

Alan R. Johnston

SIP



understanding the
**Session Initiation
Protocol**

second edition

SIP

Understanding the Session Initiation Protocol

Second Edition

For a listing of recent titles in the *Armed Forces Telecommunications Library*,
see on the back of this book.

SIP

Understanding the Session Initiation Protocol

Second Edition

Alan B. Johnston



Artech House
Boston • London
www.artechhouse.com

Library of Congress Cataloging in Publication Data

A reading record of this title is available from the Library of Congress

British Library Cataloguing in Publication Data

Johnson, Alan B.

MP3 : Understanding the Service Initiative Protocol. — 2nd ed. — (Oxford Home Information Library)

I. Computer network protocols

I. Title

MP3L72

ISBN 1-58503-015-7

Cover design by Lisa Johnson

© 2004 ARTHUR BROSSEL, INC.

885 Canton Street
Newwood, MA 02060

All rights reserved. Printed and bound in the United States of America. No part of this book may be reproduced or stored in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without permission in writing from the publisher.

All items mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Arthur BrosseL cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

International Standard Book Number: 1-58503-015-7

05987454321

References

Contents

	Foreword to the First Edition	xxii
	Preface to the Second Edition	xix
	Preface to the First Edition	xxi
I	IP and the Internet	1
1.1	Signaling Protocols	1
1.2	The Internet: Engineering Task Force	2
1.3	A Brief History of IP	3
1.4	Internet: Multimedia Protocol Stack	4
1.4.1	Physical Layer	4
1.4.2	Internet Layer	4
1.4.3	Transport Layer	5
1.4.4	Application Layer	8
1.5	Utility Applications	9
1.6	DNS and IP Addresses	10
1.7	URLs and URIs	12
1.8	Mailcast	12
1.9	ARP Representation	13
	References	14

2	Introduction to SIP	11
2.1	A Simple Service Establishment Example	17
2.2	SIP Call with Proxy Server	29
2.3	SIP Registration Example	31
2.4	SIP Presence and Instant Message Example	33
2.5	Message Transport	38
2.5.1	UDP Transport	38
2.5.2	TCP Transport	40
2.5.3	TLS Transport	40
2.5.4	SCCP Transport	41
	References	42
3	SIP Clients and Servers	43
3.1	SIP User Agents	43
3.2	Presence Agents	44
3.3	Back-to-Back User Agents	45
3.4	SIP Gateways	49
3.5	SIP Servers	47
3.5.1	Proxy Servers	47
3.5.2	Redirect Servers	52
3.5.3	Registration Servers	53
3.6	Acknowledgment of Messages	59
3.7	Reliability	56
3.8	Authentication	57
3.9	SIP DRB, Encryption	58
3.10	Multicast Support	60
3.11	Firewalls and NAT Traversal	61
3.12	Proxies and Extensions for NAT Traversal	62
3.12.1	STUN Protocol	63
3.12.2	TURN Protocol	65
3.12.3	Other SIP/SIP NAT-Related Extensions	66
	References	68

4	SIP Request Messages	71
4.1	Methods	71
4.1.1	INVITE	72
4.1.2	REGISTER	74
4.1.3	BYE	76
4.1.4	ACK	77
4.1.5	CANCEL	79
4.1.6	OPTIONS	81
4.1.7	REFER	82
4.1.8	SUBSCRIBE	86
4.1.9	NOTIFY	89
4.1.10	MESSAGE	90
4.1.11	INFO	92
4.1.12	PRACK	94
4.1.13	UPDATE	96
4.2	URI and URL Schemes Used by SIP	98
4.2.1	SIP and SIP-URIs	98
4.2.2	Telephone URIs	100
4.2.3	Presence and Instant Messaging URIs	101
4.3	Tags	102
4.4	Message Bodies	102
	References	104
5	SIP Response Messages	105
5.1	Informational	106
5.1.1	100 Trying	106
5.1.2	100 Ringing	106
5.1.3	183 Call Is Being Forwarded	107
5.1.4	183 Call-Forwarded	109
5.1.5	183 Session Progress	108
5.2	Success	112
5.2.1	200 OK	112
5.2.2	202 Accepted	112
5.3	Redirection	112
5.3.1	300 Multiple-Choices	113

5.3.2	301 Moved Permanently	113
5.3.3	302 Moved Temporarily	113
5.3.4	303 See Other	113
5.3.5	308 Alternative Service	113
5.4	Client Error	113
5.4.1	400 Bad Request	114
5.4.2	401 Unauthorized	114
5.4.3	402 Payment Required	114
5.4.4	403 Forbidden	115
5.4.5	404 Not Found	115
5.4.6	405 Method Not Allowed	115
5.4.7	406 Not Acceptable	115
5.4.8	407 Proxy Authentication Required	115
5.4.9	408 Request Timeout	116
5.4.10	409 Conflict	116
5.4.11	410 Gone	116
5.4.12	411 Length Required	116
5.4.13	412 Request Entity Too Large	117
5.4.14	413 Request-URI Too Long	117
5.4.15	415 Unsupported Media Type	117
5.4.16	416 Unsupported URI Scheme	117
5.4.17	420 Bad Extension	117
5.4.18	421 Extension Required	117
5.4.19	422 Status-Header Interval Too Small	118
5.4.20	423 Interval Too Brief	118
5.4.21	429 Use Authentication Token	118
5.4.22	429 Provide Referer Identity	118
5.4.23	430 Temporarily Unavailable	119
5.4.24	431 Dialog-Transaction Does Not Exist	119
5.4.25	432 Loop Detected	119
5.4.26	433 Too Many Hops	119
5.4.27	434 Address Incomplete	120
5.4.28	435 Ambiguous	120
5.4.29	436 Hop Here	121
5.4.30	437 Request Terminated	122
5.4.31	438 Not Acceptable Here	122

5.4.10	401 Bad Request	122
5.4.11	401 Request Fording	122
5.4.12	403 Request Underscriptable	122
5.5	Server Error	123
5.5.1	500 Server Internal Error	123
5.5.2	501 Not Implemented	124
5.5.3	502 Bad Gateway	124
5.5.4	503 Service Unavailable	124
5.5.5	504 Gateway Timeout	124
5.5.6	505 Version Not Supported	124
5.5.7	503 Message Too Large	125
5.6	Global Error	125
5.6.1	400 Bad Request	125
5.6.2	403 Decline	125
5.6.3	404 Data Not In Our System	125
5.6.4	406 Not Acceptable	125
	References	126
6	HTTP Header Fields	127
6.1	Request and Response Header Fields	128
6.1.1	Accept	128
6.1.2	Accept-Charset	129
6.1.3	Cache-Control	129
6.1.4	Connection	130
6.1.5	Content	131
6.1.6	Date	131
6.1.7	Expires	131
6.1.8	From	131
6.1.9	Organization	134
6.1.10	Received	134
6.1.11	Referer	135
6.1.12	Subject	135
6.1.13	Supported	136
6.1.14	Timestamp	136
6.1.15	To	137
6.1.16	User-Agent	137

6.2.17	Via	138
6.2	Request Header Fields	140
6.2.1	Accept	140
6.2.2	Accept-Charset	140
6.2.3	Accept-Encoding	141
6.2.4	Accept-Language	141
6.2.5	Authentication	142
6.2.6	Cache-Control	142
6.2.7	From	143
6.2.8	Host	143
6.2.9	Is-Reply-To	143
6.2.10	Join	143
6.2.11	Priority	144
6.2.12	Proxy	145
6.2.13	Proxy-Authenticate	145
6.2.14	Proxy-Require	145
6.2.15	P-Content-Disposition	145
6.2.16	P-Handled-Identity	147
6.2.17	P-Preferred-Identity	147
6.2.18	Min-Forward	147
6.2.19	Reason	147
6.2.20	Refer-To	148
6.2.21	Referred-By	148
6.2.22	Reply-To	148
6.2.23	Require	150
6.2.24	Reject-Contact	150
6.2.25	Require-Disposition	151
6.2.26	Require	151
6.2.27	Response-Key	152
6.2.28	Route	152
6.2.29	Route	152
6.2.30	Service-Expires	153
6.2.31	Subscription-State	153
6.3	Response Header Fields	153
6.3.1	Authentication-Info	153
6.3.2	From-Info	154

6.3.3	MIME-Expires	154
6.3.4	MIME-SE	154
6.3.5	Proxy-Authorization	155
6.3.6	Referer	155
6.3.7	Unsupported	155
6.3.8	Warning	156
6.3.9	WWW-Authenticate	156
6.3.10	Retry	156
6.4	Message Body Header Fields	158
6.4.1	Allow	158
6.4.2	Content-Encoding	158
6.4.3	Content-Disposition	158
6.4.4	Content-Language	158
6.4.5	Content-Length	159
6.4.6	Content-Type	159
6.4.7	Expires	160
6.4.8	MIME-Version	160
	References	160
F	Related Protocols	162
7.1	SDP—Session Description Protocol	163
7.1.1	Protocol Version	165
7.1.2	Origin	165
7.1.3	Session Name and Information	166
7.1.4	URI	166
7.1.5	E-Mail Address and Phone Number	166
7.1.6	Connection Data	166
7.1.7	Bandwidth	167
7.1.8	Time, Repeat Times, and Time Zones	167
7.1.9	Encryption Keys	167
7.1.10	Media Annotations	168
7.1.11	Attributes	168
7.1.12	Use of SDP in SIP	169
7.2	RTP—Real-Time Transport Protocol	171
7.3	RTP Audio-Video Profile	174
7.4	PSTN Protocols	176

7.4.1	Client-Assisted Signaling	176
7.4.2	ISUP Signaling	176
7.4.3	ISDN Signaling	176
7.5	SIP for Telephones	177
7.6	Universal Plug and Play Protocol	178
	References	
8	Comparison to H.323	181
8.1	Introduction to H.323	181
8.2	Example of H.323	184
8.3	Versions	187
8.4	Comparison	187
8.4.1	Functional Differences	188
8.4.2	Examples of Each Protocol	198
8.5	Conclusions	191
	References	191
9	Windows and MGCP	198
9.1	IP Mobility	193
9.2	SIP Mobility	194
9.3	MGCP Architecture and SIP	201
9.4	MGCP Header Fields	203
9.4.1	Service-Route	203
9.4.2	Path	203
9.4.3	Other M-Headers	203
9.5	Future of SIP and Windows	204
	References	204
10	Call Flow Examples	208
10.1	SIP Call with Authentication, Proxies, and Record-Route	207
10.2	SIP Call with Suspend and Standby Proxies with Called Party Busy	214
10.3	SIP to PSTN Call Through Gateway	218
10.4	PSTN to SIP Call Through Gateway	221

18.5	Fossil Seeds	279
18.6	H.323 to SIP Call	280
18.7	SIP Window Call Flow	279
18.8	Call Setup Example with Two Parties	284
18.9	SIP Presence and Instant Message Example	290
	References	288
11	Future Directions	301
11.1	SIP, SIP-INVOLVING, and SIP-PLUS Working Group Design Teams	261
11.1.1	SIP and Streaming Impairments Design Team	262
11.1.2	Conferencing Design Team	262
11.1.3	Application Interaction Design Team	263
11.1.4	Emergency Calling Design Team	263
11.2	Other SIP Work Areas	263
11.2.1	Emergency Preparedness	263
11.2.2	Globally Reachable Contact URIs	263
11.2.3	Service Examples	263
11.3	SIP Instant Message and Presence Work	264
	References	264
	Appendix A: Changes in the SIP Specification from RFC 2543 to RFC 3261	303
	About the Author	311
	Index	323

Foreword to the First Edition

The Internet now challenges the close to \$1 trillion world telecom industry. A revolution in communications is taking place on the Internet. At its core are new communication protocols that would be inspired by the centralized-control systems of circuit-switched networks used in telecommunication.

The Internet and the World Wide Web can be technically defined only by their protocols. Similarly, IP telephony and the wider family of IP communications are defined by several key protocols, most notably by the Session Initiation Protocol, or SIP.

The previously closed class of telecommunication is now wide open to web developers because of SIP and its relation to the web HTTP 1.1 protocol and the e-mail SMTP protocol. IP communications include video/voice, presence, instant messaging, mobility, conferencing, and even games. We believe many other communication areas are yet to be invented. The integration of all types of communication on the Internet may represent the next “killer application” and generate yet another wave of Internet growth.

As explained in this book, SIP is a clear relative of the HTTP 1.1 and SMTP protocols. This represents a revolution in communications because it abandons the telecom signaling and control models developed for telephony over many years in favor of Internet and web-based protocols. Users and service providers obtain not only seamless integration of telephony and conferencing with many other World Wide Web and messaging applications, but also benefit from new forms of communications, such as presence and instant messaging.

Mobility can also be managed across various networks and devices using SIP. Location management is now under user control, so that incoming “calls” can be routed to any network and device that the called party may prefer. Users

may even roam across the globe to utilize services provided and maintain not only their URL “number”, but also their personal tailored services and preferences. The end user gains control over all possible preferences, depending on various parameters such as who the other party is, what network he is on and what devices he is using, as well as time of day, subject, and other variables.

The new dimension in communications called “presence” enables users for the first time to indulge in “polls calling” by first scouting presence and preferences of the other party, before making a call. In its turn, presence can trigger location- and time-dependent user preferences. Users may want to be contacted in different ways, depending on their location and type of network access.

E-commerce will also benefit from IP communications. Extremely complex telecom applications, as found in call centers, have become even more complex when integrated with e-mail and web applications for e-commerce. Such applications, however, are quite straightforward to implement using SIP, due to its seamless structure with the web and e-mail. For example, both call routing and e-mail routing to agents—based on various criteria such as queue length, skill set, time of day, customer ID, the web page the customer is looking at, and customer history—can be reduced to simple XML scripts when using SIP and another IETF standard, the Call Processing Language (CPL). These examples are in no way exhaustive, but are mentioned here as a way of introduction.

This book starts with a short summary of the Internet, the World Wide Web, and its core protocols and addressing. Though familiar to many readers, these chapters provide useful focus on issues for the topics ahead. The introduction to SIP is made easy and understandable by examples that illustrate the protocol architecture and message details. Finally, in the case of the book, a methodical and complete explanation of SIP is provided. We refer the reader to the Table of Contents for a better overview and navigation through the topics.

Alan Johnston has made significant contributions toward the use of SIP for communications over the Internet. I had the privilege of watching Alan in meetings with some of the largest telecom vendors as he went methodically line by line over hundreds of call flows, which were then submitted as an Internet Draft to the Internet Engineering Task Force (IETF) and implemented in commercial systems. Alan combines in this book his expertise and methodical approach with page turning narrative and a distinct sense of humor.

I could not help reading the book manuscript page by page, since everything from Internet basics, protocols, and SIP itself is explained so well, in an attractive and concise manner.

Henry Samueli
Distinguished Member of Engineering
WorldCom
Richardson, Texas
July 2000

Preface to the Second Edition

Much has changed in the 2.5 years since the first edition of *SIP: Understanding the Session Initiation Protocol* was published. In 2001, SIP was a relatively unknown quantity, an upstart in the voice over IP (VoIP) and multimodal communications industry. Today, SIP is seen as the future of call signaling and telephony. It has been widely deployed by service providers and enterprises and is used casually every day by users of the dominant PC operating systems. The full range of possibilities enabled by SIP is just now being glimpsed, and many more possibilities are yet to come.

One reason for this rapid acceptance is that SIP is an incredibly powerful call control protocol. It allows intelligent end points to implement the entire suite of telephony, Private Branch Exchange (PBX), Class, and Center services without a service provider, and without a controller or switch, for example.

The biggest driver for SIP on the Internet, however, has less to do with SIP's signaling and call control capabilities. Instead, it is due to the extension of SIP that turns it into a powerful "websterian" protocol that leverages mobility and presence to allow users to communicate using different devices, codes, and services anywhere they are connected to the Internet. SIP applications provide support for presence—the ability to find out the status or location of a user without attempting to set up a session.

Another major change in the past few years is the adoption of wireless SIP to enable multimodal IP communications. As described in the chapter on wireless, SIP is now being used both by its standard form over 802.11 wireless networks and by planned commercial Third Generation Partnership Project (3GPP) networks in the coming years. SIP is ideally suited for this key application.

Since 2004 SIP has also grown in terms of the specification itself. Initially, SIP was described by a single RFC with a few related RFCs and a couple of RFC extensions. In Chapter 6 alone, more than 20 SIP-related RFCs are referenced. This book attempts to put all these documents together and provide a single reference for the protocol and all its extensions. Even SIP headers and responses that were standardized in the past but are now reserved (deprecated) are listed in this text, providing useful context and background. Many others are discussed that are in the final stages of standardization prior to publication as RFCs, providing an open-data reader's view of the future of the protocol.

In closing, I again thank my colleagues in the Internet Engineering Task Force (IETF) and at IRI for all their contributions to the development of this protocol—it has been a privilege to be a part of a group of people that have created the SIP industry. Finally, I'd like to tip my hat to two of the key inventors of SIP who continue to develop and propel its implementation: Henning Schulzrinne and Jonathan Rosenberg.

Preface to the First Edition

When I began looking into the Session Initiation Protocol (SIP) in October 1998, I had prepared a list of a half-dozen protocols relating to Voice over IP and Next-Generation Networking. It was only a few days into my study that my list narrowed to just one: SIP. My background was in telecommunications, so I was familiar with the complex suite of protocols used for signaling in the Public Switched Telephone Network. It was readily apparent to me that SIP would be revolutionary in the telecommunications industry. Only a few weeks later I remember describing SIP to a colleague as the “OSI of future telephony”—quite a bold statement for a protocol that almost no one had heard of, and that was not even yet a proposed standard!

Nearly 2 years later, I have continued to work almost exclusively with SIP since that day in my position with WorldCom, giving seminars and teaching the protocol to others. This book grew out of those seminars and my work on what was Internet Drafts.

This revolutionary protocol was also the discovery of a radical standards body—the Internet Engineering Task Force (IETF). Later, I attended my first IETF meeting, which was for me a career-changing event. To interact with this dedicated band of engineers and developers, who have quickly taken the Internet from obscurity into one of the most important technological developments of the last 20th century, for the first time was truly exciting.

Just a few short years later, SIP has taken the telecommunications industry by storm. The industry gives concrete measurements after measurements of SIP product and service support from established vendor startups, and from established carriers. As each new group and company joins the dialog, the protocol has been able to adapt and grow without becoming unwieldy or overly complex. It

the future, I believe that SIP, along with a TCP/IP stack, will find its way into practically every intelligent electronic device that has a need to communicate with the outside world.

With my telecommunications background, it is not surprising that I rely on telephone examples and analogies throughout this book to explain and illustrate SIP. This is also consistent with the probability that telecommunication is the first widely deployed use of the protocol. SIP stacks will soon be in multi-media PCs, laptops, palmtops, and in dedicated SIP telephones. The protocol will be used by telephone switches, gateways, wireless devices, and mobile phones. One of the key features of SIP, however, is its flexibility; as a result, the protocol is likely to be used in a whole host of applications that have little or nothing to do with telephony. Quite possibly one of these applications, such as instant messaging, may become the next “killer application” of the Internet. However, the operation and concepts of the protocol are unchanged regardless of the application, and the telephone analogies and examples are, I feel, easy to follow and comprehend.

The book begins with a discussion of the Internet, the RTP, and the Internet Multimedia Protocol stack, of which SIP is a part. From there, the protocol is introduced by examples. Next, the elements of a SIP network are discussed, and the details of the protocol in terms of message types, headers, and response codes are covered. In order to make up a complete telephony system, related protocols, including Session Description Protocol (SDP) and Real-Time Transport Protocol (RTP), are covered. SIP is then compared to another signaling protocol, H.323, with the key advantages of SIP highlighted. Finally, the future direction of the evolution of the protocol is examined.

Two of the recurring themes of this book are the simplicity and stateless nature of the protocol. Simplicity is a hallmark of SIP due to its use of coded, highly readable messages, and its simple transactional model with few exceptions and special conditions. Statelessness relates to the ability of SIP servers to store minimal (or not) information about the state or existence of a media session in a network. The ability of a SIP network to use stateless servers that do not need to record transactions, keep logs, fill and empty buffers, etc., is, therefore, a critical step in the evolution of communication systems. I hope that these two themes become apparent as you read this book and learn about this exciting new protocol.

The text is filled with examples and sample SIP messages. I had to invent a whole set of IP addresses, domain names, and URLs. Please note that they are all fictional—do not try to send anything to them.

I would first like to thank the group of current and former engineers at WorldCom who shared their knowledge of this protocol and gave me the opportunity to author my first Internet-Draft document. I particularly thank Henry Stenrich, Steve Dentona, Dean Willis, and Matt Carlson. I also thank

Robert Sparks, who I first met at the first session on SEP that I ever presented. Throughout the whole 5-hour session I kept wondering about the guy with the pony tail who seemed to know more than me about this brand new protocol. Robert and I have spent countless hours discussing fine points of the protocol. In addition, I would like to thank him for his expert review of this manuscript prior to publication—it is a better book due to his thoroughness and attention to detail. I also thank everyone on the IETF SEP list who has assisted me with the protocol and added to my understanding of it.

A special thanks to my wife Lisa for the terrific career advice and the good figures throughout the book.

Finally, I thank my editor Jon Workman, the series editor and reviewer, and the whole team at Artech for helping me in this, my first adventure in publishing.

1

SIP and the Internet

The Session Initiation Protocol (SIP) is a new signaling, presence and instant messaging protocol developed to set up, modify, and tear-down multimedia sessions, requests and delivery processes and instant messages over the Internet [1]. This chapter covers some background for the understanding of the protocol. SIP was developed by the IETF as part of the Internet Multimedia Conferencing Architecture, and was designed to dovetail with other Internet protocols such as Transmission Control Protocol (TCP), Transmission Layer Security (TLS), User Datagram Protocol (UDP), Internet Protocol (IP), Domain Name System (DNS), and others. This organization and these related protocols will be briefly introduced. Related background topics such as Internet uniform resource indicators (URIs) and uniform resource locators (URLs), IP multicast routing, and Augmented Backus-Naur Form (ABNF) representations of protocol messages will also be covered.

1.1 Signaling Protocols

This book is about the Session Initiation Protocol. As the name implies, the protocol allows two end points to establish media sessions with each other. The main signaling functions of the protocol are as follows:

- Location of an end point;
- Contacting an end point to determine willingness to establish a session;
- Exchange of media information to allow session to be established;
- Modification of existing media sessions.

- Tear-down of existing media sessions.

SIP has also been extended to request and deliver presence information (on-line/off-line status and location information such as that contained in a "buddy" list) as well as instant message sessions. These functions include:

- Publishing and updating of presence information.
- Requesting delivery of presence information.
- Presence and other event notifications.
- Transporting of instant messages.

While most of the examples discuss SIP from a telephony perspective, there will be many non-telephony uses for SIP. SIP will likely be used to establish a whole set of session types that bear almost no resemblance to a telephony call.

1.2 The Internet Engineering Task Force

SIP was developed by the IETF. To quote "The Ties of the IETF" (2): "The Internet Engineering Task Force is a loosely self-organized group of people who make technical and policy contributions to the engineering and evolution of the Internet and its technologies." The two document types used within the IETF are Internet-Drafts (I-Ds) and Request for Comments (RFCs). I-Ds are the working documents of the group; anyone can author one on any topic and submit it to the IETF. There is no formal membership in the IETF; anyone can participate. Every I-D contains the following paragraph on the first page: "Internet-Drafts are 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.'"

Internet standards are achieved by the IETF as the Request for Comments series of numbered documents. As changes are made in a protocol, or new revisions come out, a new RFC document with a new number is issued, which "obsoletes" the old RFC. Some I-Ds are cited in this book; I have tried, however, to restrict this to mature documents that are likely to become RFCs by the time this book is published. A standard begins its life as an I-D and then progresses to an RFC once there is consensus and there are working implementations of the protocol. Anyone with Internet access can download any I-D or RFC at no charge using the World Wide Web, ftp, or e-mail. Information on how to do so can be found on the IETF Web site: <http://www.ietf.org>.

The IETF is organized into working groups, which are chartered to work in a particular area and develop a protocol to solve that particular area. Each

working group has its own archive and mailing list, which is where most of the work gets done. The IETF also meets three times per year.

1.3 A Brief History of SIP

SIP was originally developed by the IETF Multi-Party Multimedia Session Control Working Group, known as MMUSIC. Version 1.0 was submitted as an Internet-Draft in 1997. Significant changes were made to the protocol and resulted in a second version, version 2.0, which was submitted as an Internet-Draft in 1998. The proposal achieved Proposed Standard status in March 1999 and was published as RFC 2543 [3] in April 1999. In September 1999, the SIP working group was established by the IETF to coast the growing interest in the protocol. An Internet-Draft containing bug fixes and clarifications to SIP was submitted beginning in July 2000, referred to as RFC 2543 "bis". This document was eventually published as RFC 3261 [4], which obsoletes (or replaces) the original RFC 2543 specification. In addition, several SIP extension RFC documents have been published.

The popularity of SIP in the IETF has led to the formation of other SIP-related working groups. The Session Initiation Protocol Investigation (SIPPING) working group was formed to investigate applications of SIP, develop requirements for SIP extensions, and publish best current practice (BCP) documents about the use of SIP. This working group also publishes SIP event package RFCs. The SIP for Instant Messaging and Presence Leveraging Extensions (SIMPLE) working group was formed to standardize related protocols for presence and instant messaging applications. Other working groups that make use of SIP include PSTN and Internet Interworking (PINT) working group and Service in the PSTN/IN requiring Internet Services (SPRINT) working group.

To advance from Proposed Standard to Draft Standard, a protocol must have multiple independent implementations and limited operational experience. Since the early days of RFC 2543, SIP interoperability test events, called SIPs (formally called "boilerfests"), have been held a few times per year. For the latest information about SIPs, visit the SIP Forum Web site <http://www.sipforum.org>. (Note that the SIP Forum is a marketing/promotion organization for SIP and does not have any standardization function.) The final level, standard, is achieved after operational success has been demonstrated [5]. With the documented interoperability from SIPs, SIP will reach its Draft Standard status in the future.

SIP incorporates elements of two widely used Internet protocols: Hypertext Transfer Protocol (HTTP) used for Web browsing and Simple Mail Transfer Protocol (SMTP) used for e-mail. From HTTP, SIP borrowed a

client-server design and the use of URLs and URIs. From SMTP, SIP borrowed a text-encoding scheme and header style. For example, SIP uses SMTP headers such as To, From, Date, and Subject. The interaction of SIP with other Internet protocols such as IP, TCP, UDP, and DNS will be described in the next section.

1.8 Internet Multimedia Protocol Stack

Figure 1.1 shows the four-layer Internet Multimedia Protocol stack. The layers shown and protocols identified will be discussed.

1.8.1 Physical Layer

The lowest layer is the physical and link layer, which could be an Ethernet local area network (LAN), a telephone line (V.90 or 96k modem) carrying Point-to-Point Protocol (PPP), or a digital subscriber line (DSL) carrying asynchronous transfer mode (ATM), or even a wireless IEEE 802.11 network. This layer performs such functions as signal exchange, frame synchronization, and physical interface specifications.

1.8.2 Internet Layer

The next layer in Figure 1.1 is the Internet layer. Internet Protocol (IP) [5] is used at this layer to route a packet across the network using the destination IP address. IP is a connectionless, best-effort packet delivery protocol. IP packets can be lost, delayed, or received out of sequence. Each packet is routed on its own, using the IP header appended to the physical packet. Most IP address examples in this book use the older version of IP, version 4 (IPv4). IPv6



Figure 1.1 The Internet Multimedia Protocol stack.

addresses are four octets long, usually written in so-called “dotted decimal” notation (for example, 207.134.3.51). Between each of the dots is a decimal number between 0 and 255. At the IP layer, packets are not acknowledged. A checksum is calculated to detect corruption in the IP header, which could cause a packet to become misrouted. Corruption or errors in the IP payload, however, are not detected; a higher layer must perform this function if necessary. IP uses a single-octet protocol number in the packet header to identify the transport layer protocol that should receive the packet.

IP version 6 (IPv6) [7] was developed by the IETF as a replacement for IPv4. It has been steadily gaining support and is supported now by a number of operating systems. The biggest initial networks of IPv6 will likely be the wireless telephony carriers who care the most important advantage of IPv6 over IPv4—there is a much enlarged addressing space. IPv6 increases the addressing space from 32 bits in IPv4 to 128 bits, providing for over 4 billion IPv6 addresses. There are many other improvements in IPv6 over IPv4, including security. An IPv6 address is typically written as a sequence of eight hexadecimal numbers separated by colons. For example, 8:0:0:0:aaaa:bbbb:cccc:dddd is an IPv6 address written in this format. It is also common to drop sequences of zeros with a single double colon. The same address can then be written as ::aaaa:bbbb:cccc:dddd.

IP addresses used over the public Internet are assigned in blocks by the Internet Assigned Number Association (IANA). As a result of this centralized assignment, IP addresses are globally unique. This enables a packet to be routed across the public Internet using only the destination IP address. Various protocols are used to route packets over an IP network, but they are outside of the scope of this book. Subnetting and other aspects of the structure of IP addresses are also not covered here. There are other excellent sources [8] that cover the entire suite of TCP/IP protocols in more detail.

3.4.3 Transport Layer

The next layer shown in Figure 1.1 is the transport layer. It uses a two-octet port number from the application layer to deliver the datagrams or segments to the correct application layer protocol at the destination IP address. Some port numbers are dedicated to particular protocols—these ports are called “well-known” port numbers. For example, HTTP uses the well-known port number of 80, while SMTP uses the well-known port number of 2580 or 5883. Other port numbers can be used for any protocol, and they are assigned dynamically from a pool of available numbers. These so-called “ephemeral” port numbers are usually in the range 49152 to 65535. There are three commonly used transport layer protocols, Transmission Control Protocol, Transmission Layer Security, and User Datagram Protocol, which are described in the following sections.

14.2.1 TCP

TCP [9] provides reliable, connection-oriented transport over IP. TCP uses sequence numbers and positive acknowledgments to ensure that each block of data, called a segment, has been received. Lost segments are retransmitted until they are successfully received. Figure 1.2 shows the message exchange to establish and use down a TCP connection. A TCP server “listens” on a well-known port for a TCP request. The TCP client sends a SYN (synchronization) message to open the connection. The SYN message contains the initial sequence number the client will use during the connection. The server responds with a SYN message containing its own initial sequence number, and an acknowledgment number, indicating that it received the SYN from the client. The client completes the three-way handshake with an ACK or a DATA packet with the ACK flag set to the server acknowledging the server’s sequence number. Note that the connection is open, either client or server can send data in DATA packets called segments.

Each time a sender transmits a segment, it starts a timer. If a segment is lost in transit, this timer will expire. Detecting a lost segment, the sender will retransmit the segment until it receives the acknowledgment. The FIN message closes the TCP connection. The sequence of four messages shown in Figure 1.2 shows the connection. The ephemeral port numbers used in the connection are then free to be used in establishing other connections. TCP also has built-in mechanisms for flow control. During the SYN processing, a window size



Figure 1.2 Opening and closing a TCP connection.

representing the initial maximum number of unacknowledged segments is sent, which starts at 1 and increases exponentially up to a maximum limit. When network congestion and packet loss occur, the window resets back to 1 and gradually ramps back up again to the maximum limit. A TCP segment header contains 26 bytes. Forward segments are detected by a checksum covering both the TCP header and payload.

14.2.2 UDP

UDP [10] provides unreliable transport across the Internet. It is a low-effort delivery service, since there is no acknowledgment of sent datagrams. Most of the complexity of TCP is not present, including sequence numbers, automatic alignment, and window sizes. UDP does direct received datagrams with a checksum. It is up to higher layer protocols, however, to detect this datagram loss and initiate a retransmission, if desired.

14.2.3 TLS

TLS [11] is based on the Secure Sockets Layer (SSL) protocol first used in Web browsers, and it uses TCP for transport. TLS is commonly used today on the Internet for secure Web sites using the secure HTTP (HTTPS) URI scheme.

The TLS protocol has two layers: the TLS Transport protocol and the TLS Handshake protocol. The TLS Transport protocol is used to provide a reliable and private transport mechanism. Data sent using the TLS Transport protocol is encrypted so that a third party cannot intercept the data. A third party also cannot modify the transported data without one of the parties discovering this. The TLS Handshake protocol is used to establish the connection, negotiate the encryption keys used by the TLS Transport protocol, and provide authentication.

The key negotiation scheme selects an encryption algorithm and generates a session key based on a secret passed between the two sides. During the handshake, the parties exchange certificates, which can be used for authentication.

The cryptographic computations for a TLS connection are not trivial, and the multiple round trips needed to open a connection as shown in Figure 1.3 can add to message latency. Also, certificate verification can introduce processing delays. However, TLS transport has clear security advantages over UDP or TCP. Many operating systems already support TLS due to its wide use in secure Web browsers and servers.

14.2.4 SCTP

The Stream Control Transport Protocol (SCTP) [12] is similar to TCP in that it provides reliable stream-based transport. However, it has some advantages over TCP transport for a message-based protocol. First, it has built-in message segmentation, so that individual messages are separated at the transport layer.



Figure 1.2 Opening of a TCP connection.

Another advantage is that SCTP avoids the so-called “head-of-line (HoL)” problem of TCP. This is a common TCP problem in which a dropped segment with a large window causes the entire window’s worth of segments to wait in a buffer (i.e., be blocked) until the dropped segment is retransmitted.

SCTP also supports multihoming, so if one of a pair of host belonging across fails, the other can immediately begin receiving the message without encountering DNS or other database lookups.

SCTP is a next-layer-3 transport protocol that requires operating system-level support to be used, which will initially delay its use in the Internet. Also, note that the advantages of SCTP over TCP only occur during packet loss. In a loss-free network, the performance of the two is identical.

1.4.4 Application Layer

The top layer shown in Figure 1.1 is the application layer. This includes signaling protocols such as SIP and media transport protocols such as Real-time Transport Protocol (RTP), which is introduced in Section 1.1. Figure 1.1 includes H.323, introduced in Chapter 8, which is an alternative signaling protocol to SIP developed by the International Telecommunication Union (ITU). Session Encryption Protocol (SEP), described in Section 1.1, is shown above SIP in the protocol stack because it is carried in a SIP message body. SIP, RTP,

SMTP, FTP, and Telnet are all examples of application layer protocols. Because SIP can use any transport protocol, it is shown interacting with both TCP, TLS, and UDP in Figure 1.1. The use of TCP, TLS, SCTP, and UDP transport for SIP will be discussed in the next chapter.

1.5 Utility Applications

Two utility applications are also shown in Figure 1.1 as users of UDP. The most common use of the DNS (well-known port number 53) is to resolve a symbolic name (such as domain1.com, which is easy to remember) into an IP address (which is required by IP to route the packet). Also shown is the Dynamic Host Configuration Protocol (DHCP). DHCP allows an IP device to download configuration information upon initialization. Common fields include a dynamically assigned IP address, DNS addresses, subnet masks, maximum transmission unit (MTU), or maximum packet size, and server addresses for mail and Web browsing. Figure 1.4 shows the layer interaction for processing a request. At the top, a URL from the user layer is input to the application layer. URLs, described



Figure 1.4 Request processing in the Internet Protocol stack.

later in this chapter, are names used to represent resources, hosts, or files on the Internet. The application passes the generated request (for example, a HTTP GET request, which requires a Web page download), the URL, and the port number to the transport layer. The transport layer uses a utility to resolve the domain name extracted from the URL into an IP address. The IP address, along with the request, depending on the transport layer protocol used, and port-number identifying the transport protocol are then passed to the Internet layer. The Internet layer then passes the packet to the physical layer along with a media access control (MAC) address for routing, in the case of a LAN. A response is generated by reversing the above steps. A response received at the physical layer flows back up the layers, with the header information being stripped off and the response data passed upwards towards the user. The main difference is that no utility is used in response processing.

1.6 DNS and IP Addresses

Domain Name Service [13] is used in the Internet to map a symbolic name (such as `www.amazon.com`) to an IP address (such as `199.101.102.100`). DNS is also used to obtain information needed to route e-mail messages and, in the future, SIP messages. The use of names instead of numerical addresses is one of the Internet's greatest strengths because it gives the Internet a human, friendly feel. Domain names are organized in a hierarchy. Each level of the name is separated by a dot, with the highest-level domain on the right side. (Note that the dots in a domain name have no correspondence to the dots in an IP address written in dotted-decimal notation.) General top-level domains are shown in Table 1.1 (see <http://www.icann.org/bld> for the latest list). There is also a set of country domains such as `us` (United States), `uk` (United Kingdom), `ca` (Canada), and `au` (Australia). Each of these top-level domains has just one authority that assigns that domain to a user or group.

Once a domain name has been assigned, the authority places a link in their DNS server to the DNS server of the user or group who has been assigned the domain. For example, when `company.com` is allocated to a company, the authoritative DNS server for the top-level `com` domain `com` for `company.com` contains the IP address of the company's DNS server(s). A name can then be further qualified by entries in the company's DNS server to point to individual servers in their network. For example, the company's DNS server may contain entries for `www.company.com`, `ftp.company.com`, and `smtp.company.com`. A number of types of DNS record types are defined. The DNS records used to resolve a host name into an IP address are called address records, or *A records*. Other types of records include *CNAME* (or canonical name or alias records), *MX* (or mail exchange records), *SRV* (or Service records, used by SIP and other

Table 11
General Top-Level Domains

Domain	Description
com	Company
net	Network
int	Internet
org	Not-for-profit organization
edu	Education or college
gov	U.S. government
mil	U.S. military
arpa	ARPA
info	Information
biz	Business
museum	Museum
name	Name
pro	Professional
nom	For foreign identity
coop	Cooperation

protonic), and TXT (no free-form text records). Another type of DNS record is a PTR, or pointer record, used for reverse lookups. Reverse lookups are used to map an IP address back to a domain name. These records can be used, for example, in generating server logs that show not only the IP addresses of clients served, but also their domain name. Web browsing provides an example of the use of the DNS system. Another type of DNS record is known as a Naming Authority Pointer (NAPTR) record that can be used by a protocol known as ENUM [14] to map global telephone numbers into Internet URLs.

When a user types in a Web address, such as `www.apple.com`, the name must be resolved to an IP address before the browser can send the request for the index Web page from the Apple Home Web server. The Web browser first launches a DNS query to the IP address for its DNS server, which has been manually configured or set up using DHCP. If the DNS server happens to have the name's *A* record stored locally (cached) from a recent query, it will return the IP address. If not, the DNS root server will then be queried to locate the authoritative DNS server for Apple Home, which must contain the *A* records for the `apple.com` domain. The HTTP GET

request is then sent to that IP address, and the Web browsing session begins. There is only one authoritative DNS server for a domain, and it is operated by the owner of the domain name. Due to a very efficient caching scheme built into DNS, a DNS request often does not have to travel all the way to this server. DNS is also used by an SMTP server to deliver an e-mail message. An SMTP server with an e-mail message to deliver initiates a DNS request for the MX record of the domain name in the destination e-mail address. The response to the request allows the SMTP server to connect the destination SMTP server and transfer the message. A similar process has been proposed for locating a SIP server using SRV, or service, DNS records.

1.7 URLs and URIs

Uniform resource locators [15] are names used to represent addresses or locations in the Internet. URLs are designed to encompass a wide range of protocols and resource types in the Internet. The basic form of a URL is *scheme:specifier*—for example, `http://www.artechhouse.com/search/search.html`. The token `http` identifies the scheme or protocol to be used, in this case HTTP. The specifier follows the “/” and contains a domain name (`www.artechhouse.com`), which can be resolved here as an IP address and a file name (`/search/search.html`). URLs can also contain additional parameters or qualifiers relating to transport. For example, `telnet://host.company.com:24` indicates that the Telnet Protocol should be used to access `host.company.com` using port 24. New schemes for URLs for new protocols are easily constructed, and dozens have been defined, such as `mailto`, `tel`, and `https`. The `sip` and `sips` schemes will be introduced in Section 2.2 and are described in detail in Section 4.2.

Most protocols reference URLs, but with SIP we mainly reference URIs. This is due to the mobility aspects of SIP, which means that a particular address (URI) is not tied to a single physical device but instead is a logical entity that may move around and change its location in the Internet. However, the terms URL and URI are often used almost interchangeably in other systems.

1.8 Multicast

In normal Internet packet routing, or unicast routing, a packet is routed to a single destination. In multicast routing, a single packet is routed to a set of destinations. Single LAN segments running a protocol such as Ethernet offer the capability for packet broadcast, where a packet is sent to every node on the network. Scaling this to a larger network with routers is a recipe for disaster, as

broadcast traffic can quickly cause congestion. An alternative approach for this type of packet distribution is to use a packet reflector that receives packets and forwards copies to all destinations that are members of a broadcast group. This also can cause congestion in the form of a "packet storm" [16]. For a number of years, the Internet Multicast Backbone Network (MBONE), an overlay of the public Internet, has used multicast routing for high-bandwidth broadcast services. Participants who wish to join a multicast session send a request to join the session to their local MBONE router. This router will then begin to broadcast the multicast session on that LAN segment. Additional requests to join the session from others in the same LAN segment will result in no additional multicast packets being sent, since the packets are already being broadcast. If the router is not aware of any multicast participants on its segment, it will not forward any of the packets. Routing of multicast packets between routers uses special multicast routing protocols to ensure that packet traffic on the backbone is kept to a minimum. Multicast Internet addresses are reserved in the range 224.0.0.0 to 239.255.255.255. Multicast transport is always UDP, since the handshake and acknowledgments of TCP are not possible. Certain addresses have been defined for certain protocols and applications. The scope or extent of a multicast session can be limited using the time-to-live (TTL) field in the IP header. This field is decremented by each router that forwards the packet, which limits the number of hops the packet takes. SIP support for multicast will be discussed in Section 3.18. Multicast is slowly becoming a part of the public Internet as service providers begin supporting it.

1.9 ABNF Representation

The meta-language Augmented Backus-Naur Form [17] is used throughout RFC 3261 [1] to describe the syntax of SIP, as well as other Internet protocols. An example construct used to describe a SIP message is as follows:

```

SIP-message = Request / Response

```

This is read: A SIP message is either a request or a response. `SIP-message` on the left side of the "equals" sign represents what is being defined. The right side of the "equals" sign contains the definition. The "/" is used to mean logical OR (note: older versions of ABNF used "V" in place of "/" to grammar). Next, `Request` and `Response`s are defined in a similar manner using ABNF:

```

Request = Request-Line * message-header ; CRLF ; message-body ;

```

`Request-Line` will be defined in another ABNF statement. Elements enclosed in () are treated as a single element. The "*" means the element may be

repeated, separated by at least one space. The minimum and maximum numbers can be represented as x^y , which means a minimum of x and maximum of y . Since the default values are 0 and infinity, a solitary "*" (as in this example) indicates any number is allowed, including none. CR/LF is defined as a carriage return/line feed, or the ASCII characters that are written in Internet textual notation as `#x0D` and `#x0A`. Other common ABNF representations include `SP` for space (ASCII 0x20). A message body is optional in a Request, and is enclosed in square brackets `[]` to indicate this. Comments in the ABNF begin with a semicolon `;` and continue to the end of the line. Lines continue the same ABNF definition when they are indented. Tokens are defined in ABNF as any set of characters besides control characters and separators. Display names and other components of a SIP header that are not used by the protocol are considered tokens; they are simply parsed and ignored. For example:

```
transport < "udp" / "tcp" / "tls" / "sctp" >
```

This ABNF defines four possible values, which may be used in a transport parameter. In this text, few references to ABNF will be made. Instead, SIP messages and elements will be introduced by descriptions and examples rather than by using ABNF. The ABNF for SIP is in Section 25 of RFC 3261.

References

- [1] Litterer, R., et al., "A Brief History of the Internet," The Internet Society, <http://www.internetsociety.org/history/brief.html>.
- [2] Malkin, G., and the IETF Secretariat, "The Use of the HTTP—A Guide for New Arrivals of the Internet Engineering Task Force," <http://www.ietf.org/rfc.html>.
- [3] Handley, M., et al., "SIP Session Initiation Protocol," RFC 2543, 1999.
- [4] Rosenberg, J., et al., "SIP Session Initiation Protocol," RFC 3261, 2002.
- [5] Bushon, S., "The Internet Standards Process Revisited 3," RFC 2026, 1998.
- [6] "Stream Protocol," RFC 791, 1980.
- [7] Davie, S., and S. Handley, "Internet Protocol, Version 6 (IPv6) Specification," RFC 1889, 1995.
- [8] White, R., *A Guide to the TCP/IP Protocol Suite*, Microsoft, MA: Graphix House, 1998.
- [9] "Transmission Control Protocol," RFC 793, 1981.
- [10] Postel, J., "User Datagram Protocol," RFC 768, 1980.
- [11] Davie, S., et al., "The UDP Protocol Version 1.0," RFC 2296, 1998.
- [12] Savari, E., et al., "Stream-Control Transmission Protocol," RFC 2960, 1999.
- [13] Manning, B., "DNS NSID RR," RFC 1448, 1992.
- [14] Edman, P., "E-MAIL Header and CRLF," RFC 1846, 2000.

-
- [14] Bennett-Lee, T., L. Blanton, and M. McCull, "Uniform Resource Locators," RFC 1738, 1994.
 - [15] Harsanyi, G., D. Guck, and J. Park, *IP Telephony Packet-Based Multimedia Communication Systems Handbook*, McGraw-Hill/Wiley, 2000, Chapter 8.
 - [16] Crockett, D., "Standard for the Format of ARPANET Internet Text Messages," RFC 822, 1982.

2

Introduction to SIP

Often the best way to learn a protocol is to look at examples of its use. While the terminology, structure, and format of a new protocol can be confusing at first read, an example message flow can give a quick grasp of some of the key concepts of a protocol. The example message exchanges in this chapter will illustrate SIP as defined by RFC 3261 [1].

The first example shows the basic message exchange between two SIP devices to establish and tear-down a session. The second example shows the message exchange when a SIP proxy server is used. The third example shows SIP registration. The fourth example shows a SIP presence and instant message example. The chapter concludes with a discussion of SIP message transmission using UDP, TCP, TLS, and SCTP.

The examples will be introduced using call flow diagrams between a called and calling party, along with the details of each message. Each arrow in the figures represents a SIP message, with the arrowhead indicating the direction of transmission. The thick lines in the figures indicate the media stream. In these examples, the media will be assumed to be RTP [2] packets containing media, but it could be another protocol. Details of RTP are covered in Section 7.2.

2.1 A Simple Session Establishment Example

Figure 2.1 shows the SIP message exchange between two SIP-enabled devices. The two devices could be SIP phones, hand-helds, palmtops, or cell phones. It is assumed that both devices are connected to an IP network such as the Internet and know each other's IP address.

Since SIP is a text-encoded protocol, this is actually what the SIP message would look like “on the wire” as a UDP datagram being transported over, for example, Ethernet.

The fields listed in the `INVITE` message are called header fields. They have the form `Header: Value-CRLF`. The first line of the request message, called the start line, lists the method, which is `INVITE`, the Request-URI, then the SIP version number (2.0), all separated by spaces. Each line of a SIP message is terminated by a CRLF. The Request-URI is a special form of SIP URI and indicates the resource to which the request is being sent, also known as the request target. SIP URIs are discussed in more detail in later sections.

The first header field following the start line shown is a `Via` header field. Each SIP device that originates or forwards a SIP message stamps its own address in a `Via` header field, usually written as a host name that can be resolved into an IP address using a DNS query. The `Via` header field contains the SIP version number (2.0), a “/”, then `UDP` for UDP transport, a space, then the hostname or address, a colon, then a port number, in this example the “well-known” SIP port number 5060. Transport of SIP using TCP, UDP, TLS, and SCTP and the use of port numbers are covered later in this chapter. The `branch` parameter is a transaction identifier. Responses relating to this request can be correlated because they will contain this same transaction identifier.

The next header field shown is the `Max-Forwards` header field, which is initialized to some large integer and decremented by each SIP server, which receives and forwards the request, providing simple loop detection.

The next header fields are the `To` and `From` header fields, which show the originator and destination of the SIP request. When a name label is used, as in this example, the SIP URI is enclosed in brackets and used for routing the request. The name label could be displayed during dialing, for example, but is not used by the protocol.

The `Call-ID` header field is an identifier used to keep track of a particular SIP session. The originator of the request creates a locally unique string, then usually adds an “S” and its host name to make it globally unique. In addition to the `Call-ID`, each party in the session also contributes a random identifier, unique for each call. These identifiers, called tags, are included in the `To` and `From` header fields as the session is established. The initial `INVITE` shown contains a `From` tag but no `To` tag.

The next agent that generates the initial `INVITE` to establish the session generates the unique `Call-ID` and `From` tag. In the response to the `INVITE`, the next agent answering the request will generate the `To` tag. The combination of the local tag (contained in the `From` header field), remote tag (contained in the `To` header field), and the `Call-ID` uniquely identifies the established session, known as a “dialog.” This dialog identifier is used by both

parties to identify this call because they could have multiple calls set up between them. Subsequent requests within the established session will not this dialog identifier, as will be shown in the following examples.

The `seq` header field shown in the CSeq, an `in-dialog` sequence. It contains a number, followed by the method name, `INVITE` in this case. This number is incremented for each new request sent. In this example, the `seq` header field value is initialized to 1, but it could start at another integer value.

The Via header fields plus the `Max-Forwards`, `To`, `From`, `Call-ID`, and `CSeq` header fields represent the minimum required header field set in any SIP request message. Other header fields can be included as optional additional information, or information needed for a specific request type. A `Contact` header field is also required in this `INVITE` message, which contains the SIP URI of Tada's communication device, known as a user agent (UA); this URI can be used to route messages directly to Tada. The optional `State-Id` header field is present in this example. It is not used by the protocol, but could be displayed during dialing to aid the called party in deciding whether to accept the call. The same sort of useful prioritization and screening commonly performed using the `State-Id` and `From` header fields in an e-mail message is also possible with a SIP `INVITE` request. Additional header fields are present in this `INVITE` message, which contain the media information necessary to set up the call.

The `Content-Type` and `Content-Length` header fields indicate that the message body is SDP [8] and contains 158 bytes of data. The basis for the byte count of 158 is shown in Table 2.1, where the CR LF at the end of each line is shown as a 000 and the byte count for each line is shown on the right-hand side. A blank line separates the message body from the header field line, which ends with the `Content-Length` header field. In this case, there are seven lines of SDP data describing the media attributes that the caller Tada desires for the call. This media information is needed because SIP makes no assumptions about the type of media session to be established—the caller must specify exactly what type of session (audio, video, gaming) that he wishes to establish. The SDP field names are listed in Table 2.2, and will be discussed detail in Section 7.1, but a quick review of the lines shows the basic information necessary to establish a session.

Table 2.2 includes the:

- Connection IP address (192.168.1.10)
- Media format (audio)
- Port number (5060)
- Media transport protocol (RTP)
- Media encoding (PCMu Law)

The 180 Ringing is an example of a SIP response message. Responses are numerical and are classified by the first digit of the number. A 180 response is an “informational class” response, identified by the first digit being a 1. Informational responses are used to convey contextual information about the progress of the call. Many SIP response codes were based on HTTP version 1.1 response codes with some extensions and additions. Anyone who has ever browsed the World Wide Web has likely received a “404 Not Found” response from a Web server when a requested page was not found. 404 Not Found is also a valid SIP “class error class” response to a request in an unknown case. The other classes of SIP responses are covered in Chapter 5.

The response code number in SIP also determines the way the response is interpreted by the server or the user. The reason phrase, Ringing in this case, is suggested in the standard, but any text can be used to convey extra information. For instance, 180 Hold your horses. I’m trying to make him up! is a perfectly valid SIP response and has the same meaning as a 180 Ringing response.

The 180 Ringing response has the following structure:

```
SIP/1.0 180 Ringing
Via: SIP/2.0/UDP lab-high-voltage.org;branch=invite;method=invite
To: B. Marconi <mailto:marconi@lab-high-voltage.org>;tag=654321
From: M. Tola <mailto:tola@lab-high-voltage.org>;tag=12345
Call-ID: 1234567890lab-high-voltage.org
CSeq: 1 INVITE
Contact: <mailto:marconi@lab-high-voltage.org>
Content-Length: 0
```

The message was created by copying many of the header fields from the INVITE message, including the Via, To, From, Call-ID, and CSeq, then adding a response status line containing the SIP version number, the response code, and the reason phrase. This approach simplifies the message processing for responses.

The Via header field contains the original branch parameter but also has an additional received parameter. This parameter contains the local IP address that the request was received from (199.101.102.100), which typically is the same address that the URI in the Via resolves using DNS (lab.high-voltage.org).

Note that the To and From header fields are not reversed in the response message as one might expect them to be. Even though this message is sent to Marconi from Tola, the header fields read the opposite. This is because the To and From header fields in SIP are defined to indicate the direction of the request, not the direction of the message. Since Tola initiated this request, all responses will read To: Marconi From: Tola.

request. This exchange of media information allows the media session to be established using another protocol, RTP in this example.

```

100 200msec;method=invite;media=audio;seq=1000;0
Via: SIP/2.0/UDP 192.168.1.100;branch=1000;method=invite
Max-Forwards: 70
To: 0;branch=1000;method=invite;seq=1000;0
From: 1000@192.168.1.100;tag=1000;branch=1000;seq=1000;0
Call-ID: 1000@192.168.1.100;branch=1000
CSeq: 1 ACK
Content-Length: 0
  
```

The returned response, CSeq, has the same number as the INVITE, but the method is set to ACK. At this point, the media session begins using the media information carried in the SIP messages. The media session takes place using another protocol, typically RTP. The branch parameter in the Via header field contains a new transaction identifier than the INVITE, since an ACK sent to acknowledge a 200 OK is considered a separate transaction.

