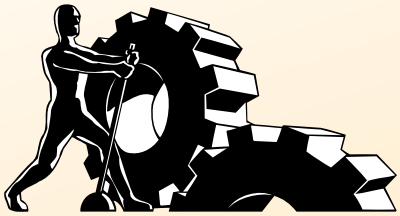
Check for Updates
Make sure you have the latest information!



TidBITS Publishing Inc.

Take Control of Your v1.7

Wi-Fi Security

Adam Engst &
Glenn Fleishman

$10

Help Catalog Feedback Blog Order Print Copy

Table of Contents

READ ME FIRST 4

Updates and More 4
Basics 5
What's New in Version 1.7 6
What Was New in Version 1.6 6

INTRODUCTION 7

WI-FI SECURITY QUICK START 9

DETERMINE YOUR SECURITY RISK 11

Evaluate the Likelihood of Attack 12
Determine Your Liability 15
Calculate Lost Opportunity 24
What You Should Do 25

PREVENT ACCESS TO YOUR WIRELESS NETWORK 27

Use Secure Settings 27
Ignore These Sops to Security 30
Watch out for WEP Encryption 32
Use Wi-Fi Protected Access (WPA or WPA2) 35
Enable Guest Access 44

SECURE YOUR DATA IN TRANSIT 46

Encrypt Email Passwords 47
Encrypt Specific Files and Messages 49
Encrypt Sessions and Data Sequences with SSL/TLSwith SSL/TLS 55
Encrypt Data Streams with SSH 60
Encrypt All Data with a VPN 66

PROTECT YOUR SYSTEMS 68

Get Paranoid 68
Install Antivirus Software 71
Assign Private Addresses for Passive Protection 73
Enable an Active Firewall 75

SECURE SMALL OFFICE WI-FI 80

Three Security Options 80
Use a Shared Key 81

Use WPA2 Enterprise Logins 82
Use a VPN 88

GLOSSARY 93

APPENDIX A: PASSWORD ADVICE 101

Generate Three Passwords 101
Learn to Create a Highly Secure Password 103

ABOUT THIS BOOK 105

Ebook Extras 105
About Glenn 105
About Adam 106
Authors' Acknowledgments 106
Shameless Plugs 106
About the Publisher 107
Production Credits 107

COPYRIGHT AND FINE PRINT 108

FEATURED TITLES 109

Read Me First

Welcome to *Take Control of Your Wi-Fi Security*, version 1.7, published in November 2010 by TidBITS Publishing Inc. This book was written by Glenn Fleishman and Adam C. Engst, and it was edited by Tonya Engst.

This book is devoted to helping you most effectively secure your home and office wireless network under Mac OS X and Windows using common networking hardware.

Copyright © 2010, Glenn Fleishman and TidBITS Publishing Inc. All rights reserved.

If you have an ebook version of this title, please note that if you want to share it with a friend, we ask that you do so as you would a physical book: "lend" it for a quick look, but ask your friend to buy a new copy to read it more carefully or to keep it for reference.

Discounted classroom and Mac user group copies are also available.

UPDATES AND MORE

You can access extras related to this book on the Web (use the link in Ebook Extras, near the end of the book; it's available only to purchasers). On the ebook's Take Control Extras page, you can:

- Download any available new version of the ebook for free, or purchase any subsequent edition at a discount.
- Download various formats, including PDF and—usually—EPUB and Mobipocket. (Learn about reading this ebook on handheld devices at http://www.takecontrolbooks.com/device-advice.)
- Read postings to the ebook's blog. These may include new information and tips, as well as links to author interviews. At the top of the blog, you can also see any update plans for the ebook.
- Get a discount when you order a print copy of the ebook.

BASICS

Here are a few "rules of the road" that will help you read this book:

- **Path syntax:** We occasionally use a *path* to show the location of a file or folder in your file system. For example, Mac OS X stores most utilities, such as Terminal, in the Utilities folder. The path to Terminal is: /Applications/Utilities/Terminal.

  The slash at the start of the path tells you to start from the root level of the disk. You will also encounter paths that begin with ~ (tilde), which is a shortcut for any user's home directory. For example, if a person with the user name joe wants to install fonts that only he can access, he would install them in his ~/Library/Fonts folder, which is just another way of writing /Users/joe/Library/Fonts.

- **Menus:** When we describe choosing a command from a menu in the menu bar, we use an abbreviated description. For example, the abbreviated description for the menu command that creates a new 802.1X connection in Internet Connect is "File > New 802.1X Connection."

- **Wi-Fi, AirPort, and wireless networking:** *Wi-Fi* is an industry term that encompasses four short-range, unlicensed, radio technologies: 802.11n, 802.11g, 802.11b, and 802.11a. Apple calls 802.11b "AirPort" and 802.11g and 802.11n "AirPort Extreme." Regardless of the term, it's all wireless networking. For many more definitions and a further explanation of how the standards work, read *Take Control of Your 802.11n AirPort Network* (focused on newer 802.11n gear from Apple). For more details, see the Glossary.

- **Adapters and gateways:** A standard wireless network has two distinct components: a wireless network adapter (or wireless card) and a wireless gateway (or wireless router). The *wireless network adapter* is attached to or inserted into a computer and connects to a *wireless gateway,* which in turn manages the entire wireless network and shares your Internet connection.

WHAT'S NEW IN VERSION 1.7

This update includes the following changes:

- The ebook still covers Mac OS X 10.5 Leopard, but minor changes for 10.6 Snow Leopard are now incorporated. It also still refers to older versions of Windows, but some bits of information about Windows now include Windows 7.
- The ebook is updated to be aware that iPhone OS is now iOS and that the iPad has joined Apple's line of mobile devices.
- Determine Your Liability was updated to discuss the latest state of affairs, with numerous small changes.
- In Use Wi-Fi Protected Access (WPA or WPA2), we've added clarifying details and updated the discussion generally, plus we've made several revisions to move our emphasis to the newer WPA2 security standard.
- The topic "Share via Devicescape" is now called Share via Easy WiFi, and it now describes the latest details on this Web service that helps you manage and share access to your Wi-Fi hotspot or network.
- Overall, we've added a few snippets about recent security flaws in Wi-Fi and related protocols, and we've moved the discussions forward to reflect a late-2010 point of view.

WHAT WAS NEW IN VERSION 1.6

We created this new version to update the ebook in a variety of areas:

- The ebook is now fully updated for Mac OS X 10.5 Leopard.
- You'll find an updated discussion of Wi-Fi Protected Setup (WPS), a simpler way of securing Wi-Fi networks. See Use Wi-Fi Protected Access (WPA or WPA2).
- The ebook covers how to Use Apple's guest networking.
- Secured Web sites now discusses Enhanced Validation Web sites, something you may be interested in knowing about next time you do online banking.

Introduction

Just because you're paranoid doesn't mean they're not out to get you.

–Internet security saying

Networking wasn't supposed to be like this. When computer networks were invented, no one anticipated hundreds of millions of naïve users. Nor did they expect crackers, viruses, worms, spam, or spyware. But that's where we've ended up. Most people are clueless about security, and few people devote any time to making their systems secure.

The biggest security risk comes from the fact that computers are all networked these days: to each other and to the Internet. Want a totally secure computer? Make sure it isn't connected to the Internet, or to any other computer, and put it in a locked room with an armed guard checking identification on those who enter. Not very useful, eh?

Wireless networking, because it makes connecting computers so simple, makes proper security even more critical. Before wireless networking, you could rely on a locked door to restrict access to your Ethernet jacks, and thus to your network. But now, transmissions over wireless networks—because they go through locked doors, along with walls, ceilings, floors, and other obstructions—are easily intercepted by consumer-level equipment just like the gear you use to connect your computers and access point.

So anyone in range of your wireless network can connect to it, and, unless you've taken appropriate precautions, wreak all sorts of havoc. And, unfortunately, understanding the reality of wireless security is nowhere near as simple as setting up a wireless network to start.

Our goal in *Take Control of Your Wi-Fi Security* is to bring clarity to the topic; to help you decide how worried you should be about security problems; and to help you to lock down your network, protect your data in transit, and secure your systems against attack.

Before we get started, we want to mention a few important caveats:

- We're writing this book for individual users with wireless networks at home and for people who run small to medium-sized office networks (from 2 to 50 people), not for veteran network administrators who manage large institutional networks.
- Security, whether you're talking about protecting your car, your home, or your wireless network, is hard, mostly because it's a battle with another human being. Locking your door with a simple knob lock stops amateur thieves, but keeping more experienced thieves out requires a strong deadbolt. And if you live where burglary is likely, or if you have especially valuable property, you have to think about if multiple locks, alarm systems, or bars on the windows are also necessary. Unfortunately, the kind of people who break into networks are usually much smarter than garden-variety thieves, and as a result, the security measures you must take are commensurately more complicated. So, our apologies up front, but some sections of this book are inherently quite technical.
- Because every network uses different hardware, software, and configurations, we can't give exact, foolproof, step-by-step instructions for every task we explain. That said, by the time you finish reading this book, you should have the background necessary to configure the networking hardware and software you do have (or are willing to purchase) to the level of security you want to achieve.

We've been using and writing about networking for more than 40 years combined, and we've both set up and maintained numerous wired and wireless networks over that time. And over those years of networking computers together, we've experienced the seedier side of the industry: attacks on our networks via the Internet, password thefts, wireless snoopers, and more. We've shared our experience in many articles and public presentations, and now we look forward to sharing it with you.

Tip: For more about wireless networking, check out these other ebooks written by Glenn: *Take Control of Your 802.11n AirPort Network*, *Take Control of iPad Networking & Security*, and *Take Control of iPhone and iPod touch Networking & Security, iOS 4 Edition*.

Wi-Fi Security Quick Start

You can read this title in the order shown here, or you can click a link to jump to a topic immediately. That said, if you're new to the topic of security, we encourage you to read Determine Your Security Risk first to get a sense of how concerned you should be about security.

Determine how worried you should be about security:

- Learn about the three Ls of security: likelihood of attack, liability in the event of loss, and lost opportunity. See Determine Your Security Risk.
- Figure out where you stand on the continuum of people who should be concerned about security. See What You Should Do.

Lock down your wireless network:

- Take care of three basic security measures with the configuration of your Wi-Fi router. Read Use Secure Settings.
- Discover which widely used security mechanisms won't prevent determined attackers. See Ignore These Sops to Security and Watch out for WEP Encryption.
- Turn on wireless security that is guaranteed to keep intruders out. Find directions in Use Wi-Fi Protected Access (WPA or WPA2) and Simplify with Wi-Fi Protected Setup. Also, be sure to read Appendix A: Password Advice.
- Consider setting up special access features for guests; see Use Apple's guest networking and Share via Devicescape.
- If you need to protect more than just a home computer or two, be sure to read Secure Small Office Wi-Fi for additional details.

Protect your data in transit:

- Keep miscreants from discovering your passwords and reading your communications. Consult Encrypt Email Passwords and Encrypt Specific Files and Messages.

- Armor your Internet sessions inside protected tunnels to keep snoopers from listening to your traffic. Read Encrypt Sessions and Data Sequences with SSL/TLS, Encrypt Data Streams with SSH, and Encrypt All Data with a VPN.

Secure your computers:

- Protect Your Systems from viruses, spyware, and crackers.

Determine Your Security Risk

Security is something we tolerate, not embrace. Your comfort level with security may vary enormously depending on your background and location. Growing up in rural New York State in the early 1980s, Adam left his car keys in his elderly Dodge Colt when it was parked at home. No one lived within a mile; cars driving by were infrequent, easily seen, and usually announced by the family dog; and a rusty Dodge Colt wasn't worth much.

Living in a populous suburb of Seattle a decade later, Adam not only didn't leave his keys in his shiny, red Honda Civic when it was parked in the driveway, he also locked the doors. Adam's behavior changed—more paranoid or more realistic, take your pick—because of a different evaluation of the three Ls of security: likelihood, liability, and lost opportunity. You can get a better idea of where you stand in terms of likelihood, liability, and lost opportunity by answering these questions:

- **Likelihood:** How likely is it that someone will break into your wireless network or *sniff* (monitor) the traffic going across your wireless connection? (See Evaluate the Likelihood of Attack, next page, for more info.)
- **Liability:** What is the potential liability if someone breaks in to your network, either to monitor your traffic or to use your connection for other purposes, including illegal ones? (Read Determine Your Liability, ahead, for details.)
- **Lost opportunity:** How much money and effort are you willing to expend on the security of your wireless network? (See Calculate Lost Opportunity, further ahead, for details.)

In the rest of this chapter, we help you answer those questions. We don't want to turn you into a tic-ridden paranoiac. Instead, we want to present a fair discussion of the risks and potential outcomes when you rely on wireless networks.

Note: Adam wrote an article for *TidBITS* that discusses a theory of privacy and why most people don't give privacy much attention. Read the article at http://db.tidbits.com/article/5951.

EVALUATE THE LIKELIHOOD OF ATTACK

When thinking about the likelihood of attack, consider two variables: your location and the desirability of the item you're trying to secure. In rural New York, at the end of a dirt road, the likelihood of someone stealing an old, rusty car was low, both because of the remote location and because no one wanted the car anyway. Flash forward 10 years to when Adam had an attractive new car in a busy Seattle neighborhood with lots of strangers driving by, and the likelihood of theft rose significantly.

In terms of wireless networks, Adam lives far enough from the population center in Ithaca, New York, that he and his wife Tonya aren't worried about the potential of a snooper: it would be difficult for someone to access their wireless network without parking in their driveway. In contrast, Glenn and his wife Lynn reside in a moderately dense part of Seattle. He can pick up three to five networks from his living room. So, for Adam the likelihood of a security breach is low, whereas for Glenn it's moderately high. If your company is in an office building that holds other companies, the likelihood of someone accessing your network is probably very high.

Warning! *Newer wireless equipment using multiple antennas—all the new 802.11n gear included—can retrieve and transmit signals two to four times farther than older network gear. This change could increase the likelihood of your network being seen.*

Consider your location

First consider where you use wireless networks, because location is the primary variable when determining the likelihood that someone would try to connect to your network and snoop. It's likely that you use wireless networks in one or more of the locations in **Table 1** (next page). And when we say "use wireless networks," we're talking about either your own network or networks run by others, because when you access someone else's wireless network, you're still at some level of risk.

Table 1: Likelihood of Snooping in Different Locations		
Location	**Details**	**Likelihood of Snooping**
Rural/far away	In your home and far from other houses	Extremely low
Long-range	Over a long-range, point-to-point link with a wireless ISP or neighbor	Low, due to the directional nature of most point-to-point links
Dense urban or suburban	In your home in a dense urban area or with at least several other houses close by	Moderately high, particularly if you have high-tech neighbors, but actual attacks are unlikely
Mixed-use	In a mixed-use residential and commercial neighborhood	Moderately high, since businesses are more attractive targets and are more likely to use wireless networks
Metro Wi-Fi network	Receiving a Wi-Fi signal in your home from a city-wide network or neighborhood hot zone	Moderately high, because many users will be lulled into a false sense of security; or very low if the network operator builds in robust security, which some do
Public-space neighborhood	In a neighborhood near a public park or where people can park on the street	High, since community networks receive constant use by a diverse, anonymous population
Office building	In an office building with multiple businesses or a nearby parking lot within line of sight	Very high, due to proximity and the attractiveness of targets
Roaming	While on the road in airports, cafés, hotels, and other locations	Moderately high, due to ease of monitoring and likelihood of viruses on other computers on the same network

Unless you fall into the rural/far away category, or the less common long-range category, there's a non-trivial likelihood that someone

could access your unprotected wireless network or watch your traffic without your knowledge.

Pay special attention to the roaming category. Even if you protect your own network, using your computer while connected to untrusted networks can still put your data at risk. Whether the network is free or for-fee, you have no control over the network-based security precautions, and everyone else using the network may have the ability to see your data in transit on a wired or wireless link.

Determine the desirability of your data

Although location is the most important variable in evaluating the likelihood of attack, you can't ignore the desirability of your data.

If you're a home user who uses the wireless network mostly to connect to the Internet for browsing the Web and sending and receiving email, your data is, and pardon our bluntness, quite dull. It's possible that someone sifting through your network traffic could pull out your credit card or bank account number, if you use online banking, but it's quite unlikely since we know of no banking Web sites that fail to use SSL/TLS to protect the contents of each transaction. The most likely concern if you're a home user is that your passwords could be captured and used to break into other machines or to send spam through your connection. This is, in fact, happening more these days; for details, read Adam's *TidBITS* article "Change Your Passwords: Email Account Hacking on the Rise," at http://db.tidbits.com/article/11376.

On the other hand, if you run a small business and you frequently transmit customer credit card numbers between computers on your wireless network (perhaps between databases, or even as part of a backup solution), the desirability of your data is significantly higher than that of a home user's data. A snooper could steal hundreds, if not thousands, of credit card numbers fairly quickly, which is a much better haul than trying to sniff a single home user's credit card number. A ring of thieves stole as many as 200 million credit card numbers from the parent company of clothing retailer TJ Maxx from 2005 to 2006 via a weakness exposed from a wireless network.

While the small business scenario is realistic, think of the desirability (to the right customer) of classified government information. In 2005, when Adam gave a presentation at Los Alamos National Labs (a government research organization involved largely with national security),

he learned that wireless networks are entirely forbidden at Los Alamos because the value of their data is too high to risk the possibility of snooping.

DETERMINE YOUR LIABILITY

Though Adam's income as a teenager was low, the potential liability if the old car had been irretrievably stolen was also low—the car wasn't worth much, and he had alternate transportation from his parents and the school bus. In Seattle, however, not only was the Honda Civic worth a great deal more, but it was also his only form of automotive transportation, and being forced to rely on sketchy public transit for grocery shopping and the like would have been a huge hassle.

When thinking about the liability of having a wireless network cracked, a business might be concerned about whether someone in its parking lot could access its confidential data, such as invoices as they pass between employee computers and a central database, email containing information that might interest competitors, or even customer credit card numbers that might be encrypted over a link between a company Web server and a customer's browser, but totally unprotected on a local network. The liability of having that data stolen could be the bankruptcy of the company.

Warning! *A revision to credit-card security standards for all merchants requires that equipment that processes cards use newer wireless security if any Wi-Fi networks are involved, as well as a laundry list of other requirements. If you're handling cards and unaware of these requirements, contact your merchant bank for details. If you're found out of compliance, you could lose your ability to process cards, and have to pay huge fines to the credit-card issuers.*

If you work as a sole proprietor or for a company in a number of fields related to bank and credit-card information, health care, and a few other industries, laws in the United States, the European Union, and elsewhere may provide for financial penalties and criminal prosecution for failing to evaluate the liability of your network and take sufficient action. These laws can affect a home worker processing medical claims as a contractor and the world's largest health maintenance organizations alike.

Home users have fewer worries, of course, but do face a potential financial liability should an attacker steal passwords used for logging in to an online banking account or bill-paying system, for instance.

But liability isn't just about the theft of data, and the concerns break down into three categories:

- Access liability (see below): What happens if someone uses my wireless network to share my Internet connection?
- Network traffic liability: What happens if someone is able to eavesdrop on my wireless network traffic?
- Computer intrusion liability: What happens if someone on my wireless network breaks into my computer?

Access liability

Do you want to give unknown people access to your wireless network?

This question doesn't come up with wired networks, since no one installs Ethernet jacks on the outside of a house or office. Since many people believe strongly in the sharing ethic, asking this question isn't unreasonable. Although intentionally allowing people to connect to your network does increase the risk that your connection will be used to bad ends, the fact that you're aware of the possible presence of outsiders on your network means that you're probably also more aware of the security considerations that their presence engenders.

Tip: The guest network feature in Apple's March-2009-and-later base station models allows you to provide visitors with an open network or a password-protected network that's separate from your own network. It's a great way to provide Internet access without risking exposure of your data or computers. See Enable Guest Access.

A few problems can arise whenever someone accesses your wireless network for friendly, or at least benign, purposes:

- **Performance:** If you have a cable- or DSL-based Internet connection, performance isn't likely to be a concern most of the time, assuming the unknown visitor isn't uploading or downloading vast quantities of data. However, you may be unintentionally violating your ISP's terms of service or acceptable use policy by

sharing your connection. In the worst-case scenario, your Internet connection could be shut off for that violation.

- **Cost:** If you pay for traffic on your Internet connection, letting unknown people share your Internet connection could result in a nasty and unexpected bill. In the United States, some broadband providers offer unlimited use, but many charge or shut off your access when more than a certain number of gigabytes of data per month are transferred. Outside the United States, it's more likely that an ISP caps all traffic on a daily or monthly basis.

  The MiFi and other cellular routers that share Internet access from a mobile broadband network over Wi-Fi are now common. However, most cell companies meter usage, capping 3G usage at 5 GB per month, and then charge for excess use or throttle your service if you exceed that amount. With an account for which you pay overages (often $50 per GB), you could run up a high bill by accident. (Virgin Mobile is the only service that doesn't throttle or meter. They have a $40 plan that offers unlimited data on Sprint Nextel's network for a 30-day period with no contract.)

- **Mischief and abuse:** Although we'd like to assume that anyone accessing an open wireless network would use it responsibly, the possibility for abuse does exist. An unknown visitor could use your wireless network to send spam or launch an Internet worm attack, for instance, and while neither would likely hurt you personally all that badly, your ISP might shut you down for being the source of the abuse. In India, in 2008, terrorists used open residential Wi-Fi networks to claim credit for bombings, leading to real problems for the innocent people operating those networks.

Your mission, then, is to determine how concerned you are about each of these possible access-related scenarios, after which, of course, you can read the rest of this book to address your concerns.

Network traffic liability

You likely believe that most of your private data sits on your Mac, that you transmit and receive only limited amounts of sensitive information, and that someone would have to listen at a specific time to capture those bits. The reality of the situation is that we all transmit and receive quite a lot of sensitive data that people with common

equipment and widely available software could extract easily from an unprotected network.

