RADVISION

TECHNICAL OVERVIEW

Session Initiation Protocol (SIP)

NOTICE

For further information contact RADVISION or your local distributor or reseller.

April, 2005

http://www.radvision.com

CONTENTS

Session Initiation Protocol (SIP)

RADVISION

Session Initiation Protocol (SIP)

TECHNICAL OVERVIEW

Introduction

The Session Initiation Protocol (SIP) is a signaling protocol for initiating, managing and terminating voice and video sessions across packet networks. Borrowing from Internet protocols, such as HTTP and SMTP, SIP is text-encoded and highly extensible. SIP can be extended to accommodate features and services such as call control services, mobility and interoperability with existing telephony systems. SIP is being developed by the SIP Working Group, within the Internet Engineering Task Force (IETF). The protocol is published as RFC 3261. This section describes the key constituents of SIP.

SIP Entities

A SIP network is composed of five types of logical SIP entities. Each entity has specific functions and participates in SIP communication as a client (initiates requests), as a server (responds to requests), or as both. One “physical device” can have the functionality of more than one logical SIP entity. For example, a network server working as a Proxy server can also function as a Registrar at the same time.

The logical SIP entities are:

- User Agent
- Proxy server
- Redirect server
- Registrar server
- Back-to-Back User Agent (B2BUA)

User Agent

In SIP, a User Agent (UA) is the endpoint entity. User Agents initiate and terminate sessions by exchanging requests and responses.

RFC 3261 defines the User Agent as an application, which contains both a User Agent client and User Agent server, as follows:

- User Agent Client (UAC)—a client application that initiates SIP requests.
- User Agent Server (UAS)—a server application that contacts the user when a SIP request is received and that returns a response on behalf of the user.

Some of the devices that can have a UA function in a SIP network are workstations, IP-phones, telephony gateways, call agents, and automated answering services.

Proxy Server

A Proxy server is an intermediary entity that acts as both a server and a client, for the purpose of making requests on behalf of other clients. Requests are serviced either internally or by passing them on, possibly after translation, to other servers. A Proxy interprets, and, if necessary, rewrites a request message before forwarding it.

Redirect Server

A Redirect server is a server that accepts a SIP request, maps the SIP address of the called party into zero (if there is no known address) or more new addresses and returns them to the client. Unlike Proxy servers, Redirect Servers do not pass the request onto other servers.

Registrar

A Registrar is a server that accepts REGISTER requests for the purpose of updating a location database with the contact information of the user specified in the request.

B2BUA

A B2BUA is a logical entity that receives a request, processes it as a User Agent Server (UAS) and, in order to determine how the request should be answered, acts as a User Agent Client (UAC) and generates requests. A B2BUA must maintain call state and actively participate in sending requests and responses for dialogs in which it is involved. The B2BUA has tighter control of the call than a Proxy—for example, a Proxy cannot disconnect a call or alter the messages.

Messages

The following sections deal with SIP messages.

Message Types

There are two types of SIP messages:

- Requests—sent from the client to the server.
- Responses—sent from the server to the client.

Requests

Table 1 *Request Methods Example*

Method	Description
INVITE	Initiates a call, changes call parameters (re-INVITE).
ACK	Confirms a final response for INVITE.
BYE	Terminates a call.
CANCEL	Cancels searches and “ringing”.
OPTIONS	Queries the capabilities of the other side.
REGISTER	Registers with the Location Service.
INFO	Sends mid-session information that does not modify the session state.

Responses

Response messages contain numeric response codes. The SIP response code set is partly based on HTTP response codes. There are two types of responses and six classes:

- **Response Types**
  - Provisional (1xx class)—provisional responses are used by the server to indicate progress, but they do not terminate SIP transactions.
  - Final (2xx, 3xx, 4xx, 5xx, 6xx classes)—final responses terminate SIP transactions.

- **Classes**
  - 1xx = Provisional—request received, continuing to process the request.
  - 2xx = Success—the action was successfully received, understood and accepted.
  - 3xx = Redirection—further action needs to be taken to complete the request.
  - 4xx = Client Error—the request contains bad syntax or cannot be fulfilled at this server.
  - 5xx = Server Error—the server failed to fulfill an apparently valid request.
  - 6xx = Global Failure—the request cannot be fulfilled at any server.