This message exchange shows that SIP is an end-to-end signaling protocol. A SIP network, or SIP server, is not required for the protocol to be used. Two end points running a SIP protocol stack and knowing each other's IP addresses can use SIP to set up a media session between them. Although less obvious, this example also shows the client-server nature of the SIP protocol. When Teds originates the INVITE request, he is acting as a SIP client. When Marcos responds to the request, he is acting as a SIP server. After the media session is established, Marcos originates the ACK request and acts as the SIP client, while Teds acts as the SIP server when he responds. This is why a SIP-enabled device must contain both SIP server and SIP client software—during a typical session, both are needed. This is quite different from other client-server Internet protocols such as HTTP or FTP. The Web browser is always an HTTP client, and the Web server is always an HTTP server, and similarly for FTP. In SIP, an end point will switch back and forth during a session between being a client and a server.

In Figure 2.1, a BYE request is sent by Marcos to terminate the media session:

```

BYE 1000@192.168.1.100;branch=1000;seq=1000;0
Via: SIP/2.0/UDP 192.168.1.100;branch=1000;method=bye
Max-Forwards: 70
To: 1000@192.168.1.100;tag=1000;method=bye;seq=1000;0
From: 1000@192.168.1.100;tag=1000;method=bye;seq=1000;0
Call-ID: 1000@192.168.1.100;branch=1000
CSeq: 1 BYE
Content-Length: 0
  
```

The Via header field in this example is populated with Maroon's host address and contains a new transaction identifier since the BYE is considered a separate transaction from the INVITE or ACK transactions shown previously. The From header field reflects that this request is originated by Maroon, as they are derived from the messages in the previous transaction. Teles, however, is able to identify the dialog using the presence of the same local and remote tag and Call-ID in the INVITE, and tear-down the correct media session.

Notice that all the branch IDs shown in the example so far begin with the string `af38244c`. This is a special string that indicates that the branch ID has been calculated using strict rules defined in RFC 3261 and is as a result usable as a transaction identifier.¹

The confirmation response to the BYE is a 200 OK:

```

200 OK 200 OK
Via: SIP/2.0/UDP 192.168.1.100;branch=af38244c-192.168.1.100;transport=UDP
From: Maroon <192.168.1.100@192.168.1.100>
To: Teles <192.168.1.100@192.168.1.100>
Content-Length: 0

```

The response reflects the Call-ID of the original request: `af38244c`.

3.2 SIP Call with Proxy Server

In the SIP message exchange of Figure 2.3, Teles knew the IP address of Maroon and was able to send the INVITE directly to that address. This will not be the case in general—an IP address cannot be used like a telephone number. One reason is that IP addresses are often dynamically assigned due to the shortage of IPv4 addresses. For example, when a PC dials in to an Internet service provider (ISP) modem bank, an IP address is assigned using DHCP to the PC from a pool of available addresses allocated to the ISP. For the duration of the session, the IP address does not change, but it is different for each dial-in session. Even for an “always-on” Internet connection such as a DSL line, a different IP address can be assigned after each reboot of the PC. Also, an IP address does not uniquely identify a user, but identifies a user as a particular physical IP network. You have one IP address at your office, another at home, and still another when you log on remotely when you travel. Ideally, there would be one address

1. This string is needed because branch IDs generated by user agents prior to RFC 3261 may have concatenated branch IDs which are not suitable as transaction identifiers. In this case, a dialog must continue its own-transaction identifier using the To-tag, From-tag, Call-ID, and Call-ID-tag.

that would identify you wherever you are. In fact, there is an Internet protocol that does exactly that, with e-mail. SMTP uses a host or system independent name (an e-mail address) that does not correspond to a particular IP address. It allows e-mail messages to reach you regardless of what your IP address is and where you are logged on to the Internet.

In addition, a request routed using only IP addresses will reach only one end point—only one device. Since communication is typically user-to-user instead of device-to-device, a more useful addressing scheme would allow a particular user to call another particular user, which would result in the request reaching the target user regardless of which device they are currently using, as if they have multiple devices.

SIP uses e-mail-like names for addresses. The addressing scheme is part of a family of Internet addresses known as URIs. SIP URIs can also handle telephone numbers, transport parameters, and a number of other items. A full description, including examples, can be found in Section 4.2. For now, the key point is that a SIP URI is a name that is resolved to an IP address by using SIP proxy server and DNS lookups at the time of the call, as will be seen in the next example. Figure 2.1 shows an example of a more typical SIP call with a type of SIP server called a “proxy server.” In this example, the caller Schenckinger calls Heisenberg through a SIP proxy server. A SIP proxy operates in a similar way to a proxy in HTTP and other Internet protocols. A SIP proxy does not set up or terminate sessions, but sits in the middle of a SIP message exchange, receiving messages and forwarding them. This example shows one proxy, but there can be multiple proxies in a signaling path.

SIP has two broad categories of URIs: ones that correspond to a user, and ones that correspond to a single device or end point. The user URI is known as an address of record (AOR) and a request sent to an address of record will require database lookups and server and future operations and can result in the request being sent to one or more end devices. A device URI is known as a contact, and typically does not require database lookups. An address of record URI is usually used in To and From header fields, as this is the general way to reach a person and is suitable for storing in address books and for remembering missed calls. A device URI is usually used in a Contact: header field and is associated with a particular user for a shorter period of time. The method of relating (or binding) a contact URI with an address of record URI will be discussed in Section 2.3.

Because Schenckinger does not know exactly where Heisenberg is currently logged on, and which device they are currently using, a SIP proxy server is used to route the INVITE. First, a DNS lookup of Heisenberg’s SIP URI domain name (`intell.ch`, `ch`) is performed, which returns the IP address of the proxy server (`proxy.marlets.ch`, which handles the domain). The INVITE is then sent to that IP address.




Figure 2.2 SIP call example with proxy server.

```

user-agent: sip-proxy, helmsberg@helms.ch; dir=00000000
Via: SIP/2.0/UDP 100.201.100.201; branch=00000000000000000000 To
Max-Forwards: 70
To: Hamburg; dir=00000000, helmsberg@helms.ch; dir=00000000
From: H. Helmsberg; dir=00000000000000000000, user-agent=
00000000
Call-ID: 1000000000.100.201.100
CSeq: 1 INVITE
Subject: Where are you exactly?
Contact: helmsberg@helms.ch; dir=00000000
request-type: application/sdp
Content-Length: 100

V=0
o=helmsberg0001 1000000000 1000000000 IN 100 100.201.100.201
s=0000000000
c=IN
m=0 100 100.201.100.201
a=media: 0 0 0 0 0 0
a=rtpmap: 0 none/0 0

```

The proxy looks up the SIP URI in the Request-URI (`dir=00000000, helmsberg@helms.ch; dir=00000000`) in its database and locates Helmsberg. This completes the two-step process:

1. DNS lookup by user agent to locate the IP address of the proxy; database lookup is performed by the proxy to locate the IP address;

- The **INVITE** is then forwarded to Heisenberg's IP address with the addition of a second **Via** header field stamped with the address of the proxy:

```

INVITE sip:marion.heisenberg@505.201.300.201 SIP/2.0
Via: SIP/2.0/UDP proxy.max.oh.de;branch=0;transport=UDP
Via: SIP/2.0/UDP 100.201.100.201:5055;branch=0;transport=
User-Agent: 0
To: Heisenberg <sip:marion.heisenberg@505.201.300.201>
From: marion <sip:marion@505.201.100.201>
Call-ID: 100000.100.201.100
Seq: 1 20000
Content-Type: application/sdp
Content-Length: 120

m=0 0 0
m=audio 0 0 1 100000000 2P 24 100.201.100.201
a=rtpmap 0 100 0
c=IN 24 100.201.100.201
t=0 0
s=00000 0 0 0 0 0 0
m=video 0 0 0 0 0

```

From the presence of two **Via** header fields, Heisenberg knows that the **INVITE** has been routed through a proxy server. Having received the **INVITE**, a **180 Ringing** response is sent by Heisenberg to the proxy:

```

200 OK 100 Ringing
Via: SIP/2.0/UDP proxy.max.oh.de;branch=0;transport=UDP
Via: SIP/2.0/UDP 100.201.100.201:5055;branch=0;transport=
To: Heisenberg <sip:marion.heisenberg@505.201.300.201>
From: M. Marion <marion@505.201.100.201>
Call-ID: 100000.100.201.100
Seq: 1 20000
Content-Type: application/sdp

```

Again, this response contains the **Via** header fields, and the **To**, **From**, **Call-ID** and **Seq** header fields from the **INVITE** request. The response is then sent to the address in the first **Via** header field, `proxy.max.oh.de` (i.e. the port number listed in the **Via** header field 5055, in this case). Notice that the **To** header field now has a tag added to it to identify this particular dialog. Only the first **Via** header field contains a related **req** parameter, since the second **Via** header already contains the literal IP address in the URI. The **Content** header field contains the device URI of Heisenberg.

The proxy receives the response, checks that the first **Via** header field has its own address (`proxy.max.oh.de`), uses the transaction identifier in the

Via header, then removes that Via header field, then forwards the response to the address in the next Via header field: IP address 100.101.102.103, port 5000. The resulting response sent by the proxy to Schroedinger is:

```
HTTP/1.0 200 OK
Via: SIP/2.0/UDP 100.101.102.103; branch=00000000000000000000000000000000
To: schroedinger <sip:schro@schroedinger.com>; tag=114450
From: E. Schrodinger <e.schro@schroedinger.com>; tag=1
Call-ID: 100000-101-101-101
Contact: <sip:schro@schroedinger.com>; <tel:100-100-100>
Content-Length: 0
```

The use of Via header fields in routing and forwarding SIP messages reduces complexity in message forwarding. The response required a database lookup by the proxy to be routed. The response requires no lookup because the routing is included in the message in the Via header fields. Also, this means that responses come back through the same set of proxies as the request. The call is accepted by Heisenberg, who sends a 200 OK response:

```
HTTP/1.0 200 OK
Via: SIP/2.0/UDP proxy.example.com; branch=00000000000000000000000000000000
(Proxy-Require: 001-101-101-101)
Via: SIP/2.0/UDP 100.101.102.103; branch=00000000000000000000000000000000
To: schroedinger <sip:schro@schroedinger.com>; tag=114450
From: E. Schrodinger <e.schro@schroedinger.com>; tag=1
Call-ID: 100000-101-101-101
CSeq: 1 INVITE
Contact: <sip:schro@schroedinger.com>; <tel:100-100-100>
Content-Type: application/sdp
Content-Length: 100

v=0
o=schroedinger 2000000000 2000000000 IN 100 100 101 100 101
s=Phone Call
c=IN 100 100 101 100 101
t=0
m=audio 100 100 100 0
a=CCMID: 0 00000000
```

The proxy forwards the 200 OK message to Schroedinger after removing the first Via header field:

```
HTTP/1.0 200 OK
Via: SIP/2.0/UDP 100.101.102.103; branch=00000000000000000000000000000000
To: schroedinger <sip:schro@schroedinger.com>; tag=114450
From: E. Schrodinger <e.schro@schroedinger.com>; tag=1
Call-ID: 100000-101-101-101
CSeq: 1 INVITE
Contact: <sip:schro@schroedinger.com>; <tel:100-100-100>
Content-Type: application/sdp
```

```
Content-Length: 129
```

```
To:
Helsinkiberg 2000042125.2000041224 DE 204.200.201.100.201
From: 204.200.201.100.201
Call-ID:
Message-ID: 2000042125.2000041224
```

The presence of the Contact header field with the SIP URI address of Helsinkiberg in the 200 OK allows Schroedinger to send the ACK directly to Helsinkiberg, bypassing the proxy. (Note that the Request-URI is set to Helsinkiberg's Contact URI and not the URI in its To header field.) This request, and all future requests continue to use the tag in the To header field:

```
ACK 204.200.201.100.201.2000042125.2000041224 SIP/2.0
Via: SIP/2.0/UDP 204.200.201.100.201; branch=2000042125.2000041224
Max-Forwards: 70
To: Helsinkiberg <tag=2000042125.2000041224@204.200.201.100.201>
From: Schroedinger <tag=2000042125.2000041224@204.200.201.100.201>
Call-ID: 2000042125.2000041224
Msg: 1 ACK
Content-Length: 0
```

This shows that the proxy server is not really “in the call.” It facilitates the two-end-points locating and contacting each other, but it can drop out of the signaling path as soon as it no longer adds any value to the exchange. A proxy server can finish further managing its state through its learning a Branch-Record header field, which is described in Section 6.2.12. In addition, it is possible to have a proxy server that does not retain any knowledge of the fact that there is a session established between Schroedinger and Helsinkiberg (inferred to as call state information). This is discussed in Section 3.3.1. Note that the media is always end-to-end and not through the proxy.

In SIP the path of the signaling message is totally independent of the path of the media. In telephony, this is described as the separation of control channel and bearer channel.

The media session is ended when Helsinkiberg sends a BYE message:

```
BYE tag=2000042125.2000041224@204.200.201.100.201 SIP/2.0
Via: SIP/2.0/UDP 204.200.201.100.201; branch=2000042125.2000041224
Max-Forwards: 70
To: Helsinkiberg <tag=2000042125.2000041224@204.200.201.100.201>
From: Helsinkiberg <tag=2000042125.2000041224@204.200.201.100.201>
Call-ID: 2000042125.2000041224
Content-Length: 0
```

Note that Heisenberg's CSeq was initialized to 2000. Each SIP device maintains its own independent CSeq number space. This is explained in more detail in Section 8.1.3. The Request-URI is set to Schwaninger's Contact URI. Schwaninger answers with a 200 OK response:

```

200 OK 200 OK
To: SIP/24.8200@192.168.1.102;tag=12345678901234
From: Heisenberg <heisenberg@001.001.001.001>
Call-ID: 1234567890123456789012345678901234
CSeq: 2000 OK
Content-Length: 0
  
```

Not discussed in the previous example is the question of how the database accessed by the proxy contained Heisenberg's current IP address. There are many ways this could be done using SIP or other protocols. The mechanism for accomplishing this using SIP is called registration and is discussed in the next section.

2.3 SIP Registration Example

In this example, shown in Figure 2.3, Heisenberg sends a SIP REGISTER request to a type of SIP server known as a registrar server. The SIP registrar server receives the message uses the information in the request to update the database used by proxies to route SIP requests. Contained in the REGISTER message To header is the SIP URI address of Heisenberg. This is Heisenberg's “well-known” address, perhaps printed on his business card or published on a Web page or in a directory. Also contained in the REGISTER is a Contact URI which represents the current device (and its IP address) that the user




Figure 2.3 SIP registration example

Heisenberg is currently using. The registrar binds the SIP URI of Heisenberg and the IP address of the device in a database that can be used, for example, by the proxy server in Figure 2.2 to locate Heisenberg. When a proxy server with access to the database receives an HTTPREQ request addressed to Heisenberg's URI (i.e., an incoming call), the request will be proxied to the Contact's URI of the currently registered device.

This registration has no real counterpart in the telephone network, but it is very similar to the registration a wireless phone performs when it is turned on. A cell phone sends its identity to the base station (BS), which then forwards the location and phone number of the cell phone to a home location register (HLR). When the mobile switching center (MSC) receives an incoming call, it consults the HLR to get the current location of the cell phone. Further aspects of SIP mobility are discussed in Chapter 9.

The REGISTER message sent by Heisenberg to the SIP registrar server has the form:

```
REGISTER sip:registrar.example.de SIP/2.0
Via: SIP/2.0/UDP 192.168.1.100:5060;branch=1
Max-Forwards: 70
To: Heisenberg <sip:heisenberg@registrar.example.de>
From: Heisenberg <sip:heisenberg@registrar.example.de>
;tag=1431
Call-ID: 200200.192.168.100
CSeq: 1 REGISTER
Contact: <sip:heisenberg@registrar.example.de>
Content-Length: 0
```

The Request-URI in the start line of the message contains the address of the registrar server. In a REGISTER request, the To header field contains the URI that is being registered. In this case `sip:heisenberg@registrar.example.de`. This results in the To and From header fields usually being the same, although an example of third-party registration is given in Section 4.3.2. The SIP URI in the Contact address is stored by the registrar.

The registrar server acknowledges the successful registration by sending a 200 OK response to Heisenberg. The response echoes the Contact information that has just been stored in the database and includes a To tag:

```
SIP/2.0 200 OK
Via: SIP/2.0/UDP 192.168.1.100:5060;branch=1
To: Heisenberg <sip:heisenberg@registrar.example.de>;tag=1431
From: Heisenberg <sip:heisenberg@registrar.example.de>;tag=1431
Call-ID: 200200.192.168.100
CSeq: 1 REGISTER
Contact: <sip:heisenberg@registrar.example.de>;tag=1431
Content-Length: 0
```

The CONTACT URI is encoded along with an expires parameter, which indicates how long the registration is valid, which in this case is 1 hour (3600 seconds). If Hildeberg wants the registration to be valid beyond that interval, he must send another REGISTER request within the expires interval.

Registration is typically automatically performed on initialization of a SIP device and at regular intervals determined by the expires interval chosen by the register server. Registration is an additive process—more than one device can be registered against a SIP URI. If more than one device is registered, a proxy may forward the request to either or both devices, either in a sequential or parallel search. Additional register operations can be used to clear registrations or retrieve a list of currently registered devices.

2.4 SIP Presence and Instant Message Example

This example shows how SIP is used in a presence and instant messaging application. Presence information can be thought of as the state of a user or device at a particular instant. It can be as simple as whether a particular user is signed in or not, whether they are active on their machine, or idle or away. For a mobile device, presence information can include the actual location in terms of coordinates, or in general terms such as “in the office,” “on travel,” or “in the lab.” Presence information can even include information about the status or mood of the user, whether they are working, relaxing, or socializing. For all these examples, a presence protocol is mainly concerned about establishing subscriptions or long-term relationships between devices about transferring status information, and the delivery of that information. The actual information transferred, and how that information is presented to the user is application dependent. In terms of the SIP protocol, SUBSCRIBE is used to request status or presence updates from the presence server (or “presencing”), and NOTIFY is used to deliver that information to the requester or presence “watcher.” SIP presence uses the SIP Event extensions [4] and instant message extensions [5].

In this example, Chelychev wishes to communicate with Polson. The message flow is shown in Figure 2.4. To find out the status of Polson, Chelychev subscribes to Polson’s presence information by sending a SUBSCRIBE message to Polson. The request looks like:

```
REGISTER sip:192.168.1.100:5060 sip:192.168.1.100:5060
Via SIP/2.0/UDP 192.168.1.100:5060;branch=0;msg=1
From: Chelychev<chelychev@cs.cmu.edu>
To: Polson<polson@cs.cmu.edu>
Call-ID: 200211141040344000000000
CSeq: 1112 REGISTER
Allow: INVITE, OPTIONS
```




Figure 11: TCP connection establishment example.

```

client: seq, window, options, len, source, destination, window
server: seq, acknowledgment, window, transport, len
len: sequence
window: length.

```

In this example, TCP found a the window for the UDP message indicated in the `WIN` header field and is the `Connection-Exp` parameter in the `Connection-URL`. This request also contains `all:seq` and `all:seq-server` header fields which are used to indicate capabilities. In this example, Client is indicating support for window sizes indicated in the `all:seq` header field and the parameter subscriptions in the `all:seq-server` header field. In this example, the client is creating a dialog for an application that is

Note that this NOTIFY is sent within the dialog established with the SUBSCRIBE—it uses the same dialog identifier (Call-ID, local and remote tags)—and the request is sent to the Contact URI provided by Chobychew in the subscription request. The Expires: 1; header field indicates that the subscription has been authorized and active and that it will expire in 1 hour unless refreshed by Chobychew (using another SUBSCRIBE request).

The Common Event and Instant Message Presence Information Data Format (CPIM PEF) [8] XML message body contains the status information that Felicia is currently off-line (010000).

Chobychew sends a 200 OK response to the NOTIFY to confirm that it has been received:

```
HTTP/1.0 200 OK
Via: SIP/2.0/UDP 0100-00000001@csd.ty.ing.nyu.edu
Content-Disposition: event-refetch; id=1;
To: R. S. Chobychew <rs.chobychew@csd.ty.ing.nyu.edu>;
From: R. Feliciano <rfel@csd.ty.ing.nyu.edu>;
Call-ID: 0001114004004400001100
Msg: 200 NOTIFY
Content-Length: 0
```

Later, when Felicia does sign in, this information is provided in a second NOTIFY (omitting the change in status):

```
NOTIFY sip:rs@csd.ty.ing.nyu.edu; expires=3600
Via: SIP/2.0/UDP 0100-00000001@csd.ty.ing.nyu.edu
Content-Disposition:
Subscription-To:
To: R. S. Chobychew <rs.chobychew@csd.ty.ing.nyu.edu>;
From: R. Feliciano <rfel@csd.ty.ing.nyu.edu>;
Call-ID: 0001114004004400001100
Msg: 200 NOTIFY
Allow: ACK, BYE, CANCEL, INFO, NOTIFY, SUBSCRIBE, MESSAGE
Allow-Events: presence
Contact: sip:rs@csd.ty.ing.nyu.edu; expires=3600
Expires: 3600
Content-Type: application/identity+xml; charset=utf-8
Content-Length: 100
<xml version="1.0" xmlns:app="http://www.ietf.org/rfc/rfc4474.xml" >
  <app:presence xmlns="http://www.ietf.org/rfc/rfc4474.xml" >
    <app:status >01000000</app:status>
    </app:presence>
  </xml>
```

```
>/invite>
>/proceed>
```

The expiration time indicated in the Subscription-State header field indicates that 30 minutes have passed since the subscription was established. The CPBID-PIDF XML message body now indicates that Polina is no-line (upns) and can be reached via the URI sip:s.polina.gonzalez@att.proband13ty.org;transport=udp.

Chabychan receives a copy of the NOTIFY with a 200 OK response:

```
HTTP/1.1 200 OK
Via: SIP/2.0/UDP att-proband13ty.org;branch=1666666666;seq=31.1.1
To: P. L. Chabychan <sip:chabychan@att.com;tag=141>
From: E. Polina <sip:polina.gonzalez@att.com;tag=2140>
Call-ID: 1666666666-1666666666
CSeq: 202 NOTIFY
Content-Length: 0
```

Now that Chabychan knows that Polina is no-line, he sends an instant message to her using the Contact URI from the NOTIFY:

```
MESSAGE sip:s.polina.gonzalez@att.proband13ty.org SIP/2.0
Via: SIP/2.0/UDP att-proband13ty.com;branch=1666666666
To: E. Polina <sip:s.polina.gonzalez@att.proband13ty.com>
From: P. L. Chabychan <sip:chabychan@att.com;tag=141>
Call-ID: 1666666666-1666666666
CSeq: 16 MESSAGE
Allow: ACK, INFO, OPTIONS, SIP, NOTIFY, MESSAGE, MESSAGE
Content-Type: text/plain
Content-Length: 7

IM 166666
```

Notice that this MESSAGE is sent outside the dialog. Instant messages sent using the MESSAGE method in SIP are like page messages—they are not part of any dialog. As a result, each message contains a new Call-ID and From tag. The 200 OK response is used to acknowledge receipt of the instant message:

```
HTTP/1.1 200 OK
Via: SIP/2.0/UDP att-proband13ty.com;branch=1666666666;seq=31.1.1
To: E. Polina <sip:s.polina.gonzalez@att.proband13ty.com;tag=2140>
From: P. L. Chabychan <sip:chabychan@att.com;tag=141>
Call-ID: 1666666666-1666666666
CSeq: 16 MESSAGE
Content-Length: 0
```

Phone answers with a reply, which is also sent outside of any dialog, with a new Call-ID and FROM tag (an instant message response is never sent in a 200 OK reply to a MESSAGE request):

```
MESSAGE sip:chalsky@voicelabs.com; via SIP/1.0
Via SIP/2.0/UDP 66.209.64.221; branch=6620964221
Max-Forwards: 70
To: P. L. Chalsky <sip:chalsky@voicelabs.com>
From: M. Robinson <mailto:robinson@voicelabs.com>
Call-ID: 1604412429010
Msg: 114 MESSAGE
Allow: ACK, BYE, CANCEL, FORK, NOTIFY, PURSUE, REFER
Content-Type: text/plain
Content-Length: 0

well, hello there to you, too
```

which receives a 200 OK reply:

```
SIP/1.0 200 OK
Via SIP/2.0/UDP 66.209.64.221; branch=6620964221
From: P. L. Chalsky <mailto:chalsky@voicelabs.com>
To: M. Robinson <mailto:robinson@voicelabs.com>
Call-ID: 1604412429010
Msg: 114 MESSAGE
Content-Length: 0
```

Other protocol packages define what sets of information that can be requested by watchers from protocol servers.

2.5 Message Transport

As discussed in Chapter 1, SIP is a layer four, or application layer, protocol in the Internet Multimedia Protocol stack shown in Figure 1.1. RFC 3261 defines the use of TCP, UDP, or TLS transport. An extension document describes how SCTP can be used. How a SIP message is transported using these four protocols will be described in the following sections. The comparison of SIP for transport over low bandwidth connections, such as wireless, is discussed in Chapter 8.

2.5.1 UDP Transport

When using UDP, each SIP request or response message is carried in a single UDP datagram or packet. For a particularly large message body, there is a compact form of SIP that uses space by representing some header fields with a single

changes. This is discussed in Chapter 6. Figure 2.5 shows a SIP BYE request exchange during an established SIP session using UDP.

The source port is chosen from a pool of available port numbers (above 49152), or sometimes the default SIP port of 5060 is used. The lack of handshaking or acknowledgments in UDP transport means that a datagram could be lost and a SIP message along with it. The checksum, however, enables UDP to discard crossed datagrams, allowing SIP to assume that a received message is complete and error-free. The reliability mechanisms built into SIP to handle message retransmissions are described in Section 3.3. The reply is also sent to port 5060, as the port number listed in the top VLN header field.

UDP provides the simplest transport for user agents and servers and allows them to operate without transport layer state. However, UDP offers no congestion control. A series of loss packets on a heavily loaded IP link can cause serious problems, which in turn produce more loss packets and can push the link into congestion collapse. Also, UDP may only be used for SIP when the message (and its response) is known to be less than the Message Transport Unit (MTU) size of the IP network. For simple SIP messages, this is not a problem. However, for large messages containing multiple message bodies and large header fields, this can be a problem. In this case, TCP must be used, since SIP does not support fragmentation at the SIP layer.




Figure 2.5 Transmission of SIP messages using UDP.

2.5.2 TCP Transport

TCP provides a reliable transport layer, but at a cost of complexity and transmission delay over the network. The use of TCP for transport in a SIP message exchange is shown in Figure 2.6. This example shows an INVITE sent by a user agent at 100.101.102.103 to a type of SIP server called a redirect server at 200.201.202.203. A SIP redirect server does not forward INVITE requests like a proxy, but looks up the destination address and forwards the address in a redirection class (3xx) response. The 3xx Redirect Temporarily response is acknowledged by the user agent with an ACK message. Not shown in this figure is the next step, where the INVITE would be re-sent to the address returned by the redirect server. As in the UDP example, the well-known SIP port number of 5060 is chosen for the destination port, and the source port is chosen from an available pool of port numbers. Before the message can be sent, however, the TCP connection must be opened between the two end points. This transport layer diagram exchange is shown in Figure 2.6 as a single arrow, but it is actually a three-way handshake between the end points as shown in Figure 3.2. Once the connection is established, the messages are sent in the stream.

The Content-Length header field is critical when TCP is used to transport SIP, since it is used to find the end of one message and the start of the next. When TCP or another stream-based transport is used, Content-Length is a required header field in all requests and responses.

To send the 3xx Redirect Temporarily response, the server typically opens a new TCP connection in the reverse direction, using 5060 for the port listed in the top Via header field as the destination port.¹ The acknowledgment ACK is sent in the TCP stream used for the INVITE. Because this concludes the SIP session, the connection is then closed. If a TCP connection closes during a dialog, a new one can be opened to send a request within the dialog, such as a BYE request to terminate the call in session.

As previously mentioned, TCP provides reliable transport and also congestion control. It can also transport SIP messages of arbitrary size. The disadvantages of TCP include the setup delay in establishing the connection and the need for servers to maintain this connection state at the transport layer.

2.5.3 TLS Transport

New to RFC 3261 is support of TLS version 1.0 [8] in the SIP specification. SIP can use TLS over TCP as for encrypted transport with the additional capabilities

1. There is work underway to enable TCP connection reuse [7]. This would allow responses to be sent over the same connection as the request, and also allow multiple connections and dialogs to share a TCP connection, for example between two proxy servers.




Figure 2.8 Transmission of SIP messages using TLS.

of authentication. In Section 4.2.1 the secure SIPURI scheme (*sipga*) will be discussed, which uses TLS transport. The default SIP port number for TLS transport is port 5061.

If TLS is used between two parties, each party may have a certificate allowing mutual authentication. However, if a client does not have a certificate, TLS can be used in conjunction with another authentication mechanism, such as SIP digest, to allow mutual authentication.

The SIP use of TLS takes advantage of both the encryption and authentication services. However, the encryption and authentication is only useful on a single hop. If a SIP request takes multiple hops (i.e., includes one or more proxy servers), TLS is not useful for end-to-end authentication. S/MIME encryption, described in Section 3.7 solves this problem.

SIP proxies must support TLS and will likely use TLS for long-lived connections.

2.5.4 SCTP Transport

An extension to SIP defines the use of SCTP [9] with SIP to provide reliable stream-based transport with some advantages over TCP transport for a

message-based protocol such as SIP. First, it has built-in message segmentation, so that individual SIP messages are separated at the transport layer. With TCP, the SIP protocol must use the Content-Length calculation to delineate messages. If a TCP connection is being shared by a number of SIP transactions and dialogs, the “read at line blocking” problem discussed in Section 1.4.1.4 can cause the buffer to contain valid SIP messages that could be processed by the server while the accumulation takes place. Due to its message-level delimiters, SCTP is able to continue to forward messages to the application layer while simultaneously requesting a retransmission of a dropped message. Note that this is only a problem when multiple applications are multiplexed over a single TCP connection. An example of this is a TCP link between two signaling proxy servers. For a user agent or proxy TCP connection, this is usually not a problem unless the user has many simultaneous dialogs established.

SCTP also supports multihoming, so if one of a pair of load balancing SIP proxies fails, the other can immediately begin sending the messages without even requiring a DNS or other database lookup.

The SIP usage of SCTP is described in [10], which defines the syntax for the transport to use by URI parameter.

References

- [1] Rosenberg, J., et al., “SIP: Session Initiation Protocol,” RFC 3261, 2002.
- [2] Schulzrinne, H., et al., “RTP: A Transport Protocol for Real-Time Applications,” RFC 3235, 2002.
- [3] Handley, M., and V. Jacobson, “SIP: Session Description Protocol,” RFC 3263, 1998.
- [4] Beeth, A., “Session Initiation Protocol (SIP)-Specific Event Notifications,” RFC 3863, 2005.
- [5] Campbell, R., et al., “Session Initiation Protocol (SIP) Extension for Instant Messaging,” RFC 3428, 2002.
- [6] Sugan, H., et al., “Common Presence and Instant Messaging (CPIM) Presence Information Data Format,” IETF Instant-Data, Work in Progress, 2002.
- [7] Maly, R., “Requirements for Conversation State in the Session Initiation Protocol (SIP),” IETF Instant-Data, Work in Progress, 2002.
- [8] Davie, T., et al., “The TLS Protocol Version 1.0,” RFC 2246, 1999.
- [9] Stewart, R., et al., “Stream-Control Transmission Protocol,” RFC 2960, 2000.
- [10] Rosenberg, J., H. Schulzrinne, and G. Casavola, “The Stream Control Transmission Protocol as a Transport for the Session Initiation Protocol,” IETF Instant-Data, Work in Progress, 2002.

3

SIP Clients and Servers

The client-server nature of SIP has been introduced in the example message flows of Chapter 2. In this chapter, the types of clients and servers in a SIP network will be introduced and defined.

3.1 SIP User Agents

A SIP-enabled end-device is called a **SIP user agent** [1]. One purpose of SIP is to enable sessions to be established between user agents. As the name implies, a user agent takes direction or input from a user and acts as an agent on their behalf to set up and tear-down media sessions with other user agents. In most cases, the user will be a human, but the user could be another protocol, as in the case of a gateway (described in the next section). A user agent must be capable of establishing a media session with another user agent.

A UA must maintain state on calls that it initiates or participates in. A minimum call state set includes the local and remote tags, Call-ID, local and remote CSeq header fields, along with the reason set and any state information necessary for the media. This information is used to ensure the dialog information used for reliability. The remote CSeq storage is necessary to distinguish between a re-INVITE and a retransmission. A re-INVITE is used to change the session parameters of an existing or pending call. It uses the same Call-ID, but the CSeq is incremented because it is a new request. A retransmitted INVITE will contain the same Call-ID and CSeq as a previous INVITE. Even after a call has been terminated, call state must be maintained by a user agent for at least 32 seconds in case of lost messages in the call tear-down.

User agents already discard an ACK for an unknown dialog. Requests to an unknown URI receive a 404 Not Found Response. A user agent receiving a request for an unknown dialog responds with a 481 Dialog/Transaction Does Not Exist Response from an unknown dialog are also already discarded. These silent discards are necessary for security. Otherwise, a malicious user agent could gain information about other SIP user agents by spamming false requests or responses.

Although not required to understand every response code defined, a minimal implementation must be able to interpret any unknown response based on the class (first digit of the number) of the response. That is, if an undefined 499 Strong Phrase of the 500 response is received, it must be treated as a 499 Client Error.

A user agent responds to an unsupported request with a 501 Not Implemented response. A SIP UA must support UDP transport and also TCP if it sends messages greater than 1080 bytes in size.

A SIP user agent contains both a client application and a server application. The two parts are a user agent client (UAC) and user agent server (UAS). The UAC initiates requests while the UAS generates responses. During a user agent, a user agent will usually operate as both a UAC and a UAS.

A SIP user agent must also support SDP for media description. Other types of media description protocols can be used in bodies, but SDP support is mandatory. Details of SDP are in Section 7.1.

A UA must understand any extensions listed in a Require header field in a request. Unknown header fields may be ignored by a UA.

A UA should advertise its capabilities and features in any request it sends. This allows other UAs to learn of them without having to make an explicit capabilities query. For example, the methods that a UA supports should be listed in an Allow header field. SIP extensions should be listed in a Supported header field. Message body types that are supported should be listed in an Accept header field.

3.2 Presence Agents

A presence agent (PA) [2] is a SIP device that is capable of receiving subscription requests and generating state notifications as defined by the SIP Event specification [3]. An example of a presence agent is in the example of Section 2.4. A presence agent supports the presence Event package [4], responds to SUBSCRIBE requests, and sends NOTIFY requests.

A presence agent can collect presence information from a number of devices. Presence information can come from a SIP device registering, a SIP device publishing presence information [4], or from many other non-SIP sources.

A *proxy server* is a server that sometimes acts as a proxy agent and supplies proxy information and other times acts as a proxy, forwarding SUBSCRIBE requests to another proxy agent.

A *proxy agent* first authenticates a subscription request. If the authentication passes, it establishes a dialog and sends the notifications over that dialog. The subscription can be refreshed by resending new SUBSCRIBE requests.

3.3 Back-to-Back User Agents

A *back-to-back user agent (B2BUA)* is a type of SIP device that receives a SIP request, then reformulates the request and sends it out as a new request. Responses to the request are also reformulated and sent back in the opposite direction. For example, a B2BUA device can be used to implement an anonymous service in which one SIP UA can contact another without either party knowing the other party's URI or IP address, or any other information. To achieve this, an anonymous B2BUA would reformulate a request with an entirely new From, Via, Contact, Call-ID, and SDP media information, also removing any other SIP header fields that might contain information about the calling party. The responses received would also change the Contact and SDP media information from the called party. The modified SIP would pass to the RUA in the network, which would forward RTP media packets from the called party to the calling party and vice-versa. In this way, neither end point learns any identifying information about the other party during the session establishment. (Of course, the calling party needs to know the called party's URI in order for the call to take place.)

Sometimes RUAAs are employed to implement other SIP services. However, they break the end-to-end nature of an Internet protocol such as SIP. Also, a B2BUA is a call-manchef single point of failure in a network, which means that it can still reduce the reliability of SIP services over the Internet. The relayed media suffers from increased latency and increased probability of packet loss, which can reduce the quality of the media session. Geographic distribution of B2BUAs can reduce these effects, but the problem of selecting the best B2BUA for a particular session is a very difficult one since the source and destination IP address of the media are not known until the session is actually established (with a 200 OK).

The most common form of B2BUA present in SIP networks is application layer gateways (ALG). Some firewalls have ALG functionality built in, which allows a firewall to permit SIP and media traffic while still maintaining a high level of security. The use of ALGs is described in Section 3.11.

3.4 SIP Gateways

A *SIP gateway* is an application that interfaces a SIP network on a network utilizing another signaling protocol. In terms of the SIP protocol, a gateway is just a

special type of user agent, where the user agent acts on behalf of another protocol rather than a human. A gateway terminates the signaling path and can also terminate the media path, although this is not always the case. For example, a SIP to H.323 gateway terminates the SIP signaling path and converts the signaling to H.323, but the SIP user agent and H.323 terminal can exchange RTP media information directly with each other without going through the gateway. An example of this is described in Section 10.6.

A SIP to Public Switched Telephone Network (PSTN) gateway terminates both the signaling and media paths. SIP can be translated into, or interwork with, various PSTN protocols such as Integrated Services Digital Network (ISDN), ISDN User Part (ISUP), and other Circuit Associated Signaling (CAS) protocols, which are briefly described in Section 7.4. A PSTN gateway also converts the RTP media stream in the IP network into a standard telephone trunk or line. The conversion of signaling and media paths allows calling to and from the PSTN using SIP. Examples of these gateways are described in Sections 10.3 and 10.4. Figure 2.1 shows a SIP network connected via gateways with the PSTN and a H.323 network. There is work underway to standardize the SIP/H.323 interworking function [5].

In the figure, the SIP network, PSTN network, and H.323 networks are shown as clouds, which abstract the underlying details. Shown connecting to the SIP cloud are SIP IP telephones, SIP-enabled PCs, and corporate SIP gateways with standard telephones. The clouds are connected by gateways. Shown



Figure 2.1 SIP network with gateways.

attached to the H.323 network are H.323 terminals and H.323-enabled PCs. The PSTN cloud connects to ordinary analog black telephones (so-called because of the original color of their shells), digital PSTN telephones, and corporate private branch exchanges (PBXs). PBXs connect to the PSTN using shared trunks and provide line interfaces for either analog or digital telephones.

Gateways are sometimes decomposed into a media gateway (MG) and a media gateway controller (MGC). An MGC is sometimes called a call agent because it manages call control protocols (signaling), while the MG manages the media connection. This decomposition is transparent to SIP, and the protocols used to decompose a gateway are not described in this book.

Another difference between a user agent and a gateway is the number of users supported. While a user agent typically supports a single user (although perhaps with multiple lines), a gateway can support hundreds or thousands of users. A PSTN gateway could support a large corporate customer, or an entire geographic area. As a result, a gateway does not REGISTER every user it supports in the same way that a user agent might. Instead, a new-SIP protocol can be used to inform providers about gateways and make it possible. One protocol that has been proposed for this is the Telephony Routing over IP (TRIP) protocol [4], which allows an interdomain routing table of gateways to be developed. Another protocol called Telephony Gateway Registration Protocol (TGREP) [7] has also been developed to allow a gateway to register with a proxy server within a domain.

3.6 SIP Servers

SIP servers are applications that accept SIP requests and respond to them. A SIP server should not be confused with a user agent server or the client server nature of the protocol, which describe operations in terms of clients (originators of requests) and servers (originators of responses to requests). A SIP server is a different type of entity. The types of SIP servers discussed in this section are logical entities. Actual SIP server implementations may contain a number of server types, or may operate as a different type of server under different conditions. Because servers provide services and licenses to user agents, they must support both TCP, TLS, and UDP for transport. Figure 3.2 shows the interaction of user agents, servers, and a location service. Note that the protocol used between a server and the location service or database is not in general SIP and is not discussed in this book.

3.6.1 Proxy Servers

A SIP proxy server receives a SIP request from a user agent or another proxy and acts on behalf of the user agent in forwarding or responding to the request. A




Figure 3.2 SIP user agent, server, and location-service interaction.

proxy is not a B2BUA since it is only allowed to modify requests and responses according to strict rules set out in RFC 3261. These rules preserve the end-to-end transparency of the SIP signaling, while still allowing a proxy server to perform valuable services and functions for user agents.

A proxy server typically has access to a database or a location service to aid it in processing the request (determining the user agent). The interface between the proxy and the location service is not defined by the SIP protocol. A proxy can use any number of types of database to aid in processing a request. Databases could contain SIP registrations, presence information, or any other type of information about where a user is located. The example of Figure 3.2 introduced a proxy server as a facilitator of SIP message exchange providing user location services to the caller.

A proxy does not need to understand a SIP request in order to forward it—any unknown request type is assumed to use the non-INVITE transaction model. A proxy should not change the order of header fields or to generate, modify or delete header fields.

A proxy server is different from a user agent or gateway in three key ways:

1. A proxy server does not have requests: it only responds to requests from a user agent. (A CANCEL request is an exception to this rule.)
2. A proxy server has no media capabilities.
3. A proxy server does not parse message bodies: it relies exclusively on header fields.

Figure 2.2 shows a common network topology known as the IP Tunnel. In this topology, a pair of user agents in different domains establishes a session using a pair of proxy servers, one in each domain. The request refers to the object located by the originating user agent message. In this configuration, each user agent is configured with suitable authentication servers, which handle all requests. This proxy server typically will authenticate the user agent and may pull up a profile of the user and apply customized routing services. In an intranet environment, IPW HTTP queries will be used to locate a proxy server in the other domain. This proxy, sometimes called an Internet proxy, may apply Internet routing services on behalf of the other party. The proxy also handles the user agent registration information for the user, and can route the request to the called party. In general, direct IP requests will be sent directly between the two user agents, unless one or both parties have a Server-Service header field.

A proxy server can be either active or passive. A passive proxy server processes each IP request as requests listed solely on the message contents. Once the message has been parsed, processed, and forwarded or responded to, no information about the message is stored—no state information is stored. It handles proxy server requests, a message, and does not use any IP classes. Note that the standard loop detection using Via header fields described in RFC 2616 has been deprecated (removed) in RFC 2616b because of the use of content-length:0 as a means for header fields all requests.



Figure 2.2 IP Tunnel

A stateful proxy server keeps track of requests and responses received in the past and uses that information in processing future requests and responses. For example, a stateful proxy server starts a timer when a request is forwarded. If no response to the request is received within the timer period, the proxy will retransmit the request, relaying the user agent of this task. Also, a stateful proxy can require user agent authentication, as described in Section 3.8.

The most common type of SIP proxy is a transaction stateful proxy. A transaction stateful proxy keeps state about a transaction but only for the duration that the request is pending. For example, a transaction stateful proxy would keep state when it receives an `INVITE` request until it receives a `200 OK` or a final failure response (e.g., `404 Not Found`). After that, it would discard the state information. This allows a proxy to perform useful search services but minimize the amount of state storage required.

One such example of a search service is a proxy server that receives an `INVITE` request, then forwards it to a number of locations at the same time. This “forking” proxy server keeps track of each of the outstanding requests and the response to each, as shown in Figure 2-4. This is useful if the location service or database lookup system multiple possible locations for the called party that need to be tried.

In the example of Figure 2-4, the `INVITE` contains:

```
INVITE sip:marco@192.0.2.1:5060 SIP/2.0
Via: SIP/2.0/UDP 192.0.2.1:5060 ;branch=64666667900
Max-Forwards: 70
To: <sip:marco@192.0.2.1>
From: A. B. Buchanan <sip:buchanan@192.0.2.1>;tag=7612667
Call-ID: 846666666
Seq: 1 INVITE
Subject: Buchanan's Convention
Contact: <sip:marco@192.0.2.1>
Content-Type: application/sdp
Contact-Language: ...
```

(SIP not shown)

The `INVITE` is received by the standard S&F proxy server, which forks to two user agents. Each user agent begins dialing, sending two parallel responses back to Buchanan. They are:

```
SIP/2.0 180 Ringing
Via: SIP/2.0/UDP 192.0.2.1:5060;branch=64666667900
To: <sip:marco@192.0.2.1>;tag=7612667
From: A. B. Buchanan <sip:buchanan@192.0.2.1>;tag=7612667
Call-ID: 846666666
Seq: 1 INVITE
Contact: <sip:marco@192.0.2.1>
Content-Language: ...
```



```

To: <mailto:marcel@bluewin.ch>, <mailto:chris@cs.cmu.edu>
From: J. M. Rosenberg <mailto:marcel@bluewin.ch>, S. M. Jones <mailto:smj@cs.cmu.edu>
Call-ID: 1449931793
Seq: 1 20000
Contact: <mailto:marcel@bluewin.ch>, <mailto:smj@cs.cmu.edu>
Content-Type: application/sdp
content-length: ...

(more not shown)

```

The *forking proxy server* sends a **CANCEL** to the second UA to stop that phone ringing. If both UAs had answered, the forking proxy would have forwarded both 200 OK responses back to the caller who then would have had to choose which one, probably accepting one and sending a BYE to the other one.

A *successful proxy* usually sends a 200 Trying response when it receives an INVITE. A *successful proxy server* sends a 200 Trying response. A 200 Trying response received by a proxy is never forwarded—it is a single hop only response. A proxy handling a TCP request must be successful, since a user agent will assume reliable transport and rely on the proxy for retransmissions on any UDP hops in the signaling path.¹

The only limit to the number of proxies that can forward a message is controlled by the **Max-Forwards** header field, which is decremented by each proxy that touches the request. If the **Max-Forwards** count goes to zero, the proxy discards the message and sends a 484 Too Many Hops response back to the originator.

A SIP *session timer* [4] has been proposed to limit the time period over which a successful proxy must maintain state information. In the initial INVITE request, a **Session-Expires** header field indicates a timer interval after which successful proxies may discard state information about the session. User agents must tear down the call after the expiration of the timer. The caller can send re-INVITEs to refresh the timer, enabling a “keep alive” mechanism for SIP. This solves the problem of how long to carry state information in cases where a BYE request is lost or misdelivered, or in other security cases described in later sections. The details of this implementation are described in Section 6.1.2).

3.1.2 Reflect Servers

A *reflect server* was introduced in Figure 2.6 as a type of SIP server that responds to, but does not forward requests. Like a proxy server, a reflect server

1. TCP usually provides end-to-end reliability for applications. In SIP, however, TCP only provides single-hop reliability. End-to-end reliability is only achieved by a chain of TCP hops or TCP hops intermixed with UDP hops and successful proxies.

uses a database to location service to look up a user. The location information, however, is sent back to the caller in a redirection class response (CRCC), which, after the SCCP, concludes the transaction. Figure 3.5 shows a call flow that is very similar to the example of Figure 2.2, except the server uses redirection, instead of proxying to assist Selma/Singer locate Hansberg.

The SIP/RTSP from Figure 3.5 contains:

```

SIP/2.0 200 OK (application/sdp)
Via: SIP/2.0/UDP 194.108.201.100:5060; branch=1941082010
Max-Forwards: 70
To: Hansberg, sip:selma@singer.de
From: H. Selma/Singer <tel:+4930-800049601, sip:194-108-201-100@194-108-201-100>
Call-ID: 88088.100.100.100
CSeq: 1 SIP/2.0
Subject: Where are you exactly?
Contact: sip:selma@singer.de; uri=...
Content-Type: application/sdp
Content-Length: 150

v=0
o=8808888004 880888811 8808888100 IN 194.108.201.100
s=Phone Call
t=0
m=audio 194.108.201.100/100
a=sendrecv
a=rtpmap:0 PCMU/8000
  
```

The redirection response to the SIP/RTSP is sent by the redirect server:



Figure 3.5 Example with redirect server.

```

EPLV/O 200 BYE=200BYE
Flg: 81FV2,8,80CF 100,201,100,200, 0000 000000000000000000
To: 801000000 +010-800000, 8010000000000000_00+ 000-000000
From: H. Schenkling@ip-192.168.1.100 (tag=0111111)
Call-ID: 80100_100_201_100
Seq: 1 20V20
Contact: sip-werner@schenkling000.100.168.100
Contact-Length: 0

```

Schenkling acknowledges the response

```

ACK sip-werner@schenkling000.100.168.100
Flg: 81FV2,8,80CF 100,201,100,200, 0000 000000000000000000
Via-Forwards: To
To: 801000000 +010-800000, 8010000000000000_00+ 000-000000
From: H. Schenkling@ip-192.168.1.100 (tag=0111111)
Call-ID: 80100_100_201_100
Seq: 1 ACK
Contact-Length: 0

```

Notice that the ACK response carries the same branch ID as the INVITE and the 200 response. This is because an ACK to a non-200 final response is considered to be part of the same transaction as the INVITE. Only an ACK sent in response to a 200 OK is considered a separate transaction with a unique branch ID. Also, an ACK to a non-200 final response is a hop-by-hop response, not an end-to-end response as discussed in Section 2.6.

This exchange completes this call attempt, as a new INVITE is generated with a new Call-ID and sent directly to the location obtained from the Contact header field in the 200 response from the reflect server:

```

EPLV/O 200 BYE=200BYE
Flg: 81FV2,8,80CF 100,201,100,200, 0000 000000000000000000
Via-Forwards: To
To: 801000000 +010-800000, 8010000000000000_00+
From: H. Schenkling@ip-192.168.1.100 (tag=01111)
Call-ID: 80-01-00-00-00
Seq: 1 20V20
Contact: sip-werner@schenkling000.100.168.100
Contact-type: application/sdp
Contact-Length: 100

```

```

V=0
schenkling000 100004010 100004400 30 100 100-201-100-201
P-Expires: Call
call 0
P=CF 100 100-201-100-201
expires: 180S 000/000 0
s=0000000 0 00000000

```

The call then proceeds in the same way as Figure 2.2, with the message being identical. Note that in Figure 2.5, a 180 Ringing response is not seen

Instead, the 200 OK response is sent right away. Since 2xx informational responses are optional, this is a perfectly valid response by the UAS if Helsing responded to the ringing immediately and accepted the call. In the PSTN, this scenario is called *fast answer*.

3.3.3 Registration Servers

A SIP registration server was introduced in the example of Figure 2.3. A registration server, also known as a registrar, accepts SIP REGISTER requests and other requests under a 504 Not Implemented response. The contact information from the request is then made available to other SIP servers within the same administrative domain, such as proxies and redirect servers. In a registration request, the To header field contains the name of the resource being registered, and the Contact header fields contain the alternative addresses or aliases. The registration server creates a temporary binding between the Address of Record (ORIG URI) in the To and the device URI in the Contact Contact.

Registration servers usually require the registering user agent to be authenticated, using means described in Section 3.6, so that incoming calls cannot be hijacked by an unauthorized user. This could be accomplished by an unauthorized user registering someone else's SIP URI to point to their own phone. Incoming calls to that URI would then ring the wrong phone. Depending on the header fields present, a REGISTER request can be used by a user agent to review a list of current registrations, clear all registrations, or add a registration URI to the list. These types of requests are described in Section 4.1.2.

For full registration security, TLS must be used as HTTP Digest does not provide the needed integrity protection.

3.6 Acknowledgment of Messages

Most SIP requests are end-to-end messages between user agents. That is, proxies between the two user agents simply forward the messages they receive and rely on the user agents to generate acknowledgments or responses.

There are some exceptions to this general rule. The CANCEL method (used to terminate pending calls or watches and discussed in detail in Section 4.1.5) is a hop-by-hop request. A proxy receiving a CANCEL immediately sends a 200 OK response back to the sender and generates a new CANCEL, which is then forwarded to the next hop. (The order of sending the 200 OK and forwarding the CANCEL is not important.) This is shown in Figure 4-4.

Other exceptions to this rule include 3xx, 4xx, 5xx, and 6xx responses to an INVITE request. While an ACK to a 2xx response is generated by the end point, a 3xx, 4xx, 5xx, or 6xx response is acknowledged as a

hop-by-hop basis. A proxy server receiving one of these responses immediately generates an ACK back to the sender and forwards the response to the next hop. This type of hop-by-hop acknowledgment is shown in Figure 4.2.

ACK messages are only sent to acknowledge responses to INVITE requests. For responses to all other request types, there is no acknowledgment. A lost response is detected by the UAC when the request is retransmitted.

3.7 Reliability

SIP has reliability mechanisms defined, which allow the use of unreliable transport layer protocols such as UDP. When SIP uses TCP or TLS, these mechanisms are not used, since it is assumed that TCP will retransmit the message if it is lost and inform the client if the server is unreachable.

For SIP transport using UDP, there is always the possibility of messages being lost or even received out of sequence, because UDP guarantees only that the datagram is error-free. A UAC validates and parses a SIP request to make sure that the UAC has not violated by creating a request missing required header fields or other syntax violations. Reliability mechanisms in SIP include:

- Retransmission timers
- Increasing command sequence/Call-ID numbers
- Positive acknowledgments

SIP timer T1 is started by a UAC or a stateful proxy server when a new request is generated or sent. If no response to the request (as identified by a request containing the identical local tag, remote tag, Call-ID, and Call-IDeq) is received when T1 expires, the request is re-sent. If a provisional (information class 1xx) response is received, the UAC or stateful proxy server ignores T1 and starts a new longer timer T2. No retransmissions are sent until T2 expires.

After a request is retransmitted, the timer period is doubled until T1 is reached. After that, the retransmissions occur at T2 intervals. This capped exponential backoff process is continued until a maximum of 10 retransmissions at increasing intervals are sent. A stateful proxy server that receives a retransmission of a request discards the retransmission, and continues its retransmission schedule based on its own timer. Typically, it will resend the last provisional response.

