but the callee is currently unavailable (e.g., not logged in or logged in in such a manner as to preclude communication with the callee). The response may indicate a better time to call in the Retry-After header. The user may also be available elsewhere (unbeknownst to this host), thus, this response does not terminate any searches. The reason phrase SHOULD indicate the more precise cause as to why the callee is unavailable. This value SHOULD be setable by the user agent.

#### 7.4.11 481 Invalid Call-ID

The server received a BYE or ACK request with a Call-ID value it does not recognize.

#### 7.4.12 482 Loop Detected

The server received a request with a Via path containing itself.

### 7.5 Server Failure 5xx

5xx responses are failure responses given when a server itself has erred. They are not definitive failures, and SHOULD NOT terminate a search if other possible locations remain untried.

#### 7.5.1 500 Server Internal Error

The server encountered an unexpected condition that prevented it from fulfilling the request.

#### 7.5.2 501 Not implemented

The server does not support the functionality required to fulfill the request. This is the appropriate response when the server does not recognize the request method and is not capable of supporting it for any user.

#### 7.5.3 502 Bad Gateway

The server, while acting as a gateway or proxy, received an invalid response from the upstream server it accessed in attempting to fulfill the request.

#### 7.5.4 503 Service Unavailable

The server is currently unable to handle the request due to a temporary overloading or maintenance of the server. The implication is that this is a temporary condition which will be alleviated after some delay. If known, the length of the delay may be indicated in a Retry-After header. If no Retry-After is given, the client SHOULD handle the response as it would for a 500 response.

Note: The existence of the 503 status code does not imply that a server must use it when becoming overloaded. Some servers may wish to simply refuse the connection.

#### 7.5.5 504 Gateway Timeout

The server, while acting as a gateway, did not receive a timely response from the upstream server (e.g., a location server) it accessed in attempting to complete the request.

### 7.6 Global Failures 6xx

6xx responses indicate that a server has definitive information about a particular user, not just the particular instance indicated in the Request-URI. All further searches for this user are doomed to failure and pending searches SHOULD be terminated.

#### 7.6.1 600 Busy

The callee's end system was contacted successfully but the callee is busy and does not wish to take the call at this time. The response may indicate a better time to call in the Retry-After header. If the callee does not wish to reveal the reason for declining the call, the callee should use status code 680 instead.

### 7.6.2 603 Decline

The callee's machine was successfully contacted but the user explicitly does not wish to participate. The response may indicate a better time to call in the Retry-After header.

### 7.6.3 604 Does not exist anywhere

The server has authoritative information that the user indicated in the To request field does not exist anywhere. Searching for the user elsewhere will not yield any results.

### 7.6.4 606 Not Acceptable

The user's agent was contacted successfully but some aspects of the session profile (the requested media, bandwidth, or addressing style) were not acceptable.

A "606 Not Acceptable" reply means that the user wishes to communicate, but cannot adequately support the session described. The "606 Not Acceptable" reply MAY contain a list of reasons in a Warning header describing why the session described cannot be supported. These reasons can be one or more of:

**606.1 Insufficient Bandwidth:** The bandwidth specified in the session description or defined by the media exceeds that known to be available.

**606.2 Incompatible Protocol:** One or more protocols described in the request are not available.

**606.3 Incompatible Format:** One or more media formats described in the request is not available.

**606.4 Multicast not available:** The site where the user is located does not support multicast.

**606.5 Unicast not available:** The site where the user is located does not support unicast communication (usually due to the presence of a firewall).

Other reasons are likely to be added later. It is hoped that negotiation will not frequently be needed, and when a new user is being invited to join a pre-existing lightweight session, negotiation may not be possible. It is up to the invitation initiator to decide whether or not to act on a "606 Not Acceptable" reply.

## 8 SIP Message Body

The session description body gives details of the session the user is being invited to join. Its Internet media type MUST be given by the Content-type header field, and the body length in bytes MUST be given by the Content-Length header field. If the body has undergone any encoding (such as compression) then this MUST be indicated by the Content-encoding header field, otherwise Content-encoding MUST be omitted.

### 8.1 Body Inclusion

For a request message, the presence of a body is signaled by the inclusion of a Content-Length header. A body may be included in a request only when the request method allows one.

For response messages, whether or not a body is included is dependent on both the request method and the response message's response code. All 1xx informational responses MUST NOT include a body. All other responses MAY include a payload, although it may be of zero length.

### 8.2 Message Body Length

If no body is present in a message, then the Content-Length header MAY be omitted or set to zero. When a body is included, its length in bytes is indicated in the Content-Length header and is determined by one of the following:

1. Any response message which MUST NOT include a body (such as the 1xx responses) is always terminated by the first empty line after the header fields, regardless if any entity-header fields are present.
2. Otherwise, a Content-Length header MUST be present (this requirement differs from HTTP/1.1). Its value in bytes represents the length of the message body.

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

## 9 Examples

### 9.1 Invitation to Multimedia Conference

The first example invites `schooler@vlsi.cs.caltech.edu` to a multicast session.

#### 9.1.1 Request

```
C->S: INVITE schooler@vlsi.cs.caltech.edu SIP/2.0
      Via: SIP/2.0/UDP 239.128.16.254 16
      Via: SIP/2.0/UDP 131.215.131.131
      Via: SIP/2.0/UDP 128.16.64.19
      From: mjh@isi.edu (Mark Handley)
      Subject: SIP will be discussed, too
      To: schooler@cs.caltech.edu (Eve Schooler)
      Call-ID: 62729-27@oregon.isi.edu
```

```
       Content-type: application/sdp
       CSeq: 4711
       Content-Length: 187

       v=0
       o=user1 53655765 2353687637 IN IP4 128.3.4.5
       s=Mbone Audio
       i=Discussion of Mbone Engineering Issues
       e=mbone@somewhere.com
       c=IN IP4 224.2.0.1/127
       t=0 0
       m=audio 3456 RTP/AVP 0
```

The Via fields list the hosts along the path from invitation initiator (the first element of the list) towards the invitee. In the example above, the message was last multicast to the administratively scoped group `239.128.16.254` with a ttl of `16` from the host `131.215.131.131`.

The request header above states that the request was initiated by `mjh@isi.edu`. The Via header indicates that is was initiated from the host `128.16.64.19`. `schooler@cs.caltech.edu` is being invited; the message is currently being routed to `schooler@vlsi.cs.caltech.edu`.

In this case, the session description is using the Session Description Protocol (SDP), as stated in the Content-type header.

The header is terminated by an empty line and is followed by a message body containing the session description.

#### 9.1.2 Reply

The called user agent, directly or indirectly through proxy servers, indicates that it is alerting ("ringing") the called party:

```
S->C: SIP/2.0 180 Ringing
      Via: SIP/2.0/UDP 239.128.16.254 16
      Via: SIP/2.0/UDP 131.215.131.131
      Via: SIP/2.0/UDP 128.16.64.19 1
      From: mjh@isi.edu
      Call-ID: 62729-27@128.16.64.19
      Location: sip://es@jove.cs.caltech.edu
      CSeq: 4711
```

A sample reply to the invitation is given below. The first line of the reply states the SIP version number, that it is a "200 OK" reply, which means the request was successful. The Via headers are taken from the request, and entries are removed hop by hop as the reply retraces the path of the request. A new authentication field MAY be added by the invited user's agent if required. The Call-ID is taken directly from the original request, along with the remaining fields of the request message. The original sense of From field is preserved (i.e., it is the session initiator).

