



A Sip of SIP

Avshalom Houri
Lotus Software - IBM SWG
November 2003

Abstract:

SIP (Session Initiation Protocol) is a standard protocol that covers a wide variety of applications. The document explains why SIP is a powerful real-time collaboration protocol. SIP is currently being adopted by many companies and is evolving more rapidly than any other existing real-time collaboration protocol. This document provides background information on SIP and touches the wide variety of applications whose implementation is either made easier or made possible, via SIP. This document concludes with a set of questions and answers on SIP.

Please send comments on this document to:
avshalom@il.ibm.com

Disclaimer

THIS DOCUMENTATION IS PROVIDED FOR REFERENCE PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS DOCUMENTATION, THIS DOCUMENTATION IS PROVIDED "AS IS" WITHOUT ANY WARRANTY WHATSOEVER AND TO THE MAXIMUM EXTENT PERMITTED, IBM DISCLAIMS ALL IMPLIED WARRANTIES, INCLUDING WITHOUT LIMITATION THE IMPLIED WARRANTIES OF MERCHANTABILITY, NONINFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE, WITH RESPECT TO THE SAME. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION, DIRECT, INDIRECT, CONSEQUENTIAL OR INCIDENTAL DAMAGES, ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS DOCUMENTATION OR ANY OTHER DOCUMENTATION. NOTWITHSTANDING ANYTHING TO THE CONTRARY, NOTHING CONTAINED IN THIS DOCUMENTATION OR ANY OTHER DOCUMENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF THE APPLICABLE LICENSE AGREEMENT GOVERNING THE USE OF THIS SOFTWARE.

Copyright

Under the copyright laws, neither the documentation nor the software may be copied, photocopied, reproduced, translated, or reduced to any electronic medium or machine-readable form, in whole or in part, without the prior written consent of IBM, except in the manner described in the documentation or the applicable licensing agreement governing the use of the software.

© Copyright IBM Corporation 2003

All rights reserved.

Lotus Software

IBM Software Group

One Rogers Street

Cambridge, MA 02142

US Government Users Restricted Rights - Use, duplication or disclosure restricted by GS ADP Schedule Contract with IBM Corp.

List of Trademarks

Lotus, , Notes Mail, Sametime and Lotus Workplaces are trademarks or registered trademarks of Lotus Development Corporation and/or IBM Corporation, in the United States, other countries, or both.

AOL is registered service marks of America Online, Inc.

MSN is registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Other company, product, and service names may be trademarks or service marks of others.

Table of Contents:

1	Executive Summary	3
2	One Protocol to... ..	4
2.1	General Notification & Presence	5
2.2	Conferencing.....	6
2.2.1	Awareness in Conferences.....	8
2.3	IP Telephony & VOIP	8
2.3.1	SIP Phones	9
2.3.2	SIP PBXs	9
2.3.3	A SIP Phone System	10
2.4	Mobile Devices	11
3	Questions & Answers	12
4	Summary	17
5	References.....	18
6	Acknowledgments.....	18
7	Appendices.....	18
7.1	The IETF.....	18
7.2	SIP History.....	19
7.3	SIP Related IETF Working Groups	21
7.4	SIP Concepts.....	21

1 Executive Summary

The Session Initiation Protocol (SIP) protocol standard was created by the IETF in 1996. The original intention of the protocol was to enable establishment of media (as audio or video) sessions between users. Since 1996 the protocol has evolved significantly and now it covers a wide range of real-time collaboration functionalities.

IBM is in the process of creating a SIP infrastructure as part of the Lotus Workplace (LWP) project. The SIP infrastructure that is built for LWP will enable IBM to build various SIP applications.

This document describes in high level why SIP is an important and powerful protocol. Following are the key points:

- Standard – SIP is a standard set of protocols that was created by the Internet Engineering Task Force (IETF). This is the same organizations that created TCP/IP, HTTP and many of the other protocols that actually enabled the Internet. A standard protocol has many benefits, which are described in this document. Among these benefits are: security, standard APIs, and much more.
- One protocol – Compared to other standards, SIP covers a very wide range of applications and features. This enables the implementation of almost all the real-time collaboration aspects of Lotus Workplace (LWP) using a single protocol. For IBM in general, and for LWP in particular, SIP was the natural selection.

- Widely adopted – SIP is being recognized as the de-facto protocol for real-time collaboration by several standards bodies and by many companies. It is being adopted as the protocol for the enormous markets of Voice over IP (VoIP) and telephony.
- Evolves quickly – SIP working groups are the most active groups in the IETF. Many new features are being added to the protocol. Some of these features are crucial for a real-time collaboration protocol. Among these features are: multi-way conferencing, emergency preparedness, and more.

This document provides detail for the claims above by:

- Giving a general, not too technical, background to the standards activity around SIP and other real-time collaboration protocols. The document shows that the activity around SIP is stronger and more promising than any other protocol.
- Describing a wide variety of scenarios and applications that are either enabled and/or made much easier to implement via SIP.
- Finishing with a set of questions and answers of "Everything you wanted to know about SIP, but were afraid to ask... ". The intention is to answer important questions that are not covered in other parts of the document.

2 One Protocol to...

If we will examine the real-time collaboration solutions that are being used by enterprises today, we will find one or more of the following solutions:

- Regular PBX
- A presence, IM and conferencing solution like Sametime™
- An email system. E.g. Lotus Notes™
- An electronic Calendar that is used by every employee.
- Most employees have mobile phone.

Replacing or enhancing the various solutions that are listed above with a single standard protocol, will benefit the enterprise in the following ways:

- Standard – A standard protocol is used instead of proprietary protocols. See section 31, (Questions & Answers) for more details on the benefits of using a standard.
- Integration – Using a single protocol instead of a collection of protocols will create a seamless real-time collaboration solution. Integration between various components will become more trivial.