Table 2 *Response Code Example*

Number	Meaning
100	Trying
180	Ringing
200	OK
300	Multiple choices
301	Moved permanently
302	Moved temporarily
400	Bad request
401	Unauthorized
403	Forbidden
408	Request time-out
480	Temporarily unavailable
481	Call/Transaction does not exist
482	Loop detected

Table 2 *Response Code Example*

Number	Meaning
500	Server error
600	Busy everywhere
603	Decline
604	Does not exist anywhere
606	Not acceptable

Message Parts

SIP messages are composed of the following three parts:

- Start line
- Headers
- Message body

Start Line

Every SIP message begins with a Start Line. The Start Line conveys the message type (method type in requests, and response code in responses) and the protocol version. The Start Line may be either a Request-line (requests) or a Status-line (responses), as follows:

- The Request-line includes a Request-URI, which indicates the user or service to which this request is being addressed.
- The Status-line holds the numeric Status-code and its associated textual phrase.

Headers

SIP header fields convey message attributes that provide additional information about the message. They are similar in syntax and semantics to HTTP header fields (in fact, some headers are borrowed from HTTP) and thus always take the format:

```
<name>:<value>
```

Headers can span multiple lines. Some SIP headers such as Via, Contact, Route and Record-Route can appear multiple times in a message or, alternatively, can take multiple comma-separated values in a single header occurrence.

Message Body

A message body is used to describe the session to be initiated (for example, in a multimedia session this may include audio and video codec types and sampling rates), or alternatively it may be used to contain opaque textual or binary data of any type which relates in some way to the session. Message bodies can appear both in request and in response messages. SIP makes a clear distinction between signaling information, conveyed in the SIP Start Line and headers, and the session description information, which is outside the scope of SIP.

Possible body types include:

- SDP—see Session Description Protocol (SDP)
- Multipurpose Internet Mail Extensions (MIME)
- Others—to be defined in the IETF and in specific implementations

Message Samples

The following samples show the message exchange between two User Agents for the purpose of setting up a voice call. SIP user alice@radvision.com invites SIP user bob@acme.com to a call for the purpose of discussing lunch. Alice sends an INVITE request containing an SDP body. Bob replies with a 200 OK response also containing an SDP body.

Request Message

Table 3

Request Message line	Description
INVITE sip:bob@acme.com SIP/2.0	Request line: Method type, Request-URI (SIP address of called party), SIP version.
Via: SIP/2.0/UDP 172.20.1.1:5060; branch=z9hG4bK-2f059	Identifies the location where the response is to be sent.
Max-Forwards:70	Limits the number of hops the request will make on the way to its destination.
From: Alice A. <sip:alice@radvision.com>;tag=1121137	User originating this request. Includes a unique tag.
To: Bob B. <sip:bob@acme.com>	User being invited, as specified originally.
Call-ID: 2388990012@alice_ws.radvision.com	Globally unique ID of this call.
CSeq: 1 INVITE	Command sequence. Identifies transaction.
Contact:<sip:alice@pc33.radvision.com>	Direct route to contact Alice in further requests.
Subject: Lunch today.	Call subject and/or nature.
Content-Type: application/SDP	Type of body—in this case SDP.
Content-Length: 182	Number of bytes in the body.
	Blank line marks end of SIP headers and beginning of body.
v=0	Version of SDP.
o=Alice 53655765 2353687637 IN IP4 128.3.4.5	Owner/creator and session identifier, session version address type and address.
s=Call from Alice.	Session subject.
c=IN IP4 alice_ws.radvision.com	Connection information.
M=audio 3456 RTP/AVP 0 3 4 5	Media description: type, port, possible formats caller is willing to receive and send.