In addition, the Location header gives details of the host where the user was located, or alternatively the relevant proxy contact point which should be reachable from the caller's host.

```
S->C: SIP/2.0 200 OK
      Via: SIP/2.0/UDP 239.128.16.254 16
      Via: SIP/2.0/UDP 131.215.131.131
      Via: SIP/2.0/UDP 128.16.64.19 1
      From: mjh@isi.edu
      To: schooler@cs.caltech.edu
      Call-ID: 62729-27@128.16.64.19
      Location: sip://es@jove.cs.caltech.edu
      CSeq: 4711
```

The caller confirms the invitation by sending a request to the location named in the Location header:

```
C->S: ACK schooler@jove.cs.caltech.edu SIP/2.0
      From: mjh@isi.edu
      To: schooler@cs.caltech.edu
      Call-ID: 62729-27@128.16.64.19
      CSeq: 4711
```

## 9.2 Two-party Call

A two-party call proceeds as above. The only difference is

For two-party Internet phone calls, the response must contain a description of where to send data to. In the example below, Bell calls Watson. Bell indicates that he can receive RTP audio codings 0 (PCMU), 3 (GSM), 4 (G.723) and 5 (DVI4).

```
C->S: INVITE watson@boston.bell-telephone.com SIP/2.0
      Via: SIP/2.0/UDP 169.130.12.5
      From: a.g.bell@bell-telephone.com (A. Bell)
      To: watson@bell-telephone.com (T. A. Watson)
      Call-ID: 187602141351@worcester.bell-telephone.com
      Subject: Mr. Watson, come here.
      Content-type: application/sdp
```

```
       Content-Length: ...

       v=0
       o=bell 53655765 2353687637 IN IP4 128.3.4.5
       c=IN IP4 135.180.144.94
       m=audio 3456 RTP/AVP 0 3 4 5

S->C: SIP/2.0 200 OK
       From: a.g.bell@bell-telephone.com (A. Bell)
       To: watson@bell-telephone.com
       Call-ID: 187602141351@worcester.bell-telephone.com
       Location: sip://watson@boston.bell-telephone.com
       Content-Length: ...

       v=0
       o=watson 4858949 4858949 IN IP4 192.1.2.3
       c=IN IP4 135.180.161.25
       m=audio 5004 RTP/AVP 0 3
```

Watson can only receive PCMU and GSM. Note that Watson's list of codecs may or may not be a subset of the one offered by Bell, as each party indicates the data types it is willing to receive. Watson will send audio data to port 3456 at 135.180.144.94, Bell will send to port 5004 at 135.180.161.25.

## 9.3 Aborting a Call

If the caller wants to abort a pending call, it sends a BYE request.

```
C->S: BYE schooler@jove.cs.caltech.edu
       From: mjh@isi.edu
       To: schooler@cs.caltech.edu
       Call-ID: 62729-27@128.16.64.19
```

### 9.3.1 Redirects

Replies with status codes "301 Moved Permanently" or "302 Moved Temporarily" SHOULD specify another location using the Location field.

```
S->C: SIP/2.0 302 Moved temporarily
       Via: SIP/2.0/UDP 131.215.131.131
       Via: SIP/2.0/UDP 128.16.64.19
       From: mjh@isi.edu
```

```
       To: schooler@cs.caltech.edu
       Call-ID: 62729-27@128.16.64.19
       Location: sip://239.128.16.254;ttl=16;transport=udp
       Content-length: 0
```

In this example, the proxy located at `131.215.131.131` is being advised to contact the multicast group `239.128.16.254` with a ttl of `16` and UDP transport. In normal situations, a server would not suggest a redirect to a local multicast group unless, as in the above situation, it knows that the previous proxy or client is within the scope of the local group. If the request is redirected to a multicast group, a proxy server SHOULD query the multicast address itself rather than sending the redirect back towards the client as multicast may be scoped; this allows a proxy within the appropriate scope regions to make the query.

#### 9.3.2 Alternative Services

An example of a "350 Alternative Service" reply is:

```
S->C: SIP/2.0 350 Alternative Service
       Via: SIP/2.0/UDP 131.215.131.131
       Via: SIP/2.0/UDP 128.16.64.19
       From: mjh@isi.edu
       To: schooler@cs.caltech.edu
       Call-ID: 62729-27@128.16.64.19
       Location: recorder@131.215.131.131
       Content-type: application/sdp
       Content-length: 146

       v=0
       o=mm-server 2523535 0 IN IP4 131.215.131.131
       s=Answering Machine
       i=Leave an audio message
       c=IN IP4 131.215.131.131
       t=0 0
       m=audio 12345 RTP/AVP 0
```

In this case, the answering server provides a session description that describes an "answering machine". If the invitation initiator decides to take advantage of this service, it should send an invitation request to the answering machine at `131.215.131.131` with the session description provided (modified as appropriate for a unicast session to contain the appropriate local address and port for the invitation initiator). This request SHOULD contain a different Call-ID from the one in the original request. An example would be:

```
C->S: INVITE mm-server@131.215.131.131 SIP/2.0
       Via: SIP/2.0/UDP 128.16.64.19
```

```
       From: mjh@isi.edu
       To: schooler@cs.caltech.edu
       Call-ID: 62729-28@128.16.64.19
       Content-type: application/sdp
       Content-length: 146

       v=0
       o=mm-server 2523535 0 IN IP4 131.215.131.131
       s=Answering Machine
       i=Leave an audio message
       c=IN IP4 128.16.64.19
       t=0 0
       m=audio 26472 RTP/AVP 0
```

Invitation initiators MAY choose to treat a “350 Alternative Service” reply as a failure if they wish to do so.

### 9.3.3 Negotiation

An example of a “606 Not Acceptable” reply is:

```
S->C: SIP/2.0 606 Not Acceptable
       From: mjh@isi.edu
       To: schooler@cs.caltech.edu
       Call-ID:62729-27@128.16.64.19
       Location: mjh@131.215.131.131
       Warning: 606.1 Insufficient bandwidth (only have ISDN),
         606.3 Incompatible format,
         606.4 Multicast not available
       Content-Type: application/sdp
       Content-Length: 50

       v=0
       s=Lets talk
       b=CT:128
       c=IN IP4 131.215.131.131
       m=audio 3456 RTP/AVP 7 0 13
       m=video 2232 RTP/AVP 31
```

In this example, the original request specified 256 kb/s total bandwidth, and the reply states that only 128 kb/s is available. The original request specified GSM audio, H.261 video, and WB whiteboard. The audio coding and whiteboard are not available, but the reply states that DVI, PCM or LPC audio could be supported in order of preference. The reply also states that multicast is not available. In such a case, it might be appropriate to set up a transcoding gateway and re-invite the user.

