Companies mentioned:
Alcatel (www.alcatel.com)
AOL (www.aol.com)
Avaya (www.avaya.com)
Blackberry (www.blackberry.com)
Broadsoft (www.broadsoft.com)
Cisco (www.cisco.com)
IBM Lotus (www.lotus.com)
Microsoft (www.microsoft.com)
Nortel (www.nortel.com)
Siemens (www.icn.siemens.com)
Yahoo (www.yahoo.com)

Client Convergence: Unifying Real-Time Communications

Business Communications Review, December 2004

Irwin Lazar

Disparate real-time communications will rapidly integrate into a unified, presence-enabled infrastructure capable of streamlining enterprise communications methods and processes.

Enterprises are rapidly replacing legacy communications with IP-enabled applications that leverage a common network infrastructure for connectivity. IP-based phone systems are replacing traditional circuit-switched PBXs. IP-based videoconferencing systems, both room-based and desktop-based, are replacing traditional systems based on ISDN. Instant messaging, both public services and enterprise-provided applications, have embedded themselves into the corporate communications landscape

In parallel, Blackberrys, PDAs, smart phones, and/or cell phones have become standard issue in the enterprise, either as stand-alone devices or devices running softphone applications enabling remote access to enterprise calling systems.

While there may be some commonality among some systems (such as shared address books), in most cases, systems and applications exist and operate in isolated silos, often with no integration or communications between either the applications or the departments responsible for their administration. Individuals wishing to communicate must play the role of “Communications Detective” to determine what device or system to use to communicate with another person. The net result of the multitude of new ways to communicate is that enterprise communications have become more complicated than simplified.

Converging Communications
Given this current reality, it should come as no surprise that vendors and service providers see an opportunity in providing a means to converge disparate corporate communications systems and applications into a single, unified infrastructure that is capable of refocusing communications away from devices and systems and instead focusing on communications among people.

Client convergence requires all individuals to have a common application that they can use to interface with all communications systems. Solutions must be able to support multiple forms of simultaneous real-time communications across a variety of devices and/or applications. Individuals must be able to share and obtain presence information, either by manual preference selection or via the converged client’s interaction with existing scheduling and calendaring applications. Systems should also be extensible, supporting interaction with other corporate applications as well as customization of the user interface to allow seamless integration with existing corporate portals.

Presence: Can You Find Me Now?
A key requirement to support converged communications is “presence,” which is simply the ability to convey real-time information about a user’s current activity and/or availability. This information can include not only a user’s current physical location, but also the forms of communication that they are able to use, such as mobile phone, instant messaging, audio phone or desktop videoconference. Presence information may refer to a general location (e.g. “office”) or a specific location, e.g., relayed by a mobile phone with GPS (global positioning satellite) capabilities.

While presence capabilities are commonly thought of simply as allowing one user to see the current status of another user, the real advantage of presence-based applications is that they serve to virtualize user devices such that individuals no longer need to play the previously described role of “communications detective” and can instead focus on communicating directly with each other, allowing the presence system to be used to direct communications sessions to available and or preferred devices. For example, when a user activates his or her cell phone when stepping off an airplane, the phone automatically would communicate with a presence application, notifying the application that cell phone communications are now possible. This information would be available to anyone trying to reach the user, or could result in all incoming calls being directed to the cell phone per user-defined preferences.

Presence is not a new concept; most instant message systems have had presence capabilities since their inception. For example, systems such as AOL Instant Messenger, Yahoo Messenger, and Microsoft Messenger keep track of a user's availability by monitoring keyboard activity. A long gap in user input triggers the IM client to notify those making session requests that the user is away from the keyboard. A return to the keyboard may be configured to automatically change the user's status back to "available". Users may also manually set their current location information

Advanced presence features allow integration among a variety of systems. For example, users of Microsoft Exchange or Lotus Notes can tie their calendars into presence systems, allowing individuals to see when other users are in a meeting and for how long. This integration allows for far more intelligent status messages rather than the ubiquitous "I'm away from my desk right now" messages that any user of IM applications has undoubtedly experienced.

