

SIP Call Control Services

Status of this Memo

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

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

To learn the current status of any Internet-Draft, please check the “lid-abstracts.txt” listing contained in the Internet-Drafts Shadow Directories on ftp.is.co.za (Africa), nic.nordu.net (Europe), munnari.oz.au (Pacific Rim), ds.internic.net (US East Coast), or ftp.isi.edu (US West Coast).

Distribution of this document is unlimited.

Copyright Notice

Copyright (c) The Internet Society (1998). All Rights Reserved.

Abstract

This document describes the org.ietf.sip.call extensions to the Session Initiation Protocol (SIP). The document also describes how standard telephony services can be implemented in SIP.

Contents

1 Terminology	2
2 Introduction	2
3 Headers	2
3.1 Accept-Location	3
3.2 Also	3
3.3 Call-Disposition	4
3.4 Location	5
3.5 Replaces	6
3.6 Requested-By	7
4 Status Code Definitions	7
4.1 181 Queued	7
5 ISDN and Intelligent Network Services	8
5.1 Call Redirection and “Number”-Translation Services	8
5.2 Camp-on	9
5.3 Call Screening	9

5.4	Directed Call Pickup	9
5.5	Directed Call Pickup with Barge-In	10
5.6	Outgoing Call Routing	10
5.7	End-System Services	10
5.8	Billing Features	11
5.9	User-to-User Signaling	11
5.10	Operator Services	11
5.11	Multipoint Control Unit (MCU) Services	11
5.12	Fully-Meshed Conferences	13
6	Acknowledgements	13

1 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 [1] and indicate requirement levels for compliant SIP implementations.

2 Introduction

This document describes the `org.ietf.sip.call` extensions to the Session Initiation Protocol (SIP) [2]. When using the extensions described here, the client **MUST** include the extension name in a **Require** header.

3 Headers

	type	ACK	BYE	INV	OPT	REG
Accept-Location	R	-	-	o	o	-
Also	R	-	-	o	o	-
Also	r	-	-	o	o	-
Call-Disposition	R	-	o	o	-	-
Replaces	R	-	-	o	o	-
Replaces	r	-	-	o	o	-
Requested-By	R	-	-	o	o	-
Requested-By	r	-	-	o	o	-

Table 1: 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.

3.1 Accept-Location

The **Accept-Location** request header allows the caller to provide hints to proxy and redirect servers. It uses the same parameters as the **Location** header (Section 3.4).

3.2 Also

The **Also** request and response header advises the recipient to issue **INVITE** requests to the addresses listed. Each of these invitations **SHOULD** contain a **Requested-By** header that contains the **From** field of the message containing the **Also** field. The **Also** header **MUST** only be processed by the calling or called user agent, not by any intermediate proxy or redirect servers.

If a message contains both **Also** and **Replaces**, the invitations requested in the **Also** header **MUST** be completed, successfully or not, before the terminations requested in the **Replaces** header.

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

Example header:

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

If *A*, *B* and *C* are end points, the following is a typical scenario:

```
A→B:  INVITE B SIP/2.0
      Call-ID: 19971214T123503.33@A
      Also: C
      From: A
B→A:  SIP/2.0 200 OK
      Call-ID: 19971214T123503.33@A
      From: A
B→C:  INVITE C SIP/2.0
      Call-ID: 19971214T123503.33@A
      From: B
      Requested-By: A
C→B:  SIP/2.0 200 OK
      Call-ID: 19971214T123503.33@A
      From: B
```

If *G* is a group address with members *X* through *Z*, a group invitation may proceed as follows:

```
A→G:  INVITE G SIP/2.0
      From:  A
      Call-ID: 19971214T124523.00@A

G→A:  SIP/2.0 200 OK
      From:  A
      Call-ID: 19971214T124523.00@A
      Also:  X, Y, Z

A→X:  INVITE X SIP/2.0
      From:  A
      Call-ID: 19971214T124523.00@A
      Requested-By:  G

A→Y:  INVITE Y SIP/2.0
      From:  A
      Call-ID: 19971214T124523.00@A
      Requested-By:  G

A→Z:  INVITE Z SIP/2.0
      From:  A
      Call-ID: 19971214T124523.00@A
      Requested-By:  G
```

The Also header makes it possible to create full meshes (generalized “three-way” calling) and supports the resolution of group addresses. Unlike the Location header, which enumerates alternatives to be tried, the Also header lists addresses to be all invited.

