Connective Middleware for Voice and Integrated Communications

Submission to TERENA Networking Conference 2004

Ben Teitelbaum < *ben@internet2.edu* >
3025 Boardwalk, Suite 200
Ann Arbor, MI 48108, USA

Keywords

Voice over IP (VoIP), Session Initiation Protocol (SIP), Peer-to-Peer (P2P), Connective Middleware, SIP for Instant Messaging and Presence Leveraging Extensions (SIMPLE)

Abstract

To engage Internet2 users in an exploration of peer-to-peer, real-time communications, these users require connectivity. Superior networking connectivity alone is not sufficient. Directory-enabled call and presence routing middleware must be deployed to connect applications with applications and ultimately users with users. Internet2 is working to grow this connectivity and to encourage large-scale experimentation at the edge. The best new communications applications will combine peer-to-peer design principles with new, centrally-managed enterprise services. Identity management, authentication, location services, and rich presence are all best implemented by the enterprise. Several new Internet2 initiatives are surveyed, along with some of the challenges ahead.

1 Introduction

Throughout the history of communications networks, the highest value applications have been those that support generic interpersonal communication and collaboration. Electronic mail, telephony, and instant messaging have long been "killer applications", enabling everything from frivolous gossip to collaborative research among Nobel Prize winning physicists. Technology and business elites have historically been slow to recognize the value of these general purpose person-to-person communications, tending instead to emphasize content and relatively narrow applications (*e.g.* distance learning, remote medical consultation) [1].

Several recent Internet2 initiatives are working to enable next-generation campus communications services through the development and deployment of connective middleware that is both application and media neutral. These initiatives are building upon Internet2's unique confluence of superior network connectivity, growing foundation of core middleware services, and youthful demographic to increase the scalability, security, reliability, and richness of voice, video, and instant messaging applications.

2 Advanced VoIP

Voice is the primary means of real time human communication and, for most users, the only real time network communications medium in use. Any serious effort at advancing interpersonal communications applications must address voice. Unfortunately, too much attention has been on matching the fidelity and feature set of plain old telephone service (POTS) at a lower cost, while not enough has been on improving the experience and the value of voice communications. The transition to voice over IP (VoIP) provides a unique opportunity to make voice not only cheaper, but better. Voice communications is central to all collaboration; it cannot be considered a solved problem. Instead, the research and education networking community must work to promote *advanced* voice applications.

The Internet2 VoIP [2] and PIC [3] working groups are working to improve upon POTS along a number of dimensions[1], including: fidelity (*e.g.* wideband and binaural audio), privacy, addressing, and integration with other media (*e.g.* instant messaging, video, and, most interestingly and subtly, presence). Rather than promote a specific vision for a next-generation user communications experience, these groups are working towards the deployment of core connective middleware that will allow users to explore a large variety of applications, human-computer interfaces, and media types.

[1] Although the transition to VoIP is not without trade-offs, the success of mobile telephony has shown that users are willing to make concessions to benefit from new functionality (*e.g.* sacrifice fidelity and reliability to achieve mobility).

