WHITE PAPER

Texas Instruments

IP Phone Security:
Packet Filtering Protection Against Attacks

By Atul Verma
Engineering Manager,
IP Phone Solutions
Communications Infrastructure
and Voice Group
averma@ti.com

Abstract

As IP phones continue to proliferate in the traditional voice communications marketplace, they are becoming enticing targets for hackers. For the marketplace momentum of IP phones to continue to grow, it is crucial to incorporate capabilities and tools to protect against current known threats while also being flexible enough to protect against threats that will emerge in the future.

This paper provides an overview of some of the security vulnerabilities of IP phones and outlines Texas Instruments' approach to thwart some of these threats.

Introduction

The advantages of a converged voice and data network are apparent. Cost savings is one of the primary benefits driving growth of Voice over Internet Protocol (VoIP) services in general and IP phones in particular. Other advantages include integrated applications, ease of deployment of new services and features, and high-quality voice.

As VoIP technology matures and IP phones offer more sophisticated and powerful applications, they become increasingly vulnerable to various forms of attacks. IP phones, which mostly run POSIX-compatible operating systems, are the phone equivalent of computers connected to the Internet. In enterprise environments, IT managers painstakingly ensure that their computers and networks are secure. IP phones will require similar attention.

The problem is even more serious in residential environments because most consumers are not vigilant about updating the software or firmware that runs their phone systems. As a result, residential systems are usually susceptible to an onslaught of ever-changing threats. Furthermore, while computers have the luxury of anti-virus software that, in many cases, automatically updates itself, IP phones are not designed with such capabilities. Therefore, there is a rapidly growing need to design IP phone hardware and software that provides resiliency against potential service threats.

IP Phone Vulnerabliities

Malicious users and their programs can exploit the vulnerabilities in IP phones at several levels. In fact, some of these vulnerabilities are not unique to IP phones. Most computer systems connected to the Internet suffer from these deficiencies as well. Some techniques commonly used to attack IP phone systems are:

Exploiting Bugs in Software

It is very common for hackers to exploit bugs in various libraries and operating system components to cause the target system to crash. Unchecked array bounds or uninitialized variables are usually the reasons for these vulnerabilities. Although these weaknesses do not surface during normal operations, they leave the system open to attack. For example, a crafted packet could cause an IP phone's buffers to overflow, or it could cause the device to execute unintended code. The end result is a system that either behaves abnormally or crashes and is entirely inoperable.

Exploiting Poorly Administered Systems

Careless practices such as unprotected configuration files, predictable file names, unsecured telnet access and open IP ports can render systems vulnerable to attack. These weaknesses make it easy for malicious users to affect the operations of the system or to gain control of it.

Exploiting Application Protocols

This involves exploiting protocol behaviors at the application level to cause undesirable system behavior. For example, with Session Initiation Protocol (SIP) it is possible for hackers to hijack an IP phone's registration and either prevent the delivery of incoming calls or redirect them to another endpoint on the Internet. Another tactic known as request spoofing impersonates a person's identity. Also, a proxy server can be impersonated and outgoing calls intercepted. There are vulnerabilities in the media path as well. For instance, voice RTP streams can be snooped and recorded. In addition, sensitive data such as PIN codes can be extracted when they are transmitted as in-band DTMF digits.

Exploiting IP Protocol

Certain classes of attacks take advantage of the Internet Protocol. Collectively, these have been characterized as Denial-of-Service (DoS) attacks even though other types of attacks, including those described above, can result in denial of legitimate service as well. Hackers or malicious programs typically exploit layer 3 and layer 4 of the IP stack to launch DoS attacks. Some well-known attacks in this category are ARP floods, TCP SYN floods, SYN fragments, ICMP floods, UDP floods, LAND and Jolt. The distributed versions of these attacks are known as Distributed-Denial-of-Service (DDoS) attacks. These attacks, in which unsuspecting host systems are harnessed into a collaborative launch, are particularly fatal and can clog data pipes in no time.

