



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

The VoIP Dilemma
GIAC Security Essentials
Practical Assignment
Version 1.4b
Option 1

Fernando Robles
06/30/2004

Table of Contents

Abstract	3
Introduction	3
Expectations.....	3
The Physical Dilemma.....	4
The Rogue Dilemma	5
The Equipment Dilemma.....	6
The Power Dilemma.....	7
The Encryption Dilemma	8
The Protocol Dilemma.....	10
A Legal Dilemma	10
Conclusion	11
References.....	12

Abstract

This paper intends to discuss the general security concerns that need to be considered during the design and implementations of a VoIP (Voice over Internet Protocol) converged infrastructure. Physical, operating system, power, confidentiality, and protocol issues will be covered to the extent that the reader should recognize the many facets of securing such an all encompassing technology.

Introduction

It is finally here, that promised nirvana of a fully converged, multi-service network, where one infrastructure and one support team can save a company untold thousands of dollars and provide ubiquitous telephony services to the enterprise. Now armed with the ability to roll out phone services wherever and whenever needed, information technology departments are finding out that managing the phone service is a bit more involved than the point-and-click adds, moves, and changes that VoIP (Voice over Internet Protocol) vendors promised us.

One of the most important (and basic) considerations, security, was never brought to the sales presentations. The implications of failing to address these fundamental issues range from the benign to severe; from annoying glitches in performance to a total loss of service in the enterprise. Now that this mission critical-service is in Information Technology's domain, it is imperative that a company's security policy is updated to reflect this vital network service.

Expectations

People have always identified phone service as a constant; it will always be there. No matter what kind of environmental phenomenon or disaster, when someone picks up a phone they are answered with a comfortably familiar dial tone. Other services such as cable, pale in comparison with the stoic reliability of POTS (plain old telephone service). Telecom switches, and by extension PBX's (private branch extensions), are usually large, expensive, complicated, and utilize proprietary management interfaces. A relative few understand the complex signaling that occurs on the PSTN (public switched telephone network). This security by obscurity has given POTS an ambiguity that has protected it from serious attack beyond toll fraud and general phreaking (short for phone freak, a hacker of the telephone system).

Considering these expectations and the fact that users generally don't expect (or often experience) this kind of reliability of their data networks, availability of the network becomes of paramount importance for the successful implementation of packetized voice. Availability encompasses several areas; physical security, power concerns, network architecture and design, and interruption of services via DoS (denial of service) attacks among other things. All of these aspects must be considered equally. Any vulnerability discovered may impact an organizations vital communications, both voice and data.

The Physical Dilemma

Physical component security is the often overlooked aspect of an organizations security profile, but when the network is now asked to fulfill the business critical tasks of telephony as well as data communications, physical security is of paramount importance. If an attacker has physical access to a communications server, router, or firewall he (or she) pretty much owns that piece of hardware. Not only that, they may be able to compromise the integrity or the availability of the data. Physical security encompasses more than just the communications infrastructure components; access to the work premises, employee security, etc. are other issues that need to be considered (but are beyond the scope of this paper).

As stated above, access to all infrastructure components must be controlled. This includes all servers, switches, routers, firewalls, management devices, etc. Of course there are a myriad of ways to do this, but the best one will have an auditable trail built into them.

Loss of power is an effective DoS attack. Mitigate this vulnerability by ensuring all wiring closet's power outlets are also secured. Make certain critical equipment (even in the closets) are on there own circuits. The last thing you need is someone plugging in a hair dryer and tripping a breaker, possibly rendering part of the network unavailable.

What about the phones themselves? Since it is an input device, it isn't feasible or sensible to control physical access to telephones. But you definitely want the ability to control which telephones can access the network, and what services they may use. With quality of service, scalability, manageability, and security concerns, best practices dictate that the voice network be segmented from the data network.

Most network devices today support VLAN (virtual local area network) technology making the same physical network appear as many virtual segments. IP telephones on disparate parts of the network will appear to be on the same local segment. If done properly, this segmentation augments security by keeping

highly vulnerable workstations and servers, and causal users from affecting the telephony devices in the voice VLAN.

The Rogue Dilemma

The threats of rogue telephones are numerous; toll fraud, eavesdropping, and the ability to impersonate other telephones (and possibly retrieve sensitive communications). Removing hubs, turning off unused switch ports, fixing the configuration of used switch ports, and filtering services between endpoints and segments will help mitigate the risk of rogue devices on the network.