3.3 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 to a proxy or redirect server that it should resolve the address to a list of group members and return a 300 (Multiple Choices) response. The list of group members is contained in an Also header.

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 4.1). A pending call be terminated by a SIP BYE request.

do-not-respond: The do-not-respond directive indicates to the callee that it should not issue a response, informational or final. This may be used to send invitations to multicast groups.

Maybe the combination of do-not-respond and all could be used for group invitations to larger lists?

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

Example:

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

3.4 Location

This document adds extension parameters to the Location header.

```
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        = Internet media type [ "/" subtype ]
feature-list     = "voice-mail" | "attendant" | "permanent"
```

```
extension-attribute | "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
```

class: The class parameter indicates 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 address. The "permanent" flag indicates that this address is a new, permanent address. When used for registration, the server SHOULD return a 301 status instead of 302.

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. Media types can be the standard Internet media types ("audio", "video", "text", "application"), optionally followed by a subtype (e.g., "text/html"). In addition, the type "application/email" is defined.

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.

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. This is meant for indicating general communication capability, sufficient for the caller to choose an appropriate address.

```
Location: sip://watson@worchester.bell-telephone.com ;q=0.7
        ;service=IP,voice-mail
        ;media=audio+video+application/x-wb ;duplex=full
Location: rtsp://tape.bell-telephone.com?watson482178 ;q=0.6
        ;service=IP,voice-mail
        ;media=audio+video ;duplex=full
Location: phone://1-415-555-1212 ;q=0.5
        ;service=ISDN;mobility=fixed;
        language=en,es,iw
Location: phone://1-800-555-1212 ;q=0.2
        ;service=pager;mobility=mobile;
        duplex=send-only;media=text; priority=urgent;
        description="For emergencies only"
Location: mailto:watson@bell-telephone.com ;q=0.1
        ;media=application/email
Location: http://www.bell-telephone.com/~watson ;q=0.1
        ;service=text/html
```

A 301 or 302 response MAY contain additional information in human-readable form, e.g., as **Content-Type: text/html**. It is up to the server issuing the **Location** header to ensure consistency between the content of the **Location** header and the response entity.

3.5 Replaces

The **Replaces** request and response header is analogous to the **Also** header (Section 3.2, except that it asks the recipient to issue a **BYE** to the addresses listed, with the same **Call-ID** as the request containing the

Replaces header. The special address “*” indicates that the recipient should replace all existing legs of the call within the **Call-ID**. If a message contains both **Also** and **Replaces**, the invitations requested in the **Also** header **MUST** be completed, successfully or not, before the terminations requested in the **Replaces** header.

For example, with entities *A*, *B* and *M* (an MCU):

```
A→B:  INVITE B SIP/2.0
      Call-ID: 19971214T123503.33@A
      From:   A
```

```
B→A:  SIP/2.0 200 OK
      Call-ID: 19971214T123503.33@A
      From:   A
```

```
M→B:  INVITE B SIP/2.0
      Call-ID: 19971214T123503.33@A
      From:   M
      Replaces: *
```

```
B→M:  SIP/2.0 200 OK
      Call-ID: 19971214T123503.33@A
      From:   M
```

```
B→A:  BYE A SIP/2.0
      Call-ID: 19971214T123503.33@A
      From:   B
      Requested-By: M
```

```
A→B:  SIP/2.0 200 OK
      Call-ID: 19971214T123503.33@A
      From:   A
```

3.6 Requested-By

The **Requested-By** request header is only used in requests triggered by **Also** or **Replaces**. It contains the URI of the entity that issued the request containing the **Also** header. The URI is taken from the **From** header of the request. For example, if *A* sends an invitation to *B* containing **Also: C**, *B* issues an invitation to *C* with **Requested-By: A**.

