State of SIP

Henning Schulzrinne
Dept. of Computer Science
Columbia University
New York, New York
(sip:)schulzrinne@cs.columbia.edu

SIP Summit – Richardson, TX

May 1st, 2001

Overview

- A brief history
- Service models
- SIP design principles
- Extensions in progress
- Potential hazards

Whence SIP?

Feb. 1996: earliest Internet drafts

Feb. 1999: *Proposed Standard*

March 1999: RFC 2543

April 1999: first SIP bake-off

November 2000: SIP accepted as 3GPP signaling protocol

December 2001: 6th bake-off, 200+ participants

March 2001: 7th bake-off, first time outside U.S.

SIP years

Year	development	trade rags
1996-1998	R&D	“academic exercise’, “distraction from H.323”
1999	standard & skunk works	“what does SIP stand for again?”
2000	product development	“SIP cures common cold!”
2001	pioneer deployment	“Where are the SIP URLs?”
2002	kmart.com/sip	SIP product comparisons

SIP developments

- working towards eco-system of interoperable solutions
- device configuration
- service architectures
- emergency services
- benchmarking

The SIP eco-system

- need whole suite of things:
  - SIP stacks
  - phones (soft + hard) $\surd$, but still $$$
  - proxies, redirect, location servers $\surd$
  - services: conferencing, unified messaging
  - test tools
  - service creation tools
- 8th SIP interoperability test event in August, now close to 60 companies
- basic features work, with fancy things on the way

The blessing of delay

- bad news: both cable modem and 3G are being delayed
- 3G not before 2004
- good news: makes it more viable to go to (close-to) all-IP solution immediately
- in 3G, R5, instead of R3/4 first

Device configuration

- need to plug in store-bought phone, without more than personalization
- limited user interface
- configuration from local (visited) network and from home network
- don’t want current PBX single-vendor tie-ins
- cannot rely on California-style upgrades
- notifications of new configurations ➠ SUBSCRIBE/NOTIFY

