SPECIAL REPORT

IP Telephony Security: An Overview

By Joel Shore

Sponsored Exclusively By:

CISCO SYSTEMS
THIS IS THE POWER OF THE NETWORK. NOW.

This Special Advertising Section Produced By:

NetworkWorld
www.nwfusion.com

Table of Contents

IP communications, the convergence of data, voice, audio, and video into a single, high-performance network, represents a major opportunity for businesses large and small. By eliminating the need to maintain separate telephone and data infrastructures, extraordinary benefits are often achieved. PBX equipment, required for a traditional analog telephone system, is no longer necessary, leading to immediate savings. New kinds of applications that combine telephony with other types of data are for the first time possible. With only one network to maintain, administrative overhead is reduced substantially. Deploying IP telephony in regional or branch offices leverages the power of the Internet, further reducing communications costs.

Business benefits accrue quickly in an IP telephony environment. In addition to the new application types made possible by the union of telephony with network-based data, removal of the old, stand-alone PBX infrastructure creates immediate and substantial savings in administrative overhead and system maintenance costs.

Though the benefits of migrating a business' telephone system to the IP network are clear, doing so is not without challenges. Above all else, the telephone network is sacred, requiring that availability be guaranteed. The network infrastructure must be robust enough to support the considerable additional demands placed on it. Security must prevent theft of service, authenticate users, and repel a range of attacks from outside and inside the firewall. A comprehensive security strategy encompasses three key areas of focus: protection, privacy, and control.

The security situation is dire and getting worse. According to the Aberdeen Group, enterprises in the United States spent more than $3.5 billion managing security vulnerabilities in 2002. Of this, fully 92 percent was in the form of IT staff time with the other eight percent spent on technology solutions.

This paper examines the migration of telephony from traditional analog technology to the IP network, provides an overview of the two technologies, looks at the demands placed on the network, and scrutinizes the impact on security and methods for dealing with them.

About the Author

Joel Shore is Editorial Director of Reference Guide, a professional-services firm that provides market analysis and custom editorial services to technology vendors, and solutions strategies to businesses. His opinion column in Accela Communications' *Solutions Integrator* newsletter appears biweekly.

Previously, Joel was Editor-In-Chief of ITworld.com and host of the television webcast *IDC Live*. He was the cofounder and longtime director of the *Computer Reseller News* Test Center, the award-winning publisher of product reviews. Prior to CRN, Joel was a senior systems analyst at a major bank and director of point-of-sale systems for two national retail chains. A frequent speaker at industry events, he is the author of more than a dozen books on personal computing, and has appeared on numerous television and radio programs.

Benefits of IP Telephony

IP telephony (IPT) is a misnomer; the phrase suggests that telephony is separate and distinct from other IP-based applications. This is not at all the case. The sole job of a network is to carry bits. Various methods, or protocols, are used to do so. Once popular protocols, such as SPX and IPX dominated corporate networks. But with the rise of Internet and its ubiquitous World Wide Web application, IP (Internet Protocol) has risen to a position of supremacy. Like e-mail applications that transmit messages around the world, or a word-processing application simply saving a file to the server down the hall, telephony is just one more application running on the IP network. It is, after all, a stream of bits no different than a spreadsheet, an mp3 music file, or, unfortunately, a virus bent on damage.

It is specifically because telephony is just another IP application that extraordinary achievements are possible. The migration of telephony to the IP network and its melding with other sources of data provide a basis for new types of applications:

- When receiving a call, a photo of the caller can be retrieved from a database and displayed on the telephone desk set's color display panel.
- Users often away from their desks, such as real estate agents, can enter their schedule so that calls to their desk are forwarded to various other numbers based on time of day, with the caller completely unaware.
- Employees, as they arrive for work each day, log on to the phone system, ensuring their authentication and collecting time-sheet data for the payroll application.
- Through a unified messaging system, e-mail messages can be converted to audio format and "read" to employees who may be traveling and without access to a computer.
- An IP telephone deskset automatically registers itself when plugged into an Ethernet port, eliminating the need for programming or rewiring when an employee moves from one department to another.
- Videoconferencing, once the realm of large, bulky equipment that operated over dial-up or ISDN lines with often unsatisfactory results benefits greatly from the high-performance capabilities of the IP network and IP-enabled video gear.