The following subsections will describe the SIP based technologies that can replace existing real-time collaboration solutions. The following SIP based technologies are described:

- General Notification and Presence – Subscribe to a state of an object such as another user, and get notifications on status changes.
- Instant Messaging – Use SIP messaging solutions to communicate via IM with other employees in the company and talk with colleagues in other companies.
- Conferencing – Use the SIP session initiation to create one to one or multi-party media conferences.
- SIP Telephony & VOIP – Use SIP phones and SIP PBXs for intra and inter-company communications.
- Mobile devices – Use SIP based mobile architectures and services to enable mobile employees to be more connected to the company intranet while being out of the site.

2.1 General Notification & Presence

Recently, subscription and notification capabilities have been added to SIP. A protocol has been defined for enabling SIP users to subscribe on an object, and get a notification when the status of the object has changed. Using the concept of an *event package* it is possible to define the protocol details of subscription and notifications for different objects.

The SIMPLE working group at the IETF has defined an important event package – the presence event package. The presence event package describes how to use of SIP for awareness on other users and generation of a presence list (i.e., the buddy list™ as trademarked by AOL). To implement the presence list, a new SIP server component is introduced – the *presence server*. The presence server is responsible for getting and serving subscription requests from users. It provides this service by forwarding status change notifications to all appropriate subscribers.

Figure 2.1 shows a presence server and two user agents, where each user agent represents a different user. Alice subscribes on the status of Bob (1). The presence server will remember that Alice is interested in Bob, and once Bob has changed his status (2), Alice will get a notification (3).

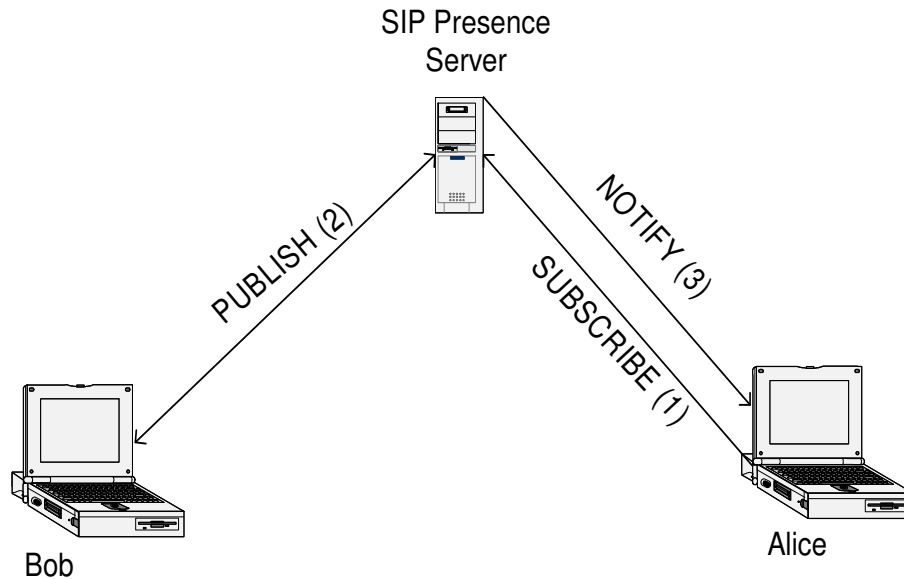


Figure 2.1 Subscription and notification

Later in this section we will see how the SIP notification technology can enrich the features and capabilities of other applications.

2.2 Conferencing

The original role of SIP was to enable establishment of media sessions. It does this by providing one basic and very important concept and two core services:

- SIP URI – A uniform resource identifier as bob@example.com. The SIP URI is the logical identifier of a user.
- Registrar – A SIP *user agent* (client) can send a REGISTER message to the registrar to advertise its availability. The register operation actually associates the SIP URI of the user with one or more access points, known as *contacts* in the SIP terminology. When a caller wants to establish a session with the callee, the information in the registrar is used to find the callee.
- Session management – To place a call to another user, a SIP user agent sends an INVITE message to the server, who will forward it to the target user. This message and those that follow allow for negotiation of media types, such as audio and video.

After the session is established, the media can be exchanged directly (peer to peer) between the session participants. The session can be one to one or have multiple participants.

The following are diagrams that show several types of SIP conferences.




Figure 2.2 Two way media session

Figure 2.2 shows a SIP session established between two parties. The INVITE messages have established the session parameters through the server and now the media can go peer to peer.