“The best way to secure ... SIP-enabled IP phones is to put them on private, non-Internet-addressable addresses and to make sure there is good perimeter security for the LAN through the use of firewalls and intrusion detection systems” (Cisco). All phones should use RFC 1918 (private, non-routable) addressing and avoid the use of NAT (network address translation). This reduces the possibility of traffic crossing into the public network and reduces the likelihood that hackers may enumerate your voice network. Since most IP telephones use DHCP (dynamic host control protocol) to determine their address information, telephones may be susceptible to DoS (denial of service) triggered by IP address starvation. This can be mitigated by introducing a dedicated DHCP server into the voice VLAN, so if a DoS attack is initiated on the DHCP server in the data VLAN, the DHCP server in the voice VLAN will be unaffected. The ability of rogue devices being introduced into the voice VLAN is also reduced, though not entirely.

Statically mapping MAC addresses to IP addresses on the DHCP server, thereby each phone will always boot with the same address, will make it much harder for a rogue device to be plugged into the network as well. Any unknown device (that is not spoofed) will not receive an address. It also makes it much less likely that IP address starvation will occur and reduce the possibility of man-in-the-middle attacks. It is assumed and recommended that any auto-registration function has been turned off on the call processing servers. When new phones need to be added to the network they can be added manually, or if auto-registration is required, remember to turn it back off when finished.

Filtering should be used to control access between segments. Unknown devices (spoofed) in one segment should not be allowed access to other segments. Conversely, a rogue call processing server should not be allowed to receive connection requests from endpoints in other segments. Also, even though voice traffic is on its own VLAN, there is still interaction with devices on the data VLAN. It is imperative that a stateful or application aware firewall be used between the segments to control access. According to Cisco Systems' SAFE blueprint, the following is the most that should be allowed to travel between the segments:

- The voice-mail system when placed in the data segment connecting to the call-processing manager in the voice segment
- IP phones in a voice segment connecting to the call-processing manager in another voice segment for call establishment control and configuration.
- IP phones in the voice segment connecting to the voice-mail system when placed in the data segment
- IP phones in the voice segment browsing resources via the proxy server in the voice segment
- Users in the data segment browsing the call-processing manager in the voice segment
- Proxy server in the voice segment accessing resources in the data segment (Halpern).

By preventing rogue devices from joining the network and controlling access to essential services, you reduce the risk of toll fraud. By requiring user authentication you can almost eliminate it, although users are never too keen to remember another password.

The Equipment Dilemma

Many call processing servers today are based on general purpose operating systems. These operating systems, not specifically designed for one purpose, generally have many attack vectors and vulnerabilities that are exploitable. These vital servers must therefore be built with availability in mind. Systems and components should be made as redundant as possible, be they call-processing servers, switch fabrics, trunk lines, gateways, etc. Physical access must be controlled, redundant and backup power must be available.

The servers are most probably hardened by the manufacturer, but if not this must also be done. According to the president of Core Competance, Inc. "... Server hardening is a process of eliminating overly permissive defaults and unnecessary and potentially exploitable features (Piscitello). Hardening of the operating system includes the removal of unnecessary services, applications and features, the applying of service packs and patches to account for vulnerabilities that were previously discovered. By using a vulnerability scanning tool, you can detect all current vulnerabilities on your systems, and most tools also document procedures to mitigate these vulnerabilities. Maintenance of the servers is an ongoing task as new vulnerabilities are discovered and exploited.

HIDS (host-based intrusion detection system) is an excellent addition to any server. It provides automatic mitigation of monitored events which can prove to be valuable especially while the OS vendor is scrambling to provide vulnerability patches. If deploying HIDS, be sure to include all servers, including email servers (integrated messaging), management servers and proxies.

High availability encompasses server redundancy, network redundancy and capacity, and backup power. When designing self-healing networks, hopefully bandwidth considerations were taken into account when the original path becomes unavailable. The backup paths may need to carry double the traffic (and double the voice traffic, a congestion and delay sensitive application).

The Power Dilemma

Power is also another sensitive issue. Remember that most people are accustomed to phone service that survives loss of power and blackouts. The expectation has been set. But now the telephones are connected to RJ-45 jacks that may be powered from an Ethernet switch plugged into premise wiring instead of drawing power from the telephone company.

The distributed nature of voice over IP also means there are numerous devices on the network that also may need backup power. Careful considerations need to be made on the size and capacity of power backup systems. It's not enough just to backup the PBX anymore. Every telecom closet from the phone to the core to the gateways need some level of backup power. Every server that is responsible for any of the voice services also needs special attention. When sizing backup requirements it isn't good enough just to have enough power to shut down the servers gracefully. Now the network is tasked to deliver this lifeline service even in the advent of a total blackout. Luckily today's systems can be designed to supply almost limitless backup power. Options include stand alone UPS's (uninterruptible power supplies), or racks of batteries powering complete circuits within the enterprise, with gas or diesel generators just a relay away.

