SIP: Session Initiation Protocol

Corso di ***Applicazioni Telematiche***

A.A. 2006-07 – Lezione n.7

Ing. Salvatore D'Antonio

Università degli Studi di Napoli Federico II
Facoltà di Ingegneria

Session Initiation Protocol (SIP)

- SIP is the IETF's standard for establishing VoIP connections
- It is an application layer control protocol for creating, modifying and terminating sessions with one or more participants
- The architecture of SIP is similar to that of HTTP
- Requests are generated by the client and sent to the server
- The server processes the requests and then sends a response to the client
- A request and the associated responses make a *transaction*

SIP services

- The services that SIP provides include [RFC2543]:
  - User Location: determination of the end system to be used for communication
  - Call Setup: ringing and establishing call parameters at both called and calling party
  - User Availability: determination of the willingness of the called party to engage in communications
  - User Capabilities: determination of the media and media parameters to be used
  - Call handling: the transfer and termination of calls

3

SIP components

- A SIP-based System consists of two components:
  - **User Agents:** A user agent is an end system acting on behalf of a user.
    - There are two parts to it: a client and a server:
      - The client portion is called the User Agent Client (UAC)
      - The server portion is called User Agent Server (UAS)
      - The UAC is used to initiate a SIP request while the UAS is used to receive requests and return responses on behalf of the user
  - **Network Servers:** There are 3 types of servers within a network
    - Registration server:
      - receives updates concerning the current locations of users
    - Proxy server:
      - upon receiving requests, forwards them to the next-hop server, which has more information about the location of the called party
    - Redirect server:
      - upon receiving requests, determines the next-hop server and returns the address of the next-hop server to the client instead of forwarding the request

4

SIP sessions: an example

- A client is inviting a participant for a call
- A SIP client creates an INVITE message for saldanto@unina.it, which is normally sent to a proxy server.
- This proxy server tries to obtain the IP address of the SIP server that handles requests for the requested domain. The proxy server consults a Location Server to determine this next hop server. The Location server is a non SIP server that stores information about the next hop servers for different users.
- On getting the IP address of the next hop server, the proxy server forwards the INVITE to the next hop server.
- After the User Agent Server (UAS) has been reached, it sends a response back to the proxy server.
- The proxy server in-turn sends back a response to the client.
- The client then confirms that it has received the response by sending an ACK.
- In this case, we assume that the client's INVITE request was forwarded to the proxy server. However, if it is forwarded to a redirect server, then the redirect server returns the IP address of the next hop server to the client. The client then directly communicates with the UAS

5

SIP addresses

- SIP uses Uniform Resource Identifiers (URIs). At least, SIP URIs and SIPS URIs are supported, although others (such as TEL URL) are commonly supported.
  - sip:miguel.a.garcia@ericsson.com
  - sips:miguel.a.garcia@ericsson.com
  - tel:+358-9-299-3553
  - sip:proxy.atlanta.com:5060
  - sip:another-proxy.biloxi.com;transport=UDP
- SIP and SIPS URIs must include a host name, and may include username, may include port numbers, may include parameters
- Address space is unlimited
- Non SIP/TEL URIs are also valid under certain circumstances: HTTP, IM, PRES, MAILTO…

6

Still on addresses…

- This decoupling between the globally unique user-level identifier and device network addresses supports *personal mobility*, the ability of a single user to use multiple devices, and deals with the practical issue that many devices acquire their IP address temporarily via DHCP
- The proxy typically also performs call routing functions, for example, directing unanswered calls to voice mail or an auto-attendant. The SIP proxy plays a role somewhat similar to an SMTP Mail Transfer Agent (MTA), but naturally does not store messages.
- Proxies are not required for SIP; user agents can contact each other directly

7

SIP routing and Domain Name System

- SIP clients use DNS to route requests and find the next hop to route the request
  - By looking into a NAPTR (Naming Authority Pointer) record in DNS
  - By looking into a SRV (Services) record in DNS
  - By looking into A (IPv4) or AAAA (IPv6) records in DNS

