
VoIP Security

By Ken Camp

With the dramatic increase in VoIP deployments during 2005, there have been numerous surveys about security concerns. A few major threats to VoIP appear consistently in every survey. Falling prey to a Denial of Service (DoS) attack is most often the greatest fear. Worms and viruses are another major concern. Eavesdropping on calls raises anxieties for many companies deploying VoIP. In addition, VoIP spam gets a lot of mention in the press, but so far hasn't become a real problem.

VoIP security is a large topic, and the subject of many books and papers. This article provides a simple overview into the primary areas of focus and concern when addressing VoIP security. To keep things focused, we'll explore threats and dangers in the context of three fundamental security concerns—confidentiality, integrity, and availability. We'll cover VoIP network threats from the outside and the inside. In addressing how to deal with the threats, we'll stick to three fundamental precepts—prevention, detection, and reaction.

Characteristics of Security: Confidentiality, Integrity, and Availability

Confidentiality most often relates to ownership or control of the data. Data confidentiality is always at risk to malicious users who might discover, disclose, inappropriately monitor, or copy propriety corporate data. Protecting information privacy in VoIP networks is important. Keeping telephone conversations private is expected, just like in the public telephone network. In the public phone network, eavesdropping is difficult and requires a court order. In VoIP networks, eavesdropping through means such as packet sniffing may be technically easier to accomplish than public phone network eavesdropping. A fairly inexperienced attacker can potentially capture traffic. Free programs from the Internet, such as Audacity, may be able to reassemble these packets into an audio stream.

Integrity

Eavesdropping isn't the only threat. Some malicious users try to disrupt business traffic by degrading the integrity of systems. Integrity relates to preventing unauthorized or erroneous changes to network information. VoIP system integrity can be damaged by the insertion or use of false data. Integrity attacks frequently lead to modification, removal, repudiation, or misuse of system data.

Protecting integrity is frequently tied to user authentication procedures. System integrity requires that unauthorized users cannot make changes. Passwords should only be reset by authorized administrators or end users—only authorized administrators make configuration changes to the VoIP service infrastructure. The ability to make configuration changes is the attacker’s playground. If an attacker can configure the network, malicious activity can continue completely unnoticed. If an attacker gains configuration access, it’s no longer your network...it’s the attacker’s network.

Any system disruption or data corruption threatens system integrity. Malicious intruders aren’t the only threat to integrity. Legitimate employees make errors, sometimes taking unauthorized actions that cause problems. Disgruntled employees might purposely take harmful action. It’s important that user access levels provide permission to access only the resources needed to perform their jobs.

Availability

Availability is simply the loss of access, either to the network itself or resources on the network. Data might be unavailable because of corruption or destruction (an integrity breach). Congestion on the network may cause delays in accessing information. Any problem that makes the network resource difficult to use or access degrades availability.

Threats and Attacks

Attacks may come from outside or they may come from within. Attacks from the outside tend to be focused and malicious in nature. Attacks from within, although sometimes malicious, are often caused by user error or improper system configuration.

From the outside, pings, probes, and scans hit the network constantly. These are viewed as reconnaissance attempts, as outside users try to learn about your network. Worms and viruses are a constant and growing threat to business networks. As these problems commonly arrive via email, they are often admitted to the network, then cleaned, purged, or quarantined in an antivirus system inside the network. VoIP deployment creates a new potential attack vector in the gateways, VoIP servers, and unified messaging systems. In a fully converged network, attackers may attack the data network only in an attempt to get at the VoIP services within. Once VoIP services have been deployed, scanning from outside may increase as intruders probe for vulnerable services. VoIP services may open new TCP/IP ports for calling or signaling traffic.

Internal attacks are dangerous because they originate from a trusted segment of the network. Although they may come from malicious users, these attacks often come from inappropriate use of the network or some unexpected condition. Internal attacks are often the simple result of some new software program being installed somewhere on the network. A newly added, but misconfigured, server often creates unexpected conditions that may look like a scan or attack against the network.

What Is Being Attacked and Where Are the Targets?

To protect the network, you need to understand the targets. Some targets are logical. Others are physical targets. Logical targets are frequently user accounts. User accounts carry an associated set of accessible resources. Guest accounts typically have very limited permissions, but a domain administration account might have permissions to completely reconfigure the network. Running processes—whether in servers, routers, or other network elements—represent another logical target. Often the target is a program being run—its data, stack pointers, registers, and so on. Malicious attacks attempt to overflow memory buffers, creating unexpected results.