American Power Conversion, Corp states that to properly size the typical battery-based UPS, four factors need to be considered:

- The total power required in watts
- The run time required in minutes
- The level of redundancy or fault tolerance desired
- The voltages and receptacles required (APC).

Total power is self-explanatory. Remember to include any component that is part of the VOIP path, call-processing servers, powered switches, core switches, gateways, firewalls, proxies and routers. Also, include any network protection/security devices for this extended runtime as well. The last thing you want is a network that is functioning but not protected. As with network components and servers, UPS's may be designed with fault-tolerances built in.

Major components may be equipped with an N+1 redundancy, adding an extra level of confidence to the enterprise. Of course the requisite number of receptacles and voltages will dictate model decisions.

When deciding on run-time required, you might want to think that more is always better. But take into account that if all the power is out, it also means that the building's lights, HVAC, desktops, etc. are also generally unavailable. The possibility that most business operations will come to a complete stand-still, buying twenty-four hours of runtime may be overkill. Generally, one hour of runtime is a good starting point but every enterprise must determine its own requirements.

Many, if not most, wiring closets do not have adequate backup-power available to them, especially when factoring in newer, current-hungry, power-over-Ethernet switches. *Business Communications Review* reports in an article about last years regional blackouts; "...These (wiring closets) did not necessarily have backup power in a data-only environment, but they must have backup at least equivalent to whatever's in the MDF and datacenter" (Krapf). Alas, many of these closets will not have adequate AC power to them as well. Careful consideration needs to be given to all the wiring closets if a high-availability voice network is to be successfully implemented. Also, with the small, cramped quarters that many wiring closets tend to be, the addition of powered switches and large UPS's may dramatically effect ventilation, another concern that may affect equipment longevity and availability.

One last caveat when it comes to a total loss of power; digital communication lines don't receive life-line power from the central office. If the power is out, and the PBX or VOIP system is connected via digital service only, calls outside of the network will fail. For this reason it is imperative that a small number of standard analog trunk lines be incorporated into any VOIP design. This will ensure that the enterprise can still make important and/or emergency calls when needed.

The Encryption Dilemma

Eavesdropping and packet replay attacks may be mitigated by using encryption. Even though you might have the voice on its own segment, the intrepid hacker may still be able to 'listen' to data from the data segment. It was previously mentioned that separating the voice and data networks using VLAN's was a crucial first step in securing VOIP infrastructure. Though this will stymie most users, the ardent hacker will find a way to this data. Tools such as DSNIFF, allow your switched Ethernet port to behave like a hub or shared medium, which enables them to listen to traffic on other segments.

After taking a tcpdump of all the traffic now visible to them, the intrepid hacker may use a tool such as VOMIT (voice over misconfigured Internet telephones) which takes the VOIP data stream and encodes it into a very playable .WAV file. Allowing PC-based phones on the data segments increases the chances of eavesdropping as well. Voice data streams, originally relegated to the voice VLAN, now are routinely traversing the data VLAN's making it easier for packet sniffing and replaying.

Encryption of the voice streams and headers are ways to mitigate eavesdropping and replay attacks. Encryption doesn't thwart a hacker's ability to sniff the data; it just makes it extremely difficult to make heads or tails of it. Encryption can be session based, as in SRTP (secure real-time protocol), or can be transport based, an example being IPsec VPN's (IP security virtual private networks). Used in conjunction with strong authentication mechanisms, message integrity is ensured as well as higher level of non-repudiation.

One problem with encryption is the performance and bandwidth penalties that arise with its use. A big attraction of voice over IP is the bandwidth savings over leased line facilities. The traditional voice call always took up 64Kb of bandwidth, allowing 24 simultaneous calls over a typical T1 leased line. By utilizing codecs (coder/decoder or compression techniques) that offer the right combination of bandwidth and quality, VOIP calls can be compressed calls to as little as 8Kb, substantially increasing the capacity of expensive leased lines. But using encryption can eliminate most of the bandwidth savings, the price of security!

IP telephones also need to support encryption in their firmware. As we touched on previously, many IP phones are susceptible to ARP attacks, spoofing, and eavesdropping. Without endpoint encryption support, Help Net Security states; "... Where VoIP handsets do not support the secure RTP protocol necessary to protect traffic ... it should be assumed that all communications could be intercepted" (Allsop). It would behoove the organization to carefully look at available IP telephones and assure themselves that the devices can participate in a secure environment (and to make sure existing devices are updated and/or upgraded to support such features such as SRTP).

Performance considerations relate to the endpoint's ability to encrypt and decrypt data in real-time and quality of service concerns. An IP telephone may not have problems encrypting/decrypting one call, but media gateways may become overwhelmed when supporting a multitude of encrypted traffic. These devices need to be scaled properly and additional hardware may be necessary to offload the main processors of encryption/decryption duties. Encryption may also introduce delay or jitter into the transmission, both anathemas to the quality of service that voice demands. Quality of service can be preserved with careful engineering of the voice VLAN's QoS attributes, or by over-engineering the solution.

