*By Patricia Hawkins, Eric Johansson and Edward Steinfeld,
Automata International Marketing.*

Securing Internet Enabled Devices

Now is the time, before you begin the design for an Internet connected system, to learn about the importance of network security and how to reduce the risk of attack. You should also consider involving a security expert in your design and review process.

INTRODUCTION

Imagine your VP calls you to his office. You are informed that the VP has just received a call from the head of the supermarket chain you outfitted with the very latest supermarket electronics all joined by a network. Apparently, one of the store managers, on opening the store this morning walked into a blast of hot air and the smell of rotting food.

On investigation, you discover that the freezers were turned off at 2 AM and the heat was turned on full blast at 2:30 AM. The security cameras show that nobody was in the store or had even been near the store at that time.

Congratulations. Your customer has been hacked. Now they want to know how to keep it from happening again.

If you did not design security into the system from the start, you may find it hard to put in after the fact. Your network topology, the firmware you used, the systems you have chosen, may all be working against you. You may need to ask the supermarket staff to change the way they work - in ways that they find awkward and inconvenient.

None of this would be happening if you had designed security in from the start.

Now, before you begin the design for an Internet connected system, is the time to learn about the importance of network security and how to reduce the risk of attack. You should also consider involving a security expert in your design and review process.

WHY INTERNET-ENABLED REFRIGERATORS?

A system that is not connected to the Internet at all is certainly safest from outside attack. As technology moves forward, that becomes less and less of an option.

An increasing number of manufacturers are heading towards Internet enabled appliances and devices. The reasons for doing this are many but include lower cost of manufacture, lower cost of service (by remote access), and coordinated management with other devices. In the example of a supermarket chain, that is monitoring the health of all of its refrigerator units and controlling energy usage over the Internet.

There are active projects to make the products in a supermarket Internet aware. The MIT Auto-ID center is working on a system to replace the UPC barcode on products with a code that contains reference to Internet-resident information. These barcodes will help the store manager to manage inventory and target marketing to consumers.

While these bar codes can provide a great deal of information to both the retailer and the consumer. Imagine waving a package of chicken in front of a home refrigerator and having a series of recipes from the manufacturer's Web site appear on the screen in front of you. This means that the store, appliance manufacturer, as well as the product vendor must be Web-enabled.

The future is Internet enabled and with great promises brings risks. It is important to understand those risks and manage them using network security techniques.

Candidates for attackers (or the wolf at the door)

It is also important to understand who are your attackers and why, because who they are affects how they behave, and determines the points at which they can access the system - and those details in turn affect security design. In most circumstances, the attackers are current or former employees, competitor's agents or script kiddies.

Depending on whose numbers you use, somewhere between 45 and 55 percent of all attacks come from people inside the organization. Internal attacks are generally not sophisticated; they usually boil down to exploitation of bad practices, such as sharing of passwords, or use of publicly known passwords. This means it is important to think about good internal network security, so that most internal attacks can only affect small areas of the network, or a limited number of devices. It is important to design security to make it is easy to follow good practices.

This leaves the other 55 to 45 percent of all attacks in the hands of much more sophisticated attackers. For the most part, the challenge is keeping the more sophisticated attacker off your network in the first place. Pay attention to and secure all entry points to the network. This means paying attention to physical security on the local network hardware as well as electronic entry points such as other network connections and the firewall.

The case of a competitor's agent is probably more the stuff of the tabloid newspaper than reality, but not enough so that you can ignore the possibility. When dealing with a professional attacker, the main goal is to discourage them from attacking, instead forcing them to try a different and - one hopes - a more visible route. The tools available are: VPNs, encryption, firewall, and strong authentication. These are sufficient to discourage the professional attacker.

When discouraged, an attacker will most likely turn to social engineering to gain information and a foothold on the network. Social engineering is a polite term for fooling someone inside the organization to reveal information about the network. It can be as simple as a phone call claiming to be the vendor trying to repair the system, but unable to access the system. "The password I have is ' Magic beans', is that the right one?" All too often, a person inside the organization reveals the real password without a second thought. Cultural consciousness of security is also important when protecting a system.