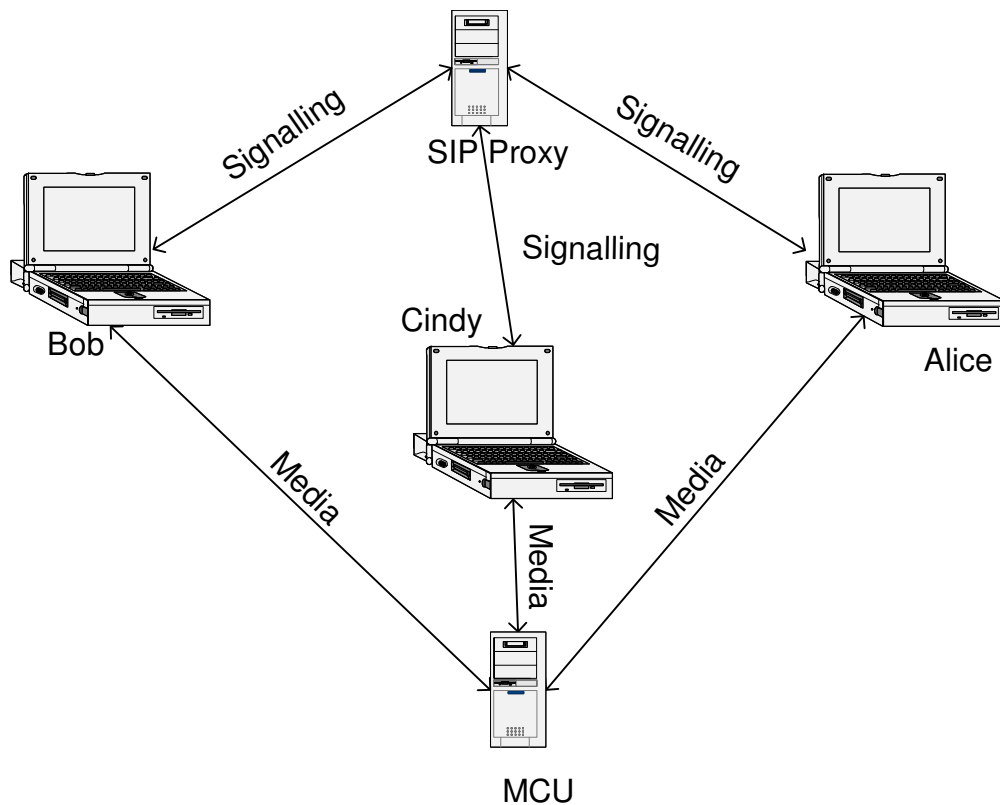


Figure 2.3 Three way media session

When a third party is added to the call, there is a need for merging the media streams of the participating users. One option is for a server to act as the Multipoint Control Unit

(MCU). The central server is responsible for receiving media streams from all conference participants, mixing the streams, and sending it to the other participants.

2.2.1 Awareness in Conferences

Using SIP presence technology to add the ability for awareness of other participants in the conference adds a whole new dimension to conferences. Imagine that, in addition to the voices and/or video, it would be possible to see the names of the participants, see the name of the person currently speaking, and much more. The experience of a conference becomes richer and a bit closer to the "real thing".

In addition to the presence list that can be shown as part of a conference, there is a need for conference policy. Conference policy determines the membership and media policy of the conference. This is a very complex issue and a new working group (XCON – Centralized Conference) has been established in the IETF especially for this purpose.

Note that there are existing solutions today, such as Sametime™ meeting technology, that enable a presence list in a conference and some form of conference control. However, these technologies use several proprietary protocols. Adding new features involves in depth knowledge of the proprietary protocols and the relationship between them. SIP enables all these features using a single standard protocol.

2.3 IP Telephony & VOIP

SIP can be regarded as the enabler protocol for telephony and voice over IP (VoIP) services. The following features of SIP play a major role in the enablement of IP telephony and VoIP:

- Location and status of users – Using the SIP registration and presence server it is possible to locate the contact points of other users. It is also possible to consult the capabilities that a user may attach to each contact in order to select the best contact for a given operation. The subscription and notification mechanism provided by the presence server adds the dimension for receiving constant notification on the status of users. This enables, for example, selection of the best time to call a user.
- Who is in the call? – In a phone conference call, SIP enables awareness of the call participants. Adding indicators of 'who is speaking now' eliminates the embarrassment of guessing the speaker's identity...
- User centric – In the "old" phone systems the logical unit of operation is the phone extension, while in the new SIP world the logical unit is the user. Calls can be created and directed according to the various registrations of the user, not to a specific extension that is assigned to the user. This different perception opens a very wide range of applications in the SIP world that could not be available in the "old" telephony world.
- Media negotiation – The inherent SIP mechanisms that enable negotiation of the media used in a call, enable selection of the appropriate codex for establishing a call between the various devices. This way, less advanced devices can participate in the call, provided the appropriate codex is selected.

- Tel URIs and PSTN gateways – Work done in the IPTEL (IP Telephony) group of the IETF enables mapping between regular telephone numbers and SIP URIs. Additional work has been done in creating gateways between the IP world and the legacy PSTN world. All this work is an enabler for merging the IP and the old telephony world.

2.3.1 SIP Phones



Figure 2.4 A Cisco 7960 SIP Phone

SIP phones are new kind of phones, which are very much like computers built to be phones. Consider having a phone on which you can see a presence list and in the presence list it is possible to see who is available and who is not available. It will also be possible to use a wide range of services that are not available in the standard or even the smart phones of today. Some examples: use an 'alert me' service in order to receive notification when a user is available, initiate a call when the user is available, and many other services that are bound only by our imagination (and some development).

2.3.2 SIP PBXs

The Private Branch eXchange (PBX) is a telephone switching center that is owned by a private business, compared to one that is owned by the common carrier or Telephone Company. Similar to the flexibility of the SIP phones, a SIP PBX can have a wide range of possibilities that are not possible in the legacy PBXs available today. In addition, using a standard protocol for the PBX enables a company to switch PBX providers without the need to replace all the end devices as is done today. Consider all those "smart" phones that you need to replace when switching a PBX...

Figure 2.5 illustrates a SIP PBX system. A SIP *proxy* that sits in the front of the system has the knowledge to route calls to the appropriate devices according to several sources. One source can be the calendar of the user specifying in which room the user is in at specific times and whether he/she can be disturbed. Other companies may use an automatic tag system that is worn by users and the call is routed to the nearest phone system.