As the number of applications run on the IP network grows, corresponding economies of scale are realized and the business value of the enterprise network increases. With separate voice and data infrastructures eliminated, so too are duplicate expenses and, in some cases, the need to maintain separate IT and telecommunications staffs. Wall ports at each worker's desk no longer need be designated "voice" or "data." IPT, after all, is data. The IP network allows network resources and clients to be distributed over wide geographical areas and connected with great flexibility; anywhere network connection is suitable for IPT.

According to the Telecommunications Industry Association, the long-anticipated migration to convergence technologies is well underway. Spending on enterprise voice and data communications equipment — driven largely by the move to IP — is forecast to reach $122 billion in 2007, a 6.7 percent compound annual growth rate from 2004 to 2007. IP revenue for audioconferencing, videoconferencing, Web conferencing, follow-me services, unified messaging and instant messaging were $1.5 billion in 2003, more than twice the $696 million of 2002, and are projected to reach $11.4 billion by 2007, a 66.5 percent compound annual growth rate.

A Tale of Two Technologies

Though the task of both traditional and IPT are to originate, route, and maintain voice-based communications sessions, the similarities end there. To understand the capabilities of each and the unique security needs IPT presents, an brief overview of the two technologies is in order.

Traditional telephony, running on its own dedicated network, relies on private branch exchange (PBX) technology. A PBX device, installed at the customer's premises, provides two capabilities. It lets a company's employees communicate with each other by dialing the appropriate extension. Equally important, it provides a means for a business to share trunk lines (outside lines) among employees. Doing so avoids the monumental expense of providing each employee with an outside line, a practice would be wasteful. It is through these trunk lines that calls are made to and received from the outside world. The business, or subscriber in phone industry terms, has its PBX connected to the pub-

lic switched telephone network (PSTN) in order to place calls to and receive them from outside parties. A business with operations — and PBX devices — in two cities may choose to link them via a direct link known as a tie line.

More than 125 years in the making, traditional telephone technology represents the highest standards of availability, reliability, and sonic quality. Pick up the telephone and a dial tone is always present. Indeed, it is rare for anyone to remember the last time telephone service was unavailable. The PSTN is highly reliable, and able to route and maintain calls with only the rarest of failures.

The bits that constitute an IPT conversation hitch their ride across the enterprise network (corporate intranet or the at-large Internet) on IP. Operating at Layer 3 (the network layer) of the seven-layer OSI (International Standards Organization) reference model, IP is a mechanism used to route packets on a network — and no more. The network layer chooses the path between the transmission endpoints (routing) and provides connectivity.

IP is oblivious to the bits inside the packets it transports. Like the Postal Service, which routes mail to the destination specified on an envelope regardless of its content, IP routes packets based on their destination address. What's inside isn't important — at least not to IP. But something else is needed to differentiate the contents of one packet from another, ensuring that bits from a telephony conversation are not interpreted as a database query, e-commerce transaction, or e-mail message. To handle IPT packets, two protocols are commonly used, H.323 and Session Initiation Protocol (SIP).

H.323 is an International Telecommunications Union (ITU) standard for sending real-time audio, video, multimedia, and other data over packet-switched networks. Adopted in 1996 and updated since then, the H.323 signaling protocol identifies several multimedia entities — endpoint, gateway, multipoint conferencing unit and gatekeeper — and defines their interaction. Well-suited for voice over IP, H.323 is an "umbrella" standard; it refers to other standards (H225, H.245, Q.931, and others) to describe its actual protocol.