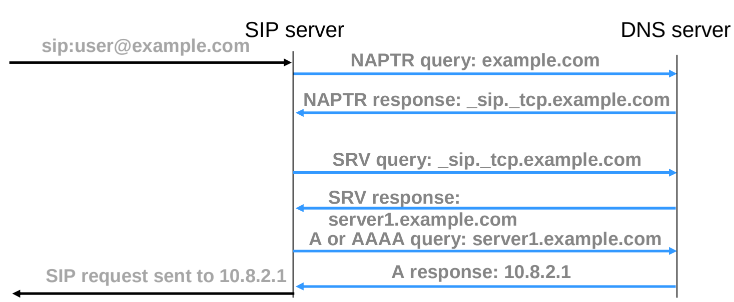

8

SIP: contacting the proxy

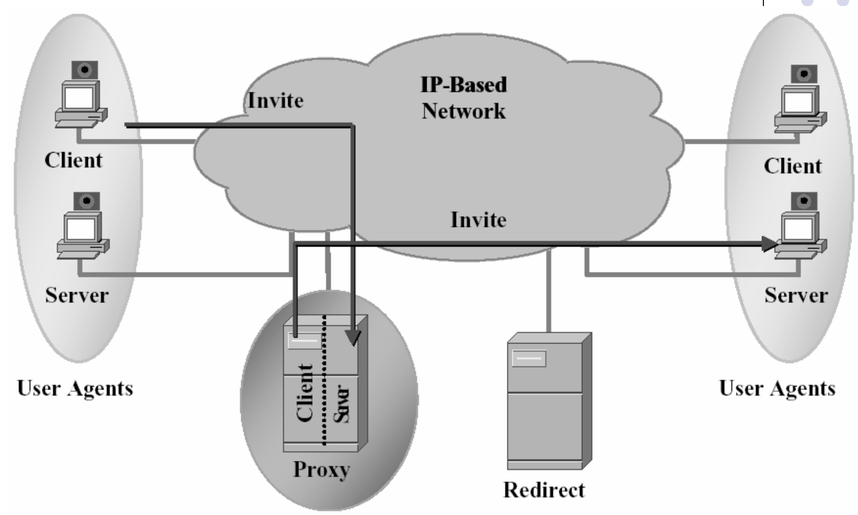

9

SIP: proxy responding to client

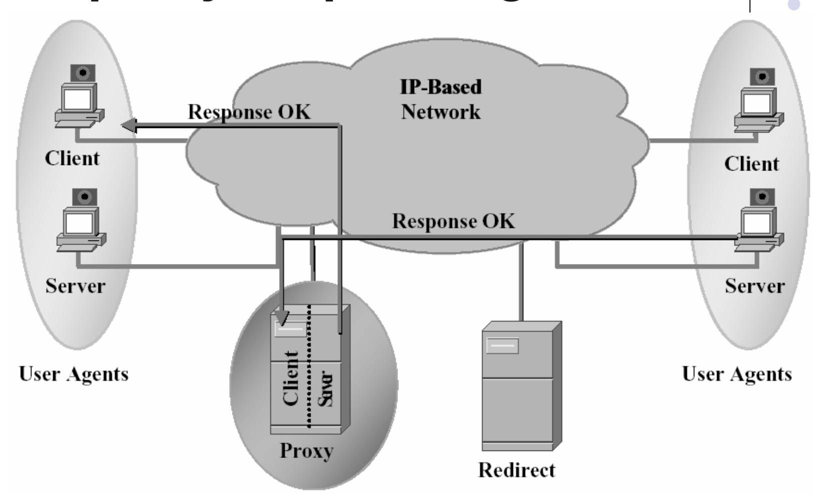

10

SIP: client acknowledgment

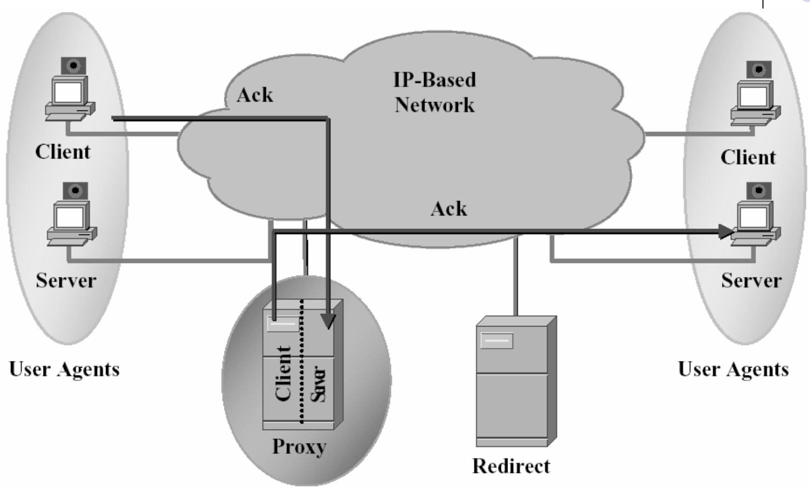

11

SIP: contact the redirect server

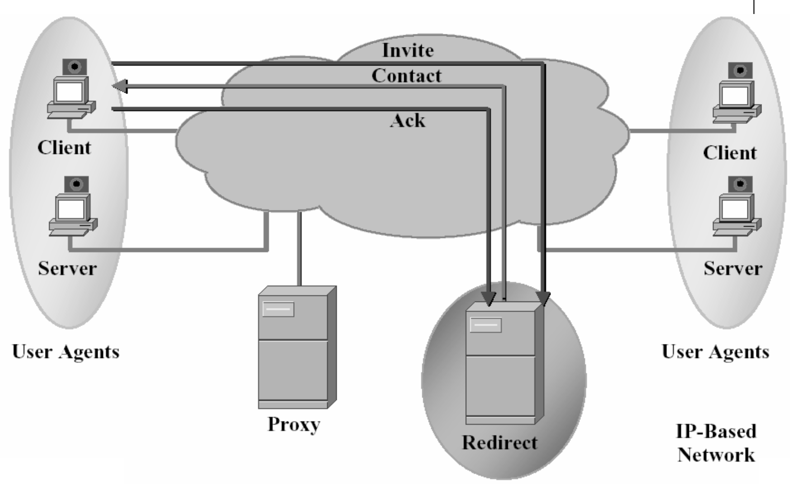

12

SIP: contacting the user

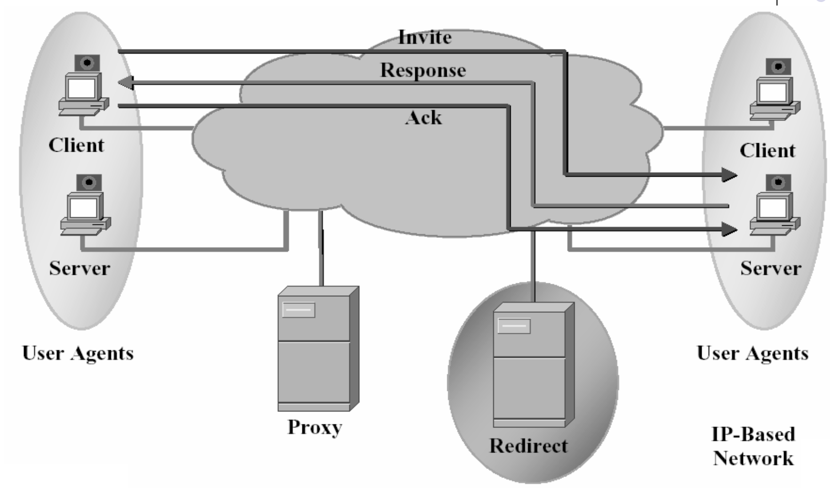

13

SIP registration

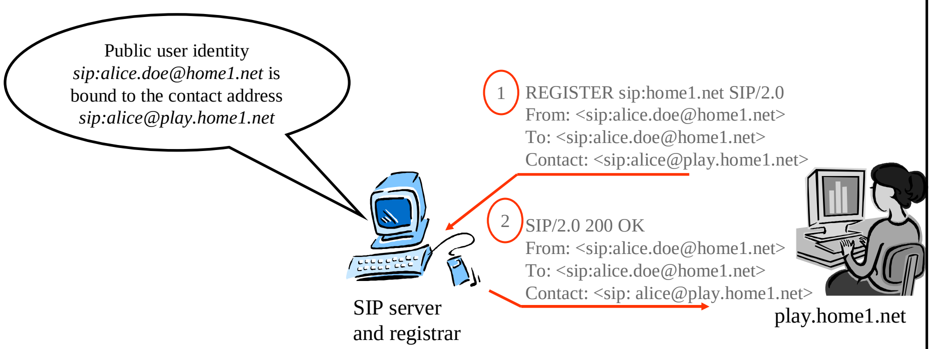

14

Routing: SIP server in proxy mode

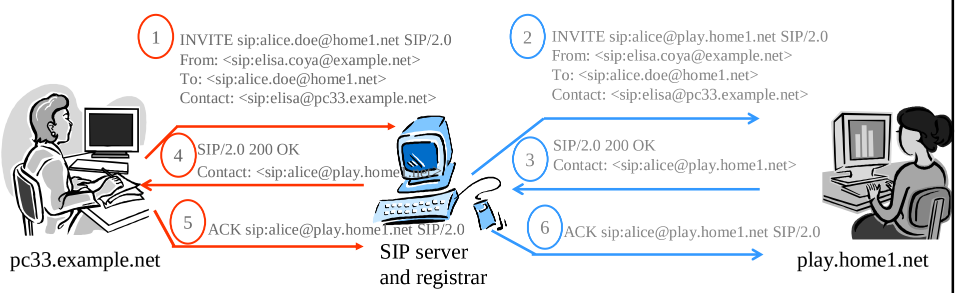