For an INVITE request, the retransmission scheme is slightly different. After a provisional (1xx) response is received, the INVITE is never retransmitted. However, a proxy may discard transaction state after 3 minutes.

A successful proxy must store a forwarded request or generated response message for 32 seconds. An example message flow involving two user agents, a successful proxy, two lost messages (labeled by "X" in Figure 3-6), and three retransmissions is shown in Figure 3-6. In this example, the OPTIONS sent by the successful proxy arrives at the UAS is lost. As a result, it is retransmitted when the proxy's T1 timer expires and no response is received. The 200 OK forwarded by the proxy is also lost. When timer T1 in the UAC expires without a final response from the proxy, the OPTIONS is retransmitted. When the proxy receives the retransmitted OPTIONS, it deduces that the 200 OK was lost and resends it. The proxy recognizes the 200 OK as a retransmission and does not forward it.

Suggested default values for T1 and T2 are 500 ms and 4 seconds, respectively. Timer T1 is suggested to be an estimate of the round-trip time (RTT) in the network. Larger values are allowed but not desirable, because this will generate more message retransmissions. See Table 4 in RFC 3261 [1] for a summary of SIP timers.

Note that gaps in CSeq number do not always indicate a lost message. In the authentication examples in the next section, not every request (and hence CSeq) generated by the UAC will reach the UAS if authentication challenges occur by proxy in the path.

3.8 Authentication

Authentication in SIP takes two general forms. One is the authentication of a user agent by a proxy, relay, or registration server. The other is the



Figure 3-6 SIP reliability example.

authentication of a user agent by another user agent. Mutual authentication between proxies or a proxy and a user agent is also possible using certificates.

A proxy or redirect server might require authentication to allow a user agent to access a service or feature. For example, a proxy server may require authentication before forwarding an INVITE to a gateway or invoking a service. A registration server may require authentication to prevent flooding (as described previously). User agents can authenticate each other to verify who they are communicating with, since From header fields are easily forged.²

SIP supports both a simple lightweight authentication scheme and a very robust scheme. The simple scheme is based on HTTP Digest [9] and uses a simple challenge/response mechanism and a shared secret between the two servers. The other schemes involve cryptographic means to exchange and verify certificates.

Using HTTP Digest authentication, a proxy requiring authentication replies to an unauthenticated INVITE with a 401 Proxy Authentication Required response (containing a Proxy-Authenticate header field with the form of the challenge). After sending an ACK for the 401, the user agent can then resend the INVITE with a Proxy-Authentication header field containing the credentials. This request is usually sent using the same Call-ID but an incremented CSeq count. User agents, proxies, or redirect servers typically use 401 Unauthenticated response to challenge authentication containing a WWW-Authenticate header field, and expect the user agent to resubmit an INVITE to contain an Authorization header field containing the user agent's credentials. A user agent's credentials are usually an escaped username and password, as in the example of Section 10.3. The Authentication-Info-URI header field also allows a user agent to authenticate a proxy server using HTTP Digest. However, HTTP Digest does not provide integrity protection—for this, TLS or S/MIME must be used.

A call flow involving both proxy and user agent authentication is shown in Figure A.7.

HTTPS is used as transport, as discussed in Section 2.5.3, the two servers can exchange and verify certificates for authentication.

2. The ability to forge From header fields is present in SMTP, where it is usually a *feature*. A preference setting in an e-mail program sets your name and e-mail address, which need not correspond to the address or domain that is used to send the message. This allows a user to send multiple e-mail addresses from the same e-mail account by simply changing the From address before sending a message. Only a detailed examination of a full set of SMTP header fields will determine the actual user name and address.




Figure 3.7: SIP authentication call flow.

3.9 SIPSEC Encryption

While authentication is used as a means of access control and identity confirmation, encryption is used for privacy. SIP messages intercepted during media setup reveal considerable information, including:

- Both parties' SIP URIs and IP addresses
- The fact that the two parties have established a call
- The IP addresses and port numbers associated with the media, allowing eavesdropping.

Presence information can reveal even more private information, such as:

- The user's geographic location
- The user's current activity level

- The “*mess*” or other personal information.

As such, strong encryption and privacy mechanisms have been built into SIP that work in both single-domain networks and across the public Internet. The use of intermediary devices such as proxies and B2BUAs also make SIP security very important.

SIP supports the encryption of both message bodies and message header fields. The encryption of message bodies makes it more difficult for an eavesdropper to listen in, if the message body contains SIP information, for example. Also, an untrusted third party, knowing all the SIP information could guess the RTP SSRC number and steal unsecured media to either party, so-called media stripping. Message bodies in presence and instant message messages also contain private information that is only needed by the user agents, not intermediary devices in the network.

Encryption can be done hop-by-hop or end-to-end. TLS can be used for the hop-by-hop but only S/MIME [10] allows end-to-end encryption. S/MIME allows UAs to discover if a third party (a proxy or B2BUA) is modifying SIP messages or bodies between them. Note that allowed proxy header field modification (such as deletion of TLS header fields in response routing) is allowed and not considered a security violation.

3.10 Multicast Support

SIP support for UDP multicast has been mentioned in previous sections. There are two main uses for multicast in SIP.

SIP registration can be done using multicast, by sending the REGISTER message to the well-known “All SIP Servers” URI `sip:sip.nonat.net` at IP address 224.0.0.175.

The second use for multicast is to send a multicast session invitation. This effectively allows a conference call to be established with a single request. An INVITE with a partially defined Request-URI can be sent using multicast. For example, a multicast INVITE could be sent to `sip:*@acfl.com`, which would invite all AECI employees with SIP phones receiving the request to respond. Responses to a multicast request are also sent by multicast. To limit congestion, only a limited set of responses is allowed by the standard.

A proxy can forward a unicast INVITE request to a multicast address. The use of multicast is recorded in a SIP message using the `multicast` parameter in the Via header field, as discussed in Section 6.1.4.

However, due to the limited implementations of multicast, these features are rarely used.

2.11 Firewalls and NAT Interaction

Most corporate LANs or networks connect to the public Internet through a firewall. A firewall is filtering software usually in a router or hub that is used to protect the LAN behind it from various kinds of attacks and unwanted worms. Firewalls are also increasingly being used in home networks routers and wireless hubs and in PCs themselves. Sometimes they are used to prevent users behind the firewall accessing certain resources in the Internet. In the simplest deployment, a firewall can be thought of as a one-way gate: It allows outgoing packets from the Internet to the Internet, but blocks incoming packets from the Internet unless they are responses to queries. Only certain types of requests from the Internet will be allowed to pass through the firewall, such as HTTP requests to the corporate Web server, SMTP e-mail messages, or DNS queries to the authoritative DNS for the corporate domain. The firewall does this by keeping track of TCP connections opened and filtering ports.

Firewalls pose a particularly difficult challenge to SIP systems. Because SIP can use TCP and a well-known port, configuring a firewall to pass SIP is not too difficult. This does not help the media path, however, which uses RTP over UDP on various ports and will be blocked by most firewalls. A firewall or a proxy that controls the firewall needs to understand SIP, be able to pass an INVITE request and 200 OK response, extract the IP addresses and port numbers from the SDP, and open up “pin holes” in the firewall to allow this traffic to pass. The hole can then be closed when a BYE is sent or a session timer expires. An alternative is an ALG—a SIP ALG that is trusted by the firewall. The firewall then allows SIP and RTP traffic, which terminates on the ALG and blocks all other traffic. The authentication and security policies of allowing or denying SIP systems are then controlled by the SIP ALG instead of in the firewall itself.

Network address translation (NAT) also causes serious problems for SIP. A NAT can be used to conserve IPv4 addresses, or can be used to hide the IP address and LAN structure behind the NAT. It is used on a router or firewall that provides the only connection of a LAN to the Internet, a so-called stub network. A NAT allows nonunique IP addresses to be used internally within the LAN. When a packet is sent from the LAN to the Internet, the NAT changes the nonglobally unique address (usually addresses in the range 10.x.x.x, 172.16.x.x, or 192.168.x.x) to a globally unique address from a pool of available addresses. Addresses can also be statically assigned. This means that every node on the network does not have to have a globally unique IP address. Responses from the Internet are translated back to the nonunique address. A NAT, however, is not completely transparent to higher layers. For a signaling protocol such as SIP, a NAT can cause particular problems.

Because responses in SIP are routed using VLN header fields, a device behind a NAT will stamp its non-routable private IP address in its VLN header field of messages that it originates. When the request is forwarded outside the network by the NAT, the UDP and IP packet headers will be rewritten with a temporarily assigned global Internet address. The NAT will keep track of the binding between the local address and the global address so that incoming packets can have the UDP and IP headers rewritten and routed correctly. However, IP addresses in a SIP message, such as VLN and Contact header fields, or IP addresses in SIP message bodies will not be rewritten and will not be routable.

To partially solve the message routing problem, SIP has a mechanism for detecting if a NAT is present in a SIP message path. Each proxy or user agent that receives a request checks the received IP address with the address in the VLN header field. If the addresses are different, there is a NAT between them. The non-routable VLN header field is fixed with a received-with tag containing the actual global IP address. Outside the NAT, the request is routed using the received IP address. Inside the NAT, the VLN address is used. This does solve the message response routing problem (except when the port number is also wrong), but not the media problem.

Another problem with NATs is the time span of the NAT address binding. For a TCP connection, this is not an issue—the binding is maintained as long as the connection is open. For a UDP SIP session, the time period is determined by the application. If a binding were removed before a RTT was sent or a session timer had expired, the connection would effectively be closed and further signaling impossible. A keep-alive mechanism may be needed to refresh this binding.

A SIP ALG consistent with the NAT solves many of these problems. The ALG would rewrite the media IP addresses in the SIP messages and would not allow the NAT to remove the address binding until a RTT was sent or a session timer had expired. However, NATs are often deployed deep inside a service provider's network that is not associated with providing SIP service, and hence has no incentive or upgrading the NATs to allow this service to work.

Even without ALGs or upgrades to NATs, it is possible to use SIP to establish a media session through many types of NATs. The protocols described in the next section allow a SIP client to discover the presence and type of NATs between it and the public Internet, learn its public IP address, and possibly fix the incorrect addresses in the SIP and SDP messages.

3.12 Protocols and Extensions for NAT Traversal

A detailed analysis of various scenarios and solutions to NAT traversal has been done [11]. Some of the results are summarized in this section along with two

non-SIP protocols: Simple Traversal of UDP through NATs (STUN) and Traversal Using Relay NAT (TURN), which aid in NAT traversal. Finally, some extensions to SIP and SIP to enable NAT traversal are discussed.

1.2.2 STUN Protocol

The STUN protocol [12] allows a client to discover the presence and type of NATs between the client and the public Internet. In addition, a client can discover the mapping between the private IP address and port number and the public IP address and port number. Typically, a service provider will sponsor a STUN server in the public Internet, with STUN clients being embedded in end-devices, which are possibly behind a NAT.

A STUN server can be located using DNS SRV records using the service provider's domain as the lookup. STUN typically uses the well-known port number 3498. STUN is a binary-encoded protocol with a 20-byte header field and possibly additional attributes.

Since a STUN client uses the protocol to learn public IP addresses, some security is necessary or an impostor STUN server could provide incorrect public IP addresses and block or intercept communications destined for the client. As a result, the first step with STUN is for the client to contact the server and negotiate a shared secret (usually a username and password) over a secure link, which will be used to format STUN requests. The initial STUN connection is opened using TLS over TCP. The client verifies the certificate of the server to make sure that it is connected to the proper server. The client then sends a Shared Secret Request packet to the server and receives the shared secret to be used.

The client is then ready to use STUN to determine the presence and type of NATs between the client and the server. Using the shared secret, the client sends one or more Binding Request packets using UDP to the server. These packets must be sent from the same IP address that the client will use for the other protocol, since this is the address translation information that the client is trying to discover. The server returns Binding Response packets, which tell the client the public IP address and port number from which it received the Binding Request. Since the client knows the private IP address and port number it used to send the Binding Request, it knows the mapping between the private and public address space being performed by the NAT. Of course, if the Binding Response packets indicate the same address and port number as the request, this tells the client that no NATs are present.

Figure 1.8 shows a SIP client behind a NAT obtaining its public IP address using STUN and setting up a session as a result.

The type of NAT can be determined by sending multiple Binding Request packets with different attributes. As a result, the symmetric type can be determined




Figure 8.1 IPsec operations for NAT traversal and flow.

as for one of those listed in Table 8.1. ITUN can also be used to refresh the NAT address binding to keep it valid during the application lifespan.

For all the cases described in Table 8.1 except the symmetric NAT, ITUN provides the client with the public IP address needed to establish a media session using SIP. The IP addresses obtained using NAT are used in the SIP:Content-Disposition and in the SIP media information in the SDP, for example, embedded in a successful media session with sessions in the public Internet.

However, to traverse a symmetric NAT on the certain topologies such as communication between two user agents behind NATs, signaling and media relay procedures such as TURN is required, which is described in the next section.

Table 2.1
Types of NATs

Type	Translation
Internet	Non-unique, one-to-one public IP address
Full cone NAT	Constant mapping between private IP address and/or public IP address
Restricted cone NAT	Constant mapping, but an outgoing packet is needed to open a connection path
Symmetric NAT	Different public IP address mapping is used based on destination IP address

2.10.2 TURN Protocol

The TURN protocol [13] allows a client to obtain a transport IP address and port that it can receive packets sent from a single IP address in the public Internet. For some NAT topologies such as a client behind a symmetric NAT, using a relay located in the public Internet is the only approach that allows communication to take place.

Similar to STUN, a TURN client can use DNS SRV records for the domain of the service provider. TURN uses an identical syntax to STUN and reuses the Shared Secret Response and Shared Secret Response packets to establish the shared secret. The client can also use Binding Request and Binding Response packets to detect and categorize the NAT in the path. A client uses an Allocate Request to request a relay IP address and port number be returned in an Allocate Response packet.

Requests for transport addresses to be used for UDP must be sent using UDP while requests for transport addresses for TCP transport must be sent using TCP.

TURN, due to its triangular routing, will result in increased packet latency and increased probability of packet loss. TURN should only be used when it is the only approach available (i.e., the IP addresses obtained using STUN will not work). Note that TURN can be used instead of a B2BIU to provide an asymmetric service.

Figure 2-9 shows a SIP client behind a NAT using STUN and TURN to obtain a transport IP address and setting up a session as a result. Note that the resulting SIP and RTP messages are received through the TURN server.

In some cases, the use of STUN and TURN may result in the use of a relay when, in fact, the two clients can communicate successfully without one. This case is where both clients are behind the same NAT.




Figure 18: IP and NAT/NAT64 interaction diagram.

3.6.2 Other NAT/NAT64-Related Extensions

Other IP and NAT extensions have been developed to address some of the problems associated with NATs. They are described in the following sections and can be used in conjunction with protocols such as STUN and TURN to enable NAT traversal. Each extension either is a *plug-in* for the complete NAT traversal problem. Instead, they represent an optimization or improvement used along with STUN and TURN.

As mentioned previously, using TCP transport for IP enables NAT traversal easier. In particular, it is easier to keep the private/public address binding for a TCP connection than it is with UDP. This is a particular problem for

registrations. If UDP is used for registration behind a NAT, either STUN packets or special registrations are needed (as often, so every minute). It is better to register using TCP, but only if the TCP connection can be kept open and incoming requests are routed over that open connection. A SIP extension has been proposed to solve this problem [14].

An example Via header field in a REGISTER request contains a parameter `transport`, which indicates that the UAC supports this extension:

```
Via: SIP/2.0/SIP/udp;branch=0;branch=0;branch=0;branch=0;branch=0
```

The next hop server that supports the extension would record both the received IP address and port number in the Via header field for example:

```
Via: SIP/2.0/SIP/udp;branch=0;branch=0;branch=0;branch=0;branch=0;
  rport=9600;to_c:1.2.3.4;to_p:9600;to_t:udp
```

This TCP connection would then be kept open and responses to this user would be sent back to the client at address 192.0.2.2 port 21121. An incoming INVITE or other request would be routed over this open TCP connection.

The ability to specify in SIP a “symmetric” RTP session [15] has been proposed as a solution for some NAT traversal problems. Since a client behind a NAT can usually successfully send RTP packets to another client in the public Internet, in a symmetric mode, RTP sent in the other direction could be sent to the address and port that RTP was received from. For example, consider an SIP offer to an INVITE sent by the UAC, which supports this symmetric-oriented media extension:

```
V=0
m=audio 49154 RTP/AVP 0 99
c=IN
rtpmap 0 PCMU/8000
rtpmap 99 PCMA/8000
a=symmetry:active IN IP4
```

If the UAC supports this extension, it will wait for RTP packets to be received from the client behind the NAT before sending. The answer SIP will be:

```
V=0
m=audio 192.0.2.2 RTP/AVP 0 99
c=IN
rtpmap 0 PCMU/8000
rtpmap 99 PCMA/8000
a=symmetry:active IN IP4 client-public.org
```

```

>>> 100 200 OK (application/javascript)
  
```

The direction address in the answer tells the UAC that the UAS supports the symmetric RTP extension. The UAS will then send its RTP packets to the IP address and port number received from the UAC which overrides the unreachable 10.1.2.23 address and 48173 port number in the SDP.

Finally, it is assumed that RTP Control Protocol (RTCP) packets (see Section 7.2) are received on some port higher than the RTP port number. If RTUN or TURN is used to obtain port numbers, this consecutive numbering may not be possible if the NAT is mapping port numbers as well. Also, RTCP may need to be sent to an entirely different IP address. A solution to this is an explicit RTCP port number and possibly IP address in SDP extension [36]. This allows RTCP to still be received in these cases. For example, a UAC behind a NAT could include the attribute:

```

>>> 100 200 OK (application/javascript)
  
```

A UAS supporting the extension would send RTCP packets to this IP address and port number instead of using the RTP IP address (connection fixed) and port number plus one. If just the port number needs to be changed, the address information can be omitted.

```

>>> 100 200 OK
  
```

Note that without this SDP extension, the SIP and RTP sessions will still work, but no RTCP would be exchanged.

References

- [1] Rosenberg, J., et al., "SIP: Session Initiation Protocol," RFC 3261, 2002.
- [2] Rosenberg, J., "A Session Event Package for the Session Initiation Protocol (SIP)," IETF Internet-Draft, Work in Progress, January 2003.
- [3] Bush, A., "Session Initiation Protocol (SIP)-Specific Event Notifications," RFC 3262, 2002.
- [4] Campbell, B., et al., "SDP/LE: Presence Publication Mechanism," IETF Internet-Draft, Work in Progress, February 2003.
- [5] Madsen, H., and C. Agfloh, "Session Initiation Protocol (SIP)-H.323 Interworking Requirements," IETF Internet-Draft, Work in Progress, February 2003.
- [6] Rosenberg, J., H. Salama, and M. Spivey, "Telephony Routing over IP (TRIP)," RFC 3269, 2002.
- [7] Bergman, M., et al., "A Telephony Gateway Registration Protocol (TGRP)," IETF Internet-Draft, February 2003.

-
- [8] Deacon, S., and J. Rosenberg, "The SIP Session Timer," IETF Internet-Draft, Work in Progress.
 - [9] Frank, J., et al., "HTTP Authentication: Basic and Digest Access Authentication," RFC 2617, 1999.
 - [10] Randall, R., "SIP-URI Version 3 Usage Specification," RFC 2605, 1999.
 - [11] Rosenberg, J., K. Maki, and S. Ito, "NAT and Firewall Traversal and Solutions for SIP," IETF Internet-Draft, Work in Progress, March 2004.
 - [12] Rosenberg, J., et al., "SIP-URI—Simple Universal of User Datagram Protocol (UDP) Through Network Address Translation (NAT)," RFC 2608, 1999.
 - [13] Rosenberg, J., et al., "Traversal Using Relay NAT (TURN)," IETF Internet-Draft, Work in Progress, March 2005.
 - [14] Yin, D., "Connection-Oriented Media Transport in SIP," IETF Internet-Draft, Work in Progress, May 2004.
 - [15] Rosenberg, J., J. Weinberger, and H. Schulzrinne, "The Extension to the Session Initiation Protocol (SIP) for Symmetric Response Routing," RFC 2983, 2001.
 - [16] Hatakeyama, C., "RTCP Profiles for SIP," IETF Internet-Draft, Work in Progress, September 2002.

4

SIP Request Messages

This chapter covers the types of SIP requests called methods. Six are described in the SIP specification document RFC 3261 [1]. Seven more methods are described in separate RFC documents. After discussing the methods, this chapter concludes with a discussion of SIP URIs and URIs, tags, and message bodies.

4.1 Methods

SIP requests or methods are considered “verbs” in the protocol, since they request a specific action to be taken by another user agent or server. The `INVITE`, `REGISTER`, `BYE`, `ACK`, `CANCEL`, and `OPTIONS` methods are the original six methods in SIP. The `REFER`, `MESSAGE`, `NOTIFY`, `MESSAGE`, `OPTIONS`, `INFO`, and `PRACK` methods are described in separate RFCs.

Note that a proxy does not need to understand a request method in order to forward the request. A proxy treats an unknown method as if it were an `OPTIONS`; that is, it forwards the request to the destination if it can. This allows new features and methods useful for user agents to be introduced without requiring support from proxies that may be in the middle. A user agent receiving a method it does not support replies with a 501 Not Implemented response. Method names are case sensitive and conventionally use all uppercase for visual clarity to distinguish them from header fields, which can both upper- and lower-case.

4.1.3 INVITE

The INVITE method is used to establish media sessions between user agents. In telephony, it is similar to a *startup message* in ISDN or an *initial address message* (IAM) in ISUP. (ISDN protocols are briefly introduced in Section 7.4.) Responses to INVITEs are always acknowledged with the ACK method described in Section 4.1.4. Examples of the use of the INVITE method are described in Chapter 7.

An INVITE usually has a message body containing the media information of the call. The message body can also contain other session information such as quality-of-service (QoS) or security information. If an INVITE does not contain media information, the ACK contains the media information of the UAC. An example of this call flow is shown in Figure 4.1. If the media information contained in the ACK is not acceptable, then the called party must send a BYE to cancel the session—a CANCEL cannot be sent because the session is already established. A media session is considered established when the INVITE, 200 OK, and ACK messages have been exchanged between the UAC and the UAS. A successful INVITE request establishes a dialog between the two user agents, which terminates until a BYE is sent by either party to end the session, as described in Section 4.1.3.

A UAC that originates an INVITE to establish a dialog creates a globally unique Call-ID that is used for the duration of the call. A Call-ID counter is initialized (which need not be set to 1, but must be an integer) and incremented for each new request for the same Call-ID. The To- and From-headers are populated with the source and local addresses. A From-tag is included in the



Figure 4.1 INVITE with an SDP message body.

INVITE, and the UAS includes a To tag in any response, as described in Section 4.3. A To tag in a 200 OK response to an INVITE is used in the To header field of the ACK and all future requests within the dialog. The combination of the To tag, From tag, and Call-ID is the unique identifier for the dialog.

An INVITE sent for an existing dialog references the same Call-ID as the original INVITE and contains the same To and From tags. Sometimes called a re-INVITE, the request is used to change the session characteristics or refresh the state of the dialog. The CSeq sequential sequence number is incremented so that a UAS can distinguish the re-INVITE from a retransmission of the original INVITE.

If a re-INVITE is refused or fails to map way, the session continues as if the INVITE had never been sent. A re-INVITE may also be sent by a UAC until a final response to the initial INVITE has been received—instead, an UPDATE request can be sent, as described in Section 6.1.13. There is an additional case where two user agents simultaneously send re-INVITEs to each other. This is handled in the same way with a forked-ACKbar header. This condition is called *glare* in telephony, and occurs when both ends of a trunk group seize the same trunk at the same time.

An Expires header in an INVITE indicates to the UAS how long the call request is valid. For example, the UAS could have an unanswered INVITE request displayed as a screen for the duration of specified in the Expires header. Once a session is established, the Expires header has no meaning—the expiration of the timer does not terminate the media session. Instead, a `Duration-Expires` header (Section 6.2.26) can be used to place a time limit on an established session.

An example INVITE request with a SDP message body is shown below:

```

INVITE sip:alice@lab.org;user=alice SIP/2.0
Via: SIP/2.0/UDP lab.org;branch=0;seq=1;method=INVITE
Max-Forwards: 70
To: Alice-FLORIAN@lab.org;user=alice
From: Charles@lab.org;user=Charles;lab.org
Content-Type: application/sdp
Call-ID: 1234567890@lab.org;seq=1
CSeq: 1 INVITE
Subject: Test Session
Contact: sip:charles@lab.org;user=charles
Content-Disposition: hold
Content-Length: 134

v=0
o=charles 280447204 280447204 IN IP4 lab.org
s=Phone Call
t=0 0
m=audio 44100 1/0

```

```

>>> SIP/2.0 200 OK
>>> Content-Type: text/plain

```

In addition to the required headers, this response contains the optional `Stateless` header. Note that this `Request-URI` contains a phone number. Phone numbers appear in SIP URIs as described in Section 4.2.

The mandatory and header field in an INVITE request are shown in Table 4.1.

4.2 REGISTER

The REGISTER method is used by a user agent to notify a SIP network of its current `Contact` URI (IP address) and the URI that should have requests sent to this `Contact`. As mentioned in Section 2.3, SIP registration bears some similarity to cell phone registration on initialization. Registration is not required to enable a user agent to use a proxy server for outgoing calls. It is necessary, however, for a user agent to register so receive incoming calls from parties that serve that domain unless some non-SIP mechanism is used by the location service to populate the SIP URIs and `Contacts` of end-points. A REGISTER request may contain a message body, although its use is not defined in the standard. Depending on the use of the `Contact` and `Expires` headers in the REGISTER request, the registrar server will take different action. Examples of this use shown in Table 4.2. If no `expires` parameter or `Expires` header is present, a SIP URI will expire in 1 hour. The presence of an `Expires` header sets the expiration of `Contacts` with no `expires` parameter. If an `expires` parameter is present, it sets the expiration time for that `Contact` only. Non-SIP URIs have no default expiration time.

The `CSeq` is incremented for a REGISTER request. The use of the `Request-URI To`, `From`, and `Call-ID` headers in a REGISTER request is slightly different than for other requests. The `Request-URI` contains only the

Table 4.1
Mandatory Headers from INVITE Request

Call-ID
CSeq
From
To
Via
Contact
Max-Forwards

Table 6.2
Types of REGISTER and Contact Headers

Request Headers	Register Action
Contact: * Expires: 0	Cancel all registrations
Contact: sip:private@domain.edu.it; expires=30	Add Contact to current registra- tion; registration expires in 30 minutes
Contact: sip:private@domain.edu.it Expires: 30	Add Contact to current registra- tion; registration expires in 30 minutes
Contact: sip:private@domain.edu.it; expires=0	Add all Contacts to registrations in preference order (top), first one expires in 0 minutes, suc- ceed in 30 minutes
Contact: sip:private@domain.edu.it Expires: 30	
Contact: sip:private@domain.edu.it; registration=ipof	Add Contacts to current registra- tion using specified preference IP requests
Contact: multi:private@domain.edu.it; ipof:1	Should be rejected; SIP URI expires (old values default) multi URI does not expire
No Contact header present	Remove all current registrations in response

domains of the register server with no user portion. The REGISTER request may be forwarded or denied until it reaches the authoritative register server for the specified domain. The To header contains the SIP URI of the AOR of the user agent that is being registered. The From contains the SIP URI of the sender of the request, usually the same as the To header. It is recommended that the same Call-ID be used for all registrations by a user agent.

A user agent sending a REGISTER request may receive a 302 redirection or 303 failure response containing a Contact header of the location to which registrations should be sent.

A third-party registration occurs when the party sending the registration request is not the party that is being registered. In this case, the From header will contain the URI of the party submitting the registration on behalf of the party identified in the To header. Chapter 3 contains an example of a first-party

registration. An example third-party registration request for the user `Example` is shown below:

```
REGISTER sip:example.com SIP/2.0
Via: SIP/2.0/UDP 204.204.204.204;branch=1000000000
Max-Forwards: 70
To: sip:example.com
From: sip:example.com;branch=1000000000
Call-ID: 2000-0000-0000-0000-0000-0000-0000-0000
Seq: 1 00000000
Contact: sip:example.com;branch=1000000000
Expires: 3600
Contact-Header: 0
```

The mandatory headers in a REGISTER request are shown in Table 4.3.

4.3 BYE

The BYE method is used to terminate an established media session. In telephony, it is similar to a release message. A session is considered established if an INVITE has received a success-class response (200) or an ACK has been sent. A BYE is sent only by user agents participating in the session, never by proxies or other third parties. It is an unilateral method, so responses are only generated by the other user agent. A user agent responds with a 200 BYE or a transactional 404 Not Found to a BYE for an unknown dialog.

It is not recommended that a BYE be used to cancel pending INVITEs because it will not be treated like an CANCEL and may not reach the same set of user agents as the INVITE. An example BYE request looks like the following:

```
BYE sip:example.com SIP/2.0
Via: SIP/2.0/UDP 204.204.204.204;branch=1000000000
Max-Forwards: 70
To: sip:example.com;branch=1000000000
From: sip:example.com;branch=1000000000
Call-ID: 2000-0000-0000-0000-0000-0000-0000-0000
```

Table 4.3
Mandatory Headers for REGISTER Request

Call-ID
Seq
From
To
Via
Max-Forwards

```

CSeq: 47 BYE
Content-Length: 0

```

The mandatory headers in a BYE request are shown in Table 4-4.

4.4.4 ACK

The ACK method is used to acknowledge final responses or 2xx-4xx responses. Final responses to all other requests are never acknowledged. Final responses are defined as 2xx, 3xx, 4xx, 5xx, or 6xx class responses. The CSeq number is never incremented for an ACK, but the CSeq method is changed to ACK. This is so that a UAS can match the CSeq number of the ACK with the number of the corresponding 2xx-4xx.

An ACK may contain an application/sdp message body. This is permitted if the initial 2xx-4xx did not contain a SDP message body. If the 2xx-4xx contained a message body, the ACK may not contain a message body. The ACK may not be used to modify a media description that has already been sent in the initial 2xx-4xx; a re-2xx-4xx must be used for this purpose. SDP in an ACK is used in some interworking scenarios with other protocols where the media characteristics may not be known when the initial 2xx-4xx is generated and sent. An example of this is described in Section 18.6.

For 2xx responses, the ACK is end-to-end, but for all other final responses it is done on a hop-by-hop basis when useful proxies are involved. The end-to-end nature of ACKs to 2xx responses allows a message body to be transported. An ACK generated in a hop-by-hop acknowledgment will contain just a single Via header with the address of the proxy server generating the ACK. The difference between hop-by-hop acknowledgments to a response and end-to-end acknowledgments is shown in the message fragments of Figure 4-2.

A hop-by-hop ACK carries the same branch ID as the 2xx-4xx since it is considered part of the same transaction. An end-to-end ACK uses a different branch ID as it is considered a new transaction.

Table 4-4
Mandatory Headers in BYE Request

CSeq: 47
Content-Length: 0
From: sip:1000@10.10.10.10
To: sip:1000@10.10.10.10
Via: SIP/2.0/UDP;branch=1000

Table 6.1
Mandatory headers from SCC Request

to:CS-uri
From
To
PLN
Max-Forwards

4.1.5 CANCEL

The CANCEL method is used to terminate pending searches or call attempts. It can be generated by either user agents or proxy servers provided that a 2xx response containing a tag has been received, but no final response has been received. A user agent can use the method to cancel a pending call attempt it had earlier initiated. A holding proxy can use the method to cancel pending parallel branches after a successful response has been passed back to the UAC. CANCEL is a hop-by-hop request and receives a response generated by the next transfer element. The difference between a hop-by-hop request and an end-to-end request is shown in Figure 4.3. The CSReq is not implemented for this method as that proxy and user agents can match the CSReq of the CANCEL with the CSReq of the pending INVITE to which it corresponds.



Figure 4.3 End-to-end versus hop-by-hop requests.

The branch ID for a CANCEL matches the INVITE that it is canceling. A CANCEL only has meaning for an INVITE since only an INVITE may take several seconds (or minutes) to complete. All other SIP requests complete immediately (that is, a UAS must immediately generate a final response). Consequently, the final result will always be generated before the CANCEL is received.

A proxy receiving a CANCEL forwards the CANCEL to the same set of locations with pending requests that the initial INVITE was sent to. A proxy does not wait for responses to the forwarded CANCEL requests, but responds immediately. A user agent confirms the cancellation with a 200 OK response to the CANCEL and replies to the INVITE with a 487 Request Terminated response.

If a final response has already been received, a user agent will need to send a BYE to terminate the session. This is also the case in the case condition where a CANCEL and a final response cross in the network, as shown in Figure 4-4. In this example, the CANCEL and 200 OK response messages cross between the proxy and the UAS. The proxy will reply to the CANCEL with a 200 OK, but then also forwards the 200 OK response to the INVITE. The 200 OK response to the CANCEL sent by the proxy only means that the CANCEL request was received and has been forwarded—the UAC must still be prepared



Figure 4-4 Race condition in call cancellation

to receive further final responses. No 487 response is sent in this scenario. The session is canceled by the UAC sending an *ACK* then a *BYE* in response to the 200 *OK*.

Since it is a hop-by-hop response, a *CANCEL* message contains a message body. An example *CANCEL* request contains:

```
CANCEL sip:123456789@cs.toronto.edu;via;glo;diff2.0
Via: SIP/2.0/UDP 10.200.1.100;branch=2000000000
   ;branch=2000000000.2000000000
Max-Forwards: 70
To: Susan Heston <123456789@cs.toronto.edu>;via;glo
From: Susan Heston <123456789@cs.toronto.edu>;tag=2000000000
Call-ID: 20000-2000000000-00
Msg: 21432 2000000000
Contact: 123456789 <>
```

The mandatory header fields in a *CANCEL* request are shown in Table 4.6.

4.5.6 OPTIONS

The *OPTIONS* method is used to query a user agent or server about its capabilities and discover its current availability. The response to the request lists the capabilities of the user agent or server. A proxy server generates an *OPTIONS* request. A user agent or server responds to the request as it would to an *INVITE* (i.e., if it is not accepting calls, it would respond with a 4xx or 5xx response). A success class (2xx) response can contain *Allow*, *Accept*, *Accept-Encoding*, *Accept-Language*, and *Supported* headers indicating its capabilities.

An *OPTIONS* request may not contain a message body. A proxy determines if an *OPTIONS* request is for itself by examining the *Request-URI*. If the *Request-URI* contains the address of the proxy, the request is for the proxy. Otherwise, the *OPTIONS* is for another proxy or user agent and the request is forwarded. An example *OPTIONS* request and response contains:

Table 4.6
Mandatory Headers in a *CANCEL* Request

Call-ID
From
To
Via
Max-Forwards

```

OPTIONS 444 444@cam.ac.uk, via SIP/2.0
Via: SIP/2.0/UDP cam.ac.uk; branch=1; seq=1
Content-Disposition: session
Max-Forwards: 70
To: 4444@cam.ac.uk; call-id=1000
From: J.C. Marshall <jc@cam.ac.uk>; branch=1; seq=1
cseq=1000
Call-ID: 1000@cam.ac.uk; branch=1; seq=1
seq: 1 OPTIONS
Content-Length: 0

SIP/2.0 200 OK
Via: SIP/2.0/UDP cam.ac.uk; branch=1; seq=1001
Content-Disposition: session; seq=1001
To: <jc@cam.ac.uk>; branch=1; seq=1001
From: J.C. Marshall <jc@cam.ac.uk>; branch=1; seq=1001
cseq=1001
Call-ID: 1000@cam.ac.uk; branch=1; seq=1001
seq: 1 OPTIONS
Allow: INVITE, OPTIONS, ACK, BYE, CANCEL, REFER
Accept-Contact: *; *; IP
Contact: <jc@cam.ac.uk>
...

via
via...

```

The mandatory header in an `OPTIONS` request is the same as Table 4.3.

4.1.3 REFER

The `REFER` method [2] is used by a user agent to request another user agent to access a URI or UUI resource. The resource is identified by a URI or UUI in the required `Refer-To` header field (see Section 6.2.10). Note that the URI or UUI can be any type of URI: `sip`, `sip-ur`, `tel`, `mailto`, `urn` and so forth. When the URI is a `sip` or `sip-ur` URI, the `REFER` is probably being used to implement a call transfer service. `REFER` can also be used to implement peer-to-peer call control.

A `REFER` request can be sent either inside or outside an existing dialog. A typical call flow is shown in Figure 4.5. In this example, a UAC sends a `REFER` to a UAS. The UAS, after performing whatever authentication and authorization, decides to accept the `REFER` and responds with a 202 Accepted response. Note that this response is sent immediately without waiting for the triggered request to complete. This is important because `REFER` uses the non-INVITE method state machine, which requires an immediate final response, unlike an INVITE which may take several seconds (or even minutes) to complete. Since the `Refer-To` URI in this example is a `sip` URI, the UAC sends an INVITE using the `Refer-To` URI as the `Refer-To` URI. This



Figure 4.4. HTTP 303 example and flow.

HTTP is successful when it receives a 303-See other response. This successful response is conventional, but to the UAC using a NOTIFY method (described in Section 4.4.2). The message body of the NOTIFY contains a special copy of the final response to the original request. In this case, it contains the server's 303-See other. This part of a NOTIFY message is described in the NOTIFY-TOUCH header field in Section 4.4.2.2.2 [2].

An example of a NOTIFY message is shown below:

```

NOTIFY 200 OK, http://www.snoopy.com/ http://www.snoopy.com/
HTTP/1.1 200 OK (text/html)
Date: Wed, 07 Aug 2002 08:00:00 GMT
Server: Apache/1.3.3.6
Content-Length: 12
Content-Type: text/html

```

Another example of the use of HTTP as a “push” eWeb page is shown in Figure 4.5. In this UAC sends a NOTIFY to the UAS with a NOTIFY-TO set to an HTTP URL, as a Web page. This causes the UAS to send a 303 See other and then sends HTTP-GET requests to the Web server identifying the



Figure 4.7 RF Interfering in radio-relay station.

URL. After the Web page has loaded, the URL sends a `POST` consisting a body and `HTTP/1.0 200 OK`.

A `POST` and the `RF` request triggered by the `POST` may contain the file format, by header field just function (URL), which contains information about what requested (the request).

Figure 4.7 shows an abstract use of `POST` to implement a common `POST` or `POST` request between a remote node (X) and the server. The connection is assumed to be in a state of a session with the transmitter. The connection phase: the transmitter on field, then sends an `HTTP` to another pump, called the remote target. After the session is established between the transmitter and the remote target, the transmitter then gets the target on field. Now the transmitter has set on field system. The transmitter then sends a `POST` to the transmitter which causes the transmitter to generate a new `POST` (called a 'body' part) `POST` to the target. The successful `POST` 'updates' the existing session between the transmitter and the remote target. When the transmitter session notification that the remote was successful, the session between the transmitter and the transmitter is terminated with a `BYE`. This application can 'request' header fields in the `POST` to URL. That is, create `RF` header field, request and propagated to the URL which set then application the

subscription. If that time period expires before the triggered request has completed, both sides terminate the subscription, with the caller sending a final notification as discussed in the next section.

The subscription is terminated when the transfer target (the party that accepted the REFER) sends a final notification to NOTIFY with `Subscription-State: terminated; reason=success`. Usually, this is after the transfer target has received a final response to the triggered request. However, a transfer target that does not wish to establish a subscription and provide a final result of the REFER may send an immediate NOTIFY indicating that the subscription has been terminated. Each REFER now creates a separate subscription. If more than one REFER is sent within a dialog, the resulting notifications (and subscriptions) are identified by an ID parameter in the Event header field. The ID parameter is optional in refer triggered NOTIFYs except when multiple REFERs have been accepted, in which case it is mandatory.

The optional `Referred-By` header field can be included in a REFER request. Table 4.7 lists the mandatory header fields in a REFER request.

4.4.3 SUBSCRIBE

The `SUBSCRIBE` method [5] is used by a user agent to establish a subscription for the purpose of receiving notifications (via the NOTIFY method) about a particular event. A successful subscription establishes a dialog between the UAC and the UAS. The subscription request contains an `Expires` (see Section 6.4.7) header field, which indicates the desired duration of the existence of the subscription. After this time period passes, the subscription is automatically terminated. The subscription can be renewed by sending another `SUBSCRIBE` within the dialog before the expiration time. A server accepting a subscription returns a 200 OK response also containing an `Expires` header

Table 4.7
Mandatory Header Fields for REFER

To
From
Call-ID
Dialog
Contact
Max-Forwards
ToS
Refer-To

field. The replication does not do this as the request, as the server may choose the format, but it may not lengthen the format. There is no “`REPLICATION`” method used in SMTP—instead a `REPLICATE` with `Exp: rfc822` (compensate for omission of a subscription and hence the ability. If omitted subscription (after server returns not any continuation request) will result in a final `RCPTTO` indicating that the subscription has been terminated (see Section 4.1.4 for `RCPTTO`).

A 240 response to a `CONTINUE` does not indicate whether the subscription has been continued—it merely means it has been continued by the server.

The basic workflow is shown in Figure 4.8. The client sends a `STARTXID` (if), which is successful, and sends `RCPTTO` as the requested event name in the server. Before the expiration of the subscription time, the client sends `RCPTTO` to re-activate the subscription and hence receive more notifications.

Note that a client must be prepared to receive a `RCPTTO` before sending a 240 OK response to the `RCPTTO` (if). Also, during holding a subscription for prepared to receive `RCPTTO` from multiple servers (the `RCPTTO` will have



Figure 4.8 Example events and responses of flow.

different To tags and hence will establish separate dialogs, although only one 200 OK response to the SUBSCRIBE may be received.

An example SUBSCRIBE request is shown below:

```
SUBSCRIBE sip:101@cs.cmu.edu:5060;transport=udp; expires=300
Via: SIP/2.0/UDP proxy.cs.cmu.edu;branch=1
  ;seq=20-2000000000000000000
Via: SIP/2.0/UDP proxy.cs.cmu.edu;branch=1
  ;seq=20-2000000000000000000
Max-Forwards: 69
To: sipsip:101@cs.cmu.edu:5060;transport=udp
From: Thomas George <tdg@cs.cmu.edu>;branch=1
  ;seq=1-1000000000000000000
Call-ID: 1412 SUBSCRIBE
Allow-events: dialog
Contact: *tel/*cs.cmu.edu/*sip/*tdg@cs.cmu.edu;branch=1
  ;seq=1-1000000000000000000
```

The type of event subscription is indicated by the required **Event** header field (see Section 6.4.7) in the SUBSCRIBE request. Each application of the SIP Event framework [5] defines a package with a unique event tag. Each package defines the following things:

- Default subscription expiration interval
- Expected SUBSCRIBE message bodies
- What events cause a NOTIFY to be sent, and what message body is expected in the NOTIFY
- Whether the NOTIFY contains complete state or increments (default)
- Maximum notification rate

A protocol called PSTN and Internet Interworking (PINT) [6] defined methods SUBSCRIBE, NOTIFY, and UNSUBSCRIBE, which have a similar semantics to SIP. A server can distinguish a PINT SUBSCRIBE request from a SIP SUBSCRIBE by the absence of an **Event** header field in the PINT request.

A server should indicate which event packages it supports by listing them in an **Allow-Events** header field.

If a SUBSCRIBE refresh is sent within a dialog but receives a 481 Dialog Does Not Exist response, this means that the server has already terminated the subscription. The client should consider the dialog and subscription terminated and send a SUBSCRIBE to establish a new dialog and subscription.

An event template package is a special type of package that can be applied to any other package including statistical, access policy, and subscription lists. The application of a template package to a package is shown by

representing the package and template package names with a "." as in `presence-relinfo`, which is the application of the `matcher` info template package to the `presence` package. Table 4.8 lists the current set of SIP event and template packages.

Table 4.9 lists the mandatory header fields in a `REGISTER` request. Packages are standardized in the `SIPPING` or `SIMPLE.WG` based on the requirements in [5].

4.5.2 NOTIFY

The `NOTIFY` method [5] is used by a user agent to convey information about the occurrence of a particular event. A `NOTIFY` is always sent within a dialog

Table 88
Event Packages and Template Packages

<code>conference</code>	Conference information including participant list, join information, and so forth [2]
<code>dialog</code>	Dialog state/identification information [8]
<code>message-summary</code>	Message notification, used for message-routing indicators based with <code>relinfo</code> [8]
<code>presence</code>	Presence information [8]
<code>register</code>	Register state explicit subscription-oriented <code>NOTIFY</code> [2]
<code>call</code>	Event registration state [8]
<code>relinfo</code>	Relinfo information template package [12]

Table 89
Mandatory Header Fields for a `REGISTER`

To
From
CALL-ID
CDR
Max-Forwards
Min
Contact
Event
all low-entropy

Table 4.10
Mandatory Header Fields for a MESSAGE

To
From
Call-ID
CSeq
Max-Forwards
Via
Contact
Event
Subscription-State
Expires

participants engaged in a “conversation.” MESSAGEs may be sent within a dialog or outside a dialog, but they do not establish a dialog by themselves. The actual message content is carried in the message body as a MIME attachment. All UAs that support the MESSAGE method must support plain/text format; they may support other formats such as message/rfc822 [14] or text/html, or many others.

A MESSAGE request normally receives a 200 OK response to indicate that the message has been delivered to the final destination. An Instant Message response should not be sent in the message body of a 200 OK, but rather a separate MESSAGE request sent to the original sender. A 202 Accepted response indicates that the request has reached a store-and-forward device and will likely eventually be delivered to the final destination. In neither case does the 2xx response indicate that the message content has been rendered to the user.

A MESSAGE request may use the Im (Instant Message) URI scheme [14] in a Request-URI, although a client should try to resolve to a sip or sipso.

An example MESSAGE call flow is shown in Figure 4.5.

Note that the MESSAGE method is not the only application of instant messaging with SIP. It is also possible to use SIP to establish an instant message session in a completely analogous way that SIP is commonly used to establish a media session. An INVITE could be used to establish the session with a SDP body that describes the instant message protocol to be used directly between the two users. IM sessions have been proposed using SIP MESSAGE, Common Presence and Instant Messaging (CPIM) [14], and even Jabber [15]. An example call flow is shown in Figure 4.10.



Figure 4-10 SIP-based storage example.



Figure 4-11 Using SIP to establish a multimedia-capable session.

An example SIP-based session is shown below:

```

SIP/2.0 200 OK (application/sdp)
Via: SIP/2.0/UDP 192.168.1.100:5060;branch=1
Content-Type: application/sdp
Content-Disposition: inline
Content-Length: 100

```

```

To: <tel:12345678901010>
From: "G. L. Blandinet" <gibland@csd.sri.com>
Call-ID: 2047137170244240
CSeq: 1414 MESSAGE
Method: FLDP INV
Contact: <gibland@csd.sri.com>
Content-Type: text/plain
Content-Length: 0

```

8. 86

Table 4.11 lists the mandatory header fields in a MESSAGE request.

4.3.10 INFO

The INFO [36] method is used by a user agent to send call signaling information to another user agent with which it has an established media session. This is different from a re-INVITE since it does not change the media characteristics of the call. The request is end-to-end, and is never initiated by proxies. A proxy will always forward an INFO request—it is up to the UAS to check to see if the dialog is valid. INFO requests for unknown dialog receive a 401. Transaction/Dialog Does Not Exist responses.

An INFO method typically contains a message body. The contents may be signaling information, a midcall event, or some sort of stimulus. INFO has been proposed to carry certain PSTN midcall signaling information such as ISUP USR messages.

The INFO method always increments the CSeq. An example INFO method is:

```

INFO <gibland@csd.sri.com> sip:12345678901010 SIP/2.0
Via: SIP/2.0/UDP csd.sri.com; branch=1000; seq=1414
Content-Length: 0
Max-Forwards: 70

```

Table 4.11

Mandatory Header Fields for a MESSAGE

To
From
Call-ID
CSeq
Max-Forwards
Via

```

To: John FROTHING (John.Frothing@Example.com) [001-1111]
From: J.C. Howell (jc.howell@example.com) [001-1111]
Content-Type: multipart/mixed
Call-ID: 1847789482@Example.com
Call-Info: 1847789482@Example.com
Call-Reason: INVITE

```

```

300 OK (application/javascript)

```

The mandatory headers in an IMPP request are shown in Table 4.12.

4.3.2 PRACK

The PRACK [17] method is used to acknowledge receipt of reliably transported provisional responses (200X). The reliability of 200X, 300X, 400X, 500X, and 600X responses to INVITEs is achieved using the ACK method. However, in cases where a provisional response, such as 180 Ringing, is critical to determining the call state, it may be necessary for the receipt of a provisional response to be confirmed. The PRACK method applies to all provisional responses except the 180 Ringing response, which is never reliably transported.

A PRACK is generated by a UAC when a provisional response has been received containing a Ring reliable response number (see Section 4.3.1B) and a Support header: 100rel header. The PRACK echoes the number in the Ring and the CSeq of the response in a PRACK header. The message flow is as shown in Figure 4.5. In this example, the UAC sends the 180 Ringing response reliably by including the Ring header. When no PRACK is received from the UAC after the expiration of a timer, the response is retransmitted. The receipt of the PRACK confirms the delivery of the response and stops all further retransmissions. The 200 OK response to the PRACK stops retransmissions of the PRACK request. The call completes when the UAC sends the ACK in response to the 200 OK.

Table 4.12
Mandatory Headers in an IMPP Request

Mandatory Headers
Call-ID
CSeq
From
To
Via
Max-Forwards


```

Call-ID: 88888888-00000000-00000000-0000
Seq: 114
CSeq: 1 INVITE
Contact: <mailto:0>

From: sip:sara@exam.techlearningplc.com; SIP/2.0
Via: SIP/2.0/TCP 10.0.0.1:5060; branch=0; maddr=10.0.0.1
       ; branch=0; maddr=10.0.0.1
Max-Forwards: 70
To: ExamTech <mailto:sara@exam.techlearningplc.com>; o=ExamTech
From: ExamTech <mailto:sara@exam.techlearningplc.com>; o=ExamTech
Call-ID: 88888888-00000000-00000000-0000
Seq: 1 18488
Msg: 114 1 INVITE
Contact: <mailto:0>

SIP/2.0 200 OK
Via: SIP/2.0/TCP 10.0.0.1:5060; branch=0; maddr=10.0.0.1
       ; branch=0; maddr=10.0.0.1
To: ExamTech <mailto:sara@exam.techlearningplc.com>; o=ExamTech
From: ExamTech <mailto:sara@exam.techlearningplc.com>; o=ExamTech
Call-ID: 88888888-00000000-00000000-0000
Seq: 1 18488
Contact: <mailto:0>

```

The mandatory header fields in a **PRACK** request are shown in Table 4.13.

4.3.3 UPDATE

The **UPDATE** method [18] is used to modify the state of a session without changing the state of the dialog. A session is established in SIP using an **INVITE** request (see Section 4.1.2) in an offer/answer manner (see Section 7.1.2.2). Typically, a session offer is made in the **INVITE** and an answer made in a response to

Table 4.13
Mandatory Header Fields in a **PRACK** Request

Mandatory Header Fields
Call-ID
Seq
From
To
Via
Max-Forwards
Msg

Table 6.11
Mandatory Header Fields for an INVITE

To
From
Call-ID
CSeq
Route-Forwarded
Via
Contact

4.2 SIP and URI Schemes Used by SIP

SIP supports a number of URI and URL schemes including `sip`, `sip+`, `tel`, `presence`, and `im` for SIP, secure SIP, telephony, presence, and instant message URIs as described in the following sections. In addition, other URI schemes can be present in SIP header fields as listed in Table 6.15.

4.2.1 SIP and SIP URIs

The addressing scheme of SIP URIs and URLs has been previously mentioned. SIP URIs are used in a number of places including the `To`, `From`, and `Contact` headers, as well as the `Request-URI`, which indicates the destination. SIP URIs are similar to the `mailto` URI [29] and can be used in hypertext

Table 6.12
Common URI Schemes Used in SIP Messages

Scheme	Meaning	Use
<code>sip</code>	SIP URI	Most common <code>To</code> , <code>From</code> , <code>Contact</code> , <code>Request-URI</code> , and so forth [2]
<code>sip+</code>	Secure SIP URI	Same as SIP, provides end-to-end encryption using TLS [3]
<code>tel</code>	Telephone URI	Represents a telephone number [2, 30], usually in <code>To</code> or <code>From</code> [31]
<code>presence</code>	Presence URI	Used to represent the URI of a presence agent [4]
<code>im</code>	InstantMessage URI	Instant Message client [34]
<code>mailto</code>	Email URI	Can be included in <code>Contact</code> or registration registration response [23]
<code>http</code>	Web URL	Can be used in some headers such as <code>WWW-Info</code> , <code>Geo-Info</code> , and so forth [29]

subject, and priority to be set for the request. Additional headers can be specified, separated by a “;”. The header name/body indicates that the contents of a message body for an SIP/1.0 request is being specified in the URI.

If the parameter `name` option is present, then the username portion of the URI can be interpreted as a telephone number. This allows additional parameters in the username portion of the URI. This allows the parameters and structure of a tel URI [19] to be present in the user part of the SIP URI as described in the next section.

The sip URI scheme has the same structure as the sips URI but begins with the sip scheme name. Note that a sip URI is not equivalent to a sips URI with `transport=tls`, since the sip URI does not have the same security requirements as the sips URI. The requirement is that TLS transport is used end-to-end for the SIP path. The only exception is hop between the final proxy and the URS, which may use another security mechanism besides TLS (IPSec, for example).

Not shown in the example is the `host` route parameter `!r`, which can be present in sip or sips `Record-Route` and `Route` URIs to indicate that the proxy server identified by the URI supports host routing.