SIP is a newer protocol, developed specifically for carrying voice over IP. Introduced by the Internet Engineering Task Force (IETF) in 1999, SIP since has been expanded to encompass video and instant messaging. SIP's signaling protocols are task specific: locate a called party, establish a session, negotiate sampling-rate and codec parameters, and tear down the session when the communicating parties hang up. A separate group of media transport protocols handles the digitization, encoding, decoding, and packetization of the caller's voice.

H.323 and SIP are different means to the same end: enablement of multimedia applications via IP. Each approaches those tasks from a different direction, however. H.323 was defined under the auspices of the ITU, a telecommunications industry organization. The IETF created SIP, which approaches those tasks from the IP perspective.

Why two protocols for IPT? H.323 arrived three years before SIP — and before IPT itself. As a result, H.323 has been widely adopted and deployed. SIP, designed with IPT in mind, is likely to eventually supplant H.323 as it finds new uses, such as phone-based instant messaging. Nevertheless, broad, industry-wide support for the large — and still rapidly growing — installed base of H.323 products ensures that they are guaranteed against obsolescence.

To initiate a call, several steps take place. First, the system's signaling protocol attempts to locate the called party. This can be done in various ways, depending on the destination, whether it is within the same enterprise system, at a branch office that supports IPT, or on a non-IPT phone connected to the PSTN. Once the destination party is located, its availabil-

ity is determined. Parameters to be used during the call are negotiated, a necessary step because products from different manufacturers often incorporate proprietary technology. The media transport protocol, such as the Realtime Transport Protocol (RTP), carries the actual digitized human speech. It performs its digitization, encoding, and packetizing in accordance with the specification set forth in the negotiated parameters.

Once the call is established, the signaling protocol stands by, ready to step in, if necessary. The signaling protocol might be called upon to change various call parameters and is responsible for tearing down the call upon its conclusion. Because of their varying tasks, signaling information and voice are likely to follow different paths. Whereas the signaling protocol might be routed through several signaling-related servers, the voice signal usually is sent on a direct path between the call participants.

IP desktop phones look and operate much like their analog precursors but, depending on manufacturer and model, might contain additional features such as a color display panel. Depending on the applications installed on the network, that panel might be called on to identify an incoming call by displaying the caller's photo or perhaps the internal employee phone directory. All the features associated with the traditional telephony system are present: three-way calling, call transfer, redial, hold, call forwarding, and voice mail access. IP telephony devices plug into an Ethernet port in exactly the same manner as a networked computer or printer, using an RJ-45 plug instead of the common RJ-11 telephony plug.

Security Issues for Traditional Systems

In January 1878, the first telephone exchange began operating in New Haven, Conn. Though security was not initially a concern, it quickly became increasingly important as businesses and government came to rely on the technology. With the widespread adoption of the telephone, PBX, voice mail systems, and later, dedicated voice networks, an array of threats exists.

- **Toll fraud.** The theft of long-distance telephone service often is accomplished through unauthorized access to the PSTN via a trunk (outside line) on a PBX device or voice mail system. In the 1970s, toll fraud became something of a sporting event, practiced by college students who came to be called "phone phreaks." With the advent of Tough-Tone technology, an industry sprung up building hand-held devices that emulated the Dual Tone Multi-Frequency sounds associated with each button on the telephone keypad.
- **Loss of privacy**. Eavesdropping on conversations ranges from the simple act of surreptitiously picking up an extension phone to the more complex scheme of intercepting calls at the PBX or at virtually any point between parties. With conversations transmitted as an unencrypted analog signal, anyone with a handset and pair of alligator clips could potentially tap in to a call. Even digital signals can be tapped and translated for eavesdropping.
- **Unauthorized access.** Tampering with voice systems, user identities, and telephone configurations provides a path for access to corporate phone systems. The ability to impersonate specific users and bypass password security can lead to the interception of voice mail messages.
- **Denial of service.** Flooding call-processing equipment with more incoming calls that it can handle often leads to a situation where the caller is denied service. It can happen legitimately: When a radio station contest spurs thousands of listeners to dial in simultaneously, most hear nothing more than an "all circuits are busy" error message. Similarly, the malicious overloading of call-processing equipment is easily accomplished. All it takes is a few modem-equipped PCs and a dialer application to continuously redial the same number. The resulting flood of incoming calls has the effect of denying access by legitimate users.