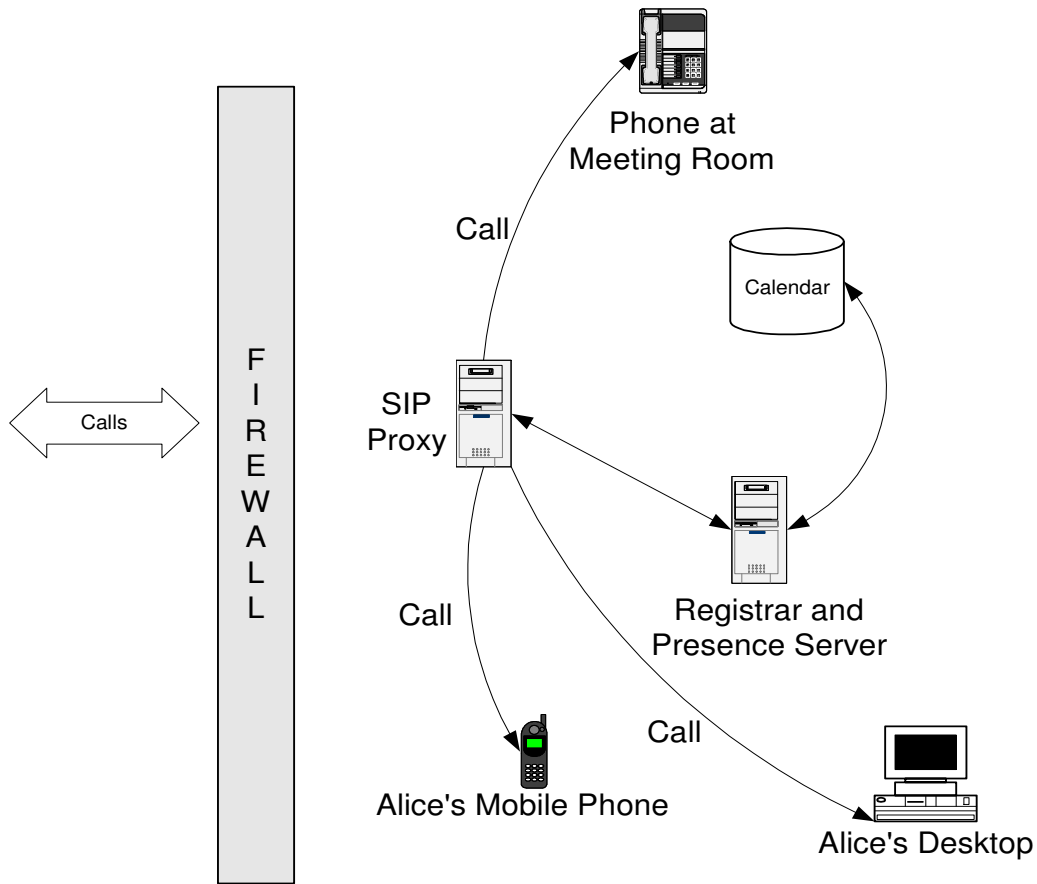


Figure 2.5 SIP PBX

2.3.3 A SIP Phone System

Figure 2.6 shows a bigger picture of a SIP-based phone system and a PSTN gateway. We can see the PSTN gateways enabling routing of calls that are initiated in the PSTN system to the SIP system, and vice versa.