4.2.2 Telephone URIs

The telephone URI scheme, tel, [20] can be used to represent a resource identified by a telephone number. Telephone numbers can be of two general forms, local or global. A local number is only valid in a particular geographic area and has only local significance. If the number is used outside of this area, it will either fail or return the wrong resource. A global telephone number, also called an E.164 number, is one that is, in principle, valid anywhere. It contains enough information about the country, region, and locality so that the PSTN network is capable of routing calls to the correct resource. An example of a local phone number is

```
tel:01149300000000000000
```

which indicates a call to directory assistance valid only within country code 1 and area code 314 as identified in the required phone-context parameter. An example of a global phone number is

```
tel:+12126921212
```

Global phone numbers always begin with the “+” identifier followed by the country code. In this case, followed by the remaining telephone digits.

A tel URI can also contain some characters and information about dialing strings and patterns. For example:

4.3 Tags

A *tag* is a cryptographically random number with at least 32 bits of randomness, which is added to To and From headers to uniquely identify a dialog. The examples of Chapters 2 and 18 show the use of the tag header parameter. The To header in the initial INVITE will not contain a tag. A caller must include a tag in the From header, although a RFC 2543 user agent generally will not do so as it was optional in that specification. Including 100 Trying, all responses will have a tag added to the To header. The sending or reception of a response containing a From tag creates an early dialog. A tag returned in a 200 OK response is then incorporated as a dialog identifier and used in all future requests for this Call-ID. A tag is never copied across calls. Any response generated by a proxy will have a tag added by the proxy. An ACK generated by either a user agent or a proxy will always copy the From tag of the response in the ACK request.

If a UAC receives responses containing different tags, this means that the responses are from different UASs, and hence the INVITE has been forked. It is up to the UAC as to how to deal with this situation. For example, the UAC could establish separate sessions with each of the responding UAS. The dialogs would contain the same From-Call-ID, and CSeq, but would have different tags in the To header. The UAC also could BYE existing legs and establish only one session.

Note that tags are not part of the To or From URI but are part of the header and always placed outside any "<".

4.4 Message Bodies

Message bodies in SIP may contain various types of information. They may contain SIP information, which can be used to convey media information or QoS or even security information.

The optional Content-Disposition header is used to indicate the intended use of the message body. If not present, the function is assumed to be media, which means that the body describes a media session. Besides media, the other defined function is text, which means that the message body should be presented to the user or otherwise used or displayed. This could be used to pass a small JPEG image file as URI.

The format of the message body is indicated by the Content-Type header described in Section 6.4.5. If a message contains a message body, the message must include a Content-Type header. All user agents must support a Content-Type of application/sdp. The encoding scheme of the message body is indicated in the Content-Encoding header. If not

specified, the encoding is assumed to be text/plain. The specification of a Content-Encoding scheme allows the message body to be compressed.

The Content-Length header contains the number of bytes in the message body. If there is no message body, the Content-Length header should still be included but has a value of 0. Because multiple SMTP messages can be sent in a TCP stream, the Content-Length count is a reliable way to detect when one message ends and another begins. If a Content-Length is not present, the URC must assume that the message body continues until the end of the UDP stream, or until the TCP connection is closed, depending on the transport protocol.

Message bodies can have multiple parts if they are encoded using Multipart Internet Mail Extensions (MIME) [2]. Message bodies in SMTP, however, should be small enough so that they do not exceed the UDP MTU of the network. Proxies may reject requests with large message bodies with a 413 Request Entity Too Large response, since processing large messages can lead to a server.

As mentioned in the previous section, SMTP carries message bodies the same way that it sends email attachments. It is possible to carry multiple message bodies within a single SMTP message. This is done using a multipart MIME body. The Content-Type is listed as multipart/mime, and a parameter is defined, which is used by the parser to separate the message. Any SMTP request or response that can contain a message body may carry a multipart MIME body. An example is in SMTP (see Section 7.5) in which an SMTP/HTTP carries both a SMTP message body (application/vnd) and an encapsulated HTTP message (application/zip). An example multipart MIME is

```

MIME-Version: 1.0
To: <mailto:example@example.com>
From: <mailto:example@example.com>
Subject: Example
Content-Type: multipart/mixed; boundary="boundary"
Content-Length: ...

--boundary--

Content-Type: application/vnd
Content-Length: ...

vol
<mailto:example@example.com>
application/zip
  
```


- [8] Rosenberg, J., and H. Schulzrinne, "An IIS/NTT Initial Dialog Event Package for the Session Initiation Protocol (SIP)," IETF Internet-Draft, Work in Progress, March 2003.
- [9] Maly, R., "A Message Summary and Message Waiting Indication Event Package for the Session Initiation Protocol (SIP)," IETF Internet-Draft, Work in progress, March 2003.
- [10] Rosenberg, J., "A Presence Event Package for the Session Initiation Protocol (SIP)," IETF Internet-Draft, Work in Progress, January 2003.
- [11] Rosenberg, J., "A Session Initiation Protocol (SIP) Event Package for Registrations," IETF Internet-Draft, October 2002.
- [12] Rosenberg, J., "A Weather Information Event Template Package for the Session Initiation Protocol (SIP)," IETF Internet Draft, Work in Progress, January 2003.
- [13] Campbell, R., et al., "Session Initiation Protocol (SIP) Extension for Instant Messaging," RFC 3429, 2002.
- [14] Crockett, D., et al., "Common Presence and Instant Messaging (CPIM)," IETF Internet-Draft, Work in Progress, August 2002.
- [15] Sparks, R., "Enabling Jabber Messaging Sessions with the Session Initiation Protocol," IETF Internet-Draft, Work in Progress, October 2002.
- [16] Dromi, S., "The SIP/SIP2 Method," RFC 2879, 2000.
- [17] Rosenberg, J., and H. Schulzrinne, "Reliability of Provisional Responses in Session Initiation Protocol (SIP)," RFC 3363, 2002.
- [18] Rosenberg, J., "The Session Initiation Protocol (SIP) UPDATE Method," RFC 3311, 2002.
- [19] Schulzrinne, H., and A. Vitebsky, "The use of URIs for Telephone Calls," IETF Internet-Draft, February 2000.
- [20] Hoffman, P., I. Meunier, and J. Zvejbick, "The multi-URI Scheme," RFC 2448, 1998.
- [21] Belding, R., et al., "Hypertext Transfer Protocol — HTTP/1.1," RFC 2616, 1999.
- [22] Reed, N., and N. Borenstein, "Multipurpose Internet Mail Extension (MIME) Part One: Format of Internet Message Bodies," RFC 2045, 1996.

5

SIP Response Messages

This chapter covers the types of SIP response messages. A SIP response is a message generated by a UAC or a SIP server to reply to a request generated by a UAC. A response may contain additional header fields containing information needed by the UAC. Or, it may be a simple acknowledgment to prevent retransmissions of the request by the UAC. Many responses direct the UAC to take specific additional steps. The responses are discussed in terms of structure and class. Then, each response type is discussed and examined in detail.

There are six classes of SIP responses. The first five classes were borrowed from HTTP; the sixth was created for SIP. The classes are shown in Table 5.1.

If a particular SIP response code is not understood by a UAC, it must be interpreted by the class of the response. For example, an unknown 500 Server: Unplugged response must be interpreted by a user agent as a 500 Server: Failed response.

The reason phrase is for human consumption only—the SIP protocol uses only the response code in determining behavior. Thus, a 200 Call Failed is interpreted the same as 200 OK. The reason phrases listed here are the suggested ones from the RFC documents. They can be used to convey more information, especially in failure class responses—the phrase is likely to be displayed to the user. Some response codes were borrowed from HTTP, perhaps with a slightly different reason phrase.¹ However, not all HTTP response codes are valid in SIP, and some even have a different meaning. New response codes

1. Not all HTTP response codes are suggested in SIP. Only the response codes described in RFC 3261 and supporting RFCs are suggested in SIP.

Table 5.1
SIP Response Codes

Class	Description	Action
1xx	Informational	Indicates status of call prior to completion. If for informational/provisional responses.
2xx	Success	Request has succeeded. If for an INVITE, ACK should be sent, otherwise, retransmission of request.
3xx	Redirection	Server has returned possible locations. The client should retry request at another server.
4xx	Client error	Request has failed due to an error by the client. The client may retry the request if retransmission is appropriate.
5xx	Server failure	Request has failed due to an error by the server. The request may be retried at another server.
6xx	Global failure	Request has failed. Request should not be retried again at this or other servers.

created for SIP typically start at 2000 to try to avoid collisions with HTTP response codes.

Unless otherwise referenced, the responses described here are defined in RFC 3261 [3].

5.1 Informational

The informational class of responses (1xx) are used to indicate call progress. Informational responses are end-to-end responses and may contain message bodies. The exception to this is the 100 Trying response, which is only a hop-by-hop response and may not contain a message body. Any number of informational responses can be sent by a UAS prior to a final response (2xx, 3xx, 4xx, 5xx, or 6xx class response) being sent. The first informational response received by the UAC confirms receipt of the INVITE and stops retransmission of the INVITE, as described in Section 5.3. For this reason, servers receiving 100 Trying responses minimize INVITE retransmissions in the network. Further informational responses have no effect on INVITE retransmissions. A stateful proxy receiving a retransmission of an INVITE will resend the last provisional response sent to date. Informational responses are optional—a UAS can send a final response without first sending an informational response. While final responses to an INVITE receive an ACK to confirm

receipt, provisional responses are not acknowledged, except using the FRACK method described in Section 4.1.12.

All provisional responses with the exception of 100 Trying must contain a Content-Type URI and value and all Request-Header headers involved in the request. However, a RFC 2543 implementation will not do this, as it was not mandated in that document.

5.1.1 100 Trying

This special case response is only a hop-by-hop response. It is never forwarded and may not contain a message body. A forwarding proxy must send a 100 Trying response, since the extended needs being performed may take a significant amount of time. This response can be generated by either a proxy server or a user agent. It only indicates that some kind of action is being taken to process the call—it does not indicate that the user has been located. A 100 Trying response typically does not contain a To tag.

5.1.2 100 Ringing

This response is used to indicate that the INVITE has been received by the user agent and that ringing is taking place. This response is important in interworking with telephony protocols, and it is typically mapped to messages such as an ISDN Progress or ISUP Address Complete Message (ACM) [2]. When the user agent answers immediately, a 200 OK is sent without a 100 Ringing; this scenario is called the “fast answer” case in telephony.

A message body in this response could be used to carry QoS or security information, or to convey ring tone or animations from the UAS to the UAC.

A UA normally generates its own ringing tone or security ringing indication, unless a `Alert-Info-URI` header field (see Section 6.1.10) is present.

5.1.3 181 Call Is Being Forwarded

This response is used to indicate that the call has been handed off to another endpoint. This response is sent when this information may be of use to the caller. Also, because a forwarding operation may take longer for the call to be answered, this response gives a status for the caller.

5.1.4 183 Call Is Being Forwarded

This response is used to indicate that the 181/183 has been received, and will be processed in a queue. The reason phrase can be used to indicate the estimated wait time or the number of callers in line, as shown in Figure 5.1. A message body in this response can be used to carry music on hold or other media.




Figure 5.1 Call processing example with call processing center.

5.5.5 183 Session Progress

The 183 Session Progress response indicates that information about the progress of the session (call state) may be present in a message body or media stream. Unlike a 180 Ringing response, a 183 is an end-to-end response and does establish a dialog (must contain a To tag and Contact). Unlike a 180, 181, or 182 response, however, it does not convey any specific information about the status of the URFITS. A typical use of this response is to allow a UAC to hear ring tone, busy tone, or a recorded announcement in calls through a gateway into the PSTN. This is because call progress information is carried in the media stream in the PSTN. A one-way media connection or trunk is established from the calling party's telephone switch to the called party's telephone switch in the PSTN prior to the call being answered. In SIP, the media session is established after the call is answered—after a 200 OK and ACK have been exchanged between the UAC and URS. If a gateway used a 180 Ringing response instead, no media path would be established between the UAC and the gateway, and the caller would never hear ring tone, busy tone, or a recorded announcement (e.g., "The number you have dialed has changed, the new number is ...") since these are all heard in the media path prior to the call being answered. Figure 5.2 shows an example where a SIP caller does not hear a recorded announcement coming from the PSTN. Figure 5.3 shows the use of the 183 Session Progress allowing an early media session to be established prior to the call being answered. The PSTN interworking scenario in Chapter 10 shows this in detail.




Figure 5.2 LDAP interaction without early media.




Figure 5.3 LDAP interaction with early media.

5.2 Success

Success class responses indicate that the request has succeeded or has been accepted.

5.2.1 200 OK

The 200 OK response has two uses in SIP. When used to accept a session invitation, it will contain a message body containing the media properties of the UAS (called party). When used in response to other requests, it indicates successful completion or accept of the request. The response stops further retransmissions of the request. In response to an OPTIONS, the message body may contain the capabilities of the server. A message body may also be present in a response to a REGISTER request. For 200 OK responses to CANCEL, INFO, MESSAGE, SUBSCRIBE, NOTIFY, and REFER, a message body is not permitted.

5.2.2 202 Accepted

The 202 Accepted response [4] indicates that the UAS has received and understood the request, but that the request may not have been authorized or processed by the server. It is commonly used in responses to SUBSCRIBE (see Section 4.1.8) and REFER (see Section 4.1.7), and sometimes MESSAGE (see Section 4.2.18) methods.

5.3 Redirection

Redirection class responses are generally sent by a SIP server acting as a redirect server in response to an INVITE, as described in Section 5.3.2. A UAS, however, can also send a redirection class response to implement certain types of call forwarding features. There is no requirement that a UAC receiving a redirection response must proxy the request to the specified address. The UAC can be configured to automatically generate a new INVITE upon receipt of a redirection class response without requiring user interaction. In addition, proxies may also automatically send an ACK to a callee and proxy the INVITE to the new location provided in the Contact: URI of the redirection. To prevent looping, the server must not return any addresses contained in the request Via header field, and the client must check the address returned in the Contact: header field against all other addresses tried in an earlier call attempt. Note that this type of transaction looping is different from request looping.

5.2.1 300 Multiple Choices

This redirection response contains multiple Contact header fields, which indicates that the location server has returned multiple possible locations for the `uri` or `uris` URI in the Request-URI. The order of the Contact header fields is assumed to be significant. That is, they should be tried in the order in which they were listed in the response.

5.2.2 301 Moved Permanently

This redirection response contains a Contact header field with the new permanent URI of the called party. The address can be used and used in future INVITE requests.

5.2.3 302 Moved Temporarily

This redirection response contains a URI that is currently valid but that is not permanent. As a result, the Contact header field should not be cached across calls unless an Expires header field is present, in which case the location is valid for the duration of the time specified.

5.2.4 303 See Proxy

This redirection response contains a URI that points to a proxy server who has authoritative information about the calling party. The caller should forward the request to the proxy for forwarding. This response could be sent by a UAS that is using a proxy for incoming call screening. Because the proxy makes the decisions for the UAS on acceptance of the call, the UAS will only respond to INVITE requests that come from the screening proxy. Any INVITE request received directly would automatically receive this response without user intervention.

5.2.5 304 Alternative Service

This response returns a URI that indicates the type of service that the called party would like. An example might be a reference to a voicemail server.

5.4 Client Error

This class of response is used by a server or UAS to indicate that the request cannot be fulfilled as it was submitted. The specific client error response or the presence of certain header fields should indicate to the UAC the nature of the error and how the request can be reformulated. The UAC should not retransmit the

5.4.4 401 Unauthorized

This response is used to deny a request without giving the caller any reasons. It is sent when the server has understood the request, found the request to be correctly formatted, but will not service the request. This response is not used when authentication is required.

5.4.5 404 Not Found

This response indicates that the user identified by the a `Uri` or a `Uri` URI in the `Request-URI` cannot be located by the server, or that the user is not currently signed on with the user agent.

5.4.6 405 Method Not Allowed

This response indicates that the server or user agent has received and understood a request but is not willing to fulfill the request. An example might be a `POST` request sent to a user agent. An `Allow` header field (Section 6.4.1) must be present to inform the UAC what methods are acceptable. This is different from the case of an unknown method, in which a `501 Not Implemented` response is returned. Note that a proxy will forward request types it does not understand unless the request is targeted to the proxy server (i.e., the `Request-URI` is the URI of the proxy server).

5.4.7 406 Not Acceptable

This response indicates that the request cannot be processed due to a requirement in the request message. The `Accept` header field in the request did not contain any options supported by the UAS.

5.4.8 407 Proxy Authentication Required

This request sent by a proxy indicates that the UAC must first authenticate itself with the proxy before the request can be processed. The response should contain information about the type of credentials required by the proxy in a `Proxy-Authenticate` header field. The request can be retransmitted with the proper credentials in a `Proxy-Authorization` header field. Unlike in HTTP, this response may not be used by a proxy to authenticate another proxy.

```
HTTP/1.1 407 Proxy Authentication Required
Via: HTTP/1.1/0.9 31204166.10001102.000:7000-30000-2000000001
From: 30000-20000-2000000001@31204166.10001102.000
To: Internet; url=http://www.ietf.org/mail/ietf
X-123-456 30000-20000-2000000001-31204166.10001102
```

```

Class: 1 201108
From: Jonathan Levan <levan@cs.cmu.edu> [mailto:levan@cs.cmu.edu]
Subject: "SIP: REGISTER (RFC 3261) - REGISTER (RFC 3261)"
Response: 200 OK (application/javascript)
Content-Length: 0

```

5.4.8 408 Request Timeout

This response is sent when an Expires header field is present in an REGISTER request, and the specified time period has passed. This response could be sent by a floating proxy or a user agent. The request can be retried at any time by the UAC, perhaps with a longer time period in the Expires header field or no Expires header field at all. Alternatively, a successful proxy user could this response after the request transaction times out without receiving a final response.

5.4.9 409 Conflict

This response code has been removed from RFC 3261 but is defined in RFC 2545. It indicates that the request cannot be processed due to a conflict in the request. This response is used by registrars to reject a registration with a conflicting action parameter.

5.4.10 410 Gone

This response is similar to the 404 Not Found response but conveys the hint that the requested user will not be available at this location in the future. This response could be used by a service provider when a user cancels their service.

5.4.11 411 Length Required

This response code has been removed from RFC 3261 but is defined in RFC 2545. This response can be used by a proxy to reject a request containing a message body but no Content-Length header field. A proxy that takes a UDP request and forwards it as a TCP request could generate this response, since the use of Content-Length is more critical in TCP requests. However, the response code is not very useful since a proxy can easily calculate the length of a message body in a UDP request (it is until the end of the UDP packet) but cannot with a stream-oriented transport such as TCP. In this case, a missing Content-Length header field would cause the message body to go on indefinitely, which would generate a 513 Response Too Large response instead of a 411 Length Required.

5.4.13 403 Request Entity Too Large

This response can be used by a proxy to reject a request that has a message body that is too large. A proxy offering compression could use positively generate this response to save processing large requests.

5.4.14 404 Request-URI Too Long

This response indicates that the Request-URI in the request was too long and cannot be processed correctly. There is no maximum length defined for a Request-URI in the SIP standard document.

5.4.15 405 Unsupported Media Type

This response sent by a user agent indicates that the media type contained in the `Content-Type` request is not supported. For example, a request for a video conference to a PSTN gateway that only handles telephone calls will result in this response. The response should contain header fields to help the UAC reformulate the request.

5.4.16 406 Unsupported URI Scheme

The 406 Unsupported URI Scheme response is new to RFC 3261 and is used when a UAC uses a URI scheme in a Request-URI that the UAS does not understand. For example, if a request URI contains a non-SIP (e.g., sip) scheme that a proxy does not understand, it would return a 406 response. Since all SIP elements must understand the sip scheme, the request should be rejected using a 400 call in the request set.

5.4.17 408 Bad Extension

This response indicates that the extension specified in the Request-URI header field is not supported by the proxy or user agent. The response should contain a Supported header field (Section 6.1.12) listing the extensions that are supported. The UAC could resubmit the same request without the extension in the Request-URI header field (Section 6.2.12) or submit the request to another proxy or user agent.

5.4.18 409 Extension Required

The 409 Extension Required response indicates that a server requires an extension to process the request that was not present in a Supported header field in the request. The required extension should be listed in a Required header field in the response. The client should retry the request

adding the extension to a supported header field, or by the request as a different error that may not require the extension.

5.4.28 422 Session Timer Interval Too Small

The 422 *Session Timer Interval Too Small* response [4] is used to reject a request containing a *Session-Expires* header field (Section 6.2.26) with too short an interval. The ability to reject short durations is important to prevent excessive re-INVITE or UPDATE traffic. The minimum allowed interval is indicated in the required *Min-SE* header field (Section 6.3.4). The responder may reject the request without the *Session-Expires* header field or with a value less than or equal to the specified minimum.

5.4.29 423 Interval Too Brief

The 423 *Interval Too Brief* response is returned by a registrar that is rejecting a registration request because the requested expiration time on user or proxy contacts is too brief. The response must contain a *Min-Expires* header field (see Section 6.3.4) listing the minimum expiration interval that the registrar will accept. A client reporting a too short interval can conceivably lead a registrar server with registration refresh requests—this response allows a registrar to protect against this.

5.4.30 424 Use Authentication Tokens

The 424 *Use Authentication Tokens* response [5] is used by a UAS that is requiring the use of an *Authentication Information Body* (AIB) [6]. The AIB is a *SMIME* body that is an encrypted *message/sip* or *message/sipfrag* body. As a minimum, an AIB must contain the *From*, *Date*, and *Call-ID* header fields, and they should contain the *To*, *Contact*, and *Content* header fields as well. (The inclusion of these header fields is to prevent out-of-path attacks and hijacking.) The SIP request should contain a *Content-Type: message/sipfrag* header field and a *Content-Disposition: sdn* handling optional header field.

For more information on *SMIME* signatures, see Section 3.7.

5.4.32 425 Provide Referrer Identity

The 425 *Provide Referrer Identity* response [7] is used to request that a *Referred-By* header field be re-used with a valid *Referred-By*

security token. The security token is carried as an SMIME message body. The recipient of this error message (the UA that received and accepted the REFER) should reply this request back to the originator of the REFER by including it in a NOTIFY. The sender of the REFER can then generate the Refreshed-By security token and include it in the REFER, which would then be copied into the triggered request.

5.4.23 408 Temporarily Unavailable

This response indicates that the request has reached the correct destination, but the called party is not available for some reason. The reason phrase should be modified for this response to give the caller a better understanding of the situation. The response should contain a Retry-After header indicating when the request may be able to be fulfilled. For example, this response could be sent when a telephone has its ringer turned off, or a “do not disturb” button has been pressed. This response can also be sent by a redirect server.

5.4.24 401 Dialog Transaction Does Not Exist

This response indicates that a request referencing an existing call or transaction has been received for which the server has no records or state information.

5.4.25 402 Loop Detected

This response indicates that the request has been looped and has been sent back to a proxy that previously forwarded the request. Each server that forwards a request adds a Via header with its address to the top of the request. A branch parameter is added to the Via header, which is a hash function of the Request-URI and the To, From, Call-ID, and CSeq number. A second part is added to the branch parameter if the request is being forked. The reason the branch parameter must be checked is to allow a request to be sent back to a proxy, provided that the Request-URI has changed. This could happen with a call forwarding feature. In this case, the Via headers would differ by having different branch parameters.

5.4.26 403 Too Many Hops

This response indicates that the request has been forwarded the maximum number of times as set by the Max-Forwards header in the request. This is



Figure 5.4: Gateway dialing to the PSTN with SIP.

location, but this 400 response requires a choice by the caller, which is why it is classified as a client error class response. A server configured to receive this response must take user registration privacy into consideration; otherwise a rogue or general Request-URI could be used by a rogue user agent to try to discover tel:p or tel:psn URIs of registered users.

5.4.29 406 Busy Here

This response is used to indicate that the user agent cannot accept the call at this location. This is different, however, from the 503 *Busy Everywhere* response, which indicates that the request should not be tried elsewhere. In general, a 406 *Busy Here* is sent by a URS unless it knows definitively that the user cannot be contacted. This response is equivalent to the busy tone in the PSTN.

5.4.20 483 Request Terminated

This response can be sent by a user agent that has received a CANCEL request for a pending INVITE request. A 200 OK is sent to acknowledge the CANCEL, and a 200 is sent in response to the INVITE.

5.4.21 484 Not Acceptable Here

This response indicates that some aspect of the proposed session is not acceptable and may contain a Warning header field indicating the exact reason. This response has a similar meaning to 404 Not Acceptable, but only applies to one location and may not be true globally as the 404 response indicates.

5.4.22 485 Bad Event

The 485 Bad Event response [4] is used to reject a subscription request or notification containing an Event package that is unknown or not supported by the UAS. The response code is also used to reject a subscription request that does not specify an Event package, assuming that the server does not support the PINT protocol (see Section 4.1.6).

5.4.23 486 Request Pending

The 486 Request Pending response is used to resolve accidental simultaneous re-INVITEs by both parties in a dialog. Since both INVITEs seek to change the state of the session, they cannot be processed at the same time. While a user agent is awaiting a final response to a re-INVITE, any re-INVITE request received must be replied to with this response code. This is analogous to the “glace” condition in telephony in which both ends seize a trunk at the same time. The reconciliation algorithm in SIP is for the user agent to generate a delay (randomly selected within a range determined by δ) the user agent send the initial INVITE or send then retry the re-INVITE, assuming that another re-INVITE has not been received in the meantime. In this way, one side or the other will “win” the race condition and have the re-INVITE processed.

An example is shown in Figure 5.5.

5.4.24 488 Request Undecipherable

The 488 Request Undecipherable response is used when an S/MIME message body can not be deciphered because the public key is unavailable. If the UAS does not support S/MIME, no message body will be present in the response. If the UAS does support S/MIME, the response will contain a message

can be used to identify the type of failure. The client can retry the request again at this server after several seconds.

5.5.2 501 Not Implemented

This response indicates that the server is unable to process the request because it is not supported. This response can be used to decline a request containing an unknown method. A proxy, however, will forward a request containing an unknown request method. Thus, a proxy will forward an unknown `MESSAGE-INDIRECT` request, assuming that the UAS will generate this response if the method is not known.

5.5.3 502 Bad Gateway

This response is sent by a proxy that is acting as a gateway to another network, and indicates that some problem in the other network is preventing the request from being processed.

5.5.4 503 Service Unavailable

This response indicates that the requested service is temporarily unavailable. The request can be retried after a few seconds, or after the expiration of the `Retry-After` header field. Instead of generating this response, a loaded server may just refuse the connection. This response code is important in that its receipt triggers a new DNS lookup to locate a backup server to obtain the desired service. The set of SIP DNS procedures for locating SIP servers is defined in [9].

5.5.5 504 Gateway Timeout

This response indicates that the request failed due to a timeout encountered in the other network to which the gateway connects. It is a server error class response because the call is failing due to a failure of the server in accessing resources outside the SIP network.

5.5.6 505 Version Not Supported

This response indicates that the request has been refused by the server because of the SIP version number of the request. The detailed semantics of this response have not yet been defined because there is only one version of SIP (version 2.0) currently implemented. When additional version numbers are implemented in the future, the mechanisms for dealing with multiple protocol versions will need to be detailed.

5.5.7 511 Message Too Large

The 511 *Message Too Large* response is used by a URS to indicate that the request size was too large for it to process.

5.6 Global Error

This response class indicates that the server knows that the request will fail wherever it is tried. As a result, the request should not be sent to other locations. Only a server that has definitive knowledge of the case identified by the Request-URI in every possible instance should send a global error class response. Otherwise, a client error class response should be sent. A *Retry-After* header field can be used to indicate when the request might be successful.

5.6.1 500 Internal Server Error

This response is the definitive version of the 405 *Busy Here* client error response. If there is a possibility that the call to the specified Request-URI could be answered in other locations, this response should not be sent.

5.6.2 500 Decline

This response has the same effect as the 405 *Busy Everywhere* but does not give away any information about the call state of the server. This response could indicate the called party is busy, or simply does not want to accept the call.

5.6.3 504 Does Not Exist Response

This response is similar to the 404 *Not Found* response but indicates that the user in the Request-URI cannot be found anywhere. This response should only be sent by a server that has access to all information about the user.

5.6.4 506 Not Acceptable

This response can be used to implement user session negotiation capability in SIP. This response indicates that some aspect of the desired session is not acceptable to the URS, and as a result, the session cannot be established. The response may contain a *Session-req* header field with a numerical code describing exactly what was not acceptable. The request can be revised with different media session information. An example of simple negotiation with SIP is shown in Figure 5.6. If more complicated negotiation capability is required, another protocol should be used.




Figure 8.8: Session registration with QIP.

References

- [1] Rosenberg, J., et al. "QIP:Session Initiation Protocol." RFC 3261, 2002.
- [2] Camarillo, G., et al. "Integrated Services Digital Network (ISDN) User Part (ISUP) to Session Initiation Protocol (SIP) Mapping." RFC 3398, 2002.
- [3] Bredt, A. "SIP Specific Events." RFC 3261, 2001.
- [4] Davidson, S., and J. Rosenberg. "Session States in the Session Initiation Protocol (SIP)." IETF Internet-Draft, Work in Progress, November 2002.
- [5] Peterson, J. "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)." IETF Internet-Draft, Work in Progress, February 2003.
- [6] Peterson, J. "SIP Authenticated Identity Body (AIB) Format." IETF Internet-Draft, Work in Progress, February 2003.
- [7] Sparks, R. "The SIP Refresh By Mechanism." IETF Internet-Draft, Work in Progress, February 2003.
- [8] Avshalom, E., *Introduction to Telecommunication Network Engineering*, Norwood, MA: Artech House, 1999.
- [9] Rosenberg, J., and H. Schulzrinne. "Session Initiation Protocol (SIP): Locating SIP Servers." RFC 3263, 2002.

6

SIP Header Fields

This chapter describes the header fields present in SIP messages. In RFC 2543, there were four categories of SIP header fields: general, request, response, and entity. RFC 3261 removed this distinction since it was not used by the protocol. However, the header fields discussed in this chapter are categorized as request and response, request only, response only, and message body header fields. Except as noted, header fields are defined in the SIP specification RFC 3261 [1].

SIP header fields in most cases follow the same rules as HTTP header fields [2]. Header fields are defined as `Header: field`, where `Header` is the case-insensitive token (not conventionally lower case with some capitalization) used to represent the header field name, and `field` is the case-insensitive set of tokens that contain the information. Except when otherwise noted, their order in a message is not important. Header fields can occur on one or multiple lines as long as the line begins with at least one space or horizontal tab character. Unrecognized header fields are ignored by proxies. Many common SIP header fields have a compact form, where the header field name is denoted by a single lowercase character. These header fields are shown in Table 6-1. Header fields can be either end-to-end or hop-by-hop. Hop-by-hop header fields are the only ones that a proxy may insert or, with a few exceptions, modify. Because SIP typically involves end-to-end contact, most header fields are end-to-end. The hop-by-hop header fields that may be inserted by a proxy are shown in Table 6.2.

Table 6.1
Compact Form-to-SIP Header Fields

Header Field	Compact Form
accept-encodings	0
allow-branches	10
Call-ID	1
Contact	00
CONTACT-ADDRESS	0
Contact-Expires	1
CONTACT-ID	0
Event	0
EX-100	0
Expires-To	0
REFERED-TO	0
Refer-ToContact	0
REFERED	0
To	0
Via	0

6.1 Request and Response Header Fields

This set of header fields can be present in both requests and responses.

6.1.1 Alert-Info

The Alert-Info header field can be used to provide a “directive ring” service. If present in an INVITE, the UAC may use the URI in such an alert tone to be used in place of the default alerting tone—that is, it would be forwarded to the called party. If present in a 180 Ringing response, the UAC may use the URI in such a ring-back tone to be rendered to the calling party. In both cases, the URI is fetched and rendered without user intervention, so careful policy rules are necessary to avoid unwanted sounds and notes being generated.

One use is for a trusted proxy to insert the header field with a local (to the domain of the user agent’s URI) URI. This then allows for very simple policy in the user agent in deciding whether or not to render.

An example is shown here:

parameters, with the header field parameters outside the `< >`, even if no display name is present.

There are three additional parameters defined for use in Contact header fields: `q`, `mailto`, and `expires`. They are placed at the end of the URI and separated by semicolons. The `q`-value parameter is used to indicate relative preference, which is represented by a decimal number in the range 0 to 1. The `q` value is not a probability, and there is no requirement that the `q` values for a given list of Contacts add up to 1. (The `mailto` parameter defined in RFC 2143 has been deprecated and is not used in RFC 3261. It was only used in registration Contact header fields, and is used to specify proxy or redirect operation by the server.) The `expires` parameter indicates how long the URI is valid and is also only used in registrations. The parameter `order` contains an integer number of seconds as a date in SIP form (see Section 6.1.4). Examples are shown in Table 6.3.

The CONTACT header field may contain a feature tag [+], which can be used to indicate the capabilities of the device identified by the Contact URI. For example, the feature tag `audio` is used to indicate that the URI in the Contact header field is a Conference URI, and that the dialing is associated with a focus. A focus is a SIP user agent that hosts a particular instance of a conference, called a "bridge" or MCU in other protocols. The presence of the `audio` feature tag can be used by a SIP user agent that supports advanced

TABLE 3
Examples of Contact Header Fields

Header Field	Meaning
<code>Contact: sip:john@lisp.example.com</code>	A single SIP URI without a display name.
<code>Contact: sip:12345@678.lisp.example.com; q=0.5</code>	A display name with the URI is enclosed in <code>< ></code> ; the display name is treated as a token and ignored.
<code>Contact: 06_Paradise; sip:paradise@lisp.example.com; "Paradise" <mailto:paradise@lisp.example.com></code>	Two URIs are listed, the second being a non-SIP URI with a display name enclosed in quotes.
<code>Expires: 3600; Contact: sip:12345@678.lisp.example.com; Expires: 1800</code>	The complete form of the header field is used for a single URI. The URI contains a port number and a SIP parameter enclosed within the <code>< ></code> . An expires header field parameter uses a SIP date enclosed in the quotes.

conferencing features to include certain call control operations [5] or subscribe to the conference package [6].

Other feature tags are listed in Table 6.4. The compact form is in

6.3.5 CSeq

The command sequence CSeq header field is a required header field in every request. The CSeq header field contains a decimal number that increases for each request. Usually, it increases by 1 for each new request, with the exception of CANCEL and ACK requests, which use the CSeq number of the INITIAL request to which it refers.

The CSeq count is used by UACs to determine out-of-sequence requests or to differentiate between a new request (different CSeq) or a retransmission (same CSeq). The CSeq header field is used by UACs to match a response to the request it references. For example, a UAC that sends an INITIAL request then a CANCEL request can tell by the method in the CSeq of a 200 OK response if it is a response to the last action or cancellation request. Examples are shown in Table 6.3.

Each user agent maintains its own command sequence number space. For example, consider the case where user agent 1 establishes a session to user agent 2 and initializes its CSeq count to 1. When user agent 2 initiates a request (such as a INITIAL or INFO), or even ACK, it will initialize its own CSeq space, totally independent of the CSeq count used by user agent 1. The examples of Chapter 10 show this behavior of CSeq.

6.3.6 Date

The Date header field is used to convey the date when a request or response is sent. The format of a SIP date is based on HTTP dates, but allows only the

Table 6.4
Global Feature Tags

Feature tag	Meaning
audio	Supports audio
video	Supports video
image	Supports images
text	Supports messaging
call	Supports call
voluntary	Is voluntary user
conference	Is a proxy, conference user

Table 6.4
Client-side Field/Comment

Header Field	Meaning
From: 1 200108	The current sequence number has been initialized to 1 for this initial request.
From: 411 200108	The current sequence number is set to 411 for this 300th request.
From: 4147 200108	This was the last request by the user agent for this destination since the Client was initialized to 400, or the previous request generated by the Call Center or 300E or other request-waiting logic. Client 40000 uses.

preferred Internet date standard referenced by RFC 1123 [7]. To keep user agent data and time logic simple, SIP only supports the use of the GMT time zone. This allows time entries that are stored in data form rather than record form to be easily converted into delta seconds without requiring knowledge of time zone offsets.

A DATE example is shown here:

```
Date: Fri, 12 Oct 2000 23:45:00 GMT
```

6.3.7 Encryption

The **Encrypt** header field was defined in RFC 2543 but is not included in RFC 3261. Instead, encryption using **SRTP** is defined as discussed in Section 3.8.

6.3.8 From

The **From** header field is a required header field that indicates the originator of the request. It is one of two addresses used to identify the dialog. The **From** header field contains a URI, but it may not contain the transport, method, or ttl URI parameters. A **From** header field may contain a tag, used to identify a particular call. A **From** header field may contain a display name, in which case the URI is enclosed in $\langle \rangle$. If there is both a URI parameter and a tag, then the URI including any parameters must be enclosed in $\langle \rangle$. Examples are shown in Table 6-6.

A **From** tag was optional in RFC 2543 but is mandatory to include in RFC 3261.

Table 6.8
Examples of from-header lines

Header Field	Meaning
From: <alice@acme.com>@acme.com <alice@11.11.11>	A single SIP URI with a tag
From: Bob@acme.com <alice@11.11.11>@acme.com; <bob@11.11.11>	A source SIP URI with multiple tags
From: "Alice <alice@acme.com>" <alice@11.11.11>@acme.com <11.11.11@11.11.11>	Using the compact form of the header field, with parentheses quotes along with a SIP URI with parameters inside the quotes
From: null < >	A null URI without a display name or tag; the <>-> is required, mandated by RFC 3261.4

6.5.3 Organization

The Organization header field is used to indicate the organization to which the originator of the message belongs. It can also be located by parties as a message is passed from one organization to another. Like all SIP header fields, it can be used by parties for making routing decisions and by user agents for making call answering decisions.

An example is below:

```
Organization: BCE
```

6.5.4 Record-Route

The Record-Route header field is used to force routing through a proxy for all subsequent requests in a session between two user agents. Normally, the presence of a Contact header field allows user agents to send messages directly bypassing the proxy chain used in the initial request (which probably involved database lookups to locate the called party). A proxy inserting its address into a Record-Route header field overrides this and forces future requests to include a Record-Route header field containing the address of the proxy that forces this proxy to be included.

A proxy, such as a firewall control proxy, wishing to implement this inserts the header field containing its own URI, or adds its URI to an already present Record-Route header field. The URI is constructed so that the URI resolves

back to the proxy server. The UAC copies the Record-Route header field into the 200 OK response to the request. The header field is forwarded unchanged by proxies back to the UAC. The UAC then uses the Record-Route proxy list plus a Contact header field if present in the 200 OK for use in a Route header field in all subsequent requests. Because Record-Route is bidirectional, messages in the reverse direction will also traverse the same set of proxies. Chapter 10 contains an example of the use of the Record-Route and Route header fields. The Lr parameter is new to RFC 3261 and indicates that the proxy server supports “local routing.” Older RFC 2546 compliant proxy servers create Record-Route URIs that instead of the Lr parameter often contain the media parameters with an address or host that routes to that proxy server.

Example use:

```
record-route: <udp-proxy.example.com>,
  <udp://192.0.2.100.example.com>
Record-Route: <udp://192.0.2.100>
```

6.5.5 Retry-After

The Retry-After header field is used to indicate when a resource or service may be available again. In 503 Service Unavailable responses, it indicates when the server will be available. In 404 Not Found, 400 Busy Everywhere, and 401 Unauthorized responses, it indicates when the called user agent may be available again.

The header field can also be included by proxy and redirect servers in responses if a route registration was removed with a Retry-After header field indicating when the user may sign on again. The contents of the header field can be either an integer number of seconds or a SIP date. A duration parameter can be used to indicate how long the resource will be available after the time specified. Examples of this header field are shown in Table 6.7.

6.5.6 Subject

The optional Subject header field is used to indicate the subject of the media session. It can be used by user agents to do simple call screening. The contents of the header field can also be displayed during alerting to aid the user in deciding whether to accept the call. The compact form of this header field is a. Some examples are:

```
Subject: Steve just calls about SIP
a: SIP 200 66666; 20077
```


Table 6.1
Examples of Retry-After header field

Header Field	Meaning
Retry-After: 3000	Request can be retried again in 1 hour
Retry-After: Sat, 20 May 2006 08:00:00 GMT	Request can be retried after the date listed
Retry-After: 3000	Request can be retried after 1 hour
Retry-After: Sat, 20 Feb 2006 12:00:00 GMT (after a 408 error)	Request can be retried after the specified date for 30 minutes

6.1.3 Supported

The **Supported** header field is used to list one or more options implemented by a user agent or server. It is typically included in responses to OPTIONS requests. If no options are implemented, the header field is not included. If a UAC has an option in a Supported header field, proxies or UASs may use the option during the call. If the option must be used as supported, the Required header field is used instead. Table 6.8 shows the current set of defined features tags.

An example of the header field is

```
Supported: call-transfer
```

6.1.4 Timestamp

The **Timestamp** header field is used by a UAC to mark the exact time a request was generated in some numerical time format. A UAS must echo the

Table 6.8
Extension Feature Tags

Tag	Meaning
call-transfer	SIP Forking [2]
call-info	Additional call parameters [2]
call-rtcp	Push header field [2]
call-transfer	Reliable provisional responses (PRACK) support [2]
call-transfer	Replace call control parameter [17]
call-transfer	Session Timer feature [2]

header field in the response to the request and may add another numerical time entry indicating the amount of delay. Unlike the Date header field, the time format is not specified. The most accurate time format should be used, including a decimal point. Examples are shown in Table 6-9.

6.18.10 To

The To header field is a required header field in every SIP message used to indicate the recipient of the request. Any responses generated by a user agent will contain this header field with the addition of a tag. (Note that an RFC 2543 client will typically only generate a tag if more than one Via header field is present in the request.) Any response generated by a proxy must have a tag added to the To header field. A tag added to the header field in a 200 OK response is used throughout the call and incorporated into the dialog. The To header field is never used for routing—the Request-URI is used for this purpose. An optional display name can be present in the header field, in which case the SIP URI is enclosed in < >. If the URI contains any parameters or extension parameters, the URI must be enclosed in < > even if no display name is present. The compact form of the header field is t. Examples are shown in Table 6-18.

6.18.11 User-Agent

The User-Agent header field is used to convey information about the user agent originating the request. Based on the HTTP header field of the same name [2], it can contain manufacturer information, software version, or comments. The field may contain multiple values, with the ordering assumed to be from most general to most specific. This information can be used for logging or for generating a specific response for a specific user agent. For security reasons, this header field may be suppressed. For example, an attacker probing a UA for vulnerabilities could learn the particular vendor and software level that is susceptible to a particular attack and reuse that attack against other UAs that have the same software as identified by the User-Agent header field.

Examples include:

Table 6-9
Example of Timing Header Field

Header Field	Meaning
Timing-Info: 120.12	Client has elapsed 120 seconds for the request.
Timing-Info: 120.12 - 10	This header field from the response has the delay time added by the client.

Table 6.11
Examples of Via header field

Header Field	Meaning
To: sip:1042@paris42.org;branch=bar;tag=123456789	A single SIP URI with a tag and without a display name
To: Thomas Watson -sip:1042@paris42.org;branch=bar	A display name (local), with the sip: URI enclosed in <>
To: "Tom W." <1042@paris42.org>;branch=bar	A display name (quoted) along with the SIP URI enclosed within <>
To: <1-114-555-1110@paris42.org>;branch=bar;tag=123456789	Both a URI parameter and tag are used, as URI is enclosed in <>

```
From: Agustin <1042@paris42.org>
Via: SIP/2.0/TCP;branch=bar;tag=123456789
```

6.1.2 Via

The required *Via* header field is used to record the SIP route taken by a request and is used to route a response back to the originator. A user agent generating a request records its own address in a *Via* header field in the request. While the ordering of most SIP header fields is not significant, the *Via* header fields order is significant because it is used to route responses. A proxy forwarding the request adds a *Via* header field containing its own address to the top of the list of *Via* header fields. A proxy adding a *Via* header field always includes a branch tag consisting a cryptographic hash of the To, From, Call-ID header fields and the Request-URI. A proxy or user agent generating a response to a request copies all the *Via* header fields from the request in order into the response, then routes the response to the address specified in the top *Via* header field. A proxy receiving a response checks the top *Via* header field to ensure that it matches its own address. If it does not, the response has been misrouted and should be discarded. The top *Via* header field is then removed, and the response forwarded to the address specified in the next *Via* header field. A simplified *Via* header field is shown in Figure 6.1.

Via header fields contain protocol name and version number and transport (SIP/2.0/UDP, SIP/2.0/TCP, etc.) and may contain path information and parameters such as *branch*, *method*, *media*, and *ttl*. A *received* tag is added to a *Via* header field if a user agent or proxy receives the request from a different address than that specified in the top *Via* header field. This indicates that a NAT or firewall proxy is in the message path. If



Figure 6.8 Verifying decision 194.

present, the received tag is used in response routing. (The hidden parameter, deprecated in RFC 3261, was used to indicate the Via header field has been overwritten.) A branch parameter is added to Via header fields by UAs and proxies, which is composed as a hash function of the Request-URI and the To-From, Call-ID and CSeq number. A second pair is added to the branch parameter if the request is being forked as shown in Figure 3-4. The middle and left parameters are used for median transport and have a similar meaning as the equivalent SIP URI parameters. The compact form of the header field is v. Examples are given in Table 6.11.

Table 6.11
Examples of SIP header fields

Header Field	Meaning
Via SIP/2.0/SIP/UDP 192.168.1.100:5050 (branch=164484773)	IP4 address using unicast UDP transport (enclosed part of list)
Via SIP/2.0/SIP/TCP 192.168.1.100:5050 (branch=164484773)	Domain name using TCP transport and port number 5050
Via SIP/2.0/SIP/UDP 192.168.1.100:5050 (branch= 164484773;CSeq=1)	Proxy added via header field with branch
V: SIP/2.0/SIP 192.168.1.100:5050 (branch= 164484773;CSeq=1)	Compare form with domain name using SIP, find exact location of being proxy
Via SIP/2.0/SIP 192.168.1.100 (remote=192.168.1.100;branch=164484773)	The address is a multicast address. IP4 address is roughly unique. Request has been forwarded through a proxy which changed the From: field to a globally unique one
Via SIP/2.0/SIP 192.168.1.100 (remote=192.168.1.100;branch=164484773)	The address is a multicast address specification made with a qualified IP

6.2 Request Header Fields

This set of header fields can only be present in a request.

6.2.1 Accept

The **Accept** header field is defined by HTTP [2] and is used to indicate acceptable message formats/media types [10] in the message body. The header field declares media types using the format *type/sub-type* (usually used in the format: *type/sub-type; q=qvalue*). If not present, the assumed acceptable message body format is *application/sdp*. A list of media types can have preferences set using *q* value parameters. The wildcard "*" can be used to specify all media-types. Examples are given in Table 6.12.

6.2.2 Accept-Contact

The **Accept-Contact** [4] header field specifies in which URIs the request may be processed. Some additional parameters are also defined for Contact:

Table 6.10
Examples of Accept Header Field

Header Field	Meaning
Accept: application/sdp	This is the default assumption if no Accept header field is present
Accept: */*	Accept all not existing schemes
Accept: application/sdp;q=0.9,application/xhtml+xml;q=0.8,application/xhtml+xml;q=0.7	Use SIP (application/sdp) as preferred, use XML (application/xhtml+xml) as second

header fields such as `media`, `display`, and `language`. This header field is part of the caller preferences extensions to SIP, which have been defined to give some control to the caller in the way a proxy server processes a call. The compact form is as follows.

Some examples follow:

```
Accept-Contact: *; Language=en
or: *; media=video
```

6.2.3 Accept-Encoding

The `Accept-Encoding` header field, defined in HTTP [1], is used to specify acceptable message body encoding schemes. Encoding can be used to create a SIP message with a large message body fit inside a single UDP datagram. The use of `q` value parameters can set preferences. If none of the listed schemes are acceptable to the UAC, a 406 Not Acceptable response is returned. If not included, the assumed encoding will be text/plain. Examples include:

```
Accept-Encoding: text/plain
Accept-Encoding: */*
```

6.2.4 Accept-Language

The `Accept-Language` header field, defined in HTTP [1], is used to specify preferences of language. The language specified can be used for reason phrases in responses, informational header fields such as `Subject`, or in message bodies. The HTTP definition allows the language tag to be made of a primary tag and an optional subtag. This header field would also be used by a proxy to route to a human operator in the correct language. The language tags are registered by IANA. The primary tags are ISO-639 language abbreviations. The use of `q` values allows multiple preferences to be specified. Examples are shown in Table 6.11.

6.2.7 Event

The **Event** header field is used in a **SUBSCRIBE** (see Section 4.1.8) or **NOTIFY** (see Section 4.1.9) methods to indicate which event package is being used by the method. In a **SUBSCRIBE**, it lists the event package to which the client would like to subscribe. In a **NOTIFY**, it lists the event package that the notification contains more information about. Currently defined event packages are listed in Table 4.5. The compact form is 0.

An example follows:

```
Event: message
; value
```

6.2.8 Hide

The **Hide** header field was defined in RFC 2543 but has been deprecated from RFC 3261. It was intended to be used by user agents or proxies to request that the next hop proxy encrypts the **Via** header fields to hide message routing path information. Encrypted **Via** headers were identified with the **Address** **Via** parameter. However, the security provided and the mechanism requiring next hop was made the value of this header field minimal.

6.2.9 In-Reply-To

The **In-Reply-To** header field is used to indicate the **Call-ID** that this request references or is returning. For example, a missed call could be returned with a new **Invite** and the **Call-ID** from the missed **Invite** copied into the **In-Reply-To** header field. This allows the UAS to determine that this is not an unrelated call, which could be used to override call screening logic, for example. Examples of this header field are as follows:

```
In-Reply-To: id-43-32-44-@example.com
In-Reply-To: id-43-32-44-@example.com, id-43-32-44-@example.com
```

6.2.10 Join

The **Join** header field [4] is used in an **Invite** to request that the dialog (or chat) be joined with an existing dialog (or chat). The parameter of the **Join** header field identifies the dialog by the **Call-ID**. To **req** and **From** tag is a similar way to the **Require** header field.

If the **Join** header field references a point-to-point dialog between two user agents, the **Join** header field is effectively a request to turn the call into a conference call. If the dialog is already part of a conference, the **Join** header field is a request to be added into the conference. An example call flow is shown in Figure 6.2 in which a one-way call is turned into a conference call.


```
Priority: 1000000000
```

6.2.12 Priority

The **Priority** header field [14] is used by a UAC to request varying degrees and types of priority. Currently defined tags include critical, Amaster, id, session, user, or none.

An example follows:

```
Priority: header-critical-critical
```

6.2.13 Proxy-Authentication

The **Proxy-Authentication** header field is to carry the credentials of a user agent in a request to a server. It can be used in reply to a 407 Proxy Authentication Required response containing challenge information, so it can be sent first, without waiting for the challenge if the form of the challenge is known (e.g., if it has been cached from a previous call). The authentication mechanism for SIP digest is described in Section 3.8. A proxy receiving a request containing a **Proxy-Authentication** header field searches for its own realm. If found, it processes the proxy. If the credentials are correct, any remaining proxies are kept in the request when it is forwarded to the next proxy. An example of this is in Figure 6.3.

Examples are shown in Table 6.15.

6.2.14 Proxy-Require

The **Proxy-Require** header field is used to list features and extensions that a user agent requires a proxy to support in order to process the request. A 420 Bad Extension response is returned by the proxy listing any unsupported features in an **Unsupported** header field. Because proxies by default ignore header fields and features they do not understand, the use of a **Proxy-Require** header field is needed for the UAC to be certain that the feature is understood by the proxy. If the support of this option is desired but not required, it is listed in a **Supported** header field instead. An example is:

```
Proxy-Require: time
```

6.2.15 P-OSIP-Auth-Token

The **P-OSIP-Auth-Token** header field [15] is used to transport an Open Session Protocol (OSP) token [16] with a SIP INVITE request. A gateway or proxy server receiving a token can verify the token and use this information about accepting the INVITE or rejecting the call. This approach is suitable for a clearinghouse model of VoIP carrier interconnection.




Figure 6-3 Multiparty conferencing example.

Table 6-9
Example of Proxy-Application header field

Header Field	Meaning
<pre>Proxy-Application: CiscoSIP User-Agent: "Cisco SIP" Version: "0.0.0.0" Contact: "sip:bob@192.168.1.100" Origin: "sip:192.168.1.100" Response: "200 OK" </pre>	<p>This proxy application header field contains the metadata of Contact. Because was supplied by the SIP server located at the IP specified the response contains the original contents and protocol, no escaping is present.</p>

An example is

```
Proxy-Application: CiscoSIP
User-Agent: "Cisco SIP"
Version: "0.0.0.0"
Contact: "sip:bob@192.168.1.100"
Origin: "sip:192.168.1.100"
```

6.2.16 P-Asserted-Identity

The **P-Asserted-Identity** header field [17] is used between trusted intermediaries (provided to assert the identity of a user agent that has been authenticated using some means such as those described in Section 3.8. A UA receiving a request from a proxy that it trusts will typically render the value in a **P-Asserted-Identity** header field to the user as a "Verified Caller ID" as opposed to a From header value which is unverified. A proxy receiving a **P-Asserted-Identity** from another proxy that it does not trust will remove the header field.

An example is:

```
P-Asserted-Identity: <tel:+12125551212>
```

6.2.17 P-Preferred-Identity

The **P-Preferred-Identity** header field [17] is used by a user agent to tell a trusted intermediary which identity it would prefer be asserted on its behalf when more than one identities are associated with that user agent.

An example is:

```
P-Preferred-Identity: <tel:+12125551212>
```

6.2.18 Max-Forwards

The **Max-Forwards** header field is used to indicate the maximum number of hops that a SIP request may take. The value of the header field is decremented by each proxy that forwards the request. A proxy receiving the header field with a value of zero discards the message and sends a 483 Too Many Hops response back to the originator.