Routing: SIP server in redirect mode

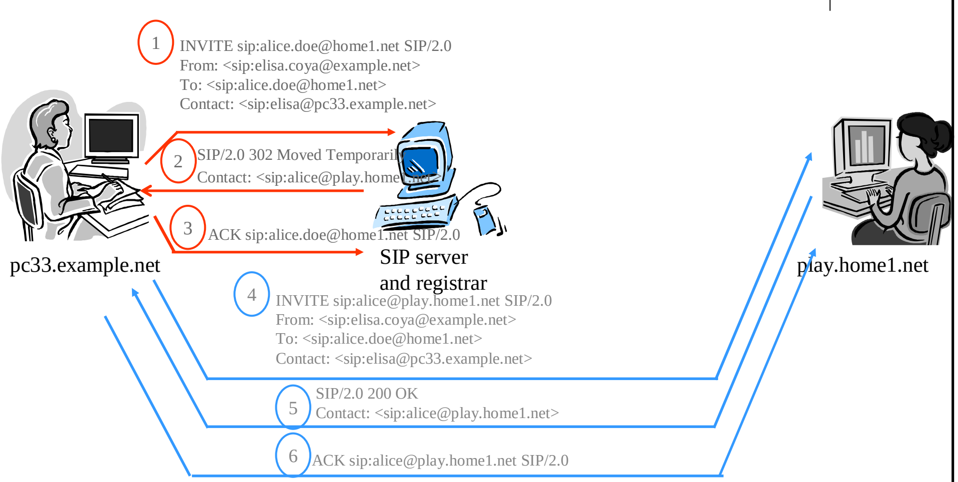

Proxy traversal

- A request can traverse any number of proxies, but typically at least two, namely one *outbound proxy* in the caller's domain and the *proxy* in the callee's domain.
- For reliability and load balancing, a domain can use any number of proxies. A client identifies a proxy by looking up the DNS SRV record enumerating primary and fall-back proxies for the domain in the SIP URI.
- Session setup messages and media generally traverse independent paths, that is, they only join at the originating and terminating client.

17

Stateless and stateful SIP proxies

There are several types of SIP proxies, depending on the state they keep:

1. **Stateless proxy:** a proxy that does not keep any state when forwarding requests and responses.
2. **Transaction stateful proxy, or stateful proxy:** a proxy that stores state during the duration of the transaction.
3. **Call stateful proxy:** a proxy that stores all the state pertaining to a session (e.g., from INVITE to BYE). A call stateful proxy is always a transaction stateful proxy, but not the other way round.