Device configuration

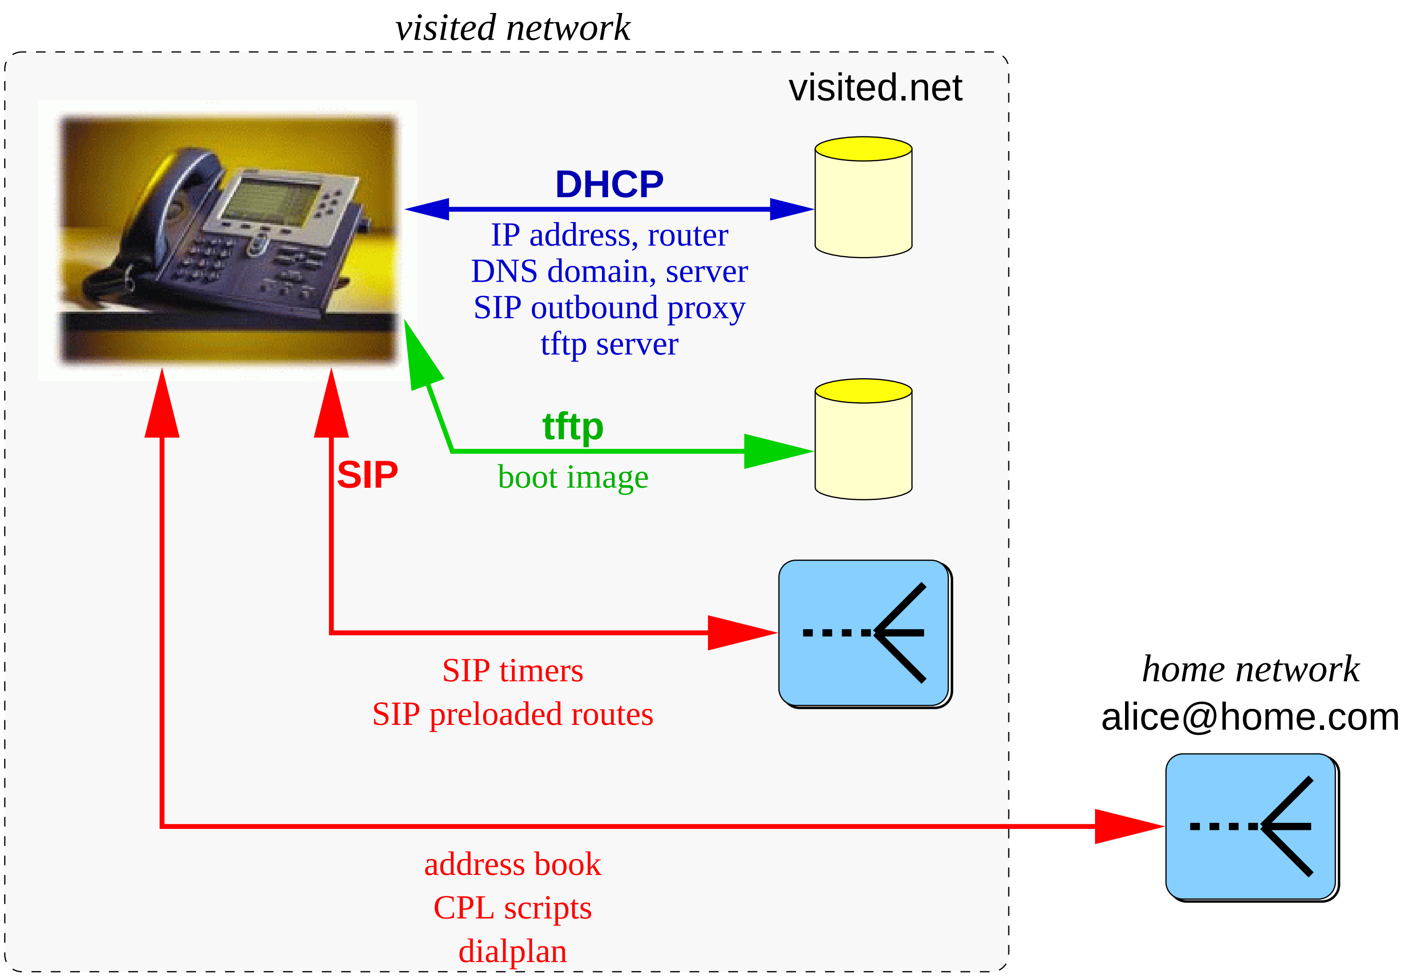

Service architectures

- single soft switch or proxy ➠ cluster of (cheap?) PCs
- “the Google model” $\leftrightarrow$ multiprocessor Tandem-style server with bulletproofing
- SIP is inherently suited for distributing services
- tools: third-party call control, redirect servers, proxies
- allow both chain-of-servers and central-coordinator model

Interworking with 911

Problems for interworking with current 911:

- identify public safety answering point (PSAP)
- but: gateway can be anywhere in the U.S
- need national database (e.g., SIP redirect server) that can return 10-digit E.164 number
- determine location – smart Ethernet sockets? SNMP?
- identify caller location ➠ IETF WG