Max-Forwards is a mandatory header field in requests generated by a RFC 3261 compliant UA. However, an RFC 2543 UA generally will not include the header field field. The suggested initial value is 70 hops.

An example is:

```
max-forwards: 70
```

6.2.19 Reason

The **Reason** header field [18] can be used in BYE and CANCEL messages to indicate the reason why the session or call attempt is being terminated. It can carry a SIP response code or a Q.858 cause value (from an ISUP REL message, for example).

For example, a failing proxy could include the following header field in a CANCEL sent to a kg after one kg has answered the call:

```
Example: SIP <user@host> <uri> <call identifier> <method>
```

6.2.8 Refer-To

The **Refer-To** header field [29] is a required header field in a **REFER** request, which contains the URI or URIs resource that is being referenced. It may contain any type of URI from a *sip* or *sips* or a *tel* URI or a *mailto* or *mailto*-URI. For a *sip* or *sips* URI the URI may contain a method or escaped header fields. For example, the following **Refer-To** header field:

```
Refer-To: <tel:+12025551234; uri=sip:example.com>
sip:[escaped]example.com; uri=sip:[escaped]example.com
```

contains an escaped **Refer-To** header field. The resulting **INVITE** message generated by this **Refer-To** header field would have a **Request-URI** of *sip:[escaped]example.com* and a

```
Refer-To: <tel:+12025551234; uri=sip:example.com>
```

header field. Note that the characters *;* and *>* are replaced by their hex equivalents *%3B* and *%3E*. In the next example, the header field contains a method:

```
Refer-To: <tel:+12025551234; uri=sip:example.com>[escaped]
```

would cause a **REFER** request to be sent instead of an **INVITE**, which is the default method if none is present. An example of the **Refer-To** header field in compact form with an HTTP URI is

```
to: <http://www.example.com>
```

6.2.9 Refered-By

The **Refered-By** header field [30] is an optional header field in a **REFER** request and a request triggered by a **REFER**. It provides the recipient of a triggered request information that the request was generated as a result of a **REFER** and the originator of the **REFER**. This information can be presented to the user or have policy applied in deciding how the UA should handle the request. An assigned **Refered-By** has the form:

```
Refered-By: <http://example.com>
```

However, as this header field could be modified or fabricated, a more secure usage involves the addition of a **Refered-By** security token. The token is carried as a message body whose content id (CID) is indicated in the **Refered-By** header field. The token is an S/MIME signature over a message/*sip*/*frag*, which contains, as a minimum, the **From**, **To**, and

Call-ID, **Refer-To**, and **Refered-By** header fields from the REFER request. An example part of a REFER request is shown here:

```
Refered-By: sip:referrer@example.com
Content-Disposition: referred-to=application
Content-Type: multipart/mixed; boundary=unique-boundary-1
Content-Language: appropriate value

--unique-boundary-1

Content-Type: multipart/mixed;
 protocol=application/sipful-signature;
 attachment; boundary=diverse2
Content-ID: 41446021-49e6f8d5d5f8e2ec-0001
Content-Language: appropriate value

--DIVERS2
Content-Type: multipart/mixed
Content-Disposition: multipart-voice
From: sip:referrer@example.com
Date: Thu, 14 Feb 2002 14:21:32 GMT
Call-ID: 41446021-49e6f8d5d5f8e2ec-0001
Refer-To: sip:referred-to@example.com
Refered-By: sip:referrer@example.com
Content-ID: 41446021-49e6f8d5d5f8e2ec-0002-0001

--DIVERS2
Content-Type: application/sipful-signature; name=voice.g7
Content-Disposition: attachment; filename=voice.g7
Content-Language: appropriate value

appropriate signature goes here

--diverse2

--unique-boundary-1
```

An unassigned **Refered-By** header field may be rejected responding that the **Refered-By** entity value be included using the **229 Provide Referer Identity** response code (see Section 5.4.22). An example using the compact form is:

```
229 (229) PROVIDE-IDENTITY: SIP
```

6.2.22 Reply-To

The **Reply-To** header field is used to indicate a sip or sipuri URI, which should be used in replying to this request. Normally, this URI is present in the **From** header field (the **Contact** is not used as it is only assumed valid for the duration of the dialog). However, in some cases, the **From** cannot be populated

with this information, so the URI in this header field should be used instead of the From URI.

An example is:

```
Reply-To: <mailto:alice@lyllabs.com>
```

4.2.3 Replaces

The `Replaces` header field [11] is used in SIP call control applications. A user agent in an established dialog receiving another INVITE with a `Replaces` header field that matches the existing dialog must accept the INVITE, terminate the existing dialog with a BYE, and transfer all resources and state from the existing dialog to the newly established dialog.

If the `Replaces` header field matches no dialog, the INVITE must be rejected with a 481 Dialog Does Not Exist response.

In addition, `Replaces` has one application in pending dialogs. A UAC that has sent an INVITE but has not yet received a final response may receive an INVITE containing a `Replaces` header field that matches the pending INVITE. The UAC must terminate the pending dialog with a CANCEL (and be prepared to send a ACK and RFC if a 200 OK eventually arrives) and accept the new INVITE.

For an INVITE containing both a `Require: replaces` and `Replaces` header field, this results in the return of one of the following set of responses:

- 200 (if a match is found)
- 481 (if no match is found)
- 410 (if `Replaces` is not supported).

Figure 4.4 shows a call flow using `Replaces` to implement a feature called “call pickup.” Figure 4.7 shows the use of `Replaces` in an “assisted transfer” example.

This example `Replaces` header field:

```
Replaces: <tel:5555555555@bluewin.ch>, <tel:4444444444> from: <tel:4444444444>
```

would match the dialog identified by:

```
To: <tel:5555555555@bluewin.ch>
From: <tel:4444444444@bluewin.ch>
Call-ID: <5555555555@bluewin.ch>
```

4.2.4 Request-Contact

The `Request-Contact` [4] header field specifies the URIs in which the request may be processed; some additional parameters are also defined for



Figure 10-11 Client getting out line using proxy server

Certain header fields such as `cookie`, `referer`, and `user-agent` often used in this header field. This header field, along with `accept-encoding` and `accept-language` are part of the MP cache performance extensions. The complete form is `<code>[Example]</code>`.

```

[Example-Header]: value, value2, value3, ...
] ; *no-cache
  
```

11.29 Request-Dispatcher

The `request-dispatcher` header field is used to request a response from an other proxy or server, or further until no parallel (linking) requests. An example is

```

request-dispatcher: request
  
```

11.30 Require

The `require` header field is used to the decision and extension that a UAC requires a URI to support in order to process the request. A 404 (not found) or response is returned by the UAC listing any unsupported features

is an unsupported header field. If support or use of a feature is desirable but not required, the unsupported header field is used instead. See Table 6.3 for a list of feature tags.

An example is

```
Require: rfc3261
```

6.2.6 Response-Key

The Response-Key header field was defined in RFC 3263 but was deprecated in RFC 3261 along with all SIP-based exceptions in favor of SIPSURI encryption.

6.2.7 Route

The Route header field is used to provide routing information for responses. RFC 3261 introduces two types of routing: strict and loose routing, which have similar meanings to the IP routing modes of the same name. In strict routing, a proxy must use the first URI in the Route header field to rewrite the Request-URI, which is then forwarded. In loose routing, a proxy does not rewrite the Request-URI, but either forwards the request to the first URI in the Route header field or it may forward the request to another loose routing element. In loose routing, the request must reach every server in the Route list (but may also reach other servers) before it may be routed based on the Request-URI. In strict routing, the request must only reach the set of servers in the Route header field with the Request-URI being rewritten at each hop. A proxy or UAC can tell if the next element in the route set supports loose routing by the presence of a *Lr* parameter. An example is:

```
Route: <tel:12345678901234567890> ;Lr
```

Chapter 18 contains an example of the use of the Record-Route and Route header fields. Examples of Route header fields constructed from the example Record-Route header fields in Section 6.2.12 are:

```
Route: <url:sip:1234567890123456789012345678901234567890> ;Lr
       <tel:12345678901234567890> ;Lr
```

```
Route: <tel:12345678901234567890> ;Lr
```

6.2.8 RSeq

The RSeq header field [10] is used within a response to a PROCEED request to reliably acknowledge a provisional response that contained a CSeq header field. The RSeq header field echoes the CSeq and the RSeq from the provisional

request. The reliable sequence number is incremented for each response sent reliably. A call flow is shown in Figure 4.13. An example is:

```
3000: 2142003 2 201708
```

4.2.20 Session-Expires

The `Session-Expires` header field [21] is used to specify the expiration time of the session. To extend the session, either UA can send a re-INVITE or UPDATE with a new `Session-Expires` header field. At the expiration of the interval in the `Session-Expires` header, either UA may send a BYE and call-related parties may destroy any state information. A proxy may shorten the expiration time by reducing the interval in the header field in its proxy response. A URS confirms the session timer by including the `Session-Expires` header field in the response to the request. A URS may also shorten the interval by reducing the interval. An example is:

```
Session-Expires: 3000
```

4.2.21 Subscription-State

The `Subscription-State` header field [9] is a required header field in a NOTIFY request. It indicates the current state of the subscription. Values defined include active, pending, or terminated. Additional parameters include expires, reason, and retry-after. Values defined for the reason parameter include deactivating, giveup, preconditions, nonrecursion, rejected, and timeout.

An example is:

```
Subscription-State: terminated;reason=terminated
```

4.3 Response Header Fields

These header fields are present only in responses.

4.3.1 Authentication-Info

The `Authentication-Info` header field can be inserted in responses when performing mutual authentication using HTTP Digest. In mutual HTTP Digest as described in Section 3.8, the server challenges the client to provide a shared secret, which the client then provides in a request of the request containing an `Authorization` or `WWW-Authenticate` header field. To do mutual authentication, the server would then provide an `Authentication-Info` header field containing either a new nonce or a

response is an `opaque` parameter. The response auth digest is calculated by the server on the SIP response using the same algorithm as the successful request authentication and using the same shared secret (client's username and password). In this way, the server proves that it also knows the client's secret, providing mutual authentication. The credentials are carried in the `opaque` parameter in the header field.

An example is:

```
Authorization: Digest username="40403@jabber.linux.it",opaque=""
```

4.3.2 Error-Info

The `Error-Info` header field is used in failure response to convey more information about an error. A UAC receiving the header field in a failure response may fetch and render the URI to the user. The header field can be used to give the client the details of how the error can be presented to the user. For example, a client with a graphical interface will likely display the reason phrase on the response, which should provide very specific information about the failure. However, an audio-only UA does not have this capability (although a text-to-speech synthesizer could be used to provide this capability). Instead, an audio-only UA could fetch the URI and play the resulting audio stream to the user.

If the URI is a `http` or `rtsp` URI, the UA may treat the `Error-Info` as a `Contact` in a redirection response, which would result in a SIP session established to play the recording.

An example is:

```
Error-Info: http://www.example.com/recordings/1234
```

4.3.3 Min-Expires

The `Min-Expires` header field is used in a `423 Interval Too Small` response (Section 5.4.20) from a registrar rejecting a `REGISTER` request in which one or more `Contacts` have an expiration time that is too short. The header field contains an integer number of seconds that represents the minimum expiration interval that the registrar will accept. A client receiving this header field can update the expiration intervals of the registration request accordingly and resend the `REGISTER` request.

An example is:

```
Min-Expires: 1200
```

4.3.4 Min-SE

The `Min-SE` header field [11] is a required header field in a `423 Session Timer Interval Too Small` response (Section 5.4.19). The response

may also be present in an `INVITE` or `UPDATE` containing a `Session-Expires` header field. It contains an integer number of seconds.

An example is

```
Session-Expires: 400
```

6.25 Proxy-Authenticate

The `Proxy-Authenticate` header field is used in a `407 Proxy-Authenticate` or `401 Required authentication challenge` by a proxy server to a UAC. It contains the name of the challenge so that the UAC may formulate credentials in a `Proxy-Authorization` header field in a subsequent request. Examples are shown in Table 6.16.

6.26 Server

The `Server` header field is used to convey information about the URS generating the response. The use and content of the header field are similar to the `User-Agent` header field in Section 6.1.36. An example is

```
Server: Apache/2.0.46 (Ubuntu)
```

6.27 Unsupported

The `Unsupported` header field is used to indicate features that are not supported by the server. The header field is used in a `421 Bad Extension` response to a request containing an unsupported feature listed in a `Require` header field. Because multiple features may have been listed in the `Require` header field, the `Unsupported` header field indicates all the unsupported features—the rest can be assumed by the UAC to be supported. See Table 6.8 for a list of feature tags.

An example is

```
Unsupported: foo,bar
```

Table 6.16
Example of Proxy-Authenticate Header Field

Header Field	Meaning
Proxy-Authenticate: Digest realm="example.com", nonce="8c48f6c2d1614e54f2e107061d7692c6", opaque=""	HTTP digest challenge header field

6.1.8 Warning

The `Warning` header field is used in a response to provide more specific information than the response code alone can convey. The header field contains a three-digit warning code, a warning agent that indicates what server issued the header field, and warning text enclosed in quotes used for display purposes. Warning codes in the 1000 and 2000 range are specific to HTTP [5]. The SIP standard defines 12 new warning codes in the 3000 class. The breakdown of the class is shown in Table 6.17. The complete set of defined warning codes is listed in Table 6.18.

Examples are

```
Warning: 300 proxy "Incompatible transport protocol"
Warning: 300 00000000000000000000 "Unauthorized access denied"
```

6.1.9 WWW-Authenticate

The `WWW-Authenticate` header field is used in a 401. Unauthorized authentication challenge by a user agent or registrar server to a UAC. It contains the name of the challenge or that the UAC may formulate credentials in a `Proxy-Authorization` header field in a subsequent request. SIP supports HTTP digest authentication mechanisms. Examples are shown in Table 6.19.

6.1.10 Req

The `Req` header field [10] is used in provisional (2xx class) responses to 300/370s to request reliable transport. The header field may only be used if the 300/370 response contained the `supported`: `rel1000` header field. If present in a provisional response, the UAC should acknowledge receipt of the

Table 6.17
SIP Warning Codes

Warning Code Range	Error Type
300, 310, 320	SIP keywords
300	Network-related
300, 310, 340	Reserved for future use
370	URI parameter
300	Reserved
300	Miscellaneous

6.4 Message Body Header Fields

These header fields contain information about the message body or body.

6.4.1 Allow

The `Allow` header field is used to indicate the methods supported by the user agent or proxy server sending the response. The header field must be present in a 405 Method Not Allowed response and should be included in a positive response to an OPTIONS request. An example is:

```
Allow: DELETE, ACK, BYE, INFO, OPTIONS, CANCEL
```

6.4.2 Content-Encoding

The `Content-Encoding` header field is used to indicate that the listed encoding scheme has been applied to the message body. This allows the URS to determine the decoding scheme necessary to interpret the message body. Multiple listings in this header field indicate that multiple encodings have been used in the sequence in which they are listed. Only encoding schemes listed in an `Allow-Encoding` header field may be used. The compact form is `n-Fram: plus include`.

```
Content-Encoding: text/html  
n-1234
```

6.4.3 Content-Disposition

The `Content-Disposition` header field is used to describe the function of a message body. Defined values include `attachment`, `icon`, `alert`, and `transfer`. The value `attachment` indicates that the message body contains information to describe a media session. The value `transfer` indicates that the message body should be displayed or otherwise rendered for the user. If a message body is present in a request or a 2xx response without a `Content-Disposition`, the function is assumed to be `transfer`. For all other response classes with message bodies, the default function is `transfer`. An example is:

```
Content-Disposition: attachment
```

6.4.4 Content-Language

The `Content-Language` header field [2] is used to indicate the language of a message body. It contains a language tag, which identifies the language.

```
Content-Language: en
```

8.4.3 Content-Length

The `Content-Length` is used to indicate the number of octets in the message body. A `Content-Length: 0` indicates no message body. As described in Section 2.4.2, this header field is used to separate multiple messages sent within a TCP stream. If not present in a UDP message, the message body is assumed to continue to the end of the datagram. If not present in a TCP message, the message body is assumed to continue until the connection is closed. The `Content-Length` octet count does not include the CR/LF that separates the message header fields from the message body. It does, however, include the CR/LF at the end of each line of the message body. An example octet calculation is in Chapter 2. The `Content-Length` header field is not a required header field to allow dynamically generated message bodies, where the `Content-Length` may not be known a priori. The compact form is 2. Examples include:

```
Content-Length: 0
Content-Length: 1, 207
```

8.4.4 Content-Type

The `Content-Type` header field is used to specify the Internet media type (IMT) in the message body. Media types have the headline form `type/subtype`. If this header field is not present, `application/octet-stream` is assumed. If an `Accept` header field was present in the request, the response `Content-Type` must contain a listed type, or a 415 `Unsupported Media Type` response must be returned. The compact form is `v`. Specific MIME types that are commonly used are listed in Table 8.20.

`Content-Disposition` [26] can be used to provide a URI in place of actual MIME message body. An example is:

Table 8.20
Content-Types Present in HTTP Requests and Responses

Content-Type	Use
<code>application/javascript</code>	HTTP in HTML, XML, or XHTML requests [22]
<code>application/javascript-javscript</code>	HTTP request in SCRIPT in older implementations [22]
<code>application/javascript-javascript</code>	HTML script [26]
<code>application/javascript-javascript+xml</code>	XML content in HTML [26]
<code>application/javascript-javascript+xml</code>	XML [26]
<code>text/css</code>	Plain text
<code>text/html</code>	HTML text
<code>application/javascript-javascript+xml</code>	Encapsulated HTTP in HTML, XML, or XHTML [27]


```
Content-Type: message/rfc822-headers; charset="UTF-8"
Host: example.com *
```

The compact form is as follows: Example.com

```
example.com;type=application/whp
m: foo@bar.com
```

6.4.7 Expires

The Expires header field is used to indicate the time interval in which the request or message contents are valid. When present in an **HTTP/1.1** request, the header field sets a time limit on the completion of the **HTTP/1.1** request. That is, the UAC must receive a final response (non-2xx) within the time period as the **HTTP/1.1** request is successfully associated with a 408 Request Timeout response. Once the session is established, the value from the Expires header field in the original **HTTP/1.1** has no effect—the **Session-Expires** header field (Section 6.2.18) must be used for this purpose. When present in a **REGISTER** request, the header field sets the time limit as the **URL** in **Contact** header fields that do not contain an expires parameter. Table 4.3 shows examples of the Expires header field in registration requests. The header field is not defined for any other method types. The header field may contain a **MF** date or a number of seconds. Examples include:

```
Expires: 00
Expires: Fri, 26 Apr 2008 00:00:00 GMT
```

6.4.8 MIME-Version

The **MIME-Version** header field is used to indicate the version of **MIME** protocol used to construct the message body. **MF**, like **HTTP**, is not considered **MIME**-compliant because parsing and semantics are defined by the **MF** standard, not the **MIME** Specification [29]. Version 1.0 is the default value. An example is:

```
MIME-Version: 1.0
```

References

- [1] Rosenberg, J., et al. "SIP: Session Initiation Protocol." RFC 3261, 2002.
- [2] Fielding, R., et al. "Hypertext Transfer Protocol — HTTP/1.1." RFC 2616, June 1999.
- [3] Bach, A. "Session Initiation Protocol (SIP)-Specific Event Notifications." RFC 3853, 2005.

- [4] Rosenberg, J., H. Schulzrinne, and P. Kyzivat, "Call Performance and Call Capabilities for the Session Initiation Protocol (SIP)," IETF Internet-Draft, Work in Progress, March 2003.
- [5] Johnson, A., and G. Larkin, "Session Initiation Protocol (SIP) Call Control - Conferencing for User Agents," IETF Internet-Draft, Work in Progress, April 2003.
- [6] Rosenberg, J., and H. Schulzrinne, "A Session Initiation Protocol (SIP) Event Package for Conference State," IETF Internet-Draft, Work in Progress, June 2003.
- [7] Braden, R., "Requirements for Internet Hosts, Applications and Support," RFC 1123, 1989.
- [8] Maly, R., and D. Peate, "The Session Initiation Protocol (SIP) 'Join' Header," IETF Internet-Draft, Work in Progress, March 2003.
- [9] Willis, D., and B. Horowitz, "Session Initiation Protocol (SIP) Extension: Header Field for Registering Non-Subsequent Contacts," RFC 3347, 2003.
- [10] Rosenberg, J., and H. Schulzrinne, "Reliability of Provisional Responses in Session Initiation Protocol (SIP)," RFC 3363, 2003.
- [11] Maly, R., B. Stigs, and R. Doss, "The Session Initiation Protocol (SIP) 'Replace' Header," IETF Internet-Draft, Work in Progress, March 2003.
- [12] Dromi, S., and J. Rosenberg, "Session Timers in the Session Initiation Protocol (SIP)," IETF Internet-Draft, Work in Progress, November 2002.
- [13] Ford, J., "Media Type Registration Procedures," RFC 4969, 2006.
- [14] Peterson, L., "A Proxy Mechanism for the Session Initiation Protocol (SIP)," RFC 3325, 2003.
- [15] Johnson, A., et al., "Session Initiation Protocol (SIP) Extension for an SIP Authorization Token," IETF Internet-Draft, Work in Progress, February 2003.
- [16] European Telecommunications Standards Institute, "Telecommunications and Internet Protocol Homologation Over Networks (TIPHON), Open Session Protocol (OSP) for Inter-domain Pricing, Authentication, and Usage Exchange," Technical Specification TS 34.11, Series 3.1.0.
- [17] Jennings, C., J. Peterson, and M. Watson, "Private Extensions to the Session Initiation Protocol (SIP) for Anonymity within Traversal Networks," RFC 5025, 2007.
- [18] Schulzrinne, H., D. Doss, and G. Casella, "The Reason Header Field for the Session Initiation Protocol (SIP)," RFC 3326, 2003.
- [19] Sparks, R., "The Session Initiation Protocol (SIP) Info Method," RFC 3503, 2003.
- [20] Sparks, R., "The SIP Refered-by Mechanism," IETF Internet-Draft, Work in Progress, February 2003.
- [21] Dromi, S., and Rosenberg, J., "Session Timers in the Session Initiation Protocol (SIP)," IETF Internet-Draft, Work in Progress, November 2002.
- [22] Handley, M., and V. Jacobson, "SIP: Session Description Protocol," RFC 3261, 1998.
- [23] Sparks, R., "Internet Media Type:multipart/mixing," RFC 3426, 2003.
- [24] Rosenberg, J., and H. Schulzrinne, "An (SIP) Event Package for the Session Initiation Protocol (SIP)," IETF Internet-Draft, Work in Progress, March 2003.
- [25] Rosenberg, J., and H. Schulzrinne, "A Session Initiation Protocol (SIP) Event Package for Conference State," IETF Internet-Draft, Work in Progress, June 2003.

- [16] Aggarwal, H., et al., "Classroom Presence and Instant Messaging (CPIM) Presence Information Data Format," IETF Instant-Data, Work in Progress, 2002.
- [17] Zinnerman, E., et al., "MIME Media Types for SIP and QSIG Objects," RFC 3284, December 2001.
- [18] Chen, J., "A Mechanism for Contact Reflection in Session Initiation Protocol (SIP) Messages," IETF Instant-Data, Work in Progress, February 2005.
- [19] Ford, M., and M. Hennecke, "Multipurpose Internet Mail Extension (MIME). Part One: Format of Internet Message Bodies," RFC-1521, 1996.

7

Related Protocols

The Session Initiation Protocol (SIP) is one part of the protocol suite that makes up the Internet Multimedia Conferencing architecture, as shown in Figure 1.1. In this chapter, other related Internet protocols mentioned or referenced in other sections are introduced, along with details on the use of the protocol with SIP. This is by no means a complete discussion of multimedia communication protocols over the Internet. First, SDP, the media description language, will be discussed. Then the RTP and RTCP media transport protocols will be discussed. The application of RTP/AVP profiles that link SIP and RTP will then be discussed. The chapter concludes with a brief discussion of signaling protocols in the PSTN. The H.323 protocol will be discussed and compared to SIP in the next chapter. The chapter concludes with a discussion of SIP-T and UTPP protocols.

7.1 SDP—Session Description Protocol

The Session Description Protocol, defined by RFC 2327 [1], was developed by the IETF MMUSIC working group. It is more of a descriptive syntax than a protocol in that it does not provide a full-range media negotiation capability. The original purpose of SDP was to describe multimedia sessions set up over the Internet's multimedia backbone, the MBONE. The first application of SDP was by the experimental Session Announcement Protocol (SAP) [2] used to post and receive announcements of MBONE sessions. SAP messages carry a SDP message body, and was the template for SIP's use of SDP. Even though it was

designed for multicast, SDP has been applied to the more general problem of describing general multimedia sessions established using SIP.

As seen in the examples of Chapter 3, SDP contains the following information about the media session:

- IP Address (IPv4 address or host name)
- Port number (used by UDP or TCP for transport)
- Media type (audio, video, interactive whiteboard, and so forth)
- Media encoding scheme (PCM A-Law, MPEG II video, and so forth)

In addition, SDP contains information about the following:

- Subject of the session
- Start and stop times
- Contact information about the session

Like SIP, SDP uses text encoding. An SDP message is composed of a series of lines, called fields, whose names are abbreviated by a single lower-case letter, and are in a required order to simplify parsing. The set of mandatory SDP fields is shown in Table 2.1. The complete set is shown in Table 7.1.

SDP was not designed to be easily extensible, and parsing rules are strict. The only way to extend or add new capabilities to SDP is to define a new attribute type. However, unknown attribute types can be silently ignored. A SDP parser must not ignore an unknown field, a missing mandatory field, or an out-of-sequence line. An example SDP message containing many of the optional fields is shown here:

```

v=0
i=SESSION DESCRIPTION INFORMATION IN SIP 41.76.2.2
m=0 0 audio
l=TEXT INFORMATION WILL COVER THIS NEW SIP MESSAGE
a=range:0-1000000000,0-1000000000
a=range:0-1000000000,0-1000000000
p=0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
c=IN IP4 10.10.10.10
a=range:0-1000000000,0-1000000000
a=range:0 0-1000000000
a=range:0 0-1000000000
a=range:0 0-1000000000

```

The general form of a SDP message is:

```

parameter1 parameter1 ... parameterN

```

Table 11
SDP field codes (their required order)

Field	Name	Mandatory/Optional
v=	Protocol version number	m
o=	Sender/owner and session identifier	m
s=	Service name	m
i=	Service information	o
p=	Session/Resource identifier	o
l=	Email address	o
g=	Phone number	o
c=	Connection information	o
b=	Bandwidth information	o
t=	Time session starts and stops	o
m=	Support lines	o
r=	Time zone conversion	o
k=	Encryption key	o
z=	Attributes list	o
u=	Media information	o
a=	Media attributes	o

The line begins with a single lower-case letter *v*. There are never any spaces between the letter and the *v*, and there is exactly one space between each parameter. Each field has a defined number of parameters. Each line ends with a CRLF. The individual fields will now be discussed in detail.

2.5.1 Protocol Version

The *v=* field contains the SDP version number. Because the current version of SDP is 0, a valid SDP message will always begin with *v=0*.

2.5.2 Origin

The *o=* field contains information about the originator of the session and session identifier. This field is used to uniquely identify the session. The field contains

```

<network-type> <address> <type> <connection-address>
<address>

```

The `<network-type>` parameter contains the originator's logic or host or - if none. The `<connection-id>` parameter is a Network Time Protocol (NTP) [1] timestamp or a random number used to ensure uniqueness. The `<version>` field is a numeric field that is increased for each change to the session, also recommended to be a NTP timestamp. The `<network-type>` is always IS for Internet. The `<address-type>` parameter is either IP4 or IP6 for IPv4 or IPv6 address either in dotted decimal form or a fully qualified host name.

2.2.3 Session Name and Information

The `<n>` field contains a name for the session. It can contain any non-zero number of characters. The optional `<id>` field contains information about the session. It can contain any number of characters.

2.2.4 URI

The optional `<u>` field contains a uniform resource indicator (URI) with more information about the session.

2.2.5 E-Mail Address and Phone Number

The optional `` field contains an e-mail address of the host of the session. If a display name is used, the e-mail address is enclosed in `<n>`. The optional `<pn>` field contains a phone number. The phone number should be given in globalized format, beginning with a +, then the country code, a space or -, then the local number. Either space or - are permitted as spaces in SDP. A comment may be present in `<D>`.

2.2.6 Connection Data

The `<cn>` field contains information about the media connection. The field contains

```

<network-type> <address> <type> <connection-address>

```

The `<network-type>` parameter is defined as IS for the Internet. The `<address-type>` is defined as IP4 for IPv4 addresses, IP6 for IPv6 addresses. The `<connection-address>` is the IP address that will be sending the media packets, which could be either multicast or unicast. If multicast, the `<connection-address>` field contains

```

base-addr = address-base - base - multi-addr - multi-addr - multi-addr - id -
  address

```

where *id* is the identifier value, and *multi-addr* may or may not indicate how many anonymous multi-addr addresses are included starting with the *base-multi-addr*.

3.3.7 Bandwidth

The optional *bw* field contains information about the bandwidth required. It is of the form:

```

bw = modifier : bandwidth - value

```

The *modifier* is either *CT* for conference total or *AS* for application specific. *CT* is used for multi-addr sessions to specify the total bandwidth that can be used by all participants in the session. *AS* is used to specify the bandwidth of a single *addr*. The *bandwidth-value* parameter is the specified number of kilobytes per second.

3.3.8 Time, Repeat Times, and Time Zones

The *tz* field contains the start time and stop time of the session.

```

tz = start-time / stop-time

```

The times are specified using NTP timestamps. For a scheduled session, a *stop-time* of *zero* indicates that the session goes on indefinitely. A *start-time* and *stop-time* of *zero* for a scheduled session indicates that it is permanent. The optional *rr* field contains information about the repeat times that can be specified in either in NTP or in days(0), hours(0), or minutes (0). The optional *dst* field contains information about the time zone offset. This field is used if a recurring session spans a change from daylight-savings to standard time, or vice versa.

3.3.9 Encryption Keys

The optional *key* field contains the encryption key to be used for the media session. The field contains

```

key = method / secret / id - key

```

The *method* parameter can be *clear*, *base64*, *url*, or *prompt*. If the method is *prompt*, the key will not be carried in SDP; instead, the user will be

prompted as they join the encrypted session. Otherwise, the key is sent in the `encrypt=on-key` parameter.

7.1.8 Media Assessments

The optional `media` field contains information about the type of media session. The field contains

```
media=audio/transport-format-list
```

The `media` parameter is either `audio`, `video`, `application`, `data`, `telephone-event`, or `control`. The `port` parameter contains the port number. The `transport` parameter contains the transport protocol, which is either RTP/RTCP or scp. (RTP/RTCP stands for Real-time Transport Protocol [4] / audio video profiles [5], which is described in Section 7.1.) The `format-list` contains more information about the media. Usually, it contains media payload types defined in RTP audio/video profiles. More than one media payload type can be listed, allowing multiple alternative codes for the media session. For example, the following `media` field lists three codes:

```
media=audio RTP/AVP 0 1 3 98
```

One of these three codes can be used for the media session. If the intention is to establish three media channels, three separate `media` fields would be used. For non-RTP media, Internet media types should be listed in the `format-list`. For example,

```
media=application/x-rtsp video-rtsp
```

could be used to specify the application/sub-media type.

7.1.9 Attributes

The optional `attrs` field contains attributes of the preceding media session. This field can be used to extend SDP to provide more information about the media. If not fully understood by a SDP user, the `attributes` field can be ignored. There can be one or more `attributes` fields for each media payload type listed in the `media` field. For the RTP/RTCP example in Section 7.1.8, the following three `attributes` fields could follow the `media` field:

```
attribute: 0 PCM/16000
attribute: 1 PCM/12000
attribute: 3 PCM/8000
attribute: 98 AAC
```

Other attributes are shown in Table 7.2. Full details of the use of these attributes are in the standard document [3]. The details of the RBC (Internet Low Bit Rate Codec) are in [8].

7.3.12 Use of SDP in SIP

The use of SDP with SIP is given in the SDP Offer Answer RFC 3264 [7]. The default message body type in SIP is `application/sdp`. The calling party lists the media capabilities that they are willing to receive in SDP in either an `INVITE` or in an `ACR`. The called party lists their media capabilities in the `200 OK` response to the `INVITE`. More generally, offers or answers may be in `INVITEs`, `PRACKs`, or `UPDATEs` or in reliably sent `180s` or `250s` responses to these methods.

Because SDP was developed with scheduled multimedial sessions in mind, many of the fields have little or no meaning in the context of dynamic sessions established using SIP. In order to maintain compatibility with the SDP protocol, however, all required fields are included. A typical SIP use of SDP includes the version, origin, subject, time, connection, and one or more media and attribute fields as shown in Table 7.1. The origin, subject, and time fields are not used by SIP but are included for compatibility. In the SDP standard, the subject field is a required field and must contain at least one character, suggested to be `xyz` if there is no subject. The time field is usually set to `0x0 0`.

SIP uses the connection, media, and attribute fields to set up sessions between user agents. Because the type of media session and codes to be used are part of the connection negotiation, SIP can use SDP to specify multiple alternative media types and to selectively accept or decline those media types. When multiple media codes are listed, the caller and called party's media fields must be aligned—that is, there must be the same number, and they must be listed in the same order. The offer answer specification, RFC 3264 [7], recommends that an attribute containing `accepttype` be used for each media field [7]. A media session is declined by setting the port number to zero for the corresponding media field in the SDP response. In the following example, the caller Tedd wants to set up an audio and video call with two possible audio codecs and a video codec in the SDP carried in the initial `INVITE`:

```

>>>
V=0
O=TEDD 192.168.1.100 5060 5060 IN IP4 192.168.1.100-5060
S=
C=0
M=audio 49170 RTP/AVP 0 4
a=accepttype:0 H261/0x00
a=accepttype:0 H263/0x00

```

Table 12:
SIP Attributes/Values

Attribute	Value
ac/lenap	POSTCAP list
a-cat	Category of the session
ak/ep/ak	Endpoint of session
a-to	Name of tool used to create SIP
as/line	Length of line in milliseconds for each packet
a-ec/only	Facsimile-only mode
a-ec/none	Text and media mode
a-ec/only	Text-only mode
as/ed	Orientation for redifined session
a-type	Type of conference
as/ed	Cluster and media subject and a formation field
as/ling	Language for the session-description
as/ing	Default language for the session
a-frame/size	Maximum video frame size in frames per second
as/quality	Suggests quality of encoding
a-keep	Format transport
as/nd	Media identification-grouping
as/direction	Direction for symmetric mode
as/tp	Explicit RTP port (and address)
as/rtptime	Initial RTT

© 2005 IETF RFC 4923

© 2005 IETF RFC 4923

The values are referenced by the RTP/RTCP profile numbers 4, 5, and 12. The called party Muxed answers the call, chooses the second codec for the first media field and declines the second media field, only wanting a PCM A-law media session.

© 2005

© 2005 IETF RFC 4923

© 2005

© 2005 IETF RFC 4923

```

c=0-0
m=audio 44444 RTP/AVP 0
s=123456789 123456789
m=video 0 RTP/AVP 31

```

If this *audioonly* call is not acceptable, then Tada would send an *ACK* then a *BYE* to cancel the call. Otherwise, the audio session would be established and RTP packets exchanged. As this example illustrates, unless the number and order of media fields is maintained, the calling party would not know for certain which media sessions were being accepted and declined by the called party.

One party in a call can temporarily place the other on hold (i.e., suspending the media packet sending). This is done by sending an *INVITE* with identical *SDP* to that of the original *INVITE* but with an *answeronly* attribute present. The call is made active again by sending another *INVITE* with the *answeronly* attribute present. (Note that older RFC 2543-compliant UAs may initiate hold using *c=0, 0, 0, 0*.) For further examples of *SDP* use with SIP, see the *SIP Offer Answer Examples* document [5].

7.2 RTP—Real-Time Transport Protocol

Real-Time Transport Protocol [6] was developed to enable the transport of real-time packets (such as log files, video, or other information) over IP. RTP is defined by RTP Proposed Standard RFC 3550 (which updates RFC 1889). RTP does not provide any quality of service over the IP network—RTP packets are handled the same as all other packets in an IP network. However, RTP allows for the detection of some of the impairments introduced by an IP network, such as:

- Packet loss
- Variable transport delay
- Out-of-sequence packet arrival
- Asymmetric routing

As shown in the protocol stack of Figure 1.1, RTP is an application layer protocol that uses UDP for transport over IP. RTP is not text encoded, but uses a 64-bit-oriented header similar to UDP and IP. RTP version 0 is only used by the net audio tool for H.323E benchmarks. Version 1 was a pre-RFC implementation and is not in use. The current RTP version 2 packet header is shown in Figure 7.1. RTP was designed to be very general; most of the headers are only loosely defined in the standard; the details are left to profile documents. The 12 bytes are defined as:



Figure 2.8 RTP packet header

- **Version (V):** This 2-bit field is set to 2, the current version of RTP.
- **Padding (P):** If this bit is set, there are padding bytes added to the end of the packet to make the packet a fixed length. This is done automatically and if the media stream is encrypted.
- **Extension (X):** If this bit is set, there is one additional extension following the header (giving a total header length of 14 octets). Extensions are defined by certain payload types.
- **CSRC count (CC):** This 4-bit field contains the number of contrib source identifiers (CSRC) are present following the header. This field is only used by mixers that take multiple RTP streams and output a single RTP stream.
- **Marker (M):** This single bit is used to indicate the start of a new frame in video, or the start of a talkspurt in silence-suppressed speech.
- **Payload Type (PT):** This 7-bit field defines the codec in use. The value of this field matches the profile number listed in the SDP.
- **Sequence Number:** This 32-bit field is incremented for each RTP packet sent and is used to detect missing/lost or sequence packets.
- **Timestamp:** This 32-bit field indicates in relative terms the time when the payload was sampled. This field allows the receiver to remove jitter and to play back the packets at the right interval assuming sufficient buffering.
- **Synchronization Source Identifier (SSRC):** This 32-bit field identifies the sender of the RTP packet. At the start of a session, each participant chooses a SSRC number randomly. Should two participants choose the same number, they each choose again until each party is unique.
- **CSRC (Contributing Source Identifier):** There can be some or up to 15 instances of this 32-bit field in the header. The number is set by the CSRC Count (CC) header field. This field is only present if the RTP packet is being sent by a mixer, which has received RTP packets from a number of sources and sends out combined packets. A non-mixing conference bridge would utilize this header.

RTP allows detection of a lost packet by a gap in the Sequence Number. Packets received out of sequence can be detected by out-of-sequence

disruptive situations. Note that RTP allows detection of these transport-related problems but leaves it up to the codec to deal with the problem. For example, a video codec may compensate for the loss of a packet by repeating the last video frame, while an audio codec may play background noise for the interval. Variable delay or jitter can be detected by the `Timestamp` field. A continuous bit rate codec such as PCM will have a linearly increasing `Timestamp`. A variable bit rate codec, however, which sends packets at irregular intervals, will have an irregularly increasing `Timestamp`, which can be used to play back the packets at the correct interval.

The RTP Control Protocol (RTCP) is a related protocol also defined in RFC 3550 that allows participants in an RTP session to send each other quality reports and statistics, and exchange some basic identity information. The four types of RTCP packets are shown in Table 7.3. RTCP has been designed to scale to very large conferences. Because RTCP traffic is all overhead, the bandwidth allocated to these messages remains fixed regardless of the number of participants. That is, the more participants on a conference, the less frequently RTCP packets are sent. For example, in a basic two-participant audio RTP session, the RTP/AVP profile states that RTCP packets are to be sent about every 5 seconds for four participants. RTCP packets can be sent every 18 seconds. Sender reports (SR) or receiver reports (RR) packets are sent the most frequently, with the other packet types being sent less frequently. The use of reports allows feedback on the quality of the connection, including information such as

- Number of packets sent and received
- Number of packets lost
- Packet jitter

Table 7.3
RTP Packet Types

Packet Type	Name	Description
SR	Sender report	Sent by a participant that both sends and receives RTP packets
RR	Receiver report	Sent by a participant that only receives RTP packets
SDP	Session description	Carries information about the participant in the session, including e-mail address, phone number, and host
BYE	Bye	Sent to terminate the RTP session
APP	Application specific	Reflects a particular profile
RRTP	Extended report	Extended report and summary

In a multimedia session established with SIP, the information needed to select codecs and send the RTP packets to the right location is carried in the SDP message body. Under some scenarios, it can be desirable to change codecs during an RTP session. An example of this relates to the transport of dual tone multiple frequency (DTMF) digits. A user who uses codecs that is optimized for transmitting vocal sounds will not transport the superimposed sine waves of a DTMF signal without introducing significant noise, which may cause the DTMF digit receiver to fail to detect the digit. As a result, it is useful to switch to another codec when the sender detects a DTMF tone. Because a RTP packet contains the payload type, it is possible to change codecs “on the fly” without any signaling information being exchanged between the user agents. On the other hand, switching codecs in general should probably not be done without a SIP signaling exchange (re-INVITE) because the call could fail if one side switches to a codec that the other does not support. The SIP re-INVITE message exchange allows this change in media session parameters to fail without causing the established session to fail.

The use of random numbers for CSRC provides a minimal amount of security against “media spoofing” where a literally malicious third party tries to break into a media session by sending RTP packets during an established call. Unless the third party can guess the CSRC of the branched sender, the receiver will discard a change in CSRC number and either ignore the packets or inform the user that something is going on. This behavior for RTP streams, however, is not universally accepted, because in some scenarios (wireless hand-off, announcement server, call center, and so forth) it might be desirable to send media from multiple sources during the progress of a call.

RTP supports encryption of the media. In addition, RTP can use IPsec [8] for authentication and encryption.

1.3 RTP Audio Video Profiles

The use of profiles enables RTP to be an extremely general media transport protocol. The current audio video profiles defined by RFC 3551 are listed in Table 7.4. The profile document makes the following specifications for RTP:

- UDP is used for underlying transport
- RTP port numbers are always even, the corresponding RTCP port number is the next highest port, always an odd number
- No header extensions are used

For each of the profiles listed in Table 7.4, the profile document lists details of the codec, or a reference for the details is provided. Payloads in the

Table 7.4
H.264/AVC Audio-Visual Profile Types

Profile	Codec	Clock	Description
0	PCM1	9000	ITU G.711 PCM μ-Law 8000Hz 8Kbps
1	HE1	9000	HEP Audio-4 8Kbps
2	CE1	9000	ITU BT1.4 ADPCM-Audio 32 Kbps
3	CE2	9000	European-CEM-Audio 32 Kbps
4	CE4	9000	ITU ADPCM-Audio 32 Kbps
6	CE4	9000	ITU ADPCM 8Kbps
7	LPC	9000	Experimental LPC Audio
8	PCM4	9000	ITU G.711 PCM μ-Law Audio-64 Kbps
9	CE3	9000	ITU G.723 Audio
10	LC	44100	Linear PCM-Audio 3855 Kbps
11	LC	44100	Linear PCM-Audio Stereo-Audio 14112 Kbps
14	HE4	9000	MPEG-4a MPEG-4 Audio-Only
15	CE5	9000	ITU G.728 Audio 16 Kbps
20	CE6	9000	16K-Audio
24	J8C	9000	J8C-Audio
26	H1	9000	H1-Audio
31	H21	9000	ITU G.201 Audio
32	HE5	9000	MPEG-4a and MPEG-4 Video
33	HE6	9000	MPEG-4 transport stream-video
dynamic	HEC	—	Internal low-bit rate H.264 (H)
dynamic	AVC	—	Adaptive Multi-Rate Codec (A)

range 36–127 can be defined dynamically during a session. The minimum profile support is defined as 0 (PCM1) and 5 (HE1). The document recommends dynamically assigned profile numbers, although 9004 and 9005 have been registered for use of the profile and can be used instead. The standard also describes the process of engineering new profile types with DANA. There are other references for a detailed description of many of these audio-codes [18] and video-codes [14].

The information in the first three columns of Table 7.4 is also contained in the SDP `avpt=` field, which is why the attribute is optional.

2.4 PSTN Protocols

Three types of PSTN signaling protocols are mentioned in this text: Circuit Associated Signaling (CAS), ISDN (Integrated Services Digital Network), and ISUP (ISDN User Part). They will be briefly introduced and explained. How these protocols work in the PSTN today are covered in other references [10].

2.4.1 Circuit-Associated Signaling

This type of signaling is the oldest, currently used in the PSTN today. The signaling information goes for some miles circuit to the voice path, with digits and characters represented by audio tones. These are the tones that used to be barely discernible on long-distance calls before ring tone is heard. The tones are called multi-frequency (MF) tones. They are similar to the tones used to signal between a telephone and a central office switch, which use DTMF tones. Long post dial delay is introduced because of the time taken to convert long strings of digits. Also, CAS is susceptible to fraud, as fraudulent tones can be generated by the caller to make free telephone calls. This type of signaling is common in trunk circuits between a central office and a corporation's private branch exchange (PBX) switch.

2.4.2 ISUP Signaling

ISDN User Part is the protocol used between telephone switches in the PSTN for call signaling. It is used over a dedicated packet-switched network that uses Signaling System #7 (SS7) for transport. This signaling method was developed to overcome some of the delay and security problems with CAS. There are examples of ISUP signaling in the call flow example of Chapter 10. The adoption of this out-of-band signaling protocol was the first step taken by telecommunications carriers away from circuit-switched networks and towards packet-switched networks. The final step will be moving the bearer channels onto a packet-switched network.

2.4.3 ISDN Signaling

Integrated Services Digital Network (ISDN) signaling was developed for all-digital telephone connections on the PSTN. The most common types of interfaces are the basic rate interface (BRI) and the primary rate interface (PRI). A BRI can contain two 64-Kbps B-channels for either voice or data and a 16-Kbps D-channel for signaling. BRI can be used over conventional telephone lines but requires an ISDN telephone or terminal adapter. PRI uses a 1.544-Mbps link called a T-1 or a DS-1, which is divided up into 23 B-channels and one D-channel, with each channel being 64 Kbps. The H.323 protocol, described in

Chapter 8, treats a subset of the EBN Q.931 signaling protocol used over the D-channel.

3.5 SIP for Telephones

SIP for Telephones (SIP-T) is a framework for SIP interworking with the PSTN [13]. It includes two approaches: translation and encapsulation. Both approaches will be discussed.

Translation is the direct mapping between PSTN protocols and SIP. The mapping between various PSTN protocols such as ISUP [14], Q.931 [15], and others have been defined. Examples of SIP interworking with PSTN protocols including EBN and CAS are in the SIP/PSTN Call Flow document [16]. In this approach, as much of the information that is common to each protocol are mapped between them, with the remaining values being set to configurable defaults. In this approach, a SIP call from a PSTN gateway is indistinguishable from a SIP call from a native device, and is handled such by the protocol. However, since not every single parameter in a PSTN signaling message has a counterpart (or has any meaning) in SIP, some information is lost if the call routes back to a PSTN termination point.

Encapsulation is another approach that is only useful for SIP/PSTN gateways. Using this approach, first, PSTN-to-SIP translation is done to construct the appropriate SIP message, then the PSTN protocol message is encapsulated and included with the SIP message as a message body. If the SIP message is received by another SIP/PSTN gateway, the resulting PSTN signaling message is reconstructed from both the SIP message and the encapsulated PSTN message that was received by the other gateway. This approach offers the possibility of transparency (i.e., no loss of PSTN information as a call is carried across a SIP network). However, this only works in a network in which only one variation of PSTN protocol is used. Unlike Internet protocols, PSTN protocols vary by region and are not compatible without a special type of PSTN outside, which converts one message format to another. There are many dozens of protocol variants used throughout the world.

Another disadvantage of encapsulation is that if the PSTN message bodies must be encrypted if they are transported over the public Internet, or used in a network with native SIP devices. This is because private information can be carried in PSTN messages because PSTN protocols assume a different trust model than an Internet protocol such as SIP. To prevent accidental disclosure of this information, the message bodies must be encrypted by the originating gateway and decrypted by the terminating gateway, which adds significant processing requirements and call setup delay.

Encapsulated PSTN messages are carried as MIME bodies, which have been standardized for both ISUP and Q.931 [17].

3.8 Universal Plug and Play Protocol

The Universal Plug and Play Protocol (UPnP) [18] is a protocol developed to allow configuration and peer-to-peer networking of intelligent IP devices. UPnP works together with DHCP or Auto IP to configure and set up devices in small networks typical of home and small office networks [19]. The protocol has six basic functions: addressing, discovery, description, control, events, and presentation.

Some SIP clients use UPnP to gain configuration information and to talk to firewalls and NATs that are UPnP enabled to open firewalls and learn private/public IP address mappings. For example, if a router or wireless bridge that encompasses a firewall and NAT function is UPnP enabled, an authorized SIP client also UPnP enabled can automatically manage NAT and firewall traversal.

In the future, proponents of UPnP see it playing a key role in home network configurations where multiple intelligent IP devices are connected in a home or small office. However, the protocol is unlikely to be used in studios and large enterprises.

UPnP references a number of RTP standards such as DHCP, HTTP Multicast UDP (HTTPMU), and Auto IP. It also uses W3C protocols such as XML and SOAP.

References

- [1] Handley, M., and V. Jacobson, "SIP Session Description Protocol," RFC 2543, 1998.
- [2] Handley, M., C. Perkins, and E. Weibel, "Session Announcement Protocol," RFC 2974, 2000.
- [3] Mill, D., "Network Time Protocol Version 3: Specification, Implementation, and Analysis," RFC 1883, 1995.
- [4] Schulzrinne, H., et al., "RTP: A Transport Protocol for Real-Time Applications," RFC 1889, 1995.
- [5] Schulzrinne, H., "RTP Profile for Audio and Video Conferences with Minimal Control," RFC 2633, 1999.
- [6] Daini, A., and S. Anderson, "RTP Payload Format for G.723 Speech," IETF Internet-Draft, Work in Progress, March 2003.
- [7] Rosenberg, J., and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)," RFC 2634, 1999.
- [8] Johnston, S., and R. Sparks, "SIP Offer Answer Example," IETF Internet-Draft, Work in Progress, June 2003.
- [9] Kent, S., and B. Atkinson, "Security Architecture for the Internet Protocol," RFC 2401, 1998.

-
- [10] Anandaram, T., *Introduction to Telecommunication Network Engineering*, Norwood, MA, Artech House, 1999.
 - [11] Schiphorst, R., *Videoconferencing and Videoconferencing Technology and Standards*, 2nd ed., Norwood, MA, Artech House, 1999.
 - [12] Seifberg, L., et al., "Real-Time Transport Protocol Hybrid Format and File Storage Format for the Adaptive Multi-Rate (AMR) and Adaptive Multi-Rate Wideband (AMR-WB) Audio Codes," RFC 3267, June 2002.
 - [13] Yamani, A., and J. Peterson, "Session Initiation Protocol for Telephones (SIP-T): Context and Architecture," RFC 3261, 2002.
 - [14] Camarillo, G., et al., "Integrated Services Digital Network (ISDN) User Part (ISUP) to Session Initiation Protocol (SIP) Mapping," RFC 3268, 2002.
 - [15] Hird, J., et al., "Interworking between SIP and QSIG," IETF Internet-Draft, Work in Progress, April 2003.
 - [16] Johnson, A., et al., "Session Initiation Protocol (SIP) Call Flow," IETF Internet-Draft, Work in Progress, April 2003.
 - [17] Zammit, L., et al., "MIME Media Types for ISUP and QSIG Objects," RFC 3294, December 2002.
 - [18] SIP specifications are available at <http://www.ietf.org>.
 - [19] Schmidt, K. F., "Devices That Stay Together, Work Together," *IEEE Magazine*, September 2003, pp. 65-70.

8

Comparison to H.323

This chapter compares SIP to another IP telephony signaling protocol: the ITU recommendation H.323, entitled “Packet-Based Multimedia Communications.” H.323 is introduced and explained using a simple call flow example. H.323 and SIP are then compared.

8.1 Introduction to H.323

H.323 [1] is an umbrella recommendation that covers all aspects of multimedia communication over packet networks. It is part of the H.32x suite¹ of protocols that describes multimedia communication over BDN, broadband ATM,² telephone (PSTN), and packet (IP) networks, as shown in Table 8-1. Originally developed for video conferencing over a single LAN segment, the protocol has been extended to cover the general problem of telephony over the Internet. The first version was approved by the ITU in 1996 and was adopted by early IP telephony networks because there were no other standards. Version 2 was adopted in 1998 to fix some of the problems and limitations in version 1. Version 3 was adopted in 1999 and includes modifications and extensions to enable

1. In this chapter, the use of an *x* instead of a digit does not imply that all digits (0–9) in the range are included. In this case H.32x does not include H.321 to H.329, which have yet to be defined.
2. In this context, broadband means transported over an asynchronous transfer mode network. In anticipation of the universal deployment of ATM networks by carriers, the ITU developed a suite of protocols to support conventional telephony over ATM networks. For example, Q.2931 is the extension of Q.931 BDN over ATM. Today, the term is used to mean high bandwidth connections—faster than modem speeds.

Table 8.1
It's a little family of protocols

Protocol	Title
H.320	Communication-over GDN networks
H.321	Communication-over broadband GDN(MM) networks
H.322	Communication-over ATM with guaranteed QoS
H.323	Communication-over ATM with non-prioritized QoS (IP)
H.324	Communication-over PSTN (TDM nodes)

communications over a larger network. Version 4 was adopted in 2000 with some major changes to the protocol. Version 5 is currently under revision as the ITU-T. Currently, most deployed systems use Version 2. H.323 has been designed to be backward compatible, so, for example, a version 1 compliant end point can communicate with a version 2 gatekeeper and a version 4 end point.