Most of the data sent or received over a wired or wireless network is transmitted *in the clear* to anyone able to join or plug into the network. In the clear means that the data is sent in a form that a human being can intercept and then either read directly or convert easily into usable data.

Here's a list of what you *might* be sending or receiving in the clear:

- Your email account password
- The text of all email messages sent and received, including via so-called *push* email, such as mail sent to an iPhone, BlackBerry, or other smartphones
- The contents of any documents sent or received as attachments
- Your Twitter direct messages, unless your Twitter client is set to work only via an encrypted connection
- The location and contents of any Web pages viewed
- The authentication tokens used for some Web services, social networking sites, and other sites that don't protect the token with SSL/TLS; for more details, see Glenn's *TidBITS* article "Sidejack Attack Jimmies Open Gmail, Other Services," at http://db.tidbits.com/article/9129. (Many sites affected by this vulnerability have started using SSL/TLS to block it.)
- Your user name and password for any non-secure Web sites (sites that don't use SSL/TLS)
- Your FTP user name and password when sent via plain FTP
- Files transmitted via FTP or FTP over SSH (but not SFTP or FTPS)
- Unprotected WebDAV file-transfer sessions
- The text of instant messages you send or receive, if an encrypted session isn't active
- The contents of any music or other files you send or receive using LimeWire, BitTorrent, or other peer-to-peer file sharing programs,

unless you've restricted transfers to only peers that have encryption enabled

- The IP addresses and port numbers of any connections you make
- The complete contents, including passwords, of telnet sessions but not SSH sessions
- Timbuktu remote control or file transfer sessions over a regular connection, or unprotected VNC remote control sessions (including those via Apple Remote Desktop)

These items are *never* sent in the clear:

- The contents of encrypted sessions—including email, WebDAV, terminal, Twitter, and Web transactions—using SSH, SCP, SSL/TLS, or a VPN (described in Secure Your Data in Transit)
- Your POP-based email account's password if your ISP uses APOP
- Web services tokens for email and other applications if the operator is smart enough to use SSL/TLS to protect Web cookies (Google now does with Gmail)
- Timbuktu Pro, VNC (including Apple Remote Desktop), or pcAnywhere passwords
- Timbuktu Pro or VNC sessions when using Remote Access (SSH)
- LogMeIn and GoToMyPC for Mac remote control sessions
- Files transferred using Dropbox
- Files and account information transmitted via SFTP (Secure FTP) or FTPS (FTP over SSL/TLS)
- Voice conversations and instant messages sent via Skype, which encrypts all traffic
- Instant messages sent using iChat and the Jabber server in Mac OS X Server, or among iChat users with MobileMe accounts
- Any secure SSL/TLS Web pages (their URLs begin with `https`)
- The contents of any email message or file encrypted with PGP or similar public-key encryption technology

Warning! *Even if you close your network through the means that we describe in Prevent Access to Your Wireless Network, you may still expose data to network crackers and others who can penetrate the basic methods of preventing access. We talk about securing the contents of what you're sending in Secure Your Data in Transit.*

Each item that you might transfer in the clear falls into one of three categories: account access information (user names and passwords), information that could be used to track your online steps, and content related to what you say and do:

- **Access information:** Most important is account-access information, which, when stolen, presents two types of risk:
  - First, since most people tend to use the same passwords in multiple places, having your email password stolen could compromise a more sensitive use of that password, such as logging in to your online banking account.
  - Second, attackers often use a password to one account to break into another account, working their way ever deeper into a computer with the eventual goal of stealing data, causing damage, or using the computer to run an automated program that attacks other computers. In this respect, protecting your passwords isn't something you do just for your own benefit, it's something you do for the benefit of everyone who may be affected if the attacker takes out a server that you use.
- **Online movements:** Similarly, information that tracks online movement doesn't worry either of us, since as journalists, we can always claim we visited a Web site for research purposes. (That might not work if it's a site we visited 700 times.) But it doesn't take much imagination to see how the fact that a politician had frequented certain sex sites could ruin his career, or how a political activist in a repressive country could be jailed for the online company he or she keeps. Again, you probably have a decent idea of whether your online movements could be in any way damaging.
- **Content:** We're pretty transparent people (well, not literally), so there isn't much that we would say or do online that we would worry about someone else reading. We might be embarrassed if the wrong person read the wrong document, but that's it. But what

if that document were posted on a widely read mailing list or Web site? Even for us, that could be a problem, and other people might have data that could get them fired, damage their businesses, humiliate them publicly, or cause lawsuits or divorces. You probably have a pretty good sense of whether or not you're at risk from the things you say or do.

In general, because anything you send or receive could be intercepted and read (text) or used (files and programs), you must accept the notion that everything could be examined or stolen if you're in a location where other people might be able to connect to the network you're using, or even the network your recipient is using.

So what's your liability? Obviously it depends on the data you're transferring, but no one wants their passwords in other people's hands, and we strongly encourage everyone to take some basic precautions that we outline in Secure Your Data in Transit. And if you're more concerned, that section also has solutions for even the most anxious.

Computer intrusion liability

The final form of liability you should consider when thinking about security for your wireless network is what happens if someone uses your wireless network to break into your computer.

Note: Protecting computers from intrusion via your wireless network isn't fundamentally different from protecting them from intrusion via your wired Internet connection. However, many anti-intrusion programs trust computers on the same local network more than computers on the rest of the Internet, and you must make sure your settings reflect your degree of risk.

We see several types of concerns here:

- **Data theft:** If someone can gain remote access to a computer and its files, she could easily steal sensitive files. All it takes is a few minutes of inattention, or a misconfigured setting, for someone to copy files from your computer. Glenn found this out back in 1994, when his Unix server's password file was stolen (but the passwords weren't cracked, at least) and in 2005 when a Brazilian cracking team almost gained access to one of his Linux servers via Web traffic analysis software—but instead, they just disabled its access to the Internet.

More recently, Adam was irritated with himself after his ISP asked if he knew that anyone could see files on one of his Macs via AppleShare. And an iPad owner at a conference discovered he could use GoodReader to browse via WebDAV the contents of folders being shared unintentionally by other iPad users running the Quickoffice app.

- **Data damage:** You may never know if someone has stolen files from your computer, but you'll certainly realize if he instead vandalized your system and deleted all your files. Worse, some attacks focus on more subtle destruction or manipulation that you wouldn't notice at all. If someone were to tweak Excel spreadsheets with hundreds of numbers in them, could you tell?

- **Sidejacking:** Sidejacking was first demonstrated in August 2007 at a hacking conference. You could be sidejacked if a Web site sends your browser an unsecured cookie containing a session token. The purpose of this token is to identify you to the Web server as you move from page to page. A sidejacking attack sniffs this unprotected cookie, and then inserts it in the attacker's browser so he can access the same site as you. Even sites that offered a secure login, such as Google, had not secured this cookie at that time.

  A sidejacker can't pull out a credit card number or even change your password (except on poorly designed sites), but he could insert malicious virus code on your social Web pages or email a virus to everyone in your address book.

  In October 2010, a Firefox extension called Firesheep was released. Firesheep automates sidejacking for popular services like Facebook and Twitter over open Wi-Fi networks, such as hotspots. What Firesheep enables is scary, and it should provoke even more rapid change. (The *TidBITS* article explaining the risk is at http://db.tidbits.com/article/11701.)

  Since 2007, many organizations moved toward offering SSL/TLS versions of their Web sites, which protects both your login, the contents of your browsing session, and the cookie used to identify you for a session. Google has been particularly aggressive, even offering Google Search with encryption (https://www.google.com/).

Tip: You can automatically surf encrypted versions of sites, even content sites like *The New York Times*, by downloading and installing the Firefox extension HTTPS Everywhere, developed by the Electronic Frontier Foundation and The TOR Project (https://www.eff.org/https-everywhere).

- **Exploitation:** Some attacks focus on known bugs in software that allow a remote program or person to infiltrate your computer and take control of some of your software or the entire operating system. Once the attacker has established that level of control, he can install software that acts on his or her command, turning your computer into what's called a *zombie*. Most attacks are aimed at Microsoft Windows or specific Windows software from Microsoft, such as Outlook or Internet Information Server. Over the years, many different bugs have been found that allow attackers to take over a machine; equally as problematic are worms and viruses that may cause damage, replicate themselves, turn the infected computer into a zombie, or all three.

  Zombie attacks against other computers may use a *distributed denial-of-service* (*DDoS*) approach, where the attacker tries to overwhelm a computer by sending it huge amounts of data from dozens or thousands of zombies. DDoS attacks prevent normal operation of a site or company, and could result in millions of dollars of lost revenue or damage. They can be difficult or even impossible to shut down. The country of Estonia faced a DDoS attack in spring 2007 after officials relocated a Soviet-era memorial statue with meaning to local ethnic Russians. Zombies are also highly involved in sending spam that sells actual products or that links to fraudulent activity. The latter including *phishing,* in which notes are sent that appear to come from banks or other financial institutions.

  Microsoft keeps patching known holes, but many Windows users don't download and install these security patches, leaving their computers open to further exploitation and infection.

Note: Glenn once spent a full day watching his network be saturation-bombed with garbage traffic before he could convince an ISP from whose network the attack was launched that he had a serious problem. Glenn finally, with informal advice from the FBI, suggested he might have to sue the ISP and mentioned the FBI; that apparently spurred the ISP to shut down the offending DSL customer (who was likely the victim of an attack that had turned his computer into a zombie).

The liability for each of these three scenarios—data theft, data damage, and exploitation—is fairly severe; but luckily it's easy to take simple precautions that significantly reduce the likelihood of anything bad happening. Protect Your Systems offers the necessary advice.

CALCULATE LOST OPPORTUNITY

We're stretching a little to find a third L. Lost opportunity, or opportunity cost, is the cost in time, money, and effort in achieving a security goal. The "opportunity"—time, money, and effort—is "lost" because if you spend it on security, you can't spend it on something else.

Obviously, Adam could have reduced the risk with his old car in rural New York even further by locking the doors and keeping the keys in the house, but it wasn't worth doing. The low likelihood and minimal liability made the effort unnecessary. Once he had a new car in an area where car theft was more likely, though, Adam was willing to expend more effort and money: locking the doors when the car was in the driveway, buying and using a brake lock when parking in a Park-and-Ride, and so on. More money and hassle would have provided even greater security—an electronic car alarm, for instance, or a house with a locking garage.

The same situation applies with wireless networking. You can use every available security technique (even the ones that are easily broken) and in doing so, you'll find that accessing your network is a royal pain. But it will be annoying for would-be attackers as well. Some effort likely makes sense, such as using a password that is relatively easy to type but not easily guessed, and organizations that need even more security can consider a variety of techniques

for keeping snoopers out and for protecting data traversing a network that work well, though at a price.

Obviously, the amount you're willing to spend on wireless network security relates directly to your liability. The higher your liability, the more you should be willing to spend, and the more effort you should be willing to expend.

Don't fall into the trap of assuming that a low likelihood of attack means that you can avoid spending time, money, and effort on your wireless network security. It's all about liability, and if the result of an attack could be highly damaging to you, reconsider your willingness to pay for security accordingly.

WHAT YOU SHOULD DO

Let's combine likelihood, liability, and lost opportunity for a number of sample users to evaluate your real-world risks and determine which sections of this book are most important for you to read:

- If you're a home user with no immediate neighbors or nearby public spaces, and if you don't believe your data is particularly sensitive, you don't have much to worry about. At most, read Protect Your Systems to see if you want to take steps to prevent anyone from attacking your computers over the Internet.
- If you're a home user and live in a state, province, or country in which you could be at fault in some manner for allowing an open network to be used, read Prevent Access to Your Wireless Network.
- If you're a home user in an urban environment, you should definitely read Prevent Access to Your Wireless Network and Encrypt Email Passwords. If you're concerned about the sensitivity of your data, read the rest of Secure Your Data in Transit. It's also worth reading Protect Your Systems just in case.
- If you maintain a wireless network in a business, you should read this entire book, thinking hard about your company's risk factors as you go. In particular, in Secure Your Data in Transit, consider how far you want to go to protect your organization's sensitive data. Also important is Protect Your Systems because your data is probably more attractive to electronic thieves than the data of a home user.

Lastly, be sure to read Secure Small Office Wi-Fi carefully, since it contains information specific to small office needs.

- If you regularly use wireless networks while traveling, be sure to read Secure Your Data in Transit. The more sensitive your data, the more seriously you should consider the approaches in that chapter. The introduction of sidejacking means you need to be aware of whether you're connected securely to a Web site for email or other purposes.

Prevent Access to Your Wireless Network

Wireless networks weren't originally designed to be very secure. The only encryption available initially, Wired Equivalent Privacy (WEP), was supposed to work as well as those locks you find on old bathroom doors that can be picked with a paperclip. The designers assumed most people wouldn't have the interest in getting in. When Wi-Fi became popular, so did cracking techniques and tools, busting WEP's never-strong encryption. Further, most people buying Wi-Fi after the first wave weren't early adopter geek tech-heads. So security options, when available, weren't turned on.

As cracks and flaws evolved, so did replacement technologies: Wi-Fi Protected Access (WPA and WPA2), WEP's replacements. You can now reliably secure home and small business networks without much fuss. Even small businesses now can achieve corporate-level security without much cost or complexity.

In this chapter we first look at three easy things you can do immediately to enhance your network's security. We then look at common mistakes and techniques that don't provide any real security, and we run through how to secure your network with assurance.

We also offer a few suggestions on enabling safe access for visitors.

USE SECURE SETTINGS

The first task in preventing access to your network is changing three default settings that—when left as they come from the factory—make cracking significantly easier. Connect to your wireless gateway with its management software and then:

- Change the admin password to one that's not obvious but that you'll remember.
- Verify that remote administration from outside the network is off.
- Change the network name.

Everyone Is Looking for Security

One of the most visited Web pages on Glenn's Wi-Fi Networking News site is "Change Your Linksys WRT54G Admin Password Right Now!" from June 2, 2004—even in 2010 (http://wifinetnews.com/archives/2004/06/change_your_linksys_wrt54g_admin_password_right_now.html). That's partly because searching on Google for "linksys wrt54g password" brings up that page (the WRT54G is one of the best-selling Wi-Fi routers of all time). Linksys's current generation of configuration software tries hard to convince you to change the administrative password.

Of these three simple configuration changes, changing the admin password is by far the most important. If you fail to do that, anyone who could connect to your network could also guess the default password (it's usually admin or public) and then gain control of your wireless gateway. Changing the password also prevents someone from plugging into the gateway or your Ethernet network and reconfiguring the device. You wouldn't be locked out forever; a factory reset would blow away any settings changes an attacker made, but you don't want to end up in that situation.

Verifying that remote network administration is turned off is also important. This setting is disabled on most routers when you first unbox them; this is a big change from just a few years ago. If remote administration is enabled, anyone on the Internet can attempt to access the management interface for your wireless gateway (and if you've failed to change the admin password as well, you're doubly at risk). There are legitimate reasons to enable remote administration of a wireless gateway over the Internet—for example, you might be in charge of configuring gateways at several remote locations—but unless you have a reason, keep that option turned off.

Note: We've found that router makers have become much smarter about remote access, presumably because of previous flaws or concerns. The latest routers we've looked at usually offer several options for setting which features can be changed remotely. A few let you use an SSL/TLS connection to make sure a password isn't sniffed while you're administering the router.

Finally, changing the network name is always worthwhile, partly because it says to any would-be crackers that you know enough to do it (many people don't) and thus your network is more likely to be properly secured. But the main reason to change the network name is because the WPA2 encryption method uses the network name when generating encryption keys. Because so many people fail to change their network names, keys generated with those default names are significantly weaker and more prone to being cracked than keys generated for networks with unusual names.

Kudos to Apple

Apple's Wi-Fi base stations create networks that include the base station's unique MAC (Media Access Control) address in the network name by default. Not only does that ensure that every AirPort base station will create a differently named network by default, but also the hexadecimal MAC address is sufficiently ugly that users are more likely to change it than they are with other defaults, like `linksys` (Linksys gateways) or `tsunami` (old Cisco gateways).

Warning! *Another reason to change the network name is to prevent a device you use on your network from automatically joining another network with the same name, which may be undesirable. Most operating systems—including Apple's iOS—remember the names of networks you've said you want to join in the future, and thus attempt to join them whenever they see them. (The exception is that if you join a network that uses encryption, most operating systems won't automatically join a same-named network with different encryption settings or keys.)*

Now let's look at two techniques you may have heard a lot about—closing your network and restricting access by MAC address. These techniques used to be commonly recommended, but you should be aware that they are almost completely useless against anyone with a few simple tools. After that, we turn our attention to a pair of encryption methods: the older and completely broken WEP, and the newer WPA and WPA2, which offer real security.

IGNORE THESE SOPS TO SECURITY

To increase security in the real world, a person might remove a street number from a house or take the company name off the front door. Others might put up large "No Trespassing!" signs. These approaches don't prevent burglars from breaking in, and they're analogous to several common approaches used to secure a wireless network.

We shouldn't be entirely negative, though, since you may wish to discourage casual access to a network without worrying about preventing a serious attack. Closing your network and restricting access by MAC address may be futile for foiling a determined attacker, but they will prevent a casual passerby from sharing your Internet connection.

Don't bother closing your network

If a network is "open," its wireless access point continuously broadcasts the network's name, helping wireless adapters find the network. Most access points offer an option to turn this off; your access point may allow you to "close" the network or "disable broadcast name." No matter what the terminology, a closed network's name won't appear in the list of available networks in ordinary client software.

Don't be lulled into a false sense of complacency. Although a closed network offers protection from the casual observer, many sniffer programs that monitor wireless networks—from commercial down to open-source freeware—can easily see the name of a closed network whenever a legitimate user connects to it (**Figure 1**). If no one ever connects, it remains hidden, but that's hardly useful.

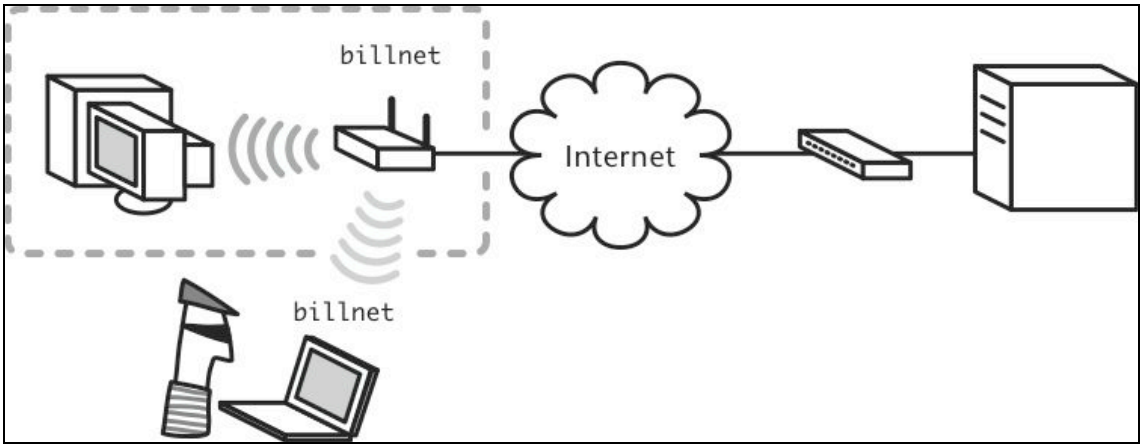


Figure 1: Closing your network prevents casual users from seeing the network name. But with only a little effort, a snooper can determine the name.

In short, if you don't want average people connecting to your network, there's nothing wrong (other than extra hassle) with closing it, but the only people you're keeping out are those who almost certainly weren't a security risk anyway. That may be worthwhile to you, if your main goal is to prevent non-savvy passersby from sharing your Internet connection, but you shouldn't consider it real security.

Ignore MAC address access controls

Another mostly useless way to restrict access to a network is to allow only specific network adapters to connect (**Figure 2**). Like all Ethernet network adapters, a Wi-Fi adapter is identified by its *MAC (Media Access Control) address,* a unique serial number assigned to every network adapter.

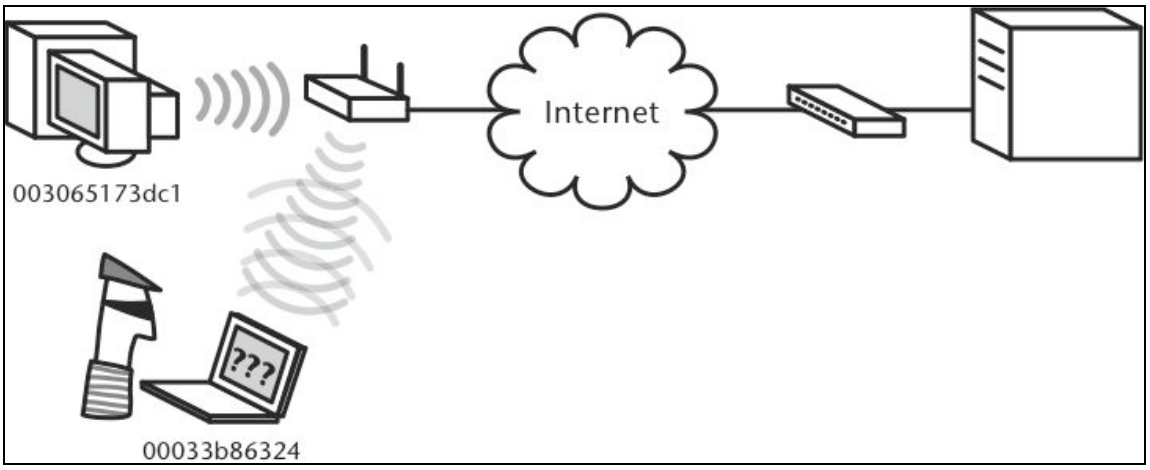


Figure 2: Restricting access by MAC address helps keep unauthorized computers out of your wireless network. The attacker's computer isn't allowed to connect because its MAC address isn't authorized.

