Volume 10 Issue 01 Published, February 15, 2006 ISSN 1535-864X DOI: 10.1535/itj.1001.02

Intel® Technology Journal

Converged Communications

Session Initiated Protocol (SIP) Evolution in Converged Communications

More information, including current and past issues of Intel Technology Journal, can be found at:
http://developer.intel.com/technology/itj/index.htm

Session Initiation Protocol (SIP) Evolution in Converged Communications

Mohammad Kolbehdari, Digital Enterprise Group, Intel Corporation
Dave Lizotte, Information Technology, Intel Corporation
Glen Shires, Digital Enterprise Group, Intel Corporation
Scott Trevor, Information Technology, Intel Corporation

Index words: VoIP, Presence, TDM, PSTN, IP PBX, SIP, SIMPLE, RTP, Converged Communications, Interoperability, Security, Enterprise Networking, Mobility, IETF

ABSTRACT

With the telephony business world focusing on Voice over IP (VoIP), numerous IP Private Branch Exchange (PBX) vendors are offering rich new product lines supporting VoIP capabilities, presence integration, and other enhanced multimedia and location-based services. In order to achieve true convergence of these technologies, PBX, phone and network vendors will need to design and support products that will interoperate with each other based upon industry standards. Session Initiation Protocol (SIP) is being widely adopted in the industry [1] as a signaling protocol because it enables data and voice convergence for devices and applications across a wide range of industry sectors. SIP enables voice, video, Instant Messaging (IM) and other media, and facilitates presence and location-based services. SIP's extensibility and versatility enable rapid innovation of new, rich features, and rapid deployment, and it has become a market enabler for VoIP PBX's and IP telephony devices and applications. However, SIP's extensibility has also introduced interoperability challenges as vendors differentiate by extending beyond baseline SIP specifications. Implementation of rich features in a standard, interoperable manner requires ongoing standardization and industry consensus.

INTRODUCTION

Converged communication improves person-to-person business communication with the unification of numerous messaging modalities including voice, video, IM and Presence, and integration with data applications and collaboration tools.

SIP is emerging as a key enabler of converged communication offering an interoperable protocol for such requirements and flexibility for many services' and vendors' networks. SIP differs from other communication protocols by its strong industry support, multi-vendor integration at the application layer, modularity, and common standards. Since SIP is an application-layer protocol, it is transparent of the underlying data link layer topology, which simplifies its deployment.

In this paper we focus on the emerging capabilities of SIP within real-time communication technologies while addressing challenges within the areas of interoperability, security, and enterprise network integration. We describe how the seamless integration of presence-based SIP, VoIP, Mobile IP, SIP mobility support, unified communications, and applications can lead to converged communications.

SIP AND CONVERGED COMMUNICATION

Most legacy telephony devices are dumb endpoints with no memory or processors built into them, as all of the intelligence was located within the proprietary PBX processor software. In contrast, with VoIP PBX solutions, all endpoints have processors and memory thus creating superior intelligence out at the network edge.

SIP is a signaling protocol for the establishment of communication sessions between these smart endpoints. It is commonly used to initiate voice, video, and IM sessions and can also be used to convey presence, location, and other information. SIP has emerged as a key protocol with strong industry support for the deployment of IP-based telephony. In addition to the rich media session and information that it can convey, SIP offers these additional benefits:

- Converged Network: Using a single network for voice and data reduces cost and simplifies management.

- Mobility: SIP provides a user with a logical identity regardless of the device type he is currently using or the device's physical location. This allows users to roam and to switch between devices (such as from a handheld to a computer SIP phone), while remaining reachable through a single address: callers do not need to try numerous phone numbers.
- Enhanced Audio Quality: The traditional public telephone network (PSTN) only transmits a small portion (200 Hz to 3.4 kHz) of the full range of human speech (80 Hz to 10 kHz). Its limited frequency response and dynamic range is the reason people "sound different" over the phone, and must resort to saying "S" as in "Sam," "F" as in "Frank." In contrast, VoIP phones can use a wideband codec to capture and reproduce audio with a much wider frequency response (for example, 50 Hz to 7 kHz) and dynamic range, resulting in a dramatically clearer call with minimal distortion of speech.
- Integrated Presence: A communications session is commonly initiated after identifying the user's availability or willingness to communicate. The publication of the user's presence information can also be used to determine the appropriate type of session to initiate (for example: while in a meeting, the user may prefer to receive an IM, whereas while driving the user may prefer voice).