### 9.4 OPTIONS Request

A caller Alice can use an OPTIONS request to find out the capabilities of a potential callee Bob, without "ringing" the designated address. In this case, Bob indicates that he can be reached at three different addresses, ranging from voice-over-IP to a PSTN phone to a pager.

```
C->S: OPTIONS bob@example.com SIP/2.0
      From: alice@anywhere.org (Alice)
      To: bob@example.com (Bob)
      Accept: application/sdp

S->C: SIP/2.0 200 OK
      Location: sip://bob@host.example.com ;service=IP,voice-mail
                ;media=audio ;duplex=full ;q=0.7
      Location: phone://1-415-555-1212 ; service=ISDN;mobility=fixed;
                language=en,es,iw ;q=0.5
      Location: phone://1-800-555-1212 ; service=pager;mobility=mobile;
                duplex=send-only;media=text; q=0.1
```

Alternatively, Bob could have returned a description of

```
C->S: OPTIONS bob@example.com SIP/2.0
      From: alice@anywhere.org (Alice)
      To: bob@example.com (Bob)
      Accept: application/sdp

S->C: SIP/2.0 200 OK
      Content-Length: 81
      Content-Type: application/sdp

      v=0
      m=audio 0 RTP/AVP 0 1 3 99
      m=video 0 RTP/AVP 29 30
      a:rtpmap:98 SX7300/8000
```

## 10 Compact Form

When SIP is carried over UDP with authentication and a complex session description, it may be possible that the size of a request or reply is larger than the MTU. To reduce this problem, a more compact form of SIP is also defined by using alternative names for common header fields. These short forms are NOT abbreviations, they are field names. No other abbreviations are allowed.

| short field name | long field name | note |
|---|---|---|
| c | Content-Type | |
| e | Content-Encoding | |
| f | From | |
| i | Call-ID | |
| l | Content-Length | |
| m | Location | from “moved” |
| s | Subject | |
| t | To | |
| v | Via | |

Thus the header in section 9.1 could also be written:

```
INVITE schooler@vlsi.caltech.edu SIP/2.0
v:SIP/2.0/UDP 239.128.16.254 16
v:SIP/2.0/UDP 131.215.131.131
v:SIP/2.0/UDP 128.16.64.19
f:mjh@isi.edu
t:schooler@cs.caltech.edu
i:62729-27@128.16.64.19
c:application/sdp
l:187

v=0
o=user1 53655765 2353687637 IN IP4 128.3.4.5
s=Mbone Audio
i=Discussion of Mbone Engineering Issues
e=mbone@somewhere.com
c=IN IP4 224.2.0.1/127
t=0 0
m=audio 3456 RTP/AVP 0
```

Mixing short field names and long field names is allowed, but not recommended. Servers MUST accept both short and long field names for requests. Proxies MUST NOT translate a request between short and long forms if authentication fields are present.

# 11 SIP Transport

SIP is defined so it can use either UDP or TCP as a transport protocol.

## 11.1 Achieving Reliability For UDP Transport

### 11.1.1 General Operation

SIP assumes no additional reliability from IP. Requests or replies may be lost. A SIP client SHOULD simply retransmit a SIP request periodically with timer $T1$ (default value of $T1$: once a second) until it receives a response, or until it has reached a set limit on the number of retransmissions. The default limit is 20.

SIP requests and replies are matched up by the client using the Call-ID header field; thus, a server can only have one outstanding request per call at any given time.

If the reply is a provisional response, the initiating client SHOULD continue retransmitting the request, albeit less frequently, using timer $T2$. The default retransmission interval $T2$ is 5 seconds.

After the server sends a final response, it cannot be sure the client has received the response, and thus SHOULD cache the results for at least 30 seconds to avoid having to, for example, contact the user or user location server again upon receiving a retransmission.

### 11.1.2 INVITE

Special considerations apply for the INVITE method.

1. After receiving an invitation, considerable time may elapse before the server can determine the outcome. For example, the called party may be "rung" or extensive searches may be performed, so delays between the request and a definitive response can reach several tens of seconds. If either caller or callee are automated servers not directly controlled by a human being, a call attempt may be unbounded in time.

   It is undesirable to retransmit the INVITE request, as this introduces additional network traffic. The retransmission interval would have to be no more than about a second, since the callee would encounter a "dead" voice path if the "200 OK" response is lost.

2. It is possible that the invitation request reaches the callee and the callee is willing to take the call, but that the final response (200 OK, in this case) is lost on the way to the caller. If the session still exists but the initiator gives up on including the user, the contacted user has sufficient information to be able to join the session. However, if the session no longer exists because the invitation initiator "hung up" before the reply arrived and the session was to be two-way, the conferencing system should be prepared to deal with this circumstance.

3. If a telephony user interface is modeled or if we need to interface to the PSTN, the caller will provide "ringback", a signal that the callee is being alerted. Once the callee picks up, the caller needs to know so that it can enable the voice path and stop ringback. The callee's response to the invitation could get lost. Unless the response is transmitted reliably, the caller will continue to hear ringback while the callee assumes that the call exists.

4. The client has to be able to terminate an on-going request, e.g., because it is no longer willing to wait for the connection or search to succeed. One cannot rely on the absence of request retransmission, since the server would have to continue honoring the request for several request retransmission periods, that is, possible tens of seconds if only one or two packets can be lost.

The first problem is solved by indicating progress to the caller: the server returns a provisional response indicating it is searching or ringing the user.

The second and third problems are solved by having the server retransmit the final response at intervals of $T3$ (default value of $T3$ = 2 seconds) until it receives an ACK request for the same Call-ID and CSeq or until it has retransmitted the final response 10 times. The ACK request is acknowledged only once. If the request is syntactically valid and the Request-URI matches that in the INVITED request with the same Call-ID, the server answers with status code 200, otherwise with status code 400.

Fig. 8 and 9 show the client and server state diagram for invitations.

### 11.2 Connection Management for TCP

A single TCP connection can serve one or more SIP transactions. A transaction contains zero or more provisional responses followed by exactly one final response.

The client MAY close the connection at any time. Closing the connection before receiving a final response signals that the client wishes to abort the request.

The server SHOULD NOT close the TCP connection until it has sent its final response, at which point it MAY close the TCP connection if it wishes to. However, normally it is the client's responsibility to close the connection.

If the server leaves the connection open, and if the client so desires it may re-use the connection for further SIP requests or for requests from the same family of protocols (such as HTTP or stream control commands).

## 12 Behavior of SIP Servers

This section describes behavior of a SIP server in detail. Servers can operate in proxy or redirect mode. Proxy servers can "fork" connections, i.e., a single incoming request spawns several outgoing (client) requests.

A proxy server always inserts a Via header field containing its own address into those requests that are caused by an incoming request. To prevent loops, a server MUST check if its own address is already contained in the Via header of the incoming request.