Security basics

Security helps defend against attackers by providing a series of obstacles between the attacker and their goal. Various defenses will work in most situations to protect against attacks. These defense mechanisms keep the attacker from doing harm by keeping the attacker off the network, protecting network traffic from being read, and by preventing the attacker from masquerading as an authorized user.

Keeping an attacker off your network is quite important. If an attacker cannot get on the network, they cannot do anything. If an attacker does get on the network, tools such as end-to-end encryption can keep them from being able to read the network traffic. Additionally, there are tools that allow only authorized users access to services on the network.

Security can never be perfect

Anyone who promotes security measures as a guarantee against intrusion or attack is overstating what network security measures can do. A realistic assessment of network security measures shows that security measures can give only two assurances. They raise the level of expertise needed to intrude and give you time to notice an intrusion in progress. Remember too that security measures are only as good as the people using them are.

When security measures are put in place, they make it more difficult to intrude on a network. The better the security measures, the more expert an intruder needs to be. Make the measures tough enough and intruders will try to attack a different part of the system to get around the security measures.

Put a simple padlock and hasp on a door and anyone with a crowbar can get in. Put a dead bolt lock on the door and anyone with a battering ram can break down the door. Put in a metal reinforced door with multiple dead bolt locks and it is time to start looking at the windows for in easier way in.

However, all the security measures in the world do absolutely no good if somebody leaves the door open or unlocked.

Diamonds or Rhinestones?

In designing a secure system, you will need to know how much security to put in place. Determining the level of security is a balancing act, where one balances the cost of security against the consequences and the likelihood of a loss. For example, people wear seat belts because the consequences of having a car accident without a seatbelt is very severe even though the likelihood of having a car accident is fairly low.

In considering consequences of attack, you need to look at the value of what you are protecting - are you protecting diamonds or rhinestones? If you are protecting a big vault of gem-quality diamonds, you want to apply strong security measures. If you are protecting a little box of rhinestones, then security measures do not need to be as strict.

You, as a device manufacturer, need to think through the consequences of an attack. For example, a supermarket's vendor price list -- what it pays for its produce

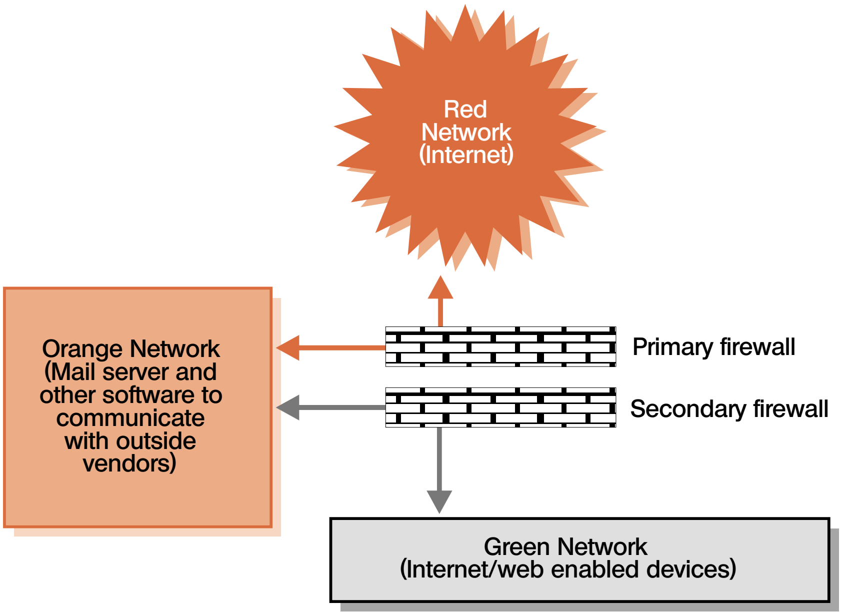
Red Network (Internet)
Orange Network (Mail server and other software to communicate with outside vendors)
Primary firewall
Secondary firewall
Green Network (Internet/web enabled devices)

Figure 1. Network security architecture using three levels of security.