Without proper security measures, any part of a network — including the voice systems — can be susceptible to attack or unauthorized activity. Network attacks can create many problems, including service disruption, breaches in confidentiality or communications integrity, and damage to public image and customer confidence. Depending on the type of attack and what is compromised, consequences vary from mildly annoying to completely debilitating, and recovery costs can range from a small amount to profit-impacting sums.

Security Issues for IP Systems

Migrating telephony to an IP network in no way places the telephony function at any greater risk than before. Different risks, perhaps, but not greater risk. All networks, whether used for analog PBX telecommunications IP data communications, demand aggressive, comprehensive, and well-maintained security. With a properly designed network that places telephony on dedicated segments, and with firewalls and routers correctly configured, IPT is as safe as the PBX-based solutions to which we are accustomed.

With IPT based on a series of well-known standards and protocols that use widely documented ports, equally well-known weaknesses are ripe for exploitation. This is not to say that the modern enterprise should steer clear of IPT; on the contrary, with appropriate security measures implemented and kept up to date, IPT provides savings, economies of scale, and new ways of doing business that are simply not possible with older communications technology.

Within this universe of IPT security, an array of threats can exploit vulnerabilities.

- **Snooping/eavesdropping.** In theory, eavesdropping on an IPT conversation is possible, though doing so would require significant expertise. With authenticated access to the specific virtual LAN (VLAN) to which the IPT data packets are assigned and the ability to defeat the Layer 2 protection mechanisms that defend against eavesdropping, spoofing, and other specializations, it might be possible to intercept unencrypted RTP packets. RTP is essentially an unsealed envelope that has no integral mechanism for confidentiality.

  One way in which an attempt to intercept RTP packets might occur is to insert a tap into an appropriate Ethernet switch and extract traffic. By mirroring the switch port traffic, it can be captured and replayed later. However, with properly designed IPT equipment running on a network that has appropriate security measures, the interception of voice traffic should not be possible, even when direct access to network switches is available.

  As a last resort, an enhanced protocol, Secure RTP (SRTP), is available and adds confidentiality, message authentication, and replay protection. With SRTP, snooping remains possible, but decoding encrypted packets becomes impossible. SRTP is not appropriate for everyday use, requiring significant processing overhead for encryption and decryption of data packets. It is a technology best reserved for special cases when the privacy and security of every bit is required.

- **Denial of service.** H.323 and SIP, the two signaling protocols that IPT uses, can become a target. Similar to attacks against Web sites, an IP assault typically floods an IPT device, server, or software application in an effort to outstrip its available resources. The result might be unwanted disconnects or false busy signals. RTP, the protocol which carries the digitized voice signal, also might be the target of so-called data-flooding attacks. In addition to IPT-specific denial of service attacks, the core network operating system is vulnerable to an array of assaults, just as it is in a traditional data-only network. Protection against flooding can be accomplished with TCP interception at Internet gateways. Firewalls, with properly configured threshold settings, manage the flow of packets and may shut down that flow when necessary. In IPT, further protection is available with the digital certificates incorporated into an increasing array of devices, including telephone handsets. These certificates are used to ensure that device processes only those commands that originate from trusted sources.

Denial of service does not always require an overwhelming flood of packets to do its work. It also can be carried out with pinpoint precision. By sending a forged password-change control pocket to an IP telephone, the genuine user, unaware of the password change, suddenly will be denied access. To protect such an occurrence, digital certificates commonly are used to authenticate management and configuration traffic. A network or telephony administrator can reset the password, but time and productivity are lost in the meantime. Should such an attack become widespread, restoring access to affected users quickly can mushroom in a major undertaking.