911

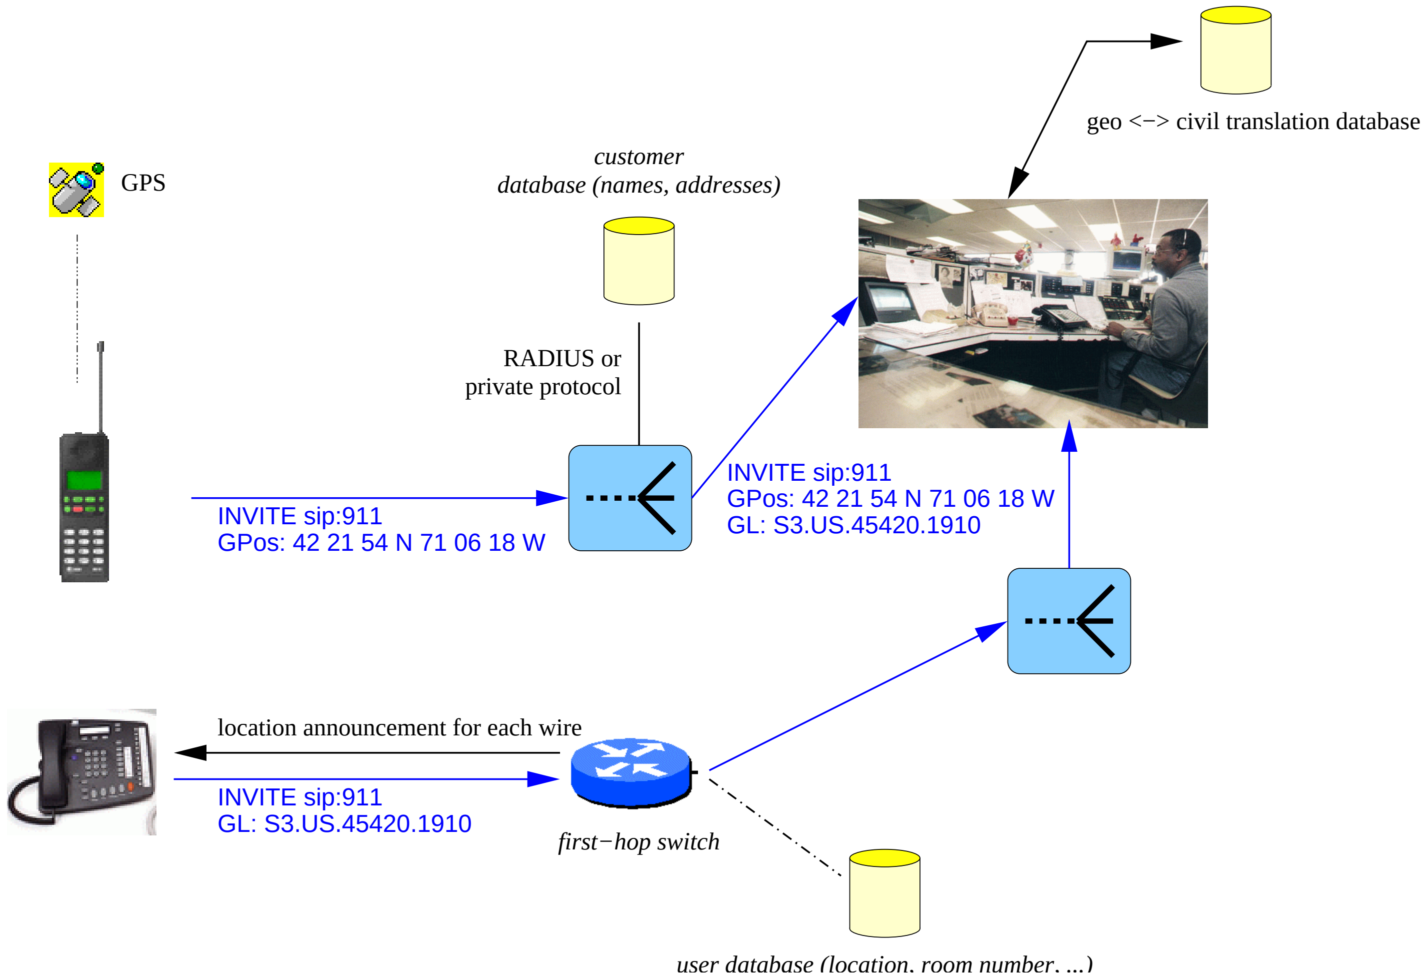
geo <−> civil translation database
GPS
customer
database (names, addresses)
RADIUS or
private protocol
INVITE sip:911
GPos: 42 21 54 N 71 06 18 W
GL: S3.US.45420.1910
INVITE sip:911
GPos: 42 21 54 N 71 06 18 W
location announcement for each wire
INVITE sip:911
GL: S3.US.45420.1910
first−hop switch
user database (location, room number, ...)