Response Message

Table 4

Response Message line	Description
SIP/2.0 200 OK	Status line: SIP version, response code, reason phrase.
Via: SIP/2.0/UDP 172.20.1.1:5060; branch=z9hG4bK-2f059	Copied from request.
From: Alice A. <sip:alice@radvision.com>;tag=1121137	Copied from request.
To: Bob B. <sip:bob@acme.com>;tag=17462311	Copied from request. Includes unique tag to identify call-leg.
Call-ID: 2388990012@alice_ws.radvision.com	Copied from request.
CSeq: 1 INVITE	Copied from request.
Contact:<sip:bob@172.20.1.77>	Direct route to contact Bob.
Content-Type: application/SDP	
Content-Length: 200	
	Blank line marks end of SIP headers and beginning of the body.
v=0	Version of SDP.
o=Bob 4858949 4858949 IN IP4 192.1.2.3	Owner/creator and session identifier, session version address type and address.
s=Lunch	Session subject.
c=IN IP4 machine1.acme.com	Connection information.
m=audio 5004 RTP/AVP 0 3	Description of media streams the receiver of the call is willing to accept.

Entity Interaction

This section describes the interaction between SIP entities in various common session initiation scenarios.

Session Establishment and Termination

Figure 1 shows the interaction between two user agents during trivial session establishment and termination.

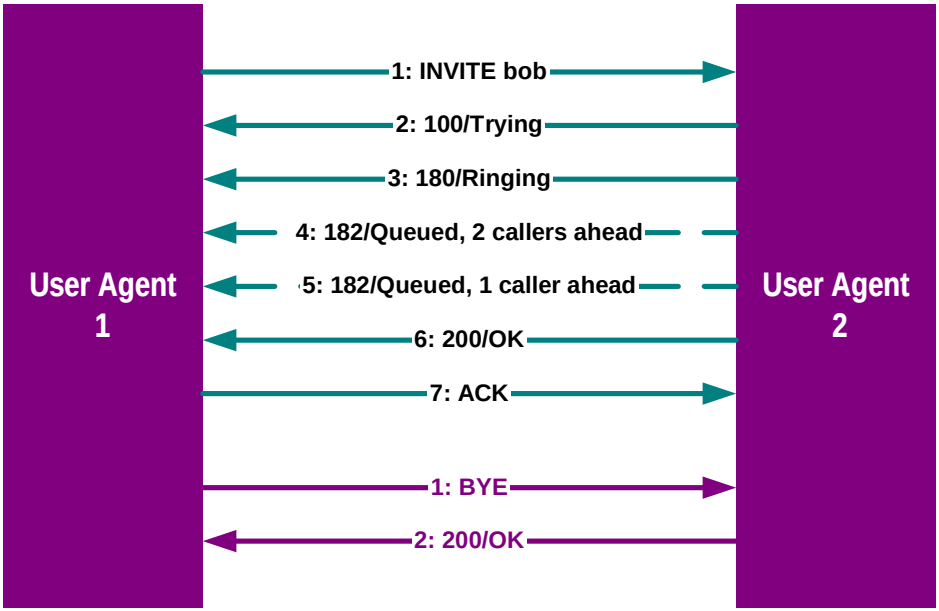

Figure 1 *SIP Session Establishment and Call Termination*

Session Establishment

Call Flow

1. UA1 sends an INVITE message to Bob's SIP address: sip:bob@acme.com. This message also contains an SDP packet describing the media capabilities of the calling terminal.
2. UA2 receives the request and immediately responds with a 100-Trying response message.

3. UA2 starts “ringing” to inform Bob of the new call. Simultaneously a 180 (Ringing) message is sent to the UAC.
4. UA2 sends a 182 (Queued) call status message to report that the call is behind two other calls in the queue.
5. UA2 sends a 182 (Queued) call status message to report that the call is behind one other call in the queue.
6. Bob picks up the call and the UA2 sends a 200 (OK) message to the calling UA. This message also contains an SDP packet describing the media capabilities of Bob’s terminal.
7. UA1 sends an ACK request to confirm the 200 (OK) response was received.

Session Termination

The session termination call flow proceeds as follows:

1. The caller decides to end the call and “hangs-up”. This results in a BYE request being sent to UA2.
2. UA2 responds with 200 (OK) message and notifies Bob that the conversation has ended.