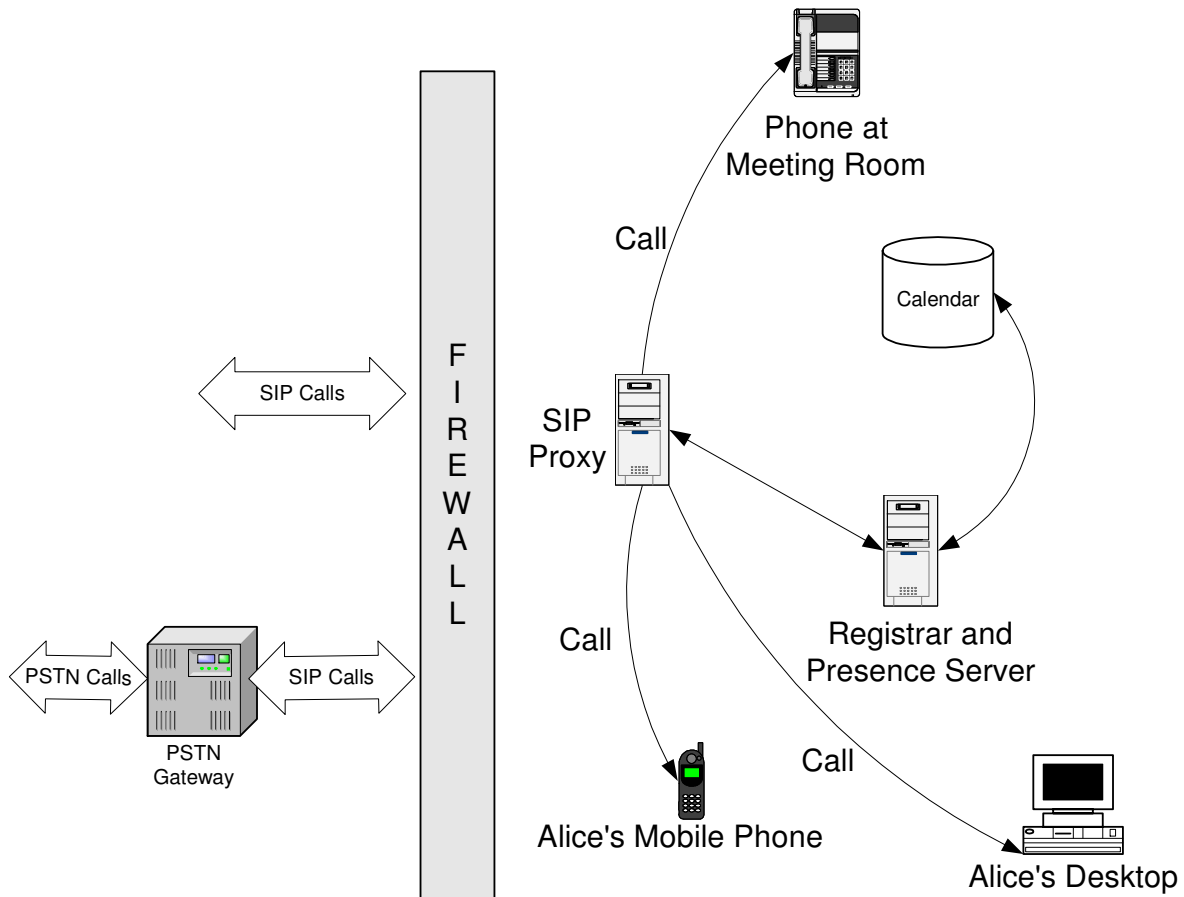


Figure 2.6 SIP Phone System and a PSTN gateway

2.4 Mobile Devices

The 3GPP (3rd Generation Partnership Project – www.3gpp.org) is an organization whose goal is to "produce globally applicable Technical Specifications and Technical Reports for a 3rd Generation Mobile System based on evolved GSM core networks and the radio access technologies that they support..." This organization has adopted SIP as one of the major protocols to be used in the architectures defined by the organization. The organization also has a liaison agreement with the IETF in order to be able to synchronize work on standards in both organizations.

The 3GPP organization has created the IMS architecture [IMS], which is based on the SIP standard. From the document:

"In order to achieve access independence and to maintain a smooth interoperation with wireline terminals across the Internet, the IP multimedia subsystem attempts to be conformant to IETF "Internet standards". Therefore, the interfaces specified conform as far as possible to IETF "Internet standards" for the cases where an IETF protocol has been selected, e.g. SIP."

The 3GPP2 (3rd Generation Partnership Project 2 – www.3gpp2.org) is another organization that deals with goals for mobile systems, but for a different set of network protocols (Who said that the standard world is perfect?). The 3GPP2 organization like 3GPP is also basing its architecture on SIP. The dependencies list of 3GPP2 on the IETF work in SIP can be seen in [3GGP2-DEP]

Another standards organization that deals with mobile systems is the OMA (Open Mobile alliance – www.openmobilealliance.org). The mission of the organization is as follows:

The mission of the Open Mobile Alliance is to grow the market for the entire mobile industry by removing the barriers to global user adoption and by ensuring seamless application interoperability while allowing businesses to compete through innovation and differentiation.

The OMA is also considering SIP as one of the major protocols for the services it defines.

In short, the bottom line is that organizations that are defining architecture and services for next generation mobile systems are looking at SIP as the major protocol to be used in these systems. Adopting SIP now enables better preparedness for the new mobile systems.

3 Questions & Answers

Everything you wanted to know about SIP but were afraid to ask... ☺

Some of the material here may repeat what is said above, but is provided here in the format of questions and answers to clarify specific issues.

(Q) Why is a standard protocol better than a proprietary protocol?

(A) It is assumed that most readers will agree that standards are very important for the industry, and that the industry should adopt them whenever possible. For those who are not yet convinced, the following are some reminders:

- More security – When a protocol undergoes standardization, it is scrutinized by the same community that created it. This is even more so at the IETF. Hence, the security of the protocol can be verified only once and not for every proprietary protocol. In addition, a single company will rarely have the expertise and resources to do the same level of security inspection for the protocol as can be done by the standards community.
- More interoperability – Standards prevent walled gardens, or at least the "official" excuse for having a walled garden... Walled gardens are, for example, the presence and IM communities of AOL, Yahoo, MSN, etc. Users of these services get very good service but have no interoperability with their family, colleagues, and friends who use the other communities. Companies justify this lack of interoperability as the result of no interoperability standard.
- More features – When a standard is available, groups that develop features (e.g., third party developers) can design and implement the feature only once. There is

no need to design and implement the feature for every proprietary protocol, not to mention the need to design interoperability of the feature with the various proprietary protocols.

- More freedom – When a company buys a solution based on a standard, the company can purchase additional features and solutions from third party developers that develop for the standard and not for the product of a specific company. In addition, a company that buys a standards-based solution is not "caged" by the supplier of the solution. The company can always switch to another solution provider, without changing clients and thus minimize re-deployment costs.

(Q) Isn't a proprietary protocol more efficient?

(A) It is true that proprietary protocols are usually more efficient than standard protocols. There are several reasons for this, which are beyond the scope of this document. However, the diminished efficiency of the standard solution is outweighed by the benefits of the standard. For example, today, no one would consider not using HTTP because it is not as efficient as a proprietary solution could have been.

(Q) There are already several widely used standards in the real-time collaboration area. Why do we need another one?

(A) Let's review the status of the real-time collaboration protocols, prior to the acceptance of SIP.

- Proprietary presence and IM solutions – Current presence and IM communities like Sametime, AOL, Yahoo, MSN and ICQ use proprietary protocols. Most of these communities are also referred to as "walled gardens," meaning the users are getting very good services, but they can not see peek outside of the garden wall.
- ITU solutions – Complicated solutions for multimedia and Voice over IP (VoIP) such as H.323 and T.120.
- Wireless Village – The Wireless Village protocol was created by the wireless village standard organization (now part of the Open Mobile Alliance organization (www.openmobilealliance.org), and is a presence and IM protocol designed for mobile devices and wireless operators. However, no public deployment of Wireless Village by a mobile operator has been declared as of the date of this document.
- Telephony solutions – The PSTN and the mobile world use several protocols for their operations. These protocols include: SS7, GSM etc.