SIP benchmarking

- provide guidance to operators – “how many proxy servers do I need?”
- workload hard to characterize: 20 INVITE/REGISTER (call center) vs. 6 REGISTER/INVITE (mobile)
- initial workload: INVITE/200, INVITE/480, REGISTER, for TCP or UDP
- separate issue: overload behavior, protocol implementation robustness
- on-going effort; draft soon

Challenges and obstacles

- scalable device configuration
- PSTNv3
- “walled garden”
- service infrastructure

Potential obstacles

- SIP as transport – for legacy signaling
  - due to proxies, UDP not designed for volume data
  - doesn't add significant value
- NATs and firewalls – can engineer around them, but ugly
  - leads to IP-over-HTTP solutions, defeating firewall
  - proxy boxes outside NATs

“Walled garden” model

- 3G wireless carriers adopting SIP, but used to closed services
- SIP users should be able to use any proxy for services, not just carrier service
- typical users have many identities (and, thus, servers):

work	hgs@cs.columbia.edu
travel	schulzrinne@yahoo.com
home	henning@schulzrinne.leonia.nj.us
professional	h.g.schulzrinne@ieee.org

- hard to prevent: SIP can use any port number
- if not, requires draconian restrictions on IP packets, not just filtering port 5060 (SIP port)
- also, services may be split across servers

So I want to build a SIP network. . .

Ready for trials, but probably not quite for shrink-wrap status:

- installation and operation still requires fair amount of expertise
- lots of web and email experts, few SIP experts
- needs some external infrastructure: DHCP and SRV, possibly AAA
- inconsistent configuration for Ethernet phones (being worked on)
- SIP phones still more expensive than analog phones ➠ hard to justify PBX replacement (incremental cost)
- no just-download or ship-with-OS “soft” clients

Need for service infrastructure

- need carriers that offer SIP gateways
- without having to provide SS7 connectivity
- with *outbound* PSTN calling
- with *inbound* calls and *number portability* – need to be able to keep old PSTN numbers
- either IP Centrex model or in-house servers – like ISP services for email or web
- for commercial-grade conferences, need nailed-up Internet connectivity, orderable (at least) by web page – across providers!
- PBX revenue already decreasing

Why aren't we junking switches right now?

What made other services successful?

email: available within self-contained community (CS, EE)

web: initially used for local information

IM: instantly available for all of AOL

All of these ...

- work with bare-bones connectivity ($\geq$ 14.4 kb/s)
- had few problems with firewalls and NATs
- don't require a reliable network

Why aren't we junking switches right now?

Telephone services are different:

- reliability expectation 99.9% $\nearrow$ 99.999%
- PC not well suited for making/receiving calls – most residential handsets are cordless or mobile
- business sets: price incentive minor for non-800 businesses
- services, multimedia limited by PSTN interconnection
- initial incentive of access charge bypass fading (0.5c/min.)
- international calls only outside Western Europe and U.S.

Standardization

- SIP working group is one of the most active in IETF
- located in “transport” area, but really an application
- about 80 active Internet drafts related to SIP
- typically, 400 attend WG meetings at IETF
- but few drafts are working group items
- 80-20% – 80% of the technical work takes 20% of the time