Call Redirection

Figure 2 shows a simple call redirection scenario.

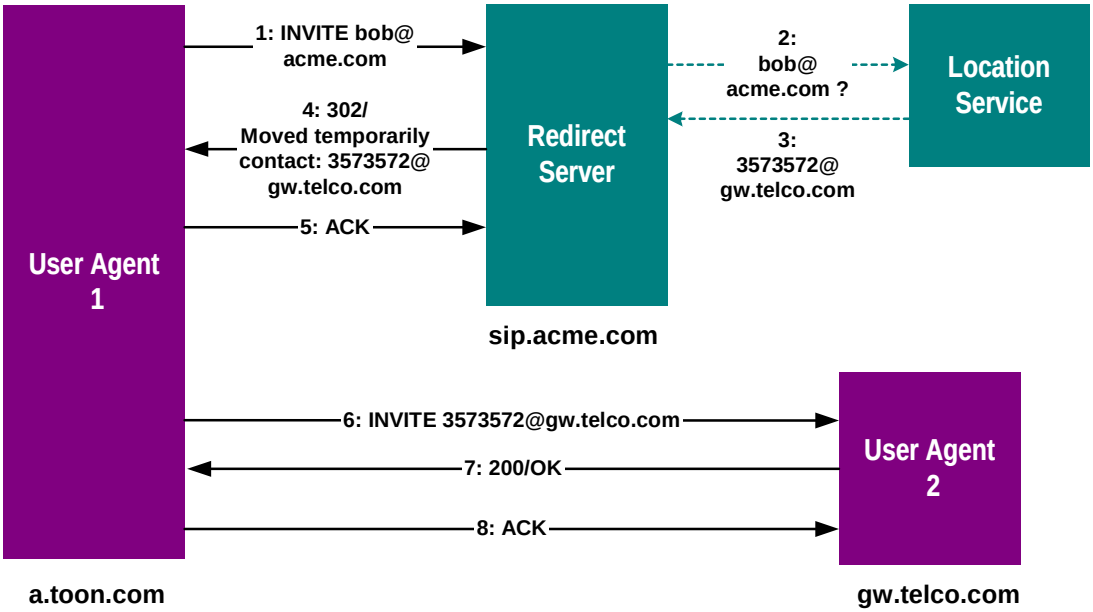

Figure 2 *Simple Call Redirection Using a Redirect Server*

Call Flow

1. First a SIP INVITE message is sent to bob@acme.com, but finds the Redirect server sip.acme.com along the signaling path.
2. The Redirect server looks up Bob's current location in a Location Service using a non-SIP protocol (for example, LDAP).
3. The Location Service returns Bob's current location: SIP address 3573572@gw.telco.com.
4. The Redirect Server returns this information to the calling UA using a 302 (Moved Temporarily) response. In the response message it enters a contact header and sets the value to Bob's current location, 3573572@gw.telco.com.
5. The calling UA acknowledges the response by sending an ACK message.
6. The calling UAC then continues by sending a new INVITE directly to gw.telco.com.

7. gw.telco.com is able to notify Bob's terminal of the call and Bob "picks up" the call. A 200 (OK) response is sent back to the calling UA.
8. The calling UA acknowledges with an ACK message.

Call Proxying

Figure 3 shows call set-up between two User Agents with the assistance of an intermediate Proxy server.