Interoperability and Standardization

The full potential of VoIP can only be realized if calls are connected over an IP network end-to-end, rather than relying on gateways through the PSTN to VoIP networks. Not only do these gateways add cost and latency, but they block the rich features that VoIP can provide, features that include enhanced audio quality, video, IM, presence, and application sharing. These features require end-to-end IP connections, commonly referred to as direct-IP peering. Interoperability is required for direct-IP peering.

The rapid convergence to SIP is a strong step towards interoperability, but the SIP specification alone is not sufficient. SIP's advantages include simplicity and extensibility. However, its extensibility has resulted in implementations that extend SIP in incompatible and non-interoperable ways. SIP and many related protocols have been developed through the Internet Engineering Task Force (IETF) standards organization.

The primary SIP specification, RFC 3261, specifies how sessions are created, modified, and terminated, and it defines its use with registration and proxy servers. However, many additional IETF specifications are commonly used in conjunction with RFC 3261; for example, additional specifications are used to define how session capabilities are formatted and negotiated, how firewalls and NATs are traversed, and to clearly define certain transition states and call features. Additional specifications are used to define the codecs for various media types and the transport they flow across, typically RFC 3550 Real-time Transport Protocol (RTP). Other specifications are used to define presence and IM features. Numerous other specifications are used to define security, identity, and authorization; with additional specifications defining the underlying transport protocols.

Development of these specifications has been, and will continue to be, done through IETF working groups including the "SIP" working group to develop the primary protocols, the "SIPPING" working group to determine and document new requirements, and the "SIP for IM and Presence Leveraging Extensions" (SIMPLE) working group to define IM and presence applications. As the protocols continue to evolve, they are in various stages of becoming a standard. Some are full Internet Standards; others are in the stable but not yet at the ratified "Request For Comments" (RFC) stage, while others are rapidly evolving Internet Drafts. Specifications developed by other organizations are also commonly used, including some codec specifications developed by International Telecommunication Union (ITU) [5] and some transport protocol specifications developed by the Institute of Electrical and Electronics Engineers (IEEE).

Clearly, consensus throughout the industry is needed for consistent implementation and a rich level of interoperability to be achieved. The greatest roadblock or challenge for an enterprise attempting to provide converged communications seems to occur within the integration of different vendor products into a seamless solution. The lack of interoperability of vendor products can cause a project budget to increase to address the integration efforts. SIP functions as a signaling protocol, but does not support the enhanced capabilities vendors are providing within their products.

Despite lack of ratification, many extensions have been widely implemented in numerous products. These extensions can however compromise the interoperability between vendor's solutions. When it comes to integrating rich presence and other key value-added features between vendor products, it has become clear that SIP interoperability is not enough. Open source implementations of SIP stacks, such as reSIProcate, have helped achieve interoperability among various implementations, but since the stacks are not a complete implementation of a SIP product, they alone are not enough to ensure rich interoperability.

Some industry organizations [7] are beginning to address these concerns. For example, the 3rd Generation Partnership Project (3GPP) has created the IP Multimedia

Subsystem (IMS) specification which is a standardized implementation based on SIP. The GSM Association (GSMA) has a series of interoperability trials based on SIP and IMS. The SIP Forum has begun work on specifying best industry practices for SIP and also coordinates SIP Interoperability test events (SIPit). Any vendor with an implementation of SIP is encouraged to attend these bi-annual events where implementations are tested against each other and problems are often resolved on-site. Although specific vendor results are not made public, SIPit has succeeded in eliminating many of the hurdles of interoperability.

Development and Deployment Benefits of SIP

The SIP specification was originally published in 1999 as IETF RFC-2543 [2] and updated in 2002 as RFC-3261 [3]. Whereas signaling protocols such as H.323 utilize a single administrative domain architecture, SIP can be peer to peer and across domains. By providing an effective communication method between peers, SIP enables innovation of client features without requiring deployment of additional network infrastructure.