IETF is getting slower

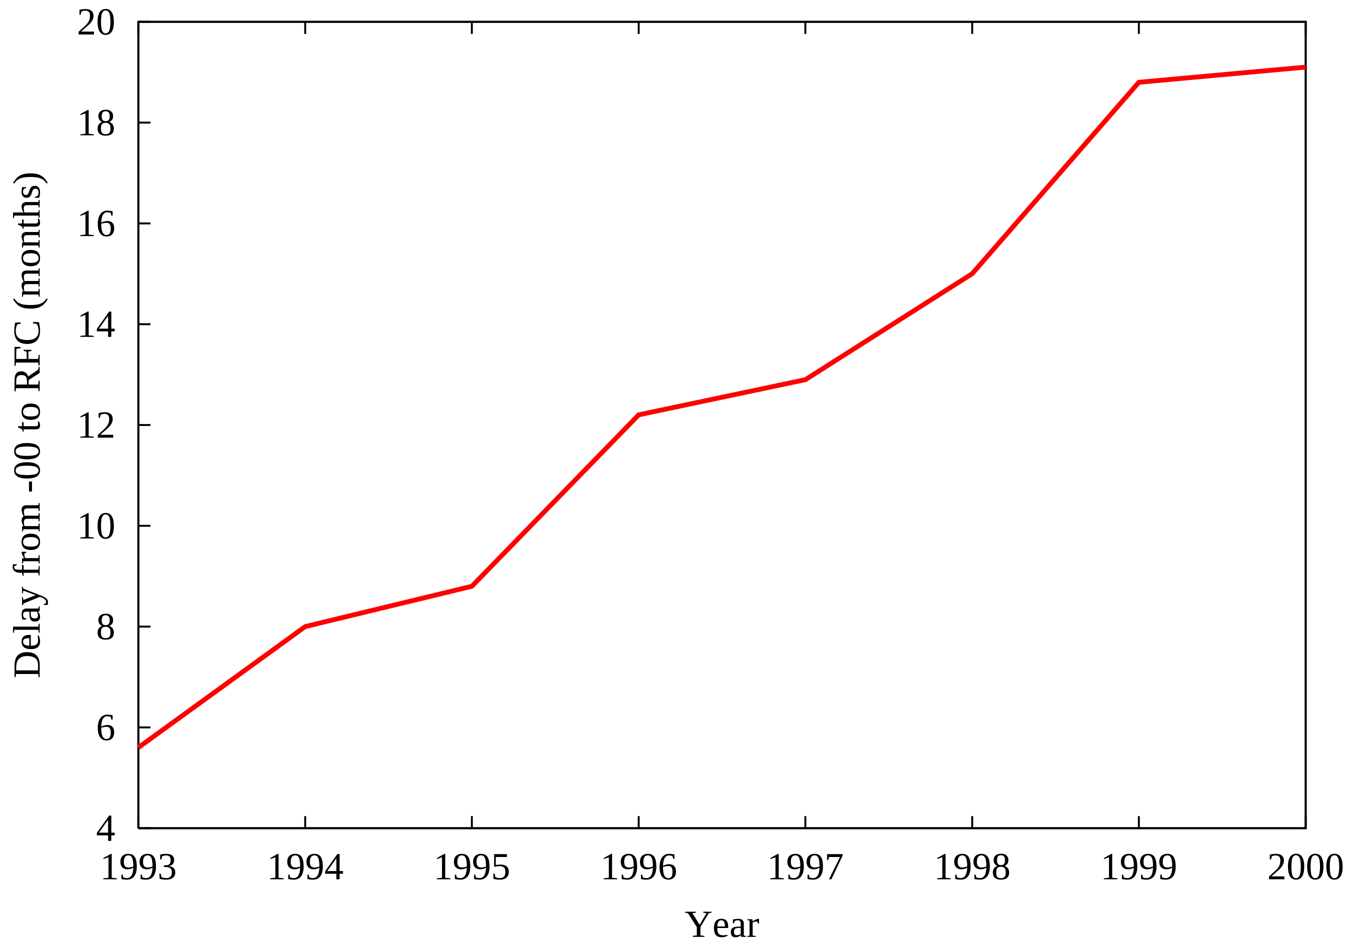

Standardization

- interaction with resource reservation
- caller preferences (“no mobile phones, please”)
- interoperation with ISUP (“SIP-T”)
- call transfer and third-party control
- conferencing: central server, end system, full mesh
- server benchmarking and scaling
- requirements for deaf users
- call processing language: coordination with iCal

Standardization challenges

- keep complexity in check
- remove, rather than add, features to base spec: Via hiding, PGP encryption
- new crypto security: S/MIME

SDPng

- current version of SDP limited in functionality
- e.g., negotiation of capability sets difficult
- can't group media (pick one or the other)
- MMUSIC developing SDPng
- XML-based description of capabilities and actual codecs and addresses used

Conclusion

- SIP maturing – base stable, extension in progress
- avoid creating PSTN replica
- leverage, not inhibit, Internet flexibility
- significant deployment challenges remain

For more information. . .

SIP: http://www.cs.columbia.edu/sip

RTP: http://www.cs.columbia.edu/~hgs/rtp

Papers: http://www.cs.columbia.edu/IRT