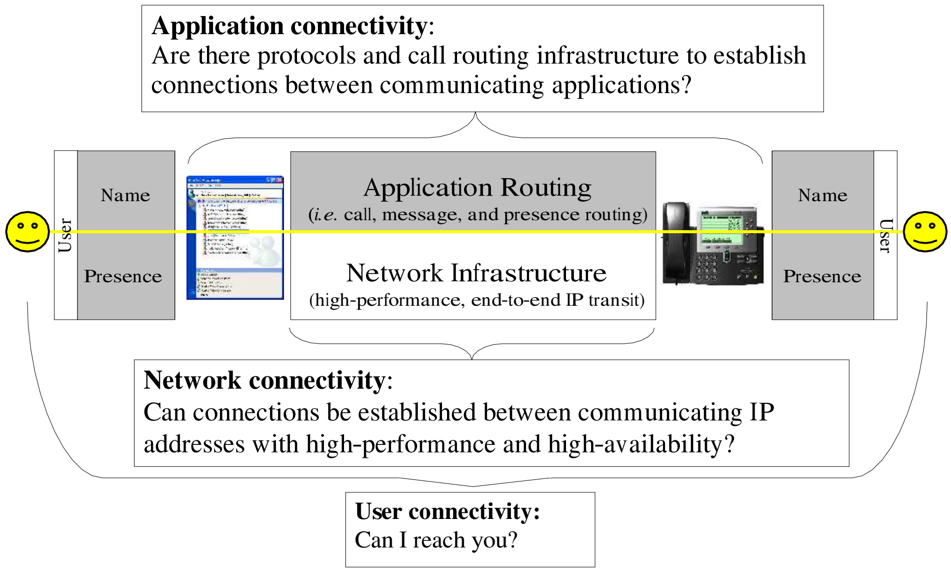

Figure 1: User-to-User Connectivity Requires New Middleware

3 Connectivity

Internet2 member universities provide communications services to 3.8 million students [4, 5]. Eager to adapt to new modes of communication (*e.g.* mobile telephony, instant messaging), these users are already opting out of campus communications offerings (*e.g.* residential telephony). To engage this demographic in a large scale exploration of peer-to-peer (P2P) real-time communications, users must be able to *connect* to each other.

The research and education networking community has historically focused on core networking infrastructure (*e.g.* Abilene, GÉANT, NLR) and services (*e.g.* IPv4, IPv6, multicast). More recently, this community has embraced the need for core middleware to discover resources (*e.g.* GRID) or broker inter-institutional authentication and authorization [12, 13]. We argue that the community also needs to embrace *connective middleware* to broker user-to-user communications sessions.

Currently, Internet2 users enjoy superior IP connectivity. But this alone is insufficient. New core middleware is required to route calls between communicating applications.

The combination of good networking and application connectivity is still not sufficient. To support the goal of connecting users with each other, users require names (*e.g.* +1-234-567-8900 or bob@bigu.edu) and presence services, which at a minimum provide a binding between a user name and a set of IP address/port pairs where the user can be reached.

The layering of these different levels of connectivity is shown in Figure 1.

4 SIP

The IETF Session Initiation Protocol (SIP) [6] is a signalling protocol for creating, modifying, and terminating real-time internet media sessions. SIP is an open standard and it leverages a number of widely-deployed open standards (*i.e.* SDP, RTP, MIME, DNS) to accomplish its goals. SIP supports traditional telephony routing features as well as instant messaging and presence. SIP also enjoys strong and broad industry support, ranging from core network equipment providers to desktop applications software vendors to the 3G wireless standards bodies, who, through the 3GPP, have standardized upon SIP for call control signalling.

SIP signalling, routing infrastructure, and services represent exactly the kind of connective middleware that can provide the necessary application-level and user-level connectivity. Because SIP is media-neutral and, after call establishment, purely peer-to-peer, rich SIP connectivity throughout Internet2 will enable experimentation with a great variety of applications, codecs, user interfaces, and media types. Toward this end, Internet2 members institutions are deploying SIP-based call routing infrastructure [7] and experimenting with SIP-based presence services and integrated communications applications [3].

5 SIP.edu

To grow SIP connectivity and to build a community of universities that works together to deploy new enterprise SIP services, Internet2 has launched the SIP.edu initiative [7]. The initial SIP.edu architecture specifies connective middleware that provides inbound SIP voice reachability and promotes the use of the email address, as a "one-stop" electronic identity. The architecture makes heavy use of existing campus infrastructure (shown in gray in figure 2). Specifically, the architecture leverages the legacy campus private branch exchange (PBX), the domain name service (DNS), and the campus Person directory. New components (shown in white in figure 2) include: a SIP-PRI gateway to the campus PBX, SRV [8] and NAPTR [9] records in the campus DNS zone file, and a SIP proxy/registrar server integrated with the campus Person directory.