As can be seen from the above list, the world of real-time collaboration protocols is divided. In addition, there is no single base protocol that can be used by all other protocols. A single base protocol should supply a basic registration, locator, and awareness service that is used by all the other protocols. Having a basic protocol gives us the ability to dramatically enhance services, for example, for conferences and other types of media oriented services. As a very naïve example, if I want to set up a call with

someone, being able to see the status of the other user can be very helpful in setting up the call.

SIP stands for Session Initiation Protocol. The intent was to enable initiation of a media session with another user. As part of the session establishment the need for a way to locate the other user was obvious; therefore, a registrar in which users register themselves was added to the protocol. Over the years, a lot of other capabilities were added to the protocol, arriving at a protocol that now covers almost all aspects of real-time collaboration. The following is a partial list of current SIP capabilities:

- Registration and locator – Enabling users to register themselves and other users to find them.
- Session establishment – Using SIP for setting up a media session, including media negotiation, call control, call forwarding, and more.
- Events notification – Subscription and notification of events. Of particular interest is the subscription and notification of presence information.
- Device capabilities – The registration model of SIP enables registering multiple devices by the same user (e.g., PC and mobile phone). In addition, SIP enables the definition of different capabilities for the devices, thus enabling the selection of the appropriate device when a session is established.
- Interfacing with "old" phone systems – A lot of work has been done towards creating gateways between the SIP world and PSTN world, including mapping between phone numbers and SIP identifiers, and more. These enhancements to SIP should enable easy interface between SIP and PSTN.
- Mobile world – The following standard organizations that deal with the mobile world have either adopted SIP as the de-facto protocol (3GPP, 3GPP2), or one of the basic protocols (OMA). The 3GPP and 3GPP2 organizations have defined an IP architecture that is based on the SIP protocol (IMS and MMD). The OMA is in the process of defining services over SIP.

(Q) Are there standard APIs that are built on top of the SIP protocol?

(A) Yes. The JCP (Java Community Process – jcp.org) have already defined several APIs for SIP. The following are two examples (JSR – Java Specification Request):

- JSR 32 – A low level Java SIP API (<http://jcp.org/en/jsr/detail?id=32>). This API specification provides a standard portable interface to share information between SIP Clients and SIP Servers, providing call control elements enabling converged-network applications.
- JSR 116 – An API that defines a high-level extension API for SIP servers. It enables SIP applications to be deployed and managed based on the servlet model. See <http://jcp.org/en/jsr/detail?id=116> for more details.

The importance of the APIs can not be exaggerated. APIs adds a new dimension to the standard. An application that exposes the standard API can be easily extended by plugging in components that use the API. For example JSR 116, enables extending the SIP server by writing servlets that will implement the required extensions.

(Q) Should we invest in SIP now?

(A) A short answer would be: the sooner the better. A few facts:

The following companies are busy developing SIP infrastructures and/or SIP devices. Most of these companies already have SIP software and devices ready:

- IBM
- Microsoft
- Cisco
- DynamicSoft
- Sprint
- Nokia
- Dozens of other companies.

The SIP standardization efforts are tremendous. The core protocol is already in a stable state and there are efforts to finalize the specifications for additional services and multi-way conferences, emergency preparedness, and much more. Dozens of companies are highly engaged in these efforts. SIP also is responsible for the most active groups in the IETF (more than 150 active drafts at any given moment). In addition, as mentioned above, the organizations that are responsible for creating architectures and services for the mobile world are using SIP as the de-facto standard for their work.

The SIP Forum (www.sipforum.org), which deals with SIP products and deployment, has 34 companies as full members, including: IBM, Microsoft, Nokia, Ericsson, Hughes, Indigo, and more.

The readiness of SIP enables the creation of many new applications that were previously hard to build using the old protocols. Among these applications:

- SIP phones that have a presence list (buddy list™)
- Voice over IP applications that use the SIP session management and the capability to interface with PSTN gateways.
- SIP PBXs that use the following capabilities of SIP: session establishment and redirection, SIP registrar and more.

(Q) Can SIP support multiple devices for a user?

(A) One of the main design principles in SIP is its ability to support multiple devices for a user. A user can register multiple contacts for his/her address-of-record (AOR). Another user that needs to access the user can send the request through the server where the user is registered and the server can fork the request or try the various contacts channels for the user, one by one.

Another important point is the addition of user capabilities to user contact. User capabilities are hints to the capabilities of the contact. For example, indicating the devices that can support video stream will enable other users that need to set up a video session with the user to select only those contacts for the session setup.

(Q) Firewalls and NATs do not provide a good experience when deploying real-time services. Does SIP handle NATs and Firewalls?

(A) Yes.

NAT is a network component that translates between external and internal IP addresses. For example a company that wants to hide its IP addresses (for security or other reasons) can use a NAT that translates the IP of an outgoing connection to an external IP address. The problems occur when there are incoming connections that specify the internal IP address of a server behind a NAT; in these cases, the NAT does not know how to route the IP packets.

Firewalls are used for protecting a network from any incoming and/or outgoing packet connections. Only those packets allowed by the firewall administrator will pass through the firewall. SIP messages, and audio/video UDP packets in particular, will not pass through the firewall unless explicitly allowed by the firewall administrator.

The MIDCOM (Middlebox Communication) working group of the IETF has been working on this issue for a while. The emerging solution of the group is to enable applications behind a NAT or firewall (or any middlebox) to talk with the middlebox and notify the middlebox of the application needs for communication. A SIP proxy for example, will tell the middlebox that it should accept SIP messages coming from the outside network. This is a very simplistic explanation of the solution, but the point is that solutions exist and will be part of NATs and firewalls. See also <http://www.checkpoint.com/press/2002/voip061102.html> as an example of an existing firewall that has a solution for SIP.

(Q) What is the SIP security model?