SIP utilizes a standard call control mechanism to set up, manage, and tear down a communication session. It is the first protocol that can run over reliable and unreliable transport protocols, has request routing capabilities for performance and control, and is extensible. Carried within the SIP message body, SIP utilizes the Session Description Protocol (SDP) to describe the session parameters such as call attributes, Real-time Transport Protocol (RTP) and payload format, and User Datagram Protocol (UDP) port selection, while also negotiating and exchanging media capabilities such as audio codec selection, video, or shared applications. A typical session initiated by SIP is a packet stream of the RTP, a standardized packet format for delivering audio and video over IP. SIP defines basic transactions and is extensible, scalable, and allows for supplementary information to be carried within the payload allowing devices to make intelligent call-handling decisions and invoke other application-level services such as IM and Presence. SIP is the first protocol to enable multi-user sessions regardless of the media content.

SIP is similar to the Hyper Text Transport Protocol (HTTP) in the way that messages are constructed. This allows developers to easily create SIP applications using common programming languages and Web services, to re-use code and tools, and to more easily debug applications.

SIP Re-uses Existing Internet Features

Prior to SIP, a typical VoIP PBX used signaling protocols such as H.323, H.245, and H.225 for the call setup, control, and teardown of a voice call. With IP telephony, the H.323 protocol suite had to be revised, as the absence of a standard for VoIP resulted in incompatible products. Only a portion of the H.323 architecture is used for VoIP when it comes to audio calls. One of the drawbacks of H.323 is that it will first establish the session then negotiate the capabilities and features for that session. The H.323 protocol provides only a numbering scheme for identities or addresses, thus does not provide the scalability and flexibility of the more versatile URI-based addressing. SIP's URI-based addressing allows callers to use either the URI names (which might be the same as the recipient's e-mail address) or mapping to a numeric dialing plan.

With the emergence of converged communications, the SIP protocol offers that attractiveness of re-using existing Internet features for real-time, mobile, and seamless collaboration. SIP uses a large selection of protocols that are already being utilized by applications for the Web, Internet, and IP-based networks. IP networks route differently than traditional PSTN telephony networks. The basics of IP routing is to route a packet to a desired destination or intermediate point that can make further routing decisions based upon the final destination's IP address. Because a user typically does not know the IP address of the end user they are trying to communicate with, the use of a Domain Name System (DNS) is utilized. Utilizing VoIP telephony with SIP, a user's identity is defined by a Uniform Resource Identifier (URI) based upon his/her IP address, username or phone number, and host name, and not on a distinct telephone number tied to only one location, as is done in traditional telephone systems. SIP uses DNS procedures to resolve a SIP URI and locate the appropriate SIP registrar for the call recipient. Based upon these services end users can utilize one identity name and become reachable anywhere on the network upon which they reside.

When a user wants to place a call, a SIP invite is initiated. The SIP communication session will determine the end device to be contacted and the user's availability and willingness to communicate. The session will also negotiate the media and other capabilities, such as the audio codec to be utilized and may even renegotiate any additional features or capabilities needed during the session. SIP also provides the ability for either endpoint (or an intermediate proxy) to tear down and terminate the call.

SIP is utilized in various architectural components including User Agents (UA endpoints), registrars, proxy, and redirect servers. A SIP registrar allows for all SIP user agents to register and authenticate to the network as an active user capable of placing calls. The registrar also acts as a repository for SIP URL/URI's and other identity information. A SIP proxy server can perform application-level routing of SIP requests and SIP responses for the requested home location services. If a proxy cannot identify the request it will send it to a redirect server. The redirect server does not forward SIP requests but points the proxy to contact another server that might know where the requested INVITE or UA resides. When utilizing a non-SIP endpoint, such as a legacy digital or analog device, the call will utilize a SIP gateway to act like a UA to allow for the protocol translation to non-SIP networks such as H.323, Media Gateway Control Protocol (MGCP), and Public Switched Telephone Networks (PSTNs).