A simplified call-flow from Alice (using a "soft phone" on her laptop) to Bob at BigU would proceed as follows:

1. To initiate a call to Bob, Alice types "bob@bigu.edu" into her SIP user agent (UA);
2. Alice's UA performs a DNS SRV resolution on `_sip._udp.bigu.edu` discovering the name and IP address of BigU's SIP proxy (`proxy.bigu.edu`)[2];
3. Alice's UA sends the SIP message `INVITE sip:bob@bigu.edu` to BigU's SIP proxy;
4. The proxy consults the campus Person directory, learning Bob's extension;

[2]The UA may first issue a NAPTR request, though few do at this time.

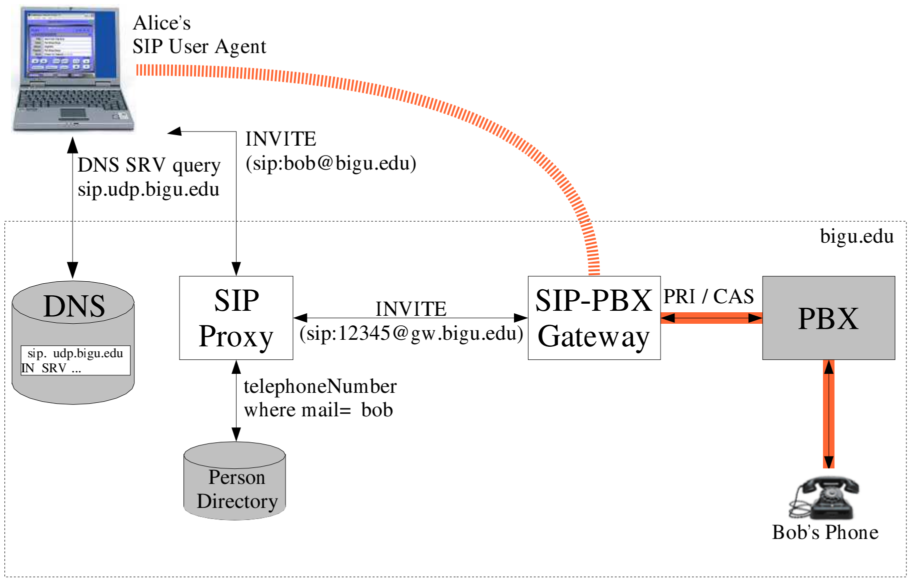

Figure 2: Initial SIP.edu Architecture

5. The proxy re-writes the INVITE to `INVITE sip:12345@gw.bigu.edu` and sends it to the gateway `gw.bigu.edu`;

6. The gateway rings Bob's desk phone (ext. 12345) through the PBX and gateways the VoIP media to/from the PBX during the call.

Thus far, four universities have implemented the initial SIP.edu architecture, providing SIP reachability to nearly 100,000 Internet2 users. To make it a "no-brainer" for a university to implement SIP.edu, the project has secured corporate sponsorship and has published a "SIP.edu Cookbook" [10]. Work is underway to integrate several common SIP proxies with the widely-deployed LDAP eduPerson object class [11].

SIP.edu participants are also beginning to support user agent registration, allowing pure IP calls. In the example above, the proxy's call routing logic can be changed to ring Bob's IP phones first, if any are registered, and, otherwise, ring his PBX extension. In this manner, calls may be signalled and transported purely over IP, opening the door for users to experiment with advanced VoIP features like wideband audio and multimedia integration.

Voice is clearly moving from a carrier service offering to an enterprise IT function (like email). As the SIP.edu community grows, it will function as a support group and forum for those who are building and operating campus SIP services. It will also leverage new inter-institutional trust federations [12, 13] to authenticate calls and provide for interoperability with the PSTN.

6 Enterprise Presence and Location Services

The best new communications applications will combine peer-to-peer design principles with new, centrally-managed enterprise services. Because the enterprise controls the physical and networking environments of its users (at least for most of the day) and has extensive directory assets that describe the roles, personal attributes, and relationships of its users, the enterprise can add tremendous value to communications services in ways that carriers and applications service providers (ASPs) cannot. Presence and location services are excellent examples.