However, MAC addresses can be easily *spoofed* (faked to seem like a different address); for more information, see the Wikipedia entry on MAC spoofing at http://en.wikipedia.org/wiki/MAC_spoofing.

This flexibility, combined with the fact that MAC addresses are sent in the clear even on encrypted networks, means a cracker can easily see MAC addresses in use and then assign one of those addresses to her equipment.

As with a closed network, restricting access by MAC address will keep honest people honest, but it won't do squat against a determined intruder. Worse, if you restrict access to specific MAC addresses, you'll

find it annoying to allow a visitor to access your network, since you'll have to enter his laptop's MAC address into your Wi-Fi gateway manually and then reboot the gateway.

A simplified security setup system called Wi-Fi Protected Setup (WPS) improves how MAC restriction works by allowing an individual computer to request to join a network, and then assigning that computer a WPA encryption key. WPS uses MAC addresses just for identity, but without the key that's exchanged on that first connection, a cracker can't gain access with the MAC address alone. We write more about WPS in Simplify with Wi-Fi Protected Setup.

Enough with those sops! Let's move on to other options. While closed networks and adapter address limitations don't do much good, there is hope. Instead of trying to hide or restrict access, you can use technology that requires users to enter a password to join a network; that same password is used to scramble all the data passing over the network. Without the password, no one can connect to the network or intercept the data.

WATCH OUT FOR WEP ENCRYPTION

WEP (Wired Equivalent Privacy) was the only way to secure a Wi-Fi network from 1999 to 2003, but it is now considered so broken as to be entirely unreliable. (Don't worry, there is another option, which we cover ahead shortly in Use Wi-Fi Protected Access (WPA or WPA2); you can skip there now if you don't want to learn about WEP.)

The developers of 802.11b, released back in 1999, intended WEP to offer an equivalent level of privacy to what could be found on a standard wired network. To compromise a wired network, an attacker generally needs to break in to a room and install a network-sniffing program that watches traffic traveling over the wire.

WEP was designed to act merely as a locked door, to keep intruders from penetrating to the wireless network traffic itself; other measures were supposed to bolster this initial line of defense. WEP basically encrypts all the data that flows over a wireless network, preventing attackers from eavesdropping on network traffic (**Figure 3**).

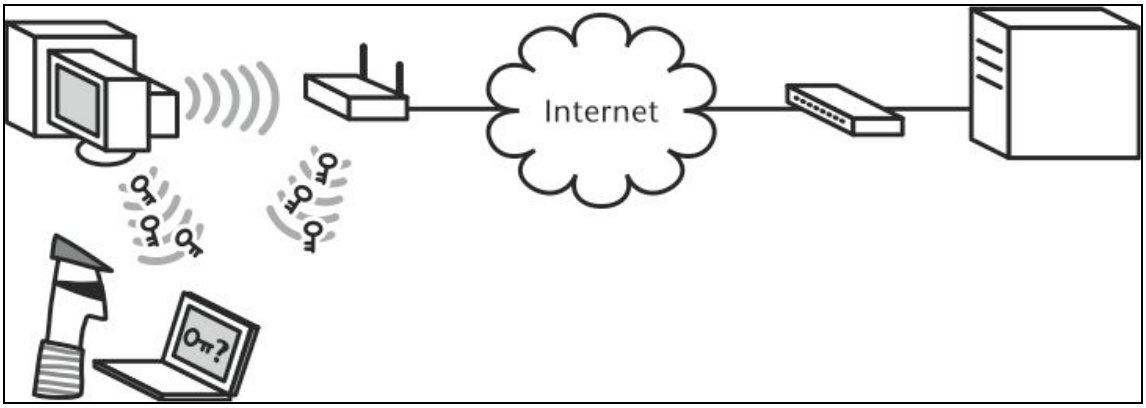


Figure 3: Turning on WEP prevents attackers from eavesdropping on network traffic. The attacker can't understand WEP-encrypted traffic without the appropriate WEP key.

Unfortunately, even this relatively minimal protection was crippled due to several poor design decisions made on the cryptographic front—some due to the political hot potato that was encryption in the United States in the mid-1990s, some just bad in retrospect—and because some options were built in but never enabled by most or all manufacturers. Also, although WEP still offers some protection, most people don't turn WEP on because it's a pain to use.

WEP is considered *broken,* which in cryptography means it is so trivial to extract a supposedly secret key that there's effectively no protection gained from using it. In 2008, researchers perfected earlier techniques and reduced WEP cracking to under a minute in some cases.

WEP works by using a "shared secret": one to four encryption keys per network shared by everyone on the network. Your wireless network adapter uses the encryption key to encode all traffic before it leaves your computer. Then, when the data arrives, the access point uses the key to decode it into its original form.

While shared keys are used by many systems, including the latest and greatest options to protect Wi-Fi, they're effective only when two conditions are met:

- Ease of use: They're easy to remember and change, making it more likely that they won't be written down or left in place too long.
- Effective key security: They provide real security.

WEP fails on both fronts, as we explain next.

Note: Apple's Mac OS 8.6, 9.x, and X (all versions) allow only a single WEP key to be used on a network instead of up to four, and use something called "Shared" instead of "Open authentication," which renders WEP even more susceptible to cracking. Apple's WPA implementation, discussed later, is just fine.

Ease of use

WEP keys are hard to create, enter, and change, compared to most other password systems. Users or administrators must enter the WEP key manually (and tediously) on each computer they want to connect to a WEP-protected network. A few systems tried to ease this entry by storing the key on a removable USB flash drive or by creating and exchanging a key when you press a special button. Those options required you to have the right devices from a single maker.

Worse, with all but Apple's Wi-Fi hardware and Mac OS X, the WEP key is often expressed in the base-16 hexadecimal numbering system in which the letters A through F represent 10 to 15 as a single digit. Most users haven't the slightest idea of how to deal with hex (reasonably enough—that's what computers are for!). If you combine user confusion with the tedium of inventing and entering strings of hexadecimal numbers, you can see why WEP is annoying to use.

Note: Some larger businesses used a form of WEP that allowed a server to create a password and assign one uniquely to each computer connected to the network, as well as to sync changing the password every few minutes. That turned out to be simpler to use, but only slightly more effective in protecting networks.

You enable WEP in an access point by thinking up a sequence of 10 or 26 hexadecimal digits; some gateways will make up a sequence for you. These correspond to 40 or 104 bits of encryption. (There's an extra 24 bits that isn't part of the key, which is why WEP is often called 64-bit or 128-bit encryption.)

Some access points have a feature in which you type a passphrase of 5 or 13 letters and numbers and then the access point translates that into hexadecimal digits for you. All adapters and access points must use the same length of key on a single network. But this shortcut

makes a WEP key even more vulnerable to cracking by reducing the number of possible keys you might have entered.

Apple's method of securing WEP is unique: The company developed its own algorithm to turn a passphrase into hexadecimal numbers. This provided greater randomness, and thus better protection, than simply looking up the hexadecimal value for each letter as other systems employed. But that small improvement is no longer valuable.

Effective key security

The problem with using shared keys might have been mitigated if WEP provided strong encryption, but WEP's encryption weaknesses are only exacerbated by the difficulty in entering and changing keys. WEP suffers from a now well-understood problem relating to how encryption keys are made even harder to crack. WEP takes a number called an *IV* (initialization vector), a 24-bit sequence in WEP's case, that's combined with the encryption key before a packet is secured. IVs were supposed to increase the difficulty in cracking, but due to poor implementations and a problem with an underlying algorithm, WEP is relatively easy to crack.

Rather than use random IVs and prevent those numbers from being reused during any given series of packets, many WEP implementations used just one or two numbers (like the 24-bit equivalent of 0 or 1), reused IVs, or even used them in a predictable sequence. Further, some IVs turn out to be particularly *weak,* meaning that when they appear, their use reveals even more information about the key. These IV weaknesses coupled with other problems have made it a simple matter to crack the WEP key on any network. Fortunately, there's a replacement in wide use, which we discuss next.

USE WI-FI PROTECTED ACCESS (WPA OR WPA2)

With all WEP's flaws, you'd hope that some of the giant brains that develop wireless networking standards would close WEP's gaping security holes. And they did. Starting in 2003, successive waves of security improvements hit the street based on standards created by the IEEE engineering group. This group is responsible for many networking and computer standards.

Background and history

The IEEE developed 802.11i. The 802.11i standard replaces WEP with *TKIP* (Temporal Key Integrity Protocol), a method that fixes all WEP's faults and adds a second, government-grade encryption option called *AES-CCMP* (Advanced Encryption System Counter Mode CBC-MAC Protocol, whew!).

The IEEE doesn't test equipment, however; testing is done by the Wi-Fi Alliance, an industry trade group that uses laboratory testing to certify whether a given device conforms to standards that include and go beyond the IEEE's guidelines. In this case, the Wi-Fi Alliance has released two security standards: WPA and WPA2:

- *WPA* (Wi-Fi Protected Access) came out in late 2003 as an interim measure since the IEEE's group was taking so long to complete 802.11i—it's a consensus-voting based organization. WPA includes just TKIP support, enabling older devices to enhance security.
- WPA2 is based on the completed 802.11i standard. The IEEE approved 802.11i as a final standard in mid-2004. The Wi-Fi Alliance updated WPA to WPA2 by early 2005, changing to the AES-CCMP method and making a few minor changes, such as improved handoff for Wi-Fi phones as they move among multiple access points in the same network.

Note: The 802.11i spec was rolled up into a new collection of Wi-Fi specifications called 802.11-2007. It's no longer referred to separately, except historically.

Many devices support WPA, WPA/WPA2 mixed mode, and WPA2. Those modes allow just TKIP keys, either TKIP or AES-CCMP keys, or just AES-CCMP keys, respectively.

Warning! *802.11n networks cannot use WEP or TKIP. They can use only WPA2's AES-CCMP. On a network that's set to be compatible with 802.11n and earlier 802.11b, g, or a standards, you can use mixed WPA/WPA2 modes.*

Because all hardware shipping with the Wi-Fi label since 2004 (and some from 2002 and 2003) can use WPA2, we recommend you use only that mode for the greatest security, in case flaws are found in

TKIP. If you still have some WPA-only devices (from 2003 or earlier), you could use the mixed WPA/WPA2 mode.

Note: The Wi-Fi Alliance has created a set of target dates starting in 2011 for when new devices will no longer be certified if they support WEP or TKIP.

WPA and WPA2 also come in two flavors based on the kind of network you're running: Personal and Enterprise. *WPA/WPA2 Personal,* designed for home networks, uses a single passphrase shared among all users; *WPA/WPA2 Enterprise* requires a network server that handles logins, and provides a unique encryption key to each user on the network automatically without the user seeing the key.

WPA vs. WPA2 terminology: *Because WPA2 has replaced WPA, we refer generally to WPA2, unless we're talking about a feature exclusive to WPA, such as the use of the TKIP key type.*

WPA2 fixes

Although WPA2 Personal also uses a shared key, it is enormously more straightforward than WEP because WPA2 allows you to make up and use a *passphrase*—a sequence of letters, numbers, and punctuation that can be just like a normal phrase you might enter. Passphrases must be between 8 and 63 characters long. (A simpler method involves even less effort: see Simplify with Wi-Fi Protected Setup, a few pages ahead.)

All WPA2 drivers must support the same conversion from passphrase into encryption-key material—the underlying bits that are used to scramble packets—so there's no compatibility issue between Mac OS X and Windows, or among any operating systems.

WPA2 Personal has its own key weakness: short keys based on words found in any dictionary in any language can be discovered through brute force relatively quickly. Making a key at least 20 characters long—pick your favorite obscure song lyric, for instance—overcomes this problem, as does choosing random characters for shorter keys. One cracking service, advertised as a service for checking your key's strength, charges $17 or $34 (price based on speed) to run a cluster of computers that test 136 million dictionary words against a network name you provide.

Changing the network name (SSID) from its default also increases security because WPA2 Personal uses the network name to create the actual key, and because some crackers compile pre-computed keys based on common access point names and dictionary words. One group of crackers has a multi-terabyte-sized database of a million words encrypted against 1,000 common SSIDs.

Do you like pain? *You can enter a WPA2 key in hexadecimal form, if you like pain. While WEP keys are no longer than 26 hexadecimal digits, WPA2 keys are 64 hexadecimal digits long! Devices that support WPA2 must handle passphrase entry in order to gain Wi-Fi Alliance certification. Thus, it's almost impossible that you would need to enter the hex key. At least, we hope so, because it's nearly impossible to enter 64 hex digits in a row correctly.*

Wi-Fi adapters that don't support WPA2 can use only WEP, and while some access points—including both models of Apple's 2007 AirPort Extreme Base Station—can support WEP and WPA/WPA2 on the same network, we don't recommend it. The best option is to upgrade to devices that support at least plain WPA2. Check manufacturers' Web sites for information on updating.

Note: Windows XP Service Pack 2, Windows Vista and 7, and Mac OS X 10.3.3 and later all support both WPA and WPA2.

Turn on WPA2

Enabling WPA2 security is straightforward in every gateway we've used. But, before you get started, make sure you are ready:

- For office security setup, first read Secure Small Office Wi-Fi, where we explain using WPA2 Enterprise, which is more appropriate for company security.
- If you plan to use Wi-Fi Protected Setup (WPS) to connect from clients (or if you don't know what WPS is) see Simplify with Wi-Fi Protected Setup, a few pages ahead. Some kinds of WPS require that WPA2 be configured first.

Launch the configuration software or a Web browser to connect to your router before proceeding; you'll find steps ahead for Apple base stations, Linksys routers, and other routers.

Note: A very small number of base stations have WPA2 and WPS modes. When set to WPA2, you set and enter keys by typing. When set to WPS, keys can be handled only automatically. You must choose WPA2 mode if any device you want to connect lacks WPS support.

Configure an Apple base station

To set up WPA2 on an Apple base station, follow these steps:

1. With the base station selected in AirPort Utility, click the Manual Setup button (at the bottom of the window).
2. In the AirPort pane, click the Wireless button.
3. From the Wireless Security pop-up menu, choose WPA2 .Personal. (If you need backward compatibility with 2003-or-earlier computer models, choose WPA/WPA2 Personal.)
4. Enter a passphrase in the Wireless Password field and re-enter it in the Verify Password field. We recommend selecting the Remember checkbox beneath to store the password in your keychain.
5. Click Update to save the settings and restart the base station.

Configure most Linksys routers

Most Linksys routers share a similar configuration screen regardless of model. Older units may need to have their firmware updated to see the same choices. For instance, with a WRT610N (a simultaneous dual-band 802.11n router released in 2008), follow these steps, once for each radio band supported:

1. In the Web-based configuration system, after logging in, click the Wireless pane, and then the Wireless Security tab beneath it.
2. From the Security Mode pop-up menu, choose WPA2 Personal. (The WPA Personal mode is for older devices and offers fewer compatibility options.)
3. From the Encryption pop-up menu choose either:
    - AES: Only WPA2 compatible adapters may connect.
    - WPA-TKIP/WPA2-AES: Any WPA- or WPA2-supporting adapter may connect. Choose this option only if you truly need backward compatibility.

4. In Passphrase, enter an 8 to 63 character sequence.
5. Click Save Settings to restart the router with the new security method.

Configure other routers

Other Wi-Fi routers will more closely resemble Linksys routers' options than Apple's. They may label WPA2 Personal as WPA2-PSK (Preshared Key) or provide an option to enter a 64-digit hexadecimal number. You are typically presented with the choice, as with Linksys, of TKIP, AES, or TKIP/AES-CCMP, with the same tradeoffs.

Connect with WPA2

When you use a WPA2-protected network, your computers' network adapters must be set to be compatible with the settings of the network you choose to join. Typically, you choose a network from a list, which both Mac OS X and Windows XP/Vista/7 identify as having some form of protection:

- Windows clearly states the kind of protection the network uses.
- With Mac OS X 10.5 Leopard or 10.6 Snow Leopard, hold down the Option key before clicking the AirPort menu (in the menu bar), and then hover over a network name to see what encryption it uses.

To join a network, choose it from the AirPort menu in Mac OS X or the wireless networks dialog box in any recent version of Windows. After selecting the network, enter the encryption key when prompted.

Tip: You can create network profiles for commonly used networks where you store the WPA2 key for automatic logins. Glenn writes extensively about connecting from Mac OS X and Windows and setting up such profiles in *Take Control of Your 802.11n AirPort Network*.

The One Tiny Problem with WPA

The only known vulnerability for modern networks affects just WPA and its TKIP key type. Researchers in Germany discovered that they could use a flaw from WEP that was incorporated into the improved TKIP system to modify and insert very short sequences of data into a network when used with a router that supports Wi-Fi Multimedia (802.11e). WMM is used to improve the quality of voice calls over a Wi-Fi network. For more detail, see Glenn's *TidBITS* article, "A Crack in Wi-Fi Security and How To Fix It," at http://db.tidbits.com/article/9846.

The researchers' work relied on a lot of separate pieces and it didn't allow long, arbitrary amounts of data to be inserted, nor did it allow decryption of data passing over the network.

This is one reason to choose AES-CCMP exclusively, as that stronger key type can't be broken with this particular flaw. But it's also likely that other flaws will come to light so long as TKIP is in wide use.

Simplify with Wi-Fi Protected Setup

The Wi-Fi Alliance has taken WPA2 one step further by eliminating even the need to create and enter a passphrase; security happens with just a click or two. *WPS* (Wi-Fi Protected Setup) lets computers and other Wi-Fi–equipped devices join a WPA2-enabled Wi-Fi network without any password, or by using a short PIN code that's generated by the Wi-Fi software, and which you enter once.

Problems with WPS Adoption and Implementation

We're disappointed in the support WPS has seen as we write this, about 3 years after it was formally introduced. WPS requires updated operating system software, drivers, and hardware, and we haven't seen as many examples of WPS in action we'd hoped.

Apple has built a closed infrastructure in which WPS works with Mac OS X 10.4.8 Tiger or later using any of Apple's Wi-Fi gear introduced in 2007 or later. Apple extends WPS to allow simple controlled access, too; you can let a computer join a network for just 24 hours, for instance. In testing, we've been unable to make Apple's WPS implementation work with any other company's routers, or any PCs running Windows.

What can be confusing with WPS is that depending on your base station, you may have to configure WPA2 first in order to use WPS, or you may have to put your base station into a WPS-only mode in which you can't enter a WPA2 passphrase by hand. Other base stations offer a hybrid mode with some assistance.

For instance, Apple requires that you set up WPA2 encryption first (as explained earlier in this section), at which point you can add clients via WPS. With a newer Linksys router, by contrast, WPS is a default security mode (in the Web-based configuration system's Wireless pane's Basic Wireless Settings tab). Linksys then offers the encryption key that it generated so you can enter it by hand on non-WPS clients, or use one of several methods to connect via WPS.

Once you figure out whether or not WPA2 needs to be enabled first, and then make sure it's turned off if necessary, you can start setting up a WPS connection.

There are three ways that you can set up a protected connection using WPS, depending on your hardware: by the base station advertising its availability, by the client initiating the connection, and by the client requesting a PIN from the base station.

Advertise availability

Apple's approach (which we haven't yet seen elsewhere) has a base station initiate the process by advertising its availability. You tell the base station you want to attach a client via WPS, and then turn to the client to obtain the necessary information to enter in the base station management software:

1. Either:

    - Connect to your Wi-Fi gateway through its administrative interface (stand-alone software or a Web browser, depending on the model) and tell it you want to add a computer or device with WPS. With an Apple Wi-Fi gateway, use AirPort Utility (found in /Applications/Utilities). Where appropriate, choose whether the next machine that joins is automatically secured via WPS, or whether a PIN is required.
    - Press a physical button on the router to trigger a WPS mode (not available on Apple base stations).

2. From the device you want to add to the network, choose your Wi-Fi network.

3. This step depends on which method you chose during Step 1:

    - With the "next machine that joins" approach (always the case if you pressed a button on the base station), the device exchanges information with the gateway, receives a WPA2 key, and joins the network.

    - With the client PIN approach, the client device displays a short code. You then enter that code into a field in the gateway's administrative interface. If you enter it correctly, the device receives the WPA2 key, and joins the network.

Client initiates

Another approach, found in gear from Linksys and others, requires a client to take action to which the base station then responds. It works like this:

1. From the client device, either:

    - Press a physical button or a button in a piece of software that initiates a WPS request.

    - Use software on the client computer or other device to create a WPS PIN and make a note of the number.

2. Connect to your Wi-Fi gateway through its administrative interface (stand-alone software or a Web browser, depending on the model).

3. In the base station's wireless security area, either click a button to confirm the WPS process, or enter the client's PIN.

Client requests a base station PIN

There's even *one more approach*, found in Linksys gear, in which the client software asks for a base station PIN:

1. From the device you want to add to the network, choose your Wi-Fi network. You're prompted for the base station PIN.

2. Connect to your Wi-Fi gateway through its administrative interface (stand-alone software or a Web browser, depending on the model) from another computer.

3. Find the part of the interface that provides the base station's PIN.

4. Enter the base station PIN on the client.

One last reminder

There's always an option to use a passphrase, too, for older devices that don't support WPS. As WPS appears on a greater number of gateways, computers, and other devices, you might never have to enter a WPA2 passphrase again. But unless manufacturers can settle on compatible ways of using WPS, its promise may go unfulfilled.

ENABLE GUEST ACCESS

It's routine for us to have visitors come to our home or office who need Wi-Fi access. For Glenn's brother–in-law, Michael, he just types in the network password and stores it; Mike's OK. But if someone you don't know as well wants to use the Internet, or, heaven forbid, you have friends of your children who want on temporarily, you might want to have a slightly different policy.