Physical targets are most often the network infrastructure—servers, workstations, routers, and so on. Domain controllers, VoIP service elements, and management platforms are also crucial physical elements of service delivery, making them key targets.

Who Is Attacking and Why?

It's important to understand that people perpetrate direct attacks. These types of attacks are very different than the perceived attacks caused by a misconfigured system. Some attackers may attack for the challenge. Professional criminals seeking financial gain are becoming far more common. The rise in identity theft problems online provides evidence of this increase in criminal involvement. In some cases, corporate raiders may be attempting to find a competitive edge through malicious attacks or hacking. Some malicious intruders want to cause damage by corrupting data. For some businesses or organizations, attackers may be spies or terrorists looking for information that can be leveraged for political gain.

Attackers don't play by any defined set of rules. There is no honor among thieves. They have a single objective—gaining access, control, or information by any means. Some attacks may be incredibly sophisticated technical efforts. Many attacks employ subtle, non-technical techniques:

Technical Attacks

Technical attacks often focus on the data leakage problem. Their goal is to intercept. Wiretapping is a technical attack but requires physical access to the wiring. Packet capture, in contrast, requires little technical skill. It can easily be accomplished with a free program. Intruders will attempt to breach access controls and circumvent security measures in order to access network resources.

Other forms of technical attack might include password guessing, perhaps using a dictionary attack. Theft of electronic media or *dumpster diving* can yield a wealth of proprietary information.

Worms and viruses aren't the only malware delivered in email. Spyware and keyloggers are often inserted without the end user's awareness, creating another data stream the malicious intruder can use to learn about the network, to learn passwords, or to steal proprietary information.

Reconnaissance attacks, penetration testing, and the exploitation of known vulnerabilities are becoming more common every day. Many tools to accomplish these attacks are downloadable from the Internet and freely available to anyone.

Non-Technical Attacks

Non-technical attacks focus not on technology, but on people. These types of attacks are frequently described as *social engineering*. It's human nature to be helpful. This truth can make it easy for attackers to misrepresent themselves to gain unauthorized access.

An intruder doesn't need much information to call the enterprise Help desk posing as an employee and requesting help in resetting a password. One very common technique is for intruders to masquerade as third-party vendors providing system support. In either of these cases, it's natural for support staff to want to be helpful, perhaps easing unauthorized access to the network. Don't underestimate human targets. They're always the weakest link in network security.

DoS Attacks

A DoS attack is always among the top fears. DoS attacks exhaust network resources. Because there are so many different resources available, the attack vectors can vary widely. IP addresses, network bandwidth, and processor memory are the most common attack vectors. Although DoS attacks take on many different forms, they all focus on starving out those resources critical to network operations.

These attacks can take on several forms. Buffer overflows exhaust system memory or CPU capacity, creating unexpected conditions. Some attacks simply attempt to consume all available bandwidth, degrading the ability of the network to deliver traffic. Routing and DNS attacks can lose or misdirect packets and disrupt information.

Distributed DoS Attacks

Distributed DoS (DDoS) attacks represent an especially troublesome variation. DDoS attacks involve several elements and are made of up *botnets*. Usually an intruder or worm delivers some form of malicious program, the *bot*, to the network. This bot program is often called a *zombie*. It's a software agent that gives outside control of the user's system to a *master* or *handler*, controlling the actions of the botnet. The bots will often log in to some chat channel and wait for a command from the handler. When an attack is launched, the users who have been infected with *bots* often don't even know that their systems are being used to attack some other victim.

A VoIP network presents many attack points for a DoS or DDoS attack. Every VoIP endpoint or phone contains some form of call agent software. Trunking gateways, signaling gateways, access gateways to other networks, and media and application servers are all server-based resources on the network. They all represent potential attack points. These VoIP services are often installed on servers running general-purpose operating systems (OSs) such as Windows servers. As with any Windows server, these systems are vulnerable to many Windows exploits.