We define an "$A$–$B$ proxy" as a proxy that receives SIP requests over transport protocol $A$ and issues requests, acting as a SIP client, using transport protocol $B$. If not stated explicitly, rules apply to any

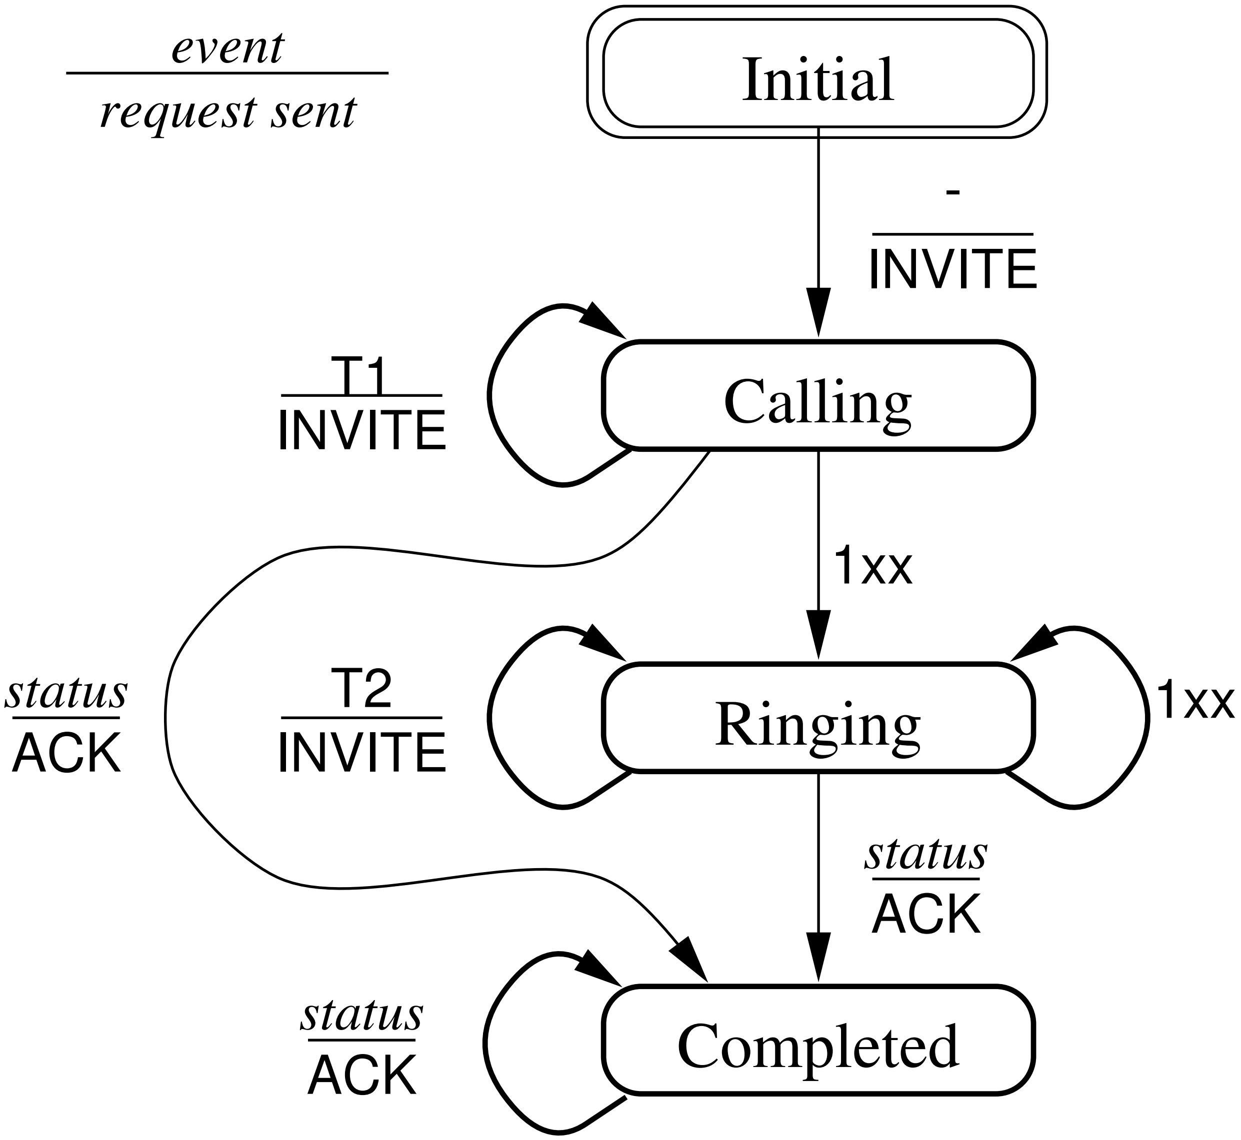


Figure 8: State transition diagram of client for INVITE method

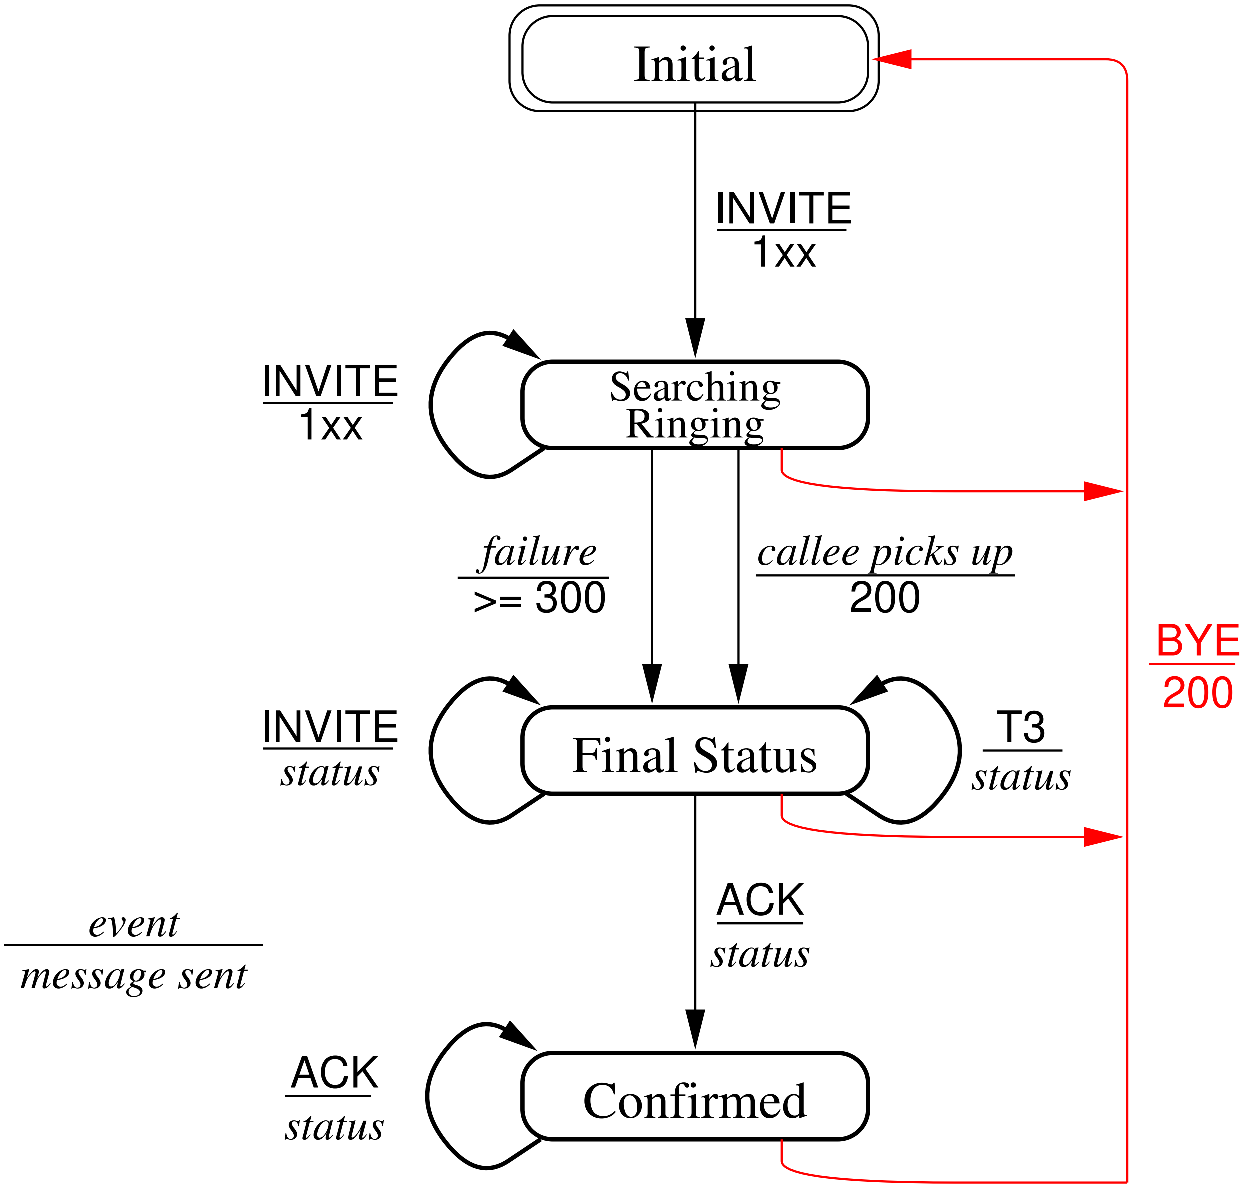


Figure 9: State transition diagram of server for INVITE method

combination of transport protocols. For conciseness, we only describe behavior with UDP and TCP, but the same rules apply for any unreliable datagram or reliable protocol, respectively.

The detailed connection behavior for UDP and TCP is described in Section 11.

### 12.1 Redirect Server

A redirect server does not issue any SIP requests of its own. It can return a response that refuses or redirects the request. After receiving an INVITE request, a redirect server proceeds through the following steps:

1. If the INVITE request cannot be answered immediately (e.g., because a location server needs to be contacted), it returns one or more provisional responses.
2. Once the server has gathered the list of alternative locations or has decided to refuse the call, it returns the final response. This ends the SIP transaction.

The redirect server maintains transaction state for the whole SIP transaction.

### 12.2 User Agent Server

Servers in user agents behave similarly to redirect servers, except that they may also accept a call.

### 12.3 Proxies Issuing Single Unicast Requests

Proxies in this category issue at most a single unicast request for each incoming SIP request, that is, they do not "fork" requests. Servers may choose to always operate in the mode described in Section 12.4.

The server can forward the request and any responses. It does not have to maintain any state for the SIP transaction. Reliability is assured by the next redirect server in the server chain.

A proxy server SHOULD cache the result of any address translations and the response to speed forwarding. After the cache entry has been expired, the server cannot tell whether an incoming request is actually a retransmission of an older request, where the TCP side has terminated. It will treat it as a new request.

### 12.4 Proxy Server Issuing Several Requests

All requests carry the same Call-ID. For unicast, each of the requests has a different (host-dependent) Request-URI. For multicast, a single request is issued, likely with a host-independent Request-URI. A client receiving a multicast query does not have to check whether the *host* part of the Request-URI matches its own host or domain name. To avoid response implosion, servers SHOULD NOT answer multicast requests with a 404 (Not Found) status code. Servers MAY decide not to answer multicast requests if their response would be 5xx.

The server MAY respond to the request immediately with a "100 Trying" or "180 Ringing" response; otherwise it MAY wait until either the first response to its requests or the UDP retransmission interval.