Use Apple's guest networking

In March 2009, Apple added a new option to offer guests access: a separate virtual network that runs on the same hardware as your main network, but which appears with a different network name, and which can have a unique (or no) password. Guest networking provides only Internet access to clients using any platform, not just Mac OS X. Hardware and services on the main local wireless network are separated from the guest network, protecting its security.

Using guest networking requires an AirPort Extreme Base Station or Time Capsule released in March 2009 or later.

The guest network uses both bands at the same time and can't be configured with a separate name for each band, which is possible for the main network. Password protection is optional.

Note: The base station must be configured as a gateway, sharing addresses. It won't work if it's set to bridge mode: in AirPort Utility, in the Internet pane, in Internet Connection, examine the Connection Sharing pop-up menu. If you have it set to Off (Bridge Mode), then guest networking is unavailable.

To configure this option:

1. Launch AirPort Utility, which you'll find in /Applications/Utilities/.
2. Select your base station from the list at left.
3. Click Manual Setup.
4. Click the Guest Network button in the main pane.
5. Select the Enable Guest Network checkbox, and enter a network name. It's prefilled for you by Apple with your computer's logged-in user name.
6. To allow guests to use Bonjour networking or other local network services only on the guest network, select the Allow Guest Network Clients to Communicate with Each Other checkbox.
7. You can optionally set password protection by choosing WPA/WPA2 Personal or WPA2 Personal from the pop-up menu and entering a password. The options here are set the same as in Turn on WPA2, earlier.
8. Click Update to restart the router and activate the guest network.

Share via Easy WiFi

Devicescape Software aims to help you avoid entering passwords and keys for Wi-Fi hotspots and home or small networks. You can use Devicescape's Easy WiFi system to give your friends, and friends of your friends, access to your network without giving them a key, and you can revoke access at will. (Not that you'd ever pull the rug out from under your actual friends, of course.)

Sign up for a free account at http://www.easywifi.com/, and enter information about networks you have, including the security pass-phrases. You can then invite ("share") friends to use the network. They also need free Easy WiFi accounts, and they must install simple desk-top or mobile device software that automates the login process.

When friends arrive at your network, the Easy WiFi desktop software automatically connects their computers. On any iOS device, your friends must first launch the free Easy Wi-Fi app.

If you later want to remove someone's access or "unshare" your network, log in to the EasyWiFi site to make the change.

Secure Your Data in Transit

In the previous chapter, we explained how to prevent people from accessing your wireless network. However, you may still find yourself in circumstances in which you want protection but restricting access to the network won't help:

- You're sharing your local wireless network without encryption or using a shared network on which encryption can't be enabled.
- You're using a public wireless network in a location such as a coffee shop, hotel, airport, or community networking hotspot with a laptop, smartphone, iPad, or other mobile device.
- Your employer won't let you use any network except its wired network without encrypting communications to and from your computer.

You have an alternative: you can encrypt the data before it leaves your machine, and have it decrypted only when it arrives at its destination. By creating end-to-end links using strong encryption standards, you can keep your data completely safe from prying network sniffers. Even if people can join your network and reach the Internet—hijacking your link—they still can't see your data. Encrypting your data in transit is a lot more difficult than setting up a closed WPA2-protected network, but it provides more security.

We look at five popular categories and methods of securing data in transit, ranging from simple password protection up to full network encryption of all data, summarized in **Table 2**, coming up.

Tip: An added bonus of encrypting data from end to end is that the data you send and receive becomes completely unreadable by others not just on your wireless network, but also on every Internet link between your computer and the destination machine. That's why large organizations generally require their employees to use encryption technology for all communications.

Table 2: Methods of Securing Data in Transit		
Method	**What It's Good For**	**Who Should Use It**
Encrypt Email Passwords	Prevents snoopers from discovering your passwords and using them to access other accounts	Everyone
Encrypt Specific Files and Messages	Good for protecting a particular file or email message that's especially sensitive	Anyone with sensitive data to communicate with another person
Encrypt Sessions and Data Sequences with SSL/TLS	Protecting specific pieces of data, such as credit card numbers, being transferred between programs (like a Web browser and a Web server) or entire sessions	Everyone for banking and shopping, where it's a requirement, along with anyone who has access to SSL/TLS-enabled client and server programs
Encrypt Data Streams with SSH	Protecting data that travels between your computer and your server via a particular protocol	People interested in protecting an entire class of traffic, such as a remote terminal session or FTP
Encrypt All Data with a VPN	Protecting every last bit of traffic that travels to and from your computer to a remote destination	Business and government users who regularly handle sensitive data, or anyone concerned about transferring confidential data over public wireless networks

ENCRYPT EMAIL PASSWORDS

Even if you aren't worried that people might read your email, you should worry about protecting your account passwords (**Figure 4**).

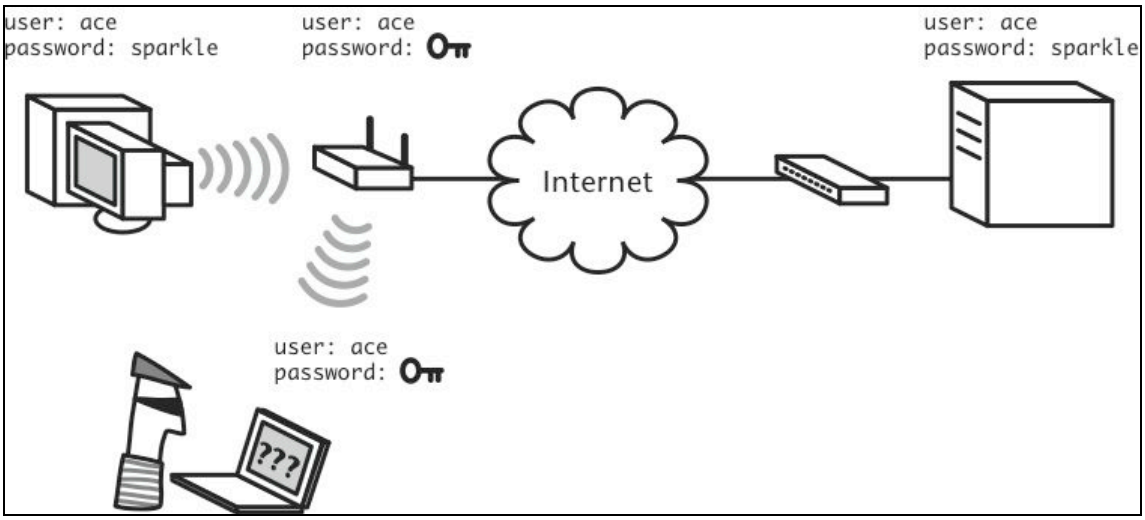


Figure 4: Encrypting email passwords prevents them from being stolen. In this case, the attacker can't read any passwords protected by APOP and SMTP AUTH.

Identification Is Often Insecure

Many Web sites use user names and passwords merely for identification and thus don't use secure pages when asking for those passwords. Since these passwords are easily stolen, make sure they're different from your email passwords and passwords to sensitive information. See Appendix A: Password Advice.

There are two primary methods of encrypting just your password for receiving email—APOP and CRAM-MD5—both of which require support in your email client and your mail server, though most modern mail programs support these options. For sending email, there's just one possibility, but it's a good one. Your ISP or network administrator may have already enabled these various methods on the server side, requiring that you just set an option in your email program.

POP3 with APOP

APOP (Authenticated POP) protects your password when you retrieve inbound email from a *POP (Post Office Protocol)* server. Instead of sending your password in the clear, APOP sends a unique, per-session token that the server uses to confirm that your email program knows the correct password. The token can't be reused after that session is over, or used to recover the original password.

APOP protects only your password; it has no effect on the contents of email or on sending email. We recommend it as a sensible minimum-security precaution; most home users won't need the more significant protections described in the rest of this section.

IMAP with CRAM-MD5

What if you use IMAP instead of POP for retrieving your email? A technology called *CRAM-MD5* can encrypt your IMAP password; if your server supports it, then your email client should automatically use it (you must enable it manually on a per-account basis in Apple Mail). However, CRAM-MD5 isn't particularly secure, which means that for most people, the only way to use IMAP securely is to use it with SSL/TLS to encrypt all the IMAP traffic, described in Encrypt Sessions and Data Sequences with SSL/TLS. Ask your ISP or network administrator if you can use IMAP over SSL/TLS.

SMTP AUTH

SMTP AUTH (the AUTH part is actually an SMTP command), or *Authenticated SMTP,* identifies you to your SMTP server when you want to send outgoing messages. Authenticated SMTP has become commonplace in this age of spam, because if an SMTP server requires SMTP AUTH, that prevents a random spammer from sending spam through that server. SMTP AUTH typically uses the same user name and password that you use for checking mail via POP or IMAP.

Unfortunately, your SMTP AUTH login information isn't encrypted by default, which could itself constitute a security hole. The best way to use SMTP AUTH is to couple it with an SSL/TLS connection. Sensible ISPs that require SMTP AUTH also offer SSL/TLS. For more details, read Encrypt Sessions and Data Sequences with SSL/TLS.

ENCRYPT SPECIFIC FILES AND MESSAGES

Although we strongly suggest protecting your passwords, there is a middle ground between encrypting passwords and encrypting all your data: using content encryption on specific files and email messages. This approach lets you protect the pieces of content that you feel are the most sensitive.

Content encryption makes it almost impossible that anyone other than your intended recipient could read the file or email message, even if

they obtained access to a mail server between you and your recipient. That's because when you encrypt content manually on your end, it usually requires the recipient to decrypt it with a manual action on the other end.

Tip: For a full description of exactly how to send and receive encrypted messages in Apple Mail, read *Take Control of Apple Mail in Snow Leopard*.

Beware of Physical Access

There are three things to be worried about when you're encrypting specific pieces of content, both related to physical access.

Should a bad guy manage to access your computer while your account is logged in, he may be able to read any content that was encrypted by a key stored in an open password manager or, on a Mac, an unlocked keychain. You can usually set such tools to lock themselves after some period of time to reduce this risk.

If a bad guy obtains access to your recipient's machine, all bets are off with regard to whether or not your content can be read. Your recipient may have decrypted it and stored the plain text version already, or she may have weak passwords that enable the attacker to use information on the machine (the name of the hard disk, for instance) to guess the necessary password and decrypt your message.

And, much as we hate to say it, a bad guy who really wants a piece of information very well may go low tech and resort to physical violence. It's extremely unlikely, of course, but if you're trading in highly sensitive information, don't assume that encryption will provide all the protection you need. For an illustrated guide to that concept, see http://xkcd.com/538/.

The most popular software that encrypts the contents of messages or entire files is called PGP (Pretty Good Privacy). PGP uses public-key cryptography to secure a message so that only the intended recipient can read it (**Figure 5**).

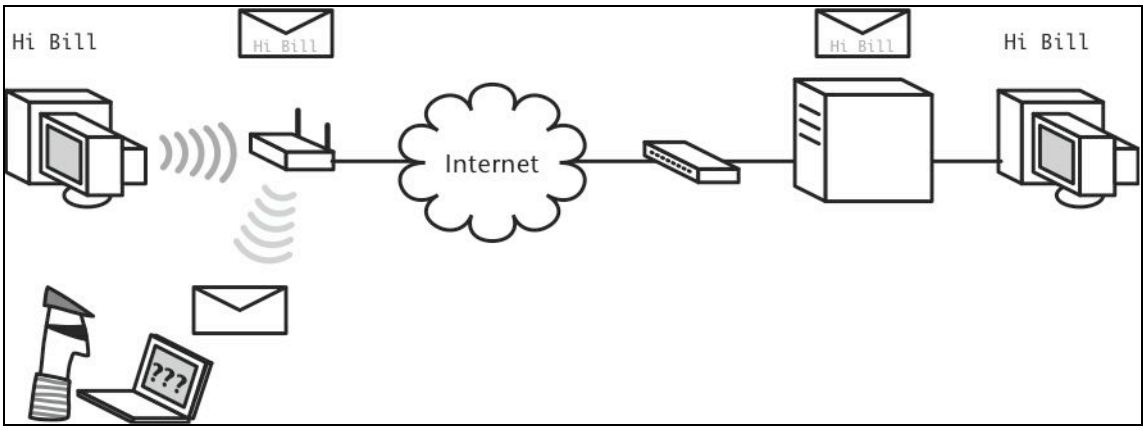


Figure 5: Encrypting email messages and specific files with PGP prevents crackers from reading just those messages or files. In this case, although the cracker can see an encrypted copy of the message, he can't read the contents.

Note: PGP (the corporation, now part of Symantec) used to offer affordable desktop and free (non-commercial use) software. That's all gone; the company's products are now all aimed at the enterprise.

An Open-Source Alterative to PGP

An open-source alternative to PGP is GPG or GNU Privacy Guard (http://www.gnupg.org/); a Mac version with various graphical interface tools and email client plug-ins is available (http://macgpg.sourceforge.net/). GPG works with most files created by newer versions of PGP and vice versa, but read the GPG FAQ if you plan to use the two together.

Understand public-key cryptography

In public-key cryptography systems, each user generates a pair of keys, one public, and one private. Using combinations of those keys, users can sign files or messages to prove that they sent them and can encrypt files or messages so only the intended recipient can open them. The keys work together like puzzle pieces—if someone encrypts something with your public key, only your private key can open it. And, if you sign something with your private key, only your public key can verify that you signed it. As an example, assume that Glenn and Adam set up PGP so they can exchange encrypted drafts of this title without concern about industrial spies from other publishers sneaking a look at the drafts. (In reality, we're nowhere near that paranoid.)

The first step in public-key cryptography is to generate a public key and a private key. Along with the keys, you must generate a passphrase that enables you to decrypt and unlock your own private key when you want to use it. Adam and Glenn both run through these steps, so they each have public and private keys, and then they share their public keys with each other (see Distributing keys). Now, here's how they can use their keys.

Adam wants to send Glenn an extremely important email message regarding the book schedule. Adam's not worried about someone else seeing this message, but he does want to make certain that Glenn believes it comes from him, and hasn't been forged by some joker on the Internet. (Email forgery is remarkably simple, though doing it in such a way that it can't be traced easily or identified as a forgery is more difficult.) So Adam signs the message using his private key.

When Glenn receives the message, he can read it with no extra effort, but to verify that it did indeed come from Adam, he uses Adam's public key to check the signature. When they match, Glenn knows that the message is legitimate (**Figure 6**). Had someone used any other private key to sign the message, it wouldn't have matched with Adam's public key and the signature verification would have failed.

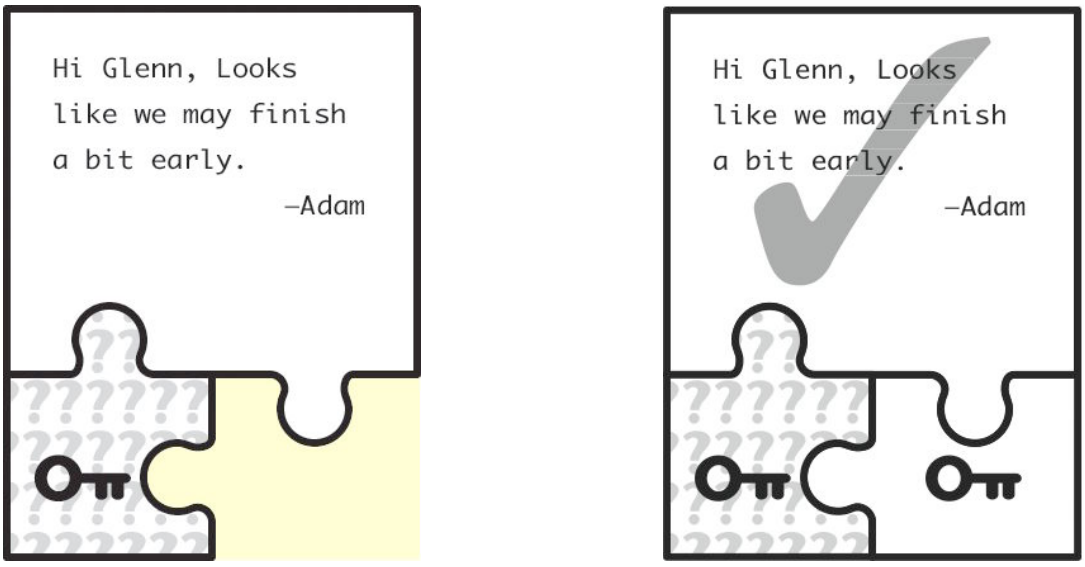


Figure 6: Adam signs the message with his private key (left), and Glenn uses Adam's public key to verify the message (right).

Next, assume that Glenn wants to send Adam a draft of the book, but because he's worried that one of his neighbors may be eavesdropping on his wireless network traffic, he decides to encrypt the file before

sending it. This time, Glenn uses Adam's public key within the PGP software to encrypt the file, and then sends the file along. When Adam receives the file, he uses his private key within the PGP software to decrypt it (**Figure 7**). If someone were to intercept the file and try to decrypt it, she couldn't because only Adam's private key can decrypt files encrypted with his public key.

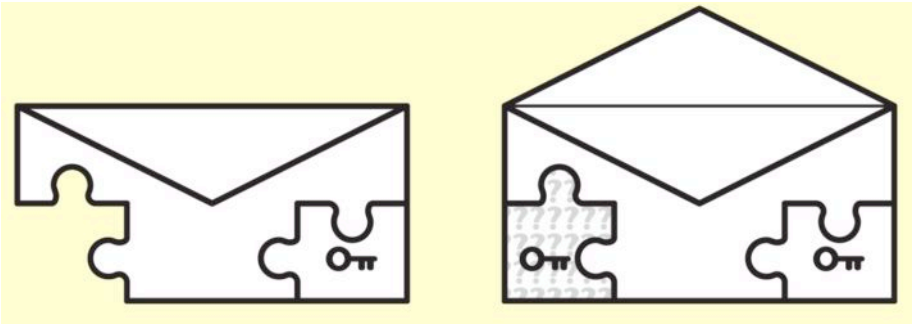


Figure 7: Glenn encrypts the file with Adam's public key (left), and Adam decrypts the file with his private key (right).

Keep your private key safe! *You can see from this example how important it is that you keep your private key safe and don't share it with anyone. If someone were to learn your private key, that person could forge your digital signature and could decrypt any encrypted information sent to you. In reality, if your private key were compromised, you'd have to revoke its public key partner to forestall anyone from using the public key.*

How PGP works

Public-key systems as implemented in PGP and in SSL/TLS don't use public keys to encrypt all the data. That would be inefficient: public-key cryptography uses a lot of computing horsepower—as much as 1,000 times more than simpler but strong encryption systems in which the keys are insecure, and must be kept private at all times.

Instead of encrypting everything with a public key, PGP and SSL/TLS generate a strong session password, used uniquely for a single document or Web session, and then encrypt just that password using a public key.

Thus a PGP-protected document starts with a public-key header, which contains the session key. With a private key in hand, PGP software or your browser can retrieve the session key, which then can be used to decrypt the document.

Thanks to this approach, PGP documents can also be read by multiple recipients by using each recipient's public key to encrypt the session key and storing all those encrypted session keys in the header. If the entire document (rather than just the header) were encrypted with a public key (rather than a session key), only a single recipient could decrypt it.

Distributing keys

The fact that public keys can be shared without jeopardizing encryption is what makes the public-key cryptography method unique. But sharing public keys is also an Achilles heel: how do you distribute your public key and receive keys from others for the first transaction? You can send it in email; include it in your email signature; put it on your Web site; or post it to a public directory, called a *keyserver* (such as keyserver.pgp.com, which is available from within PGP).

Although these methods all work, none are watertight, because someone bent on impersonating you could forge mail from you or post a key to a keyserver while pretending to be you. Once the fake key is out in the wild, revoking it is tricky.

The solution is to exchange keys in ways that ensure the other party's identity. For instance, Glenn and Adam could have created their public keys while at lunch together; being able to see the other person is as much verification of identity as is usually required. Slightly less sure, but more reasonable, is using a telephone or fax machine; in those cases you don't read out or write the entire public key (which is way too long for accurate transcribing).

Instead, you convey a shorter sequence of letters and numbers that verifies to the other person that the public key you've sent them is indeed yours—the sequence of letters and numbers is called a *fingerprint*. Some people put their fingerprints in their email signatures, assuming that a recipient can email them to verify identity.

Luckily, although it can be tricky to verify that a public key does indeed belong to a specific person, the worst possible outcome is that someone could distribute a new public key under your name, thus bringing documents ostensibly signed by you into question. But when somebody sends you an encrypted file or message that uses this fake public key, you can't decrypt it with your private key. This should alert you to potential problems, but your security is intact.

PGP's software combines your email address and your public key, which means that for someone to mimic you, she'd have to intercept your email and convince other people that their public key was accurate—at which point, she pretty much owns your digital identity already, so you have other problems. The PGP company keyserver also sends you email every year to confirm your identity, and if you fail to respond, it removes your entry, pruning older, unused keys from the database.

Once you have a public key for someone and have verified who she is, you can exchange messages for the life of the key. Many public keys are set to expire on a certain date for additional security.

ENCRYPT SESSIONS AND DATA SEQUENCES WITH SSL/TLS

SSL (Secure Sockets Layer)/TLS (Transport Layer Security) was initially developed to secure financial transactions on the Web, but it is now widely used to secure Internet transactions of all kinds. Every time you see an https URL, you're using SSL/TLS in your Web browser to secure the communication between your browser and the remote Web server.

What's the Difference between SSL and TLS?

SSL and TLS aren't interchangeable: TLS is a later version of SSL, although TLS supports the ultimate version of SSL that preceded the first version of TLS. A browser with TLS 1.0 or later can securely communicate with an server running SSL 3.0. Fundamental patents in SSL expired years ago, and the effort to continue its development were made under this new name.

