WHITE PAPER

SIP: THE NEXT STEP IN CONVERGED IP COMMUNICATIONS

EXECUTIVE SUMMARY

IP communications technology—the convergence of data, voice, and video onto a single network—can help organizations to reduce the costs and complexities associated with communications and to enable progressive business gains. For more than five years, Cisco Systems® has been delivering the comprehensive benefits of IP communications to enterprises and service providers alike.

Session Initiation Protocol (SIP) builds on the IP communications foundation by providing a standards-based approach to enabling IP communications with numerous devices and applications. This paper describes the benefits of SIP to enterprise customers and explains the comprehensive Cisco® roadmap for delivering SIP-based solutions—an evolving strategy that helps deliver stronger IP communications benefits to enterprise customers today and in the future. Organizations that adopt Cisco solutions can reap the benefits of IP communications and establish a solid foundation to support new and emerging SIP networking technologies.

INTRODUCTION

A project manager receives an e-mail message informing him of a crucial change in a project currently underway. He needs to quickly rally his project team to deal with the new development.

He clicks on the “Project Team” icon in his address book and sends an instant message inviting the team to a rich-media conference to commence immediately. The icons next to each team member’s name indicate that five of the six team members are available. The sixth team member will receive the message on his cell phone, alerting him to join the conference in progress when he can.

Following the URL provided, each team member connects to a meeting that includes audio and Web conferencing capabilities, enabling all participants to access and update the master project planning document. Within minutes, the five team members are participating in the conference, even though two are on the road (both received the instant message via wireless connections to their PDAs). With a click of the mouse, the manager decides to record the meeting for the benefit of the sixth team member.

After discussion, the team members can begin taking care of tasks related to the project change. The manager e-mails a recording of the meeting to the sixth team member, who will be able to listen to the whole conversation as soon as he is available, and can view changes to the master planning document. He’ll be informed of the changes and the resultant tasks, without the manager having to recap the entire meeting.

Only 15 minutes have elapsed from the time the project manager received word of the crucial change.

Figure 1

An Integrated Communications Scenario

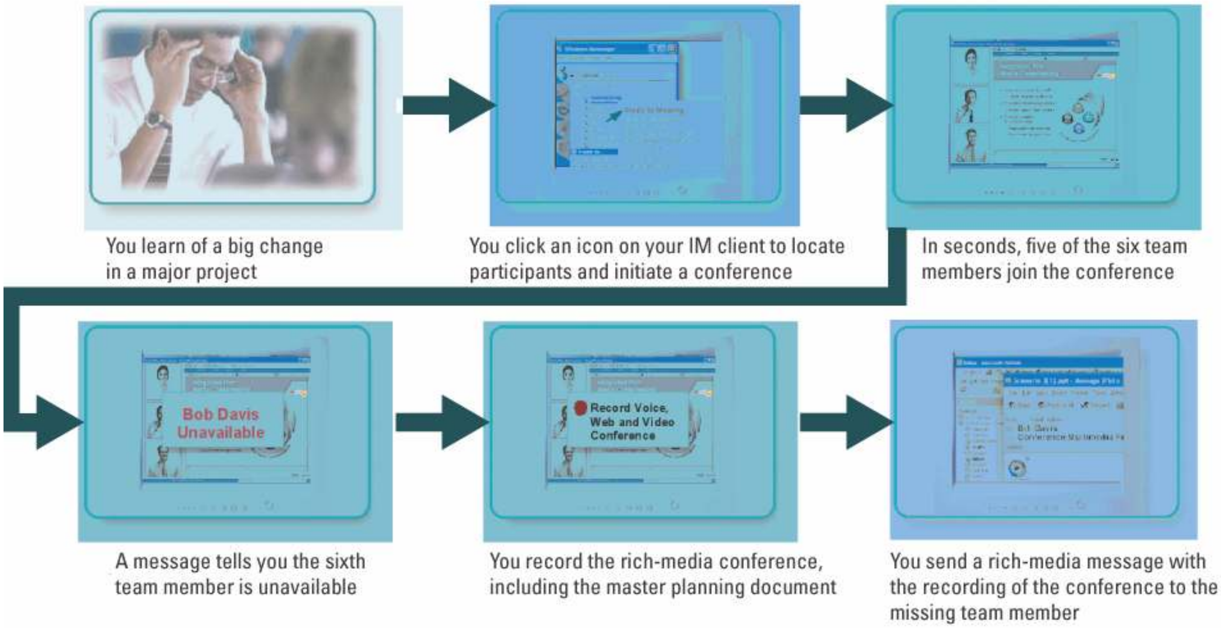

IP COMMUNICATIONS AND SIP

Converged IP network technology is a reality. Cisco has been delivering these types of productivity benefits for years, with solutions built on Cisco AVVID (Architecture for Voice, Video and Integrated Data), a blueprint for building secure, high-performance converged IP networks.

SIP is a peer-to-peer, multimedia signaling protocol that integrates with other Internet services, such as e-mail, Web, voice mail, instant messaging, multiparty conferencing, and multimedia collaboration. When used with an IP infrastructure, SIP helps to enable rich communications with numerous multivendor devices and media. SIP can set up individual voice or conference calls, videoconferences and point-to-point video-enabled calls, Web collaboration and chat sessions, or instant messaging sessions between any number of SIP enabled endpoints, including IP phones, PCs, laptops, personal digital assistants (PDAs), and cell phones. In the opening scenario, the participants could be using end devices from any number of different vendors, and if the devices supported the necessary SIP applications with sufficient attention paid to implementation, the rich-media conference call would work perfectly.

SIP is an IETF standard that promises to open up IP communications networks to new hardware and software players, giving enterprises more options and flexibility in building converged networks. At one time, enterprises that employed time-division multiplexing (TDM)-based PBXs had to rely on the PBX vendor to supply any required features and functions; now, converged IP networks and SIP open up the application development process, allowing applications from independent software vendors with expertise in specific vertical markets. This process is enabled by the approach the IETF has taken to SIP which is defining the base-level functions required for interoperability, but leaving room for differentiation at the application level.