A DoS attack against a VoIP network will disrupt the delivery of calls, presenting several manifestations. VoIP phones may not be able to register with the network. Internal calls might work but calls to the Public Switched Telephone Network (PSTN) might fail or vice versa. Users might not get a dial tone when picking up the phone, calls might be blocked due to lack of available resources, the caller might hear nothing but silence after dialing, or users might be able to dial, but calls may just fail to complete. Even if the VoIP network allows successful calling, conversations may be disconnected midstream. Calls could simply disconnect when the answering party picks up the phone. Network congestion could increase delay and jitter, making conversations unintelligible.

An overloaded network, under DoS attack might add 2500 milliseconds (ms) delay to every packet. This network would still work fine for email and Web services, but voice service would be unusable.

Viruses and Worms

Worms are self-replicating programs and are similar to viruses. A virus usually attaches itself to some other program; a worm is self-contained. Worms tend to propagate by exploiting file transmission capabilities of computers, most often email or network file sharing. Worms often attempt other undesirable actions. They may delete files or send documents via email. Some recent worms have carried other malicious code in their payload.

Worms tend to consume all available network resources, creating a DoS situation. Domain controllers and Active Directory (AD) servers are often targets, but VoIP resources bring a new rich target set to the network. You can expect new worms and viruses to exploit potential vulnerabilities in VoIP servers, endpoints (softphones and hardware phones), and the VoIP protocols (SIP, H.323, SCCP, and so on).

Eavesdropping and the Man in the Middle

There is a fear with VoIP that an intruder might be able to insert into the network between authorized endpoints, allowing a man-in-the-middle attack. Eavesdropping is simplified because the intruder either appears to be on the network or is listening in transparently. In network terminology, this security breach is called promiscuous listening. The intruder might be able to redirect traffic through spoofing or broadcasting alternate addresses.

Intruders eavesdropping on the network might also be able to alter call detail records (CDRs), modifying call setup information or corrupting billing data. An attacker might spoof SIP responses and redirect the caller to a rogue SIP address that intercepts the call. In practice, these exploits are achievable, but as yet, not widely seen in real-world VoIP implementations.

Spam Over Internet Telephony


Spam is a troublesome problem for anyone who uses email today. Users are bombarded with a stack of annoying, unsolicited junk mail daily. Imagine those messages filling not only your email inbox but also your voicemail inbox as spam moves to VoIP systems. Spam over Internet Telephony (SPIT) resembles spam in that it's essentially unsolicited junk phone calls or messages. SPIT has the potential to fill voice mailboxes with junk messages, just as spam fills your email inbox.

For now, documented cases of VoIP spam remain isolated and few, but SPIT could clearly pose a major headache. Although real-world examples remain slim in this area, it's certainly an area that gets a lot of press coverage. The largest real concern in the future is perhaps seeing SPIT used as a form of DoS attack.

Defense Mechanisms: Prevention, Detection, and Reaction

Preventive measures are the first line of network defense and can range from locking the server room door to setting up high-level security policies. Some preventative measures are simple and seem quite obvious. Other techniques are quite sophisticated. The best place to start is by identifying what needs to be protected.

1. First, protect the network. There are several problems to which networks are vulnerable. The classic problem is a DoS attack, but address spoofing is another common problem in networks under attack. To protect the network against increasingly sophisticated attacks, you need to build layers of defense in the network, each adding to the overall security and defense of VoIP services and other critical resources.
2. Second, protect the services. Every network service will present its own set of security requirements based on the intended use of the service. An internal service will have a much different set of security requirements than a service designed for external use. If VoIP is used only for internal calls within the company, simply protecting the internal servers from external access with access control lists (ACLs) might be adequate.

 Don't locate both internal and external servers on the same server hardware if possible. External traffic should be isolated using an outside-facing demilitarized zone (DMZ) segment of the network. Use firewalls and ACLs between networks to allow only authorized communications to pass.

3. Third, protect the protection. Be sure to protect the security and management platforms. It may be smart to create a special management segment of the network to house these critical services. As the management segment oversees the entire network, systems on this segment may need access to the entire network. It's important to ensure that the reverse is not left as a default access configuration. Management servers should not be accessible to anyone outside the network administration staff. Log servers should only permit authorized users to view log contents. Intrusion detection systems (IDSs) shouldn't be seen or touched by anyone outside the network security staff.