SSL/TLS solves the "shared secret" problem in WEP and WPA2 Personal by using an *ecosystem* version of public-key cryptography: different elements work together to form a complete solution, just as symbiotic pieces of an ecosystem allow plants to bloom, insects to multiply, and animals to thrive, but only together. (See Encrypt Specific Files and Messages, earlier.)

Instead of requiring that people agree in advance on a secret (the encryption key in WEP or WPA2) or requiring that public keys be

published, an SSL/TLS-equipped browser and server use trusted third parties, known as *certificate authorities,* to confirm each other's identity. The authorities bootstrap the capability of a Web browser to accept a key in a safe manner that isn't vulnerable to violations of trust.

Tip: To see what authorities Apple says you trust, in Mac OS X 10.5 Leopard and 10.6 Snow Leopard, launch Keychain Access (found in /Applications/Utilities), and click the System Roots item in the Keychains list at the upper left.

This trust can be a weakness: If a party convinces a certificate authority to believe in a fraudulent entity, that trust could be misused to sucker users, especially in conjunction with *phishing,* where a spam email message is used to convince someone to visit a Web site that's not what it appears to be.

Typically, SSL/TLS is used for session-based interactions, like sending credit card information via a Web form. But it can be used to encrypt email sessions (sending and receiving, including the entire contents of email from the client to the server), for FTP (in a form known as FTP over SSL/TLS or FTPS), and for many other transaction types.

SSL/TLS works with any Internet service that uses TCP for data exchange in which every packet is designed to be received or retransmitted if it doesn't reach its destination. (For example, you can use SSL/TLS for instant messaging, but not for RealAudio music playback, although there are some workarounds.) Because a third party can verify SSL/TLS, you needn't rely on trust (blind or confirmed) as you do with SSH.

You can also work with SSL/TLS where there's no certificate authority, just a *self-signed certificate*. This certificate isn't signed by another party, but for certain cases, like a private mail server you or an in-house administrator sets up, it's good enough, and saves the yearly fees required for authority-backed certificates.

Unlike SSH, in which you can connect any two arbitrary ports through a tunnel (port to port), SSL/TLS works from program to program, with the client encrypting data and the server decrypting it (**Figure 8**). But, just like SSH, all data sent over that tunnel, including passwords and all sent and received content, are securely encrypted.

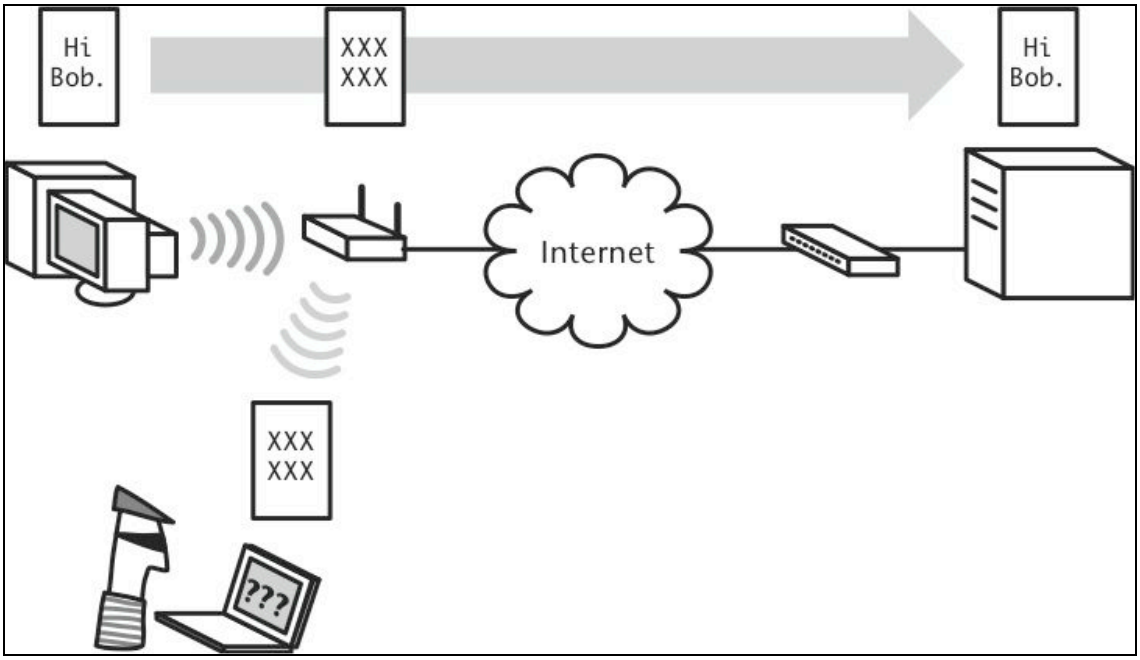


Figure 8: Encrypting select transactions with SSL/TLS protects just the contents of that transaction. In this case, the attacker can't read any of the traffic encrypted with SSL/TLS, but potentially could read other traffic sent over the connection.

Understand how SSL/TLS works

When you connect to an SSL/TLS-protected Web page, your Web browser and the remote Web server must negotiate the exchange of keys.

Your browser and operating system have a preinstalled list of digital signatures from *certificate authorities,* which are companies that confirm the validity of a given certificate being attached to a given server by IP address. When a browser connects to an SSL/TLS server, it receives the server's certificate, and uses its built-in certificate authority data or that of the operating system to confirm the certificate's validity. The browser checks that signatures match, in a manner of speaking. Browsers warn you if there's no third-party verification or there's a mismatch.

If the certificate is valid, the browser uses the public key contained within the certificate to encrypt a session key, which can be used with no worries that it was intercepted because of the third-party check. (After the certificate is exchanged, the process for creating the session works very much like PGP.)

With an email client and server, a similar transaction happens. The email client requests a secure connection; the mail server responds with certificate information; and the two exchange keys before exchanging mail.

Use SSL/TLS where possible

In contrast with PGP and SSH, when it comes to SSL/TLS, you can typically encrypt the data sent and received by any SSL/TLS-equipped application without entering passwords or any other convolutions. The client software handles communication just as it would with an unencrypted connection.

While SSL/TLS encrypts the connection without a password, this approach doesn't bypass authentication for the Web site or other service: you must still provide a user name and password to access your bank account, for instance, but that user name and password are secured from outside view.

Secured Web sites

With the Web, modern browsers and servers use SSL/TLS when necessary (assuming the webmasters have set it up properly). You can tell when you are viewing an SSL/TLS-protected Web page because there's often a little closed-lock icon in one of the extreme corners of the window (outside the content portion of the window). Also, look for the telltale sign in a URL: instead of the URL starting with `http`, it starts with `https`.

Most banking and shopping sites, as well as any site that deals with confidential, legal, medical, or government information, use encryption for any sensitive part of the session, which can include the entire session. Amazon.com, for instance, uses SSL/TLS during login, switches to unencrypted Web access while browsing, and requires an additional entry of your password and a secure connection for checkout.

Increasingly, other categories of sites, such as the modern-day versions of print publications, are offering secured versions of their sites so users don't have to worry about their Web activity being monitored by scoundrels on open networks, even if there's no security risk involving passwords or access. Examples include Wikipedia, Google Search, the *New York Times*, and the *Washington Post*.

There's an enhanced version of browser-based SSL/TLS certificate security called *EV* (Extended Validation) certificates, where a Web site

operator pays a much higher fee to a certificate authority that uses much stricter procedures to ensure the Web site operator's identity. When you're browsing the Web, you can indentify EV SSL/TLS sites because a green bar appears in the location field of your browser (**Figure 9**). Click the bar to display additional information.



Figure 9: Glenn's Washington credit union, BECU, makes sure its customers know that it is highly trusted by opting for an EV certificate for its online banking operations.

Warning! *If you're using a browser that's a few years old, you won't see any EV green bar.*

Other sites that have a lower level of concern about data sent and received use SSL/TLS only during the login or authentication process, leaving the session in the clear. This includes most webmail sites. In some cases, you can force the use of a secure connection.

For instance, in Gmail, log in on a safe network (such as at home), and click the Settings link near the upper right of the page. In the General pane, beside Browser Connection, make sure Always Use https is selected. If it isn't, select the radio button, and then scroll way down and click Save Changes.

Note: In early 2010, the new version of MobileMe's Mail webmail application added full SSL/TLS security.

Secure email connections

Many email programs support SSL/TLS, and turning on SSL/TLS simply requires that you set an option, often hidden in an advanced configuration dialog.

Unfortunately, not all mail servers support SSL/TLS, and of those that do, not all handle SSL/TLS in the same way. As a result, some SSL/TLS-capable mail servers aren't compatible with some SSL/TLS-capable email programs. To learn if SSL/TLS is an option for protecting your email, ask your ISP or system administrator, or install a mail server that is compatible with the software you've chosen to use.

Secure FTP

You can secure FTP over SSL/TLS (FTPS) in free or commercial software, as long as the FTP server handles SSL/TLS connections. It's often tricky to get FTPS working because of the necessary certificates. Many ISPs use self-signed certificates, or would prefer that you use certificate-free SFTP (SSH-based FTP), explained next.

ENCRYPT DATA STREAMS WITH SSH

SSH (Secure Shell) was created to establish encrypted terminal sessions that *tunnel*—that is, create simple end-to-end connections—between a client computer and a server computer (**Figure 10**). SSH was necessary because the remote terminal protocol (called *telnet*) sent all information in the clear as plain text, allowing any network snooper to grab data.

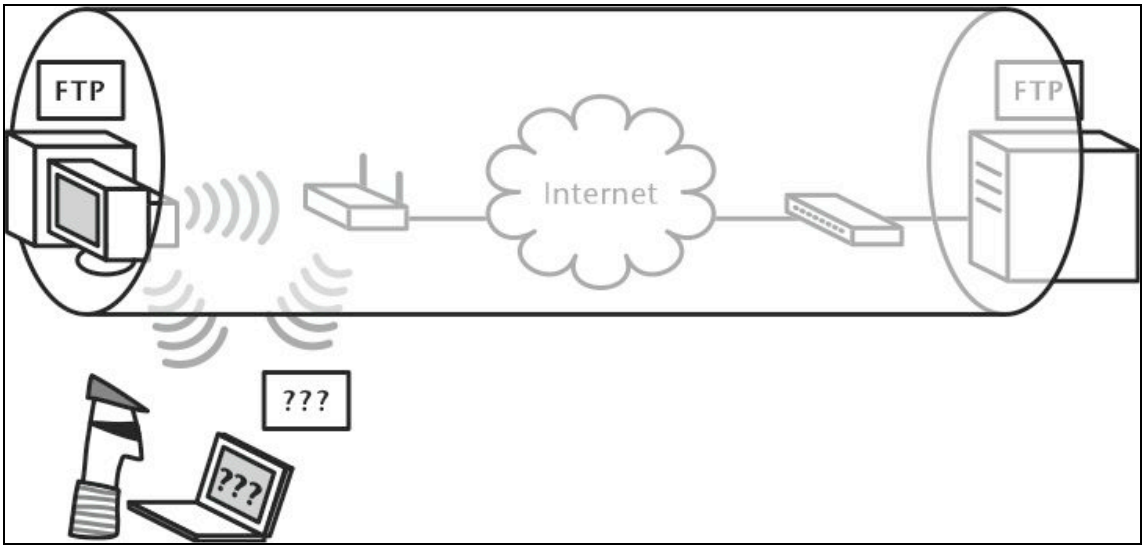


Figure 10: Creating a secure tunnel with SSH protects the data inside the tunnel. In this case, the attacker can't read any of the FTP traffic in the SSH tunnel.

SSH has expanded far beyond this original purpose. It now lets you create tunnels for any kind of TCP (though not UDP, a lower-level

Internet protocol) protocol, whether POP, SMTP, Web, or even Timbuktu Pro. It does this with a trick called *port forwarding,* which connects a local port on your computer with a remote port on a server.

With SSH, you're protecting both passwords and all the content that you send or receive via the Internet services you choose to tunnel. When used as a part of a *proxy*—an intermediary between your computer and the Internet—you can secure all Web connections and many other forms of in-the-clear connections quite simply.

SSH Best Used for Securing FTP Sessions

After years of working with SSH, we've decided that it's best used for SFTP (a special form of FTP that combines SSH with FTP for both data and the control connection), terminal sessions, and Timbuktu Pro. It's also good for setting up secure proxies where a VPN isn't an option. For email and other kinds of traffic, SSL/TLS is now simpler and superior in most respects because SSL/TLS support is baked into individual email and other applications. (See Encrypt Sessions and Data Sequences with SSL/TLS.)

Understand how SSH works

SSH encrypts the entire contents of any session, and it's considered highly secure. SSH doesn't, by default, use outside trust: the initial exchange between a server and a client to set up a trusted relationship for future sessions requires either blind faith or the use of a confirmation code, also called a fingerprint.

Finding a fingerprint: *An SSH server generates a fingerprint for its encryption key, and when you connect for the first time from a client, you can double-check that the fingerprint your client sees is identical to the one on the server. If you run your own server, you can retrieve the fingerprint yourself (see the documentation for OpenSSH). Otherwise, ask your network administrator.*

Establish Trust

SSH uses public key encryption, as described earlier, as the first step of a session. After trust is established by exchanging public keys, a session is started by encrypting a much shorter symmetric session key with a public key. The server and client can safely confirm the shorter key. A shorter key speeds encryption for real-time data transfer, and symmetric encryption is much faster than public-key encryption, so public-key encryption is generally used only to establish a random symmetric session key. This approach provides the advantages of public key cryptography without requiring all the computation of a pure-public-key implementation.

Port forwarding with SSH involves connecting a TCP port on your local computer with a port on a remote machine using an encrypted SSH tunnel as the connector. For instance, if you want to retrieve email via an SSH tunnel, you first set up the tunnel between the POP (Post Office Protocol) port (110) on your computer and the remote server's POP port. Then you configure your email program to retrieve email from IP address 127.0.0.1, which is a generic alias for your local machine, on that same port 110. SSH can establish multiple tunnels, such as POP and SMTP, with a single command.

System Administrator Access Needed for Some Changes

Using low ports (under 1024), such as 110 or 25, requires an administrator account under Mac OS X, and root-level access on other Unix and Linux systems. Because you'd probably connect software servers on your machine only if you were an administrator, this requirement isn't a real problem. You can avoid needing this level of permission if you use ports numbered 1024 or higher.

The SSH software intercepts requests for connections on that port from your mail program and forwards those connections, securely encrypted, to the mail server you specify; responses pass along the same encrypted tunnel.

Fortunately, newer applications that secure their connections with SSH avoid the complexity of setting this up. For arbitrary programs, you still need to know this level of detail, but a modern FTP client likely requires that you simply select SFTP instead of plain FTP.

The main drawback of SSH is that you must have access to a server that can run SSH or SFTP on its end of the connection—usually true for Unix systems (including Mac OS X) but not true of Microsoft Windows. (It takes two to tango in the SSH tunnel.) You may be able to avoid this problem by running your own SSH-equipped or -capable servers or by working with an ISP or a network administrator willing to set up the connection.

SFTP also gives any connected user FTP access to all files on the system to which the user would have access while logged in via a terminal session. This may or may not be appropriate for a given user. Plain FTP servers let you set special login directories for FTP users; that's an advantage of FTP over SSL (FTPS), which wraps regular FTP in an SSL/TLS tunnel.

Another issue is that SSH works well only for services that need single ports and use TCP. When you get into more complex arrangements or UDP packets—often used for streaming media—you need to pursue software or services with encryption built in, or something more comprehensive like a VPN (see Encrypt All Data with a VPN).

FTP over SSH, SFTP, and FTPS, Oh My!

There are three popular kinds of secured FTP, two of which use SSH and one uses SSL/TLS. This may lead to confusion because of how the acronyms were formed and what's actually secured.

- **FTP over SSH:** This method secures only the control channel, or the part of the FTP transaction that involves sending the user name, password, and commands. The data portion is sent in the clear. FTP over SSH is the only one of the three standards that can work with an arbitrary client and server that aren't running special FTP software, because it encrypts just the single control channel. The other two methods require coordination.
- **SFTP (Secure FTP):** This technique uses the sftp program and sftp-server software to communicate using SSH, encrypting both control and data connections. This preferred approach is widely available. Under Mac OS X, you enable the sftp-server software in the Sharing System Preferences pane by turning on the Remote Login service. SFTP can work with any server that you agree to connect to, with no prearrangement.
- **FTPS (FTP over SSL/TLS):** FTPS uses SSL/TLS to secure a connection. Not all FTP clients support this method, although it's increasingly popular. FTPS requires the client FTP software to accept and approve a certificate from the server, which can add complexity; if the certificate is signed by a third-party CA, it's generally easy, but if a self-signed certificate is used, it can be more complicated to get working.

To add confusion, there's also secure copy (scp), a Unix program generally available only on the command line that uses SSH to perform secure file transfers.

Secure all Web connections with a proxy

One of SSH's hidden strengths is using port forwarding as a security method. Normally, when you browse the Web, every page not encrypted by SSL/TLS is sent in the clear, which makes sidejacking and data interception possible. If you instead use secure proxying with SSH, you can have the security and convenience of SSL/TLS Web sessions without needing every Web server you visit to support that level of protection.

Secure proxies encrypt all Web sessions from your computer to a spot out on the Internet, securing your data across the Wi-Fi and local network links. This works by combining two separate pieces: SSH port forwarding, which enables you to forward traffic to and from your computer securely to a remote machine; and proxying, which lets Web connections to be passed on to another computer that handles the request and relays the responses back to you over the secure channel.

Proxies were originally developed to reduce bandwidth consumption by acting an as intermediary—and often as a content-filtering gatekeeper—between local users on a network and the rest of the Internet. A proxy server receives a request from a browser, and retrieves the page, either from its local cache, or from the source Web server if it doesn't have a fresh copy of that page in the local cache. (Web sites can set expiration times on pages that are meant to persist and identify pages that should expire as soon as they're sent, to avoid reposting information via a form, for instance.)

You combine these two pieces by running software or by using the command line to set up forwarding between your computer and a service provider.

Windows users can turn to a variety of companies for securing Web surfing, including Anonymizer Universal, which combines the Anonymizer software package and a subscription for $79.99 per year (http://anonymizer.com/universal/).

Secure-Tunnel offers much simpler software for Mac OS X, Windows, and Linux that hooks your computer into Secure-Tunnel's proxy servers (https://www.secure-tunnel.com/). For $9.95 per month or $99.95 per year, you can secure Web surfing and other services as well.

The added bonus with both services is that your surfing is anonymous (the Web site doesn't have to know who you are) as well as secured.

ENCRYPT ALL DATA WITH A VPN

As you've undoubtedly noticed, all the encryption solutions we've discussed so far are specific to a type of Internet service, specific files or messages, or certain software. Why not just encrypt everything? For that you need a VPN (virtual private network). *VPNs* are the ultimate solution for securing data because they create an encrypted pipe, called a *tunnel,* that carries *all* the traffic between your computer and a VPN server—it's essentially a secure extension of a network. Since the data you send or receive—email, FTP, Web, and anything else—between your computer and the VPN server is encrypted, even if someone were to break in, she couldn't decrypt the tunnel that carries your communications (**Figure 11**).

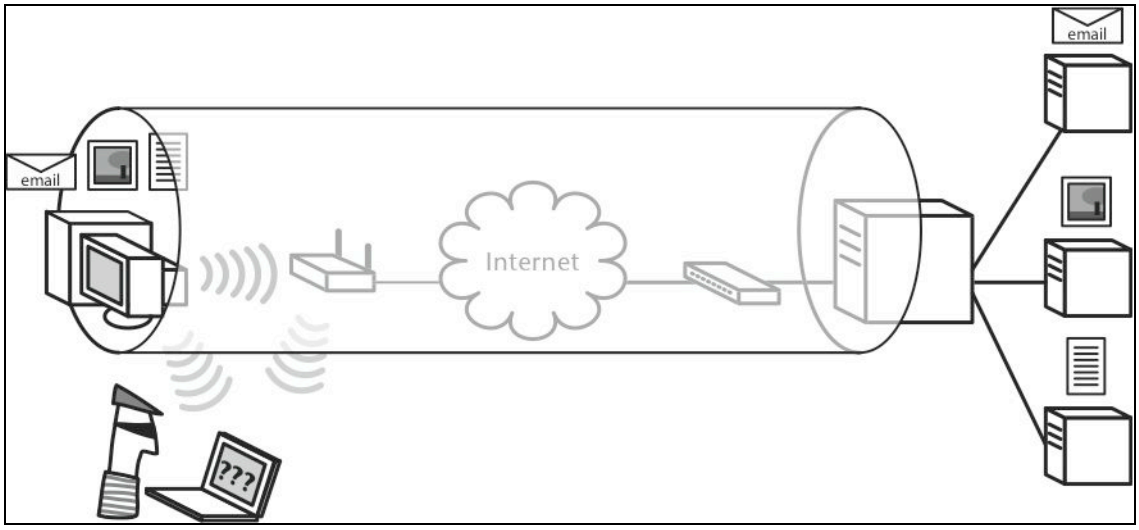


Figure 11: Encrypting all your traffic in a VPN tunnel offers the most complete level of protection. Here, the attacker can't read any traffic at all because of the encrypted VPN tunnel.

Three popular protocols are used for VPNs:

- **PPTP** (Point-to-Point Tunneling Protocol): *PPTP* is an older standard, originally developed by Microsoft, but widely available. Shorter passwords are considered weak in PPTP, among other security concerns (**Figure 12**).
- **L2TP** (Layer 2 Tunneling Protocol) over IPsec (short for "IP security"): Security experts consider *L2TP/IPsec* to be more robust than PPTP and about the same as some flavors of SSL/TLS VPNs. L2TP/IPsec clients are widely available for all platforms.