4 Status Code Definitions

This feature set defines one additional status code.

4.1 181 Queued

The called party was temporarily unavailable, but the caller indicated via a “**Call-Disposition: Queue**” directive (Section 3.3) 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.

5 ISDN and Intelligent Network Services

SIP may be used to support a number of ISDN [4] and Intelligent Network [5] 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.

This section is for information and illustration only. There are many different ways of implementing services in SIP. Since SIP only describes the behavior induced by messages, different means of implementing services will interoperate.

5.1 Call Redirection and "Number"-Translation Services

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 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.

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 user agents, proxy and redirect servers that implement custom call handling logic, with parameters including, but not limited to the set listed above. Caller preferences are expressed in the **Accept-Location** header, callee preferences in the **Location** header.

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 to appraise their “home” SIP server of their new location.

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, a host-independent address of the form *user@domain* serves as the PTN, which is translated into one or more host-dependent addresses.

5.2 Camp-on

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 with the **Retry-After** header. Also, calling and called user agents can easily record the URI of outgoing and incoming calls, so that a user can re-try or return calls with a single mouse click or automatically. **Call-Disposition: queue** allows a caller to wait until the line becomes available. This service is also known as a “camp-on” service.

5.3 Call Screening

Originating call screening (OCS) controls 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.

5.4 Directed Call Pickup

Directed call pickup allows a station user to answer calls directed to a specific address from any other address by providing the address of the line to be answered. The rung station must permit pickup. If the call has not been answered at the ringing station, regular call pickup occurs. If the call has been answered already, an error is generated.

5.5 Directed Call Pickup with Barge-In

Directed call pickup with barge-in establishes a 3-way call if the call has been answered at the original destination.

5.6 Outgoing Call Routing

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. However, the SIP `Accept-Location` header may be used by proxies and redirect servers to route calls according to caller preferences.

5.7 End-System Services

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 header) provided by the SIP INVITE request.

Malicious call identification (MCI) allows the service subscriber to control the logging (making a record) of calls 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.

5.8 Billing Features

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.

5.9 User-to-User Signaling

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

5.10 Operator Services

There are a number of services which involve three parties, for example, a secretary dialing for boss, an auto-dialer handing a call to a telemarketer, attended call transfer, operator services such as person-to-person calls and full-mesh "multicast".

Operator services can be implemented in a number of ways, combining Also, Replaces with either INVITE or BYE. The callee's end system does not have to be cognizant of the fact that this is an operator service. The services make use of the Call-ID rules that stipulate that a new INVITE for an existing Call-ID does not alert the user, but is added silently.

Figure 1 shows the example of an autodialer connecting to a prospective customer and, once the callee picks up, transferring the call to the telemarketer. (This goes to show that SIP is morally neutral.)

5.11 Multipoint Control Unit (MCU) Services

In the language of IN services, SIP supports:

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. This is described in more detail in Sections 5.11 and 5.12.

Conference calling add-on allows a user to add and drop participants once the conference is active. Participants in the SIP session accomplish this by sending INVITE and BYE requests to the parties to be added and dropped. If *A* wants *B* to drop out of a conference, it sends a BYE request with "Replaces: *".