Cisco has been instrumental in defining SIP standards. The company has been at the forefront of SIP technology since the first SIP IETF RFC was published in 1999. Cisco engineers currently co-chair both the SIP and the related SIPPING working groups, and the company has been delivering SIP-enabled products since 2000. Cisco has participated in numerous multivendor SIP interoperability and test events, and is a founding member of the SIP Forum industry group. In delivering SIP-based solutions, Cisco draws on years of experience building converged IP networks for enterprises as well as service providers—an advantage that is unique in the industry.

SIP Specifics

The core SIP specification is documented in RFC 3261 and in companion documents. These standards define the operation of SIP and how sessions are established, defined, and controlled. These specifications completely replace the original version of the SIP specification that was documented in RFC 2543.

The basic elements of a SIP system are user agents and proxy servers. The SIP user agent is software that is implemented in end-user devices and server components to manage the SIP connection. User agents include endpoints such as IP phones, SIP media gateways, conferencing servers, and messaging systems. SIP proxy servers route SIP requests to their appropriate destinations. Proxies are typically collocated with a SIP registrar, which maintains a list of contacts for specific users or accounts within a specific IP domain. SIP employs the Real-Time Protocol (RTP) to transfer packetized voice, video, and data in real time between user agents (Figure 2). This is the same RTP protocol that Cisco has employed in all of its IP communications solutions, giving the company years of experience with the protocol.

Dispelling the Myths

As is often the case with emerging technologies, some facts about SIP are overlooked and some of its capabilities are overstated. By itself, SIP is not a communications panacea—it works with many other standards to foster open, reliable, rich multimedia communications.

Important facts about SIP:

- SIP is a signaling protocol that is independent of transport protocol; it can run on top of several transport protocols, including User Datagram Protocol (UDP), TCP, and Stream Control Transmission Protocol (SCTP).
- SIP does not mandate or include specific quality of service (QoS) capabilities; it works with other protocols that perform this function.
- SIP is independent of any security protocol and may be used with several security protocols, such as Transport Layer Security (TLS) and IP Security (IPSec).
- SIP is a peer-to-peer protocol, not an IP-to-PSTN gateway control protocol such as MGCP or H.248.
- SIP provides methods to control sessions, but does not specify the applications and services that will use those sessions; as a result, SIP does not guarantee application behavior.
- SIP is independent of the media used, allowing the flexibility to initiate sessions for different media types.

Figure 2
SIP Architecture

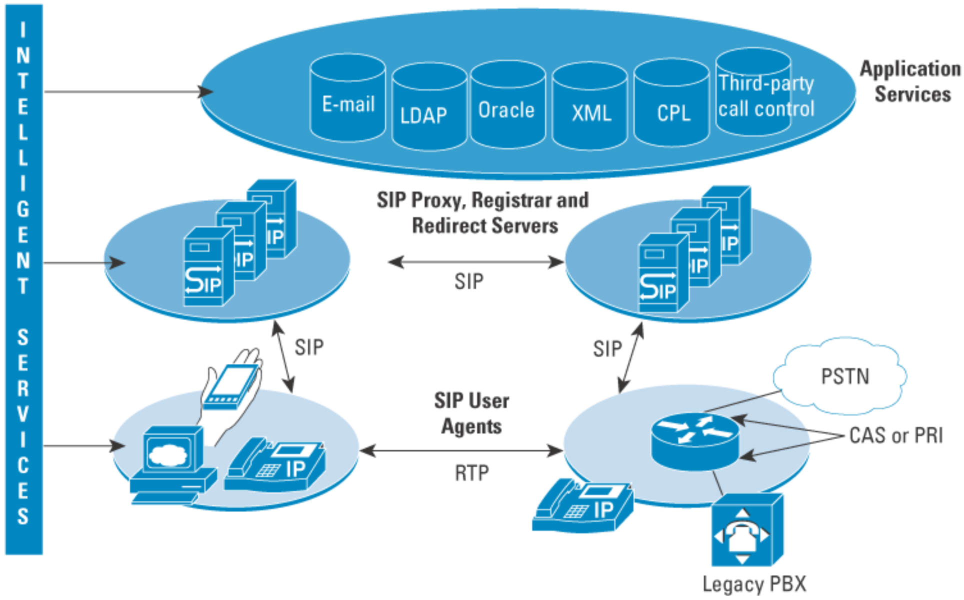

With SIP, users can be referred to in the same way that e-mail addresses are used. For example, a user could be defined as "sip: bob.jones@cisco.com", or by a telephone number format such as "14085553563". Using the database of registered contacts and locally configured policy, a SIP proxy can resolve a user address or telephone number to the proper destination IP address within its domain. SIP proxies also use IP services such as the global Domain Name System (DNS) to find SIP servers in other domains and the emerging ENUM standard to map public telephone numbers into an index to various IP addresses that DNS can recognize.

Adding Applications to SIP

SIP is attractive as a signaling standard because it can connect and control communications sessions between applications, independent of media type or the function performed by the end applications. SIP is known as a "methods-based" signaling protocol because it provides the methods to connect, signal, and control sessions. In that sense, SIP is quite different from a "functionally based" signaling protocol, such as Q.SIG, which is used not only to establish sessions, but also to define the specific features those sessions can support.

The distinction is important—it greatly affects interoperability and flexibility. As a peer-to-peer protocol, the intelligence involved in SIP-enabled applications is distributed to endpoints and other components, not centralized in a single call-control component. New features can be added without upgrading infrastructure components such as proxy servers, and developers do not require intimate knowledge of the SIP infrastructure in order to write SIP-enabled applications.

This opens up the application development process to third-party developers who can create targeted, vertically oriented applications. Internal users at a financial services company, for example, are likely to want features that are considerably different from those used by telemarketers in an outbound call center. SIP enables independent software vendors (ISVs) with expertise in each market to develop applications specific to those areas. This type of open development environment represents a dramatic shift from the traditional TDM-based PBX paradigm discussed above. By opening up the application development process to more players, SIP promises to provide more innovation in less time and at less cost.

SIP also gracefully supports protocol extensions, so that applications can support advanced features and can still interoperate with other, less functional applications. Consider the following example. Three colleagues are on a conference call. Two of them are at a headquarters location where their SIP-enabled IP phones support video capabilities; the third is at a remote office that does not support video phones. SIP will establish the conference among the three users, enabling the video portion for the two users whose equipment supports it. The third user will participate in a traditional audio call. This is a shift from the “least-common denominator” approach, in which only functions supported by all users are implemented (none of the colleagues would be able to use video).

