Connective Middleware for Cyberinfrastructure-Mediated Communications

Position paper for NSF SBE/CISE Workshop on Cyberinfrastructure and the Social Sciences

Ben Teitelbaum < *ben@internet2.edu* >
http://people.internet2.edu/~ben/

February 15, 2005

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. Government and business elites have historically been slow to recognize the value of these general-purpose person-to-person communications, tending instead to emphasize content (*e.g.* databases, video-on-demand) and narrowly-scoped collaborative applications (*e.g.* discipline-specific co-laboratories) [1].

Historically, most of NSF's cyberinfrastructure (CI) investment has been directed at building the raw computational, networking, and middleware infrastructure to support computation-intensive simulations, large-scale storage, and networked access to scientific instruments and data. These investments have focused on enabling communication sessions between a human and a machine (or set of machines) or on enabling communication solely among machines.

The "doing of science", however, is still very much a human activity, requiring the practitioner to collaborate and communicate with a rich web of local and remote colleagues. Although there has been funding for CI-mediated, human-to-human communications projects over the years—largely by DOE (*e.g.* MBONE, AccessGrid, VRVS)—most electronic communication among scientists is still through the general-purpose communications services provided by the university (specifically, email and the telephone). As individual scientists attempt to improve their personal productivity, they are increasingly turning to consumer-oriented commercial services, such as cellular telephony and "walled garden" instant messaging and presence services (*e.g.* Yahoo! IM, AIM, etc.).

In an environment of constrained budgets, Carr-ist cynicism [2] about the strategic value of information technology, and IT departments that are beleaguered by security problems and copyright enforcement mandates, it is difficult for campus IT planners to innovate. Many CIOs are happy to have their users turn to commercial services to improve their personal productivity (one less campus service to support!).

Ironically, technology developments and an unprecedented industry convergence around open-standards for real-time communication[1] have created unprecedented opportunities for campuses to offer advanced communications services. These new services are build on top of media-neutral *connective middleware* that supports the creation, modification, and termination of peer-to-peer real-time communication ses-

[1] I am referring here to the IETF's Session Initiation Protocol (SIP) [3] and related standards.

sions that can integrate voice, video, instant messaging, and other media in a single session.

Superior networking connectivity alone is not sufficient to connect users with each other. The first task of connective middleware is to map human-readable addresses (*e.g.* `+1-234-567-8900` or `bob@bigu.edu`) to application-layer addresses (*e.g.* protocols, ports, and IP-addresses) on hosts or appliances where the user is reachable. Internet2's SIP.edu initiative [4] is working to deploy this basic connective middleware and has already provided SIP-connectivity to over 150,000 users at 10 institutions.

Presence takes connective middleware to another level, enabling the exchange of user-level state information to facilitate the initiation of communication at the appropriate time and in the appropriate manner. 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; 1000 Oakbrook Drive, Suite 300; Ann Arbor, MI), activity (*e.g.* in budget meeting, over at 2:15pm EST), media capabilities (*e.g.* voice and video: not OK, instant messaging: OK), or even mood (*e.g.* pensive, bored, frustrated).

If the National Science Foundation is to make a large impact on CI-mediated communication, it must invest in *campus connective middleware* to support general-purpose, human-to-human communications. For advanced personal communications to reach their full potential, however, a number of socio-technical research problems must be addressed:

- Architectures and policies for damping unwanted, interrupting communications (*i.e.* spam)
- How can users effectively manage *ad hoc*, peer-to-peer trust fabrics[2]?
- What is the relationship between *ad hoc*, peer-to-peer trust and federated trust?
- What abstractions and metaphors can help users understand and manage privacy policies?
- What are the best user interfaces for managing privacy policies and what should the defaults be?

[2]These are needed for control of fine-grained, highly-personal presence information.

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] Nicholas G. Carr, *IT Doesn't Matter,* Harvard Business Review, May 1, 2003.

[3] 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.

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