The following pseudo-code describes the behavior of a proxy server issuing several requests in response to an incoming request. The function `request(`$r$`, `$a$`)` sends a SIP request $r$ to address $a$. `await_response()` waits until a response is received and returns the response. `close(`$a$`)` closes the TCP connection to client with address $a$. `response(`$s$`, `$l$`, `$L$`)` sends a response to the client with status $s$ and list of locations $L$, with $l$ entries. `ismulticast()` returns 1 if the location is a multicast address and zero otherwise. The variable `timeleft` indicates the amount of time left until the maximum response time has expired. The variable `recurse` indicates whether the server will recursively try addresses returned through a 3xx response. A server MAY decide to recursively try only certain addresses, e.g., those which are within the same domain as the proxy server. Thus, an initial multicast request may trigger additional unicast requests.

```
enum {INVITE,          /* request type */
  ACK, OPTIONS, BYE, REGISTER, UNREGISTER} R;
int N = 0;             /* number of connection attempts */
address_t address[];   /* list of addresses */
int done[];            /* address has responded */
location[];            /* list of locations */
int heard = 0;         /* number of sites heard from */
int class;             /* class of status code */
int best = 1000;       /* best response so far */
int timeleft = 120;    /* sample timeout value */
int loc = 0;           /* number of locations */
struct {               /* response */
  int status;          /* response status */
  char *location;      /* redirect locations */
  address_t a;         /* address of respondent */
} r;
int i;

if (multicast) {
  request(R, address[0]);
} else {
  N = /* number of addresses to try */
  for (i = 0; i < N; i++) {
    request(R, address[i]);
    done[i] = 0;
  }
}
```

```
while (timeleft > 0 && (heard < N || multicast)) {
  r = await_response();
  class = r.status / 100;

  if (class >= 2) {
    heard++;
    for (i = 0; i < N; i++) {
      if (address[i] == r.a) {
        done[i] = 1;
        break;
      }
    }
  }

  if (class == 2) {
    best = r.status;
    break;
  }
  else if (class == 3) {
/* A server may optionally recurse.  The server MUST check whether
 * it has tried this location before and whether the location is
 * part of the Via path of the incoming request.  This check is
 * omitted here for brevity. Multicast locations MUST NOT be
     * returned to the client if the server is not recursing.
     */
    if (recurse) {
      multicast = 0;
      N++;
      request(R, r.location);
    } else if (!ismulticast(r.location)) {
      locations[loc++] = r.location;
      best = r.status;
    }
  }
  else if (class == 4) {
    if (best >= 400) best = r.status;
  }
  else if (class == 5) {
    if (best >= 500) best = r.status;
  }
  else if (class == 6) {
    best = r.status;
    break;
  }
}
```

```
/* We haven't heard anything useful from anybody. */
if (best == 1000) {
  best = 404;
}
if (best/100 != 3) loc = 0;
response(best, loc, locations);

/*
 * Close the other pending transactions by sending BYE.
 */
for (i = 0; i < N; i++) {
  if (!done[i]) {
    request(BYE, address[i]);
    if (tcp) close(a);
  }
}
```

After receiving a 2xx or 6xx response, the server SHOULD terminate all other pending requests by sending a BYE request and closing the TCP connection, if applicable. (Terminating pending requests is advisable as searches consume resources. Also, INVITE requests may "ring" on a number of workstations if the callee is currently logged in more than once.)