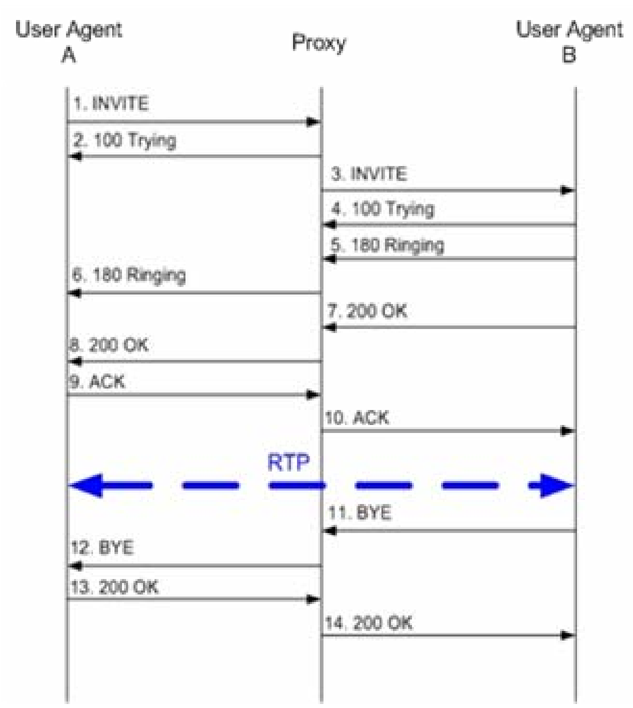

Figure 1: SIP call between user agents

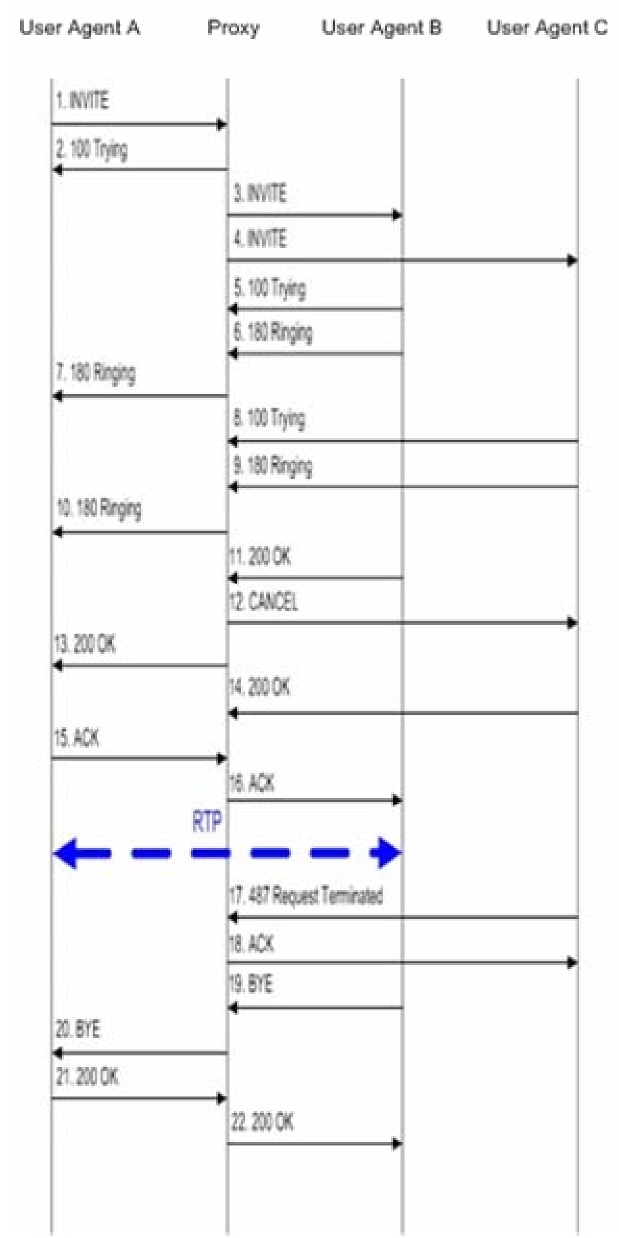

Figure 2: Multiple user agents and forking

Mobility

Mobility allows devices to stay connected even as the device moves between networks. Two methods that implement VoIP mobility are Mobile IP and SIP Mobility. Mobile IP is a network layer approach: it operates below the application layer and is applicable to most applications including VoIP applications based on H.323 and SIP, but it can add latency due to tunneling of the data stream. SIP mobility uses application layers (3 or 4) and augments existing VoIPs such as SIP or H.323 [5]. Being an application layer protocol enables SIP mobility to be deployed easily without requiring the network infrastructure support that Mobile IP requires. In addition, SIP mobility enables not just device mobility, but also personal mobility, allowing a user to easily switch between different SIP devices.

Mobile IP allows applications to use a given IP address and stay connected to devices regardless of their locations. When users with mobile devices leave the network that their device is associated with and call their home network and enter the domain of a foreign network, the foreign network uses the Mobile IP (IETF, RFC 3344) [4] to inform the home network of a Care-of-Address (CoA) to which all packets for the user's device should be sent. The Mobile IP network topology is shown in Figure 3.

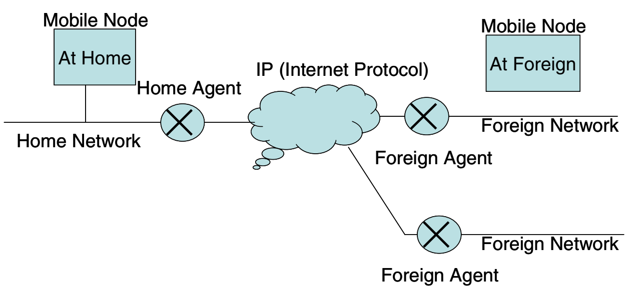

Figure 3: Typical Mobile IP topology

IP mobility is supported by transparently binding the home address of the mobile node with its CoA which is the termination point of the tunnel toward the mobile node when it is not in the home network. IP mobility binding is maintained by specialized routers known as mobile agents. Mobile IP has the following three main components: mobile node, home agent (HA), and foreign agent (FA). The three main phases of the Mobile IP process are as follows:

- Agent discovery–a mobile node discovers its FA and HA during agent discovery.
- Registration–the mobile node registers its current location with the FA and HA during registration.
- Tunneling process–a reciprocal tunnel is set up by the HA to the CoA to route packets to the mobile node as it roams.