(A) Each IETF protocol must have a security section that explains the security threats to the protocol and the possible solutions. The security section is revised by the security experts of the IETF, who provide feedback on any security holes. The core drafts of SIP have passed this review.

The actual details of the SIP security model are beyond the scope of this document.

However, we note that:

- SIP requires at least DIGEST authentication. In DIGEST authentication the user name and password are not passed in the clear but are encrypted using a challenge that is received from the server.
- The body of SIP messages can be encrypted using S/MIME. This way, only the two end points of the communication can see the message body.
- SIP enables secure communication between its components using a TLS connection. There is even a special SIP URI called SIPS (S for Secure). A SIPS URI indicates that it has traversed only secure connections (TLS) on its way to the SIP component where it currently resides.

(Q) What about XMPP (Jabber) protocol?

(A) XMPP is an effort by the Jabber organization and company (www.jabber.org and www.jabber.com) to standardize an existing presence and IM system. Jabber is open source and based on XML. The IETF has agreed to form a working group with a charter to standardize the Jabber protocol. Generally, the IETF does not allow competing standards to be created under its umbrella. However, in this case the IETF agreed due to Jabber being open source and widely used by a large community of users.

Being only TCP-based and tailored for the purpose of presence and IM via XML, XMPP is simpler than SIP. XMPP is very different from SIP in that there will be a need to use SIP and/or additional protocols in order to achieve the following:

- Session establishment
- Media negotiations
- Registration of multiple devices and capabilities
- Firewall solutions that were developed for SIP.

It might be possible to use XMPP as the presence and IM solution and SIP as a solution for other capabilities. This solution has a major drawback in that it does not use the same protocol for the registration of users and for the capabilities that are built above these registrations. For example, when implementing a SIP phone that has a list of users and their status, there will be a need to implement two protocols at the phone and two protocols at the servers. This adds the need to translate between the protocols and more.

Participating in the IETF, it is obvious how much industry attention is being given to the SIP protocols and how much attention is given to the XMPP protocols. The numbers are incomparable, where almost all the attention and activity are directed at SIP.

(Q) How much has IBM committed to SIP and in particular, what is IBM doing?

(A) IBM is committed to SIP. A new presence server is being developed based on SIP rather than the proprietary protocol used in Sametime.

The IBM SIP infrastructure will serve several purposes. It will serve as the infrastructure for Lotus Workplace and provide SIP building blocks for other SIP applications in IBM, including: a SIP stack, SIP proxy, SIP registrar, SIP presence server and more.

Of particular interest is the way the SIP infrastructure is built in relation to the WebSphere, the IBM flagship applications server. One of the main components of the SIP infrastructure in IBM is the SIP container, which is built over an existing WebSphere web container and uses the standard JCP (Java Community Process) API – JSR 116 (<http://jcp.org/en/jsr/detail?id=116>). Building SIP this way should enable full integration of the SIP technology into WebSphere.

4 Summary

In the previous sections we have shown why SIP is an important protocol for real-time collaboration. We have described how SIP can be harnessed to do various tasks that are currently done by multiple protocols, most of them proprietary.

At this point it should be obvious that SIP should be considered very seriously by IT decision makers in any enterprise that wants to prepare now for the future.

5 References

[IMS]

3rd Generation Partnership Project;
Technical Specification Group Services and System Aspects;
IP Multimedia Subsystem (IMS);
Stage 2 (Release 6)

http://www.3gpp.org/ftp/Specs/2003-09/Rel-6/23_series/23228-630.zip

[3GPP2-DEP]

3GPP2 dependencies on IETF protocols

http://www.3gpp2.org/Public_html/IETF/IETF_dependencies.cfm

[SIP-DEM]

SIP Demystified
Gonzalo Camarillo
McGraw-Hill Telecom
ISBN 0-07-137340-3

[INT-COM-SIP]

Internet Communications Using SIP
Henry Sinnreich
Alan B. Johnston
Wiley Computer Publishing
ISBN 0-471-41399-2

6 Acknowledgments

Thanks to the following people for helping in the preparation of this document:

Roger Eames
Pat Galvin
Eran Flom
Marjorie Schejter
Chani Sacharen

7 Appendices

7.1 *The IETF*

The following is a short background of the IETF (Internet Engineering Task Force). Further information is available in the IETF site (www.ietf.org) and in the TAO of the IETF <http://www.ietf.org/rfc/rfc3160.txt>.

The IETF is probably the most important standards organization for the Internet. This is the organization which created the IP protocols that enabled the creation of the Internet

and the HTTP protocol, which in turn enabled the creation of the WWW (World Wide Web).

The IETF is known for the strong security of its protocols due to a process that requires a security section in each draft. This security section is inspected by people from the security area in the IETF.

The IETF is divided into several work areas:

- Applications
- General
- Internet
- Operations and Management
- Routing
- Security
- Sub-IP
- Transport

Each area has two area directors and several working groups (WG). Each working group has at least two group chairs that are responsible for managing the group meetings and are the main contacts of the group with the area directors.

The area directors of all working areas form a group that is called the IESG – Internet Engineering Steering Group. The role of the IESG is to "steer" the work of the IETF. A draft starts as an *internet-draft* and is discussed in the mailing list and in meetings of the working group. After the group feels that it is ready, the draft passes a WGLC (Working Group Last Call). After it passes a WGLC it is passed to the IESG for final approval. When a draft is approved by the IESG it becomes an RFC (Request For Comments). When a draft reaches the state of an RFC it is considered as mature enough to be considered for wide implementation.

7.2 SIP History

Figure 7.1 is a photograph that was taken at the fifty-fourth IETF. It celebrates the release of RFC3261 which is the core standard document for SIP. Note the indication of the RFC number in the peoples' hands.