In this way, SIP supports innovation within applications that can be combined, yet still works with similar applications that support features that have not been extended. In fact, SIP extensions define how SIP discovers the feature set each endpoint supports, and how SIP establishes each call accordingly.

Built on an IP Base

The SIP communications model is based on IP and uses existing protocols and services, including DNS, RTP, the Session Description Protocol (SDP), and Transport Layer Security (TLS). While SIP can work with any underlying transport protocol, it supports more advanced applications at lower costs and with better features on an IP communications infrastructure because it can take advantage of all of the well-defined services and protocols that IP offers.

Video and multimedia applications, for example, are far easier to deploy over an IP infrastructure where applications like video are treated as just another RTP stream. While SIP could be used to support video in a circuit-switched TDM infrastructure, doing so would require a complex series of signaling gateways to convert between packet- and circuit-switched formats.

CISCO AND SIP

Few companies can claim as much SIP expertise as Cisco, whose engineers have authored and contributed to several SIP-related RFCs, including the core RFC 3261 and others that address issues ranging from Network Address Translation (NAT) device traversal to presence and security. Cisco contributors have also chaired or cochaired several relevant IETF working groups, including SIP, SIPPING, IPTEL, ENUM, and MIDCOM.

Cisco released its first SIP-enabled Cisco IOS® gateway in February 2000, and its first SIP-capable IP phone, the Cisco IP Phone 7960, in August 2000. A SIP proxy server followed in 2001. While SIP was in its early stages, customers were able to build SIP-enabled networks with Cisco products, and many early adopters employed Cisco equipment to deliver revenue-generating services while learning more about SIP.

SIP Proving Ground

Much of Cisco’s early SIP work involved proof of concept—proving that SIP worked in an IP telephony environment, with traditional features such as caller ID, call hold, call transfer, and three-way calling. Capabilities such as announcing a call transfer to the recipient before connecting the call (which involves setting up a media stream prior to connecting the call) had to be tested to make sure they would work through RTP streams and SIP-to-public switched telephone network (PSTN) gateways. Similarly, TDM out-of-band signaling (which passes information such as caller ID data) had to be passed via the SIP network.

Most of the early adopters of SIP technology were young service provider companies looking to differentiate themselves by providing IP-based services. They include Vonage, the IP-based telecommunications services provider that used Cisco SIP gateways and SIP analog telephone adapters (ATAs) to launch its service in 2001, and B2, a Swedish IP-based telephony service provider. But even companies such as Microsoft saw the promise of SIP as far back as 2001, when it launched the SIP-enabled Windows Messenger Update for its MSN .Net initiative. The Windows Messenger Update enabled MSN users to make SIP-enabled voice over IP (VoIP) calls from the MSN Messenger client, with Cisco providing the SIP infrastructure to Microsoft service provider partners.

While Vonage, B2, Microsoft, and Cisco were helping to prove the viability of SIP in the service provider realm, Cisco was proving the viability of converged IP networks in the enterprise realm with Cisco CallManager and related products. When Cisco was first implementing its IP communications strategy in the late 1990s, SIP was not ready for enterprise deployments—many of its building blocks were still in the formative stage. Instead, for session initiation, Cisco employed H.323 coupled with its own Skinny Client Control Protocol (SCCP) and with the

Media Gateway Control Protocol (MGCP), using these protocols to help companies gain experience with and reap the benefits of IP communications. Through this process, Cisco gained valuable experience in delivering enhanced voice, data, and video services over an IP infrastructure, while developing applications that take advantage of that infrastructure.

Cisco has worked with other vendors to ensure interoperability for SIP. In addition to its work with the IETF, Cisco has been a regular participant in SIP its since its inception in 1999. SIPits are weeklong, engineering-driven events coordinated by the SIP Forum, where vendors work out interoperability issues in their SIP implementations.

These parallel efforts are coming together as Cisco delivers on its strategy to incorporate SIP into Cisco CallManager and the rest of its IP Communications products.

THE CISCO SIP ROADMAP

Cisco is uniquely positioned to take advantage of all of the capabilities that SIP has to offer. With the largest SIP-enabled product line in the industry, Cisco is applying its own SIP experience to its enterprise solutions, helping to ensure that the company's implementations are interoperable. Cisco is also working with partner companies to help ensure that their SIP-enabled applications mesh well with the Cisco infrastructure. This strategy will bring new capabilities to users by opening up the application development process to new players—enterprise customers will no longer need to rely on PBX vendors for all telephony features and applications. At the same time, the open nature of SIP will enable enterprises to buy equipment and software from more suppliers, giving them more flexibility.

Adding SIP to the Enterprise

For enterprises, SIP offers at least three major benefits:

- **Interoperability**—Customers can build systems using components such as IP phones, soft phones, and collaborative applications from multiple vendors. This allows them to select the combinations of price and features that best suit them, while giving individual end users the devices they are most comfortable with and the software that works best in their environments.
- **Innovation**—The building block approach that the IETF has taken in defining SIP leaves room for innovation in terms of the applications that vendors can deliver. An early example is the implementation of presence around voice and instant messaging applications. As in the scenario in the introduction, presence data can be a powerful addition to various applications, enabling users to instantly determine whether a colleague is on the phone; out of the office but available by cell phone; online and able to accept instant messages; and so on. SIP also allows users to express their preferred means of communications.
- **Investment protection**—SIP-enabled endpoints and proxies can be added to existing IP communications infrastructures. For example, Cisco has a roadmap for SIP-enabling Cisco CallManager, Cisco Unity™ unified messaging, Cisco MeetingPlace and other Cisco IP Communications solutions. Many customers have already built highly functional converged IP networks with applications such as IP-based call processing that allow for new levels of mobility and unified voice, e-mail, and fax messaging. Others have IP-based call centers that integrate voice, e-mail, and Web customer contacts. These customers are understandably concerned that any changes to core protocols will disrupt such applications, or will take away features that users and customers have grown accustomed to.

Cisco is working diligently to ensure our products migration to SIP will not take away features or cause disruptions to existing applications. As the company works to deploy SIP on its range of IP solutions, Cisco is making sure that SIP supports the primary features and functions that customers currently enjoy. Using SIP, customers will be able to take advantage of new options in terms of the endpoints they can deploy and, over time, the applications available to them from Cisco and third-party providers.