The CoA is a temporary address that is valid while the mobile node is attached to the foreign network domain. In Mobile IP Wireless Local Area Networks (WLANs), the two mobility agents, the HA and the FA, coordinate, update, and authorize the connections and CoAs for clients from foreign networks. These connections are provided by binding the update message sent by the HA to the corresponding node. The bound message allows VoIP traffic and messages to be directly tunneled between the caller node and the mobile node. There are two key issues that can arise: the first is when roaming occurs between two foreign networks while a call is in progress between a caller node and a mobile node, and the second one is the timing requirement that must be predefined for WLANs or handset designers. The timing requirements are the period of time needed for a station to associate with an Access Point (AP), the period of time needed by a handset to associate with a foreign network, the period of time to bind to a foreign network and create a new CoA, the period of time needed to send packets directly between the Mobile Node and a Caller Node, and the period of time needed to bind update messages from an old foreign network to a new foreign network.

SIP mobility (IETF, RFC 3261) [3] supports mobility for VoIP applications by providing handoff capabilities at the application layer. The SIP mobility support protocol uses the concept of a Visited Registrar (VR) in the foreign networks. The SIP mobility support with VR features combines some of the functions of a SIP proxy server, location server, and user agent. The SIP proxy server enables SIP [6] to handle both firewall functions and Network Address Translation (NAT). SIP is designed to support roaming so that a user can be found independent of the device he/she is using and its network location. For example, with SIP, a call on a handheld phone can be transferred to a computer SIP phone. The SIP mobility approach at the micro-mobility implementation level is very similar to the concepts of foreign network and home network. In SIP mobility support, the FA of Mobile IPs is replaced by a SIP VR and foreign network. The Mobile IP HA is replaced by a combination of a SIP proxy server, a location server, and a user agent server. Advantages of SIP mobility support are of using the existing IP-based network without modification as well as being fully supported by the Windows* environment (Windows XP*) making possible a rapid deployment in the market place. Figure 4 shows VoIP and IP-based network configuration.