H.323 references a number of other ITU and IETF protocols to completely specify the requirements. Each element of the network is defined and standardized. Figure 8.1 shows the main elements: terminals, gatekeepers, gateways, and multi-point control units (MCUs). Terminals, gateways, and MCUs are network end-devices, often called end points. An end point originates and terminates media streams that could be audio, video, or data, or a combination



Figure 8.1 Elements of an H.323 network.

of all data. As a minimum, all H.323 end points must support basic G.711 PCM audio transmission. Support of video and data are optional. An H.323 gatekeeper is a server that controls a zone, which is the smallest administrative domain in H.323. If a gatekeeper is present, all end points within that zone must register with and refer to the gatekeeper on authorization decisions to place or accept a call. A gatekeeper also provides services to terminals in a zone, such as gateway location, address translation, bandwidth management, format implementation, and registration. A gatekeeper is not a required element in an H.323 network, but a terminal's capabilities without one are severely limited. A gateway is another optional element in an H.323 network. It interfaces the H.323 network with another protocol network, such as the PSTN. An MCU provides conferencing services for terminals.

Some of the protocols referenced by H.323 are shown in Table 8.2. H.225 is used for registration, admission, and status (RAS), which is used for terminal-to-gatekeeper communication. A modified subset of Q.931 is used for call setup signaling between terminals. (The H.323 usage of Q.931 is not compatible with Q.931 as used in an ISDN network.) H.245 is used for control signaling, or media negotiation and capability exchange between terminals. T.120 is used for multipoint graphic communications. H.323 audio codecs are specified in the ITU-G.7xx series. Video codecs are specified in the H.26x series. H.323 also references two RTP protocols, RTP and RTCP, for the media transport, which are described in Sections 7.2 and 7.3. The H.235 recommendation covers privacy and encryption, while H.450 covers supplementary services such as those commonly found in the PSTN (e.g., call forwarding, call hold, and call park).

Table 8.2

Protocols Referenced by H.323

Protocol	Description
H.225	Registration, admission, and status (RAS) protocol signaling
H.245	Control signaling (media control)
T.120	Multipoint graphic communications
G.7xx	Audio codecs
H.26x	Video codecs
RTP	Real-time transport protocol (RFC 3550)
RTCP	RTP control protocol (RFC 3550)
H.235	Privacy and encryption
H.450	Supplementary services

8.2 Example of H.323

Figure 8.2 shows a basic call flow (involving two terminals and a gatekeeper). The flow shows the interactions between the various elements and the various protocols used to establish the session. The call begins with an exchange of H.225.0 RAS messages between the calling terminal and the gatekeeper. All RAS messages are transported using UDP. It is assumed that both terminals have already registered with the gatekeeper using the Registration Request (RRQ) message. The calling terminal sends an Admission Request (ARQ) message to the gatekeeper containing the address of the called terminal and the type of session desired. The address could be specified as an H.323 alias, E.164 telephone number, e-mail address, or a URL. The gatekeeper knows about all calls in the zone it controls; it decides if the user is authorized to make a call and if there is enough bandwidth or other resources available. In this example, there is enough bandwidth, so the




Figure 8.2 H.323 call flow example.

gatekeeper allows the call to continue by sending an Admission Confirmation (ACF) message. The ACF indicates to the calling terminal that end-point message routing, or the direct exchange of H.225 call signaling messages with the called terminal, is to be used. Alternatively, the gatekeeper can require gatekeeper relayed signaling, where the gatekeeper acts like a proxy and forwards all signaling messages between the terminals. The gatekeeper has also translated the destination in the SIP into a transport address that was returned in the ACF.

The calling terminal is now able to open a TCP connection to the called terminal using the transport address returned in the ACF and send a Q.931 SETUP message to the called terminal. The called terminal responds with a Call Processing response to the calling terminal. The called terminal must also get permission from the gatekeeper before it accepts the call, so an RQ is sent to the gatekeeper. When it receives the ACF from the gatekeeper, the called terminal begins sending the user, and sends an ALERTING message to the calling terminal. When the user at the calling terminal answers, a CONNECT message is sent. There is no acknowledgment of messages because all these messages are sent using TCP, which provides reliable transport. These call signaling messages used in H.323 are a subset of the Q.931 recommendations that cover EBN D-channel signaling.

The next stage is the connection negotiation, which is handled by H.245 control signaling messages. A second TCP connection between the two terminals is opened by the calling terminal using the port number selected by the called terminal and returned in the CONNECT message. The TerminalCapabilityTypeSet message sent contains the media capabilities of the calling terminal, listing supported codecs. It is acknowledged with a TerminalCapabilitySetAck response from the called terminal. The called terminal also sends a TerminalCapabilityTypeSet message containing its media capabilities, which receives a TerminalCapabilityTypeSetAck response.

The H.323 protocol requires that one terminal be selected as the master with the other as the slave. This is accomplished using HostCapabilitiesExchange messages exchanged between the terminals. The message contains the terminal type of the terminal and a random number. Terminal types are hierarchical, which determines the master. If the terminal type is the same, the random number determines the master. The message is acknowledged with a HostCapabilitiesExchangeAck message.

The final phase of the call setup is the opening of two logical channels between the terminals. These channels are used to set up and control the media channel. The H.245 OpenLogicalChannels message sent in the H.245 control signaling connection contains the desired media type, including the codec that has been determined from the exchange of capabilities. It also contains a pair of addresses for the RTP and RTCP streams. The message is acknowledged with an OpenLogicalChannelsAck message.

Next, the terminals begin sending RTP media packets and also RTCP control packets using the IP addresses and port numbers exchanged in the OpenLocalAckChannel message.

Figure 8.5 shows a call teardown sequence, which either terminal may initiate. In this example, the called terminal sends an EndSessionCommand message in the H.245 control signaling channel. The other terminal responds with an EndSessionAcknowledge message in the H.245 control signaling channel, which can now be closed. The called terminal then sends a DisconnectRequest (DRQ) message and receives a DisconnectConfirmation (DCP) message from the gatekeeper. This way, the gatekeeper knows that the resources used in the call have now been freed up. A Call Detail Record (CDR) or other billing record can be written and stored by the gatekeeper. Next, a Q.931 Release Complete message is sent in the call signaling connection, which can then be closed. Finally, the other terminal sends a DRQ to the gatekeeper over UDP and receives a DCP response.

The call flows in Figures 8.2 and 8.3 show direct end-points signaling, where the calling terminal opens TCP connections to the called terminal and




Figure 8.5 H.323 call tear-down sequence.

exchanges H.225-R and H.245 messages. In the RCF response to the calling terminal, the gatekeeper can accept gatekeeper control signaling, where the call signaling and control signaling channels are opened with the gatekeeper, who then opens the channels with the called terminal. In this way, the gatekeeper stays in the signaling path and provides all signaling messages. This allows the gatekeeper to know the exact call state and be able to provide features.

8.3 Versions

There are four versions of H.323, which reflect the evolution of this protocol with a fifth underway in 2003. H.323 is fully backward compatible, so gatekeepers and terminals must support firms and mechanisms defined in all previous versions. Version 1 was approved in 1996 and was titled "Visual Telephony Systems over Networks with Non-Guaranteed Quality of Service." The example call flow in Figure 8.2 shows the version 1 call setup. Not unexpectedly given the number of messages and TCP connections, this process was very slow, sometimes taking as many as seven round trips to establish a call. While this may have been acceptable for a protocol designed for video conferencing over a single LAN segment, it is not acceptable for an IP telephony network designed to provide a similar level of service to the PSTN.

Version 2 included alternative call setup schemes to speed up the call setup. Two schemes were added to H.323, called FastStart and H.245 tunneling. FastStart is shown in Figure 8.3, in which the GetSetup message contains the TerminalCapabilitiesList byte information. This saves multiple messages and round trips. In H.245 tunneling, a separate H.245 control channel is not opened. Instead, H.245 messages are encapsulated in Q.931 messages in the call signaling channel. This saves overhead in opening and closing a second TCP connection.

Versions 3 and 4 added more features to H.323 and additional extensions. Of interest to the present discussion is the support for H.323 URLs [2], full UDP support instead of TCP, and also the standardization of the use of DNS by H.323 in Annex G. Version 5 is currently being worked on by the IPU-T and is scheduled to be finished by late 2003.

8.4 Comparison

SIP and H.323 were developed for different purposes by standards bodies with very different requirements. H.323 was developed by the IPU. Its design and implementation reflects its PSTN background and heritage, utilizing binary encoding and reusing parts of ISDN signaling. SIP, on the other hand, was




Figure 8.8 Partial Call Connections with H.323

developed by the IETF with an Internet perspective, designed to be scalable over the Internet and work in an interdomain way utilizing the full set of Internet utilities and functions.

While H.323 was deployed in early VoIP and IP videoconferencing applications, SIP with its Internet architecture is gaining traction now and is emerging as the future signaling standard for IP communications, or IP telephony is sometimes called.

8.1.1 Fundamental Differences

The first key difference is in the encoding scheme used by the protocol. SIP is a text-based protocol like HTTP and SMTP, while H.323 uses binary-encoded ASN.1 messages. The binary encoding of H.323 may result in smaller message sizes but it adds complexity to implementations. A text-based protocol such as SIP can be easily scripted and requires no tools to monitor and inspect messages—a simple packet dump from a LAN provides the ASCII-encoded SIP messages, which can be easily logged, examined, and inspected. Simple applications can be written to test and simulate SIP traffic. The text-based encoding of

SIP has made it seem more friendly to Internet and Web developers who have embraced it in developing various applications.

Another important difference is that while H.323 is exclusively a signaling protocol, SIP has both presence and instant message capability. The combination of presence and signaling in one protocol using a common universal addressing scheme, URIs, will be a very powerful driver in new applications in the future. This makes SIP an extremely powerful “non-domain” protocol that allows a user with multiple mobile endpoints to locate and communicate with another user with multiple capabilities. It is for this reason that the most exciting and innovative services are being developed with SIP rather than H.323.

Another important difference is in key vendor support and momentum. SIP has been adopted by some key players in both the PC and telecommunication industries. It is widely regarded in the industry as the signaling protocol for service establishments over IP. While it may not immediately replace many established H.323 networks used for basic services such as simple phone calls, it is being adopted or investigated by all major players in the industry, even those who currently have a vested interest in legacy H.323 systems.

SIP has also been adopted by mobile operators as the call signaling and instant message protocol for their third generation networks currently under development. This offers the promise of millions of SIP-enabled wireless devices in the next few years. SIP is also being coupled with 802.11 wireless networks for another set of mobile services.

Another important difference is the level of security in the protocol. SIP as defined in RFC 3261 has very robust security mechanisms to provide encryption, authentication using certificates, and end-to-end message integrity even in the presence of untrusted intermediary servers. SIP did not need to develop these security features because, since it is an Internet protocol, it was able to inherit the rich set of Internet security protocols such as TLS and S/MIME in very straightforward way. For example, the same security mechanisms that make it safe to enter credit card information on a secure Web page form is what allows SIP to provide secure signaling between servers.

While it is not directly related to SIP, the media capability/negotiation capabilities of SDP are quite different to that provided by H.263 for H.323. H.263 is often criticized for its complexity, but SDP is often correctly criticized for its lack of expressiveness. The ambiguous legacy of SDP has not provided SIP with a good basis for media negotiation. As a result, some vendor/media negotiation capabilities interwork poorly between implementations due to these problems. Much industry work has gone into fixing the problem, but a replacement for SDP is not yet within sight yet.

Over the years that they have coexisted, SIP and H.323 have actually become more similar in some functionality. This is natural in that each protocol seems to have adopted some of the useful features of each. For example,

while SIP had DNS and URL support right from the start, H.323 began with none but has added some support for them. Another similarity relates to conferencing. While H.323 had the concept of an MCU right from the start, SIP initially had no defined conferencing mechanisms. However, work is currently underway to define the SIP signaling equivalent, called a “focus” as described in Section 11.1.2. Additional work is underway in the IETF to standardize the media mixing functions of an MCU.

In other ways, the protocols started out in opposite ends of the spectrum and have moved towards each other. For example, SIP was initially deployed exclusively with UDP, but TCP support has grown and become more important over the years. H.323 on the other hand initially could not be used over UDP exclusively, but now has been extended to be used over UDP. Another example relates to the number of messages used to set up sessions. H.323 started with a large number of messages, then reduced that number as it evolved using FastStart. Some SIP applications that use extensions such as UPDATE and PRACK to perform precondition negotiation prior to session establishment now use many more messages to set up a session than the three in the original specification. (However, this complexity in these SIP precondition applications may prove that they, like H.323, will never achieve wide deployment, and will instead be displaced by the simpler SIP model.) Finally, H.323 started with a larger specification than SIP [2], but as SIP has grown and added functionality, it now rates the “paperweight” use in terms of more pages of specification text. As a result, most SIP versus H.323 comparisons are badly out of date due to the evolution in each protocol.

11.2 Strengths of Each Protocol

H.323 has carved out two niche areas in current deployed systems—it is widely deployed in small PSTN replacement networks for handling simple phone calls, and it dominates the IP teleconferencing market. Simple PSTN replacement networks that only originate and terminate phone calls without even basic features do not use most of the key advantages of SIP. As a result, they have little incentive to upgrade to SIP and the widespread adoption of SIP eventually makes SIP gateway ports much cheaper than H.323 ports. However, since SIP is earlier in its development cycle than H.323, this is not likely to happen for a number of years. New implementations of these types of systems will likely deploy SIP from the start, seeking to “future proof” the investment, but there is little incentive for deployed systems to upgrade. Also, the availability of some minimal SIP to H.323 signaling gateways allows both networks to work together and complex-calls [4, 5].

The teleconferencing market is dominated by PSTN EBN devices, which have been deployed since 1980. Since H.323 shares the same heritage, it

was designed to easily interwork with these existing systems. SIP was not designed to interwork with these systems and was also not designed specifically with video in mind. In fact, some key signaling components needed to do video conferencing are still not fully standardized in the IETF, although there is active work to finalize these standards. As a result, a fully functional, standards-based SIP videoconferencing system is not currently available (but will be soon). However, the deployment of videoconferencing systems has been disappointing compared to many predictions of this useful technology. This largely has to do with bandwidth and cost factors. Now that bandwidths and camera costs are falling, SIP is well positioned to handle much simpler videoconferencing systems enabled by PCs with cheap Webcams. It is likely that SIP will play a bigger role in these future software-based video systems than the current dedicated hardware videoconferencing systems. In this way, SIP may in the future be a major player in the videoconferencing market but without ever displacing the current installed base.

The major strength of SIP is that it is still relatively simple—it is an Internet protocol based on the Internet's architecture. While complex architectures have been developed using a multitude of SIP extensions, it is likely that these systems will be the exception rather than the rule. Most SIP end points will only need to support the base SIP specification and perhaps a few call control extensions to provide a wide variety of useful and innovative services.

8.5 Conclusion

This chapter has introduced H.323 and discussed and compared it to SIP. While there are some similarities between the protocols in call setup, and some niche markets that H.323 currently dominates, SIP, with its text encoding, presence and instant message extensions, and Internet architecture, is poised to be the signaling and “underneath” protocol of choice for Internet devices in the future.

References

- [1] “Text-Based Multimedia Communications System,” IETF Recommendation H.323, 2000.
- [2] Levin, G., “H.323 Unified Resource Locator (URL) Scheme Registration,” RFC 2898, 2000.
- [3] Rosenberg, J., and H. Schulzrinne, “A Comparison of SIP and H.323 for Internet Telephony,” *Research and Operating System Support for Signal Audio and Video (ROSSAV)*, Cambridge, England, July 1998.

- [4] S. Adami, M., and C. Agosti, "Session Initiation Protocol (SIP)-H.323 Interworking Requirements," IETF Internet-Draft, Work in Progress, 2003.
- [5] Agosti, C., "Session Initiation Protocol (SIP) - H.323 Interworking," IETF Internet-Draft, Work in Progress, 2005.

9

Wireless and 3GPP

The mobility features of SIP have been discussed in earlier chapters. In this chapter, these aspects will be explored further. In 2000 the Third Generation Partnership Project [1] adopted SIP as their call signaling protocol starting with Intelligent Multimedia Core Subsystem (IMS) Release 5. Since then, a number of extensions and usage documents have been authored describing their planned use of SIP; these will be discussed in this chapter. The future direction of wireless SIP will be discussed in the final section of this chapter.

9.1 IP Mobility

There are a number of different types of mobility that will be discussed in this chapter, which include terminal mobility, personal mobility, and service mobility [2]. Terminal mobility is the ability of an end device to maintain its connection to the Internet as it moves around and possibly changes its point of connection. Personal mobility is the ability to have a constant address (identify) across a number of devices. Finally, service mobility is the ability of a user to keep the same services when mobile.

Terminal mobility can be achieved by Mobile IP [3], which has been standardized in the IETF. It allows a terminal to keep the same IP address when roaming as it does in its home network. While roaming, the terminal is reachable by a "care of" address, which is registered in the home network. Packets destined for the roaming terminal are routed in the home network, then tunneled to the terminal at the "care of" address, as shown in Figure 9.1. Mobile IP has the advantage of hiding the mobile nature of the terminal from layer 3 protocols and above. These protocols can then be used without any modification.



Figure 1: Triangle routing of IP packets in Mobile IP

For example, a TCP connection can be maintained since the remote application knows constant IP address. However, Mobile IP has the disadvantage that incoming packets are not routed directly, and as a result, more efficiently. This results in increased packet latency. While this is not a problem for constant data streams such as Web browsing or e-mail, real-time media streams have strict requirements on packet latency. A solution that has been proposed [1] over the fact that approaches such as IP already has mobility support built in. In addition, IP is capable of handling some of the essential mobility aspects at the application layer. This can result in more efficient RTP packet routing and lower efficiency (in the additional overhead of IP packet encapsulation required by Mobile IP are avoided).

The result is that mobile IP devices can utilize two different architectures. One is based on the use of Mobile IP and the other utilizes the built-in mobility support in RTP. The next sections will discuss these two approaches. As discussed in the next section, SIP is ideally suited to provide both personal and voice mobility.

9.2 SIP Mobility

Personal mobility is the ability to have a constant identifier across a number of devices. A unique origin URI has exactly this property and is fundamentally supported by SIP. SIP can also support voice mobility, the ability of a user to keep the same services when mobile although some extensions and extensions have been proposed that provide this in certain architectures.

Basic personal mobility is supported by SIP using the REGISTER method, which allows a mobile device to change its IP address and point-of-connection to the Internet and still be able to receive incoming calls. As discussed in Chapters 2 and 4, registration in SIP temporarily binds a user's AOR URI with a Contact URI of a particular device. As a device's IP address changes, registration allows this information to be automatically updated in the SIP network. An end device can also move between service providers using multiple layers of registrations, in which a registration is actually performed with a Contact as an address of record with another service provider. For example, consider the user agent in Figure 9.2, which has temporarily acquired a new SIP URI with a new service provider. (The reasons for doing so could include security, local policy, NAT/firewall traversal.) The user agent then performs a double registration as shown in Figure 9.2. The first registration is with the new service provider, which binds the Contact URI of the device with the new service provider's AOR URI. The second REGISTER request is routed back to the original service provider and provides the new service provider's AOR as the Contact URI. As shown here in the call flow, when a request comes in to the original service provider's network, the INVITE is redirected to the new service provider who then routes the call to the user.

For the first registration, message containing the device URI would be:

```
REGISTER sip:capetown.com.au SIP/2.0
Via: SIP/2.0/UDP 192.3.3.1:5060;branch=capetown1111
Max-Forwards: 70
To: SIP:192.3.3.1:5060@capetown.com.au
From: SIP:192.3.3.1:5060@capetown.com.au
Call-ID: 192.3.3.1-19-07-14-00000
CSeq: 1 REGISTER
Contact: sip:192.3.3.1:5060
```

and the second registration message with the roaming URI would be:

```
REGISTER sip:capetown.com.au SIP/2.0
Via: SIP/2.0/UDP 192.3.3.1:5060;branch=capetown1111
Max-Forwards: 70
To: SIP:192.3.3.1:5060@capetown.com.au
From: SIP:192.3.3.1:5060@capetown.com.au
Call-ID: 192.3.3.1-19-07-14-00000
CSeq: 1 REGISTER
Contact: sip:192.3.3.1:5060@capetown.com.au
```

The first INVITE that is depicted in Figure 9.2 would be sent to sip:192.3.3.1:5060@capetown.com.au; the second INVITE would be sent to sip:192.3.3.1:5060@capetown.com.au; and the third INVITE would be




Figure 1.1 For understanding the IP-Understanding the Local Access Process.

was to allow the user to be notified of the user's location and allow the user to be notified of the user's location and allow the user to be notified of the user's location.

A disadvantage of this approach is that the user would have to be notified of the user's location and allow the user to be notified of the user's location. This would have to be done using some IP-based method such as a Web page sign-up, which would be coupled with the proper authentication, authorization, and accounting mechanisms.

An optimization of this is to be implemented to forward the registration information to the existing user agent back to the Home system. This has been proposed in the HTTP [4] but has yet to be adopted or standardized. It

changes to IP messages are required, just a correction adopted by engineers to recognize a naming convention and take the appropriate action. It is possible that these corrections may become standardized if the authentication and accounting systems needed to properly process such registrations are standardized in the future.

During sessions, a mobile device may also change IP address as it switches between one wireless network and another (the Mobile IP protocol is not covered—it will be discussed in the next section). Mobile IP supports this scenario as well, since an ICDTTC in a wireless LAN can be used to replace the Coarse-grain DRI and disaggregate information in the MIP. This is shown in the call flow of Figure 5.1. Here, the mobile device is on a wireless network, and MIP registers a new IP address, then performs a re-CDTTC to update the signaling and media flow to the new IP address. If the user agent successfully is able to




Figure 5.1 This flexibility using non-CDTTC.

receive media from both networks, the interruption can be almost negligible. If this is not the case, a few RTP packets may be lost as the media catches up with the signaling, resulting in a slight interruption to the call. The re-INVITE would appear as follows:

```

INVITE sip:laplace@bell.com;url=urn:uuid:1234567890 SIP/2.0
Via: SIP/2.0/UDP 10.10.10.1:5060;branch=1234567890
Call-ID: 1234567890
To: Laplace de Laplace <laplace@bell.com>;tag=1234567890
From: BellSouth PowerCo <bell@powerco.com>;ip=10.10.10.1;
Call-ID: 1234567890;branch=1234567890
CSeq: 2 INVITE
Contact: <sip:bell@10.10.10.1>
Content-Type: application/sdp
Content-Length: 100

v=0
o=bell 1234567890 1234567890 IN IP4 10.10.10.1
s=bell
c=IN IP4 10.10.10.1
t=0 0
m=audio 1234567890/0

```

which contains BellSouth's new IP address in the Via, Content-Length fields and SDP media information.

Note that both of the mobility scenarios in Figure 9.1 and 9.3 do not require cooperation between the two wireless networks. As such, this is a useful scenario in which a user agent can hand off a call between, for example, a commercial wireless network and a home or office 802.11 wireless network.

For a handoff mobility in which the actual route out (not SIP provider) that the SIP messages must traverse must change, a re-INVITE cannot be used. For example, if a proxy is necessary for NAT/traversal, then more than just the Contact URI must be changed—a new dialog must be created. The solution to this is to send a new INVITE (which creates a new dialog and a new route out) including the new set of proxies with a Replaces header (Section 6.2.2), which identifies the existing session. The call flow is shown in Figure 9.4. It is similar to that in Figure 9.3 except that a BYE is automatically generated to terminate the existing dialog when the INVITE with the Replaces is accepted. In this scenario, the existing dialog between BellSouth and Laplace includes the old visited proxy server (this proxy Record-Route'd during the initial INVITE). The new dialog using the new wireless network requires the inclusion of the new visited proxy server. As a result, an INVITE with Replaces is sent by BellSouth, which creates a new dialog that includes the new visited proxy server (which Record-Route'd but not the old visited proxy server). When Laplace accepts the INVITE, a BYE is automatically sent to terminate the old dialog that routes through the old visited proxy server that is




Figure 6-14 Illustrating NAT translation for a mobile service.

are no longer involved in the session. The resulting mobile session is established using Facebook's new IP address from the NAT in the LTE/4G.

Services in NAT can be provided in either position or in user space. If the service is resident in the user space, then there are no service mobility problems in the user space context. However, combining service mobility and personal mobility can be challenging unless each of the user's devices are identically configured with the same services. Also, multi-point resident services are only available when the end point is connected to the Internet. A terminating service such as a call forwarding service implemented in an endpoint will fail if the endpoint has temporarily lost its Internet connection. For this and other reasons, some services are implemented in the network using NAT proxy services. For these

services, service mobility for a user agent means that the same set of proxies are used to route incoming and outgoing requests when mobile.

Due to the nature of the Internet, in general, there is no reason why a user agent cannot use the same proxies when connected to the Internet at a different point. That is, a user agent that is normally in the United States that is configured to use a set of proxies in the United States can still use those proxies when roaming in Europe, for example. Perhaps the SIP hops will have a slightly higher latency due to more router hops and a call setup request may take a second or two longer to complete. However, this has no impact on the quality of the media stream as the media always flows directly between the two user agents and does not traverse the SIP proxy servers. As a result, SIP can easily support service mobility over the Internet.

However, there are some cases in which this service mobility approach will not work. For example, if a local proxy server must be traversed in order to facilitate firewall or NAT traversal, or for some other security reason, then a user agent may have to use a different first-hop proxy when roaming. In this case, service mobility is still possible provided that:

1. The roaming user agent is able to discover the necessary local proxy.
2. Both incoming and outgoing requests are routed through the home proxy in addition to any local proxies.

The first requirement is met by the DNSCP extension to SIP [5], which allows a user agent to learn of a local proxy server at the same time it learns its IP address and other IP configuration information. The second requirement is met using a "preloaded" Route header field in requests. Normally a Route header is learned in a request when a proxy requests it using a Record-Route header field. However, it is possible for a configured user agent to include a Route header field. If the Route header contains the URI of the home proxy, the request will be routed to the home proxy after the local proxies have been traversed, meeting the requirement for outgoing requests. For incoming requests, the double registration technique will result in both the home and local proxies being traversed by incoming requests. This will result in a call flow similar to Figure 9-2 but with the home proxy server forwarding the SIP/RTCP instead of reflecting.

These SIP mobility capabilities are well suited to use over a wireless network such as 802.11 in a home, office, or public space. As roaming agreements allow such wireless "hotspots" to be linked up in metropolitan areas, this will provide a wireless service. However, commercial wireless providers plan a specific-purpose wireless telephony network using SIP. For some of their business and service requirements, some SIP extensions have been developed, which will be discussed in this chapter.

These IP clients may also make use of voice codecs such as the G.723.1, which is highly robust to packet loss, which may be experienced in a heavily loaded WAN, for example.

5.1 40PP Architecture and IP

The 40PP architecture uses IP in the IP Multimedia Core Network Subsystem (IMS). A simplified architecture is shown in Figure 5-1. The elements of the IMS architecture are listed in Table 5-1. The specifications of IMS in 40PP is described in 3GPP TS 26.134 and in [7].



Figure 5-1: 40PP IMS architecture.

Table 5-1
IMS Elements

Abbreviation	Name
UE	User Equipment
I-CSCF	Interrogating - Call Session Control Function
S-CSCF	Home - Call Session Control Function
PC	Proxy Call Control
ASCS	Application Server/Call Session Control Function
ASCS	Application Server
AS	Application Server
ASCS	Application Server/Call Session Control Function
ASCS	Application Server/Call Session Control Function
ASCS	Application Server

The 3GPP architecture relies on Mobile IP instead of some of the mobility aspects of SIP described in the previous section. The reasons for doing so are primarily business related rather than technical in nature.

Another requirement of mobility systems is a keep-alive signal, which allows end points and proxies to know that a user agent still has network connectivity. On an end point to end point basis, this can be done using RTCP (see Section 7.2) reports sent periodically, even when the media is on hold or silence suppression is taking place. However, proxies do not have access to these direct end-to-end reports. Instead, the Session Timer extension [9] and re-INVITEs can be used for this purpose.

CSCF use SIP proxies that also sometimes behave as a B2BUA under certain circumstances. For example, if a P-CSCF loses the radio link to the UE (user equipment) that contains the SIP UA, it can send a RYE on behalf of the UE to tear down the session. The motivation for doing this is for networks that have a per-minute billing charge, which the out-of-contact UE would otherwise have to pay for but not have the ability to disconnect. To save bandwidth on the wireless connection, a P-CSCF removes Route, Record-Route, Path, Via, Extension-Sequence and other header fields and replaces them in the opposite direction. To prevent high bandwidth codes from being used by a UE, a P-CSCF may add the list of codes in an SDP offer or answer, preventing the codes from being used. A P-CSCF may change the To and From headers to provide privacy, which is a B2BUA function.

The Proxy CSCF provides emergency service, triggers for local services, and does telephone number normalization for the rest of the network. The P-CSCF is used as the default outbound proxy server for a UE within its home network. The Interrogating CSCF queries the HSS to determine the proper service CSCF. The ICSCF also does hiding of the S-CSCF network by removing or encrypting Via header fields. The Serving CSCF provides the services for the subscriber. It identifies the user's service profile and privileges.

The 3GPP plans to exclusively use IPv6 addresses. This is due to the number of SIP subscribers envisioned and the fact that with Mobile IP, each device may use more than one IP address at a time. SIP RFC [26] includes full support for IPv6 addresses, and an extension to SDP [8] adds IPv6 support to SDP.

The 3GPP also uses signaling compression [10] to co-encode SIP messages transmitted over a wireless link. This is primarily done to minimize latency rather than bandwidth savings. The use of signaling compression with SIP is described in [11], which defines a parameter `comp=1` (`comp`) that can be used in Via header fields and as a URI parameter that can be used, for example, in a Contact header field.

The 3GPP uses the adaptive multiplex (AMR) [12] codec for the audio encoding.

5.4 NGFP Header Fields

Some SIP header fields have been developed based on NGFP requirements and are discussed in the next sections. They are listed here instead of the main listing of SIP header fields in Chapter 6 since their use is not well standardized outside the NGFP architecture. In the future, however, their use may be standardized within the IETF and they may be used more widely.

5.4.1 Service-Route

The `Service-Route` header field [13] can be used in a `200 OK` response to a `REGISTER` request. It can be used by a registrar server to provide to the registering UA URIs to include in a `preferred-Route` header field in future requests. The `Service-Route` URIs are only valid for the duration of the registration and should be updated when the registration is refreshed.

An example is:

```
Service-Route: <tel:20202020202020202020202020202020>
```

5.4.2 Path

The `Path` header field [14] is an optional header field in `REGISTER` requests. It can be thought of as a `Record-Route` mechanism for `REGISTER` requests, which establishes a route set that is valid for the duration of the registration. The `Path` header field may be inserted by a proxy, which forwards a `REGISTER` request to a registrar server. The registrar copies the `Path` header field into the `200 OK` response to the `REGISTER`, which then provides the route set information to the user agent that is registering. In a mobile network, the `Path` header field can be used to discover and inform the user agent of the proxies that can be used to populate `preferred-Route` header fields.

An example is:

```
Path: <http://proxy1.example.com/>, <http://proxy2.example.com/>
```

5.4.3 Other P-Headers

In addition, some headers that are specific to NGFP have been defined [15]. These so-called P-Headers (which stands for proprietary, preliminary, or private) are defined in *optional* mode in an informational RFC per the SIP change process [16]. Examples of the use of these P-Headers are given in the call flow of Section 5.7. They are listed in Table 5.2.

Table 6.1
SIP Headers

Header Field	Use
P-Branch: 0.0.0.0-192	Indicates SIPs associated with the user
P-Call-Id: Party-ID	Lists the SIP of the called party
P-Contact: 1.0.0-800.0.0-0.0	Identifies the contact method
P-Connection: Name-to-Data	Identifies the access network
P-Content-Disposition: Audio	Content-carrying information
P-Content-Transfer-Encoding	More-carrying information

(From [2].)

6.6 Future of SIP and Wireless

It is clear that as IP networks become increasingly wireless, SIP will often be utilized over wireless networks. It is well suited for such use for the reasons discussed in this chapter: it has built-in mobility support when Mobile IP is not used, and can also be used with Mobile IP depending on the wireless network design.

Additional work on authentication and roaming will likely be done with SIP as the extensions developed for the IETF architecture are not specific to be useful in cross-networks.

References

- [1] Information about the IETF project, including the latest technical specifications (TS), can be found at <http://www.ietf.org>.
- [2] Abuladze, H., and S. Wolfram, "Application-Layer Mobility Using SIP," *Mobility With Computing and Communication Review (MCCR)*, Vol. 4, No. 3, July 2000.
- [3] Perkins, C., "IP Mobility Support," RFC 2002, 1996.
- [4] Field, E., et al., "Supporting Mobility for Multimedia with SIP," IETF Internet-Draft, Work in Progress, December 2000.
- [5] Abuladze, H., "Dynamic Host Configuration Protocol (DHCP) for IPv6: Option for Session Initiation Protocol (SIP) Access," RFC 1864, 2002.
- [6] Davis, A., and S. Anderson, "SIP Request Format for IBC Speech," IETF Internet-Draft, Work in Progress, March 2003.
- [7] IETF TS relating to SIP include TS 24.224 and TS 24.229. For the latest on these and other SIP-related specifications, visit <http://www.ietf.org>.

-
- [8] Deaconu, S., and J. Rosenberg. "Session Timers in the Session Initiation Protocol (SIP)." IETF Internet-Draft, Work in Progress, November 2002.
 - [9] Chen, S., G. Casavalle, and A. Bush. "Support for IPv6 in Session Description Protocol (SDP)." RFC 3446, 2002.
 - [10] Pfen, R., et al. "Signaling Compression (SigComp)." RFC 3329, 2002.
 - [11] Casavalle, G. "Compressing the Session Initiation Protocol (SIP)." RFC 3486, February 2003.
 - [12] Spohrer, J., et al. "Real-Time Transport Protocol Payload Format and File Storage Format for the Adaptive Multi-Rate (AMR) and Adaptive Multi-Rate Wideband (AMR-WB) Audio Codes." RFC 3367, June 2002.
 - [13] Wells, D., and B. Rosenheim. "Session Initiation Protocol Extension Header Field for Service Route Discovery During Registration." IETF Internet-Draft, Work in Progress, February 2003.
 - [14] Wells, D., and B. Rosenheim. "Session Initiation Protocol (SIP) Extension Header Field for Registering Non-Adjacent Contacts." RFC 3337, 2002.
 - [15] Garcia-Martin, M., E. Handberg, and D. Mills. "Private Header (P-Header) Extension to the Session Initiation Protocol (SIP) for the 3rd-Generation Partnership Project (3GPP)." RFC 3255, 2002.
 - [16] Mankin, L., et al. "Change Process for the Session Initiation Protocol (SIP)." RFC 3427, 2002.

10

Call Flow Examples

In this chapter, many of the concepts and details presented in the preceding chapters will be illustrated with examples. Each example includes a call flow diagram, a discussion of the example, followed by the message details. Each message is labeled in the figure with a message number for easy reference. For more examples of the protocol, refer to the SIP specification [1] and the SIP call flows [2, 3] documents.

The purpose of the examples in this chapter is to illustrate aspects of the SIP protocol. The incorporation scenarios with the PSTN and with an H.323 network are not intended to fully define the interworking or show a complete parameter mapping between the protocols. Likewise, simplifications such as minimal authentication and direct client-to-gateway messaging are used to make the examples more clear.

10.1 SIP Call with Authentication, Proxy, and Record-Route

Figure 10-1 shows a basic SIP call between two SIP user agents involving two proxy servers. Rather than perform a DNS query on the SIP URI of the called party, the calling SIP phone sends the INVITE request to a proxy server for address resolution. The proxy server requires authentication to perform this service and replies with a 407 Proxy Authentication Required response. Using the nonce from the challenge, the caller sends the INVITE with the caller's username and password credentials encrypted. The proxy checks the credentials, and finding them to be correct, performs the DNS lookup on the Request-URI. The INVITE is then forwarded to the proxy server listed in the DNS SRV record that handles the `language.org`




Figure 10-1 SIP-to-SIP call with authentication, proxy, and record-route.

domain. The proxy then looks up the Request-URI and locates a registration for the called party. The INVITE is forwarded to the destination UAS, a Record-Route header having been inserted to ensure that the proxy is present in all future requests by either party. This is because a direct routed SIP message to Ada would be blocked by the firewall.

The called party receives the INVITE request and sends 200 OK (ringing) and 200 OK responses, which are routed back to the caller using the Via header chain from the initial INVITE. The ACK sent by the caller includes a Route header built from the Record-Route header field in the 200 OK response. This routing skips the first proxy but includes the firewall proxy. The media session begins with the user agents exchanging RTP and RTCP packets.

The call terminates when the called party, Ada, sends a BYE, which includes a Route header generated from the Record-Route header field in the INVITE. Note that the CSeq for the called user agent is initialized to

1990. The acknowledgment of the BYE with a 200 OK response causes both sides to stop sending media packets.

- 81** `BYE sip:info@conference.com SIP/1.0` **+RequestURI**
 Via: SIP/2.0/UDP 10.20.17.10:5060; branch=0000000000
 Max-Forwards: 70
 From: Charles Babbage <info@info@conference.com> <tag=001>
 To: sip:info@conference.com
 Call-ID: 00-02-0a-04-01-a78aa47d1e1e1e1e1e1e1e1e1e1e1e1e
 CSeq: 1 BYE **+Call-Id=00-02-0a-04-01-a78aa47d1e1e1e1e1e1e1e1e1e1e1e1e**
 Contact: <info@info@conference.com> <tag=001>
 Subject: 00-00000000
 User-Agent: Conference-Kalua 1
 Content-Type: application/sdp
 Content-Length: 137
- 0-0
 info@conference.com 0000000000 0000000000 00 00 00 00 00 00 00 00
 0-0
 000 000 00 00 00 00 **info@conference.com**
 00000000 00000000 00000000 **+From number**
 00000000 00000000 **+Content-Id**
- 82** `200 OK SIP/2.0` **From: info@conference.com**
 Via: SIP/2.0/UDP 10.20.17.10:5060; branch=0000000000
 From: Charles Babbage <info@info@conference.com> <tag=001>
 To: <info@info@conference.com> <tag=001>
 Call-ID: 00-02-0a-04-01-a78aa47d1e1e1e1e1e1e1e1e1e1e1e1e
 CSeq: 1 BYE **+Content-Id=**
 From: info@conference.com
 <info@info@conference.com>
 <tag=001>, <info@info@conference.com>
 <tag=001>, <info@info@conference.com>
- 83** `200 OK sip:info@conference.com SIP/1.0`
 Via: SIP/2.0/UDP 10.20.17.10:5060; branch=0000000000
 Max-Forwards: 70
 From: Charles Babbage <info@info@conference.com> <tag=001>
 To: <info@info@conference.com> <tag=001>
 Call-ID: 00-02-0a-04-01-a78aa47d1e1e1e1e1e1e1e1e1e1e1e1e
 CSeq: 1 ACK **+Call-Id=00-02-0a-04-01-a78aa47d1e1e1e1e1e1e1e1e1e1e1e1e**
- 84** `BYE sip:info@conference.com SIP/1.0` **+SIP-URI**
 Via: SIP/2.0/UDP 10.20.17.10:5060; branch=0000000000 **+branch**
 Max-Forwards: 70
 From: Charles Babbage <info@info@conference.com> <tag=001>
 To: <info@info@conference.com>

- 820** **82000-82004** :+042:29.14.90.1:1P+
- 821** **821V2.4 140 Mapping**
 Pla: 821V2.4.820P 15.24.17.18.9040/branch/0604040200449.1
 Pla: 821V2.4.820P 15.24.17.18.9040/branch/060404020045000
 From: Charles Hillman =>ls-hillman@msl.tytl.ncsl.gov, org: tag@0041
 To: =>ls-hillman@msl.tytl.ncsl.gov, tag@0041@0041
 Cal.1-15: 06-11-06-04-01-@Wms.tytl.ncsl.gov
 CSeq: 1 200708
 Content: sfp-schedule@msl.tytl.ncsl.gov, 1.smp@msl.tytl.ncsl.gov
82000-82004 :+042:29.14.90.1:1P+
- 822** **822V2.4 140 Mapping**
 Pla: 822V2.4.820P 15.24.17.18.9040/branch/060404020045000
 From: Charles Hillman =>ls-hillman@msl.tytl.ncsl.gov, org: tag@0041
 To: =>ls-hillman@msl.tytl.ncsl.gov, tag@0041@0041
 Cal.1-15: 06-11-06-04-01-@Wms.tytl.ncsl.gov
 CSeq: 1 200708
- 823** **823V2.4 000-0E** ***CallAccepted**
 Pla: 823V2.4.820P 15.24.17.18.9040/branch/0604040200450.1
 Pla: 823V2.4.820P 15.24.17.18.9040/branch/060404020045000
 Pla: 823V2.4.820P 15.24.17.18.9040/branch/060404020045000
 From: Charles Hillman =>ls-hillman@msl.tytl.ncsl.gov, org: tag@0041
 To: =>ls-hillman@msl.tytl.ncsl.gov, tag@0041@0041
 Cal.1-15: 06-11-06-04-01-@Wms.tytl.ncsl.gov
 CSeq: 1 200708
 Content: sfp-schedule@msl.tytl.ncsl.gov, 1.smp@msl.tytl.ncsl.gov
82000-82004 :+042:29.14.90.1:1P+
 Content-Type: application/sfp
 Content-Length: 129
- end
 endto 1000004128 1000041200 08 294.1.2.3.4
 en-
 End 0
 end 294.1.2.3.4 ***CallFailure**
 eventto 20000 823V2.4P 0 ***Free number**
 smp@msl.tytl.ncsl.gov 0 2007/0808 ***Code information**
- 824** **824V2.4 000-0E**
 Pla: 824V2.4.820P 15.24.17.18.9040/branch/0604040200450.1
 Pla: 824V2.4.820P 15.24.17.18.9040/branch/060404020045000
 From: Charles Hillman =>ls-hillman@msl.tytl.ncsl.gov, org: tag@0041
 To: =>ls-hillman@msl.tytl.ncsl.gov, tag@0041@0041
 Cal.1-15: 06-11-06-04-01-@Wms.tytl.ncsl.gov
 CSeq: 1 200708
 Content: sfp-schedule@msl.tytl.ncsl.gov, 1.smp@msl.tytl.ncsl.gov
82000-82004 :+042:29.14.90.1:1P+
 Content-Type: application/sfp
 Content-Length: 129

- CallFlow-Example: 128
- V=0
 vmln 200004010 200004010 DE 30 1.2.3.4
 p=0
 call 0
 rtp 000 1.2.3.4
 rtpsize 1000 RTP/MP 0
 srtpseq 0 00000000
- 824 SIP/0.0.000 OK
 Via: SIP/0.0/UDP 10.20.17.10:5060 branch=00000000
 From: Charles Bakke <cbakke@cs.cmu.edu>
 To: <cbakke@cs.cmu.edu>
 Call-ID: 00-00-00-01-000000000000
 CSeq: 1 INVITE
 Contact: sip:cbakke@cs.cmu.edu
 Record-Route: <cbakke@10.20.17.10>
 Content-Type: application/sdp
 CallFlow-Example: 128
- V=0
 vmln 200004010 200004010 DE 30 1.2.3.4
 p=0
 call 0
 rtp 000 1.2.3.4
 rtpsize 1000 RTP/MP 0
 srtpseq 0 00000000
- 825 ACK SIP/0.0.000 OK/000.000000-000 0000.0 40000 000
 Via: SIP/0.0/UDP 10.20.17.10:5060 branch=00000000
 Max-Response: 0
 From: Charles Bakke <cbakke@cs.cmu.edu>
 To: <cbakke@cs.cmu.edu>
 Call-ID: 00-00-00-01-000000000000
 CSeq: 1 ACK
 Record-Route: <cbakke@10.20.17.10>
- 826 ACK sip:cbakke@cs.cmu.edu 0000.0
 Via: SIP/0.0/UDP 10.20.17.10:5060 branch=00000000
 Via: SIP/0.0/UDP 10.20.17.10:5060 branch=00000000
 Max-Response: 0
 From: Charles Bakke <cbakke@cs.cmu.edu>
 To: <cbakke@cs.cmu.edu>
 Call-ID: 00-00-00-01-000000000000
 CSeq: 1 BYE
- 827 BYE SIP/0.0.000 OK/1000 00000000-000 SIP/0.0
 Via: SIP/0.0/UDP 1.2.3.4:5060 branch=00000000
 Max-Response: 0

- ```

From: Alice [mailto:alice@cs.cmu.edu] <alice@cs.cmu.edu>
To: Charlie on Behalf Of <charlie@cs.cmu.edu>
Call-ID: 2004-03-24-01-07988312752014001.002
Dialog: 2004-038
Reason: <452> 486 (Not Acceptable)
From: Charlie on Behalf Of

822 BYE <alice@cs.cmu.edu>
Via: SIP/2.0/UDP 10.24.0.1:5060; branch=07988312752014001.002
Via: SIP/2.0/UDP 1.2.3.4:5060; branch=07988312752014001.002
Max-Forwards: 40
From: Alice [mailto:alice@cs.cmu.edu] <alice@cs.cmu.edu>
To: Charlie on Behalf Of <charlie@cs.cmu.edu>
Call-ID: 2004-03-24-01-07988312752014001.002
Dialog: 2004-038

823 SIP/2.0 200 OK
Via: SIP/2.0/UDP 10.24.0.1:5060; branch=07988312752014001.002
Via: SIP/2.0/UDP 1.2.3.4:5060; branch=07988312752014001.002
From: Alice [mailto:alice@cs.cmu.edu] <alice@cs.cmu.edu>
To: Charlie on Behalf Of <charlie@cs.cmu.edu>
Call-ID: 2004-03-24-01-07988312752014001.002
Dialog: 2004-038

824 SIP/2.0 200 OK
Via: SIP/2.0/UDP 1.2.3.4:5060; branch=07988312752014001.002
From: Alice [mailto:alice@cs.cmu.edu] <alice@cs.cmu.edu>
To: Charlie on Behalf Of <charlie@cs.cmu.edu>
Call-ID: 2004-03-24-01-07988312752014001.002
Dialog: 2004-038

```

## 10.2 SIP Call with Stateless and Stateful Proxies with Called Party Busy

Figure 10.2 shows an example of a SIP with a stateless proxy server and a stateful proxy server. The call is not completed because called party is busy. The called user agent initially sends a 180 Ringing response but then sends a 500 Busy Everywhere response containing a Retry-After header to indicate that the call is being rejected. The stateful proxy server sends a 180 Trying response to the INVITE, and also acknowledges the 500 Busy Everywhere response with an ACK. The stateless proxy does not send a 180 Trying and forwards the 500 Busy Everywhere and the ACK sent by the called user agent. Also note that the initial INVITE does not contain a message body.

- ```

825 INVITE sip:charlie@cs.cmu.edu;orig_uri=SIP/2.0
Via: SIP/2.0/UDP 10.24.0.1:5060; branch=07988312752014001.002

```




Figure 10.1 SIP call example with status and 180 Ringing passed without being called party.

- ```

100 200 OK (application/javascript)
From: <8000@8000>
To: <8001@8001>
Call-ID: <80008001000000000000000000000000>
Contact: <8000@8000>
Content-Length: 0
<Optional header>
<Optional Content-Length header>

801 200 OK (application/javascript)
Via: SIP/2.0/UDP 8.8.8.8:8000
To: <8001@8001>
Call-ID: <80008001000000000000000000000000>
Contact: <8000@8000>
Content-Length: 0
<Optional header>
<Optional Content-Length header>

802 200 OK (application/javascript)
Via: SIP/2.0/UDP 8.8.8.8:8000
To: <8003@8003>
Call-ID: <80008001000000000000000000000000>
Contact: <8000@8000>
Content-Length: 0
<Optional header>
<Optional Content-Length header>

803 200 OK (application/javascript)
Via: SIP/2.0/UDP 8.8.8.8:8000
To: <8001@8001>
Call-ID: <80008001000000000000000000000000>
Contact: <8000@8000>
Content-Length: 0

```



- #4** `HTTP/1.1 200 OK`  
 Via: HTTP/1.1 TCP 8.8.8.8:80 [reverse-proxy] (172.17.0.1)  
 Via: HTTP/1.1 TCP 8.8.8.8:80 [reverse-proxy] (172.17.0.1)  
 Via: HTTP/1.1 TCP 8.8.8.8:80 [reverse-proxy] (172.17.0.1)  
 Via: HTTP/1.1 TCP 8.8.8.8:80 [reverse-proxy] (172.17.0.1)  
 From: @domain ->[to] @domain@compiling.org [tag=original]  
 To: @domain ->[to] @domain@compiling.org [tag=original]  
 Set-Cookie: @domain@compiling.org [tag=original]  
 Content-Type: @domain@compiling.org [tag=original]  
 Content-Length: 0
- #5** `HTTP/1.1 200 Mapping`  
 Via: HTTP/1.1 TCP 8.8.8.8:80 [reverse-proxy] (172.17.0.1)  
 Via: HTTP/1.1 TCP 8.8.8.8:80 [reverse-proxy] (172.17.0.1)  
 Via: HTTP/1.1 TCP 8.8.8.8:80 [reverse-proxy] (172.17.0.1)  
 From: @domain ->[to] @domain@compiling.org [tag=original]  
 To: @domain ->[to] @domain@compiling.org [tag=original]  
 Set-Cookie: @domain@compiling.org [tag=original]  
 Content-Type: @domain@compiling.org [tag=original]  
 Content-Length: 0
- #6** `HTTP/1.1 200 Mapping`  
 Via: HTTP/1.1 TCP 8.8.8.8:80 [reverse-proxy] (172.17.0.1)  
 Via: HTTP/1.1 TCP 8.8.8.8:80 [reverse-proxy] (172.17.0.1)  
 From: @domain ->[to] @domain@compiling.org [tag=original]  
 To: @domain ->[to] @domain@compiling.org [tag=original]  
 Set-Cookie: @domain@compiling.org [tag=original]  
 Content-Type: @domain@compiling.org [tag=original]  
 Content-Length: 0
- #7** `HTTP/1.1 200 Mapping`  
 Via: HTTP/1.1 TCP 8.8.8.8:80 [reverse-proxy] (172.17.0.1)  
 From: @domain ->[to] @domain@compiling.org [tag=original]  
 To: @domain ->[to] @domain@compiling.org [tag=original]  
 Set-Cookie: @domain@compiling.org [tag=original]  
 Content-Type: @domain@compiling.org [tag=original]  
 Content-Length: 0
- #8** `HTTP/1.1 200 Map @domain` \*\*Mandatory to use  
 Via: HTTP/1.1 TCP 8.8.8.8:80 [reverse-proxy] (172.17.0.1)  
 Via: HTTP/1.1 TCP 8.8.8.8:80 [reverse-proxy] (172.17.0.1)  
 Via: HTTP/1.1 TCP 8.8.8.8:80 [reverse-proxy] (172.17.0.1)  
 From: @domain ->[to] @domain@compiling.org [tag=original]  
 To: @domain ->[to] @domain@compiling.org [tag=original]  
 Set-Cookie: @domain@compiling.org [tag=original]

- Call: 1 BYT5**  
**Setup-Action: Out, 5 Jul 2008 11:39:00 AM**  
**Context-Setup: 0**
- #1 ACK sip:abacklog@pt.com:5060, from SIPV.0      \*Revised part from ACK**  
**Via: SIPV.0/UDP 192.168.1.100/5060 (branch:abacklog) [P] 0**  
**Max-Forwards: 70**  
**From: abacklog sip:abacklog@pt.com;tag=9156**  
**To: abacklog sip:abacklog@pt.com;tag=9156**  
**Call-ID: abacklog@pt.com;tag=9156**  
**Call: 1 ACK**  
**Context-Setup: 0**
- #2 SIPV.0 400 Busy Everywhere**  
**Via: SIPV.0/UDP 192.168.1.100/5060 (branch:abacklog) [P] 1**  
**Via: SIPV.0/UDP 192.168.1.100/5060 (branch:abacklog) [P] 0**  
**From: abacklog sip:abacklog@pt.com;tag=9156**  
**To: abacklog sip:abacklog@pt.com;tag=9156**  
**Call-ID: abacklog@pt.com;tag=9156**  
**Call: 1 BYT5**  
**Setup-Action: Out, 5 Jul 2008 11:39:00 AM**  
**Context-Setup: 0**  
**Call-ID: abacklog 100**
- #3 SIPV.0 400 Busy Everywhere      \*Revised part from ACK response**  
**Via: SIPV.0/UDP 192.168.1.100/5060 (branch:abacklog) [P] 1**  
**Via: SIPV.0/UDP 192.168.1.100/5060 (branch:abacklog) [P] 0**  
**From: abacklog sip:abacklog@pt.com;tag=9156**  
**To: abacklog sip:abacklog@pt.com;tag=9156**  
**Call-ID: abacklog@pt.com;tag=9156**  
**Call: 1 BYT5**  
**Setup-Action: Out, 5 Jul 2008 11:39:00 AM**  
**Context-Setup: 0**
- #4 ACK sip:abacklog@pt.com:5060, from SIPV.0**  
**Via: SIPV.0/UDP 192.168.1.100/5060 (branch:abacklog) [P] 0**  
**Max-Forwards: 70**  
**From: abacklog sip:abacklog@pt.com;tag=9156**  
**To: abacklog sip:abacklog@pt.com;tag=9156**  
**Call-ID: abacklog@pt.com;tag=9156**  
**Call: 1 ACK**  
**Context-Setup: 0**
- #5 ACK SIPV.0 400 Busy Everywhere, from SIPV.0**  
**Via: SIPV.0/UDP 192.168.1.100/5060 (branch:abacklog) [P] 1**  
**Via: SIPV.0/UDP 192.168.1.100/5060 (branch:abacklog) [P] 0**  
**Max-Forwards: 60**  
**From: abacklog sip:abacklog@pt.com;tag=9156**  
**To: abacklog sip:abacklog@pt.com;tag=9156**  
**Call-ID: abacklog@pt.com;tag=9156**

Chapter 1 ACE  
 Version 1.1, August 2004

### 10.3 SIP to PSTN Call Through Gateway

In the example shown in Figure 10.3, the calling SIP phone places a telephone call to the PSTN through a PSTN gateway. The SIP phone collects the dialed digits and puts them into a SIP URI used in the Request-URI and the To header. The caller may have dialed either the globalized phone number 1-202-555-1233 or they may have just dialed a local number 555-1233, and the SIP phone added the assumed country code and area code to produce the globalized URI. The SIP phone has been preconfigured with the IP address of




Figure 10.3 SIP to PSTN call through gateway.

the PSTN gateway, so it is able to send the INVITE directly to gw1.cisco.com. The gateway initiates the call into the PSTN by selecting an SIP ISUP trunk to the next telephone switch in the PSTN. The dialed digits from the INVITE are mapped into the ISUP ISDN. The ISUP Address Complete Message (ACM) is sent back by the PSTN to indicate that the trunk has been seized. Progress tones are generated in the one-way audio path established in the PSTN. In this example, ring tone is generated by the far-end telephone switch. The gateway maps the ACM to the 183 Session Progress response containing SDP indicating the RTP port that the gateway will bridge the audio from the PSTN. Upon reception of the 183, the caller's UAC begins receiving the RTP packets sent from the gateway and presents the audio to the caller so they know that the call is progressing in the PSTN.