As a result of its pioneering SIP work and its parallel development of converged IP infrastructures, Cisco now supports the largest SIP-enabled product line in the industry (Table 1).

Table 1
Enabling the Cisco Product Lineup for SIP

Product	Description	SIP-Enabled
Cisco BTS 10200 Softswitch	The softswitch empowers service providers and large enterprises to gracefully transition to packet-based technology; serves as an interface to enhanced service and application platforms.	Yes
Cisco SIP Proxy Server	Call control software provides numerous call routing capabilities in both small and large networks.	Yes
Cisco CallManager	Market-proven IP call processor added SIP network-side support in Version 4.0; line-side support is in development to enable communications with SIP endpoints.	Yes
Cisco CallManager Express	Cisco IOS router-based call-control application for small businesses and branch locations offers SIP network support.	Yes
Cisco IOS gateways	Cisco IOS Software-based access router and gateway products have been supporting SIP trunking for enterprise and service provider customers since 2000.	Yes
Cisco PGW 2200 PSTN Gateway	Bridges legacy PSTN networks with packet networks; used by Swedish carrier B2 to launch residential IP voice service in 2002.	Yes
Cisco IP phones	Industry-leading IP phones; more than three million shipped, including many to service providers for SIP implementations.	Yes
Cisco ATA 186/188 Analog Telephone Adapter Software	ATAs are used by service providers to offer residential VoIP services since 2002.	Yes
Cisco Unity unified messaging	Unified communications and voice messaging server is used with Cisco CallManager and for integration with existing hybrid and TDM PBXs.	Yes
Cisco Survivable Remote Site Telephony (SRST)	Provides backup registrar and redirect services in the event that a SIP endpoint is unable to communicate with its primary SIP proxy; also provides a PSTN gateway.	Yes
Cisco MeetingPlace	Rich-media conferencing system includes voice and Web conferencing capabilities; ready to serve thousands of SIP endpoints.	Yes
Cisco PIX® products	Firewall products support both H.323 and SIP while gracefully addressing NAT issues.	Yes

Three-Phase Approach

Several years ago, Cisco defined a three-phase approach to introduce SIP into the Cisco AVVID enterprise solution. Cisco has completed the first two phases and is actively working on Phase 3, due to market in 2005.

Driving Development

In Phase 1, which began in 2000, Cisco introduced SIP support to its Cisco IOS Software-based gateways and firewalls, enabling enterprises and service providers to begin deploying SIP-based solutions. Cisco also added a SIP-based instant messaging client into several Cisco IP phones as a demonstration application. This client enabled an IP phone to integrate with instant messaging systems such as Microsoft Windows Messenger and Lotus Sametime, and was meant to demonstrate the strength of converged networks and to begin collecting market feedback to better understand SIP market requirements.

Given the success of its trials, Cisco is exploring the productizing of this feature such that users would be able to send instant messages to colleagues from their IP phones and to easily “escalate” the session to a phone conversation, if needed. Through this pioneering work, Cisco has gained real-world experience in how customers can employ SIP and use the capabilities that SIP enables, such as presence and instant messaging. Cisco has also discovered ways to integrate SIP-enabled applications with its own products, such as the advanced display capabilities of Cisco IP phones.

Full-Featured Solutions

In Phase 2, Cisco added SIP support to Cisco CallManager, Cisco CallManager Express, Cisco BTS 10200, and other system-level components. For example, with Cisco CallManager Version 4.0, delivered in early 2004, Cisco introduced SIP support similar to the network-side interface that it has with H.323 today, enabling customers to take advantage of SIP signaling over an all IP infra-structure to connect to core SIP networks which support various SIP services and applications as well as SIP User Agents (Figure 3). Additionally, SIP network-side support on Cisco CallManager enables users to employ hosted SIP-enabled applications from service providers, such as IP telephony and multimedia conferencing services. And because SIP is implemented natively within Cisco CallManager, as well as within all Cisco call processing agents and gateways, customers do not have to purchase additional software, servers, or network infrastructure to be assured of SIP support.

Figure 3
The Cisco SIP Enterprise Vision—Phase 2

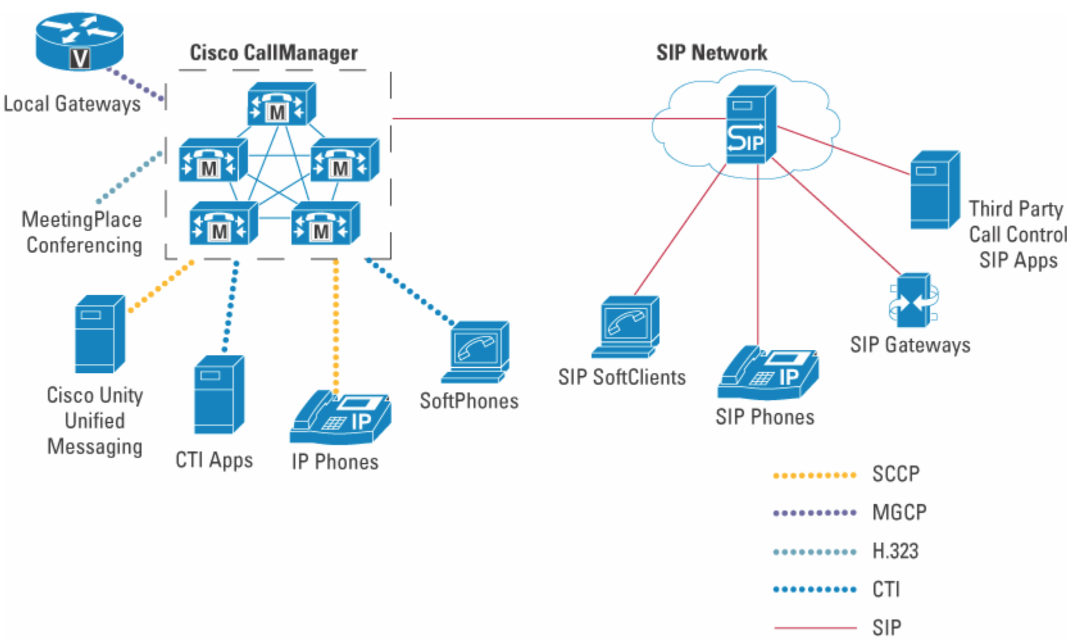

Cisco plans to build on the IP Communications solutions it has already delivered by providing full-featured, standards-based SIP solutions that are integrated with existing IP Communications solutions and services. This will provide investment protection for customers, allow them to extend the capabilities of their existing IP Communications investments, while delivering new levels of innovation with SIP enabled applications.

