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

Revised: January 15, 2004

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 and call 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 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.

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 focus on content and relatively narrow applications (*e.g.* distance learning, remote medical consultation) [1].

Several recent Internet2 initiatives have brought members together to enable next-generation campus communications services through the development and deployment of application-agnostic connective middleware. 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.

Advanced VoIP

Networking technologists have long sought "convergence"—the coexistence of voice and data on a single digital network. Despite progress in the service provider core, convergence in the last-mile and in the enterprise has occurred at a glacial pace. Too much attention has been on matching the fidelity and features 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 "solved". 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 enable exploration of a large variety of applications.

[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).

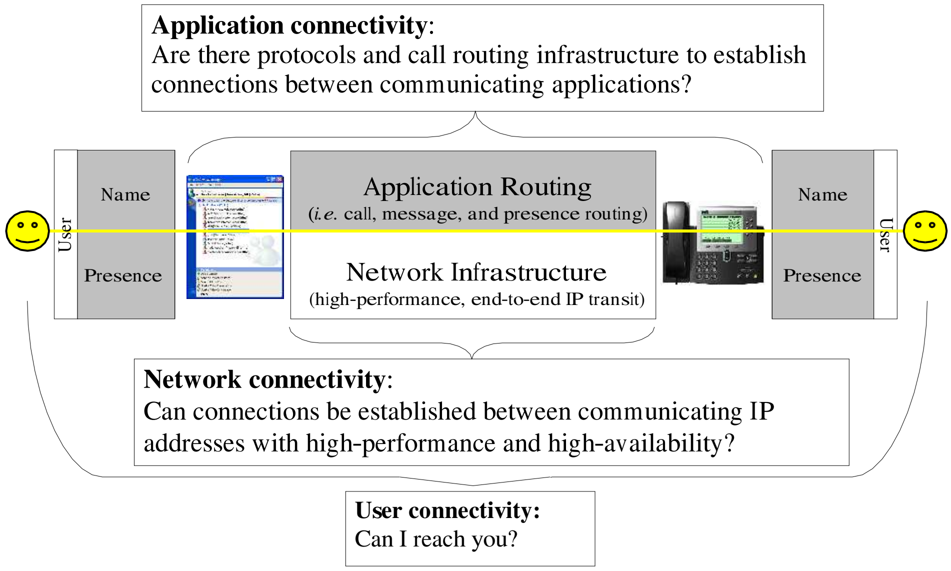

Figure 1: User-to-User Connectivity Requires New Middleware

Connectivity

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) and on meeting the bandwidth needs of researchers and educators. Often ignored and untapped has been the average user.

Internet2 member universities provide communications services to 3.8 million students [4, 5]. Eager to adapt to new modes of communication (*e.g.* cellular telephony, instant messaging), these users are already opting out of traditional campus offerings (*e.g.* residential telephony). We are working to engage this demographic in a large scale exploration of peer-to-peer (P2P) real-time communications. Such an exploration is possible, however, only if users can *connect* to each other.

Currently, Internet2 users enjoy superior IP connectivity. But this alone is insufficient. New middleware is required to route calls between communicating applications. Toward this end, Internet2 members are deploying SIP-based [6] call and presence routing infrastructure [7]. Because SIP is media-neutral, based upon open standards, 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.

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).

Presence in its full generality involves the exchange of user-level state informa-

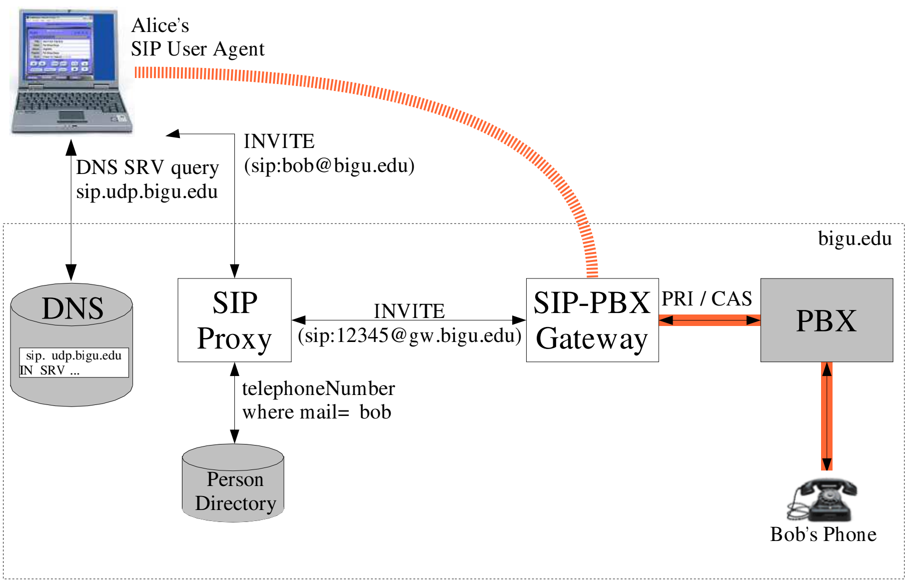

Figure 2: Initial SIP.edu Architecture

tion that facilitates communication at the appropriate time and manner. For example, presence may indicate availability (*e.g.* available, busy), location (*e.g.* in Conference Room B), or media capabilities (*e.g.* voice and video: not OK, instant messaging: OK).

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

SIP Connectivity (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 provide 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` dis-

covering 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;
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 SIP.edu, 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 phone(s) first, if any are registered, and, otherwise, ring his PBX extension. When UA registration is supported, calls are signalled and transported purely over IP and the door is wide open 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 interoperation with the PSTN.

Enterprise Presence and Location Services

The best new communications applications will combine peer-to-peer design principles with new, centrally-managed enterprise services. Today, with the exception of idle-time detection, changes to one's presence must be made manually. On the client side, users may install more sophisticated presence detectors and policy filters to determine what is published to different classes of watchers, but there are some presence services that are best provided as central services.

The Internet2 Presence and Integrated Communications (PIC) [3] working group is producing a series of participatory, social computing experiments at Internet2 meetings to explore "rich presence" services that leverage the campus LAN, physical infrastructure, and directory assets of the university to provide richer, automated user presence. This work is using the SIMPLE [14] extensions to SIP.

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

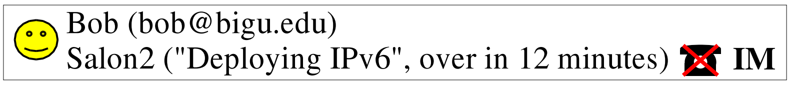

Figure 3: Rich Presence, Showing Automatic Location, Calendar, and Communications Etiquette Presence

Figure 3 shows an example buddy list entry, showing the rich presence capabilities showcased at the Fall 2003, Internet2 Member Meeting. Participants installed one of several integrated communications clients on their WiFi-enabled laptops, allowing them to initiate voice, video, and instant messaging calls to each other using SIP.edu-style addresses. Room location presence was continuously updated and cross-referenced with the meeting calendar show the name and duration of the current session in each room. Additionally, etiquette hints from the presence server indicated to clients that voice calls should not be placed to users in rooms with active sessions. The January 2004 version of this experiment will showcase automatic location presence (derived through triangulation of 802.11 signals) and room-level "internet weather" presence that will indicate to users *a priori* the voice quality that would be experienced on a call with each buddy.

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] Roster of Internet2 Universities, `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`.