The call completes when the called party answers the telephone, which causes the telephone switch to send an Answer Message (ANM) to the gateway. The gateway then cuts the PSTN audio connection through in both directions and sends a 200 OK response to the caller. Because the RTP media path is already established, the gateway echoes the SDP in the 200 but causes no changes to the RTP connection. The UAC sends an ACK to complete the SIP signaling exchange. Because there is no equivalent message in ISUP, the gateway echoes the ACK.

The call terminates when the caller sends the BYE to the gateway. The gateway maps the BYE to the ISUP Release message or RLC. The gateway sends the 200 OK to the BYE and receives a RLC from the PSTN. These two messages have no dependency on each other. If, for some reason, either the SIP or PSTN network does not respond properly, one does not want resources held in the other network as a result.

```

81 INVITE sip:+12025551234@192.168.1.100:5060 SIP/2.0
To: SIP/2.0/SIP 8.18.25.24@192.168.1.100:5060
Max-Forwards: 70
From: sip:8.18.25.24@192.168.1.100:5060;ip=8.18.25.24
Via: sip/2.0/192.168.1.100:5060;branch=0,branch=1
Call-ID: 80001410000192.168.1.100:5060_17
CSeq: 1 INVITE
Expires: 300
Contact: sip:8.18.25.24@192.168.1.100:5060;ip=8.18.25.24
Content-Type: application/sdp

v=0
o=IP 8000141000192.168.1.100 1 8.18.25.24
s=
t=0
m=audio 8000 8.18.25.24
a=sendrecv
a=rtpmap:0 PCMA/8000
a=rtpmap:8 PCMA/8000

```

© The domain name, IP, and port are fictitious.

- ```

81 100
   M=audio;O=XXX-1111;MSID=101;
   OS=0111001                                          #(lines are optional
                                                       the called party number)

82 200

83 110
   SIP/2.0 201 Session Forwarded
   Via: SIP/2.0/UDP 4.44.44.44;branch=z9hJz6454454445
   From: 4444@44.44.44.44;branch=z9hJz6454454445;to=4444
   To: 4444@44.44.44.44;branch=z9hJz6454454445;to=4444
   Call-ID: 44445445445445445445445445445445445445
   CSeq: 1 20101
   Reason: <code>300, 44, 70, 80</code>
   Content-Type: application/sdp
   Content-Length: 128

   v=0
   o=john1700 2886666666 3000000000 IN IP4 99.44.78.40
   s=
   t=
   c=IN 44.44.78.40
   m=j1700 4444 RTP/AVP 0
   a=rtpmap: 0 PCMU/8000

#(lines optional)

84 240
   2400 sdp 99.44.78.40 2401 0
   Via: SIP/2.0/UDP 4.44.44.44;branch=z9hJz6454454445
   Max-Forwards: 70
   From: 4444@44.44.44.44;branch=z9hJz6454454445;to=4444
   To: 4444@44.44.44.44;branch=z9hJz6454454445;to=4444
   Call-ID: 44445445445445445445445445445445445445
   Reason: sdp 4444;branch=z9hJz6454454445;to=4444
   Content-Length: 0

85 250 OK
   SIP/2.0 200 OK
   Via: SIP/2.0/UDP 4.44.44.44;branch=z9hJz6454454445
   From: 4444@44.44.44.44;branch=z9hJz6454454445;to=4444
   To: 4444@44.44.44.44;branch=z9hJz6454454445;to=4444
   Call-ID: 44445445445445445445445445445445445445
   CSeq: 1 2400

87 200

```


10.4 PSTN to SIP Call Through Gateway

Figure 10-4 shows a call originating from a telephone in the PSTN that terminates a SIP phone in the Internet. The compact form of SIP is used throughout the example. Note that there is no compact form for CSeq or Non-Parameters.

```

61 INVITE sip:155.214.54@voip.com;transport=udp;ttl=0
v: SIP/2.0;seq=15;2.4.2;1555-2000@155.214.54 →Compact form of header
Non-Parameters: 30
E: rfc4575 →SIP/2.0;seq=15;transport=udp;ttl=0 →includes q
t: rfc4575 →SIP/2.0;seq=15;transport=udp
L: 62-65-66-68-69;2.4.2
CSeq: 1 2000
R: SIP/2.0;155-2000@155
m: application/sdp
L: 304
end
  
```




Figure 10-4 PSTN to SIP phone through gateway.

```

<!-- 288086003 2880860033 IN 284 41.5.4.1
&E
!<@ 0
<CD 284 41.5.4.1
B-0012> 0000 RTB/MPI 0
<scripting> 0 00000000

```

84 RTB/5.0 000 001000

```

<!-- 28803.0/820 41.5.4.1 0000-0000000000000000
E <url>+45-07000000000000.com/000-000000000000
T <url>+45.00000000000000.com/000-00000000
L 40-00-00-00000.5.4.1
<Tag> 1 000000

```

84 Service Query: +00-11000

85 Location Service Response:

012-00000000-000 *Number says to MP 000

86 287700 <url>+0000000000.com 0000.0

```

<!-- 28803.0/820 00.5.4.1 0000-0000000000000000
<!-- 28803.0/820 41.5.4.1 0000-0000000000000000
000-000000000 00
E <url>+45-07000000000000.com/000-0000000000000000
T <url>+45.00000000000000.com/000-0000000000000000
L 40-00-00-00000.5.4.1
<Tag> 1 000000
R 000-000.000-000000-000
<scripting> 000000000000
L 100

```

</!--

```

<!-- 288086003 2880860033 IN 284 41.5.4.1
&E
!<@ 0

```

```

<CD 284 41.5.4.1
B-0012> 0000 RTB/MPI 0
<scripting> 0 00000000

```

87 RTB/5.0 000 011000

```

<!-- 28803.0/820 00.5.4.1 0000-0000000000000000
<!-- 28803.0/820 41.5.4.1 0000-0000000000000000
E <url>+45-07000000000000.com/000-0000000000000000
T <url>+45.00000000000000.com/000-0000000000000000
L 40-00-00-00000.5.4.1
R <url>+00000000000000.com
<Tag> 1 000000

```


- 89 SIP/2.0 200 OK**
 m: SIP/2.0/UDP 65.55.2.4.1;branch=6555241
 E: <65-55-241@msn.com>
 c: <65-55-241@msn.com>
 l: 65-55-241;2.4.1
 CSeq: 1 200OK
- 90 Warning**
- 91 SIP/2.0 200 OK**
 m: SIP/2.0/UDP 77.8.8.1;branch=778881
 M: SIP/2.0/UDP 65.55.2.4.1;branch=6555241
 E: <65-55-241@msn.com>
 c: <65-55-241@msn.com>
 l: 65-55-241;2.4.1
 CSeq: 1 200OK
 m: sip.msns.com
 c: application/sdp
 l: 123
- 9-0
 m: 20000001 20000001 IN IP4 7.8.8.8
 m:
 l: 0
 c:RTP/AVP 98 99 100
 a:media=audio
 a:range=0 65535
- 92 SIP/2.0 200 OK**
 m: SIP/2.0/UDP 65.55.2.4.1;branch=6555241
 E: <65-55-241@msn.com>
 c: <65-55-241@msn.com>
 l: 65-55-241;2.4.1
 CSeq: 1 200OK
 m: sip.msns.com
 c: application/sdp
 l: 123
- m:
 m: 20000001 20000001 IN IP4 7.8.8.8
 m:
 l: 0
 c:RTP/AVP 98 99 100
 a:media=audio
 a:range=0 65535
- 93 Content**
- 94 ACK SIP/2.0/UDP 77.8.8.1**

```

%# ECHO, BYE, BYE, BYE, BYE, BYE, BYE, BYE, BYE, BYE, BYE
Max-Forwards: 70
D: 402-45, 4384041000000, 000-000-00000-000-000000
to: callee@E, 000000000000, 000000000000-000000000000
L: 00-00-00-00000, 0, 0, 1
CSeq: 1 ACK

```

10.5 Parallel Search

In this example the caller receives multiple possible locations for the called party from a redirect server. Instead of trying the locations one at a time, the user agent implements a parallel search for the called party by simultaneously sending the INVITE to three different locations, as shown in Figure 10.5. The SIP specification gives an example of this behavior in a proxy server, which is called a *forking proxy*.

In this example the first location responds with a 404 Not Found response. The second location responds with a 180 Ringing response, while the third location returns a 180 Ringing then a 200 OK response. The caller then sends an ACK to the third location to establish the call. Because one successful response has been received, a CANCEL is sent to the second location to terminate the search. The second location sends a 200 OK to the CANCEL and a 407 Request Terminated to the INVITE. This example shows some common error cases: phrases in messages M7, M8, and M11.

```

M1 INVITE 404 000000000000, 000-000-00000-000-00000
Via: SIP/2.0/SIP 7, 8, 10, 11; 00000; branch=00000000000000000000000000000000
Max-Forwards: 70
From: J.C. Brown/1, callee@example1000.org, callee@E, callee@C
To: 4
Call-ID: 0000000000000000000000000000000000000000000000000000000000000000
CSeq: 14 INVITE
Contact: callee@example1000.org, callee@E, callee@C
Content-Length: 120

```

```

M2
0-000 200000000000 000000000000, 000-000-00000-000-00000
to:
100 0
retr 100 7, 8, 10, 11
message: 00000 000000000000 4
0-000000-0 000000000000

```

```

M3 SIP/2.0 200 Multiple Locations Call-ID: 0000000000000000000000000000000000000000000000000000000000000000
Via: SIP/2.0/SIP 7, 8, 10, 11; 00000; branch=0000000000000000000000000000000000000000000000000000000000000000 from location

```




Figure 8A Parallel search example.

```

From: J.C. Maxwell <only@exam.academy@blaze.com>
To: 180

```

```

To: only@blaze.com
Call-ID: 180

```

```

CSeq: 14 INVITE

```

```

Contact: only@blaze.com

```

```

Contact: 180@192.168.0.100

```

```

Contact: only@blaze.com

```

```

83 ACK: only@exam.academy@blaze.com
Via: SIP/2.0/UDP 192.168.0.100
Call-ID: 180

```

```

From: J.C. Maxwell <only@exam.academy@blaze.com>
To: 180

```

```

To: 180@blaze.com
Call-ID: 180

```

```

CSeq: 14 INVITE

```

```

84 200: only@blaze.com
Via: SIP/2.0/UDP 192.168.0.100

```

- Mail-Forwarder: 70
 From: J.C. Maxwell <mlp@james.meredith.org>, contact@jds.abc
 <100>
 To: <mlp@familydoctor.org>, name@jds <Tag is not signed
 <CallID is unsigned
 <CallTimestamp>
 Call-ID: 801669480000007, P.00.11
 CSeq: 99 100000
 Contact: <mlp@james.meredith.org>, contact@jds.abc
 Contact-Type: mail,location,role
 Contact-Group: 100

 <nil>
 <seq 200004011 100004001 IN 100 T, P.00.11
 <nil>
 <nil 0
 <seq 100 T, P.00.11
 <media 101-00 RTP/AVP 4
 <transport 4 DTLS/000
- 83** SIP/200 400 501 4262610-phones-offline.org;name@phones.offline.org
 Via: SIP/2.0/UDP T, P.00.11;branch=100040001
 Mail-Forwarder: 70
 From: J.C. Maxwell <mlp@james.meredith.org>, contact@jds.abc
 <10001>
 To: <mlp@familydoctor.org>, name@jds
 Call-ID: 801669480000007, P.00.11
 CSeq: 99 100000
 Contact: <mlp@james.meredith.org>, contact@jds.abc
 Contact-Type: application,role
 Contact-Group: 100

 <nil>
 <seq 200004011 100004001 IN 100 T, P.00.11
 <nil>
 <nil 0
 <seq 100 T, P.00.11
 <media 101-00 RTP/AVP 4
 <transport 4 DTLS/000
- 84** SIP/200 400 501 4262610-phones-offline.org;name@phones.offline.org
 Via: SIP/2.0/UDP T, P.00.11;branch=100040001
 Mail-Forwarder: 70
 From: J.C. Maxwell <mlp@james.meredith.org>, contact@jds.abc
 <10001>
 To: <mlp@familydoctor.org>, name@jds
 Call-ID: 801669480000007, P.00.11
 CSeq: 99 100000
 Contact: <mlp@james.meredith.org>, contact@jds.abc
 Contact-Type: application,role
 Contact-Group: 100

 <nil>
 <seq 200004011 100004001 IN 100 T, P.00.11

- From: J.C. Maxwell <mcj@cam.ac.uk>
 To: <ietf-voice@ietf.org>
 Date: 2004-08-11 10:00:00
- 87 SIP/2.0 404 The address you have requested is not available
 Via: SIP/2.0/UDP 7.8.38.12;branch=branch1
 From: J.C. Maxwell <mcj@cam.ac.uk>
 To: <ietf-voice@ietf.org>
 Call-ID: <mcj@cam.ac.uk>
 CSeq: 55 10000
- 88 ACK sip:barack@kth.se;branch=br SIP/2.0
 Via: SIP/2.0/UDP 7.8.38.12;branch=branch1
 Max-Forwards: 70
 From: J.C. Maxwell <mcj@cam.ac.uk>
 To: <barack@kth.se>
 Call-ID: <mcj@cam.ac.uk>
 CSeq: 56 ACK
- 89 SIP/2.0 200 Ringing
 Via: SIP/2.0/UDP 7.8.38.12;branch=branch1
 From: J.C. Maxwell <mcj@cam.ac.uk>
 To: <barack@kth.se>
 Call-ID: <mcj@cam.ac.uk>
 Contact: <mcj@cam.ac.uk>
 CSeq: 55 10000
- 90 SIP/2.0 200 Please wait while we locate Mr. Farber
 Via: SIP/2.0/UDP 7.8.38.12;branch=branch1
 From: J.C. Maxwell <mcj@cam.ac.uk>
 To: <barack@kth.se>
 Call-ID: <mcj@cam.ac.uk>
 Contact: <mcj@cam.ac.uk>
 CSeq: 56 10000
- 91 SIP/2.0 200 Mr. Farber is your server
 Via: SIP/2.0/UDP 7.8.38.12;branch=branch1
 From: J.C. Maxwell <mcj@cam.ac.uk>
 To: <barack@kth.se>
 Call-ID: <mcj@cam.ac.uk>
 CSeq: 56 10000
 P-Assert-Info: SIP 18

- CallFlow:** <code>./src/flow/submit-conn-req.cflow.cdf</code>
connect-type: application/ulp
CallFlow-headers: 128
- V=0**
conn: 2000040011 2000040012 IN 204 8 00 27 00
len: 8
addr: 204 8 00 27 00
srcLen: 10760 000000 4
srcConn: 4 00000000
- E11 ACK ulp/flow/submit-conn-req.cflow.cdf source: E11/V=0**
Fls: E11/V=0, S/NDF 7, 8, 20, 12, 000000, 000000+00000000
Max-Forward: 70
From: 4=0, 00000001, ulp (James Maxwell@hpl.hp.com, cambridge.ac.uk, ulp
 (tag))
To: ulp/flow/submit-conn-req.cflow.cdf (tag=001)
Call-ID: 0000000000000000, 8, 20, 12
Seq: 00, 0000
- E12 CANCEL, EIC=40, 0000, 10000000+000000, 000000, 00000000+000000 E11/V=0**
Fls: E11/V=0, S/NDF 7, 8, 20, 12, 000000, 000000+00000000 **@Cancel, seq=**
Max-Forward: 70
From: J.C. Maxwell, ulp (James Maxwell@hpl.hp.com, cambridge.ac.uk, ulp
 (tag))
To: ulp/flow/submit-conn-req.cflow.cdf (tag=001)
Call-ID: 0000000000000000, 8, 20, 12
Seq: 00, CANCEL **@Tag not recognized**
Method or no CANCEL
- E13 E11/V=0 000 00** **@CANCEL, acknowledged**
Fls: E11/V=0, S/NDF 7, 8, 20, 12, 000000, 000000+00000000
From: 4=0, 00000001, ulp (James Maxwell@hpl.hp.com, cambridge.ac.uk, ulp
 (tag))
To: ulp/flow/submit-conn-req.cflow.cdf (tag=001)
Call-ID: 0000000000000000, 8, 20, 12
Seq: 00, CANCEL
- E14 E11/V=0 400 0000000000000000** **@Final response (SMTE)**
Fls: E11/V=0, S/NDF 7, 8, 20, 12, 000000, 000000+00000000
From: 4=0, 00000001, ulp (James Maxwell@hpl.hp.com, cambridge.ac.uk, ulp
 (tag))
To: ulp/flow/submit-conn-req.cflow.cdf (tag=001)
Call-ID: 0000000000000000, 8, 20, 12
Seq: 00, 000000
- E15 ACK E11/V=0**
Fls: E11/V=0, S/NDF 7, 8, 20, 12, 000000, 000000+00000000
From: J.C. Maxwell, ulp (James Maxwell@hpl.hp.com, cambridge.ac.uk, ulp
 (tag))

The following examples are taken from [16], pp. 144–145:
 [16] L. Rosenberg, *SIP: Session Initiation Protocol*, 2nd ed., Addison-Wesley, 2005.
 ©2005 by Addison-Wesley.

10.6 H.323 to SIP Call

In this example, a H.323 terminal calls a SIP-enabled PC through a H.323/SIP gateway. The gateway does signaling translation between the protocols but allows the two end points to exchange media packets directly with each other. The full details of SIP/H.323 interworking are being developed in the SIP working group [4].

In this example, shown in Figure 10.6, the initial message exchange is between the calling H.323 terminal and the H.323 gateway. The gateway receives the H.323 alert tone as an address served by the H.323/SIP gateway. The ACK response indicates that gateway-to-end signaling is required, so the Q.931 and H.245 TCP connections are opened to the gateway, which opens TCP connections to the gateway. The calling H.323 terminal sends a Q.931 Setup message to the gateway, which forwards it to the H.323/SIP gateway. The gateway then looks up the H.323 alias and receives it to the SIP URI of the called party. It constructs an INVITE from the Setup message and forwards it to a SIP proxy, which forwards it to the called party. Note that because the Setup message does not contain any media information, the INVITE does not contain any media information either. The called party sends a 100 Ringing then a 200 OK to indicate that the call has been answered. The media information present in the SDP message body is moved by the gateway, which sends Acknowledgment and Connect messages.

Messages are sent to the gateway, which forwards them to the calling H.323 terminal. The gateway holds off sending the ACK response to the INVITE until the H.245 media-exchange is completed between the H.323 terminal and the gateway. Once that is complete, the negotiated media capabilities are returned in the ACK and the media session begins.

```

81  400
    400UNRECOGNIZED_EXTENSION

82  ACK
    200ACKNOWLEDGMENT

83  SETUP
    SD address=33334444@10.10.10.1
    CE address=33334444@10.10.10.1

84  Setup
  
```



Figure 16-6 HTTP to SIP call

```

16  sip address:192.168.1.101:5060
17  sip address:192.168.1.101:5060/transport

81  BYE
82  OK

83  BYE
84  OK

85  SIP/2.0 200 OK (application/javascript)
Via: SIP/2.0/UDP 192.168.1.101:5060;branch=100
From: sip:192.168.1.101@192.168.1.101
To: sip:192.168.1.101:5060
Call-ID: 192.168.1.101:5060
CSeq: 41130 BYE
Contact: sip:192.168.1.101:5060;transport=udp

```


- 800** sip@cs.tu-berlin.de 0
801 sip@vps1
 Fla: SIP/2.0/UDP 2.3.4.5:5044/transport=XMPP/8044
 From: <sip:from@example312-gateway.com>
 To: <sip:to@destination.com>
 Call-ID: 800711020478168094
 msg: 4142 siprr
 Contact: <cs.tu-berlin.de 0>
- 802** siprr@cs.tu-berlin.de 000011000110001.com 8001/0
 Fla: SIP/2.0/UDP 2.3.4.5:5044/transport=XMPP/8044/1
 Fla: SIP/2.0/UDP 2.3.4.5:5044/transport=XMPP/8044
 Via: SIP/2.0
 From: <sip:from@example312-gateway.com> <tag=tag1@fla1>
 To: <sip:to@destination.com> <tag=tag1>
 Call-ID: 800711020478168094
 msg: 4142 siprr
 Contact: <cs.tu-berlin.de 000011000110001.com transport=im>
- 803** sip/2.0 300 SIP/2.0
 Fla: SIP/2.0/UDP 2.3.4.5:5044/transport=XMPP/8044/1
 Fla: SIP/2.0/UDP 2.3.4.5:5044/transport=XMPP/8044
 From: <sip:from@example312-gateway.com> <tag=tag1@fla1>
 To: <sip:to@destination.com> <tag=tag1>
 Call-ID: 800711020478168094
 msg: 4142 siprr
 Contact: <cs.tu-berlin.de 000011000110001.com transport=im>
 Content-Length: 0
- 804** sip/2.0 300 siprr
 Fla: SIP/2.0/UDP 2.3.4.5:5044/transport=XMPP/8044
 From: <sip:from@example312-gateway.com> <tag=tag1>
 To: <cs.tu-berlin.de 000011000110001.com>
 Call-ID: 800711020478168094
 msg: 4142 siprr
 Contact: <cs.tu-berlin.de 000011000110001.com transport=im>
 Content-Length: 0
- 805** siprr
 Fla: SIP/2.0/UDP 2.3.4.5:5044/transport=XMPP/8044
 From: <sip:from@example312-gateway.com> <tag=tag1>
 To: <cs.tu-berlin.de 000011000110001.com>
 Call-ID: 800711020478168094
 msg: 4142 siprr
 Contact: <cs.tu-berlin.de 000011000110001.com transport=im>
 Content-Length: 0
- 806** siprr
 Fla: SIP/2.0/UDP 2.3.4.5:5044/transport=XMPP/8044
 From: <sip:from@example312-gateway.com> <tag=tag1>
 To: <cs.tu-berlin.de 000011000110001.com>
 Call-ID: 800711020478168094
 msg: 4142 siprr
 Contact: <cs.tu-berlin.de 000011000110001.com transport=im>
 Content-Length: 0
- 807** siprr
 Fla: SIP/2.0/UDP 2.3.4.5:5044/transport=XMPP/8044
 From: <sip:from@example312-gateway.com> <tag=tag1>
 To: <cs.tu-berlin.de 000011000110001.com>
 Call-ID: 800711020478168094
 msg: 4142 siprr
 Contact: <cs.tu-berlin.de 000011000110001.com transport=im>
 Content-Length: 0
- 808** siprr
 Fla: SIP/2.0/UDP 2.3.4.5:5044/transport=XMPP/8044
 From: <sip:from@example312-gateway.com> <tag=tag1>
 To: <cs.tu-berlin.de 000011000110001.com>
 Call-ID: 800711020478168094
 msg: 4142 siprr
 Contact: <cs.tu-berlin.de 000011000110001.com transport=im>
 Content-Length: 0
- 809** siprr
 Fla: SIP/2.0/UDP 2.3.4.5:5044/transport=XMPP/8044
 From: <sip:from@example312-gateway.com> <tag=tag1>
 To: <cs.tu-berlin.de 000011000110001.com>
 Call-ID: 800711020478168094
 msg: 4142 siprr
 Contact: <cs.tu-berlin.de 000011000110001.com transport=im>
 Content-Length: 0

```

CALL-ID: 20071102144710000004
seq: 4150 00000
CONTACT: +39 021807112000010001 /000 00000-000-000-
Contact-Type: application/sdp
CONTACT-INFO: 100

```

```

v=0
o=orange 20071102 20071102 00 100 0 00-10-00
s=
t=0 0
i=00 100 0 00-10-00
m=audio 1000 RTP/AVP 0
a=rtpmap:0 PCMU/8000

```

```

#21 2007/11/02 00:00:00
File: RTP_VI_0_VCF_1_1_1_1_10000000000000000000
From: sdp:orange@10.10.10.100:5060
To: +39 021807112000010001
m=audio 1000 RTP/AVP 0
c=0 0 0 0
Contact: sdp:orange@10.10.10.100:5060
CONTACT-INFO: 100
CONTACT-INFO: 100

```

```

v=0
o=orange 20071102 20071102 00 100 0 00-10-00
s=
t=0 0
i=00 100 0 00-10-00
m=audio 1000 RTP/AVP 0
a=rtpmap:0 PCMU/8000

```

#24 Contact

#25 Contact

#26 20071102144710000004

#27 20071102144710000004

#28 20071102144710000004

#29 20071102144710000004

#30 20071102144710000004

- R21** Term: *axiomatic* (21) system
- R24** Term: *axiomatic* (24) system
- R25** Term: *axiomatic* (25) system
- R26** Term: *axiomatic* (26) system
- R27** Term: *axiomatic* (27) system
- R28** Term: *axiomatic* (28) system
- R29** Term: *axiomatic* (29) system
- R30** *axiomatic* (30) system
 2013/04/11 1:28:44, 27 0000
- R31** *axiomatic* (31) system
 2013/04/11 1:28:44, 27 0000
- R32** *axiomatic* (32) system
- R33** *axiomatic* (33) system
 Cell Flow Example 182
- R34** *axiomatic* (34) system
 2013/04/11 1:28:44, 27 0000
- R35** *axiomatic* (35) system
 2013/04/11 1:28:44, 27 0000
- R36** *axiomatic* (36) system
- R37** *axiomatic* (37) system
- R38** *axiomatic* (38) system
 File: 2013/04/11 1:28:44, 27 0000
 2013/04/11 1:28:44, 27 0000
 2013/04/11 1:28:44, 27 0000

```

From: <116.100.0.100@10.11.11.100>@10.11.11.100<116.100.0.100>
To: <116.100.0.100@10.11.11.100>@10.11.11.100
Call-ID: 20021110144710040004
Seq: 4182 400
Content-Type: application/sdp
Content-Length: 120

v=0
m= 116.100.0.100 116.100.0.100 1.16.10.10
c=
t=0 0
i=00 100 1.16.10.10
a=media: 116.100.0.100 0
s=116.100.0.100/1000

E2E ACK via 116.100.0.100@10.11.11.100 116.100.0.100
Via: SIP/2.0/UDP 1.1.1.1:5060;branch=100000001.1
To: SIP/2.0/UDP 1.1.1.1:5060;branch=100000001.1
Max-Forwards: 50
From: <116.100.0.100@10.11.11.100>@10.11.11.100<116.100.0.100>
To: <116.100.0.100@10.11.11.100>@10.11.11.100
Call-ID: 20021110144710040004
Seq: 4182 400
Content-Type: application/sdp
Content-Length: 120

v=0
m= 116.100.0.100 116.100.0.100 1.16.10.10
c=
t=0 0
i=00 100 1.16.10.10
a=media: 116.100.0.100 0
s=116.100.0.100/1000

```

10.7 SIP-KDP Wireless Call Flow

An example SIP-KDP call flow is shown in Figure 10.7. Some of the SDP parameters are not shown with their actual values—instead, “...” is listed. Also, some call flows show SDP in the INVITE, but this is only needed if another SDP offer is made—in this case, the initial offer is accepted. Also, 180-OKs are not shown in response to the 180-Ringing response since it does not carry any SDP and does not need to be delivered reliably. In this call flow, the use of QoS procedures is shown (as indicated by the `QoS-parameters` header field and the QoS attributes in the SDP). After the initial offer and answer in the INVITE and 200-OK, the QoS setup is performed. When it has been done successfully, each side indicates in the UPDATE and 200-OK response in the SDP that the setup is complete. Only after that does the called UE begin ringing and sends a 180-Ringing response.


```

CallFlow: +90(1) : 0000 (1111-1111) 00001 : 0000-000-000000-
support@local
Exchange: P2P-00001000 : 000-00000
Proxy-Dialog: none;none
Allow: INVITE, ACK, CANCEL, BYE, UPDATE, MESSAGE, NOTIFY
Security-Profile: 1;none-1;ppp: sp0:1; sipdtone-ua-1-00; sp0:0:0:0:1
(ppp:0:0:0)
Content-Type: application/sdp
content-length: 150

vml
<-- 000000000 0000000010 IN 004 1 : 0000 (1111-1111) 0000
support@local
vml 0
vml 004 1 : 0000 (1111-1111) 0000
sequence: 0/0 sequence 0/ 0
total: 0
account@local none
account@remote none
authentication mandatory local none;none
i-000-000 none 000000 0000000
sequence: 0/ 000
i-000-0/ 000-000-0:1:2:3 : 00000000-0
sequence: 0/ 0000000-0000

sip:0 @ 100 00000
Via: SIP/2.0/UDP
[ : 0000 (1111-1111) 0000 ] : 0000 (authentication@local)
To: +01-01-014-000-0110
From: sip:0@100.0.0.0:0000 : 0000 (10000000)
Call-ID: 00000000000000000000000000000000
CSeq: 0 000000
content-length: 0

INVITE seq: 0-014-000-0110 0000 0
Via: SIP/2.0;user=prox via:000000.com;branch=000000000000
Via: SIP/2.0/UDP [ : 0000 (1111-1111) 0000 ] : 0000-000-000000-
[ : 0000 (10000000)@100.0.0.1]
Max-Forwards: 0
Route: sip:0@100.0.0.0:0000,com:100
Record-Route: sip:0@100.0.0.1:0000,com:100
To: via:01-014-000-0110
From: sip:0@100.0.0.0:0000 : 0000 (10000000)
P-Asserted-Identity: 'Alice' : sip:0@100.0.0.0:0000,com
P-Access-Name: sip:0 : 0000-0000-0000
Call-ID: 000-000-00-0000-000000000000
Priority: none
P-Contact: 00-000000 : 0000-0000-00-00000000
[ : 0000 (00000000)@100.0.0.1] : 0000 (0000)
Call-ID: 000-000-00-0000-000000000000
CSeq: 0 000000
content-length: 0
content-type: +90(1) : 0000 (1111-1111) 00001 : 0000-000-000000-

```

```

Suggested: 100+
Suggested: presentation
Allow: SIPTEL, ACK, CANCEL, BYE, BYACK, MESSAGE, REFER, REFER
Security-Profile: 1gnar-1gpp: qoS.L: sIpSess-02a-1-00: sipSIP101
  (SIP1-101)
Context-Type: application/obj
Context-Length: 150

end

seq= 200000004 200000010 IN IP4 1.1.1.1:5060 OUT IP4 1.1.1.1:5060
method: null
Call ID
seqID 100 1.1.1.1:5060:1111:1111:XXXX
seqIDto 40000 000000 00 00
totalCallID
sequence local none
sequence remote none
sequence mandatory local mandatory
sequence none remote mandatory
sequence 00 000
b-CSEQ=00 totalCallID=0.1.1.1.7 .00000000-0
sequence 00 totalCallID=00000000

REFER 0 100 00100
Plan: SIPTEL, SBCP none / v1 of 100000, none / none / 0000000000
Plan: SIPTEL, SBCP 1.1.1.1:5060:1111:1111:XXXX / 0000 / 0000000000
  / none / 0000000000 / 00
To: seqID=1.1.1.1:5060:1111
From: seqID=1.1.1.1:5060:1111:1111:XXXX / seqID=1111000000
Call-ID: 7000000000 / seqID=0000000000
CSeq: 10 00000
Context-Length: 0

REFER sip:seqID=1.1.1.1:5060:1111:1111:XXXX
Plan: SIPTEL, SBCP none / 1000000, none / none / 0000000000
Plan: SIPTEL, SBCP 1.1.1.1:5060:1111:1111:XXXX / 0000 / 0000000000
  / none / 0000000000 / 00
Seq-Number: 00
Sequence-Header: none / 1000000, none, / seqID=0000000000, none / 00
To: seqID=1.1.1.1:5060:1111
From: seqID=1.1.1.1:5060:1111:1111:XXXX / seqID=1111000000
P-Asserted-Party: seqID=1.1.1.1:5060:1111:1111:XXXX,
  seqID=1.1.1.1:5060:1111:1111:XXXX
P-Asserted-Party: . . .
P-URI: none
P-Handling-Party: . . .
Call-ID: 7000000000 / seqID=0000000000
CSeq: 10 00000
CSeqID: seqID=1.1.1.1:5060:1111:1111:XXXX / seqID=0000000000
Suggested: 100+
Suggested: presentation

```

```

Allow: SIPURI, ACK, CANCEL, BYE, BYACK, BYESHAKE, REFER, REFER
Security-Profile: sipuri,sgpp,spbi,signatures,sha-1,sha1,spbi,sha1
[SPICE-500]
Contact-Type: application/sdp
Contact-Uri: sip
Contact-Length: 100

```

```

vml
vml-00000000000000000000 IN IP4 1.1.1.1,111.111.111.111
vmlname vml.3
vml 0
vml 100.1.1.1:10000:111.111.111.111
vmlid: 100.70.120.100 01 04
vml.3.3
vmlCType: 100.1.1.1
vmlContact: vmlname vml
vmlContact: mandatory local vmlname
vmlContact: vml vmlname vmlname
vmlContact: 01 000
vmlContact: 01 vmlname vml.3.3.7 vmlname vml
vmlContact: 01 vmlname vmlname

```

HTTP/2 100 Baying

```

Via: SIP/2.0/UDP sipuri@domain.com branch=00000000000000000000
Via: SIP/2.0/UDP sipuri@domain.com branch=00000000000000000000
Via: SIP/2.0/UDP [1.1.1.1:111.111.111.111] rport=00000000000000000000
Content-Length: 100
To: vmlname vml.3.3.3.3.3.3
From: +111.111.111.111@domain.com; tag=00000000000000000000
Call-ID: 70000000000000000000000000000000
CSeq: 10-100000
Contact-Length: 0

```

HTTP/2 200 sipuri@domain.com SIP/2.0

```

Via: SIP/2.0/UDP sipuri@domain.com branch=00000000000000000000
Via: SIP/2.0/UDP sipuri@domain.com branch=00000000000000000000
Via: SIP/2.0/UDP sipuri@domain.com branch=00000000000000000000
Via: SIP/2.0/UDP [1.1.1.1:111.111.111.111] rport=00000000000000000000
Content-Length: 0
Max-Expires: 0
Record-Route: <udp:sipuri@domain.com>, <udp:sipuri@domain.com>,
<udp:sipuri@domain.com>
To: vml.3.3.3.3.3.3
From: <udp:sipuri@domain.com>; tag=00000000000000000000
P-Asserted-Identity: "Alice" <111.111.111.111@domain.com>,
<vml.3.3.3.3.3.3>
P-Asserted-Identity: 0000 . . .
Priority: none
P-Preferred-User-Agent: . . .
P-Call-Id-Party-Id: <udp:sipuri@domain.com>
P-Asserted-Identity: 000000000000000000000000000000000000000000000000000

```



```

!-
!- 0
!- 0-00 100 1; 0000-0000-0000-0000
!- 0000: 000 000000 00 00
!- 00 00 0
!- 00000000 0000 0000
!- 00000000 000000 0000
!- 0000 000 00000000 0000 00000000
!- 00000000 0000
!- 00000000 0000 0000 0 0 0 0 00000000
!- 00000000 00 00000000-0000

FROM: sdp (<0000> user @00000000 /PVT) requestgroup: SIPV2.0
Via: sdp/<0,UDP> (<0000:1111:0000:0000> /000) /000 requestgroup
  (<0000000000000000>)
Max-Forward: 50
Route: <sdg:0000, 00000000, com:00, <0000:00000000, com:00,
  <sdg:0000, 00000000, com:00>
To: <000:00-000-000-0000> /00-000000
From: <sdg:0000000000000000, com> /000000000000
P-Asserted-Identity: . . .
Call-ID: 000000000000000000000000
Seq: 0-0000
Contact: <sdg (<0000:1111:0000:0000> /000) requestgroup>
  000000 000000
Require: presence
Allow: SIPURI, ACK, CANCEL, BYE, OPTIONS, MESSAGE, NOTIFY, REFER
Media: 0000 IS SIPURI
Content-Length: 0

FROM: sdp (<0000> user @00000000 /PVT) requestgroup: SIPV2.0
Via: SIPV2, S/UDP (<0000:1111:0000:0000> /000) /000 requestgroup
  (<0000000000000000>)
Max-Forward: 50
Route: <sdg:0000, 000000, 000000, <0000:0000:000000, com:00>
To: <000:00-000-000-0000> /00000000
From: <sdg:0000000000000000, com> /00-0000000000
P-Asserted-Identity: . . .
Call-ID: 000000000000000000000000
Seq: 0-0000
Contact: <sdg (<0000:1111:0000:0000> /000) requestgroup>
  000000 000000
Require: presence
Allow: SIPURI, ACK, CANCEL, BYE, OPTIONS, MESSAGE, NOTIFY, REFER
Media: 0000 IS SIPURI
Content-Length: 0

FROM: sdp (<0000> user @00000000 /PVT) requestgroup: SIPV2.0
Via: SIPV2, S/UDP (<0000:1111:0000:0000> /000) /000 requestgroup
  (<0000000000000000>)
Max-Forward: 50
Route: <sdg:0000, 000000, 000000, <0000:0000:000000, com:00>
To: <000:00-000-000-0000> /00000000
From: <sdg:0000000000000000, com> /00-0000000000
P-Asserted-Identity: . . .
Call-ID: 000000000000000000000000
Seq: 0-0000
Contact: <sdg (<0000:1111:0000:0000> /000) requestgroup>
  000000 000000
Require: presence
Allow: SIPURI, ACK, CANCEL, BYE, OPTIONS, MESSAGE, NOTIFY, REFER
Media: 0000 IS SIPURI
Content-Length: 0

```

```

PLA: SIP/2.0/UDP 200.1.114.200:5060;branch=0200000000
PLA: SIP/2.0/UDP [ ]:5060;1131;1132;1133;1243;comp=sdppgroup
      (branch=0200000000)
Seq-Number: 48
To: <0>;*1-314-555-1111;tag=281334
From: sdppgroup@1418.com;tag=1418000000
P-Asserted-Identity: 704512345678901234567890
Call-ID: 281334
Contact: sdpp [ ]:5060;1131;1132;1133;1243;comp=sdppgroup
Supports: 100rel,
          presence
Allow: INVITE, ACK, CANCEL, BYE, REFER, MESSAGE, NOTIFY, UPDATE
Max-Expires: 3600;min:300;max:3600;ref:RFC3261;comp=sdppgroup
Content-Length: 0

```

```

PLA: SIP/2.0/UDP 200.1.114.200:5060;branch=0200000000;SIP/2.0
PLA: SIP/2.0/UDP 200.1.114.200:5060;branch=0200000000
PLA: SIP/2.0/UDP 200.1.114.200:5060;branch=0200000000
PLA: SIP/2.0/UDP [ ]:5060;1131;1132;1133;1243;comp=sdppgroup
      (branch=0200000000)
Seq-Number: 49
To: <0>;*1-314-555-1111;tag=281334
From: sdppgroup@1418.com;tag=1418000000
P-Asserted-Identity: 704512345678901234567890;id=
      http://aa.com/
Call-ID: 704512345678901234567890
Seq: 49;PLA
Contact: 704 [ ]:5060;1131;1132;1133;1243;comp=sdppgroup
Supports: 100rel,
          presence
Allow: INVITE, ACK, CANCEL, BYE, REFER, MESSAGE, NOTIFY, UPDATE
Max-Expires: 3600;min:300
Content-Length: 0

```

```

seq=0;tag=0
PLA: SIP/2.0/UDP 200.1.114.200:5060;branch=0200000000
PLA: SIP/2.0/UDP 200.1.114.200:5060;branch=0200000000
PLA: SIP/2.0/UDP 200.1.114.200:5060;branch=0200000000
PLA: SIP/2.0/UDP [ ]:5060;1131;1132;1133;1243;comp=sdppgroup
      (branch=0200000000)
To: <0>;*1-314-555-1111;tag=281334
From: sdppgroup@1418.com;tag=1418000000
P-Asserted-Identity: 704512345678901234567890
Call-ID: 704512345678901234567890
Seq: 49;PLA
Content-Length: 0

```

```

SIP/2.0;tag=0
PLA: SIP/2.0/UDP 200.1.114.200:5060;branch=0200000000;SIP/2.0

```

```

Vlan: 819/0, 8/80P (acct /141/0)acct, com=000000-010000000000
Vlan: 819/0, 8/80P (s: 0000/1110/0110/0011) 0043 ccompndgroup
  /0000000000000000
Tel: mail +1-314-555-2114; tag=pl1234
From: +141 6011414000000000.com; tag=pl1000000000
P-Access-Name=0000-0000-0000
  /0000-0000-0000-0000
Call-ID: Tel-819/0/0000-000000-000000
Msg: 00-0000
Context=000000; 0

```

```

819/0 > 100 00
Vlan: 819/0, 8/80P (acct /141/0)acct, com=000000-010000000000
Vlan: 819/0, 8/80P (s: 0000/1110/0110/0011) 0043 ccompndgroup
  /0000000000000000
Tel: mail +1-314-555-2114; tag=pl1234
From: +141 6011414000000000.com; tag=pl1000000000
Call-ID: Tel-819/0/0000-000000-000000
Msg: 00-0000
Context=000000; 0

```

```

819/0 > 100 00
Vlan: 819/0, 8/80P (s: 0000/1110/0110/0011) 0043 ccompndgroup
  /0000000000000000
Tel: mail +1-314-555-2114; tag=pl1234
From: +141 6011414000000000.com; tag=pl1000000000
Call-ID: Tel-819/0/0000-000000-000000
Msg: 00-0000
Context=000000; 0

```

```

***** sip: (s: 0000/0000/0000/0000) 0043 ccompndgroup 819/0-0
Vlan: 819/0, 8/80P (s: 0000/1110/0110/0011) 0043 ccompndgroup
  /0000000000000000
Req-From: 0000
From: +141 6011414000000000.com; tag=pl1000000000
  +141 6011414000000000.com; tag=pl1000000000
Tel: mail +1-314-555-2114; tag=pl1234
From: +141 6011414000000000.com; tag=pl1000000000
P-Access-Name=0000-0000-0000-0000
Call-ID: Tel-819/0/0000-000000-000000
Msg: 00-0000
msgtype: +141 6011414000000000.com; tag=pl1000000000
Supported: 0000
reqtype: 0000000000
From: 00000000-0000-0000-0000
Allow: SIP/00, ACK, INFO, SIP, OPTIONS, MESSAGE, NOTIFY, REFER
Event-ID: 000000, 0000-0000-0000-0000, 0000-0000-0000-0000
  /0000000000
Context=0000-0000-0000-0000
Context=000000; ...

```

```

end
on: 200000001 100000017 in end {+ 0000 0111-0311-0000}
P-Phone: 001
tel: 0
P-CR: 100 1; 0000-1111 0111 0311
module: 10000 000000 01 00
lang: 0 0 0
service-type: local, residential
service-type: remote, none
service-type: mandatory, local, residential
service-type: none, remote, residential
I-CPN: 01 000
url: http:// www.debian.org/~dts/compaqgroup-0000-0111-0000
I-CPN: 01 00000000 00000000-0000 00000000-0000

service-type: (+ 0000 000000 000 00) + 0111 00000 000000 000 0111
P: 0111 0000 0000 (+ 0111 0000 0000 0000 0000) 0000 00000000-0000
P: 0111 0000 0000 {+ 0000 0111 0111 0311} 0000 00000000
+ 0000-0000-00000000
Service: 00
P: 0111 0000 0000 000 0000 000 0000 000000 000 0111
P: 0000 0111 000 000 0000 0000
P: 0111 0000 0000 00000000 0000 0000 0000 0000
P-Phone: Network-Index: . . .
P-Phone: 0000 . . .
Cell-ID: 00000000000000000000
lang: 01 0000
CPN: 01 0000 1111 0111 0311 0000 0000 0000 0000
Support: local
Service: 00000000
Allow: 00000, ACK, Cancel, OK, PAUSE, MESSAGE, NEXT, UPDATE
service-type: application/vnd
Content-Length: . . .

end
on: 200000001 100000017 in end {+ 0000 0111-0311-0000}
P-Phone: 001
tel: 0
P-CR: 100 1; 0000-1111 0111 0311
module: 10000 000000 01 00
lang: 0 0 0
service-type: local, residential
service-type: remote, none
service-type: mandatory, local, residential
service-type: none, remote, residential
I-CPN: 01 000
url: http:// www.debian.org/~dts/compaqgroup-0000-0111-0000
I-CPN: 01 00000000 00000000-0000 00000000-0000

service-type: (+ 0000 000000 000 00) + 0111 00000 000000 000 0111
P: 0111 0000 0000 (+ 0000 0111 0111 0311) 0000 00000000-0000
P: 0111 0000 0000 {+ 0000 0111 0111 0311} 0000 00000000
+ 0000-0000-00000000
Service: 00
P: 0111 0000 0000 000 0000 000 0000 000000 000 0111
P: 0000 0111 000 000 0000 0000
P: 0111 0000 0000 00000000 0000 0000 0000 0000
P-Phone: Network-Index: . . .
P-Phone: 0000 . . .
Cell-ID: 00000000000000000000
lang: 01 0000
CPN: 01 0000 1111 0111 0311 0000 0000 0000 0000
Support: local
Service: 00000000
Allow: 00000, ACK, Cancel, OK, PAUSE, MESSAGE, NEXT, UPDATE
service-type: application/vnd
Content-Length: . . .

end
on: 200000001 100000017 in end {+ 0000 0111-0311-0000}
P-Phone: 001
tel: 0
P-CR: 100 1; 0000-1111 0111 0311
module: 10000 000000 01 00
lang: 0 0 0
service-type: local, residential
service-type: remote, none
service-type: mandatory, local, residential
service-type: none, remote, residential
I-CPN: 01 000
url: http:// www.debian.org/~dts/compaqgroup-0000-0111-0000
I-CPN: 01 00000000 00000000-0000 00000000-0000

service-type: (+ 0000 000000 000 00) + 0111 00000 000000 000 0111
P: 0111 0000 0000 (+ 0000 0111 0111 0311) 0000 00000000-0000
P: 0111 0000 0000 {+ 0000 0111 0111 0311} 0000 00000000
+ 0000-0000-00000000
Service: 00
P: 0111 0000 0000 000 0000 000 0000 000000 000 0111
P: 0000 0111 000 000 0000 0000
P: 0111 0000 0000 00000000 0000 0000 0000 0000
P-Phone: Network-Index: . . .
P-Phone: 0000 . . .
Cell-ID: 00000000000000000000
lang: 01 0000
CPN: 01 0000 1111 0111 0311 0000 0000 0000 0000
Support: local
Service: 00000000
Allow: 00000, ACK, Cancel, OK, PAUSE, MESSAGE, NEXT, UPDATE
service-type: application/vnd
Content-Length: . . .

end
on: 200000001 100000017 in end {+ 0000 0111-0311-0000}
P-Phone: 001
tel: 0
P-CR: 100 1; 0000-1111 0111 0311
module: 10000 000000 01 00
lang: 0 0 0
service-type: local, residential
service-type: remote, none
service-type: mandatory, local, residential
service-type: none, remote, residential
I-CPN: 01 000
url: http:// www.debian.org/~dts/compaqgroup-0000-0111-0000
I-CPN: 01 00000000 00000000-0000 00000000-0000

service-type: (+ 0000 000000 000 00) + 0111 00000 000000 000 0111
P: 0111 0000 0000 (+ 0000 0111 0111 0311) 0000 00000000-0000
P: 0111 0000 0000 {+ 0000 0111 0111 0311} 0000 00000000
+ 0000-0000-00000000
Service: 00
P: 0111 0000 0000 000 0000 000 0000 000000 000 0111
P: 0000 0111 000 000 0000 0000
P: 0111 0000 0000 00000000 0000 0000 0000 0000
P-Phone: Network-Index: . . .
P-Phone: 0000 . . .
Cell-ID: 00000000000000000000
lang: 01 0000
CPN: 01 0000 1111 0111 0311 0000 0000 0000 0000
Support: local
Service: 00000000
Allow: 00000, ACK, Cancel, OK, PAUSE, MESSAGE, NEXT, UPDATE
service-type: application/vnd
Content-Length: . . .

end

```

```

Via: SIP/2.0/SIP 192.168.1.100:5060 branch-01001@koolhaas
Via: SIP/2.0/SIP 192.168.1.100:5060 branch-01001@koolhaas
Via: SIP/2.0/SIP 1: 1999:1111:2222:3333:2043:comp@group
(branch-01001@koolhaas)
Max-Forwards: 49
Route: <url>group, koolhaas, com.ly
To: <url>:1:111-111-111:111:1:1:1:1
From: <url>:1:111-111-111-111:111:1:1:1:1:1:1
Call-ID: 764814749:14749:14749:14749
Seq: 17 000000
Contact: <url>:1: 1999:1111-2222-3333:2043:comp@group
supported: 100rel
Allow: INVITE, ACK, CANCEL, BYE, OPTIONS, MESSAGE, REFER, NOTIFY
Content-Type: application/sdp
Content-Length: ...

```

```

v0
m= 100000000 1000000000 0 0 1: 1999:1111-2222-3333
rPhone-Call
m0 0
0-00 100 1: 1999:1111-2222-3333
mmedia: 100:0 100:000 01 01
0-00 17 0
mcomp@group local, mandatory
mcomp@group 100:00 100
mcomp@group mandatory local, mandatory
mcomp@group 100:00 100:0000 0000
mcomp@group 100:0000 1, 1, 1 (mcomp@group)
0-00000 00 00100000-0000

```

```

SIP/2.0/udp 1: 1999:1111-2222-3333:2043:comp@group SIP/2.0
Via: SIP/2.0/SIP 192.168.1.100:5060 branch-01001@koolhaas
Via: SIP/2.0/SIP 192.168.1.100:5060 branch-01001@koolhaas
Via: SIP/2.0/SIP 1: 1999:1111:2222:3333:2043:comp@group
(branch-01001@koolhaas)
Max-Forwards: 49
To: <url>:1:111-111-111:111:1:1:1:1
From: <url>:1:111-111-111-111:111:1:1:1:1:1:1
Call-ID: 764814749:14749:14749:14749
Seq: 17 000000
Contact: <url>:1: 1999:1111-2222-3333:2043:comp@group
supported: 100rel
Requires: presence
Allow: INVITE, ACK, CANCEL, BYE, OPTIONS, MESSAGE, REFER, NOTIFY
Content-Type: application/sdp
Content-Length: ...