Implementing SIP on many call-processing components gives customers deployment flexibility. Enterprise accounts that need distributed call processing can deploy Cisco CallManager Express at smaller sites, with Cisco CallManager at larger sites and Cisco BTS 10200 softswitches providing a centralized dial plan and management and billing functions—all connected via SIP-enabled network links. The Cisco BTS 10200 can also provide 911 support for SIP endpoints, and has provided line-side support since early 2004. With SIP interfaces supported on Cisco Unity and Cisco Unity Express unified messaging systems, customers can tie these systems in via SIP network links as well.

Investment Protection

In Phase 3, Cisco is adding line-side support for SIP to Cisco CallManager and Cisco CallManager Express, extending SIP support to endpoints and User Agents (Figure 4), and allowing any SIP compliant IP phone to register directly with Cisco CallManager. Customers can migrate to SIP at their own pace, protecting investments in existing infrastructures and applications. For example, when Cisco CallManager SIP support is extended to SIP endpoints, customers can continue to use devices that support SCCP, H.323, and MGCP while integrating SIP endpoints providing equivalent functions that will take advantage of each protocol’s strength. Once again, SIP support will be native to Cisco CallManager and CallManager Express—no additional software or equipment will be required.

Figure 4

The Cisco SIP Enterprise Vision—Phase 3

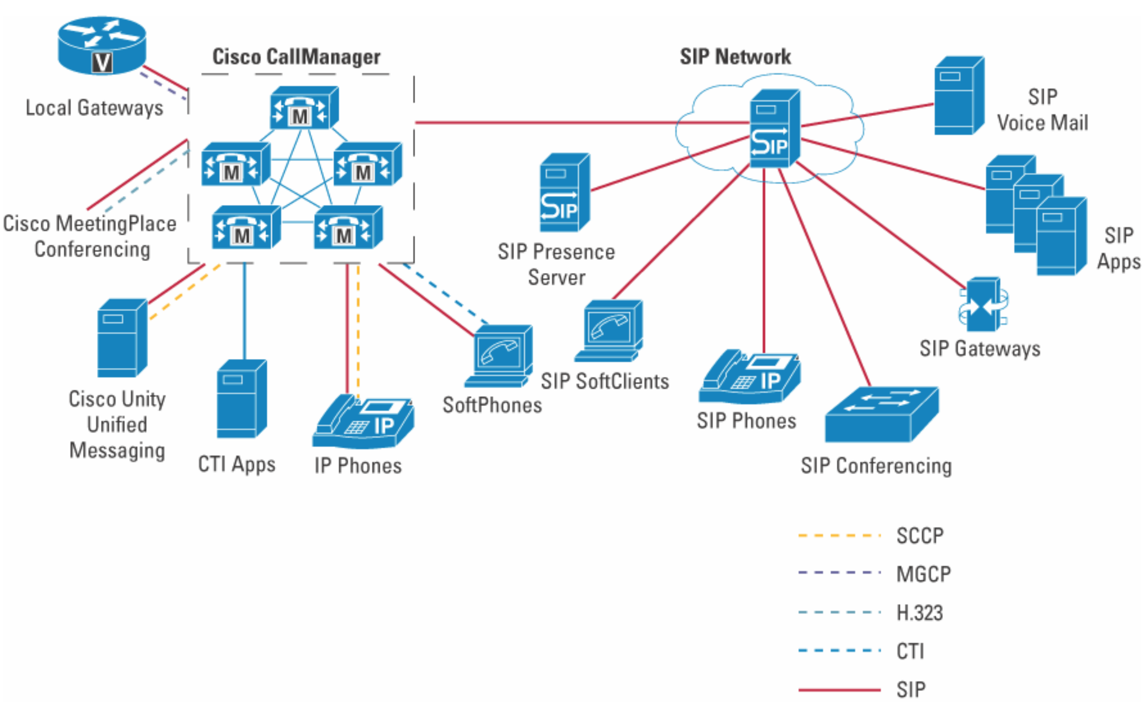

Also in Phase 3, Cisco and its partners will deliver new and innovative SIP-based applications. Initial plans will target rich-media communications applications that integrate voice, video, instant messaging, presence, and desktop applications to better serve enterprise customer needs. Cisco aims to deliver capabilities like those in the opening scenario to this paper, whereby end users can easily connect with each other, using the communications media and devices that make the most sense at the time. With numerous application developers working together to develop SIP-enabled applications, the possibilities are virtually endless.

An Army of Developers

The power and promise of SIP will truly be realized in the applications. Cisco is working with its partners to make sure that the individual needs of multiple vertical industries will be met with SIP-enabled applications that have been thoroughly tested for compatibility with the Cisco SIP-enabled IP Communications infrastructure.

Cisco is taking a systems-based approach to SIP to address interoperability and the features required throughout all of its products. This approach addresses four fundamental elements:

- **Network infrastructure**—The foundation network that carries all IP data, voice, and video traffic; includes workgroup switches, routers, and connecting links and helps to ensure proper QoS for all applications.
- **Call-processing systems**—Servers and associated equipment for call management, control, and accounting.
- **Endpoints**—IP phones, video terminals, and other user devices that connect to IP communications systems.
- **Applications—**User applications such as conferencing, unified messaging, customer contact, and Extensible Markup Language (XML)-enabled communications, as well as custom tools that extend the capabilities of IP Communications systems.

SIP AT WORK

In the Cisco vision, SIP uses the underlying Cisco IP Communications infrastructure to enable applications incorporating technologies that include presence and intelligent call routing, improving communications over multiple types of media, and ultimately enhancing user productivity. Some examples are included to illustrate this.

Find Me/Follow Me

SIP can be used to drive powerful presence capabilities that allow new levels of find me/follow me services. Many of these capabilities are being defined in the IETF SIMPLE Working Group. To date, presence capabilities have largely been limited to “off-line” or “online,” but the possibilities envisioned by SIMPLE—and Cisco—are far more sophisticated.