Figure 7.1 RFC 3261

Standing from left to right are:

- Allison Mankin – The IETF IESG transport area director who was responsible for mentoring the process.
- Robert Sparks – Dynamicsoft
- Jon Peterson – Neustar
- Jonathan Rosenberg – Dynamicsoft
- Alan Johnston – WorldCom
- Gonzalo Camarillo – Ericsson
- Henning Schulzrinne – Columbia University

Not shown in the picture, but on the back of the T-shirts, the following slogan appeared (an adaptation from Tolkien's Lord of the Ring):

*One Protocol to rule them all,
One Protocol to find them,
One Protocol to bring them all
And in the darkness bind them*

No wonder the T-shirts were black...

In 1996 the IETF merged two proposals for session initiation protocols, one by Mark Handley on Session Initiation Protocol (SIP) and the other by Henning Schulzrinc (appears in the picture above) that was named Simple Conference Invitation Protocol (SCIP).

SIP was approved as RFC 2543 (<http://www.ietf.org/rfc/rfc2543.txt>) in March 1999. In June 2002, RFCs 3261-3265 were approved by the IESG. These RFCs create the core protocol set for SIP and they were achieved after an enormous amount of work by the authors and the IETF members.

7.3 SIP Related IETF Working Groups

Following the publication of the core SIP drafts (RFCs 3261-3265), there continues to be very intensive activity in the IETF around SIP.

The following working groups at the IETF work on SIP directly and indirectly:

- SIP – The original SIP group. Now deals with extensions to the standard.
- SIPPING – Session Initiation ProPosal INvestiGation. This group is responsible for examining new requirements for the SIP protocol and deciding if an extension to the protocol is required.
- SIMPLE – SIP for Instant Messaging and Presence Leveraging Extensions. This group is responsible for the IM and presence extensions for SIP.
- XCON – Centralized CONferencing. This group is responsible for creating a set of protocols for tightly coupled multimedia conferences. Not all the protocols of this group will be SIP, but some (if not most) of them will be.

7.4 SIP Concepts

The following are some SIP concepts and terms. Some of the concepts were described above but here we give a more technical description.

- **URI** – Uniform Resource Identifier. For the purpose of this document let us say that it is a unique string that identifies a resource and a user in particular.
- **Address-of-Record:** An address-of-record (AOR) is a SIP or SIPS *URI* that points to a domain with a location service that can map the URI to another URI where the user might be available. Typically, the location service is populated through registrations. An AOR is frequently thought of as the "public address" of the user.
- **User Agent (UA)** – A logical entity that acts on behalf of the user (AKA client)
- **Registrar:** A registrar is a server that accepts REGISTER requests and places the information it receives in those requests into the *location service* for the domain it handles.

- **Contact:** Provides a SIP *URI* that can be used to contact that specific instance of the *UA* for subsequent requests.
- **Location Service:** A location service is used by a SIP redirect or proxy server to obtain information about a callee's possible location(s). It contains a list of bindings of *address-of-record* keys to zero or more contact addresses. The bindings can be created and removed in many ways; this specification defines a REGISTER method that updates the bindings.
- **Proxy, Proxy Server:** An intermediary entity that acts as both a server and a client for the purpose of making requests on behalf of other clients. A proxy server primarily plays the role of routing, which means its job is to ensure that a request is sent to another entity "closer" to the targeted user. Proxies are also useful for enforcing policy (for example, making sure a user is allowed to make a call). A proxy interprets, and, if necessary, rewrites specific parts of a request message before forwarding it.

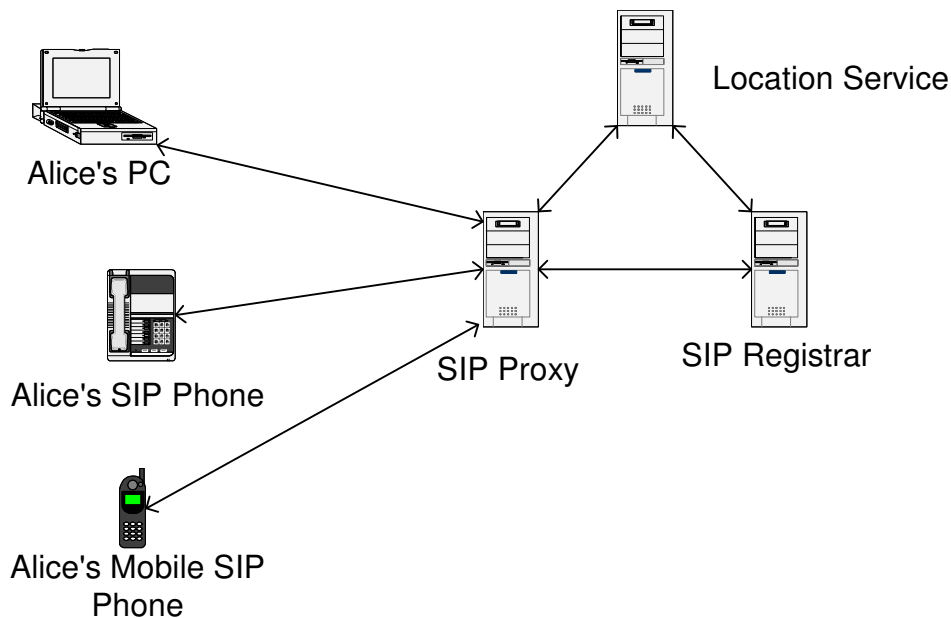


Figure 7.2 Basic SIP server

In Figure 7.2 we Alice's SIP user agents connected to a SIP server through a SIP proxy. In this case the SIP server consists of a proxy, registrar and location service.