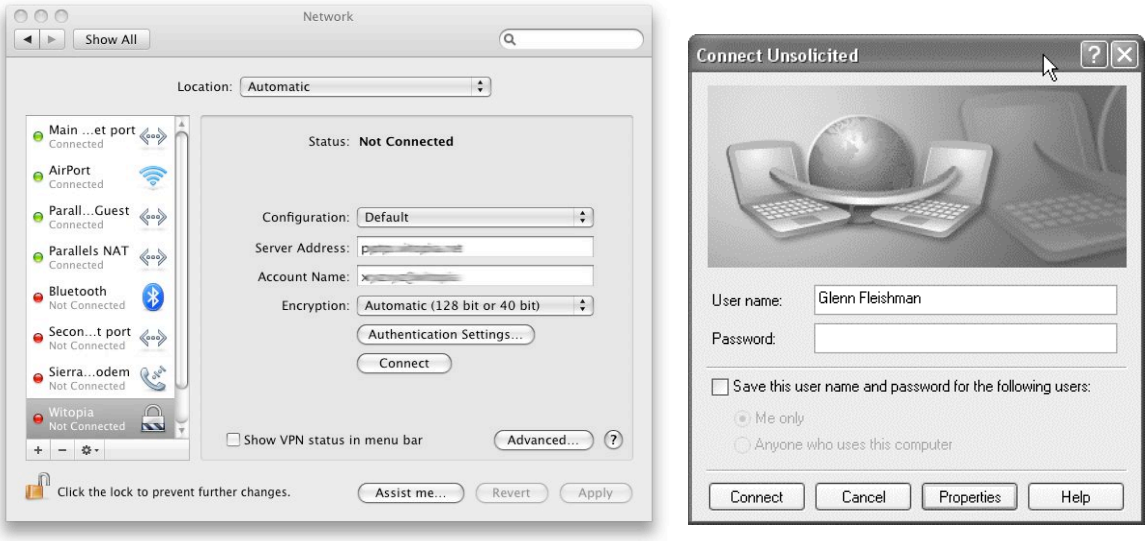


Figure 12: Connecting to a PPTP VPN in Mac OS X (left) and Windows XP (right).

VPNs using L2TP/IPsec used to have difficulty on wireless networks that created their own private addresses with NAT (Network Address Translation). However, most access points have now been upgraded to pass L2TP/IPsec VPN traffic through correctly. If your access point doesn't support L2TP/IPsec, see if the manufacturer offers a firmware upgrade, or buy an access point that can handle it.

- **SSL/TLS:** *SSL/TLS*-based VPNs require special client software, but the OpenVPN open-source project has free and well-made graphical interfaces for Mac OS X and Windows, and many flavors of Unix and Linux, among other platforms (http://openvpn.net/).

A VPN requires both a client and a server, and fortunately, several affordable options are now available, from Internet-based VPN services to low-cost VPN servers that are part of wireless gateways. We discuss these options in Secure Small Office Wi-Fi.

We recommend using a VPN if it's practical and affordable because of the combination of simplicity, complete protection, and peace of mind that it offers.

Note: IPsec used with VPNs is technically called *L2TP over IPsec,* because *IPsec* defines the connection between two points, while *Layer 2 Tunneling Protocol* carries the encrypted data over that connection. Both IPsec and L2TP have elements in common, but they have to work together to create a VPN.

Protect Your Systems

One part of security is protecting your data in transit; the other part is protecting your systems—your computers, any Internet servers you run, your wireless gateway, and so on—from online intruders. Because wireless networks potentially expose your systems to attackers who would never have the same kind of access on a wired network—unless they broke into your house or business—you need to exercise greater care when protecting your computers on wireless networks.

You can secure your computers against snooping or attack in two ways: an active firewall or network address translation. You can use them separately or, for additional security, combined. And of course, it's essential to run current antivirus software if you use Windows. But first, why worry?

GET PARANOID

You might think that you don't need to protect your computers, but, unfortunately, organized and disorganized crime has become rampant on the Internet, and these criminals need machines to do their bidding. There are seemingly hundreds of thousands amoral people or their agents out there constantly and automatically scanning large blocks of Internet addresses for weaknesses.

For the last few years, it's been only a matter of minutes after a computer first receives a *publicly routable IP address*—one that can be reached from anywhere on the Internet—before the first attack is launched against it.

These attacks focus on known bugs in software that allow a remote program or person to infiltrate your computer and take control of some of your software or the entire operating system. Once the attacker has established that level of control, he typically turns your computer into what's called a *zombie*.

Software on a zombie can be remotely directed to cause attacks on targets that are blackmailed to get the activity to stop. But more

commonly, zombies are used to send spam to people that, if read or followed to a Web site, results in viruses being installed, or scams people out of passwords and money. Zombies can even sniff local Wi-Fi traffic for login information that's used to hack webmail accounts at major email providers, which are then used to send spam.

Don't assume that attacks necessarily come from people. It's far more likely that a worm that's already taken over someone else's machine will attack your computer. Worms propagate viruses that in turn propagate worms. The virus may also cause other damage or turn the computer into a zombie for later attacks.

Attacks used to be mostly aimed at Microsoft Windows or specific Windows software from Microsoft. Microsoft has patched known holes, but many Windows users don't download and install these security patches, leaving their computers open to further exploitation and infection.

Install Windows security patches now! *If you're using Windows, stop reading right now and use Windows Update (access it in Start > All Programs > Windows Update) to install all security patches released by Microsoft! If you're a Mac or Unix user, encourage all your Windows-using friends and colleagues to do the same. If everyone would stay current on security patches, most worms would have much less impact. If you're running Windows on a Mac, make sure to keep your virtualized copy of Windows updated.*

These days, however, it's much more likely that a piece of software from another company, like Adobe's Flash software, will be targeted. For example a malicious Web site might use Flash to install files on your computer. While Flash and other programs are frequently patched, users fail in even greater numbers than with operating system patches to keep third-party software up to date.

Although some viruses exist for Macs, the number is a fraction (and a tiny fraction, at that) of those aimed at Windows, which reduces the worry for Macintosh users. Plus, most Mac users surf with Apple's Safari or Mozilla's Firefox Web browser, which, although they've had flaws exposed in the past, are considered relatively safe. To date, there has been no outbreak of a Web-page exploit that, when visited, would compromise a Mac. Researchers and black hats (crackers who sell

or exploit security holes for personal gain) have found and used such exploits, but they've never been widely implemented in the wild.

It's also difficult to force a Mac user to open an attachment in an email program unintentionally, or to convince a Macintosh email program to run malicious code attached to an email message, which are two of the primary methods by which Windows viruses spread. If you don't click on an unknown file and install it, there's no way (at the time of this writing) to get a virus on your Mac.

Like Microsoft, Apple regularly releases security updates for Mac OS X via the Software Update utility—choose Software Update from the Apple (🍎) menu; we always recommend installing them, although it can be a good idea to wait a few days for any unforeseen problems to appear and be fixed. iOS updates for the iPhone, iPod touch, and iPad are handled through iTunes, where the program downloads new software automatically, and prompts you to install fixes.

Other attacks use what's called a *denial-of-service* (DoS) approach, where the attacker sends so much data to your computer that it's overwhelmed. DoS attacks don't cause damage, per se, but they prevent normal operation and can be difficult to shut down. They are unfortunately quite frequent. A *distributed DoS* (DDoS) relies on armies of zombies, instead of one or a few connections, to flood your network. It's much harder to defeat because when one incoming connection is shut down, other zombies can join the fray.

Impact of DoS and DDoS

A DoS attack once saturated Adam's dedicated Internet connection; only calling his ISP and having it block the offending traffic fixed the problem. Also, a DDoS once hit the co-location firm at which Glenn's servers and the *TidBITS* servers are located. Despite having massive Internet connectivity and top-notch technical folks, the attack nearly shut down access for 2 hours.

The entire issue of protecting your computer becomes much more complex when you're roaming. The wireless networks' wired connections could be untrustworthy (is the Internet café's resident geek probing your system?), or a cracker at the next table could be probing your computer over Wi-Fi, or, worst of all, an infected computer within Wi-Fi range connected to the same network could be sniffing all your traffic from an adjacent building. Remember, if someone can monitor

your unencrypted network traffic and steal your passwords or even your *token*—a time-limited cookie used to handle active Web sessions—she can often use those signals to enter your machine while it's on the network.

A little precaution, such as using encrypted connections to protect your passwords and installing a firewall, goes a long way toward preventing an ocean of pain and suffering.

Always back up before you travel! *Always back up your data before you take a laptop on the road—even if you're completely safe from crackers, you may drop and break the computer while going through an airport security check or someone may steal it while you're looking the other way. Everyone loses data at some point, and those with backups suffer the least because of it. Our colleague and friend Jeff Carlson had a flight attendant pour a pitcher of water on his PowerBook some years ago. No, he didn't say anything insulting; it was just an accident, and the airline wrote him a check. Good thing he had a backup.*

For more information, see Take Control of Mac OS X Backups.

INSTALL ANTIVIRUS SOFTWARE

There are tens of thousands of viruses that can attack computers running Microsoft Windows, and a vastly smaller number of viruses have been reported for computers running the Mac OS and Unix as well. These viruses use a variety of methods of infecting computers, and although many are essentially harmless (perhaps only causing crashes due to poor programming), many others are inherently malevolent, with code that causes them to delete or corrupt files, or even erase your hard disk. Also problematic are macro viruses, which live inside documents written with programs that have some sort of scripting—they most commonly infect Microsoft Office documents due to Office's built-in scripting support.

Ways Your Computer Gets Sick

Avenues for infection include inserting an infected USB flash drive into your computer (this caused the U.S. Department of Defense to ban flash drives in November 2008, though the ban was partially lifted in February 2010), downloading an infected file from the Internet, visiting a Web site that compromises your browser, receiving and opening an infected attachment via email, being attacked over the Internet by an automated program, having your wireless adapter hacked into via a driver weakness from a nearby cracker, and more.

You can't prevent every possible way you could become infected (although exercising caution is always worthwhile), which is why antivirus software that restricts access to (and from) and constantly scans your computer is so important on Windows.

Put bluntly, if you're using a Windows computer, you will eventually be infected by a virus unless you run antivirus software that you keep up-to-date. Since so many Windows viruses appear every month, makers of antivirus software always provide an automatic update service that ensures that their software can identify and eradicate newly discovered viruses.

Numerous companies provide antivirus software, but the two most common packages for Windows are Norton AntiVirus from Symantec (http://www.symantec.com/) and McAfee's VirusScan (http://www.mcafee.com/). You can also opt for a basic free package from AVG Technologies, AVG Anti-Virus Free Edition (http://free.avg.com/); they sell a commercial package with more options. These packages are also useful if you're running Windows on a Mac through a virtualization program like VMware's Fusion (Fusion 3 comes with a free copy of VirusScan Plus).

Most antivirus packages are fairly comparable in terms of basic functions, so choose among them based on price, usability, support, and other features. We don't care which you choose, just make sure you run some form of antivirus software and keep it up to date.

Windows will warn you if you disable antivirus software, which is a nice method of ensuring consistent safety.

Sleeping Tight without Mac OS X Antivirus Software

Honestly, we don't run antivirus software on our Macs. Since we maintain nightly backups and avoid common virus vectors like pirated software and random CDs from people we don't know, the antivirus software hasn't been worth the effort in recent years.

We would change our tune if someone discovered a way to execute attachments inside a popular Mac OS X email client that would also redistribute email with the same attachment to others. When that happens, watch us run for an antivirus download.

For more information, see the *TidBITS* article "Should Mac Users Run Antivirus Software?" at http://db.tidbits.com/article/9511.

ASSIGN PRIVATE ADDRESSES FOR PASSIVE PROTECTION

Running NAT (Network Address Translation) on a gateway eliminates many types of break-ins because NAT addresses are typically private—restricted to the local network—and thus unreachable from outside the network. When a computer with a private address on the local network requests data from another machine on the Internet, the NAT gateway rewrites the request so it appears to have come from the NAT gateway, which must have a publicly reachable IP address.

Warning! *NAT doesn't protect you from Web browser exploits. If you were to visit a Web page that tried to maliciously exploit a weakness in your browser (if such a one exists), NAT would help not a whit.*

Note: In fact, you can find yourself behind two or even three layers of NAT, such as happens if you run NAT on your network's Wi-Fi router, connected to the Internet through your network's NAT-enabled broadband gateway, and that in turn connected to an ISP that runs a form of NAT on their network. Magically, this bit of Rube Goldberg magic (Heath Robinson, for Brits and Australians) generally works despite all the rewriting of addresses along the way.

Most of the time, you use NAT because your ISP has assigned you only a single IP address, and NAT enables you to use that address to provide access to a number of computers on your local network. Thus, the protection afforded by NAT is more of a side effect than its primary function.

If someone tries to attack a network protected by NAT, only the gateway is exposed. A gateway may have some vulnerability, but gateways are typically more capable than computers of resisting attacks because their software is so simple and they don't have many ports open—possibly none at all. Because gateways don't do that much, it's hard to hijack them. Some gateways monitor and even log these attacks for later forensic analysis.

Some people call NAT a "passive firewall," and many manufacturers that advertise gateways with firewalls are really offering only NAT. However, even NAT offers fairly significant protection.

Pushing through NAT

If you want to provide access to a computer on your internal network through a NAT gateway, you can open up a specific port, a process called *pass-through, port forwarding,* or *port mapping.* In essence, you're saying, "All traffic to my single NAT-protected IP address on port 25 should be directed to this computer on my network instead of being ignored by the gateway."

From a security standpoint, opening up specific ports makes more sense than a similar feature, called *DMZ Host,* in which all external traffic is directed through the gateway to a particular local computer. Using a DMZ Host is like using port mapping with all ports; it's easier to set up, but the DMZ host loses all protection from the NAT gateway.

Warning! *NAT doesn't protect you from other users on the same network, such as in a coffeehouse or hotel. Some networks enable "client isolation" to keep traffic from passing between devices connected to the network for this reason.*

NAT Security Gone in IPv6

The concept of a NAT changes entirely with a new form of IP addressing called *IPv6* (version 6); our current system is *IPv4* (version 4), which has been in use for decades. With IPv6, NAT is no longer needed, because the architecture of the network allows many trillions of unique addresses, along with improving how mobile devices maintain the same address even on different networks.

Apple accidentally revealed how dangerous this could be with a default setting in the first firmware released with the 802.11n AirPort Extreme Base Station in February 2007. IPv6 can be "tunneled" inside IPv4 networks, and the Internet has public endpoints that let you connect IPv6 networks. The Apple base station was set to connect to these tunnels automatically and to expose any IPv6 computers on the local network. Mac OS X has IPv6 built in and enabled by default, too.

Apple completely changed IPv6 tunneling and routing in the gigabit Ethernet version of the base station released in August 2007; the older model can be protected as described above, but the newer one has even finer firewall controls, and is set out of the box to be secure against IPv6 intrusion.

Cracking such a security hole before Apple's changes would have taken some real work, but without NAT in place, we'll have more vulnerability if networks aren't designed with that openness in mind.

ENABLE AN ACTIVE FIREWALL

An active firewall monitors all data entering and leaving a computer or network. Firewall software can be installed on individual computers or on a network gateway or router. Active firewalls examine inbound and outbound data and allow particular bits (and sometimes alert you) if blocked data matches certain criteria. Inside your network, using a firewall so your network's services are open only to local computers is a fine way to discourage ne'er-do-wells from wreaking havoc.

In an active firewall, you can choose to block or pass only certain protocols, only connections that use specific port numbers, IP addresses, or only specific users. In larger networks, you can combine

user authentication with a firewall to ensure that only certain people can carry out certain tasks on the network.

More advanced firewall software identifies patterns of data, and when it recognizes an attack pattern in progress, locks out the IP address the data is coming from, and optionally alerts you. Extremely expensive network firewall hardware can recognize thousands of these attack patterns.

Routers can come with *stateful packet inspection* (SPI) as an option. With SPI turned on, the router doesn't allow just random traffic in and out. Rather, it knows how connections should be set up and maintained, and rejects traffic that doesn't conform to this. There's an even more advanced form of SPI called *deep packet inspection* (DPI), more likely found in firewall software and servers, which can look at the contents of packets to make sure the data destined for a given type of service is well formed. An SPI or DPI option in a firewall can deter a lot of the nonsense that can crash or exploit computers.

Many firewalls also let you set access rules that vary by day of week and time of day. Thus, when you're paying attention to the network, it can operate at a lower level of security. This makes it easier to perform routine tasks that otherwise might be tedious with the firewall in place.

Practically every wireless gateway we've looked at claims to include a built-in firewall, although they generally just mean that they use NAT, not that they have active firewall capabilities. Refer to your manual for details on how to configure your gateway's firewall.

If you're roaming, or want more granular control, you use personal firewall software on individual computers:

- Windows users could try the built-in firewall in Windows XP, Vista, and 7 (which we discuss ahead in instructions for each version of Windows), but we recommend the more full-featured ZoneAlarm Pro, a powerful but easy-to-use package that's cheap and well supported; there's an awfully good free version from Zone Labs, too. http://www.zonelabs.com/
- Mac OS X users who want fine-grained control, which is not available in the included firewall, should consider DoorStop from Open Door Networks, or Sustainable Softworks' IPNetSentryX.

http://www.opendoor.com/doorstop/
http://www.sustworks.com/site/prod_sentryx_overview.html

Firewall Configuration Pointers

When configuring a firewall, the standard approach is to deny all inbound access, allow all outbound traffic (along with incoming responses to that outbound traffic), and then open specific holes in the firewall. That way, it's much easier to figure out what's happening in an attack, since the set of possible ways through the firewall is small. The only downside is that you must spend time determining which ports to open for unusual programs.

Do you use Roxio Retrospect or Netopia's Timbuktu Pro?
If so, you might go crazy troubleshooting connection problems with gaining access to certain remote machines, which are usually caused by the firewall. Find out which ports to open by reading the FAQs at http://www.roxio.com/enu/support/retrospect/default-mac.aspx and http://www.netopia.com/software/products/tb2/.

Enable the firewall in Mac OS X 10.3 and 10.4

In Mac OS X 10.3 and 10.4, Apple included a reasonably configurable inbound firewall. Follow these steps to turn it on:

1. Open System Preferences, and click Sharing to open the Sharing preference pane.
2. Click the Firewall button, and click Start.
3. Select the checkbox next to any services that need outside access.
4. Click New or Edit to modify the services listed if necessary.

If you want further protection, click the Advanced button and check Enable Stealth Mode, which puts Mac OS X into a kind of "silent running" state that's even harder to penetrate.

Enable the firewall in Mac OS X 10.5 and 10.6

Apple opted to swap a typical firewall that works by ports for an application-oriented firewall in 10.5 Leopard and later. This firewall lets you control which applications allow access from outside the computer.

Note: See the *TidBITS* article "Leopard Firewall Takes One Step Forward, Three Steps Back," by Rich Mogull, for more details (http://db.tidbits.com/article/9294).

To turn on this application firewall, follow these steps:

1. Open System Preferences and click the Security icon.
2. Click the Firewall button.
3. Click the lock  icon in the lower left corner, and enter an administrative account name and password.
4. Click Start to fire it up.
5. For more options, click the Advanced button in the lower right. From here, you can:
    - Hide the machine from the outside world by checking Block All Incoming Connections. This option may seem extreme, but it's a great way to go commando on an untrusted network. Network messages and VPNs using IPsec still work.
    - Allow or disallow incoming connections to specific applications.
    - Enable stealth mode, which prevents your computer from responding to requests that try to see if a device is active at a given IP address.

Enable the Windows XP firewall

Commercial firewall software may give you more options and a better interface, but the built-in firewall in Windows XP will do the job.

A wise decision: *Windows enables the Windows firewall by default, and warns you if and when there's no firewall protection enabled—such as if you turn off the firewall temporarily and forget to turn it back on.*

In Windows XP, follow these directions to set up a firewall:

1. Open Control Panel, and then double-click Network Connections.
2. Select the connection you want to secure (you can repeat this for multiple connections).

3. In the left pane, click Change Settings of This Connection under the Network Tasks area.
4. Click the Advanced tab.
5. Check Protect My Computer and Network by Limiting or Preventing Access to This Computer from the Internet.

Enable the Windows Vista firewall

In Windows Vista, follow these steps to set up a firewall:

1. Open Control Panel, and then click the Security icon.
2. Click Turn Windows Firewall On or Off.
3. If prompted by User Account Control, click Continue.
4. Click the On button, and click Apply or OK to enable.

Enable the Windows 7 firewall

In Windows 7, follow these steps to set up a firewall:

1. Open Control Panel, and then click the Windows Firewall icon.
2. In the left navigation links, click Turn Windows Firewall On or Off.
3. In the Customize Settings for Each Type of Network dialog, click Turn On Windows Firewall for each of your network locations.
4. Click OK.

Secure Small Office Wi-Fi

Small businesses used to be unable to afford the equipment, software, or know-how needed to put in place the information technology (IT) infrastructure of a big organization. That's changed. A lot of technology that was available only in server hardware or software that cost thousands or even tens of thousands of dollars has now been scaled down and made affordable. This change lets offices of as few as five people increase the security of their in-house Wi-Fi networks and the security of their mobile users using wireless networks in cafés, hotels, and at home, all without breaking the bank. In some cases, a small office of even one or two people can take advantage of these tools.

While many small offices use Wi-Fi, it's clear that many more small offices have avoided Wi-Fi or used it in limited fashions because of security concerns: You may have no dedicated IT personnel or pay consultants by the hour for assistance, and thus have been leery of installing a wireless network that could expose confidential business or customer information. In this chapter, we offer some simple and cost-effective suggestions that won't send you scuttling to the classifieds to hire an expensive staffer. Instead, you might be able to set up a secure, small-office wireless network by yourself, or at least spend only a few hours with a consultant.