[TBD: How do we cancel multicast requests? Force receivers to listen for a 200/6xx response and hope that they don't miss one?]

When operating in this mode, a proxy server MUST ignore any responses received for Call-IDs that it does not have a pending transaction for. (If server were to forward responses not belonging to a current transaction using the Via field, the requesting client would get confused if it has just issued another request using the same Call-ID.)

# 13 Third-Party Call Initiation

In some circumstances, third-party call control is required, where the calling party suggests to the called party to invite a (small) number of other parties. Third-party call control can be used to implement the following features:

**Multipoint-control unit (MCU):** Some conferences use a multipoint control unit to mix, convert and replicate media streams. While this solution has scaling problems, it is widely deployed in traditional telephony and ISDN conferencing settings, as so-called conference bridges. In a MCU-based conference, the conference initiator or any authorized member invites a new participant and indicate the address of the MCU in the Also header. The invitee then contacts the MCU using the same session description and requests to be added to the call, just like a normal two-party call.

**Telephony call initiation ("click-to-call"):** A SIP INVITE request containing two addresses in the Also header is sent to a PSTN service node that connects these two addresses by a telephone call.

**Fully-meshed small conference:** For small conferences, such as adding a third party to a two-party call, multicast may not always be appropriate or available. Instead, when inviting a new participant, the caller asks the new member to call the remaining members. TBD: Should the call-id be the same or different? Need to distinguish between new INVITE for same call and adding a party to a call. Include conference identifier?

TBD: How about just transferring an SDP description with multiple addresses?

The Also: header (Section 6.9) is used to indicate a list of parties that the callee should invite.

# 14 ISDN and Intelligent Network Services

SIP may be used to support a number of ISDN [27] and Intelligent Network [28] telephony services, described below. Due to the fundamental differences between Internet-based telephony and conferencing as compared to public switched telephone network (PSTN)-based services, service definitions cannot be precisely the same. Where large differences beyond addressing and location of implementation exist, this is indicated below. The term address implies any SIP address. (Section 1.4.1).

**Call transfer (TRA)** enables a user to transfer an established (i.e., active) call to a third party. SIP signals this via the Location header in the BYE (Section 4.2.4) method.

**Call forwarding (CF)** permits the called user to forward particular pre-selected calls to another address. Unlike telephony, the choice of calls to be forwarded depends on program logic contained in any of the SIP servers and can thus be made dependent on time-of-day, subject of call, media types, urgency or caller identity, rather than being restricted to matching list entries. This forwarding service encompasses:

**Call forwarding busy/don't answer (CFB/CFNR, SCF-BY/DA)** allows the called user to forward particular pre-selected calls if the called user is busy or does not answer within a set time.

**Selective call forwarding (SCF)** permits the user to have her incoming calls addressed to another network destination, no matter what the called party status is, if the calling address is included in, or excluded from, a screening list. The user's originating service is unaffected.

**Completion of calls to busy subscriber (CCBS)** allows a calling user encountering a busy destination to be informed when the busy destination becomes free, without having to make a new call attempt. SIP supports services close to CCBS by allowing a callee to indicate a more opportune time to call back (Section 6.25). Also, calling and called user agents can easily record the URL of outcoming and incoming calls, so that a user can re-try or return calls with a single mouse click.

**Conferencing (CON)** allows the user to communicate simultaneously with multiple parties, which may also communicate among themselves. SIP can initiate IP multicast conferences with any number of participants, conferences where media are mixed by a conference bridge (multipoint control unit or MCU) and, for exceptional applications with a small number of participants, fully-meshed conferences, where each participant sends and receives data to all other participants.

**Conference calling add-on** allows a user to add and drop participants once the conference is active.

**Conference calling meet-me (MMC)** allows the user to set up a conference or multi-party call, indicating the date, time, conference duration, conference media and other parameters. The conference session description included in the SIP invitation may indicate a time in the future. For multicast conferences, participants do not have to connect using SIP at the actual time of the conference; instead, they simply subscribe to the multicast addresses listed in the announcement. For MCU-based conferences, the session description may contain the address of the MCU to be called at the time of the conference.

**Destination call routing (DCR)** allows customers to specify the routing of their incoming calls to destinations according to

- time of day, day of week, etc.;
- area of call origination;
- network address of caller;
- service attributes;
- priority (e.g., from input of a PIN or password);
- charge rates applicable for the destination;
- proportional routing of traffic.

In SIP, destination call routing is implemented by proxy and redirect servers that implement custom call handling logic, with parameters including, but not limited to the set listed above.

**Follow-me diversion (FMD)** allows the service subscriber to remotely control the redirection (diversion) of calls from his primary network address to other locations.

In SIP, finding the current network-reachable location of a callee is left to the location service and is outside the scope of this specification. However, users may use the REGISTER method (Section 4.2.5) to appraise their "home" SIP server of their new location.

**Originating call screening (OCS)** controlls the ability of a node to originate calls. In a fashion similar to closed user groups, a firewall would have to be used to restrict the ability to initiate SIP invitations outside a designated part of the network. In many cases, gateways to the PSTN will require appropriate authentication.

**Premium rate (PRM)** allows to pay back part of the call cost to the called party, considered as an added value provider. See discussion on billing services below.

**Split charging (SPL)** allows the calling and called party being each charged for one part of the call. See discussion on billing services below.

**Universal access number (UAN)** allows a subscriber with several network addresses to be reached with a single, unique address. The subscriber may specify which incoming calls are to be routed to which address. SIP offers this functionality through proxies and redirection.

**Universal personal telecommunications (UPT)** is a mobility service which enables subscribers to be reached with a unique personal telecommunication number (PTN) across multiple networks at any network access. The PTN will be translated to an appropriate destination address for routing based on the capabilities subscribed to by each service subscriber. A person may have multiple PTNs, e.g., a business and private PTN. In SIP, the host-independent address (Section 1.4.1) of the form *user@host* serves as the PTN, which is translated into one or more host-dependent addresses.

**User-defined routing (UDR)** allows a subscriber to specify how outgoing calls, from the subscriber's location, shall be routed. SIP cannot specify routing preferences; this is presumed to be handled by a policy-based routing protocol, source routing or similar mechanisms.

Some telephony services can be provided by the end system, without involvement by SIP:

**Abbreviated dialing** allows users to reach local subscribers without specifying the full address (domain or host name). For SIP, the user application completes the address to be a fully qualified domain name.