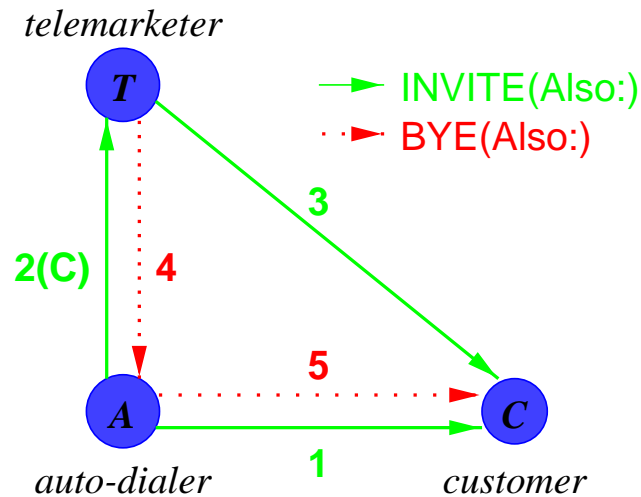


Figure 1: Telemarketer example

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.

Some conferences use a multipoint control unit (MCU) 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 indicates 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.

Parties inviting others to a conference do not have to know that the conference media is managed by an MCU. The inviting party *A* treats the MCU *M* like another participant and includes it in the **Also** list. The newly invited participant *B* invites the MCU, which in turn sends a **Replaces** header with all participants. (See example in Section 3.5). Figure 2 shows the transition from a fully-meshed conference (see below) to an MCU-based conference.

Operator-assisted dial-out: The operator calls each participant, and simply indicates the MCU in the **Also** list. The **Call-ID** and/or the address used by the operator serves as the identifier to the MCU. For example:

```
O→A:  INVITE A SIP/2.0
      From:  O
      Also:  conference176@M

A→M:  INVITE conference176@M
      Requested-By:  O
```

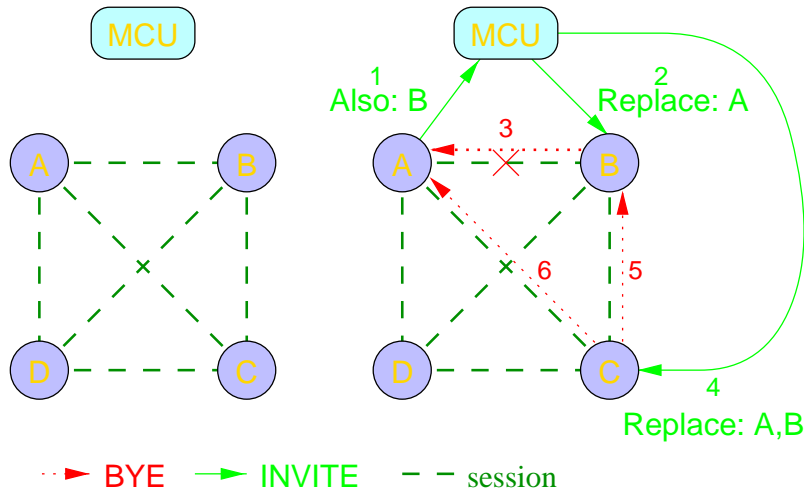


Figure 2: Transition from fully-meshed to MCU-based conference

Meet-me: The leader and participants dial into their conference at the scheduled time with an assigned conference identifier and security code.

```
A→M: INVITE conference189@M
      From: A
```

5.12 Fully-Meshed Conferences

For very 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. The Call-ID for all participants is the same.

6 Acknowledgements

Parameters of the terminal negotiation mechanism in the Location header were influenced by Scott Petrack's CMA design.

References

- [1] S. Bradner, "Key words for use in RFCs to indicate requirement levels," RFC 2119, Internet Engineering Task Force, Mar. 1997.
- [2] M. Handley, H. Schulzrinne, and E. Schooler, "SIP: Session initiation protocol," Internet Draft, Internet Engineering Task Force, Nov. 1997. Work in progress.
- [3] M. Handley and V. Jacobson, "SDP: Session description protocol," Internet Draft, Internet Engineering Task Force, Mar. 1997. Work in progress.

- [4] 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.
- [5] 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.

Full Copyright Statement

Copyright (c) The Internet Society (1998). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.