THREE SECURITY OPTIONS

In this chapter, we look at three ways that you can secure your local Wi-Fi network against snoopers and unauthorized access:

- Use a Shared Key: With a shared key, security is achieved through a single encryption key that is shared among all users on the network. (We recommend a WPA2 key, because WEP is too weak.)
- Use WPA2 Enterprise Logins: In this case, an authentication system that provides each user with a unique user name and password and then assigns a unique network encryption key to each user each time they log in.

- Use a VPN: With a virtual private network (VPN) server, all data that passes over exposed parts of a Wi-Fi network is encrypted. VPNs are useful for protecting both wired and wireless traffic on your local network and traffic from mobile users on the road.

Tip: For visitors, you might set up access in such a way that you don't have to provide a log in or a key. See Enable Guest Access.

USE A SHARED KEY

With a shared key, each user on the Wi-Fi network has the same encryption key, which means that each user can—if they run special software—see all the data that every other user on the network is passing back and forth.

For a small office, this is unimportant: your main goal is to keep unauthorized users from gaining access to your network, not to keep one network user from viewing another's traffic. (If that latter issue is important, consider application-level encryption so that SSL/TLS or IPsec is used to encrypt data as it passes between, say, a file server and a user; or use an authentication system described next.)

Although there are many provisos, if you rely on a shared key, you should follow these principles.

- **Use WPA2:** Small offices shouldn't expose themselves to risk by using WEP (see Watch out for WEP Encryption for more on this issue). Rely on the strong AES-CCMP key type that's part of WPA2.

  Replace equipment that doesn't support WPA2: *You may have to replace older gear to get the appropriate level of security. Fortunately, almost every piece of equipment sold starting in 2004 (and some from 2002 and 2003) supports WPA2.*

  If you have an important piece of hardware or an older computer that simply can't be replaced, and its Wi-Fi adapter can't be updated, see if you can attach a USB-based Wi-Fi adapter that includes drivers for WPA2. Another option for devices with Ethernet jacks would be to use an Ethernet-to-Wi-Fi adapter.

- **Choose a strong key:** WPA2 suffer from a weakness if you choose a short key comprised of words found in a dictionary. Choose long

encryption passphrases that mix letters, numbers, and punctuation of at least 20 characters.

- **Don't let users see the keys:** In most operating system software, a network administrator or other trusted employee can type in the shared key. This shared key can't be seen by users who lack administrator permissions. Specialized software might let them retrieve it, but only a determined internal hacker would try that.

Keeping a Key Private Avoids Changing It

We're not telling you to distrust your employees, but if they all know the encryption key then one of them could accidentally reveal it, or a disgruntled ex-employee could lurk outside the office and gain access. With a known, shared key, you should change it regularly or at least whenever an employees leaves or is fired. Whenever you change the shared key, every computer on the network must have its key changed at the same time to retain access.

USE WPA2 ENTERPRISE LOGINS

There's an easier, but more complicated, way to provide access to a Wi-Fi network for employees, without managing a shared key for the entire network. A system known as *802.1X,* another IEEE standard, works with user accounts and WPA2 to provide a unique encryption key for each user on each network login. In that combination, it's called *WPA2 Enterprise.*

WPA2 Enterprise also adds accountability: you know which users are on the network and can lock out particular users. More advanced systems let you set policies that might allow only certain users during working hours or provide guest credentials for 30 days.

Note: The 802.1 task group is one of the overarching networking groups in the IEEE, and the capital X in 802.1X indicates that it is the full standard. 802.1X works with Ethernet as well as Wi-Fi, but requires more advanced Ethernet switches than are found in small offices to have the same benefit of locking out access.

We first give you some background in the underlying principles of this kind of authentication, and then we provide practical advice on implementing it without great expense.

Understand 802.1X

The 802.1X protocol underlying WPA2 Enterprise is essentially a way of putting a sentry in front of a network. The sentry prevents network access until a client proves itself worthy by providing credentials that can range from a simple user name and password up to fingerprint or hand geometry, confirmed by a biometric control system.

802.1X defines three roles: a client, which is called a *supplicant*; an access point, which acts as an *authenticator* or gatekeeper; and a user database server or *authentication server,* which confirms a user's identity (**Figure 13**).

802.1X Client Support

Microsoft's Wireless Network Connection tool in Windows XP, Windows Vista, and Windows 7 supports 802.1X, as does the Internet Connect application in Mac OS X 10.3 Panther and 10.4 Tiger, and the Network preference pane (Advanced settings) in 10.5 Leopard and 10.6 Snow Leopard.

Cisco and Juniper Networks offer commercial supplicants that work on many platforms, including many versions of Windows, Unix, Linux, and Mac OS, along with handheld devices from Zaurus and Palm, or that use the Windows Mobile operating system.

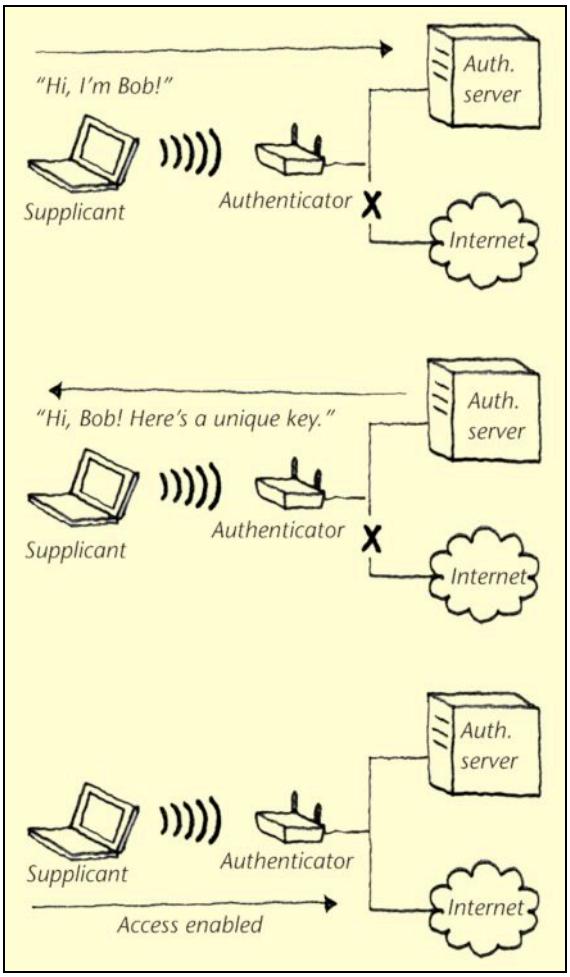


Figure 13: Let's look at how the supplicant, authenticator, and authentication server work together.

(1) First, a supplicant, having associated with an access point, sends its login credentials (like a user name and password) to the authenticator, which passes them to the authentication server. The access point doesn't give the supplicant Internet access at all.

(2) Next, the authentication server confirms the supplicant's identity and responds, again via the authenticator, providing the supplicant with a unique encryption key for using the local network.

(3) Finally, with the encryption key in hand and its credentials confirmed, the supplicant can pass traffic through the authenticator out to the Internet (or the local network, as the case may be).

When a user (supplicant) wants to join the network, she uses an 802.1X client to log in (**Figure 14**). This client is included or available for all current operating systems and handheld organizers; see 802.1X Client Support, previous page. The authenticator (access point) receives a request for a login from the supplicant. The supplicant can't access any

network resources at all except a single network port devoted to handling 802.1X logins.

The authenticator sends the user's credentials to the authentication server, often the same one used for regular network logins to access file servers and other network resources (Step 2 in **Figure 13**, previous page). The authentication server tells the authenticator that the login is valid, and the authenticator opens up network access to the supplicant (Step 3 in **Figure 13**).

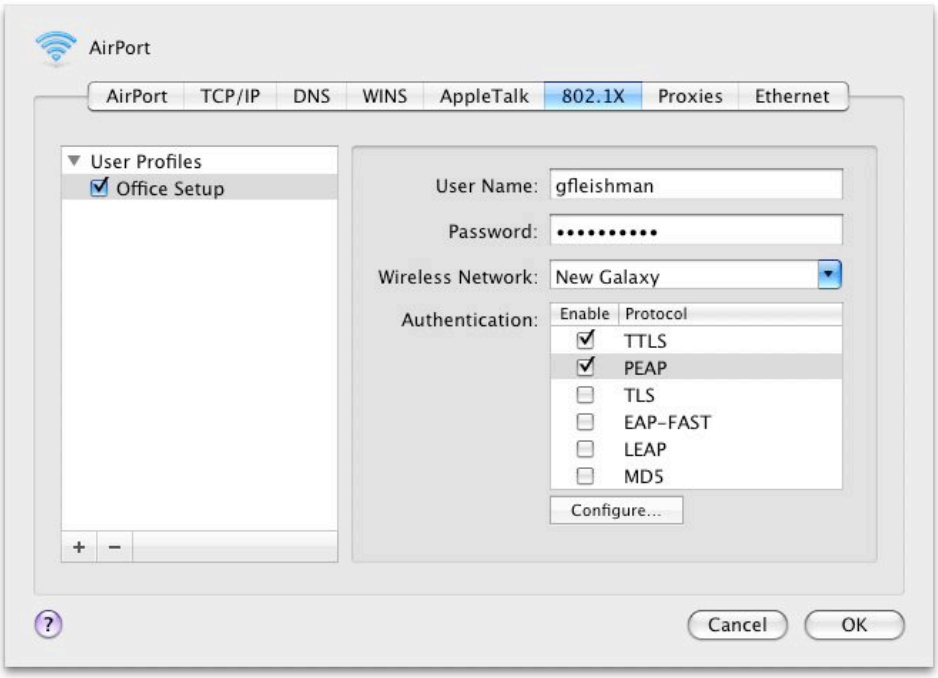


Figure 14: An 802.1X login, shown here in Mac OS X 10.5 Leopard's Network preference pane.

The 802.1X transaction has already kept the network safe from those who don't have passwords. But it gets better. Once the user has been verified, the authentication server can provide a unique WPA2 encryption key to that specific client. This way, each client on the network can have its own key.

A Small Flaw in WPA2 Enterprise

A researcher at a firm that sells network intrusion software publicized a "hole" in WPA2 Enterprise in July 2010. The hole, found in the 802.11-2007 specification itself, allows any client on a network to send messages that appear to be coming from an access point to which all the clients are connected.

To exploit this hole, however, an attacker needs a legitimate account on the network, and can use only attacks that work when broadcast packets—those sent ostensibly by the access point to every client at once—are involved. The risk of exploit seems very low, and WPA2 keys are not at risk.

Because the 802.1X process doesn't include encryption for protecting the login negotiation, there are several standards in place to layer login protection on top of basic WPA2 Enterprise. The most popular flavors include:

- **EAP-TLS** (EAP Transport Layer Security, a synonym for SSL/TLS): This technique requires that you install unique digital certificates on every computer on a Wi-Fi network, which is worthwhile only in situations that require high security.
- **PEAP** (Protected EAP): Functionally equivalent to EAP-TTLS, but built into Windows XP and later. It's the dominant flavor out there.
- **EAP-TTLS** (EAP Tunneling Transport Layer Security): This method is functionally equivalent to PEAP, but it is less widely used.

Cisco and EAP

Before these more modern flavors appeared, Cisco had its own flavor of EAP with encryption called *Lightweight EAP* or *LEAP.* LEAP is completely broken and can be cracked easily. Cisco now prefers PEAP, but the company also released a LEAP upgrade known as *EAP-FAST* (Flexible Authentication via Secure Tunneling) as a simple migration from LEAP that involves fewer network changes.

EAP-FAST still has LEAP's weakness in its least-secure mode, required for compatibility's sake. In its most secure flavor, it's no better than the more widely supported PEAP and EAP-TTLS.

Add WPA2 Enterprise to your network

The complexities of WPA2 Enterprise can be hidden from view entirely because many wireless gateways, including most inexpensive ones from Linksys and others, can work as an 802.1X authenticator. Users simply enter the domain name or IP address of the authentication server and provide a shared key that's kept secret between the access point and the server, and they're all set. This set of options is usually listed as *RADIUS* settings; RADIUS (it no longer stands for anything) is a type of user account management server.

Tip: A very small number of Wi-Fi gateways that Glenn has tested over the years have lacked RADIUS or 802.1X support. Check the manual before you buy a new gateway if you need this feature.

At one point, all 802.1X servers cost thousands of dollars whether they were standalone products or part of a RADIUS system. A few years ago, one firm developed a low-cost server that has been under active development ever since: Elektron, from Periodik Labs, is a WPA Enterprise server that runs as a separate program on an existing computer. It can tie into Windows Active Directory or Mac OS X (plain or Server) user directories.

Elektron can authenticate remote access points over the Internet as well as local Wi-Fi gateways on the same network. It runs under Mac OS X 10.4 or later and Windows XP Professional, Vista, 7, Server 2003, and Server 2008. Any standard 802.1X client works with Elektron using PEAP or EAP-TTLS—both versions are available simultaneously.

Although its price might dissuade you, Elektron is the most straightforward way to secure a network. Even with as few as 20 users, the first-year cost per user with full support is under $50—worth it for ease of administration and peace of mind (http://www.periodiklabs.com/; $750 plus $250 per year for upgrades and support, or $950 with the first year subscription included; unlimited users).

Using Windows Server or Leopard Server with RADIUS and 802.1X

Windows Server (2003 and later) and Mac OS X Server (10.5 Leopard and later) include RADIUS and 802.1X support, although it requires some expertise (more than the two solutions listed just previously) to make them work. If you already have one of those server systems running in house, see if it makes sense to enable 802.1X. The bonus is that users with accounts on the servers can use the same logins for the Wi-Fi network.

USE A VPN

We describe the nature and utility of VPNs in Encrypt All Data with a VPN. In brief, a VPN connection securely encrypts all the data entering and leaving a computer. Until a few years ago, you would have had to spend many thousands of dollars and pay for real IT expertise to purchase, configure, and maintain a VPN server or service.

Virtual private networks can improve small office network security in two basic ways:

- A VPN is essential for users who leave the security of your local network. Using a VPN is our primary recommendation for securing a long-distance connection over any network—not just from hotspots, but also from a hotel's wired broadband or from a guest connection on another organization's network.
- You can implement a VPN for wireless users connecting directly to your local network. Although this may seem like (and, in fact, may be) overkill, it provides the maximum security possible, particularly when coupled with appropriate use of WPA2.

Using VPNs on a Local Network

If you use a VPN to secure local wireless users who connect directly to your local network, you must take two steps:

- For local users on your local network to see each other's computers and network servers when using a VPN, you must host your own VPN server software or hardware, as described ahead.
- You must configure your access point to ban all inbound network traffic except that aimed at VPN ports. (Allow port 1723 for TCP and UDP for PPTP, and port 1701 for TCP and UDP for L2TP/IPsec. SSL/TLS doesn't require a special port, but uses port 443 just like a secure Web site request.)

This second step means that a VPN client can connect via the access point and to the VPN server, but all other inbound connections will be rejected. Because local clients will tunnel all traffic over that VPN connection, they have full access to your network.

VPN client software is built into Windows XP, Windows Vista, Windows 7, and Mac OS X 10.3 and later for PPTP- and IPsec-based VPNs, and you can buy a variety of VPN servers (hardware, software, and hosted services) without spending much on equipment or monthly service charges. Hosted VPN services typically include the software at no additional cost. **Table 3** summarizes the options for creating a VPN server, and we discuss each option in turn, later in this chapter.

Note: Businesses with more resources or specific needs might want to hire a corporate VPN service provider. Many companies provide such services, but prices are all over the place and some providers may require installing purchased or leased hardware on your network.

Table 3: VPN Server Options for Small Offices		
Option	**Pros**	**Cons**
Run free server software on your own server.	Free; you can configure exactly what you want	Requires more configuration and support
Run commercial server software on your own server (Mac OS X Server, Windows Server, etc.).	You can configure exactly what you want; potentially easier configuration than with free tools	Requires more configuration and support; often costs hundreds or even thousands of dollars
Subscribe to a VPN service on the Internet.	Requires little configuration or support	Monthly or yearly charge

Choose a software VPN server

Most large organizations use dedicated hardware to support hundreds, or even thousands, of VPN connections; these hardware devices rely on special chips to handle the massive computation necessary to encrypt and decrypt all network traffic. But, if you support just dozens of simultaneous users, you can use software installed on a server that might already also act as a Web or email server. If you must combine VPN and other services, monitor performance carefully to ensure that you're not being penny-wise and pound-foolish by destroying your Web performance to run your VPN.

It's conceivable that you might have the time, money, and know-how to run your own software-based VPN server, but be sure to consider the ongoing support time and effort. Here are some likely choices:

- **OpenVPN:** This free, open-source project offers client and server components that use SSL/TLS as their basis of starting a VPN tunnel, with any of a large number of encryption algorithms up to the highest level of publicly available government-grade encryption. It's available for many platforms, often in an installer form. http://openvpn.net/
- **Windows Server 2008:** This high-end Microsoft product, which costs thousands of dollars, might already be installed on your network if you're running a Windows shop and need the services that

it provides. Windows Server 2008 fully supports PPTP and L2TP/IPsec VPNs. Its advantage is that there are many Microsoft certified technicians who can consult on configuring and maintaining it. The downside, of course, is price: a 20-user version costs $1,500 to $4,000, depending on features.
https://www.microsoft.com/windowsserver2008/

- **Mac OS X Server 10.6:** Apple includes PPTP and L2TP/IPsec VPN servers in this package. The software runs only on Macintosh hardware, and includes a host of other network services, just like Windows Server 2008. The difference? Apple's server allows unlimited users and costs just $499.
http://www.apple.com/server/macosx/

Choose VPN hardware

When people talk about VPN hardware, they usually are referring to systems like the Securepoint Firewall & VPN Server Appliance (http://www.securepoint.cc/), which starts at thousands of dollars for a handful of users. At one point, we spotted a trend toward the development of inexpensive, small office–oriented VPN servers, typically embedded in broadband gateways. Unfortunately, those products must not have panned out because the category has disappeared. Buffalo and Linksys had notable products that were inexpensive and full featured; unfortunately, they're gone.

Choose a VPN service

If the VPN options listed earlier made your pocketbook ache or your eyes glaze over, consider another alternative for roaming users: subscribing to a service that offers VPN connections from wherever the user is to a secure network operations center (NOC) elsewhere on the Internet.

Typically, the least secure link in any connection is the local network: the sniffed or penetrated Wi-Fi or wired network over which traffic proceeds unencrypted out to the Internet connection. Once traffic is on the broader Internet, it's much less likely that any snoop would be able to intercept it or associate it with you—in essence, your traffic becomes more secure once it leaves the local network that you're connected to. By creating a secure client-to-NOC connection, you eliminate the most common places that people could sniff or intercept network traffic.

Secure enough? *These VPN services certainly increase your security level significantly, but if you need complete security back to your local network, you need a software- or hardware-based VPN server running on your local network. Otherwise, it's possible that your traffic could be snooped after it leaves the VPN service's NOC.*

Several companies offer VPN services for hire that are specifically designed for hotspot users. Here are two that have been around for a few years:

- **WiTopia's PersonalVPN:** This service offers OpenVPN software for an SSL/TLS connection or a PPTP connection. Pricing is yearly only: $39.95 for PPTP, $59.99 for SSL/TLS, and $69.99 for both. They have Mac- and Windows-specific software (http://www.witopia.net/).
- **PublicVPN.com:** This service costs $6.95 per month or $69.95 per year, and it deploys L2TP/IPsec and PPTP VPN tunnels and works with any standard L2TP/IPsec or PPTP client (http://www.publicvpn.com/).

Glossary

802.11: A set of wireless networking standards developed by the IEEE engineering standards body that include lettered protocols (802.11a, 802.11b, 802.11g, and 802.11n) that define the speed and spectrum used on a network, security (802.11i), and other parameters.

802.11i: A security standard from the IEEE intended to replace WEP. 802.11i encompasses two unique encryption algorithms: TKIP and AES-CCMP. 802.11i isn't the same as Wi-Fi Protected Access (WPA); rather, WPA and its successor, WPA2, were derived from different stages of the 802.11i task group's work. The 802.11i standard was rolled into 802.11-2007.

802.11-2007: This standard is a "roll-up" of previous 802.11 standards that have been cleaned up and put into one spec. The 802.11i security standard is now referred to as part of the 802.11-2007 roll-up.

802.1X: An authentication specification that allows a client to connect to a wireless access point or wired switch but prevents the client from gaining access to the Internet until it provides credentials, like a user name and password, that are verified by a separate server. In 802.1X, there are three roles: the supplicant (client), authenticator (switch or access point), and authentication server. WPA2 Enterprise is a version of 802.1X that works over Wi-Fi and provides only WPA2 encryption keys to clients.

access point: The hub of a wireless network. Wireless clients connect to the access point, and traffic between two clients must travel through the access point. Access points are often abbreviated to AP in industry literature, and you may also see them referred to as "wireless routers," "wireless gateways," and "base stations."

AES-CCMP: This extremely strong encryption standard that's part of WPA2 and 802.11-2007 is comprised of two elements. First, there's *AES,* which stands for Advanced Encryption System. Second, there's *CCMP,* which is a complex abbreviation standing for Counter-mode CBC-MAC Protocol; *CBC-MAC* stands for Cipher Block Chaining-

Message Authentication Code. You don't need to know this unless you attend highly geeky cocktail parties.

AirPort Extreme: Apple's marketing name for its 802.11g and 802.11n wireless networking technology (which includes the AirPort Express Base Station and the built-in network adapters in Macs).

AirPort: Apple's original marketing name for its 802.11b wireless networking technology. Today, AirPort refers to Wi-Fi-compatible wireless networking in general for Apple.