Firewall rules and ACLs can help determine where users can gain access, but only strong user authentication mechanisms will keep unauthorized users away from management systems. Discovery protocols such as Simple Network Management Protocol (SNMP) and Internet Control Message Protocol (ICMP) should be disabled whenever possible. Disable all services that aren't necessary on management systems. For example, don't run IIS, FTP, Telnet, or SMTP on a Windows server simply because that is the default configuration. Evaluate system requirements and only enable necessary services.


Effective network security is a combination of policies, processes, and technology. Firewalls can monitor the state of traffic and handle anomalies. Firewalls can monitor the state of VoIP sessions as well. Voice and data services should be treated as different zones of trust in the network. The goal is to control traffic flow between differing trust zones. Security policies define the rules that describe what types of traffic are allowed and what types are denied.

IDSs can monitor the flow of network packets and alerts from firewalls, routers, switches, and other systems in the network. They can identify traffic patterns either based on a digital signature or a newer heuristic approach that compares network traffic patterns with known *normal* baseline traffic patterns.

An intrusion prevention system (IPS) is a more advanced variation of this type of solution that has some ability to change network configurations in an automated manner. Instead of passively monitoring and alerting administrators, the IPS becomes an active management component. In one sense, the network becomes *self-defending*, with the ability to protect itself from some types of malicious traffic automatically.

Detection

In SANS security classes, the need for detection is often stressed as being even more important than the need for protection. Regardless of whether prevention is successful, detection systems will help you know about every event that occurs in your network. Sound detection tools minimize response time and speed mitigation. Incident management techniques can help you be proactive rather than reactive.

 For more information about SANS security classes, see <http://www.sans.org>.

Audit procedures and regular review of system logs help detect when information has been damaged, altered, or stolen; how it has been damaged, altered, or stolen; and who has caused the damage. As with prevention, detection is a combination of policies, processes, and technology. In addition to firewalls and intrusion detection and prevention technologies, syslog is a critical tool at your disposal that can be used in a variety of ways.

Syslog

The term *syslog* is used to describe both the syslog protocol and the application that sends syslog messages. The syslog protocol is very simple. The syslog sender transmits a small text message to the syslog receiver or server.

Syslog is used for network management and security auditing. Although syslog is very simple and may not be the very best tool for auditing, it has one huge advantage—it is supported in virtually every element of the network, allowing a corporate syslog server to become a central repository for audit log data. Syslog data is in plain-text format and is easily manipulated using simple tools and scripts. Many small to midsized organizations use a combination of scripts and spreadsheets to analyze syslog data when first starting out. Large networks produce much larger log files and will require more tools and processes that can scale to handle the larger data files.

Technology alone can't solve the detection problem. What you monitor, how you use log data, and how you react to incidents at the time of detection are all a critical part of the cycle of network defense. You employ detection mechanisms so that you'll know as soon as possible when an intrusion or other malicious event occurs. Network threats mutate quickly. Worms spread almost instantaneously. The threat of *zero-day* attacks will not allow for weak incident management prevention and detection processes. Effective incident management tools and processes ensure quick reaction and recovery when an event does occur.

Reaction

Layered defense is a term that is used in many ways. Just as you build security in layers, you build your incident management process layers. You take steps to prevent security breaches with firewalls and other preventative measures. You know that no matter how good your defense is, breaches will occur. You put detection tools and processes in place so that you can speed your reaction time. Your operations processes can be either reactive or proactive. How an organization responds to an incident is driven by how well prepared everyone is.

The proactive strategy is a pre-attack strategy. It requires steps to minimize existing security policy vulnerabilities and develop contingency plans. Identifying the damage that an attack will cause ahead of time requires careful analysis of weaknesses and vulnerabilities.

The reactive strategy is a post-attack strategy. It focuses on assessing and repairing any damage caused by the attack, then implementing any contingency plans developed in the proactive strategy. It's vital to document and learn from every incident, and restore business functions as quickly as possible.

Summary

Business networks are dynamic and complex systems. VoIP services add another layer of complexity to corporate networks. Security threats are real and continually on the rise. Security countermeasures come, not as a simple product to be purchased, but as a balance between policies, procedures, and technology solutions. Complexity and sophistication of both attacks and defensive strategies will continue to evolve. The malicious intruders have the upper hand in that they can use any means of attack. The best defense comes through comprehensive policies, diligent monitoring, and a cycle of ongoing review and improvement.