- might not seem too valuable on first consideration. Nevertheless, a competing supermarket that has access to that price list month after month knows exactly how to manipulate prices to put your customer out of business.

NETWORK TOPOLOGY

Whether you or someone else will be setting up overall system security, you should understand network security topology, so that your design and choice of system components can fit into an overall secure system topology. You may want to call in a security expert to help with these choices.

A secure network generally has several different levels of security, described as red, green, and orange. A red network is exposed entirely to the Internet, an orange network accesses the Internet only by means of a firewall with carefully designed access rules, and a green network has access only to the orange network - again only by means of a firewall with carefully designed access rules

Only devices that require access from the Internet - such as a mail server - should be on the orange network. All other devices should be on the green network or on a separate, unconnected network with no Internet exposure at all (See Figure 1).

When joining networks, it is important remember that the network with the weakest security sets the security standard for the entire network. If an employees at their home office needs access to the green network, this connection is frequently made using a virtual private network (VPN). Therefore, if the home office has poor or no security, it reduces the security of your green network.

Never think that a device or system is safe and secure from attack just because it is behind a firewall. Firewall rules can be badly written. A firewall may be using an old version of software that has well-known security holes. A VPN to an insecure system can put your Internet devices at an orange or red level. A classic security mistake is to use a VPN to connect a secure network to an unprotected-telecommuters' home system that is completely exposed to the Internet.

You should decide the level of security or hardening that a device or system needs based on the value of the data on it and the consequences of a successful attack - and not on your assumptions about the level of security on the surrounding network.

You've considered the various items in your system, their value, and consequences if they are attacked - corrupted, fooled with, modified, stolen, or spied on. You know you do not want the settings on the refrigerators changed, you know you want to protect the computer with the accounting data, and so forth.

ENCRYPTION

There are a number of different encryption strategies available for use on the Internet, and for use within an internal network. A complete discussion of encryption is beyond the scope of this article. Encryption is used to keep unauthorized users from reading data being sent on the network.

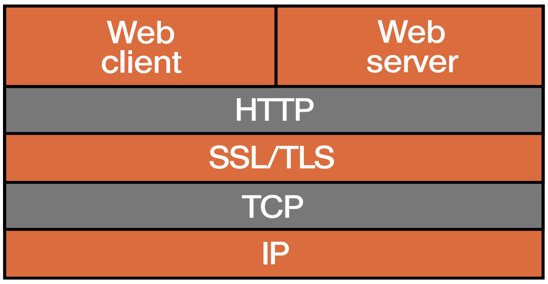

Figure 2. Network protocol layers showing position of encryption layer (SSL/TLS).

SSL/TLS

The Secure Sockets Layer (SSL 3.0) protocol and the Transport Layer Security (TLS 1.0) protocol (the successor to SSL 3.0) provide security by encrypting communications between a Web client (such as a browser) and a Web server. The encryption layer sits between the TCP protocol and HTTP protocol as shown in the diagram in Figure 2.

The SSL/TLS protocols provide for the negotiation of a secure session between an arbitrary Web client and Web server by using public key techniques and digital certificates. Digital certificates associated with a server can be used by the client to verify or authenticate the server's identity. The negotiation for the secure session begins by negotiating a certificate exchange and cipher suite between the server and the client. A cipher suite is a combination of a key exchange protocol, a bulk encryption algorithm and key size, and a hash algorithm that is used to ensure data integrity. While many different cipher suites can be negotiated, the most widely used suite by standard Web browsers is a suite that uses the RSA Key Exchange and RC4 algorithms. During the negotiation process, asymmetric encryption is used and a session-specific symmetric key is created and exchanged between the server and the client. Once the negotiations are complete, symmetric encryption (which is faster) is used for the rest of the session. Within a single secure SSL/TLS session, multiple Web pages, forms, etc. can be transmitted so that the negotiation overhead is encountered only once at the beginning of a secure transaction.