Client Convergence Applications
Current solutions for delivering client convergence rely on middleware applications that interface SIP-based voice, video and IM applications with other real-time communications, including traditional circuit-switched PBXs with SIP network interfaces. By leveraging SIP and SIMPLE to tie together multiple independent communications systems, any system that supports SIP and SIMPLE can become an element in a converged communications system. This means that even legacy circuit-switched PBXs can become part of a unified communications system, provided they have the capability to support SIP interfaces. (For more, see "SIP As The Glue.") Cell phones may also be integrated via gateways, allowing the forwarding of calls based on the user's own predefined preferences.

Middleware applications rely on SIP trunking to tie together multiple communications systems, enabling sharing of presence information between independent applications and systems via a Web-based interface. This approach enables a user to control the behavior of all forms of real-time communication available to them at any given point in time and from any location.

Typically, a Web-based client application allows users to set communications preferences, determine availability of co-workers or groups, and initiate various forms of communications regardless of system and device. The Web based client is capable of acting as a SIP-based softphone for both voice and video, as well as serving as the interface to underlying IM capabilities.

Middleware applications typically allow integration with other collaborative applications, such as those that enable document sharing, application sharing or shared whiteboarding. Many systems also support APIs or Web services protocols such as SOAP (Simple Object Access Protocol) to enable integration with existing applications such as corporate portals or custom-defined communications applications.

Emerging Trends

The approach of tying together existing and new communications systems into a common framework presents enterprises with the ability to build an open-standards based communications infrastructure that leverages best of breed components into a unified centrally accessible infrastructure. This opportunity threatens to break the lock that single-vendor voice, IM and video solutions have on the corporate desktop and communications environment (see "VOIP Doesn't Matter," *BCR*, September 2004, pp. 30-33).

Thus while previously described stand-alone applications can enable client convergence now, vendors of IP-PBX products, workgroup collaboration applications and instant messaging applications are rapidly trying to expand their products to enable them to serve as the hub for all real-time communications.

For IM and collaborative application vendors, enabling communication via voice, and supporting real-time applications such as pre-scheduled and on-demand conferencing, as well as integrating presence capabilities, allows vendors of applications and systems to directly challenge the traditional VOIP vendors on their own turf. While the shortcoming in voice-enabled IM and collaborative applications has been their inability to interface with either the PSTN or existing telecommunications systems, these concerns are being addressed as IM vendors increasingly look to partner with IP-PBX platform vendors.

For example, Microsoft in its recently announced “Istanbul” client will support SIP-based voice and video communications using its Live Communications Server platform. IBM Lotus Workplace Team Collaboration will likely add the ability to launch voice and video sessions from within its workspace. Microsoft’s move to add real-time collaborative capabilities to Office, including support for voice, video and presence in LCS 2005 will position LCS as a competitor to enterprise VOIP systems from vendors such as Cisco, Avaya, Nortel, Siemens and Alcatel. Microsoft’s plans to integrate LCS with SIP-based VOIP solutions from partners such as Broadsoft will enable it to support access to PBX-like features, as well as to PSTN gateways from within Microsoft Messenger.

VOIP vendors are responding by either delivering stand-alone products such as Siemens HiPath OpenScape, Nortel’s Multimedia Communications Server (MCS) 5100, and Avaya’s Converged Communications Server, or by embedding support for SIP/SIMPLE instant messaging and presence directly into the IP-PBX platforms. Examples include Alcatel’s OmniTouch Unified Communications and 3Com IP Conferencing. Each of these products leverages existing full-featured telecommunications systems that currently exist within the organization, coupling them with support for SIP-based communications such as voice, videoconferencing, presence and collaboration, as well as enterprise-class instant messaging.

Thus, the battle for control of the corporate communications environment is joined, with application vendors seeing an opportunity to move into voice, and voice vendors moving into enabling support for more real-time applications—and, alternatively, vendors offering middleware applications capable of interfacing any vendor’s products so long as it supports open SIP and SIMPLE standards.

For the foreseeable future, systems will likely exist in parallel, with some level of integration between VOIP solutions from major telecom systems providers and real-time voice capabilities that are directly available within collaborative and IM applications. Over time these applications will converge and telecom systems vendors will face an increasingly difficult value proposition in delivering standalone voice solutions. Rather, telecom systems providers will increasingly operate at the back end, providing call control systems and servers that may be leveraged by SIP-compliant applications to enable call policy enforcement as well as interconnection with the PSTN and legacy systems.