6.1 Presence

As described earlier, a presence service must provide at least a binding between a user name and a set of IP address/port pairs where a user can be reached. In its full generality, however, presence involves the exchange of user-level state information to facilitate communication. Users (and software agents acting on their behalf) "publish" user-level state information, while other users (and software agents) "watch", using the observed presence to inform decisions about when and how to initiate communication. Presence may indicate availability (*e.g.* available, busy), physical location (*e.g.* room 231; 3025 Boardwalk, Suite 200; Ann Arbor, MI), activity (*e.g.* in budget meeting, over at 2:15), media capabilities (*e.g.* voice and video: not OK, instant messaging: OK), or even mood (*e.g.* pensive, bored, frustrated).

"Conference Pounce" is an illustrative example application of the value of presence information. Alice urgently needs to have a 15 minute conference call with Bob and Carmen. She checks their presence and discovers that Bob is on an airplane and that Carmen is on the phone. She further discovers that Carmen has a very full schedule. Rather than "play phone tag" with Bob and Carmen or attempt to schedule a meeting with them, Alice configures her integrated communications client to watch Bob and Carmen's presence, as well as her own, and to bridge everyone automatically when they are next all available for a 15 minute voice call. Several hours later, while Bob is pulling out of a car rental lot with his cell phone turned on and while Carmen is between meetings and Alice is at her desk, all parties are automatically called and bridged for a quick conference call. Alice is able to resolve her crisis before the close of business that day.

With the exception of keyboard idle-time detection, presence today is almost entirely manual. Devices must be restarted when moved, users must select "In a meeting" from a drop-down menu, and so forth. For presence to reach its potential as a lubricant of communications, it must become more automatic.

Existing enterprise directory assets must be combined with information gleaned from an increasingly ubiquitous computing environment to provide a rich set of automatically updated presence attributes. At the same time, users must be empowered to express and install policies that determine what information about them is published to different classes of watchers. This is "*rich presence*".

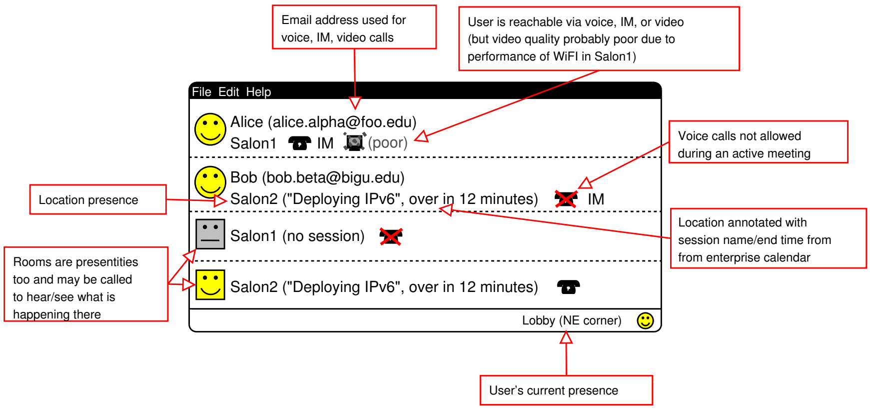

Figure 3: Rich Presence, Showing Automatic Location, Calendar, Etiquette, and "Room Weather" Presence

6.2 Rich Presence

The Internet2 Presence and Integrated Communications (PIC) [3] working group has produced a series of participatory trials at Internet2 meetings that explore rich presence and prototype future campus integrated communications service offerings.

Figure 3 shows an example contact list, showing the rich presence capabilities trialed at a January 2004 Internet2 technical workshop. Participants with WiFi-enabled laptops were invited to download and install one of several SIP-based integrated communications clients. These clients supported voice, instant messaging, and video communication to other participants, using the receiver's email address as a single, converged electronic identity (SIP.edu style). Participants were able to see not only which of their contacts were online or offline, but also for each, a current location, calendar-derived presence about what is currently happening at that location, and the expected quality of a call (based on location-specific network performance measurements). As participants used the meeting's wireless LAN infrastructure and moved from one meeting room to another, their locations were tracked by WiFi location technology. Using SIMPLE [14]—a SIP extension—a presence server served these presence attributes to clients, which applied user-directed policy control before notifying watchers of presence changes.