APOP: A protocol for protecting email passwords used with POP. APOP stands for Authenticated Post Office Protocol.

authenticate: The process of confirming the identity of someone connecting to a network.

authentication server: A back-end database server that confirms the identity of a supplicant to an authenticator in an 802.1X-authenticated network.

authenticator: The gatekeeper role in an 802.1X-authenticated network. You can think of the authenticator as a gatekeeper; access points and Ethernet switches can act as authenticators.

base station: See *wireless gateway*.

certificate authority: A trusted third party that can assure the identity of others when using security systems like SSL/TLS. A certificate authority registers the digital identity of a site or individual, and lets you confirm manually or automatically that someone you're interacting with—say, over a secure Web connection—is who he appears to be.

certificate: A computer-readable authentication credential. Certificates are typically signed by other people or certificate authorities to guarantee their authenticity.

clear text: Sensitive information like passwords sent across a network without encryption. Clear text is also commonly referred to as "in the clear."

closed network: A wireless network that doesn't advertise its network name.

DMZ: A feature in a NAT gateway that lets you expose a machine on your internal network to the outside Internet. DMZ nominally stands for *demilitarized zone,* and is sometimes also called "virtual server." It's basically port mapping for all available ports.

EAP: A standard form of generic messaging used in 802.1X, among other places. EAP stands for Extensible Authentication Protocol.

EAP-TLS: Used to create a secured connection for 802.1X by pre-installing a digital certificate on the client computer. EAP-TLS stands for Extensible Authentication Protocol-Transport Layer Security.

ESSID: Extended Service Set Identifier. See *network name*.

fingerprint: A short sequence of characters you can send someone so she can verify that a specific public key is actually your public key.

firewall: A network system that blocks malevolent or unauthorized traffic that might endanger the computers on your network.

firmware: The internal software that runs dedicated hardware devices. Upgrades to firmware are often necessary to fix problems or support new security options.

gateway: See *wireless gateway*.

hotspot: A place where you can connect to a public wireless network.

HTTP: The network protocol used by the Web, although it's also now used for many other services. HTTP stands for Hypertext Transfer Protocol.

HTTPS: The SSL/TLS-protected version of HTTP.

IMAP: A common and newer alternative to POP to receive email from a mail server and manage stored mail. IMAP defaults to storing mail on a server, in contrast to POP, which retrieves mail to your computer and then deletes it from the server. IMAP stands for Internet Message Access Protocol.

IPsec: One of the main protocols used for VPNs. IPsec stands for IP security. IPsec creates the session in a VPN connection, while L2TP carries the encrypted data between two points. The two together are called L2TP/IPsec or L2TP-over-IPsec.

key server: An Internet-based server that lets you look up other people's public keys.

L2TP: See *IPsec*.

local area network: The computers at your site, connected via Ethernet or Wi-Fi. Local area network is often abbreviated to LAN. Compare local area networks with wide area networks.

MAC address: The unique address assigned to every wireless and wired Ethernet network adapter. MAC stands for Media Access Control. Despite the fact that assigned MAC addresses are all unique, it's possible to assign one device's MAC address to another device. There are various reasons (to circumvent ISP restrictions) to clone MAC addresses.

NAT: A network service that makes it possible to share a single IP address with a network of many computers. NAT stands for Network Address Translation. Since a NAT gateway exposes only a single IP address to the outside Internet, it's useful for security, and some manufacturers may call it, somewhat incorrectly, a "firewall."

network adapter: The card or built-in hardware used in a computer or handheld device to connect to a network, whether wired or wireless.

network name: The name you give network; it appears when a wireless client displays available networks. Many manufacturers use the terms "SSID" or "ESSID" in place of network name.

open network: A wireless network that is broadcasting its name. Technically, the fact that a network is broadcasting its name is unrelated to whether or not it employs WEP or WPA encryption, but informally, an open network is often considered one that can be used by anyone, without a password.

pass-through: See *port mapping*.

personal certificate: A certificate you generate for use with SSL/TLS that doesn't have a certificate authority behind it. Personal certificates, also known as "self-signed certificates," aren't vouched for by a certificate authority, but they're good enough in cases where you're working with private SSL/TLS-enabled systems.

PEAP: A method of securing an 802.1X session within an encrypted tunnel to protect credentials used for logging in. PEAP stands for Protected Extensible Authentication Protocol.

PGP: A technology and set of programs for encrypting data. PGP stands for Pretty Good Privacy.

plain text: See *clear text*.

POP: The most common way of receiving email from a mail server on the Internet. POP defaults to retrieving email to your computer and deleting it from a server, in contrast to IMAP, which typically stores mail on the server. POP stands for Post Office Protocol.

port forwarding: See *port mapping*.

port mapping: The act of mapping a port on an Internet-accessible NAT gateway to another port on a machine on your internal network. Port mapping enables you to run a public Internet service on a machine that is otherwise hidden from the Internet by your NAT gateway. Other names for port mapping include “port forwarding,” “pass-through,” and “punch-through.”

port: Either a physical jack on a network device or a way of identifying the type of data being sent in an Internet connection. Every Internet service has its own port number.

PPTP: A Microsoft-developed protocol used for VPNs that is easily used from within Windows and Mac OS X. PPTP stands for Point-to-Point Tunneling Protocol.

pre-shared key: A passphrase used to protect your network traffic in WPA. Some manufacturers use the term “pre-shared secret” instead.

private key: The key you keep secret in public-key cryptography systems. You use your private key to decrypt encrypted data sent to you by other people, who used your public key to encrypt it. You also use your private key to sign email messages; your recipients then use your public key to verify your signature.

public key: The key you give out to the world in public-key cryptography systems. Other people use your public key when sending you encrypted data, which you can then decrypt with your private

key. You also use other people's public keys to verify the authenticity of mail messages they've signed with their private keys.

relaying: The act of sending email through your mail server when you're not connected to your local network. Spammers take advantage of mail servers that allow unrestricted relaying.

script kiddies: Wanna-be crackers who don't have the technical skills to break into computers on their own, so they use canned cracking software.

self-signed certificate: See *personal certificate*.

SMTP AUTH: A command in the SMTP protocol that provides identification to an SMTP server, so it will accept outgoing mail from you. SMTP AUTH is essentially authenticated SMTP.

SMTP: The protocol for sending email on the Internet. SMTP stands for Simple Mail Transfer Protocol.

spoofing: Replicating one device's MAC address onto another to work around restrictions that prevent only particular MAC addresses from connecting to a network. Also sometimes called "cloning."

SSH: A security system that lets you create encrypted tunnels for any Internet protocol via port forwarding. SSH stands for Secure Shell.

SSID: Service Set Identifier. See *network name*.

SSL: A security protocol that secures Internet transactions at the program level. SSL, which stands for Secure Sockets Layer, is widely used in Web browsers to protect credit card transactions, for instance. SSL is a component in EAP-TLS (Extensible Authentication Protocol-Transport Layer Security). SSL is also increasingly used for VPNs. SSL is the predecessor to TLS, but because the two standards are very similar and because TLS is backward compatible to the last version of SSL, the two terms are used interchangeably. Now called SSL/TLS.

stumbler: A software program that looks for available wireless networks in range and reports information about them.

supplicant: The client role in an 802.1X-authenticated network.

TKIP: An encryption key that was developed for the 802.11i standard part of WPA and WPA2. TKIP stands for Temporal Key Integrity

Protocol. It's nominally weaker than the government-grade AES-CCMP, and it is on its way out.

TLS: Transport Layer Security, SSL's successor. See *SSL*.

VPN: A method of creating an encrypted tunnel through which all traffic passes, preventing anyone from snooping through transmitted and received data. VPN stands for virtual private network.

WEP: An encryption system for preventing eavesdropping on wireless network traffic. WEP stands for Wired Equivalent Privacy. WEP is easily broken, and it has been replaced by WPA2.

Wi-Fi: Wi-Fi is a certification mark managed by a trade group called the Wi-Fi Alliance. Wi-Fi certification encompasses numerous standards, including 802.11a, 802.11b, 802.11g, 802.11n, WPA, WPS, WPA2, and more, and equipment must pass compatibility testing to receive the Wi-Fi mark.

Wi-Fi Protected Access: See *WPA*.

wireless gateway: We use this generic term to differentiate between a simple access point and a more-capable device that can also share an Internet connection, serve DHCP, and bridge between wired and wireless networks. You may also see the term "wireless router," or "base station."

wireless network adapter: See *network adapter*.

WPA: A modern encryption system for preventing eavesdropping on wireless network traffic that solves the problems that plague WEP. WPA stands for Wi-Fi Protected Access. WPA is a subset of the IEEE *802.11i* security standard, but WPA was released before that standard was finalized. It includes the *TKIP* encryption key.

WPA2: WPA2 is a more-advanced version of *WPA* that includes additional security measures; it's a complete implementation of the IEEE *802.11-2007* specification's security standards. WPA2 offers the *AES-CCMP*. encryption key type.

WPA/WPA2: This mixed-mode security option is found on many base stations that allows either TKIP or AES-CCMP keys to be used. It can be used to provide backward compatibility with older devices that lack WPA2 support.

WPS: A system that lets computers and other Wi-Fi–equipped devices join a Wi-Fi network without any password, or by using a short PIN code that's generated by the Wi-Fi software, and which you enter once. WPS stands for Wi-Fi Protected Setup.

zombie: A computer that has been taken over by a malevolent program that uses it to attack other computers.

Appendix A: Password Advice

We talk blithely about passwords throughout this book, and more generally, passwords are all around us. But are you picking good passwords? A bad password can be cracked easily, often just by guesswork. So here's some advice.

We can also recommend our colleague Joe Kissell's book on this subject: *Take Control of Passwords in Mac OS X*. Although it is focused on the Mac, it has broad advice applicable to everyone.

GENERATE THREE PASSWORDS

Because it's nearly impossible to remember different passwords for every possible service, we recommend using three different passwords. If you restrict yourself to three passwords and always use the same email address or user name, the likelihood of forgetting your access information for any given site or program is low:

- **Low-security:** Create a standard low-security password that's simple and easily remembered. Since it's low-security, make sure to use it *only* for Web sites that don't store personal information about you (such as your address, birth date, or credit card number). In essence, this password protects only your online identity; if someone were to guess it, they could pretend to be you in a discussion forum or the like. Don't use this password for email accounts!
- **Medium-security:** For Web sites and accounts where some personal data is at risk, create a medium-security password. It will be harder to type, since it should include upper- and lowercase letters, numbers, and punctuation.
- **High-security:** Everyone should have a highly secure password that is long, hard to type, and impossible to guess. Use it for accounts, like your bank and PayPal, where money is involved, and for programs that store other passwords. Using a longer password

won't prevent it from being stolen via an unprotected wireless transaction, but realistically, most passwords are stolen by being guessed or because someone wrote them on a Post-it note. You're also unlikely to need this password on a site that wouldn't secure the transaction, and in fact, don't use this password on sites that don't secure transactions.

For this password, you may need a few variants. For example, one bank may insist on a password with at least a dozen characters, but another may limit you to no more than eight characters. Further, one bank may insist that you use at least one non-numeric, non-alphabetic character (like an @ or #), but another might not permit those characters. So, you might have one password: `MyBlgC@tl3l3l` that can degrade to `MyBlGCat`, without you losing your mind.

1Password Can be Helpful

Glenn uses 1Password to create, manage, and protect his Web and other passwords on multiple Macs and his iPhone and iPad. He uses the 1Password desktop application as well as the 1Password app on his iPhone and iPad:

- http://agilewebsolutions.com/products/1Password ($39.95)
- http://itunes.apple.com/us/artist/agile-web-solutions/id285897621 ($9.99 or $14.99 for Pro)

1Password can generate passwords of whatever length and using whichever parameters a sites wants. It can also synchronize data between devices through a Dropbox account.

Tip: Sites like PayPal and its parent company eBay offer optional two-factor authentication for better security. One "factor" is your password; the other is a small device that fits on a key ring that generates a unique password based on the time elapsed since it was synced with a server. Without the key device, your password is useless, rendering password interception pointless.

Sharing Passwords Safely

Our recommended three-tier password approach may fall down in a situation where a family or business must share passwords needed to access joint financial information such as bank accounts, insurance policies, and so on. Unfortunately, many financial institutions restrict user name or password lengths in such a way that prevent you from using a single high-security password.

In such a case, you must record user names and passwords in a place where everyone who needs to can access them. Some people may use a notebook, but it's easily stolen or lost. A more secure approach would be to use a program like the previously mentioned 1Password. You can enter all your user names and passwords into 1Password, where they're encrypted and protected with yet another password, which you'd make sure was known to everyone who needed to know.

Be sure to discuss this situation with people in your family, since you never know when it could be important to access someone else's secure accounts (legitimately, of course) in case of severe illness or injury, or in the case of an elderly relative, mental incapacitation.

LEARN TO CREATE A HIGHLY SECURE PASSWORD

Many security experts now recommend that when you enter a really long password–often called a *passphrase* because it's composed of separate words or items—that you think of something memorable to you that no one else would know: a lyric of a song, for instance. Instead of choosing 754!#%kdja you might enter shall I compare thee 2 a summer's day?!.

The length makes it harder to crack while still rendering it memorable to you. The extra punctuation at the end (or wherever you choose to put it—even extra spaces between words would help) helps stymie efforts to crack the passphrase against a database of all poems and song lyrics.

While many systems require a short password, others like WPA2 and PGP allow dozens or even hundreds of characters. Systems with longer passphrases typically only need you to type them once per computer (to set up a secure Wi-Fi network) or once per session (each time you reboot, for instance).

Mac OS X Keychain

If you store passwords in the Keychain in Mac OS X, note that by default the Keychain uses the same password as your login, which might not be one of your more secure passwords since you must enter it so frequently. But, with the Keychain Access utility (found in /Applications/Utilities), you can change the password for your Keychain to something more secure; just choose Edit > Change Password for Keychain "username" and enter a new password.

The version of Keychain Access in 10.4 Tiger and later also has a tool that shows how secure different passwords are. Choose File > New Password Item, and then click the key icon to the right of the Password field in the New Password Item dialog. The Password Assistant window that appears suggests passwords of varying strengths, or tests ones you enter.

You needn't launch Keychain Access to use Password Assistant. Download a simple program from codepoetry that lets you invoke Password Assistant by itself (http://www.codepoetry.net/products/passwordassistant).

In theory, you could generate highly secure random passwords with Keychain Access every time you needed a password for a Web site (1Password offers this feature too). That way, if a password for one site is made public, it won't compromise your account on any other sites. The downsides are that it's more of a fuss and you're putting all your eggs in one basket, requiring an extremely solid backup strategy.

About This Book

Thank you for purchasing this Take Control book. We hope you find it both useful and enjoyable to read. We welcome your comments at tc-comments@tidbits.com.

EBOOK EXTRAS

You can access extras related to this ebook on the Web. Once you're on the ebook's Take Control Extras page, you can:

- Download any available new version of the ebook for free, or buy a subsequent edition at a discount.
- Download various formats, including PDF and—usually—EPUB and Mobipocket. (Learn about reading this ebook on handheld devices at http://www.takecontrolbooks.com/device-advice.)
- Read postings to the ebook's blog. These may include new information and tips, as well as links to author interviews. At the top of the blog, you can also see any update plans for the ebook.
- Get a discount when you order a print copy of the ebook.

ABOUT GLENN



Glenn Fleishman has written for hire since 1994, starting with *Aldus Magazine*. He contributes regularly to *Macworld, Ars Technica,* the *Economist, BoingBoing*, and the *Seattle Times*. He's also a senior editor at *TidBITS*.

Glenn spends much of his time writing about wireless networking. He has written several Take Control books, including *Take Control of Your 802.11n AirPort Network*. He edits Wi-Fi Networking News (http://wifinetnews.com/). He lives in Seattle with his wife and two sons. His older son's first word was not "Wi-Fi"; it was "book."

ABOUT ADAM



Adam C. Engst is the publisher of *TidBITS,* one of the oldest and most respected Internet-based newsletters. He has written numerous technical books, including the best-selling *Internet Starter Kit* series, and many magazine articles (thanks to Contributing Editor positions at *MacUser, MacWEEK,* and now *Macworld*).

Adam's innovations include the creation of the first advertising program to support an Internet publication (in 1992), the first flat-rate accounts for graphical Internet access (in 1993, with Northwest Nexus for *Internet Starter Kit for Macintosh*), and the Take Control ebook series. In addition, he has collaborated on several Internet educational videos and has appeared on a variety of internationally broadcast television and radio programs.

Adam's indefatigable support of the Macintosh community and commitment to helping individuals has resulted in numerous awards and recognition at the highest levels. In the annual MDJ Power 25 survey of industry insiders from 2000 through 2007, he ranked in the top five most influential people in the Macintosh industry, and he was named one of *MacDirectory's* top ten visionaries. And how many industry figures can boast of being turned into an action figure?

AUTHORS' ACKNOWLEDGMENTS

Thanks to Tonya for all she does, both in editing this title and in keeping Take Control running.

A tip of the mouse to Chris Pepper, Larry Rosenstein, and Joe Kissell for their excellent comments during our collaborative editing phase.

SHAMELESS PLUGS

If you liked this title, you'll undoubtedly like our other works:

- ***TidBITS:*** For award-winning Apple commentary and editorial from both Adam and Glenn, be sure to read *TidBITS* (http://www.tidbits.com/).

- **Wi-Fi Networking News:** Glenn writes about Wi-Fi and other wireless networking daily or nearly so at this blog that dates back to early 2001. He has tracked developments like the rollout of hot-spots worldwide, new 802.11n hardware, the municipal wireless movement, and security problems and their solutions. Visit http://wifinetnews.com/.

ABOUT THE PUBLISHER

Publishers Adam and Tonya Engst have been creating Macintosh-related content since they started the online newsletter *TidBITS,* in 1990. In *TidBITS*, you can find the latest Macintosh news, plus read reviews, opinions, and more (http://www.tidbits.com/).

Adam and Tonya are known in the Mac world as writers, editors, and speakers. They are also parents to Tristan, who thinks ebooks about clipper ships and castles would be cool.



PRODUCTION CREDITS

Take Control logo: Jeff Tolbert

Cover design: Jon Hersh

Editor in Chief: Tonya Engst

Publisher: Adam Engst

Copyright and Fine Print

Take Control of Your Wi-Fi Security
ISBN: 978-0-975950-39-5
Copyright © 2010, Glenn Fleishman and TidBITS Publishing Inc.
All rights reserved.

TidBITS Publishing Inc.
50 Hickory Road
Ithaca, NY 14850 USA
http://www.takecontrolbooks.com/

Take Control electronic books help readers regain a measure of control in an oftentimes out-of-control universe. Take Control ebooks also streamline the publication process so that information about quickly changing technical topics can be published while it's still relevant and accurate.

This electronic book doesn't use copy protection because copy protection makes life harder for everyone. So we ask a favor of our readers. If you want to share your copy of this ebook with a friend, please do so as you would a physical book, meaning that if your friend uses it regularly, he or she should buy a copy. Your support makes it possible for future Take Control ebooks to hit the Internet long before you'd find the same information in a printed book. Plus, if you buy the ebook, you're entitled to any free updates that become available.

Although the author and TidBITS Publishing Inc. have made a reasonable effort to ensure the accuracy of the information herein, they assume no responsibility for errors or omissions. The information in this ebook is distributed "As Is," without warranty of any kind. Neither TidBITS Publishing Inc. nor the author shall be liable to any person or entity for any special, indirect, incidental, or consequential damages, including without limitation lost revenues or lost profits, that may result (or that are alleged to result) from the use of these materials. In other words, use this information at your own risk.

Many of the designations used to distinguish products and services are claimed as trademarks or service marks. Any trademarks, service marks, product names, or named features that appear in this title are assumed to be the property of their respective owners. All product names and services are used in an editorial fashion only, with no intention of infringement of the trademark. No such use, or the use of any trade name, is meant to convey endorsement or other affiliation with this title.

This title is an independent publication and has not been authorized, sponsored, or otherwise approved by Apple Inc. Because of the nature of this title, it uses terms that are trademarks or registered trademarks of Apple Inc.; to view a complete list of the trademarks and of the registered trademarks of Apple Inc., you can visit http://www.apple.com/legal/trademark/appletmlist.html.

Featured Titles

Click any book title below or visit our Web catalog to add more ebooks to your Take Control collection!

Take Control of Back to My Mac (Glenn Fleishman). Make the most of all your Internet-connected Macs via Back to My Mac, with this helpful guide. $10

Take Control of Screen Sharing in Snow Leopard (Glenn Fleishman). Figure out which type of screen sharing to use when and how to get the most out of screen sharing. $10

Take Control of the Mac Command Line with Terminal (Joe Kissell) Release your inner geek and learn to harness the power of the Unix underpinnings to Mac OS X! $10

Take Control of Users & Accounts in Snow Leopard (Kirk McElhearn): Find straightforward explanations of how to create, manage, and work with—and among—user accounts. $10

Take Control of Your 802.11n AirPort Network (Glenn Fleishman): Make your AirPort network fly—get help with buying the best gear, set up, security, and more. $15

Take Control of Apple Mail in Snow Leopard (Joe Kissell): Joe gets you going and helps you get the most out of Mail. He also gives detailed directions for how to sign and encrypt messages in Mail. $15

Take Control of Sharing Files in Snow Leopard (Glenn Fleishman): Find friendly advice and steps for sharing files from your Mac, and get further ideas for using an Internet-hosted service. $10

Take Control of Passwords in Mac OS X (Joe Kissell): Create and manage strong passwords that keep your data safe without taxing your memory! $10

Take Control of iPhone and iPod touch Networking & Security (Glenn Fleishman): Learn fascinating and practical geek-level details about iOS networking and security. Covers Wi-Fi and 3G networks. $15