Organizational Preparations

For most enterprises, functionally separate groups manage voice, collaboration and calendaring systems as independent applications. As the feature sets, and in many cases the products themselves, begin to converge, enterprises will need to rethink how they manage systems, and more importantly, how management and planning of systems can be converged to deliver a seamless, integrated product to the workforce.

Therefore, enterprises should create a collaborative communications team including members from each of the aforementioned application teams, as well as resources from network planning and engineering.

The role of these teams should be to evaluate the organization's long-term trends toward collaborative and presence-based communications, enabling the development of a clear and consistent road map and architecture encompassing all affected areas and groups. Collaborative communications teams can jointly determine the business case for systems deployment as well as the method by which systems will be deployed.

Collaborative communications teams should be responsible for analyzing various trends among current providers of communications products and services, as well as additional vendors and service providers, to enable the organization to make informed decisions about how its long-term communications architecture will be defined, what products, services and technologies it will include, and how it will be implemented.

Enterprises should act quickly to establish proactive plans for integration of communications systems. As we have seen in the past with IM, PDAs, and wireless LANs, failure to proactively deliver such services means that enterprises--whether they want to or not--will have to deal with public services that are brought in by their users.

Conclusions

Disjointed forms of personal communications will rapidly converge into a unified application combining voice, instant messaging, video, collaboration and presence. The result will improve organizational efficiency, allowing individuals or groups to communicate directly with each other through a common system, regardless of device or application.

Enterprises considering deployment of converged communication must bring together disparate teams responsible for separate individual applications into a unified workgroup, capable of analyzing the best approach for enterprise communications from myriad choices, in conjunction with engineering to support network requirements for real-time any-to-any communications regardless of location, application or device.

Irwin Lazar is a senior analyst for Burton Group covering VOIP, IP-telephony, and real-time communications. He is also a Certified Information Systems Security Professional (CISSP) and is conference director for MPLScon. He can be reached at ilazar@burtongroup.com.

SIDEBAR: SIP As The Glue

Integration of real-time communications applications relies on a common and consistent interface to enable application integration. This role is being played by the Session Initiation Protocol (SIP), as defined in RFC 2543. SIP is an application-layer protocol, meaning that it resides above network protocols within the Open Systems Interconnection (OSI) model. SIP thus relies on an underlying IP-based network to provide reachability and assured packet forwarding and delivery.

SIP is simply a signaling protocol designed to provide a common framework for session establishment, modification, and termination. SIP is designed to support multimedia communications including voice and video, either in a one-to-one or one-to-many operation.

In converged communications, SIP acts as a trunking protocol, enabling the direct interconnection of independent systems, with the only requirement being a SIP network interface. Thus, a digital PBX with an IP line card supporting SIP trunking could directly connect to an IP-PBX system also supporting a SIP interface. In this configuration, the IP-PBX and the digital PBX act as gateway devices, allowing calls to be placed across systems regardless of endpoint protocols or devices.

Extending SIP the SIMPLE Way
Implementation of presence capabilities among disparate applications requires a framework that allows applications to exchange presence

information in a consistent and standard manner. The approach gathering the widest support within the real-time communications application industry is to extend SIP to support the ability to carry presence-related information.

The IETF's SIMPLE (SIP for Instant Messaging and Presence Leveraging Extensions) working group is leading the effort to define a framework for open standards for presence and instant messaging-based communications. SIMPLE is not a single protocol; rather it defines a set of protocols and requirements that extend SIP to support the requirements of both real-time instant messaging as well as presence.

SIMPLE does not limit presence information to a specific user, but rather relies on a concept known as a "presentity" which defines a single person or entity. A presentity may have multiple "presence user agents (PUA)" which are akin to devices such as cell phones or softphones. Using this model, a "presence agent (PA)" may keep track of multiple PUAs for a single presentity to allow various means of communicating presence information for specific devices for each user.

SIMPLE standards leverage the "rendezvous" feature within SIP to allow mobility through the use of dynamic SIP contact addresses. Subscribers wishing to obtain presence information about other subscribers do not need to know the actual locations of a user's presence agents, instead they simply send a SIP "SUBSCRIBE" message to the user's Address of Record (AOR) (obtained via DNS lookup) to establish a session. The AOR routes the request to the correct PA. AOR names are resolved via standard lookups of SIP URIs, so in effect a user's SIP URI will translate to the IP address of his or her presence agent.