* Other names and brands may be claimed as the property of others.

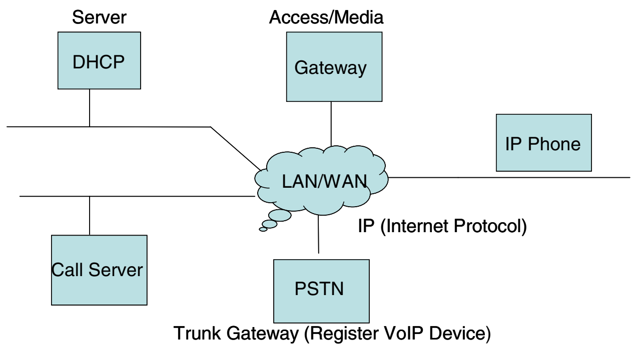

Figure 4: VoIP and IP-based LAN/WAN topology

VoIP Security

Because VoIP uses smart endpoints over a shared network, it provides many benefits and features that are not available with legacy phones over traditional PSTNs. With this flexible VoIP architecture comes additional concerns about security and privacy. However, proper system design can address these issues and in fact, make VoIP more protected, private, and manageable than the PSTN.

Smart endpoints provide a powerful tool to address security and privacy concerns. The traffic can be encrypted at the endpoints and throughout the network. Media traffic is commonly encrypted using Secure RTP (SRTP), and SIP signaling traffic is commonly encrypted using TLS (SSL) and S/MIME. All network traffic can be further encrypted using IPSEC Encapsulating Security Payload (ESP).

To store the protected encryption keys for these protocols, smart endpoints can also provide protected storage mechanisms such as the Trusted Platform Module (TPM) included in many PC platforms. Secure protocols are used for key distribution, such as the Multimedia Internet KEYing (MIKEY) and the Internet Security Association and Key Management Protocol (ISAKMP).

Smart endpoints can also provide identity authentication and attestation. Whereas traditionally identity was provided only through physical means (anyone who could physically access the phone or network could impersonate another caller), smart endpoints can provide an additional level of user authentication through the use of passwords, biometrics, or other means. Attestation to this identity can then be communicated to remote users.

CHALLENGES AND RESPONSES

VoIP in the Enterprise

VoIP technology has matured, and interoperability issues are being resolved. High-bandwidth wireless LANs and WANs have extended access to VoIP enabling more personalized and better integrated services at a reduced cost. However, for VoIP to fully deliver rich, reliable, cost-effective services, the major challenge will be to enhance the interoperability, security, QoS, and bandwidth management for wide-scale VoIP deployment. The desired goal of enterprise deployments is to deploy VoIP to that greater than 90% of enterprises by the next few years. The growing demands are an easily deployable, high-quality VoIP solution, interoperable across network and PBX environments, delivering a new level of performance to the VoIP marketplace, and integration with IP phone-based PBX technology. In some cases, PBX equipment is modular and supports VoIP interfaces such as Line Replacement Units (LRUs). The VoIP gateway is introduced to integrate VoIP into the existing enterprise TDM PBX environment. Enterprises are expected to gradually transition to VoIP because of the expense of legacy analog PBXs and phones and end-of-life of service contracts that support that equipment. There are evolutionary steps that the enterprise can take to move its global telephony environment to a full IP PBX network. A typical TDM PBX voice mail system that is currently adjunct to each distributed PBX can be migrated to the PSTN network or transitioned to an IP-based network by using SIP-based solutions. A VoIP PBX gateway and media server can be provided within the legacy PBX distributed environment to support both TDM and IP-based endpoints, thus providing a large investment protection during the transition period to a standardized IP-based telephony approach and design. New infrastructure sites can be planned to support full IP PBX solution designs with standardized call control capability. In addition, a standardized call control architecture can be developed to support the minimization of numerous legacy PBXs to a handful of centralized PBX media servers. This design method will be key, supporting zones or regions based upon the infrastructure architecture of the data network topology.