18

SIP: back-to-back user agents

- SIP user agents can initiate sessions between two other entities, acting as third-party call controllers or back-to-back user agents (B2BUAs).
- While the basic protocol mechanisms are stable, components of the SIP infrastructure are currently still under active development within the IETF and, for third-generation mobile networks, in 3GPP. Such features include support for legacy telephone features such as overlap dialing as well as advanced call routing features such as caller preferences.

19

SIP: data flow

- Media then flows directly on the shortest network path between the two terminals:
  - SIP proxies do not process media packets
- This makes it possible to route call setup requests through any number of proxies without worrying about audio latency or network efficiency
- This *path-decoupled* signaling completes the evolution of telephony signaling from in-band audio signaling to out-of-band, disassociated channel signaling introduced by Signaling System No. 7 (SS7)
- Since telephony signaling needs to configure switch paths, it generally meets up with the media stream in telephone switches; there is no such need in IP telephony

20

The SIP trapezoid

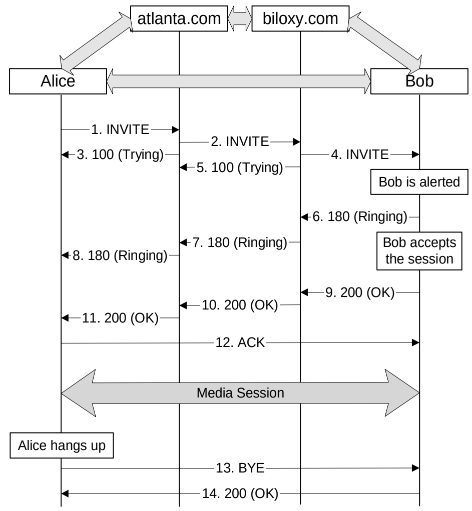

21

SIP: forking

- Just like a single phone line can ring multiple phones within the same household, a single SIP address can contact any number of SIP devices with one call, albeit potentially distributed across the network
- This capability is called *forking* and is performed by proxies. These forking proxies gather responses from the entities registered under the SIP URI and return the best response, typically the first one to pick up
- This feature makes it easy to develop distributed voicemail services and simple automatic call distribution (ACD) systems

22

Forking proxy example

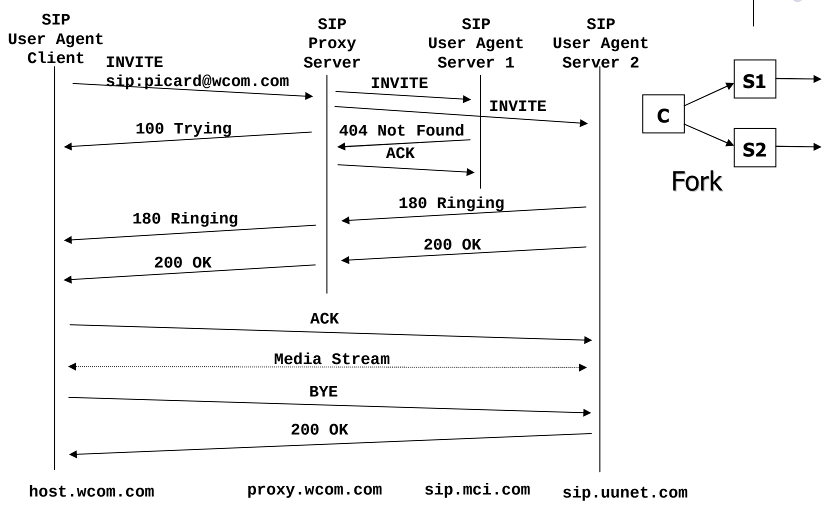

23

SIP messages

- SIP messages can be requests or responses, which only differ syntactically in their first lines
- Almost all SIP requests generate a final response indicating whether the request succeeded or why it failed, with some requests producing a number of responses that update the requestor on the progress of the request via provisional responses