Embedded Internet Encryption Software

One of the few toolkits available for device manufacturers is provided by Allegro Software Development. Their RomPager Secure product is based on the Allegro RomPager Advanced embedded Web server and adds SSL/TLS protocol capabilities. Allegro provides digital certificate services so each vendor-built device using RomPager Secure has a unique digital certificate. Any standard Web browser that includes SSL 3.0 or TLS 1.0 support can conduct a secure session with a RomPager Secure device to manage and control the device.

Allegro Software Development also provides a Web client version based on their embedded RomWebClient product. The RomWebClient Secure product includes SSL/TLS protocol and certificate capabilities so a secure HTTP-client session can be conducted with any standard Web server.

CONCLUSION

Too often security is an afterthought or thought of as an impediment to product development. As a result, security is tacked on at the end without sufficient thought to usability, utility, and functionality. By planning upfront to build a secure system, you will spend less money for more results. If you do not have in-house security experts, it is far better to bring in outside security experts early in the design process, than to wait and bring them in after disaster has struck.

URLS

There are a number of Web sites with information to assist a network designer. Some are commercial but most are sites where the information is given freely by the network community.

- Allegro Software Development Corp., a source for embedded Internet/Web/security products.
    - http://www.allegrosoft.com
- CERT Coordination Center is the major reporting center for Internet security problems, the location for security advisories, and a source of many other security resources and documents.
    - http://www.cert.org
- Firewall email list, archives, and other firewall resources.
    - http://lists.gnac.net/firewalls/
- Lance Spitz' security white papers focus on security forensics.
    - http://www.enteract.com/~lspitz/pubs.html
- Lincoln Stein's The World Wide Web Security FAQ is a good starting point.
    - http://www.w3.org/Security/Faq/www-security-faq.html
- MIT Auto-ID center, a project to connect the real world to the network.
    - http://auto-id.mit.edu/
- Rootshell whitepapers include many security classics.
    - http://rootshell.com/beta/documentation.html
- Security Focus website and email lists.
    - http://www.securityfocus.com/
- Whitehat free security tools and resources.
    - http://www.whitehats.com/ ■

Patricia Hawkins (phawkins@connact.com) is the principal of Hawkins Internet Applications LLC, specializing in the design and implementation of Internet-based systems. Her customers have included Ziplink, Adero, and Inso Systems. She has worked in software design and development since 1984, in a broad range of industries including CAD/CAM, electronic printing, and object-oriented web-enabled databases. She has collaborated with Internet Guide Service on a number of projects.

Eric Johansson (esj@inguide.com) has over 20 years of high-level system and software design experience with particular emphasis on Internet system and security design. For the past five years, Eric has headed Internet Guide Services, specializing in the design, configuration, and remediation of complex Internet-based systems. Among others, his clients have included EG&G, BBN, AllMedia Solutions, ZipLink, and Harvard Pilgrim Health Care. Prior to founding Internet Guide Service, Eric held senior-level engineering positions with Polaroid Corp., Wang Laboratories, Ziff-Davis, and Computervision.

Edward Steinfeld (edward@go-embedded.com) has more than 25 years experience in realtime and embedded computing. He began as a programmer writing code and designing hardware to test hybrid circuit boards for Picker X-ray. He has marketed embedded and realtime products to OEMs and resellers for Digital Equipment Corporation, VenturCom, Inc., and Phar Lap Software. His international experience includes a stint in Hong Kong as a Far East Channels Manager and responsibility for international OEM sales in Europe and the Pacific Rim. Ed is now providing market research, business planning, and marketing services to the embedded computing industry.