There are three critical performance issues that need to be focused on for VoIP deployment: latency–the end-to-end delay; jitter–the variable delays in each voice packet; and packet loss–the dropping of individual packets caused by network congestion. Specific values must be reached in order to ensure that the user has an experience that is the same or better than using TDM telephony. The latency, jitter, and packet loss of the VoIP system are issues regardless of the application of the network data technology. VoIP traffic is very sensitive to dropped packets, network latency, jitter, and packet loss. The key to success for reliable VoIP systems is to control those three major issues. Acceptable VoIP quality requires a latency or delay of not more than approximately 300 ms. Jitter causes irregularities in the flow and delivery of data, and although most vendors have successfully solved this issue using jitter buffers to smooth out the delivery of voice packets, excessive jitter can cause significant additional latency. While slight packet loss is typically not noticeable by users, significant packet loss results in moments of dropped audio and excessive packet loss can cause dropped calls. Queuing priorities solve many jitter and packet loss problems by ensuring timely delivery of voice packets using prioritized-packet method when the reprioritize packets are ordered on the IP network.

Network Capacity and WLAN

Seamless communication is increasingly demanded between data, voice, and video media. Whereas legacy communication technologies use a separate telephony network for voice and an IP network for data, a converged voice and data network provides richer features, integration, and multiple access options at a lower cost. However, a converged network also raises several deployment, configuration, and planning issues concerning QoS, call control, network capacity, provisioning, and architecture. Figure 4 illustrates the concept of IP-based integrated network communications. The network capacity is directly related to the throughput of the network infrastructure as well as the bandwidth requirements for typical voice, data, video, and media applications using IP-based packets. Codecs that offer voice compression help to derive as much capacity as possible by minimizing the packet size. However, to minimize latency, a new VoIP packet is typically sent every 20, 30, or 60 milliseconds, resulting in many small packets that the network must efficiently deliver. WLAN deployment has additional challenges including the variability in number and types of devices connected to an access point, and possibility of radio interference.

Infrastructure Integration

A phased approach can be used to deploy a converged communications environment in a large company. Evolutionary steps will protect current telephony investments while companies migrate to IP-based network solutions. In the interim there will most likely be a mix of older systems using proprietary protocols with newer ones based on SIP. Gateways can provide a smooth migration path for users, but may come at the cost of additional complexity and additional licensing costs. Such gateway services provide the first steps in integrating IM and presence into the telephony environment.

Client Application Integration

True convergence on the client is more than providing all the features in a single application interface; it also requires that they be ubiquitous and available across all the applications on the device; for example, the ability to quickly communicate with the authors of a document one is viewing. Users also demand the ability to quickly and seamlessly move between modes of communication. For example, a user may start out using an IM with a colleague, but then determine that a voice call would serve her better. Later, she might add a data collaboration session to the mix in order to review a document being discussed. Standards-based protocols such as SIP provide the foundation to achieve this level of integration on the client.

Influencing the Vendors

As mentioned earlier in this paper, the desire for vendors to differentiate their products on the basis of features, combined with the relatively immature state of standards makes it difficult to achieve the vision of seamless convergence of real-time collaboration capabilities. Progress is being made with the adoption of SIP by many vendors. More is needed however in several areas including interoperability between wideband audio codecs, rich presence, and direct-peering between SIP networks. Customers have perhaps some of the greatest power to influence vendors. It's important to push vendors to interoperate. Don't accept the status quo of "convergence" that is only converged into a single client interface. Look for and encourage vendor solutions that allow for interoperability both on the infrastructure and at the client/device level.

CONCLUSION

With the evolution of SIP as the standard signaling protocol for VoIP telephony, numerous application-level features and capabilities are being developed to advance mobility and productivity for businesses and their end users. Interoperability between various vendor solutions is key to enabling end users to richly communicate through the device type of their choice, regardless of their global location.