24

SIP messages

- SIP has INVITE and ACK messages which define the process of opening a reliable channel over which call control messages may be passed
- SIP makes minimal assumptions about the underlying transport protocol:
  - This protocol itself provides reliability and does not depend on TCP for reliability
- SIP depends on the Session Description Protocol (SDP) for carrying out the negotiation for codec identification
- SIP supports session descriptions that allow participants to agree on a set of compatible media types
- It also supports user mobility by proxying and redirecting requests to the user's current location

25

Summary of SIP methods

- INVITE: create a session
- BYE: terminates a session
- ACK: acknowledges a final response for an INVITE request
- CANCEL: cancels an INVITE request
- REGISTER: binds a public SIP URI to a Contact address
- OPTIONS: queries a server for capabilities
- SUBSCRIBE: installs a subscription for a resource
- NOTIFY: informs about changes in the state of the resource
- MESSAGE: delivers an Instant Message
- REFER: used for call transfer, call diversion, etc.
- PRACK: acknowledges a provisional response for an INVITE request
- UPDATE: changes the media description (e.g. SDP) in an existing session
- INFO: used to transport mid-session information
- PUBLISH: publication of presence information

26

SIP: transport solutions

- Unlike other application-layer protocols, SIP is designed to run over both reliable and unreliable transport protocols. Currently, UDP is the most common transport mechanism, but TCP and SCTP, as well as secure transport using TLS are also supported
- To achieve reliability, a request is retransmitted until it is acknowledged by a provisional or final response
- The INVITE transaction, used to set up sessions, behaves a bit differently since considerable time may elapse between the call arrival and the time that the called party picks up the phone

SIP message format

- SIP is a textual protocol, similar to SMTP and HTTP.

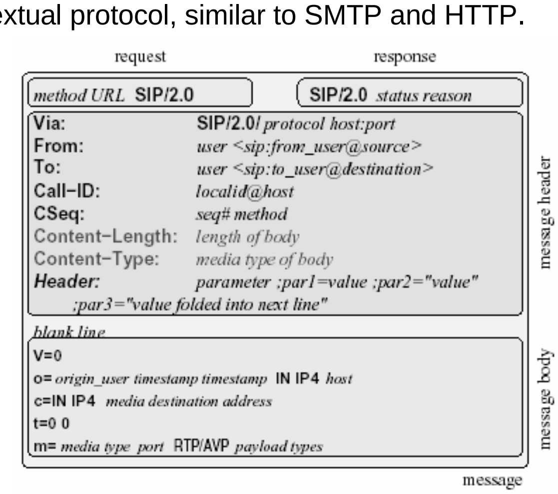

An example of a SIP request

Request Line

```
INVITE sip:miguel.a.garcia@ericsson.com SIP/2.0
```

Header

```
Via: SIP/2.0/UDP [5555::aaa:bbb:ccc:ddd];branch=z9hG4bKnashds7
Max-Forwards: 70
Route: <sip:pcscf1.visited1.net;lr>, <sip:scscf1.home1.net;lr>
From: <sip:user1_public1@home1.net>;tag=171828
To: <sip:miguel.a.garcia@ericsson.com>
Call-ID: cb03a0s09a2sdfglkj490333
Cseq: 127 INVITE
Contact: <sip:[5555::aaa:bbb:ccc:ddd]>
Content-Type: application/sdp
Content-Length: 248
```

Body

```
v=0
o=- 2987933615 2987933615 IN IP6 5555::aaa:bbb:ccc:ddd
s=-
c=IN IP6 5555::aaa:bbb:ccc:ddd
t=907165275 0
m=audio 3458 RTP/AVP 97 96 0 15
a=rtpmap:97 AMR
a=fmtp:97 mode-set=0,2,5,7; maxframes=2
a=rtpmap:96 G726-32/8000
```