As a illustrative example of how rich presence and location services might bring real value to campus communications, consider the following: A freshman is in the library desperately trying to decipher an assignment from her English Literature professor. She consults her location assistant and requests it find all students within 100 feet of her current location filtered by the attributes "freshman AND EnglishLit101". She discovers a classmate two tables away who is in the same class. After a brief instant message exchange the two freshman retreat to a table where they can discuss the assignment.

7 Conclusion and Future Work

Connective middleware based on open standards is essential to enabling a large-scale exploration of the potential of peer-to-peer real-time communications. SIP-based call routing and presence services can provide much of the needed connectivity. Although peer-to-peer design principles are highly scalable and allow users to communicate in a manner that they negotiate peer-wise, the best new applications will leverage centrally-managed enterprise services. Hence, Internet2 has launched several new initiatives to develop and deploy connective middleware that is enterprise directory-enabled and SIP-based.

There is, however, a danger in having too much, unchecked connectivity—spam. The pioneers of electronic mail might be forgiven for not having foreseen the problems of abuse, but today's pioneers of real-time internet communications could not be. With real-time communications, spam is both more costly (it demands immediate human attention—an increasingly rare commodity) and is harder to filter (an automated filter can only act upon headers, not upon the message body—by then it is too late).

The SIP.edu group is beginning to look at how SIP-spam can be prevented or mitigated against. Techniques under consideration include: white lists, black lists, federated trust relationships among SIP.edu proxies, and caller charging, including either cash micro-payments or cryptographic challenges. The group hopes to deploy one or more of these methods in its testbed.

References

[1] A. M. Odlyzko. The history of communications and its implications for the Internet, `http://www.dtc.umn.edu/~odlyzko/doc/history.communications0.pdf`, June, 2000.

[2] Internet2 Voice over IP (VoIP) Working Group, `http://voip.internet2.edu/`.

[3] Internet2 Presence and Integrated (PIC) Working Group, `http://pic.internet2.edu/`.

[4] Integrated Postsecondary Education Data System (IPEDS), National Center for Education Statistics, `http://nces.ed.gov/ipeds/`, Spring 2001 Data (Fall Enrollment 2000-01).

[5] University Members List, Internet2 web site, `http://members.internet2.edu/university/universities.cfm`, January 1, 2004.

[6] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks, M. Handley, and E. Schooler, "SIP: Session Initiation Protocol," RFC 3261, Internet Engineering Task Force (IETF), June 2002.

[7] Internet2 SIP.edu Initiative, home page, `http://voip.internet2.edu/SIP.edu/`.

[8] Gulbrandsen, A., Vixie, P. and L. Esibov, "A DNS RR for Specifying the Location of Services (DNS SRV)," RFC 2782, Internet Engineering Task Force (IETF), February 2000.

[9] Mealling, M. and R. Daniel, "The Naming Authority Pointer (NAPTR) DNS Resource Record," RFC 2915, Internet Engineering Task Force (IETF), September 2000.

[10] SIP.edu Cookbook, Baron, D. and Teitelbaum, B. (editors), `http://mit.edu/sip/sip.edu/`.

[11] EduPerson Object Class Specification, Internet2 Middleware Architecture Committee Directory Working Group (MACE-Dir), `http://www.nmi-edit.org/eduPerson/internet2-mace-dir-eduPerson-200210.pdf`, October, 2002.

[12] Shibboleth Project Home Page, `http://shibboleth.internet2.edu/`.

[13] InCommon Federation Home Page, `http://incommon.internet2.edu/`.

[14] SIP for Instant Messaging and Presence Leveraging Extensions (simple), IETF working group, `http://www.ietf.org/html.charters/simple-charter.html`.