IPT applications, most notably voice mail, are equally at risk. Fraudulent password changes and deletion of stored — and perhaps not yet listened to — messages are a real possibility. Even worse is fraudulent access, the act of stealing and listening to others' voice mail messages. With corporate CEOs, CFOs and other top-ranking executives the likely target, sensitive business and financial strategies are at risk for espionage.

- **Packet sniffing and call interception.** With traditional analog PBX-based telephony, calls can be intercepted with the simplest of technologies, a handset and pair of alligator clips. In IP-based telephony, voice content is converted to digital format and broken up into packets for transmission over the IP network. In theory, packet sniffing equipment could be used to snatch packets as they move over the network, but reassembling them is a significant undertaking, one not likely to meet with success. The closer one is to an endstation (such a desktop IP phone), the easier it is to intercept a large number of packets. Consequently, user authentication at the telephone itself is an appropriate security measure. Move further into the network, toward the backbone, and packet interception becomes far more difficult.

- **Viruses.** IP networks are already under daily assault from viruses, Trojan horses, worms, malicious scripts, and blended-threat combinations. As a result, an increasing percentage of the overall IT budget is being spent on security, diverting scarce funds away from projects meant to support the corporation's core business initiatives.

- **IP spoofing.** Just as e-mail spammers disguise their true identify by sending messages that appear to come from a legitimate source, IP spoofers gain unauthorized access to protected resources by sending messages that appear to come from the IP address of a trusted host. This is accomplished by first finding a trusted host's IP address of and then editing packet headers to use that IP address. Any network service that relies on IP address authentication to determine a message's genuineness is at risk.

- **Caller ID spoofing.** When some telephones (IP-based or not) receive a call, the phone number of the calling party is often displayed, a product known as Caller ID that is based on Automatic Number Identification technology. Indeed, it is not uncommon for Caller ID information to be used as a security test; a call from a specific pre-approved number might be used to grant systems access to an administrator calling in from home. IP spoofing and Caller ID spoofing are very different threats. If the security mechanisms are properly enabled, it is impossible, even with a spoofed source IP address, to insert a fraudulent Caller ID. Spoofing the Caller ID may be more of an issue in service-provider-based voice-over-IP networks than in enterprise IP-based telephony.

Ensuring Availability

The security of an enterprise's telephony and data networks, heretofore separate from each other, benefited from that very distinctness. With each operating independently, loss of service in one would not cause problems in the other. Although the data network was likely to — in fact, was expected to — fail from time to time, this loss of service would have no bearing on the PBX-based voice network. Likewise, should the voice network become unavailable (indeed, a rarity), this loss of service would have no impact on the enterprise data network.

When migrating an established telephony operation from its PBX-based infrastructure to the corporate IP-based data network, maintaining the highest levels of availability should be a top consideration. High availability, an achievable goal in the modern enterprise network, is more than a comprehensive security strategy. It need not depend on leading-edge technology, but

rather on the implementation of best practices, on a network design that encompasses redundancy of key components to eliminate the potential for single points of failure, and even on the deployment of back-up power to assure continued service in the face of a weather-related power outage.

A key benefit of running the corporate telephony operation on the enterprise IP-based network is availability of the highest order. The goal of 99.999% availability (the so-called "five nines") that data networks and their large IT staffs continually strive to meet is tailor-made for telephone communications. It might even represent an improvement over the availability offered by a traditional PBX infrastructure, a system that often lacks comprehensive redundancy and component backup.

Implementing Security

Given that IPT is simply another of the many services running on the enterprise network — albeit on a separate segment — its associated security can be seen as an additional layer placed atop the security that already exists for the enterprise IP data network. A chain being only as strong as its weakest link, world-class IPT security is of little value if the underlying network is not adequately secure. As such, the implementation of IPT security starts with the core network.

A successful security implementation encompasses a range of actions, not all of them technical. Developing corporatewide policies regarding access control, authentication practices, and forced password rotation form a solid foundation on which security technology can be deployed. Indeed, in an increasing number of sensitive environments, such as health organizations and government agencies, physical authentication tokens or advanced biometric authentication, such as voice or fingerprint recognition, are becoming increasingly popular.