Method
Version
Request-URI
Header Field
Header Field Name
Header Field Value

29

An example of a SIP response

Status Line

```
SIP/2.0 200 OK
```

Header

```
Via: SIP/2.0/UDP [5555::aaa:bbb:ccc:ddd];branch=z9hG4bKnashds7
Record-Route: <sip:scscf1.home1.net;lr>, <sip:pcscf1.visited1.net;lr>
From: <sip:user1_public1@home1.net>;tag=171828
To: <sip:miguel.a.garcia@ericsson.com>;tag=314159
Call-ID: cb03a0s09a2sdfglkj490333
CSeq: 127 INVITE
Contact: <sip:[5555::eee:fff:aaa:bbb]>
Content-Type: application/sdp
Content-Length: 220
```

Body

```
v=
o=- 2987933615 2987933615 IN IP6 5555::eee:fff:aaa:bbb
s=-
c=IN IP6 5555::eee:fff:aaa:bbb
t=907165275 0
m=audio 3458 RTP/AVP 97 0
a=rtpmap:97 AMR
a=fmtp:97 mode-set=0,2,5,7; maxframes=2
```

SIP Version
Reason phrase
Status code

Status codes in SIP

- 1xx - Provisional responses
- 2xx - Success
- 3xx - Redirection
- 4xx - Client Error
- 5xx - Server Error
- 6xx - Global Failures

30

Fields

- The **Record-Route** request and response header field is added to a request by any proxy that insists on being in the path of subsequent requests for the same call leg. It contains a globally reachable Request-URI that identifies the proxy server.
- Each proxy server adds its Request-URI to the beginning of the list. The server copies the Record-Route header field unchanged into the response. (Record-Route is only relevant for 2xx responses).
- The calling user agent client copies the Record-Route header into a Route header field of subsequent requests within the same call leg, reversing the order of requests, so that the first entry is closest to the user agent client. If the response contained a Contact header field, the calling user agent adds its content as the last Route header. Unless this would cause a loop, any client MUST send any subsequent requests for this call leg to the first Request-URI in the Route request header field and remove that entry.
- The **Route** request-header field determines the route taken by a request. Each host removes the first entry and then proxies the request to the host listed in that entry, also using it as the Request-URI.
- The **Via** field indicates the path taken by the request so far. This prevents request looping and ensures replies take the same path as the requests, which assists in firewall traversal and other unusual routing situations.
- The **Request-URI** is a SIP URL, which indicates the user or service to which this request is being addressed. Unlike the To field, the Request-URI MAY be re-written by proxies.

31

Session Description Protocol (SDP)

- Currently, only the Session Description Protocol (SDP) is being used, but an XML-based replacement is being discussed.

```
v=0
o=alice 2890844526 2890844526 IN IP4 client.atlanta.example.com
s=-
c=IN IP4 192.0.2.101
t=0 0
m=audio 49172 RTP/AVP 0
a=rtpmap:0 PCMU/8000
```

  - A simple audio session originated by user *alice* to be received by IP address 192.0.2.101 and port 49172 using RTP and payload type 0 (-law audio).

32

SIP: capabilities exchange

- Once a request has reached the right destination, the two parties negotiate the media streams using an offer-answer model, where the caller typically offers a capability and the callee makes a counter-proposal. Sessions can be changed in the middle of a session, e.g., to add or remove a media stream
- SIP can be extended by adding new methods, message body types or header fields
- Generally, receivers and proxies are free to ignore header fields that they do not understand, but a requestor can require that the receiver understand a particular feature by including a Require header field. If the receiver does not implement that feature, it must reject the request

33

Who is using SIP

- Third Generation Partnership Project (3GPP) uses SIP in the IP Multimedia Subsystem (IMS)
- Third Generation Partnership Project 2 (3GPP2) uses SIP in the IP Multimedia Subsystem (IMS)
- Microsoft uses SIP in Real-Time Communication in Windows XP
  - Note that earlier versions of MS Messenger had a non-standard pre-implementation of SIP for Presence and Instant Messaging
- PacketCable in Distributed Call Signalling
- The International Softswitch Consortium
- ETSI TIPHON program
- Wireless Village (now part of the Open Mobile Alliance) is looking at SIP for Instant Messaging and Presence
- America On Line (AOL) will support SIP for Instant Messaging and Presence

34