Imagine presence integrated with a calendar function. If an end user is in a meeting and receives a phone call from someone who is supposed to be in the same meeting, the call processor directly conferences the caller into the meeting; calls from others could be routed to voice mail.

Advanced applications can also recognize the hints that indicate the best method to reach a user at any given time. Next-generation (3G) wireless systems, for example, are expected to be SIP-enabled and will incorporate presence data. Such systems will be capable of recognizing when a user’s cell phone is moving, as well as its location. That presence data can be used by the SIP-enabled call processor to route incoming calls to the user’s cell phone, as in the example at the beginning of this paper.

Calling Preferences

SIP has defined extensions that deliver powerful calling preferences and capabilities, giving users new control over calls that they place and receive. When placing calls, users may choose not to complete a call if it is going to voice mail. Similarly, the user may indicate that calls from certain colleagues or customers should never go to voice mail, but be routed to an alternate contact instead.

Telecommuting and Travel

Presence and calling preference capabilities enable users who are working from home or on the road to employ the same features and functions they have at their offices. For example, a user can plug an IP softphone into any broadband outlet and begin receiving calls.

Presence-Enabled Directories

When integrated with directories, presence data can help users glean useful information from numerous applications. Consider a list of missed calls. An IP phone user reviewing a list of calls missed while away from his desk can immediately see the presence status of each of the callers—whether each is available now, on the phone, or away from the phone, for example. Such applications will greatly reduce missed calls.

Rich-Media Conferencing

SIP can deal with multiple forms of media at the same time. For example, two colleagues may be on a phone call with each other while conducting a Web conference to collaborate on a presentation. One session is on an IP phone; the other on a desktop PC. One of the users has to leave the office, but is able to transfer the phone call to her cell phone and the Web session to her PDA; SIP treats both sessions as one call.

The IETF continues to develop advanced SIP capabilities; for example, having multiple user agents cooperating on a user’s behalf. Picture a user entering a conference room where he needs to join a videoconference and a collaborative white board session. He uses his PDA to dial the videoconference number and simply checks off the white board as another media option. His SIP-enabled PDA communicates with the white board user agent to establish that session; the user is quickly online to both. When the meeting is over, the user can terminate both sessions with one click.

Interdomain Capabilities

Many companies have created IP- and SIP-based communications systems that work inside a single domain, such as within a single company. Because SIP is based on IP, it is not limited to such single-domain use. Just as e-mail can be shipped between any two companies with valid Internet addresses, SIP can initiate sessions between companies with different domains, and carry them entirely over the Internet. This requires proper security and authorization, as addressed within the SIP standards. As companies become more familiar with SIP and its capabilities, they will find that features such as call transfer and conferencing can be used between and among different organizations that support the appropriate standards.

Desktop Application Integration

Consider a Microsoft Word document that has been edited by several team members. As usual, each person’s comments are visible in the document, but a SIP-enabled document could also have contact information for each reviewer. With a simple click, a user could send an instant message to a reviewer to ask a question, or use presence data to determine whether a reviewer is available for a phone call to discuss the document.

THE CISCO ADVANTAGE

Given that SIP is intended to take advantage of other IP-based protocols, SIP-enabled applications are typically best suited to be delivered over an all-IP infrastructure. Applications that involve multimedia solutions, including video, can be delivered over an IP infrastructure—all forms of communications are treated as just another packet stream. Incorporating video and data into a TDM infrastructure on the other hand, while technically possible, involves deploying several gateways, dramatically complicating the initial deployment as well as ongoing operations.

Extensive IP Experience

Cisco has five years invested in its Cisco CallManager-based IP infrastructure and has proven the viability of converged IP networks in the enterprise (Figure 5). As it SIP-enables its IP Communications product lineup, Cisco can take advantage of the IP Communications services and internal call control capabilities already built into its products, as well as the many partners that can deliver applications and tools that work with the Cisco IP Communications infrastructure. Competitors that merely add a SIP line card to enable their call-control platform to talk SIP with certain endpoints will not be able to take advantage of the capabilities that an all-IP network provides.