The Protocol Dilemma

Up to this point we've been looking at the nuts and bolts of VOIP and the related security issues. Most of these were tangible and addressable by the enterprise. But we are at the mercy of the underlying protocols used. The varied protocols that a voice system may utilize are far from perfect. Indeed, many are still in the infancy of their development. With these growing pains come vulnerabilities that may be discovered and exploited by a nefarious attacker. Though a discussion of all the protocols is out of the scope of this paper, here is an example.

SIP (session initiation protocol) is a signaling protocol for VOIP, instant messaging, video conferencing, etc. (comparable to SS7 in the PSTN). SIP is an ASCII based protocol that sets up, maintains, and tears down calls between endpoints. The Oulu University Secure Programming Group in Finland discovered that simply injecting invalid data in the headers would cause buffer overflow conditions in IP phones and gateways. Oulu University Secure Programming Group concludes:

Although the test-material was designed as simple exercise of headers and fields in isolation, the failure rate was alarming. Only one from the sample of nine implementations survived the test-material as it is. This calls for a more comprehensive test-suite to be developed as the SIP scene matures (PROTOS).

Most vendors have since released updates addressing this vulnerability, but it shows that as we rely more on these protocols, more vulnerabilities will most likely be discovered.

A Legal Dilemma

The Justice Department is currently trying to persuade federal regulators to force IP telephony based communications companies to allow wiretaps of VoIP traffic. Federal agencies for years have had the ability to issue warrants for wiretaps of conventional POTS calls. With the advent of internet based telephony, the government is attempting to expand its reach to cover newer unregulated VoIP services. The government argues that in the Post 9/11 era it requires the ability to intercept suspected terrorist communications for pre-emptive measures.

But privacy advocacy proponents and industry leaders disagree with the Justice Department. The Associated Press reports opponents of the measure say the proposal "... isn't just unprecedented and overzealous but also dangerously

impractical ... It would chill innovation, invade privacy and drive businesses outside the United States” (Fordhal). New rules would require companies to get federal approval before rolling out certain technologies. It increases the risk that the very technologies that many corporations are building into their core business infrastructure, may also be at the regulatory mercy of the government. Not that the government may force companies to plant wiretaps on their employees. But it may impact security product development which could make products and technologies that would further secure an enterprise less available to the enterprise.

Conclusion

The challenges of securing a voice network may seem insurmountable, but in many cases much of the work may already be done. Voice over Internet Protocol, as its name implies, is a network service with many of the same security requirements demanded by a secure data infrastructure. An enterprise that has already done its due diligence may only need to address voice specific issues. Indeed, by re-examining the current infrastructure for voice security issues, existing data security is augmented. In any case, a multi-faceted security strategy will help ensure the availability of services, the successful introduction of new services, and the savings benefits of a fully converged infrastructure.

References

- Allsop, Wil. "VoIP – Vulnerability over Internet Protocol." 22 March 2004. Help Net Security. 28 June 2004.
<http://www.net-security.org/article.php?id=667>.
- "APC Solutions for Nortel's Business Communications Manager." 2003. American Power Conversion. 28 June 2004.
<http://sturgeon.apcc.com/whitepapers.nsf/169204752e0c52bf85256c77004d9e3e/59b69dca6538845085256d9e0063e7c6?OpenDocument>.
- Fordhal, Matthew. "Industry Fears Wiretap Plan Could Chill Innovation." 19 March 2003. 28 June 2004. SecurityFOCUS.
<http://www.securityfocus.com/news/8290>.
- Halpern, Jason. "SAFE: IP Telephony Security In Depth." 2003. Cisco Systems. 28 June 2004.
http://www.cisco.com/warp/public/cc/so/cuso/epso/sqfr/safip_wp.pdf.
- Krapf, Eric. "The Blackout and VOIP." Business Communications Review. Sept 2003. 28 June 2004.
<http://www.bcr.com/bcrlmag/2003/09/p10.asp>.
- Piscitello, David. "How To Harden Your Microsoft Web Server." 24 April 2003. Core Competance. 28 June 2004.
<http://www.corecom.com/external/livesecurity/hardenmws.htm>.
- "PROTOS Test-Suite: c07-sip." 17 Dec 2003. University of Oulu. 28 June 2004.
<http://www.ee.oulu.fi/research/ouspg/protos/testing/c07/sip/#h-ref15>.
- "Security in SIP-Based Networks." 2002. Cisco Systems. 28 June 2004
http://www.cisco.com/warp/public/cc/techno/tyvdve/sip/prodlit/sipsc_wp.pdf.