**Call waiting (CW)** allows the called party to receive a notification that another party is trying to reach her while she is busy talking to another calling party.

For SIP-based telephony, the called party can maintain several call presences at the same time, limited by local resources. Thus, it is up to the called party to decide whether to accept another call. The separation of resource reservation and call control may lead to the situation that the called party accepts the incoming call, but that the network or system resource allocation fails. This cannot be completely prevented, but if the likely resource bottleneck is at the local system, the user agent may be able to determine whether there are sufficient resources available or roughly track its own resource consumption.

**Consultation calling (COC)** allows a subscriber to place a call on hold, in order to initiate a new call for consultation. In systems using SIP, consultation calling can be implemented as two separate SIP calls, possibly with the temporary release of reserved resources for the call being put on hold.

**Customized ringing (CRG)** allows the subscriber to allocate a distinctive ringing to a list of calling parties. In a SIP-based system, this feature is offered by the user application, based on caller identification (From, Section 6.17) provided by the SIP INVITE request (Section 4.2.1).

**Malicious call identification (MCI)** allows the service subscriber to control the logging (making a record) of calls that received that are of a malicious nature. In SIP, by default, all calls identify the calling party and the SIP servers that have forwarded the call. In addition, calls may be authenticated using standard HTTP methods or transport-layer security. A callee may decide only to accept calls that are authenticated.

**Multiway calling (MWC)** allows the user to establish multiple, simultaneous calls with other parties. For a SIP-based end system, the considerations for consultation calling apply.

**Terminating call screening (TCS)** allows the subscriber to specify that incoming calls either be restricted or allowed, according to a screening list and/or by time of day or other parameters.

Billing features such as *account card dialing*, *automatic alternative billing*, *credit card calling (CCC)*, *reverse charging*, *freephone (FPH)*, *premium rate (PRM)* and *split charging* are supported through authentication. However, mechanisms for indicating billing preferences and capabilities have not yet been specified for SIP.

*Advice of charge* allows the user paying for a call to be informed of usage-based charging information. Charges incurred by reserving resources in the network are probably best indicated by a protocol closely affiliated with the reservation protocol. Advice of charge when using Internet-to-PSTN gateways through SIP appears feasible, but is for further study. Desirable facilities include indication of charges at call setup time, during the call and at the end of the call

*Closed user groups* (CUGs) that restrict members to communicate only within the group can be implemented using firewalls and SIP proxies.

*User-to-user signaling* is supported within SIP through the addition of headers, with predefined header fields such as Subject or Organization.

*Third-party signaling* is optionally supported within SIP (Section 6.9). Third-party signaling can be used to indicate to callees who else to invite to a call for MCU and fully-meshed conferences. Third-party signaling, combined with appropriate URLs, may be used to initiate PSTN phone calls from an Internet host.

# 15 Security Considerations

## 15.1 Confidentiality

Unless SIP transactions are protected by lower-layer security mechanisms such as SSL [**?**], an attacker may be able to eavesdrop on call establishment and invitations and, through the Subject header field or the session description, gain insights into the topic of conversation.

## 15.2 Integrity

Unless SIP transactions are protected by lower-layer security mechanisms such as SSL [**?**], an active attacker may be able to modify SIP requests.

## 15.3 Access Control

SIP requests are not authenticated unless the SIP Authorization and WWW-Authenticate headers are being used. The strengths and weaknesses of these authentication mechanisms are the same as for HTTP.

### 15.4 Privacy

User location and SIP-initiated calls may violate a callee's privacy. An implementation SHOULD be able to restrict, on a per-user basis, what kind of location and availability information is given out to certain classes of callers.

## A Minimal Implementation

### A.1 Client

All clients MUST be able to generate the INVITE and ACK requests and MUST be able to include the Call-ID, Content-Length, Content-Type, From and To headers. A minimal implementation MUST understand SDP [9]. In responses, it must be able to parse the Call-ID, Content-Length, Content-Type, Require headers. It must be able to recognize the status code classes 1 through 6 and act accordingly.

The following capability sets build on top of a minimal implementation:

**Basic:** A basic implementation SHOULD add support for the BYE method to allow the interruption of a pending call attempt. It SHOULD include a User-Agent header in its requests and indicate its preferred language in the Accept-Language header.

**Redirection:** To support call forwarding, a client needs to be able to understand the Location header, but only the SIP-URL part, not the parameters.

**Negotiation:** A client MUST be able to request the OPTIONS method and understand the 380 "Alternative Service" status and the Location parameters to participate in terminal and media negotiation. It SHOULD be able to parse the Warning response header to provide useful feedback to the caller.

**Authentication:** If a client wishes to invite callees that require caller authentication, it must be able to recognize the 401 "Unauthorized" status code, must be able to generate the Authorization request header and understand the WWW-Authenticate response header.

If a client wishes to use proxies that require caller authentication, it must be able to recognize the 407 "Proxy Authentication Required" status code, must be able to generate the Proxy-Authorization request header and understand the Proxy-Authenticate response header.

### A.2 Server

A minimally compliant server implementation MUST understand the INVITE, ACK and BYE requests. It MUST parse the generate, as appropriate, the Call-ID, Content-Length, Content-Type, From, PEP, To and Via headers. It must echo the Sequence header in the response. It SHOULD include the Server header in its responses.

# B Summary of Augmented BNF

In this specification we use the Augmented Backus-Naur Form notation described in [21]. For quick reference, the following is a brief summary of the main features of this ABNF.

"abc"
: The case-insensitive string of characters "abc" (or "Abc", "aBC", etc.);

%d32
: The character with ASCII code decimal 32 (space);

*term
: zero of more instances of term;

3*term
: three or more instances of term;

2*4term
: two, three or four instances of term;

[ term ]
: term is optional;

{ term1 term2 term3 }
: set notation: term1, term2 and term3 must all appear but their order is unimportant;

term1 | term2
: either term1 or term2 may appear but not both;

#term
: a comma separated list of term;

2#term
: a comma separated list of term containing at least 2 items;

2#4term
: a comma separated list of term containing 2 to 4 items.

### Common Tokens

Certain tokens are used frequently in the BNF this document, and not defined elsewhere. Their meaning is well understood but we include it here for completeness.

```
CR     =  %d13           ; carriage return character
LF     =  %d10           ; line feed character
CRLF   =  CR LF          ; typically the end of a line
SP     =  %d32           ; space character
TAB    =  %d09           ; tab character
LWS    =  *( SP | TAB)   ; linear whitespace
DIGIT  =  "0" .. "9"     ; a single decimal digit
```

## Changes in Version -04

Since version -03, the following changes have been made.

- The introduction has been reorganized and large parts rewritten.
- CONNECTED changed to ACK, as it applies to all responses, not just 200.
- Status code 181 (Queued) and Call-Disposition: Queue added.
- Status code 481 (Invalid Call-ID) added.
- Status code 482 (Loop Detected) added. Via description contains motivation.
- Allow phone numbers in SIP URL for easy connection to Internet telephony gateways.
- Added Also header for third-party connectivity.
- When doing parallel searches, pending searches should be aborted when one address was successful. The phone call may be ringing on a number of workstations where the user is logged in and would keep ringing.
- Added duration parameter to Retry-After to indicate how long the callee is likely to be reachable at the address given.
- Changed Sequence to CSeq for consistency with RTSP.