Figure 5
Cisco IP Business Voice Architecture

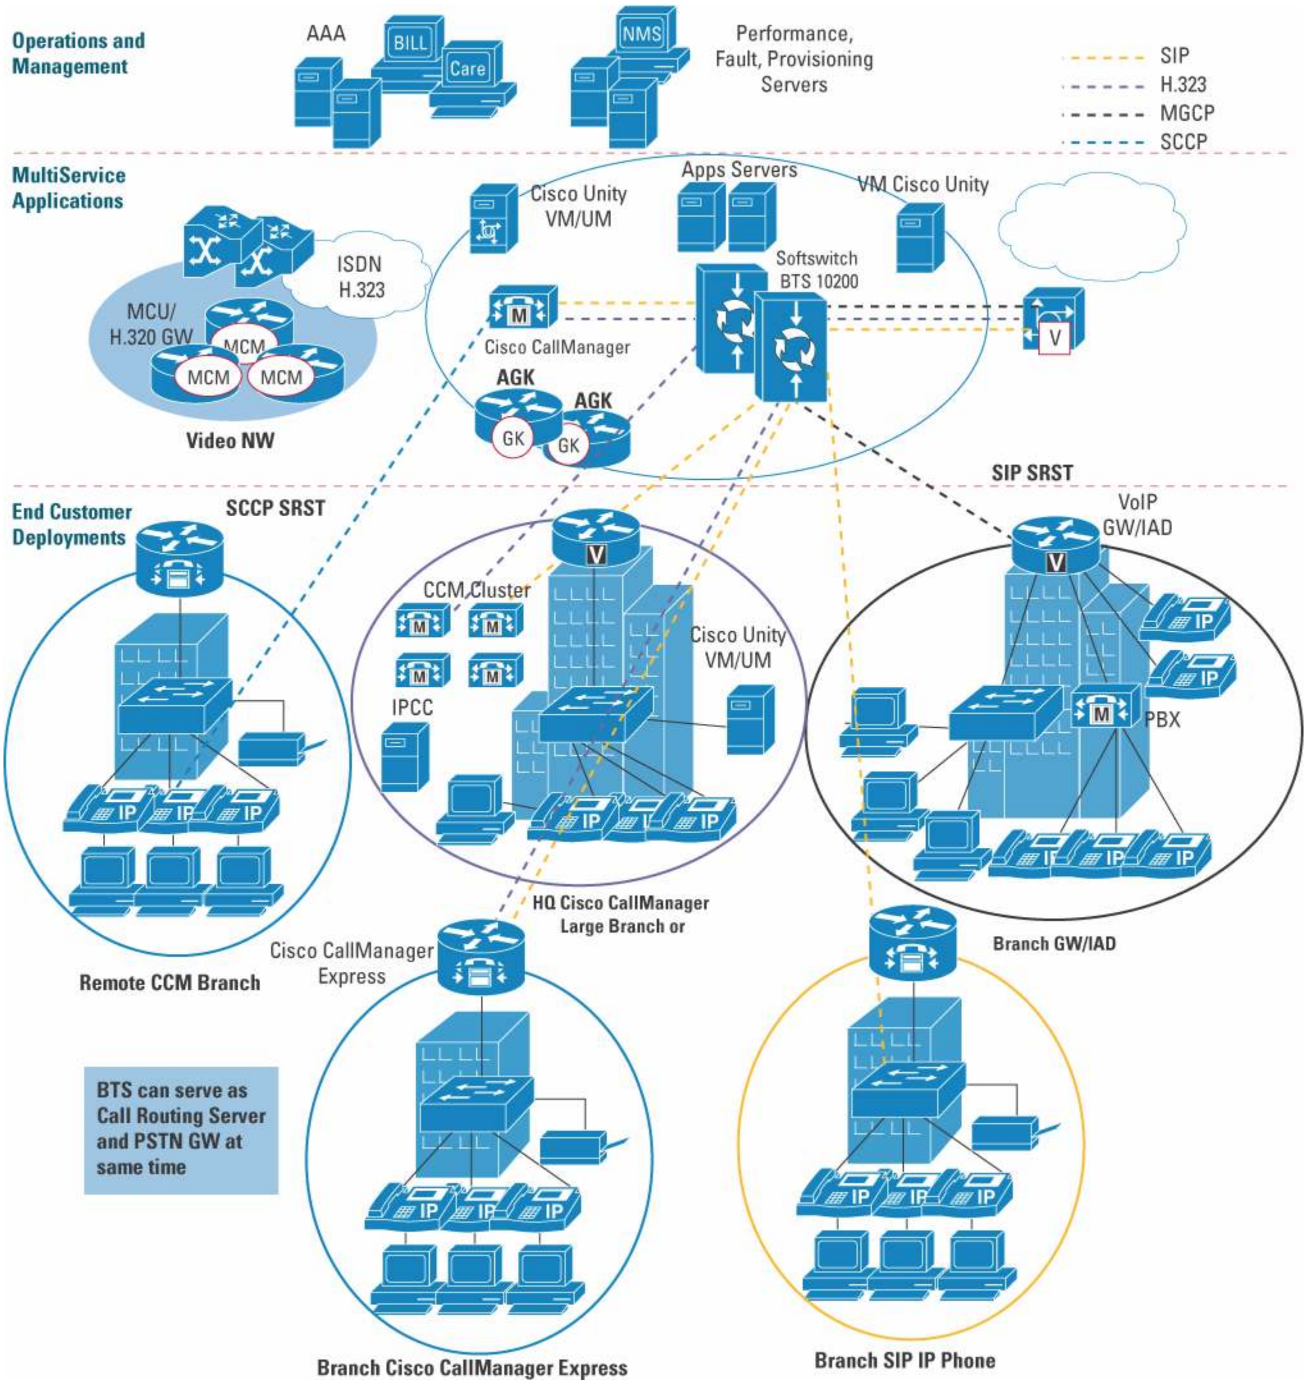

Some of these capabilities stem from the fact that IP allows for a separation between the call control infrastructure and the underlying transport medium. For example, Cisco CallManager Version 4.0 added support for video telephony, enabling users to add video capabilities to a telephone call. In practice, this capability is implemented by simply adding an additional RTP stream to an existing telephone call. The same features that are already supported for voice calls on the IP phone—such as call forwarding, three-way calling, and call transfer—work with the video stream. Supporting a video-enabled voice call on a TDM platform with all those same features would require a far more complex implementation, with protocols that enable setting up separate voice and video connections, then merging them into one logical session. Having all call processing functions delivered on an IP base makes such capabilities far simpler to implement, deploy, and support.

Unlike a PBX or hybrid environment, in an IP communications environment, telephony applications and capabilities can be integrated at the workstation or in an IP phone to support Web browser, calendaring, e-mail clients, and other applications.

SIP IMPLEMENTATION ISSUES

Implementing SIP by itself does not ensure that customers will be able to deploy a given product in a production network. The IETF continues to define numerous extensions to the core SIP standard that vendors may choose to implement. Each vendor must decide which SIP-related mechanisms its products will support and then make sure its products implement them in such a way that the products reliably interoperate with one another.

Cisco has developed a set of baseline requirements that define the SIP-related capabilities that all of its SIP-capable products will support. The baseline defines architectural considerations, security and trust models, privacy and identity handling, cross-domain operations, NAT and firewall traversal, media server and presence server requirements, caller preferences, and numerous other requirements. These requirements give coherence to Cisco SIP implementations, while offering customers Cisco SIP-enabled products that have been fully tested for interoperability.

As noted earlier, SIP and its extensions are designed as building blocks. The SIP standard only defines the signaling protocol for establishing and controlling sessions; it does not define how applications or features should be built, or how they should be delivered. All typical calling functions and features must be implemented using the building block functions embedded in end-user agents or proxies. SIP RFCs do not detail how SIP elements must be brought together to provide these functions.

This approach allows for flexibility in order to foster innovation, but also takes time and adds complexity for implementers of TDM-based systems. TDM standards typically spell out every detail, leaving nothing to interpretation, which can stifle innovation.

While SIP allows extensions to support communicating applications, the way that SIP signaling primitives are used to support a feature is a matter of best practices. Customers and vendors must plan and test applications carefully; in many cases, they must build these applications from scratch, exposing themselves to potential interoperability, manageability, and scalability issues as they deploy them. Customers should examine how vendors address various SIP implementation issues, including those that follow.