Remedies

The situation is not as grim as might be expected. Solutions are available to protect against many of these vulnerabilities. However, because the security weaknesses of IP phones can be traced to a wide spectrum of components, the solutions must also be broadly based on many different aspects of IP phone technology. In other words, there is no one-size-fits-all solution. A wide range of defenses must be mounted to counter all of the various types of attacks. A few best-practices recommendations are highlighted below:

1. Apply the latest patches. As security loopholes in commonly deployed software are exposed, vendors eventually issue patches and upgrades that should strengthen weaknesses. Attacks can be most effectively thwarted when system security is kept up to date.
2. Turn off IP services that are not being used and choose strong passwords for secure access.

3. Follow good deployment practices. For example, voice and data traffic should be partitioned on separate VLANs. This mitigates some of the problems associated with a shared medium such as Ethernet. For example, without VLANs it is fairly easy for any computer on the same LAN to snoop on voice packets using commonly available tools and programs.

4. Implement appropriate protection schemes at the firewall and router level to prevent attacks stemming from ARP spoofing, gratuitous ARP and other sources.

5. Some SIP security issues can be mitigated by using protocols such as Transport Layer Security (TLS) to secure signaling communication between various SIP components. Protecting key SIP messages such as Register and Invite with an MD5 challenge/response handshake can further buttress the security of SIP. Similarly, secure RTP (SRTP) can be employed to secure media transport and to protect against eavesdropping and other man-in-the-middle attacks.

Despite all the solutions, DoS attacks that target layer 3 and layer 4 of the IP stack are of particular significance to IP phones. These types of attacks are easy to launch and, because of the real-time sensitivity of voice services, even a mild attack can be catastrophic to a particular subscriber. But very few elegant solutions have been devised to safeguard against these threats. Most solutions rely on the crude approach of limiting the incoming traffic rates on a phone. Moreover, these approaches are implemented at higher layers of packet processing, rendering them ineffective beyond a certain data rate. So, in the case of an attack, substantial packet processing is performed at lower layers of the system before the packet reaches higher layers where a decision is made to throw away the packet. The net result is that even though attack packets are discarded, they consume valuable processing power. To compound matters, if the data rate is sufficiently high, it may end up consuming so much processing power that no real work can be accomplished.

Some might argue that corporate firewalls, access control lists on routers and switches, and other server-based intrusion-detection systems provide sufficient protection outside the IP phone environment. While these defenses at the perimeter of the network certainly are important, it is just as critical to secure networks from within and at the IP phone itself. A DoS attack emanating from a PC infected with a Trojan can seriously impact VoIP communications on that network. Furthermore, the problem is much more serious in residential environments where the network's security infrastructure may not be as formidable as that of an enterprise's VPN. Therefore, security solutions should be employed at every level from the network infrastructure to IP endpoints and in many of the entities in between.

In its latest Gigabit Ethernet IP phone solution, the TNETV1051/1052/1053, TI offers an integrated solution to protect against attacks that exploit layer 3 and layer 4 of the IP stack. This approach, which combines hardware performance with software flexibility, is discussed in the next section.

Static Packet Filter (SPF)

SPF is a programmable engine that can be configured to protect against many different types of DoS attacks. The diagram below illustrates the main functional blocks of SPF.