SIP is the standard for the establishment of multimedia sessions, including voice, video, and IM; and for conveying presence, location, and other information. SIP-based communications deliver a suite of solutions that can significantly enhance users' communication options and productivity.

ACKNOWLEDGMENTS

We are grateful for the encouragement and help that we received from our managers, particularly Mark Cooke, Alan Blair, Alex Ahmadian, Mike Powell, Michael Stanford, Jeff Morriss, and Bala Cadambi. Additional thanks go out to individuals who reviewed the paper and provided valuable feedback: Sanjay Rungta, Greg Trusley, Don Meyers, and Duncan Glendinning.

REFERENCES

[1] Regis J. Bates, *Voice-over IP (VoIP)*, McGraw-Hill, Jan. 2002.

[2] Handley, M., Schulzrinne, H., Schooler, E., and Rosenberg, J., "SIP: Session Initiation Protocol," *IETF RFC-2543*, March 1999.

[3] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and Schooler, E., "SIP: Session Initiation Protocol," *IETF RFC-3261*, June 2002.

[4] Perkins, C., "IP Mobility Support for IPv4," *IETF RFC-3344*, 2002.

[5] ITU_T "SG 16 H.323 version 1–Version 4 Standard."

[6] Avaya Technologies, White Paper "Enterprising with SIP."

[7] www.cs.columbia.edu/sip/*

AUTHORS' BIOGRAPHIES

Mohammad Kolbehdari is a senior staff architecture engineer in the Digital Enterprise Group. He received Ph.D., M.Sc., and B.Sc. degrees in Electrical Engineering from Temple University, Philadelphia, Pennsylvania, University of Mississippi, Oxford MS, and Communication University Tehran, Iran, in 1994, 1991, and 1986, respectively. He conducted research in the EM lab at Temple University prior to joining the Department of Electronics, Carleton University, in Ottawa, Canada as a research associate faculty member from 1994 to 1998. He joined Intel in 1998 and has been involved in planning and architecture development teams including those of the Pentium® 4 processor FSB and PCI Express* designs. His areas of interest are EM modeling, Power Delivery Networks, Signal Integrity, Circuit and I/O Serial Differential Design, System-level Architecture, Platform Initiatives, Clock Jitter Characterization and Measurements, and Communications and Networking Architecture. His e-mail is mohammad.kolbehdari at intel.com.

Dave Lizotte is a senior voice engineer at Intel Corporation in Folsom, California. He is currently focused on new capabilities development for Voice over IP and convergence solutions. He has in-depth experience in design and implementation of solutions incorporating voice technologies and related specifications. Dave holds a Telecommunications Engineering degree from the J.M. Perry Technical Institute, WA. He joined Intel in 1994 and can be reached at dave.a.lizotte at intel.com.

Glen Shires is a principal engineer in Digital Enterprise Group Architecture and Planning at Intel. Glen works with cross-company teams to define VoIP and SoIP (Services over Internet Protocol) end-to-end system architecture and usage models. He brings a system-wide approach to architecting platforms, having experience developing products at many levels from logic circuits and board design to architecting and developing software for major embedded and PC products. He specializes in speech, video, telephony, and Internet technologies and has been awarded five patents. His career began at Intel as an applications engineer for the Intel386™ microprocessor. Glen received his B.S. and M.S. degrees in Electrical and Computer Engineering from the University of Wisconsin-Madison. His e-mail is glen.shires at intel.com.

Scott Trevor is a senior engineer/architect within Intel's Information Technology division. Scott's current focus is developing the architecture for Intel's real-time communication capabilities as part of Intel's converged communications program. He specializes in the areas of presence infrastructure and its enablement and integration with various communication media. Scott has been working to architect, engineer, and deploy various collaboration products in the IT industry for 14 years, and has been working for Intel since 1994. Scott received his B.S. degree at Old Dominion University, Norfolk VA. His e-mail is scott.trevor at intel.com.

® ™ Pentium and Intel386 are trademarks or registered trademarks of Intel Corporation of its subsidiaries in the United States and other countries.

* Other names and brands may be claimed as the property of others.

For further information visit:

developer.intel.com/technology/itj/index.htm