Hardening the network helps secure it against all the ills data can fall prey to — denial of service, spoofing, packet sniffing, viruses, worms, and the like.

The "hardening" of the underlying data network at Layer 2 and Layer 3 is essential. Hardening includes bolstering security on routers and switches with technologies including stateful firewalls, intrusion-detection systems, and intrusion-prevention systems. Stateful firewalls are best deployed in two key locations, wherever voice and data segments or VLANs meet, and wherever stateful monitoring is needed to protect voice services.

At the LAN level, intrusion detection is an appropriate solution where network segments meet, at switches and routers, at call manager clusters, and distributed throughout a campus environment. A signature-based technology, it watches for known patterns of bits and sends an appropriate alarm to ensure that malicious packets from the Internet or other groups are kept at bay. A related technology, intrusion prevention, watches for specific behaviors rather than bit patterns.

Specific types of hardware benefit from targeted hardening practices. Hardening a router, for example, may include locking down of Simple Network Management Protocol access, closing ports by default and opening them only when required, employing Secure Shell and other approved management methods, and authenticating routing updates. Hardening procedures for a switch operating at Layer 2 of the OSI model (the data-link layer) include Address Resolution Protocol inspection or private VLANs, implementation of IP permit lists that tightly control management port access, and disabling unused ports.

Other best practices include changing the user ID and password for network hardware devices from the common default of "administrator," ensuring that all firmware and

operating system patches are applied, shutting down unneeded services, and disabling ports to mitigate the threat of back-door programs.

Segmentation of Voice and Data Segments

Best practices dictates that IPT services and devices operate on a network segments that is logically disparate from other IP services. While employing the same access, core, and distribution layers of the overall enterprise network, segmentation of IP voice services affords additional protection against operational interruptions and attacks on the IP data network. To simplify management, telephony users can be grouped into VLANs.

Although manufacturers of IPT equipment highly recommend segmentation of IP telephony services from other IP, it is not necessary to build separate IP infrastructures. Doing so would be counterproductive and needlessly expensive. Technologies that keep the Layer 3 voice and data segments separate are plentiful. These include access control, stateful firewalls, and VLANs.

Some connection points between data and voice segments are necessary. These should be enabled only for specific functions, such as the processing of IPT calls and voice mail or unified messaging system.

Unified voice mail systems generally make use of the traditional e-mail message store in the data segment for voice-message storage. They also require communication with the IPT call-processing manager to provide certain services, such as notifying users of new voice mail messages. These service operate over widely documented TCP ports and represent a significant security liability. Deployment of a properly configured stateful firewall is an essential step in reducing the threat posed by denial of service and protocol attacks against the call-processing manager.

Other situations also call for deployment of a stateful firewall. These include IP phones in one voice segment connecting to the call-processing manager in another, users in the data segment browsing the call-processing manager in the voice segment, and others.

Conclusion

IPT is a mature technology that combines the best features of the traditional telephone network with the power and speed of the enterprise data network. The combination of these two essential technologies generates substantial savings through the elimination of the stand-alone analog telephony infrastructure and through the economies of scale possible by leveraging the IP network. The union of telephony and data on a single physical network also lays the foundation for a range of groundbreaking applications.

The implementation of an enterprise IPT solution is not an undertaking to be treated lightly. With the substantial incremental network traffic that IPT generates, a careful analysis of the network design and its capacities should be undertaken, with the expectation that additional bandwidth and equipment to implement IPT on isolated segments will be required.

Beyond the issues of network design, the threat of attacks on the IP network and IPT services in particular loom large. An aggressive campaign, led by a corporate-wide security audit is an essential prerequisite to deploying IPT. Through the use of stateful firewalls, intrusion detection, and hardening of the network operating environment, combined with enforcement of corporate security policies and best practices, IPT is a safe and exciting communications platform that will pay benefits for years to come.