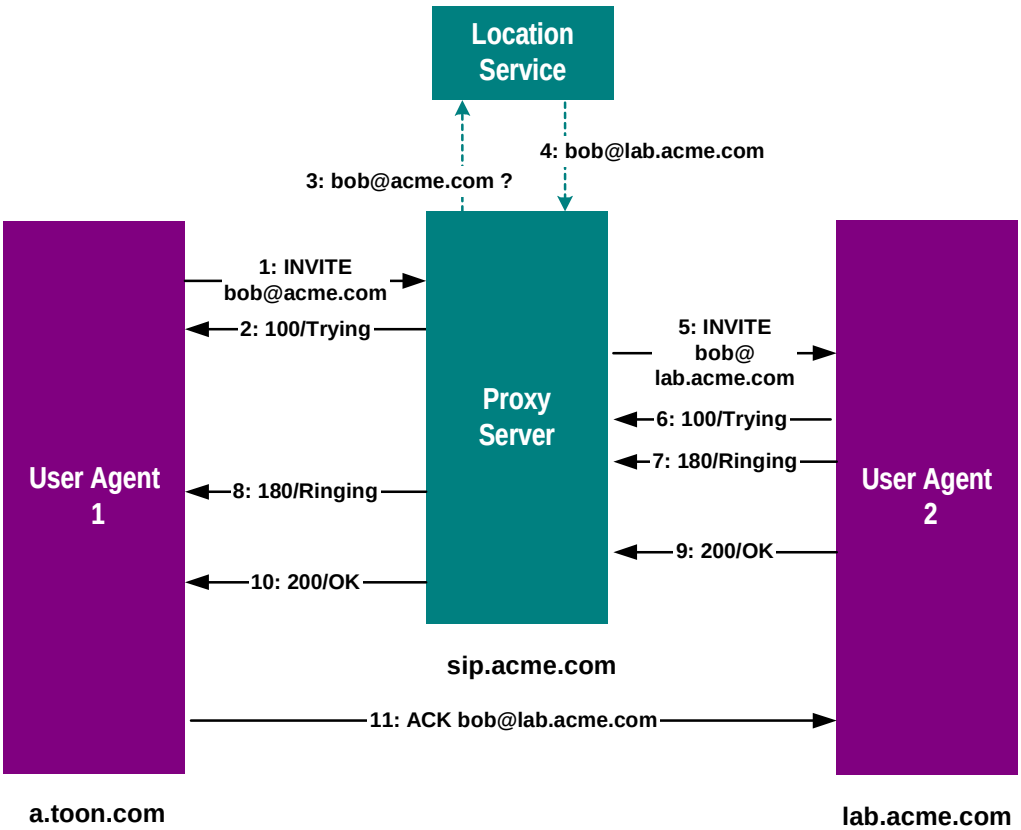

Figure 3 *Call Proxying Scenario*

Call Flow

1. An INVITE message is sent to bob@ acme.com, but finds the Proxy server sip.acme.com along the signaling path.
2. The Proxy server immediately responds with a 100 (Trying) provisional response.
3. The Proxy server looks-up Bob's current location in a Location Service using a non-SIP protocol (For example, LDAP).
4. The Location Service returns Bob's current location: SIP address bob@lab.acme.com.
5. The Proxy server decides to Proxy the call and creates a new INVITE transaction based on the original INVITE message, but with the Request-URI in the start line changed to bob@lab.acme.com. The Proxy server sends this request to the UA2.
6. The UA2 responds first with a 100 (Trying).
7. The UA2 responds with a 180 (Ringing) response.
8. The Proxy server forwards the 180 (Ringing) response back to the UA1.
9. When the call is accepted by the user (for example, by picking up the handset) UA2 sends a 200 (OK) response. In this example, UA2 inserts a Contact header into the response with the value bob@lab.acme.com. Further SIP communication will be sent directly to it and not via the Proxy server.
10. The Proxy forwards the 200 (OK) response back to the calling UAC.
11. The calling UA sends an ACK directly to UA2 at the lab (according to the Contact header it found in the 200 (OK) response).

Session Description Protocol (SDP)

SDP is the protocol used to describe multimedia session announcement, multimedia session invitation and other forms of multimedia session initiation. A multimedia session is defined, for these purposes, as a set of media streams that exist for a duration of time.

SDP Packets

SDP packets usually include the following information:

Session information:

- Session name and purpose
- Time(s) the session is active

Since the resources necessary for participating in a session may be limited, it is useful to include the following additional information:

- Information about the bandwidth to be used by the session
- Contact information for the person responsible for the session

Media information:

- Type of media, such as video and audio
- Transport protocol, such as RTP/UDP/IP and H.320
- Media format, such as H.261 video and MPEG video
- Multicast address and Transport Port for media (IP multicast session)
- Remote address for media and Transport port for contact address (IP unicast session)