Interdomain Operation

As discussed previously, one of the advantages of SIP is the ability to communicate between domains directly via the Internet. Consider arranging a multimedia conference with a business partner company. Today, you would likely do this through a third-party conferencing provider. If both you and the partner company employed SIP-enabled multimedia infrastructures, the conference could be conducted directly via the Internet, with no need to reserve conference time or pay for the service. It would be as easy as putting multiple recipients on the “To:” line of an e-mail message. While this type of capability is eminently feasible using an all-IP infrastructure, it is far more difficult in an environment that involves TDM switching, which requires numerous PSTN-to-IP gateways. In either case, it requires compatibility between the applications on both ends.

Security

Information security is a primary area of focus for Cisco. SIP takes advantage of existing IP security standards to help ensure the integrity of communications sessions. SIP supports TLS, the successor to Secure Sockets Layer (SSL), to secure the signaling channel while Secure RTP (SRTP) encrypts the media to ensure voice privacy. Together, they represent a strong security system based on established standards including Advanced Encryption Standard (AES), the U.S. government encryption standard. Cisco already supports TLS in its SIP proxies and has pledged to support TLS and SRTP in its SIP-enabled IP phones and gateways. In addition, the Cisco PIX firewall and Cisco IOS firewall products currently support SIP.

NAT Traversal

NAT is a common way to hide a private IP address from the public Internet and to extend the number of IP addresses that an enterprise can employ. NAT gateways translate a private address coming from inside an organization to a different address that is conveyed to outside IP devices. That translation can be difficult in a SIP session, since SIP needs to know the IP address of each endpoint device involved in a session. Many security component vendors address the NAT issue by examining Session Description Protocol (SDP) information, and may try to resolve addresses that were changed by NATs. This approach can cause problems in certain scenarios, particularly if the signaling information is encrypted between the client and the server.

Another solution is a protocol called STUN (Simple Traversal of UDP Through Network Address Translators), which was coauthored by Cisco. When a user sends a message to a server from inside a NAT, the server will reflect back whatever address the NAT gives it. STUN allows this reflected address to be used to establish an RTP session with the user inside the NAT, without involving any of the SIP proxies in the middle.

Solutions to the NAT issue will vary depending on the exact scenario and environment; the industry has not settled on universally accepted solutions.

Enhanced 911

Enterprise customers have legitimate concerns about how all-IP networks can support Enhanced 911 (E-911) emergency features. Cisco addresses this problem with numerous solutions. Cisco Emergency Responder dynamically identifies the location of 911 callers in an emergency, with no administration required when phones or people move from one location to another. Cisco CallManager keeps a real-time location-tracking database with enhanced routing capabilities that direct emergency calls to the appropriate Public Safety Answering Point (PSAP) based on the location of the caller. For smaller offices, Cisco SRST, a feature of voice-enabled Cisco IOS platforms that provides voice backup capabilities for small offices, also supports E-911 services by connecting to the PSTN.

Feature Support

SIP is intended to provide interoperability (enabling customers to use IP phones from one vendor while employing SIP proxies from another, for example). Because many aspects of SIP-enabled applications are not defined in the standards, much of the implementation work is left up to the vendors. Customers may find that not all features supported by one vendor’s SIP endpoint will work with another vendor’s call-control servers and proxies; they need to specify which features they will need in their implementations.

Administration

One of the benefits of a converged IP network is ease of administration. Moves, adds, and changes become far easier to implement—no PBX programming is required. Should a user move a phone from one location to another, the Cisco CallManager will immediately recognize the unique MAC address of the phone and update its database accordingly. When implementing SIP, customers need to consider how much control administrators will have over their endpoints, including whether they can enable end users to reprogram their endpoints, distribute software to endpoints, and enable different classes of features for different groups of users.

CONCLUSION

With its years of experience of delivering effective IP communications solutions to both service providers and enterprise customers, and its intimate knowledge of SIP standards and deployment issues, Cisco is uniquely positioned to deliver SIP to the enterprise.

Cisco understands that customers already have significant investments in their IP infrastructures and that they need to implement SIP in a way that protects those investments. The company is doing the work required to ensure that all of the features and applications that customers currently enjoy with their IP installations will remain as they migrate to SIP deployments.

Cisco delivers SIP support on the widest array of user agents and call-control platforms in the industry, including Cisco IP phones, Cisco CallManager, Cisco Unity unified messaging, Cisco MeetingPlace rich-media conferencing systems, Cisco BTS 10200 softswitches, and the PGW 2200 PSTN gateways. Given Cisco's longtime participation in SIP standards and interoperability processes, customers can be assured that all Cisco SIP implementations adhere strictly to industry standards and are thoroughly tested for interoperability. Using Cisco products, customers will be able to build end-to-end SIP-enabled IP networks that address all enterprise requirements and connect to IP service provider networks.

In short, customers can expect to get more benefit from SIP, in a shorter timeframe, by deploying Cisco SIP solutions.

REFERENCES

- The SIP specification (includes overview, structure, and glossary of SIP-related terms):
  http://www.ietf.org/rfc/rfc3261.txt
- IETF SIP Working Group:
  http://www.ietf.org/html.charters/sip-charter.html
- IETF SIP Working Group Supplemental Website:
  http://www.softarmor.com/sipwg/
- IETF SIPPING Working Group:
  http://www.ietf.org/html.charters/sipping-charter.html
- SIP Forum:
  http://www.sipforum.com/
- Cisco SIP resources:
  http://www.cisco.com/en/US/tech/tk652/tk701/tk587/tech_protocol_home.html
- Cisco white paper, "Security in SIP-Based Networks":
  http://www.cisco.com/en/US/tech/tk652/tk701/technologies_white_paper09186a00800ae41c.shtml

The Cisco products, service or features identified in this document may not yet be available or may not be available in all areas and may be subject to change without notice. Consult your local Cisco business contact for information on the products or services available in your area. You can find additional information via Cisco's World Wide Web server at http://www.cisco.com/. Actual performance and environmental costs of Cisco products will vary depending on individual customer configurations and conditions.

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on **the Cisco Web site at www.cisco.com/go/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

204079_ETMG_LF_07.04

Printed in the USA