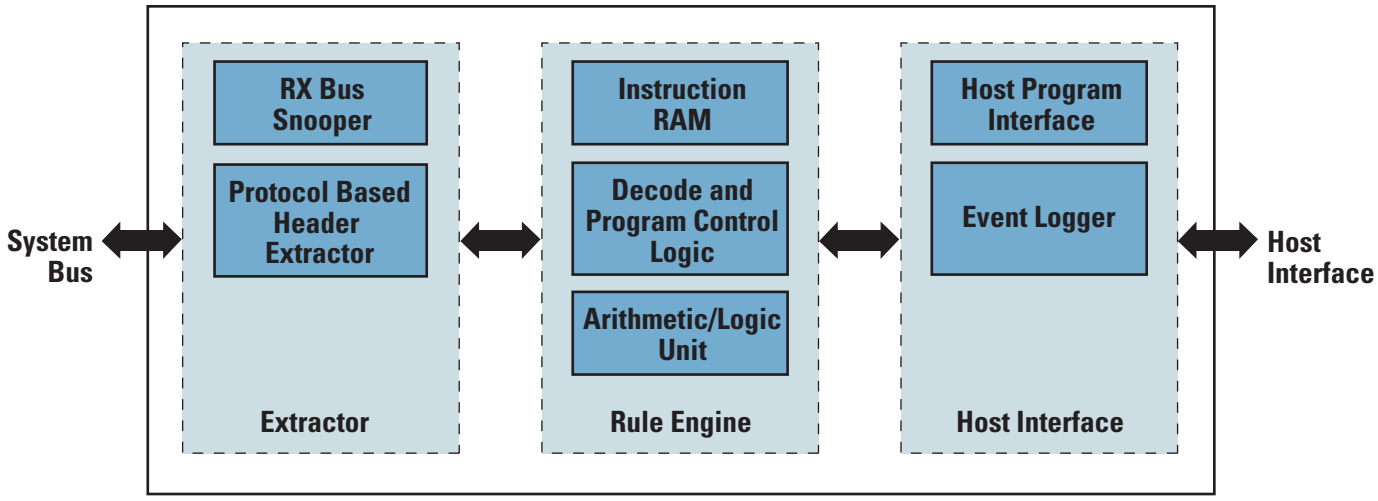

SPF Functional Block Diagram

SPF consists of three main components. First, the Extractor parses packets for analysis. Next, the Rule Engine checks the parsed packets against specific threats that it is programmed to identify. And last, the host system is able to control and configure SPF through its Host Interface.

SPF's Extractor interfaces with an Ethernet switch and processes received packets. Its protocol-aware state machine parses network packets in various formats, such as VLAN, PPoE, IP, TCP, UDP and ICMP. The location of protocol headers inside an Ethernet frame is determined, and offsets from the beginning of a packet for various layer 3 and layer 4 headers are stored in internal registers. The Extractor also refers key packet information such as IP source/destination addresses, TCP/UDP source/destination ports and ICMP type/code fields to the logger module for reporting. When an unknown protocol is encountered, the Extractor simply skips processing that packet and continues with subsequent packets. In such cases, however, SPF can be programmed to bypass extractor processing, and this task can be executed inside the programmable Rule Engine prior to packet analysis.

The Rule Engine is programmed with a specialized set of instructions that is downloaded by the host software prior to enabling SPF. The Rule Engine begins executing once the Extractor has finished parsing and loading internal registers with layer 3 and layer 4 header offsets. With the information contained in these registers, the Rule Engine processes protocol headers to identify malicious or malformed packets. The programmable nature of the Rules Engine gives the SPF immense flexibility for deploying a wide variety of checking algorithms to defend against previously identified or new threats. The Rule Engine executes its instructions for each incoming packet and determines whether or not to accept the packet. When an abnormal packet is detected, the Rule Engine instructs the external interface to drop the packet. Furthermore, it can also be programmed to limit the transfer rates of certain classes of packets. This feature is very effective for defending against flood attacks.

SPF also has excellent reporting capabilities. For example, it can record separate counts for as many as eight different drop points – instructions at which drop decisions are made. This information can be used to infer what conditions caused the packets to be dropped and at what frequency. Later, this information can be analyzed to determine which filters have been consistently attacked and whether any changes should be implemented in other network elements.

An event logger in SPF captures packet-filtering activity and can be programmed to log detailed information about dropped packets. The quantity of the information and the frequency of the logging activity can be adjusted by the host software. The information on dropped packets is stored in system memory.

A typical operational sequence for SPF would begin with the host processor downloading firmware for the Rule Engine and allocating system memory for logging information. After setting the log threshold registers and event interrupts, the host enables SPF processing. At this point, SPF begins filtering packets independently of the host. The host can disable SPF at any time and download new firmware to support different packet-filtering rules.

SPF combines hardware performance with software programmability to provide a robust and flexible defense against common DoS attacks. By implementing this solution at the lowest level of packet processing, the upper layers of system software are not burdened with receiving and discarding bad packets. The net result is an extremely resilient IP phone.

Conclusions

Any device connected to the Internet these days can be attacked by hackers armed with malicious and effective tools. IP phones are just as vulnerable as personal computers and servers, if not more so. To advance the already rapid adoption of IP telephony in the marketplace, effective security measures are needed at every level of the Internet and at strategic points in private IP-based networks. Mechanisms embedded in IP phones, such as TI's flexible and powerful Static Packet Filter engine, can be the last line of defense and, often, one of the most effective.