## C Open Issues

**Full meshes:** How about just transferring an SDP description with multiple addresses?

**H.323:** Detailed interaction with H.323 and H.245.

**TRANSACTION:** Should we have a transaction id in addition to a call ID? Call-IDs are for the end system, but a transaction ID is for a single SIP exchange. This is useful for Internet telephony, where a single call may trigger several transactions. Also, avoids BYE race condition: Proxy doing parallel search cancels pending search with BYE after one of the addresses responds with 200. Through another proxy, this BYE reaches the same end system and cancels the successful call.

# D Acknowledgments

We wish to thank the members of the IETF MMUSIC WG for their comments and suggestions. Detailed comments were provided by Jonathan Rosenberg. This work is based, inter alia, on [29, 30]. Parameters of the terminal negotiation mechanism were influenced by Scott Petrack's CMA design.

# E Authors' Addresses

Mark Handley
USC Information Sciences Institute
c/o MIT Laboratory for Computer Science
545 Technology Square
Cambridge, MA 02139
USA
electronic mail: mjh@isi.edu

Henning Schulzrinne
Dept. of Computer Science
Columbia University
1214 Amsterdam Avenue
New York, NY 10027
USA
electronic mail: schulzrinne@cs.columbia.edu

Eve Schooler
Computer Science Department 256-80
California Institute of Technology
Pasadena, CA 91125
USA
electronic mail: schooler@cs.caltech.edu

# References

[1] H. Schulzrinne, A. Rao, and R. Lanphier, "Real time streaming protocol (RTSP)," Internet Draft, Internet Engineering Task Force, Mar. 1997. Work in progress.

[2] M. Handley, "SAP: Session announcement protocol," Internet Draft, Internet Engineering Task Force, Nov. 1996. Work in progress.

[3] R. Pandya, "Emerging mobile and personal communication systems," *IEEE Communications Magazine*, vol. 33, pp. 44–52, June 1995.

[4] P. Lantz, “Usage of H.323 on the Internet,” Internet Draft, Internet Engineering Task Force, Feb. 1997. Work in progress.

[5] M. Handley, J. Crowcroft, C. Bormann, and J. Ott, “The internet multimedia conferencing architecture,” Internet Draft, Internet Engineering Task Force, July 1997. Work in progress.

[6] R. Braden, L. Zhang, S. Berson, S. Herzog, and S. Jamin, “Resource reservation protocol (RSVP) – version 1 functional specification,” Internet Draft, Internet Engineering Task Force, June 1997. Work in progress.

[7] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, “RTP: a transport protocol for real-time applications,” Tech. Rep. RFC 1889, Internet Engineering Task Force, Jan. 1996.

[8] H. Schulzrinne, A. Rao, and R. Lanphier, “Real time streaming protocol (RTSP),” Internet Draft, Internet Engineering Task Force, July 1997. Work in progress.

[9] M. Handley and V. Jacobson, “SDP: Session description protocol,” Internet Draft, Internet Engineering Task Force, Mar. 1997. Work in progress.

[10] S. Bradner, “Key words for use in RFCs to indicate requirement level,” Tech. Rep. RFC 2119, Internet Engineering Task Force, Mar. 1997.

[11] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, and T. Berners-Lee, “Hypertext transfer protocol – HTTP/1.1,” Tech. Rep. RFC 2068, Internet Engineering Task Force, Jan. 1997.

[12] C. Partridge, “Mail routing and the domain system,” Tech. Rep. RFC 974, Internet Engineering Task Force, Jan. 1986.

[13] A. Gulbrandsen and P. Vixie, “A DNS RR for specifying the location of services (DNS SRV),” Tech. Rep. RFC 2052, Internet Engineering Task Force, Oct. 1996.

[14] P. V. Mockapetris, “Domain names - implementation and specification,” Tech. Rep. RFC 1035, Internet Engineering Task Force, Nov. 1987.

[15] R. T. Braden, “Requirements for internet hosts - application and support,” Tech. Rep. RFC 1123, Internet Engineering Task Force, Oct. 1989.

[16] D. Zimmerman, “The finger user information protocol,” Tech. Rep. RFC 1288, Internet Engineering Task Force, Dec. 1991.

[17] W. Yeong, T. Howes, and S. Kille, “Lightweight directory access protocol,” Tech. Rep. RFC 1777, Internet Engineering Task Force, Mar. 1995.

[18] T. Berners-Lee, “Universal resource identifiers in WWW: a unifying syntax for the expression of names and addresses of objects on the network as used in the world-wide web,” Tech. Rep. RFC 1630, Internet Engineering Task Force, June 1994.

[19] T. Berners-Lee, R. Fielding, and L. Masinter, “Uniform resource locators (URL): Generic syntax and semantics,” Internet Draft, Internet Engineering Task Force, May 1997. Work in progress.

[20] T. Berners-Lee, L. Masinter, and M. McCahill, “Uniform resource locators (URL),” Tech. Rep. RFC 1738, Internet Engineering Task Force, Dec. 1994.

[21] D. Crocker, “Augmented BNF for syntax specifications: ABNF,” Internet Draft, Internet Engineering Task Force, Oct. 1996. Work in progress.

[22] J. C. Mogul and S. E. Deering, “Path MTU discovery,” Tech. Rep. RFC 1191, Internet Engineering Task Force, Nov. 1990.

[23] W. R. Stevens, *TCP/IP illustrated: the protocols*, vol. 1. Reading, Massachusetts: Addison-Wesley, 1994.

[24] D. Crocker, “Standard for the format of ARPA internet text messages,” Tech. Rep. Also STD0011, RFC 822, Internet Engineering Task Force, Aug. 1982.

[25] A. Vaha-Sipila, “URLs for telephony,” Internet Draft, Internet Engineering Task Force, Aug. 1997. Work in progress.

[26] L. Masinter, P. Hoffman, and J. Zawinski, “The mailto URL scheme,” Internet Draft, Internet Engineering Task Force, Oct. 1997. Work in progress.

[27] International Telecommunication Union, “Integrated services digital network (ISDN) service capabilities – definition of supplementary services,” Recommendation I.250, Telecommunication Standardization Sector of ITU, Geneva, Switzerland, 1993.

[28] International Telecommunication Union, “General recommendations on telephone switching and signaling – intelligent network: Introduction to intelligent network capability set 1,” Recommendation Q.1211, Telecommunication Standardization Sector of ITU, Geneva, Switzerland, Mar. 1993.

[29] E. M. Schooler, “Case study: multimedia conference control in a packet-switched teleconferencing system,” *Journal of Internetworking: Research and Experience*, vol. 4, pp. 99–120, June 1993. ISI reprint series ISI/RS-93-359.

[30] H. Schulzrinne, “Personal mobility for multimedia services in the Internet,” in *European Workshop on Interactive Distributed Multimedia Systems and Services*, (Berlin, Germany), Mar. 1996.

## Full Copyright Statement