```

```

v0
m= 100000000 1000000000 0 0 1: 1999:1111-2222-3333
rPhone-Call
m0 0

```



```

i=OK 200 1; 00001110 0110 0111
message: 200 OK 200 OK 01 01
i=OK 01 0
sequence-num local number
i=OK 200 200 0000 0000
sequence-number local number
order-type none none number number
sequence 01 00
sequence-ID none none 0 1 2 3 (numbered)
sequence 01 01 number-number

INFO > 100 OK
Fl: SIP/2.0/UDP 20001.2000001.com/50605-0100000000
Fl: SIP/2.0/UDP 20001.2000001.com/50605-0100000000
Fl: SIP/2.0/UDP 20001.2000001.com/50605-0100000000
Fl: SIP/2.0/UDP 1; 00001110 0110 0111; 200 200010000
  branch:0100000000
To: <01>+1-314-555-1212>; tag=01104
From: <01>20001.2000001.com>; tag=0100000000
Call-ID: 2000100001-2000000000000000
Msg: IT 010000
Content-Type: application/sdp
Content-Length: ...

v=0
o= 200000101 200000102 01 01 00000000 0000 0000
s=
t=0
i=OK 200 1; 00001110 0000 0000
message: 200 OK 200 OK 01 01
i=OK 01 0
sequence-num local number
sequence-num 20000 20000
sequence-number local number
order-type none none number number
sequence 01 00
sequence-ID none none 0 1 2 3 (numbered)
i=OK 200 01 01 number-number

INFO > 100 OK
Fl: SIP/2.0/UDP 20001.2000001.com/50605-0100000000
Fl: SIP/2.0/UDP 20001.2000001.com/50605-0100000000
Fl: SIP/2.0/UDP 1; 00001110 0110 0111; 200 200010000
  branch:0100000000
To: <01>+1-314-555-1212>; tag=01104
From: <01>20001.2000001.com>; tag=0100000000
Call-ID: 2000100001-2000000000000000
Msg: IT 010000
Content-Type: application/sdp
Content-Length: ...

v=0

```



```

*Call to 909.914.888.4144; tag=pl1234
To: +1-914-888-4144; tag=pl1234
From: +1-914-888-4144; tag=pl1234
Call-ID: 7648123456789012345678901234
Call: 20-09-2008
Contact: +1-914-888-4144; tag=pl1234
Allow: INVITE, ACK, CANCEL, BYE, PRACK, REGISTER, REFER, UPDATE
Contact-Info: 0

```

```

SIP/1.0 100 Trying
To: SIP/1.0/UDP [x.x.x.x:1111] (1111); 2008-09-20T00:00:00Z
From: SIP/1.0/UDP [x.x.x.x:1111] (1111); 2008-09-20T00:00:00Z
Contact-Info: +1-914-888-4144; tag=pl1234
To: +1-914-888-4144; tag=pl1234
From: +1-914-888-4144; tag=pl1234
Call-ID: 7648123456789012345678901234
Call: 20-09-2008
Contact: +1-914-888-4144; tag=pl1234
Allow: INVITE, ACK, CANCEL, BYE, PRACK, REGISTER, REFER, UPDATE
Contact-Info: 0

```

```

SIP/1.0 100 OK
To: SIP/1.0/UDP [x.x.x.x:1111] (1111); 2008-09-20T00:00:00Z
From: SIP/1.0/UDP [x.x.x.x:1111] (1111); 2008-09-20T00:00:00Z
Call-ID: 7648123456789012345678901234
Call: 20-09-2008
Contact: +1-914-888-4144; tag=pl1234
To: +1-914-888-4144; tag=pl1234
From: +1-914-888-4144; tag=pl1234
Call-ID: 7648123456789012345678901234
Call: 20-09-2008
Contact: +1-914-888-4144; tag=pl1234
Allow: INVITE, ACK, CANCEL, BYE, PRACK, REGISTER, REFER, UPDATE
Contact-Info: 0

```

```

SIP/1.0 100 OK
To: SIP/1.0/UDP [x.x.x.x:1111] (1111); 2008-09-20T00:00:00Z
From: SIP/1.0/UDP [x.x.x.x:1111] (1111); 2008-09-20T00:00:00Z
Call-ID: 7648123456789012345678901234
Call: 20-09-2008
Contact: +1-914-888-4144; tag=pl1234
To: +1-914-888-4144; tag=pl1234
From: +1-914-888-4144; tag=pl1234
Call-ID: 7648123456789012345678901234
Call: 20-09-2008
Contact: +1-914-888-4144; tag=pl1234
Allow: INVITE, ACK, CANCEL, BYE, PRACK, REGISTER, REFER, UPDATE
Contact-Info: 0

```

```

Client: 15 18PTE
Server to: udp:[]:1999 [aaaa:bbbb:cccc]:5000 [comp:group]
Expires: 1800s
Allow: SIPTEL, ACK, CANCEL, BYE, BLANK, MESSAGE, REFER, UPDATE
CAPABILITY: 0

SIP/1.0 200 OK
Via: SIP/2.0/UDP [a:b:c:d] [e:f:g:h]:5000 [comp:group]
Via: SIP/2.0/UDP [i:j:k:l:m]:5000 [n:o:p:q:r]:5000 [comp:group]
[branch:0] [branch:1]
From: SIP/2.0/UDP [a:b:c:d] [e:f:g:h]:5000 [comp:group]
To: SIP/2.0/UDP [i:j:k:l:m]:5000 [n:o:p:q:r]:5000
Call-ID: 7654321098765432109876543210
Client: 15 18PTE
Server to: udp:[]:1999 [aaaa:bbbb:cccc]:5000 [comp:group]
Expires: 1800s
Allow: SIPTEL, ACK, CANCEL, BYE, BLANK, MESSAGE, REFER, UPDATE
CAPABILITY: 0

SIP/1.0 200 OK
Via: SIP/2.0/UDP [i:j:k:l:m]:5000 [n:o:p:q:r]:5000 [comp:group]
[branch:0] [branch:1]
From: SIP/2.0/UDP [a:b:c:d] [e:f:g:h]:5000 [comp:group]
To: SIP/2.0/UDP [i:j:k:l:m]:5000 [n:o:p:q:r]:5000
Call-ID: 7654321098765432109876543210
Client: 15 18PTE
Server to: udp:[]:1999 [aaaa:bbbb:cccc]:5000 [comp:group]
Expires: 1800s
Allow: SIPTEL, ACK, CANCEL, BYE, BLANK, MESSAGE, REFER, UPDATE
CAPABILITY: 0

ACK 8261: 8888 [aaa:bbb:ccc]:5000 [xyz:000-111111] SIP/1.0
Via: SIP/2.0/UDP [i:j:k:l:m]:5000 [n:o:p:q:r]:5000 [comp:group]
[branch:0] [branch:1]
From: SIP/2.0/UDP [a:b:c:d] [e:f:g:h]:5000 [comp:group]
To: SIP/2.0/UDP [i:j:k:l:m]:5000 [n:o:p:q:r]:5000
Call-ID: 7654321098765432109876543210
Client: 15 ACK
Server to: SIP/1.0: 8888 [aaa:bbb:ccc]:5000 [xyz:000-111111] SIP/1.0
Allow: SIPTEL, ACK, CANCEL, BYE, BLANK, MESSAGE, REFER, UPDATE
CAPABILITY: 0

```


10.8 Call Setup Example with Two Proxies

This section contains the complete message flow shown in Figure 2.2.

- ```

81 INVITE sip:1000@192.168.100.100:5060 SIP/2.0
Via: SIP/2.0/UDP 192.101.100.100:5060 (branch=0)diffserv176
max-forwards: 70
To: Helenenberg <helenenberg@192.101.100.100>
From: B. Helmschlager <hls@192.101.100.100> ; tag=1
Call-ID: 180000.100.100.100
Seq: 1 30000
MIME-Version: application/sdp
Contact: <hls@192.101.100.100>
Content-Type: application/sdp
Content-Length: 150

m=0
a=sendrecv 100000000 100000000 0 0 100 100 100 100
rPhone Call
call 0
P-OR 192.101.100.100
a=rtcp:192.101.100.100
b=ASPS 0 RTCP/0

82 INVITE sip:1000@192.101.100.100:5060 SIP/2.0
Via: SIP/2.0/UDP proxy.austriai.net (branch=0)diffserv176
Via: SIP/2.0/UDP 192.101.100.100:5060 (branch=0)diffserv176
max-forwards: 60
To: Helenenberg <helenenberg@192.101.100.100>
From: B. Helmschlager <hls@192.101.100.100> ; tag=1
Call-ID: 180000.100.100.100
Seq: 1 30000
Contact: <hls@192.101.100.100>
Content-Type: application/sdp
Content-Length: 150

m=0
P-OR 192.101.100.100
a=rtcp:192.101.100.100
b=ASPS 0 RTCP/0
a=rtcp:192.101.100.100

83 SIP/2.0 200 Ringing
Via: SIP/2.0/UDP proxy.austriai.net (branch=0) diffserv176
Via: SIP/2.0/UDP 192.101.100.100:5060 (branch=0)diffserv176
Via: SIP/2.0/UDP 192.101.100.100:5060 (branch=0)diffserv176
To: Helenenberg <helenenberg@192.101.100.100> ; tag=1
From: B. Helmschlager <hls@192.101.100.100> ; tag=1
Call-ID: 180000.100.100.100
Seq: 1 30000

```

```

Contact: <sip:walter.beiswenger@100.201.200.201>
Contact-Length: 0

#1 SIP/2.0 200 OK
Via: SIP/2.0/UDP 100.201.192.201; branch=z9hGzB7P
To: beiswenger <sip:walter.beiswenger@100.201.200.201>
From: H. Beiswenger <tel:100201200201, user=beiswenger>
Call-ID: 100201192.201.192
CSeq: 1 INVITE
Contact: <sip:walter.beiswenger@100.201.200.201>
Contact-Length: 0

#2 SIP/2.0 200 OK
Via: SIP/2.0/UDP proxy-stand.in.tum.de; branch=z9hGzB7P
Via: SIP/2.0/UDP 100.201.192.201; branch=z9hGzB7P
To: beiswenger <sip:walter.beiswenger@100.201.200.201>
From: H. Beiswenger <tel:100201200201, user=beiswenger>
Call-ID: 100201192.201.192
CSeq: 1 INVITE
Contact: <sip:walter.beiswenger@100.201.200.201>
Contact-Type: application/sdp
Contact-Length: 159

v=0
o=beiswenger 20020119 20020119 IN 201 200.201.200.201
s=beiswenger
t=0
m=audio 49170 RTP/AVP 0
a=rtpmap:0 PCMA/8000

#3 SIP/2.0 200 OK
Via: SIP/2.0/UDP 100.201.192.201; branch=z9hGzB7P
To: beiswenger <sip:walter.beiswenger@100.201.200.201>
From: H. Beiswenger <tel:100201200201, user=beiswenger>
Call-ID: 100201192.201.192
CSeq: 1 INVITE
Contact: <sip:walter.beiswenger@100.201.192.201>
Contact-Type: application/sdp
Contact-Length: 159

v=0
o=beiswenger 20020119 20020119 IN 201 200.201.200.201
s=beiswenger
t=0
m=audio 49170 RTP/AVP 0
a=rtpmap:0 PCMA/8000

#4 200 201:walter.beiswenger@100.201.192.201 SIP/2.0

```



- ```

Flg: SIP/2.0/UDP 192.161.192.201:5060 branch=6266666410
Via: SIP/2.0/UDP 192.161.192.201:5060 branch=6266666410
From: E. Asterisk@quepasa.com
To: M. Asterisk@quepasa.com
Call-ID: 192161192192192
Msg: 1 OK
Content-Length: 0

```
- 81
- ```

Flg: SIP/2.0/UDP 192.161.192.201:5060 branch=6266666410
Via: SIP/2.0/UDP 192.161.192.201:5060 branch=6266666410
From: E. Asterisk@quepasa.com
To: M. Asterisk@quepasa.com
Call-ID: 192161192192192
Msg: 200 OK
Content-Length: 0

```
- 82
- ```

Flg: SIP/2.0/UDP 192.161.192.201:5060 branch=6266666410
Via: SIP/2.0/UDP 192.161.192.201:5060 branch=6266666410
From: E. Asterisk@quepasa.com
To: M. Asterisk@quepasa.com
Call-ID: 192161192192192
Msg: 200 OK
Content-Length: 0

```

10.3 SIP Presence and Instant Message Example

This section contains the call flow details of Figure 2-4.

- 81
- ```

From: sip:alice@quepasa.com
Flg: SIP/2.0/UDP 192.161.192.201:5060
Via: SIP/2.0/UDP 192.161.192.201:5060
From: F. L. Doolittle <fl@doolittle.com>
To: M. Asterisk@quepasa.com
Call-ID: 192161192192192
Msg: 100 TRYING
Allow: INVITE, OPTIONS, BYE, CANCEL, ACK, PRACK, MESSAGE, REFER
Contact: sip:alice@quepasa.com
Content-Length: 0

```
- 82
- ```

Flg: SIP/2.0/ACK
Flg: SIP/2.0/UDP 192.161.192.201:5060
Via: SIP/2.0/UDP 192.161.192.201:5060
From: F. L. Doolittle <fl@doolittle.com>
To: M. Asterisk@quepasa.com
Call-ID: 192161192192192
Msg: 200 OK

```

- Call: 112 BRIDGE
Allow-Events: presence
Allow: ACK, UPDATE, CAPAB, RSC, NOTIFY, PRESENCE, MESSAGE
Contact: sdp://gandalf@cs.princeton.edu/; transport=tcp
FROM: 21000000
Expires: 300
Content-Length: 0
- #1 NOTIFY sdp://gandalf@cs.princeton.edu/; SIP/1.0
Via: SIP/1.0/TCP 0.0.0.0/32; protocol=try; seq=3000
; branch=0000000000
Max-Forwards: 70
To: F. L. Chaboyer <flc@caltechcsd.caltech.edu>; tag=0171
From: M. Polansky <mpolansky@caltech.edu>; tag=0140
Call-ID: 000011400000000000000000
Call: 2008 NOTIFY
Allow: ACK, UPDATE, CAPAB, RSC, NOTIFY, PRESENCE, MESSAGE
Allow-Events: dialog
Contact: sdp://gandalf@cs.princeton.edu/; transport=tcp
SUBSCRIBE-TO: STATES: active, onhold, noanswer
Event: presence
CONTEXT-TYPE: SULLOFT/STATE/ONHOLD-000
Expires: 300
Content-Length: 0
- <Call: version=1.0> received SIP/1.0
<sequence number> seq=3000; protocol=try; tag=0171
<call-id> sdp://gandalf@cs.princeton.edu/;
<seq: 0000000000>
<CSeq>
<branch=0000000000>
<CSeq>
<tag>
<Expires>
- #2 SIP/1.0 200 OK
Via: SIP/1.0/TCP 0.0.0.0/32; protocol=try; seq=3000
; branch=0000000000; seq=0000000000; tag=0171
To: F. L. Chaboyer <flc@caltechcsd.caltech.edu>; tag=0171
From: M. Polansky <mpolansky@caltech.edu>; tag=0140
Call-ID: 000011400000000000000000
Call: 2008 NOTIFY
Content-Length: 0
- #3 NOTIFY sdp://gandalf@cs.princeton.edu/; SIP/1.0
Via: SIP/1.0/TCP 0.0.0.0/32; protocol=try; seq=3000
; branch=0000000000
Max-Forwards: 70
To: F. L. Chaboyer <flc@caltechcsd.caltech.edu>; tag=0171
From: M. Polansky <mpolansky@caltech.edu>; tag=0140
Call-ID: 000011400000000000000000
Call: 2008 NOTIFY

```
ALLOW: ACK, DDPTE, DANCE, HIR, HOPPE, JUNGCHUN, KIMJONG  
ALLOW: frames: generate
```

```
CONTENT: +EIS-8-20041008101-20040411101-000-000-001-000-  
Content-Disposition: inline; filename=sample2800
```

```
EMAIL: 0000000
```

```
Content-Type: application/vnd.epson.plafont  
Content-Length: 100
```

```
email-version: 1.0; charset=utf-8; <mailto:0000000@  
<mailto:0000000@0000000.com> [mailto:0000000@0000000.com]  
<mailto:0000000@0000000.com>  
<mailto:0000000@0000000.com>  
<mailto:0000000@0000000.com>  
<mailto:0000000@0000000.com>  
<mailto:0000000@0000000.com>  
<mailto:0000000@0000000.com>  
<mailto:0000000@0000000.com>  
<mailto:0000000@0000000.com>  
<mailto:0000000@0000000.com>  
<mailto:0000000@0000000.com>
```

80 01/01 0 200 OK

```
File E:\E:\S\TCP-02\01-00000001\01-00000001
```

```
(mailto:0000000@0000000.com) [mailto:0000000@0000000.com]
```

```
To: B. L. Halsey <mailto:0000000@0000000.com> [mailto:0000000@0000000.com]
```

```
From: B. L. Halsey <mailto:0000000@0000000.com> [mailto:0000000@0000000.com]
```

```
Subject: 0000000-0000000-0000000-0000000
```

```
Content-Length: 0
```

81 000000 000 0 20041008101-20040411101-000-000-001-000

```
File E:\E:\S\TCP-02\00000001\00000001 [mailto:0000000@0000000.com]
```

```
(mailto:0000000@0000000.com)
```

```
Content-Type: TO
```

```
To: B. L. Halsey <mailto:0000000@0000000.com> [mailto:0000000@0000000.com]
```

```
From: B. L. Halsey <mailto:0000000@0000000.com> [mailto:0000000@0000000.com]
```

```
Subject: 0000000-0000000
```

```
Content-Type: 0000000
```

```
ALLOW: ACK, DDPTE, DANCE, HIR, HOPPE, JUNGCHUN, KIMJONG
```

```
CONTENT-TYPE: 0000000
```

```
Content-Length: 0
```

B. Halsey

82 01/01 0 200 OK

```
File E:\E:\S\TCP-02\00000001\00000001 [mailto:0000000@0000000.com]
```

```
(mailto:0000000@0000000.com) [mailto:0000000@0000000.com]
```

```
To: B. L. Halsey <mailto:0000000@0000000.com> [mailto:0000000@0000000.com]
```

```
From: B. L. Halsey <mailto:0000000@0000000.com> [mailto:0000000@0000000.com]
```

```
Subject: 0000000-0000000
```

```
Content-Type: 0000000
```

```
Content-Length: 0
```

```

#1  SIP/2.0;branch=200;method=INVITE;seq=11771;C
Via: SIP/2.0/UDP;skat.proxmail.ty, seq=1010
   ;branch=200;method=INVITE
Max-Forwards: 70
To: P. L. Chelidze <pl@chelidze.com>
From: B. Polonsky <bp@proxmail.ty, seq=1010
   ;branch=200>
CSeq: 11771;METHOD
Allow: INVITE, OPTIONS, MESSAGE, NOTIFY, SUBSCRIBE, UNSUBSCRIBE
Content-Type: text/plain
Content-Length: 0

```

Skat, hello there to you, too

```

#2  SIP/2.0;branch=0
Via: SIP/2.0/UDP;skat.proxmail.ty, seq=1010
   ;branch=0;method=ACK;seq=1010
To: P. L. Chelidze <pl@chelidze.com>, seq=1010;branch=0
From: B. Polonsky <bp@proxmail.ty, seq=1010
   ;branch=0>
CSeq: 1010;METHOD
Content-Length: 0

```

References

- [1] Rosenberg, J. et al. "SIP Session Initiation Protocol," RFC 3261, 2002.
- [2] Johnson, A. et al. "Basic SIP Call Flow Examples," IETF Internet-Draft, Work in Progress, April 2005.
- [3] Johnson, A. et al. "SIP/2.0 Call Flow Examples," IETF Internet-Draft, Work in Progress, April 2005.
- [4] Schierman, H., and H. Agard. "Session Initiation Protocol (SIP)-H.323 Interworking Requirements," IETF Internet-Draft, Work in Progress, February 2003.

11

Future Directions

This chapter will discuss some future areas of work in SIP-related working groups in the IETF. Instead of attempting to list and discuss a snapshot of current activity in the IETF, the reader should gather the information directly from the IETF itself. Table 11.1 lists the chapter Web pages for the three most important SIP working groups: SIP, SIPPING, and SIMPLE. The chapter page for each working group lists the deliverables of the group along with RFCs (finished documents) and Internet-Drafts (works in progress). Only Internet-Drafts that have been adopted as official working group items are listed on these Web pages—these are the documents most likely to become RFCs in the near future. The Web page also contains information about joining the working group e-mail list, which discusses the list use of Internet-Drafts.

Finally, one can search the IETF Internet-Draft archives for documents relating to SIP at <http://www.ietf.org>. However, be warned! There are many, many documents and most will likely never be published as an RFC—always check sources familiar with the working group activity before assuming that an Internet-Draft not listed on a working group chapter page is likely to ever become an Internet-Standard.

11.1 SIP, SIPPING, and SIMPLE Working Group Design Teams

From time to time working groups use a design team to work on a particular topic. The teams report regularly to the working group and have a working group chair who provides guidance. The approach has been successful in the past and will likely result in new RFCs being produced by the design teams described in the next sections.

Table 11.1
Latest Status of IETF-Related SIP Work

Working Group	Web Page URL	Area
SIP	http://www.ietf.org/html.charters/sip-charter.html	SIP protocol extensions
SIPMG	http://www.ietf.org/html.charters/sipmg-charter.html	SIP requirements, best current practice documents, and usage
SIPMLP	http://www.ietf.org/html.charters/sipmlp-charter.html	SIP for IM and presence
XCON	http://www.ietf.org/html.charters/xcon-charter.html	Serialized conferencing

Some of the referenced Internet-Drafts in this chapter may have become RFCs or have changed. Useful archives of expired Internet-Drafts are available at <http://www.ietfarchive.com/sipmg>, <http://www.ietfarchive.com/sipping>, <http://www.ietfarchive.com/sipmlp>, <http://www.ietfarchive.com/xcon>, and <http://www.ietf.org/html.charters/>.

11.1.1 SIP and Hearing Impairment Design Team

This design team is developing requirements, guidelines, and directions for SIP to support the needs of hearing-impaired users. The design team has produced a requirements document [1] that includes scenarios and requirements. The work has also produced general requirements on transcoding, which have wider scope than the hearing-impaired.

11.1.2 Conferencing Design Team

This design team is working on issues relating to tightly-coupled, multiparty conferencing. The work so far has produced a requirements [2], framework [3], and call control [4] documents that are SIPFING working group items. The framework includes voice, video, IM, and collaboration. Other items include the definition of the 3a Cocon framework [5] and the conference event package [6].

The design team is also working on areas relating to standardizing SMI managers for conference and media mixing policy. This will allow an intelligent end point or intermediate to set up and configure a multimedia conferencing service. The media policy manipulation will allow an end point to control their view and mixing of the overall conference. This work will be standardized in the new XCON (Serialized Conferencing) IETF working group. However,

Intelligent SIP end points and clients will likely support these extensions and protocols.

11.3 Application Interaction Design Team

This design team is developing a framework for application interaction using SIP. It will cover areas such as dual-tone multi-frequency (DTMF) and stimulus (button-pressing) transport and integration with SIP signaling.

11.4 Emergency Calling Design Team

This design team is developing requirements for SIP support of emergency calling (911 in the United States and 112 in other parts of the world).

11.2 Other SIP Work Areas

Some other important work areas relating to SIP are discussed in the following sections.

11.2.1 Emergency Preparedness

The IETF Internet Emergency Preparedness (IEPREP) working group has been working on requirements on how to make sure that the Internet functions properly and can provide reliable communications for emergency personnel in the event of an emergency or disaster. There is a document of requirements for SIP [7] that may produce some SIP extensions and conventions.

11.2.2 Globally Routable Contact URIs

Work is currently underway to try to provide a mechanism to guarantee that a Contact URI is globally routable. While in general a Contact URI is globally routable, in the presence of NATs, firewalls, and caching proxies, this may not be so. Contributing to the problem is that a UA may not know if its Contact URI is globally routable or not. Current work is discussing the requirements, which, when agreed upon, can be translated into a SIP extension or convention or other.

11.2.3 Service Examples

This service examples document [8] shows how intelligent end points can implement a set of service commonly found in PBX switches and in service provider Center offerings. The set of features includes call hold, call park, call

picking, auto-recall, music on hold, attended and unattended transfers, and others. All of the flows use standard SIP operations and call control extensions.

11.3 SIP Instant Message and Presence Work

In the SIMPLE working group, there has been much work in the area of instant message systems, similar to that shown in Figure 4.10 in which an IMUTE sets up an IM session between the end points. The current approach being developed was a SIP-like protocol called Message Session Relay Protocol (MSRP) [9], which allows the IM session to be established either directly end-to-end or through a relay necessary for logging or NAT/firewall traversal.

The SIMPLE working group is also working on documents to enable all aspects of a standard presence and IM client to be implemented using standard HTTP protocols. While the basic ability is provided through the defined MESSAGE, SUBSCRIBE, and NOTIFY methods, tools to upload or publish presence information, and protocols for “body list” creation and manipulation are still needed.

Other interesting work continues in the area of “rich” presence information delivery using SIP [10]. The current presence XML formats only list contacts and describe them as “open” or “closed,” essentially on-line or off-line. This work extends so that more presence information can be automatically updated based on calendar and other applications from the network.

References

- [1] Chabon, N., et al. “New Requirements for the Session Initiation Protocol (SIP) to Support of Deaf, Hard of Hearing and Speech-Impaired Individuals,” RFC 3991, 2005.
- [2] Levin, G., and R. Egan. “High Level Requirements for Tightly Coupled SIP Conferencing,” IETF Internet-Draft, Work in Progress, April 2005.
- [3] Rosenberg, J. “A Framework for Conferencing with the Session Initiation Protocol,” IETF Internet-Draft, Work in Progress, May 2003.
- [4] Johnson, A., and G. Levin. “SIP Call Control – Conferencing for User Agents,” IETF Internet-Draft, Work in Progress, April 2005.
- [5] Rosenberg, J., H. Schulzrinne, and P. Ryznar. “Caller Preferences and Caller Capabilities for the Session Initiation Protocol (SIP),” IETF Internet-Draft, Work in Progress, March 2005.
- [6] Rosenberg, J., and H. Schulzrinne. “A Session Initiation Protocol (SIP) Form Fieldset for Conference Users,” IETF Internet-Draft, Work in Progress, June 2003.
- [7] Paterson, J. “Considerations on the IETF Requirements for SIP,” IETF Internet-Draft, Work in Progress, April 2005.

-
- [8] Johnson, G., et al., "Session Initiation Protocol Service Examples," IETF Internet-Draft, Work in Progress, March 2005.
 - [9] Campbell, B., et al., "Secure Message Sessions in SIMPLE," IETF Internet-Draft, Work in Progress, April 2003.
 - [10] Schindler, H., et al., "RPSIS—Rich Presence Information Data Format for Presence Based on the Session Initiation Protocol (SIP)," IETF Internet-Draft, Work in Progress.

Appendix A: Changes in the SIP Specification from RFC 2543 to RFC 3261

In late 2001 and early 2002, the basic SIP specification RFC 2543 was rewritten and published as RFC 3261. The new RFC was backwards compatible with the old RFC except in areas in which the old RFC was deemed to be broken or buggy. Besides these fixes, the document was also rewritten from the ground up with a new structure and new content to increase the clarity of the specification. The new structure included a redefinition of the SIP state machine (SSM) into three layers: transport, transaction, and transaction user (TU). Additional timers and timer-outs were also defined. The new specification also had related topics such as SDP handling and DNS Procedures removed and published as separate specifications RFC 3264 and RFC 3263.

This appendix describes some of the key differences between the two RFCs. Many of these differences have been mentioned throughout the book—others have not been mentioned.

Some parts of RFC 2543 were removed or deprecated from RFC 3261. In the area of security, HTTP Basic authentication was removed, as it allowed the transport of passwords in the clear (HTTP Digest does not do this and is still supported in RFC 3261). Also deprecated was Pretty Good Privacy (PGP) encryption in favor of S/MIME encryption. PGP had not been implemented and the S/MIME approach is the one currently favored by the IETF to secure bodies end-to-end of the Internet. The Authentication header field (which was used for PGP encryption) was deprecated. The Response-Key and Refer header fields were also deprecated due to security problems associated with their use.

The new RFC also deprecates the use of Via header field loop detection in favor of a required Max-Forwards header field in every request. This is because of the difficulty and processing burden on proxies to parse long sets of Via headers.

Perhaps the most major change outside of security relates to the introduction of loose routing concepts in place of the strict routing used with Route headers in RFC 2543. In the old RFC, for example, a request containing three Route headers could only route to those three proxies then to the end point—no deviations were allowed. Also, a Route header could only be built from a resolved Forward-Route header field. In RFC 3261, a request containing three loose Route headers must route through those three proxies but may also visit other proxies as needed. Also, a proxied Route header may be constructed at the start of a dialog without first receiving a Forward-Route header in dialog setup. This permits much more flexible routing for services and middle applications. They can also be used to provide virtual services and hence proxy services as discussed in Chapter 3. In addition, the mechanism used is much cleaner and understandable than the old “strict” routing, and it is also backwards compatible, even in a mixed network of strict and loose routing proxies.

A major improvement in RFC 3261 comes in the security area. Besides the deprecation of HTTP Basic authentication and PGP and the addition of S/MIME for bodies, the specification also introduces the secure SIP (sips) URI scheme, which requires the use of end-to-end TLS transport (with the exception of the last hop, which may use some other encryption mechanism). This allows a UA to know that the SIP message path does not traverse any unencrypted links end-to-end through the presence of intermediary proxies.

Other changes include a new requirement that Contact URIs be globally resolvable (which is useful in certain call transfer scenarios). SIP now also includes full support for IPv6 addresses. The use of hostnames instead of literal addresses is now recommended in Via and Contact URIs. The new specification renames the “call leg” to a dialog. The definition of a dialog was also changed so that it includes the To and From tags but not URIs. This will allow future revisions of SIP to permit To and From URIs to be changed within a dialog for clarity. (For example, if A calls B but the call is forwarded to C, when C answers, the 200 OK could contain the URI of C instead of B in the To header field, showing A that C has picked up the call.)

The new specification also formally defines an “early” dialog, which is created when a response to an INVITE (or SUBSCRIBE) is received which contains a tag. In the case of forking, multiple independent early dialogs may be created for a single INVITE request. As a result, Contact and Forward-Route header fields must be returned in responses that create early dialogs.

RFC 3261 also included new header fields such as Alert-Info, Call-Info, Reply-To, In-Reply-To, and Error-Info, which add new and useful functionality to SIP.

About the Author

Alan B. Johnston is a distinguished technical member with MCI in their engineering department, and is also an adjunct assistant professor of electrical and systems engineering at Washington University in St. Louis, Missouri. He is currently working with the SEP process in their MCI Advantage product. Prior to MCI, he worked at a computer local exchange carrier (CLEC), Borealis Fiber Properties, Telcordia (formerly Bellcore), and SBC Technology Resources. Dr. Johnston is a member of the new SEP specification EPC 3001 and the co-chair of the new IEEE Circuits and Components, IET Working Group. He has a Ph.D. from Lehigh University, Bethlehem, Pennsylvania, in electrical engineering and a B.S. in engineering (first class honors) from the University of Melbourne, Melbourne, Australia, in electrical engineering. Born in Melbourne, Australia, Dr. Johnston currently resides in a national historic district in St. Louis, Missouri, with his wife Lisa, son Nolan, and daughter Nara.

Index

- 3GPP addresses, 204-2
 - illustrated, 204
 - IMSI classes, 204
 - IP addresses, 202
 - Mobile IP classes, 202
 - signaling responses, 201
- 3GPP basebands, 205-4
- 3GPP mobile cell flow, 205-04
 - defined, 205-08
 - illustrated, 205
- 300 Trying response, 109
- 301 Redirect response, 109
- 301 Call Is Being Forwarded response, 109
- 302 Call Moved response, 109
- 303 SERVICE FORGONE response, 110
- 304 OK response, 112
- 305 Accepted response, 112
- 306 Multiple Choices response, 110
- 307 Moved Permanently response, 110
- 308 Moved Temporarily response, 110
- 309 See Proxy response, 111
- 310 Authentication Service response, 111
- 400 Bad Request response, 114
- 401 Unauthorized response, 114
- 402 Payment Required response, 114
- 403 Forbidden response, 115
- 404 Not Found response, 115
- 405 Method Not Allowed response, 115
- 407 Proxy Authentication Required response, 115-16
- 408 Request Timeout response, 116
- 410 Gone response, 116
- 411 Length Required response, 116
- 412 Request Entity Too Large response, 117
- 413 Request-URI Too Long response, 117
- 415 Unsupported Media Type response, 117
- 416 Unsupported URI Scheme response, 117
- 418 Bad Relations response, 117
- 421 Extension Required response, 117-18
- 422 Session Time Interval Too Small response, 118
- 423 Interval Too Brief response, 118
- 428 See Authentication Token response, 118
- 429 Provide Reference Identity response, 118-19

- 488 Temporarily Unavailable response, 118
- 490 Biting/Temporary-Loss Conn. Not Exist. response, 117
- 492 Loop Detected response, 117
- 493 Too Many Bounces response, 118–20
- 494 Address Incomplete response, 120
- 495 Inappropriate response, 118–21
- 496 Busy Host response, 111
- 497 Request Terminated response, 122
- 498 Not Acceptable Host response, 122
- 499 Bad Host response, 111
- 499 Request: Pending response, 122
- 499 Request: Undeliverable response, 121–23
- 500 Not TopLevelMail response, 124
- 502 Mail Gateway response, 124
- 503 Service Unavailable response, 124
- 504 Gateway Timeout response, 124
- 505 Version Not Supported response, 124
- 513 Message Too Large response, 125
- 508 Busy Recipient response, 125
- 509 Mailbox Full response, 125
- 510 Not Recv'd Anywhere response, 121
- 509 Not Acceptable response, 121
- RCPT-CONTINUE header field, 140–41
- RCPT-ENVELOPE header field, 140
- RCPT header field, 140
- RCPT-LENGTH header field, 140–42
- RCV method, 77–79
 - application-level only message body, 77
 - defined, 77
 - end-to-end, 77, 78
 - example, 78
 - hop-by-hop, 77, 78
 - mandatory header, 77
 - SyntaxMethods
- acknowledgment message, 55–56
 - defined, 56
 - UA generation, 55
- Address Complete Message (ACM), 119
 - address
 - readable name, 119
 - IP, 119–20, 122
 - MAIL, 119
 - address of record (AOR), 119
 - Alert-Info header field, 120–21
 - Allow-Events header field, 120
 - Allow header field, 120
 - application design team, 119
 - application type, 1–5
 - Approved Header Name Format (AHNF), 1
 - comments, 14
 - representation, 13–14
 - rules definition, 14
 - authentication, 71–79
 - digest challenge, 75
 - forms, 71–74
 - HTTP Digest, 78
 - passwords, 78
 - multiparty, 140
 - rules, 78
 - of user agent, 78
 - Authentication-Info header field, 154–54
 - Authorization header field, 141
 - back-to-back user agents (B2B), 45
 - BVE method, 76–77
 - defined, 76
 - example, 76, 77
 - mandatory header, 77
 - SyntaxMethods
 - all files
 - SMTP, 138–46
 - with authentication, parties, and
 - read-views, 147–14
 - authentication digest, 79
 - examples, 147–51
 - HTTP, 144
 - HTTP via SMTP, 138–41
 - SMTP method, 87
 - parallel work, 139–40
 - SMTP via SMTP, 140–41
 - SMTP method, 84
 - SMTP via SMTP, 138–42
 - with reader-writer parties, 144–14
 - SMTPC100 method, 87
 - Call-ID header field, 129–30
 - Call-Info header field, 142

Domain Name Service (DNS) (continued)

- errors, 65, 111

 dual name-multiple-frequency (DTMF) digits

- 174, 263

Dynamic Host Configuration Protocol

- (DHCP), 9

 emergency-calling design team, 263 **emergency**

- calls-and, 263

- help-by-help, 60

- NUMBER, 26-28

 Escapypython header field, 133 **Expires-Info header field, 199** **Expires header field, 142** **Expires extensions, 33** **EXPIRES header field, 140** **Expires-01** **falling proxy operation, 58** **From header field, 155-56** **From-direction, 164-66** **gateway, 41-47**

- accommodation, 47

- defined, 45-50

- media (MCI), 47

- network with, 46

- SMTP, 47

- SMTP to SIP through, 121-23

- SIP to SMTP through, 124-22

 globally secure responses, 121-26 **globally readable source URLs, 143****H.323**

- callflow illustration, 184

- calls-and-flow responses, 186

- comparison, 187-91

- defined, 181

- end points, 180

- examples, 184-87

- firewall connections with, 186

- introduction to, 180-83

- network elements, 186

- video area, 190

- protocols interconnectivity, 183

- to SIP call flow, 236-39

- version, 187

 header fields, 127-60

- ACCEPT, 140

- Accept-Contact, 198-201

- Accept-Encodings, 141

- Accept-Language, 141-42

- Accept-Info, 129-30

- Allow, 136

- Allow-Events, 129

- Authenticate-Info, 153-54

- Authenticate, 142

- Call-ID, 128, 129-30

- Call-Info, 142

- contact, 127

- contact-form, 128

- Contact, 136-37

- Contact-Display-URL, 128

- Contact-Encodings, 136

- Contact-Language, 136

- Contact-Search, 129

- Contact-Type, 129-30

- CRAP, 102, 103

- Date, 132-33

- Expires, 141

- From-Info, 154

- HTTP, 143

- HTTP/1.0, 140

- From, 153-54

- In-Reply-To, 140

- isource/modified-by-param, 129

- Joins, 193-94

- Max-Forwards, 147

- message-body, 136-38

- MIME-Version, 140

- Min-Expires, 136

 MIA-01, 194-95

- OPERATION, 194

- P-Associated-URI, 147

- P-Assert-Info-Policy, 145-46

- P-Event-Related-Identity, 147

- Priority, 146-48

- Priority, 145

- Proxy-Authenticate, 145

- Proxy-Authorization, 145, 146

- Proxy-Require, 145

- RCR, 152-53

- Referer, 147-48

- Reject-Event, 194-95

- Referer-By, 148-49

- Refer-To, 148

- Reject-Contact, 198-99

- Reply, 140

- Expipg-To, 149–50
- express, 149–51
- request and response, 129–49
- REQUEST-DISPOSITION, 129
- REQUEST, 129–32
- response, 129–47
- SEQUENCE-NO, 132
- SEQ-NO-START, 129, 129
- DATA, 132
- SIZE, 129–37
- rules, 127
- SERVER, 129
- STATUS-REASON, 131
- STATUS, 129
- IDENTIFICATION-STATUS, 129
- IDENTIFY, 129
- TRANSFER, 129–37
- To, 127, 129
- Unsupervised, 129
- User-Agent, 127–28
- Via, 129–30
- Warning, 129
- WWW-Authenticate, 129
- testing/implementation design-time, 262
- HTTP Digest authentication, 54
- Hypertext Transfer Protocol (HTTP), 3, 167
 - HTTP method, 65–68
 - defined, 65
 - example, 65–66
 - mandatory headers, 66
 - see also Methods
- informational response, 109–11
 - 100 Try Again, 109
 - 101 PENDING, 109
 - 102 Call 3rd Party
 - Forwarded, 109
 - 103 Call 3rd Party, 109
 - 104 Request Progress, 109
 - defined, 109
 - see also Response(s)
- 30-Expipg-To header field, 141
- access, managing
 - allow example, 126–29
 - example, 33–38, 72
 - access, restricting, 72
 - request, 72
 - URLs, 101
 - web, 261
- Access Assigned Number Association (EAAA), 3
- Access Engineering Task Force (AETF), 2–3
 - Access Emergency Procedures (AEPs), 262
 - FAST working group, 2
 - Request for Comments (RFCs), 2, 267–69
 - EMPLD working group, 2
 - EFFING working group, 2
 - EP-related work items, 262
 - EXACT working group, 2
 - working groups, 2–3
- Access layer, 4–5
- Access Mobile Backbone Network (AMBN), 11
- Access Multimedia Protocol stack, 4–8
 - application layer, 4–8
 - illustrated, 4
 - Access layer, 4–5
 - physical layer, 4
 - transport layer, 4–8
- Access Protocol, 5–11
 - Access service providers (ASPs), 21
 - ASPTTS models, 72–74
 - defined, 72
 - example, 73–74
 - for relating/direct references, 73
 - Expipg-To header, 73
 - mandatory header, 74
 - with an ASP message body, 72
 - see also Methods
- AT, 1
 - addresses, 21–26
 - model, 25–26
 - Multimedia Core Network Subsystem (MCS), 24
- EDGE signaling, 176–77
- EDGE
 - ACM, 211
 - signaling, 176
- EDGE header field, 142–44
- EGP-PCRNALD header field, 147
 - media access control (MAC) addresses, 18
 - media gateway controller (MGC), 47
 - media gateway (MG), 47
 - message bodies, 181–4
 - forms, 182–3

message bodies (continued)

- information types, 182
- intended use, 182
- MIME encoding, 182
- number of extensions, 188
- SMTP carrying of, 182
- message-body header fields, 184–91**
 - Alt, 184
 - Content-Disposition, 188
 - Content-Encodings, 188
 - Content-Language, 188
 - Content-Length, 188
 - Content-Type, 188–89
 - Expires, 189
 - Inline-Images, 188
 - Text-Header, 188
- MIME method, 99–103**
 - 001 SMTP response, 99
 - defined, 99–101
 - example, 102–103
 - mandatory header fields, 99
 - Text-Header
- message transport, 38–41**
 - SMTP, 38–42
 - TCP, 40
 - TLS, 40–41
 - UDP, 38–39
- methods, 71–104**
 - ACL, 77–79
 - RCPT, 76–77
 - RCPTL, 77–81
 - defined, 71
 - RECV, 82–84
 - REPLY, 75–76
 - RETR, 76–81
 - save, 71
 - START, 86–88
 - OPT, 88, 92
 - START, 94–96
 - STOP, 82–86
 - START, 76–78
 - START, 86–88
 - unknown, 71
 - UPDATE, 96–98
- mid-air mobility, 187–189**
 - with INVTLS and Inplace, 189
 - with an INVTLS, 187
- NAME-Storage header field, 188**
- NAME-Storage header field, 184**

- NAME-Storage field, 184–185**
- Match IP, 189–191**
 - SMTP extension, 192
 - defined, 191
 - inplace routing, 191
- mobile**
 - IP, 191–192
 - method, 191–192
 - protocol, 194–195
 - proxy, 195
 - services, 198
 - SMTP, 189–191
- modules, 12–13**
 - support, 12
 - use, 13
- multitenant, 4, 41**
- Multipoint Extension: Mail Extension (MIME) encoding, 185**
 - example, 185–187
- network address translation (NAT), 61–62**
 - address/forward, 62
 - problems, 61, 62
 - TCP and, 62–67
 - transport, 61
 - transport protocols/extensions, 62–68
- NOTIFY method, 99–100**
 - defined, 99
 - example, 99
 - example call flow, 97
 - mandatory header fields, 99
 - SMTP response, 99
 - use within dialog, 99–100
- OPT, 86 method, 81–82**
 - defined, 81
 - example, 81–82
 - mandatory header, 81
 - message body, 81
 - Text-Header
- Original mail header field, 176**
- overlay dialog, 121**
 - parallel search example, 121–122
 - defined, 121
 - illustrated, 121
 - Text-Header
- P-BASED-EXT-TRANSFER header field, 99**
- PRI header field, 184**

- personal mobility, 194–95
 defined, 194
 support, 194
 See also Mobility
- P** headers, 203–4
- physical layer, 4
- Plain-Text Protocol (PTP)**, 4
- P-URI-Auth-Uri** header field, 149–150
- P-Service-Uri** header field, 147
- PRACK** method, 94–96
 defined, 94
 example, 95–96
 mandatory header fields, 96
 message body, 95
 reliable response, 95
 UAC generation, 94
 See also Method
- prevail** mobility, 194
- process**
 call flow example, 204–205
 Event structure and, 21
 example, 23–26
 information, 23
 packages, 26
 process, 21
 users, 45
 URLs, 101
 work, 204
- process again (PAg)**, 44–45
 information, 45
 defined, 44
 information collection, 44
- PROCEED** header field, 144–45
- Proxy** header field, 145
- probe**
 call flow example, 216–218
 forking operation, 21
 header fields transmitted by, 120
 number of, 20
 transaction method, 20
 withering method and, 21
- PROXY-Auth-Uri** header field, 145
- PROXY-Auth-Uri-Auth-Uri** header field, 145, 146
- Proxy-Authenticate** header field, 145
- proxy users**, 47–52
 users, 48
- calls with, 25–30
 defined, 47–48
 forking, 52
 message bodies, 50
 method, 49, 50
 numbers, 49
 URLs vs., 48
 See also User-Agent
- PTN and Internet Interworking (PINT)**
 process, 88
- Public Internet Telephone Network (PINTN)**, 46
 gateways, 47
 interworking, 111
 overlay-filing, 121
 protocols, 125–127
- Push** header field, 152–155
- Real-Time Transport Protocol (RTP)**, 8,
 173–174
 audio video profiles, 174–175
 Control Protocol (RTCP), 175
 development, 171
 best practice discussion, 172–173
 users, 171–172
 payload header, 172
 payload types, 173
 sessions, 174
- Reason** header field, 147–148
- Reason-Header** header field, 136–38
 collection response, 112–113
 redirection users, 52–55
 defined, 52–53
 example, 53
 redirection response, 53–54
 See also User-Agent
- REFER** method, 92–96
 responses of, 93
 defined, 92
 example call flow, 93
 example message, 93
 invite/session dialog, 92
 mandatory headers, 95
 for performing attended transfer, 93
 for pushing Web page, 94
 See also Method
- Referenced-By** header field, 144–45
- Refer-To** header field, 144

REGISTER method, 74-76

Client incremental for, 76

defined, 76

Forwarding, 75

mandatory header, 76

request errors, 24-32

defined, 24

successful registration acknowledgment,
32

responses

as acknowledgment, 31

automatic performance of, 31

defined, 31

examples, 31-32

illustration, 31

security, 31

status, 31

successful, 31

registration, admission, and status (RAS)

183

Request-Contact header field, 130-31**reliability, 26-27**

examples, 27

mechanisms, 26, 26

Reply-to header field, 130**Reply-To header field, 149-50****request and response header fields, 128-40**

Accept-Contact, 128-29

Accept-Contact, 129

Call-ID, 128-28

Contact, 128-32

CRAG, 132, 133

CRAG, 132-33

DESCRIPTION, 133

FROM, 129-29

Organization, 134

Supported-Features, 134-35

Supported-Features, 135, 136

Subject, 133

Suggest-URL, 136

Timeout, 136-37

To, 127, 128

User-Agent, 133-34

Via, 136-40

Request-Dispatch-Location header field,

150

Request for Comments (RFC), 1, 267-68

RFC 2663, 267-68

RFC 3261, 268-69

request header fields, 128-33

Accept, 128

Accept-Contact, 128-28

Accept-Dispatch-Loc, 151

Accept-Language, 141-42

Accept-Reliability, 142

cdl-Info, 142

defined, 128

Event, 143

Ex-Reply-To, 143

From, 143-44

From-Parameters, 147

P-Associated-URI, 147

P-Call-Info-Code, 146-46

P-Expires-Header-Info, 147

Priority, 144-45

Priority, 145

Proxy-Authenticate, 145, 146

Proxy-Reply-to, 145

Rel, 142-43

Response, 147-48

Reliability, 146-46

Reply-To, 148

Request-Contact, 128-28

Reply-to, 150

Reply-To, 149-50

Request-Dispatch-Loc, 150

Reply-to, 151-51

Response-By, 152

Rel, 142

Request-Expires, 151

Request-Info-Code, 150

See also Header fields**requests, See Methods****Response header field, 151-52****response code, 167**

HTTP, 167

status, 11

used, 167

response header fields, 151-57

Accept-Dispatch-Loc, 150-50

defined, 150

Event-Info, 154

From-Expires, 156

From-Info, 156-56

Proxy-Authenticate, 155

Rel, 156-57

Response, 156

Unsupported, 156

- Messaging, 98
- MSB-SubnetLocator, 156
- AcyclicRoute field
- Response-ErrorCode field, 152
- response, 107-26
 - 109 Trivial, 98
 - 109 Flushing, 109
 - 109 Call In Being
 - Forwarded, 109
 - 109 Call Granted, 99
 - 109 Session Progress, 10
 - 200 OK, 112
 - 202 Accepted, 112
 - 200 Multiple Choices, 112
 - 201 Moved Permanently, 112
 - 202 Moved Temporarily, 112
 - 205 See From, 112
 - 206 Alternative Services, 112
 - 400 Bad Request, 114
 - 401 Unauthorized, 114
 - 402 Request Required, 114
 - 403 Forbidden, 114
 - 404 Not Found, 114
 - 405 Method Not Allowed, 114
 - 406 Not Acceptable, 114
 - 407 Error Authentication Required, 114-16
 - 408 Request Timeout, 114
 - 410 Gone, 114
 - 411 Length Required, 114
 - 412 Request Entity Too Large, 114
 - 414 Request-URI Too Long, 114
 - 415 Unsupported Media Type, 114
 - 416 Unsupported URI Scheme, 114
 - 420 Bad Extension, 114
 - 421 Extension Required, 114-16
 - 422 Session Time Interval Too Small, 114
 - 423 Interval Too Brief, 114
 - 429 See Authentication Token, 114
 - 429 Provide Referrer Identity, 114-16
 - 430 Temporarily Unavailable, 114
 - 431 Dialog/Transaction Too Not Valid, 114
 - 432 Loop Detected, 114
 - 433 Too Many Requests, 114-26
 - 434 Session Incomplete, 114
 - 435 Subsequent, 114-11
 - 436 Busy Here, 114
 - 437 Request Terminated, 114
 - 438 Not Acceptable Here, 114
 - 439 Bad Branch, 114
 - 451 Request Pending, 114
 - 452 Request Unfulfillable, 114-24
 - 501 Not Implemented, 124
 - 502 Bad Gateway, 124
 - 503 Service Unavailable, 124
 - 504 Gateway Timeout, 124
 - 505 Version Not Supported, 124
 - 507 Storage Too Large, 124
 - 508 Bad Gateway, 124
 - 509 Quota Exceeded, 124
 - 510 Not Acceptable, 124
 - class, 108
 - class, 111-11
 - global, 124-26
 - international, 108-11
 - reference, 112-11
 - www, 112-25
 - www, 112
 - Retry-After field, 114, 116
 - RFC, 114, 117-18
 - RFC, 114, 116-17
 - sanitization (STP), 57
 - Route-Info field, 152
 - routing
 - header, 114
 - single, 114
 - value, 114
 - RouteInfo field, 150-57
 - www, 112-25
 - 501 Not Implemented, 124
 - 502 Bad Gateway, 124
 - 503 Service Unavailable, 124
 - 504 Gateway Timeout, 124
 - 505 Version Not Supported, 124

server error responses (continued)

513 Resource Too Large, 128

defined, 128–29

for site responses, 1

server-side header field, 175

server, 47–51

defined, 47

ports, 47–51

values, 51–55

regions, 51–52

regions, 55

Server-Error-Too-Large header field, 263

Session Announcement Protocol (SAP), 146

Session Description Protocol (SDP),

8, 163–75

applications, 163–66

attributes field, 168–69

attributes value, 170

header fields field, 167

connections field, 168–67

development, 163

media address field, 166

message type field, 167–68

field list, 165

information, 164

media announcements field, 168

media capability (capabilities) capabilities,

169

message type, 164

payload field, 168–66

protocol version field, 165

session name and information field, 166

see coding, 164

size field, 167

URI field, 166

use of, 168–71

Server-Error-Too-Large header field, 175

Session Initiation Protocol, *See SIP*

session

establishment example, 37–39

expiration, 126

SIP, 176

time, 52

signaling protocol, 1–2

Simple Mail Transfer Protocol (SMTP), 3

Simple Transport of UDP through NATs,

See STUN protocol

SIP

SMTP addresses and, 204–21

addresses, 20

field history, 3–4

calls with proxy server, 23–24

defined, 1

Event extensions, 11

extended functions, 2

gateways, 45–47

HTTP-compatibly, 147–54

header fields, 127–60

header requirements design view and, 363

Internet and, 1–14

interoperability (see *server SIPs*), 3

introduction to, 17–42

mobile operation and, 107

mobility, 104–05

popularity, 3

presence and instant message example,

73–78

in SIP, 66

registration example, 31–33

request message, 73–104

response code numbers, 23

response message, 107–20

server, 47–55

session time, 52

signaling functions, 1–2

state-based modeling, 104–09

Transport, 49

URLs, 95–100

user agent, 43–44

warning code, 110–17

Session Telephony (SIP-T), 177

SMTP message, 20–60

message process, 214–18

message process, 214–18

Simple Transport Protocol

(STTP), 7–8

defined, 7

in layer 2 transport protocol, 8

multihoming, 8–9

purpose, 41

STUN protocol, 60–65

address-resolution types, 65

class, 63

defined, 63

illustrated, 64

server, 63

Subjunctive header field, 175

Subjunctive method, 80–89

- Accept-Resume, 87
 - defined, 88
 - Error header field, 88
 - example cell flow, 87
 - mandatory header field, 88
 - value, 88
 - See also Method
- ACCEPT-RESUME-STARTS header field, 111
- accept-resume, 111
- accept-terminate, 110
- accept-terminate-headers, 110
- acc, 140
- acknowledgments, 100–101
- Telephone Gateway Registration Protocol (TGRP), 47
- Telephone Routing over IP (TRIP), 47
- THREAT-INFO header field, 136–37
- Throttle, 137, 138
- Transmission Control Protocol (TCP), 1, 3, 6–7
 - compression control, 40
 - connection, 39
 - defined, 4
 - handshaking, 40
 - spring-loading connections, 4
 - request header, 7
 - transmission illustration, 39
 - transport, 40–41
 - transport layer, 40
- Transmission Layer Security (TLS), 1, 4
 - defined, 7
 - radio authentication and, 41
 - Handshake protocol, 7
 - spring-loading connections, 4
 - see TCP, 40–41
 - transport protocol, 40
 - transport, 40–41
- transport layer, 1–4
- Traveling Using Bitly (TUT), *See* TUBN protocol
- TUBN protocol, 43–46
 - defined, 43
 - illustrated, 46
 - opens, 43
 - ringbuffer routing, 45
- uniform resource identifier (URI), 1, 11
 - ACR, 15
 - hosting, 15
 - origin, 16
 - common schemes, 16
 - device, 20
 - globally routable scheme, 20
 - port number, 19
 - URI, 19–20
 - telephone scheme, 110
 - user, 15
- uniform resource identifier (URL), 1
 - defined, 12
 - scheme, 12
 - URI, 19
 - protocol, 19
 - telephone, 100–101
- Universal Plug and Play (UPnP) Protocol, 178
- Unsupported header field, 111
- UPDATE method, 96–98
 - defined, 96
 - example, 97
 - mandatory header fields, 98
 - See also Method
- User-Agent header field, 137–38
- user-agent, 40–41
 - acknowledgments device, 44
 - acknowledgments generation, 55
 - authentication, 58
 - back-to-back (B2B), 41
 - client-server applications, 44
 - defined, 40
 - function, 40
 - group version, 40
- User-Agent Protocol (UAP), 1, 2
 - compression control and, 39
 - defined, 7
 - transmission illustration, 39
 - transport, 39–40
- see URIs, 16
- utility applications, 8–10
- URI forwarding decision *see*, 138, 139
- URI header field, 138–40
 - content, 138
 - defined, 138
 - example, 140
 - See also Header field
- warning codes, 156–57
- WORLD-WIDE header field, 136
- WWW-Authenticate header field, 136